YES YOUNG ECONOMISTS SEMINAR

HRVATSKA NARODNA BANKA
EUROSUSTAV

19th YOUNG ECONOMISTS' SEMINAR
TO THE 32nd DUBROVNIK ECONOMIC CONFERENCE

May 28 – 29, 2026, Dubrovnik, Croatia

Javier Perez-Cano

The Macroeconomic Effect of Fiscal Announcements: News-Based Evidence from European Countries

Draft version
Please do not quote

The Macroeconomic Effect of Fiscal Announcements: News-Based Evidence from European Countries

Javier Perez-Cano*

January 21, 2025

Current version

Abstract

This paper studies the macroeconomic effects of fiscal announcements in Europe, distinguishing between spending and taxes. To do so, it builds two novel news-based indicators of announcements about public spending and taxation for Spain, France, Italy, and Germany (1999–2025), and estimates their effects using pooled and country-level Bayesian VARs with recursive identification. The results reveal a pronounced asymmetry. Spending announcements are followed by higher expenditure and sizeable, persistent expansions in GDP and private demand, while tax announcements are followed by higher tax revenues and contractions in activity of a similar magnitude. Moreover, comparing announcement and implementation dynamics suggests that output responds more strongly when policy is communicated rather than when it is implemented, indicating that a substantial share of fiscal transmission operates through expectations.

Keywords: fiscal stabilization policies; macroeconomic uncertainty; sentiment analysis; large languague models clasification; euro area.

JEL Classification: C89; E60; E62.

*Email: javier.perezcano@ua.es Universidad de Alicante.

I am especially grateful to my advisor Gabriel Pérez-Quirós for his guidance and continuous support. I also thank Julien Pinter and the participants of the Working Group at the University of Alicante for their helpful comments and discussions. The project that gave rise to these results received the support of a fellowship from Fundación Ramón Areces.

1 Introduction

Governments frequently signal their intentions regarding future fiscal policy well before any concrete measure is implemented. These signals—which may take the form of specific announcements, such as a reduction in income tax or an increase in defence spending, but also more general statements, plans, parliamentary debates, ministerial leaks, or rumours about upcoming measures—are closely monitored by households, businesses, and financial markets and are therefore likely to generate immediate economic effects.

In this regard, national newspapers play a central role in disseminating both official announcements and broader fiscal intentions, allowing economic agents to observe government signals in real time and adjust expectations at the moment policies are communicated, well before their approval or implementation. Because fiscal announcements differ in their nature—particularly between public spending and taxation—agents may react asymmetrically to different types of fiscal news, with effects that may also vary across countries depending on institutional and political factors.

The empirical literature on fiscal policy shocks has used fiscal news primarily as an *identification device*—a means to date when agents learn about future policy—rather than as an object of interest in itself. Ramey (2011) uses defence news to identify when agents anticipate future spending, but the estimated multiplier still refers to the output effect of spending *eventually executed*. Similarly, Romer and Romer (2010) and Mertens and Ravn (2012) use narrative records of legislated tax changes, Ben Zeev and Pappa (2017) exploit anticipated defence shocks, and Christofzik et al. (2022) and Hayo and Mierzwa (2022) separate announcement dates from implementation dates—yet in all cases the object of interest remains the macroeconomic impact of the implemented policy change. The debate on fiscal multipliers extends to whether they vary with the exchange rate regime and trade openness (Ilzetzki et al., 2013; Corsetti et al., 2014), the state of the business cycle (Auerbach & Gorodnichenko, 2012; Ramey & Zubairy, 2018; Ghassibe & Zanetti, 2022), the monetary policy regime (Christiano et al., 2011; Cloyne et al., 2020; Klein & Winkler, 2021), or the level of public debt (Nickel & Tudyka, 2014; Huidrom et al., 2020). This paper asks a different question: does fiscal communication itself generate macroeconomic effects, regardless of whether the announced measures are subsequently implemented?

A growing body of evidence suggests that it does. Fiscal announcements move sovereign spreads and asset prices (Falagiarda & Gregori, 2015; Afonso et al., 2020), consolidation announcements reduce confidence before any policy is enacted (Beetsma et al., 2015, 2021), and the transmission of fiscal shocks depends on the clarity of government communication (Ricco et al., 2016). On the tax side, Dybowski and Adämmer (2018) document that presidential communication about tax policy in the US produces real effects, and Hayo and Mierzwa (2022) show that GDP contracts following legislative tax announcements, before the tax changes are implemented. At the household level, Gallegos et al. (2025) exploit the exogenous timing of fiscal announcements relative to survey participation in the four largest euro area economies and show that households revise their beliefs immediately: inflation perceptions rise while unemployment perceptions fall, consistent with a signaling channel through which fiscal communication reveals information about economic conditions. From a theoretical perspective, Melosi et al. (2022) show that fiscal announcements carry signaling effects—announcing a large stimulus may reveal the government's pessimism about the severity of a recession, impairing its stabilizing impact—and provide supporting evidence from Japanese fiscal packages. These findings are consistent with models of fiscal foresight in which agents adjust behaviour as soon as news about future policy becomes available (Blanchard et al., 2013; Leeper et al., 2013; Ascari et al., 2023), and with recent NLP-based evidence on fiscal announcements shocks (Latifi et al., 2024; Smets & Trienens, 2025). However, most existing studies rely on institutional sources—budget documents, legislative records, or professional forecasts—rarely distinguish between spending and tax announcements, and provide limited cross-country evidence.

This paper proceed as follows. Daily articles reporting fiscal announcements or rumors are collected from Factiva Dow Jones and processed using a classification procedure based on large language models (LLMs), which assign fiscal labels to each news item. This approach is embedded in a rapidly expanding literature that exploits textual data to extract economic signals not observable in standard statistical sources (Larsen et al., 2021; Angelico et al., 2022; Aruoba & Drechsel, 2024; Latifi et al., 2024; Granziera et al., 2025; Tillmann, 2025). In particular, each article is classified according to the direction of the fiscal signal—expansionary, contractionary, or neutral—separately for public spending and taxation, as well as by institutional level and type of communication. These dimensions are subsequently used to assign differentiated weights reflecting the relative fiscal relevance of each item. Aggregating the scores of all classified news items at quarterly frequency yields the *GovExpSent* index for public spending and the *TaxSent* index for taxation.

The fiscal announcement indices are then incorporated into a battery of Bayesian vector autoregressive (VAR) models. The models also include variables such as real government expenditure or tax revenues, real GDP, and real private consumption. The aim is to assess the macroeconomic impact of these fiscal signals independently of whether and when the announced measures are eventually implemented. Following the pioneering work of Blanchard and Perotti (2002), a large literature uses a recursive identification scheme in order to estimate the causal effects of government spending. This draws on the notion that within a quarter government spending is predetermined such that a feedback from GDP or other macro aggregates on the level of government spending should be excluded. The baseline series of fiscal sentiment lends itself to a natural extension of this identification scheme: sentiment is ordered last such that a change in sentiment cannot contemporaneously drive government spending or tax revenues. Sentiment, on the other hand, can immediately respond to economic developments. The decomposition of sentiment into exogenous and endogenous fiscal sentiment gives rise to a further extension of this identification: exogenous fiscal sentiment is ordered first, macroeconomic reactions follow, while endogenous fiscal sentiment comes last. Both pooled panel VARs and country-specific BVARs are estimated; the latter complement the pooled analysis and confirm the robustness of the results.

Equipped with the sentiment series, I also revisit the problem of fiscal foresight. Because fiscal policy is subject to decision and implementation lags, economic agents may anticipate future changes in government spending and taxes well before they materialise in the data. As shown by Leeper et al. (2013) in the context of taxation and by Ramey (2011) and Forni and Gambetti (2016) in

the context of government spending this anticipation renders the standard recursive identification of fiscal shocks invalid whenever the econometric model fails to account for the resulting non-fundamentalness of the VAR representation.

From a standard recursively-identified VAR in the spirit of Blanchard and Perotti (2002), I obtain a series of structural government spending shocks and show that fiscal sentiment predicts these shocks several quarters in advance. Building on this evidence, I condition spending shocks on prior fiscal news to disentangle announcement from implementation effects, making it possible to compare the GDP response to the act of announcing fiscal policy with the GDP response to its subsequent execution.

The results reveal a clear asymmetry between announcements regarding government expenditure and those concerning taxes. Announcements signalling higher public spending generate strong, immediate, and persistent expansions in real activity, whereas announcements of tax increases lead to more moderate but sustained contractionary effects. Several extensions show that spending announcements also raise investment, inflation, and business expectations, while tax announcements produce deflationary pressures and a deterioration in confidence. These responses are in line with standard New-Keynesian business cycle models. Furthermore, GDP responds more strongly to announcement shocks than to implementation shocks conditioned on prior information, indicating that a substantial share of fiscal transmission takes place at the communication stage. Hence, supposedly unanticipated spending shocks are in fact anticipated once information from the press is taken into account.

The paper contributes to the literature in two respects. First, it provides new evidence on the average macroeconomic effects of fiscal announcements in the euro area, exploiting the capacity of large language models to classify thousands of newspaper articles into directional fiscal signals. This contrasts with the fiscal multiplier literature, where fiscal news serves only to improve the identification of implemented policy shocks, and with existing announcement studies that rely on institutional sources and focus on individual countries or a single fiscal instrument. Second, by estimating both announcement and implementation effects within the same empirical framework, the analysis provides direct evidence on the relative importance of fiscal communication: GDP responds more strongly when policy is communicated than when it is executed, reinforcing the importance of taking into account the moment when the measure is announced.

The remainder of the paper is organised as follows. Section 2 describes the construction of the fiscal announcement indices. Section 3 presents the empirical strategy and identification approach. Section 4 discusses the main results and robustness exercises. Section 5 revisits the problem of fiscal foresight and contrasts announcement and implementation effects. Section 6 concludes.

2 Construction of Government Expenditure and Tax Sentiment Indices

This section describes the construction of the fiscal sentiment indices on government expenditure, *GovExpSent*, and taxation, *TaxSent*, based on press reports. Section 2.1 introduces the news corpus and the Factiva Dow Jones searches, together with the initial filtering and text-cleaning steps designed to remove information unrelated to each country's fiscal policy. Section 2.2 explains how LLMs classify each news item by fiscal direction, jurisdictional scope, and type of communication, and how these labels are validated against extensive manual coding. Section 2.3 outlines the scoring and aggregation scheme that converts these classifications into quarterly indices and concludes by documenting the main descriptive features of the resulting indices.

2.1 The news-corpus and text-preprocessing

National newspapers act as a privileged channel for the dissemination of government intentions on fiscal matters. By publishing immediately both official announcements and the intentions of those with the power to make or influence fiscal policy, they allow real-time monitoring of when and how fiscal measures are communicated to the public. This includes concrete policies such as a tax cut or an increase in public spending, but also more tentative signals such as plans, debates within government bodies, or rumours and leaks to the press. For ease of exposition, the term *fiscal announcement* is used throughout to encompass all these forms of communication.

One way to obtain these announcements reported in the news is through the Factiva database. Appendix B reports the Boolean queries used to retrieve news on public spending and taxation for each country. These queries combine proximity operators (e.g. `w/5`, `w/10`) with logical connectors to capture a broad set of articles reporting the announcement, discussion, approval, or rejection of fiscal measures. Following Andres-Escayola et al. (2022), the searches are conducted across several high-circulation national newspapers to avoid biases associated with single-source coverage. This constitutes the first pre-filtering step to exclude items unrelated to domestic fiscal policy and is carried out separately for spending and taxation. The first row of each panel in Table 1 reports the number of articles retrieved for each index and country.

Before assigning direction, jurisdiction, and institutional weights using LLMs, an automated pre-filter is implemented to retain only those articles that are strictly relevant to each country's fiscal policy. In doing so, the following news are excluded: items that mention taxes or government expenditure without reporting any government announcement or intention; articles focusing on other policy domains (e.g. labour markets or general macroeconomic conditions); pieces discussing economic growth, confidence indicators or broad economic trends; coverage of fiscal events outside the country of interest; event announcements or promotional content (e.g. sponsored supplements or advertorials); and exact or near-exact duplicates published across outlets or editions. This filtering step is fully automated, yet for the full universe of Spanish and German articles—and for a substantial share of the Italian and French datasets—a manual audit is conducted to verify that the algorithm performs the classification correctly. In this way, the automated procedure is calibrated to discard exclusively those items that are genuinely unrelated to fiscal policy—whether on the spending or taxation side—in each country.

Table 1. Number of news used in the construction of fiscal sentiment indices

	Country			
	Spain	Italy	France	Germany
Panel A. Government Expenditure Sentiment Index				
Initial number of news	6,106	4,853	2,381	3,416
Government Expenditure news after prefiltering	2,125	3,715	1,723	2,192
Panel B. Tax Sentiment Index				
Initial number of news	5,298	8,611	2,755	7,309
Tax news after prefiltering	3,167	3,010	2,423	4,772

Notes: The table reports the number of news items used in the construction of the two fiscal sentiment indices. The first row in each panel refers to the number of news downloaded from Factiva; the second row corresponds to the number of news items after filtering them with `gpt-4o-mini` and skimming through them to ensure that duplicates and other news items unnecessary for the analysis are discarded.

Specifically, the country–specific dataset of articles is used, and each item is submitted to the OpenAI API[1], asking the model `gpt-4o-mini` to classify it as fiscal vs. non–fiscal using a short task description and explicit labelling rules. The exact screening prompt is reported in Table A.1. Articles flagged as non–fiscal are discarded. The classifier is applied identically to the two corpora—government expenditure and taxes—and the decision is based on the full body text, not on keywords, as the literature has typically done. To ensure robustness, the API returns a strict JSON schema with (i) the binary fiscal tag and (ii) a one–sentence rationale; malformed outputs are re–queried, and near–duplicates across outlets or editions are removed prior to aggregation. This second pre-filtering process after the Boolean query in Factiva reduces the number of news items to those shown in the second row of each panel in Table 1. Following this two–step filtering process, the resulting corpus contains only those articles explicitly discussing fiscal decisions or intentions, either on the expenditure or tax side.

Before moving to the classification of fiscal news, it is useful to gain an intuitive understanding of the linguistic content of the news that will feed the sentiment indices. Figs. 1 and 2 display the most frequent terms appearing in news for government expenditure and taxation, respectively. They confirm that the corpora are dominated by fiscal and budgetary terms—such as government, budget, spending, deficit, tax, and fiscal—showing that the pre-filtering effectively isolates news explicitly referring to fiscal policy announcements.

(a) Spain — *GovExpSent*

(b) Italy — *GovExpSent*

(c) France — *GovExpSent*

(d) Germany — *GovExpSent*

Fig. 1. Word clouds for *Government Expenditure Sentiment*, by country. The larger the font size, the more often this word appears in news reflecting changes in government expenditure.

[1]I use the HTTPS interface to a hosted large language model; the code sends the article text and a concise instruction and receives a JSON response with the label. No training takes place on my side; I only perform inference.

(a) Spain — *TaxSent* **(b)** Italy — *TaxSent*

(c) France — *TaxSent* **(d)** Germany — *TaxSent*

Fig. 2. Word clouds for *Tax Sentiment*, by country. The larger the font size, the more often this word appears in news reflecting changes in taxes.

2.2 Classifying fiscal news

After the prefiltering stage, the OpenAI model `gpt-4o-mini`[2] is used to classify fiscal announcements along three dimensions—direction, jurisdiction, and commitment. All articles are submitted to the model on a single batch, but the API is instructed to evaluate each news item individually according to different classification prompts. Classification is performed separately for spending- and tax-related news using dimension-specific prompts, with each article assigned to a single category per dimension. The rationale for classifying along three dimensions rather than a single one is that not all fiscal news items carry the same informational content: an official announcement of a national tax reform conveys a qualitatively different signal than a rumour about a possible municipal spending adjustment. By scoring articles along direction, jurisdiction, and communication type, the resulting index reflects not only the balance of expansionary versus contractionary news, but also the institutional relevance and credibility of each signal.

The reliability of the automated classification is evaluated by manually coding all Spanish news items and comparing these labels with the model's predictions (see Tables A.2 and A.3). For Germany, where performance is weaker, all articles are also classified manually. For France and Italy, robustness is assessed through manual review of a representative subsample and reclassification using the `gpt-5-mini` model, which yields very similar results.

The first classification dimension concerns the fiscal direction of each news item. Using the prompt reported in Table A.4, articles are classified as Up, Down, or Neutral depending on whether they signal an increase, a decrease, or an unclear change in public spending or taxation. These categorical labels are then mapped into numerical scores of $+1$, -1, and 0, respectively, which enter the construction of the quarterly sentiment indices. The Neutral category is important to avoid forcing a directional label on articles that report fiscal discussions without a clear policy stance—for instance, general budget debates or preliminary evaluations that do not yet point in a specific direction.

The second classification dimension captures the institutional level at which the fiscal announcement is made. Since fiscal measures can originate at the national, regional, or municipal level, and central government decisions typically have broader macroeconomic relevance, news items are classified accordingly using the LLM. In this dimension, the model must select one of three mutually exclusive categories: National, when the measure is attributed to the central government, the prime minister, the ministry of finance, the national parliament, or affects a large number of regions; Regional, when it comes from regional governments; or Municipal, when the decision corresponds to municipal authorities. Table A.5 shows the exact prompt used for Italy in this classification task; the prompts used for the other countries are direct contextual adaptations of this reference version. The weighting scheme assigns a value of 1 to National items, 0.6 to Regional items, and 0.1 to Municipal items.

Beyond fiscal direction and jurisdiction, it is also necessary to distinguish news items according to the nature of the communication they convey. Fiscal information reaches the public through channels that differ in credibility and immediacy—from formal government statements to informal suggestions or speculative reports—and accounting for these differences allows me to weight more heavily those signals most likely to reflect genuine fiscal intentions. To this end, each article in the cleaned corpus is assigned to one of three categories: Announcements, corresponding to official statements or firm commitments; Recommendations, reflecting proposals by actors without executive power such as opposition parties, trade unions, employers' associations, experts or advisory

[2]The API parameters are fixed to ensure full reproducibility (temperature = 0, fixed seed, maximum token limit = 400).

bodies; and Rumours, capturing unconfirmed or speculative information. Table A.6 reports the exact prompt used. To reflect their relative relevance, Announcements and Rumours receive a weight of 1, while Recommendations are assigned a weight of 0.1; these weights are applied multiplicatively to the direction and institutional scores before aggregation. The choice of assigning equal weight to Announcements and Rumours reflects the fact that, from the perspective of expectations formation, both types of communication are likely to influence agents' behaviour: a credible rumour about a forthcoming tax increase may trigger economic adjustments just as effectively as an official announcement.

2.3 Giving score to each of the news item and obtaining the Tax and Government Expenditure Sentiment Indices

The score of each news item is obtained as the product of three components: the direction of the fiscal announcement, its jurisdictional weight, and its type–of–communication weight. Formally, for each article i published in period t I define:

$$FSI_i = D_i \times J_i \times C_i,$$

where $D_i \in \{-1, 0, +1\}$ represents the fiscal direction, $J_i \in \{1, 0.6, 0.1\}$ denotes the jurisdictional weight, and $C_i \in \{1, 0.1\}$ corresponds to the communication weight[3]. The fiscal sentiment indices capture, for each quarter t, the prevailing tone of news relating to government expenditure and taxation. Specifically:

- For the *Government Expenditure Sentiment Index*, a higher value reflects a greater presence of articles announcing or signalling increases in public spending, whereas a lower value implies more frequent references to spending cuts, freezes, or consolidation measures.
- For the *Tax Sentiment Index*, a higher value indicates that news items predominantly report tax increases, expansions in the tax burden, or the cancellation of previously announced tax cuts; conversely, a lower value suggests that coverage is mainly focused on tax reductions, postponements of planned increases, or other forms of fiscal relief.

To construct the quarterly indices, the article-level scores are first summed within each quarter and then divided by the total number of news articles published in that quarter across all topics—not only fiscal news—in the same set of newspapers used to collect the fiscal corpus. This normalisation controls for variation in overall media intensity over time and prevents periods of elevated press activity (e.g. during recessions or elections) from mechanically inflating the index. Formally, the raw quarterly index for country c and fiscal instrument $f \in \{GovExpSent, TaxSent\}$ is defined as

$$\widetilde{S}^f_{t,c} = \frac{\sum_{i \in \mathcal{I}^f_{t,c}} FSI_i}{N_{t,c}},$$

where $\mathcal{I}^f_{t,c}$ denotes the set of classified fiscal articles for instrument f in quarter t and country c, and $N_{t,c}$ is the total number of news articles published in the same quarter, country, and newspapers. Finally, the normalised series is standardised to have zero mean and unit variance over the full sample period 1999Q1–2025Q1, yielding the indices *GovExpSent* and *TaxSent* used throughout the analysis.

