

Stories that (are) Move(d by) Markets: A Causal Exploration of Market Shocks and Semantic Shifts across Different Partisan Groups

Felix Drinkall*, Stefan Zohren*, Michael McMahon†, Janet B. Pierrehumbert*‡

*Department of Engineering Science, University of Oxford

†Department of Economics, University of Oxford

‡Faculty of Linguistics, University of Oxford

felix.drinkall@eng.ox.ac.uk

Abstract

Macroeconomic fluctuations and the narratives that shape them form a mutually reinforcing cycle: public discourse can spur behavioural changes leading to economic shifts, which then result in changes in the stories that propagate. We show that shifts in semantic embedding space can be causally linked to financial market shocks – deviations from the expected market behaviour – using Granger causality tests. Furthermore, we show how partisanship can influence the predictive power of text for market fluctuations and shape reactions to those same shocks. We also provide some evidence that text-based signals are particularly salient during unexpected events such as COVID-19, highlighting the value of language data as an exogenous variable in economic forecasting. Our findings underscore the bidirectional relationship between news outlets and market shocks, offering a novel empirical approach to studying their effect on each other.

1 Introduction

Within economics and finance, a complex dialectic emerges between market events and the narratives that drive them. Narratives can drive changes in beliefs which drive behavioural changes that lead to macroeconomic shifts; but equally, narratives can be shaped by the observable behavioural changes and the economic data. Whilst the link between narratives and economic actions has been the subject of many influential works (Akerlof and Shiller, 2010; Shiller, 2017), the two-way endogeneity makes this complex relationship very hard to model. One popular approach is to focus on how specific historical narratives persisted and their effect on the world (Margalit, 2019; Tangherlini et al., 2020; Ash et al., 2024). However, developing a predictive system capable of understanding the causal influence between narratives and the economy, and vice versa, is a far more challenging task – one that this paper aims to address.

Economics has long incorporated exogenous data (Einav and Levin, 2014) – information which exists outside of the economic system itself. Linguistic behaviour from news and social media can act as a proxy for consumer beliefs, which adds nuance to predictive feature sets, contingent on the text encoding methodology. Existing efforts to model the relationship between the real world and written or spoken text have tended to focus on more interpretable and simplistic text representations, such as lexicons (Loughran and McDonald, 2011) and sentiment (Fazlija and Harder, 2022). Some efforts have centred around extracting narrative-form structures in the form of subject-verb-object triplets to understand the nature and directionality of what is contained in the text (Piper et al., 2021; Ash et al., 2024). However, narrative extraction can suffer from sparsity issues limiting its use in predictive applications. The recent progress in transformer-based LLMs potentially solves this problem by looking at the text in a distributional rather than an extractive manner – perhaps another example of the "Bitter Lesson" (Sutton, 2019). LLMs are able to encode more granular text features than extractive methods. This paper leverages the representational power of transformer-based LLMs to understand the bidirectional relationship between text and the macroeconomy.

Different narratives can circulate about the same economic event, with these differences sometimes stemming from partisan differences (Epstude et al., 2022; Hobolt et al., 2024). News outlets of various political orientations – centrist, right-wing, and left-wing – can provide contrasting interpretations of identical events (Guess et al., 2021). To better understand the relationship between news outlets and the economy, it is important to examine how different partisan groups influence and react to financial markets. Existing research typically offers broad explanations for how events, such as financial crises or elections, shape subsequent po-

litical reactions (Klüver and Spoon, 2016; Margalit, 2019). This paper looks at political groups separately and on a daily basis to understand the differences between how partisan groups respond to and precede market shocks.

Feedback loops are of great interest to researchers and practitioners attempting to predict virality. News interest can increase attention which can itself increase the news interest, entering a self-reinforcing cycle. Viral feedback loops can exist in multiple domains: tourism (Murray and Graham, 1997), social media (Metzler and Garcia, 2024), consumer products (Nguyen and Chaudhuri, 2019), and other areas of the financial markets (Semenova et al., 2024). This paper attempts to quantify feedback loops from a linguistic diachronic change perspective paving a path for other domains.

This paper makes the following contributions:

- First paper to show that shifts in semantic embedding space can be causally linked to measurable real-world shocks.
- We conduct an analysis of how partisanship affects the ability of text to predict market shocks and the reaction to market shocks.
- We show that feedback loops exist between the market and news outlets.
- We find evidence that text is an important exogenous variable in violent economic events such as COVID-19.

2 Literature Review

2.1 Text-based Forecasting

Text-based forecasting explores the predictive power of textual data such as news articles (Peng and Jiang, 2016; Zhai and Zhang, 2023; Wang et al., 2024; Rahimikia et al., 2024), social media (Arias et al., 2014; Iso et al., 2016), or corporate reports (Kogan et al., 2009). By extracting linguistic features such as keywords (Iso et al., 2016), sentiment (Fazlija and Harder, 2022), raw embedding representations (Sawhney et al., 2020) or high-density embedding clusters (Drinkall et al., 2025a,b), researchers have shown that it is possible to enhance traditional forecasting models using text-derived information. The findings of these papers imply there is some causal interaction between textual data and the economy. Yet to our knowledge, no papers have explored the bidirectional nature of this relationship; understanding the effect that markets have on news discourse.

2.2 Diachronic Shifts

Diachronic analysis (Periti and Montanelli, 2024) examines how linguistic phenomena change over time. The process begins by aggregating time-stamped information to create time-specific representations. The aggregation process can be done using averaging (Rodina et al., 2021) or clustering (Martinc et al., 2020b), and is considered a design choice dependent on what one is measuring; with averaging suiting cases which aim to detect the dominant shifts. Diachronic shifts tend to be measured using a distance metric between the point-in-time representations. There are several distance metrics that are used in the literature: Cosine Distance (CD) (Martinc et al., 2020a; Horn, 2021), Average Pairwise Distance (APD) (Giulianelli et al., 2020; Kудisov and Arefyev, 2022), Hausdorff Distance (HD) (Wang et al., 2020) etc. The standard distance metric is CD, which we use in this paper, since APD is more commonly used for polysemy detection (Keidar et al., 2022) and HD is sensitive to outliers (Wang et al., 2020). A similar conceptual metric, Cosine Similarity, is widely used in other NLP applications as a robust and efficient similarity metric (Thongtan and Phienthrakul, 2019; Reimers and Gurevych, 2019; Yamagiwa et al., 2025).

2.3 Causality

2.3.1 Causality in Time-series

Causality in time-series can be inferred using methods like transfer entropy (Barnett et al., 2009), convergent cross mapping (Ye et al., 2015) and dynamic Bayesian networks (Ghahramani, 1997); these techniques model non-linear effects but are poor at isolating individual lag contributions. Granger causality (Granger, 1969) tests whether past values of one variable can help to predict the future value of another, enabling analysis of individual lag contributions. It does not imply “causation” in a philosophical sense but identifies temporal precedence and predictability. If the inclusion of past values of variable X significantly improves the prediction of variable Y , then X is said to “Granger cause” Y . Non-linear regression models can be used in a Granger causality test to model non-linear relationships (Marinazzo et al., 2008).

