Media Multipliers: Is There a Negativity Bias in Economic News?*

Luca Gambetti†
Universitat Autònoma de Barcelona, BSE,
Università di Torino and Collegio Carlo Alberto

Nicolò Maffei-Faccioli‡
Norges Bank

Sarah Zoi§
Federal Reserve Board

December 2025

Abstract

No. We show that the stronger responsiveness of newspaper coverage to increases in unemployment relative to decreases is fully explained by the asymmetric dynamics of unemployment and the forward-looking behavior of the media. Because unemployment rises faster and more persistently than it falls, media coverage becomes more negative as outlets anticipate and reflect this asymmetry. We develop a theoretical framework demonstrating that when unemployment is asymmetric and media are forward-looking, static regressions can generate an apparent negativity bias even if coverage is symmetric. Building on this insight, we propose a new test for negativity bias based on media multipliers, defined as the media response per cumulative change in current and future unemployment. Using nonlinear local projections and measures of coverage from three major U.S. newspapers, we find that the previously documented negativity bias disappears once dynamics and expectations are taken into account.

JEL classification: C32, E32.

Keywords: News Coverage, Negativity Bias, Business Cycles, Nonlinear Local Projections.

*This paper should not be reported as representing the views of Norges Bank or the Board of Governors of the Federal Reserve System. The views expressed are those of the authors and do not necessarily reflect those of Norges Bank or of the Federal Reserve System. We thank Francesco Bianchi and Juan F. Rubio-Ramírez for insightful comments and participants at various conferences and seminars.

†Departament d'Economia i d'Historia Economica, Edifici B, Office B3-1130, Universitat Autònoma de Barcelona, 08193 Barcelona. Phone: (+34)935814569. E-mail: luca.gambetti@uab.es. Luca Gambetti acknowledges the financial support from the Spanish Ministry of Science and Innovation, through the Severo Ochoa Programme for Centres of Excellence in R&D (CEX2019-000915-S), the financial support of the Spanish Ministry of Science, Innovation and Universities through grant PGC2018-094364-B-I00, and the Barcelona School of Economics Research Network, and the financial support from the Italian Ministry of Research and University, PRIN 2017 grant J44I20000180001.

‡Norges Bank, Bankplassen 2, P.O. Box 1179 Sentrum, 0107 Oslo. Phone: (+47)40641754. E-mail: nicolo.maffei-faccioli@norges-bank.no.

§Federal Reserve Board of Governors, 20th & Constitution Ave. NW, mailstop: K-3620, 20551 Washington DC. Phone: (+1)2024408608. E-mail: sarah.zoi@frb.gov.

1 Introduction

Media coverage of economic news shapes the information available to agents, influencing their expectations and decisions.[1] This paper revisits and challenges the prevailing view that negative economic events receive disproportionately more media attention than positive ones—a negativity bias in economic reporting. If such a bias were present, it could foster an unduly pessimistic outlook on the economy, with direct implications for public sentiment and the behavior of consumers and investors. It would also cast doubt on the reliability of widely used news-based indicators of economic conditions, which implicitly assume that media coverage tracks fundamentals.[2] Our findings offer a reassuring conclusion in this respect: media coverage largely reflects economic developments in a proportional manner, rather than being systematically biased toward bad news.

Prior research has documented a negativity bias in economic news coverage using static models that regress newspaper coverage of unemployment on contemporaneous positive and negative changes in the unemployment rate (see Goidel and Langley (1995), Fogarty (2005), Soroka (2006), Soroka (2012) and Soroka et al. (2018)). In this framework, larger coefficients on unemployment increases are typically interpreted as evidence that the media responds more strongly to bad economic news.

We show that such static models are ill-suited to test for negativity bias because they overlook two key features: the asymmetric dynamics of unemployment and the forward-looking behavior of the media. The first is well established in the literature: unemployment rises sharply during recessions and declines only gradually during expansions (see, among others, Neftci (1984), Andolfatto (1997), Morley and Piger (2012), Pizzinelli et al. (2020) and Dupraz et al. (2025)). The second is part of our contribution. Using a measure constructed from articles in three major U.S. newspapers, we show that media coverage of unemployment is forward-looking: it reflects current conditions while also anticipating future unemployment, even after controlling for a range of macroeconomic and forward-looking indicators.[3] This anticipatory behavior implies that static models fail to capture the true dynamics of economic news and its relationship with unemployment.

We formalize these ideas in a theoretical framework that incorporates both ingredients. Unemployment changes follow an asymmetric AR(1) process, with increases being more persistent than decreases, while media coverage depends on the expected sum of current and future unemployment

[1] See, for example, Carroll (2003), Blinder and Krueger (2004), Binder (2017), Chahrour et al. (2021), Nimark and Pitschner (2019), Binder et al. (2025) and Link et al. (2025).

[2] See, among others, Baker et al. (2016), Caldara et al. (2020), Caldara and Iacoviello (2022), Shapiro et al. (2022) and Bybee et al. (2024).

[3] This finding relates to evidence that media-based measures of the economy generally contain useful forward-looking information; see Larsen and Thorsrud (2019), Ellingsen et al. (2022) and Bybee et al. (2024).

changes. This expected sum varies with the sign of the current change, since future unemployment trajectories differ when unemployment is rising versus falling. Intuitively, newspapers report good and bad unemployment news based not only on current developments but also on where unemployment is expected to head. Because of the underlying asymmetry in unemployment dynamics, these trajectories differ mechanically across positive and negative changes. The framework also allows the media to assign different weights to good and bad news, thereby capturing the possibility of a genuine negativity bias. Crucially, however, even in the absence of such a bias, an increase in unemployment implies a worse expected path, generating a stronger immediate media response. As a result, static regressions of media coverage on contemporaneous unemployment changes will spuriously indicate a negativity bias, even when media coverage is entirely symmetric. In short, static models conflate asymmetries in unemployment dynamics with asymmetries in media behavior.

Building on this insight, our theoretical framework suggests a natural way to test for negativity bias: media multipliers, defined as the media response per cumulative change in current and future unemployment. Unlike static regressions—which confound unemployment persistence with media responsiveness—media multipliers explicitly account for both asymmetric unemployment dynamics and the forward-looking nature of coverage. Although conceptually related to fiscal multipliers (see Ramey and Zubairy (2018)), media multipliers are distinct in that they are constructed using realized current and future unemployment changes, thereby capturing the anticipatory behavior of the media. We estimate media multipliers separately for positive and negative unemployment shocks using nonlinear local projections. This approach allows us to quantify not only the immediate effect of unemployment shocks on coverage but also how these effects evolve over time, potentially revealing different responses to good versus bad economic news. Estimation requires two key components: a measure of media coverage of unemployment and a valid instrument that isolates unexpected variation in unemployment.

To quantify media coverage, we construct two measures of economic news from articles published in three major U.S. newspapers: bad news, defined as the monthly count of articles reporting increases or high levels of unemployment, and good news, defined as the count reporting decreases or low levels. We focus on unemployment news for two main reasons. First, unemployment is a prominent cyclical indicator and central to the economic news selection process (see Fogarty (2005)). Second, this focus aligns our analysis with the existing literature on negativity bias (see, among others, Soroka (2006) and Soroka et al. (2015)). From these counts, we build a standard indicator of unemployment coverage—the difference between bad and good news—which we refer to as negativity. This measure captures net unemployment coverage, with higher values indicating that reporting leans more negative.

Negativity is highly correlated with the unemployment rate, leads its turning points, and significantly predicts future unemployment beyond a set of macroeconomic and forward-looking indicators such as consumer sentiment and stock prices. These properties show that media coverage is not only responsive to current labor market conditions but also systematically anticipates future developments.

To estimate media multipliers, we employ a nonlinear local projection framework that flexibly captures sign-dependent responses and thus allows for asymmetries in both unemployment dynamics and media behavior. A central challenge is identifying unexpected changes in unemployment. Following Angeletos et al. (2020), we define an unemployment shock as the innovation that maximizes its contribution to the volatility of unemployment at short-run frequencies within a Threshold Vector Autoregressive (TVAR) model of negativity and unemployment rate changes. This choice mirrors the assumptions of our theoretical framework, which permits asymmetric unemployment adjustments.

We use the identified shock as an instrument for the cumulative sum of current and future unemployment changes, thereby isolating the component of unemployment dynamics that the media are likely to anticipate and respond to. With this shock and our measure of coverage, we estimate media multipliers for positive and negative unemployment shocks. These multipliers trace the media response per unit of cumulative unemployment change over future horizons, revealing how media reactions evolve and whether they differ by the sign of the shock.

Our contribution is to show that stronger coverage following unemployment increases does not reflect an inherent bias, but rather a forward-looking and proportional media response to the expected unemployment trajectory. Using standard static regressions on our data, we reproduce the familiar finding of a negativity bias in media coverage, consistent with prior studies. When we instead rely on media multipliers, this asymmetry disappears: media multipliers for positive and negative unemployment shocks are statistically indistinguishable. Normalizing net unemployment coverage by the cumulative response of unemployment changes to each shock removes the apparent bias. Increases in unemployment generate larger and more persistent effects on unemployment itself, which the media anticipate and incorporate into their reporting.

While our baseline focuses on unemployment because it is a key business cycle indicator and the benchmark used in prior studies of negativity bias, we also conduct a broader robustness check using the San Francisco Fed's News Sentiment Index, which summarizes the tone of a large set of economic news articles. In a static regression, we again observe an apparent negativity bias, with news sentiment responding more strongly to negative economic shocks. However, this asymmetry disappears once we account for asymmetric dynamics and the forward-looking behavior of the media. This indicates that our finding is not specific to unemployment news but reflects a more general feature of economic news coverage. The negativity bias that emerges in static settings primarily reflects asymmetries in the economy and the timing of news responses, not a systematic tendency of the media to emphasize negative developments.

The remainder of the paper is organized as follows. Section 2 describes our news measures, their relation to the unemployment rate, and comparisons with the Michigan Survey of Consumers and other coverage measures in the literature. Section 3 introduces a simple and intuitive model of media coverage of unemployment. Sections 4 presents the main findings, and Section 5 reports a battery of robustness checks. Section 6 concludes.

2 The news indexes

This section describes the construction of our news indicators, examines their time-series properties, and presents evidence on the forward-looking nature of media coverage of unemployment.

2.1 Indexes construction

We construct two measures of newspaper coverage: one for bad news about unemployment and one for good news. These measures are based on articles from the Dow Jones Factiva database, a comprehensive archive of news. Our sample covers reporting on the U.S. economy between June 1980 and December 2019 in three major newspapers: *The New York Times*, *The Wall Street Journal*, and *The Washington Post*. We focus on these outlets because they have consistently ranked among the largest US newspapers by circulation over our sample period and all target national audiences.

We construct our indicators using an occurrence-based method, similar to those employed in the literature to measure economic policy uncertainty (Baker et al. (2016)), trade uncertainty (Caldara et al. (2020)) and geopolitical risk (Caldara and Iacoviello (2022)). We define our measure of bad news as the monthly number of articles in which the word "unemployment" appears near terms denoting an increase or high level. It is thus a measure of monthly negative coverage for unemployment. Good news is defined analogously, as the monthly number of articles in which unemployment appears near terms indicating a decrease or low level. This measure represents monthly positive coverage of unemployment.

To define the semantic groups for negative and positive articles about unemployment, we conduct a human pre-reading of 150 articles. We randomly select 15 dates and, for each date, sample 10 articles containing the word "unemployment". This procedure ensures a representative sample across our study period and allows us to construct robust semantic categories for classifying unemployment-related news. Through this process, we identify two groups of words – one associated with actual or expected increases or high levels of unemployment, and another for decreases or low levels – and determine the appropriate word distance from "unemployment". Using these groups, we develop two search queries that categorize articles as negative or positive in their coverage of unemployment. The full search queries are provided in Appendix A.

Each month, we count the number of articles meeting the two search criteria. We then refine the measures by excluding articles selected by both queries or those containing words of negation ("no" or "not") in the immediate vicinity of "unemployment". This step removes about 6% of articles that cannot be clearly assigned to either category. We acknowledge that these articles may convey information about periods of relatively stable unemployment or reflect mixed signals about the labor market. However, for the purpose of this study, which is concerned with potentially asymmetric coverage, it is of primary importance to have a clear measure of news polarization.[4] he final dataset contains 35933 bad-news articles and 22317 good-news articles on unemployment, implying substantially greater coverage of negative developments.

Using the two raw indexes, we construct our primary indicator of media coverage on unemployment for our empirical analysis: *negativity*, defined as the difference between the counts of bad and good news. This measure captures the net volume of media reporting on unemployment: positive values indicate a predominance of negative coverage, while negative values indicate more positive coverage. By comparing the frequency of negative and positive news, negativity serves as a clear and direct metric to quantify potential asymmetries in media coverage of unemployment.

A few remarks on our indicators are in order. First, it is important to clarify that our variables are not sentiment indicators.[5] Instead, they quantify the frequency of negative and positive newspaper articles related to unemployment. This choice ensures comparability with prior studies, which have traditionally relied on article counts to construct measures of negative and positive coverage of unemployment (see Soroka (2006)). Relative to the measures used in earlier work, our indicators draw on multiple newspapers and cover a longer time period. These two characteristics make our analysis more robust to possible newspaper-specific and short-sample biases.

Second, we do not distinguish between articles in which unemployment is the main topic (e.g., appears in the headline or is mentioned repeatedly) and those in which it is mentioned incidentally (i.e. the focus of the article is not necessarily unemployment, but the word is mentioned because it is somehow related to the main topic). This choice reflects the fact that unemployment is often

[4]The results presented below are robust to the inclusion of these "ambiguous" news, as shown in the robustness section.

[5]We believe it would be interesting to try to construct such an indicator and study potential differences with ours. We plan to do this in the future.

reported in connection with broader discussions of the economy, or politics. Focusing exclusively on articles mentioning unemployment in the headlines would exclude a large amount of information about this indicator which can be provided in the body of the article.

Third, a potential concern is that the overall volume of news coverage may have expanded over time. In principle, fluctuations in our indicators could therefore reflect changes in the total number of articles published, rather than genuine shifts in the balance between bad and good news. To address this, we normalize our indicators by dividing them by the total number of articles published each month across the three newspapers, including those unrelated to unemployment. The resulting normalized indicator thus represents the excess of negative over positive unemployment articles as a share of total newspaper content. Figure D.1 in the Appendix plots both the baseline and normalized indicators, which track each other closely over the sample period, with a correlation of 0.96. In the robustness section, we replicate our analysis using the normalized indicators and obtain results that are very similar to the baseline.

2.2 Descriptives

The first two upper panels of Figure 1 report our bad and good news indicators (blue lines) alongside the unemployment rate (red lines).[6] Bad news articles average 68 per month with a standard deviation of 47, while good news articles average 43 with a standard deviation of 18. Negative news about unemployment are, on average, more frequent and volatile than positive news. The most notable difference between the two indexes, however, lies in their correlation with the unemployment rate: 0.80 for bad news versus -0.24 for good news. This is also evident from visual inspection of the time series. Bad news exhibits a substantially more cyclical pattern than good news. The bad news index closely tracks the unemployment rate, featuring two major spikes corresponding to the early 1980s recession and the Great Recession. In contrast, the good news index appears largely unrelated with the unemployment rate, with notable exceptions during three periods: the late 1980s, the late 1990s, and post-2015.

A potential concern related to our news indexes is that the newspapers we consider may cover unemployment developments differently, depending on their political view. Figure D.2 in the Appendix reports our news indexes disaggregated by newspaper. Overall, the coverage of both bad and good unemployment developments is remarkably consistent across different newspapers. Indeed, all of the indexes track each other very well over the sample period. This rules out the existence of a relevant political bias in the unemployment news reporting for the newspapers considered.

[6] Data on the unemployment rate is retrieved from FRED (mnemonic: UNRATE).

