Shaping Economies Through Tax News: Analyzing the Influence on Economic Aggregates, Firm Behavior, and Financial Markets

Longji Li

October 17, 2023

Abstract

This paper focuses on the economic consequences of ”anticipated” shocks to taxes, specifically news or foresight about future shifts in tax policy. I apply an advanced topic model known as seeded Latent Dirichlet Allocation (LDA). This model offers advantages over unsupervised LDA and lexicon-based methods, providing a more nuanced interpretation of tax policy discourse. By analyzing a large corpus of news articles and presidential documents, I create two indices that separately quantify the intensity of news related to tax increases and decreases over time, starting from the Truman administration. Comparative analysis indicates that the metrics derived from seeded LDA are reliable. They can effectively forecast tax shocks as identified in the narrative-based approach by Romer and Romer (2010), as well as a commonly cited tax news measure by Leeper et al. (2012). The study incorporates both macroeconomic and firm-level perspectives, presenting empirical evidence that tax foresight significantly influences economic aggregates, corporate behavior, and financial markets. First, news of tax hikes stimulates output in the short term, and the effect of such news varies across different economic states. Second, firms with greater market power are more reactive to news of tax hikes, accelerating investment at a faster pace than their counterparts. Third, firms that are heavily reliant on government purchases show increased stock price volatility when attention to tax policy among investors is high. Finally, both stock and bond markets exhibit immediate responses when there is news of tax hikes, as these are viewed as steps toward fiscal responsibility and sustainable economic growth.

1 Introduction

Assessing the impact of fiscal policies, including taxation and government spending, is complex due to the concept of ”fiscal foresight.” This refers to how individuals or businesses change their economic behavior in anticipation of forthcoming policy changes. The lag between when a fiscal policy is first announced and when it is actually implemented often provokes such preemptive actions, thus complicating policy analysis. The significance of these anticipatory behaviors in shaping the efficacy of policy measures has been a focal point of extensive scholarly debates and investigations. Recent empirical studies, therefore, tend to distinguish between anticipated and unanticipated changes in fiscal policies.

To specifically explore the impacts of tax foresight, I constructed two distinct indices that gauge the intensity of news related to tax hikes and tax cuts over time, the sum of them measures the tax policy news in general. These indices are derived from news articles and presidential documents housed in the American Presidency Project (APP), a free online resource maintained by the University of California, Santa Barbara (UCSB). Starting with the era of President Truman, the indices generally exhibit spikes during times that correspond with either tax increase or tax reduction reform episodes, as identified by Yang (2009) and Amaglobeli et al. (2018). Following the spirit of Bybee et al. (2023), one can also think of tax news intensity as media attention paid to tax policy.

Measuring tax policy news is not an easy task. First, the quantification of text concerning tax policy is often labor-intensive and time-consuming, especially given that tax reforms are usually described in terms of being either a tax cut or a tax hike, serving multiple societal, economic, and political goals. Second, lexicon-based approaches in textual analysis are often susceptible to subjectivity as researchers need to predetermine an exhaustive list of terms based on their domain knowledge. For instance, when assessing economic activity, one could create a list of relevant terms, such as ”growth,” and each term is treated equally in the list without assigning some importance weights. Additionally, the term ”growth” could also be employed in different settings, like describing wage growth contributing to inflation. Addressing these contextual nuances using dictionary-based techniques poses significant practical challenges.

To address these hurdles, I used a topic modeling approach called Latent Dirichlet Allocation (LDA) to scrutinize a large corpus of documents for their relevance to tax policy. This allows me to automatically identify the most crucial terms for differentiating between texts and to quantify the proportion of content in each document that pertains to topics like tax hikes or tax cuts. I also conducted a comparative analysis of tax news measures using a predefined lexicon and found that LDA is more adept than the lexicon-based method at capturing the nuances and subtle variations in language, particularly as tax hikes are perceived unfavorably and often framed as necessary for political or economic reasons. The results also provide confidence that LDA is an effective tool for identifying tax-related news topics with minimal subjectivity.

Various tax policy reform indicators and news indices exist in the literature. To highlight my contribution, I compared the LDA-identified tax news measure with a commonly-cited tax foresight measure called implicit tax rate by Leeper et al. (2012). The results suggest that the LDA-based metric can not only track the progression of information related to potential changes in future taxes, thereby mirroring how economic agents continuously process and assimilate this information, but it also serves as a robust predictor for established benchmarks. In addition, Romer and Romer’s (2010) narrative-based exogenous tax shocks can be predicted by my tax news measure, indicating that tax policy changes are effectively anticipated.

I provide empirical evidence of how macroeconomic aggregates, firm-level outcomes, and asset markets evolve in the wake of tax policy news movements. To make progress, I begin with Jordà's (2005) local projection method, which sheds some light on how GDP responds to tax news on average and how its responses could differ across four different economic states defined by economic slack (Ramey and Zubairy, 2018), policy attention (An et al., 2022), policy uncertainty (Baker et al., 2016), or the monetary-fiscal policy mix (Ascari et al., 2022). The results confirm the existence of strong state-dependent tax foresight effects. To make more robust causal inferences, I adopt two sets of micro approaches by exploiting firm-level differences in market power and in exposure to government purchases of goods and services. My findings reveal that firms wielding more market power are more responsive to anticipated tax changes. Specifically, these firms ramp up investments at a faster pace than their less powerful counterparts when a tax increase is on the horizon. In addition, firms with greater exposure to government purchases experience greater stock price volatility when there is high media attention to tax news. The results are robust even when I use the tax news extracted from the Wall Street Journal as an alternative measure of media attention (Bybee et al., 2023).

Finally, the stock market, in general, reacts immediately and positively to shocks from tax hike news when the level of federal deficits is high; its response is muted in other periods. Turning to the bond market, which is often a more reliable signal of future economic trends than the stock market, I find that credit spreads narrow immediately following tax hike news. This suggests an optimistic outlook for future economic expansion. To decompose the credit spread, I consider excess bond premium and default risk. The former can be interpreted as an indicator of the overall credit supply conditions in the economy. I find that the excess bond premium declines significantly upon impact, while the default risk component does not respond. This suggests that the variation in credit spread, conditional on the tax hike news shock, is driven primarily by factors related to credit supply conditions. Overall, my findings at the asset market level suggest that news of tax hikes can amplify the competitive advantages enjoyed by firms with higher markups. This amplification occurs not only through strategic investment decisions but also through enhancements in overall credit market conditions.

This article relates to at least two strands of literature. The first focuses on the impact of fiscal foresight. Research by scholars such as Auerbach and Gale (2009), Leeper et al. (2012), Mertens and Ravn (2010, 2012), Ramey (2011), Leduc and Wilson (2013), Caggiano et al. (2015), and Chirinko and Wilson (2023) reveals marked differences in the economic outcomes of expected and unexpected fiscal changes. This highlights the importance of fiscal foresight. However, the quantitative significance of fiscal foresight is not corroborated by other studies, such as the vector autoregression analyses conducted by Blanchard and Perotti (2002) and Perotti (2012). Similarly, studies by Poterba (1988) and Romer and Romer (2010) also fail to substantiate the critical role of fiscal foresight. In addition, the existing literature generally studies the foresight effects through the lens of the macroeconomy. Based on a novel measure of foresight derived directly from tax news, my findings demonstrate that the anticipatory effects of tax policy are robust across various layers, from aggregate economic indicators to asset market behaviors and individual firm activities.

The second strand of literature is expanding to focus on sophisticated methods for textual analysis. These methods utilize sources like newspapers, annual reports, and transcripts of earnings calls to gauge various outcomes. Most recently, a study by Bybee et al. (2023) employed an unsupervised LDA model on articles from the Wall Street Journal. This study asserts that business news can serve as a predictor for a broad spectrum of economic activities and collective stock market performance. Earlier works by Larsen and Thorsrud (2019), as well as Thorsrud (2020), also used the LDA approach on

Norwegian news data to examine macroeconomic forecasting models. In 2018, Hansen et al. published what is, to my understanding, the first paper to feature LDA applications in a top-tier economics journal. They applied the LDA model to statements from the Federal Open Market Committee (FOMC) to study how individual contributors focus their attention during meetings. Dybowski and Adämmer (2018) used the correlated topic model, a variation of the conventional LDA, to analyze the macroeconomic effects of tax policy sentiment. They found that optimistic tax policy statements stimulate consumption, investment, and output, even after controlling for tax foresight using Leeper et al.'s (2012) implicit tax rate as the foresight proxy.

The paper proceeds as follows. Section 2 describes how I measure tax news. Section 3 presents the identification strategy for tax news shocks. Sections 4, 5, and 6 examine how tax news affects economic aggregates, firm-level dynamics, and asset markets, respectively. Section 7 concludes.

2 Measuring Tax News Through Advanced Topic Modeling

The field of economics has experienced a paradigm shift in the analysis of news and policy discourses through the application of machine learning algorithms. Topic modeling has emerged as a leading methodology to scrutinize the latent structures within a large corpus of documents. The method offers invaluable insights into financial news, policy debates, and prevailing economic sentiments. Scholars have leveraged topic modeling algorithms to gain a nuanced understanding of economic discourses, track market sentiment evolution, and even forecast market behavior based on evolving news narratives. The importance of topic modeling lies in its capacity to render a high-dimensional data structure into comprehensible topics that can inform economic decision-making and forecasting.

2.1 Overview of Topic Modeling

Topic modeling, as a technique, sits at the intersection of machine learning and natural language processing, aiming to discover the latent topical patterns in a collection of documents. This section serves as an overview of the methodology used in this study.

2.1.1 Classic Latent Dirichlet Allocation (LDA) Model

The Latent Dirichlet Allocation (LDA) model, originally introduced by Blei et al. (2003), posits that any document collection can be effectively represented by a finite set of underlying topics. A defining feature of LDA is that it identifies topics that permeate through all documents in the corpus, albeit in varying proportions. This makes LDA a universally applicable tool, crossing disciplinary boundaries from computational linguistics to economics. The LDA model's advantages, such as its reproducibility and automated nature, make it particularly attractive for researchers handling massive collections of text data. It streamlines the research process by reducing the time and effort involved, thereby mitigating the biases that may arise from human interpretation or lexicon-based methods.

In the domain of natural language processing, understanding the thematic structure of large textual corpora is a non-trivial task. One common approach to tackle this problem is to represent the text corpus in its "bag-of-words" form, depicted as the article-term matrix, w, as shown on the left-hand side of Figure 1. It is a $T \times V$ matrix where rows correspond to the list of T articles in the

Figure 1: Article-Term Matrix and Topic Distributions After LDA

$$
\begin{array}{c|c|c|c|c|c|}
 & V_1 & V_2 & V_3 & V_4 & V_5 \\
\hline
T_1 & w_{1,1} & w_{1,2} & w_{1,3} & w_{1,4} & w_{1,5} \\
\hline
T_2 & w_{2,1} & w_{2,2} & w_{2,3} & w_{2,4} & w_{2,5} \\
\hline
T_3 & w_{3,1} & w_{3,2} & w_{3,3} & w_{3,4} & w_{3,5} \\
\hline
T_4 & w_{4,1} & w_{4,2} & w_{4,3} & w_{4,4} & w_{4,5} \\
\hline
\end{array}
\quad \longrightarrow \quad
\begin{array}{c|c|c|c|}
 & K_1 & K_2 & K_3 \\
\hline
T_1 & \theta_{1,1} & \theta_{1,2} & \theta_{1,3} \\
\hline
T_2 & \theta_{2,1} & \theta_{2,2} & \theta_{2,3} \\
\hline
T_3 & \theta_{3,1} & \theta_{3,2} & \theta_{3,3} \\
\hline
T_4 & \theta_{4,1} & \theta_{4,2} & \theta_{4,3} \\
\hline
\end{array}
$$

corpus, and columns correspond to the vocabulary of V unique terms in the corpus. Its individual elements, $w_{t,v}$, count the number of times that the vth term appears in article t.

While the bag-of-words model substantially simplifies the complexity of the original text, it nonetheless produces an object with remarkably high dimensionality. LDA tackles this by generating a more comprehensible thematic overview of w, which can be easily interpreted by humans (as demonstrated in the right-most column of Figure 1). LDA posits that the V-dimensional vector of term counts for a specific article t, labeled as w_t, is governed by a multinomial distribution:

$$w_t \sim \text{Mult}(\Phi' \theta_t, N_t) \tag{1}$$

Here, N_t denotes the overall term count in article t and sets the scale of the multinomial distribution. In layman's terms, the expected term counts are represented by a lower-dimensional set of parameters, θ_t and $\Phi = [\phi_1, \ldots, \phi_K]'$. The kth topic in the article is defined by the V-dimensional parameter vector ϕ_k, where $\phi_{k,v} \geq 0$ for all v and $\sum_v \phi_{k,v} = 1$. In essence, a topic is characterized as a probability distribution over terms. Terms with particularly high probabilities in ϕ_k encapsulate the core theme of the topic. The reduction in dimensionality is realized by making K, the total number of topics, substantially smaller than the size of the vocabulary.

While the topics, ϕ_k, encapsulate the prevalent themes throughout the corpus, LDA views each individual article as a composite of multiple topics. The article-specific parameter vector $\theta_t = (\theta_{t,1}, \ldots, \theta_{t,K})'$ functions as a probability vector, where $\theta_{t,k} \geq 0$ for all k and $\sum_k \theta_{t,k} = 1$. In essence, θ_t quantifies the distribution of focus that article t places among the various topics. LDA incorporates a factor structure, where the topics ϕ_k act as shared factors, and θ_t measures how each article is uniquely influenced by these common factors.

In most applications, the parameters ϕ and θ can be estimated via a Bayesian approach using the posterior distributions, typically employing Gibbs sampling. The underlying logic of this estimation procedure is moderately technical but understandable. The LDA model can also be expressed as a generative model framed by a collection of sampling rules. Assuming ϕ_k and θ_t are given, the process of article creation within this model framework can be conceptualized.

Consider the t-th article composed of N_t words. The model begins by randomly selecting a topic from the list of all topics, the probabilities for which are encapsulated in the $K \times 1$ vector θ_t. The first term's topic assignment, symbolized as $z_{(t,1)}$, follows a unit multinomial (also known as categorical) distribution with parameter θ_t. Assuming that topic k is selected for the first term, the model then randomly selects the first term itself from the entire vocabulary. The likelihood of drawing any specific term is directed by the $V \times 1$ vector ϕ_k, which describes the term distribution specific to the topic. Let's refer to this first term as $x_{(t,1)}$, which also adheres to a unit multinomial distribution with parameter ϕ_k. This procedure is reiterated for each of the subsequent terms in the article, culminating after N_t iterations, at which point the article is fully generated. In terms of distribution notation, the i-th term in the article is described as:

$$x_{(t,i)} \sim \text{Mult}(\phi_{z_{(t,i)}}, 1), \quad z_{(t,i)} \sim \text{Mult}(\theta_t, 1). \tag{2}$$

The Gibbs sampling algorithm aims to find the values of ϕ_k and θ_t that most closely mimic the articles present in the actual text corpus.

The approximated proportion of topic k within article t can be represented by the frequency of its terms being allocated to topic k:

$$\hat{\theta}_{t,k} = \frac{\sum_{i=1}^{N_t} \mathbb{I}(\hat{z}_{t,i} = k)}{\sum_{q=1}^{K} \sum_{i=1}^{N_t} \mathbb{I}(\hat{z}_{t,i} = q)} \tag{3}$$

where $\hat{z}_{t,i}$ is the estimated topic assignment for each term i in article t.

Similarly, the estimated topics themselves are derived from the frequency at which each term v in the vocabulary is assigned to topic k, summed across all articles:

$$\hat{\phi}_{k,v} = \frac{\sum_{t=1}^{T} \sum_{i=1}^{N_t} \mathbb{I}(\hat{z}_{t,i} = k)\mathbb{I}(x_{t,i} = v)}{\sum_{m=1}^{V} \sum_{t=1}^{T} \sum_{i=1}^{N_t} \mathbb{I}(\hat{z}_{t,i} = k)\mathbb{I}(x_{t,i} = m)} \tag{4}$$

Lastly, the aggregate proportion of each topic k at a specific point in time τ (such as a month or a quarter) can be calculated by summing over all articles published at that time. Therefore, the estimated proportion of topic k during quarter τ is given by:

$$\hat{\theta}_{\tau,k} = \frac{\sum_{t \in \tau} \sum_{i=1}^{N_t} \mathbb{I}(\hat{z}_{t,i} = k)}{\sum_{t \in \tau} \sum_{q=1}^{K} \sum_{i=1}^{N_t} \mathbb{I}(\hat{z}_{t,i} = q)} \tag{5}$$

2.1.2 Seeded LDA Model

However, the unsupervised version of the LDA algorithm, operating within a generative probabilistic framework, often fails to distinguish between nuanced topics like a tax hike and a tax cut, which may appear similar but have drastically different implications. This limitation stems from the model's surface-level focus on word associations, influenced by uniform priors α and β in standard LDA.

In the context of Latent Dirichlet Allocation (LDA) and its variants like seeded LDA, the β matrix represents the distribution of words for each topic. Specifically, $\beta_{k,v}$ denotes the probability of encountering word v given that the topic is k. Higher values of $\beta_{k,v}$ imply that the word v is more likely to appear in documents strongly associated with topic k and is a strong representative of that topic. Conversely, lower $\beta_{k,v}$ values suggest that the word v is less likely to appear in such documents and is not particularly informative of topic k.

To overcome the limitations of standard LDA, an extended approach known as seeded or guided LDA modifies the Dirichlet prior β based on a set of seed words S_k for each topic. The formula $\beta'_{kw} = \beta_{kw} + \delta \cdot I(w \in S_k)$, where $I(\cdot)$ is an indicator function and δ is a constant indicating the strength of the seeding, is employed. By manipulating β in this manner, seeded LDA serves as a form of "guidance" or "steering" to make the model more interpretable or tailored towards specific needs. This enables the model to differentiate between nuanced topics like tax hikes and tax cuts more effectively.

It's important to note, as highlighted by Jagarlamudi et al. (2012), that while seed words guide the topic generation, they do not impose strict constraints. The model retains its flexibility to deviate if compelling evidence within the data suggests otherwise.

In conclusion, this guided LDA variant enhances the traditional algorithm's capabilities, making it invaluable for nuanced economic studies requiring differentiation between closely related policy topics.

2.1.3 Lexicon-Based Method for Measuring Tax News

In the fields of finance and economics, lexicon-based techniques for text analysis rely on an established set of terms that have been systematically classified based on their sentiment—either positive, negative, or neutral—or in relation to a specialized subject area such as economic policy uncertainty. These designated terms are given specific numerical weights or scores, which can range from simple frequency counts to more complex metrics. The purpose of this is to capture the overarching sentiment or thematic focus within a given text corpus.

One of the primary strengths of using lexicon-based approaches is the ease with which results can be interpreted. Since these methods hinge on a predetermined set of terms, the resulting analyses are usually straightforward and easily comprehensible. This level of interpretability is especially advantageous for stakeholders who may lack expertise in the area of natural language processing (NLP). Moreover, lexicon-based methods boast computational efficiency and scalability, qualities that make them well-suited for scrutinizing extensive datasets or for use-cases that necessitate real-time analytical insights.

However, these methods are not without their drawbacks. The fixed list of terms used in lexicon-based analysis can be both an asset and a limitation. While its predefined nature affords the method its scalability and ease of interpretation, it also restricts the method's flexibility and adaptability to the fluidity of language or domain-specific terminology. In essence, the method may struggle to capture the nuanced shifts in sentiment or topic, which are often vital in dynamic fields like finance and economics.