Figs. 3 and 4 summarise the quarterly news-based fiscal sentiment indices for Spain, Italy, France, and Germany. In the four countries analysed, movements in quarterly fiscal sentiment indices clearly reflect major public policy events. In Spain, *TaxSent* clearly reflects the VAT increases announced at the end of 2009 (from 16% to 18%) and in 2012 (from 18% to 21%), together with the broadening of tax bases, while *GovExpSent* reflects the sharp spending cuts announced in May 2010—reduction in civil servants' salaries, pension freeze, fall in public investment and elimination of the baby bonus—as well as the expansionary announcements made during the pandemic. In Italy, the fiscal consolidation of 2011–2012 under the "Salva-Italia" plan, which combined tax increases, pension cuts and a delay in the retirement age, corresponds to a sharp rise in the tax index and a marked fall in the expenditure index. For France, the tax increases announced in July 2012—including a new 4% tax on oil stocks—and, later, the launch of the France Relance programme on September 2020, with around 100 billion euros for investment, competitiveness and infrastructure, generate identifiable movements in the *GovExpSent* and *TaxSent* indices, respectively. Finally, in Germany, the indices also capture the excessive deficit procedure initiated in 2003 by the European Commission, which coincides with a deterioration in the expenditure-announcement indicator. They also reflect the stimulus package approved in November 2008 and the subsequent adjustments in 2009, as well as the expansionary programme launched in June 2020, which included a temporary reduction in VAT and measures to support households and investment. The fact that the indices capture well-known fiscal episodes across the four countries provides informal validation that the LLM-based classification and the aggregation procedure produce economically meaningful measures of fiscal sentiment.

Fig. A.1 shows the cross-correlations between fiscal announcement indices and their corresponding aggregates for Spain, Italy, France, and Germany. In all cases, *GovExpSent* is positively correlated with public spending at short leads, indicating that spending announcements tend to anticipate subsequent increases in government expenditure. Similarly, *TaxSent* displays weaker but still leading correlations with tax revenues, confirming that tax-related news precedes movements in revenue collection. These leading correlations are consistent with the hypothesis that fiscal news is forward-looking: newspapers report government intentions before they materialise in actual spending or revenue data, and the indices capture this anticipatory information content.

[3] To address potential concerns that these jurisdictional and communication weights could introduce arbitrariness, I also construct a highly restrictive alternative index that retains only those items classified as *national*-level and as either *announcements* or *rumours*, assigning them their directional score $(+1, 0, -1)$ and discarding all remaining observations. Re-estimating all baseline VAR models with this "high-credibility" index yields impulse responses that are virtually identical to those obtained with the full weighted series, confirming that the macroeconomic results do not depend on the specific weighting scheme used.

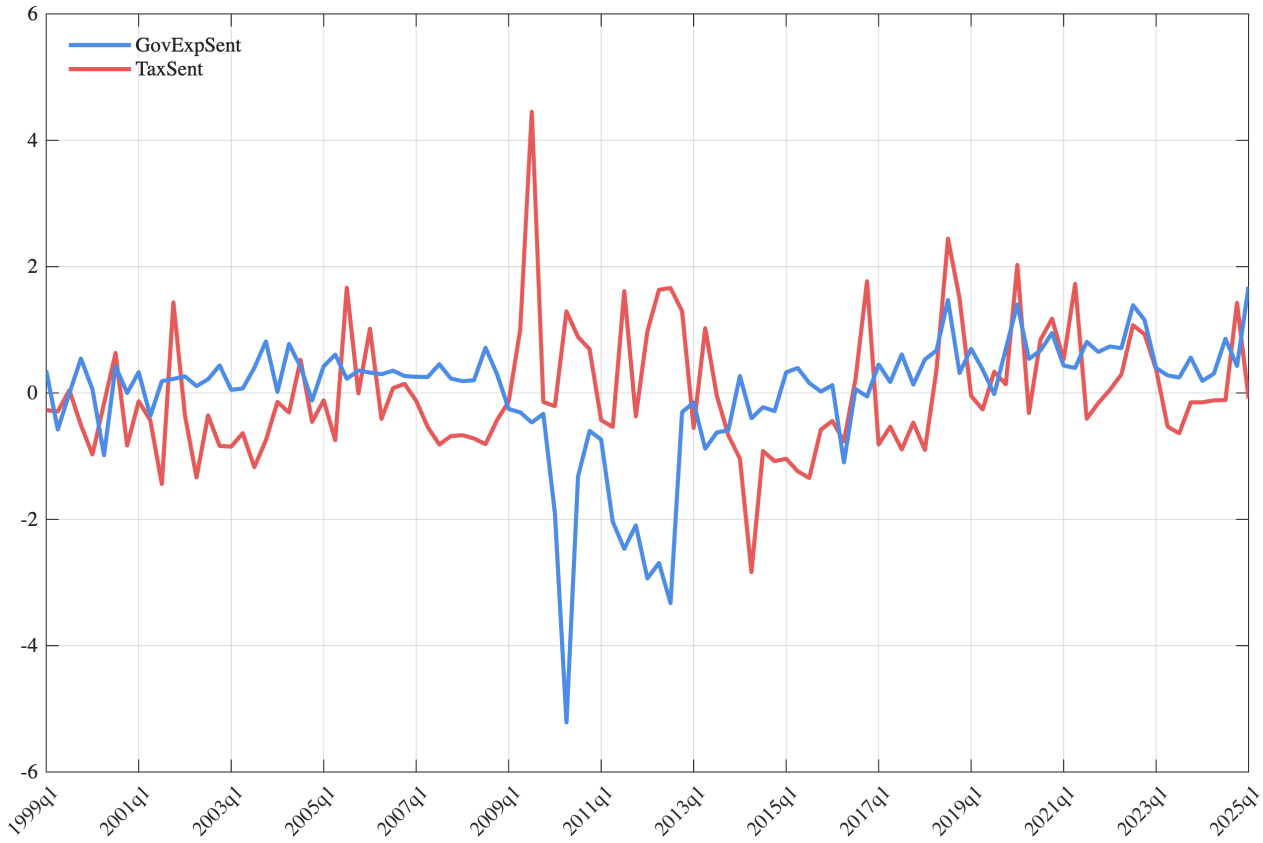

(a) Spain — *GovExpSent* (blue) vs. *TaxSent* (red).

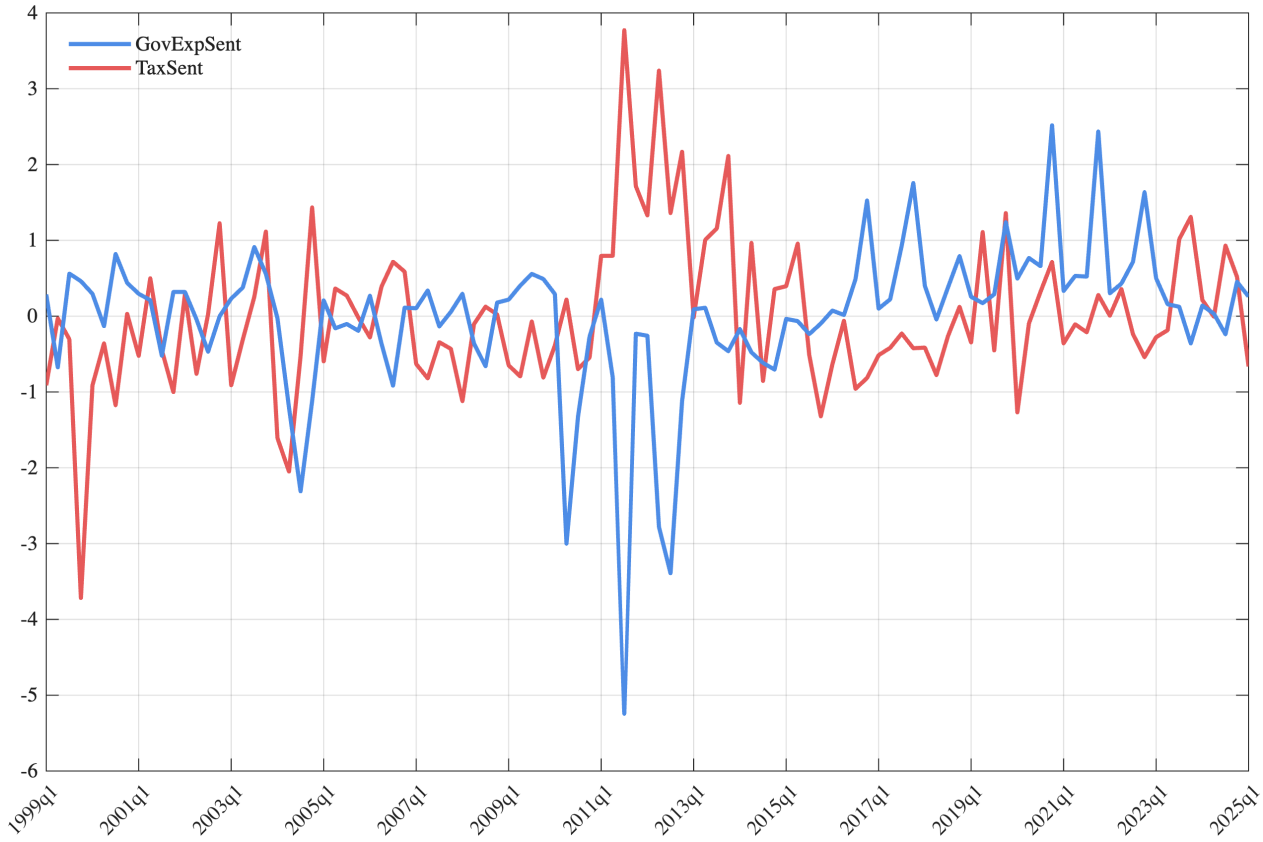

(b) Italy — *GovExpSent* (blue) vs. *TaxSent* (red).

Fig. 3. Fiscal sentiment indices for Spain and Italy. *Notes:* Indices are standardized. *GovExpSent* captures spending–increase vs. spending–cut announcements, while *TaxSent* measures the balance of tax–increase vs. tax–cut news.

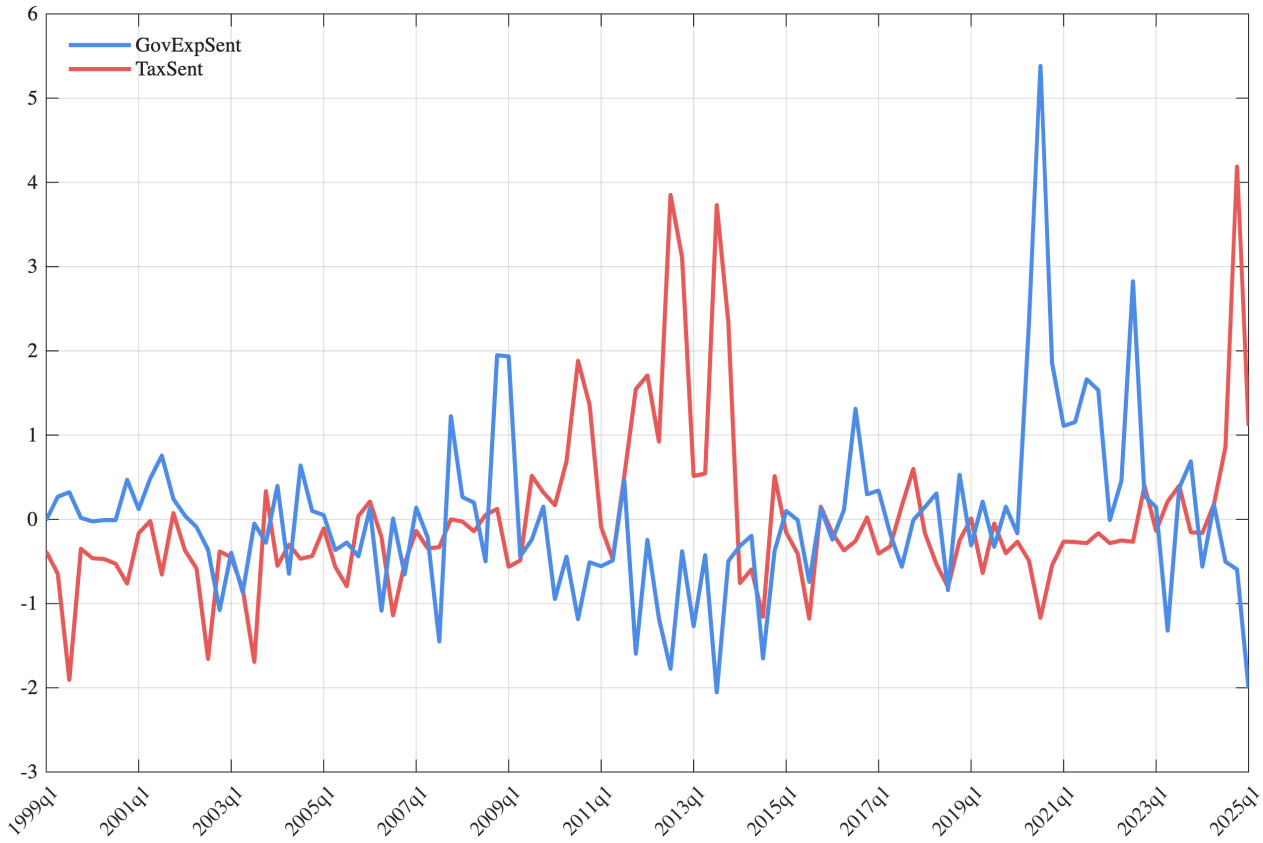

(a) France — *GovExpSent* (blue) vs. *TaxSent* (red).

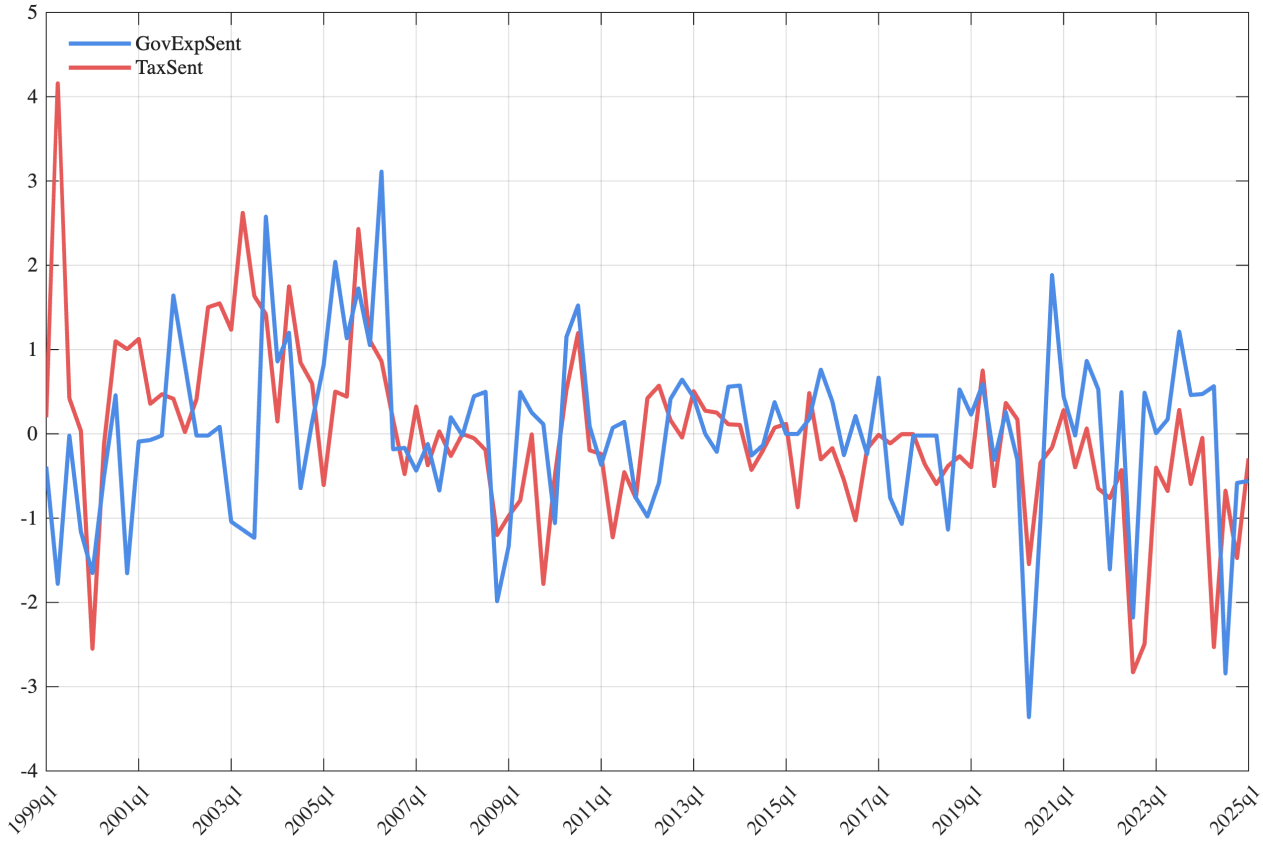

(b) Germany — *GovExpSent* (blue) vs. *TaxSent* (red).

Fig. 4. Fiscal sentiment indices for France and Germany. *Notes:* Indices are standardized. *GovExpSent* captures spending–increase vs. spending–cut announcements, while *TaxSent* measures the balance of tax–increase vs. tax–cut news.

3 Empirical strategy

This section describes the empirical strategy used to estimate the macroeconomic effects of fiscal announcement shocks. The analysis combines a pooled Bayesian panel VAR, designed to identify the average dynamic response to fiscal announcements across European countries, with country-specific SVAR models that allow for a more detailed assessment of national transmission mechanisms. Section 3.1 presents the benchmark panel VAR specification, while Section 3.2 describes the country-by-country VAR models and discusses the main estimation choices, including lag structure and prior selection. Section 3.3 details the data sources, variable definitions, and transformations. Finally, Section 3.4 outlines the identification strategy, which relies on a recursive Cholesky decomposition.

3.1 Panel SVAR model

The first and main step in the empirical analysis is the estimation of two pooled panel VARs—one for government expenditure announcements and another for tax announcements—to quantify the average macroeconomic effects of fiscal sentiment shocks. Considering a panel of $N = 4$ economies—Spain, Italy, France, and Germany—each characterised by the same $K \times 1$ vector of endogenous variables y_t observed over $t = 1, \ldots, T$ each pooled panel VAR can be specified as

$$y_t = \sum_{k=1}^{p} A^k \, y_{t-k} + C x_t + \varepsilon_t \tag{1}$$

where A^k are coefficient matrices whose dynamic coefficients are homogenous across countries, and coefficients are time-invariant, x_t is a vector of common exogenous variables, and $\varepsilon_t \sim \mathcal{N}(0, \Sigma)$ denotes the pooled reduced-form innovations. The block-diagonal structure imposes cross-country independence: each country's equations include only its own lags. Stacking the system yields

$$\begin{pmatrix} y_{1,t} \\ y_{2,t} \\ \vdots \\ y_{N,t} \end{pmatrix} = \begin{pmatrix} A^1 & \cdots & 0 \\ 0 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & A^1 \end{pmatrix} \begin{pmatrix} y_{1,t-1} \\ y_{2,t-1} \\ \vdots \\ y_{N,t-1} \end{pmatrix} + \cdots + \begin{pmatrix} A^p & \cdots & 0 \\ 0 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & A^p \end{pmatrix} \begin{pmatrix} y_{1,t-p} \\ y_{2,t-p} \\ \vdots \\ y_{N,t-p} \end{pmatrix} + \begin{pmatrix} C \\ C \\ \vdots \\ C \end{pmatrix} x_t + \begin{pmatrix} \varepsilon_{1,t} \\ \varepsilon_{2,t} \\ \vdots \\ \varepsilon_{N,t} \end{pmatrix}, \tag{2}$$

which is equivalent to estimating a single VAR for the entire panel under parameter homogeneity. These pooled panel VARs are estimated under bayesian methods using a Normal-Wishart prior.

The pooled specification offers two main advantages. First, it improves estimation efficiency by jointly exploiting the information available in my short sample of multiple countries. Second, the imposition of common coefficients delivers the average dynamic response of the main European economies to innovations in *TaxSent* and *GovExpSent*, preventing heterogeneity in size, volatility, or fiscal composition from dominating the inference.

3.2 SVAR model

The second step in the empirical analysis is the estimation of a country-by-country reduced-form VAR model

$$y_{t,c} = A_{1,c} y_{t-1,c} + \cdots + A_{p,c} y_{t-p,c} + u_{t,c}, \tag{3}$$

where $y_{t,c}$ is a $K \times 1$ vector of endogenous variables, with $c \in$ (Spain, Italy, France, Germany), $A_{i,c}$ are coefficient matrices, and $u_{t,c} \sim \mathcal{N}(0, \Sigma_c)$ is the vector of reduced-form innovations, where $\Sigma_{j,c}$ is the variance–covariance matrix with $E(u_{t,c} u'_{t,c}) = \Sigma_{j,c}$. I further assume no correlation between the residuals at all leads and lags. For each country c, two separate VARs are estimated, one including *GovExpSent* and another including *TaxSent*.

To estimate the VAR model, I again use a Bayesian approach with a Normal-Wishart prior[4] as implemented in the BEAR toolbox[5], treating the model's coefficients as random variables and deriving their posterior distributions through Markov Chain Monte-Carlo sampling. Results are robust to different lags and to different prior specifications, specifically across two widely used priors—Minnesota and Normal-Diffuse.

As mentioned before, a recursive identification scheme is adopted. Specifically, considering the K-dimensional time series $y_{t,c}$, the structural-form VAR model with p lags can be written as

$$B_{0,c} y_{t,c} = B_{1,c} y_{t-1,c} + \cdots + B_{p,c} y_{t-p,c} + w_{t,c} \tag{4}$$

where $y_{t,c}$ is the $K \times 1$ vector of endogenous variables, $w_{t,c} \sim N(0, \Sigma_{w,c})$ is the mean-zero serially uncorrelated structural error term with $\Sigma_{w,c} = E(w_{t,c} w'_{t,c}) = I_K$ being the variance–covariance matrix normalized to the identity, and $B_{0,c}$ is the non-singular matrix governing the contemporaneous interaction between the model variables.

