Journal of International Money and Finance 167 (2026) 103627

Contents lists available at ScienceDirect

Journal of International Money and Finance

journal homepage: www.elsevier.com/locate/jimf

Exchange rate narratives ☆

Vito Cormun[a,*], Kim Ristolainen[b]

[a] Santa Clara University, Santa Clara, United States
[b] University of Turku, Turku, Finland

HIGHLIGHTS

- We combine WSJ news, topic modeling, and generative AI.
- Six FX narratives since the 1970s track dollar movements.
- Narrative prominence proxies time-varying investor weights on FX drivers.
- Accounting for narrative weights strengthens FX-macro links.

ARTICLE INFO

JEL classification:
C3
C5
F3

Keywords:
Exchange rates
Big data
Textual analysis
Macroeconomic news
Wall street journal
Narrative retrieval
Scapegoat

ABSTRACT

We combine Wall Street Journal news, topic modeling, and generative AI to extract economic narratives associated with U.S. dollar fluctuations. Using data since the late 1970s, we isolate six narratives spanning fiscal and monetary policy, financial markets, geopolitical tensions, and technological change. Adding these narratives to standard exchange rate regressions substantially improves the explanatory power of macroeconomic aggregates. The narratives are consistent with shifts in investor attention across exchange rate drivers—a mechanism that generates time-varying loadings on macroeconomic aggregates and helps account for the exchange rate disconnect puzzle.

1. Introduction

In his 2017 Presidential Address to the American Economic Association, Robert Shiller argued: "While we may never be able to explain why some narratives 'go viral' and significantly influence thinking while other narratives do not, we would be wise to add some analysis of what people are talking about if we are to search for the source of economic fluctuations."

Building on this insight, this paper seeks to identify "what people are talking about" in relation to changes in the U.S. dollar exchange rate.

☆ We are grateful to Ryan Chahrour, Sanjiv Das, Pierre De Leo, Valeria Ferraro, Kris Mitchener, Gernot Müller, and Rosen Valchev for their insightful comments. We also thank seminar participants at the University of Southern California, the 2024 Annual Meeting of the Society for Economic Dynamics, the SUERF, ECB, Bank of Finland and Bank of Italy Workshop on the use of AI in economics, the annual conference of the Society for Economic Measurement, Santa Clara University, and the Kiel-CEPR Junior Workshop on Publishing in Top Macro & Finance Journals. Vito Cormun thanks the Leavey School of Business for their financial support, and Kim Ristolainen thanks the Emil Aaltonen Foundation for their financial support.

* Corresponding author.
Email addresses: vcormun@scu.edu (V. Cormun), kkrist@utu.fi (K. Ristolainen).

https://doi.org/10.1016/j.jimonfin.2026.103627

Available online 17 June 2026
0261-5606/

We focus on exchange rates for three main reasons. First, exchange rates exhibit a weak correlation with virtually any macroeconomic aggregate (Meese and Rogoff, 1983; Obstfeld and Rogoff, 2000). Second, as asset prices, exchange rates embed expectations of macroeconomic conditions (Engel and West, 2005). Media narratives often reflect these expectations by summarizing realized developments and discussing anticipated future outcomes. Third, newspaper articles often report financial market developments which, from a model standpoint, are a compelling source of exchange rate fluctuations (Itskhoki and Mukhin, 2021).

We propose a new method to extract and interpret the economic narratives—defined as the stories people tell to explain macroeconomic phenomena—embedded in news media and related to exchange rate movements. Our analysis begins with a large corpus of Wall Street Journal (WSJ) headlines. We first process this corpus using a topic model—specifically, supervised Latent Dirichlet Allocation (sLDA)—to reduce textual dimensionality and identify topics that are predictive of the monthly real exchange rate. We then use a large language model to summarize and contextualize these topics into economic narratives, which are represented as directed acyclic graphs (DAGs).[1]

Narratives improve our understanding of exchange rate fluctuations. We find no single narrative that can account for exchange rate changes over our sample period. Instead, we isolate six largely non-overlapping narratives, each related to a specific time period and economic theme. Individually, each narrative appears associated with distinct macroeconomic and financial variables in a manner consistent with established theoretical drivers of exchange rate fluctuations. Together, they offer a new anatomy of exchange rate fluctuations.

Narrative coverage helps reconcile the apparent disconnect between exchange rates and macroeconomic aggregates. We use narrative attention as a proxy for scapegoat effects (Bacchetta and van Wincoop, 2004, 2013) and show that conditioning on narratives substantially increases the explanatory power of standard exchange rate regressions. These findings are consistent with a mechanism in which investor attention shifts across exchange rate drivers over time.

The paper proceeds in four parts. In the first part, we describe the data and modeling approach. Our corpus consists of headlines from WSJ articles published between January 1976 and August 2019. The Wall Street Journal has become a central source in recent work on news and economic analysis. For instance, Bybee et al. (2024) applies an unsupervised LDA model to WSJ text to recover a set of latent topics, each defined as a probability distribution over co-occurring terms. Like principal component analysis in time series econometrics, topic models reduce dimensionality—but instead of decomposing variation in numerical data, they uncover latent semantic structure summarized into topics. From these, one can compute topic attention—defined as the share of words assigned to a topic at a given time—and test whether it correlates with macroeconomic outcomes. Bybee et al. (2024), for example, finds that attention to a "recession risk" topic strongly correlates with several macroeconomic aggregates.

Our analysis builds on this approach, but we focus specifically on exchange rates and use *supervised* Latent Dirichlet Allocation (sLDA) (Blei and McAuliffe, 2007). Unlike unsupervised LDA, sLDA estimates topics to maximize predictive power with respect to an outcome variable—in our case, the real exchange rate of the dollar against the currency basket of the other G7 countries. Conceptually, this parallels econometric approaches that identify latent factors by maximizing their explanatory power for a target variable, while additionally recovering the associated textual structure.

In the second part, we present the topic model results. We estimate 180 topics, balancing semantic coherence and interpretability, and show that topic attention captures substantial variation in exchange rate fluctuations. Because interpreting such a large number of topics is impractical, we group them into six broader "metatopics" using hierarchical clustering. These metatopics exhibit distinct time-series patterns, coherent economic themes, and a hump-shaped evolution—consistent with the epidemiological features of narratives emphasized by Shiller (2017).

To interpret the economic content of the metatopics, we prompt a large language model to summarize the representative headlines associated with each metatopic. The resulting summaries are expressed as directed acyclic graphs (DAGs), describing the main causal pathways embedded in the headlines. While these graphs provide an interpretable characterization of the economic narratives covered by the media, they should *not* be taken as evidence of causal relationships estimated econometrically. Rather, they likely reflect preferences for simplified explanations whereby correlation is framed as causation (Spiegler, 2016; Eliaz and Spiegler, 2020).

In the third part, we detail the evolution of these narratives over time. In the late 1970s and 1980s, exchange rate dynamics are associated with fiscal and monetary policy—captured by metatopics labeled "Historical Economic Policies and Energy Markets" and "Global Finance and Monetary Policies," including {Carter, Reagan, oil, debt, interest rates} among the most prominent terms, and featuring DAGs where exchange rates are driven by U.S. policies.

In the 1990s, narrative focus shifts mainly toward technological innovation. The metatopic "Technological Advancements and Financial Reporting" captures news about rising productivity, is represented by the DAG: *technological innovation → higher productivity → dollar appreciation*, and strongly correlates with changes in U.S. total factor productivity.

In the 2000s, attention turns to equity markets, with "Corporate Transactions and Stock Market" and "Economic Trends and Corporate Strategy" metatopics reflecting dot-com and post-dot-com dynamics. We find that their attention series correlate with movements in stock prices and the VIX.

Finally, in the most recent decade, we identify a dominant narrative around trade tensions, captured by "International Business and Financial News" (with terms such as {business, Trump, banking, exchange}), and a DAG including global uncertainty and flight-to-safety as drivers of dollar fluctuations.

In the last part of the paper, we discuss the implications of our results for models of exchange rate determination. In doing so, we view our findings through the lens of the scapegoat theory of exchange rate fluctuations developed by

[1] Directed acyclic graphs are network representations describing causal pathways among variables (see Eliaz and Spiegler, 2020).

Bacchetta and van Wincoop (2004, 2013). The theory posits that when some exchange rate drivers are not observable, markets assign excessive weight to observable indicators that happen to move with the exchange rate. These “scapegoats” evolve over time, resulting in time-varying relationships between exchange rates and macro aggregates.

We perform two additional exercises that provide supporting evidence for the scapegoat theory. First, we use narrative attention as a proxy for scapegoat effects, and compare exchange rate regressions that include narratives interacted with macroeconomic aggregates, with regressions that include macro aggregates alone. Including narrative attention substantially increases explanatory power: the adjusted R^2 rises from 1% to 7% for monthly exchange rate changes, and from 7% to 28% for annual changes. Second, we test whether the attention to a narrative evolves as predicted by the theory. By regressing narrative attention on the interaction between exchange rate changes and macro aggregates, we find that a narrative receives more coverage precisely when the joint realization of macro variables and exchange rates fits the story embedded in that narrative.

Overall, our results support the scapegoat mechanism and suggest that, once investor attention is accounted for, exchange rates may be more closely linked to macroeconomic aggregates than previously thought. Yet whether the media merely reflects or actively shapes the transmission of macroeconomic shocks remains an open question.

Literature review. Our paper relates to three main strands of the literature: the use of textual data in economics, the role of narratives in macroeconomic fluctuations, and the literature on exchange rate determination.

The first strand concerns the growing use of text data in economic research. Surveys by Gentzkow et al. (2019) and Ash and Hansen (2023) document the rapid expansion of textual analysis methods across economics. Topic models in particular have only recently gained traction. Early applications include Hansen et al. (2018), who apply LDA to FOMC communications, and Larsen and Thorsrud (2019) and Thorsrud (2020), who use news-based topic models to improve macroeconomic forecasting. More recently, Bybee et al. (2024) applies LDA to Wall Street Journal articles to measure economic conditions and forecast activity. Complementary work by Agarwal et al. (2024) constructs country-specific measures of media narrative sentiment toward China and shows that cross-country dispersion in how the same information is framed causally affects cross-border portfolio flows.

Our approach builds on this literature but differs in two important respects. First, we employ supervised LDA, which estimates topics to maximize predictive power with respect to an outcome variable—in our case, the real exchange rate—rather than relying on unsupervised topic extraction. Second, we focus specifically on exchange rate fluctuations, where the link between macroeconomic information and asset prices is known to be particularly weak. Complementary work incorporates news data using researcher-curated measures rather than statistical topic models—for example, Chahrour et al. (2021) uses Factiva article tags to study sectoral news exposure, and Baker et al. (2016) constructs economic policy uncertainty indices based on keyword counts.

The second strand concerns the economic role of narratives. In *Narrative Economics*, Shiller argues that contagious economic narratives can shape macroeconomic outcomes, motivating efforts to measure narratives empirically. Several recent contributions build on the theoretical framework of Eliaz and Spiegler (2020, 2024), in which narratives are represented as simplified causal models that economic agents use to interpret complex environments. Empirical applications include experimental evidence on causal reasoning in narratives (Charles and Kendall, 2022), analyses of inflation narratives (Andre et al., 2023; Heddaya et al., 2025), and studies of competing narratives shaping sentiment on social media (Macaulay and Song, 2023).

We complement this literature by proposing an automated procedure to recover narrative representations from large text corpora. Specifically, we combine supervised topic modeling—which extracts economically relevant textual structure—with large language models that synthesize representative headlines into interpretable narrative summaries expressed as directed acyclic graphs. While related work has used topic models or LLMs separately, our contribution lies in integrating these tools to recover narratives directly linked to exchange rate fluctuations.

The third strand is the extensive literature on exchange rate determination. A recent synthesis is provided by Maggiori (2022). Our results relate most directly to two themes in this literature. First, they are consistent with the view that exchange rate fluctuations are linked to changes in macro aggregates, as emphasized in classic models (Dornbusch, 1976; Frankel, 1979) and in more recent work such as Schmitt-Grohé and Uribe (2022). In particular, we find a significant role for fiscal policy in explaining dollar movements during the late 1970s and early 1980s, consistent with studies including Ravn et al. (2012); Kim and Roubini (2008); Monacelli and Perotti (2010), and Jiang (2021).

Second, our findings relate to research emphasizing instability and nonlinearities in exchange rate relationships (Frankel and Rose, 1995; Rossi, 2013). They are also consistent with the scapegoat theory of exchange rate determination (Bacchetta and van Wincoop, 2004, 2013), which posits that when key exchange rate drivers are imperfectly observed, market participants place shifting emphasis on observable indicators that appear to explain exchange rate movements. Our narrative attention measures provide a new empirical proxy for such time-varying salience. Conceptually, this approach parallels survey-based studies such as Cheung and Chinn (2001) and Fratzscher et al. (2015), which document that traders’ perceived importance of macroeconomic variables changes over time. Whereas those studies rely on survey data, we extract comparable information directly from news text.

2. Topic model estimation

We begin by describing the topic model, the data cleaning procedure, and the model performance.

2.1. Supervised LDA

Latent Dirichlet Allocation (LDA), introduced by Blei et al. (2003), is a probabilistic topic model widely applied in fields such as political science and psychology. It represents each document as a mixture of latent topics and each topic as a distribution over words. Formally, each topic k is characterized by a probability vector $\beta_k \in \Delta^{V-1}$, where V denotes the size of the vocabulary and Δ^{V-1} is the $(V-1)$-dimensional probability simplex. Each document d is associated with a topic distribution $\theta_d \in \Delta^{K-1}$, where K is the number of topics and Δ^{K-1} denotes the $(K-1)$-dimensional simplex. The elements $\theta_{k,d}$, for $k = 1, \ldots, K$, represent the proportion of topic k in document d, satisfying $\sum_{k=1}^{K} \theta_{k,d} = 1$.

This representation maps each document from the original V-dimensional term space to a lower-dimensional K-dimensional topic space, thereby achieving dimensionality reduction.

A key feature of LDA is the use of Dirichlet priors for both the topic—word distributions β_k and the document—topic distributions θ_d. Specifically, $\beta_k \sim \mathrm{Dir}(\eta)$ and $\theta_d \sim \mathrm{Dir}(\alpha)$, where η and α are hyperparameters governing the concentration of the distributions. Higher concentration parameters lead to more diffuse topic or word distributions, whereas lower concentration parameters induce sparsity, with documents (or topics) dominated by a smaller number of topics (or words). The central inference problem in LDA consists of approximating the posterior distributions of β_k and θ_d given the observed corpus, the number of topics K, and the prior parameters α and η.

The Supervised LDA (sLDA) model, introduced by Blei and McAuliffe (2007), extends the standard LDA framework to incorporate an external response variable Y_d for each document. Unlike traditional LDA, which is unsupervised and focuses solely on discovering latent topics, sLDA is designed to predict an external variable based on the document's topic proportions. In sLDA, the document-topic proportions θ_d are linked to the response variable through a linear regression model, where the aim is to infer latent topics that not only describe the documents but also have predictive power for the external outcome.

This supervised extension of LDA improves predictive accuracy. Blei and McAuliffe (2007) show that the topics estimated using sLDA have significantly higher predictive power compared to those estimated using standard LDA. The reason for this improvement lies in the supervised nature of sLDA. By introducing the external variable Y_d, sLDA aligns the topic discovery process with the prediction task, ensuring that the topics it uncovers are not only latent structures in the text but also directly related to the response variable.

A practical advantage of sLDA over unsupervised LDA is that it can distinguish between news that carries opposite implications for the exchange rate. An unsupervised topic model groups words by co-occurrence alone, so, for example, coverage of monetary policy tightening and easing would load on the same topic, conflating episodes of dollar appreciation and depreciation. By forming topics jointly with the exchange rate outcome, sLDA is more likely to distinguish such episodes into distinct narratives. In this sense, supervision is not merely a device for improving fit: it is what allows the model to recover economically distinct narratives that an unsupervised decomposition would blur together.