To illustrate the application of lexicon-based methods in practical research, Lyziak and Sheng (2022) study can be considered as a case in point. In their research, they employ a news-based index to examine the extent of disagreement among experts in the U.S. concerning inflation forecasts, as presented in Wall Street Journal articles. The methodology involves collecting quarterly data on articles containing at least one term from four pre-established categories: 'Expert,' 'Inflation,' 'Forecast,' and 'Direction.' An article is deemed to capture expert disagreement on inflation expectations if it discusses at least two contrasting viewpoints about future inflation trends.

In a subsequent section of this paper, I implement a methodology similar to that of Lyziak and Sheng (2022), albeit on a different textual dataset. The objective is to determine whether the outcomes generated using this lexicon-based approach significantly diverge from results acquired through LDA.

2.2 Tax-Related Textual Data Retrieval

2.2.1 The American Presidency Project

The text documents used in this study are sourced from the APP. This comprehensive archive contains various presidential documents, including but not limited to spoken addresses and remarks, news conferences and press briefings, interviews, statements, proclamations, and even tweets (specifically those made by former President Donald Trump). The documents in the APP are organized by presidents and presented in chronological order (See Figure 2).

The APP is free to access, non-commercial and allows automated data retrieval. Commercial databases like Factiva, ProQuest, or LexisNexis severely prohibit such an action, and authorized bulk download of computer-readable news texts from these databases is very expensive and not friendly to independent researchers. Most importantly, first-hand news about tax policy directly released from the president is more relevant to the research question in this paper since information about new tax legislations is generally conveyed in

Figure 2: The American Presidency Project (APP) Website Homepage

Note: https://www.presidency.ucsb.edu/presidents

presidential speeches and statements (Romer and Romer, 2010). Major news outlets such as the WSJ generally quote the president's words in their articles, and presidents could also direct news attention towards top issues that are on the administration's agenda, helping the forward-looking public form expectations (Eshbaugh-Soha, 2013). For instance, the APP archives a statement made by former President Donald Trump on April 21, 2017, during which he commented on signing an executive order and memorandums related to the regulation of the financial services industry. In this statement, he noted, "And we'll be having a big announcement on Wednesday having to do with tax reform. The process has begun long ago, but it really formally begins on Wednesday. So go to it." The WSJ and Reuters promptly cited this remark in their coverage on the same day. Therefore, the APP seems to be an appropriate source for obtaining tax legislation news.

To shed light on the insights the seeded LDA model is expected to glean about future tax policy changes, I offer a chronological sequence of documents from former President Donald Trump, ranging from the initial discussion of tax reform to its eventual legislative enactment:

1. Interview with Bill O'Reilly of Fox News (Date: February 03, 2017)

- O'Reilly: Well, that's good. Let's get to the bottom of this. In 2017, can Americans expect a **tax cut**?
- The President: I think so, yes, and I think before the end of the year. I would like to say yes.

2. The President's Weekly Address (Date: April 28, 2017; 1st 100 days in office)

- Our country is going up, and it's going up fast. Our companies are doing better—they just announced fantastic profits—all because of what's happened in this rather short period of time. And that's just the beginning. We're putting in a massive **tax cut** for the **middle class** and for **business**. It's going to have an enormous effect.

3. Remarks in an Exchange with Reporters (Date: September 24, 2017, few days before unified framework for tax reform unveiled)

- Q: Did the Big Six finalize your tax plan yet, do you know?
- The President: Yes, we have a tax plan that's totally finalized. I think it will be terrific. I think it's going to go through, and it's a very—it will be the largest **tax cut** in the history of our country.

4. Tweets of October 25, 2017

- "Working hard on the biggest **tax cut** in U.S. history. Great support from so many sides. Big winners will be the **middle class business** & **JOBS**"

5. Remarks Prior to a Cabinet Meeting (Date: November 20, 2017)

- I want to congratulate the House of Representatives for passing a vital and historic **tax cut** last week, and I'm very hopeful the Senate will do the same very soon. We're going to give the American people a huge **tax cut** for Christmas. Hopefully, that will be a great, big, beautiful Christmas present.

6. Remarks on Signing Legislation on Tax Reform (Date: December 22, 2017)

- As you know, 3.2 trillion in **tax cut** for American families, including the doubling of the standard deduction and the doubling of the child **tax credit**. The typical family of four earning $75,000$ will see an income **tax cut** of more than $2,000$—many much higher than that—slashing their **tax bill** in half. And they're going to start to see that. Because we're signing today, they're going to start to see that in February. The numbers will speak.

To capture potential policy changes and disclosures prior to the formal inauguration, this study also includes documents by the president-elect. These documents can provide insights into expected or revealed policy shifts. For instance, a notable example (see Figure 3) is the intensive tweeting by Trump concerning potential tax changes leading up to his inauguration.

The collected corpus from the APP contains about 65,000 documents over 1945 to 2019, comprising a total of around 75,000,000 words. After the corpus is processed by following conventional procedures in the field of natural language processing (NLP), the seeded LDA model can step in and calculate the proportions of tax cut topic and tax hike topic for each document.

2.2.2 Seed Words Selection for Guided LDA Model

The lexicon for identifying "tax hike" and "tax cut" topics is meticulously chosen based on two authoritative sources:

1. The narrative analysis by Romer and Romer (2009), offering an Act-by-Act breakdown of postwar tax changes.
2. Yang's chronology (2009) on federal income tax policies, also featuring a granular summary of tax legislations.

Figure 3: Former President Trump's Tweets About Potential Tax Changes Before His Inauguration

December 04, 2016 11:41:47	The U.S. is going to substantialy reduce taxes and regulations on businesses, but any business that leaves our country for another country, **Retweets:** 13697 **Favorites:** 53412
December 04, 2016 11:49:06	fires its employees, builds a new factory or plant in the other country, and then thinks it will sell its product back into the U.S. ...... **Retweets:** 11997 **Favorites:** 46329
December 04, 2016 11:57:41	without retribution or consequence, is WRONG! There will be a tax on our soon to be strong border of 35% for these companies ...... **Retweets:** 12287 **Favorites:** 48467
December 04, 2016 12:05:35	wanting to sell their product, cars, A.C. units etc., back across the border. This tax will make leaving financially difficult, but..... **Retweets:** 9654 **Favorites:** 40464
December 04, 2016 12:21:01	these companies are able to move between all 50 states, with no tax or tariff being charged. Please be forewarned prior to making a very ... **Retweets:** 10315 **Favorites:** 42184
December 04, 2016 12:23:27	expensive mistake! THE UNITED STATES IS OPEN FOR BUSINESS **Retweets:** 14892 **Favorites:** 58188

Both sources provide a wealth of direct quotations from presidents, standard phrases commonly found in tax policy narratives, legislative motivations, and aggregate revenue change estimates for each tax act. Leveraging this rich data, reforms are categorized into either tax hikes or tax cuts. Subsequently, NLP techniques are employed to distill the most frequently occurring bigrams—two-word sequences—from each category of texts.

The decision to focus on bigrams over unigrams or trigrams is a calculated one. Unigrams lack the contextual depth required for this analysis, while trigrams are less frequently encountered in the source narratives, making them less representative. Over-reliance on either could introduce undesirable artifacts such as overfitting and semantic misspecification due to excessive subjective judgment.

Given these considerations, the top three bigrams for each topic are selected to serve as seed words in the LDA model's inference process. This minimizes human intervention, thereby preserving the objectivity of the model. The selected seed words are summarized in Table 1 below.

Table 1: Seed Words for Guided LDA Model

Topic	**Seed Words**
Tax Cut	tax_cut, tax_reduction, economic_recovery
Tax Hike	tax_increase, deficit_reduction, budget_deficit

2.3 Tax Topic Extraction

To effectively use LDA for topic modeling, the number of topics (k) must be predetermined. Choosing an appropriate value for k is crucial, as setting it too low can result in overly broad topics that hinder analysis, while setting it too high can lead to redundant or overly specific topics. One approach to determining the optimal value for k is to use the CV (C_V) Coherence Score to evaluate the

Figure 4: Optimum Number of Topics Based on CV (C_V) Coherence Score

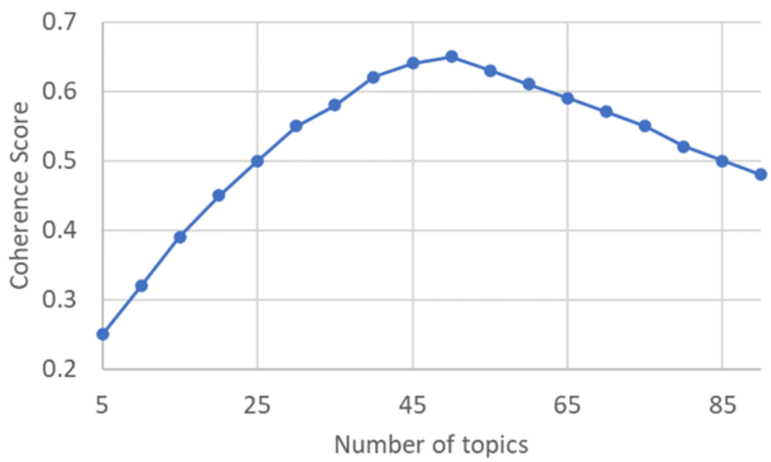

coherence of the topics generated by the LDA model. This involves iteratively running the LDA model with different values of k and calculating the corresponding CV (C_V) Coherence Scores. A higher coherence score indicates that the generated topics are more coherent and interpretable, making them better suited for downstream analysis. The rule of thumb is using the number of topics that yields the highest score before flattening out/declining, and therefore $40 \leq k \leq 50$ could be a good range (see Figure 4). Within this range, I can test the LDA model by incrementally increasing the number of topics and manually inspecting the resulting topics. I find the optimal k should be 46, which creates the most interpretable and relevant topics to my research question.

Figure 5 gives the tax topics estimated by the seeded LDA. The most probable vocabulary for each topic is generally in line with major postwar tax policy objectives. According to Yang (2009), raising revenues (or reducing deficits) and stimulating (or promoting) economic growth are the dominated themes in postwar tax events, and only three were motivated to contain inflation (thus inflation-related terms are assigned lower probabilities) in the early 1950s and late 1960s. However, a major issue underlying the LDA model and other pop-

Figure 5: Word Cloud of the Seeded-LDA-Identified Tax Topics

(a) Tax cut topic

(b) Tax hike topic

ular textual analysis algorithms with the bag-of-words assumption is that the model is likely to classify the sentences such as "there will be no new tax increases" as a signal about future tax hikes. It should also be noted that LDA focuses on word co-occurrences only. Therefore, for example, one can see that both "economic_recovery" and "recovery" are in the same topic word set, even though the former is semantically included in the latter.

2.4 Quantifying Tax News

With the two identified tax topics and the corresponding proportions for each corpus document, I can construct two time series of tax news, according to the equation of $\hat{\theta}_{\tau,k}$ in the previous section: 1) the aggregate tax news intensity (ATNI); and 2) the net tax news intensity (NTNI). The former is capturing the tax news topic attention over time and defined as the sum of the topic proportions of tax hike and tax cut topics, and the latter is capturing the potential direction of future tax changes and is a difference given by the topic proportion of the tax hike topic minus that of the tax cut topic.

Figure 6 shows (from top to bottom): 1) the proportions of the tax hike and tax cut topics; 2) the ATNI; and 3) the NTNI over time from Truman to Trump. Tax episodes (shaded red & green bars) are identified either by Yang (2009) or taken from Amaglobeli et al. (2018), and are defined as the number of months elapsed between when a tax proposal was first officially announced by the president and when it was signed into law. Although the time series increases during identified legislative lags, it also spikes several times outside of them. This indicates that presidents gave further tax policy speeches which might have constituted the beginning of a legislative lag or early signals of tax changes.

For example, the period from January 1987 to December 1988 shows elevated levels of "Hike News" but no actual "Hike Episodes" under President Ronald Reagan. Several factors could account for this phenomenon. Firstly, the late 1980s were marked by growing concerns over the federal budget deficit. Although the Tax Reform Act of 1986 had simplified the tax code, it did not necessarily alleviate budgetary concerns, possibly leading to discussions about potential tax hikes. The Act was revenue-neutral, meaning it was designed not to affect the overall level of taxation but to redistribute it. As a result, while the tax code became simpler, the Act did not significantly boost federal revenue or cut the deficit. In this context, the late 1980s saw ongoing discussions and debates about other ways to address the deficit issue, which could include tax hikes. Secondly, Ronald Reagan, known for his anti-tax stance, was in his second term and not running for reelection, allowing for more open dialogue about tax hikes among policymakers. These unproductive discussions about tax hikes eventually took concrete form when George H.W. Bush assumed the presidency from 1989 to 1993. Contradicting his 1988 campaign promise of "Read my lips: no new taxes," Bush consented to a 1990 budget agreement that featured tax hikes aimed at shrinking the federal budget deficit, along with spending reductions and revenue enhancements.

Figure 6 also implies that tax policy has been more intensively and consistently discussed since the 1980s, as evidenced by higher news spikes and slower reversion to the zero line. This trend may be explained by the structural shifts in tax policy objectives, the high policy priority of using tax instruments, and the notion that tax cuts are more often and openly discussed by presidents than tax hikes. Before the 1980s, the main tax policy objective was to fight high inflation (see Figure 7) or to finance wars (e.g., for the Korean and Vietnam Wars financing) by triggering a series of tax increases. Since the 1980s (the 2000s especially), however, the government has been relying on tax policy to fight recessions or promote long-run growth. This priority relies on using either demand-side tax cuts such as the large-scale tax rebate, increases in standard deductions, and individual general tax credits aimed to boost private spending or supply-side measures such as massive corporate income tax reductions as well as reductions in top individual rates.

It should be noted that although deficit issues are also debated and deficit-driven tax hike acts are enacted accordingly, economic growth and tax cut policy still dominate after the 1980s regardless of how deficits deteriorate. Figure 8 shows the deficit trend over time, which may shed some light on the tax policy priority shifts.

2.5 Placing the LDA-Identified Tax News Information Flow Within the Existing Measures

Figure 9 demonstrates the placement of my measure within the current body of research on documenting tax policy changes. Yang (2009), Romer and Romer (2010), Mertens and Ravn (2012), and Amaglobeli et al. (2018) utilize narrative-based approaches to analyze tax policy. Yang (2009) delivers an extensive his-

Figure 6: Seeded-LDA-Identified Tax Policy News Index (Monthly)

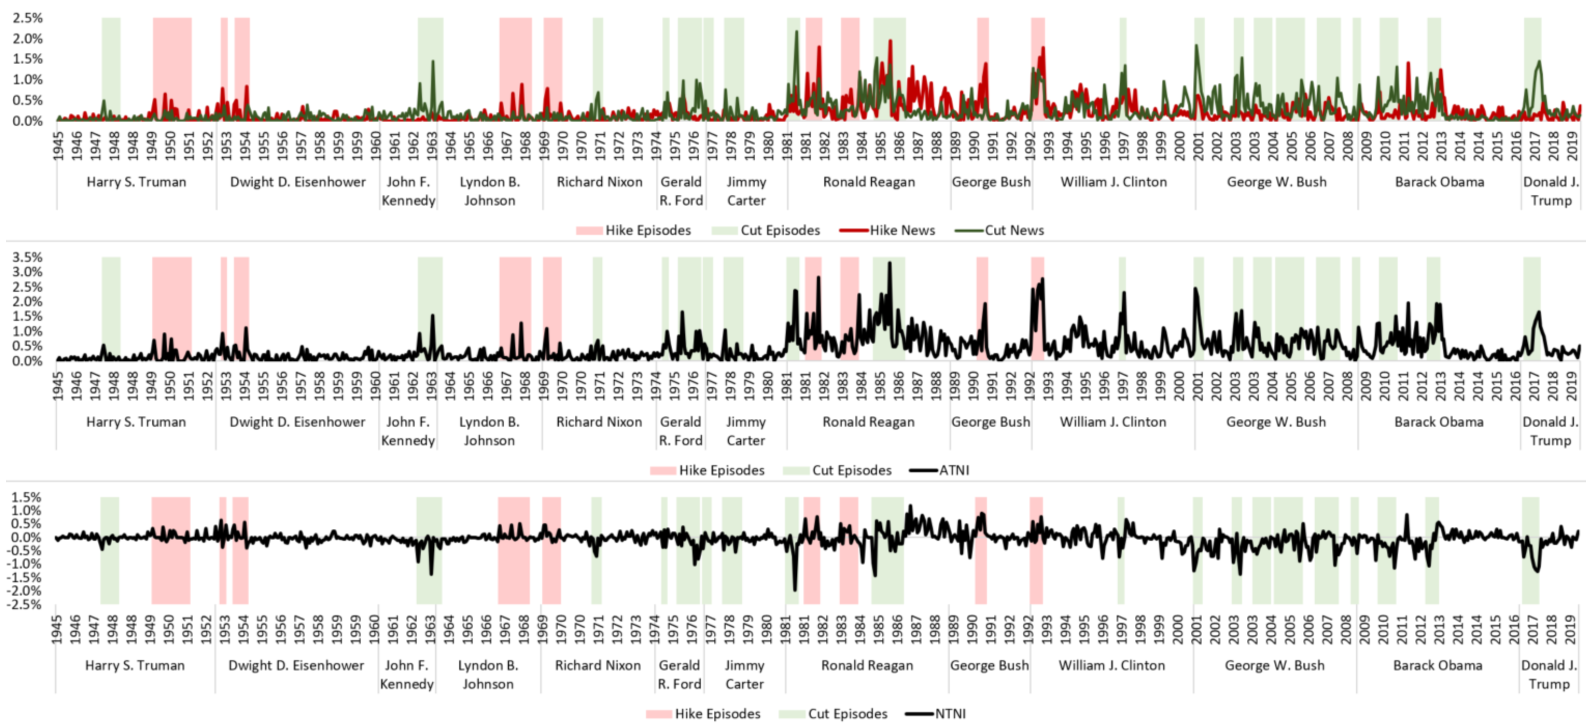

Figure 7: Inflation Rate

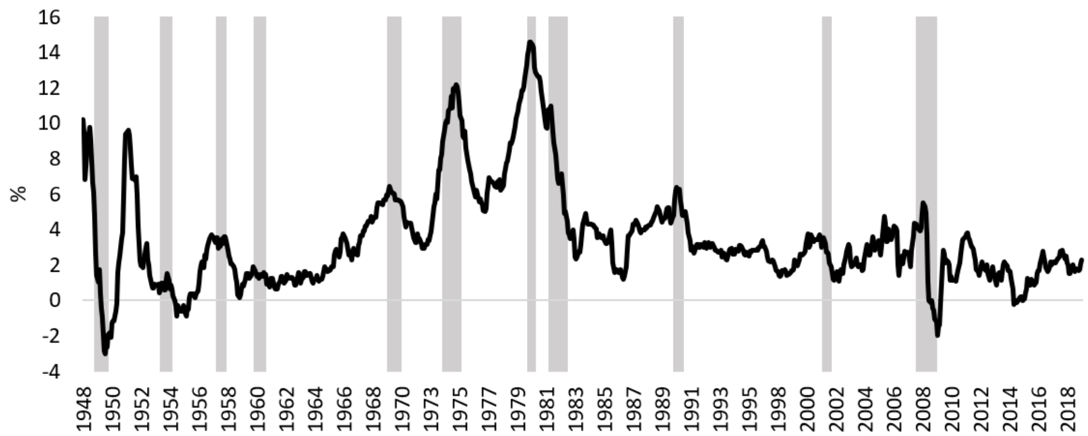

Note: Shaded areas denote recessions.

