

Exchange Rate Narratives*

Vito Cormun[†]

Kim Ristolainen[‡]

November 27, 2024

Abstract

Leveraging *Wall Street Journal* news, recent developments in textual analysis, and generative AI, we estimate a narrative decomposition of the dollar exchange rate. Our findings shed light on the connection between economic fundamentals and the exchange rate, as well as on its absence. From the late 1970s onwards, we identify six distinct narratives that explain changes in the exchange rate, each largely non-overlapping. U.S. fiscal and monetary policies play a significant role in the early part of the sample, while financial market news becomes more dominant in the second half. Notably, news on technological change predicts the exchange rate throughout the entire sample period. Finally, using text-augmented regressions, we find evidence that media coverage explains the unstable relationship between exchange rates and macroeconomic indicators.

JEL Codes: C3, C5, F3

Keywords: Exchange rates, big data, textual analysis, macroeconomic news, Wall Street Journal, narrative retrieval, scapegoat

*We are grateful to Ryan Chahrour, Sanjiv Das, Pierre De Leo, Valeria Ferraro, and Rosen Valchev for their insightful comments. We also thank participants at 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, and Santa Clara University. Vito Cormun thanks the Leavey School of Business for their financial support, and Kim Ristolainen thanks Emil Aaltonen Foundation for their financial support.

[†]Santa Clara University, vcormun@scu.edu.

[‡]University of Turku, kkrist@utu.fi.

Introduction

We propose a novel approach to explain time series fluctuations using newspaper text. The procedure allows us to estimate potentially complex relations between news and a time series of interest. When adopted to explain the dollar exchange rate, it reveals a narrative decomposition that gives novel insights into the sources of exchange rate fluctuations.

While newspaper articles contain information about various economic topics and macroeconomic indicators, in this paper, we focus on exchange rates. We do so because (i) of the widely documented inability of macroeconomic time series to explain exchange rate fluctuations, both in a contemporaneous and in a forecasting sense (Meese and Rogoff, 1983), (ii) being an asset price, the exchange rate should be particularly sensitive to news (Engel and West, 2005), and (iii) newspaper articles often report financial market news which is a compelling source of exchange rate fluctuations (Itskhoki and Mukhin, 2021).

Our results reveal that there is no single narrative that can explain the last 50 years of exchange rate changes. Instead, we find as many as six narratives that are largely non-overlapping. Taken in isolation, each narrative relates to several time series and articles in a way consistent with proposed sources of exchange rate changes. Together, they provide an anatomy of the fluctuations with novel insights underlying the apparent disconnect between the exchange rate and the macroeconomy.

The paper is divided into four parts. In the first part, we describe our dataset and the model that we estimate. We use headlines from the *Wall Street Journal* (WSJ) articles published since January 1976. WSJ articles have been increasingly adopted for economic analysis. Recently, Bybee et al. (forthcoming) estimates a topic model – the unsupervised Latent Dirichlet Allocation (LDA) model – to reduce the corpus of the WSJ into a relatively small number of topics, where each topic is a distribution of terms that co-occur in articles. The unsupervised LDA model is similar to dimensionality reduction techniques used with time series, like principal component analysis (PCA), but instead of focusing on variance in numerical data, it identifies latent topics that explain the distribution of terms in a dataset. From topic mapping to words, one can compute topic attention over time, defined as the proportion of words assigned to a topic at each point in time, and use it to study the relation with economic time series. For instance,

Bybee et al. (forthcoming) finds that news attention to a topic about recession risk captures useful information on future economic outcomes, above and beyond common indicators such as stock prices, interest rates, and measures of economic uncertainty.

Our analysis shares a similar spirit to Bybee et al. (forthcoming), with the key difference that our primary focus is the relation between news and exchange rate changes. For this reason, we implement *supervised* LDA. Supervised LDA is an extension of unsupervised LDA in which topics are estimated to maximize their predictive power of an object of interest. This method has been developed in Blei and McAuliffe (2007), and applied in fields other than economics, for example, to estimate those topics from movie or product reviews that best explain variations in movie or product ratings. We target the monthly real exchange rate of the dollar vis-a-vis the trade-weighted G6 currency. To draw a comparison with time series econometrics, supervised LDA is similar to the “max-share” shock identification method applied in the business cycle literature using Vector Autoregressions, such as in the seminal work of Uhlig (2003), or in the more recent works of Angeletos et al. (2020) and Chahrour et al. (2024). As in supervised LDA, the max-share shock identification method decomposes the variation in the target variable by finding the latent predictors that best explain its variation. However, unlike sLDA, the outcome of the method consists of time series only, without the additional information from the word distributions.

In the second part of the paper, we assess the performance of the topic model and compare it against the one of standard economic time series. To estimate the topic model, we set the number of topics to 180 based on metrics capturing the balance between interpretability and distinctiveness of the topics. We then run Random Forest regressions with monthly topic attention series as independent variables, and contemporaneous and future changes in the real exchange rate as dependent variables. We find that topics explain most of the current and future exchange rate variations, with R^2 values ranging from 61% to 88% depending on the specification. Next, we demonstrate that topics provide a more accurate description of exchange rate changes than standard macroeconomic time series. We do so by using the FRED-MD dataset constructed in McCracken and Ng (2016), which consists of a set of 120 monthly economic indicators. The Random Forest regressions reveal that the R^2 values are smaller than

the ones estimated using topics. In addition, we estimate the Random Forest model using both topics and economic indicators, and find topics to have the largest explanatory power.

In the third part of the paper, we select and interpret the most important topics. The selection step comprises two parts. First, we select topics with the highest explanatory power. We verify that using sixteen topics delivers R^2 s that are almost unchanged relative to the model with 180 topics. Second, we cluster topics into six metaserries or metatopics. We do so by implementing a hierarchical cluster analysis based on the word distribution associated with each topic. The resulting metatopics contain topics with similar word distributions, thereby capturing similar news. When we plot metatopic attentions through time, we find that they feature bell-shaped dynamics. Each metatopic is elevated at a distinct period of time while its frequency is roughly zero elsewhere, resulting in six non-overlapping narratives.

Having reduced the number of topics to six metatopics, we turn to their interpretation. To do so, we exploit both the word and the time series dimension of each metatopic. Starting with the word distributions, we ask ChatGPT to assign a title to each metatopic and to the topics it maps into. In addition, we construct the word clouds for each metatopic. Both the metatopic labels and the word clouds give us some idea of the narrative that each metatopic captures. We observe that this is one key advantage of constructing latent variables from text data, as opposed to working with factors generated from available time series. Next, we examine the relationship between metatopic attention series and several business cycles and policy indicators. We find that four out of six metatopics explain a significant fraction of the business cycle, as well as of U.S. monetary and fiscal policy. Since these metatopics also explain exchange rate fluctuations, they point to a connection between the exchange rate and economic fundamentals. On the other hand, the remaining two metatopics are more closely tied to financial market outcomes, such as stock prices and the VIX, and less connected to fundamentals. Although metatopics are linked to time series, they are only salient during specific periods. Thus, our findings suggest significant non-linearity in the relationship between time series aggregates and exchange rates, resulting in the disconnect documented in the literature.

Taken together, the word analysis and the time series analysis point to the following narrative decomposition of the exchange rate. From 1976 to the late 80s, exchange rate changes were

primarily due to the fiscal and monetary policy of the U.S. The metatopic labeled “Historical Economic Policies and Energy Markets” loads heavily on words such as {Carter, Reagan, oil, debt}, and explains changes in the U.S. government debt, the Federal Funds rate, and the unemployment rate. The metatopic “Global Finance and Monetary Policies” loads on words such as {report, rates, prices, international}, and explains the rise of the Federal Funds rate, and the dollar appreciation during the Volcker-Greenspan era. Similarly, the dollar patterns of the 90s are also associated with changes in macroeconomic fundamentals. However, they are not linked to news on fiscal or monetary policy, instead, news on technological advancements appears more prominent. The metatopic “Technological Advancements and Financial Reporting” peaks during that time and contributes to the rise of industrial production and the fall of the unemployment rate. During the early 2000s, we estimate a tighter relationship between news about financial markets and changes in the exchange rate. Metatopics including “Corporate Transactions and Stock Market”, and “Economic Trends and Corporate Strategy,” become more relevant. These metatopics largely explain the stock price behavior during the dot-com bubble, pointing to a connection between the exchange rate and equity prices. Finally, we find that the metatopic “International Business and Financial News” explains the bulk of exchange rate changes of the last decade. This metatopic loads on words such as {business, duty, exchange, Trump}, particularly explains the dollar appreciation during the European sovereign debt crisis, as well as changes in the unemployment rate.

Motivated by the scapegoat theory of [Bacchetta and van Wincoop \(2004, 2013\)](#), in the last part of the paper, we offer an explanation to the unstable relationship between exchange rate and macroeconomic indicators. According to the theory, when investors are unable to observe certain exchange rate-relevant economic indicators, they assign varying levels of importance to the indicators they can observe — such as government deficits, technological changes, or shifts in equity prices — essentially making these factors scapegoats for explaining exchange rate fluctuations. To test this theory, we run text-augmented linear regressions proxying investors’ importance weights with metatopic attention series. We find large increases in adjusted R^2 relative to the model without metatopic attentions, with values rising from 1.2% to 7.3% in the regression predicting monthly exchange rate changes, and from 7% to 28% in the

regression predicting annual changes in the monthly exchange rate. These results underscore the value of incorporating investor attention into exchange rate models, and offer a more nuanced understanding of the factors driving exchange rate fluctuations.

Literature Review This paper relates to three strands of the literature: (1) the use of text as data, (2) the measurement and the importance of narratives for economic fluctuations, and (3) the broad literature on exchange rate determination.

[Gentzkow et al. \(2019\)](#), and [Ash and Hansen \(2023\)](#), review the literature on the language models applied to economics. Within this class, topic models have only recently begun to be explored. [Hansen et al. \(2018\)](#), the first paper estimating a topic model, uses LDA on FOMC statements, [Larsen and Thorsrud \(2019\)](#) and [Thorsrud \(2020\)](#) apply LDA to Norwegian news data and analyze macroeconomic forecasting models, [Bybee et al. \(forthcoming\)](#) estimates LDA on WSJ articles. To the best of our knowledge, sLDA has not been applied in the context of economic analysis. Besides applying sLDA, we provide an estimation routine that enables the extraction of relevant information from news sources to explain a given time series of interest.¹

In *Narrative Economics*, [Shiller \(2017\)](#) argues that contagious narratives can drive macroeconomic events. Since then, the literature has sought to estimate narratives from text data. For example, [Andre et al. \(2023\)](#) conducts open-ended surveys to estimate narratives underlying inflation, [Goetzmann et al. \(2022\)](#) measures narratives about financial crashes in news media, and [Macaulay and Song \(2023\)](#) measures how news coverage of narratives affects sentiment on social media. We complement this literature by providing a narrative decomposition of exchange rate fluctuations.