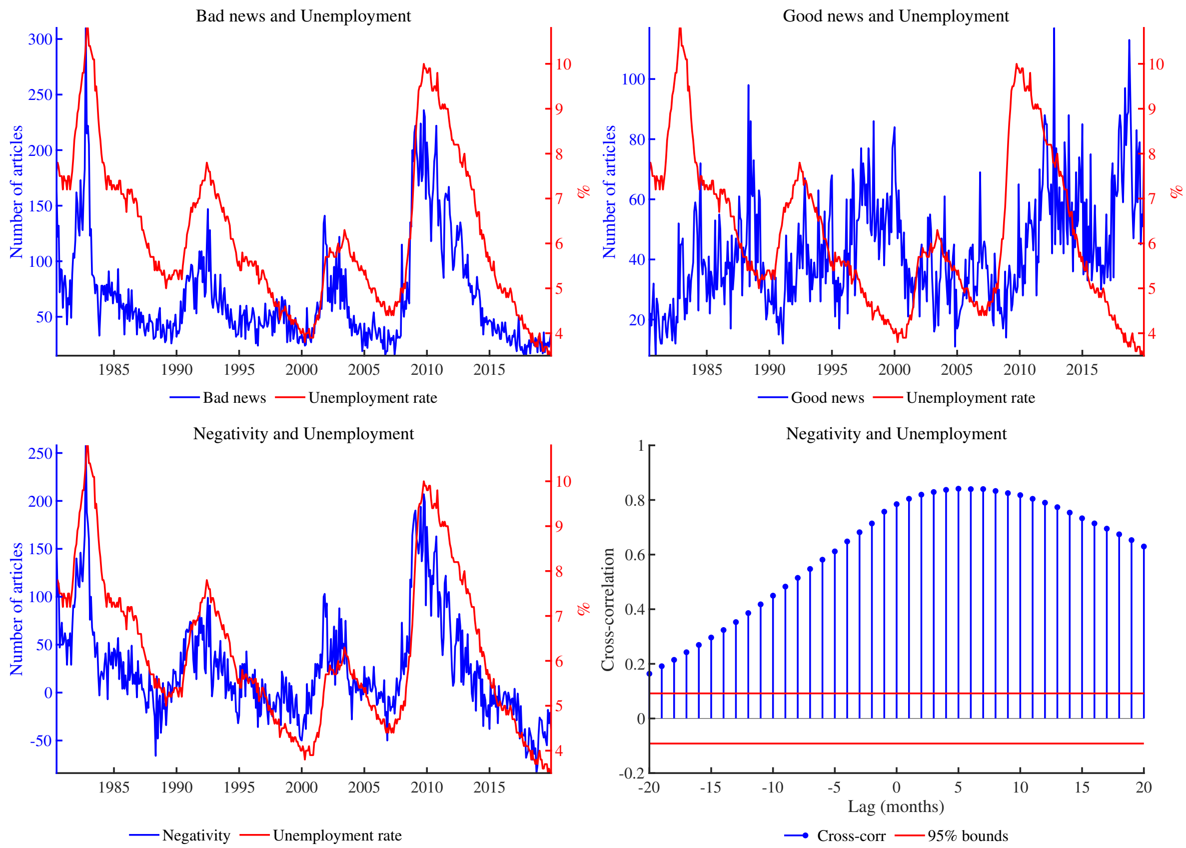

Figure 1: Bad news, good news, negativity and unemployment.

We report our measure of the net volume of unemployment coverage in the bottom-left panel of Figure 1. Negativity has a correlation of 0.78 with the unemployment rate and a historical mean of 25, statistically different from zero. A notable feature of negativity is that it leads the unemployment rate, anticipating changes in this variable. The bottom-right panel of 1 shows lag-lead cross-correlations between these two variables. The vertical axis reports the correlation between the unemployment rate and the lags (positive values) and leads (negative values) of negativity. The peak correlation occurs at five lags, suggesting that the articles we consider may have informational content about future developments in the unemployment rate.

To formally test the predictive content of negativity for unemployment, we conduct Granger causality tests. Our null hypothesis is that the coefficients on the lags of negativity are jointly equal to zero. We consider two models: a baseline including lags of unemployment, annual PCE inflation, the growth rate of industrial production, the S&P500 index and consumer sentiment from the Michigan Survey, and an extended model including all the previous variables and the San Francisco FED News Index. We perform the tests for both models using 4 and 6 lags of all regressors, resulting in four specifications. Table 1 reports the results. The null hypothesis

	F-stat	p-value	Lags	n. obs
Baseline Model	9.46	0.00	4	459
	6.28	0.00	6	457
Extended Model	8.99	0.00	4	459
	6.24	0.00	6	457

Notes: Granger causality test for the null hypothesis that the first four and six lags of negativity don't predict current unemployment. The baseline model includes lags of unemployment, annual PCE inflation, the growth rate of IP, the stock market's return and the Consumer Sentiment Index. The extended model includes all the variables in Model 1 and the News Sentiment Index of the San Francisco FED.

Table 1: Granger causality test

is rejected for all specifications, suggesting that our measure of net coverage, negativity, contains information about future realizations of unemployment that is not captured by standard economic or forward-looking variables. This observation supports the view that media are not merely reacting to unemployment figures but are forward-looking in their news reporting process.

2.3 Comparisons with other measures of information

The relatively weak procyclicality of the good news indicator is noteworthy, as one might expect its pattern to closely mirror that of bad news with opposite sign. To further investigate this asymmetry and validate our measures, we compare our news indexes with alternative indicators of economic news perceptions. In particular, we use data from the *Michigan Survey of Consumers*, which provides insights into how consumers perceive and process economic information. This comparison allows us to assess the consistency of our measures with the information reported by survey respondents, potentially shedding light on the observed behavior of our good news indicator.

The survey provides a wide variety of variables that reflect agents' information and expectations about the current and future state of the economy. The variable NEWS in the survey corresponds to the percentage of individuals who recently heard of any favorable or unfavorable changes in business conditions. Question A6 of the questionnaire asks the following: *"During the last few months, have you heard of any favorable or unfavorable changes in business conditions?"*. There are two possible answers: *"Yes"* and *"No, haven't heard"*. If the individual answers *"Yes"*, then the second question is A6a: *"What did you hear?"*, which is an open-ended question. The Michigan Survey provides few variables constructed on the basis of the type of answer to these two questions. Among those, we focus on the following variables: "Favorable: employment" and "Unfavorable: unemployment", corresponding to answers to question A6a which are specifically related to positive evaluations of employment and negative evaluations of unemployment.

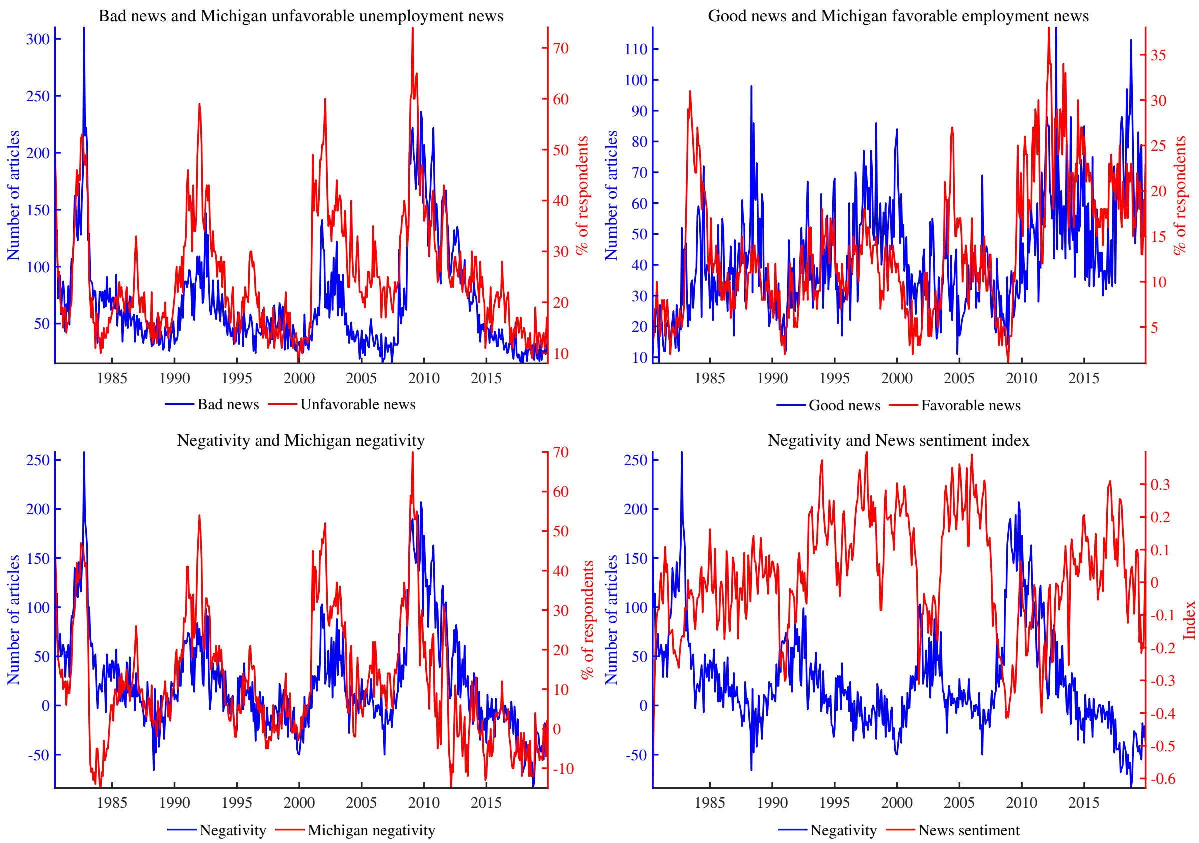

Figure 2: Negativity, Michigan negativity and the news sentiment index from the San Francisco FED

While our indicators of bad and good news represent *objective* measures of the amount of negative and positive published news items related to unemployment figures, the corresponding two variables from the Michigan Survey represent the *subjective* information that the agents perceive from the media or alternative sources of information. In principle, agents' subjective information may not coincide with our measures of objective information. For example, agents may get informed mainly through other channels (TV, social networks, etc.) or they may be rational inattentive even in information-rich environments (see Sims (2003), Nimark and Sundaresan (2019)).

The first row of Figure 2 plots our bad and good news indexes together with the corresponding measures in the *Michigan Survey of Consumers*, namely the "Unfavorable: unemployment" and "Favorable: employment" items of NEWS. Despite their conceptual differences, both indexes track the corresponding variables of the Michigan Survey extremely closely over the sample considered. The correlation between "Unfavorable: unemployment" and bad news is 0.70, and the correlation between "Favorable: employment" and good news is 0.49. The bottom-left panel of Figure 2 reports negativity together with its counterpart constructed using the variables of the Michigan survey. Their correlation is 0.65. Overall, our indexes are consistent with the survey measures, suggesting that newspaper information is a relevant channel for agents' information, and that our indexes capture information relevant to agents' economic outlook.

Finally, we further validate our negativity measure by comparing it, in the bottom-right panel of Figure 2, to the News Sentiment Index developed by the Federal Reserve Bank of San Francisco, which reflects the tone of macroeconomic news more broadly. While our measure is not a sentiment indicator – rather, it captures the net volume of negative versus positive unemployment-related coverage – we find a strong negative correlation of -0.61 between the two series. This is notable given their conceptual differences: the San Francisco Fed index aggregates sentiment across a broad set of economic topics and closely aligns with consumer sentiment, whereas our measure is narrowly focused on unemployment. The similarity between the two suggests that the intensity of unemployment coverage moves closely with broader shifts in economic sentiment. This finding not only lends credibility to our measure, but also highlights the central role of unemployment news in shaping the tone of overall macroeconomic reporting, consistent with Fogarty (2005)'s finding that unemployment is a primary focus of economic news coverage.

3 A simple model of unemployment news coverage

In this section, we develop a stylized theoretical framework to formalize the concept of negativity bias in economic news coverage. This setup illustrates the limitations of the static methods commonly used in the literature, and motivates our empirical approach based on media multipliers. Our framework includes three key features: the asymmetric dynamics of unemployment, the forward-looking nature of media coverage, and the possibility of (negativity) bias in media reporting.

We begin by modeling changes in unemployment with a sign-dependent threshold AR(1) process:

$$\Delta U_t = \phi_1\Big(I^+_{t-1} \times \Delta U_{t-1}\Big) + \phi_2\Big(I^-_{t-1} \times \Delta U_{t-1}\Big) + \varepsilon_t \tag{1}$$

where $0 \leq \phi_i < 1$ $(i = 1, 2)$, $I^+_{t-1} = 1$ if $\Delta U_{t-1} > 0$ and zero otherwise, $I^-_{t-1} = 1$ if $\Delta U_{t-1} < 0$ and zero otherwise, and $\varepsilon_t \sim iid(0, \sigma^2_\varepsilon)$ denotes the unemployment shock, namely the innovation to changes in unemployment orthogonal to past information. The coefficients ϕ_1 and ϕ_2 capture the persistence of unemployment changes in a sign-dependent manner: when the previous change ΔU_{t-1} is positive, persistence is governed by ϕ_1, whereas when it is negative, persistence is governed by ϕ_2. This specification provides a simple way to account for the asymmetric dynamics of unemployment documented in prior studies (see, among others, Neftci (1984), Andolfatto (1997), Morley and Piger (2012), Pizzinelli et al. (2020) and Dupraz et al. (2025)), which feature a steep rise–slow decline pattern. In our framework, this asymmetry corresponds to the case $\phi_1 > \phi_2$. When $\Delta U_{t-1} > 0$, a relatively large value of ϕ_1 implies that positive changes in unemployment are strongly propagated forward: an increase in unemployment today raises the likelihood of further increases tomorrow, producing the steep and rapid surge observed in recessions. By contrast, when $\Delta U_{t-1} < 0$, a smaller value of ϕ_2 implies that negative changes are less persistent: decreases in unemployment today have a weaker effect on tomorrow's changes, so the recovery unfolds only gradually.

At each time t, the media observe a positive or negative change in the unemployment rate and decide how many good and bad stories to publish. In what follows, coverage M_t is defined as the difference between bad and good news articles, though the same reasoning applies to any measure of unemployment coverage. Consistent with prior literature (Soroka, 2006; Soroka et al., 2015), we assume that changes in economic variables – rather than their levels – are the primary drivers of media reporting.[7] To capture the forward-looking behavior of the media, we postulate that this decision is not based solely on current unemployment, but also on the expectations about its future trajectory. In particular, when setting the number of stories, the media consider the current change and the expected cumulative change over h horizons ahead:

$$M_t = m_0^+\left(I_t^+ \times E\left[\sum_{j=0}^{h} \Delta U_{t+j} \middle| \Delta U_t > 0\right]\right) + m_0^-\left(I_t^- \times E\left[\sum_{j=0}^{h} \Delta U_{t+j} \middle| \Delta U_t < 0\right]\right) + v_t. \tag{2}$$

where $v_t \sim iid(0, \sigma_v^2)$ is an error term capturing idiosyncratic variation in media coverage and I_t^+ and I_t^- are indicator functions which take value 1 if the unemployment change at t is positive or negative, respectively. The parameters m_0^+ and m_0^- are the impact media multipliers associated with positive and negative unemployment changes, respectively. For a given size and persistence of an unemployment increase, a higher value of m_0^+ means that the media publish more bad stories. Negativity bias in economic news coverage corresponds to $m_0^+ > m_0^-$: increases in unemployment generate more coverage than decreases of comparable magnitude and persistence. When evaluated at horizon $h = 0$, equation (2) describes the same media response function underlying the static regression models employed in the literature (see Soroka (2006)). The key extension is that media coverage also incorporates prospective changes in the unemployment rate for $h > 0$, reflecting a forward-looking component absent in static models.

For simplicity, we assume in what follows that media have full-information rational expectations: they know the process for unemployment in (1) and its parameters, and make their expectations accordingly. This assumption is not essential for our argument. What matters is that expectations reflect the asymmetric persistence of unemployment, namely that positive and negative shocks

[7]Nonetheless, we show in the robustness section that our empirical results are the same when using unemployment in levels instead of changes.

differ in their perceived persistence. Any expectation-formation rule that captures this asymmetry would lead to the same qualitative implications. To set M_t, the media have to forecast future changes in unemployment, conditioning on the current change being positive or negative. Computing these forecasts analytically becomes extremely cumbersome as the horizon h increases, as it requires to forecast the state variable (i.e. to compute the probability of each state) for each horizon. For clarity of exposition, we make a simplifying assumption and treat each state as absorbing over the entire forecasting horizon. This allows us to deliver the main insights in a intuitive way, while we defer to Appendix B for the solution of the full model without this assumption using numerical methods, which corroborates the conclusions obtained from the simplified framework. Under this assumption, we can express the media function (2) as:

$$M_t = m_0^+ \left(I_t^+ \times E\left[\sum_{j=0}^{h} \Delta U_{t+j} \middle| \{\Delta U_{t+j} > 0\}_{j=0}^{h} \right] \right) + m_0^- \left(I_t^- \times E\left[\sum_{j=0}^{h} \Delta U_{t+j} \middle| \{\Delta U_{t+j} < 0\}_{j=0}^{h} \right] \right) + v_t. \tag{3}$$

where $v_t \sim iid(0, \sigma_v^2)$ and orthogonal to ΔU_{t+j} for any j.[8]

The j-step ahead conditional forecast depends on the sign of the current change:

$$E\left(\Delta U_{t+j} \middle| \{\Delta U_{t+j} > 0\}_{j=0}^{h} \right) = \phi_1^j \Delta U_t \tag{4}$$

$$E\left(\Delta U_{t+j} \middle| \{\Delta U_{t+j} < 0\}_{j=0}^{h} \right) = \phi_2^j \Delta U_t \tag{5}$$

and the expected cumulative sum of current and future unemployment changes up to horizon h is:

$$E\left(\sum_{j=0}^{h} \Delta U_{t+j} \middle| \{\Delta U_{t+j} > 0\}_{j=0}^{h} \right) = \sum_{j=0}^{h} \phi_1^j \Delta U_t \tag{6}$$

$$E\left(\sum_{j=0}^{h} \Delta U_{t+j} \middle| \{\Delta U_{t+j} < 0\}_{j=0}^{h} \right) = \sum_{j=0}^{h} \phi_2^j \Delta U_t \tag{7}$$

[8]This equation is equivalent to:

$$M_t = m_0^+ \left(I_t^+ \times E\left[\sum_{j=0}^{h} \Delta U_{t+j} \middle| \varepsilon_t > 0, \varepsilon_{t+1} = \cdots = \varepsilon_{t+h} = 0 \right] \right) + m_0^- \left(I_t^- \times E\left[\sum_{j=0}^{h} \Delta U_{t+j} \middle| \varepsilon_{t+1} = \cdots = \varepsilon_{t+h} = 0 \right] \right) + v_t.$$

Substituting the two forecasts in the media function (3) we obtain:

$$M_t = m_0^+ \sum_{j=0}^{h} \phi_1^j \left(I_t^+ \times \Delta U_t \right) + m_0^- \sum_{j=0}^{h} \phi_2^j \left(I_t^- \times \Delta U_t \right) + v_t. \tag{8}$$

3.1 The static regression

Suppose that the econometrician estimates the static regression:

$$M_t = \beta_1 \left(I_t^+ \times \Delta U_t \right) + \beta_2 \left(I_t^- \times \Delta U_t \right) + v_t, \tag{9}$$

where $\beta_1 = m_0^+ \sum_{j=0}^{h} \phi_1^j$ and $\beta_2 = m_0^- \sum_{j=0}^{h} \phi_2^j$.