Figure 8: Federal Surplus Share of GDP

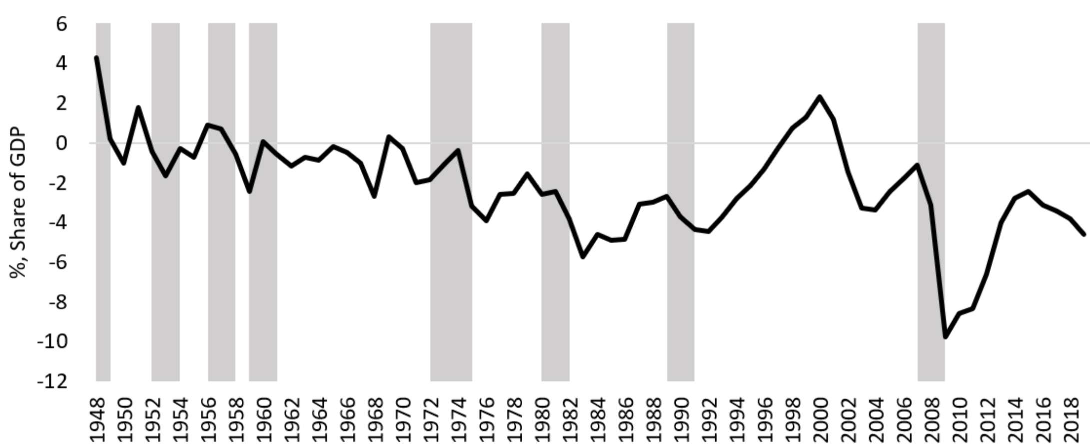

Note: Shaded areas denote recessions.

Figure 9: General Form of News Information Flow

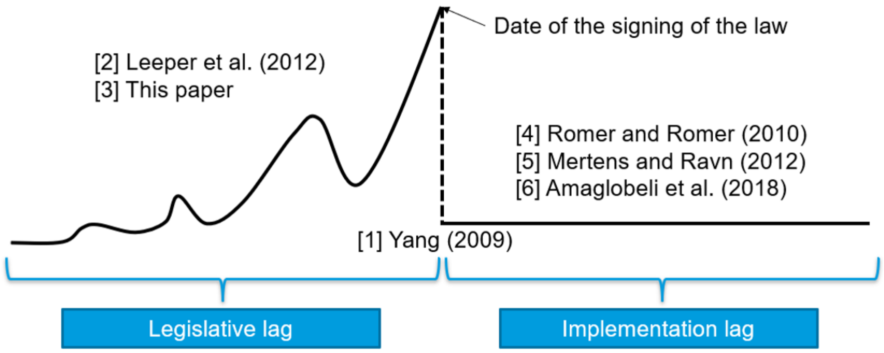

torical overview of the evolution and progression of federal income tax policy in the United States from 1947 to 2009, covering reform motivations, legislative timelines, crucial tax-related provisions, and anticipated shifts in total revenue. Romer and Romer (2010) conducts a groundbreaking study that examines the relationship between fiscal policy and economic growth by investigating historical tax records to identify substantial exogenous tax policy changes, ascertaining the occurrence of tax changes based on actual liability alterations, and measuring each change's size relative to GDP. Mertens and Ravn (2012) elaborates on Romer and Romer (2010)'s analysis of the macroeconomic impacts of tax policy, providing a deeper comprehension of how expectations influence economic results by concentrating on tax changes that have been enacted but not yet enforced. Amaglobeli et al. (2018) supplements Romer and Romer (2010) and Mertens and Ravn (2012) by broadening the range of countries examined and breaking down tax changes into rate and base effects for various taxes, such as corporate income tax (CIT) and personal income tax (PIT).

When analyzing the effects of tax changes on the macroeconomy, firm dynamics, and asset price movements, it is important to consider the dynamic evolution of information flow. While measures such as implementation or announcement dates provide a useful starting point, they fail to capture the ongoing process of information dissemination and the subsequent adjustments in expectations that occur well before the actual enactment of a tax policy. Furthermore, announcement date is not uniformly defined across various studies. For instance, Romer and Romer (2010) consider the date when the final tax bill is signed as the announcement date. On the other hand, Amaglobeli et al. (2018) identify the announcement date as the moment when the government first formally proposes the tax plan. This could be through an official statement from the Prime Minister or the Ministry of Finance, or the introduction of draft legislation like a Budget bill to Parliament. My measure of tax news not only spikes on the dates pinpointed by these prior studies but also effectively captures the flow of relevant information well in advance of the formal tax proposal. The dynamic evolution of information flow can be visualized as a solid curved line in Figure 9, reflecting the continuous transmission and absorption of information by economic agents over time. As new information emerges, economic agents continually update their expectations about the potential tax policy changes.

In order to track the progression of information related to potential changes in future taxes, Leeper et al. (2012) employs the concept of the implicit tax rate. This rate is determined by calculating the yield difference between a one-year tax-exempt municipal bond and a one-year taxable Treasury bond. An increase in the implicit tax rate could signal the possibility of a future tax increase. In the United States, municipal bonds are not subject to federal taxes. This distinct treatment of municipal and Treasury bonds offers valuable insights for detecting news about tax modifications. Within the context of Blanchard and Perotti (2002)'s VAR framework, Leeper et al. (2012) discovers that a significant and considerable growth in output follows an increase in the implicit tax rate. Although the construction of the implicit tax rate is straightforward, Schwert (2017) argues that, since the financial crisis in 2007, the historical relationship between these bonds has broken down due to an increase in municipal

Figure 10: NTNI vs. Implicit Tax Rate

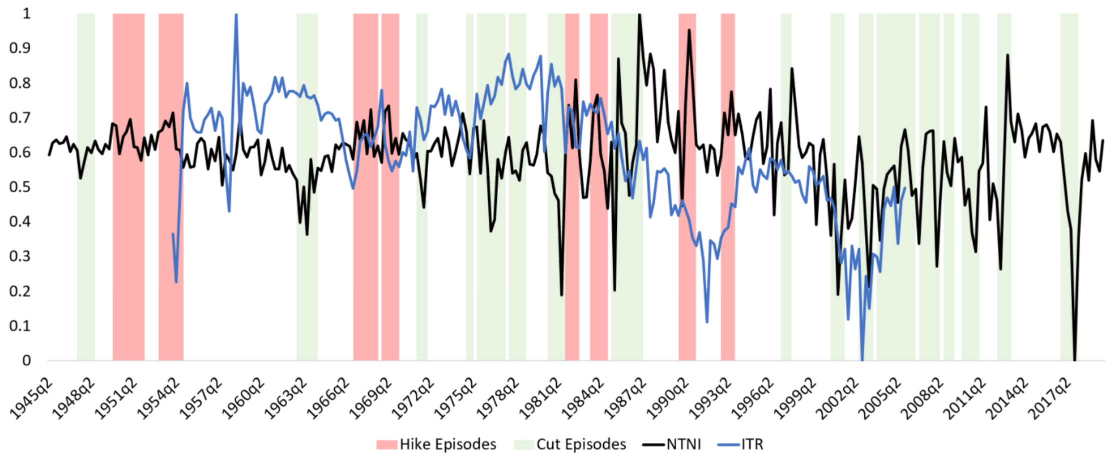

Note: ITR data is obtained directly from Leeper et al. (2012) and its coverage is narrower than NTNI.

default risk. Kueng (2018) also points out that there are other factors than expected tax rates affect municipal yield spreads. My measure overcomes these limitations by directly extracting tax news from a collection of news documents with the help of LDA algorithm. Figure 10 shows NTNI and implicit tax rate (ITR). Both are normalized to a 0-1 scale for a straightforward comparison.

To assess the predictive power of my tax news measure in relation to future tax changes, I conduct predictive regressions following the approach of Leeper et al. (2012) and Romer and Romer (2010) as outlined by Fantozzi and Muscarnera (2021). I estimate the following regression equation:

$$\Delta\tau_t = \alpha + \beta \cdot \text{NTNI}_{t-j} + \sum_{i=1}^{4} \gamma_i \cdot \text{Controls}_{t-i} + u_t \tag{6}$$

where $j = 1, 2, \ldots, 4$, $\Delta\tau$ could be the implicit tax rate from Leeper et al. (2012) or narrative tax changes (expected change in tax liabilities at time of implementation) from Romer and Romer (2010), and the controls include lags of the dependent variable, output growth, tax revenue growth, government spending growth, change in debt-to-GDP ratio, inflation, change in interest rate, and change in unemployment rate. The predictive power of NTNI on the implicit tax rate or narrative measures is determined by the F-statistic on the exclusion of NTNI.

Table 2: Predictive Power of NTNI

	j			
Source	1	2	3	4
Leeper et al. (2012)	14.92	8.02	5.96	2.35
Romer and Romer (2010)	5.28	13.26	38.07	27.57

Table 2 shows that NTNI is a strong predictor for both the implicit tax rate and narrative tax changes based on the typical rule of thumb that an F-statistic of more than 10 (Stock and Yogo, 2002). The NTNI can forecast a tax change one quarter earlier than the implicit tax rate and at least two quarters ahead of the narrative tax shocks identified by Romer and Romer (2010).

2.6 Comparative Analysis of Tax News Measures Using a Pre-defined Lexicon

Are humans or computers better at identifying tax hike and tax cut news in a large corpus of documents? Many of the prior papers use subjectively determined word lists to produce news-based indexes and the results are robust within their research settings or questions (Loughran and McDonald (2011); Baker et al. (2016); Łyziak and Sheng (2022)). For this approach to be effective, the process must be transparent, and the resulting lists should be reasonably exhaustive. Making the lists exhaustive precludes the potential for p-hacking a list down to the most ex post powerful words or having managers simply avoid identified words in crafting their future documents (Loughran and McDonald, 2020). To see if the lexicon-based method works well for my research question, I follow Łyziak and Sheng (2022) but with some adjustments given the fact that our focuses are different. It should be noted that I am not making a conclusion here that LDA is superior to conventional textual analysis methods. One approach should not fit all, and the choice needs to be case-specific.

I decompose each article into paragraphs and then keep those that contain "tax". If an article is not tax-related, then no paragraph will be retained. With tax-related articles, I then follow routine pre-processing steps such as stop word removal and lemmatization. To determine whether a news article is mainly about tax hike or tax cut, I count, for each article, the number of paragraphs that contain at least one term in tax hike term set and the number of paragraphs that contains at least one term in tax cut term set. An article is then coded as tax hike (tax cut) if the number of tax-hike paragraphs is greater (less) than that of tax-cut paragraphs. Lastly, I obtain the average number of coded articles within each month weighted by the absolute difference between the number of tax hike and tax cut paragraphs and denote the two resultant series as tax hike news index and tax cut news index. Alternatively, I can simply count the monthly number of coded articles.

Table 3 summarizes the key terms extracted from the same documents (i.e., presidential speeches compiled by Yang (2009) and Romer and Romer (2009)) I used for seed word selection in the previous section. Unlike those terms explicitly indicating future inflation path (e.g., rise, fall, or unchanged) in Łyziak and Sheng (2022), term selection for implying a tax hike or cut policy change should be mainly based on term co-occurrence. For example, the term "loophole" is likely to occur in a speech about tax hikes. Clearly LDA method shares the same logic as the lexicon-based one, but the former can minimize the risk of subjectivity and determine the importance of each term to the tax topic rather than treating them equally.

Figure 11 presents a comparison between lexicon-based measures and LDA-identified measures for tax news. The correlation between the two indexes is

Table 3: Term Sets for Tax Hike and Tax cut

Tax Hike	**Tax Cut**
additional revenue	small business
additional tax	burden
balance budget	create job
budget deficit	economic recovery
cut deficit	economic growth
defense spending	incentivize
deficit cut	increase employment
deficit reduction	increase investment
fair balance	investment
fair share	job
fairness	lower tax
federal revenue	middle class
increase revenue	rate reduction
increase tax	recession
tax loophole	reduce rate
military spending	reduce unemployment
new spending	relief
raise revenue	slow growth
raise tax	stimulate
reduce debt	stimulus
reduce deficit	strengthen economy
revenue increase	surplus
spending control	tax break
social responsibility	tax burden
social security	tax credit
tax hike	tax cut
tax increase	tax decrease
tax revenue	tax rebate
war	tax reduction

0.71 for ”Hike” and 0.77 for ”Cut,” respectively, in the lexicon-based and LDA-identified measures, indicating that both methods are effective at detecting significant changes in tax-related news. While spikes in tax reform episodes can be observed, the LDA-identified measures appear to capture more intense spikes, particularly in the period before the 1970s. During this time, tax policy objectives were primarily focused on combating high inflation or financing wars (e.g., for the Korean and Vietnam Wars financing), leading to a series of tax increases. It can be inferred from this that LDA is more adept at capturing the nuances and subtle variations in language, particularly as tax hikes are perceived unfavorably and often framed as necessary for political or economic reasons. Ultimately, the results of the lexicon-based approach provide confidence that LDA is an effective tool for identifying tax-related news topics with minimal subjectivity.

3 Tax News Shock Identification

In accordance with the methodology proposed by An et al. (2022), tax news shocks can be easily identified by conducting a regression analysis of tax news with lags of fiscal variables, other macroeconomic variables, and the tax news itself. The regression residuals are then referred to as tax news shocks. Specifically, I perform a regression of the NTNI on an extensive set of lagged variables, such as output growth, tax revenue growth, government spending growth, alterations in the debt-to-GDP ratio, inflation, changes in interest rates, fluctuations in unemployment rates, economic policy uncertainty (Baker et al., 2016), and the NTNI itself. I then acquire the regression residuals (i.e., unforeseen information about future tax changes) as tax news shocks. For clarity in subsequent empirical analyses, I standardize the obtained news shock series to have a mean of zero and a standard deviation of one. Figure 12 illustrates the (standardized) new shocks over time, with spikes corresponding to tax reform episodes. In my sample, prominent instances of extensive tax cut episodes include the Economic Recovery Tax Act of 1981, the Economic Growth and Tax Relief Reconciliation Act of 2001, and the Tax Cuts and Jobs Act of 2017. On the other hand, noticeable cases of large tax increase episodes comprise the Revenue Act of 1950, the Revenue and Expenditure Control Act of 1968, the Deficit Reduction Act of 1984, and the Omnibus Budget Reconciliation Act of 1990.

It is worth noting that the tax news shock, as identified within the SVAR framework by Blanchard and Perotti (2002), exhibits a strong correlation (approxi-

Figure 11: Lexicon-Based Results (Bottom Panel) vs. LDA Results

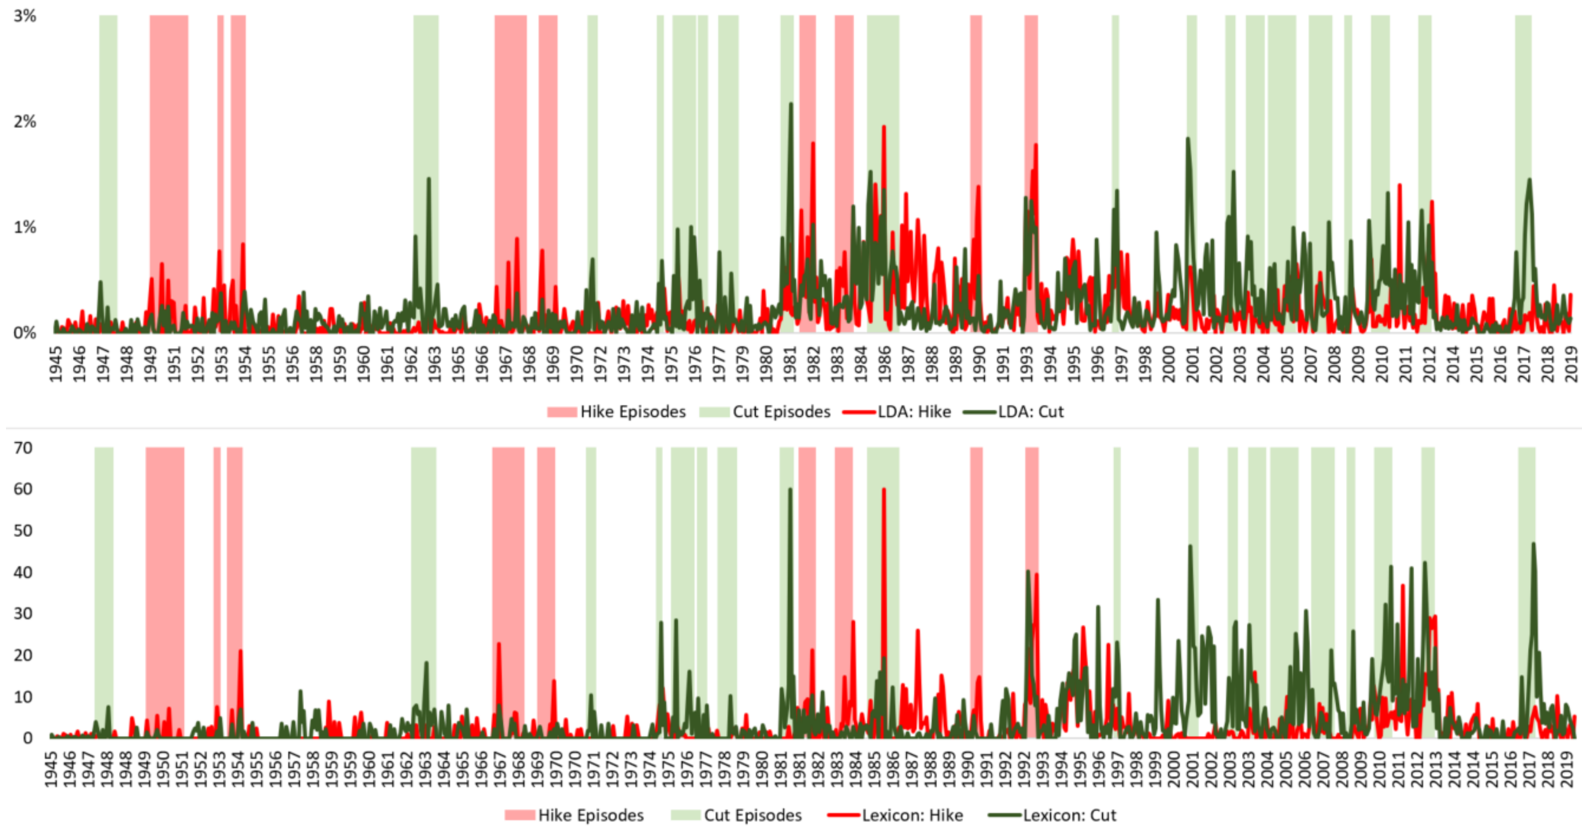

Figure 12: Tax News Shock Over Time

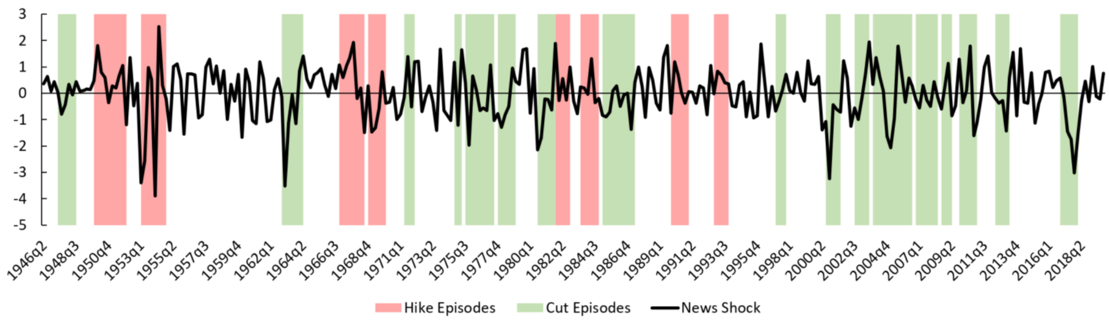

mately 0.9) with the tax news shock identified through regression residuals. This correlation provides reassurance regarding the validity of the tax news shock as identified in accordance with the methodology proposed by An et al. (2022). Details can be found in the Appendix.