2.3.2 Causality in Text

There are three main research areas that explore “causal” interactions in NLP. Firstly, autoregres-

sive generation which underpins modern generative LLMs is sometimes referred to as "causal language modelling" (Radford et al., 2018, 2019; Brown et al., 2020), but this area is not related to the identification of causal relationships with exogenous data. Secondly, causal relation extraction (Ning et al., 2018; Dasgupta et al., 2018; Yang et al., 2021) identifies causal relationships between clauses. The final causal NLP research area which is of interest to this paper creates a link between text and observable phenomena in the real world. Text-derived features can be used in causal inference where text is the treatment variable – what causes the outcome – for a temporally static outcome. Past work has found success in medical records (Zeng et al., 2021), politics (Fong and Grimmer, 2016), and mental health (Zhang et al., 2020). This research direction is mature and has developed solutions that encode text with topic models (Roberts et al., 2020) as well as text embeddings (Veitch et al., 2019). However, there has been little work exploring the predictive causal relationship between text and time-series information. Existing methods do not link changes in an LLM-derived semantic space with observable time-series, instead focusing on keywords (Balashankar et al., 2019; Mariano et al., 2022), lexicons, and LDA topics (Kang et al., 2017), making our framework novel. Further to this, no techniques measure the temporal causality of observable time-series on text.

3 Datasets

3.1 News Dataset

We use the news dataset from Spliethöver et al. (2022), which spans a large time period from 2000 to 2022 and includes the political orientation labels from AllSides.com. The news outlets included in this paper, alongside their partisan label, are listed in §B. We use data from 01/07 to 02/22 to ensure sufficient coverage.

Text Representation. Let S be the set of synonyms for the terms $\{prices, inflation, labor, growth\}$. Let \mathcal{O} be the set of news outlets, and let t index each day in the dataset. Related words were isolated by identifying the top 50 most similar words in the vocabulary of *all-mpnet-base-v2*, and then removing problematic polysemic words.

Article Selection and Embedding. For each day t and outlet $o \in \mathcal{O}$, all articles were stored that referenced at least one synonym in S . The em-

bedding $\mathbf{e}_{t,o,k}$ of the k -th article from outlet o on day t , was obtained using the *all-mpnet-base-v2* model (Reimers and Gurevych, 2019). If an article exceeded the model’s maximum sequence length, it was split into chunks, and the chunk embeddings were averaged to form a single article representation $\mathbf{e}_{t,o,k}$. We picked the token number chunking methodology over sentence splitting as it is less likely to overweight a section within a paragraph. If sentence embeddings were averaged, then smaller sentences would overweighted in the final representation. An encoder-based model is a natural choice for this task as opposed to a large generative model because bidirectional attention models often perform better on embedding tasks (Song et al., 2020; Muennighoff et al., 2023; Lee et al., 2025).

Daily Outlet Embeddings. Let $N_{t,o}$ be the number of relevant articles published by outlet o on day t . The *daily outlet embedding* $\mathbf{E}_{t,o}$ for outlet o on day t was calculated by averaging over the articles:

$$\mathbf{E}_{t,o} = \frac{1}{N_{t,o}} \sum_{k=1}^{N_{t,o}} \mathbf{e}_{t,o,k}$$

provided $N_{t,o} > 0$. If $N_{t,o} = 0$, then the embeddings were forward filled:

$$\mathbf{E}_{t,o} = \mathbf{E}_{t-1,o}$$

This was to make sure that in the outlets that were more sparse, the significant jumps in the time-series were related to semantic shifts, rather than whether economic articles appeared or not. This preprocessing step did not affect the majority of outlets that published an article every day.

News Feature. The *news feature* for outlet o on day t is the cosine distance between the daily embeddings of consecutive days:

$$\text{NewsFeature}_{t,o} = 1 - \frac{\mathbf{E}_{t,o} \cdot \mathbf{E}_{t-1,o}}{\|\mathbf{E}_{t,o}\| \|\mathbf{E}_{t-1,o}\|}$$

This measure quantifies the day-to-day semantic shift for each outlet. We used this inline with the diachronic literature outlined in §2.2 and because it can be interpreted as a textual shock, mirroring the market shocks that we introduce in §3.2. We only consider outlets where there is an article published on 25% of the days. The similarity in data format allows us to create a symmetric comparison between news and market data.

3.2 Market Shocks

To study the effect of evolving market beliefs on narrative, and vice versa, we ideally need high-frequency data. Unfortunately, most economic indices are reported monthly or quarterly, which would miss the transmission of short-term shocks to the economy (Chen and Du, 2015). News cycles can last a matter of days (Castillo et al., 2014), motivating the need for more granular time series. We therefore focus on market data which has the advantages that (i) it is updated at high frequency on trading days (we use daily data), and (ii) it is understood to reflect the beliefs and information of market participants and their desire to make money means that new information should be reflected quickly in asset prices. Specifically, we use daily asset-price-derived indicators from Cieslak and Pang (2021) that capture four distinct types of economic news: *growth*, *monetary*, *common premium*, and *hedging premium* shocks. These shocks are derived by first purging stock prices (S&P 500) and bond yields of different maturities (two-, five- and ten-year) of their movements that are predictable from systematic relationships between the asset prices. By exploiting economic reasoning, the remaining variation in asset prices can be decomposed into the four distinct types of news:¹

- **Growth** (ω_g): Positive growth news raises stock prices and bond yields at short-to-medium maturities.
- **Monetary (tightening)** (ω_m): Positive monetary news lowers stock prices and increases short-to-medium bond yields, with a weaker effect at longer maturities.
- **Common premium** (ω_{cp}): An increase in the discount rate risk premium (often linked to inflation) boosts bond prices, reducing yields, and lowers stock prices, creating negative co-movement between yield changes and stock returns. This effect is stronger at the longer maturities.
- **Hedging premium** (ω_{hp}): Elevated future cash flow risk depresses stock prices and raises long-term bond yields (safe haven effect), generating positive co-movement be-

¹Formally, Cieslak and Pang (2021) identify the shocks via sign restrictions on a structural vector autoregression (SVAR). The procedure is standard in economics and explained in §C.

tween yield changes and stock returns. This effect is stronger at the longer maturities.

Growth and monetary policy shocks mainly affect shorter maturities, reflecting mean reversion around recent economic news, whereas risk-premium shocks manifest at longer maturities, driven by longer-term uncertainty. The resultant time-series encapsulate structural deviations from the expected market behaviour which should reflect new beliefs (howsoever arrived at) as earlier information is already reflected in the market prices. The result is four daily orthogonal time-series, each representing the market's evolving interpretation of economic events via their effects on both stock and bond markets.

4 Methodology

The goal of the paper is to determine how, if at all, market time-series affect the narrative in the media and narratives in the media drive market prices. That is, whether there is any information in the textual media time-series that is predictive of market time-series, or vice versa. To do this we devised an evaluation tool using Granger causality tests, the steps of which are outlined in the following section:

1. **Lagged Feature Creation:** For each target variable y_t , we construct lagged predictors $y_{t-\ell}$ and x_{t-L} , $\ell \in \{1, \dots, L\}$.
2. **Estimation and Prediction:** We fit both the base and enhanced models on the available training data and obtained out-of-sample predictions $\hat{y}_{B,t}$ and $\hat{y}_{E,t}$, for every news outlet. The enhanced regression model predicts $\hat{y}_{E,t}$ using $y_{t-\ell}$ and x_{t-L} , whereas the base model predicts $\hat{y}_{B,t}$ using only $y_{t-\ell}$.
3. **Error Evaluation:** We compute errors $e_{B,t}$ and $e_{E,t}$.
4. **Statistical Testing:** We calculate $d_t = e_{B,t} - e_{E,t}$ and apply paired tests to determine whether differences in prediction errors are significant.

