Macroeconomic news: A literature survey and methodological guidelines

Marcel Garz, Jönköping University[1]

November 2023

Abstract

This article surveys the economic literature on the role of mass media and social media for the dissemination of news about the economy. Focusing on studies that use media content data, two key insights emerge. First, identical macroeconomic facts may receive different levels of attention and can be framed in different ways, depending on political and psychological factors. Second, information spread by mass media and on social media may affect economic outcomes independent of the facts on the ground. These and other findings in the literature have helped to refine various theories in behavioral economics, finance, macroeconomics, and other fields. The article also identifies untapped research potential and formulates specific recommendations for future studies, especially in terms of underutilized sources of media content data, the application of computational methods (e.g., transformer models, image classification, emotion recognition), and econometric designs supporting causal inference.

Keywords: computational linguistics; computer vision; inflation; media; GDP; unemployment

JEL classifications: C1; D8; D9; E00; G00; L82

[1] Marcel Garz, Jönköping University, P.O. Box 1026, SE-55111 Jönköping, Sweden. Email: marcel.garz@ju.se.
Declarations of interest: none

1. Introduction

Signals available in everyday life may help individuals to learn about some but not all economic variables. For example, prices observed while shopping for groceries could allow individuals to obtain a rough idea about overall price trends, and the length of the queue at the local job center may be a crude indicator of the unemployment rate. However, for the most part, it is not feasible for individuals to compile information about macroeconomic variables. Collecting the data necessary to accurately measure a country's economic activity is expensive and skill-intensive, which is why most governments maintain statistics agencies that gather, aggregate, and communicate information about macroeconomic data to the public. While these agencies have their owns means of communication, such as press conferences, websites, and social media accounts, most people receive news about the macroeconomy from mass media (Blinder and Krueger, 2004).

The study of media coverage of the economy was not initiated in economics, but the topic caught interest in neighboring disciplines first: In political science, macroeconomic news drew scholars' attention because of its relevance for voting (e.g., Goidel and Langley, 1995; Hetherington, 1996), while researchers in communications wanted to understand how journalists and editors select, process, interpret, and present macroeconomic data in the news (e.g., Harrington, 1989; Blood and Phillips, 1997).

The first papers in economics analyzing observational media data on the topic were published around 2010 (e.g., Alsem et al., 2008; Hollanders and Vliegenthart, 2011; Larcinese et al., 2011). These papers and subsequent research investigate links to many variables that economists find important, such as asset prices, consumer confidence, inflation expectations, and macroeconomic fundamentals. While this research applies various perspectives, methods, and data sources, it essentially deals with three groups of questions: i) how news media and, more recently, social media transmit information about the economy to households, investors, and firms; ii) how mediated information affects economic activity; and iii) how measures of media coverage and social media can be used to improve macroeconomic forecasting. As discussed throughout this article, the answers to these questions have had a substantial impact on policymaking and economic theory, for instance, regarding endogenous attention in financial markets, investor sentiment, and expectations formation.

Previous surveys of the media-related literature focus on the political economy of media (DellaVigna and Gentzkow, 2010; Prat and Strömberg, 2013) and media economics at large (Anderson et al., 2015). A survey related to macroeconomic news is missing so far, especially one that covers recent findings pertaining to social media as well as methodological advances for the analysis of texts and images. For these reasons, this article reviews the economic literature on macroeconomic news. While this article uses a structured approach and clear criteria for the inclusion of papers, it is not the goal to systematically summarize the evidence on any specific research question. Instead, this article surveys the literature to discuss the advantages and disadvantages of relevant research practices, identifies untapped research potential, and provides methodological guidelines.[2]

The next section defines the scope of this survey and its boundaries. With a focus on quantitative empirical studies working with content data from mass media and social media, the section describes trends in research topics and discusses the relationship of relevant theories to each other. Section 3 examines commonly used data sources and corresponding challenges regarding data access, replication, and legal/ethical issues. Section 4 proposes a typology of media measures based on the categories volume, topics, and tone; it debates methodological options to create these measures, including keyword-based classification, content analysis by human coders, natural language processing, and machine learning; and the section provides recommendations for the application of novel tools from computational linguistics (e.g., transformer models) and computer vision (e.g., imagine classification, emotion recognition). Section 5 examines econometric approaches used in the surveyed literature, emphasizing the tension between theory of endogenous news coverage and time-series analysis. This section provides practical suggestions for empirical designs that can be used to disentangle cause and effect in research on macroeconomic news. Finally, Section 6 concludes with a summary of the main findings and key recommendations for future research.