4 Macroeconomic Effects of Tax News Shock

4.1 Model Specification

Recent fiscal literature increasingly used Jordà's (2005) local projections (LP). Compared to SVARs, LP allows a more parsimonious specification since not all variables are required to be included in all equations, and it does not impose dynamic restrictions embedded in SVARs. In addition, LP also easily accommodates state dependence or nonlinearities (Auerbach and Gorodnichenko, 2013; Ramey and Zubairy, 2018; Born et al., 2020).

I begin by estimating a linear LP model given by:

$$\tilde{y}_{t+h} = \alpha_t + \phi_h(L)\text{control}_{t-1} + \beta_h \text{shock}_t + \epsilon_{t+h} \tag{7}$$

for $h = 0, 1, \ldots, 20$, where $\tilde{y}_{t+h}$ is the response of an outcome variable of interest (e.g., the cumulative change in the log of real GDP) at horizon $t+h$ to a tax (hike) news shock at time t, and control_{t-1} includes a rich set of predetermined control variables including lags of output growth, tax revenue growth, government spending growth, inflation, change in debt-to-GDP ratio, change in interest rate, and change in unemployment rate. In addition, following Ramey and Zubairy (2018), I also include lags of tax news shocks to control for any serial correlation in the news variable. $\phi_h(L)$ is a lag polynomial of order 4. The impulse response for H periods is obtained from a sequence of estimates β_h.

I also investigate whether the output effect of tax news shock is dependent on the state of the economy by using a state-dependent LP as follows:

$$\begin{aligned} \tilde{y}_{t+h} = I_{t-1}[\alpha_{A,t} + \phi_{A,h}(L)\text{control}_{t-1} + \beta_{A,h}\text{shock}_t] + \\ (1 - I_{t-1})[\alpha_{B,t} + \phi_{B,h}(L)\text{control}_{t-1} + \beta_{B,h}\text{shock}_t] + \epsilon_{t+h} \end{aligned} \tag{8}$$

where I_t is a dummy variable that indicates the state {A,B} of the economy in the previous period. $\beta_{A,h}$ and $\beta_{B,h}$ measure the state-dependent responses of variable $\tilde{y}_{t+h}$ at time $t+h$ to a tax news shock at time t. Following Ramey and Zubairy (2018), I use the Newey-West correction for standard errors (Newey and West, 1987).

To imply whether the state-dependent responses are statistically different to tax news shocks, I test at each year of the forecast horizon the following hypothesis:

$$H_0 : \beta_{A,h} - \beta_{B,h} = 0$$

This hypothesis can be tested with a simple two-sided t-test. A similar approach is applied by Klein (2017) to test whether the effects of fiscal consolidations in high and low private-debt states are different.

Specifically, the paper takes into account four meticulously chosen economic states, each serving as a lens through which the complexities of fiscal policy can be more fully understood:

1. Economic slack: This state captures the cyclical nature of the economy. Utilizing the unemployment rate as an indicator, this study examines how the effectiveness of tax news shocks could vary during economic downturns compared to booms. The unemployment rate threshold that defines slack is following Ramey and Zubairy (2018).

2. General tax policy attention: The sum of tax cut and tax hike news, indicating the overall level of media attention on tax policy (An et al. (2022); Bybee et al. (2023)).

3. Economic policy uncertinty: Drawing from the influential work of Baker et al. (2016), this state helps understand how uncertainty can act as a moderating factor in the effectiveness of tax news shocks.

4. Fiscally- or monetary-led regime: Following the classification by Ascari et al. (2022), this state distinguishes periods dominated by fiscal policy actions from those characterized by monetary policy interventions. This differentiation allows the study to explore whether the impact of tax news shocks is conditioned by the prevailing policy focus. Fiscal regime during the Great Inflation (1960Q1-1979Q2) and monetary regime during the Great Moderation (1984Q1-2007Q2).

4.2 Linear Results

I begin by fitting a linear local projection model to quarterly U.S. data spanning from 1947Q1 to 2019Q4. Tax news shocks incorporated in the model are identified using regression residuals. Figure 13 demonstrates that output initially rises following a positive one-standard-deviation shock to tax news (i.e., tax hike news shock), peaking five quarters after the shock, and then declines. This observation aligns with the findings of Leeper et al. (2013) and Mertens and Ravn (2012), which suggest that anticipated tax increases can be expansionary in the short term. This occurs because individuals and firms are motivated to shift production to the anticipation period when taxes are expected to be lower, despite the different tax news measures used in their studies (either implicit tax rate or narrative tax shocks). In a recent paper, Herrera and Rangaraju (2019) examine the influence of federal tax news (represented by the implicit tax rate) on state economic activity. They also discover that tax hike news not only results in greater aggregate output growth but also leads to increased personal income and employment growth across the majority of states. Christofzik et al. (2022) report similar results for Germany. Alesina et al. (2015) argue that even such 'unanticipated' changes may have been informally anticipated, and the effects of tax shifts can be partially offset by anticipation of tax shifts in the opposite direction.

However, the actual implementation of tax changes introduces a contrasting dynamic. Since a significant amount of production and consumption has been shifted to the anticipation period, the actual period of tax change sees a decrease in these activities. The reason is that individuals and firms have already made many of their purchases and investments during the anticipation period, and now they are likely scaling back. As a result, there is a drop in demand and economic activity, which is reflected in the negative impact on GDP. Also, the tax increase itself, when it does come into effect, can have a contractionary effect on the economy. Higher taxes mean that both consumers and businesses have less net income to spend and invest, which can decrease aggregate demand and further slow down the economy. Therefore, these results combined indicate a sort of 'boom-bust' cycle induced by the announcement and subsequent implementation of a tax increase.

Figure 13: Output Response to a Tax News Shock

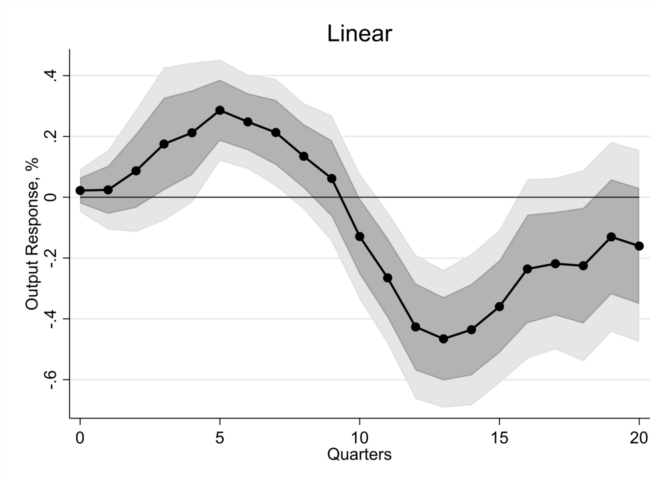

Note: Impulse response to a positive one-standard-deviation shock to tax news; Shaded areas are the 90 and 68 (darker) percent confidence intervals constructed by the Newey-West method.

Economic activities rest on households and firms making billions of consumption and investment decisions. Such decisions are not only directly influenced by economic realities (e.g., household's disposable income), but also indirectly driven by economic perceptions (e.g., perceptions that could respond to remarks and evaluations by the president regarding tax policy issues). Eventually, these decisions amount to macroeconomic aggregates of investment and consumption.

Figure 14 plots the consumption and nonresidential investment responses using the linear specification. The expansionary response to a tax hike news shock is indicative of forward-looking households accelerating their consumption until the tax hikes are realized. These results complement the work by Herrera and Rangaraju (2021) who, using implicit tax rate as a measure of tax foresight, also estimate an increase in consumption growth for households when future tax hikes are expected. Baker et al. (2019) find that car sales rise by more than 8% in the months before a 1% increase in the sales tax rate, pointing out that tax foresight translates into significant effects on durable

Figure 14: Consumption and Investment Responses to a Tax News Shocks

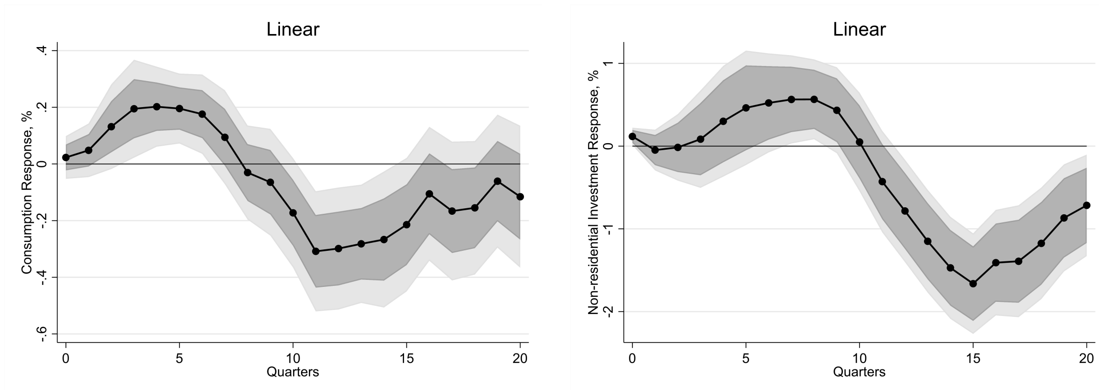

(a) Consumption response (b) Investment response

Note: Impulse response to a positive one-standard-deviation shock to tax news; Shaded areas are the 90 and 68 (darker) percent confidence intervals constructed by the Newey-West method.

consumption. Regarding investment, I find a short-run rise in the expenditures by firms on capital such as tools, machinery, and factories when tax hikes are anticipated, which empirically supports House and Shapiro (2006)'s argument that there are strong incentives to alter the timing of investment in response to anticipated tax changes.

4.3 State-Dependent Results

4.3.1 Tax News Shock During Slack States

Fgure 15 shows the response by running the state-dependent model where I distinguish the economy by slack states measured by the unemployment rate following Ramey and Zubairy (2018). It clearly shows the effect of a tax hike news shock on output is state-dependent, with a strong output expansion during low-unemployment states. In contrast, anticipation effects in bad times are muted in earlier. The results imply that tax foresight may significantly hinder the effects of countercyclical tax hike policies that aim to cool down the econ-

Figure 15: Output Response to a Tax Hike News Shock: Slack States (Red: Unemployment Rate $\geq$ 6.5)

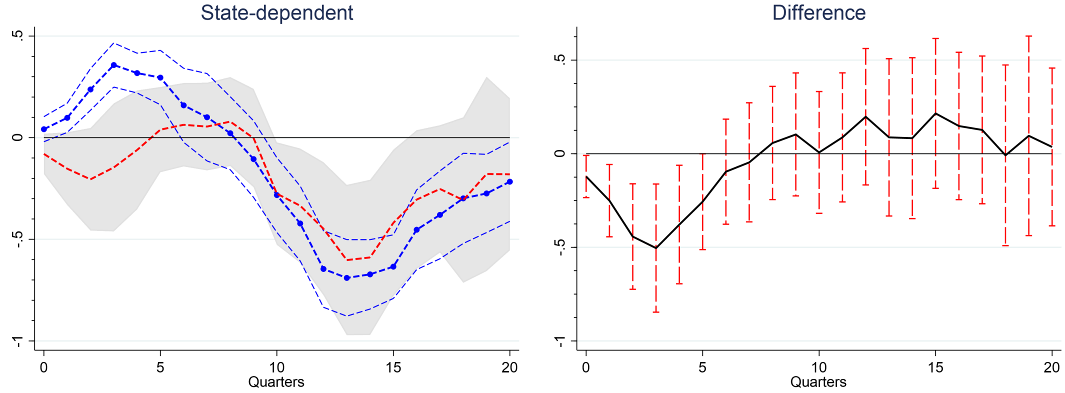

Note: The blue dashed line with circles is the response in low-unemployment states and the red dashed line is the response in high-unemployment states; 90 percent confidence intervals (e.g., red error bars and shaded areas) constructed by the Newey-West method are shown; The right panel shows the estimated difference between state-dependent responses.

omy. The right panel of the figure shows the respective differences $\beta_{A,h} - \beta_{B,h}$ for output at each quarter of the forecast horizon. A negative value indicates that the response in bad times is lower than in good times. The error bars indicate statistical significance at the 90 percent level, and therefore it complements the results in the left panel.

The findings are consistent with the research conducted by Hayo and Mierzwa (2022). They gather information on proposed tax alterations during the initial phase of the lawmaking process (for example, when a measure was initially announced in a white paper or as part of a parliamentary discussion), which may not always correspond to the ultimate figures. Moreover, they discovered that the impact of formulating tax increase legislation on US production is notably positive during prosperous periods.

4.3.2 Tax News Shock During Different States of Tax News Attention

It would also be worth studying whether the output effects of a tax news shock could differ, conditional on the level of media attention. A recent paper by An et al. (2022) investigates the role of inattention in shaping the impact of monetary policy on macroeconomic variables and finds that inattention, measured as proportion of forecasters who do not revise their forecasts of a target economic variable such as GDP growth, significantly amplifies the effects of monetary policy shocks, implying that the economic variables are more responsive to changes in monetary policy when agents are inattentive. Additionally, the results suggest that inattentive agents generate more persistent responses to monetary policy shocks, which can have important implications for the conduct of monetary policy.

Regarding the use of news as a measure of attention, Bybee et al. (2023) apply the unsupervised LDA model to the full text of WSJ articles for 1984-2017, summarize business news into interpretable topical themes (one of them is the tax theme), and define them as categorical news/media attentions. Although my source of news is different from theirs, major news outlets such as the WSJ generally quote the president's words in their articles, and presidents may also direct media attention towards top issues that are on the administration's agenda (Eshbaugh-Soha, 2013). The correlation between my ATNI and their WSJ general tax attention index is strong (about 0.7) and spikes are coincident in both. Therefore, I can study, with my tax news data, whether tax news shocks have differential output effects under different levels of attention.

Figure 16 clearly shows that, when media attention on tax is low (i.e., less than the sample median), a tax hike news shock can have a more significant output boosting effect during the anticipation period and a larger contraction when tax changes become effective in comparison to high-attention states. This falls

Figure 16: Output Response to a Tax News Shock: States of Tax News Attention (Red: High Attention States)

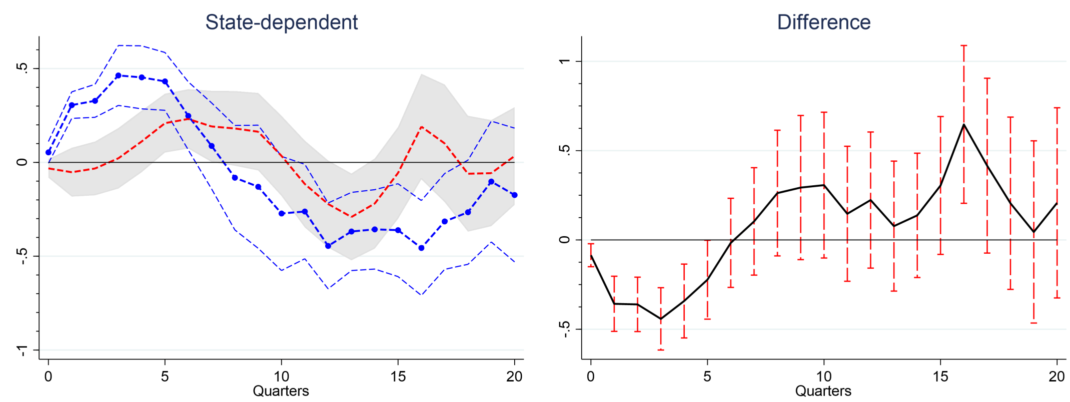

Note: The blue dashed line with circles is the response in low-attention states and the red dashed line is the response in high-attention states; 90 percent confidence intervals (e.g., red error bars and shaded areas) constructed by the Newey-West method are shown; The right panel shows the estimated difference between state-dependent responses.

in line with An et al. (2022)'s findings that monetary policy shocks have larger real effects when the degree of inattention is higher. The right panel displaying the difference in coefficients (i.e., high-attention minus low-attention) further corroborates this observation.

An et al. (2022) adopt a different specification of the local projection model than mine and identify periods of high attention as the episodes in which attention is above the 90th percentile and low attention as those below the 10th percentile. Specifically, their model is shown as follows:

$$y_{t+h,t} = \beta_{\alpha,h} MPS_t + \phi_{\alpha,h}(L) z_{t-1} + \beta_{b,h} MPS_t \times I_t + \phi_{b,h}(L) z_{t-1} \times I_{t-1} + \gamma_h I_t + \alpha_h + \epsilon_{t+h} \tag{9}$$

where y measures the cumulative cyclical component, through Hamilton's (2018) filter, of real GDP. z denotes a set of controls. The coefficient $\beta_{b,h}$ of the interaction term between attention (i.e., I) and monetary policy shock (i.e., MPS) shows the extent to which attention changes the effect of MPS. I modify this

Figure 17: Output Response to a Tax News Shock: States of Tax News Attention (Red: High Attention States Defined by Following An et al. (2022))

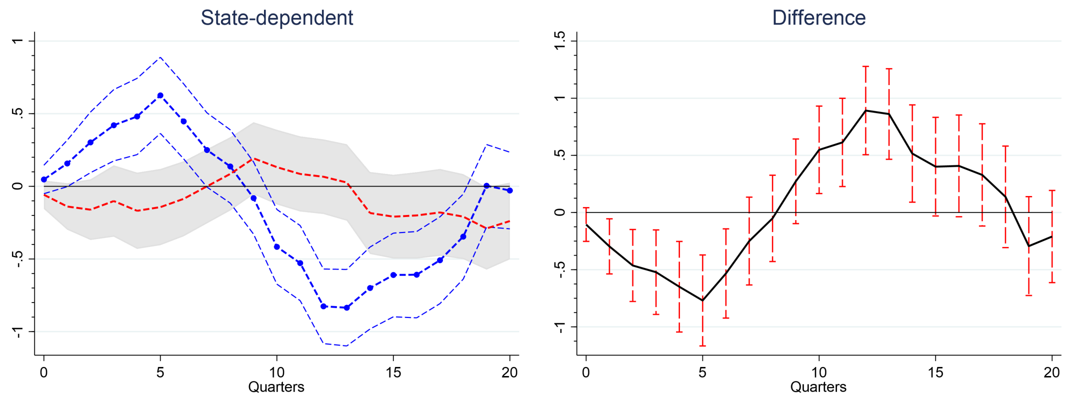

Note: The blue dashed line with circles is the response in low-attention states and the red dashed line is the response in high-attention states; 90 percent confidence intervals (e.g., red error bars and shaded areas) constructed by the Newey-West method are shown; The right panel shows the estimated difference between state-dependent responses.

empirical model by replacing MPS and I with my tax news shocks and news attention, respectively. I also borrow the same set of control variables from my previous local projection model.

Figure 17 shows that an unexpected tax hike news shock during low-attention states brings up real GDP. This impact reaches its maximum after five quarters and subsequently transitions into a negative effect. In contrast, tax news shock does not have a sizeable real impact with high attention. As further shown in the right panel, the differential responses or real GDP between high and low attention states are statistically significant and economically meaningful. These results are consistent with those of An et al. (2022), even though their attention measure and shock variable totally differ from mine.