This methodology allows us to isolate and quantify the effect of additional explanatory features relative to a baseline that depends solely on past values of the target variable.

4.1 Model Specification

We assess the predictive value of text-based features by comparing two regression model configurations: a *base model* and an *enhanced model*. The base model incorporates only lagged values of the target variable, while the enhanced model adds lagged text-derived or macroeconomic predictors, depending on the domain of the target variable.

Base Model. Let y_t be the target variable at time t , and suppose the base model uses L lagged observations of y as predictors. We denote the parameter set of the base model by θ_B :

$$y_t = f_B(y_{t-1}, y_{t-2}, \dots, y_{t-L} \mid \theta_B) + \varepsilon_{B,t} \quad (1)$$

Enhanced Model. In the enhanced model, we introduce a single additional predictor x_{t-L} (e.g., news outlet cosine distances or market shocks) at the same lag level. We denote the parameter set of the enhanced model by θ_E :

$$y_t = f_E(y_{t-1}, \dots, y_{t-L}, x_{t-L} \mid \theta_E) + \varepsilon_{E,t} \quad (2)$$

4.2 Regression Models

We employ two regression models: *Linear Regression* (Linear) and *Kernel Ridge Regression* (KRR) with a radial basis function (RBF) kernel. Comparing the two allows us to determine whether the relationship between the input features and the target variable is linear or non-linear in nature.

Linear Regression. In the linear case, we use an OLS loss function that identifies a parameter vector β that minimizes the squared error:

$$\min_{\beta} \|\mathbf{y} - \mathbf{X}\beta\|^2 \quad (3)$$

where $\mathbf{X} \in \mathbb{R}^{n \times p}$ denotes the matrix of input features, and $\mathbf{y} \in \mathbb{R}^n$ represents the target values.

Kernel Ridge Regression. In the non-linear case – as in [Marinazzo et al. \(2008\)](#) – we use an RBF kernel $K(\mathbf{x}_i, \mathbf{x}_j)$ that projects inputs into a high-dimensional space. The KRR solution α is found by solving:

$$\min_{\alpha} \|\mathbf{y} - \mathbf{K}\alpha\|^2 + \lambda \alpha^\top \mathbf{K}\alpha \quad (4)$$

where $\lambda > 0$ controls the trade-off between data fidelity and model complexity.

4.3 Error Metrics

Let $\hat{y}_{B,t}$ and $\hat{y}_{E,t}$ be the predictions from the base and enhanced models, respectively, for time t . We denote the prediction errors as:

$$e_{B,t} = \hat{y}_{B,t} - y_t \quad \text{and} \quad e_{E,t} = \hat{y}_{E,t} - y_t$$

We use MSE error to determine the performance of the model configurations:

$$\text{MSE} = \frac{1}{T} \sum_{t=1}^T (e_{M,t})^2 \quad (5)$$

where $M \in \{B, E\}$ (base or enhanced), and T is the total number of test observations.

4.4 Statistical Significance Test

We evaluate if $e_{E,t}$ is significantly lower than $e_{B,t}$ with the following null hypothesis:

$$d_t = e_{B,t} - e_{E,t} \quad H_0 : \mathbb{E}[d_t] \leq 0$$

against the one-sided alternative $H_1 : \mathbb{E}[d_t] > 0$. We conduct the following paired t-test on the sample $\{d_t\}_{t=1}^T$:

$$T_{\text{stat}} = \frac{\bar{d}}{\sqrt{\widehat{\text{Var}}(d)/T}}$$

where $\bar{d} = \frac{1}{T} \sum_{t=1}^T d_t$ and $\widehat{\text{Var}}(d)$ is the sample variance of d_t . A statistically significant positive mean difference $\bar{d} > 0$ provides evidence that including the additional features reduces prediction errors on average.

4.5 P-hacking

P-hacking arises when a large number of hypothesis tests are conducted, increasing the risk of false positives. Let $\{p_i\}_{i=1}^N$ denote the collection of p -values from N individual tests, each compared against a standard threshold $\alpha = 0.05$. Define

$$\hat{p} = \frac{1}{N} \sum_{i=1}^N \mathbf{1}\{p_i < \alpha\}$$

as the empirical proportion of observations deemed significant, to see if this proportion truly exceeds the 5% expected by chance, we perform a one-sided binomial test under the null hypothesis that $\hat{p} = 0.05$. We conclude that a group-level effect is present only if the binomial test p-value is below α , indicating that \hat{p} is significantly above 0.05.

Moreover, to determine whether a specific pair is significant in §5.2 and §5.3, we use a Bonferroni-adjusted threshold $p_i < \frac{\alpha}{M}$, where M is the total number of tests (Bonferroni, 1936) across the partisan groups. This procedure reduces the likelihood of false discoveries. In fact, there are multiple studies that suggest that the Bonferroni adjustment is too conservative and it promotes false negative results (Perneger, 1998; Anderson et al., 2000; Nakagawa, 2004). Given this strict definition of significance, if a pair does emerge as significant we can be more confident that the effect is real.

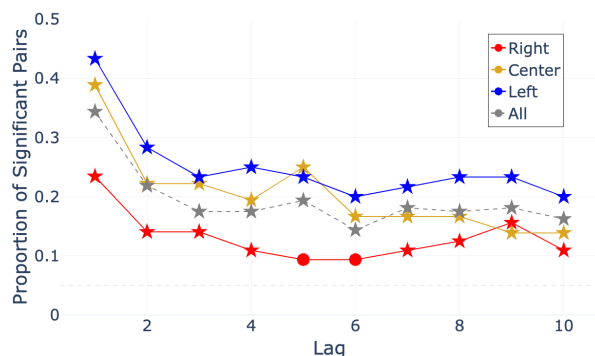
5 Results

Unless specified the following results are calculated using a static temporal split, where the training set constitutes 70% of the samples, the test set made up the final 30%, and the validation was created using the final 10% of the training set. The first and most elementary result is outlined in Tab. 1, which suggests that while all of the configurations exceed the significance threshold, the non-linear regression model identifies more significant pairs than the linear regression model. As a result, we use the non-linear KRR regression model for the remainder of the results in this section.

	KRR	Linear
$y = \text{text}, x = \text{econ}$	19.5*	6.69*
$y = \text{econ}, x = \text{text}$	6.34*	6.46*

Table 1: Proportion (%) of significant pairs across all political groups, news outlets and market shocks. An * indicates that the effect is significant. Full data - §A.

The first contribution on decomposing the interrelations between narrative and market data is




(a) Sign. proportion at different lags for $\{y = \text{text}, x = \text{econ}\}$.

our finding that market shocks have a greater influence on changes in text than the other way around. This is hardly surprising, but it does imply that news outlets are more reactive than predictive and that the markets are generally efficient (with some exceptions explored below) – pricing in changes before news outlets change their narrative. However, this table averages over the more granular partisan group effects.