As shown in Kilian and Lütkepohl (2017), standard estimation methods allow for consistent estimation of the reduced-form parameters of the pooled and country-by-country VARs. From these estimates, the structural-form parameters can be recovered by imposing exclusion restrictions on selected elements of the inverse of the contemporaneous-impact matrix. In this paper,

[4] Following standard practice, I set overall tightness to 0.2, cross-variable weight to 0.5, lag decay to 1.5, and the AR coefficient to 0.8, which allows me to capture persistence without overfitting. The results are robust to variations in these hyperparameters.

[5] The BEAR toolbox is a "comprehensive (Bayesian Panel) VAR toolbox for forecasting and policy analysis" developed by the European Central Bank.

identification is achieved through a recursive scheme based on a Cholesky decomposition, which requires an economically meaningful ordering of the variables. Section 3.4 discusses the economic rationale underlying the contemporaneous restrictions imposed in both the pooled panel and country-level VARs.

3.3 Data

In the benchmark specification, two pooled panel VAR systems are estimated: one for government expenditure announcements and another for tax announcements. The panel consists of quarterly observations for the four countries of interest—Spain, Italy, France, and Germany—observed over the same sample period. For each country c, the vectors of endogenous variables are given by

(i) Government expenditure model

$$y_{t,c}^{GovExp} = \begin{bmatrix} GovExp_{t,c} & GDP_{t,c} & Cons_{t,c} & GovExpSent_{t,c} \end{bmatrix} \tag{5}$$

(ii) Tax model

$$y_{t,c}^{Tax} = \begin{bmatrix} TaxRev_{t,c} & GDP_{t,c} & Cons_{t,c} & TaxSent_{t,c} \end{bmatrix} \tag{6}$$

All log variables are multiplied by 100, while the sentiment indicators $GovExpSent_t$ and $TaxSent_t$ are standardised. The VAR lag length is set to $p = 4$, which allows the system to capture medium-run propagation consistent with fiscal policy transmission. In all specifications two exogenous regressors are also included: a linear time trend and a COVID dummy equal to one in 2020Q1–2020Q3 and zero otherwise. The estimation frequency, as mentioned earlier, is quarterly, and the data spans 1999Q1-2025Q1[6]. To ensure that the results are not driven by the baseline model, several extended specifications are estimated by augmenting the VAR with additional macroeconomic variables, as reported in Table A.7.

The core macroeconomic series—real GDP, real consumption, real government expenditure, real investment, inflation, tax revenues, trade balance, and interest rates—are obtained from Eurostat, business expectations are taken from the European Commission Business and Consumer Surveys and 10-year sovereign yields are obtained from the ECB. All series are quarterly, seasonally adjusted, and transformed as described above.

3.4 Identification

The objective in this section is to justify the recursive identification proposed in the present work. To this end, the analysis draws on the extensive literature on the identification of exogenous fiscal policy shocks proposed by Blanchard and Perotti (2002) and later applied by Galí et al. (2007), Ramey (2011), Born and Müller (2012), Auerbach and Gorodnichenko (2012), Ilzetzki et al. (2013), Blanchard et al. (2013), Ramey and Zubairy (2018), Latifi et al. (2024), and Smets and Trienens (2025), among others.

Following this literature, the order of variables established in Eq. 5 and Eq. 6 implies that government expenditure and tax revenues are predetermined within the quarter; that is, it is assumed that changes in GDP and consumption—as well as the additional variables included in the extensions—do not affect public expenditure or tax revenues contemporaneously. This assumption, widely accepted in the literature, is supported by the institutional calendar of fiscal policy: parliamentary approval deadlines, delays in budget execution and the rigidity of the legislative process mean that announced changes take time to implement.

Based on this institutional structure, the next step is to justify the order assigned to the fiscal sentiment indicators. In my central approach, the focus is not on the response of the fiscal policy implemented, but on understanding how fiscal communication—captured through news—impacts economic activity. It is crucial to understand that the indicators reflect expectations, opinions, and debates that evolve in tandem with the economy. In this sense, the cross-correlation functions shown in Figs. 5 and 6 reinforce the idea that fiscal news does not behave like exogenous innovations, but rather responds to the economic cycle: in most countries, growth and business expectations deteriorate, as well as investment, before an increase in fiscal sentiment indices. Therefore, *TaxSent* and *GovExpSent* are placed last in the Cholesky ordering.

This structure implies that the identification does not rely on assuming that the sentiment series are exogenous with respect to economic activity; in particular, it does not require removing announcements that are themselves motivated by the state of the business cycle. Instead, exogeneity is imposed by recursive ordering: since sentiment is ordered last, shocks to GDP, consumption or fiscal variables can move sentiment contemporaneously, but sentiment cannot contemporaneously affect those variables. Therefore, what remains—after applying Cholesky ordering—is the component of sentiment orthogonal to contemporary macroeconomic conditions, interpretable as a pure fiscal news shock.

In this framework, a clear asymmetry in the transmission of announcements is expected: an exogenous increase in *GovExpSent*, associated with news anticipating a more expansionary public-spending stance, should precede subsequent increases in spending itself and in economic activity, whereas an exogenous increase in *TaxSent*, linked to expectations of higher taxes, should anticipate contractionary effects on output and consumption, as well as an increase in tax revenues.

[6] In the case of the Tax model, and given that for some German variables we only have data from 2002 onwards, we use the sample period 2002q1-2025q1.

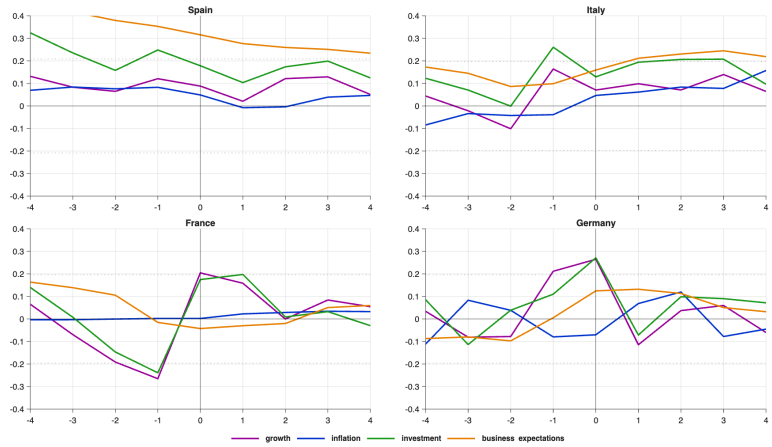

Fig. 5. Cross-correlations between *GovExpSent* and macroeconomic variables (output growth, inflation, investment growth and business expectations) for Spain, Italy, France and Germany. The vertical line marks the quarter of the government-expenditure news announcement (lag 0); negative (positive) lags correspond to quarters before (after) the announcement.

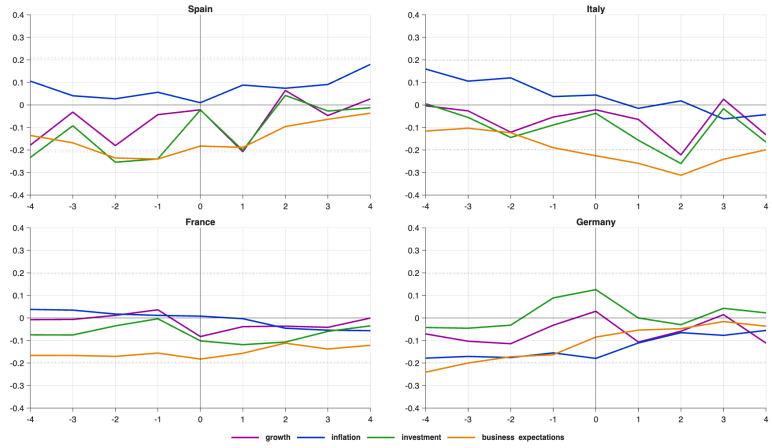

Fig. 6. Cross-correlations between *TaxSent* and macroeconomic variables (output growth, inflation, investment growth and business expectations) for Spain, Italy, France and Germany. The vertical line marks the quarter of the tax-news announcement (lag 0); negative (positive) lags correspond to quarters before (after) the announcement.

4 Macroeconomic Effects of Fiscal Announcements

This section analyses the macroeconomic effects of fiscal announcements. Section 4.1 presents the main impulse–response evidence from the panel and country-level VARs, while Section 4.2 and 4.3 assesses the stability of these results under alternative specifications and detrending procedures. Lastly, Section 4.4 splits fiscal sentiment into exogenous (purged) and endogenous components to assess the role of systematic stabilisation in driving the baseline effects.

4.1 Results

Figs. 7 and 8 show the responses of endogenous variables to a positive shock of one-standard-deviation in *GovExpSent* and *TaxSent*, respectively; that is, they give us the average macroeconomic impacts of announcing higher government expenditure or taxes across the main euro area economies. All figures also include probability bands that cover 68% of all draws[7]. A first inspection of the responses reveals two stylised facts. First, fiscal announcements have significant and persistent macroeconomic effects, anticipating future movements in both fiscal variables and real activity. Second, and particularly noteworthy, the responses show marked asymmetry: news of increases in public spending generates substantial and lasting expansions in GDP and consumption, while news of tax increases leads to more moderate and less persistent contractions. This asymmetry is present in the panel model as well as in the country-by-country estimation.

Fig. 7 displays the panel VAR impulse responses of macroeconomic aggregates to a one-standard-deviation shock in the *GovExpSent* index, capturing the average macroeconomic effects of announcing higher government expenditure. This shock leads to a gradual and statistically significant increase in public spending in the panel. The adjustment occurs gradually, in line with a news-driven mechanism, in which announcements precede measures whose budgetary impact only materialises after a delay of several quarters. GDP reacts clearly, immediately and persistently: economic activity expands rapidly, with an increase of around 0.4–0.6%, and remains steadily above the reference level over the medium term. Private consumption closely follows this trend, showing even greater expansion in magnitude and persistence.

[7]Probability bands covering 90% of all draws are not reported, as they are nearly identical to the 68% bands.

Fig. 8 reports the panel VAR impulse responses of macroeconomic aggregates to a one-standard-deviation shock in the *TaxSent* index, capturing the average macroeconomic effects of announcing tax increases or the withdrawal of previously announced reductions. This shock generates a clear and persistent increase in effective tax revenues, confirming the index's ability to anticipate future fiscal movements. The effect on GDP is contractionary, significant and lasting: output falls immediately after the announcement and continues to decline gradually, reaching around 0.2–0.3% below the reference level in the medium term. Private consumption follows a very similar contractionary path, highlighting that tax announcements generate a persistent drag on aggregate economic activity.

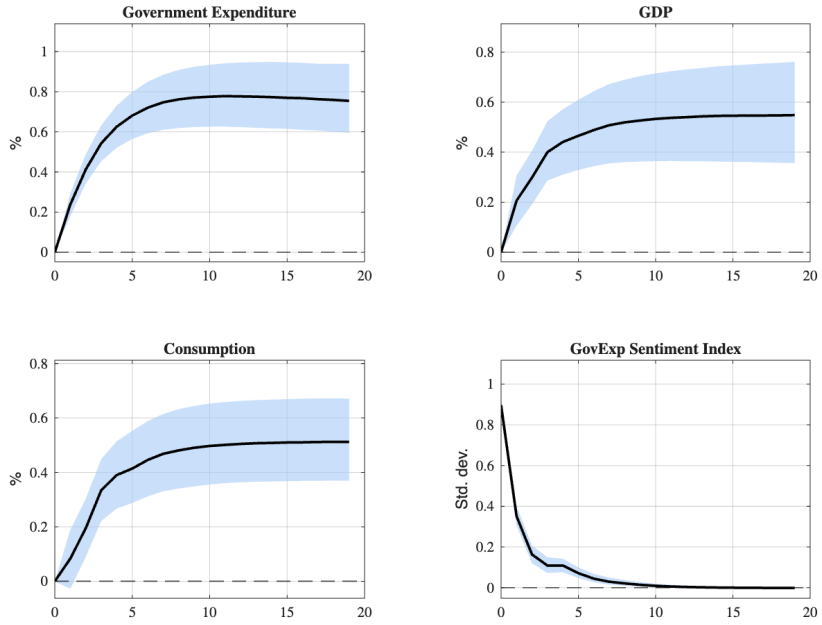

Fig. 7. Responses to a government expenditure sentiment shock (*GovExpSent*). *Notes:* Impulse responses of government expenditure, GDP, consumption, and *GovExpSent* to a one–standard–deviation increase in government expenditure sentiment, estimated from a Bayesian pooled VAR. Shaded areas denote 68% credible intervals.

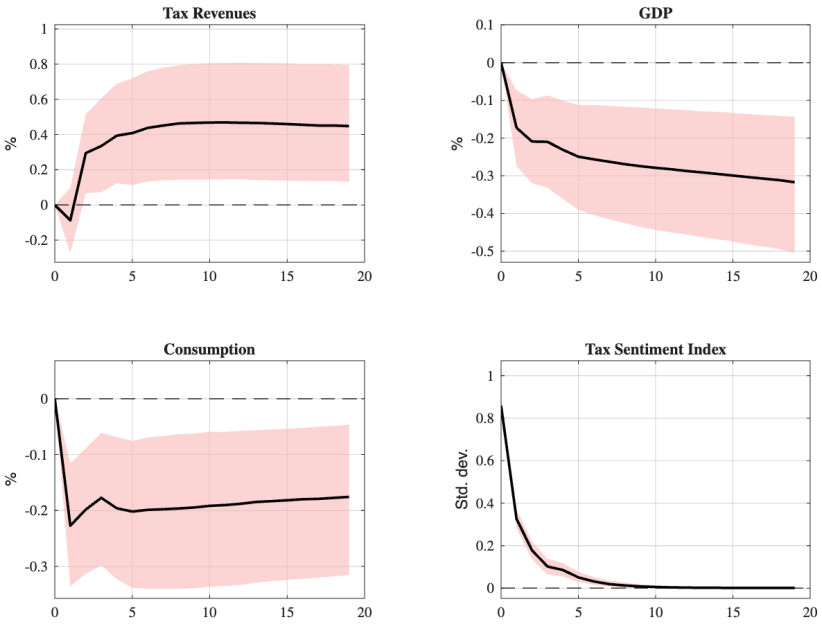

Fig. 8. Responses to a tax sentiment (*TaxSent*) shock. *Notes:* Impulse responses of tax revenues, GDP, consumption, and *TaxSent* to a one–standard–deviation increase in tax sentiment, estimated from a Bayesian pooled VAR. Shaded areas denote 68% credible intervals.

The impulse-response functions estimated at the country level broadly confirm the results of the panel VAR. Following a shock in *GovExpSent*, as shown in Fig. A.2, public spending increases gradually in all countries, with more intense responses in Spain, Italy and France, although more moderate in Germany, while GDP and private consumption react positively with very similar dynamics across economies. In the case of shocks in *TaxSent*, as shown in Fig. A.3, the responses are weaker and more heterogeneous: although tax revenues increase in all countries, the contractionary effects on GDP and consumption are more visible in Spain and Italy and more muted—and often statistically uncertain—in France and Germany.

4.2 Alternative specifications

To assess robustness of baseline results and explore effects on additional macroeconomic aggregates, the panel VAR is extended by sequentially incorporating additional variables[8]. First, private investment is added, ordered after GDP and consumption but before the sentiment indicator. As shown in Figs. A.4 and A.5, announcements of higher public spending generate a sizable and persistent increase in investment, closely aligned with the expansion in GDP and consumption, while news of higher taxes lead to a pronounced and sustained contraction in investment of comparable magnitude, consistent with an expectations-driven investment channel.

Second, the specification is augmented with additional fiscal and monetary variables—tax revenues or public expenditure, inflation, and the short-term interest rate—ordered after activity and before the sentiment indicator. Figs. A.6 and A.7, show that spending-related news generate persistent inflationary pressures and a delayed increase in interest rates, consistent with New-Keynesian predictions, whereas tax-related announcements are associated with deflationary dynamics and largely unchanged interest rates.

Third, business expectations are incorporated, ordered after consumption. As reported in Figs. A.8 and A.9, news of higher public spending triggers a rapid and significant improvement in expectations, while announcements of higher taxes lead to an immediate deterioration in the economic outlook. The open-economy dimension is then explored by adding the trade balance as a percentage of GDP. Figs. A.10 and A.11 show that spending announcements temporarily worsen the external balance, consistent with demand-driven import responses, whereas tax announcements are associated with a gradual improvement driven by weaker domestic demand.

Finally, to address the concern that fiscal sentiment may partly reflect episodes of sovereign stress, the panel VAR is extended to include 10-year sovereign yields. As shown in Figs. A.12 and A.13, the results remain virtually unchanged relative to the baseline specification in terms of sign, magnitude, and persistence.

4.3 Robustness

As a robustness exercise, the analysis considers two alternatives to the use of log-level variables in the baseline panel VAR models incorporating either *GovExpSent* or *TaxSent*. First, a detrending procedure, in the spirit of Gordon and Krenn (2010), Ramey (2016), Ramey and Zubairy (2018), Latifi et al. (2024), and Smets and Trienens (2025), is applied. Specifically, the trend component of real GDP is extracted by fitting a quadratic trend to log real GDP, and I construct for each macroeconomic variable $x_t \in \{\text{GDP, consumption, government expenditure, tax revenues}\}$ the transformation

$$100 \times (\ln x_t - \ln y_t^{\text{trend}})$$

where y_t^{trend} denotes the estimated trend in real GDP. Then, the model incorporating *GovExpSent* and the one including *TaxSent* are re-estimated, using these detrended series. The results based on the quadratic trend in GDP are reported in Fig. A.14 and Fig. A.15. An unexpected increase in government expenditure sentiment leads to higher government spending and private consumption, along with GDP expansion. Conversely, news about higher taxes generate a sustained rise in tax revenues, while GDP and private consumption contract persistently. All responses look similar to the results based on the log-level variables presented in the previous section.

Second, following Hamilton (2018), Ilori et al. (2022), and Latifi et al. (2024), the quadratic trend is replaced with the Hamilton filter. In this case, the y_t^{trend} is obtained by applying the Hamilton filter to log real GDP and use the resulting cyclical components of GDP, consumption, government expenditure, and tax revenues as inputs in all VAR specifications with *GovExpSent* and *TaxSent*. The results based on the GDP trend derived from the Hamilton filter are shown in Figs. A.18 and A.19, where it can be seen that, for both sentiment indices, the responses closely resemble those obtained in the previous section using log-level variables, as well as those from the Gordon detrending procedure, meaning that news about higher government expenditure generate an expansion in economic activity, whereas news about higher taxes produce contractionary effects.

In line with the panel VAR evidence, the country-by-country results reported in Figs. A.2 and A.3 lead to the same qualitative conclusions, although the magnitude and timing of the responses vary across countries. Importantly, these cross-country differences do not alter the overall interpretation of the results, as the direction and persistence of the estimated effects remain consistent with the pooled analysis.

4.4 Exogenous vs. endogenous fiscal announcements

A potential concern with news-based fiscal announcement indices is that part of their variation may reflect systematic fiscal stabilisation—that is, governments (or the policy debate reported in the press) respond to the state of the business cycle—rather than genuinely exogenous shifts in the fiscal stance (Romer & Romer, 2010; Ramey, 2011; Cloyne, 2013). Following the literature, this issue is addressed by decomposing fiscal sentiment into *exogenous* and *endogenous* components, where the former captures policy variation not intended to offset other macroeconomic shocks, while the latter reflects countercyclical stabilisation motives.

In this setting, however, the short and often ambiguous nature of journalistic texts makes a semantic classification of each announcement into "exogenous" versus "endogenous" categories difficult to implement. Instead, a simple *purging* strategy is adopted that

[8] Table A.7 explains these alternative specifications. Moreover, augmented country-level VARs reported in the Appendix confirm the pooled panel evidence reported in this section.

operationalises the same intuition using only observable macro-fiscal information. Specifically, for each country c, an ARX-type regression of the announcement index on its own lags and on lagged macroeconomic and fiscal controls is estimated:

$$Sent_{t,c} = \alpha_c + \sum_{k=1}^{4} \phi_k \, Sent_{t-k,c} + \sum_{k=1}^{4} \Gamma_k' X_{t-k,c} + u_{t,c}, \tag{7}$$

where $Sent_{t,c}$ denotes either $GovExpSent_{t,c}$ or $TaxSent_{t,c}$, and $X_{t,c}$ includes (depending on the specification) government expenditure or tax revenues, GDP, and consumption. The residual $u_{t,c}$ is interpreted as the *unexpected* (orthogonal) component of fiscal news—namely, the part of announcements that is not predictable from recent macro-fiscal conditions—whereas the fitted component captures the *systematic* (endogenous) part that co-moves with the state of the economy by construction. That is, this decomposition yields two series per country and per instrument

(i) an "exogenous" (purged) sentiment shock, $Sent_{t,c}^{exo} \equiv \widehat{u}_{t,c}$; and,

(ii) an "endogenous" sentiment component, $Sent_{t,c}^{endo} \equiv \widehat{Sent}_{t,c} \equiv Sent_{t,c} - \widehat{u}_{t,c}$

The baseline panel BVAR is then re-estimated replacing the raw sentiment index with these components in order to assess whether the macroeconomic effects are primarily driven by the purged (orthogonal) variation. The results are reported in Figs. A.16 and A.17. For *GovExpSent*, the expansionary responses of spending, GDP and consumption remain broadly unchanged when using the purged innovation, suggesting that the baseline effects are not mechanically driven by countercyclical variation. For *TaxSent*, instead, the macro responses are much weaker and less precisely estimated once purged, consistent with tax-related news being more state-contingent and noisier in the press, so that ARX purging may remove part of its economically relevant information content.