Our decomposition sheds light on the determinants of the exchange rate. A recent summary of this literature is in the Handbook chapter of [Maggiori \(2022\)](#), and new promising evidence of a connection with fundamentals is [Engel and Wu \(2024\)](#). Closely related to our work is the idea that monetary policy shapes exchange rates, as pointed out by [Dornbusch \(1976\)](#) and [Frankel](#)

¹Complementary research incorporates news text into macro-finance analyses by relying on carefully selected researcher inputs instead of statistical models. For instance, [Chahrour et al. \(2021\)](#) examines the impact of newspapers' sectoral coverage on macroeconomic activity. Instead of analyzing the article text directly, they associate article tags from Factiva with firms and their sectors to quantify news coverage. Similarly, [Baker et al. \(2016\)](#) constructs indices of economic policy uncertainty from news articles by counting the frequency of keywords curated by researchers.

(1979), or more recently by Schmitt-Grohé and Uribe (2022). Furthermore, we also find that fiscal policy contributed to the dollar movements of the late 70s and early 80s, consistent with the works examining the exchange rate response to fiscal policy, such as Ravn et al. (2012), Kim and Roubini (2008), Monacelli and Perotti (2010), and Jiang (2021). More generally, the idea that the exchange rate reacts to news has been emphasized in several studies. Engel and Frankel (1984), Anderson et al. (2003), Stavrakeva and Tang (2020) examine how announcements of measures of economic aggregates affect the exchange rate. Chahrour et al. (2024) shows that noisy news about U.S. TFP explains a large fraction of exchange rate fluctuations. In addition to providing evidence supporting the aforementioned literature, we find that the relation between exchange rate and economic fundamentals is nonlinear (see Rossi, 2013, for a review), and the scapegoat theory of Bacchetta and van Wincoop (2004, 2013) can be a plausible explanation.

1 Topic Model and Estimation

This section describes the supervised LDA model and the data filtering procedure.

1.1 Supervised LDA

Latent Dirichlet Allocation (LDA), introduced by Blei et al. (2003), is a popular dimensionality reduction technique widely applied in various fields such as political science and psychology. LDA reduces documents to a random mixture of latent topics, where each topic k is characterized by a probability vector $\beta_k \in \Delta^{V-1}$ over the V unique terms in the dataset. Each document d can be summarized by a topic distribution $\theta_d \in \Delta^{K-1}$, where K is the number of topics. The elements of θ_d are the topic proportions $\theta_{k,d}$, with $k = 1, \dots, K$, which represent the contributions of each topic k to document d , such that $\sum_{k=1}^K \theta_{k,d} = 1$. This allows for a dimensionality reduction, transforming each document from the original V -dimensional space of words to a lower-dimensional K -topic space.

A key feature of LDA is its assumption of Dirichlet priors on both the topic-word distributions β_k and the document-topic distributions θ_d . Specifically, $\beta_k \sim \text{Dir}(\eta)$ and $\theta_d \sim \text{Dir}(\alpha)$, where η and α are hyperparameters that govern the concentration of topics and words. A higher

concentration leads to documents (or topics) dominated by a small number of topics (or words), whereas a lower concentration allows for a broader mixture. The core inference problem in LDA is approximating the posterior distributions over β_k and θ_d , given 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 in tasks where an external variable needs to be forecasted, such as document classification or regression tasks. [Blei and McAuliffe \(2007\)](#) shows 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.

In contrast, LDA's topics are purely focused on representing the data's internal structure (i.e., the word distributions) and are not guaranteed to be relevant for any external prediction task. This makes LDA less effective for applications where the goal is not just to discover latent topics but also to predict an outcome variable. Since sLDA directly models the relationship between topics and the response variable, it better captures the relevant patterns needed for accurate predictions.

Appendix [A](#) describes the estimation algorithm. In our analysis, the response variable is the monthly trade-weighted real exchange rate between the U.S. and the other G7 countries. To determine the number of topics, we estimate several models with the topic number ranging from 20 to 220 and evaluate them based on interpretability measures. Figure [12](#) in Appendix [A](#) shows that a value of 180 strikes a good balance between semantic coherence, which measures

the degree to which words within a topic frequently co-occur across documents, and exclusivity, which quantifies the uniqueness of words to a given topic.

1.2 *The Wall Street Journal Data Set*

We leverage a comprehensive text corpus comprising all article titles 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.

We prioritize article titles over full texts, observing that titles encapsulate the primary message of a news article succinctly. To enhance the quality of our topics, we automate the elimination of words that could obscure the interpretation. This includes stopwords and specific names including countries, cities, organizations, or individuals. The refined corpus comprises 707,984 news titles and a vocabulary of 14,995 unique words. A detailed description of the text collection and pre-processing steps is in Appendix B.

Before sLDA model estimation, we convert the raw text into a data feature matrix (DFM) or a bag-of-words representation. In this matrix, each row corresponds to a distinct article title, columns are pegged to unique words from the corpus, and individual elements denote word frequencies within respective titles. Given that we use monthly real exchange rates, we aggregate news titles to monthly frequency, combining all titles published within a specific month into a singular row of the DFM. We do so because retaining an article title-level dataset (with a constant monthly real exchange rate value for each title) would inaccurately imply a potential direct relationship between *every* title and the exchange rate, whereas our assumption posits this relationship to be pertinent to a subset of articles from that month.

2 Topic Performance and Selection

This section gathers results from the sLDA model, estimates Random Forest models using topics, and compares the performance of the sLDA with the one of standard macroeconomic indicators.

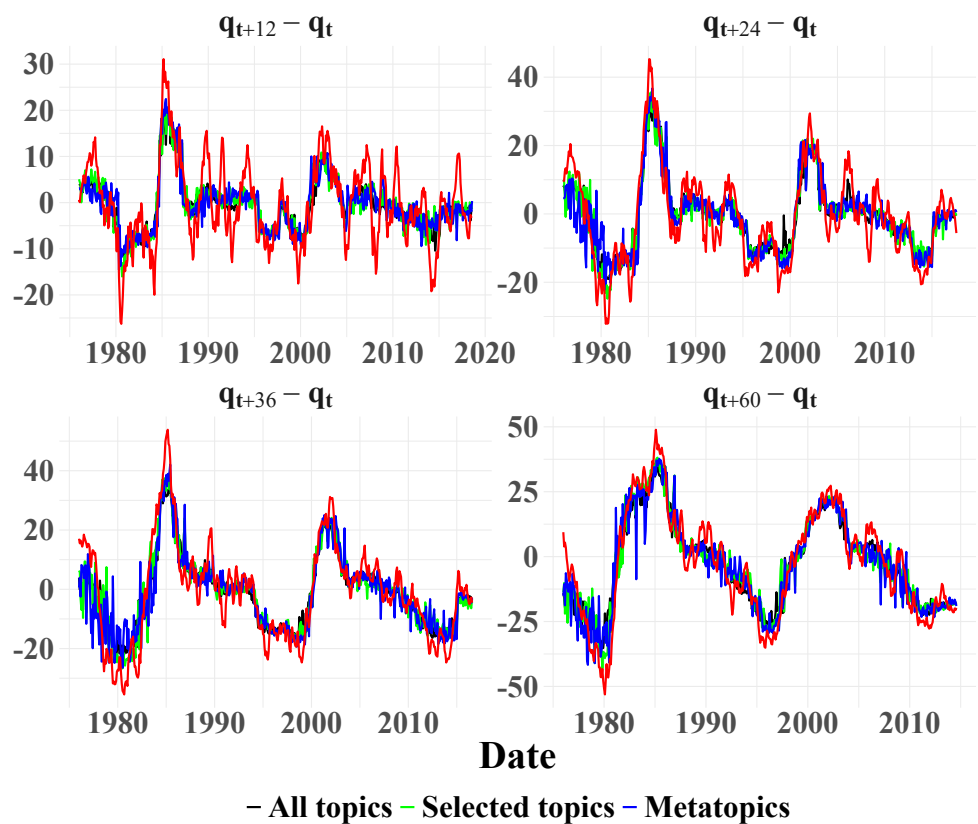
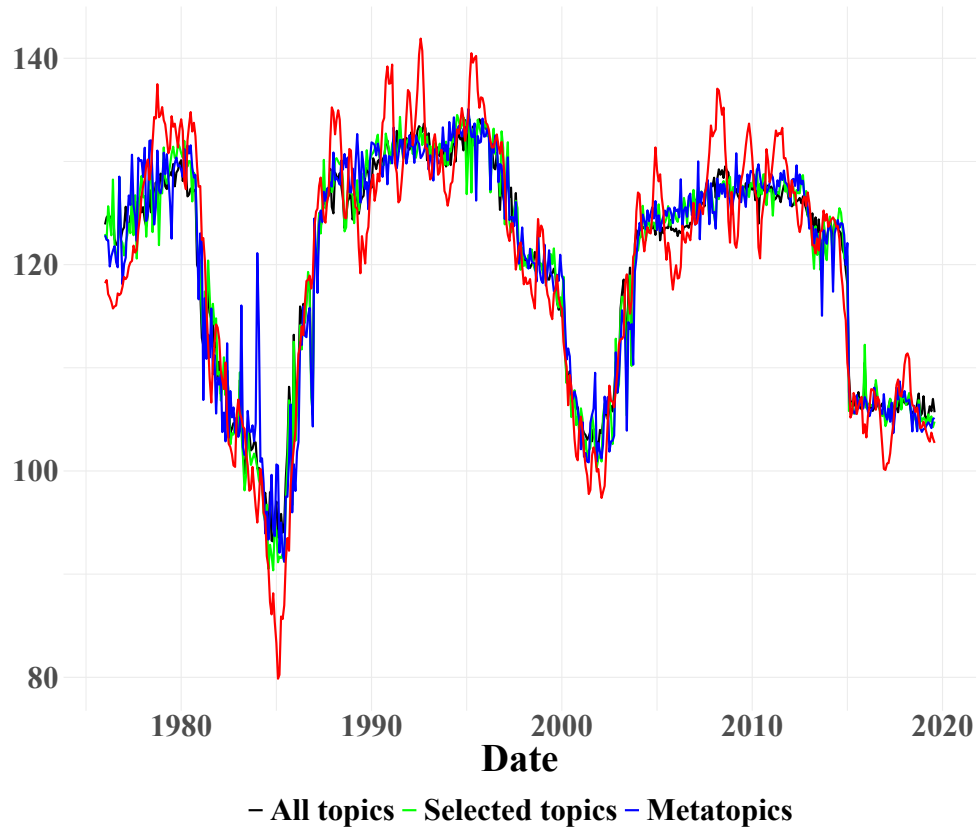


Figure 1: Actual and predicted exchange rate values from Random Forest models

Note: The figure shows the prediction of the Random Forests in (1) using all topic proportions (black line), sixteen topic proportions (green line), and six metatopic proportions (blue line).