If unemployment dynamics are the same in the two regimes, $\phi_1 = \phi_2$, estimating $\hat{\beta}_1 > \hat{\beta}_2$ implies $m_0^+ > m_0^-$ – a negativity bias in unemployment coverage. In this case, the static regression estimated by the literature would be sufficient to asses the presence of the bias. However, if the dynamics of unemployment changes differ between regimes, the econometrician cannot identify m_0^+ and m_0^- from a static specification. For example, when $\phi_1 > \phi_2$, an estimate of $\hat{\beta}_1 > \hat{\beta}_2$ becomes inconclusive: it may arise either from genuine negativity bias ($m_0^+ > m_0^-$) or simply from asymmetric persistence in unemployment dynamics ($\phi_1 > \phi_2$ with $m_0^+ = m_0^-$). In other words, an apparent structural bias captured through a static regression could merely be an artifact of underlying differences in the persistence of unemployment changes. Empirical estimates of the process in equation (1) indeed reveal a substantially larger autoregressive coefficient for positive changes compared to negative, underscoring the importance of accounting for asymmetries in unemployment to avoid spurious conclusions about media behavior.[9]

In sum, when $\phi_1 > \phi_2$, a larger estimated coefficient on positive unemployment changes in the static regression does not necessarily imply negativity bias ($m_0^+ > m_0^-$). The econometrician cannot identify the true media multipliers because the media themselves are forward-looking: they anticipate that increases in unemployment tend to persist more than decreases and accordingly adjust their coverage. As a result, media reporting will naturally mirror the underlying asymmetry in unemployment dynamics, even in the absence of any intrinsic bias toward negative news.

3.2 Media multipliers

Given the limitations of the static regression, we propose an alternative strategy to estimate m_0^+ and m_0^- which arises naturally from our framework.

[9] In our sample, $\hat{\phi}_1 = 0.26$ and $\hat{\phi}_2 = 0.02$.

When the change in unemployment is positive, the expectation over its future trajectory can be written as:

$$E\left(\sum_{j=1}^{h}\Delta U_{t+j}\,\middle|\{\Delta U_{t+j}>0\}_{j=0}^{h}\right)=\sum_{j=1}^{h}\Delta U_{t+j}-\sum_{j=1}^{h}e^{+}_{j,t+j} \tag{10}$$

where $e^{+}_{j,t+j}$ is the j-th step ahead forecast error conditional on $\{\Delta U_{t+j}>0\}_{j=0}^{h}$.

When the change is negative, the expectation is:

$$E\left(\sum_{j=1}^{h}\Delta U_{t+j}\,\middle|\{\Delta U_{t+j}<0\}_{j=0}^{h}\right)=\sum_{j=1}^{h}\Delta U_{t+j}-\sum_{j=1}^{h}e^{-}_{j,t+j} \tag{11}$$

where $e^{-}_{j,t+j}$ is the j-th step ahead forecast error conditional on $\{\Delta U_{t+j}<0\}_{j=0}^{h}$.[10]

3.2.1 Impact media multipliers

We replace the two expectations in (2) to obtain:

$$M_t = m_0^{+}\left(I_t^{+}\times\sum_{j=0}^{h}\Delta U_{t+j}\right)+m_0^{-}\left(I_t^{-}\times\sum_{j=0}^{h}\Delta U_{t+j}\right)+\xi_t \tag{12}$$

where $\xi_t = v_t - I_t^{+}m_0^{+}\sum_{j=1}^{h}e^{+}_{j,t+j} - I_t^{-}m_0^{-}\sum_{j=1}^{h}e^{-}_{j,t+j}$.

This equation cannot be estimated using OLS since the regressors are correlated with the forecast errors, but it can be estimated using an instrumental variable (IV) approach. A valid instrument z_t has to be correlated with $\sum_{j=1}^{h}\Delta U_{t+j}$ and uncorrelated with the forecast errors $e^{-}_{j,t+j}$ and $e^{+}_{j,t+j}$ for every j. The unemployment shock ε_t in equation (1) satisfies both conditions: by construction, it affects future unemployment and, by the *iid* assumption, is uncorrelated with future forecast errors. Such shock captures variation in unemployment changes that media are likely to anticipate and respond to.

Notably, equation (12) nests the static regression model in (9) for $h = 0$. Since the relevant horizon h for media responses is not known a priori, we treat it as an object to be determined empirically: by estimating the multipliers across different horizons, we can directly evaluate the sensitivity of the negativity bias to the assumed temporal structure. We call m_0^{+} and m_0^{-} the *impact media multipliers*, as they capture the contemporaneous response of the media to positive and negative changes in unemployment.

[10] For instance, $e^{+}_{1,t+1}=\varepsilon_{t+1}$, $e^{+}_{2,t+2}=\phi_1\varepsilon_{t+1}+\varepsilon_{t+2}$, $e^{-}_{1,t+1}=\varepsilon_{t+1}$ and $e^{+}_{2,t+2}=\phi_2\varepsilon_{t+1}+\varepsilon_{t+2}$.

3.2.2 Dynamic media multipliers

While estimating $m_0^+ = m_0^-$ would suggest the absence of contemporaneous bias in media reporting, it would not necessarily preclude the possibility of dynamic bias. In fact, the media might exhibit a delayed asymmetric response, publishing a greater volume of news in subsequent periods following a positive compared to a negative change in unemployment.

To formalize this concept, we take the cumulative sum of both sides of equation (3):

$$\sum_{i=0}^{k} M_{t+i} = \sum_{i=0}^{k} m_i^+ \sum_{j=0}^{h} \phi_1^j \left(I_t^+ \times \Delta U_t \right) + \sum_{i=0}^{k} m_i^- \sum_{j=0}^{h} \phi_2^j \left(I_t^- \times \Delta U_t \right) + v_t \tag{13}$$

A *dynamic negativity bias* implies that $m_i^+ > m_i^-$ for any i and for a given forecasting horizon h.

While this type of bias has been overlooked in previous studies, we believe its examination is important for a comprehensive analysis. Dynamic media multipliers allow us to trace the full temporal response of media reporting to unemployment shocks, and to relate it to the dynamics of unemployment itself – similar to how dynamic fiscal multipliers are used to characterize the propagation of fiscal policy over time (see Ben Zeev et al. (2023) and Ramey and Zubairy (2018)).

To investigate this potential asymmetry, we first denote with m_k^+ and m_k^- the cumulated sums $\sum_{i=0}^{k} m_i^+$ and $\sum_{i=0}^{k} m_i^-$ associated with a forecasting horizon of h periods ahead for unemployment. We then test for the presence of a *dynamic bias* by estimating the cumulative media multipliers $m_{k,h}^+$ and $m_{k,h}^-$ for different horizons k. As for impact multipliers, we can instrument $\sum_{j=0}^{h} \Delta U_{t+j}$ with the unemployment shock ϵ_t. A negativity bias implies that, as k increases, the difference $m_k^+ - m_k^-$ gets larger.

In the next section we explain in more detail how to estimate the multipliers and we present our empirical results.

4 Empirical estimates of media multipliers

We begin by replicating the results of the standard static regression used in the literature to document the negativity bias. The purpose of this exercise is to ensure that our findings are consistent with the existing literature despite the use of different indicators of unemployment news coverage and a different time span. We specify the same regression estimated by Soroka (2006):

$$M_t = \beta_1 \left(I_t^+ \times \Delta U_t \right) + \beta_2 \left(I_t^- \times \Delta U_t \right) + \gamma \mathbf{x}_{t-1} + v_t, \tag{14}$$

Variable	Estimate	*t-stat*
$I_t^+ \times \Delta U_t$	**27.41**	**2.22**
$I_t^- \times \Delta U_t$	2.53	0.21
M_{t-1}	**0.44**	**9.16**
M_{t-2}	**0.33**	**6.48**
M_{t-3}	**0.20**	**3.88**
M_{t-4}	-0.04	-0.92
Constant	0.33	0.21
Observations	468	
R-squared	0.807	
F-statistic	322	

Table 2: Regression of negativity (M_t) on its first four lags and contemporaneous positive/negative changes in unemployment (ΔU_t)

where M_t is our measure of media coverage, negativity, and $\mathbf{x}_{t-1}$ is a vector of controls including a constant and four lags of the dependent variable. I_t^+ and I_t^- are the indicators defined above for positive and negative changes in unemployment. Our sample spans the period 1980:06-2019:12.

The results of the regression for negativity are displayed in Table 2. The estimated coefficient associated with an increase in unemployment is larger than the one associated with a reduction, namely $\hat{\beta}_1 > \hat{\beta}_2$, and only the former is significant. That is, increases in unemployment receive a stronger media response than decreases of the same magnitude.

If our analysis stopped here, we would conclude that economic news reporting exhibits a negativity bias, consistent with previous findings in the literature. However, as we argued in the previous section, this regression is ill-suited to identify asymmetries in economic news reporting. The observed asymmetry in media coverage ($\hat{\beta}_1 > \hat{\beta}_2$) might simply reflect unaccounted-for asymmetries in unemployment dynamics ($\phi_1 > \phi_2$), rather than a systematic bias in media behavior ($m_0^+ > m_0^-$).

In the next subsection, we address this concern by directly estimating the media multipliers, which explicitly control for asymmetries in unemployment dynamics and the forward-looking behavior of the media.

4.1 Impact media multipliers

To assess whether the contemporaneous reaction of the media to an increase in unemployment is larger than to a decrease, we estimate the following model for $h = 0, ..., 12$:

$$M_t = m^+_{0,h}\Big(S^+_t \times \sum_{j=0}^{h} \Delta U_{t+j}\Big) + m^-_{0,h}\Big(S^-_t \times \sum_{j=0}^{h} \Delta U_{t+j}\Big) + \gamma^+_h(S^+_t \times \mathbf{x}_{t-1}) + \gamma^-_h(S^-_t \times \mathbf{x}_{t-1}) + u_t \quad (15)$$

where M_t is negativity and $\sum_{j=0}^{h} \Delta U_{t+j}$ is the cumulative change in the unemployment rate over the forecast horizon h. S^+_t and S^-_t are indicator variables equal to 1 when the unemployment shock is strictly positive or negative, and zero otherwise. These indicators capture the sign of unexpected changes in the unemployment rate. $m^+_{0,h}$ and $m^-_{0,h}$ are impact media multipliers for positive and negative unemployment shocks, respectively. These capture the contemporaneous response of negativity per one-percentage point cumulative change in unemployment over h months, allowing for differences in the media's response to shocks of comparable magnitude and persistence.

We instrument $S^+_t \times \sum_{j=0}^{h} \Delta U_{t+j}$ with a positive unemployment shock, $\varepsilon^+_t = S^+_t \times \varepsilon_t$, while $S^-_t \times \sum_{j=0}^{h} \Delta U_{t+j}$ with a negative shock, $\varepsilon^-_t = S^-_t \times \varepsilon_t$.

To obtain the unemployment shock, we estimate a threshold VAR (TVAR) with changes in unemployment and negativity. The threshold is defined by whether the change in unemployment at $t-1$ is strictly positive or negative.[11] The model includes six lags of each variable as suggested by an average of the AIC and BIC criteria. This specification is arguably the simplest and most parsimonious to capture the asymmetric behavior of unemployment highlighted in the theoretical framework of section 3. In the robustness section, we show that our main findings do not depend on this particular choice: including additional variables, altering the lag structure or using alternative identification strategies yields the same conclusions.

We identify an unemployment shock as the innovation that maximizes the contribution to the volatility of unemployment over short-run frequencies (see Angeletos et al. (2020)), defining the frequency domain between 12 and 36 months. We use the resulting estimated positive and negative shocks, $\hat{\varepsilon}^+_t$ and $\hat{\varepsilon}^-_t$, as instruments for $S^+_t \times \sum_{j=0}^{h} \Delta U_{t+j}$ and $S^-_t \times \sum_{j=0}^{h} \Delta U_{t+j}$. The vector of controls $\mathbf{x}_{t-1}$ includes a constant, six lags of the dependent variable and the unemployment rate change. We include this vector in both the first- and second-stage regressions. To allow the coefficients to vary with the sign of the shock, we interact these controls with the indicators S^+_t and S^-_t. This added flexibility is consistent with the assumptions used to identify the shock in the TVAR and is necessary for correctly computing state-dependent multipliers, as emphasized by Ramey and

[11] Additional details on the model are laid out in Appendix C.

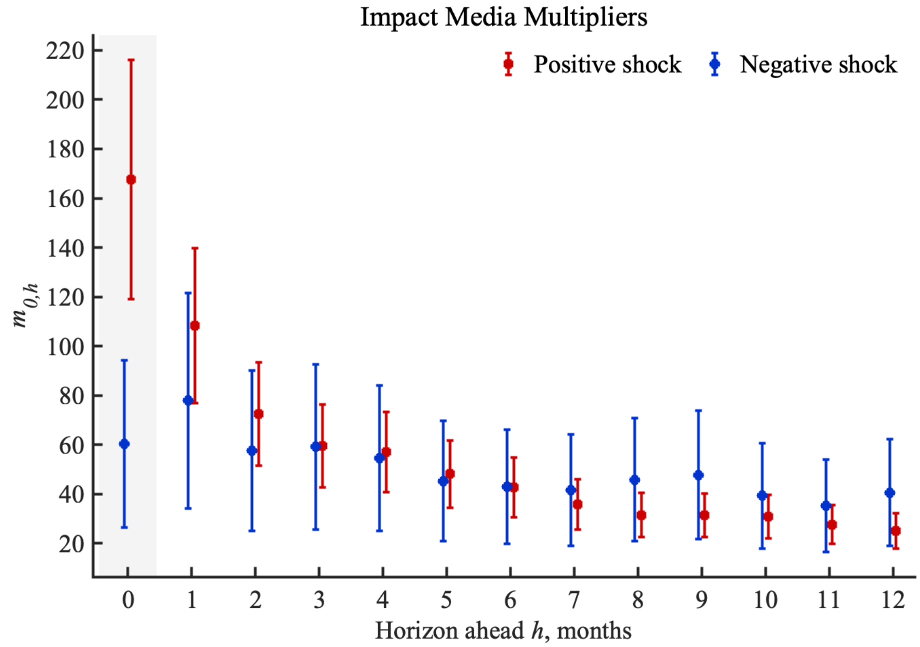

Notes: Red squares and blue dots denote point estimates for positive and negative impact media multipliers, $\hat{m}^{+}_{0,h}$ and $\hat{m}^{-}_{0,h}$ respectively. Error bars represent 95% confidence bands. Gray shaded areas indicate horizons where the difference between estimated positive and negative multipliers, $\hat{m}^{+}_{0,h} - \hat{m}^{-}_{0,h}$, is statistically significant at the 5% level based on the HAC-based test of Ramey and Zubairy (2018).

Figure 3: Estimated impact media multipliers for positive and negative unemployment shocks

Zubairy (2018) and Ben Zeev et al. (2023).

Figure 3 reports the estimated media multipliers for positive and negative unemployment shocks when the media discounts future unemployment realizations up to horizon h, $\hat{m}^{+}_{0,h}$ and $\hat{m}^{-}_{0,h}$, together with 95% confidence bands. The x-axis denotes the horizon ahead h, in months, of the cumulative unemployment change, $\sum_{j=0}^{h} \Delta U_{t+j}$. The y-axis denotes the contemporaneous response of negativity per one-percentage point cumulative change in unemployment. Gray shaded areas mark horizons for which the difference between the two estimated impact multipliers, $\hat{m}^{+}_{0,h} - \hat{m}^{-}_{0,h}$, is significantly different from zero.

For $h = 0$, the regression nests the static specification in equation (14), with two key extensions: (i) we instrument for positive and negative changes in unemployment, and (ii) we allow the control variables to vary by the sign of the shock. These modifications, however, do not meaningfully alter the conclusions drawn from the simpler regression. When only information about current unemployment is included in the model, the evidence of a negativity bias remains robust: media coverage responds more strongly to increases in unemployment than to decreases. Using a HAC-based test as in Ramey and Zubairy (2018), we reject the null hypothesis that the media multipliers for positive and negative unemployment shocks are equal at the contemporaneous horizon, confirming a statistically significant difference with a p-value of 0.001. The only notable difference with respect to the static regression is that the response of negativity to decreases in unemployment becomes statistically significant when instrumented with the unemployment shock.