2. Research topics and relevant theories

The survey covers papers published in peer-reviewed, indexed journals in economics that analyze observational data on (social) media content related to the economy at large or any of the "big three" macroeconomic indicators, GDP, inflation, and unemployment.[3] Figure 1 shows that the publication of these studies follows an upward trend over time, with the earliest contribution published in 2008 and a concentration of research in the past six years. Table 1 categorizes the assessed studies by their main topic of interest, the type of media (e.g., newspapers, social media) and content (e.g., unemployment, inflation) under investigation, the main source of the content data, the type of analysis used to create which media measure, the main econometric approach, and whether the research design offers a causal interpretation.

[2] See Grant and Booth (2009) for a discussion of the different purposes of literature surveys vs. systematic reviews.
[3] Specifically, Scopus and Web of Science were used to obtain a broad list of potentially relevant studies, by searching for English-language articles that include the following combination of keywords in their title or abstract: (economy OR gpd OR inflation OR unemployment) AND (media OR news OR twitter OR facebook). This search retrieved 1,647 articles (after removing 392 duplicates), which were screened for inclusion by reading the abstracts and – in case of doubt – the data and methods sections. Studies focused on economic policy, central bank communication, and financial news about individual companies are not included. Studies using web data (e.g., search traffic or web traffic), survey data on media use, or data on media access/freedom – but no data on media content – are not included. In addition, studies focusing on indices of economic policy uncertainty (e.g., Baker et al., 2016) are also excluded because macroeconomic news is just one of many elements used to create these indices but not of central interest.

Considering the main topic of interest, most research can be categorized into studies that investigate the role of macroeconomic news for i) asset prices, ii) consumer confidence, iii) forecasting/nowcasting of macroeconomic variables, and iv) inflation expectations. These studies usually treat macroeconomic news as a "right-hand side variable" that is used to explain the outcome variable of interest.

Regarding category i), studies have examined a variety of assets, including stock prices (e.g., Reed, 2016), equity premia (Adämmer and Schüssler, 2020), risk premia (Fisher et al., 2022), and oil prices (Brandt and Gao, 2019). This strand of research usually motivates its focus on media with financial theories of endogenous attention and theories of investor sentiment, where news affects the chances that investors exert costly effort to learn about fundamentals and develop beliefs about future developments that are not exclusively based on facts on the ground.

Research focused on consumer confidence usually relies on survey data to study the relationship between media and consumer perceptions of macroeconomic developments. Often, these studies use agenda-setting theory to justify why and how news coverage affects perceptions (e.g., Hollanders and Vliegenthart, 2011; Garz, 2018). Accordingly, an important mechanism of whether consumers update their perceptions is the degree of salience an issue or topic receives in the media.

A central question in the fore- and nowcasting literature is whether media coverage or social media content contain information that is not included in macroeconomic variables, and how this information can be best used to improve the predictive performance of empirical models (e.g., Ardia et al., 2019; Thorsrud, 2020; Barbaglia et al., 2023). Consequently, studies focused on fore- and nowcasting are relatively agnostic about possible theoretical mechanisms driving any observed relationships. These studies usually do not attempt to offer a causal interpretation of their results either. The primary goal is to improve predictions to guide decision and policymaking.

The research focused on inflation expectations has been informed by and has contributed to theoretical models of expectation formation, where (a lack of) media coverage plays a crucial role in explaining why households have partial, rigid, or sticky information (e.g., Dräger, 2014; Lamla and Lein, 2014; Lei et al., 2015; Larsen et al., 2021). An important finding in this strand of literature is that volume, topics, and tone of news coverage all matter for the formation of expectations, highlighting the necessity to create measures of macroeconomic news that capture both quantitative and qualitative characteristics of media coverage.

In addition, several studies listed in Table 1 consider news as a "left-hand side variable", investigating how economic, political, and other factors shape media coverage of the economy. This literature is dominated by two themes. First, the political landscape – especially the party affiliation of the government – affects to what extent and how the media covers news about the economy (Larcinese et al., 2011; Lott and Hassett, 2014). The causes of partisan bias in news coverage can theoretically be found in the demand side (i.e., ideological preferences of news consumers) and the supply side of the market (i.e., ideological preferences of journalists, editors, and media owners). However, it remains challenging to empirically disentangle these explanations. The second theme deals with asymmetric responses of news outlets to good and bad developments. For instance, an increase in the unemployment rate by one percentage point usually receives more media attention than a one-point decrease, which is in line with prospect theory and negativity bias (Casey and Owen, 2013; Garz, 2014).