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

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

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.

2.1 Random Forest Estimation

We evaluate the explanatory power of the topic model using 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 on a set of explanatory variables X_t . In our analysis, we estimate the model as follows:

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

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

where q_t is the natural logarithm of the monthly real exchange rate between the U.S. and the other G7 countries. 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,

²Since we pool monthly titles together before estimating the topic model, the estimated monthly topic proportions $\hat{\theta}_{k,t}$ are equal to the estimated topic proportions $\hat{\theta}_{k,d}$, where d is the set of titles in a given month.

which consist of data points not included in the training of a specific tree.

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 of the exchange rate at 12, 24, 36, and 60 month horizons. The OOB R^2 is the average R^2 from 5000 trees.³ The table reveals that topic proportions explain a significant portion of the exchange rate fluctuations, with OOB R-squared values ranging from 61% to 88%.

Figure 1 illustrates predicted and actual exchange rate values for the model specifications in (1). 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 dollar movements are more likely to be reported and thus captured by the model.

2.2 Comparison with macroeconomic indicators

We seek to compare the performance of the topic model against the performance of standard macroeconomic indicators. To do so, we use the monthly dataset constructed by McCracken and Ng (2016), comprising many U.S. macroeconomic indicators. After excluding exchange rates from the variable set, we end up with 120 time series. We transform the variables to ensure stationarity, and estimate two sets of Random Forest regressions. The first set includes only the macroeconomic indicators, whereas the second set augments the model with the topic series.

The second and third columns of Table 1 reveal that the R^2 s from using the macroeconomic indicators are smaller than the ones produced by topics. Moreover, adding topics to the model leads to a non-trivial increase in the R^2 s. For instance, the R^2 on the exchange rate rises from 79% to 90%.

Figure 2 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 are new housing permits divided by region.

³We 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.

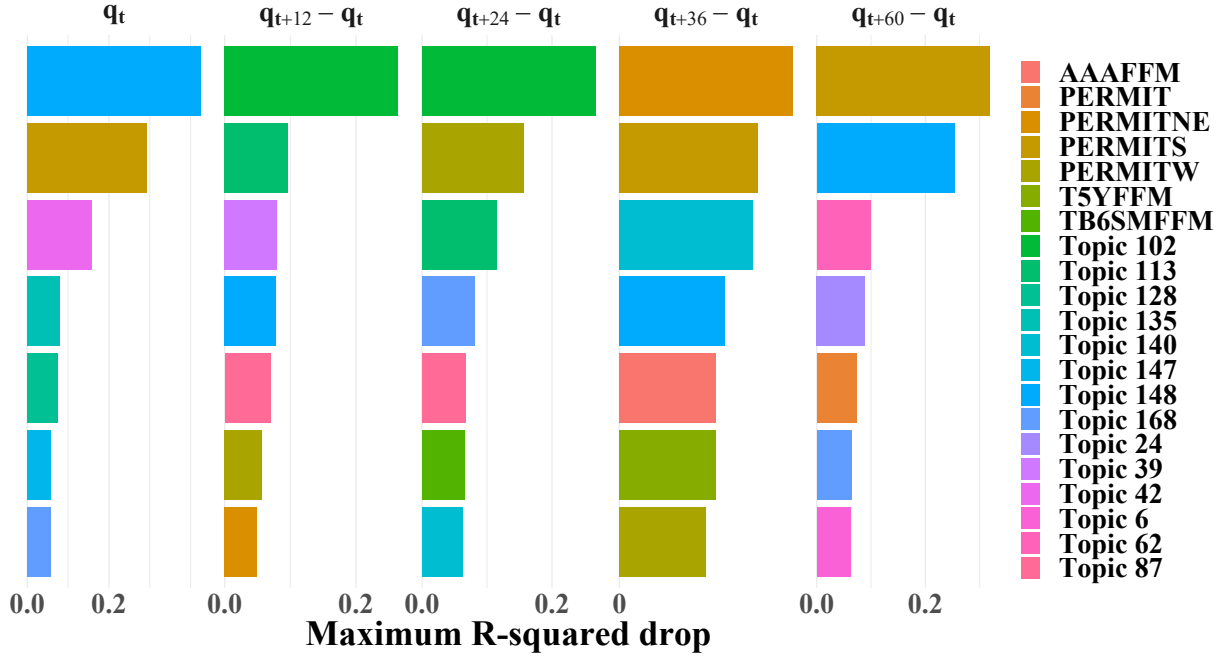


Figure 2: Predictor importance ranking from Random Forest models

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 Minus Federal Funds Rate, PERMIT is the number of new housing permits in the U.S., PERMITNE, PERMITS, and PERMITW are the number 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, TB6SMFFM is the spread between 6-month Treasury and the Federal Funds Rate

Overall, results suggest that topics span a larger amount of exchange rate fluctuations than time series. Next, we are going to provide a parsimonious procedure to interpret these topics.

3 Topic Interpretation

In this Section, we interpret the outcome of the topic model. Examining 180 topics is impracticable. Thus, we reduce the number of topics in two steps. In the first step, we select those topics with the highest explanatory power. We choose a number of topics such that the goodness of fit is not compromised. Second, we cluster topics with similar word distributions. We end up with six metatopic series, which we then examine in relation to both the exchange rate and other macroeconomic indicators.



Figure 3: Dendrogram

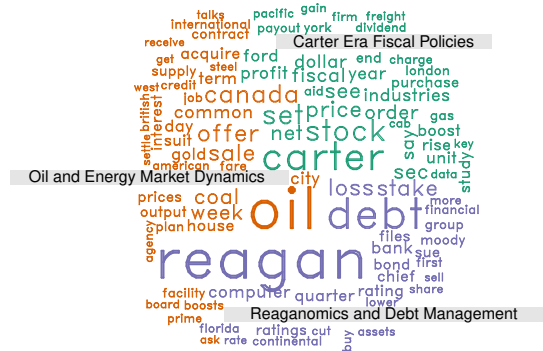
Note: The figure illustrates the outcome of a cluster analysis on the word distributions of the topics. Topic and metatopic titles are constructed using ChatGPT4. The prompt and the ChatGPT-generated metatopic descriptions are in Appendix C.

3.1 Topic selection and clustering

We select the top seven topics with the largest R^2 drop in each specification, resulting in sixteen distinct topics. We run the Random Forest regressions in (1) with the sixteen topics only, and report the values in column (iv) of Table 1. Reassuringly, we observe no drop (up to the second digit) in the OOB R^2 of the exchange rate regression, and a very minimal drop in the other specifications.



MT1: Global Finance and Monetary Policies



MT2: Historical Economic Policies and Energy Markets



MT3: International Business and Financial News



MT4: Corporate Transactions and Stock Market



MT5: Economic Trends and Corporate Strategy



MT6: Technological Advancements and Financial Reporting

Figure 4: 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 with the maximum deviation of its topic-word probability from the average topic-word probability across the topics in the same metatopic.

Next, we further reduce the topic number by clustering 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 Figure 3. Figure 3 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 titles and the aggregation nodes. We group topics into six clusters of topics, or metatopics. These metatopics include a range of one to four topics, depending on the similarities between word distributions.

Figure 4 displays the word clouds for each metatopic, emphasizing the most prominent words within each topic. Both the dendrogram and word clouds illustrate that the six metatopics are quite distinct from one another, with some being easier to interpret than others. For example, MT1: Global Finance and Monetary Policies, and MT2: Historical Economic Policies and Energy Markets, are composed of topics associated with monetary and fiscal policies, while MT4: Corporate Transactions and Stock Market, and MT5: Economic Trends and Corporate Strategy, relate more to stock market dynamics, and MT6: Technological Advancements and Financial Reporting, appears focused on technological advancements. On the other hand, MT3, labeled "International Business and Financial News," is less straightforward. Although it includes terms such as {business, Trump, banking, exchange}, it also encompasses words like {duty, art, mansion, drinking}, making its interpretation more challenging.

3.2 A narrative decomposition of the exchange rate

After clustering topics into metatopics, we can study their time series properties. We compute metatopic attention series as the weighted sum of their estimated topic proportions, where the weights are the topic's average R^2 drops across the specifications in (1). Doing so reduces the loss of explanatory power 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.

Figure 5 shows metatopic attentions over time. Three patterns emerge. First, each metatopic series behaves like an epidemic curve, lasting about five to fifteen years and closely resembling the properties of narratives discussed in Shiller (2020). Second, each metatopic reaches its peak attention at a different time, with minimal overlap between them. This suggests that metatopics represent distinct, individual events rather than a continuous stream of news related to the same

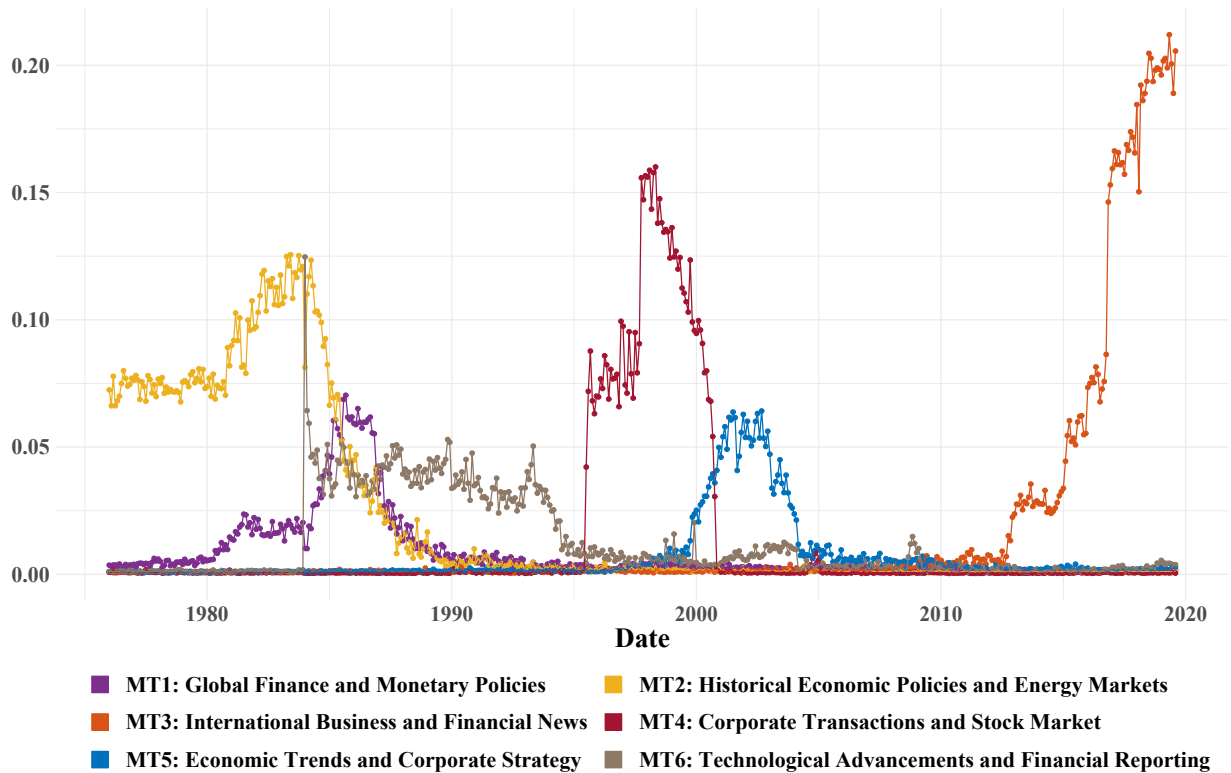


Figure 5: Weighted metatopic attention series

Note: The figure plots the six metatopic attention series, each constructed as the weighted sum of topic proportions. For each metatopic, comprising T topics, we compute a topic weight as the topic's average R^2 drop across the specifications in (1), divided by the sum of the average R^2 drops of all T topics.

economic theme. Third, and perhaps surprisingly, we do not observe any attention peak during the years around the Great Recession.