Appendix A provides details on the estimation algorithm. In our analysis, the response variable is the monthly trade-weighted real exchange rate between the United States and the other G7 economies. To select the number of topics, we estimate models with the number of topics ranging from 20 to 220 and evaluate them using interpretability metrics. Specifically, Fig. 10 in Appendix A indicates that $K = 180$ provides a reasonable balance between semantic coherence—measuring the extent to which words within a topic tend to co-occur across documents—and exclusivity—which captures how distinctive the words are to a given topic.

2.2. The Wall Street Journal data set

We leverage a comprehensive text corpus comprising all headlines published in *The Wall Street Journal* from January 1976 to August 2019. The data comes from the ProQuest Historical Newspapers database, and we use the ProQuest Text and Data Mining (TDM) tool for data extraction.

To improve topic interpretability, we preprocess the text by removing stopwords as well as specific names, including countries, cities, organizations, and individuals, that could obscure broader thematic structure. The resulting corpus contains 707,984 headlines and a vocabulary of approximately 15,000 unique terms. We provide a detailed description of the data collection and preprocessing in Appendix B.

Prior to estimating the sLDA model, we convert the text corpus into a document-feature matrix (DFM) that records word counts for each headline. Rows correspond to individual headlines, columns to vocabulary terms, and entries represent word frequencies. Because we examine the real exchange rate at a monthly frequency, we aggregate headlines by month, combining all titles published within a given month into a single document. This aggregation avoids mechanically associating each individual headline with the same monthly exchange-rate realization and is consistent with the assumption that only a subset of news items in a given month is relevant for exchange-rate dynamics.

Before presenting the model results, we briefly motivate our focus on headlines rather than full article texts. Headlines are intended to summarize the central informational content of articles in a highly condensed form, which may help concentrate salient keywords relevant for topic identification. By contrast, full article texts often include background or stylistic material common across articles that could blur topic distinctions. Moreover, because our empirical design aggregates news at the monthly level, using full texts would generate very long composite documents containing multiple unrelated themes, which could weaken the time variation in estimated topic proportions. Since our objective is to capture the narratives to which readers are exposed rather than the precise timing of information arrival, headline framing is plausibly a relevant feature of the data, though we do not formally test this conjecture.

2.3. Topic performance and selection

After estimating the topic model, we evaluate its performance with Random Forest (Breiman, 2001). Random Forest is an ensemble learning technique that constructs a collection of B individual regression trees, denoted as T_b, to predict a dependent variable based

Table 1
Adjusted Out-Of-Bag R^2 from Random Forest models.

Dep. variable	Out-Of-Bag R^2				
	(i) All topics	(ii) FRED	(iii) FRED + topics	(iv) Selected topics	(v) Metatopics
q_t	0.88	0.79	0.90	0.88	0.85
$q_{t+12} - q_t$	0.61	0.58	0.65	0.58	0.53
$q_{t+24} - q_t$	0.80	0.72	0.82	0.78	0.71
$q_{t+36} - q_t$	0.80	0.80	0.86	0.80	0.73
$q_{t+60} - q_t$	0.87	0.78	0.89	0.87	0.84

Note: The table reports the adjusted Out-Of-Bag R^2 values from Random Forests predicting the natural logarithm (and its future change) of the trade-weighted monthly real exchange rate between the U.S. and the other G7 countries during the period from January 1976 to August 2019. The predictors are (i) the topic proportions, (ii) 120 time series from the dataset maintained by McCracken and Ng (2016), (iii) both topic proportions and time series indicators, (iv) the sixteen topics with the largest explanatory power, and (v) six metatopic attention series constructed based on the outcome of a cluster analysis on the topic word distributions.

on a set of explanatory variables X_t. In our analysis, we estimate the model as follows:

$$q_t = RF(X_t) = \frac{1}{B}\sum_{b=1}^{B} T_b(X_t)$$

$$q_{t+h} - q_t = RF(X_t) = \frac{1}{B}\sum_{b=1}^{B} T_b(X_t), \quad \text{for } h = 12, 24, 36, 60 \tag{1}$$

where q_t is the natural logarithm of the monthly real exchange rate between the U.S. and the other G7 countries, expressed as the price of the foreign consumption basket in terms of the U.S. basket, so that an increase indicates a real depreciation of the dollar. The explanatory variables are the estimated monthly topic proportions $\hat{\theta}_{k,t}$.[2]

The Random Forest model operates by generating multiple individual decision trees, each trained on a randomly selected subset of the predictor variables and bootstrapped samples from the original dataset. A key feature of this method is the use of out-of-bag (OOB) samples, which consist of data points not included in the training of a specific tree.[3]

Column (i) of Table 1 reports the OOB R^2 of the Random Forest models predicting the current level of the exchange rate and the change in the exchange rate at 12, 24, 36, and 60 month horizons. The OOB R^2 is the average R^2 from 5000 trees.[4] Topic proportions explain a substantial portion of the exchange rate fluctuations, with OOB R^2 values ranging from 61% to 88%. Fig. 11 in Appendix C illustrates predicted and actual exchange rate values for the model specifications in (1).

Overall, the topic model predicts exchange rate changes fairly well. However, both the table and the figure reveal that higher frequency changes are somewhat harder to predict. This may be because high-frequency changes are less newsworthy, whereas major episodes leading to large and long-lasting dollar movements are more likely to be reported (Nimark, 2014).

A natural concern is that the predictive performance reported in Table 1 is high by construction: sLDA forms topics jointly with the exchange rate, so high R^2 values simply confirm that the model worked as intended. We therefore emphasize that the predictive results serve a purely instrumental purpose—they guide topic selection and validate that the dimensionality reduction retains exchange-rate-relevant variation. Furthermore, none of our economic conclusions rest on this predictive relationship. The supervision target is the real exchange rate alone, whereas the narrative interpretations we draw are validated against variables the topic model never observes—including the unemployment rate, government debt, total factor productivity, and equity-market volatility (Section 3.4)—and against the theory-consistent-versus-placebo comparison in Section 4.2. Because the topics are not estimated to fit any of these series, their alignment with narrative attention cannot be an artifact of supervision.

2.4. Comparison with macroeconomic indicators

A natural question is whether topics encapsulate different information from that contained in standard time series. To address this question, we compare the predictive performance of the topic model with that of a large number of macroeconomic indicators. Specifically, we use the monthly dataset constructed by McCracken and Ng (2016), which includes 120 U.S. macroeconomic indicators excluding exchange rates. The variables are transformed to ensure stationarity, and we estimate two sets of Random Forest models: one using only macroeconomic indicators and another augmenting these indicators with the topic series.

[2] Because headlines are pooled by month prior to estimation, the monthly topic proportions $\hat{\theta}_{k,t}$ are obtained from the aggregated document containing all headlines in month t.

[3] Because both exchange rates and topic proportions are highly persistent, a natural concern is that predictive relationships may reflect shared low-frequency variation rather than economically meaningful structure. While the out-of-bag (OOB) procedure mitigates overfitting by evaluating predictions on observations not used during training, it does not isolate exogenous sources of variation and therefore does not, by itself, rule out predictability driven by common trends or slowly evolving latent factors. For this reason, we interpret predictive R^2 as descriptive evidence of predictive structure rather than as evidence of causal or structural relationships, and complement it with the narrative, macroeconomic, and scapegoat-based validation exercises that follow.

[4] We use the Ranger package in R to estimate the models, and the tuneRanger package for sequential model-based optimization to tune the number of variables available for splitting at each node, the minimum number of samples per node, and the fraction of observations used in each bootstrap sample.

Table 1 shows that topic-based models deliver predictive performance broadly comparable to that obtained using standard macroeconomic indicators, while augmenting the macro specification with topics leads to a noticeable improvement in fit. For example, the R^2 for the exchange rate increases from 79% to 90%.

In addition, Appendix Fig. 12 presents the results for the augmented model and ranks the variables based on the drop in R^2 computed from the predictions excluding one variable at a time. Out of thirty-five entries, macroeconomic indicators appear only in seven cases. Three of these cases include interest rate spreads, whereas the remaining four include new housing permits across U.S. regions.

Taken together, these results suggest that topics capture information related to macroeconomic conditions while also containing additional variation not fully summarized by standard macroeconomic indicators. Next, we turn to the interpretation of these topics.

3. Topic interpretation

In this section we interpret the outcome of the sLDA. Since examining 180 topics is impractical, we focus only on the subset of topics with the highest explanatory power. Furthermore, we cluster topics with similar word distributions. This procedure results in six clusters or metatopics, with minimal loss in the explanatory power of exchange rate fluctuations. We then examine the prevalence of each metatopic over time and find that they capture distinct macroeconomic narratives.

3.1. Topic selection and clustering

To reduce the number of topics examined, we select the top seven topics with the largest R^2 drop in each specification in (1). Doing so ensures that the resulting subset of sixteen topics leads to little change in the OOB R^2 values, as reported in column (iv) of Table 1.

Among the sixteen remaining topics, we exploit their word distributions and group those topics with similar news content. Specifically, we implement a hierarchical cluster analysis of the word distributions β_k, and report the outcome in the form of a dendrogram in Fig. 1. We describe the clustering procedure in Appendix E.

Fig. 1 also shows topic and metatopic labels, which we assign by prompting ChatGPT with the topic-word distribution. Starting from right to left, the dendrogram shows the topic labels and the aggregation nodes. We focus on the node with six clusters of topics, or metatopics, as it strikes a good balance between the number and the semantic distinction across clusters. The resulting metatopics include a range of one to four topics, depending on the similarities between word distributions.

Fig. 2 displays the word clouds for each metatopic, highlighting the most prominent terms. Together with the dendrogram, the word clouds underscore that the six metatopics are largely distinct.

Fig. 1. Dendrogram. *Note*: The figure illustrates the outcome of a cluster analysis on the word distributions of the topics. Topic and metatopic titles are constructed using ChatGPT-4. The prompt and the ChatGPT-generated metatopic descriptions are in Appendix G.

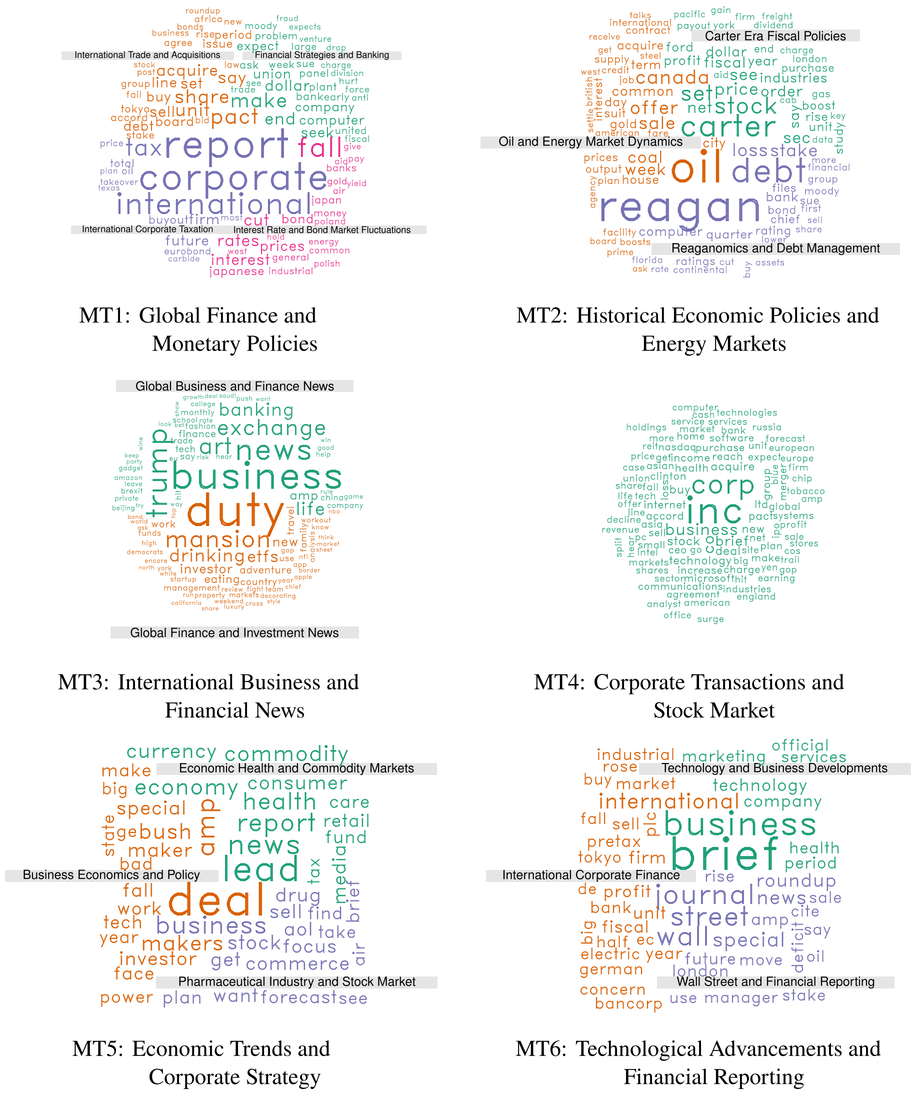

Fig. 2. Metatopic comparison word clouds. *Note*: The figure shows the comparison word clouds of the metatopic word distributions. In a comparison word cloud, the size of each word is mapped to the maximum deviation of its topic-word probability from the average topic-word probability across the topics in the same metatopic.

For example, MT1 (Global Finance and Monetary Policies) and MT2 (Historical Economic Policies and Energy Markets) are composed of topics about monetary and fiscal policy. In contrast, MT4 (Corporate Transactions and Stock Market) and MT5 (Economic Trends and Corporate Strategy) are more closely associated with stock market dynamics. MT6 (Technological Advancements and Financial Reporting) appears centered on news about technological changes.

MT3, labeled "International Business and Financial News", is less straightforward. Given its prominent terms – such as {business, Trump, banking, exchange} – and its composition of topics related to global economic developments, this metatopic likely reflects narratives tied to U.S. foreign policy, rather than domestic policy or technological innovation.

3.2. Metatopic prevalence over time

After clustering topics into metatopics, we study their time series properties. We compute metatopic attention series, that is, the share of words assigned to each metatopic, as the weighted sum of its estimated topic proportions. Formally, the attention series of metatopic j is

$$\hat{\Theta}_{j,t} = \sum_{k \in \mathcal{M}_j} \omega_{k,j} \hat{\theta}_{k,t} \tag{2}$$

where $\mathcal{M}_j$ denotes the set of topics that are grouped into metatopic j, and the weight $\omega_{k,j}$ is the topic k's average R^2 drops across the specifications in (1), normalized such that $\sum_{k \in \mathcal{M}_j} \omega_{k,j} = 1$. We use these weights so as to minimize the loss of explanatory power

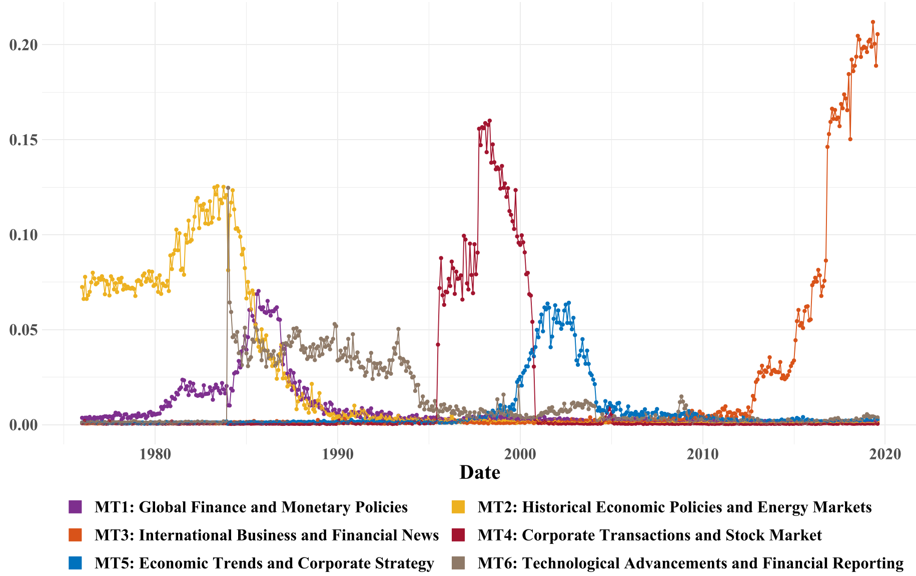

Fig. 3. Weighted metatopic attention series. *Note*: The figure plots the six metatopic attention shares, each constructed as the weighted sum of topic proportions as in Eq. (2).

resulting from aggregating topics into metatopics. Column (v) in Table 1 reports the OOB R^2 when we use metatopics as predictors. The values are only a few percentage points smaller than those from the model with the selected sixteen topics. For example, the OOB R^2 is 85% in the real exchange rate regression with the six metatopic attentions, relative to 88% in the regression with all 180 topics.