However, as the forecasting horizon h increases and we control for expected future realizations of unemployment, the gap between the media multipliers quickly narrows, with the asymmetry disappearing entirely by $h = 3$ months ahead. Consistently, we do not reject the null hypothesis $m^+_{0,h} = m^-_{0,h}$ for any $h > 0$, indicating that the media's responses to positive and negative unemployment changes of comparable magnitude and persistence are statistically indistinguishable.

This dynamic pattern is fully consistent with the predictions of our theoretical framework. When unemployment dynamics are asymmetric ($\phi_1 > \phi_2$), increases in unemployment are more persistent and therefore signal a larger cumulative deterioration in labor market conditions. If the media respond to this expected path but the econometrician controls only for the contemporaneous change ($h = 0$), the estimated impact media multiplier for positive shocks, $\hat{m}^+_{0,0}$, will appear larger than that for negative shocks, $\hat{m}^-_{0,0}$. This leads to wrongly conclude in favor of a negativity bias even when the media sensitivities are symmetric ($m^+_0 = m^-_0$). As the forecasting horizon h increases and the estimation explicitly accounts for the asymmetry in the expected path of unemployment, the apparent bias is absorbed, and the difference between $\hat{m}^+_{0,h}$ and $\hat{m}^-_{0,h}$ vanishes.

Our results therefore indicate that the apparent bias towards negative economic events is not intrinsic to media behavior, but reflects the interaction between asymmetric unemployment dynamics and the media's forward-looking coverage. Once these features are properly accounted for, the evidence of a negativity bias disappears. This finding is novel and contrasts with previous studies that relied on static models unable to capture the dynamics of unemployment (see Soroka et al. (2018) for a review).

4.2 Dynamic media mulipliers

So far, we have shown that the impact media multiplier is symmetric once we account for asymmetries in unemployment dynamics and the media's forward-looking behavior. Nevertheless, symmetry in the impact response does not preclude the possibility that media coverage evolves differently over time following positive versus negative unemployment shocks. In other words, the future time path of media coverage may still differ depending on the sign of the shock.

To examine this, we compute the dynamic media multiplier, which captures the cumulative media response to a given cumulative unemployment rate change, separately for positive and negative shocks. Specifically, we estimate the following non-linear local projection for $i = 0, ..., k$

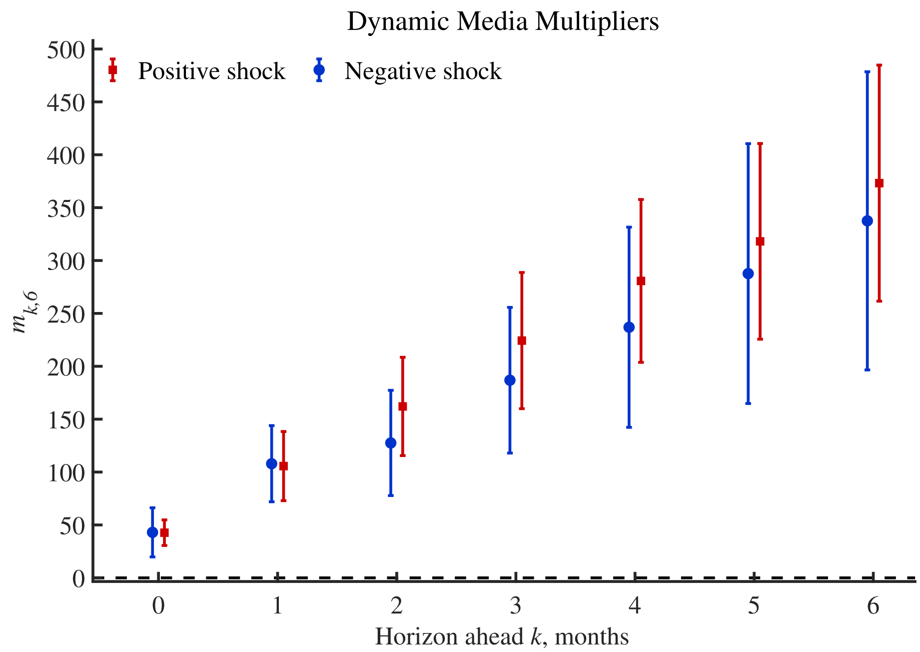

Notes: Red squares and blue dots denote point estimates for positive and negative impact media multipliers, $\hat{m}^+_{k,h}$ and $\hat{m}^-_{k,h}$ respectively. Error bars represent 95% confidence bands. Gray shaded areas indicate horizons where the difference between estimated positive and negative multipliers, $\hat{m}^+_{k,h} - \hat{m}^-_{k,h}$, is statistically significant at the 5% level based on the HAC-based test of Ramey and Zubairy (2018).

Figure 4: Estimated dynamic media multipliers for positive and negative unemployment shocks

and a given h:

$$\sum_{i=0}^{k} M_{t+i} = m^+_{k,h}\Big(S^+_t \times \sum_{j=0}^{h} \Delta U_{t+j}\Big) + m^-_{k,h}\Big(S^-_t \times \sum_{j=0}^{h} \Delta U_{t+j}\Big) + \gamma^+_{k,h}(S^+_t \times \mathbf{x}_{t-1}) + \gamma^-_{k,h}(S^-_t \times \mathbf{x}_{t-1}) + u_{t+k} \quad (16)$$

where $m^+_{k,h} = m^+_{0,h} + \ldots + m^+_{k,h}$ and $m^-_{k,h} = m^-_{0,h} + \ldots + m^-_{k,h}$. In the previous subsection we showed that $m^+_{0,h} = m^-_{0,h}$ for $h > 0$. Now, we are interested in testing $m^+_{i,h} = m^-_{i,h}$ for $i = 1, \ldots, k$. In the baseline, we set $h = 6$ as this represents a sufficiently forward-looking horizon at which the estimated impact media multipliers no longer differ by shock sign. Nonetheless, we show in the robustness section that the results are robust to alternative choices of h and to allowing the cumulative unemployment window to move with the projection horizon, that is replacing ΔU_{t+j} with ΔU_{t+j+i} in equation (16).

Figure 4 reports the estimated dynamic media multipliers for positive and negative unemployment changes across k horizons, $\hat{m}^+_{k,6}$ and $\hat{m}^-_{k,6}$. The x-axis denotes the horizon k, in months, over which the cumulative sum of media coverage, $\sum_{i=0}^{k} M_{t+i}$, is computed. Consistent with the findings for the impact multipliers, we find no significant difference between the dynamic media multipliers following positive and negative unemployment shocks for any horizon considered. This suggests that the media response is symmetric not only in its impact, but throughout its entire dynamics.

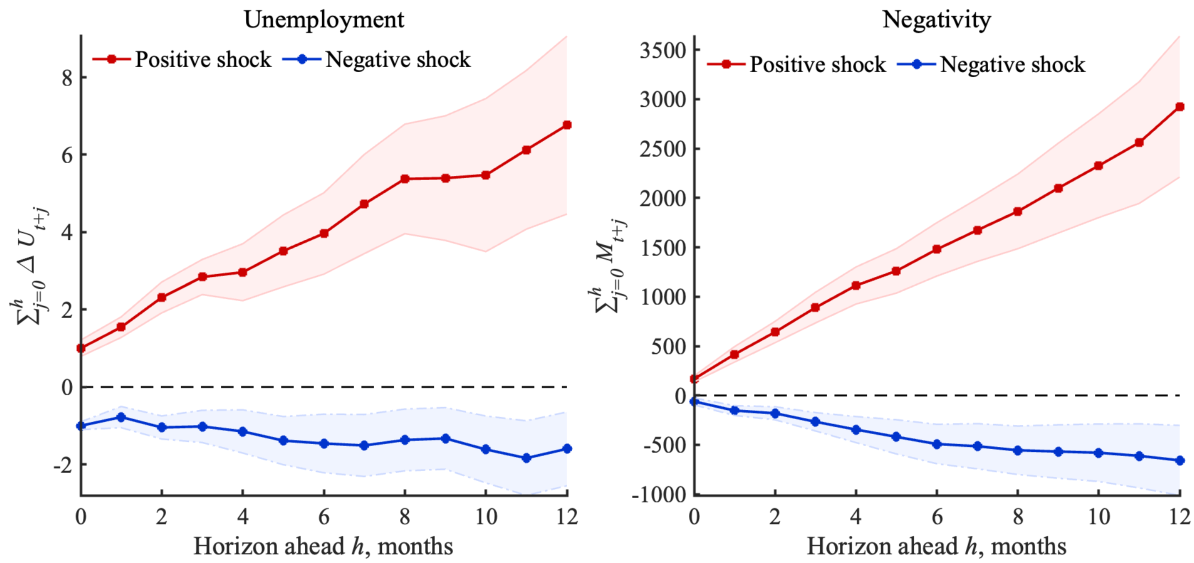

Notes: Red squares and blue dots denote point estimates for positive and negative unemployment shocks, respectively. Shaded areas represent 95% confidence bands.

Figure 5: Estimated impulse responses for positive and negative unemployment shocks

We conclude that no contemporaneous nor dynamic negativity bias is present in the media coverage of unemployment.

4.3 Media coverage reflects the dynamics of unemployment

To uncover the mechanisms behind the baseline media multiplier estimates, we separately examine the cumulative responses of negativity and of the unemployment rate change to positive versus negative unemployment shocks, in the spirit of Ben Zeev et al. (2023). To this end, we estimate the following regression for $j = 0, ..., h$:

$$\sum_{j=0}^{h} Y_{t+j} = \beta_h^+\left(S_t^+ \times \hat{\epsilon}_t^+\right) + \beta_h^-\left(S_t^- \times \hat{\epsilon}_t^-\right) + \gamma_h^+(S_t^+ \times \mathbf{x}_{t-1}) + \gamma_h^-(S_t^- \times \mathbf{x}_{t-1}) + u_{t+h} \tag{17}$$

where Y_t is either the unemployment rate change or negativity. The control variables $\mathbf{x}_{t-1}$ and indicator functions S_t^+ and S_t^- are the same used in the previous regressions.

Figure 5 illustrates the cumulative responses of both the unemployment rate change and negativity to identified positive and negative unemployment shocks, estimated using Equation (17).[12] The left panel shows that increases in unemployment generate a more persistent and pronounced rise in

[12] We report estimated (non-cumulative) impulse responses for the unemployment rate change and negativity in Figure D.3. The results are consistent with those obtained from the cumulative responses.

the unemployment rate compared to decreases. The right panel demonstrates that media negativity closely tracks this asymmetry: net coverage increases more sharply and persistently following positive unemployment shocks, but exhibits a smaller response to negative shocks. This pattern mirrors the underlying dynamics of unemployment itself and underscores our main result: what appears to be a media negativity bias reflects the asymmetric nature of unemployment dynamics and the forward-looking behavior of the media, rather than an inherent disproportionate focus on bad news. Unemployment rises faster and more persistently than it falls, leading to more negative news as media anticipate and reflect this asymmetry.

We show in Figure D.4 in the Appendix that both bad and good news respond to unemployment shocks and jointly shape the overall response of negativity. This rules out the interpretation that the result is driven solely by movements in bad news and underscores that good news is also an integral part of the adjustment. Bad news rises and good news falls after adverse shocks, and vice versa. This evidence strengthens our interpretation that the observed dynamics reflect the asymmetric nature of unemployment fluctuations rather than a simple bias toward bad news.

5 Robustness

This section presents several sensitivity checks. All related figures are reported in Appendix D.

5.1 Alternative model specifications

We first consider the robustness of our main findings to alternative lag and variable specifications in the local projection equation (15) and in the TVAR used to identify unemployment shocks. In the baseline specification, we set the lag length of the TVAR–and of the local projection–to 6, based on the average of the BIC and AIC criteria to ensure parsimoniousness. We re-estimate the model using the lag length suggested by the AIC (8 lags), which favors longer lag structures relative to the BIC. The results remain qualitatively unchanged (see Panel (a) of Figure D.5), confirming that our main findings are not sensitive to the choice of higher lags.

We also assess the robustness of our results by computing the impact media multiplier using each newspaper in the sample separately. Panels (b), (c), and (d) of Figure D.5 report the multipliers for the New York Times, Washington Post, and Wall Street Journal, respectively. Overall, the results are broadly consistent with the baseline. At the same time, some interesting heterogeneity emerges across outlets: the New York Times exhibits a significant bias up to three horizons ahead before it dissipates. The Washington Post shows little evidence of bias at any horizon h, with significance only at the 10% level on impact. The Wall Street Journal closely mirrors the baseline pattern. This heterogeneity may reflect differences in how forward-looking the newspapers are.

As discussed in Section 2, our negativity measure excludes articles captured by both queries and does not adjust for the overall volume of news coverage in each newspaper. Figure D.6 reports results based on two alternative measures: panel (a) is based on negativity including ambiguous articles and panel (b) on negativity normalized by the total number of articles published each month. The results are very similar to our baseline.

We also assess the robustness of our results to including more variables in the model. In the baseline specification, we use unemployment changes and negativity, as this provides the most intuitive framework for analyzing the media response to positive and negative unemployment shocks. To examine whether our findings are driven by the simplicity of this setup, we augment the system of variables in the TVAR, and consequently the controls in the local projection, with additional macroeconomic and forward-looking indicators: industrial production growth, employment growth, and stock price growth. We then re-identify unemployment shocks in the same way as in the baseline, as the shock that maximizes the contribution to the short-run volatility of the unemployment rate change. The results remain fully consistent with our main findings, confirming that our conclusions are not sensitive to the specification of the model (see Panel (e) of Figure D.5). Because adding an increasing amount of variables makes the TVAR less parsimonious, a complementary approach is to enrich the set of controls in both stages of the local projection only, while keeping the identification of unemployment shocks based on the more parsimonious TVAR. Panel (f) of Figure D.5 uses the baseline unemployment shocks and estimates the local projection with industrial production growth, employment growth, stock price growth, consumer sentiment and PCE inflation as controls. Again, the results are consistent with the baseline.

As a final check, we evaluate our baseline findings with respect to alternative identifications with the max-share approach of the TVAR model. First, we vary the frequency window used in the max-share approach. While the baseline specification focuses on fluctuations within the 12–36 month range, we extend the window to 12–96 months (8 years), consistent with the business-cycle frequency window used in Angeletos et al. (2020). The resulting impact media multipliers remain very similar to the baseline (see Panel (a) of Figure D.7), suggesting that our findings are not sensitive to the choice of frequency window. Second, we modify the target variable in the identification stage while maintaining the same max-share approach. Instead of identifying shocks based on changes in unemployment, we re-estimate the TVAR using industrial production growth and employment growth as target variables (see Panel (b) of Figure D.7) in the larger TVAR with industrial production growth, employment growth, unemployment changes and stock price growth.

In the local projection, we now instrument either industrial production growth or employment growth with the respective shocks. In these two cases, the media multipliers are of opposite sign, since industrial production and employment growth increases refer to positive news, and should thus reduce our negativity measure. Both alternative specifications yield qualitatively similar results to the baseline. Although the media multiplier are statistically different on impact for industrial production shocks, these discrepancies dissipate once we account for the forward-looking behavior of the media and the asymmetric dynamics in industrial production growth. This robustness aligns with the evidence in Angeletos et al. (2020), which shows that maximizing with respect to different macroeconomic indicators often captures the same underlying business-cycle dynamics.

5.2 Alternative definitions of media multipliers

We now examine the robustness of our findings to alternative definitions of the media multipliers.

In our specification of the dynamic media multipliers, we fixed the forecasting horizon to $h = 6$. Figure D.8 shows that our results are not sensitive to this choice. The first three panels show the dynamic media multipliers computed for $h = 3$, $h = 9$, and $h = 12$. In all cases, there is no statistically significant difference between the media multipliers of positive and negative unemployment shocks. The last panel of Figure D.8 reports results when the cumulative unemployment window is allowed to shift with the projection horizon, replacing ΔU_{t+j} with ΔU_{t+j+i} in equation (16). For impact multipliers, this alternative allows to account for the forward-looking behavior of the media in a dynamic media multiplier setting. Once again, the estimates show no meaningful difference between positive and negative dynamic media multipliers, regardless of specification.

We also summarize the media response using a standard cumulative media multiplier, in line with the fiscal policy literature (see Ramey and Zubairy (2018); Ben Zeev et al. (2023)). Specifically, we set $k = h$ in equation (16), cumulating both media coverage and unemployment changes over the same horizon. This specification abstracts from the forward-looking behavior of the media, but provides a useful summary measure of the overall response of negativity per one-percentage point cumulative change in unemployment. Figure D.9 shows that the cumulative multipliers for positive and negative shocks closely track each other beyond the impact horizon. Consistent with the dynamic estimates, any apparent asymmetry at $h = 0$ disappears once future unemployment realizations are included, and the cumulative media response becomes symmetric across horizons.

Finally, we re-estimate the media multipliers using the unemployment rate in *levels*. In this specification, the right-hand side of the local projection includes the contemporaneous unemployment rate and its lags. Everything else is identical to the baseline. Figure D.10 reports the impact, dynamic, and cumulative multipliers. The results remain qualitatively unchanged, confirming that our main conclusions are robust to modeling unemployment in levels or differences.