Altogether, Figure 5 reveals that a few narratives explain exchange rate fluctuations, except during the Great Recession, which given (i) the potentially highly nonlinear mapping between metatopic attentions and the exchange rate, and (ii) the fact that topic proportions sum to one, the absence of a metatopic peak indicates the presence of a multitude of narratives explaining exchange rate fluctuations around those years.

Metatopics and the exchange rate Figures 6 and 7 plot the prediction of the (demeaned) real exchange rate from the Random Forest model estimated with the metatopic series, along with the metatopic contributions to the predicted value.⁴ A few observations are in order. First, at any

⁴We 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 determine the total contribution of each metatopic to the model's prediction.

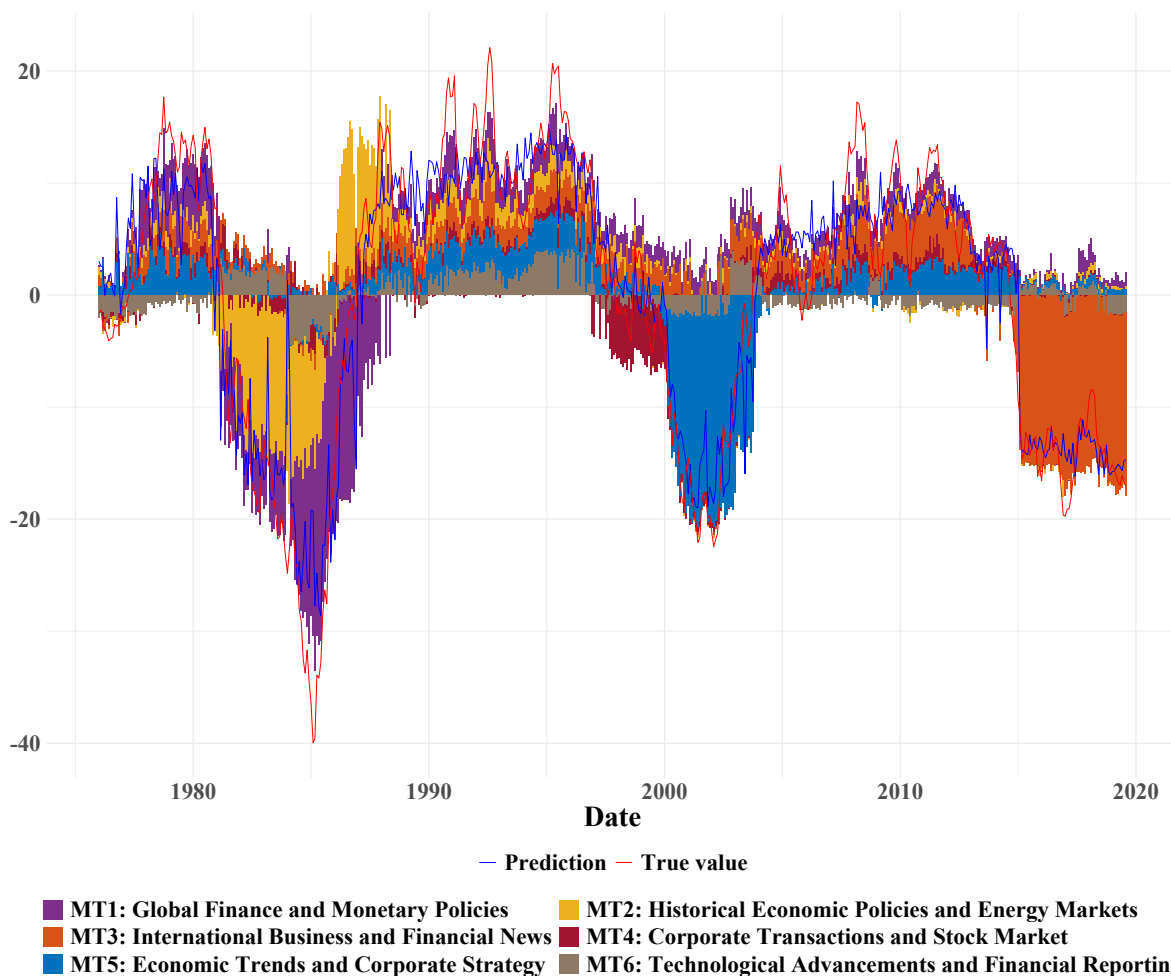


Figure 6: Metatopic predictions and contributions to the real exchange rate

Note: The figure shows the predicted value (blue line) and the actual log of the real exchange rate (red line) from a Random Forest model using six metatopic proportions as predictors. 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.

given time, there is usually one, or at most two metatopics explaining exchange rate changes. For instance, the exchange rate appreciation of the early 80s is mostly explained by MT1 and MT2, associated with U.S. fiscal and monetary policies. The exchange rate fluctuations during the dot-com bubble are mostly driven by MT4 and MT5, metatopics about stock price fluctuations. Perhaps an exception is MT6, which spans exchange rate fluctuations for a greater part of the sample, consistent with recent evidence by [Chahrour et al. \(2024\)](#) showing that changes in expectations of future TFP drive a significant portion of exchange rate fluctuations. Second, results are stable across the various prediction horizons that we consider. This is reassuring since the real exchange rate series is highly persistent and our findings might be resulting from spurious relationships. A notable difference between Figures 6 and 7 lies in the contribution

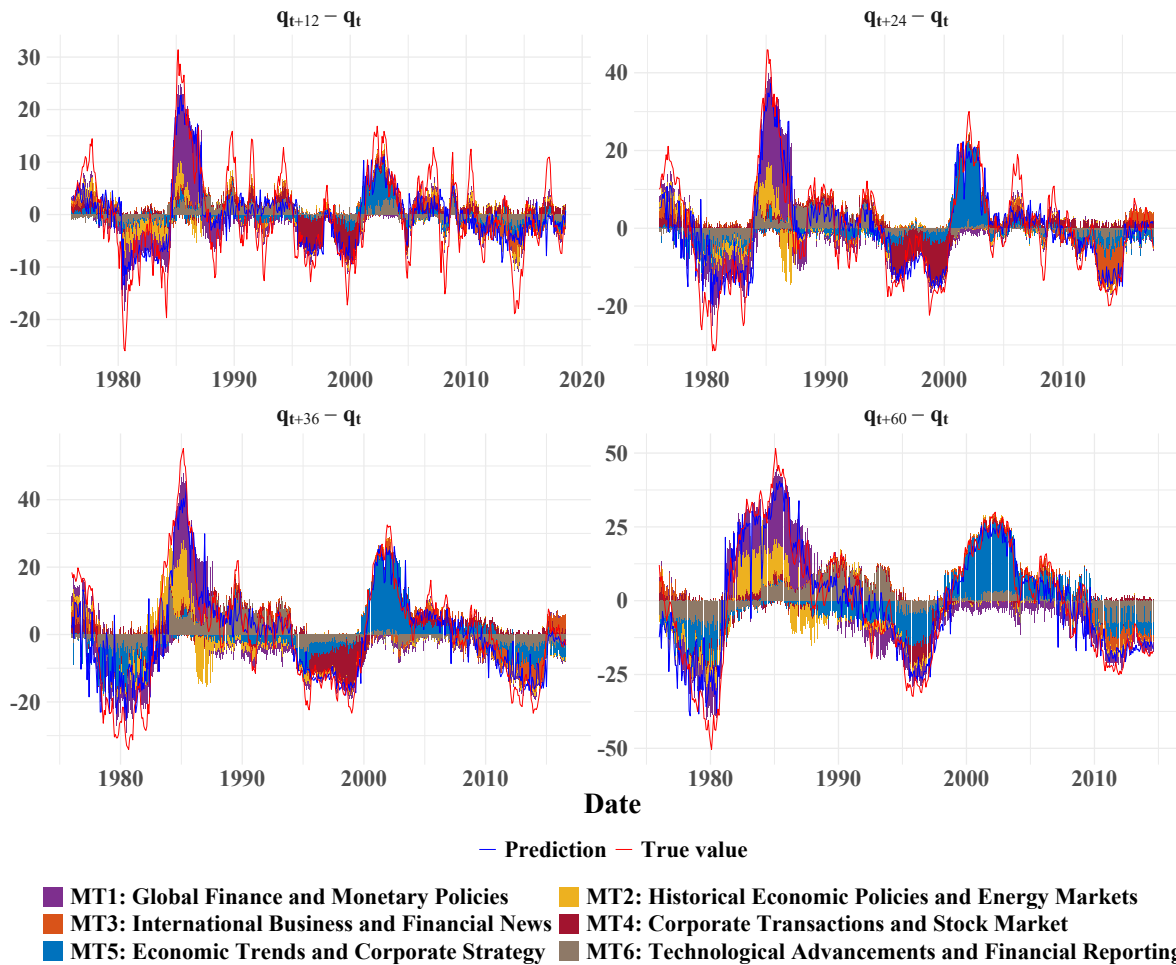


Figure 7: Metatopic predictions and contributions to real exchange rate changes

Note: The figure shows the predicted value (blue line) and the actual log change of the real exchange rate (red line) from Random Forest models using six metatopic proportions as predictors. 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.

of MT3. In the contemporaneous regression (Figure 6), MT3 plays a significant role, but it becomes far less important when predicting future exchange rate movements (Figure 7). One possible explanation is that MT3 may be spuriously related to the exchange rate, resulting in an unstable statistical association. This would align with our findings from the word analysis, where MT3 lacked a clear economic interpretation.