5.1 Partisan Differences


Fig. 1 outlines a more nuanced story. Fig. 1a shows that the effect that market shocks have on news discourse approximates a negative exponential. It is also clear that the left-wing and centrist outlets are more reactive to market shocks than right-wing outlets. This relative difference is persistent across time but is more pronounced at shorter time horizons. Fig. 1b shows that at longer time horizons (5-10 days), news outlets do contain some useful predictive information about the market. This suggests that when a significant shift occurs in how a news outlet writes about the economy, there is a delay of around a week before a market correction. We also see that the different political groups display different distributions of significant pairs across the lags - the same as Fig. 1a. The differences between the partisan groups are less pronounced but the results suggest that there is a longer delay for the market to react to a news semantic change from a centrist outlet and that right-wing outlets only have an effect on the market at lags 5-7. Some of the significant results

Fig. 2 shows the proportion of significant pairs for each market shock, aggregated across all lags for each political orientation. It is clear that different political groups react to and precede different




(b) Sign. proportion at different lags for $\{y = \text{econ}, x = \text{text}\}$.

Figure 1: The proportion of significant pairs for when $\{y = \text{text}, x = \text{econ}\}$ and $\{y = \text{econ}, x = \text{text}\}$. Significant proportions are marked with a star.



(a) $\{y = \text{text}, x = \text{econ}\}$. News outlets responding to market shocks.



(b) $\{y = \text{econ}, x = \text{text}\}$. News outlet changes preceding market shocks.

Figure 2: Significant proportion for different market shocks for each political groups across all lags. An * denotes a significant effect inline with §4.4 and §4.5, and the shade of blue represents the size of the group effect, with darker colours pertaining to larger effects.

market shocks. Similar to Fig. 1, the overall level of significant pairs is high in left-wing news outlets. Fig. 2b shows that there is not much partisan skew in the proportion of news outlets that respond to growth shocks, but that there are large differences in the response to other market shocks. Centrist outlets are particularly reactive to common risk premium shocks (ω_{cp}) - a shock related to inflation risk. Fig. 2a shows that left-wing outlets exhibit a higher proportion of significant pairs for ω_{cp} . The left-wing rate of 10.7% is significantly higher than the centrist and right-wing rates, providing a predictive opportunity for researchers predicting this market shock. It seems that, in general, news outlets contain more information useful for predicting changes in ω_{cp} than any other shock. This could be because news about inflation is more likely to influence consumer behaviour because it affects

everyone, which in turn affects market movements. In contrast, news outlets of all political groups contained little information useful for predicting growth shocks. Some of the reporting effects were weak, despite being significant. It is likely because the news is not always affected by or causal for the overall market. Even for apparently obvious causal pairs, the results show that the causal pair is not significant for all outlets all of the time.

5.2 Confounders

Confounders represent variables that cause both the treatment and outcome variables to change (VanderWeele, 2019). This can result in false-positive causal pairs being identified. To test whether any of the significant pairs, identified using the Bonferroni correction, were affected by confounding variables, we tested for Granger causality with the other news outlets, market shocks, emotion scores aggregated over all of the news outlets and market volatility. We used the Bonferroni correction to determine the significance level. This resulted in a modest 5.36% of significant pairs having a confounding variable. The most common confounder type was the news emotion scores, suggesting that overall emotion can sometimes drive both changes in news discourse and market trends. Given the low confounder identification rate, it is clear that for the majority of significant pairs, the mutual interaction is direct. The full analysis is outlined in §D.

5.3 Feedback Loops

Identifying market shock and news outlet pairs that are susceptible to feedback loops can help researchers identify potential viral behaviour. Circular causality was identified by using the Bonferroni adjusted p-values to infer a significant degree of predictive causality, if both variables in the pair were deemed to cause one another, irrespective of lag, via the methodology in §4.5, then circular causality was flagged. Although there were not many pairs that exhibited circular causality in our study, the pairs with feedback loops tended to involve ω_{cp} shocks: 4.22% of ω_{cp} pairs displayed feedback loops. Monetary policy ω_m shocks also showed a modest presence of circular dynamics with 1.56% of all pairs containing feedback loops, exceeding the minimal 0.2% seen in hedging risk premium ω_{hp} shocks. Growth-related shocks ω_g exhibited no feedback loops, which may be explained by the fact that news outlets provide limited information for growth shocks (Fig. 2).

There are a number of reasons why feedback loops might tend to form with ω_{cp} shocks. Firstly, inflation-linked shocks often ignite widespread media coverage, amplifying investors' sensitivity to any hint of price instability. Secondly, frequent and sometimes conflicting potential shock-inducing events linked to inflation – such as commodity prices, central bank signals, or wage growth – provide fertile ground for repeated bouts of speculation. Thirdly, repeated episodes of speculation can quickly magnify market moves, drawing additional coverage and further intensifying the original inflation narrative. There are fewer potential shock-inducing events for ω_m shocks, which could explain the lower amount of feedback loops compared to ω_{cp} shocks. Hedging premium shocks ω_{hp} - which arise when markets are concerned with long-horizon cash flow - are more likely to occur in rare episodic crises when there is a demand for safe assets. Fewer shock-inducing events reduce the likelihood of feedback loops. The shock type is also less intuitive, meaning that news may be less likely to cover such shocks.

5.4 Temporal Effect

To examine how news outlets react to or precede market shocks during different time-periods, we used a rolling evaluation framework. Each window spanned 365 days of data, with the last 90 days in each window forming the test set; consecutive windows were separated by 180 days. Within each window v , let MSE_B^v and MSE_E^v be the test-set mean-squared errors of the base (B) and enhanced (E) models, respectively. We define the MSE improvement of model E over B for window v as Δ_v . Let $\bar{\Delta}$ denote the average improvement across all test windows. Figure 3 plots the deviation $\Delta_v - \bar{\Delta}$ for each test set. Points above zero represent windows in which the enhanced model's improvement over the base model exceeded the average improvement, whereas points below zero indicate windows where the enhanced model's improvement fell short.

The most striking outlier occurs during the 03/20 to 05/20 test window, highlighting an atypical change in the performance differential between the base and enhanced models. The result shows that information from text improved the ability of models to predict market shocks during COVID-19, whereas market shocks were less useful in predicting changes in news discourse. Financial markets during this time exhibited unusual behaviour, which meant that fewer movements could be en-




Figure 3: Deviation in MSE improvement of the enhanced model over the base model, $\Delta_v - \bar{\Delta}$, for each test window. Different orientations are indicated with the following colors: **left**, **right**, **center**, and **all**.

dogenously explained using historical market information resulting in more structural shocks. It also suggests that during this time the markets were highly reactive to outside news, showing that text is particularly useful during market events which are being mostly driven by changes exogenous to the market, such as COVID-19. Fig. 3a also shows that market information was worse at predicting semantic changes in news outlets during early 2020, which could be because news was being driven by the pandemic rather than the market shocks.

6 Conclusion

In this paper, we find that left-wing and centrist outlets exhibit a stronger and faster response to market changes, whereas right-wing outlets show a smaller effect. Our novel methodology demonstrates, that shifts in semantic embedding space can be causally linked to measurable real-world shocks. By applying a bidirectional Granger-causality framework to text and economic indicators, we uncover how partisanship shapes both the predictability of market shocks and the reactivity of news discourse. We also identify some feedback loops between market shocks and news outlets. Additionally, we provide strong evidence that text-based features serve as essential exogenous variables in unexpected events, such as COVID-19. Our work underscores the value of semantic representations for understanding

and anticipating complex economic phenomena.