5.3 Alternative measures of media coverage

Although our baseline focuses on unemployment because it is a key business cycle indicator and was used in prior studies of negativity bias, we also perform a broader validation exercise using the San Francisco Fed's News Sentiment Index (NSI), which summarizes the tone of a large set of economic news articles, not related to unemployment only.

We begin by estimating a simple static regression of the news sentiment index on a constant, its four lags and positive and negative changes in industrial production. We report the results in Table D.1. Consistent with earlier evidence, news sentiment declines significantly following negative changes in industrial production but shows no significant response to positive changes, indicating an apparent negativity bias when using a static specification.

Next, we estimate our impact media multipliers using industrial production shocks identified through the same TVAR framework, where we replace media negativity with NSI. Figure D.11 reports the estimated impact multipliers (Panel (a)) and the cumulative impulse responses (Panel (b)). The multipliers differ significantly for the first three months following the shock, but converge thereafter. Once we account for the asymmetric dynamics of industrial production and the forward-looking behavior of the media, the difference between responses to positive and negative shocks disappears. As for unemployment, industrial production responds more strongly and persistently to negative shocks than to positive, and the media reflect this asymmetry in the short run.

Together with our unemployment results, this exercise confirms that the disappearance of the apparent negativity bias is not specific to unemployment coverage. Rather, it reflects a more general mechanism: when economic asymmetries and expectation-driven reporting are properly accounted for, media tone responds symmetrically to good and bad economic news.

6 Concluding remarks

Using nonlinear time series techniques and two measures of newspaper coverage of bad and good economic events, we provide novel empirical evidence on the absence of a negativity bias in economic news coverage. News coverage is more responsive to negative than positive economic developments because bad economic shocks have larger and more persistent effects on economic variables than good shocks, and media reflect these dynamics. Not taking this asymmetry into account may lead to misleading conclusions about the way media cover bad and good economic events.

References

Andolfatto, D. (1997). Evidence and theory on the cyclical asymmetry in unemployment rate fluctuations. *Canadian Journal of Economics*, 709–721.

Angeletos, G.-M., Collard, F., & Dellas, H. (2020). Business-cycle anatomy. *American Economic Review*, *110*(10), 3030–70.

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

Ben Zeev, N., Ramey, V. A., & Zubairy, S. (2023, May). Do government spending multipliers depend on the sign of the shock? *AEA Papers and Proceedings*, *113*, 382–87.

Binder, C. (2017). Federal reserve communication and the media. *Journal of Media Economics*, *30*(4), 191-214.

Binder, C., Frank, P., & Ryngaert, J. M. (2025). The causal effect of news on inflation expectations. *NBER Working Papers, 34088*.

Blinder, A., & Krueger, A. (2004). What does the public know about economic policy, and how does it know it? *Brookings Papers on Economic Activity*, *35*(1), 327-397.

Bybee, L., Kelly, B., Manela, A., & Xiu, D. (2024). Business news and business cycles. *The Journal of Finance*, *79*(5), 3105-3147.

Caldara, D., & Iacoviello, M. (2022). Measuring geopolitical risk. *American Economic Review*, *112*(4), 1194–1225.

Caldara, D., Iacoviello, M., Molligo, P., Prestipino, A., & Raffo, A. (2020). The economic effects of trade policy uncertainty. *Journal of Monetary Economics*, *109*, 38–59.

Carroll, C. D. (2003). Macroeconomic expectations of households and professional forecasters. *the Quarterly Journal of economics*, *118*(1), 269–298.

Chahrour, R., Nimark, K. P., & Pitschner, S. (2021, December). Sectoral media focus and aggregate fluctuations. *American Economic Review*, *111*(12), 3872-3922.

Dupraz, S., Nakamura, E., & Steinsson, J. (2025). A plucking model of business cycles. *Journal of Monetary Economics*, *152*, 103766.

Ellingsen, J., Larsen, V. H., & Thorsrud, L. A. (2022). News media versus fred-md for macroeconomic forecasting. *Journal of Applied Econometrics*, *37*(1), 63-81.

Fogarty, B. J. (2005). Determining Economic News Coverage. *International Journal of Public Opinion Research*, *17*(2), 149-172.

Goidel, R. K., & Langley, R. E. (1995). Media coverage of the economy and aggregate economic evaluations: Uncovering evidence of indirect media effects. *Political Research Quarterly*, *48*(2), 313–328.

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

Link, S., Peichl, A., Pfaeuti, O., Roth, C., & Wohlfart, J. (2025). Attention to the macroeconomy. *CEBI working paper series 23-09*.

Morley, J., & Piger, J. (2012). The asymmetric business cycle. *Review of Economics and Statistics*, *94*(1), 208–221.

Neftci, S. (1984). Are economic time series asymmetric over the business cycle? *Journal of Political Economy*, *92*(2), 307-28.

Nimark, K. P., & Pitschner, S. (2019). News media and delegated information choice. *Journal of Economic Theory*, *181*, 160–196.

Nimark, K. P., & Sundaresan, S. (2019). Inattention and belief polarization. *Journal of Economic Theory*, *180*, 203–228.

Pizzinelli, C., Theodoridis, K., & Zanetti, F. (2020, August). State dependence in labor market fluctuations. *International Economic Review*, *61*(3), 1027-1072.

Plagborg-Møller, M., & Wolf, C. K. (2021). Local projections and vars estimate the same impulse responses. *Econometrica*, *89*(2), 955–980.

Ramey, V. A., & Zubairy, S. (2018). Government spending multipliers in good times and in bad: evidence from us historical data. *Journal of political economy*, *126*(2), 850–901.

Shapiro, A. H., Sudhof, M., & Wilson, D. J. (2022, None). Measuring news sentiment. *Journal of Econometrics*, *228*(2), 221-243.

Sims, C. A. (2003). Implications of rational inattention. *Journal of monetary Economics*, *50*(3), 665–690.

Soroka, S. N. (2006). Good news and bad news: Asymmetric responses to economic information. *The journal of Politics*, *68*(2), 372–385.

Soroka, S. N. (2012). The gatekeeping function: Distributions of information in media and the real world. *The Journal of Politics*, *74*(2), 514–528.

Soroka, S. N., Daku, M., Hiaeshutter-Rice, D., Guggenheim, L., & Pasek, J. (2018). Negativity and positivity biases in economic news coverage: Traditional versus social media. *Communication Research*, *45*(7), 1078–1098.

Soroka, S. N., Stecula, D. A., & Wlezien, C. (2015). It's (change in) the (future) economy, stupid: economic indicators, the media, and public opinion. *American Journal of Political Science*, *59*(2), 457–474.

Appendix

A - News indexes

We construct our bad and good news indexes using newspaper articles from *Dow Jones Factiva*. We focus our search to three major US newspapers by circulation, namely *The Wall Street Journal*, *The New York Times* and *The Washington Post*, and to news related to the US economy over the time period from June 1980 to December 2019.

For each newspaper, we look for all the articles, in a given month, in which the word "unemployment" appears within a predetermined distance, in any order, to one denoting an increase or high level or to one denoting a decrease or low level. More specifically, we first define two semantic groups, one for the words qualifying a *bad news* and another for words qualifying a *good news*.

We choose the sets of words in each group, as well as the word distance from "unemployment" after a human pre-reading of 150 articles. We select 15 dates at random and, for each date, we randomly sample 10 articles in which the word "unemployment" appears. We list separately all words that appear together "unemployment" when describing increases or high levels and decreases or low levels and build two search queries.:

- *Bad news group*. The words included in this group have one of the following roots: "high-", "increas-", "ris-", "rose-", "soar-", "rais-", "up-", "climb-", "accelerat-".
- *Good news group*. The words included in this group have one of the following roots: "down-" or "low-" or "slow-" or "decreas-", "drop-", "fall-", "fell-", "slip-", "declin-".

We classify an article as a *bad news* item if the word "unemployment" appears within a 5-word distance to a word belonging to semantic bad news group, but not within a 1-word distance to a word in semantic good news group or to a negation ("no" or "not"). Symmetrically, we define an article as a *good news* item if the word "unemployment" appears within a 5-word distance to a word belonging to semantic good news group, but not within a 1-word distance to a word in semantic group bad news or to a negation ("no" or "not").

We choose the 5-word distance criteria to maximize the probability that the corresponding word in the bad news group or in the good news group is related to the word "unemployment" and not to other words. We obtain very similar results if we restrict this criteria to 4- or 3-word distance.

We clean our two measures of bad and good news by substracting from both the number of articles that can be classified as belonging to both groups according to our criteria. In fact, this class of articles cannot be clearly classified as positive or negative, either because these articles deliver mixed signals about unemployment,[13] so that their resulting tone is *neutral*, or because the word "unemployment" is incidentally mentioned close to a word in the bad news group and the good news group, even if the article does not include direct information about unemployment (e.g. articles reporting presidential talks close to the elections). The articles belonging to this last category represent on average 6% of total articles over the period considered. After cleaning the measures, the number of all *bad news* in a given month is the value of the bad news index for that month, while the number of all *good news* in a given month is the value of the good news index for that month.

B – Model

In this Appendix, we solve the model presented in Section 3 using numerical methods, relaxing the simplifying assumption that the state is absorbing.

Unemployment dynamics. The change in unemployment follows a threshold AR(1):

$$\Delta U_t \;=\; \phi_1\big(I^+_{t-1} \times \Delta U_{t-1}\big) \;+\; \phi_2\big(I^-_{t-1} \times \Delta U_{t-1}\big) \;+\; \varepsilon_t, \tag{18}$$

where $0 \leq \phi_i < 1$ $(i = 1, 2)$, $I^+_{t-1} = \mathbf{1}\{\Delta U_{t-1} > 0\}$, $I^-_{t-1} = \mathbf{1}\{\Delta U_{t-1} < 0\}$, and ε_t is i.i.d. with mean zero and variance σ^2_ε.

Media coverage. When setting the number of stories, the media consider the current change and the expected cumulative change over h horizons ahead:

$$M_t \;=\; m^+\left(I^+_t\; E\left[\sum_{j=0}^{h} \Delta U_{t+j} \,\middle|\, \Delta U_t > 0\right]\right) \;+\; m^-\left(I^-_t\; E\left[\sum_{j=0}^{h} \Delta U_{t+j} \,\middle|\, \Delta U_t < 0\right]\right) \;+\; \nu_t, \tag{19}$$

where ν_t is i.i.d. with mean zero and variance σ^2_ν. The parameters m^+_0 and m^-_0 are the impact media multipliers for positive and negative current changes, respectively.

We consider the case in which $\Delta U_t > 0$, $I^+_t = 1$ and $I^-_t = 0$. The case $\Delta U_t < 0$ is analogous. The

[13] For example, on the 12th of March 2010, *The Wall Street Journal* writes "[...] initial claims for unemployment insurance dropped to 462,000 in the week ended March 6th, down 6,000 from the week before. Meanwhile, the number of people collecting unemployment checks rose 37,000 to 4.6 million in the week ending Feb. 27".

one-step ahead forecast is:

$$
\begin{aligned}
E(\Delta U_{t+1}|\Delta U_t > 0) &= E\left[I_t^+\phi_1\Delta U_t + I_t^-\phi_2\Delta U_t + \varepsilon_{t+1}|\Delta U_t > 0\right] \\
&= \phi_1\Delta U_t
\end{aligned}
$$

since $E[\varepsilon_{t+1}|\Delta U_t > 0] = 0$. The two-step ahead forecast is:

$$
\begin{aligned}
E(\Delta U_{t+2}|\Delta U_t > 0) &= E\left[(I_{t+1}^+\phi_1 + I_{t+1}^-\phi_2)\Delta U_{t+1} + \varepsilon_{t+2}|\Delta U_t > 0\right] \\
&= E\Big[\big(I_{t+1}^+\phi_1 + I_{t+1}^-\phi_2\big)\big(\phi_1\Delta U_t + \varepsilon_{t+1}\big) + \varepsilon_{t+2}|\Delta U_t > 0\Big]
\end{aligned}
$$

The three-step ahead forecast is:

$$
\begin{aligned}
E(\Delta U_{t+3}|\Delta U_t > 0) &= E\left[(I_{t+2}^+\phi_1 + I_{t+2}^-\phi_2)\Delta U_{t+2} + \varepsilon_{t+3}|\Delta U_t > 0\right] \\
&= E\left[(I_{t+2}^+\phi_1 + I_{t+2}^-\phi_2)\Big((I_{t+1}^+\phi_1 + I_{t+1}^-\phi_2)(\phi_1\Delta U_t + \varepsilon_{t+1}) + \varepsilon_{t+2}\Big) + \varepsilon_{t+3}|\Delta U_t > 0\right]
\end{aligned}
$$

For the h-step ahead:

$$
\begin{aligned}
E(\Delta U_{t+h} \mid \Delta U_t > 0) &= E\Bigg[\ \Big(\prod_{i=0}^{h-1}(\phi_1 I_{t+i}^+ + \phi_2 I_{t+i}^-)\Big)\,\Delta U_t \\
&\qquad + \sum_{i=1}^{h-1}\Big(\prod_{j=h-i}^{h-1}(\phi_1 I_{t+j}^+ + \phi_2 I_{t+j}^-)\Big)\,\varepsilon_{t+h-i} \ + \ \varepsilon_{t+h}\ \Bigg|\ \Delta U_t > 0\Bigg].
\end{aligned}
$$

As the forecasting horizon grows, the analytical expression for the forecast of unemployment changes become increasingly complicated due to the compounding of the conditional probabilities of future states and the expectation of the cross-products between the indicator variables and shocks. However, the model can be easily solved using numerical methods.

We select a grid $\mathcal{G}$ of initial values for ΔU_t. For each $x \in \mathcal{G}$, we simulate L independent forward paths of length H for the process

$$
\Delta U_{t+h}^{(\ell)} \ = \ \phi_1\,\mathbf{1}\{\Delta U_{t+h-1}^{(\ell)} > 0\}\,\Delta U_{t+h-1}^{(\ell)} \ + \ \phi_2\,\mathbf{1}\{\Delta U_{t+h-1}^{(\ell)} < 0\}\,\Delta U_{t+h-1}^{(\ell)} \ + \ \varepsilon_{t+h}^{(\ell)}, \quad h = 1, \ldots, H,
$$

starting from $\Delta U_t^{(\ell)} = x$, with shocks $\{\varepsilon_{t+h}^{(\ell)}\}$ i.i.d.. The Monte Carlo estimator of the h-step conditional expectation is then

$$\widehat{E}_L\big[\Delta U_{t+h} \mid \Delta U_t = x\big] \;\equiv\; \frac{1}{L}\sum_{\ell=1}^{L} \Delta U_{t+h}^{(\ell)}(x),$$

which converges to $E[\Delta U_{t+h} \mid \Delta U_t = x]$ by the law of large numbers. This yields a simulated expectation surface on the grid $\mathcal{G} \times \mathcal{H}$, with $\mathcal{H} = \{1, \ldots, H\}$.

We then simulate R samples of length T for the unemployment process. For each realization $\Delta U_t^{(r)}$ with $t = 1, ..., T$ and $r = 1, ..., R$, we choose the corresponding expectations h-step ahead on the grid $\mathcal{G}x\mathcal{H}$, $\widehat{E}_L\big[\Delta U_{t+h} \mid \Delta U_t^{(r)}\big]$ and construct media coverage as:

$$M_t^{(r)} = m^+ \left(\mathbf{1}\{\Delta U_t^{(r)} > 0\}\; \hat{E}_L\left[\sum_{h=0}^{H} \Delta U_{t+h} \middle| \Delta U_t^{(r)} > 0\right]\right) + m^- \left(\mathbf{1}\{\Delta U_t^{(r)} < 0\}\; \hat{E}_L\left[\sum_{h=0}^{H} \Delta U_{t+h} \middle| \Delta U_t^{(r)} < 0\right]\right) + v_t$$

Throughout all our exercises, we assume $m^+ = m^- = \alpha$, i.e. no media bias.

We set $(\phi_1, \phi_2, \sigma_\varepsilon) = (0.26,\ 0.02,\ 0.17)$, equal to the OLS estimates of the threshold AR(1) in (18). We choose $H = 6, T = 500$, $L = 10000$, $R = 10000$, and set the coverage scale to $\alpha = 100$. Figure 6 reports the Monte Carlo expectation surface $\widehat{E}[\Delta U_{t+h} \mid \Delta U_t = x]$ over the grid $\mathcal{G} \times \mathcal{H}$. The red areas correspond to positive values and the blue areas to negative values. Under this calibration, only a narrow region with $\Delta U_t < 0$ yields negative expected changes; even there, $\widehat{E}[\Delta U_{t+h} \mid x]$ turns positive by $h \geq 2$. By $h \approx 4$ the surface is essentially flat at about 0.02, regardless of the initial x. Overall, in this simple framework, and for empirically plausible (ϕ_1, ϕ_2), expected future changes in unemployment are positive for most initial conditions and horizons. Moreover, the expected changes in unemployment at horizons $h = 1, ..., 3$ for $\Delta U_t > 0$ are much larger, in absolute value, then for $\Delta U_t < 0$. Even if media discount these expectations symmetrically (i.e., $m^+ = m^- = \alpha$), this already implies (i) a higher volume of bad stories and ii) a stronger response to an increase in unemployment than to an equally sized decrease.