Metatopics and macroeconomic indicators The word analysis reveals that the metatopic word distributions are different from each other, suggesting that metatopics identify distinct narratives. The time series analysis shows that these narratives explain a large portion of exchange rate fluctuations. Thus, metatopic attentions might be associated with different

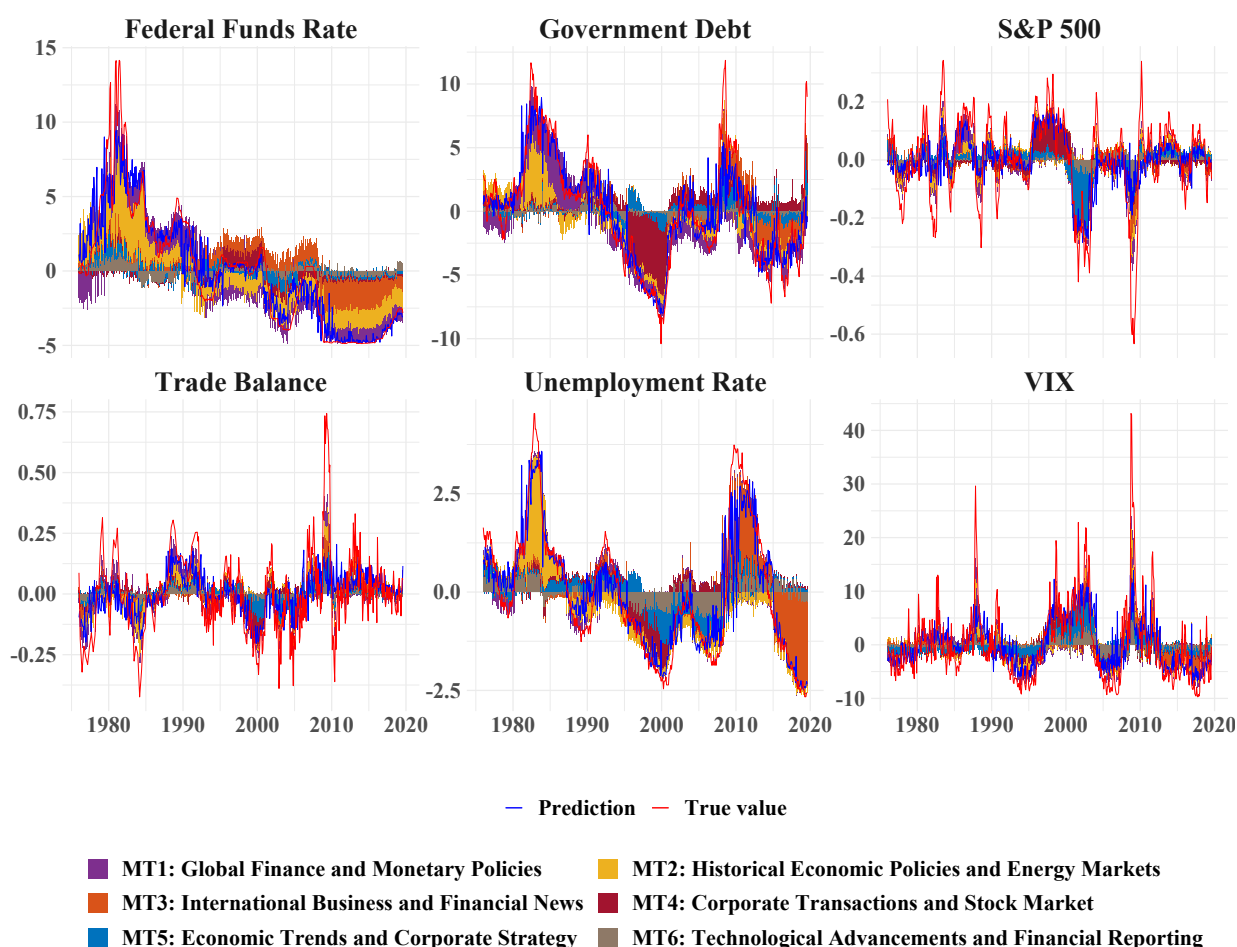


Figure 8: Metatopic predictions of monthly indicators

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.

sources of exchange rate fluctuations, some related to fundamentals such as changes in fiscal and monetary policy, while others associated with news on financial market changes not necessarily driven by fundamentals. To gain further insights, we estimate the relation between metatopic attentions and several macroeconomic indicators. We expect metatopics to be differently related to macroeconomic time series in such a way consistent with the outcomes of the word analysis.

We consider the following six monthly macroeconomic indicators: the Federal Funds rate, the future annual growth in the real government debt, the annual growth rate of the S&P 500, the annual change in the trade balance as a fraction of GDP, the unemployment rate, and the

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

Table 2: OOB R^2 from Random Forests using metatopics proportions as predictors

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 \(2012\)](#). All series are demeaned.

VIX. We estimate six Random Forest models with the macroeconomic variable of interest as a dependent variable, and metatopics as predictors.⁵

Figure 8 shows the metatopic predictions and contributions, whereas Table 2 reports the OOB R^2 s. Metatopics appear to explain a sizeable fraction of the fluctuations in the indicators we consider, with important heterogeneity across indicators and over time. To further inspect the degree of heterogeneity, Figure 9 shows metatopic importances divided by decade, where we compute metatopic importances as the Mean Squared Error losses obtained after excluding a metatopic from the estimated model.

Results from the time series analysis support the interpretation of the word analysis. Overall, we find a substantial degree of heterogeneity among the relations between metatopics and the macroeconomic aggregates, indicating that metatopics capture distinct narratives. Starting from MT1: Global Finance and Monetary Policies and MT2: Historical Economic Policies and Market Dynamics, these metatopics explain a large fraction of the Federal Funds Rate, Unemployment rate, and government debt during the 80s, suggesting that they capture news related to monetary and fiscal policy during that period. The primary predictors of stock prices are MT4: Corporate Transactions and Stock Market and MT5: Economic Trends and Corporate Strategy, both explaining the boom and bust of the stock market around the early 2000s. MT3:

⁵We take the *future* annual growth in the government debt as it is most consistent with the forward-looking nature of news. In addition, we construct the trade balance to GDP ratio by linearly interpolating the log of quarterly GDP.

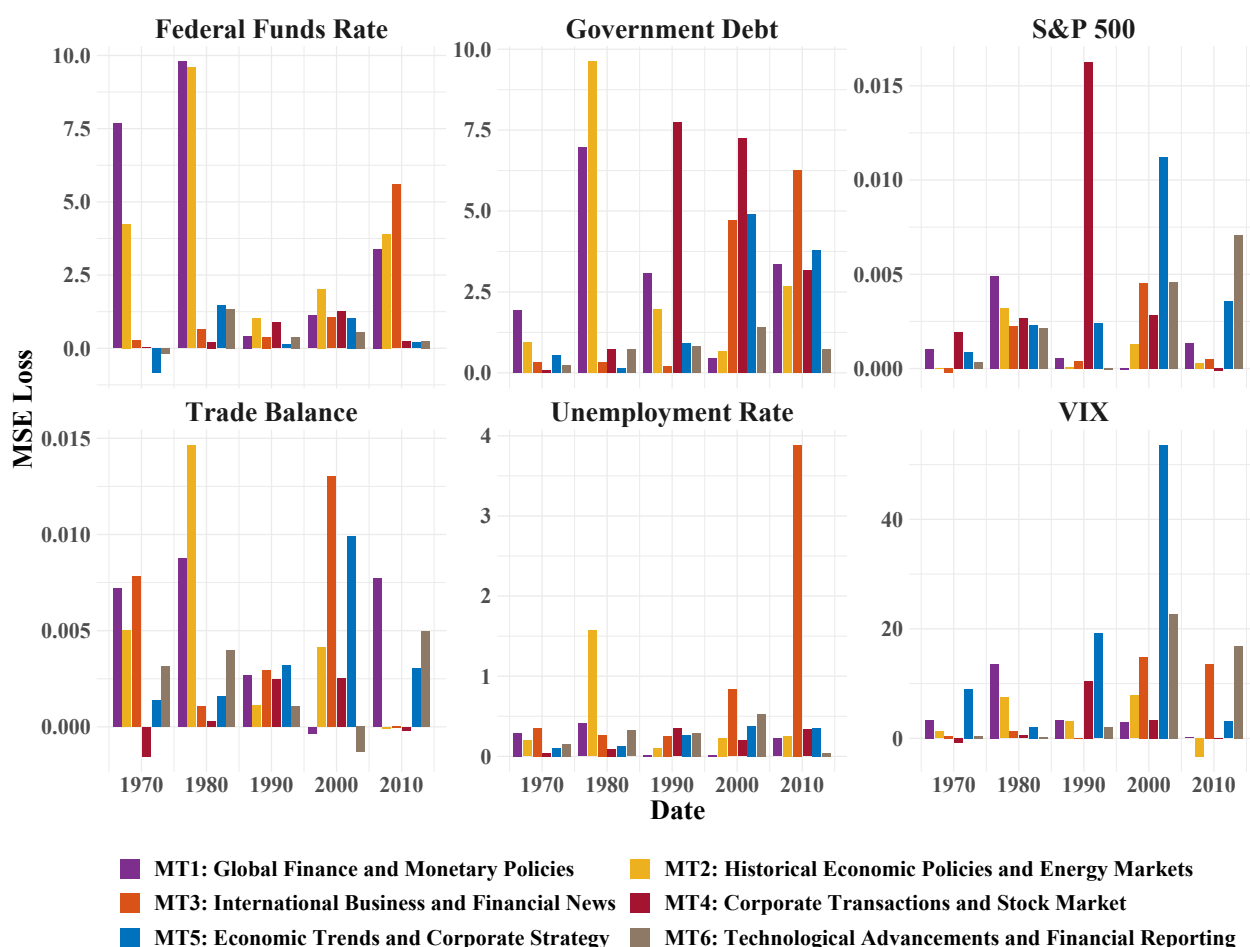


Figure 9: Metatopic Mean Squared Error losses divided by decade

Note: The figure shows the Mean Squared Error losses of excluding a metatopic from the prediction of a Random Forest. The MSE losses are computed over a rolling sample including each decade 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.

International Business and Financial News predicts the pattern of the Federal Funds Rate and economic activity in the latter part of the sample, as well as fluctuations in the VIX. This suggests that MT3 may capture news related to broader economic developments, including geopolitical events that impact economic activity, beyond just financial markets.