Figure 2 summarizes the main theories applied in and refined by the surveyed literature. As the figure indicates, these theories have diverse disciplinary backgrounds, ranging from communications (agenda-setting theory) over finance (endogenous attention, investor sentiment) to macroeconomics (expectations formation) and behavioral economics (negativity bias). The theories have common elements, despite their different origins and purposes. Agenda-setting theory (e.g., McCombs and Shaw, 1972), sticky information theory (e.g., Mankiw and Reis, 2002), rational inattention theory (e.g., Sims, 2003), and theories of endogenous attention in finance (e.g., Kacperczyk et al., 2016) all postulate that audiences – be it consumers or investors – are more likely to update their information, the greater the salience of a topic. Similarly, theories of negativity bias (e.g., Kahneman and Tversky, 1979; Rozin and Royzman, 2001) and theories of investor sentiment (e.g., Baker and Wurgler, 2008) postulate that the content of the message determines how audiences update their information. Both postulations are directly linked to the most important measures of macroeconomic news in the media: the volume of content captures *how much* attention is devoted to an issue, the topic captures *what* is brought to the agenda, and the tone measures *how* the information is portrayed. In addition, all theories listed in Figure 2 implicitly acknowledge the possibility that (social) media content has purely psychological consequences, i.e., effects not strictly based on the facts on the ground.

3. Media types and data sources

Most of the surveyed research investigates media coverage in newspapers (26 studies). Others examine television (Xu et al., 2022), online news sites (Jiang and Hong, 2021), social media (Reed, 2016; Angelico et al., 2022; Carnazza, 2023), or a mix of these (15 papers; cp. Table 1). Researchers typically assume that their selection of media outlets allows for the creation of proxies that are representative of (part of) the media landscape in a given context. There may not be strong theoretical arguments in favor of one or another media type, but researchers usually select the outlets in their sample based on the availability of content data. These data are from digital news archives (28 papers), content analyses providers (8 papers), public online resources (6 papers), and analytics service providers (4 papers).

Prominent examples of digital news archives are LexisNexis and Dow Jones Factiva, which offer access to content from media outlets in numerous countries. There are also multiple examples of studies using archives located and specialized in the country under investigation. Digital news archives can be a resource of rich data, usually in the form of raw article text and sometimes images, based on which researchers can create measures for their econometrics analyses (i.e., volume, topics, tone). Creating these measures can be challenging though. First, most archives operate on a commercial basis and require a license or subscription. This requirement excludes researchers that do not have institutional access and cannot afford the (often prohibitively expensive) fees for individual access. The restricted access to proprietary content data also poses challenges for the replication of results. Second, digital news archives typically provide the researcher with a browser-based search interface, which can be convenient but creates problems if the interface imposes restrictions on how much content can be downloaded. For instance, depending on the type of license, LexisNexis and Factiva limit the download of articles to a few hundred at a time, which makes comprehensive investigations unfeasible. Third, some archives automatically classify content by using proprietary, non-transparent algorithms. For example, Factiva offers tags that are indicative of the topic of an article (e.g., "economics" or "inflation") or referenced locations (e.g., "Switzerland" or "New York"). However, it may remain unclear how exactly these tags are assigned and how accurate and reliable they are. In addition, the provider may change the process at any time, which threatens the replicability of findings that are based on those classifications.

Several studies analyze data obtained from content analyses providers. A prominent example is Media Tenor International, whose data are used by Dräger (2014), Tausch and Zumbuehl (2018), Lamla et al. (2020), and others. Content analyses providers employ human coders that analyze media coverage and social media content according to pre-defined codebooks and based on the standards of quantitative content analysis (e.g., Krippendorff, 2004). The advantage is that the resulting data can be analyzed without much processing. In addition, human coders might be better able to catch certain issues than machine-based approaches, especially when language subtleties matter or when the level of complexity is high. The approach can be research-efficient, as the content analysis is conducted only once – including recruitment and training of coders – while the resulting data can be used in multiple studies. However, the standardized process limits the range of research questions for which the data hold value. The researcher also needs to trust that the quality standards of content analyses are met (e.g., intercoder reliability). Unless content analyses providers share their datasets as part of a research collaboration, they need to be purchased. In either case, the publication of data for replication is usually not permitted.

The discussion of advantages and disadvantages is similar for media data obtained from analytics service providers, such as Thomson Reuters News Analytics (Calomiris and Mamaysky, 2019), Thomson Reuters

MarketPsych (Alomari et al., 2021), and RavenPack News Analytics (Brandt and Gao, 2019). Those providers may offer access to raw media data, including article texts, social media posts, or television transcripts. However, their focus of service is on easy-to-use machine-learning and natural language processing tools. Researchers can use these tools in a time-efficient manner to create media metrics for their analyses, such as sentiment scores and emotion indicators. However, data access again requires a license or subscription, the data may not be shared for replication, and the automated procedures used to create those metrics often remain non-transparent.