To the best of my understanding, these findings present novel contributions to the existing body of knowledge on tax policy and inattention. When individuals face resource constraints, such as limited attention, they may struggle to thoroughly analyze all relevant information. As a result, they employ simplified decision-making strategies called heuristics, rationally allocating their inattention. Heuristics, such as the anchor heuristic (Blaufus et al., 2013), reduce complexity and cognitive load, thus mitigating the information-processing costs borne by decision-makers (Gigerenzer and Gaissmaier, 2011). In the context of taxation, the anchor heuristic is a reasonable approach, wherein individuals establish an initial value, or anchor, to estimate their tax burden. However, this heuristic has a drawback in that the anchor is consistently given excessive weight, and additional information is not adequately incorporated, leading to tax biases.

The selection of the anchor is often influenced by either the initial information individuals encounter (Hogarth and Einhorn, 1992) or information perceived as most important (Yadav, 1994). Subsequently, the anchor value is insufficiently adjusted based on subsequent or less significant information. Tax-related news plays a crucial role as part of this information, with news outlets frequently highlighting the magnitude of potential tax reforms in terms of rate cuts or increases, typically in their headlines or initial paragraphs. Indeed, Blaufus et al. (2013) and Amberger et al. (2023) argue that the readily available tax rate, rather than the tax base (e.g., deductible expenses), acts as the anchor. Thus, decision-makers who are rationally inattentive may base their judgments on tax rates while neglecting tax bases. This tendency can be attributed to the greater complexity involved in assessing the economic consequences of altering the tax base, compared to evaluating the impact of changing tax rates. These disparities in complexity are relevant to tax-related decisions since decision-makers often encounter the need to evaluate substantial amounts of information within limited timeframes. Consequently, overweighting the effects of tax rates relative to tax bases in tax decisions could lead to tax-rate bias, misperceive the actual tax burden, and alter the macroeconomic implications of tax policy changes when inattention towards tax policy prevails.

Figure 18: Output Response to a Tax News Shock: States of EPU (Red: High EPU States)

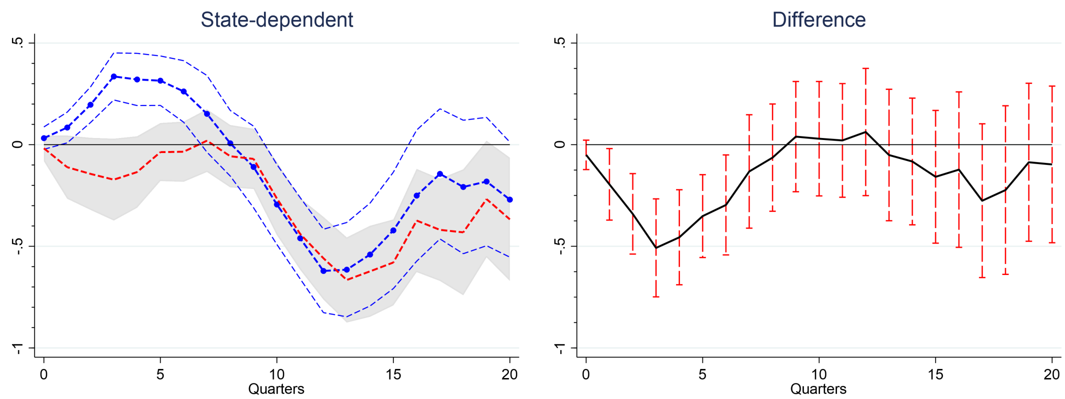

Note: The blue dashed line with circles is the response in low-EPU states and the red dashed line is the response in high-EPU states; 90 percent confidence intervals (e.g., red error bars and shaded areas) constructed by the Newey-West method are shown; The right panel shows the estimated difference between state-dependent responses.

4.3.3 Tax News Shock During Different States of Economic Uncertainty

Inspired by Jerow and Wolff (2022) who study the uncertainty-dependent output effects of government spending shocks, I empirically show that, as shown in Figure 18, a tax hike news shock only has a short-term expansionary output effect when uncertainty is low (i.e., less than the sample median). During periods of heightened uncertainty, the household elects to save via the risk-free asset, thus significantly reducing the value of capital. This decline in the value of capital translates to a decline in net worth for entrepreneurs. Depleted net worth results in less favorable borrowing terms and thus, a decline in the entrepreneurs' conversion of capital from raw to effective which is available for production. Therefore, anticipation effects of a tax hike could be muted to some extent when uncertainty is high.

4.3.4 Tax News Shock During Different States of Fiscal and Monetary Policy Mix

A recent paper by Ascari et al. (2022) studies how the impulse response of output to a shock to Ramey's (2011) defense news variable, a measure of anticipated government spending, differs during fiscally-led regime (the Great Inflation from 1960Q1 to 1979Q2) and monetary-led regime (the Great Moderation from 1984Q1 to 2007Q2), and they find that output increases during the Great Inflation but decreases during the Great Moderation. The fiscal authority held the dominant role in shaping economic policy from the late 1950s up until the time Paul Volcker was appointed as Chairman of the Federal Reserve. While a significant shift in monetary policy occurred towards the end of 1979 with Volcker's appointment, the fiscal authority did not adjust its approach until the end of 1981, following the election of Ronald Reagan. Since that point, the monetary authority has taken the lead in guiding economic policy (Bianchi and Ilut, 2017).

I investigate how the anticipation effects of tax foresight could differ in these two periods. I run the linear LP regression separately for each sample period, and Figure 19 shows that a tax hike news shock has a very strong and positive short-term anticipation effect on output in a fiscally-led regime, and the effect is muted in a monetary-led regime. During the period of high inflation, purchasing power of money could be eroded. Individuals and businesses are more likely to accelerate spending and investments in anticipation of tax hikes, as they sought to preserve their purchasing power and avoid the higher taxes that might further erode their real income. In addition, in a fiscally-led regime when inflation is high, government often adjust tax policies to counter high inflation. This may increase the likelihood of tax foresight, as people have more opportunities to anticipate and respond to potential tax changes.

Figure 19: Output Response to a Tax News Shock: Fiscal and Monetary Policy Mix

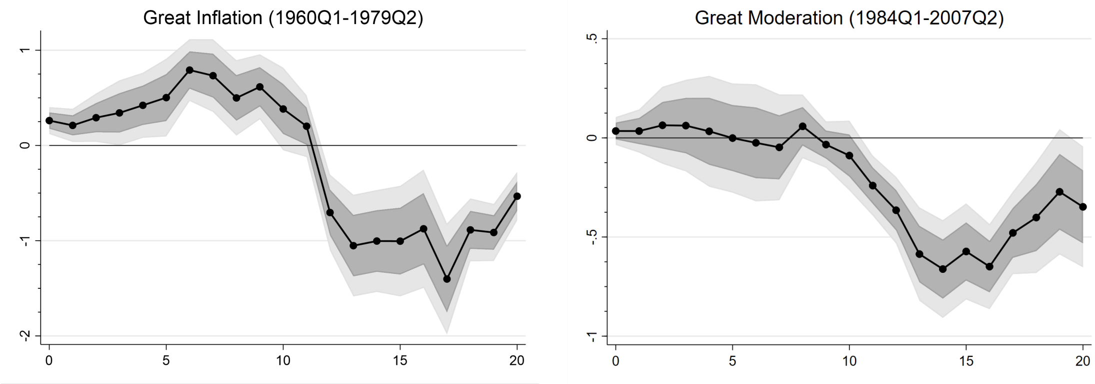

(a) (b)

Note: Shaded areas are the 90 and 68 (darker) percent confidence intervals constructed by the Newey-West method.

In contrast, one of the key features of the Great Moderation was lower and more stable inflation rates compared to the preceding Great Inflation period. With less concern about the erosion of purchasing power, households and firms were less motivated to alter their actions in anticipation of tax hikes. In addition, during the Great Inflation, the government used fiscal actions to adjust the economy and the Federal Reserve was supposed to support the policy by preventing an increase in market interest rates (Meltzer, 2005; Bianchi and Ilut, 2017). However, during the Great Moderation, improved monetary policy and the central role of the Federal Reserve in managing the economy led to better management of inflation expectations and greater economic stability (Blanchard et al., 2010), therefore reducing the need for people to adapt their behavior in response to anticipated tax changes.

4.3.5 Confidence Intervals Using Bootstrap Method

Instead of using the Newey-West method, I adopt the approach proposed by An et al. (2022) for computing confidence intervals, which employs the bootstrap method. It is possible that the data does not adhere to a normal distribution, or there may be uncertainty regarding whether the sample population conforms to such a distribution. In such instances, when determining confidence intervals, bootstrap sampling offers a viable solution. This technique involves randomly sampling the data with replacement, a process that is iterated multiple times (2,000 times according to An et al. (2022)). Each sampled result is treated as an individual data point. After 2,000 iterations, I obtain 2,000 data points, and the distribution of these points can be regarded as an approximation of the population distribution. By virtue of the central limit theorem, it is likely that these 2,000 points adhere to a normal distribution, allowing me to employ the formula for a normal distribution to calculate the confidence interval reliably. Although the confidence intervals now become wider, the statistical significance of response differences remains.

5 Firm-Level Outcomes and Tax News Shocks

In this section, I combine macroeconomic data with micro firm-level data to estimate the dynamic effects of tax news shocks on firms' sales and investment. It is reasonable to expect heterogeneous effects of tax news shocks at the firm level. In fact, Crouzet and Mehrotra (2020) have suggested that the largest 1 percent of companies are less sensitive to aggregate macroeconomic shocks. In this study, I examine the asymmetric effects by firms' market power measured by price markups.

There is growing evidence that price markups, market concentration, and cor-

Figure 20: CI: Bootstrap Method

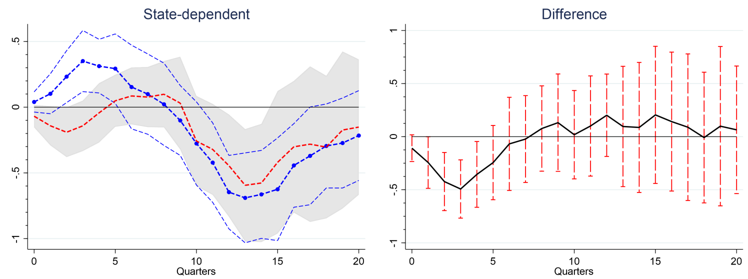

(a) Output response to a tax hike news shock: slack states (red: unemployment rate ≥ 6.5)

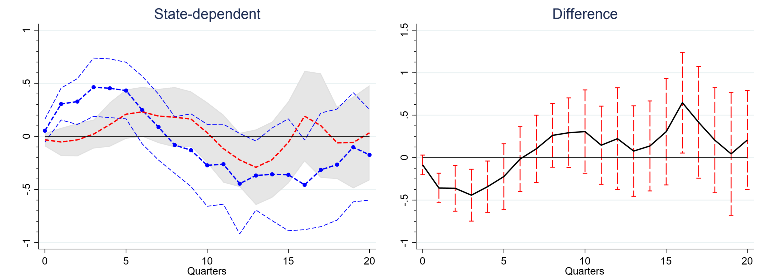

(b) Output response to a tax news shock: states of tax news attention (red: high attention states)

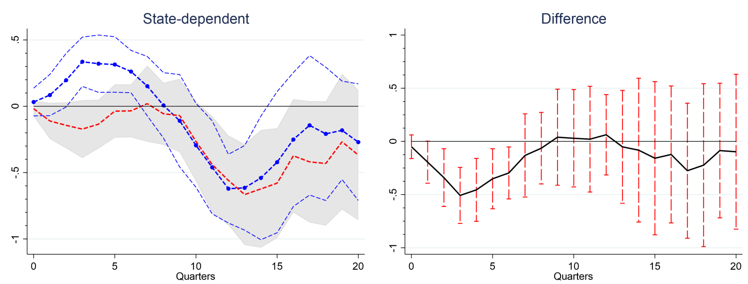

(c) Output response to a tax news shock: states of EPU (red: high EPU states)

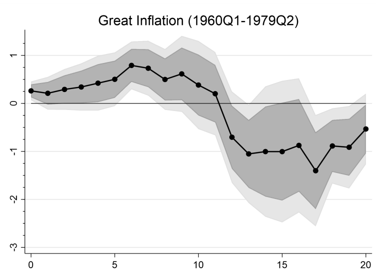

(d)

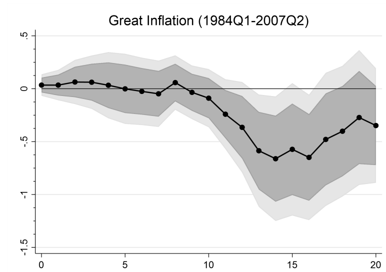

(e)

porate profit rates have increased in the United States (De Loecker and Warzynski (2012); De Loecker et al. (2020); Díez et al. (2021)). Rising market power has the potential to affect the transmission of monetary policy, depending on market structure and the source of firms' market power (Baqaee et al. (2021)). At the most basic level, a higher markup, as the flip side of a lower elasticity of demand, could affect, and under certain conditions, mitigate or amplify the response of a firm's output to shocks such as monetary policy shocks that affect input costs (Syverson (2018)). Furthermore, in the presence of financial frictions, higher corporate profits associated with higher markups might protect firms from shocks to external funding conditions, allowing them to continue financing their working capital or non-pledgeable innovation-enhancing investments (Aghion et al. (2019)). This could also alter a firm's output response to monetary policy shocks. The existing literature largely overlooks the interaction between firm dynamics and tax shocks, particularly through the perspective of market power. I aim to make a small contribution by examining the potential preemptive actions that firms might undertake in anticipation of tax changes.

5.1 Firm-Level Financial Data

I obtain firm-level quarterly variables from Compustat, a panel of publicly traded U.S. firms. The dataset contains detailed and high-quality balance sheet and income statement items. However, it excludes privately held firms. After the data cleaning (including deflating financials using sector-level price indices) based on the literature, I obtain 1,180,799 observations for 16,532 firms over the period of 1961 to 2019.

5.2 Firm-Level Price Markups

A firm's markup is defined as the ratio of the price (P) to the marginal cost (MC). Estimating firm markups is empirically challenging for many reasons, one of which being that most firm-level databases do not include information on firm-level prices. Here, I follow De Loecker and Warzynski (2012), who derive the following expression for the markup ($\mu_{i,t}$) from the firm's cost-minimization problem:

$$\mu_{i,t} \equiv \frac{P_{i,t}}{MC_{i,t}} = \frac{\left(\frac{\partial F_{i,t}(\cdot)}{\partial v_{i,t}}\right)\left(\frac{v_{i,t}}{F_{i,t}(\cdot)}\right)}{\frac{P_{i,t}^{v_{i,t}} v_{i,t}}{P_{i,t} Q_{i,t}}} = \frac{\beta_{i,t}^{v}}{\alpha_{i,t}^{v}} \tag{10}$$

where i and t are the subindexes for the firm and year considered, $F_{i,t}(\cdot)$ is the firm's production function, and $v_{i,t}$ refers to any flexible input. The firm's markup is thus estimated as the ratio of the output elasticity of the variable input considered ($\beta_{i,t}^{v}$) to the expenditure share of that input ($\alpha_{i,t}^{v}$). While the latter can be readily computed, the former needs to be estimated, which in turn requires estimating a production function. After cleaning the Compustat dataset, markups are estimated at the annual level using the same code as Díez et al. (2021), who, in turn, apply the De Loecker and Warzynski (2012) methodology. I use the baseline markup in these papers, which relies on a Cobb-Douglas production function estimated using the control function approach of Ackerberg et al. (2015) and considering the cost of goods sold (COGS) as the variable input. Estimation details can be found in the Appendix.

As Figure 21 reports, markups of U.S. (public) firms have increased by a sales-weighted average of about 50% during the last forty years, and the rise in average markups is associated with especially large increases among the highest markup firms. As argued by Díez et al. (2021), high-markup firms are also highly performing firms. Conditional on size, they are 20% more productive, report 3% higher profits, and are more likely to spend on intangible assets.

Figure 21: Evolution of Estimated Markups

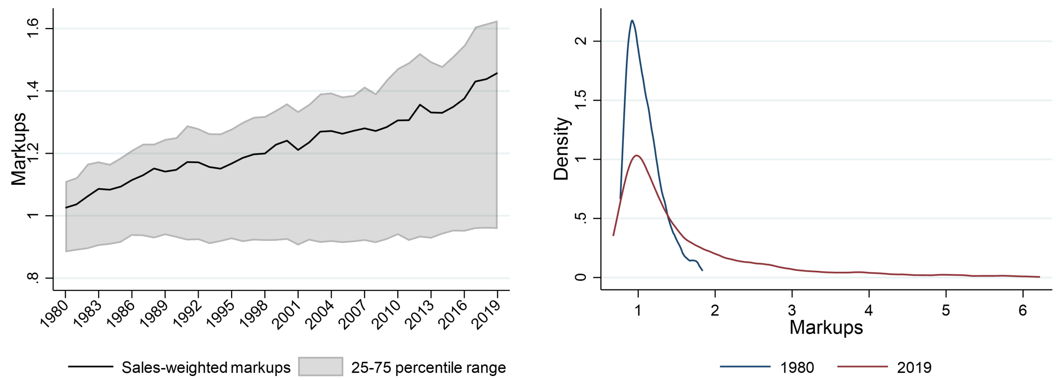

Therefore, the effects of tax news shocks are likely to be heterogeneous among the forward-looking firms with different levels of market power.

5.3 Empirical Framework

To estimate the dynamic effects of tax news shocks on firm's real sales and investment, I again use the local projection method. This approach is quite flexible since it allows estimating impulse response functions on firm-level panel data while controlling for a constellation of fixed effects. Moreover, it is particularly suitable for estimating non-linear effects in the response of the key variables of interest.

Following Duval et al. (2021) and Cloyne et al. (2023), I begin with the basic unconditional specification as follows:

$$y_{i,t+h} - y_{i,t-1} = \alpha_i^h + \beta^h NS_t + \rho^h(L)X_{i,t-1} + \epsilon_{i,t+h} \qquad (11)$$

where $h = 0, \dots, 12$, t refers to quarters, $y_{i,t}$ is the log of real sales (or investment), and NS_t denotes tax news shocks. α_i^h are firm fixed effects, included to absorb the effect of time-invariant firms' characteristics. X is a vector of controls including four lags of the dependent variable, macroeconomic controls (such as output growth, tax revenue growth, government spending growth, employment growth, inflation, debt ratio, and interest rate), and tax news shocks, as well as time-varying firm characteristics (such as size, age, leverage, asset liquidity and tangibility, Tobin's q). β^h denotes the response to a tax (hike) news shock at each horizon (quarter) h.

Estimating the unconditional response of firms' real sales and real investment could provide indirect reassurance regarding the validity of my tax news shocks and compare the response of firms to the aggregate output (and investment) response obtained using national account data in the previous section. In addition, the average impulse response could also provide a benchmark against which I can evaluate the contribution of the response of each group (i.e., low- and high-markup firms).

The second specification allows the response to vary across firms depending on their markup (or exposure to government purchases) level. I closely follow the approach adopted by Duval et al. (2021) and Cloyne et al. (2023) to estimate non-parametrically the heterogeneous effects of tax news shocks across different levels of firm's market power, and estimate the following specification with a forecast horizon of 12 quarters:

$$y_{i,t+h} - y_{i,t-1} = \alpha_i^h + \sum_{g \in G} \beta_g^h \mathbb{1}(i \in g) NS_t + \rho^h(L) X_{i,t-1} + \epsilon_{i,t+h} \tag{12}$$

where $\mathbb{1}$ is an indicator variable, which takes value 1 if the firm characteristic falls in a particular bin of the distribution (i.e., $g \in G$), which I will refer to as the firm's group. Specifically, firms are put into three bins: bottom 25 percentile, top 25 percentile, and middle (between the 25th and 75th percentiles) of the distribution of markup levels. To mitigate endogeneity concerns stemming from the fact that firm markups might respond to tax news shocks, I use for each firm its (time-invariant) average over the entire sample period. Following Cloyne et al. (2023), I include firm fixed effects, which not only absorb any sector fixed effect but also allow me to exploit within-firm variation. As suggested by Cloyne et al. (2023), standard errors are clustered by firm and time using the approach in Driscoll and Kraay (1998) for dealing with possible serial correlation in the forecast errors, which is potentially a feature of local projections.