Finally, the word analysis revealed that MT6: Technology Advancements and Financial Reporting might be related to changes in total factor productivity. The monthly specifications in Figure 8 show that MT6 predicts the unemployment rate throughout the sample, and during the 90s, especially. Unfortunately, U.S. measures of TFP are only available at a quarterly frequency. Therefore, we aggregate the monthly metatopic series to quarterly, and estimate a

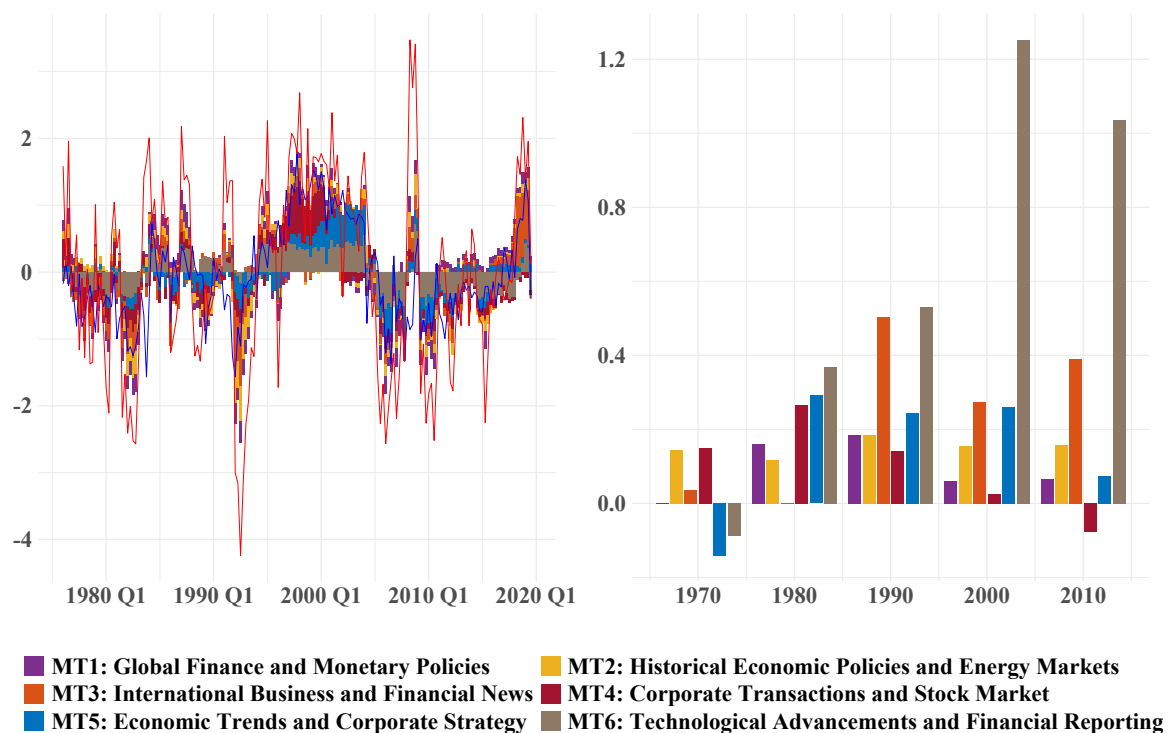


Figure 10: Metatopic predictions, contributions and MSE losses of quarterly TFP

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 to 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.

Random Forest model with quarterly TFP. More specifically, we estimate the future annual log change in the quarterly utilization-adjusted TFP constructed by Fernald (2012). We focus on future changes because news articles are likely to report about changes in TFP that have not yet materialized. Figure 10 presents the results, confirming that MT6 predicts changes in TFP. Unlike the more isolated contributions of other metatopics, MT6 appears to consistently explain variations in TFP across almost the entire sample.

4 A narrative-based explanation of the disconnect

Despite the descriptive nature of our analysis, our findings provide new insights into the sources of exchange rate fluctuations and the documented disconnect between exchange rates and macroeconomic fundamentals. When considered in isolation, some of our results, are consistent with previous studies linking exchange rates to macroeconomic or financial indicators. Examples of these studies are [Jiang \(2021\)](#), showing that the U.S. fiscal cycle predicts currency returns, [Dornbusch \(1976\)](#), emphasizing the role of monetary policy, [Chahrour et al. \(2024\)](#), highlighting the impact of productivity news, [Cormun and De Leo \(2024\)](#), which focuses on the relationship between exchange rates and the VIX, and [Brandt et al. \(2006\)](#), examining equity prices as source of exchange rate fluctuations. However, the relationship between these sources and exchange rate changes appears to be weak or unstable. In this section, we argue that one reason could be the influence of the viral spread of source-related narratives, and offer a narrative-based explanation of the exchange rate disconnect.

4.1 Back to the fundamentals

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). Considering the evidence of nonlinearity presented earlier, and the duration of the narratives, we use a rolling window of 120 months for our regressions.

Figure 11 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 reveals significant parameter instability. Notably, the adjusted R^2 values from the rolling window regressions are generally much higher than those from the

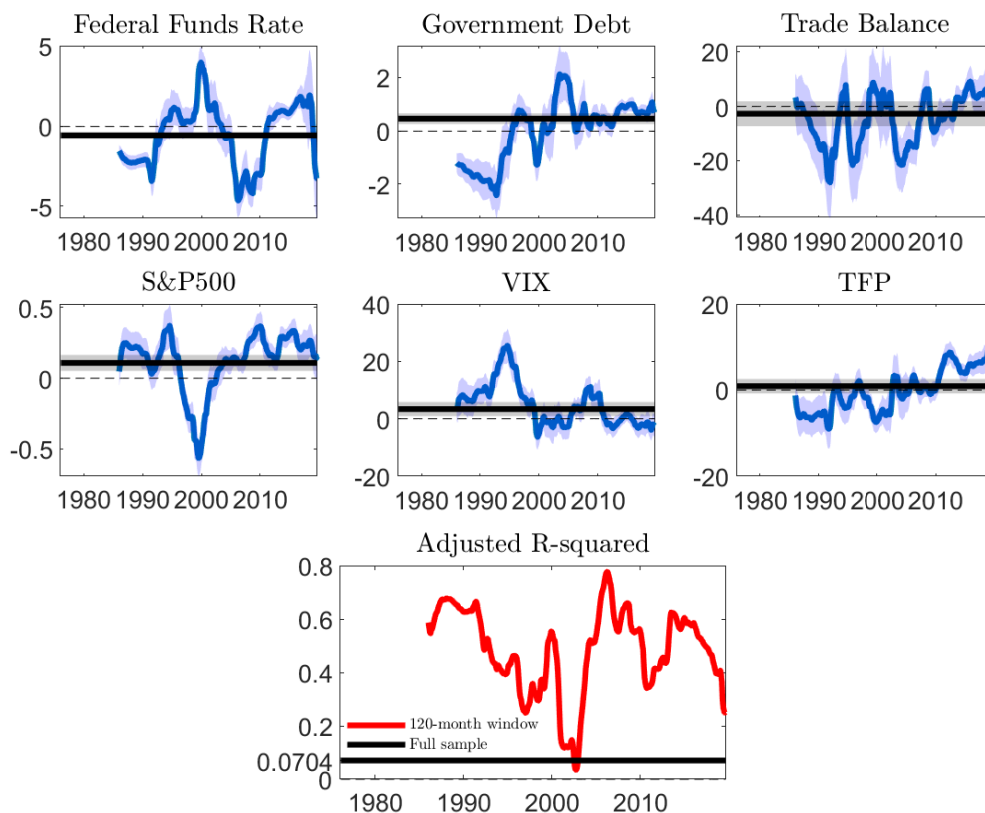


Figure 11: Rolling window regression coefficients and adjusted R^2 values

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 estimate. The shaded areas are 95% confidence intervals.

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 Powerful text-augmented regressions

Next, we use the scapegoat theory of [Bacchetta and van Wincoop \(2004, 2013\)](#) to rationalize how media coverage can give rise to the unstable relation between the exchange rate and macroe-

conomic indicators documented above. According to the theory, at times, some macroeconomic factors receive unusually large weight and thus become scapegoats of exchange rate movements. [Bacchetta and van Wincoop \(2013\)](#) shows that in a simple model with unknown and constant parameters, these scapegoats arise because of agents' rational confusion as they make inference on the true parameters of the model conditioning on observables.

Starting with a standard present value equation for the exchange rate, the change in the real exchange rate can be expressed as

$$\Delta q_t \approx (1 - \lambda)\mathbf{f}_t'\beta + (1 - \lambda)b_t + \lambda\mathbf{f}_t'E_t\beta \quad (2)$$

where \mathbf{f}_t is the vector of observable indicators, b_t is one or more unobservable indicators, and β is the vector of structural parameters. The last term includes $E_t\beta$, which is agents' expected value of the structural parameters and arises from the unknown parameters assumption. It captures the scapegoat effects, as it represents the importance that investors attach to different macroeconomic factors to explain exchange rates at a given time.⁶

Note that since λ is the discount factor and is close to one, the first two terms are close to zero, and the effect of changes in fundamentals depends mainly on agents' expected parameter values and how these expectations change with the observable indicators.

[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 40-60 financial market participants and ask them to rank the importance of six macroeconomic variables as drivers of a country's exchange rate. They use the results of the survey as proxies for the scapegoat parameters $E_t\beta$ and find evidence supporting the theory.

We estimate our version of the scapegoat test by proxying the vector of scapegoat parameters with metatopic attention series over our sample period. We conjecture that metatopic attention series could too serve as proxies for the importance that investors give to the factors driving the exchange rate. While there could be several interesting ways to rationalize this, a natural one is to assume that changes in the factors \mathbf{f}_t become more newsworthy when investors attach a

⁶[Bacchetta and van Wincoop \(2013\)](#) shows that innovations in the unobservable factors affect $E_t\beta$ resulting in "excessive weight" attributed to observable indicators.