Future Work. Building on our findings, several directions emerge for future investigation. A natural extension involves applying our methodology across international and multilingual contexts, where cultural differences could influence the sensitivity of media discourse to market shocks. Additionally, adopting richer embedding techniques or domain-adapted language models may further illuminate subtle shifts in discourse, particularly during black-swan events such as pandemics or geopolitical crises. Incorporating more granular lag structures that are shorter than 1 day could yield finer temporal insights into the interplay between markets and news discourse. There is also potential to expand our causal framework by integrating unstructured shocks or continuing our confounder analysis in §D to create a causal graph that captures interdependencies among multiple shocks and outlets. Finally, evaluating the robustness of our approach in real-time forecasting environments could pave the way for practical applications in finance, policymaking, and risk management.

7 Limitations

While our methodology establishes a novel framework for linking semantic embedding shifts to real-world financial market shocks, some limitations should be acknowledged. First, although we employed multiple regression models (e.g., linear, kernel-based), further scrutiny with deep learning architectures or ensemble methods could yield different insights. Second, our focus on English-language news articles limits the generalisability to other linguistic or cultural contexts where the relationship between media discourse and markets might follow different dynamics. Third, we focused on identifying the presence and strength of causal relationships and did not cover directionality, representing an opportunity for future research. Finally, we relied on predefined partisan labels and consolidated market shock definitions, but these categorisations can be fluid and may not capture all subtleties in political orientation or economic behaviour.

Acknowledgments

The research was funded by the Alan Turing Institute, the UK Defence Science and Technology Laboratory, and the European Research Council (Consolidator Grant Agreement 819131). The first

author was funded by the Economic and Social Research Council of the UK via the Grand Union DTP. This work was supported in part by a grant from the Engineering and Physical Sciences Research Council (EP/T023333/1). We are also grateful to the Oxford-Man Institute of Quantitative Finance, the Oxford e-Research Centre, and the Department of Economics for their support.

References

- George A. Akerlof and Robert J. Shiller. 2010. *Animal Spirits: How Human Psychology Drives the Economy, and Why It Matters for Global Capitalism*. Princeton University Press.
- David R Anderson, Kenneth P Burnham, and William L Thompson. 2000. [Null Hypothesis Testing: Problems, Prevalence, and an Alternative](#). *The Journal of Wildlife Management*, pages 912–923.
- Marta Arias, Argimiro Arratia, and Ramon Xuriguera. 2014. [Forecasting with Twitter Data](#). *ACM Trans. Intell. Syst. Technol.*, 5(1).
- Elliott Ash, Germain Gauthier, and Philine Widmer. 2024. [Relatio: Text Semantics Capture Political and Economic Narratives](#). *Political Analysis*, 32(1):115–132.
- Ananth Balashankar, Sunandan Chakraborty, Samuel Fraiberger, and Lakshminarayanan Subramanian. 2019. [Identifying Predictive Causal Factors from News Streams](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 2338–2348, Hong Kong, China. Association for Computational Linguistics.
- Lionel Barnett, Adam B Barrett, and Anil K Seth. 2009. [Granger Causality and Transfer Entropy are Equivalent for Gaussian Variables](#). *Physical review letters*, 103(23):238701.
- C.E. Bonferroni. 1936. *Teoria Statistica Delle Classi e Calcolo Delle Probabilità*. Pubblicazioni del R. Istituto Superiore di Scienze Economiche e Commerciali di Firenze. Seeber.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language Models are Few-Shot Learners](#). In *Advances in Neural Information Processing Systems*

- 33 (*NeurIPS 2020*). NeurIPS 2020, December 6–12, Virtual Conference.
- Carlos Castillo, Mohammed El-Haddad, Jürgen Pfeffer, and Matt Stempeck. 2014. [Characterizing the Life Cycle of Online News Stories Using Social Media Reactions](#). In *Proceedings of the 17th ACM Conference on Computer Supported Cooperative Work & Social Computing, CSCW '14*, page 211–223, New York, NY, USA. Association for Computing Machinery.
- Shu-Heng Chen and Ye-Rong Du. 2015. [Granularity in Economic Decision Making: An Interdisciplinary Review](#). *Granular Computing and Decision-making: Interactive and Iterative Approaches*, pages 47–71.
- Anna Cieslak and Hao Pang. 2021. [Common Shocks in Stocks and Bonds](#). *Journal of Financial Economics*, 142(2):880–904.
- Tirthankar Dasgupta, Rupsa Saha, Lipika Dey, and Abir Naskar. 2018. [Automatic Extraction of Causal Relations from Text using Linguistically Informed Deep Neural Networks](#). In *Proceedings of the 19th Annual SIGdial Meeting on Discourse and Dialogue*, pages 306–316, Melbourne, Australia. Association for Computational Linguistics.
- Felix Drinkall, Janet B. Pierrehumbert, and Stefan Zohren. 2025a. [Forecasting Credit Ratings: A Case Study where Traditional Methods Outperform Generative LLMs](#). In *Proceedings of the Joint Workshop of the 9th Financial Technology and Natural Language Processing (FinNLP), the 6th Financial Narrative Processing (FNP), and the 1st Workshop on Large Language Models for Finance and Legal (LLMFinLegal)*, pages 118–133, Abu Dhabi, UAE. Association for Computational Linguistics.
- Felix Drinkall, Janet B. Pierrehumbert, and Stefan Zohren. 2025b. [When Dimensionality Hurts: The Role of LLM Embedding Compression for Noisy Regression Tasks](#).
- Liran Einav and Jonathan Levin. 2014. [Economics in the Age of Big Data](#). *Science*, 346(6210):1243089.
- Kai Epstude, Daniel A. Effron, and Neal J. Roese. 2022. [Polarized Imagination: Partisanship Influences the Direction and Consequences of Counterfactual Thinking](#). *Philosophical Transactions of the Royal Society B: Biological Sciences*, 377(1866):20210342.
- Bledar Fazlija and Pedro Harder. 2022. [Using Financial News Sentiment for Stock Price Direction Prediction](#). *Mathematics*, 10(13).
- Christian Fong and Justin Grimmer. 2016. [Discovery of treatments from text corpora](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1600–1609, Berlin, Germany. Association for Computational Linguistics.
- Renee Fry and Adrian Pagan. 2011. [Sign Restrictions in Structural Vector Autoregressions: A Critical Review](#). *Journal of Economic Literature*, 49(4):938–960.
- Zoubin Ghahramani. 1997. [Learning Dynamic Bayesian Networks](#). *International School on Neural Networks, Initiated by IIASS and EMFCSC*, pages 168–197.
- Mario Giulianelli, Marco Del Tredici, and Raquel Fernández. 2020. [Analysing Lexical Semantic Change with Contextualised Word Representations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3960–3973, Online. Association for Computational Linguistics.
- C. W. J. Granger. 1969. [Investigating Causal Relations by Econometric Models and Cross-spectral Methods](#). *Econometrica*, 37(3):424–438.
- Andrew M. Guess, Pablo Barberá, Simon Munzert, and JungHwan Yang. 2021. [The Consequences of Online Partisan Media](#). *Proceedings of the National Academy of Sciences*, 118(14):e2013464118.
- Sara B. Hobolt, Katharina Lawall, and James Tilley. 2024. [The Polarizing Effect of Partisan Echo Chambers](#). *American Political Science Review*, 118(3):1464–1479.
- Franziska Horn. 2021. [Exploring Word Usage Change with Continuously Evolving Embeddings](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*, pages 290–297, Online. Association for Computational Linguistics.
- Hayate Iso, Shoko Wakamiya, and Eiji Aramaki. 2016. [Forecasting Word Model: Twitter-based Influenza Surveillance and Prediction](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pages 76–86, Osaka, Japan. The COLING 2016 Organizing Committee.
- Dongyeop Kang, Varun Gangal, Ang Lu, Zheng Chen, and Eduard Hovy. 2017. [Detecting and Explaining Causes From Text For a Time Series Event](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2758–2767, Copenhagen, Denmark. Association for Computational Linguistics.
- Daphna Keidar, Andreas Opedal, Zhijing Jin, and Mrinmaya Sachan. 2022. [Slangvolution: A Causal Analysis of Semantic Change and Frequency Dynamics in Slang](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1422–1442, Dublin, Ireland. Association for Computational Linguistics.
- Heike Klüver and Jae-Jae Spoon. 2016. [Who Responds? Voters, Parties and Issue Attention](#). *British Journal of Political Science*, 46(3):633–654.
- Shimon Kogan, Dmitry Levin, Bryan R. Routledge, Jacob S. Sagi, and Noah A. Smith. 2009. [Predicting Risk from Financial Reports with Regression](#).