5 Fiscal Foresight, Announcements, and Policy Implementation

The literature on government spending shocks emphasises that fiscal foresight invalidates the recursive Blanchard and Perotti (2002) identification, because shocks recovered from a Cholesky ordering can be predictable once agents observe information about forthcoming fiscal policy (Ramey, 2011, 2016). This issue is particularly relevant when comparing the macroeconomic effects of fiscal announcements with those of subsequent spending implementation[9]. If announcements convey advance information about future spending, then the innovation to observed government expenditure in a standard Blanchard–Perotti VAR may partly reflect news that has already been incorporated by private agents, rather than genuinely unanticipated policy. A natural way to address this concern is to incorporate a news-based announcement index as a proxy for the information set available at the time of the announcement and to condition the identification of spending shocks on this information[10].

Against this backdrop, the analysis presented in this section proceeds in two stages. First, it is assessed whether the announcement index predicts recursively identified government spending shocks. Second, this information is used to compare the macroeconomic effects of fiscal news from those of realised spending, by identifying an implementation shock conditional on prior announcements.

Following Latifi et al. (2024), government spending shocks are first identified from a panel VAR including government expenditure, GDP, and consumption, adopting the standard Blanchard–Perotti recursive scheme with expenditure ordered first. Importantly, the VAR is estimated without including the sentiment index, which is introduced only in the subsequent forecasting exercise. To assess robustness to the treatment of the data, three alternative shock series are constructed: (i) using log-levels, (ii) after Gordon-style detrending, and (iii) after applying the Hamilton filter. This yields three corresponding measures of structural government spending shocks, one for each country–transformation of the endogenous variables.

The predictive content of fiscal sentiment is then assessed using local projections in the spirit of Jordà (2005). Specifically, the government spending shock at horizon $t+h$ is regressed on fiscal sentiment at time t:

$$shock_{t+h}^{k} = \alpha_h + \beta_h \, GovExpSent_t + \sum_{k=1}^{4} \gamma_k' X_{t-k,c} + \varepsilon_{t+h}, \tag{8}$$

where X_t contains contemporaneous and lagged realisations of government expenditure, consumption, and GDP as control variables. As the dependent variable is the result of a structural identification, it should be orthogonal to these controls. Nevertheless, these variables are included as controls. Eq. 8 is estimated for each treatment $k \in \{\text{log-levels, Gordon, Hamilton}\}$ and 68% and 90% confidence bands are reported based on Newey–West standard errors. For cross-country comparability, the identified shock series are standardised within each country over the estimation sample.

Fig. 9 plots the estimated $\hat{\beta}_h$ as a function of the horizon h for the log-levels transformation[11]. Across countries, the estimates indicate that a shift towards a more expansionary fiscal sentiment in t predicts an increase in subsequently identified government spending shocks at medium horizons, suggesting that government expenditure shocks identified from a Blanchard–Perotti style VAR can indeed be anticipated. These findings underline that news-based fiscal announcements capture forward-looking information about forthcoming spending decisions.

[9] In this section, the analysis is restricted to government expenditure, since tax changes are proxied by tax revenues, which does not allow for a clean identification of discretionary tax policy shocks.

[10] This approach relies on the maintained assumption that the news-based announcement index perfectly captured all anticipated changes in government spending so that the identified shocks will consist of (i) anticipated changes in government spending that are captured by the econometrician in the news variable and (ii) unanticipated government spending (Ramey, 2011).

[11] The corresponding estimates for the alternative transformations, which are qualitatively and quantitatively similar to the log-levels transformation, are reported in Fig. A.20 and Fig. A.21.

Building on the evidence that the news-based index helps to mitigate fiscal foresight, it becomes possible to distinguish the macroeconomic effects of fiscal news from those of realised spending. Since the macroeconomic impact of fiscal announcements has already been documented in the previous section, the remaining task is to estimate the effects of the effective implementation of public spending, net of announcement effects. To do so, a VAR is estimated using the same set of variables as in the analysis of fiscal announcements, but ordering *GovExpSent* before public expenditure in the recursive decomposition. Under this specification, the shock to public expenditure excludes, as far as possible, the component of fiscal adjustment that has already been announced and anticipated, capturing only the unanticipated variation conditional on the information available in the press.

As shown in Fig. 10, the response of GDP to a fiscal announcement shock is stronger and more persistent than that associated with the subsequent implementation of public spending. This pattern suggests that a substantial share of the macroeconomic adjustment occurs at the announcement stage, when agents revise expectations about future fiscal conditions and adjust behaviour accordingly. By contrast, once advance information is accounted for, the GDP response to a shock in implemented spending is more moderate and displays a flatter dynamic profile, consistent with a lower marginal informational content at the time of execution.

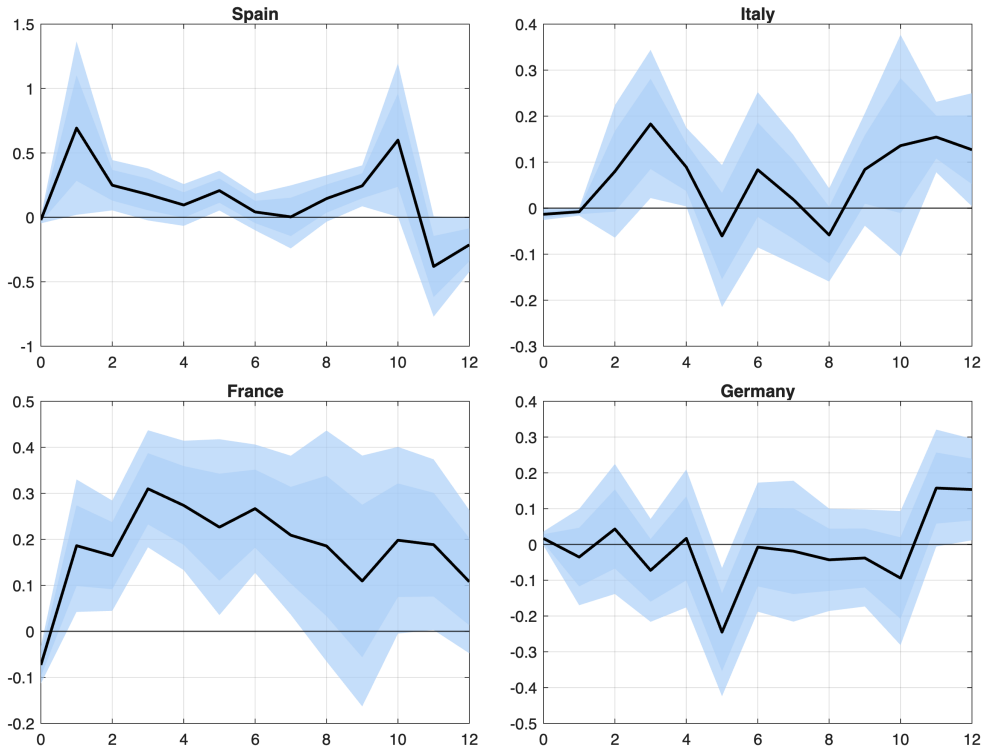

Fig. 9. Response of the government expenditure shock to fiscal sentiment (log-levels). *Notes:* The figure shows the estimated $\hat{\beta}_h$ from Eq. 8. Shaded areas denote 68% and 90% confidence bands based on Newey–West standard errors.

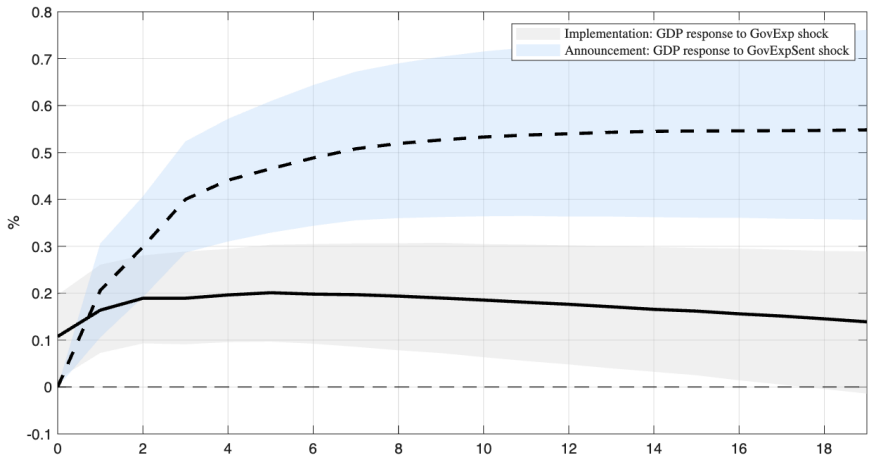

Fig. 10. GDP responses to fiscal announcements and implementation. *Notes:* Impulse responses of GDP to a one–standard–deviation shock in government expenditure sentiment (*GovExpSent*) and to a government expenditure shock (*GovExp*) conditional on prior fiscal information. Shaded areas denote 68% credible intervals.

6 Conclusions

This paper has examined the earliest stage of the fiscal policy-making process: the public communication of governments' intentions. The analysis uses a comprehensive set of fiscal announcements reported in the main national newspapers in Spain, France, Italy, and Germany to capture the signals communicated by governments before fiscal measures are approved or implemented. Using state-of-the-art, manually validated language models, I construct two quarterly indicators that summarise the direction of fiscal announcements reported in the press and allow for the analysis of both the macroeconomic impact of fiscal announcements and that of their subsequent implementation.

Two main conclusions can be drawn. First, fiscal announcements have significant macroeconomic effects when the information is made public: announcements of increased spending anticipate subsequent increases in spending and generate persistent expansions in GDP, consumption, investment, inflation and business expectations, while announcements of tax increases raise revenue and produce contractions of comparable magnitude, as well as deflationary pressures. Future research should analyse how these regularities depend on the cyclical context, political and institutional factors, and possible spillover effects between economies. Hence, a change in fiscal sentiment has macroeconomic effects consistent with standard New-Keynesian business cycle models.

Secondly, the comparison of announcement and implementation effects within the same empirical framework reveals that GDP responds more strongly at the communication stage than at the execution stage, consistent with forward-looking behaviour on the part of households and firms who adjust their decisions as soon as fiscal signals become available (Ramey, 2011). This finding carries a direct policy implication: the economic impact of fiscal policy begins well before any measure is approved or executed, so that the way in which governments communicate their intentions—and the manner in which these signals reach economic agents through the press—is itself a relevant dimension of fiscal policy design. More broadly, the paper illustrates how large language models can be used to extract structured economic information from unstructured media sources at scale, opening avenues for the real-time monitoring of fiscal communication and for the study of announcement effects in other policy domains.

References

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

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

Galí, J., López-Salido, J. D., & Vallés, J. (2007). Understanding the effects of government spending on consumption. *Journal of the European Economic Association*, *5*(1), 227–270.

Gordon, R. J., & Krenn, R. (2010). *The end of the great depression 1939–41: Policy contributions and fiscal multipliers* (NBER Working Paper No. 16380). National Bureau of Economic Research.

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.

Christiano, L., Eichenbaum, M., & Rebelo, S. (2011). When is the government spending multiplier large? *Journal of Political Economy*, *119*(1), 78–121.

Ramey. (2011). Identifying government spending shocks: It's all in the timing. *Quarterly Journal of Economics*, *126*(1), 1–50.

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

Born, B., & Müller, G. J. (2012). Government spending shocks in quarterly and annual time series. *Journal of Money, Credit and Banking*, *44*(2-3), 507–517.

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

Blanchard, O., Lorenzoni, G., & L'Huillier, J.-P. (2013). News, noise, and fluctuations: An empirical exploration. *American Economic Review*, *103*(7), 3045–3070.

Cloyne, J. (2013). Discretionary tax changes and the macroeconomy: New narrative evidence from the united kingdom. *American Economic Review*, *103*(4), 1507–28.

Ilzetzki, E., Mendoza, E. G., & Végh, C. A. (2013). How big (small?) are fiscal multipliers? *Journal of Monetary Economics*, *60*(2), 239–254.

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

Corsetti, G., Meier, A., & Müller, G. J. (2014). What determines government spending multipliers? *Economic Policy*, *27*(72), 521–565.

Nickel, C., & Tudyka, A. (2014). Fiscal stimulus in times of high debt: Reconsidering multipliers and twin deficits. *Journal of Money, Credit and Banking*, *46*(7), 1313–1344.

Beetsma, R., Cimadomo, J., Furtuna, O., & Giuliodori, M. (2015). The confidence effects of fiscal consolidations. *Economic Policy*, *30*(83), 439–489.

Falagiarda, M., & Gregori, W. D. (2015). The impact of fiscal policy announcements by the italian government on the sovereign spread: A comparative analysis. *European Journal of Political Economy*, *39*, 288–304.

Forni, M., & Gambetti, L. (2016). Government spending shocks in open economy vars. *Journal of International Economics*, *99*, 68–84.

Ramey. (2016). Chapter 2 - macroeconomic shocks and their propagation. In J. B. Taylor & H. Uhlig (Eds.). Elsevier.

Ricco, G., Callegari, G., & Cimadomo, J. (2016). Signals from the government: Policy disagreement and the transmission of fiscal shocks. *Journal of Monetary Economics*, *82*, 107–118.

Ben Zeev, N., & Pappa, E. (2017). Chronicle of a war foretold: The macroeconomic effects of anticipated defence spending shocks. *The Economic Journal*, *127*(603), 1568–1597.

Kilian, L., & Lütkepohl, H. (2017). *Structural vector autoregressive analysis*. Cambridge University Press.

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

Hamilton, J. D. (2018). Why you should never use the hodrick-prescott filter. *The Review of Economics and Statistics*, *100*(5), 831–843.

Ramey & Zubairy. (2018). Government spending multipliers in good times and in bad: Evidence from us historical data. *Journal of Political Economy*, *126*(2), 850–901.

Afonso, A., Jalles, J. T., & Kazemi, M. (2020). The effects of macroeconomic, fiscal and monetary policy announcements on sovereign bond spreads. *International Review of Law and Economics*, *63*, 105924.

Cloyne, J., Jordà, Ò., & Taylor, A. M. (2020, September). *Decomposing the fiscal multiplier* (Working Paper No. 2020-12). Federal Reserve Bank of San Francisco.

Huidrom, R., Kose, M. A., Lim, J. J., & Ohnsorge, F. L. (2020). Why do fiscal multipliers depend on fiscal positions? *Journal of Monetary Economics*, *114*, 109–125.

Beetsma, R., Furtuna, O., Giuliodori, M., & Mumtaz, H. (2021). Revenue- versus spending-based fiscal consolidation announcements: Multipliers and follow-up. *Journal of International Economics*, *131*, 103455.

Klein, M., & Winkler, R. (2021). The government spending multiplier at the zero lower bound: International evidence from historical data. *Journal of Applied Econometrics*, *36*(6), 744–759.

Larsen, V. H., Thorsrud, L. A., & Zhulanova, J. (2021). News-driven inflation expectations and information rigidities. *Journal of Monetary Economics*, *117*, 507–520.

Andres-Escayola, E., Ghirelli, C., Molina, L., Pérez, J. J., & Vidal, E. (2022). *Using newspapers for textual indicators: which and how many?* (Working Papers No. 2235). Banco de España.

Angelico, C., Marcucci, J., Miccoli, M., & Quarta, F. (2022). Can we measure inflation expectations using twitter? *Journal of Econometrics*, *228*(2), 259–277.

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.

Ghassibe, M., & Zanetti, F. (2022). State dependence of fiscal multipliers: The source of fluctuations matters. *Journal of Monetary Economics*, *132*, 1–23.

Hayo, B., & Mierzwa, S. (2022). Legislative tax announcements and gdp: Evidence from the united states, germany, and the united kingdom. *Economics Letters*, *216*, 110548.

Ilori, A. E., Paez-Farrell, J., & Thoenissen, C. (2022). Fiscal policy shocks and international spillovers. *European Economic Review*, *141*, 103969.

Melosi, L., Morita, H., & Zanetti, F. (2022). *The signaling effects of fiscal announcements* (Working Paper No. 2022-38). Federal Reserve Bank of Chicago.

Ascari, G., Beck-Friis, P., Florio, A., & Gobbi, A. (2023). Fiscal foresight and the effects of government spending: It's all in the monetary-fiscal mix. *Journal of Monetary Economics*, *134*, 1–15.

Aruoba, S. B., & Drechsel, T. (2024). *Identifying monetary policy shocks: A natural language approach* (NBER Working Paper No. 32417). National Bureau of Economic Research.

Latifi, A., Naboka-Krell, V., Tillmann, P., & Winker, P. (2024). Fiscal policy in the bundestag: Textual analysis and macroeconomic effects. *European Economic Review*, *168*, 104827.

Gallegos, J. E., García-Miralles, E., Kataryniuk, I., & Párraga Rodríguez, S. (2025). *Fiscal announcements and households' beliefs: Evidence from the euro area* (Documentos de Trabajo No. 2543). Banco de España.

Granziera, E., Larsen, V. H., Meggiorini, G., & Melosi, L. (2025). *Speaking of inflation: The influence of fed speeches on expectations* (Norges Bank Working Paper No. 10). Norges Bank.

Smets, F., & Trienens, L. (2025). *Identifying fiscal surprise and news shocks: A natural language approach. effects on u.s. business cycles, inflation, and the yield curve* [Preliminary version], Bank of International Settlement (BIS) and University of Goettingen, Chair for Econometrics.

Tillmann, P. (2025). Monetary policy and the fiscal talk in the bundestag.

Appendix

Table A.1. Prompts used in the pre–filtering stage: fiscal vs. non–fiscal classification

Index	Prompt used for fiscal relevance filtering
GovExpSent	"**Nature of the measure**: expenditure–related / non–expenditure–related. A news item should be classified as **government–expenditure–related** when it refers, even indirectly, to decisions, plans or policies that involve public spending or the use of state resources in a broad sense. This includes not only explicit announcements of expenditure but also any measure implying a budgetary impact, a reallocation of funds, an expansion of public programmes, or the creation of mechanisms financed with public money. Typical examples of expenditure-related measures include the following: (i) Increases or decreases in spending items (investment, transfers, subsidies, social spending, health, education, defence, pensions, public employment, infrastructure); (ii) Programmes, funds, or plans (national, regional, or European) that entail public disbursements, even if the exact figures or funding sources are not specified; (iii) Any political or institutional announcement that implies or presupposes public expenditure, even if not explicitly mentioned (e.g., new plan to improve schools, strengthening of the health system, aid to firms affected by the crisis). A news item should be classified as **non–government–expenditure–related** only when there is no plausible connection to public expenditure or the use of state funds, and the article refers exclusively to regulatory, legal, political or macroeconomic issues (e.g. costless legislative reforms, monetary policy, parliamentary debates without budgetary measures, labour market or environmental policies without public funding). → **Inclusion rule:** in case of doubt, if there is any reasonable possibility that the measure involves a public–spending component or an allocation of state resources, it should be classified as fiscal."
TaxSent	"**Nature of the measure**: tax–related / non–tax–related. A news item should be classified as **tax–related** when it refers, even indirectly, to decisions, plans or policies affecting the tax burden or government revenues in a broad sense. This includes not only explicit changes in rates but also any measure that modifies the tax obligation, the tax base, exemptions/deductions/allowances, deadlines, or the administration and enforcement of taxes. Typical examples of tax-related measures include the following: i) Creation, elimination or modification of taxes (personal income tax, VAT, corporate tax, excise duties, green taxes), social contributions, fees, or special levies; i) Changes in rates, brackets, exemptions, deductions, credits, surcharges, withholdings or assessed bases; iii) Tax amnesties, payment deferrals or instalment plans, tax holidays, cadastral revaluations with revenue effects, anti–fraud measures, or new reporting obligations expected to affect revenue or the effective burden; iv) Explicit announcements about not raising or not cutting taxes, extensions of tax benefits or temporary schemes, and national or EU–wide tax harmonisation or coordination with effects on the tax burden. A news item should be classified as **non–tax–related** only when there is no plausible connection to taxation or revenue, and the article refers solely to regulatory, legal, political or macroeconomic aspects without fiscal instruments (e.g. sectoral reforms without tax implications, monetary policy, labour market or environmental measures without taxes or tax benefits, subsidies or spending unrelated to tax changes). → **Inclusion rule:** in case of doubt, if there is any reasonable possibility that the measure alters government revenue or the tax burden (directly or indirectly), it should be classified as tax–related."

Notes: Each prompt was supplied to the OpenAI API as part of the JSON–based classification task. The model `gpt-4o-mini` used these rules to determine whether each article was fiscally relevant before inclusion in the sentiment indices. Prior to classification, all articles were translated into Spanish to ensure consistent application of the rules.