Fig. 3 shows the evolution of metatopic attentions over time. Three patterns emerge. First, metatopic attentions are hump-shaped, lasting about five to fifteen years and closely resembling the epidemiological properties of narratives highlighted in Shiller (2020). Second, attention peaks are different across metatopics, suggesting that metatopics represent distinct, individual events, rather than a continuous stream of news related to the same economic theme. Third, we do not estimate any attention peak during the years around the Great Recession. Given that, at each point in time, topic proportions must sum to one, the absence of a metatopic peak is consistent with a multitude of narratives related to exchange rate fluctuations around those years.

Metatopics and the exchange rate. Figs. 4 and 5 plot the Random Forest predictions of the demeaned real exchange rate based on the metatopic series, together with the estimated contributions of individual metatopics to the predicted values.[5] At any given time, only a limited number of metatopics contribute materially to the exchange rate prediction. For example, the dollar appreciation in the early 1980s coincides with increased prominence of MT1 and MT2, which broadly capture media coverage related to U.S. fiscal and monetary policy developments. By contrast, exchange rate fluctuations around the dot-com period appear more closely associated with MT4 and MT5, metatopics linked to stock market developments. One partial exception is MT6, which shows a more persistent association with exchange rate movements over the sample. This pattern is broadly consistent with recent evidence in Chahrour et al. (2024), suggesting that revisions in expectations about future productivity may play an important role in exchange rate dynamics. Finally, Fig. 5 indicates that these qualitative patterns remain broadly stable across the alternative differencing horizons we consider.

3.3. Exchange rate narratives

While topic models are effective at uncovering latent thematic structure in text, they are not designed to identify relationships among events or sequences of economic mechanisms, whether stated explicitly or implicitly. Such narrative structures can help interpret topic-based findings and may inform structural macroeconomic modeling. For example, Andre et al. (2023) manually constructs DAGs from open-ended survey responses on post-Covid inflation, and embeds them in a New Keynesian framework to study inflation dynamics.

To characterize the narratives associated with each metatopic, we first identify representative headlines. We rank headlines according to the share of words assigned to topic k, defined as

$$n_{d,k} = \sum_{n=1}^{N_d} \frac{\mathbb{I}(z_{d,n} = k)}{N_d},$$

[5] We compute metatopic contributions by tracing the path from the root to the leaf node in each decision tree. At each node, the metatopic responsible for the split contributes an amount equal to the change in the predicted value at that node. These contributions are summed along the path for each metatopic, and we average them across all trees in the Random Forest to obtain the overall contribution of each metatopic to the model's prediction. See Appendix F for further details.

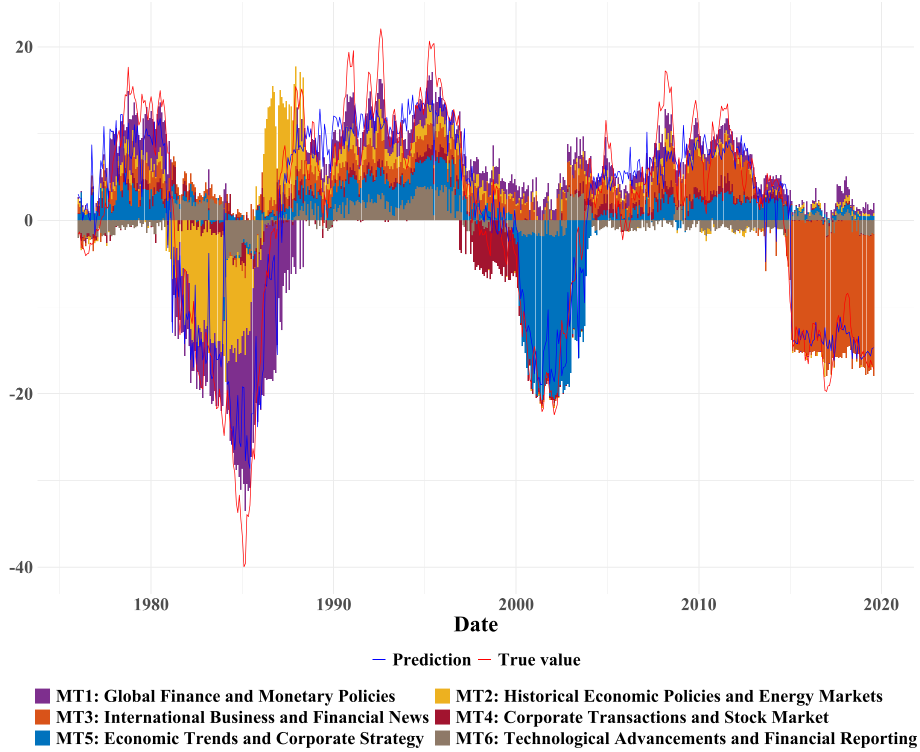

Fig. 4. Metatopic predictions and contributions to the real exchange rate. *Note*: The figure shows the log real exchange rate (red line), the Random Forest prediction based on six metatopic proportions (blue line), and the corresponding metatopic feature contributions to the prediction (see Appendix F for details).

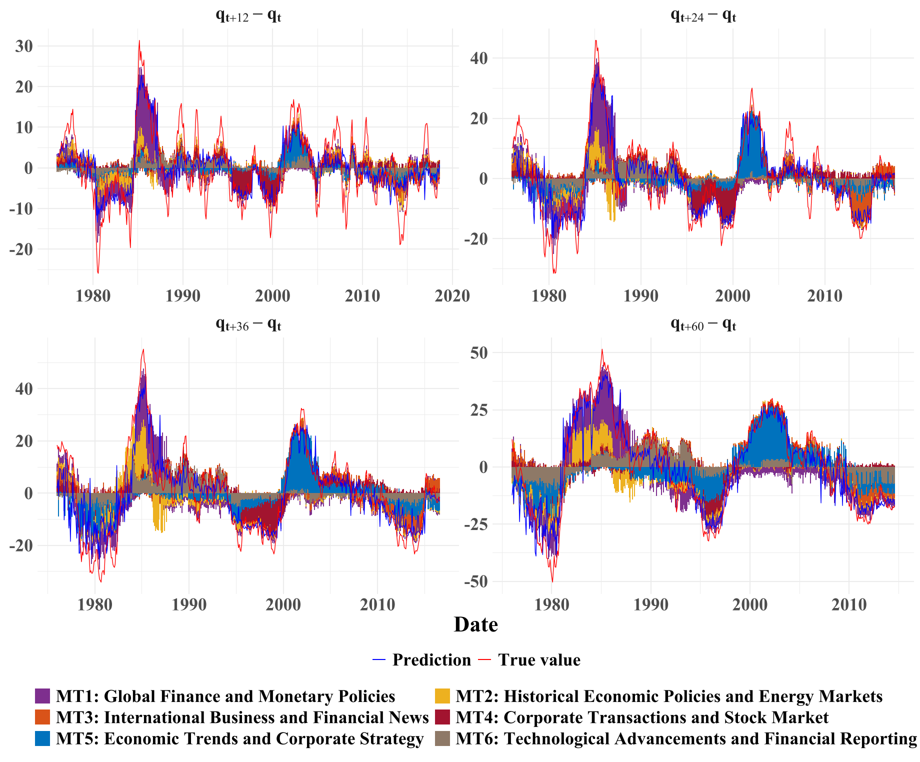

Fig. 5. Metatopic predictions and contributions to real exchange rate changes. *Note*: The figure shows changes in the log real exchange rate (red lines), Random Forest predictions based on six metatopic proportions (blue lines), and the corresponding metatopic feature contributions (see Appendix F for details).

Table 2
Narrative Descriptions and their corresponding DAGs.

Narrative Description	DAG Representation
MT1: Global Finance and Monetary Policies Tightening of U.S. monetary policy, reflected in higher interest rates, leads to capital inflows and drives dollar appreciation.	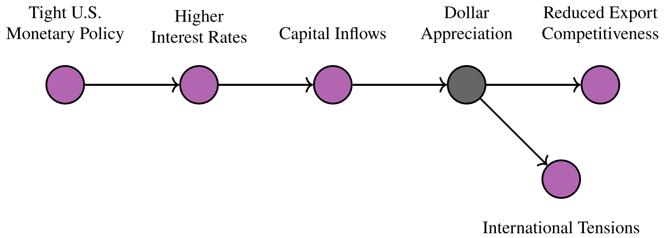
MT2: Historical Economic Policies and Energy Markets The first DAG, reflecting the Carter era, shows how rising inflation expectations eroded confidence, triggered capital outflows, and led to dollar depreciation. The second, aligned with Reagan's policies, illustrates how fiscal expansion raised interest rates, attracted capital inflows, and drove dollar appreciation.	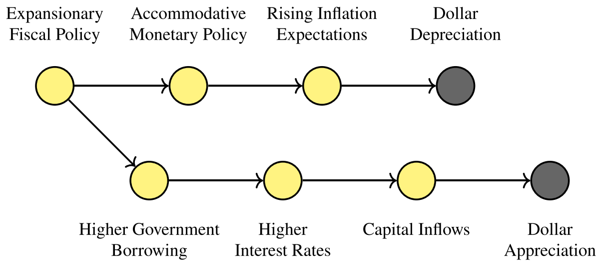
MT3: International Business and Financial News Trade disputes and geopolitical tensions create uncertainty, driving capital into U.S. safe-haven assets. This leads to dollar appreciation through shifts in risk sentiment.	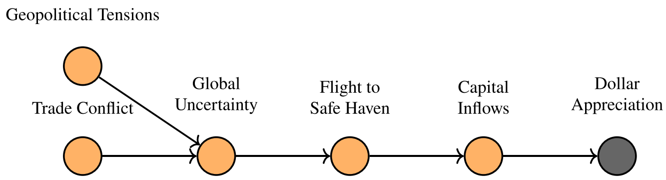
MT4: Corporate Transactions and Stock Market During the late 1990s, tech sector growth and equity booms attracted global capital, contributing to dollar appreciation.	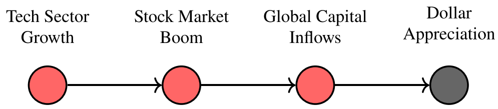
MT5: Economic Trends and Corporate Strategy In the early 2000s, global crises and shifts in risk appetite influenced capital flows and dollar strength. Demand for safe-haven assets increased.	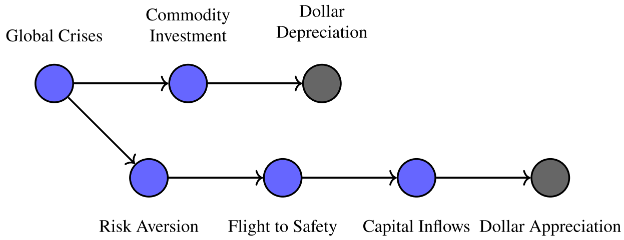
MT6: Technological Advancements and Financial Reporting Mid-1990s technological innovation boosted productivity and attracted global investment, supporting dollar appreciation.	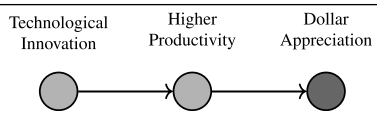

where N_d denotes the number of words in headline d and $z_{d,n}$ is the topic assignment for the n-th word. Higher values of $n_{d,k}$ indicate stronger topic representation in headline d. Aggregating across topics, we retain the 100 most representative headlines for each metatopic.

We then use a large language model (LLM) to generate narrative summaries and corresponding DAG representations based on these selected headlines. Specifically, we prompt ChatGPT-4 using the metatopic label together with the representative headlines as contextual input. We employ the Deep Research functionality within ChatGPT-4 to provide additional contextual grounding, allowing the model to reference abstracts and related bibliographic information when available. This feature is used to enhance narrative coherence and consistency, while the extracted narratives remain anchored in the selected headline set.[6]

[6] The LLM also reports the share of headlines broadly consistent with the generated narrative. These prevalence measures range from roughly 30% to over 70% across metatopics, suggesting moderate narrative concentration within headline clusters. Appendix H reports the prompt and examples of the headlines used in the narrative extraction.

Table 2 presents the resulting narrative summaries alongside their DAG representations. These DAGs provide a structured characterization of the economic narratives reflected in media coverage and complement the latent thematic structure identified by the topic model. They should be interpreted as descriptive summaries of narrative framing rather than as evidence of structural causal relationships.

In the same prompt, we also ask the LLM to indicate a few headlines that are most closely related to the narrative. For example, in the monetary-policy-related DAG of MT1, the LLM highlights the headline “Banks’ Foreign-Exchange Revenues Dive After Wrong Guessing on Dollar’s Strength” (July 20, 1984) as illustrative of tensions between market expectations and realized U.S. monetary policy that could plausibly affect foreign exchange markets. In this episode, tighter-than-expected monetary policy coincided with foreign-exchange losses and subsequent portfolio adjustments. Similarly, in the technology-related appreciation narrative of MT6, the headline “Roadway Services Inc.: Trucking Concern to Start Heavy Air Freight Business” (May 13, 1993) is interpreted as reflecting a shift toward technology-intensive logistics, potentially signaling productivity improvements that may attract capital inflows and support dollar appreciation.

The LLM also uncovers divergent narratives related to similar events. For example, in MT2, it highlights two different pathways originating from an expansionary fiscal policy, depending on the Federal Reserve’s reaction. Similarly, in the context of global crises, it distinguishes between alternative paths driven by expectations of growth versus risk aversion. It also helps clarify otherwise diffuse clusters such as MT3 (International Business and Financial News). By emphasizing headlines related to the U.S.–China trade tensions—e.g., “Beijing Calls U.S. Bully as Tariffs Kick In” (September 25, 2018) and “U.S.-China Trade Talks Hit a Bump” (March 9, 2019)—the extracted narratives suggest geopolitical developments as a recurring theme associated with recent dollar movements.

Notably, narratives often rely on headlines that do not mention exchange rates explicitly. Yet, they are plausibly linked to exchange rate expectations and investor behavior. In this sense, narratives need not be explicitly *about* exchange rates to be relevant for exchange-rate dynamics. Overall, the LLM-assisted summaries help organize economically interpretable narrative channels from sparse but information-rich headline data, while remaining descriptive rather than causal in nature.

3.4. Narratives and macroeconomic conditions

To assess whether the extracted narratives capture the economic developments they purport to describe, we examine how metatopic attention aligns with standard macroeconomic indicators.

We focus on six key monthly indicators: the Federal Funds rate, the future annual growth in real government debt, the past annual growth of the S&P 500, the annual change in the trade balance-to-GDP ratio, the unemployment rate, and the VIX. These variables capture a broad range of macroeconomic conditions frequently discussed in financial media and policy commentary and are plausibly related to the estimated narratives.[7]

We estimate the relationship between metatopic attention and these indicators using Random Forest models. For each target variable, we compute predicted series based on metatopic attention shares and evaluate both overall predictive performance and variable importance across metatopics.

Fig. 6 reports the predicted and actual series together with metatopic contributions, while Table 3 presents the corresponding out-of-bag R^2 statistics.

Metatopic attention accounts for a sizeable share of variation in several indicators, with notable heterogeneity across variables and over time. To further explore this heterogeneity, Fig. 7 reports Mean Squared Error losses obtained after excluding individual metatopics, computed separately for each non-overlapping decade.