- In *Proceedings of Human Language Technologies: The 2009 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 272–280, Boulder, Colorado. Association for Computational Linguistics.
- Artem Kudisov and Nikolay Arefyev. 2022. [BOS at LSCDiscovery: Lexical Substitution for Interpretable Lexical Semantic Change Detection](#). In *Proceedings of the 3rd Workshop on Computational Approaches to Historical Language Change*, pages 165–172, Dublin, Ireland. Association for Computational Linguistics.
- Chankyu Lee, Rajarshi Roy, Mengyao Xu, Jonathan Raiman, Mohammad Shoeybi, Bryan Catanzaro, and Wei Ping. 2025. [NV-Embed: Improved Techniques for Training LLMs as Generalist Embedding Models](#).
- Tim Loughran and Bill McDonald. 2011. [When is a Liability not a Liability? Textual Analysis, Dictionaries, and 10-Ks](#). *The Journal of finance*, 66(1):35–65.
- Yotam Margalit. 2019. [Political Responses to Economic Shocks](#). *Annual Review of Political Science*, 22:277–295. First published as a Review in Advance on January 23, 2019.
- Maisonnave Mariano, Delbianco Fernando, A. Tohmé Fernando, E. Milios Evangelos, and Gabriela Maguitman Ana. 2022. [Causal Graph Extraction from News: A Comparative Study of Time-series Causality Learning Techniques](#). *PeerJ Computer Science*, 8.
- Daniele Marinazzo, Mario Pellicoro, and Sebastiano Stramaglia. 2008. [Kernel Method for Nonlinear Granger Causality](#). *Physical review letters*, 100(14):144103.
- Matej Martinc, Petra Kralj Novak, and Senja Pollak. 2020a. [Leveraging Contextual Embeddings for Detecting Diachronic Semantic Shift](#). In *Proceedings of the Twelfth Language Resources and Evaluation Conference*, pages 4811–4819, Marseille, France. European Language Resources Association.
- Matej Martinc, Syrielle Montariol, Elaine Zosa, and Lidia Pivovarov. 2020b. [Capturing Evolution in Word Usage: Just Add More Clusters?](#) In *Companion Proceedings of the Web Conference 2020*, WWW ’20, page 343–349, New York, NY, USA. Association for Computing Machinery.
- H. Metzler and D. Garcia. 2024. [Social Drivers and Algorithmic Mechanisms on Digital Media](#). *Perspectives on Psychological Science*, 19(5):735–748.
- Niklas Muennighoff, Nouamane Tazi, Loic Magne, and Nils Reimers. 2023. [MTEB: Massive Text Embedding Benchmark](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2014–2037, Dubrovnik, Croatia. Association for Computational Linguistics.
- Michael Murray and Brian Graham. 1997. [Exploring the Dialectics of Route-based Tourism: The Camino de Santiago](#). *Tourism Management*, 18(8):513–524.
- Shinichi Nakagawa. 2004. [A Farewell to Bonferroni: The Problems of Low Statistical Power and Publication Bias](#). *Behavioral ecology*, 15(6):1044–1045.
- Hang T. Nguyen and Malika Chaudhuri. 2019. [Making new products go viral and succeed](#). *International Journal of Research in Marketing*, 36(1):39–62.
- Qiang Ning, Zhili Feng, Hao Wu, and Dan Roth. 2018. [Joint Reasoning for Temporal and Causal Relations](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2278–2288, Melbourne, Australia. Association for Computational Linguistics.
- Yangtuo Peng and Hui Jiang. 2016. [Leverage Financial News to Predict Stock Price Movements Using Word Embeddings and Deep Neural Networks](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 374–379, San Diego, California. Association for Computational Linguistics.
- Francesco Periti and Stefano Montanelli. 2024. [Lexical Semantic Change through Large Language Models: a Survey](#). *ACM Comput. Surv.*, 56(11).
- Thomas V Perneger. 1998. [What’s Wrong with Bonferroni Adjustments](#). *BMJ*, 316(7139):1236–1238.
- Andrew Piper, Richard Jean So, and David Bamman. 2021. [Narrative Theory for Computational Narrative Understanding](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 298–311, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. [Improving Language Understanding by Generative Pre-Training](#). Technical report, OpenAI. OpenAI Technical Report.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language Models are Unsupervised Multitask Learners](#). *OpenAI Blog*, 1(8):9. OpenAI Technical Report.
- Eghbal Rahimikia, Felix Drinkall, et al. 2024. [Re \(Visiting\) Large Language Models in Finance](#). Available at SSRN.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence Embeddings using Siamese BERT-Networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.