Variable	$q_t - q_{t-1}$		$q_t - q_{t-3}$		$q_t - q_{t-12}$	
	(I)	(II)	(I)	(II)	(I)	(II)
(Intercept)	-0.029 [-0.115, 0.057]	0.276** [0.169, 0.384]	-0.079 [-0.255, 0.097]	0.750*** [0.534, 0.967]	-0.270 [-0.650, 0.110]	2.904*** [2.458, 3.351]
Federal Funds rate	-0.206** [-0.295, -0.117]	0.274* [0.125, 0.424]	-0.631*** [-0.813, -0.450]	0.731** [0.430, 1.032]	-2.225*** [-2.643, -1.807]	1.805*** [1.175, 2.435]
Government Debt	0.223** [0.133, 0.313]	0.332*** [0.228, 0.435]	0.641*** [0.457, 0.824]	0.903*** [0.695, 1.111]	1.997*** [1.567, 2.427]	3.926*** [3.435, 4.416]
Trade Balance	-0.124 [-0.211, -0.036]	-0.178** [-0.267, -0.089]	-0.214 [-0.392, -0.036]	-0.455** [-0.635, -0.275]	-0.458 [-0.854, -0.062]	-1.749*** [-2.148, -1.351]
S&P 500	-0.065 [-0.168, 0.039]	0.055 [-0.063, 0.172]	0.307 [0.094, 0.520]	0.618*** [0.381, 0.856]	1.646*** [1.213, 2.079]	2.579*** [2.127, 3.032]
VIX	0.019 [-0.083, 0.121]	0.013 [-0.097, 0.123]	0.277 [0.068, 0.486]	0.335 [0.113, 0.558]	1.141*** [0.724, 1.558]	1.661*** [1.243, 2.080]
TFP (Interpolated)	0.076 [-0.013, 0.165]	-0.226** [-0.338, -0.115]	0.096 [-0.086, 0.277]	-0.478** [-0.702, -0.253]	0.417 [0.022, 0.811]	0.513 [0.056, 0.971]
(MT1+MT2) × FFR	-	-7.398*** [-9.557, -5.240]	-	-21.532*** [-25.873, -17.191]	-	-47.254*** [-57.114, -37.394]
MT2 × Government Debt	-	-4.841** [-7.215, -2.468]	-	-8.966* [-13.734, -4.199]	-	-63.583*** [-74.066, -53.101]
MT3 × Trade Balance	-	-0.388 [-2.939, 2.163]	-	10.000* [4.876, 15.124]	-	32.428 [10.952, 53.905]
(MT4+MT5) × S&P 500	-	-7.128** [-9.924, -4.331]	-	-19.247*** [-24.872, -13.622]	-	-36.953*** [-46.656, -27.250]
(MT4+MT5) × VIX	-	-2.783 [-5.697, 0.130]	-	-9.787* [-15.641, -3.933]	-	-40.323*** [-52.172, -28.475]
MT6 × TFP	-	23.747*** [18.062, 29.431]	-	42.958*** [31.541, 54.374]	-	-26.119 [-48.327, -3.911]
Number of observations	523	523	521	521	512	512
Adjusted R ²	0.012	0.073	0.030	0.117	0.070	0.284

Table 3: Text-augmented exchange rate regressions

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, the h -month future change in the log of government debt and of 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. * p-value < 0.1, ** p-value < 0.05, *** p-value < 0.01.

higher weight to them.

Table 3 reports results from the regression model estimated with and without the metatopic interactions. We test three versions of equation (2) using monthly, quarterly, and annual exchange rate changes, with the interactions between metatopics and macroeconomic series disciplined by the results of Figure 8-10. The coefficients of the benchmark models without metatopic attentions are often not significant. In contrast, the models with the metatopic attentions interacted with the macroeconomic variables show coefficient estimates significantly different from zero in most cases. Moreover, we observe a large increase in the adjusted R² relative to the benchmark specifications, rising from 7% to 28% in the model predicting the annual change of the exchange rate, and from 1.2% to 7.3% in the regression predicting the monthly change

in the exchange rate, in line with the values obtained by [Fratzscher et al. \(2015\)](#).⁷

Overall, we have demonstrated that using metatopic attention series as proxies for importance weights proves particularly useful in disciplining the unstable relationship between exchange rates and economic indicators shown in Figure 11. The results in table 3 offer evidence of an exchange rate reconnect, supporting the scapegoat theory. However, the success of the text-augmented regressions may also be explained by alternative theories, in which the importance of economic time series *depends on* their media coverage, or more generally, by environments where news selection affects the business cycle. For example, in [Nimark and Pitschner \(2019\)](#) and [Chahrour et al. \(2021\)](#) agents face attention constraints and delegate the collection of information to news publishers. In such environments, news media can significantly alter the propagation of structural shocks, thereby influencing the importance of economic fundamentals.

5 Conclusion

News media represents a novel source of information for economists and is not a substitute for traditional macroeconomic indicators. Market agents may overreact or underreact to current and future economic developments depending on the media coverage of these events. Moreover, the information conveyed through news media is structured differently than that of conventional time series data. News is organized by salience and relevance, acting as a natural filter that can enhance economic analysis. Additionally, because news media comes in the form of text, it offers valuable insights for interpreting latent factors in economic models.

In this paper, we have presented a simple yet innovative approach to extracting relevant information from media outlets. Focusing on the dollar exchange rate, we have leveraged a topic model to distill a large word-document corpus into six distinct exchange rate narratives. These narratives, which capture multiple potential sources of exchange rate fluctuations, prove to be highly informative. Our results point to fundamental drivers of exchange rates, though the relationship is nonlinear. For example, fiscal and monetary policy influence dollar fluctuations primarily at the beginning of our sample period, while stock market dynamics are more

⁷[Fratzscher et al. \(2015\)](#), Table 4, reports an adjusted R^2 of 8.3% for the constant parameter model (CP-MS) predicting the EUR/USD monthly exchange rate changes from 2001 and 2011.

influential in the late 1990s and early 2000s. Technology, or expectations surrounding it, plays a more consistent role across the entire sample.

Much work remains, particularly in understanding the connection between the outcomes of topic models or Natural Language Processing methods and traditional macroeconomic models. In this paper, we have demonstrated one such mapping, applying the scapegoat theory of exchange rates and using metatopic attention series as proxies for the importance investors place on macroeconomic indicators. Given that scapegoat effects are likely to be observed more broadly in macroeconomics and finance beyond exchange rates, our strategy may be successful in other contexts as well.

Declaration of generative AI and AI-assisted technologies in the writing process During the preparation of this work the authors used ChatGPT in order to label the estimated topics. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

References

- ANDERSON, T. G., T. BOLLERSLEV, F. X. DIEBOLD, AND C. VEGA (2003): “Micro effects of macro announcements: Real-time price discovery in foreign exchange,” *American economic review*, 93, 38–62.
- ANDRE, P., I. HAALAND, C. ROTH, AND J. WOHLFART (2023): “Narratives about the Macroeconomy,” .
- ANGELETOS, G.-M., F. COLLARD, AND H. DELLAS (2020): “Business-cycle anatomy,” *American Economic Review*, 110, 3030–3070.
- ASH, E. AND S. HANSEN (2023): “Text algorithms in economics,” *Annual Review of Economics*, 15, 659–688.
- BACCHETTA, P. AND E. VAN WINCOOP (2004): “A scapegoat model of exchange-rate fluctuations,” *American Economic Review*, 94, 114–118.
- BACCHETTA, P. AND E. VAN WINCOOP (2013): “On the Unstable Relationship Between Exchange Rates and Macroeconomic Fundamentals,” *Journal of International Economics*, 91, 18–30.
- BAKER, S. R., N. BLOOM, AND S. J. DAVIS (2016): “Measuring economic policy uncertainty,” *The quarterly journal of economics*, 131, 1593–1636.
- BLEI, D. AND J. MCAULIFFE (2007): “Supervised topic models,” *Neural Information Processing Systems*, 21.
- BLEI, D. M., A. Y. NG, AND M. J. JORDAN (2003): “Latent dirichlet allocation,” *The Journal of Machine Learning Research*, 3, 993–1022.
- BRANDT, M. W., J. H. COCHRANE, AND P. SANTA-CLARA (2006): “International risk sharing is better than you think, or exchange rates are too smooth,” *Journal of monetary economics*, 53, 671–698.
- BREIMAN, L. (2001): “Random Forests,” *Machine Learning*, 45, 5–32.
- BYBEE, L., B. KELLY, A. MANELA, AND D. XIU (forthcoming): “Business News and Business Cycles,” *Journal of Finance*.
- CHAHROUR, R., V. CORMUN, P. DE LEO, P. A. GUERRÓN-QUINTANA, AND R. VALCHEV (2024): “Exchange rate disconnect revisited,” Tech. rep., National Bureau of Economic Research.

- CHAHROUR, R., K. NIMARK, AND S. PITSCHNER (2021): “Sectoral media focus and aggregate fluctuations,” *American Economic Review*, 111, 3872–3922.
- CORMUN, V. AND P. DE LEO (2024): “Shocks and exchange rates in small open economies,” *Working Paper*.
- DORNBUSCH, R. (1976): “Expectations and exchange rate dynamics,” *Journal of political Economy*, 84, 1161–1176.
- ENGEL, C. AND J. FRANKEL (1984): “Why interest rates react to money announcements: An explanation from the foreign exchange market,” *Journal of Monetary Economics*, 13, 31–39.
- ENGEL, C. AND K. D. WEST (2005): “Exchange Rates and Fundamentals,” *Journal of Political Economy*, 113, 485–517.
- ENGEL, C. AND S. P. Y. WU (2024): “Exchange Rate Models are Better than You Think, and Why They Didn’t Work in the Old Days,” Tech. rep., Working Paper.
- 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,” *The American economic review*, 69, 610–622.
- FRATZSCHER, M., D. RIME, L. SARNO, AND G. ZINNA (2015): “The scapegoat theory of exchange rates: the first tests,” *Journal of Monetary Economics*, 70, 1–21.
- GENTZKOW, M., B. KELLY, AND M. TADDY (2019): “Text as data,” *Journal of Economic Literature*, 57, 535–574.
- GOETZMANN, W. N., D. KIM, AND R. J. SHILLER (2022): “Crash narratives,” Tech. rep., National Bureau of Economic Research.
- HANSEN, S., M. McMAHON, AND A. PRAT (2018): “Transparency and Deliberation within the FOMC: a Computational Linguistic Approach,” *The Quarterly Journal of Economics*, 133, 801–870.
- ITSKHOKI, O. AND D. MUKHIN (2021): “Exchange rate disconnect in general equilibrium,” *Journal of Political Economy*, 129, 2183–2232.
- JIANG, Z. (2021): “US Fiscal cycle and the dollar,” *Journal of Monetary Economics*, 124, 91–106.
- KIM, S. AND N. ROUBINI (2008): “Twin deficit or twin divergence? Fiscal policy, current account, and real exchange rate in the US,” *Journal of international Economics*, 74, 362–383.