The patterns broadly align with the narrative interpretations of the metatopics, and emphasize important heterogeneity. For instance, MT1 (Global Finance and Monetary Policies) and MT2 (Historical Economic Policies and Energy Markets) are most closely associated with movements in interest rates, unemployment, and government debt during the 1980s. In contrast, MT4 (Corporate Transactions and Stock Market) and MT5 (Economic Trends and Corporate Strategy) are more strongly associated with stock-market fluctuations around the dot-com cycle. MT3 (International Business and Financial News) shows stronger associations with later-sample measures of economic activity and financial volatility, suggesting links to broader macroeconomic and geopolitical developments.

Finally, word-level analysis suggests that MT6 (Technological Advancements and Financial Reporting) may capture narratives related to productivity developments. Because total factor productivity (TFP) data are available only at a quarterly frequency, we aggregate the monthly metatopic series accordingly, and estimate a Random Forest model for future annual log changes in the quarterly utilization-adjusted TFP from Fernald (2012). Fig. 8 shows that MT6 is consistently associated with variation in TFP, especially during the 1990s.

Taken together, these results indicate that narrative attention extracted from news headlines aligns with distinct macroeconomic developments in ways consistent with their economic interpretation. This supports the view that media narratives provide a useful complementary lens for tracking macroeconomic expectations, financial conditions, and structural economic developments over time.

4. A narrative-based explanation of the disconnect

Narratives provide insights into exchange rate fluctuations and the documented disconnect between exchange rates and macroeconomic aggregates. Taken individually, they are consistent with existing evidence linking exchange rates to specific macroeconomic

[7] We use the *future* annual growth in government debt to reflect the forward-looking nature of news reporting. The trade balance-to-GDP ratio is constructed using linearly interpolated quarterly GDP.

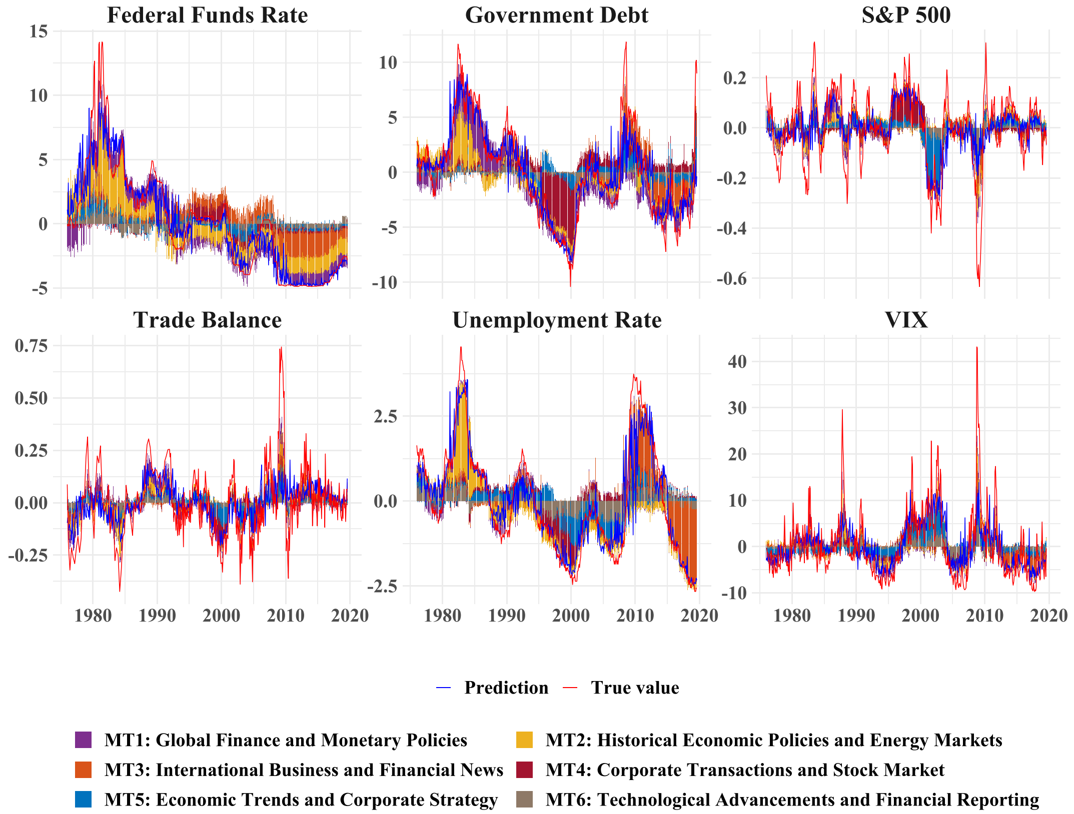

Fig. 6. Metatopic predictions of monthly indicators. *Note*: The figure shows the predicted (blue line) and actual (red line) values from Random Forest models using six metatopic proportions as predictors. The predicted series are the Federal Funds rate, the future annual log change in the real U.S. government debt, the past annual log change in the S&P 500, the past annual change of the U.S. trade balance to GDP ratio where the GDP is interpolated from the quarterly GDP series, the U.S. unemployment rate, and the VIX. All series are demeaned. For each metatopic, we compute its contribution by averaging the change in the predicted value at a tree node due to the metatopic, across all nodes.

Table 3
OOB R^2 from Random Forests using metatopic proportions as predictors.

Variable	Out-Of-Bag R^2
Fed Funds Rate	0.844
Government Debt	0.750
S&P 500	0.368
Trade Balance	0.299
Unemployment Rate	0.810
VIX	0.410
TFP (quarterly)	0.284

Note: The table reports the adjusted Out-Of-Bag R^2 values from Random Forests using six metatopic proportions as predictors. The predicted series are the Federal Funds Rate, the future annual log change in the real U.S. government debt, the past annual log change in the S&P 500, the past annual change of the U.S. trade balance to GDP ratio where the GDP is interpolated from the quarterly GDP series, the U.S. unemployment rate, the VIX, and the future annual change of quarterly utilization-adjusted TFP from Fernald (2012).

or financial indicators. Examples include the predictive role of the U.S. fiscal cycle (Jiang, 2021), the importance of monetary policy expectations (Dornbusch, 1976), the role of equity prices (Brandt et al., 2006), the association with financial market volatility (Cormun and De Leo, 2025), and the impact of productivity news (Chahrour et al., 2024).

At the same time, the reduced-form relationship between exchange rates and these aggregates often appears weak or unstable in the data. One interpretation, consistent with survey evidence in Cheung and Chinn (2001), is that investors attach time-varying importance to different sources of exchange rate fluctuations. Building on this insight, we propose a narrative-based perspective on the exchange rate disconnect: shifts in the prominence of different economic narratives may reflect changing investor attention

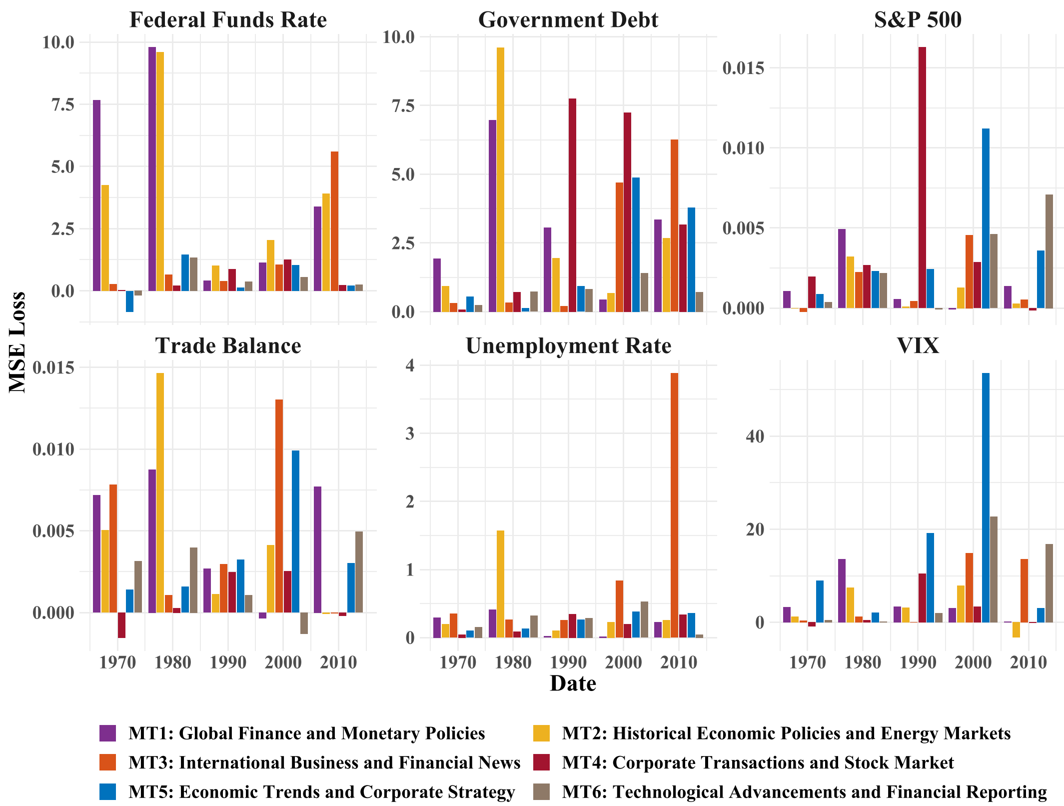

Fig. 7. Metatopic Mean Squared Error losses divided by decade. *Note*: The figure shows the Mean Squared Error losses from excluding a metatopic in the prediction of a Random Forest. The MSE losses are estimated separately for non-overlapping decades from January 1976 to August 2019. The predicted series are the Federal Funds rate, the future annual log change in the real U.S. government debt, the past annual log change in the S&P 500, the past annual change of the U.S. trade balance to GDP ratio where the GDP is interpolated from the quarterly GDP series, the U.S. unemployment rate, and the VIX.

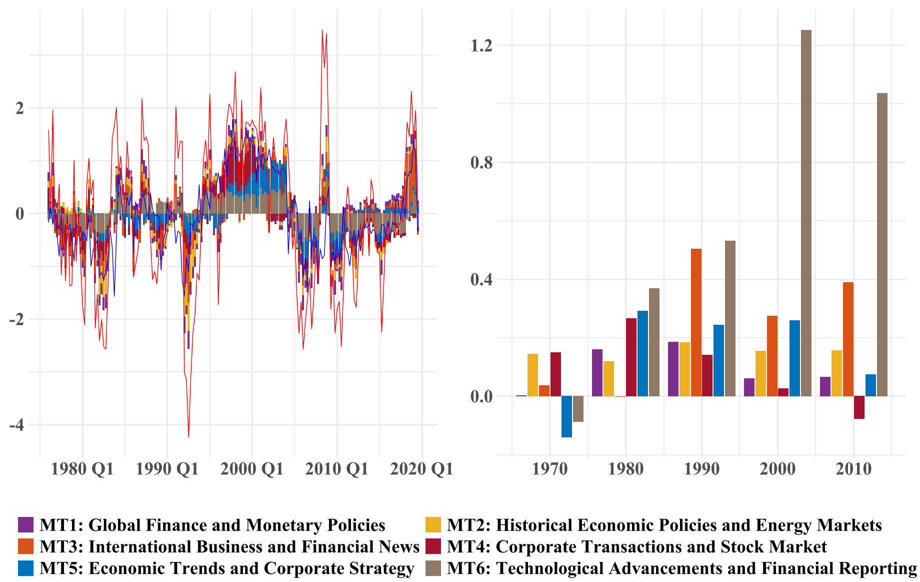

Fig. 8. Metatopic predictions, contributions and MSE losses of quarterly TFP. *Note*: The left panel of the figure shows the predicted value (blue line) and the future annual log change of quarterly utilization-adjusted TFP by Fernald (2012) (red line) from a Random Forest model using six metatopic proportions as predictors. The TFP series is demeaned and metatopic proportions are aggregated quarterly. For each metatopic, we compute its contribution by averaging the change in the predicted value at a tree node due to the metatopic, across all nodes. The right panel shows the Mean Squared Error losses of excluding a metatopic from the prediction. The MSE losses are computed over a rolling sample including each decade from January 1976 to August 2019.

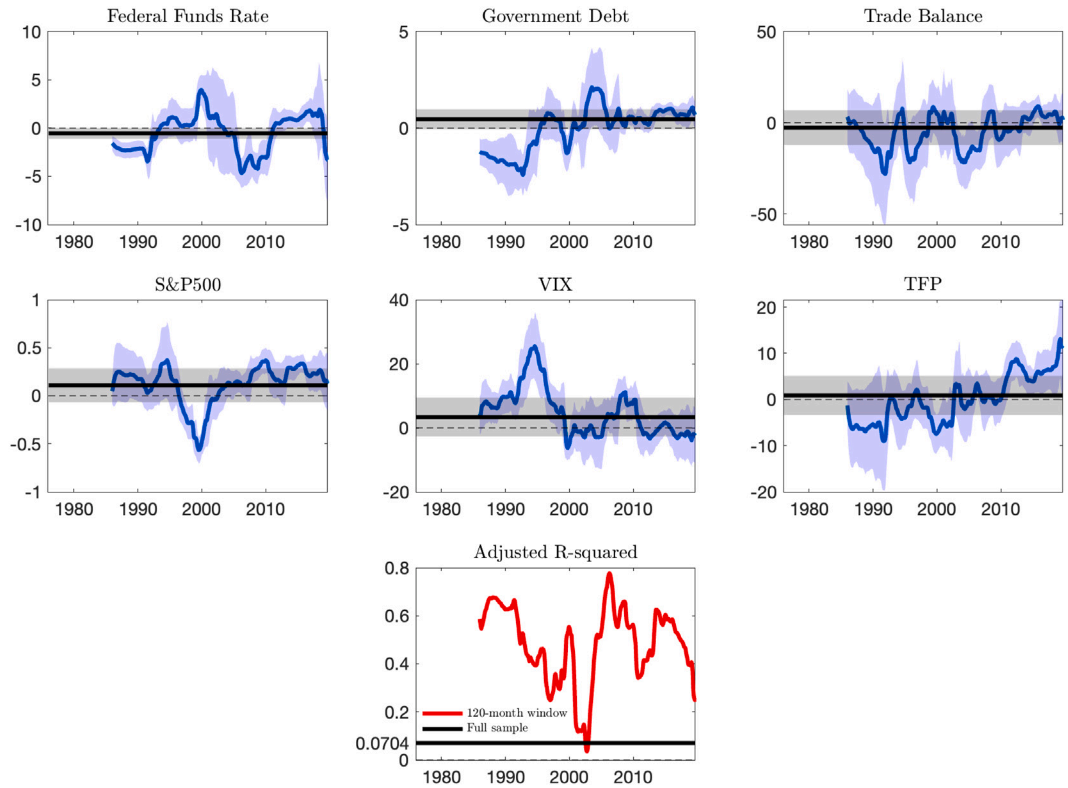

Fig. 9. Rolling window regression coefficients and adjusted R^2 values. *Note*: The figure reports the coefficients (blue lines) of a 120-month rolling window linear regression of the annual log change in the real exchange rate on the Federal Funds rate, the annual log change in the S&P 500 index, the annual change in the VIX and in the trade balance-to-GDP ratio, the future annual change in real government debt, and the future annual log change in total factor productivity (linearly interpolated from quarterly data). All series are standardized. The red line indicates the adjusted R^2 values. The black lines are the full sample estimates. The shaded areas are 95% confidence intervals.

across macro drivers, thereby tightening the observed link between exchange rates and macro aggregates over time. We show that this perspective naturally connects to scapegoat models of exchange rates (Bacchetta and van Wincoop, 2004).

4.1. *Evidence of parameter instability*

To begin, we observe that since metatopics explain dollar exchange rate movements and are connected to various macroeconomic indicators, we should observe a direct relationship between the exchange rate and these indicators, without the need for topic estimation. To test this, we estimate a linear regression of the annual log change in the real exchange rate on the Federal Funds rate, the annual log change in the S&P 500 index, the annual change in the VIX and in the trade balance-to-GDP ratio, the future annual change in real government debt, and the future annual log change in total factor productivity (linearly interpolated from quarterly data). Given the persistence of narratives presented earlier, we estimate rolling-window regressions using a 120-month window.