- Margaret E Roberts, Brandon M Stewart, and Richard A Nielsen. 2020. [Adjusting for Confounding with Text Matching](#). *American Journal of Political Science*, 64(4):887–903.
- Julia Rodina, Yuliya Trofimova, Andrey Kutuzov, and Ekaterina Artemova. 2021. [ELMo and BERT in Semantic Change Detection for Russian](#). In *Analysis of Images, Social Networks and Texts*, pages 175–186, Cham. Springer International Publishing.
- Ramit Sawhney, Shivam Agarwal, Arnav Wadhwa, and Rajiv Ratn Shah. 2020. [Deep Attentive Learning for Stock Movement Prediction From Social Media Text and Company Correlations](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8415–8426, Online. Association for Computational Linguistics.
- Valentina Semenova, Dragos Gorduza, William Wildi, Xiaowen Dong, and Stefan Zohren. 2024. [Wisdom of the Crowds or Ignorance of the Masses? A Data-Driven Guide to WallStreetBets](#). *Journal of Portfolio Management*, 50(4).
- Robert J Shiller. 2017. [Narrative Economics](#). *American Economic Review*, 107(4):967–1004.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2020. [MPNet: Masked and Permuted Pre-training for Language Understanding](#). In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS '20*, Red Hook, NY, USA. Curran Associates Inc.
- Maximilian Splethöver, Maximilian Keiff, and Henning Wachsmuth. 2022. [No Word Embedding Model Is Perfect: Evaluating the Representation Accuracy for Social Bias in the Media](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Rich Sutton. 2019. [The Bitter Lesson](#). Acc.: 03/02/25.
- Timothy R. Tangherlini, Shadi Shahsavari, Behnam Shahbazi, Ehsan Ebrahimpzadeh, and Vwani Roychowdhury. 2020. [An Automated Pipeline for the Discovery of Conspiracy and Conspiracy Theory Narrative Frameworks: Bridgegate, Pizzagate, and Storytelling on the Web](#). *PLoS ONE*, 15(6):e0233879. Published: June 16, 2020.
- Tan Thongtan and Tanasanee Phienthrakul. 2019. [Sentiment Classification Using Document Embeddings Trained with Cosine Similarity](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics: Student Research Workshop*, pages 407–414, Florence, Italy. Association for Computational Linguistics.
- Tyler J. VanderWeele. 2019. [Principles of Confounder Selection](#). *European Journal of Epidemiology*, 34:211–219.
- Victor Veitch, Dhanya Sridhar, and David M. Blei. 2019. [Adapting Text Embeddings for Causal Inference](#). In *Conference on Uncertainty in Artificial Intelligence*.
- Benyou Wang, Emanuele Di Buccio, and Massimo Melucci. 2020. [University of Padova @ DIACR-Ita](#). In *Proceedings of the Evaluation Campaign of Natural Language Processing and Speech Tools for Italian (EVALITA 2020)*, Marrakech, Morocco. CEUR-WS. CEUR Workshop Proceedings.
- Xinlei Wang, Maike Feng, Jing Qiu, Jinjin Gu, and Junhua Zhao. 2024. [From News to Forecast: Integrating Event Analysis in LLM-Based Time Series Forecasting with Reflection](#). In *Neural Information Processing Systems*.
- Hiroaki Yamagiwa, Momose Oyama, and Hidetoshi Shimodaira. 2025. [Revisiting Cosine Similarity via Normalized ICA-transformed Embeddings](#). In *Proceedings of the 31st International Conference on Computational Linguistics*, pages 7423–7452, Abu Dhabi, UAE. Association for Computational Linguistics.
- Jie Yang, Soyeon Caren Han, and Josiah Poon. 2021. [A Survey on Extraction of Causal Relations from Natural Language Text](#). *Knowledge and Information Systems*, 64:1161 – 1186.
- Hao Ye, Ethan R Deyle, Luis J Gilarranz, and George Sugihara. 2015. [Distinguishing Time-delayed Causal Interactions Using Convergent Cross Mapping](#). *Scientific Reports*, 5:14750.
- Jiaming Zeng, Michael Francis Gensheimer, Daniel L. Rubin, Susan Athey, and Ross D. Shachter. 2021. [Uncovering Interpretable Potential Confounders in Electronic Medical Records](#). *Nature Communications*, 13.
- Shuang (Sophie) Zhai and Zhu (Drew) Zhang. 2023. [Read the News, Not the Books: Forecasting Firms’ Long-term Financial Performance via Deep Text Mining](#). *ACM Trans. Manage. Inf. Syst.*, 14(1).
- Justine Zhang, Sendhil Mullainathan, and Cristian Danescu-Niculescu-Mizil. 2020. [Quantifying the Causal Effects of Conversational Tendencies](#). *Proceedings of the ACM on Human-Computer Interaction*, 4(CSCW2):1–24.

A Linear vs. Non-linear

Tab. 4 displays a more granular comparison between the linear and non-linear models outlined in §4.2.

B News Outlets

We use the articles from the following outlets:

- **Center:**
 - Heavy, CNBC, Business Insider, Reuters, Wall Street Journal, USA TODAY, BBC News, Forbes, FiveThirtyEight




Figure 4: The proportion of significant pairs when $\{y = \text{text}, x = \text{econ}\}$ and $\{y = \text{econ}, x = \text{text}\}$ for all news outlets. Comparing the **Non-linear** and **Linear**.

• **Left:**

- The New Yorker, Newsweek, CNN, The New York Times, CBS News, Vox, Rolling Stone, The Washington Post, The Guardian, Slate, Daily Beast, Vice, Politico, HuffPost, MSNBC, ABC News

• **Right:**

- The Gateway Pundit, PJ Media, The Western Journal, Breitbart News, Washington Times, ZeroHedge, New York Post, Reason, Fox News, Newsmax, Red State, The Epoch Times, Townhall, Orange County Register, The Daily Caller, National Review, Bizpac Review, TheBlaze

The political categories come from Spliethöver et al. (2022), which used AllSides.com tags.

C SVAR Framework

To generate the financial market shock time series, (Cieslak and Pang, 2021) fit an unrestricted vector autoregression on the daily changes in bond yields (2-year, 5-year, and 10-year maturities) and daily

log stock returns. Let $\mathbf{y}_t \in \mathbb{R}^4$ denote the vector of observed variables on day t :

$$\mathbf{y}_t = \begin{pmatrix} \Delta y_{2\text{yr},t} \\ \Delta y_{5\text{yr},t} \\ \Delta y_{10\text{yr},t} \\ \Delta \ln(S_t) \end{pmatrix}$$

where $\Delta y_{m\text{yr},t}$ is the daily yield change at maturity $m \in \{2, 5, 10\}$ and $\Delta \ln(S_t)$ is the daily log return of the stock index S_t . The structural VAR of lag order L can be written in the following form:

$$\mathbf{y}_t = \boldsymbol{\mu} + \sum_{j=1}^L A_j \mathbf{y}_{t-j} + \mathbf{u}_t$$

where $\boldsymbol{\mu} \in \mathbb{R}^4$ is a constant term, each $A_j \in \mathbb{R}^{4 \times 4}$ is a coefficient matrix for lag j , and $\mathbf{u}_t \in \mathbb{R}^4$ is the reduced-form error term with covariance matrix $\Sigma = \mathbb{E}[\mathbf{u}_t \mathbf{u}_t^\top]$.

Identifying Structural Shocks. Following standard practice for structural identification, the covariance matrix Σ is factorised, with a lower-triangular Cholesky decomposition:

$$\Sigma = PP^\top$$

where $P \in \mathbb{R}^{4 \times 4}$ is invertible and maps the structural shocks $\boldsymbol{\varepsilon}_t \in \mathbb{R}^4$ to the reduced-form errors:

$$\mathbf{u}_t = P \boldsymbol{\varepsilon}_t, \quad \boldsymbol{\varepsilon}_t \sim \mathcal{N}(\mathbf{0}, I_4)$$

Sign restrictions (Fry and Pagan, 2011) are then imposed on the impulse responses, which amounts to applying an orthonormal rotation Q to P , yielding

$$\mathbf{u}_t = P Q \tilde{\boldsymbol{\varepsilon}}_t$$

where $Q \in \mathbb{R}^{4 \times 4}$ satisfies $QQ^\top = I_4$ and $\tilde{\boldsymbol{\varepsilon}}_t$ are the rotated structural shocks. Among all permissible rotations, the impulse whose response satisfied the sign restrictions for each of the four shocks (*growth*, *monetary*, *common premium*, and *hedging premium*) is selected.