Table A.2. Classification of public spending news in Spain: GPT-4o mini vs. manual classification

	Outcome			
	Increase	Decrease	Neutral	Total
GPT-4o-mini	771	1,127	227	2,125
	36.3%	53.0%	10.7%	
Manual classification	800	1,070	255	2,125
	37.6%	50.4%	12%	
Disagreement (GPT vs Manual)	−29	57	−28	
	-1.3 p.p.	+2.6 p.p.	-1.3 p.p.	

Notes: The top row of each method reports the number of news items; the row in *small italics* below reports percentages over the total for that method. The row "Disagreement (GPT vs Manual)" shows the absolute difference in the number of news items and in percentage points (p.p.).

Table A.3. Classification of tax news in Spain: `gpt-4o-mini`vs. manual classification

	Outcome			
	Increase	Decrease	Neutral	Total
GPT-4o mini	1,492	1,617	116	3,167
	47.1%	51.0%	2.9%	100%
Manual classification	1,643	1,375	149	3,167
	51.9%	43.4%	4.7%	100%
Disagreement (GPT vs Manual)	−151	300	−149	
	-4.8 p.p.	+9.5 p.p.	-4.7 p.p.	

Notes: The top row of each method reports the number of news items; the row in *small italics* below reports percentages over the total for that method. The row "Disagreement (GPT vs Manual)" shows the absolute difference in the number of news items and in percentage points (p.p.).

Table A.4. Prompts used in the direction classification

Index	Prompt used for direction of fiscal announcement
TaxSent	"**Direction of fiscal policy:** increase / decrease / neutral. **Increase (Up).** The news item should be classified as an increase in taxation when it reports the announcement of new taxes or an increase in the tax burden (e.g., higher personal income tax, VAT, excise duties, social contributions, new levies, or broader tax bases). It should also be classified as an increase when: (i) the postponement of tax reforms that would have implied a tax cut is announced; (ii) it is explicitly stated that taxes will *not be reduced*; or (iii) the *extension, renewal, or maintenance* of existing taxes is confirmed. **Decrease (Down).** The news item should be classified as a decrease in taxation when it reports the announcement of a tax cut or reduction in the tax burden (e.g., lower tax rates, elimination of taxes, new deductions, rebates, or lower social contributions, reductions in VAT or energy-related taxes). It should also be classified as a decrease when: (i) an existing reduction is *extended or renewed*; (ii) it is stated that tax rates will *not be raised*; or (iii) a reform that would have increased taxes is *postponed or cancelled*. **Neutral.** If the direction cannot be clearly inferred from the text, it should be classified as Neutral."
GovExpSent	"**Direction of fiscal policy:** increase / decrease / neutral. **Increase**. The news item should be classified as an increase in government expenditure when it reports the announcement of higher public spending (e.g., creation of new budget lines, higher transfers, pension increases, public wage rises, greater investment in infrastructure, health, education, or defence, or extensions of subsidies or social benefits). It should also be classified as an increase when: (i) extraordinary spending measures that were expected to expire are *maintained or extended*; or (ii) it is announced that planned spending cuts *will not take place*. **Decrease**. The news item should be classified as a decrease in government expenditure when it reports cuts, reductions, or freezes in public spending (e.g., lower investment, reductions in social transfers or benefits, limits on subsidies, lower public wages, or cuts in health or education). It should also be classified as a decrease when: (i) an existing austerity or consolidation policy is *extended*; or (ii) it is announced that spending increases *will not be implemented* as previously expected. **Neutral**. If the direction cannot be clearly inferred from the text, it should be classified as Neutral."

Notes: These prompts were provided to the `gpt-4o-mini` model to classify the fiscal direction of each article after the pre–filtering stage. Each news item is labelled as Increase, Decrease, or Neutral within its respective corpus (*TaxSent* or *GovExpSent*).

Table A.5. Prompt used for the classification of fiscal announcement jurisdiction

Prompt used for jurisdiction classification
"**Jurisdiction of the announcement: national / regional / municipal. National:** measures adopted by the central Government of Italy, the national Parliament, or the ministries (for example, the Ministry of Economy and Finance, the Ministry of Infrastructure, the Ministry of Defence, etc.) that affect the entire country or several regions simultaneously. This category also includes announcements made by the President of the Council of Ministers or by the Council of Ministers itself. **Regional:** announcements issued by regional governments (*Giunte regionali*) or regional presidents (*Presidenti di Regione*) that apply only within a specific region (for example, Lombardia, Veneto, Lazio, Sicilia, Toscana, Emilia-Romagna, etc.). Measures approved by Regional Councils that exclusively affect their own territory are also included. **Municipal:** local or communal decisions or announcements made by municipalities (*Comuni*), mayors (*Sindaci*), or municipal councils that affect only a specific city or municipality."

Notes: These prompts were provided to the `gpt-4o-mini` model to classify the jurisdiction of each article after the pre–filtering stage. Each news item is labelled as National, Regional, or Municipal within its respective corpus (*TaxSent* or *GovExpSent*).

Table A.6. Prompt used for the classification of the type of fiscal communication

Prompt used for type–of–communication classification
"**Type of communication: announcement / recommendation / rumour.** **Announcement:** official communication by a competent authority (President, ministers, regional councillors, mayors, etc.). Firm political commitments, even if referring to measures planned for future implementation, are also considered announcements. **Recommendation:** proposals or demands made by actors without direct executive power (opposition parties, trade unions, employers' associations, experts, central banks, or advisory bodies). **Rumour:** unconfirmed leaks or speculations, ongoing negotiations, or expressions such as "the government is considering", "plans to", or "could approve", without official confirmation by a competent authority."

Notes: These prompts were provided to the `gpt-4o-mini` model to classify the type of communication of each article after the pre–filtering stage. Each news item is labelled as Announcement, Recommendation, or Rumour within its respective corpus (*TaxSent* or *GovExpSent*).

Table A.7. VAR specifications for panel Government Expenditure and Tax models

Panel A. Government Expenditure VAR

Baseline

$[\,GovExp_t \quad GDP_t \quad Cons_t \quad GovExpSent_t\,]$

Domestic business cycle

$[\,GovExp_t \quad GDP_t \quad Cons_t \quad Investment_t \quad GovExpSent_t\,]$

Fiscal–monetary

$[\,GovExp_t \quad GDP_t \quad Cons_t \quad InterestRate_t \quad Inflation_t \quad TaxRevenues_t \quad GovExpSent_t\,]$

Expectations

$[\,GovExp_t \quad GDP_t \quad Cons_t \quad BusinessExp_t \quad GovExpSent_t\,]$

Open–economy

$[\,GovExp_t \quad GDP_t \quad Cons_t \quad TradeBalance_t \quad GovExpSent_t\,]$

Sovereign–risk

$[\,GovExp_t \quad GDP_t \quad Cons_t \quad Stress_t \quad GovExpSent_t\,]$

Panel B. Tax VAR

Baseline

$[\,TaxRev_t \quad GDP_t \quad Cons_t \quad TaxSent_t\,]$

Domestic business cycle

$[\,TaxRev_t \quad GDP_t \quad Cons_t \quad Investment_t \quad TaxSent_t\,]$

Fiscal–monetary

$[\,TaxRev_t \quad GDP_t \quad Cons_t \quad InterestRate_t \quad Inflation_t \quad GovernmentExpenditure_t \quad TaxSent_t\,]$

Expectations block

$[\,TaxRev_t \quad GDP_t \quad Cons_t \quad BusinessExp_t \quad TaxSent_t\,]$

Open–economy

$[\,TaxRev_t \quad GDP_t \quad Cons_t \quad TradeBalance_t \quad TaxSent_t\,]$

Sovereign–risk

$[\,TaxRev_t \quad GDP_t \quad Cons_t \quad Stress_t \quad TaxSent_t\,]$

Notes: This table reports the endogenous variable sets used in each panel VAR specification. Panel A lists the models for government expenditure sentiment ($GovExpSent_t$), while Panel B reports the analogous models for tax sentiment ($TaxSent_t$). All real variables are expressed as log() × 100, whereas the interest rate and inflation are in natural units, and the sentiment indicators are standardised. The VAR lag length is $p = 4$ in all specifications.

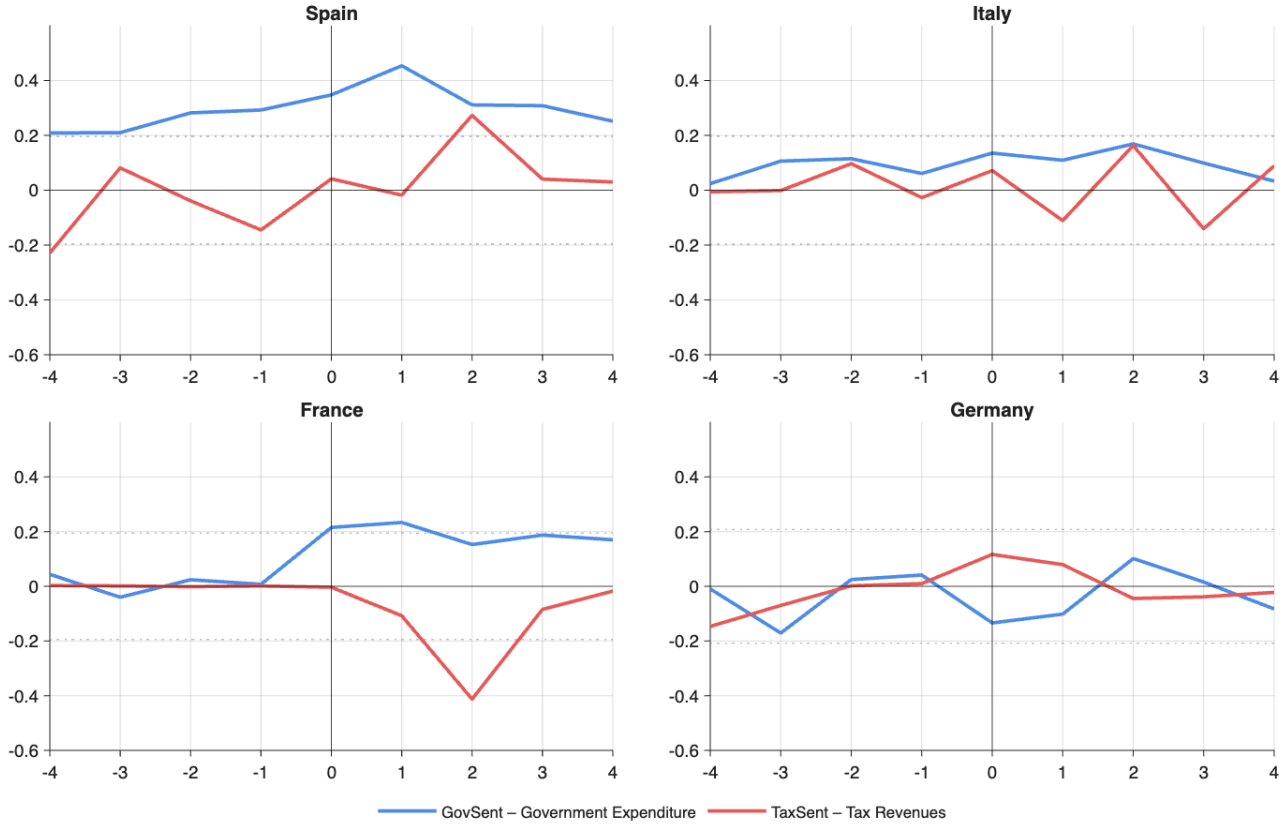

Fig. A.1. Cross-correlations between fiscal sentiment indices and fiscal aggregates. *Notes:* The figure plots the correlation between the fiscal announcement indices in quarter t and their corresponding fiscal aggregates in $t \pm k$ for $k = 0, \ldots, 4$ across Spain, Italy, France, and Germany.

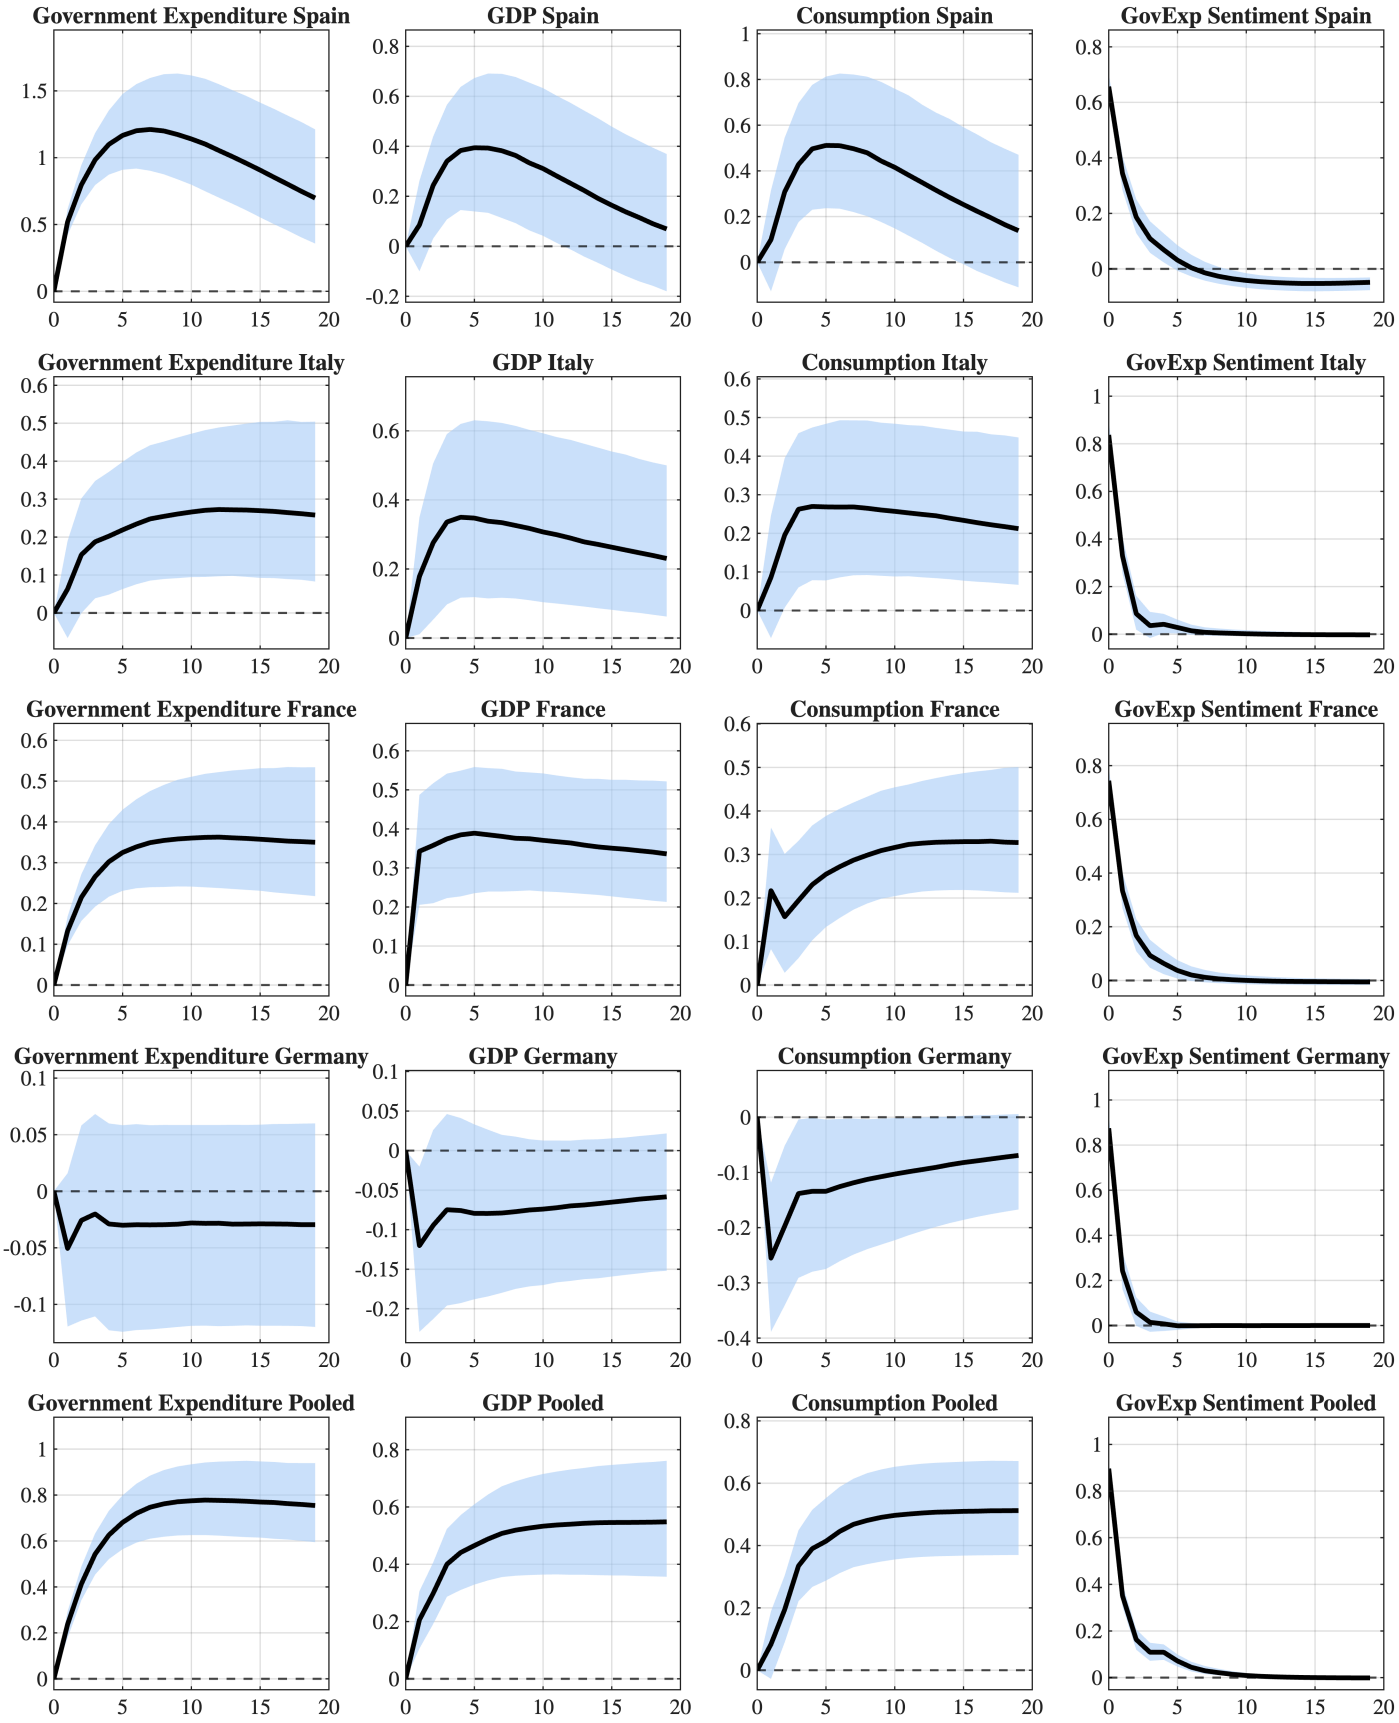

Fig. A.2. Responses to a government expenditure sentiment (*GovExpSent*) shock. *Notes:* The figure shows the responses of government expenditure, real GDP, private consumption, and the government expenditure sentiment index to a one–standard–deviation increase in the *GovExpSent* indicator. The panels display results for Spain, Italy, France, Germany, and for the pooled VAR that combines the four economies. All responses are derived from recursively identified Bayesian VAR models with four lags and a Normal–Wishart prior, where the variables are ordered as government expenditure, GDP, private consumption, and *GovExpSent*. The shaded areas cover 68% of all draws.