There are pros and cons of using total assets, property, plant & equipment (PPE), and capital expenditures (CapEx) as measures of firm-level investments. Total assets can be a useful measure of a firm's investment level as it includes all the assets that the firm has acquired over time. This can give an overall view of the firm's investment level, but it may not always provide a clear picture of the specific investments made by the firm. For example, some assets may not necessarily be related to investment in productive capacity. PPE is a more specific measure of investment as it focuses on tangible assets such as buildings, machinery, and equipment. This measure is useful for firms that heavily rely on physical assets to produce goods or services. However, it may not capture investments in intangible assets, which can be equally important for some firms. CapEx is another useful measure of investment as it captures the amount of money spent by the firm on acquiring or improving its assets. This measure can provide a more detailed view of the firm's investment level and the specific investments made by the firm. However, it may not capture investments in intangible assets or investments made through other means such as mergers and acquisitions. Since each of these measures has its strengths and weaknesses, I will consider all these three investment measures to analyze the firm-level effects of tax foresight.

5.4 Unconditional Results

Figure 22 below present the impulse response of real sales, real total assets, real PPE, and real CapEx to a positive one-standard-deviation shock to tax news (i.e., a tax hike news shock). Anticipation effects of a tax hike news shock on all the variables of interest are statistically significant and positive. Sales increase in the anticipation of potential tax hike. The firm-level investment (i.e., measures such as total assets, PPE and CapEx) also accelerates after a tax news shock, confirming House and Shapiro (2006)'s argument that there are strong incentives by firms to alter the timing of investment in response to tax changes.

The reaction of these financial metrics to the tax news shock shares a certain degree of consistency with the findings from aggregate macroeconomic data. However, caution should be exercised in interpreting these results due to the dataset's specific nature: it only includes publicly listed firms. Even though my findings may not paint the complete picture of the broader economy, the significant and relatively enduring anticipation effects in the sample when compared to macroeconomic responses point to the influential role of these firms in the economy.

One primary motive for firms boosting investments before a tax hike is to accelerate deductions. By investing in assets or making capital expenditures, firms can increase their tax deductions, effectively reducing taxable income at the lower pre-hike rates (Auerbach and Hassett, 1992). This strategy enables companies to maximize the tax benefits associated with their investments. Moreover, inflation is another factor that also can drive firms to increase investments before a tax hike. Higher taxes may lead to increased prices as businesses may pass on the tax burden to consumers, ultimately causing inflation (Romer and Romer, 2010). By investing in assets or increasing inventory such as pro-

Figure 22: Response of Real Sales and Total Assets to Tax News Shocks: Average Effect

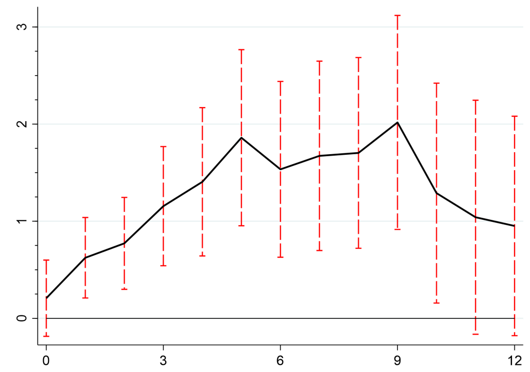

(a) Real sales (%)

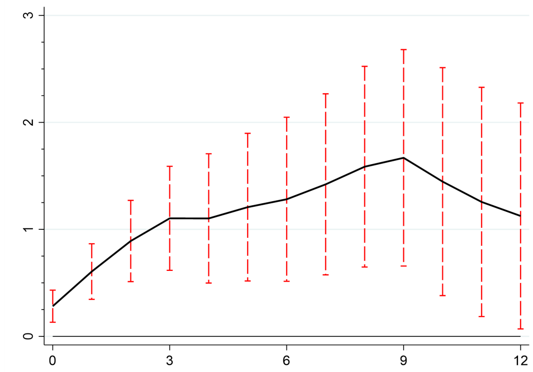

(b) Real total assets (%)

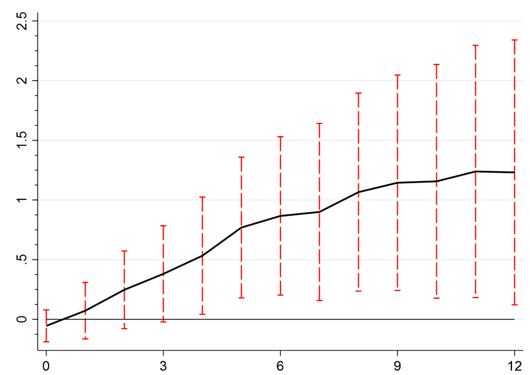

(c) Real PPE (%)

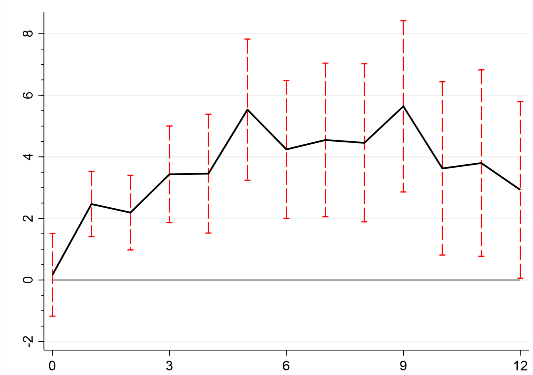

(d) Real CapEx (%)

Note: The solid lines indicate the impulse response functions to a one-standard-deviation tax hike news shock; The vertical dashed lines denote 90 percent confidence bands constructed using the standard errors clustered by firm and time and adjusted for potential serial correlation using the Driscoll-Kraay method; The x-axis denotes quarterly time; The y-axis denotes percentage change.

duction inputs, firms can hedge against potential inflation and protect their purchasing power.

Simultaneously, firms may attempt to increase sales in the short term in anticipation of a tax hike. By generating higher revenues (for example, through promotions that may stimulate consumer demand) at the existing lower tax rates, companies can maximize after-tax profits, positioning themselves advantageously for when the tax increase is implemented (Yagan, 2015). In addition, businesses might try to move income from future periods to the present, recognizing more income at the lower tax rates currently in force before the tax hike. This can be accomplished through various strategies, such as changing accounting methods or accelerating revenue recognition (Auerbach and Hassett, 2002).

5.5 High- vs. Low-Markup Firms

In this section, the influence of a tax hike news shock on firms with differing market power is explored. Drawing from De Loecker, Eeckhout, & Unger's (2020) research on the implications of market power, it is clear that high-markup firms usually hold a significant advantage over their less powerful counterparts due to their inherent capabilities and resources. As suggested by the results (see Figure 23), these firms typically respond more positively to the anticipation of a tax hike policy change than those with less market power. The explanation for this behavior can be rooted in the fact that high-markup firms often possess extensive financial resources, including cash reserves and access to credit. Bates et al. (2009) found that U.S. firms hold significantly more cash than before, which implies that they have more financial flexibility. Jhang et al. (2019) also find that high market power firms can earn abnormal economic rent easily, keep abundant cash or resources on hand and face fewer constraint problems. This allows them to act quickly in anticipation of a tax hike, potentially capitalizing on opportunities that such policy changes may present. Moreover, high-markup firms are better positioned to identify and capitalize on investment opportunities due to their more extensive industry knowledge, better access to information, including the assessment of the potential impact of tax policy changes, and stronger networks than low-markup firms (Haltiwanger et al., 2017). This helps the former make strategic decisions about where to invest in anticipation of a tax hike change. In contrast, low-markup firms may lack the necessary information and resources to make such strategic decisions. My findings show that high-markup firms react to a tax hike news shock more strongly than the low-markup firms.

Interestingly, I find that the difference in sales response between high and low markup firms to a tax hike news shock seems to be relatively weak in comparison to the other variable responses. This may be because consumers often operate within budget constraints and may want to make purchases before the tax hike takes effect to maximize their purchasing power and maintain their standard of living, irrespective of the firm's market power (Klenow and Malin, 2011).

In short, these findings suggest that the response of firms to a tax hike news shock is shaped by their level of market power, with the investments by high-markup firms reacting more positively than low-markup firms. Policymakers should take these differences into account when designing tax policies that affect firms of different sizes and market power.

Figure 23: Response to a Tax Hike News Shock: High- vs. Low-Markup Firms

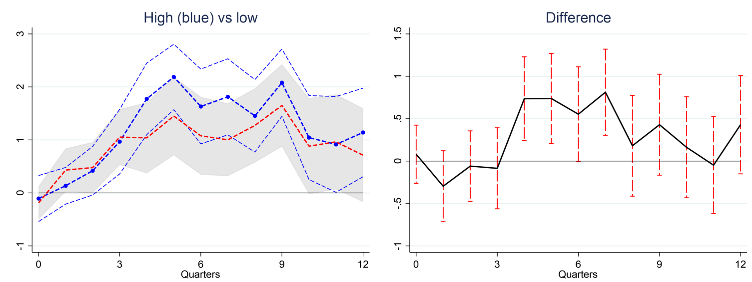

(a) Real sales (%)

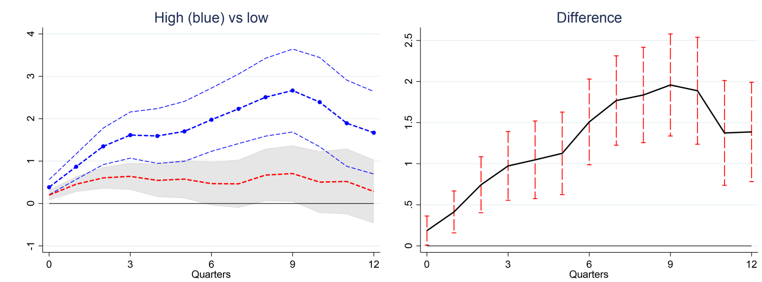

(b) Real total assets (%)

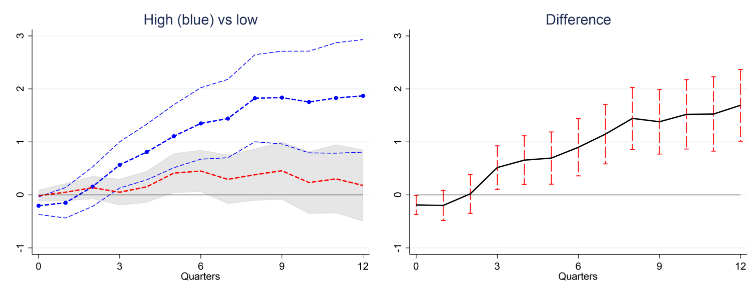

(c) Real PPE (%)

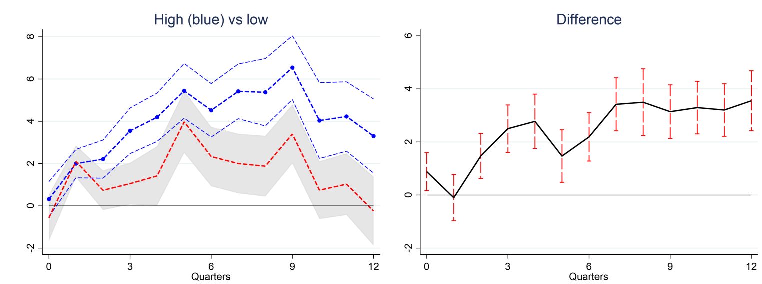

(d) Real CapEX (%)

Note: The blue dashed line with circles is response for high-markup firms and the red dashed line is response for low-markup firms; 90 percent confidence intervals are shown; The x-axis denotes quarterly time; The y-axis denotes percentage change.

6 Tax News and Asset Markets

6.1 Tax News Shock and Asset Prices

Compared to monetary policy, fiscal policy's potential impact on asset markets has received less attention. Discussions have mainly focused on the effects of fiscal policy on macroeconomic variables such as GDP, consumption, and investment, little consideration has been given to its influence on asset markets. This section aims to shed light on how investors in the stock and bond markets perceive and factor in unexpected tax news shocks when making investment decisions.

If markets are efficient, asset prices reflect all information currently available to market participants, especially news concerning the future paths of relevant variables. This hypothesis led Beaudry and Portier (2006) to include stock prices in a VAR in order to capture agents' expectations about future changes in productivity, while Fisher and Peters (2010) use stock prices of government defense suppliers to identify news about future government spending. If the market is efficient in terms of fiscal policy, stock prices should promptly incorporate any changes in fiscal policy as soon as this information becomes public. Ardagna (2009) conducted a study using a panel of OECD countries (including the US) from 1960 to 2002, demonstrating that stock market prices surge during periods of substantial fiscal tightening and decline in times of very loose fiscal policy. Moreover, her research reveals that the results depend on countries' initial fiscal conditions and the nature of fiscal consolidations. Fiscal adjustments implemented during years with high government deficit levels, achieved through reductions in government spending, and leading to a permanent and significant decrease in government debt, are associated with larger increases in stock market prices.

Figure 24: Response of S&P 500 Composite Index to a Tax Hike News Shock

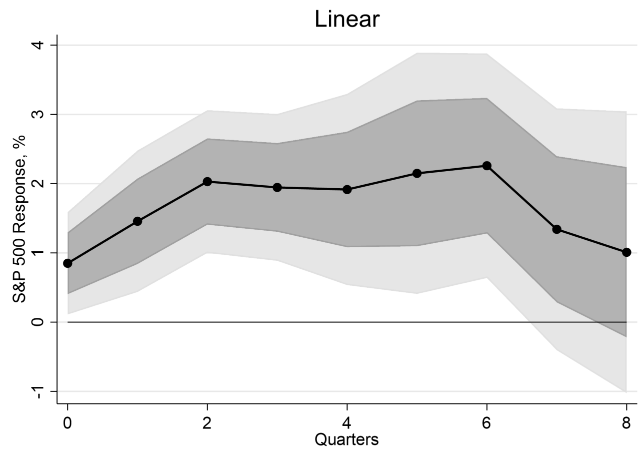

Note: Impulse response to a positive one-standard-deviation shock to tax news; Shaded areas are the 90 and 68 (darker) percent confidence intervals constructed by the Newey-West method.

Building on the local projection model in the previous section, I find that the contemporary unanticipated tax news shock has an immediate and long-lasting impact on stock market return, and this evidence is line with the efficient market hypothesis in a way that a tax news shock was not foreseen by investors, hence generating an immediate reaction in stock market. Figure 24 clearly shows that the S&P 500 composite index reacts positively to information indicating a more restrictive or less expansionary tax policy and may imply that investors prefer fiscal discipline instead of a lax fiscal policy and fiscal consolidation measures that help reduce deficits are perceived as good news.

To see if such a conjecture is sensible, I distinguish the state of economy into high and low federal deficit (as percentage of GDP) based on the sample median, create a state-dependent impulse response chart below, and find that, as shown in Figure 25, the stock market reacts immediately and positively to a tax hike news shock only when deficits are high. These findings are consistent with the literature (e.g., Ardagna (2009); Afonso and Sousa (2011); Stoian and

Figure 25: Response of S&P 500 Composite Index to a Tax Hike News Shock: States of Federal Deficits (Red: High Deficits States)

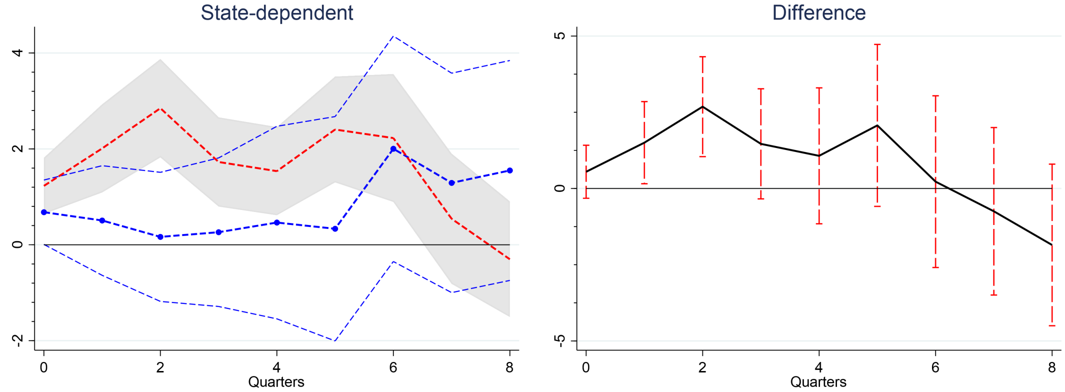

Note: The blue dashed line with circles is the response in low-deficits states and the red dashed line is the response in high-deficits states; 90 percent confidence intervals constructed by the Newey-West method (e.g., red error bars and shaded areas) are shown; The right panel shows the estimated difference between state-dependent responses.

Iorgulescu (2020)).

Movements in credit spreads are thought to contain important signals regarding the evolution of the real economy and risks to the economic outlook, a view supported by the insights from the large literature on the predictive content of credit spreads for economic activity.

Gilchrist and Zakrajšek (2012) construct a credit spread index (henceforth, GZ spread) with high information content for future economic developments that is built from the bottom up, using secondary market prices of senior unsecured bonds issued by a large representative sample of U.S. non-financial firms, and they find that the spread is a highly significant predictor (i.e., lower GZ spread higher future economic activity) of output growth (see Figure 26). In addition, as shown by Philippon (2009), the ability of the bond market to signal future economic activity is more accurate than the stock market. Therefore, I investigate whether the bond market could react positively to a tax hike news shock in anticipation of deficit reduction promoting sustainable economic growth. As

Figure 26: GZ Credit Spread Extended by Giovanni et al. (2016)

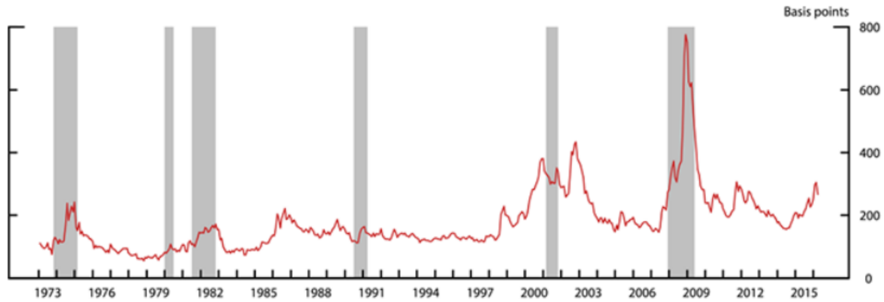

Note: The shaded vertical bars represent the NBER-dated recession. Giovanni et al. (2016)'s calculations based on: Center for Research in Security Prices (CRSP); CRSP/Compustat Merged Database, Wharton Research Data Services (WRDS), and Bank of America Merrill Lynch Bond Indices.

documented by Ji and Qian (2015), a tax hike that improves the government's balance sheet increases the possibility of a government bailout in case of a banking crisis, and such a reduction in the systemic risk can then reduce the risk premium charged by banks, and hence lower the credit spread. Figure 27 shows that investors in the corporate bond market immediately react to a tax hike news shock and expect an acceleration in real GDP, evident by a decrease in the GZ spread on impact.