To show that a static regression can produce $\hat{\beta}_+ > \hat{\beta}_-$ even when the media have no intrinsic bias ($m^+ = m^- = \alpha$), we repeat the Monte Carlo exercise for $\phi_1 \in \{0.10,\ 0.26,\ 0.50\}$ while varying ϕ_2 on the grid $\{0.01, \ldots, \phi_1\}$. For each simulated sample ℓ we estimate

$$M_t^{(\ell)} \;=\; \beta_0 \;+\; \beta_+ \Delta U_t^{(\ell)}\, I_t^{+(\ell)} \;+\; \beta_- \Delta U_t^{(\ell)}\, I_t^{-(\ell)} \;+\; e_t^{(\ell)},$$

and compute the sample bias $\hat{\beta}_+^{(\ell)} - \hat{\beta}_-^{(\ell)}$. Figure 7 plots the average bias, $(\hat{\beta}_+ - \hat{\beta}_-)^L$, computed as the Monte Carlo mean of the bias, $\frac{1}{L}\sum_{\ell=1}^{L} \hat{\beta}_+^{(\ell)} - \hat{\beta}_-^{(\ell)}$, against ϕ_2 and for $\phi_1 \in \{0.10,\ 0.26,\ 0.50\}$. The larger the difference in the persistence $\phi_1 - \phi_2$, the larger the estimated bias. When $\phi_2 \simeq \phi_1$,

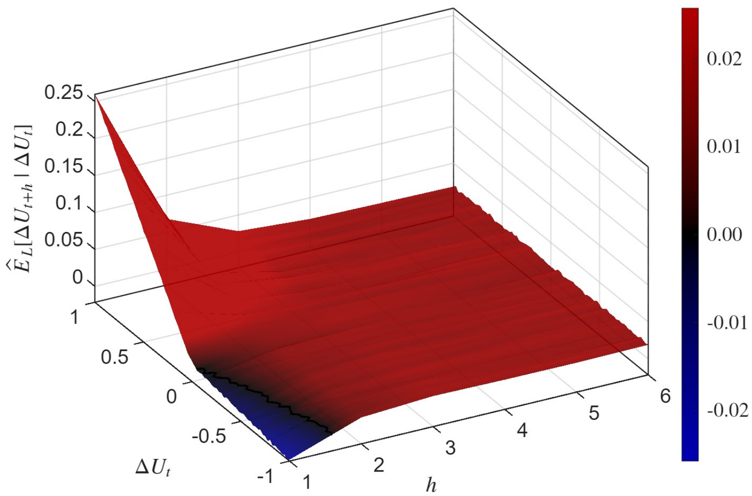

Figure 6: Conditional expectation surface $E[\Delta U_{t+h} \mid \Delta U_t]$. The blue area corresponds to negative and the red area to positive values for expected changes in unemployment.

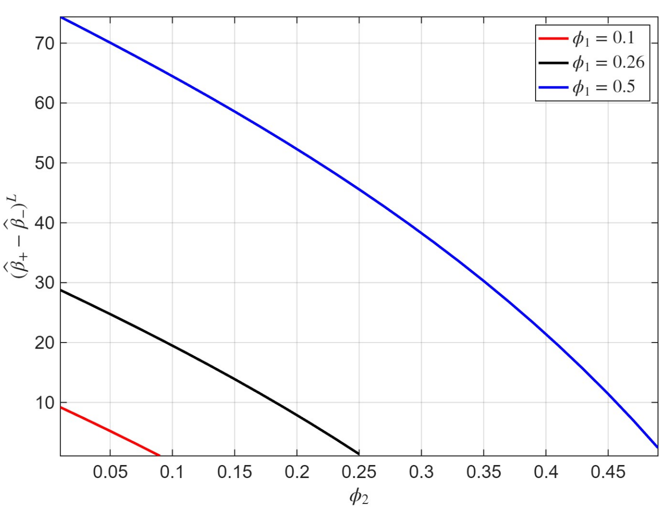

Figure 7: Estimated mean bias, $(\hat{\beta}_1 - \hat{\beta}_2)^L$, as a function of ϕ_2, for $\phi_1 = \{0.10, 0.26, 0.50\}$. The bias is computed as the mean of the differences in OLS coefficients estimated over R = 10,000 samples of length $T = 500$ for each pair (ϕ_1, ϕ_2)

the bias shrinks to zero. We conclude that, with state-dependent persistence, a static specification suggests a negativity bias in coverage even when none is present in the data-generating process. This reinforces the intuition we illustrated analytically in Section 3 with a simplified version of this model.

C - A Threshold VAR for unemployment and media dynamics

This Appendix describes the Threshold VAR model employed to derive the unemployment shock $\hat{\epsilon}_t$, subsequently used in the estimation of the media multipliers.

Let $y_t = [\Delta U_t;\ M_t]$, a vector including changes in unemployment and negativity. We model y_t as follows:

$$y_t = I^+_{t-1}[a + A(L)]y_{t-1} + I^-_{t-1}[b + B(L)]y_{t-1} + \varepsilon_t \tag{20}$$

where $\varepsilon_t \sim WN(0, \Sigma)$ is a vector of white noise residuals. $A(L) = A_1 + A_2L + ... + A_pL^{p-1}$ and $B(L) = B_1 + B_2L + ... + B_pL^{p-1}$ are matrix polynomials in the lag operator L, and a and b are vectors of constant terms. Again, I^+_{t-1} and I^-_{t-1} are indicator variables taking value 1 when ΔU_{t-1} is larger and smaller than 0, respectively. Thus, $A(L)$ are the VAR parameters governing the dynamics of the system when the first lag of the unemployment rate change is positive, while $B(L)$ are the VAR parameters in place when the change is negative.

We identify the unmployment shock as the one explaining most of the volatility in unemployment over short-run frequencies (see Angeletos et al. (2020)), defining the frequency domain between 12 and 36 months. With this model specification, the sign of the unemployment shock becomes the relevant state for the impulse response functions.

This empirical model is subject to a potential caveat. We derive the impulse responses below for the two regimes under the assumption that the regime set on impact stays in place for the whole horizon, in line with the simplifying assumption in the theoretical model of Section 3. That is, the sign of the shock determines the relevant state for the VAR dynamics. This is a reasonable assumption when the impulse responses of the change in unemployment are sufficiently persistent. This assumption is not the same for the nonlinear local projection, for which the state can change at each horizon. As shown below, the two approaches deliver very similar results.

Figure 8 shows the responses of unemployment and negativity to positive and negative unemployment shocks. The first column reports the impulse responses and the impact media multipliers, while the second column shows the corresponding asymmetry index – that is, the difference between the responses to positive and negative shocks. The first two rows show the impulse responses of unemployment and negativity, whereas the last row reports the impact multipliers.

Positive unemployment shocks trigger a larger and more persistent reaction of the media relative to negative shocks. The right panel indicates that the difference in the responses is positive and statistically significant. Likewise, the response of unemployment changes to a positive shock is significantly more persistent than the response to a negative one (first row). Results for the media multipliers are identical to those in the baseline: at horizon $h = 0$, a bias appears, but it vanishes

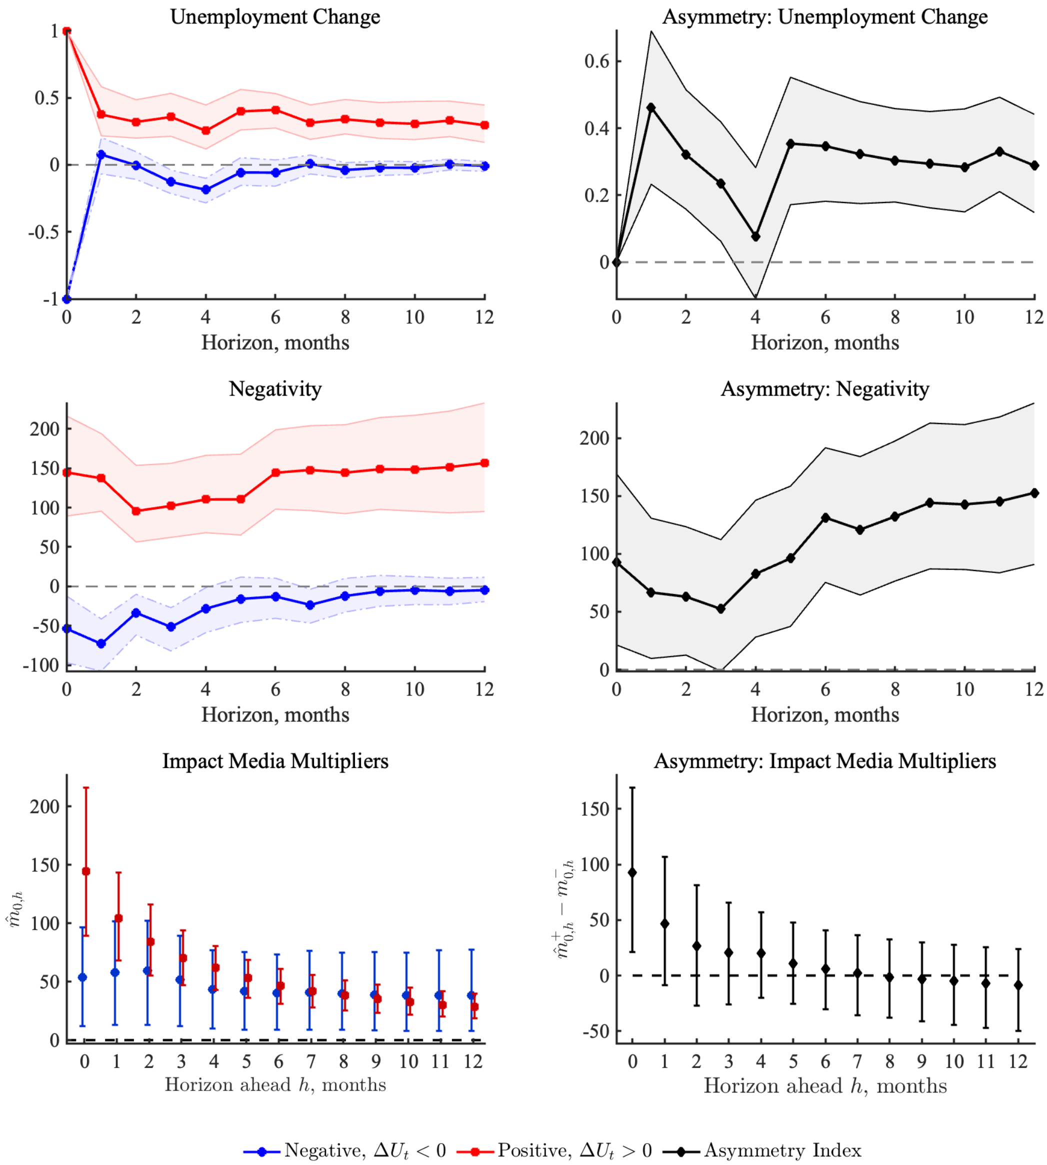

Figure 8: Estimated impulse responses for positive and negative unemployment shocks in the TVAR

once we control for asymmetries in unemployment.

Figure 9 plots the cumulative impulse responses at different horizons. When we cumulate the IRFs of unemployment changes and negativity, we find that, despite differences in the impact responses, the overall evolution of media negativity and unemployment is remarkably similar across positive and negative shocks. In other words, after a positive (negative) unemployment shock, both the unemployment rate and negativity increase (decrease) at a comparable rate. The cumulative responses closely mirror those obtained in Figure 5 using local projections.

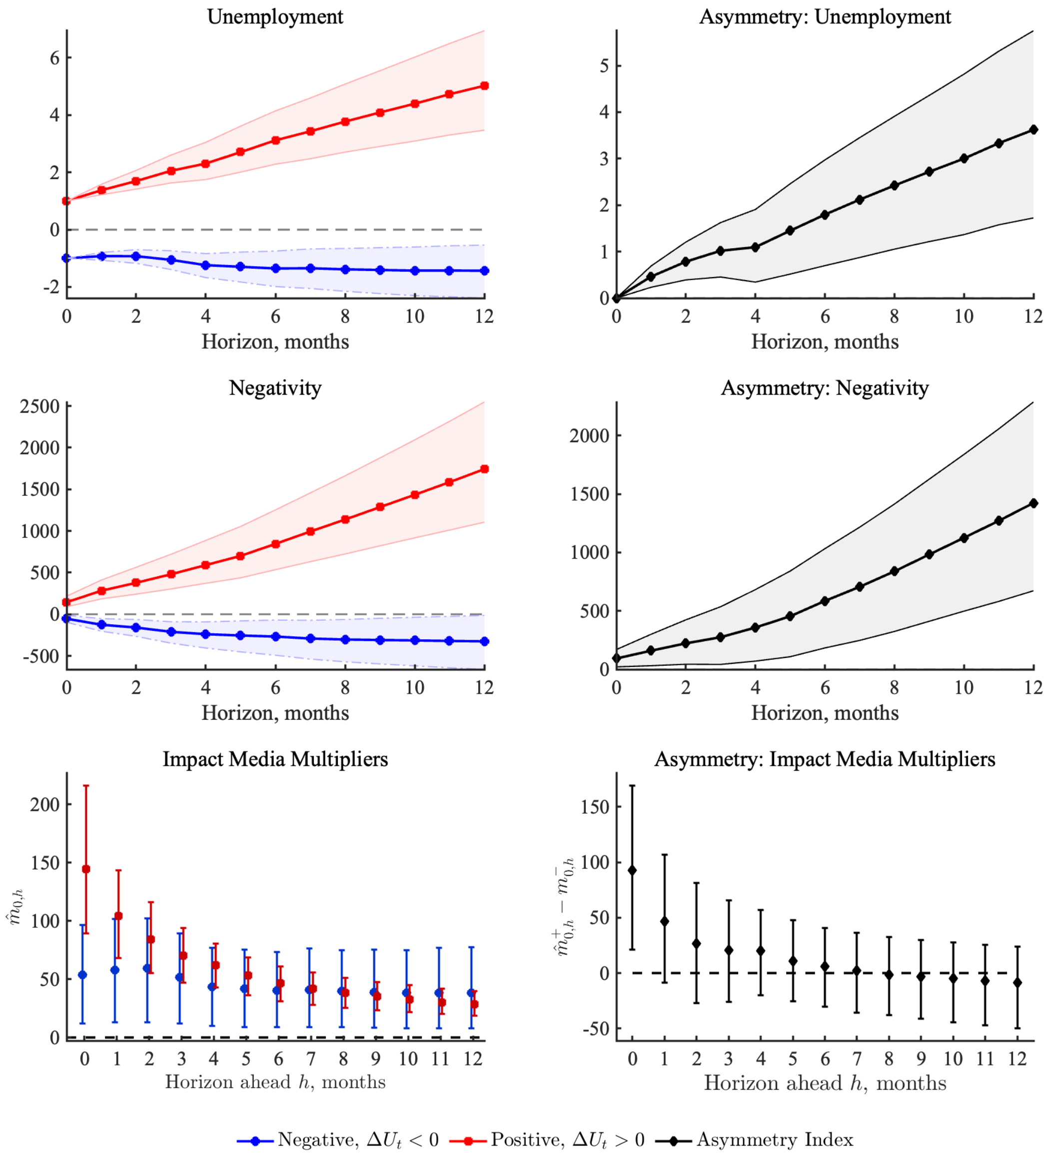

Figure 9: Estimated cumulative impulse responses for positive and negative unemployment shocks in the TVAR

Overall, the impulse responses and impact media multipliers obtained from the TVAR model closely track those derived from the nonlinear local projection. Both approaches include the same set of control variables, and the similarity in results reinforces the robustness of our baseline findings. While the two methods differ – local projections offering greater flexibility, and TVAR capturing regime-dependent dynamics – it is not surprising that they yield comparable responses. In linear settings, previous research has shown that local projections and VARs often produce similar IRFs when estimated on the same system (see Plagborg-Møller and Wolf (2021)). Nonetheless, we provide a simple proof for the equality of media multipliers under the two nonlinear approaches below.

LP–IV and TSVAR identity

Consider the dynamic LP–IV regression in (16):

$$\sum_{i=0}^{k} M_{t+i} \;=\; m^{+}_{k,h}\left(S^{+}_{t}\sum_{j=0}^{h}\Delta U_{t+j}\right) \;+\; m^{-}_{k,h}\left(S^{-}_{t}\sum_{j=0}^{h}\Delta U_{t+j}\right) \;+\; \gamma^{+}_{k,h}\big(S^{+}_{t}x_{t-1}\big) \;+\; \gamma^{-}_{k,h}\big(S^{-}_{t}x_{t-1}\big) \;+\; u_{t+k}.$$

Let $Z^{\pm}_{t} = S^{\pm}_{t}\hat{\epsilon}_{t}$ be the sign-specific instrument built from the TSVAR-identified unemployment shock $\hat{\epsilon}_{t}$, $Y^{(k)}_{t} = \sum_{i=0}^{k} M_{t+i}$ and $X^{\pm}_{t}(h) = S^{\pm}_{t}\sum_{j=0}^{h}\Delta U_{t+j}$.