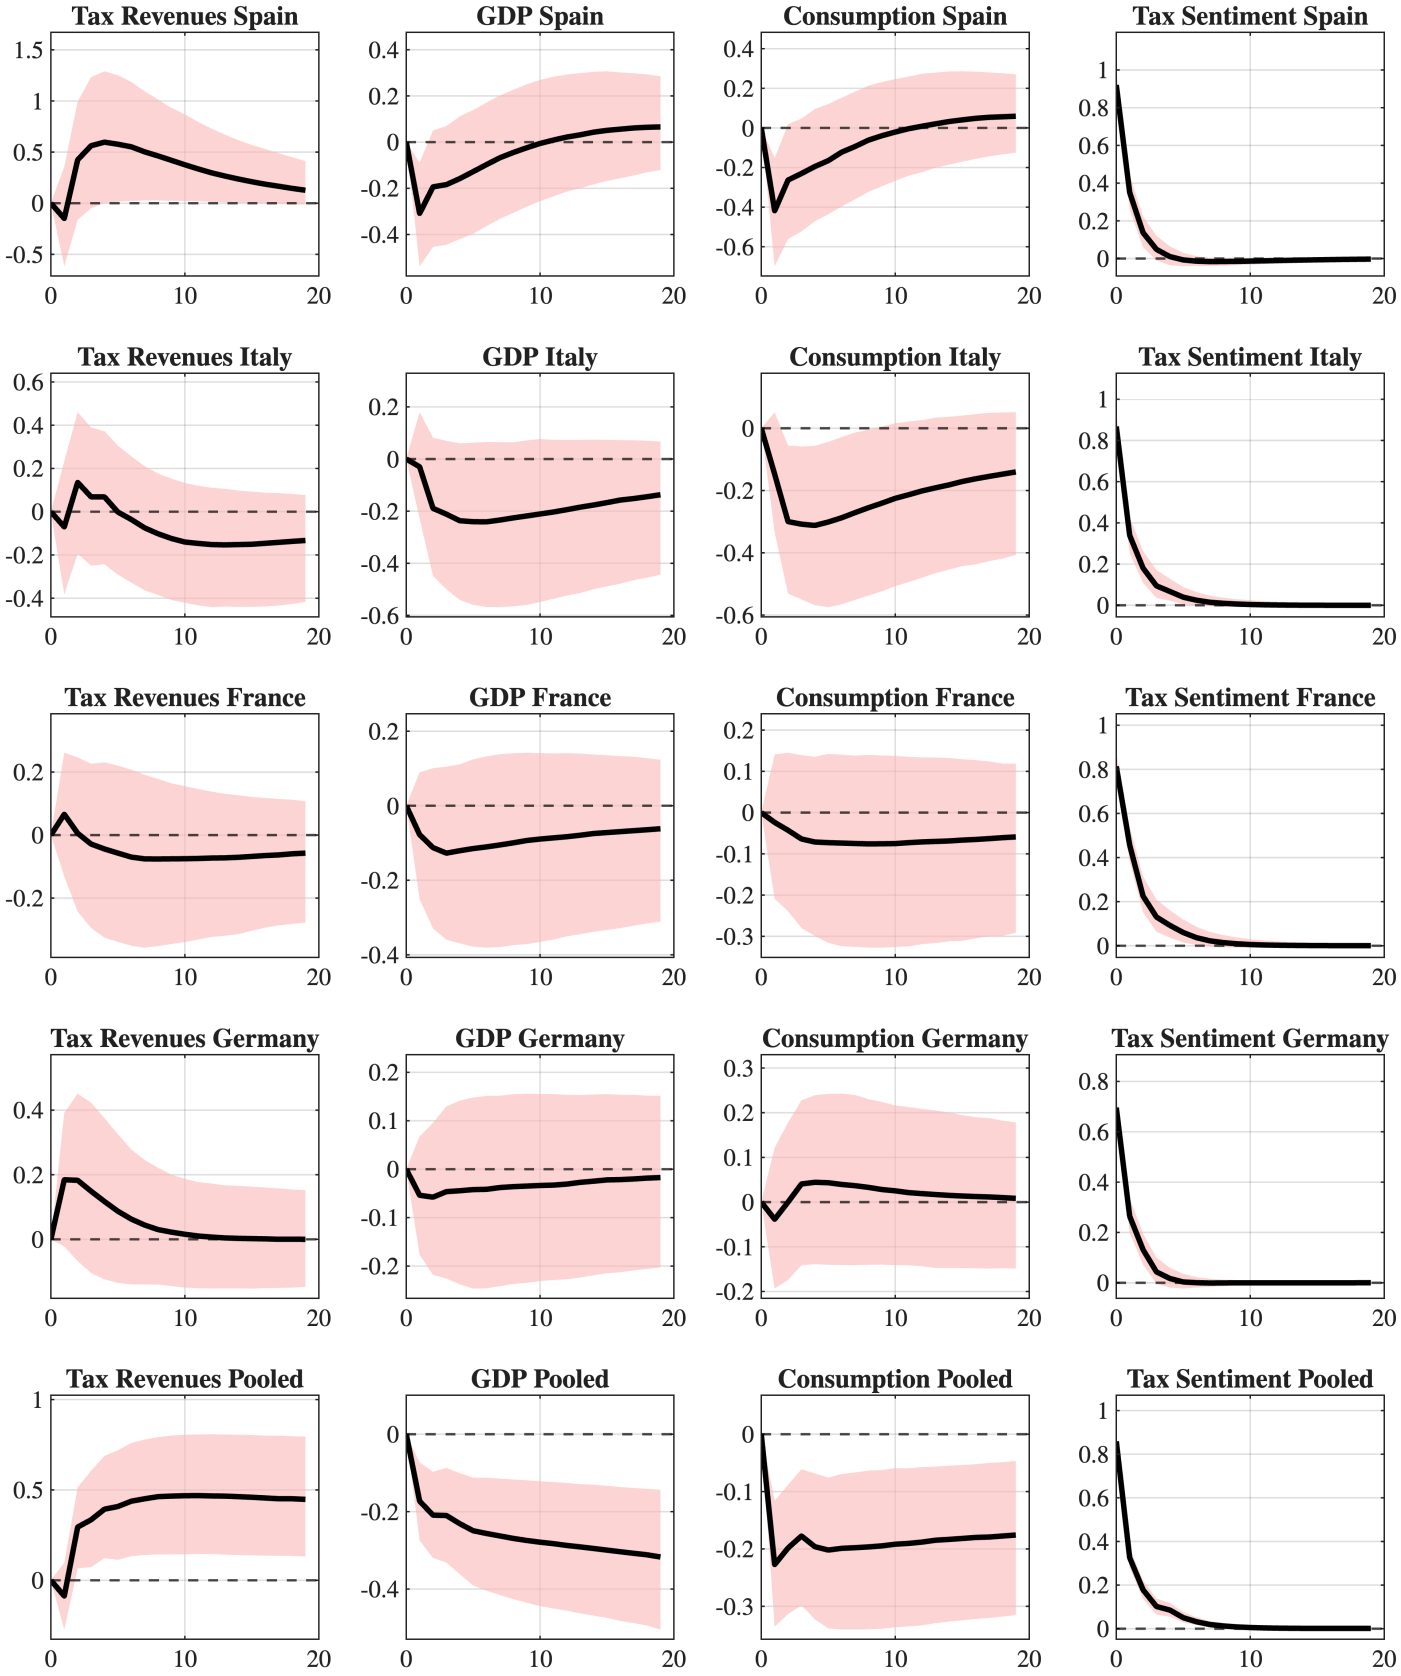

Fig. A.3. Responses to a tax sentiment (*TaxSent*) shock. *Notes:* The figure shows the responses of tax revenues, real GDP, private consumption, and the tax sentiment index to a one–standard–deviation increase in the *TaxSent* indicator. The panels display results for Spain, Italy, France, Germany, and for the pooled VAR that combines the four economies. All responses are derived from recursively identified Bayesian VAR models with four lags and a Normal–Wishart prior, where the variables are ordered as tax revenues, GDP, private consumption, and *TaxSent*. The shaded areas cover 68% of all draws.

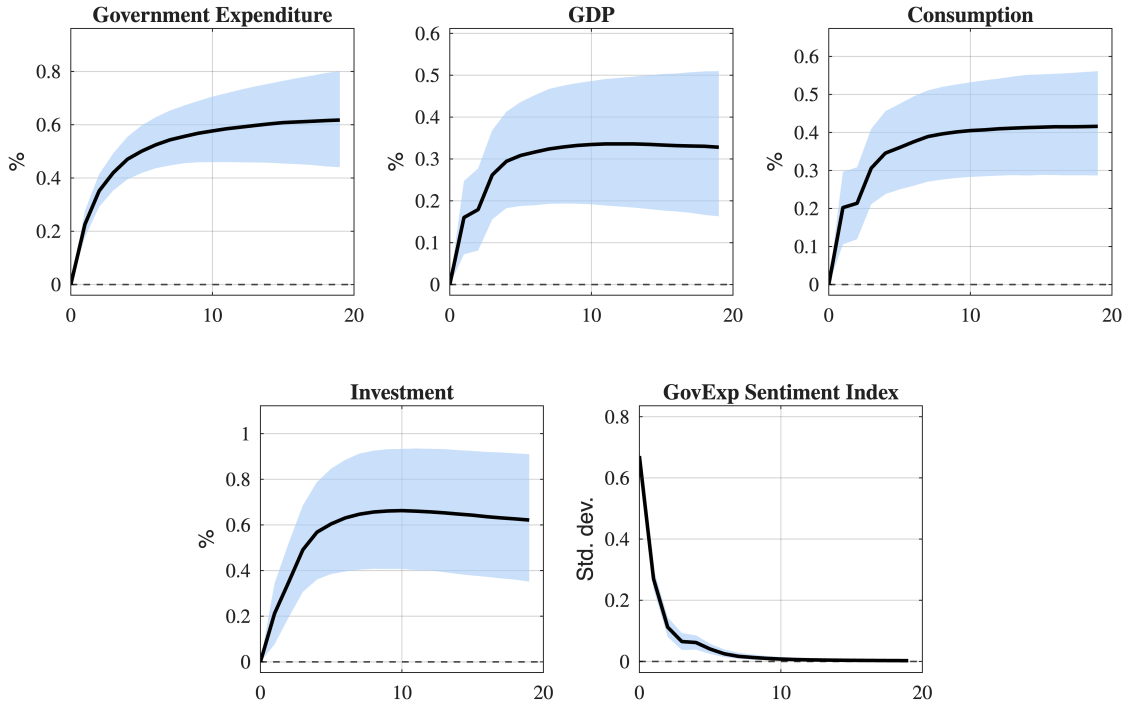

Fig. A.4. Response of the domestic business cycle block variables to a *GovExpSent* shock in the panel VAR. *Notes:* The figure reports the impulse–response functions of government expenditure, real GDP, private consumption, investment and the government expenditure sentiment index to a one–standard–deviation increase in the *GovExpSent* indicator for Spain, Italy, France, and Germany. Shaded areas denote 68% posterior credibility sets.

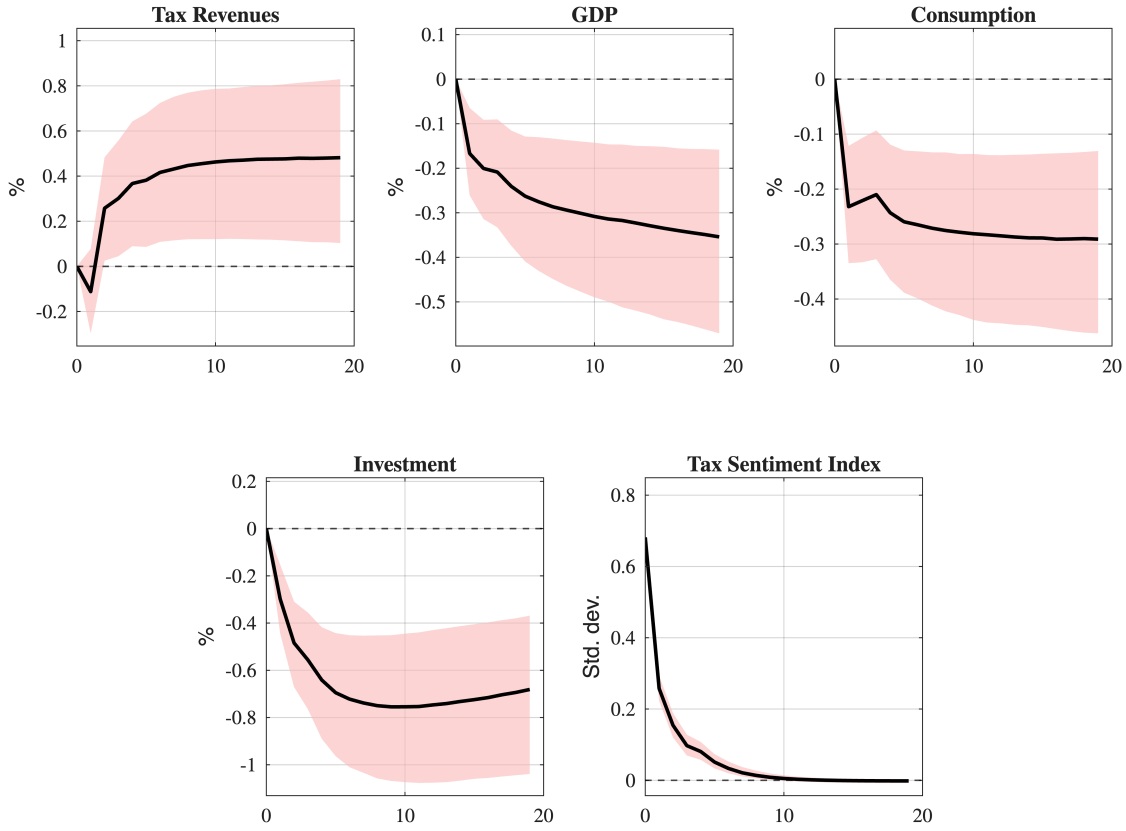

Fig. A.5. Response of the domestic business cycle block variables to a *TaxSent* shock in the panel VAR.. *Notes:* The figure reports the impulse–response functions of tax revenues, real GDP, private consumption, investment and the tax sentiment index to a one–standard–deviation increase in the *TaxSent* indicator for Spain, Italy, France, and Germany. Shaded areas denote 68% posterior credibility sets.

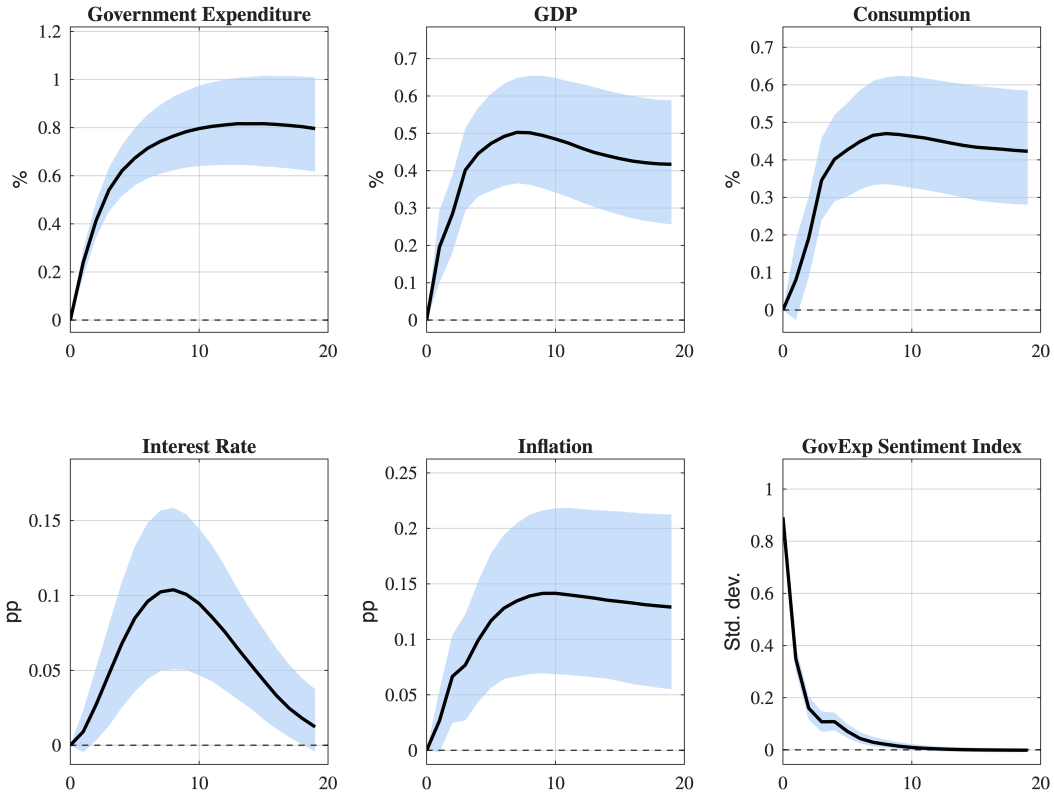

Fig. A.6. Response of the fiscal-monetary block variables to a *GovExpSent* shock in the panel VAR. *Notes:* The figure reports the impulse–response functions of government expenditure, real GDP, private consumption, inflation, interest rate and the government expenditure sentiment index to a one–standard–deviation increase in the *GovExpSent* index. Shaded areas denote 68% posterior credibility sets.

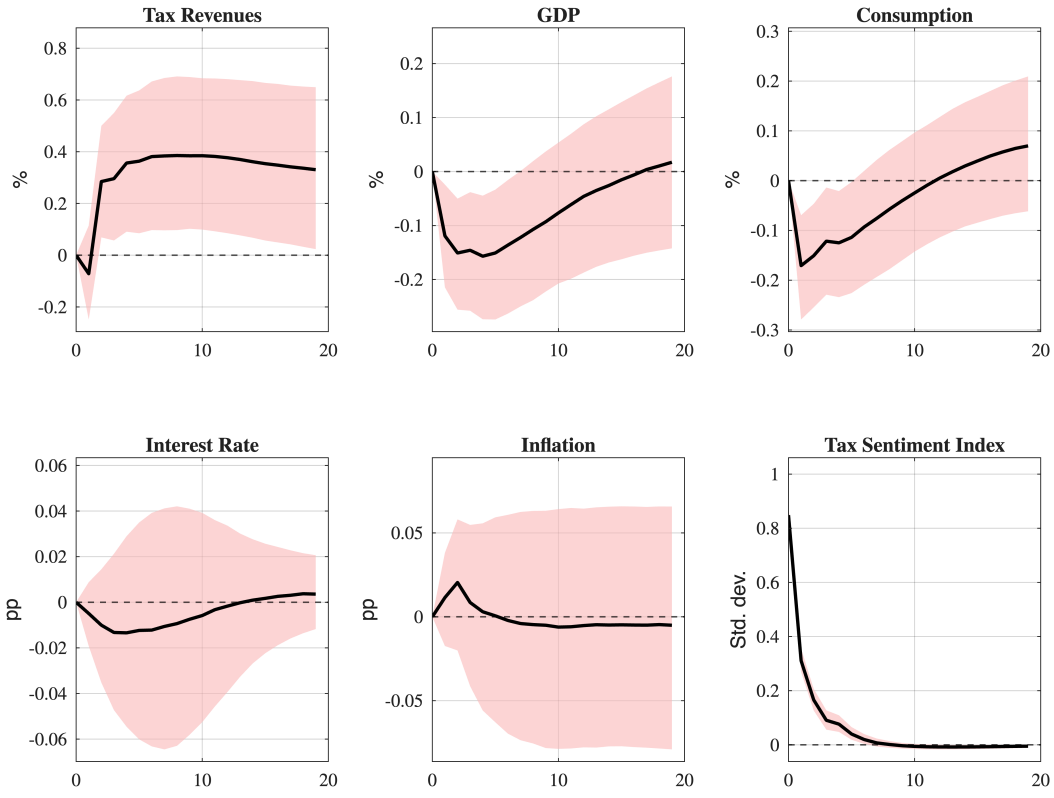

Fig. A.7. Response of the fiscal-monetary block variables to a *TaxSent* shock in the panel VAR. *Notes:* The figure reports the impulse–response functions of tax revenues, real GDP, private consumption, inflation, interest rate and the tax sentiment index to a one–standard–deviation increase in the *TaxSent* index. Shaded areas denote 68% posterior credibility sets.

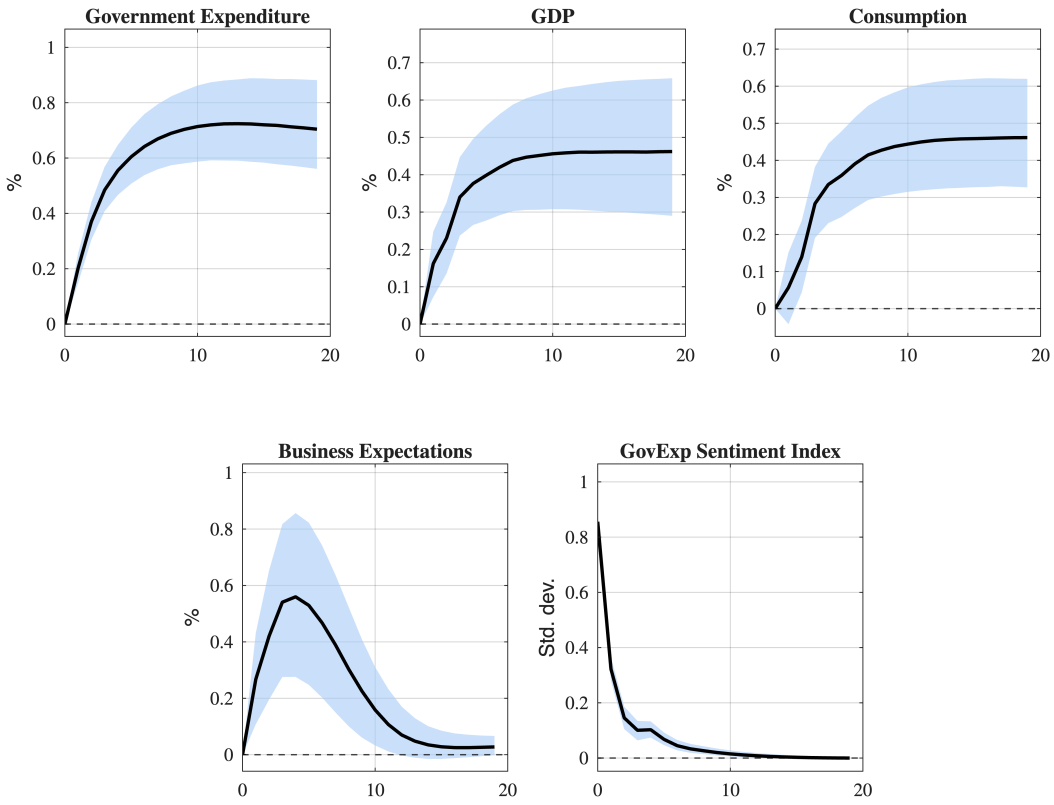

Fig. A.8. Response of the expectations block variables to a *GovExpSent* shock in the panel VAR. *Notes:* The figure reports the impulse–response functions of government expenditure, real GDP, private consumption, business expectations and the government expenditure sentiment index to a one–standard–deviation increase in the *GovExpSent* index. Shaded areas denote 68% posterior credibility sets.