The GZ spread can be further decomposed into two components: a component capturing the usual countercyclical movements in expected defaults (i.e., a component that captures default risk of individual firms), and a component representing the cyclical changes in the relationship between measured default risk and credit spreads—the so-called excess bond premium that captures investor attitudes toward corporate credit risk–that is, credit market sentiment. In effect, the EBP tries to capture the variation in the average price of bearing U.S. corporate credit risk–above and beyond the compensation that investors in the corporate bond market require for expected defaults, and it can also be interpreted as an indicator of the overall credit supply conditions in the economy. As argued by Gilchrist and Zakrajsek (2012), the EBP is signifi-

Figure 27: Response of GZ Credit Spread to A Tax Hike News Shock

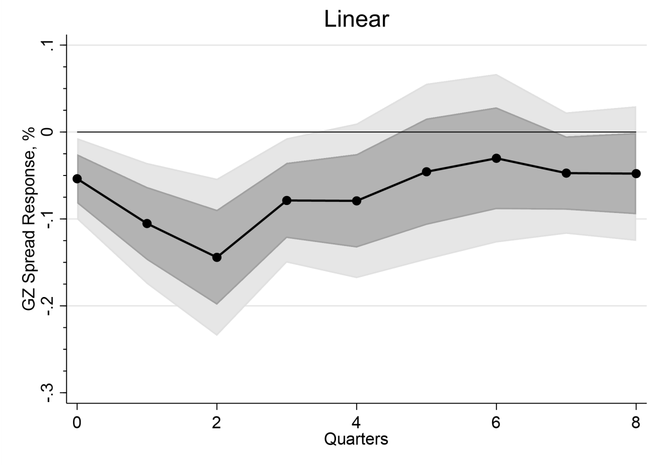

Note: Impulse response to a positive one-standard-deviation shock to tax news; Shaded areas are the 90 and 68 (darker) percent confidence intervals constructed by the Newey-West method.

cantly more informative–in both economic and statistical terms–about future economic activity than a component of expected defaults. I examine the behavior of the excess bond premium, default risk, market value of bank equity, and lending standards in response to a tax hike news shock by replacing each of these indicators in the local projection specification discussed above in place of the GZ spread. The market value of U.S. commercial bank's equity and the Senior Loan Officer Opinion Survey of Bank Lending Practices (SLOOS) are the two indicators for balance sheet conditions of intermediaries extending loans. The SLOOS measures the net percentage of domestic respondents tightening standards for commercial and industry loans. I use the net percentage applicable for loans to medium and large firms. Specifically, the net percentage measures the fraction of banks that reported having tightened ("tightened considerably" or "tightened somewhat) minus the fraction of banks that reported having eased ("eased considerably" or "eased somewhat").

Figure 28 shows that the excess bond premium declines significantly on impact and it is similar to the behavior of the GZ spread. Interestingly, the default

Figure 28: Responses to A Tax Hike News Shock

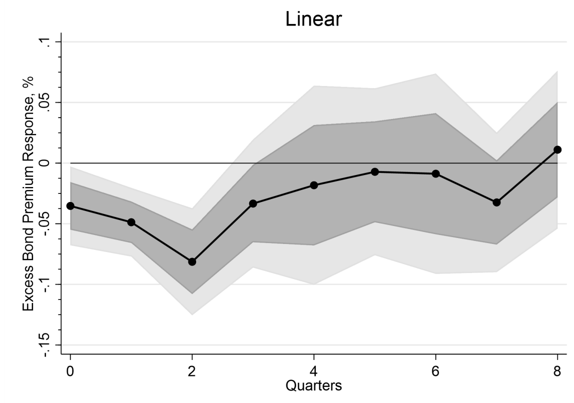

(a)

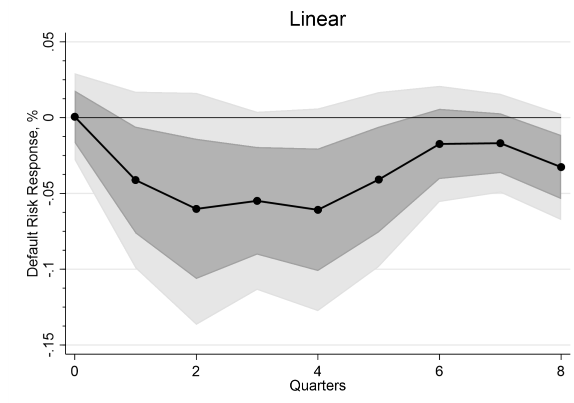

(b)

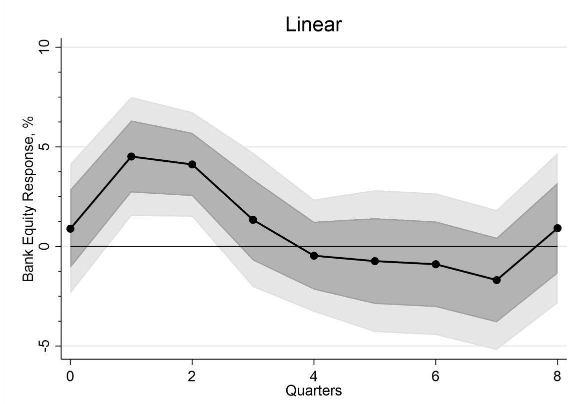

(c)

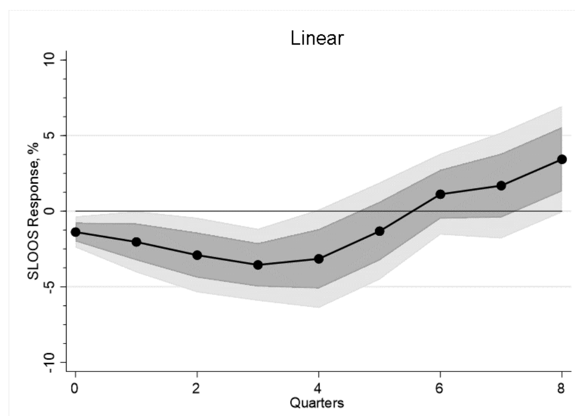

(d)

Note: Impulse response to a positive one-standard-deviation shock to tax news; Shaded areas are the 90 and 68 (darker) percent confidence intervals.

risk component of the GZ spread is, in contrast to the excess bind premium, not reacting significantly in response to the tax news shock. This observation suggests that the variation in the GZ credit spread conditional on the tax news shock is driven by factors mostly related to credit supply conditions. The dynamic response of bank equity is strong and significantly positive. The bank equity response is consistent with the notion that it reflects increased profitability or higher asset valuation in the balance sheet of intermediaries. The response of the SLOOS variable suggests an immediate and significant relaxation of lending standards. The findings related to the joint response of the excess bond premium, bank equity and lending standards are consistent with the evidence reported in Gilchrist and Zakrajsek (2012), where higher profitability of the U.S. financial corporate sector is associated with a reduction in the excess bond premium. Taken together, these results support the hypothesis that balance sheet and more generally credit supply conditions are an important transmission channel for tax news shocks to the real sector.

The results also echo the findings in the previous section. High-markup firms, due to their substantial financial resources and superior information access, are better equipped to respond to tax hike news shocks positively. Meanwhile, news shocks may improve credit supply conditions, which again tend to favor high-markup firms. This is because these firms are often better positioned to access and utilize credit, given their financial robustness, credibility, and resources. This improvement in credit supply conditions may stimulate these firms to augment their investments in anticipation of a tax hike. As such, tax hike news shocks may serve to further consolidate the advantageous position of high-markup firms, not just by presenting strategic investment opportunities but also by enhancing the overall credit conditions.

6.2 Tax News Attention and Stock Price Volatility

An increasing body of empirical evidence suggests that investor attention plays a crucial role in influencing asset prices (Da et al., 2011). Studies have shown that fluctuations in investor attention over time have significant implications for market dynamics. High levels of attention have been associated with buying pressures and sudden price reactions (Barber and Odean, 2008; Barber et al., 2009), while low levels of attention have been linked to underreactions to important announcements (Dellavigna and Pollet, 2009). This highlights the importance of investor attention in driving price movements and the incorporation of new information into asset valuations. Andrei and Hasler (2014) examine the relationship between attention to news and return volatility. They find that heightened attention to news (proxied by Google search data on financial and economic news) can contribute to increased volatility in stock prices. Building on this premise, the present study focuses specifically on the impact of tax news attention on stock volatility.

It is important to consider the underlying factors that may contribute to the relationship between tax news attention and stock volatility. One such factor could be the differential revenue exposure of firms to government contracts, which can vary across industries and companies. The potential for tax reforms, such as tax hikes, can significantly impact investor attention and subsequently contribute to stock volatility for firms highly exposed to government spending.

For instance, increased attention to tax-related news is likely to occur as investors closely monitor any proposed major changes in tax policy. In this context, firms with substantial revenue dependence on government contracts become particularly susceptible to the effects of tax-related news attention. A tax hike, when combined with other austerity measures such as reduced government spending, can adversely affect these firms' revenue streams to a greater extent than the less exposed ones.

The anticipation of lower government spending, in conjunction with tax increases, heightens investor concerns regarding the future profitability and growth prospects of government-contract-dependent firms. Consequently, heightened attention to tax-related news may lead to increased stock volatility for these firms, as investors adjust their expectations and trading strategies based on the potential impact of tax reforms.

By incorporating the revenue exposure aspect into the analysis of tax news attention and stock volatility, this study recognizes the heterogeneity of firms' vulnerability to tax-related changes. It further contributes to our understanding of how investor attention to tax news can interact with specific industry dynamics, government spending, and tax reforms to shape stock price movements and overall market volatility.

I closely follow Baker et al. (2016) and use their compiled firm-level option-implied stock price volatility and preferred firm-level exposure to government purchases of goods and services that may lessen the scope for reverse causality. Tax cut and hike news intensity obtained from the previous section will be used as an indirect proxy for measuring investor attentions. All variables are at quarterly frequency.

Table 4 displays results from regressing firms' stock price volatility on tax news attention. Columns (1) and (2) contain the key results. I add a full set of firm and time fixed effects to control for unobserved factors that differ across firms and unobserved common factors that vary over time. I interact aggregate tax news attention, tax hike news attention, and tax cut attention with Baker et al. (2016)'s measure of exposure to government purchases. Column (1) tests whether implied volatility at firms with greater exposure to government purchases covaries more strongly with general tax news attention. I find very strong evidence for this. The coefficient of 0.245 on the "Agg. Tax News Attention * Intensity" suggests that for every 1% increase in general tax news attention a firm with a 50% government revenue share would see its stock volatility rise by about 0.12%. By decomposing news attention into tax cuts and hikes, column (2) clearly suggests that volatility at highly exposed firms responds to tax hike news attention much stronger than to cut news, implying that investors may anticipate reduced government spending alongside the tax increase. Darrat (2008) explores the causal relation between government spending and taxation and finds that raising taxes (working primarily through aroused public awareness, which can incite public demands to curtail unnecessary expenditures) provokes spending cuts. This expectation of lower revenue potential can lead to heightened concerns about the financial health and future prospects of these firms, resulting in increased stock volatility. Columns (3) to (7) presents a range of additional robustness results. Columns (3) and (4) add EPU and tax EPU that are interacted with the exposure intensity, yielding similar results on the relationship between attention and volatility but suggesting that none of the EPU measures has a large and statistically significant impact on volatility. Columns (5) and (6) consider realized volatility and 182-day implied volatility, and the results are still robust. Column (7) replaces my aggregate tax news attention with the WSJ-based general tax news attention from Bybee et al. (2023), and still the coefficient of the attention-exposure interaction term is line with my expectation. While tax news attention is not a major focus of Bybee et al.'s (2023) study, their findings reveal that the news attention given to "oil drilling" and "oil market" topics in the WSJ closely coincide with the volatility in crude oil prices.

Table 4: Option-Implied Stock Price Volatility and Tax News Attention

	(1) 30-d imp	(2) 30-d imp	(3) 30-d imp	(4) 30-d imp	(5) Realized vol	(6) 182-d imp	(7) 30-d imp
Agg. Tax News Attention × intensity	0.245***						
	(0.0418)						
Federal purchases/GDP × intensity	-32.76**	-40.48***	-42.12***	-41.63***	-32.26**	-37.52***	-24.73***
	(13.32)	(13.32)	(12.54)	(12.58)	(15.83)	(12.51)	(8.066)
Hike News Attention × intensity		0.371***	0.362***	0.366***	0.386***	0.325***	
		(0.0702)	(0.0710)	(0.0718)	(0.0884)	(0.0677)	
Cut News Attention × intensity		0.150***	0.123**	0.134**	0.172***	0.123**	
		(0.0486)	(0.0568)	(0.0549)	(0.0613)	(0.0478)	
Log(EPU) × intensity			0.0732				
			(0.0774)				
Log(Tax EPU) × intensity				0.0254			
				(0.0423)			
WSJ Tax News Attention × intensity							0.370***
							(0.118)
Constant	-1.064***	-1.057***	-1.061***	-1.058***	-1.244***	-1.074***	-0.760***
	(0.0136)	(0.0135)	(0.0161)	(0.0140)	(0.0161)	(0.0126)	(0.00638)
Firm and time effects	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Observations	136578	136578	136578	136578	136578	136578	136578

Note: Standard errors in parentheses. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$. The sample contains 136,578 observations on 5,460 firms from 1996 to 2012. The dependent variable is the natural log of firm-level implied volatility, averaged over all days in the quarter. Intensity is the firm's exposure to federal purchases of goods and services. Federal purchases-to-GDP ratio is from NIPA tables. Log(EPU) is the log of the newspaper-based economic uncertainty index. Log(Tax EPU) is the log of tax policy uncertainty index. All regressions weighted by the firm's average sales in the sample period. Standard errors based on clustering at the firm level.

7 Conclusion

This study delves into the intricate relationship between fiscal policies, particularly tax changes, and economic behavior under the prism of "fiscal foresight." Utilizing a novel approach based on seeded LDA to quantify tax news intensity, the paper sheds light on how anticipatory behaviors significantly influence economic aggregates, firm-level outcomes, and asset markets.

One of the key contributions of this research is the construction of two distinct indices gauging the intensity of news related to tax hikes and tax cuts. These indices, which harness data from the American Presidency Project, capture the ebb and flow of media attention to tax policies, thereby acting as robust indicators of anticipatory economic behaviors. The LDA-based method not only proves to be more adept than lexicon-based approaches at capturing the nuances in language and subjectivity surrounding tax policy discussions, but also predicts well the existing measures of foresight and tax shocks, thus strengthening the credibility and advantages of the metrics used in the study.

Empirically, this paper confirms significant tax foresight effects. It presents new findings showing that firms with greater market power respond more to anticipated tax changes by quickly adjusting their investments in anticipation of upcoming tax increases. Additionally, firms highly exposed to government purchases experience increased stock price volatility with rising tax news intensity. At the asset market level, the study reveals that investors view tax hike news as a sign of future improvement in the government's balance sheet, contributing to more sustainable economic growth. This perception could lead to an increase in credit supply due to increased possibility of a government bailout in case of a banking crisis.

Looking ahead, my analysis underscores the importance of the government meticulously planning both the timing of tax reforms. For instance, a tax increase aimed at cooling an overheated economy could inadvertently become pro-cyclical if the legislation process is prolonged, allowing forward-looking firms and households ample time to take preemptive actions before the tax changes are implemented.

8 Appendix

8.1 Tax News Shock Identification Using SVAR

The identification of (unanticipated) tax revenue shocks is based on Blanchard and Perotti (2002)'s methodology, a seminal work on fiscal policy structural VAR (SVAR) analysis. The main idea behind their approach is to exploit legislative or decision lags (more than a quarter) that are commonly seen in fiscal policy enactment, so that discretionary fiscal actions (i.e., fiscal adjustments made in response to an unexpected movement in output) could be eliminated if using quarterly data. I augment their three-variable SVAR model to identify tax news shocks by including the NTNI as the last variable following Leeper et al. (2013).

The basic reduced-form 4-lag VAR specification, according to Blanchard and Perotti (2002), is

$$Y_t = C(L)Y_{t-1} + U_t \tag{13}$$

where $Y_t = [T_t, G_t, y_t]'$ is a vector in the logarithms of real per-capita federal net tax revenue, real per-capita government spending, and real per-capita GDP, respectively. $C(L)$ is a lag polynomial and U_t is a vector of reduced-form residuals. I use quarterly data and allow for a linear time trend.

The reduced-form residuals U_t are linearly linked to the underlying structural shocks ε_t:

$$AU_t = B\varepsilon_t \tag{14}$$

where $\mathbb{E}[\varepsilon_t] = 0$, $\mathbb{E}[\varepsilon_t \varepsilon_t'] = I$, and $\mathbb{E}[\varepsilon_t \varepsilon_s'] = 0$ for $t \neq s$. The A matrix imposes the recursive structure, while the diagonal B matrix orthogonalizes the effects of innovations. With 14, the structural form of the VAR can be obtained as follows:

$$AY_t = AC(L)Y_{t-1} + AU_t = AC(L)Y_{t-1} + B\varepsilon_t \tag{15}$$

I obtain the reduced-form residuals $U_t = [u_t^T \; u_t^G \; u_t^y]'$ by estimating 13, and then express the reduced-form residuals as:

$$u_t^T = \alpha_y^T u_t^y + \alpha_G^T \varepsilon_t^G + \varepsilon_t^T \tag{16}$$

$$u_t^G = \alpha_y^G u_t^y + \alpha_T^G \varepsilon_t^T + \varepsilon_t^G \tag{17}$$

$$u_t^y = \alpha_T^y u_t^T + \alpha_G^y u_t^G + \varepsilon_t^y \tag{18}$$

where α_y^T and α_y^G measure the automatic effects of economic activity on spending and taxes under the existing fiscal policy, as well as discretionary fiscal adjustments in response to unanticipated movements in output within the quarter. α_G^T and α_T^G capture how the structural shocks to government spending (i.e., ε_t^G) and tax revenue (i.e., ε_t^T) affect tax revenue and government spending, respectively, within the quarter.

Blanchard and Perotti (2002) exploit the institutional information that fiscal policymakers in the U.S. generally take more than a quarter to respond to an output shock (i.e., it takes time for policymakers to realize a shock to GDP and make fiscal decisions accordingly). As a result, with the use of quarterly data, α_y^T and α_y^G measure only the automatic feedback from real GDP to government spending and tax revenue (i.e., automatic output elasticity of fiscal policy). From their estimates, $\alpha_y^T = 2.08$ and $\alpha_y^G = 0$, as they could not find any automatic feedback from economic activity to government spending.

From Equations 16, 17, and 18, it is evident that u_t^T and u_t^G are correlated with ε_t^y. Therefore, Blanchard and Perotti (2002) compute cyclically adjusted reduced-form residuals for the fiscal variables as follows:

$$u_t^{(T,ca)} \equiv u_t^T - \alpha_y^T u_t^y = \alpha_G^T \varepsilon_t^G + \varepsilon_t^T \tag{19}$$

$$u_t^{(G,ca)} \equiv u_t^G - \alpha_y^G u_t^y = \alpha_T^G \varepsilon_t^T + \varepsilon_t^G \tag{20}$$

and use them as instruments to estimate α_T^y and α_G^y in Equation 18. However, $u_t^{(T,ca)}$ and $u_t^{(G,ca)}$ may still correlate with each other, though neither correlates with ε_t^y. Hence, to identify the system, the ordering of the fiscal variables needs to be decided. If tax decisions come first, then $\alpha_G^T = 0$. If spending decisions come first, then $\alpha_T^G = 0$. Both assumptions give similar results.