Fig. 9 displays the regression coefficients, their 95% confidence intervals, and the adjusted R^2 values, with the x-axis marking the end dates of the rolling windows. Comparison with full-sample estimates indicates substantial parameter instability. Notably, the adjusted R^2 values from the rolling window regressions are generally much higher than those from the full-sample estimates.

This parameter instability aligns with our exchange rate narrative decomposition. For instance, we find that the Federal Funds rate and changes in government debt are highly predictive of the exchange rate early in the sample, coinciding with peaks in attention to MT1 and MT2. Similarly, stock prices become highly predictive during the dot-com bubble, corresponding to variations in MT4 and MT5.

4.2. *Scapegoat theory*

Next, we show that the scapegoat theory of Bacchetta and van Wincoop (2004, 2013) can rationalize the unstable relationship between the exchange rate and the macroeconomic indicators documented above. The theory posits that, due to imperfect information regarding the parameters driving the exchange rate, investors occasionally attribute excessive weight to specific macroeconomic factors, effectively making them “scapegoats” for exchange rate movements. We find that when this occurs, media coverage of these factors rises, thereby improving the fit of exchange rate regressions based on macroeconomic indicators.

The model builds on a standard present value equation for the exchange rate, which, as Engel and West (2005) shows, gives rise to the following relationship between exchange rates and macroeconomic variables:

$$q_t = \xi \mathbf{f}_t' \Theta + \xi b_t + (1-\xi)\mathbf{f}_t' \mathbb{E}_t \Theta \tag{3}$$

where $\xi = (1-\lambda)/(1-\rho_b \lambda)$, λ is the discount factor, b_t is an unobservable variable that follows an AR(1) process with persistence parameter ρ_b, $\mathbf{f}_t$ is a vector of observable indicators, and Θ is a vector of structural, unknown parameters. Since the discount factor λ is close to one, the first two terms are negligible; consequently, the effect of changes in fundamentals depends primarily on investors' expected parameter values and how these expectations evolve with observable indicators.

Scapegoats emerge from investors' rational confusion as they infer the true model parameters conditional on observables. Assuming that the factors $\mathbf{f}_t$ follow random walks, the solution to the signal extraction problem yields the following process for expectations:

$$\mathbb{E}_t \Theta = A_t \mathbb{E}_{t-1} \Theta + (I - A_t)\Theta + Z_t \tilde{\Delta} \mathbf{f}_t \epsilon_t^b, \tag{4}$$

where $\Delta \mathbf{f}_t = \mathbf{f}_t - \rho_b \mathbf{f}_{t-1}$, $Z_t \equiv P_{t-1}(\tilde{\Delta}\mathbf{f}_t' P_{t-1} \tilde{\Delta}\mathbf{f}_t + \sigma_b^2)^{-1}$ and P_t is the conditional variance of the signal innovation.[8] The interaction between the change in factors $\tilde{\Delta}\mathbf{f}_t$, and the innovation in the unobserved component of the exchange rate ϵ_t^b, captures the scapegoat effect. When a shock to b_t occurs, investors attribute "excessive" weight to those indicators that happen to change simultaneously.

Fratzscher et al. (2015) tests the scapegoat theory by proxying the vector of scapegoat parameters using data for 12 exchange rates over the 2000–2011 period. They survey up to 60 financial market participants, asking them to rank the importance of six macroeconomic variables as drivers of a country's exchange rate. Using the survey results as proxies for the scapegoat parameters $\mathbb{E}_t \Theta$, they find evidence supporting the theory.

Similarly, we use narrative shares $\hat{\Theta}_{j,t}$ as proxies for the vector of time-varying investor attention weights $\mathbb{E}_t \Theta$. We conjecture that narrative shares capture the importance investors attach to specific exchange rate drivers. While there are several ways to rationalize this link, a natural interpretation is that changes in the factors $\mathbf{f}_t$ become more newsworthy—and thus receive greater coverage—when investors attach a higher weight to them.

Table 4 reports results from the regression model estimated with and without the metatopic interactions. We test three versions of Eq. (3) using monthly, quarterly, and annual exchange rate changes, with the interactions between metatopics and macroeconomic series informed by the results in Figs. 6–8.[9] For instance, we interact policy-related indicators such as the Federal Funds rate and government debt with policy-related narratives MT1 and MT2. Similarly, we interact stock market variables, including the S&P 500 and the VIX, with MT4 and MT5—narratives related to equity prices and cross-border capital flows.

Most of the coefficients of the benchmark models without metatopic attention are not statistically significant. In contrast, the models interacting metatopic attention with macroeconomic variables yield coefficient estimates that are significantly different from zero in most cases. Moreover, we observe a large increase in the adjusted R^2 relative to the benchmark specifications, rising from 7% to 28% in the model predicting the annual exchange rate change, and from 1% to 7% in the regression predicting the monthly change. These improvements are comparable to those reported by Fratzscher et al. (2015).[10]

Further evidence on the scapegoat theory. Eq. (3) highlights how the fit of exchange rate regressions depends on the time-varying weights investors assign to macro indicators. Table 4 provides evidence supporting this channel. Next, we seek to further test the theory's predictions by examining the evolution of the attention weights themselves. According to Eq. (4), a variable becomes a scapegoat (receiving higher weight) when its development aligns with unexpected changes in the unobservable component of the exchange rate.

While we lack a direct measure of the unobservable component, we show in Appendix I that a calibrated version of the scapegoat model predicts that the result holds if one replaces unexpected changes with *actual* exchange rate changes. In other words, the theory predicts that investors systematically attribute higher weights to macro variables whenever they see a movement in such a variable in a direction that can explain the actual exchange rate, not just its unobservable component. Thus, the interaction between the change in the macro variable and the exchange rate should affect investors' weights. This stronger prediction yields an empirically implementable test.

Specifically, we regress the narrative attention series on the interaction between the change in the exchange rate and the narrative-related macroeconomic indicator. Formally, the regression reads:

$$\hat{\Theta}_{j,t} = \zeta_0 + \zeta_1(\Delta^h q_t \times f_{j,t}) + \nu_t \tag{5}$$

where $\hat{\Theta}_{j,t}$ is the estimated attention to narrative j, $f_{j,t}$ is the macro variable associated with that narrative, and $\Delta^h q_t$ is the h-horizon change in the exchange rate (where a negative value indicates dollar appreciation). We set $h = 12$ months, consistent with our earlier evidence that narratives primarily capture medium-to-low frequency exchange rate fluctuations. Table 5 reports the results. Note that we use the same macroeconomic variables and conventions as in Table 4. Specifically, we take annual changes of the macro variables

[8] Appendix I shows how to derive Eq. (4) which closely follows Appendix A of Bacchetta and van Wincoop (2013). Furthermore, note that the random walk assumption can be relaxed, all we need for parameter learning is that the processes for unobservable and observable variables have different persistence.

[9] Bacchetta and van Wincoop (2009) show that Eq. (3) also holds in differences. Thus, we estimate the equation using several differencing horizons.

[10] Fratzscher et al. (2015), Table 4, reports an adjusted R^2 of 8.3% for the constant parameter model (CP-MS) predicting EUR/USD monthly exchange rate changes between 2001 and 2011.

Table 4
Text-augmented exchange rate regressions.

Variable	$q_t - q_{t-1}$		$q_t - q_{t-3}$		$q_t - q_{t-12}$	
	(I)	(II)	(I)	(II)	(I)	(II)
(Intercept)	−0.03	0.28**	−0.08	0.75**	−0.27	2.90**
	[−0.12, 0.06]	[0.16, 0.39]	[−0.34, 0.19]	[0.41, 1.09]	[−1.38, 0.84]	[1.50, 4.31]
Federal Funds rate	−0.21**	0.27*	−0.63**	0.73	−2.23*	1.81
	[−0.31, −0.10]	[0.12, 0.43]	[−0.93, −0.34]	[0.28, 1.18]	[−3.38, −1.07]	[0.02, 3.59]
Government Debt	0.22*	0.33**	0.64**	0.90**	2.00*	3.93***
	[0.11, 0.34]	[0.19, 0.47]	[0.34, 0.94]	[0.52, 1.28]	[0.86, 3.14]	[2.67, 5.18]
Trade Balance	−0.12	−0.18*	−0.21	−0.45**	−0.46	−1.75**
	[−0.21, −0.03]	[−0.27, −0.08]	[−0.44, 0.01]	[−0.68, −0.23]	[−1.29, 0.38]	[−2.52, −0.98]
S&P 500	−0.06	0.05	0.31	0.62	1.65	2.58**
	[−0.18, 0.06]	[−0.08, 0.19]	[−0.06, 0.67]	[0.23, 1.01]	[0.28, 3.01]	[1.35, 3.81]
VIX	0.02	0.01	0.28	0.34	1.14	1.66*
	[−0.09, 0.12]	[−0.10, 0.12]	[0.00, 0.55]	[0.06, 0.61]	[0.10, 2.18]	[0.70, 2.62]
TFP (Interpolated)	0.08	−0.23**	0.10	−0.48*	0.42	0.51
	[−0.02, 0.17]	[−0.34, −0.11]	[−0.15, 0.35]	[−0.74, −0.22]	[−0.58, 1.41]	[−0.49, 1.52]
(MT1 + MT2) × FFR	–	−7.40***	–	−21.53***	–	−47.25*
		[−9.77, −5.03]		[−27.72, −15.35]		[−71.70, −22.81]
MT2 × Government Debt	–	−4.84*	–	−8.97	–	−63.58***
		[−7.35, −2.33]		[−15.46, −2.47]		[−85.54, −41.63]
MT3× Trade Balance	–	−0.39	–	10.00***	–	32.43
		[−2.32, 1.54]		[6.45, 13.55]		[11.56, 53.30]
(MT4 + MT5) × S&P 500	–	−7.13***	–	−19.25***	–	−36.95**
		[−9.74, −4.51]		[−26.08, −12.42]		[−54.98, −18.93]
(MT4 + MT5) × VIX	–	−2.78	–	−9.79*	–	−40.32**
		[−5.87, 0.30]		[−15.32, −4.25]		[−58.11, −22.54]
MT6 × TFP	–	23.75***	–	42.96**	–	−26.12
		[16.14, 31.35]		[24.65, 61.26]		[−69.65, 17.41]
Obs.	523	523	521	521	512	512
Adjusted R^2	0.01	0.07	0.03	0.12	0.07	0.28

Note: The table presents the estimated coefficients of the exchange rate changes on macroeconomic indicators and their interaction with metatopic attention series. The indicators are the Federal Funds rate and the h-month future change in the log of government debt and TFP, where the monthly log of TFP is interpolated from the quarterly utilization-adjusted TFP of Fernald (2012). Other indicators include the h-month past change in the trade balance to GDP ratio (with GDP interpolated from the quarterly indicator), as well as the h-month past changes in the log of the S&P 500 index and the log of the VIX. We use $h = 3$ for predicting monthly and quarterly changes, and $h = 12$ for the annual exchange rate change. All indicators are standardized. The sample period is from 1976m1 to 2019m8. One-standard deviation intervals are reported in brackets and computed using Newey-West standard errors. * p-value < 0.1, ** p-value < 0.05, *** p-value < 0.01.

(excluding the Federal Funds rate), and in three instances, the dependent variable is the sum of two attention series, reflecting the fact that narratives can relate to multiple macro variables.

We find that the interaction coefficients (ζ_1) are statistically significant for 4 out of 6 specifications, with the coefficient for the technology-related narrative being marginally insignificant. As a robustness check, we re-estimated Eq. (5) using all possible "mismatched" combinations (e.g., regressing Fiscal Policy attention on the interaction between the exchange rate and Total Factor Productivity). While we find statistical significance in 30% of these placebo cases—likely reflecting general equilibrium correlations between macro variables—the performance of the theory-consistent baseline is superior. The baseline specification yields significant results in 67% of cases (4 out of 6), more than double the "hit rate" of the placebo assignments. This differential suggests that while general volatility matters, the specific economic logic connecting a narrative to its correct macro variable provides substantially higher explanatory power.

Notably, the *sign* of these coefficients provides further validation for our interpretation of the attention series. Most DAGs in Table 2 imply a specific direction of comovement between the exchange rate and macro variables. Eq. (5) assesses whether narrative attention responds in a manner consistent with this comovement.

For instance, the monetary policy narrative (MT1) implies a negative comovement between the exchange rate and the Federal Funds rate. Thus, the attention that investors place on narrative MT1 should *decrease* whenever the exchange rate *depreciates* while the FFR rises; in other words, ζ_1 should be *negative*. Indeed, all estimated coefficients—including those for which we fail to reject the null—exhibit the sign consistent with the underlying causal chain.

For narratives where the DAG allows for ambiguous effects (i.e., Fiscal Expansion and Global Crises), the estimated signs reveal which causal chain prevails in media coverage. In particular, the results suggest that the Fiscal Expansion narrative is dominated by the interest rate channel (leading to appreciation), while the Global Crisis narrative is dominated by the flight-to-safety channel (also leading to appreciation).

Table 5
Determinants of narrative attention.

Variable	MT1 + MT2	MT2	MT3	MT4 + MT5	MT4 + MT5	MT6
Intercept	2.773***	2.268***	1.769***	2.044***	2.029***	1.232***
	[2.167, 3.378]	[1.794, 2.743]	[1.086, 2.452]	[1.478, 2.610]	[1.465, 2.594]	[1.004, 1.460]
$\Delta^{12}q\ \times$ Federal Funds Rate	−1.133***	–	–	–	–	–
	[−1.514, −0.752]					
$\Delta^{12}q\ \times$ Government Debt	–	−1.715***	–	–	–	–
		[−2.234, −1.197]				
$\Delta^{12}q\ \times$ Trade Balance	–	–	0.059	–	–	–
			[−0.124, 0.243]			
$\Delta^{12}q \times$ S&P 500	–	–	–	−0.609*	–	–
				[−0.920, −0.298]		
$\Delta^{12}q \times$ VIX	–	–	–	–	−0.472*	–
					[−0.730, −0.213]	
$\Delta^{12}q\ \times$ TFP (Interpolated)	–	–	–	–	–	−0.151
						[−0.297, −0.004]
Obs.	512	512	512	512	512	512
Adjusted R^2	0.096	0.207	−0.002	0.028	0.014	0.007

Note: The table reports estimated coefficients from regressions of narrative attention series on the interaction between the annual change in the real exchange rate and the narrative-related macroeconomic indicator, as in Eq. (5). The macroeconomic indicators are the Federal Funds rate and the annual future change in the log of government debt and TFP, where the monthly log of TFP is interpolated from the quarterly utilization-adjusted TFP of Fernald (2012), the annual past change in the trade balance to GDP ratio (with GDP interpolated from the quarterly indicator), the annual changes in the log of the S&P 500 index and the log of the VIX. All indicators are standardized. The sample period is from 1976m1 to 2019m8. One-standard deviation intervals are reported in brackets and computed using Newey-West standard errors. * p-value < 0.1, ** p-value < 0.05, *** p-value < 0.01.

4.3. Discussion

Overall, we have demonstrated that using metatopic attention series as proxies for time-varying importance weights proves useful in disciplining the unstable relationship between exchange rates and economic indicators shown in Fig. 9. The results in Table 4, together with those in Table 5, provide evidence of an exchange rate reconnect that is consistent with the scapegoat theory. These findings, however, warrant a word of caution regarding interpretation.

The regressors in Table 4 are not exogenous disturbances, and the results should therefore not be interpreted as identifying structural causal effects of macroeconomic variables on exchange rates. Rather, they document systematic comovement between exchange rates and macro aggregates, conditional on time-varying narrative attention. We interpret the directed causal chains implicit in media narratives as a reflection of preferences for simplified explanations, whereby correlation is framed as causation (Spiegler, 2016; Eliaz and Spiegler, 2020). The DAGs in Table 2 formalize these narrative causal attributions; they are not representations of causal relationships estimated in the empirical analysis.