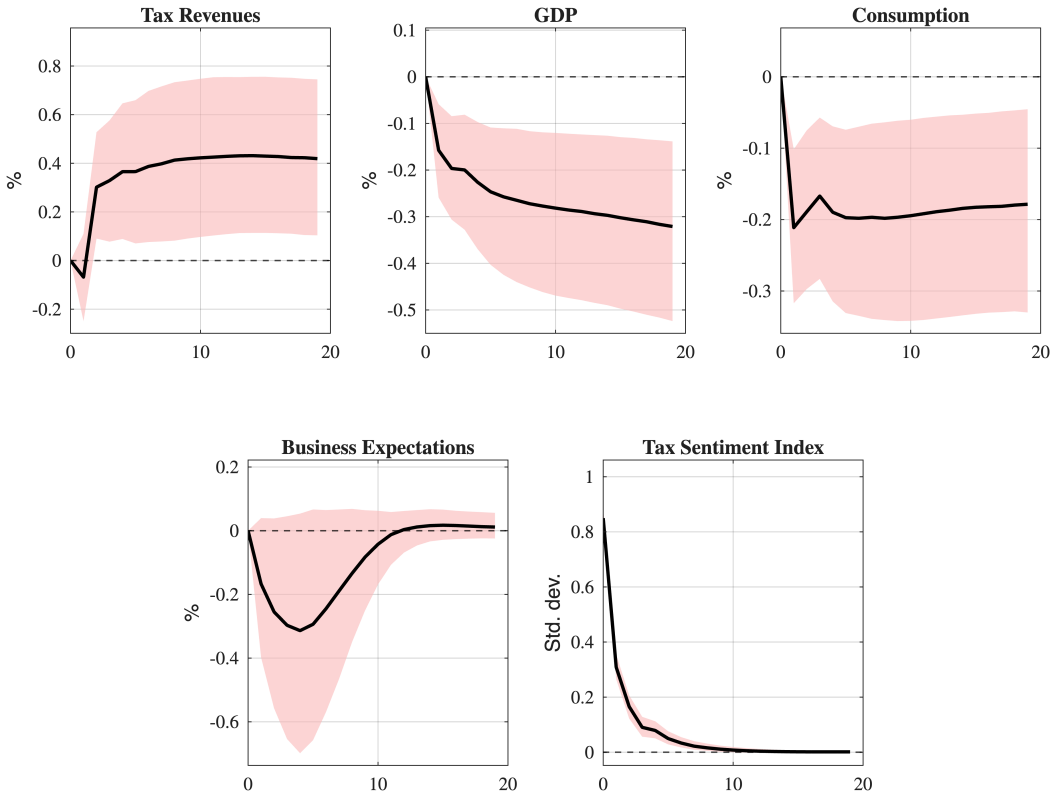

Fig. A.9. Response of the expectations block variables to a *TaxSent* shock in the panel VAR. *Notes:* The figure reports the impulse–response functions of tax revenues, real GDP, private consumption, business expectations and the tax sentiment index to a one–standard–deviation increase in the *TaxSent* index. Shaded areas denote 68% posterior credibility sets.

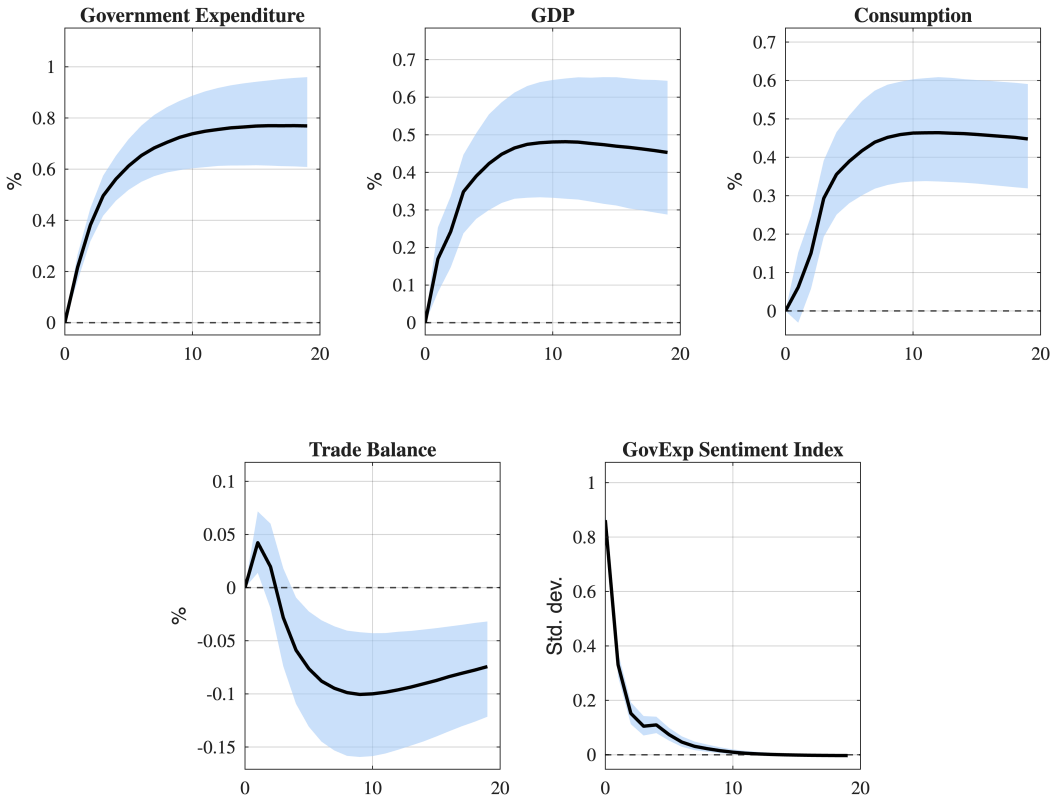

Fig. A.10. Responses of the open-economy block variables to a *GovExpSent* shock in the panel VAR. *Notes:* The figure reports the impulse–response functions of government expenditure, real GDP, private consumption, trade balance as a percentage of GDP and the government expenditure sentiment index to a one–standard–deviation increase in the *GovExpSent* index. Shaded areas correspond to 68% posterior credibility sets.

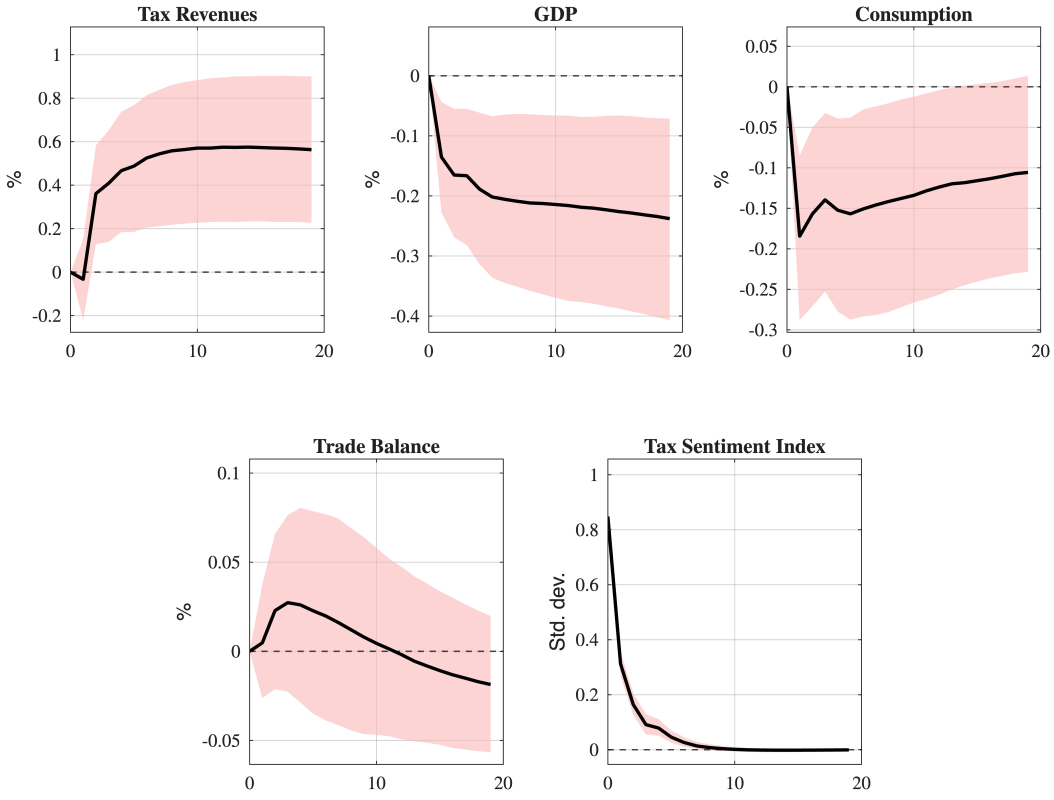

Fig. A.11. Response of the open-economy block to a *TaxSent* shock in the panel VAR. *Notes:* The figure reports the impulse–response functions of tax revenues, real GDP, private consumption, trade balance as a percentage of GDP and the tax sentiment index to a one–standard–deviation increase in the *TaxSent* index. Shaded areas correspond to 68% posterior credibility sets.

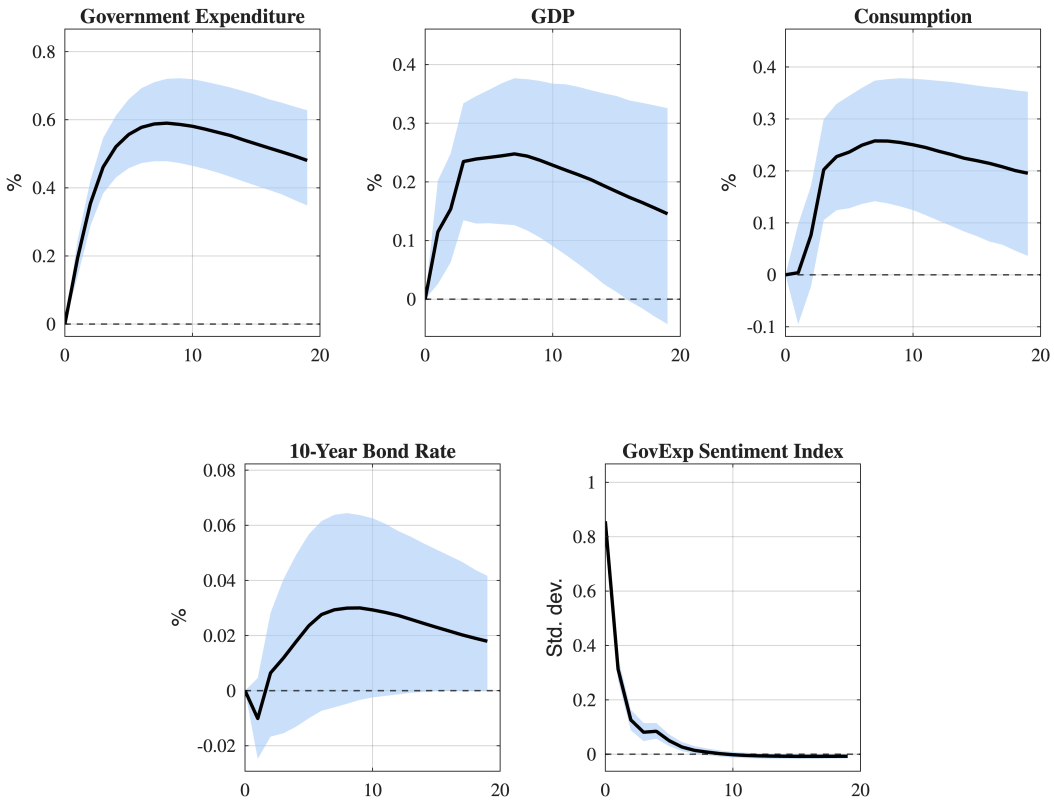

Fig. A.12. Responses of the sovereign–risk block variables to a *GovExpSent* shock in the panel VAR. *Notes:* The figure reports the impulse–response functions of government expenditure, real GDP, private consumption, the 10-year sovereign yield, and the government expenditure sentiment index to a one–standard–deviation increase in the *GovExpSent* index. Shaded areas correspond to 68% posterior credibility sets.

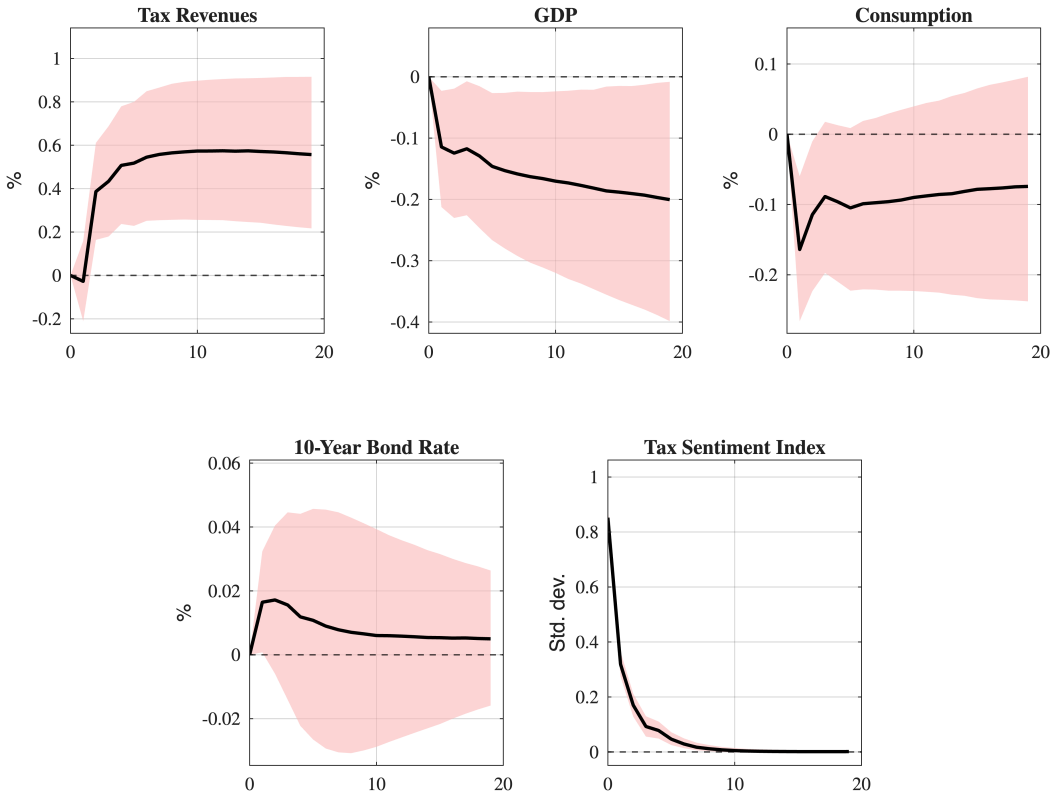

Fig. A.13. Responses of the sovereign–risk block variables to a *TaxSent* shock in the panel VAR. *Notes:* The figure reports the impulse–response functions of tax revenues, real GDP, private consumption, the 10-year sovereign yield, and the tax sentiment index to a one–standard–deviation increase in the *TaxSent* index. Shaded areas denote 68% posterior credibility sets.

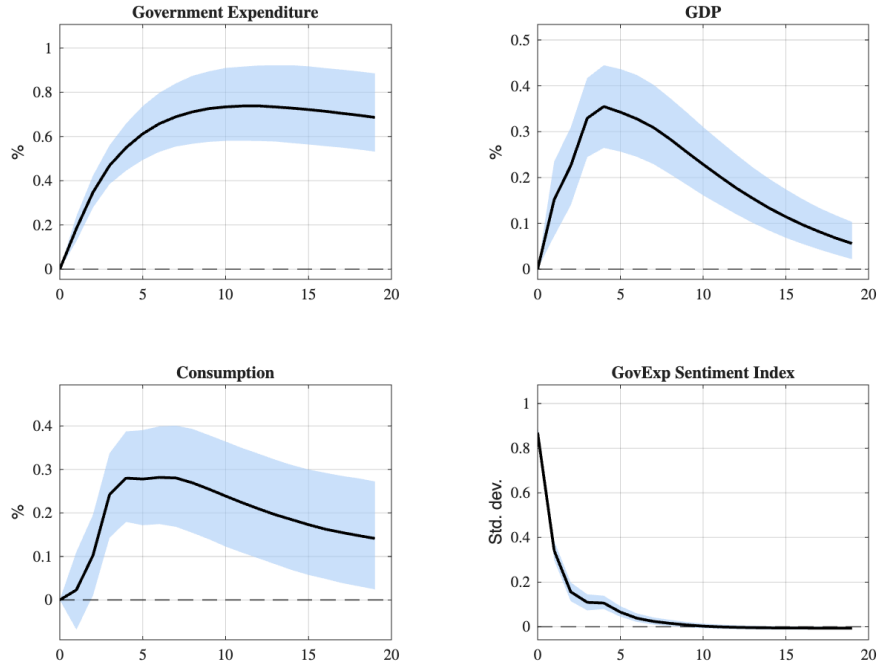

Fig. A.14. Responses of the baseline Gordon-detrended variables to a *GovExpSent* shock in the panel VAR. *Notes:* Impulse–response functions of government expenditure, real GDP, private consumption, and the *GovExpSent* index to a one–standard–deviation increase in government expenditure sentiment, estimated from a recursively identified Bayesian panel VAR with four lags and a Normal–Wishart prior. Macroeconomic series are Gordon-detrended. Shaded areas denote 68% credible intervals.

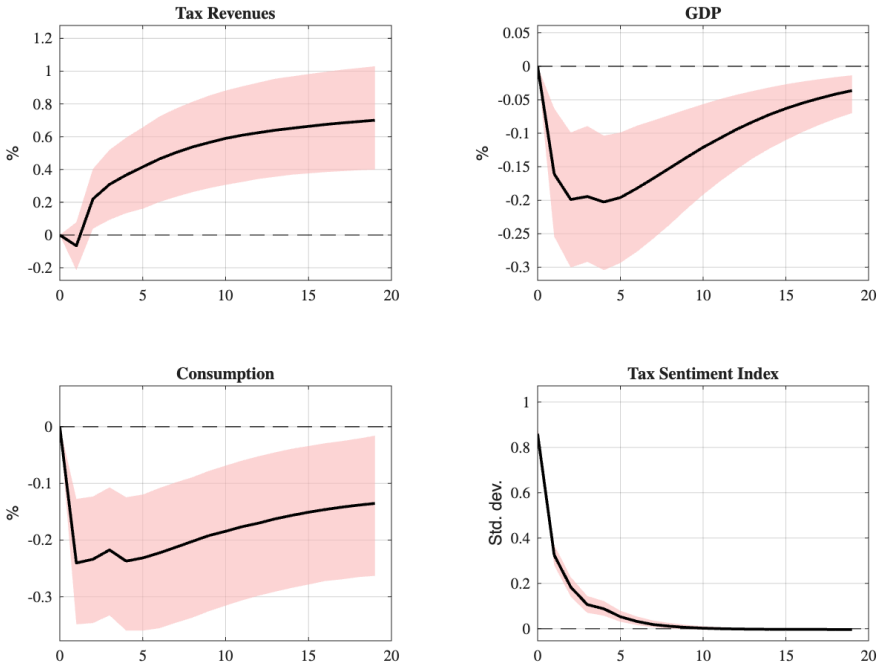

Fig. A.15. Responses of the baseline Gordon-detrended variables to a *TaxSent* shock in the panel VAR. *Notes:* Impulse–response functions of tax revenues, real GDP, private consumption, and the *TaxSent* index to a one–standard–deviation increase in tax sentiment, estimated from a recursively identified Bayesian panel VAR with four lags and a Normal–Wishart prior. Macroeconomic series are Gordon-detrended. Shaded areas denote 68% and 90% credible intervals.