To identify the structural tax news shocks, I augment the original SVAR with my quarterly tax news index, N_t (i.e., the NTNI). Following Leeper et al. (2013)'s strategy, tax news shocks have no contemporaneous effect on tax revenues, government spending, and output. Consequently, the tax change news equation added last to the system of reduced-form errors is as follows:

$$u_t^N = \alpha_T^N u_t^T + \alpha_G^N u_t^G + \alpha_y^N u_t^y + \varepsilon_t^N \tag{21}$$

where α_T^N, α_G^N, and α_y^N capture the contemporaneous dependence of tax news on tax revenues, government spending, and output, respectively. Similarly, α_T^N, α_G^N, and α_y^N can be estimated by following the same instrumental variable methods in (7) and (8). Specifically, $u_t^T - \alpha_y^T u_t^y$, $u_t^G - \alpha_y^G u_t^y$, and $u_t^y - \alpha_T^y u_t^T - \alpha_G^y u_t^G$

are used as instruments for estimating α_T^N, α_G^N, and α_y^N, respectively.

Equation 14 can be written in a matrix form, and all the α's can be estimated by the steps mentioned above. The elements on the diagonal of the B matrix are the standard deviations of the ε's (i.e., a 1-standard-deviation shock to each variable). Six restrictions together exactly identify the structural shocks of the system.

$$\begin{bmatrix} 1 & 0 & -2.08 & 0 \\ 0 & 1 & 0 & 0 \\ -\alpha_T^y & -\alpha_G^y & 1 & 0 \\ -\alpha_T^N & -\alpha_G^N & -\alpha_y^N & 1 \end{bmatrix} \begin{bmatrix} u_t^T \\ u_t^G \\ u_t^y \\ u_t^N \end{bmatrix} = \begin{bmatrix} \sigma^T & 0 & 0 & 0 \\ \alpha_T^G & \sigma^G & 0 & 0 \\ 0 & 0 & \sigma^y & 0 \\ 0 & 0 & 0 & \sigma^N \end{bmatrix} \begin{bmatrix} \varepsilon_t^T \\ \varepsilon_t^G \\ \varepsilon_t^y \\ \varepsilon_t^N \end{bmatrix} \tag{22}$$

8.2 Firm-level Price Markup Estimation

I obtain firm-level productivity by closely following Díez et al. (2021). Under the assumption that firms within an industry share the same technology, we estimate an industry-specific (at NAICS 2-digit level) Cobb-Douglas production:

$$q_{it} = \beta_v v_{it} + \beta_k k_{it} + \omega_{it} + \epsilon_{it} \tag{23}$$

where all variables are in logs, q_{it} denotes the log of real sales, v_{it} is the log of real variable input, k_{it} represents the log of real capital stock, ω_{it} refers to productivity, and ϵ_{it} stands for the error term that includes measurement error and unexpected shocks. β_v represents the output elasticity of the variable input and β_k denotes the output elasticity of capital.

Upon estimating the input-output elasticity, we can recover the firm-level productivity estimates as residuals with available firm-level financials. The usual endogeneity concern consists of the potential simultaneity bias resulting from the possibility of correlation between the input choice and the productivity. Following Díez et al. (2021), the methodology addresses this concern through the control function approach by assuming that the demand for the variable input, v, depends on productivity: $v_{it} = f(\omega_{it}, k_{it})$. Inverting it yields $\omega_{it} = f^{-1}(v_{it}, k_{it})$ and thus the production function can then be written as follows:

$$q_{it} = \beta_v v_{it} + \beta_k k_{it} + f^{-1}(v_{it}, k_{it}) + \epsilon_{it} = \phi(v_{it}, k_{it}) + \epsilon_{it} \tag{24}$$

where ϕ can be estimated using any consistent non-parametric estimator.

In the second stage, the method assumes that productivity follows a first-order Markov process: $\omega_{it} = E(\omega_{it}|\omega_{it-1}) + \xi_{it}$, where ξ_{it} stands for an innovation shock to the productivity process. Then solving for ξ_{it} and replacing ω_{it} with the first-stage estimates can have:

$$\hat{\xi}_{it} = \hat{\phi}_{it} - \beta_v v_{it} - \beta_k k_{it} - E\left(\hat{\phi}_{it-1} - \beta_v v_{it-1} - \beta_k k_{it-1}\right) \tag{25}$$

With standard GMM procedures, β_v and β_k can be recovered. By assuming that the variable input v responds to current productivity shocks but its lagged values do not, the following moment condition can be formed:

$$E\left(\begin{bmatrix} \xi_{it} v_{it-1} \\ \xi_{it} k_{it} \end{bmatrix}\right) = 0 \tag{26}$$

from which β_v and β_k can be obtained (both β's are industry-specific). Finally, firm-level markup is obtained as:

$$\mu_{it} = \frac{\beta_s^v}{\alpha_{it}^v} \tag{27}$$

where β_s^v is the output elasticity of the flexible input v in industry s and α_{it}^v is the expenditure share of flexible input v by firm i in period t. Therefore, price markups are the deviation between the elasticity of output with respect to a variable input and that input's share of total revenue.

References

1. Afonso, A., & Sousa, R. M. (2011). What are the effects of fiscal policy on asset markets? *Economic Modelling*, 28(4), 1871-1890.

2. Aghion, P., Akcigit, U., Cage., & Kerr, W.R. (2019). Taxation, corruption, and growth. *European Economic Review*, 111, 238-262.

3. Aghion, P., Farhi, E., & Kharroubi, E. (2019). Monetary policy, product market competition and growth. *Economica*, 86(343), 431-470.

4. Amaglobeli, M. D., Crispolti, M. V., Dabla-Norris, M. E., Karnane, P., & Misch, F. (2018). Tax policy measures in advanced and emerging economies: a novel database. International Monetary Fund.

5. Amberger, H. J., Eberhartinger, E., & Kasper, M. (2023). Tax-rate biases in tax decisions: Experimental evidence. *Journal of the American Taxation Association*, 45(1), 7-34.

6. Andrei, D., & Hasler, M. (2015). Investor attention and stock market volatility. *The review of financial studies*, 28(1), 33-72.

7. An, Z., Abo-Zaid, S., & Simon Sheng, X. (2023). Inattention and the impact of monetary policy. *Journal of Applied Econometrics*.

8. Ardagna, S. (2009). Financial markets' behavior around episodes of large changes in the fiscal stance. *European Economic Review*, 53(1), 37-55.

9. Ascari, G., Beck-Friis, P., Florio, A., & Gobbi, A. (2022). Fiscal foresight and the effects of government spending: It's all in the monetary-fiscal mix. *Journal of Monetary Economics*.

10. Auerbach, A. J., & Hassett, K. A. (1992). Tax policy and business fixed investment in the United States. *Journal of Public Economics*, 47(2), 141-170.

11. Auerbach, A.J., & Gorodnichenko, Y. (2012). Measuring the output responses to fiscal policy. *American Economic Journal: Economic Policy*, 4(2), 1-27.

12. Auerbach, A.J., & Gorodnichenko, Y. (2013). Fiscal Multipliers in Recession and Expansion. In *University of Chicago Press* (pp. 63-102).

13. Auerbach, A. J., & Gale, W. G. (2009). Activist fiscal policy to stabilize economic activity (No. w15407). *National Bureau of Economic Research.*

14. Barber, B. M., & Odean, T. (2008). All that glitters: The effect of attention and news on the buying behavior of individual and institutional investors. *The review of financial studies*, 21(2), 785-818.

15. Barber, B. M., Odean, T., & Zhu, N. (2009). Systematic noise. *Journal of Financial Markets*, 12(4), 547-569.

16. Baker, S. R., Bloom, N., & Davis, S. J. (2016). Measuring economic policy uncertainty. *The quarterly journal of economics*, 131(4), 1593-1636.

17. Baker, S. R., Kueng, L., McGranahan, L., & Melzer, B. T. (2019). Do household finances constrain unconventional fiscal policy?. *Tax Policy and the Economy*, 33(1), 1-32.

18. Barro, R.J., & Redlick, C.J. (2011). Macroeconomic effects from government purchases and taxes. *The Quarterly Journal of Economics*, 126(1), 51-102.

19. Baqaee, D., Farhi, E., & Sangani, K. (2021). The supply-side effects of monetary policy (No. w28345). National Bureau of Economic Research.

20. Bates, T.W., Kahle, K.M., & Stulz, R.M. (2009). Why do US firms hold so much more cash than they used to?. *The Journal of Finance*, 64(5), 1985-2021.

21. Beaudry, P., & Portier, F. (2006). Stock prices, news, and economic fluctuations. *American Economic Review*, 96(4), 1293-1307.

22. Bianchi, F., & Ilut, C. (2017). Monetary/fiscal policy mix and agents' beliefs. *Review of economic Dynamics*, 26, 113-139.

23. Bishop, C.M., & Nasrabadi, N.M. (2006). Pattern recognition and machine learning (Vol. 4, No. 4, p. 738). New York: Springer.

24. Blanchard, O., & Perotti, R. (2002). An empirical characterization of the dynamic effects of changes in government spending and taxes on output. *Quarterly Journal of Economics*, 117(4), 1329-1368.

25. Blanchard, O., Dell Ariccia, G., & Mauro, P. (2010). Rethinking macroeconomic policy. *Revista de Economía Institucional*, 12(22), 61-82.

26. Blei, D.M., Ng, A.Y., & Jordan, M.I. (2003). Latent Dirichlet allocation. *Journal of Machine Learning Research*, 3(Jan), 993-1022.

27. Blaufus, K., Bob, J., Hundsdoerfer, J., Kiesewetter, D., & Weimann, J. (2013). Decision heuristics and tax perception–An analysis of a tax-cut-cum-base-broadening policy. *Journal of Economic Psychology*, 35, 1-16.

28. Blinder, A.S., & Watson, M.W. (2016). Presidents and the US economy: An econometric exploration. *American Economic Review*, 106(4), 1015-1045.

29. Born, B., Müller, G.J., & Pfeifer, J. (2020). Does austerity pay off?. *Review of Economics and Statistics*, 102(2), 323-338.

30. Bybee, L., Kelly, B.T., Manela, A., & Xiu, D. (2023). Business news and business cycles, Journal of Finance.

31. Cloyne, J., Ferreira, C., Froemel, M., & Surico, P. (2023). Monetary Policy, Corporate Finance, and Investment. *Journal of the European Economic Association.*

32. Crouzet, N., & Mehrotra, N.R. (2020). Small and large firms over the business cycle. *American Economic Review*, 110(11), 3549-3601.

33. Christofzik, D. I., Fuest, A., & Jessen, R. (2022). Macroeconomic effects of the anticipation and implementation of tax changes in Germany: evidence from a narrative account. *Economica*, 89(353), 62-81.

34. Da, Z., Engelberg, J., & Gao, P. (2011). In search of attention. *The journal of finance*, 66(5), 1461-1499.

35. DellaVigna, S., & Pollet, J. M. (2009). Investor inattention and Friday earnings announcements. *The journal of finance*, 64(2), 709-749.

36. De Loecker, J., & Warzynski, F. (2012). Markups and firm-level export status. *American Economic Review*, 102(6), 2437-2471.

37. De Loecker, J., Eeckhout, J., & Unger, G. (2020). The rise of market power and the macroeconomic implications. *The Quarterly Journal of Economics*, 135(2), 561-644.

38. Driscoll, J.C., & Kraay, A.C. (1998). Consistent covariance matrix estimation with spatially dependent panel data. *Review of Economics and Statistics*, 80(4), 549-560.

39. Díez, F.J., Fan, J., & Villegas-Sánchez, C. (2021). Global declining competition?. *Journal of International Economics*, 132, 103492.

40. Duval, M.R.A., Furceri, D., Lee, R., & Tavares, M.M. (2021). Market power and monetary policy transmission. *International Monetary Fund.*

41. Dybowski, T. P., & Adämmer, P. (2018). The economic effects of US presidential tax communication: Evidence from a correlated topic model. *European Journal of Political Economy*, 55, 511-525.

42. Eshbaugh-Soha, M. (2013). Presidential influence of the news media: The case of the press conference. *Political Communication*, 30(4), 548-564.

43. Fisher, J.D., & Peters, R. (2010). Using stock returns to identify government spending shocks. *The Economic Journal*, 120(544), 414-436.

44. Fantozzi, D., & Muscarnera, A. (2021). A News-based Policy Index for Italy: Expectations and Fiscal Policy.

45. Gertler, M., & Gilchrist, S. (1994). Monetary policy, business cycles, and the behavior of small manufacturing firms. *Quarterly Journal of Economics*, 109(2), 309-340.

46. Gilchrist, S., & Zakrajšek, E. (2012). Credit spreads and business cycle fluctuations. *American Economic Review*, 102(4), 1692-1720.

47. Gigerenzer, G., & Gaissmaier, W. (2011). Heuristic decision making. *Annual review of psychology*, 62, 451-482.

48. Görtz, C., Tsoukalas, J.D., & Zanetti, F. (2022). News shocks under financial frictions. *American Economic Journal: Macroeconomics*, 14(4), 210-243.

49. Guajardo, J., Leigh, D., & Pescatori, A. (2014). Expansionary austerity? International evidence. *Journal of the European Economic Association*, 12(4), 949-968.

50. Haltiwanger, J., Jarmin, R.S., & Miranda, J. (2017). Who creates jobs? Small versus large versus young. *Review of Economics and Statistics*, 99(2), 187-198.

51. Hamilton, J. D. (2018). Why you should never use the Hodrick-Prescott filter. *Review of Economics and Statistics*, 100(5), 831-843.

52. Hansen, S., McMahon, M., & Prat, A. (2018). Transparency and deliberation within the FOMC: A computational linguistics approach. *The Quarterly Journal of Economics*, 133(2), 801-870.

53. Hayo, B., & Mierzwa, S. (2022). Legislative tax announcements and GDP: Evidence from the United States, Germany, and the United Kingdom. *Economics Letters*, 216, 110548.

54. Herrera, A. M., & Rangaraju, S. K. (2019). The quantitative effects of tax foresight: Not all states are equal. *Journal of Economic Dynamics and Control*, 107, 103726.

55. House, C.L., & Shapiro, M.D. (2006). Phased-in tax cuts and economic activity. *American Economic Review*, 96(5), 1835-1849.

56. Hogarth, R. M., & Einhorn, H. J. (1992). Order effects in belief updating: The belief-adjustment model. *Cognitive psychology*, 24(1), 1-55.

57. Huang, A.H., Lehavy, R., Zang, A.Y., & Zheng, R. (2018). Analyst information discovery and interpretation roles: A topic modeling approach. *Management Science*, 64(6), 2833-2855.

58. Jagarlamudi, J., Daumé III, H., & Udupa, R. (2012). Incorporating lexical priors into topic models. In *Proceedings of the 13th Conference of the European Chapter of the Association for Computational Linguistics* (pp. 204-213).

59. Jerow, S., & Wolff, J. (2022). Fiscal policy and uncertainty. *Journal of Economic Dynamics and Control*, 145, 104559.

60. Jhang, S.S.S., Ogden, J.P., & Suresh, N.C. (2019). Operational and financial configurations contingent on market power status. *Omega*, 88, 91-109.

61. Ji, K., & Qian, Z. (2015). Does tax policy affect credit spreads? Evidence from the US and UK. *Journal of Macroeconomics*, 43, 318-329.

62. Jordà, Ò. (2005). Estimation and inference of impulse responses by local projections. *American Economic Review*, 95(1), 161-182.

63. Kalcheva, I., Plečnik, J.M., Tran, H., & Turkiela, J. (2020). (Un)intended consequences? The impact of the 2017 tax cuts and jobs act on shareholder wealth. *Journal of Banking & Finance*, 118, 105860.

64. Kaplan, S.N., & Zingales, L. (1997). Do investment-cash flow sensitivities provide useful measures of financing constraints? *Quarterly Journal of Economics*, 112(1), 169-215.

65. Klein, M. (2017). Austerity and private debt. *Journal of Money, Credit and Banking*, 49(7), 1555-1585.

66. Klenow, P.J., & Malin, B.A. (2011). Microeconomic evidence on price-setting. In *Handbook of Monetary Economics* (Vol. 3, pp. 231-284). North-Holland.

67. Leeper, E.M., Walker, T.B., & Yang, S.C.S. (2013). Fiscal foresight and information flows. *Econometrica*, 81(3), 1115-1145.

68. Larsen, V. H., & Thorsrud, L. A. (2019). The value of news for economic developments. *Journal of Econometrics*, 210(1), 203-218.

69. Łyziak, T., & Sheng, X. S. (2018). Disagreement in consumer inflation expectations. *Journal of Money, Credit and Banking*.

70. Mertens, K., & Ravn, M.O. (2012). Empirical evidence on the aggregate effects of anticipated and unanticipated U.S. tax policy shocks. *American Economic Journal: Economic Policy*, 4(2), 145-181.

71. Meltzer, A. H. (2005). Origins of the great inflation. Federal Reserve Bank of St. Louis Review, 87(2), 145-75.

72. Mertens, K., & Ravn, M.O. (2014). A reconciliation of SVAR and narrative estimates of tax multipliers. *Journal of Monetary Economics*, 68, S1-S19.

73. Newey, W.K., & West, K.D. (1987). Hypothesis testing with efficient method of moments estimation. *International Economic Review*, 28(3), 777-787.

74. Pástor, Ľ., & Veronesi, P. (2020). Political cycles and stock returns. *Journal of Political Economy*, 128(11), 4011-4045.

75. Owyang, M.T., Ramey, V.A., & Zubairy, S. (2013). Are government spending multipliers greater during periods of slack? Evidence from twentieth-century historical data. *American Economic Review*, 103(3), 129-134.

76. Ramey, V.A. (2019). Ten years after the financial crisis: What have we learned from the renaissance in fiscal research? *Journal of Economic Perspectives*, 33(2), 89-114.

77. Ramey, V.A., & Zubairy, S. (2018). Government spending multipliers in good times and in bad: evidence from US historical data. *Journal of Political Economy*, 126(2), 850-901.

78. Romer, C.D., & Romer, D.H. (2009). A narrative analysis of postwar tax changes. Unpublished paper, University of California, Berkeley (June).

79. Romer, C.D., & Romer, D.H. (2010). The macroeconomic effects of tax changes: estimates based on a new measure of fiscal shocks. *American Economic Review*, 100(3), 763-801.

80. Shienbaum, K.E., & Shienbaum, E. (1982). Public perceptions of presidential economic performance: From Johnson to Carter. *Presidential Studies Quarterly*, 12(3), 421-427.

81. Stoian, A., & Iorgulescu, F. (2020). Fiscal policy and stock market efficiency: An ARDL Bounds Testing approach. *Economic Modelling*, 90, 406-416.

82. Stock, J. H., & Yogo, M. (2002). Testing for weak instruments in linear IV regression.

83. Syverson, C. (2018). Changing market structure and implications for monetary policy. In Remarks Made at the 2018 Jackson Hole Symposium.

84. Thorsrud, L. A. (2020). Words are the new numbers: A newsy coincident index of the business cycle. *Journal of Business & Economic Statistics*, 38(2), 393-409.

85. Wagner, A.F., Zeckhauser, R.J., & Ziegler, A. (2018). Company stock price reactions to the 2016 election shock: Trump, taxes, and trade. *Journal of Financial Economics*, 130(2), 428-451.

86. Yadav, M. S. (1994). How buyers evaluate product bundles: A model of anchoring and adjustment. *Journal of Consumer Research*, 21(2), 342-353.

87. Yang, S. C. S. (2009). A chronology of federal income tax policy: 1947-2009.

88. Yagan, D. (2015). Capital tax reform and the real economy: The effects of the 2003 dividend tax cut. *American Economic Review*, 105(12), 353.