The interaction results in Table 5 provide empirical support for this interpretation. We find that attention to a narrative increases when movements in the associated macroeconomic variable align with observed exchange rate changes. That is, narratives receive greater weight precisely when the joint realization of macro variables and exchange rates fits the causal chain implied by the narrative. Thus, while the underlying movement in the macro aggregates can result *from* exchange rate fluctuations and not vice versa, the narrative attention pattern can be naturally explained by agents attributing exchange rate movements to macroeconomic developments that provide a coherent ex-post explanation.

In this sense, Table 5 documents a mechanism of causal attribution consistent with the behavioral evidence in Charles and Kendall (2022) and with models of simplified causal reasoning such as Eliaz and Spiegler (2020).

A related concern arises from the distinction between exchange rate fluctuations due to changes in the expected path of interest rate differentials—often labeled "fundamental" fluctuations—and those operating through time variation in risk premia. One might therefore ask whether the reconnect we document should be interpreted as evidence of a tighter link between exchange rates and fundamentals in this narrow sense. In the baseline scapegoat model, we abstract from risk premia for expositional convenience, not because we view them as empirically irrelevant. The framework can be extended to environments with time-varying, potentially unobservable risk premia, giving rise to analogous scapegoat-type mechanisms. To explore this empirically, Appendix J implements a decomposition in the spirit of Engel (2016), separating the real exchange rate into the component driven by expected interest rate differentials and the component driven by expected currency returns, reflecting variation in risk premia. Re-estimating the regressions in Table 4 for each component, we find small differences in the explanatory power across the two sets of regressions, with, if anything, higher R^2 values in the regressions involving the risk premium component. Moreover, we find no systematic difference in the significance of the narrative coefficients. We conclude that the sources of variation captured by narrative attention are not confined to a single transmission channel, but are consistent with a large body of work in which macroeconomic conditions and risk premia are jointly determined.

Finally, while our baseline interpretation views narratives primarily as proxies for time-varying information and attention, the empirical success of the text-augmented regressions is also compatible with theories in which media coverage itself affects economic dynamics. In models such as Nimark and Pitschner (2019) and Chahrour et al. (2021), agents face attention constraints and delegate information collection to news producers. In such environments, news selection can influence the propagation of structural shocks, implying that media narratives may not only reflect but also shape the connection between exchange rates and macroeconomic conditions.

5. Conclusion

News media represents a valuable source of information for economists that complements, rather than substitutes, traditional macroeconomic indicators. Media coverage organizes information by salience and perceived relevance, offering a structured lens through which market participants interpret economic developments. Because this information is conveyed in textual form, it can provide useful insights for interpreting latent factors and evolving expectations that may not be fully captured by conventional time-series data.

In this paper, we have presented a simple yet novel approach to extracting economically meaningful information from media coverage. Focusing on the dollar exchange rate, we combine topic modeling with large language model–based narrative extraction to distill a large corpus of news headlines into six exchange rate narratives. These narratives capture multiple macroeconomic transmission processes involving exchange rate movements, and help organize the heterogeneous evidence linking exchange rates to macroeconomic and financial indicators.

Much work remains, particularly in integrating narrative-based information into standard macroeconomic frameworks. Here we have taken a step in that direction by relating narrative attention measures to the scapegoat theory of exchange rates and using metatopic attention series as empirical proxies for time-varying investor focus on macroeconomic indicators. Since similar attention-driven mechanisms may arise in other areas of macroeconomics and finance, the approach developed here may prove useful beyond the exchange rate context.

Declaration of generative AI and AI-assisted technologies in the writing process

During the preparation of this work the authors used ChatGPT to label the estimated topics and extract narratives. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Supervised LDA

In supervised Latent Dirichlet Allocation (sLDA), each document is associated with both latent topics and an observed response variable, which in our case is the natural logarithm of the monthly real exchange rate. The model assumes that both the topic-word and document-topic distributions follow Dirichlet distributions: the word distribution for each topic, $\beta_k \sim \mathrm{Dir}(\eta)$, and the topic distribution for each document, $\theta_d \sim \mathrm{Dir}(\alpha)$. In our model, we set the prior parameters $\eta = 2$ and $\alpha = 0.1$.

Given a corpus with D documents, each containing N_d words, and a vocabulary of V unique words, the generative process for each document d proceeds as follows:

For each document d:

1. Draw topic proportions $\theta_d \sim \mathrm{Dir}(\alpha)$, where θ_d is a K-dimensional vector representing the distribution of topics in document d, with K being the number of topics.
2. For each word n:
   - Draw a topic assignment $z_{d,n} \sim \mathrm{Mult}(\theta_d)$, where $z_{d,n}$ represents the topic assigned to the n-th word in the document.
   - Draw the word $w_{d,n} \sim \mathrm{Mult}(\beta_{z_{d,n}})$, where $\beta_k \sim \mathrm{Dir}(\eta)$ is the word distribution for topic k. This means that each word is drawn from the word distribution β corresponding to its assigned topic $z_{d,n}$.
3. Draw the response variable Y_d, given the topic assignments:

   $$Y_d \sim N(\phi' \bar{z}_d, \sigma^2)$$

   where $\bar{z}_d = \frac{1}{N_d} \sum_{n=1}^{N_d} z_{d,n}$ is the average topic assignment across all words in the document, and $N(\phi' \bar{z}_d, \sigma^2)$ is the Gaussian likelihood for the response variable, with parameters ϕ and variance σ^2.

Given the above generative process, the joint probability distribution for the observed words W, the topic assignments Z, the topic-word distributions β, the document-topic distributions θ, and the observed response variable Y is:

$$P(\theta, \beta, Z, W, Y) \propto \left[\prod_{k=1}^{K} P(\beta_k|\eta) \prod_{d=1}^{D} P(\theta_d|\alpha) \prod_{n=1}^{N_d} P(z_{d,n}|\theta_d) P(w_{d,n}|\beta_{z_{d,n}}) \right] P(Y_d|\phi, \bar{z}_d, \sigma^2)$$

To estimate the model parameters $\{\theta_d, \beta_k, \phi, \sigma^2\}$, we use Gibbs sampling. In each iteration, we alternate between:

1. Sampling the topic assignments $z_{d,n}$ for each word in the corpus, and
2. Updating the model parameters based on these assignments.

After each iteration:

- The document-topic distribution θ_d is estimated as:

$$\theta_d = \frac{n_{dk} + \alpha}{N_d + K\alpha}$$

where n_{dk} is the count of words in document d assigned to topic k.
- Similarly, the topic-word distribution β_k for each topic k is estimated as:

$$\beta_k = \frac{m_{kv} + \eta}{M_k + V\eta}$$

where m_{kv} is the count of times word v is assigned to topic k, M_k is the total number of words assigned to topic k, and V is the size of the vocabulary.
- The regression parameters ϕ and the variance σ^2 are updated based on the document's estimated topic proportions θ_d and the response variable Y_d.

After a predefined number of iterations, which we set to 100, the Gibbs sampling process produces estimates for the posterior distributions of θ_d, β_k, ϕ, and σ^2, as well as the inferred topic assignments $z_{d,n}$.

A.1. Model performance

To guide the selection of the number of topics K, we estimate the model for a range of topics, from 20 to 220, and compute the values of semantic coherence and exclusivity for each model.

Semantic coherence measures the degree to which the most probable words in a topic frequently co-occur in documents, and is computed as:

$$C(T) = \sum_{i<j} \log \frac{D(w_i, w_j) + \epsilon}{D(w_j)},$$

where $D(w_i, w_j)$ is the number of documents in which both words w_i and w_j co-occur, and $D(w_j)$ is the number of documents that contain w_j. ϵ is a small constant added to avoid division by zero.

Exclusivity, on the other hand, captures how unique the top words of a topic are compared to other topics. A high exclusivity score indicates that the words within a topic are distinct from those in other topics, reducing overlap across topics. Exclusivity is computed as:

$$E(T) = \frac{1}{|T|} \sum_{w \in T} \left(\frac{p(w \mid z)}{\sum_{k \neq z} p(w \mid k)} \right)$$

where $p(w \mid z)$ is the probability of word w in topic z, and $\sum_{k \neq z} p(w \mid k)$ is the sum of the probabilities of word w across all other topics $k \neq z$.

Fig. 10 reports the results, showing how both metrics vary across different topic numbers. In our baseline, we choose $K = 180$ which gives a large exclusivity number without compromising much on semantic coherence.

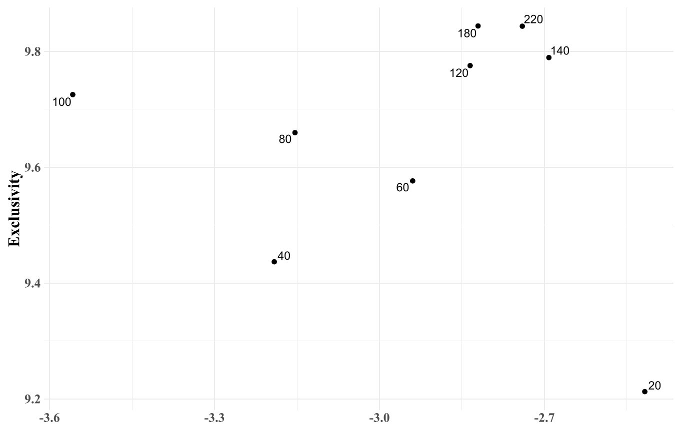

Fig. 10. Topic model performance for different topic numbers. *Note*: The figure reports the estimated values of semantic coherence and exclusivity for models with different numbers of topics, from 20 to 220.

Appendix B. Text collection and cleaning

The newspaper text data was collected from ProQuest Historical Newspapers using the TDM tool between 2022-03-04 and 2022-03-28. The corpus includes the titles of all available features, news, commentary, editorial, and Front Page/Cover Stories published in the Wall Street Journal between 1889-07-08 and 2022-02-05. The full corpus included 4,323,637 unique titles.

Before creating the data feature matrix that we use to estimate the sLDA model, we filter headlines as follows:

1. We include all the distinct titles published from January 1976 to August 2019, consistent with the availability of the monthly real exchange rate.
2. We remove titles with the following patterns: *No Title* OR *Dividends Rep* OR *ted* OR *Stocks Ex-Dividend—Stockholder Meeting Brief:* OR *Corrections Amplifications:* OR *Corrections & Amplifications:* OR TITLE BEGINS *REVIEW* OR TITLE BEGINS *REVIEW amp; OUTLOOK (Editorial):* OR TITLE BEGINS *Business Brief:* OR *Theater:* OR *Dividend News:* OR *Sports* OR *Film* OR *Letters to the Editor:* OR *Co.* TITLE ENDS OR *Corp.* TITLE ENDS OR *Inc.* TITLE ENDS OR *Bookshelf:* OR *Opera:* OR *Gardening:* OR *Seeing Stars:* OR *Thinking Things Over:* OR *Art* OR TITLE BEGINS *Books:* OR *Television:* OR *financial briefing book:* OR *reporter's notebook.*
3. We use the spaCy library (an open-source Python library developed for part-of-speech tagging) to perform the following:
   (a) Remove words with the following entity categorizations: CARDINAL, DATE, EVENT, FAC, GPE, LANGUAGE, LAW, LOC, MONEY, NORP, ORDINAL, ORG, PERCENT, PERSON, PRODUCT, QUANTITY, TIME and WORK OF ART;
   (b) Keep adjectives (ADJ), nouns (NOUN) and verbs (VERB);
   (c) Transform remaining words into their lemma form.
4. We remove words that include numbers or punctuation, a list of unwanted words (stopwords, weekdays, months and numbers) and one character words.
5. We remove titles with fewer than 10 words.

Finally, we keep the top 15,000 words according to the tf-idf. The final corpus includes 707,984 titles with the following breakdown (Table 6):

Table 6
WSJ headlines corpus breakdown by text type.

Type	Commentary	Feature	Front Page/Cover Story	News	Total
N	34,999	125,521	1493	545,971	707,984

Appendix C. Predictions of the random forest

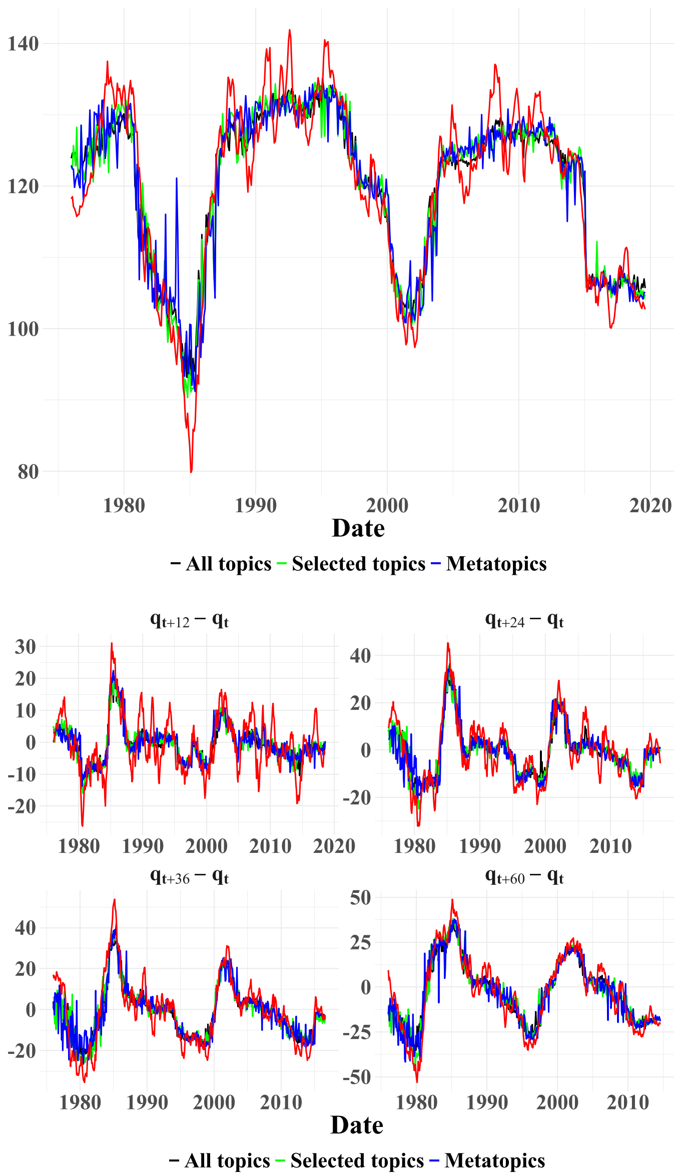

Fig. 11. Actual and predicted exchange rate values from Random Forest models. *Note*: The figure shows the log real exchange rate (red lines) and the Random Forest prediction in (1) using all topic proportions (black lines), sixteen topic proportions (green lines), and six metatopic proportions (blue lines).

Appendix D. Random forest with macroeconomic indicators

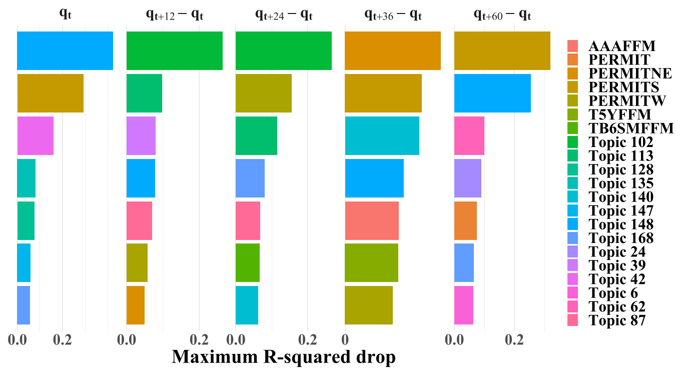

Fig. 12. Predictor importance ranking from Random Forest models. *Note*: The figure reports the predictors that cause the largest drop in R^2 when excluded from the models in (1). The predictors include 180 topics and 120 time series from McCracken and Ng (2016). AAAFFM is the spread between Moody's Aaa Corporate Bond and the Federal Funds Rate, PERMIT is the number of new housing permits in the U.S., PERMITNE, PERMITS, and PERMITW are the numbers of new housing permits in the northeast, in the south, and in the west of the U.S., T5YFFM is the spread between the 5-year Treasury and the Federal Funds Rate, and TB6SMFFM is the spread between the 6-month Treasury and the Federal Funds Rate.