We can rewrite $Y^{(h)}_{t}$ and $X_{t}(h)$ as:

$$Y^{(h)}_{t} = \Big(C^{+}_{M}(h)S^{+}_{t} + C^{-}_{M}(h)S^{-}_{t}\Big)\hat{\epsilon}_{t} + \nu_{t} \tag{21}$$

$$X_{t}(h) = \Big(C^{+}_{\Delta U}(h)S^{+}_{t} + C^{-}_{\Delta U}(h)S^{-}_{t}\Big)\hat{\epsilon}_{t} + \zeta_{t} \tag{22}$$

with $C^{\pm}_{M}(h) = \sum_{i=0}^{h}\mathrm{IRF}^{\pm}_{M}(i)$, $C^{\pm}_{\Delta U}(h) = \sum_{i=0}^{h}\mathrm{IRF}^{\pm}_{\Delta U}(i)$, $E[\eta_{t}|\hat{\epsilon}_{t}] = 0$ and $E[\zeta_{t}|\hat{\epsilon}_{t}] = 0$.

We show that, for $k = h$, the LP–IV dynamic multipliers and the TSVAR multiplier equate:

$$m^{\pm}_{h,h} \;=\; \frac{C^{\pm}_{M}(h)}{C^{\pm}_{\Delta U}(h)} \;=\; \mu^{\pm}_{h} \tag{23}$$

Proof. The IV estimators of $m^{+}_{h,h}$ and $m^{-}_{h,h}$ solve the moment conditions:

$$E\Big[Z^{+}_{t}\big(Y^{(h)}_{t} - m^{+}_{h,h}X^{+}_{t}(h) - m^{-}_{h,h}X^{-}_{t}(h)\big)\Big] = 0, \tag{24}$$

$$E\Big[Z^{-}_{t}\big(Y^{(h)}_{t} - m^{+}_{h,h}X^{+}_{t}(h) - m^{-}_{h,h}X^{-}_{t}(h)\big)\Big] = 0. \tag{25}$$

Because $S^{+}_{t}S^{-}_{t} = 0$, the cross-sign products vanish: $E[Z^{+}_{t}X^{-}_{t}(h)] = E[Z^{-}_{t}X^{+}_{t}(h)] = 0$, so (24) and (25) decouple by sign. We focus on the case $S^{+}_{t} = 1$ ($\hat{\epsilon}_{t} > 0$). The case $S^{-}_{t} = 1$ is analogous with the sign flipped.

Substitute (21) into $E[Z^{+}_{t}Y^{(h)}_{t}]$:

$$E[Z^{+}_{t}Y^{(h)}_{t}] = E\Big[S^{+}_{t}\hat{\epsilon}_{t}\big(S^{+}_{t}C^{+}_{M}(h) + S^{-}_{t}C^{-}_{M}(h)\big)\hat{\epsilon}_{t}\Big] + E[S^{+}_{t}\hat{\epsilon}_{t}\eta_{t}] = C^{+}_{M}(h)E\Big[S^{+}_{t}\hat{\epsilon}^{2}_{t}\Big], \tag{26}$$

where we used $S^{+}_{t}S^{-}_{t} = 0$, the law of iterated expectations to write: $E[S^{+}_{t}\hat{\epsilon}_{t}\eta_{t}] = E[S^{+}_{t}\hat{\epsilon}_{t}|\hat{\epsilon}_{t}]E[\eta_{t}|\hat{\epsilon}_{t}]$

and $E[\eta_t \mid \hat{\epsilon}_t] = 0$.

$$E[Z_t^+ X_t^+(h)] = S_t^+ \hat{\epsilon}_t \big(S_t^+ C_{\Delta U}^+(h) + S_t^- C_{\Delta U}^-(h)\big)\hat{\epsilon}_t\big] + E[S_t^+ \hat{\epsilon} \cdot S_t^+ \xi_t] = C_{\Delta U}^+(h) E\big[S_t^+ \hat{\epsilon}^2\big], \tag{27}$$

where we used $S_t^+ S_t^- = 0$, the law of iterated expectations to write: $E[S_t^+ \hat{\epsilon}_t \zeta_t] = E[S_t^+ \hat{\epsilon}_t | \hat{\epsilon}_t] E[\zeta_t | \hat{\epsilon}_t]$ and $E[\zeta_t \mid \hat{\epsilon}_t] = 0$.

Because $S_t^+ S_t^- = 0$, the moment (24) reduces to $E[Z_t^+ Y_t^{(h)}] = m_{h,h}^+ E[Z_t^+ X_t^+(h)]$, hence

$$m_{h,h}^+ = \frac{C_M^+(h)\, E[S_t^+ \hat{\epsilon}_t^2]}{C_{\Delta U}^+(h)\, E[S_t^+ \hat{\epsilon}_t^2]} = \frac{C_M^+(h)}{C_{\Delta U}^+(h)} = \mu_h^+. \tag{28}$$

□

D - Additional tables and figures

Variable	Estimate	t-stat
$I_t^+ \times \Delta y_t$	0.00008	0.01
$I_t^- \times \Delta y_t$	**0.0255**	**2.82**
M_{t-1}	**1.0324**	**22.26**
M_{t-2}	**-0.2619**	**-3.93**
M_{t-3}	0.0168	0.25
M_{t-4}	**0.1160**	**2.53**
Constant	0.0076	1.56
Observations	470	
R-squared	0.84	
F-statistic	418	

Table D.1: Regression of News Sentiment on its first four lags and current positive/negative changes in industrial production

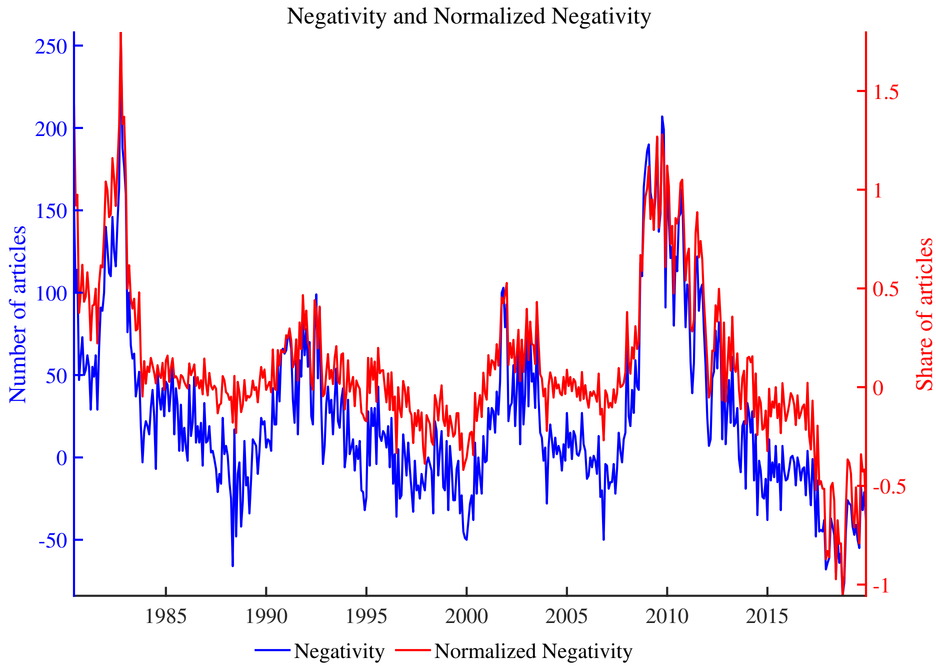

Figure D.1: Negativity and Normalized Negativity.

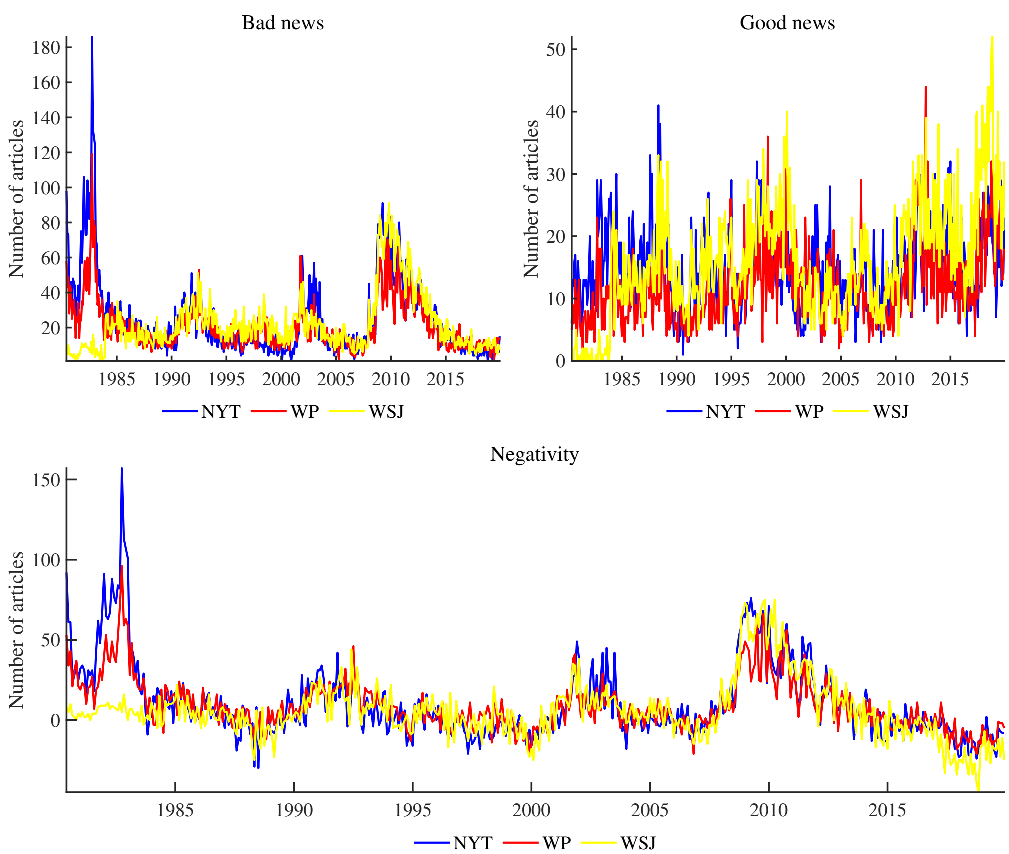

Figure D.2: Bad news, good news and negativity by newspaper.

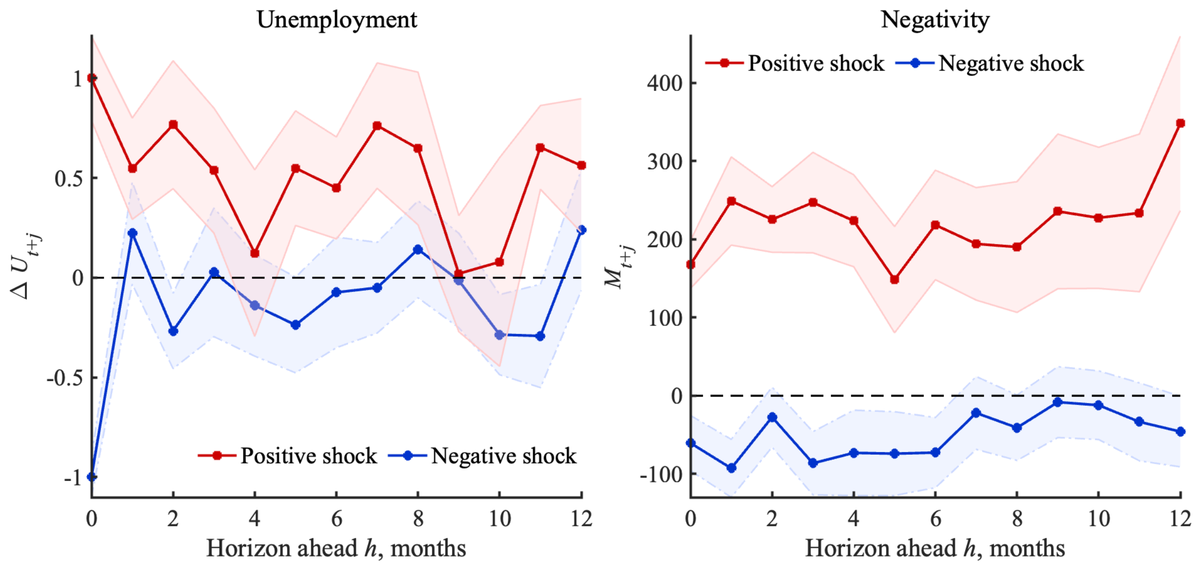

Figure D.3: Estimated impulse responses for positive and negative unemployment shocks

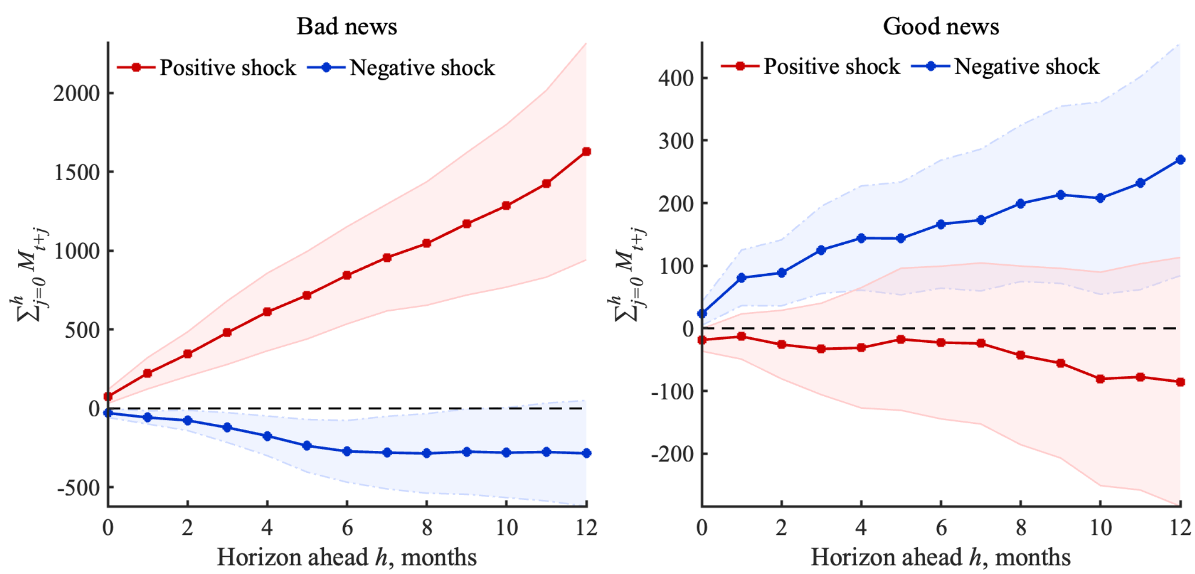

Figure D.4: Estimated impulse responses for positive and negative unemployment shocks