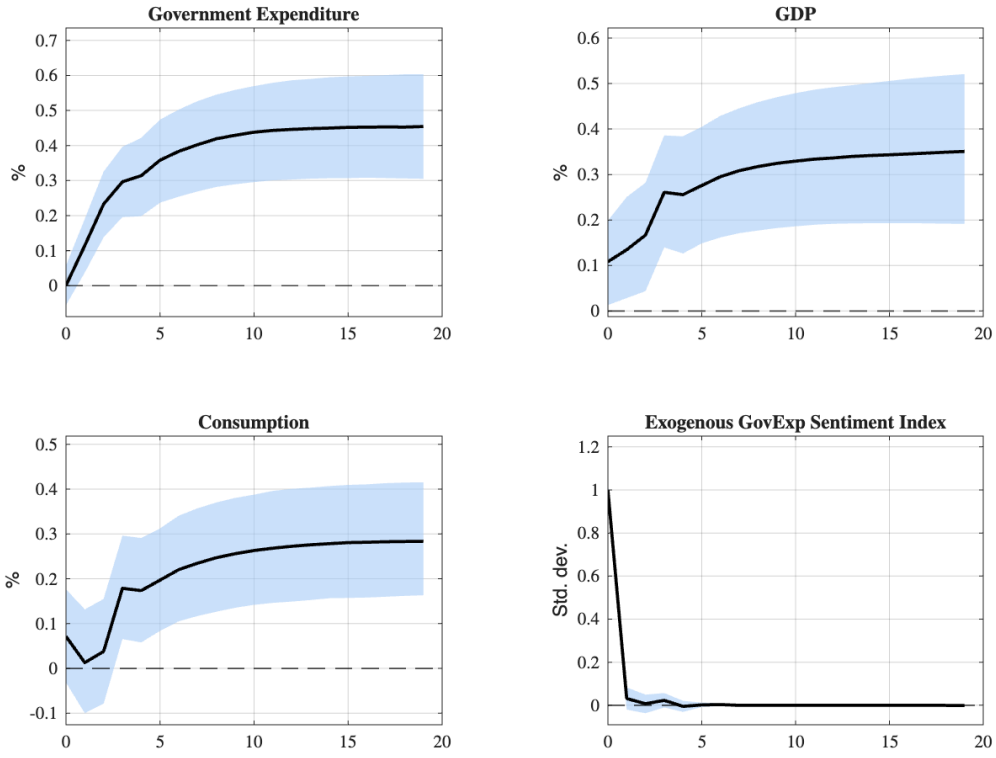

(a) Exogenous (purged) shock.

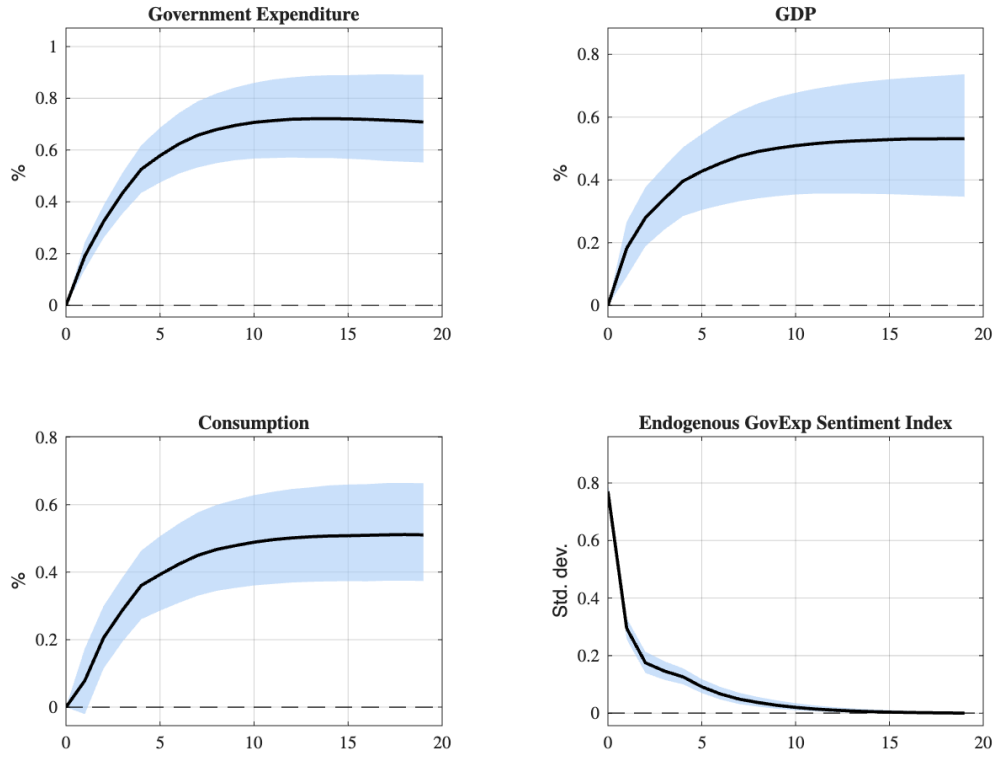

(b) Endogenous (systematic) component.

Fig. A.16. Purging decomposition for government expenditure announcements (*GovExpSent*). *Notes:* Impulse responses from the panel BVAR replacing *GovExpSent* with its exogenous (ARX residual) and endogenous (fitted) components. Shaded areas denote 68% credible intervals.

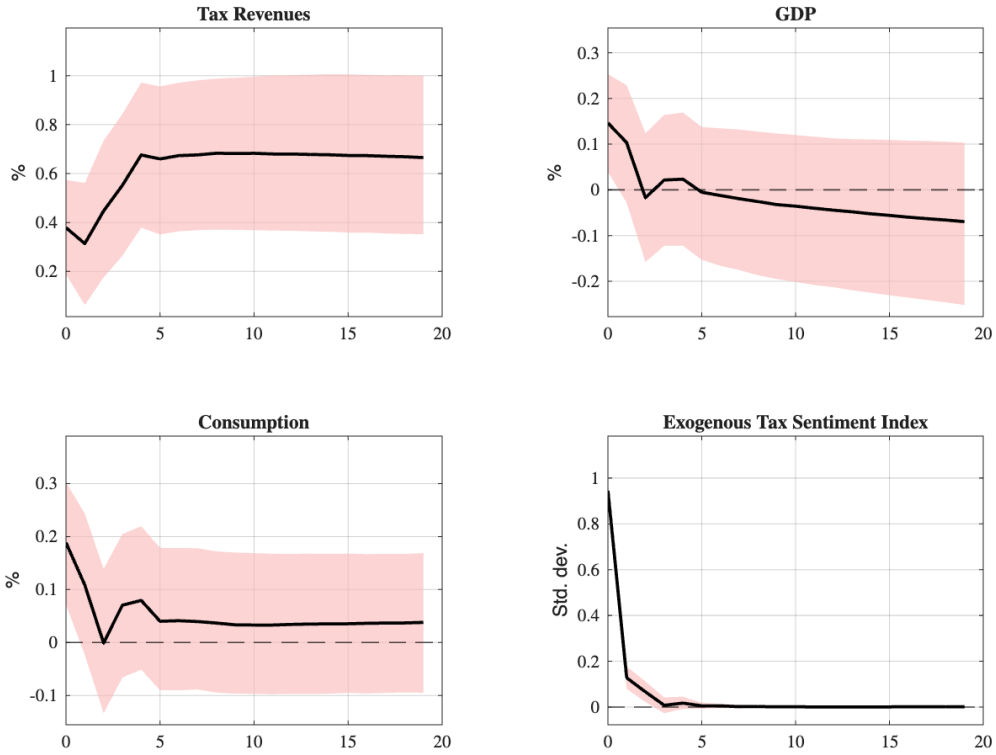

(a) Exogenous (purged) shock.

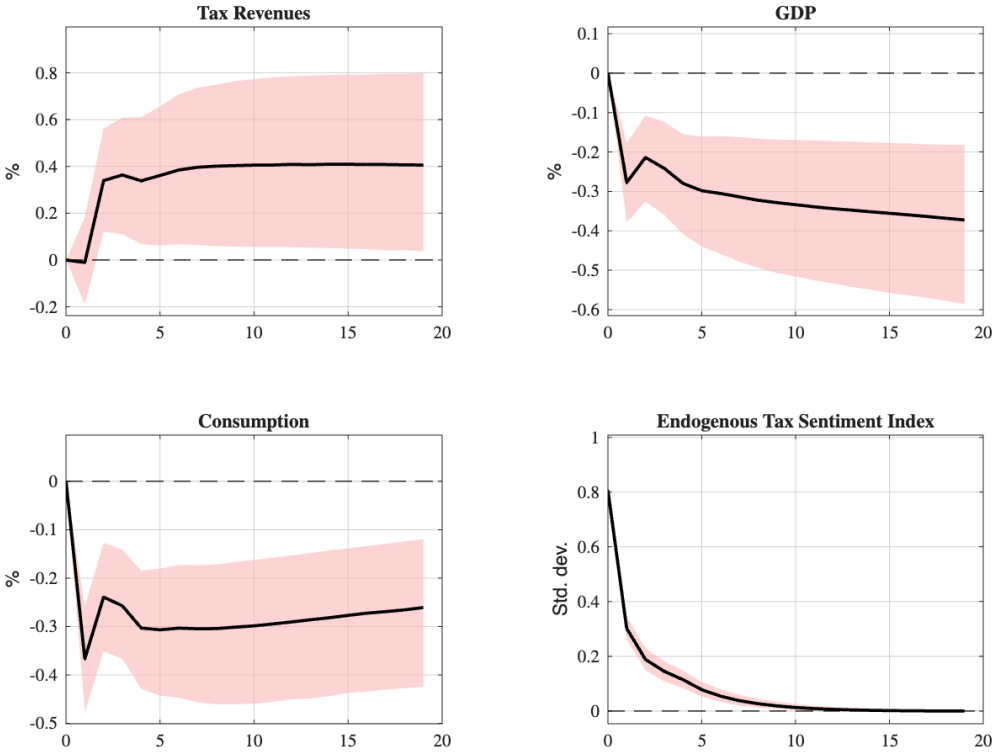

(b) Endogenous (systematic) component.

Fig. A.17. Purging decomposition for tax announcements (*TaxSent*). *Notes:* Impulse responses from the panel BVAR replacing *TaxSent* with its exogenous (ARX residual) and endogenous (fitted) components. Shaded areas denote 68% credible intervals.

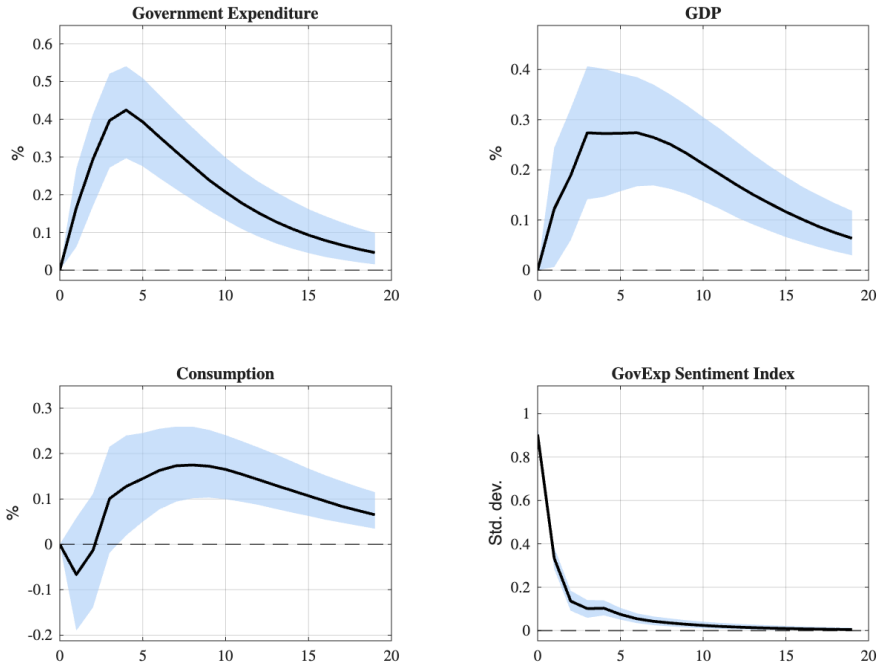

Fig. A.18. Responses of the baseline Hamilton-filtered variables to a *GovExpSent* shock in the panel VAR. *Notes:* Impulse–response functions of government expenditure, real GDP, private consumption, and the *GovExpSent* index to a one–standard–deviation increase in government expenditure sentiment, estimated from a recursively identified Bayesian panel VAR with four lags and a Normal–Wishart prior. Macroeconomic series are expressed in cyclical form using the GDP trend component from the Hamilton filter. Shaded areas denote 68% credible intervals.

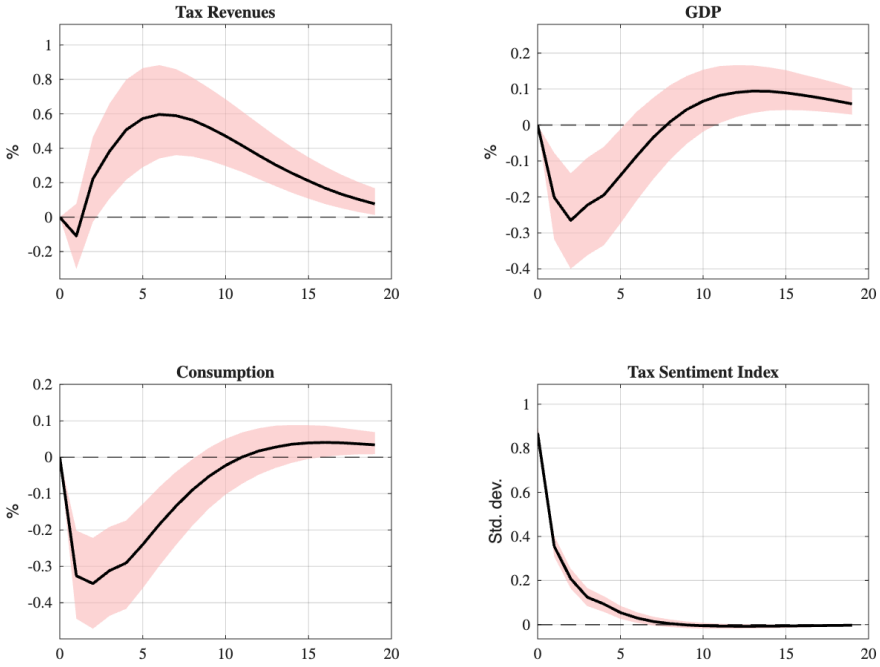

Fig. A.19. Responses of the baseline Hamilton-filtered variables to a *TaxSent* shock in the panel VAR. *Notes:* Impulse–response functions of tax revenues, real GDP, private consumption, and the *TaxSent* index to a one–standard–deviation increase in tax sentiment, estimated from a recursively identified Bayesian panel VAR with four lags and a Normal–Wishart prior. Macroeconomic series are expressed in cyclical form using the GDP trend component from the Hamilton filter. Shaded areas denote 68% credible intervals.

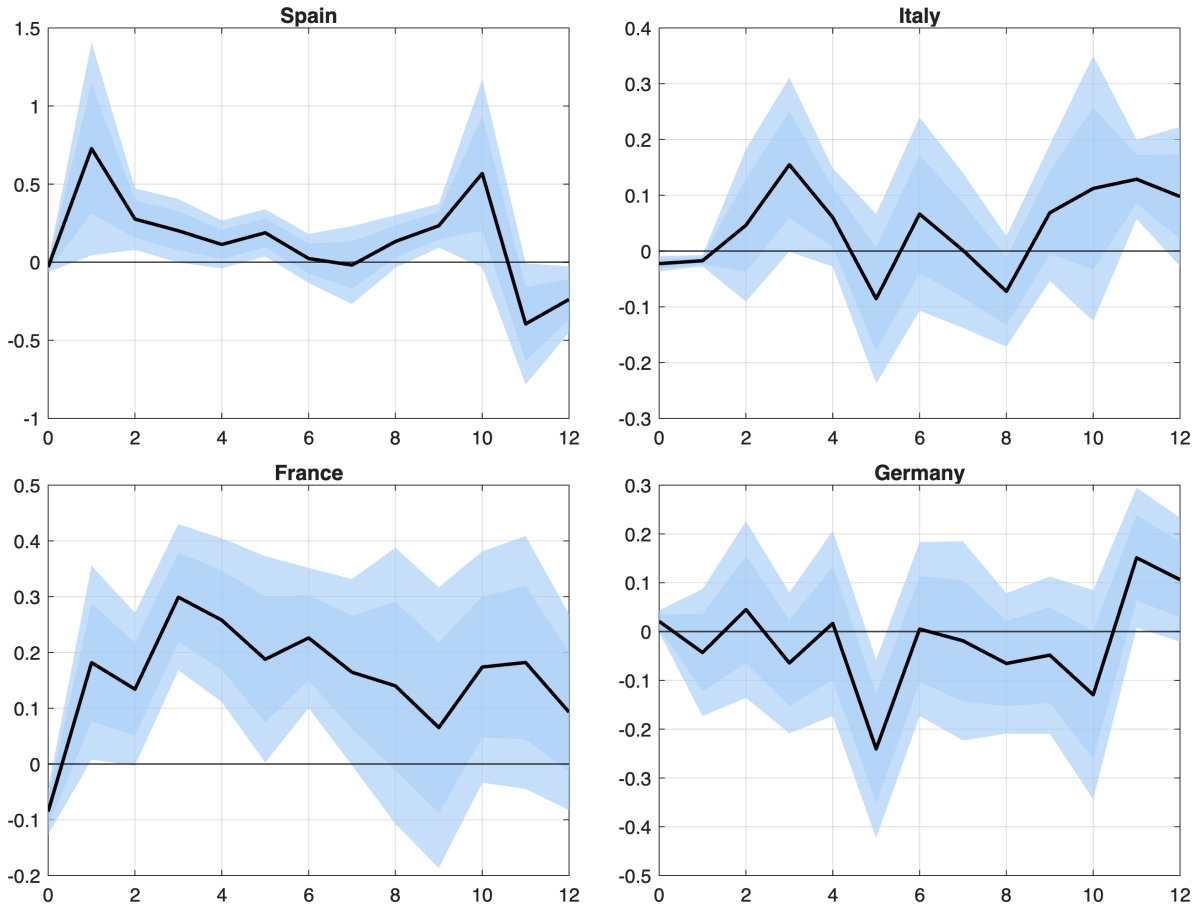

Fig. A.20. Response of the government expenditure shock to fiscal sentiment (Gordon detrending). Notes: The figure shows the estimated $\hat{\beta}_h$ from Eq. 8. Shaded areas denote 68% and 90% confidence bands based on Newey–West standard errors.

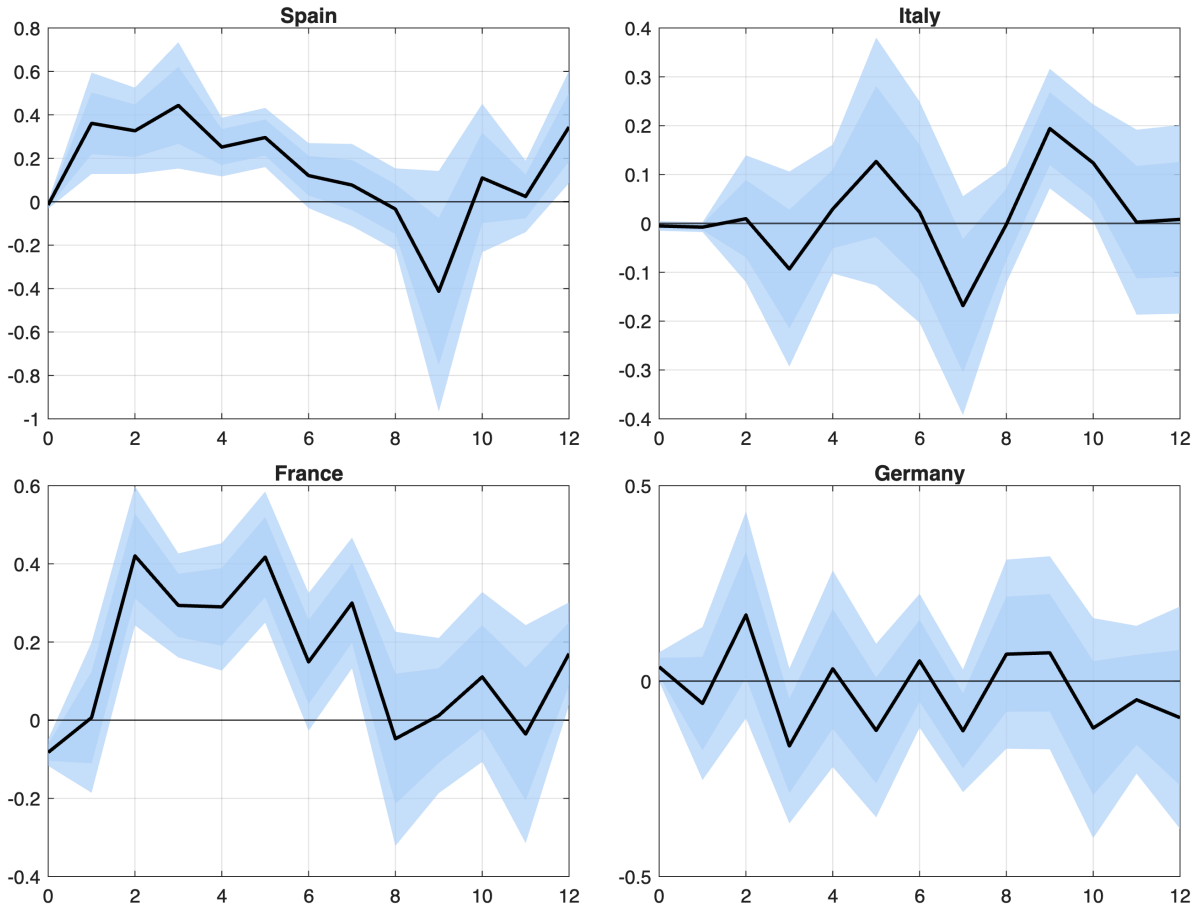

Fig. A.21. Response of the government expenditure shock to fiscal sentiment (Hamilton filtering). Notes: The figure shows the estimated $\hat{\beta}_h$ from Eq. 8. Shaded areas denote 68% and 90% confidence bands based on Newey–West standard errors.