Appendix E. Forming metatopics via hierarchical clustering

To systematically group topics from a topic model into higher-level metatopics, we apply hierarchical clustering to the topic-word distributions β_k. Let K denote the total number of topics, and let each topic $k = 1, \dots, K$ be represented by a word probability distribution:

$$\beta_k = (\beta_{k,1}, \beta_{k,2}, \dots, \beta_{k,V}) \in \mathbb{R}^V,$$

where V is the vocabulary size and $\beta_{k,v}$ denotes the probability of word v in topic k.

To ensure comparability across dimensions, we standardize each topic-word vector by centering and scaling each word dimension across topics. Let $\tilde{\beta}_k$ denote the resulting standardized vector for topic k, where each element $\tilde{\beta}_{k,v}$ is the standardized value of word v in topic k.

We then compute pairwise Euclidean distances between the standardized topic vectors:

$$d(k, m) = \sqrt{\sum_{v=1}^{V} (\tilde{\beta}_{k,v} - \tilde{\beta}_{m,v})^2}.$$

The pairwise distance matrix is:

$$\mathbf{D} = [d(k, m)] \in \mathbb{R}^{K \times K}.$$

Given the distance matrix, we perform agglomerative hierarchical clustering using the following procedure:

1. Each topic k starts as its own singleton cluster. Let C_i denote a cluster of topic indices at any step in the algorithm, so initially $C_i = \{k\}$ for each k.
2. At each step, the two clusters with the smallest pairwise distance are merged.
3. The distance between two clusters C_i and C_j is computed using the complete linkage method:

$$d(C_i, C_j) = \max_{k \in C_i, m \in C_j} d(k, m).$$

4. Steps 2 and 3 are repeated until all topics are merged into a single dendrogram representing the hierarchical structure.

The final topic clusters, or metatopics, are obtained by cutting the dendrogram at a chosen level, yielding J disjoint groups of topics. Let $\mathcal{M}_1, \dots, \mathcal{M}_J$ denote these metatopics, where each $\mathcal{M}_j \subseteq \{1, \dots, K\}$ is a non-empty set of topic indices. The collection $\{\mathcal{M}_1, \dots, \mathcal{M}_J\}$ forms a partition of the topic index set.

Appendix F. Feature contributions in random forest predictions

We use the featureContrib() function from the `tree.interpreter` package to decompose the Random Forest predictions into metatopic contributions. The additive decomposition approach method decomposes a Random Forest model's predictions into individual contributions from each explanatory variable. The Random Forest model is an ensemble of B regression trees, denoted as T_b, used to predict the natural logarithm of the monthly real exchange rate, q_t, based on the metatopic prevalences X_t. The ensemble prediction is given by:

$$\hat{q}_t = \frac{1}{B}\sum_{b=1}^{B} T_b(X_t). \tag{6}$$

Each tree T_b partitions the feature space into regions corresponding to leaf nodes, with each internal node representing a decision based on a specific feature (in this case, a metatopic prevalence). As an input X_t traverses a tree from the root to a leaf, it follows a unique path determined by the feature-based decisions at each node.

To quantify the contribution of each metatopic j to the prediction $\hat{q}_t$, we examine the changes in the predicted value as X_t progresses through the nodes of each tree. In a decision tree, a parent node is any non-terminal node that applies a decision rule to split the data into two or more subsets. Each subset corresponds to a child node that inherits from the parent. In the context of random forests, which are ensembles of decision trees, parent nodes define intermediate decision points within individual trees and determine the conditions under which feature contributions propagate along a path from the root to a leaf node. Specifically, for a given tree T_b, let $\mathcal{P}_b(X_t)$ denote the set of nodes that X_t traverses. For each node $n \in \mathcal{P}_b(X_t)$, associated with metatopic j, we define the contribution $\phi_{j,t}^{(b,n)}$ as the change in the predicted value caused by the split at node n:

$$\phi_{j,t}^{(b,n)} = v_n - v_{\mathrm{parent}(n)}, \tag{7}$$

where $v_{\mathrm{parent}(n)}$ is the prediction value at the parent node of n, and v_n is the prediction value at node n. The total contribution of metatopic j in tree T_b is then the sum of $\phi_{j,t}^{(b,n)}$ over all nodes $n \in \mathcal{P}_b(X_t)$ associated with metatopic j:

$$\phi_{j,t}^{(b)} = \sum_{n\in\mathcal{P}_b(X_t),\,\mathrm{feature}(n)=j} \phi_{j,t}^{(b,n)}. \tag{8}$$

Averaging over all trees in the forest, the overall contribution of metatopic j to the prediction $\hat{q}_t$ is:

$$\phi_{j,t} = \frac{1}{B}\sum_{b=1}^{B} \phi_{j,t}^{(b)}. \tag{9}$$

Thus, the predicted exchange rate $\hat{q}_t$ can be decomposed into the sum of the baseline prediction ϕ_0 (typically the average prediction over all instances) and the contributions from each metatopic:

$$\hat{q}_t = \phi_0 + \sum_{j=1}^{N} \phi_{j,t}. \tag{10}$$

In this decomposition, $\phi_{j,t}$ represents the influence of metatopic j on the prediction at time t.

Appendix G. Prompt for topic labels

LLM prompt for labeling topics

I have an Excel file representing the output of a topic model applied to a corpus of news articles spanning 1976 to 2020. The second sheet contains topic-word distributions (beta vectors), with each row corresponding to a topic and columns representing words with associated probabilities. We are using this data to study long-term (3-10 year) narratives that may explain changes in the US dollar exchange rate. From the 180 topics, we selected a subset of 16 topics (topics 6, 24, 39, 42, 62, 87, 96, 102, 113, 117, 128, 135, 140, 147, 148, and 168). These 16 topics have also been grouped into 6 metatopics via hierarchical clustering of their beta vectors:

1. (39, 6, 102, 113)
2. (24, 62, 148)
3. (42, 135)
4. (87)
5. (96, 147, 168)
6. (117, 128, 140)

Can you help us assign meaningful, interpretable labels to each of the 16 topics and the 6 metatopics, based on the topic-word distributions?

Appendix H. Narrative retrieval: Prompt and headlines

LLM prompt for narrative retrieval

```
For the following task do not use any information from our previous chats. I am analyzing six
separate topics from a topic model based on WSJ news article titles. The labels of the topics are
the following:

  1. Global Finance and Monetary Policies
  2. Historical Economic Policies and Energy Markets
  3. International Business and Financial News
  4. Corporate Transactions and Stock Market
  5. Economic Trends and Corporate Strategy
  6. Technological Advancements and Financial Reporting.

I attached an excel file with the most representative headlines for each topic. Please summarize
the main causal relationships expressed or implied in these titles. Use the publication date
information in the analysis as well. Try to recover the main text or abstract related to a title and
use that instead of the title for the analysis. Focus on causal claims involving economic variables,
especially the dollar exchange rate. Use the language of causal inference and directed acyclic graphs
(DAGs).
Your output should include:

  1. A label and a short summary (1 sentence) of what this topic is about in terms of causal
     relationships.
  2. A list of 1-3 inferred causal pathways represented as DAGs.
  3. What is the share of articles that imply the summary.
  4. Extract 5-8 titles that represent the main narrative/causal claim. Avoid vague associations.
     If the titles are too weak or don't contain clear causal content, say so.
```

Table 7 reports the top titles that the LLM considered most representative of the economic narratives.

Table 7
Representative headlines of exchange rate narratives from the LLM.

MT1: Global Finance and Monetary Policies
1980-12-04: Interest Rates, Concern On Poland Cause Drop For Soybeans, Grain. 1981-06-12: French Minister Urges Firm European Protest Over U.S. Interest Rates. 1983-12-27: Foreign-Loan Rule Isn't Seen Hurting U.S. Bank Earnings. 1984-07-20: Banks' Foreign-Exchange Revenues Dive After Wrong Guessing on Dollar's Strength. 1985-01-02: U.S. Dollar Ends Trading in 1984 By Setting Highs.
MT2: Historical Economic Policies and Energy Markets
1976-03-24: "Ex-Maker of Kepone, Allied Chemical Sued On Pollution Charges." 1977-02-28: "Major Potash Firm Dropped From Case Tied to Price-Fixing." 1977-07-07: "New England Electric Unit Rate Boosts Cut In FPC Aide's Ruling." 1977-12-01: "Alberta Lease Sale For Oil, Natural Gas Sets Record Prices." 1979-03-22: "Carter Urges Increase In Authority of HUD In Housing-Bias Cases." 1979-07-16: "Amerada Hess Agrees To Follow Council's Anti-Inflation Guides."
MT3: International Business and Financial News
2017-05-24: "China's World: Big Brother Tightens Control By Mining Foreign Firms' Data." 2018-04-20: "Iceland Rethinks Tech Boom's Cost — Cheap power draws businesses, but worries emerge of risks to environment." 2019-01-02: "High Emerging-Markets Leverage Sparks Anxiety." 2019-02-23: "Saudi Arabia Seeks Closer Ties to Beijing — Prince Mohammed signs raft of deals, including pact for petrochemical refinery." 2019-03-09: "U.S.-China Trade Talks Hit a Bump — Outcome of Trump-Kim summit stokes concern in Beijing on setting meeting of presidents." 2019-07-17: "Wall Street Finds a New Way To Bet on Climbing Debt Risk."
MT4: Corporate Transactions and Stock Market
1997-10-14: "AlliedSignal Inc.: Agreement Is Announced To Acquire Astor Holdings." 1997-11-20: "Intel CEO Says Asia Must Invest More In Information Technology for Growth." 1998-03-13: "Stock-Market Merger Talk Hits Europe." 1998-05-14: "Tech-Sector Strength Boosts Small-Cap Issues; C-Phone More Than Triples Value, Onsale Soars." 1998-07-16: "Netscape Will Let PC Users Try Free Upgrade of Internet Software."

(*continued on next page*)

Table 7 (*continued*)

MT5: Economic Trends and Corporate Strategy
2002-10-09: "How the Dollar Thrives Amid Global Woes."
2002-12-05: "El Paso Ex-Trader Indicted Over Energy Reports."
2003-04-18: "Soybeans See Contract Highs On Solid Demand, Low Stocks."
2003-05-19: "Hog Prices Reach New Highs On Active Speculative Buying."
2003-05-29: "Live-Cattle Futures Hit Highs As Canada Beef Ban Continues."
2003-09-10: "Gold Rally Keeps on Shining Despite Stocks' Upward Trend."
MT6: Technological Advancements and Financial Reporting
1984-02-02: "London Shares Drop; Tokyo Issues Are Up."
1984-07-27: "Wall Street Spurs Share Prices to Rise In Tokyo, London."
1985-10-24: "U.S. Currency Gains As Dollar Purchase Offsets Intervention."
1993-05-13: "Roadway Services Inc.: Trucking Concern to Start Heavy Air Freight Business."
1993-06-15: "Barr Laboratories Inc.: Two Generic Drugs Return After More Studies for FDA."
2003-10-13: "International Business Machines: New Software Will Allow Links."

Note: The headlines reported here exclude the subject or section tags that often precede WSJ titles in the raw data. These tags are retained in the sLDA estimation because they are informative to readers and can help classify narratives. As a robustness check, we re-estimated the topic model after removing such tags. The resulting metatopics and their time-series prevalence remain similar to the baseline results.

Appendix I. Scapegoat theory: Equilibrium equations and calibration

The model consists of two equilibrium equations: the interest parity condition, and a second equation for the interest rate differential—common to several open economy models (see Engel and West, 2005)—that reads:

$$r_t - r_t^* = \frac{1-\lambda}{\lambda}\left(q_t - \mathbf{f}_t'\Theta - b_t\right) \tag{11}$$

Assume that the variables in $\mathbf{f}_t$ are random walks with i.i.d. error terms and standard deviation σ_f, whereas b_t evolves as an AR(1) process with persistence ρ_b and error term ϵ_t^b with standard deviation σ_b. Combining the interest rate parity equation with Eq. (11) yields Eq. (3).

Table 8
Baseline parameter Calibration.

(a) Externally Set Parameters		
	Description	Value
Θ	True parameter value	5
σ	Initial uncertainty	2.5
λ	Discount factor	0.97
N	Number of observed factors	6

(b) Targeted Moments and Calibrated Parameters					
	Description	Value	Moment	Target	Model
σ_b	Std. dev. unobserved fundamental	$3.40 \cdot 10^{-2}$	Autocorr. $\Delta^{12}q$	0.95	0.91
ρ_b	Autocorr. of unobserved fundamental	0.99	Std. dev. $\Delta^{12}q$	$8.91 \cdot 10^{-2}$	$8.77 \cdot 10^{-2}$
σ_f	Std. dev. observed fundamentals	$1.04 \cdot 10^{-5}$	R^2 of q on $\mathbf{f}$	0.07	0.09

Note: One period in the model represents one month. The target R^2 is from the specification in annual changes reported in Table 4. In the model, we regress annual changes of q on annual changes of $\mathbf{f}$.

Signal extraction problem. Agents form expectations about the parameter vector Θ. Since they observe both the exchange rate and the interest rate differential, the signal can be written as $y_t = \mathbf{f}_t'\Theta + b_t$. Following Hamilton (1994), Section 13.8, the Kalman filter yields:

$$\mathbb{E}_t\Theta = \mathbb{E}_{t-1}\Theta + K_t(\tilde{\Delta} y_t - \tilde{\Delta}\mathbf{f}_t'\mathbb{E}_{t-1}\Theta) \tag{12}$$

where the Kalman gain is $K_t \equiv P_{t-1}\tilde{\Delta}\mathbf{f}_t(\tilde{\Delta}\mathbf{f}_t' P_{t-1}\tilde{\Delta}\mathbf{f}_t + \sigma_b^2)^{-1}$. Given initial uncertainty $P_1 = \sigma I_N$, the uncertainty is updated as $P_t = A_t P_{t-1}$, where $A_t \equiv I_N - P_{t-1}\tilde{\Delta}\mathbf{f}_t(\tilde{\Delta}\mathbf{f}_t' P_{t-1}\tilde{\Delta}\mathbf{f}_t + \sigma_b^2)^{-1}\tilde{\Delta}\mathbf{f}_t'$. Rearranging terms in Eq. (12), yields Eq. (4).

Calibration. Following Bacchetta and van Wincoop (2013), we calibrate the structural parameters λ and Θ, as well as the initial belief uncertainty σ to the values used in their benchmark parametrization. We set the number of observed indicators N to 6, representing the six narratives we estimate. Finally, we estimate the parameters of the shock processes ($\sigma_f, \sigma_b, \rho_b$) via Simulated Method of Moments to match the volatility and persistence of the annual change in the exchange rate in our sample, together with the R^2 of the exchange rate on fundamentals, which we estimate to be equal to 7% (see Table 4). We do so via Monte Carlo simulation of 1000 samples of size 512 as in the data. Table 8 reports the parameter values and the median moments simulated from the model.

Simulation of regression (5). We estimate the model prediction of the coefficient ζ_1 in Eq. (5). We do so by estimating the regression in 1000 samples of size 512 with a burn-in sample of 10 years. While we follow the same variable transformation as in the data, we also take the 12-horizon difference of the expected weights $\mathbb{E}_t\Theta$ to remove the long run trend that emerges from the underlying learning. Finally, we compute Newey-West standard errors.

The simulation yields a median p-value of 0.29%, confirming that the scapegoat model predicts a significant relationship between narrative attention and the interaction of fundamentals with the exchange rate change. In addition, we verify that the p-value remains in the 1% rejection region as we change the value of initial uncertainty σ or the differencing horizon.