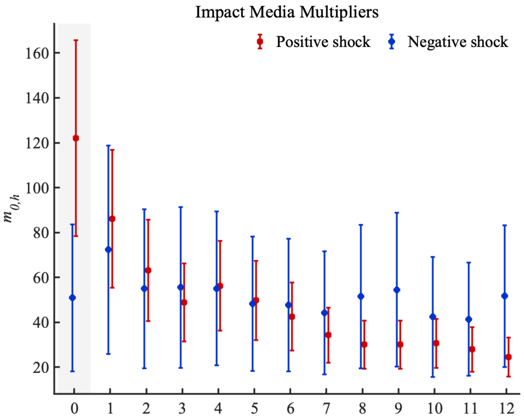

(a) 8 lags

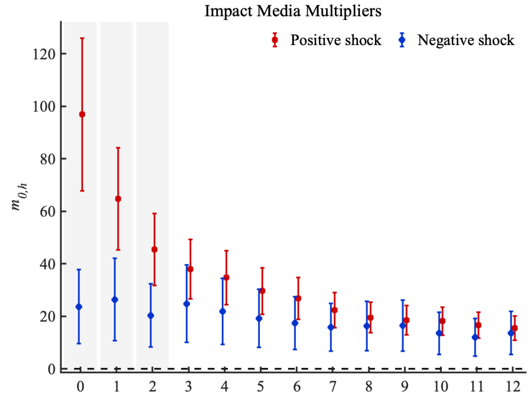

(b) New York Times

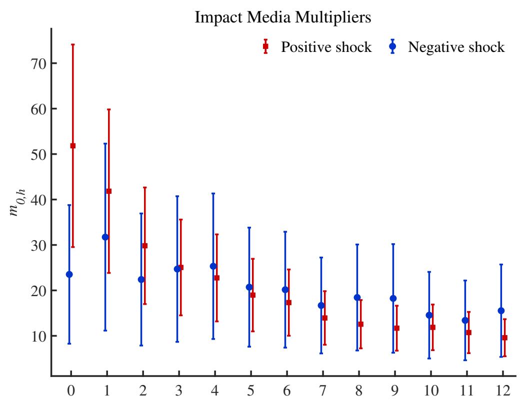

(c) Washington Post

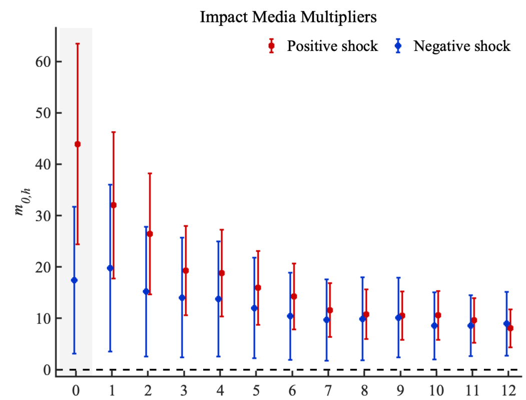

(d) Wall Street Journal

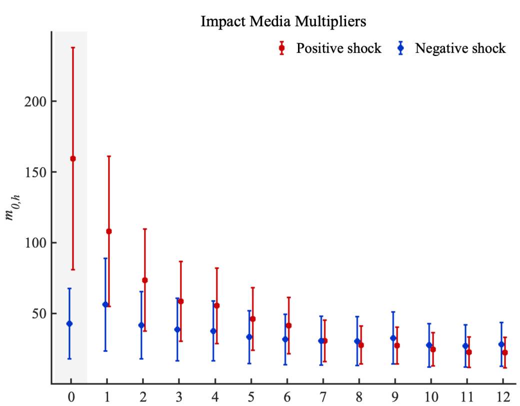

(e) TVAR and LP with additional variables

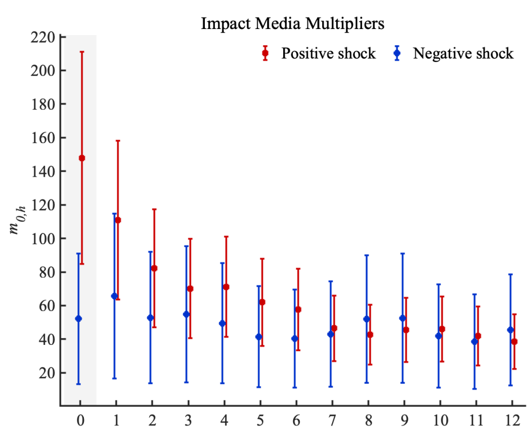

(f) LP only with additional variables

Figure D.5: Impact media multipliers under different specifications

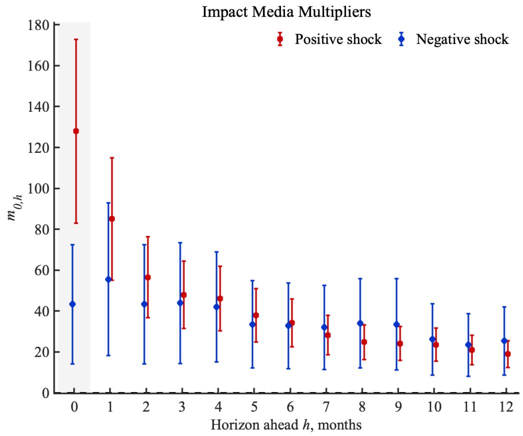

(a) Including ambiguous articles

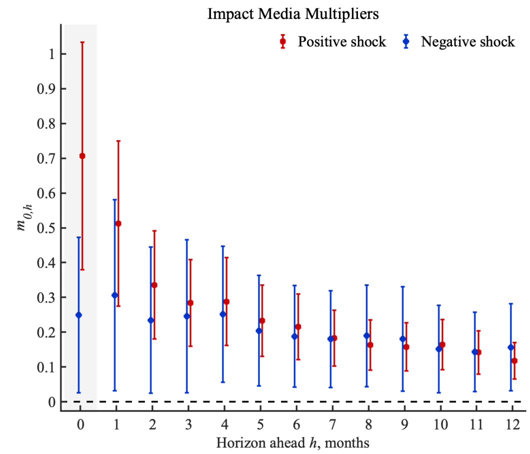

(b) Normalized Negativity

Figure D.6: Impact media multipliers under different TVAR shock identification

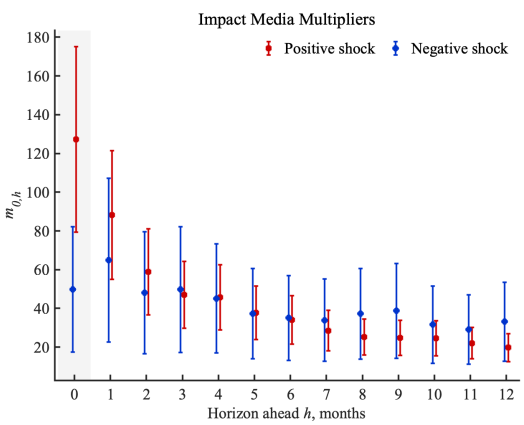

(a) Unemployment shock - wider frequency domain

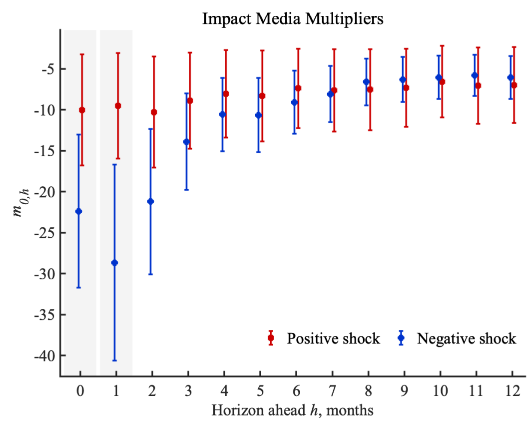

(b) Industrial production shock

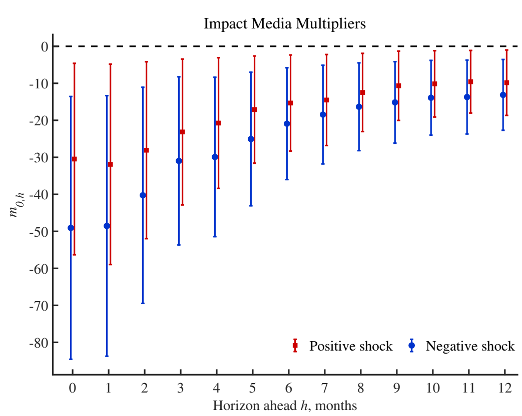

(c) Employment shock

Figure D.7: Impact media multipliers under different TVAR shock identification

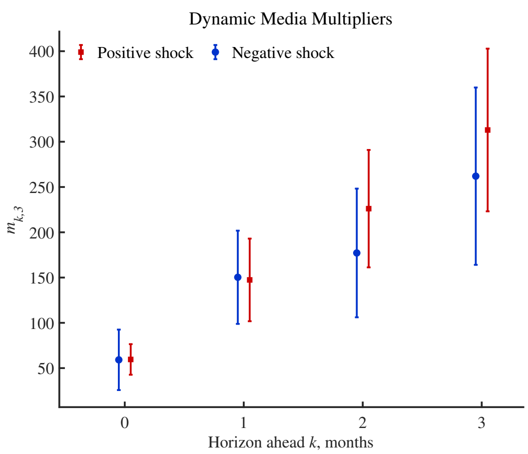

(a) $h = 3$

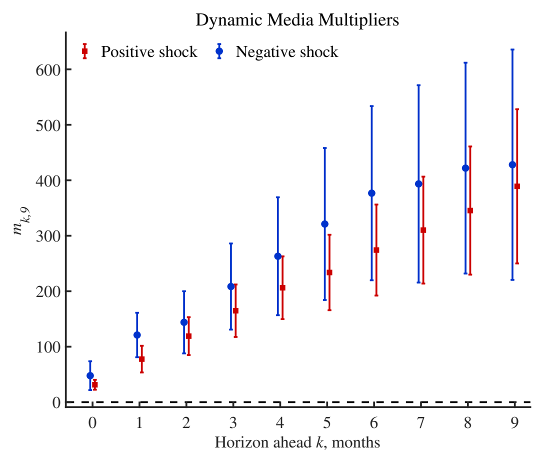

(b) $h = 9$

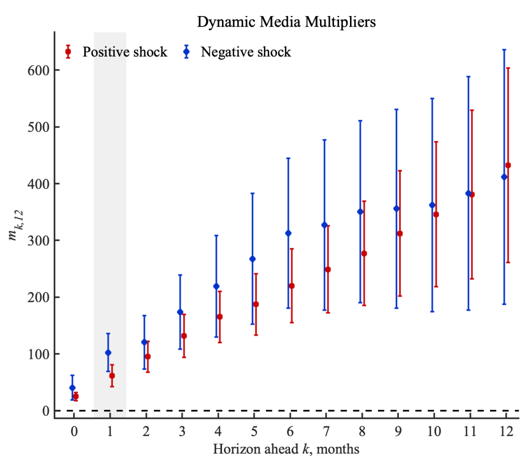

(c) $h = 12$

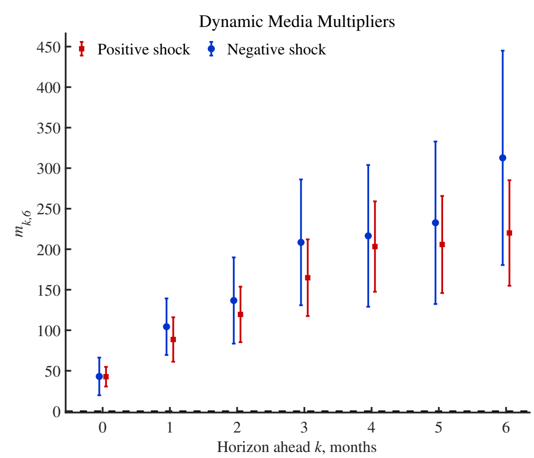

(d) Defining the sum $\sum_{i=0}^{k} \Delta U_{t+h+i}$, with $h = 6$

Figure D.8: Alternative dynamic media multipliers

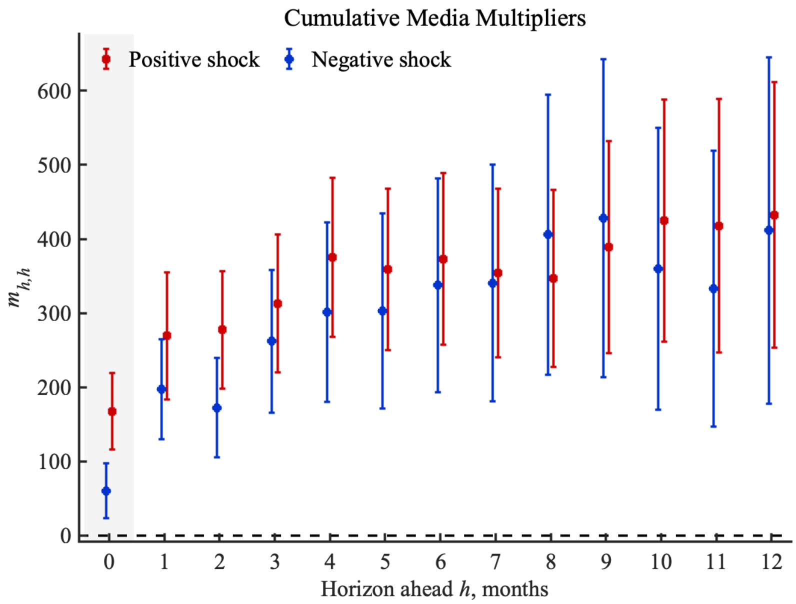

Figure D.9: Cumulative media multipliers for positive and negative unemployment shocks

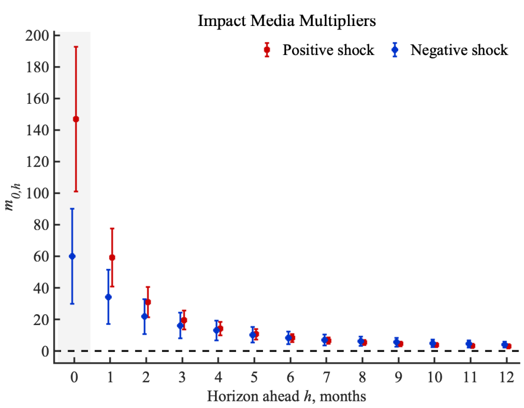

(a) Impact media multiplier

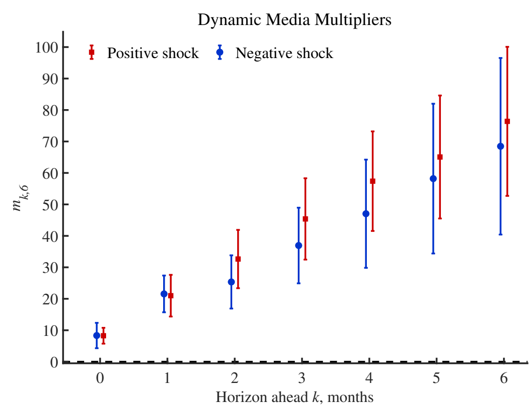

(b) Dynamic media multiplier

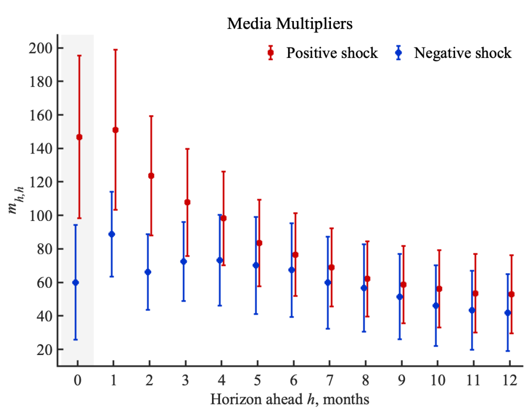

(c) Cumulative media multiplier

Figure D.10: Media multipliers with unemployment in levels

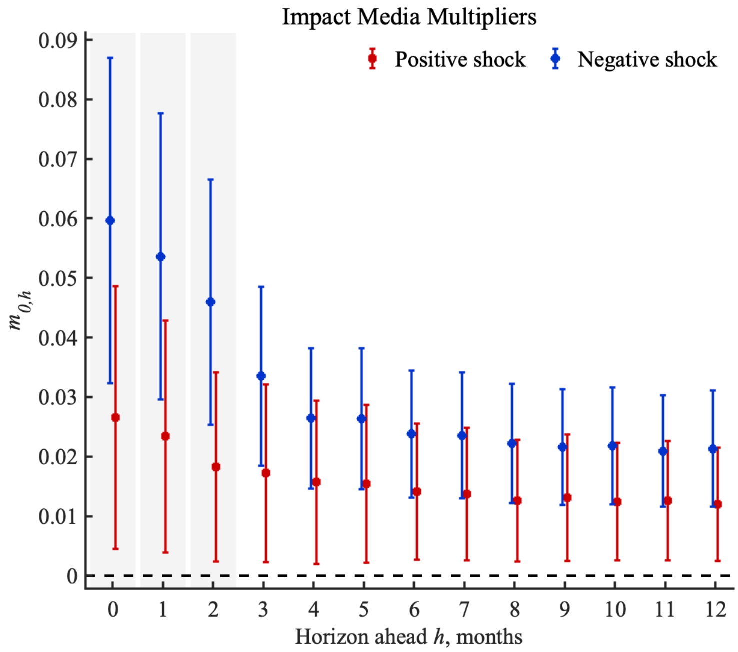

(a) Impact media multiplier

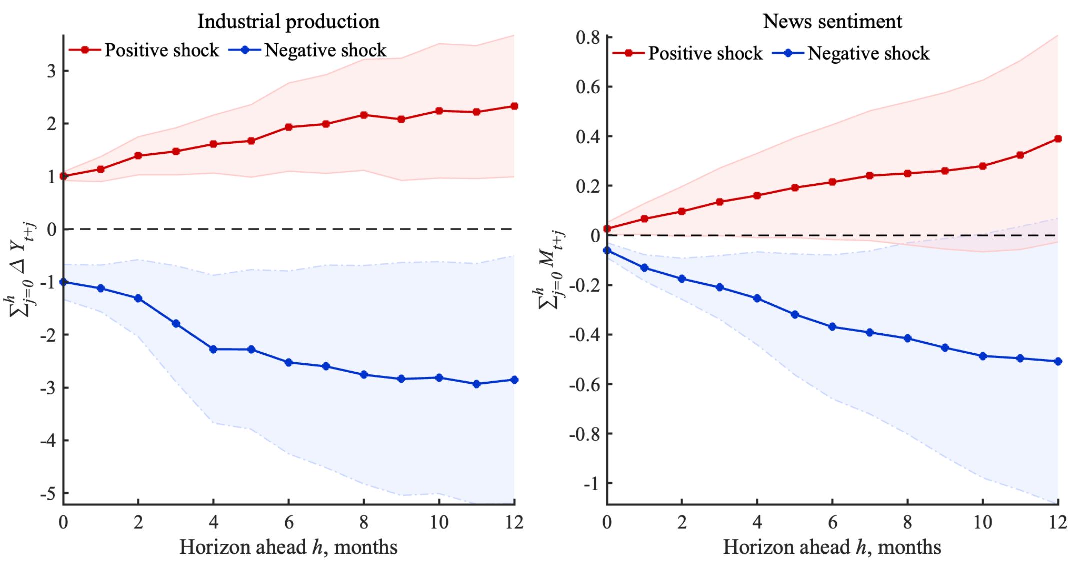

(b) Impulse responses

Figure D.11: San Francisco FED News sentiment