D Confounders Analysis

To understand the nature of the significant pairs that were identified using the Bonferroni correction, we tested whether the causal pairs had any confounding variables that caused both the treatment and

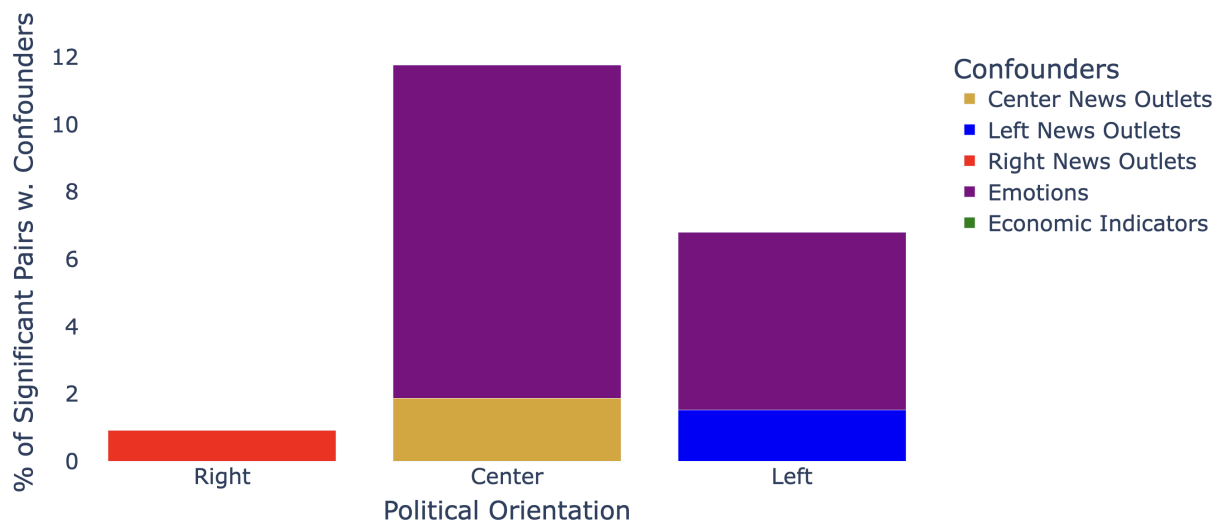


Figure 5: Proportion of significant pairs where $\{y = \text{text}, x = \text{econ}\}$ that have a confounding variable by orientation. The category of the confounding variables are denoted by the colour of the bar section.

outcome variables. The confounders were identified by whether a single variable, of any lag between 1 and 10 days, significantly improved the performance of both the treatment (x) and outcome (y) variables using a Bonferroni correction. The following variables were tested: the cosine distance of all of the other news outlets, the other structural shocks, VIX volatility index (which is grouped with the Economic Indicators in Fig. 5), and an emotion score for the economic news. The emotion score was calculated by first applying a [pre-trained emotion classification model](#) to each news article that related to the economy, which assigned probability scores to a set of predefined emotion labels. For each day d , the emotion score for label ℓ was computed as the average predicted probability across all articles published on that day:

$$E_{\ell,d} = \frac{1}{N_d} \sum_{i=1}^{N_d} p_i(\ell),$$

where $p_i(\ell)$ is the predicted probability of emotion ℓ for article i , and N_d is the total number of articles published on day d . This aggregation produced a time series of daily average emotion scores, which were then tested for their role as potential confounders in the causal relationships between the treatment and outcome variables. The potential confounding variables had to have a reporting frequency of 1 day, which limited the total number of variables that were trialled. The reporting frequency is generally very low in economics (CPI, GDP, employment, etc.). The decision was to either include these broader economic features and

accept the sparsity (which will likely lead to type 2 errors), aggregate the daily indicators to match their reporting frequency (leading to a very small dataset) or to exclude them and find a proxy that accounts for them in some way. We went for the last option by including other market indicators and features from the news, but we accept that it is not a perfect solution to a challenging problem. We believe that we have captured some of the major confounders that were feasible to include.

Fig. 5 shows the resulting percentage of total significant pairs that had a confounder. Firstly, there are no examples where there is a confounding news outlet from another partisan group, which shows that each partisan group has limited influence on outlets with different political outlooks. Secondly, there were no economic indicators that were confounding variables largely due to the fact that none of the economic indicators were deemed as causal for the other economic indicators. This could be explained by the Efficient Market Hypothesis, which states asset prices can be explained by all available information, meaning that the other economic variables do not add any predictive information. What is interesting is that the news emotion score is often a confounding variable, causing changes in both economic indicators and the news outlets. This suggests that not everything is priced in and that market shocks and news outlets can react to overall emotions within the news. However, right-wing outlets do not react to news emotion scores, showing a distinct difference between partisan groups.

E Licenses

The news data from [Spliethöver et al. \(2022\)](#) was not published with a license, but the original source of the data - [Commoncrawl News](#) - has a permissive Apache 2.0 license. The market shock data from [Cieslak and Pang \(2021\)](#) was shared by the authors.

F Model Setup

To simplify the experimental setup and limit the number of observations that could distort the statistical significance of our findings we used the default parameters for both the Linear and KRR models, except we used an *RBF* kernel instead of a *linear* kernel for the KRR regressor. If we had done hyperparameter tuning, then the statistical significance corrections would have been more strict. We used the [sklearn](#) package for both regressors, using the [LinearRegression](#) and [KernelRidge](#) modules.

G Dataset Statistics

Fig. 6 shows the 30-day rolling average of article counts after the filtering outlined in §3. The spike in 2022 can be attributed to the fact that [Spliethöver et al. \(2022\)](#) used a web scrape from 2022 to create the dataset. The data came from a [Commoncrawl News](#) WARC file, which contains the news available on the day of the web scrape. While the publish date is being displayed on Fig. 6, older articles might not be available. The absolute number of articles is not overly important so long as there is sufficient coverage. Cosine distance means that there is a value for every day - including a value of zero on days where no articles were published.

The outlet counts are displayed in Tab. 2. This represents all of the articles from 01/07 to 02/22. Some of the outlets should evidently have more articles, however companies like [The New York Times](#) have recently attempted to combat web scraping by restricting access to bots. To make sure that we

have sufficient information contained in the text-based time-series, we only consider outlets which have published articles on 25% of the days, including the rolling window framework.




Figure 6: 30-day rolling average of articles by day.

Outlet	Article Count
Vox	382,666
Newsweek	246,740
HuffPost	192,210
CNBC	185,890
Daily Beast	131,291
Reuters	130,678
Breitbart News	129,580
Slate	125,184
Wall Street Journal	123,964
Washington Times	113,836
Newsmax	108,015
New York Post	106,676
Vice	92,927
The Washington Post	90,056
The New Yorker	89,003
CNN	76,320
PJ Media	65,282
Rolling Stone	52,896
The Guardian	44,671
USA TODAY	38,355
ABC News	37,414
MSNBC	35,942
The Epoch Times	30,693
Red State	28,435
Townhall	26,345
Forbes	25,488
CBS News	23,202
Orange County Register	22,529
BBC News	21,683
National Review	20,297
ZeroHedge	18,497
Heavy	16,618
Bizpac Review	16,492
Reason	16,407
FiveThirtyEight	16,036
The Daily Caller	15,768
TheBlaze	10,778
Fox News	10,133
Business Insider	9,911
The Western Journal	9,666
The Gateway Pundit	7,240
Politico	5,468
The New York Times	5,021

Table 2: Article counts per news outlet