Appendix J. Exchange rate decomposition

Let $\lambda_{t+1} \equiv \Delta q_{t+1} - (r_t - r_t^*)$ be the excess return on the foreign exchange. Assuming stationarity of the real exchange rate and the interest rate differential, the exchange rate can be decomposed into a UIP component (q^{UIP}) and an expected excess-return component (q^{λ}):

$$q_t = \underbrace{-\sum_{k=0}^{\infty} \mathbb{E}_t(r_{t+k} - r^*_{t+k})}_{q_t^{UIP}} \underbrace{-\sum_{k=0}^{\infty} \mathbb{E}_t \lambda_{t+k+1}}_{q_t^{\lambda}} . \tag{13}$$

Following Engel (2016), we estimate a VECM with the vector of variables Y_t including the nominal exchange rate, the CPI differential, and the interest rate differential. The sample goes from August 1978 to August 2019. The VECM reads:

$$\Delta Y_t = \mathbf{C_0} + \mathbf{G} Y_{t-1} + \mathbf{C}(L) L \Delta Y_{t-1} + U_t,$$

where the polynomial matrix $\mathbf{C}(L)$ is unrestricted, and the matrix $\mathbf{G}$ takes the form:

$$\mathbf{G} = \begin{bmatrix} g_{11} & -g_{11} & g_{13} \\ g_{21} & -g_{21} & g_{23} \\ g_{31} & -g_{31} & g_{33} \end{bmatrix},$$

which restricts the cointegrating vectors. As in Engel (2016), we set the number of lags to three, but we observe that the results are largely insensitive to this choice.

Using expectation measures implied by the VECM, we compute the two exchange rate components on the right of Eq. (13), and re-estimate the set of regressions in Table 4 separately for each exchange rate component. Table 9 (a) and (b) report the results.

Table 9
Additional text-augmented exchange rate regressions.

(a) Text-augmented exchange rate regressions using q^{UIP}

Variable	$q_t^{UIP} - q_{t-1}^{UIP}$		$q_t^{UIP} - q_{t-3}^{UIP}$		$q_t^{UIP} - q_{t-12}^{UIP}$	
	(I)	(II)	(I)	(II)	(I)	(II)
(Intercept)	0.00	0.06	−0.01	0.17*	−0.09	0.40
	[−0.04, 0.04]	[0.02, 0.10]	[−0.10, 0.09]	[0.07, 0.26]	[−0.36, 0.18]	[0.06, 0.74]
Federal Funds rate	0.04	0.12**	−0.01	0.27**	−0.48	0.43
	[−0.03, 0.11]	[0.06, 0.18]	[−0.17, 0.15]	[0.16, 0.39]	[−0.80, −0.16]	[0.06, 0.80]
Government Debt	−0.05	−0.01	0.01	0.06	1.03***	1.00***
	[−0.09, −0.00]	[−0.04, 0.02]	[−0.06, 0.08]	[0.00, 0.11]	[0.76, 1.29]	[0.75, 1.25]
Trade Balance	0.01	−0.01	0.13	0.10	0.22	0.22
	[−0.03, 0.04]	[−0.05, 0.03]	[0.01, 0.24]	[−0.02, 0.22]	[−0.08, 0.52]	[−0.02, 0.46]
S&P 500	−0.15**	−0.15*	0.00	0.06	0.32	0.47
	[−0.22, −0.08]	[−0.23, −0.06]	[−0.12, 0.12]	[−0.09, 0.21]	[0.02, 0.63]	[0.18, 0.76]
VIX	−0.10*	−0.12	−0.10	−0.13	0.30	0.31
	[−0.16, −0.04]	[−0.19, −0.05]	[−0.21, −0.00]	[−0.26, 0.00]	[0.12, 0.49]	[0.10, 0.51]
TFP (Interpolated)	−0.04	−0.13*	−0.05	−0.17	0.10	0.41
	[−0.09, 0.01]	[−0.21, −0.06]	[−0.16, 0.05]	[−0.31, −0.02]	[−0.14, 0.35]	[0.16, 0.67]
(MT1 + MT2) × FFR	–	−1.11	–	−4.43	–	−16.64***
		[−2.70, 0.48]		[−7.64, −1.22]		[−23.05, −10.22]
MT2 × Government Debt	–	−1.84	–	−1.52	–	6.98
		[−4.25, 0.57]		[−4.79, 1.75]		[−0.15, 14.11]
MT3× Trade Balance	–	−0.21	–	−0.60	–	1.85
		[−0.69, 0.26]		[−1.81, 0.61]		[−3.33, 7.02]
(MT4 + MT5) × S&P 500	–	0.17	–	−4.10**	–	−9.68
		[−0.75, 1.09]		[−5.96, −2.25]		[−16.69, −2.67]

(continued on next page)

Table 9 (*continued*)

(a) Text-augmented exchange rate regressions using q^{UIP}						
(MT4 + MT5) × VIX	–	−0.50 [−0.32, 1.33]	–	−0.64 [−2.15, 0.87]	–	−1.64 [−4.76, 1.49]
MT6 × TFP	–	6.90** [4.11, 9.68]	–	8.57 [3.14, 14.01]	–	−27.47*** [−37.36, −17.58]
Number of observations	488	488	486	486	477	477
Adjusted R^2	0.01	0.03	0.01	0.03	0.17	0.28

(b) Text-augmented exchange rate regressions using q^{λ}

Variable	$q_t^{\lambda} - q_{t-1}^{\lambda}$		$q_t^{\lambda} - q_{t-3}^{\lambda}$		$q_t^{\lambda} - q_{t-12}^{\lambda}$	
	(I)	(II)	(I)	(II)	(I)	(II)
(Intercept)	−0.06 [−0.15, 0.02]	0.19* [0.08, 0.31]	−0.20 [−0.46, 0.07]	0.46 [0.11, 0.80]	−0.85 [−1.91, 0.21]	1.78 [0.32, 3.24]
Federal Funds rate	−0.26*** [−0.37, −0.16]	0.13 [−0.02, 0.28]	−0.70** [−1.00, −0.41]	0.33 [−0.10, 0.77]	−2.49*** [−3.39, −1.59]	0.62 [−1.14, 2.38]
Government Debt	0.27** [0.16, 0.38]	0.34*** [0.21, 0.47]	0.66** [0.36, 0.95]	0.84** [0.49, 1.20]	1.45 [0.40, 2.50]	2.92** [1.69, 4.15]
Trade Balance	−0.14 [−0.23, −0.04]	−0.18* [−0.28, −0.08]	−0.32 [−0.52, −0.11]	−0.54** [−0.76, −0.33]	−0.53 [−1.26, 0.20]	−1.71** [−2.49, −0.92]
S&P 500	0.09 [−0.04, 0.22]	0.21 [0.06, 0.37]	0.38 [0.00, 0.76]	0.60 [0.18, 1.03]	1.92 [0.59, 3.25]	2.56** [1.30, 3.81]
VIX	0.13 [0.02, 0.24]	0.13 [0.02, 0.25]	0.40 [0.11, 0.69]	0.46 [0.14, 0.77]	0.89 [−0.12, 1.89]	1.36 [0.38, 2.34]
TFP (Interpolated)	0.13 [0.04, 0.22]	−0.09 [−0.20, 0.02]	0.20 [−0.03, 0.43]	−0.27 [−0.54, 0.00]	0.62 [−0.28, 1.52]	0.54 [−0.39, 1.47]
(MT1 + MT2) × FFR	–	−6.09** [−8.56, −3.62]	–	−16.45** [−22.97, −9.93]	–	−30.44 [−51.63, −9.26]
MT2 × Government Debt	–	−2.98 [−5.93, −0.03]	–	−5.99 [−13.02, 1.05]	–	−57.04*** [−75.16, −38.92]
MT3× Trade Balance	–	−0.05 [−1.82, 1.72]	–	10.45*** [7.18, 13.72]	–	28.46 [9.74, 47.17]
(MT4 + MT5) × S&P 500	–	−7.31*** [−9.86, −4.77]	–	−15.14** [−21.77, −8.52]	–	−29.75 [−48.71, −10.78]
(MT4 + MT5) × VIX	–	−3.24 [−6.11, −0.38]	–	−8.60 [−14.16, −3.04]	–	−35.88** [−52.87, −18.89]
MT6 × TFP	–	16.84** [9.93, 23.75]	–	33.38** [16.48, 50.28]	–	−12.02 [−53.39, 29.34]
Obs.	488	488	486	486	477	477
Adjusted R^2	0.03	0.07	0.05	0.10	0.09	0.25

Note: The table reports estimated coefficients from regressions in which the two components of the exchange rate decomposition—the expected cumulative interest rate differential (Panel A) and the expected cumulative excess currency return (Panel B)—are used as dependent variables. The macro indicators are the Federal Funds rate and the h-month future change in the log of government debt and TFP, where the monthly log of TFP is interpolated from the quarterly utilization-adjusted TFP of Fernald (2012). Other indicators include the h-month past change in the trade balance to GDP ratio (with GDP interpolated from the quarterly indicator), as well as the h-month past changes in the log of the S&P 500 index and the log of the VIX. We use $h = 3$ for predicting monthly and quarterly changes, and $h = 12$ for the annual exchange rate change. All indicators are standardized. One-standard deviation intervals are reported in brackets and computed using Newey-West standard errors. * p-value < 0.1, ** p-value < 0.05, *** p-value < 0.01.

Data availability

The authors do not have permission to share data.

References

Agarwal, I., Chen, W., Prasad, E.S., 2024. Beyond the Fundamentals: How Media-Driven Narratives Influence Cross-Border Capital Flows. National Bureau of Economic Research.
Andre, P., Haaland, I., Roth, C., Wohlfart, J., 2023. Narratives about the macroeconomy. CESifo Working Paper.
Ash, E., Hansen, S., 2023. Text algorithms in economics. Annu. Rev. Econ. 15, 659–688.
Bacchetta, P., van Wincoop, E., 2004. A scapegoat model of exchange-rate fluctuations. Am. Econ. Rev. 94, 114–118.
Bacchetta, P., van Wincoop, E., 2009. On the unstable relationship between exchange rates and macroeconomic fundamentals. NBER Working Paper.
Bacchetta, P., van Wincoop, E., 2013. On the unstable relationship between exchange rates and macroeconomic fundamentals. J. Int. Econ. 91, 18–30.
Baker, S.R., Bloom, N., Davis, S.J., 2016. Measuring economic policy uncertainty. Q. J. Econ. 131, 1593–1636.
Blei, D., McAuliffe, J., 2007. Supervised topic models. Neural Inf. Process. Syst. 21.
Blei, D.M., Ng, A.Y., Jordan, M.I., 2003. Latent Dirichlet allocation. J. Mach. Learn. Res. 3, 993–1022.
Brandt, M.W., Cochrane, J.H., Santa-Clara, P., 2006. International risk sharing is better than you think, or exchange rates are too smooth. J. Monet. Econ. 53, 671–698.
Breiman, L., 2001. Random forests. Mach. Learn. 45, 5–32.
Bybee, L., Kelly, B., Manela, A., Xiu, D., 2024. Business news and business cycles. J. Finance. 79, 3105–3147.
Chahrour, R., Nimark, K., Pitschner, S., 2021. Sectoral media focus and aggregate fluctuations. Am. Econ. Rev. 111, 3872–3922.
Chahrour, R., Cormun, V., De Leo, P., Guerrón-Quintana, P.A., Valchev, R., 2024. Exchange rate disconnect revisited. Tech. rep., National Bureau of Economic Research.
Charles, C., Kendall, C., 2022. Causal narratives. NBER Working Paper.
Cheung, Y.-W., Chinn, M.D., 2001. Currency traders and exchange rate dynamics: a survey of the US market. J. Int. Money Finance 20, 439–471.
Cormun, V., De Leo, P., 2025. Shocks and exchange rates in small open economies. Working Paper.

Dornbusch, R., 1976. Expectations and exchange rate dynamics. J. Polit. Econ. 84, 1161–1176.
Eliaz, K., Spiegler, R., 2020. A model of competing narratives. Am. Econ. Rev. 110, 3786–3816.
Eliaz, K., Spiegler, R., 2024. News media as suppliers of narratives. arXiv preprint arXiv:2403.09155.
Engel, C., 2016. Exchange rates, interest rates, and the risk premium. Am. Econ. Rev. 106, 436–474.
Engel, C., West, K.D., 2005. Exchange rates and fundamentals. J. Polit. Econ. 113, 485–517.
Fernald, J.G., 2012. A quarterly, utilization-adjusted series on total factor productivity, Working Paper 2012-19, Federal Reserve Bank of San Francisco, available at https://www.frbsf.org/economic-research/files/wp12-19bk.pdf.
Frankel, J.A., 1979. On the mark: a theory of floating exchange rates based on real interest differentials. Am. Econ. Rev. 69, 610–622.
Frankel, J.A., Rose, A.K., 1995. Empirical research on nominal exchange rates. In: Handbook of International Economics, vol. 3. Elsevier, pp. 1689–1729.
Fratzscher, M., Rime, D., Sarno, L., Zinna, G., 2015. The scapegoat theory of exchange rates: the first tests. J. Monet. Econ. 70, 1–21.
Gentzkow, M., Kelly, B., Taddy, M., 2019. Text as data. J. Econ. Lit. 57, 535–574.
Hamilton, J.D., 1994. Time Series Analysis. Princeton University Press.
Hansen, S., McMahon, M., Prat, A., 2018. Transparency and deliberation within the FOMC: a computational linguistic approach. Q. J. Econ. 133, 801–870.
Heddaya, M., Tan, C., Voigt, R., Zeng, Q., Zentefis, A., 2025. A century of inflation narratives, Available at SSRN 5259107.
Itskhoki, O., Mukhin, D., 2021. Exchange rate disconnect in general equilibrium. J. Polit. Econ. 129, 2183–2232.
Jiang, Z., 2021. US fiscal cycle and the dollar. J. Monet. Econ. 124, 91–106.
Kim, S., Roubini, N., 2008. Twin deficit or twin divergence? Fiscal policy, current account, and real exchange rate in the US. J. Int. Econ. 74, 362–383.
Larsen, V.H., Thorsrud, L.A., 2019. The value of news for economic developments. J. Econom. 210, 203–218.
Macaulay, A., Song, W., 2023. Narrative-driven fluctuations in sentiment: evidence linking traditional and social media. Tech. rep., Bank of Canada.
Maggiori, M., 2022. International macroeconomics with imperfect financial markets. In: Handbook of International Economics, vol. 6. Elsevier, pp. 199–236.
McCracken, M.W., Ng, S., 2016. FRED-MD: a monthly database for macroeconomic research. J. Bus. Econ. Stat. 34, 574–589.
Meese, R.A., Rogoff, K., 1983. Empirical exchange rate models of the seventies: do they fit out of sample? J. Int. Econ. 14, 3–24.
Monacelli, T., Perotti, R., 2010. Fiscal policy, the real exchange rate and traded goods. Econ. J. 120, 437–461.
Nimark, K.P., 2014. Man-bites-dog business cycles. Am. Econ. Rev. 104, 2320–2367.
Nimark, K.P., Pitschner, S., 2019. News media and delegated information choice. J. Econ. Theory 181, 160–196.
Obstfeld, M., Rogoff, K., 2000. The six major puzzles in international macroeconomics: is there a common cause? NBER Macroecon. Annu. 15, 339–390.
Ravn, M.O., Schmitt-Grohé, S., Uribe, M., 2012. Consumption, government spending, and the real exchange rate. J. Monet. Econ. 59, 215–234.
Rossi, B., 2013. Exchange rate predictability. J. Econ. Lit. 51, 1063–1119.
Schmitt-Grohé, S., Uribe, M., 2022. The effects of permanent monetary shocks on exchange rates and uncovered interest rate differentials. J. Int. Econ. 135, 103560.
Shiller, R.J., 2017. Narrative economics. Am. Econ. Rev. 107, 967–1004.
Shiller, R.J., 2020. Narrative Economics: How Stories Go Viral and Drive Major Economic Events. Princeton University Press.
Spiegler, R., 2016. Bayesian networks and boundedly rational expectations. Q. J. Econ. 131, 1243–1290.
Thorsrud, L.A., 2020. Words are the new numbers: a newsy coincident index of the business cycle. J. Bus. Econ. Stat. 38, 393–409.