- LARSEN, V. H. AND L. A. THORSRUD (2019): “The value of news for economic developments,” *Journal of Econometrics*, 210, 203–218.
- MACAULAY, A. AND W. SONG (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*, Elsevier, vol. 6, 199–236.
- MCCRACKEN, M. W. AND S. NG (2016): “FRED-MD: A monthly database for macroeconomic research,” *Journal of Business & Economic Statistics*, 34, 574–589.
- MEESE, R. A. AND K. ROGOFF (1983): “Empirical Exchange Rate Models of the Seventies: Do They Fit Out of Sample?” *Journal of International Economics*, 14, 3–24.
- MONACELLI, T. AND R. PEROTTI (2010): “Fiscal policy, the real exchange rate and traded goods,” *The Economic Journal*, 120, 437–461.
- NIMARK, K. P. AND S. PITSCHNER (2019): “News media and delegated information choice,” *Journal of Economic Theory*, 181, 160–196.
- RAVN, M. O., S. SCHMITT-GROHÉ, AND M. URIBE (2012): “Consumption, government spending, and the real exchange rate,” *Journal of Monetary Economics*, 59, 215–234.
- ROSSI, B. (2013): “Exchange rate predictability,” *Journal of economic literature*, 51, 1063–1119.
- SCHMITT-GROHÉ, S. AND M. URIBE (2022): “The effects of permanent monetary shocks on exchange rates and uncovered interest rate differentials,” *Journal of International Economics*, 135, 103560.
- SHILLER, R. J. (2017): “Narrative economics,” *American economic review*, 107, 967–1004.
- (2020): *Narrative economics: How stories go viral and drive major economic events*, Princeton University Press.
- STAVRAKEVA, V. AND J. TANG (2020): “A fundamental connection: Exchange rates and macroeconomic expectations,” *FRB of Boston Working Paper*.
- THORSRUD, L. A. (2020): “Words are the new numbers: A newsy coincident index of the business cycle,” *Journal of Business & Economic Statistics*, 38, 393–409.
- UHLIG, H. (2003): “What moves real GNP?” *Working Paper*.

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, $\beta_k \sim \text{Dir}(\eta)$, and the topic distribution for each document, $\theta_d \sim \text{Dir}(\alpha)$. 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:

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

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

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

Figure 12 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.

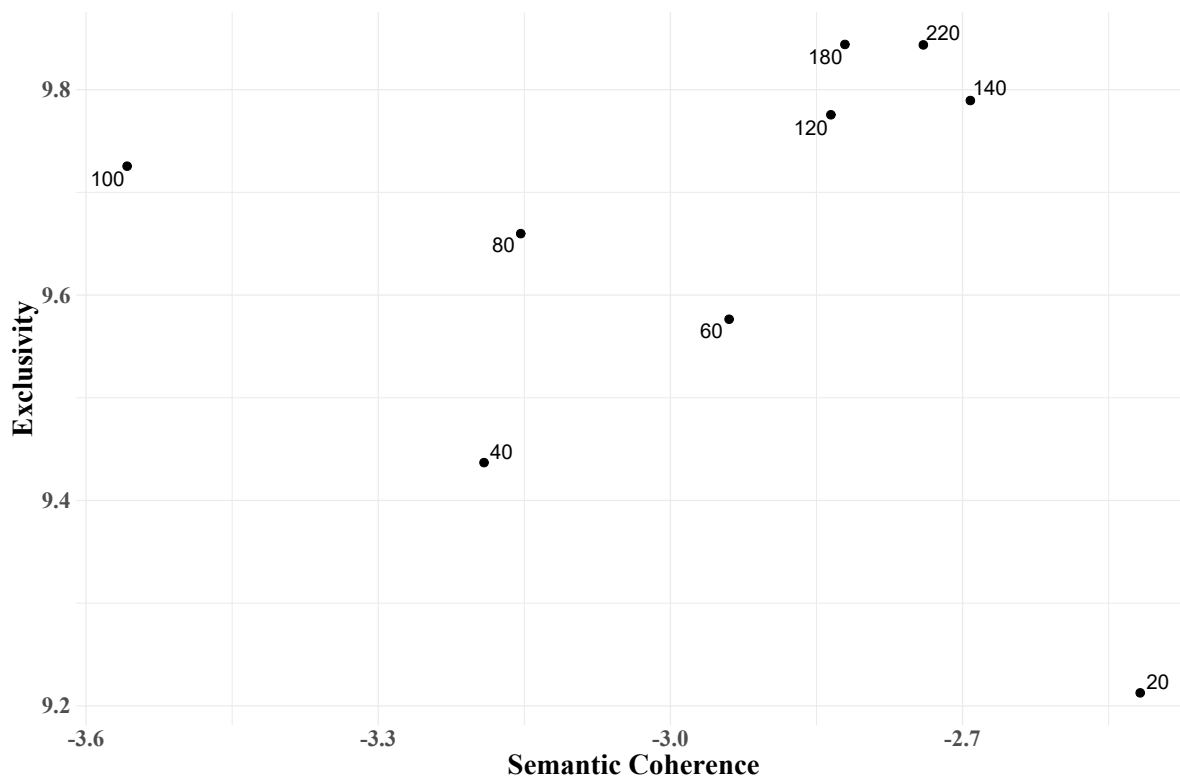


Figure 12: Topic model performances for different topic numbers

Note: The figure reports the estimated values of semantic coherence and exclusivity for models with different numbers of topics, from 20 to 220.

B Text collection and cleaning

The newspaper text data was collected from Proquest Historical Newspapers using the TDM tool between 2022-03-28 and 2022-03-04. 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 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 (open-source Python library developed for part-of-speech tagging) to perform the following:
 - (a) Remove words with the following entity categorization: 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 (stop-words, weekdays, months and numbers) and one character words.
5. We remove titles with less 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:

Type	Commentary	Feature	Front Page/Cover Story	News	Total
N	34,999	125,521	1,493	545,971	707,984

Table 4: WSJ headlines corpus breakdown by text type

C Topic words and labels

ChatGPT4 prompt November 14th 2023

Prompt: *I have an excel file that represents the output of a topic model. More specifically the file includes the beta vectors for each topic e.g. topic-word distributions. We are explaining/predicting exchange rate movements with the topics and have selected a subset of the 180 topics that includes 16 topics. We have also formed*

6 metatopics from these 16 topics with the topics being in the same metatopics if their beta vectors are similar to each other (basically a hierarchical clustering analysis with the beta vectors). We would like to label the 16 individual topics and the 6 metatopics. Can you help us with this by looking at the topic-word probabilities in the excel?

ChatGPT: Yes, I can certainly assist you with naming the topics and metatopics based on their topic-word distributions from your Excel file. To proceed, please upload the Excel file containing the beta vectors for each topic. I'll analyze these distributions and suggest names for the individual topics and metatopics based on the most significant words in each topic. Keep in mind that the naming process is somewhat subjective and is influenced by the interpretation of the words' significance within the context of your data. My suggestions will be based on the most prominent words in each topic, considering their relevance and possible thematic connections. If you have any specific context or domain knowledge that might help in this process, feel free to share it as well.

Prompt: A little bit of background for the file. Sheet number 2 has the topic-word distributions. The first column is the number of the topic. We are only interested in topics 148, 42, 128, 168, 147, 135, 102, 113, 39, 87, 96, 117, 140, 62, 6, 24. We would like to understand whether these topics relate to long-lasting (3-10 year) narratives explaining changes in the dollar exchange rate at different horizons. The period of the articles is from 1976 to 2020.

ChatGPT: I have successfully loaded your Excel file. The second sheet contains the topic-word distributions with each row representing a topic and columns representing words with their associated probabilities. Given your interest in specific topics (148, 42, 128, 168, 147, 135, 102, 113, 39, 87, 96, 117, 140, 62, 6, 24) and the aim to identify long-lasting narratives related to US exchange rate movements, I'll focus on these topics. The process involves examining the highest probability words in each topic to infer possible themes, which can be then labeled accordingly. Let's start by extracting the relevant topics and their top words to facilitate the naming process. Based on the top words in each topic, here are some suggested names for the individual topics. These suggestions are derived from the most prominent words in each topic, keeping in mind the context of long-lasting narratives influencing US exchange rates: 1. Topic 6: "Financial Strategies and Banking" 2. Topic 24: "Carter Era Fiscal Policies" 3. Topic 39: "International Trade and Acquisitions" 4. Topic 42: "Global Business and Finance News" 5. Topic 62: "Oil and Energy Market Dynamics" 6. Topic 87: "Corporate Transactions and Stock Market" 7. Topic 96: "Economic Health and Commodity Markets" 8. Topic 102: "International Corporate Taxation" 9. Topic 113: "Interest Rate and Bond Market Fluctuations" 10. Topic 117: "Technology and Business Developments" 11. Topic 128: "International Corporate Finance" 12. Topic 135: "Global Finance and Investment News" 13. Topic 140: "Wall Street and Financial Reporting" 14. Topic 147: "Business Economics and Policy" 15. Topic 148: "Reaganomics and Debt Management" 16. Topic 168: "Pharmaceutical Industry and Stock Market" For the metatopics, we would need to know which topics are grouped together under each metatopic. This grouping is essential for naming the metatopics accurately, as each would represent a broader

theme encompassing multiple topics. If you have this grouping available, please provide it so I can suggest names for the metatopics as well.

Prompt: Thanks! The metatopics are the following: 1 (39, 6, 102,113), 2 (24,62,148), 3 (42,135), 4 (87), 5 (96,147,168), 6 (117,128,140).

ChatGPT: Thank you for providing the groupings for the metatopics. Based on the individual topic names I suggested and their respective groupings, here are some proposed names for the six metatopics:

- 1. Metatopic 1 (Topics 39, 6, 102, 113): "Global Finance and Monetary Policies" - This metatopic combines themes related to international trade, banking strategies, corporate taxation, and interest rate dynamics, reflecting a broad focus on global financial and monetary policies.*
- 2. Metatopic 2 (Topics 24, 62, 148): "Historical Economic Policies and Energy Markets" - This grouping brings together topics related to fiscal policies during the Carter era, oil and energy market dynamics, and Reagan-era debt management, highlighting historical economic policies and their influence on energy markets.*
- 3. Metatopic 3 (Topics 42, 135): "International Business and Financial News" - This metatopic covers global business and finance news, including market insights and investment news, representing a focus on international business and financial journalism.*
- 4. Metatopic 4 (Topic 87): "Corporate Transactions and Stock Market" - As a standalone metatopic, this focuses on corporate transactions, stock market activities, and business dealings, emphasizing corporate finance and market dynamics.*
- 5. Metatopic 5 (Topics 96, 147, 168): "Economic Trends and Corporate Strategy" - This metatopic combines economic health, business economics, and pharmaceutical industry dynamics, indicating a blend of economic trends and corporate strategic movements.*
- 6. Metatopic 6 (Topics 117, 128, 140): "Technological Advancements and Financial Reporting" - This grouping combines technology and business developments, international corporate finance, and Wall Street reporting, highlighting the intersection of technology, corporate finance, and financial journalism.*

These names are suggestions based on the top words in each topic and their thematic connections. The actual names may vary depending on the deeper context and narrative connections you've observed in your analysis.