Finally, several studies rely on media content collected directly from websites and other public web resources (e.g., Bortoli et al., 2018; Goshima et al., 2021; Carnazza et al., 2023). Some researchers manually compile content into a database (e.g., Reed, 2016), while others use tools that automatically crawl, scrape, and parse information on websites (e.g, Jiang and Hong, 2021). With the latter approach, researchers potentially have access to vast amounts of information, facilitating the construction of large-scale media measures in an efficient and flexible manner. However, the terms and conditions of many websites prohibit screen scraping, posing challenges in obtaining research ethics approval and potentially leading to legal issues. Some journals might not be willing to publish research that uses scraped data, depending on editorial policies and the applicable jurisdiction. At the time of writing this survey, the issue is subject to ongoing discussions in academia and among legal experts; see Luscombe et al. (2022), AEA (2023), and Krotov and Johnson (2023).

The uncertain implications of data scraping can be avoided by exploring other options of accessing media data online. In the context of macroeconomic news, an underutilized approach is to collect data via application programming interfaces (APIs) that many news outlets and social media maintain. Examples include the New York Times API, Meta's Content Library API for access to historical data from Facebook and Instagram, as well as Twitter's API. While there is a certain risk associated with providers reconfiguring or disabling their interface, as illustrated by the uncertainty surrounding Twitter under Elon Musk, it remains preferable to use APIs instead of screen scraping when collecting media data online.

Another alternative to obtain content data on macroeconomic news are open-source media collections, such as GDELT (Leetaru and Schrodt, 2013) and Mediacloud (Roberts et al., 2021). These collections have not been used in the surveyed studies but examples from other areas in economics illustrate their value for applied research (Campante and Yanagizawa-Drott, 2018; Manacorda and Tesei, 2020). Specifically, the access to information on billions of news stories published by hundreds of thousands of news sites all over the world and over long periods makes it possible to study macroeconomic news on a large scale. Additionally, those collections support investigations in contexts that could not be studied otherwise due to the lack of data, such as developing countries whose media are not covered by traditional news archives. GDELT, Mediacould, and similar providers maintain browser-based search interfaces that support the retrieval of data without any technical skills. On the downside, the documentation of these collections may lack transparency, and replication issues could arise if the provider decides to make changes to their data curation processes. In that sense, Mediacould is perhaps the more reliable choice, as the database is maintained by a consortium of academic institutions.

4. Processing and analysis of media content

This section categorizes the media variables employed in the surveyed studies into measures of volume, topics, and tone, based on which the advantages and disadvantages of methods used to construct them are discussed. In addition, this section examines the application of tools from computation linguistics and computer vision, which have been neglected in the literature but have the potential to create richer, more nuanced measures of macroeconomic news.

4.1 Volume

In its most simple form, the volume of news can be measured by counting news items, such as the number of reports, words, social media posts, or occurrences of relevant keywords. These measures offer a straightforward approach to capturing how much media attention an issue receives, or how salient the issue is. However, a challenge that is particularly relevant when using data from news archives and open-source media collections relates to content duplication. In some cases, duplicates appear for technical reasons, in other cases they exist because media outlets owned by the company share parts of their content. This issue deserves attention as the wrong handling of duplicates can introduce severe measurement error in volume indices.

Measurement error may also be caused by archiving inconsistencies. For commercial and technical reasons, many archives and media collections do not cover all outlets over the same period. For example, if Newspaper A is archived during 1998–2023, while Newspaper B is archived during 2001–2005 and 2008–2023, simply aggregating counts of relevant news items yields a distorted picture of the true reporting. Unfortunately, some archives and collections do not provide lists of included outlets and covered periods, in which case it is advisable to normalize counts of interest by the overall number of archived items. If the archived outlets and periods are documented, a normalization of counts can still be useful because the overall amount of news varies over the course of the year due to seasonality in advertising and reader demand. However, in that case the researcher may also want to create a hand-curated list of outlets that are consistently archived during the investigation period.

4.2 Topics

Several studies listed in Table 1 use Latent Dirichlet Allocation (LDA) to model the topics addressed in media coverage or social media posts (Adämmer and Schüssler, 2020; Thorsrud, 2020; Larsen et al., 2021; Angelico et al., 2022). LDA is a popular method to classify or filter media content. It relies on a bag-of-words approach, where documents are stored as word vectors that count the occurrence of each word but ignore their original order in the document and their relationship to each other (Gentzkow et al., 2019). LDA retrieves latent topics by identifying those words that often appear jointly in the same document. For each word, the probability of belonging to a topic is calculated. Researchers may then rank words according to this probability and come up with labels for different topics based on the top-ranked words. The approach also allows for the computation of topic weights that may vary over time and across units. These weights can be used to test which topics are particularly relevant for a given outcome variable. LDA can be conveniently implemented via open-source packages available for Stata, R, Python, and other common applications.

LDA requires the researcher to specify the number of topics, which can be difficult if there is no theory guiding this choice, as metrics used to evaluate the fit of different candidate models (e.g., perplexity score, semantic coherence, exclusivity) often contradict each other (e.g., Roberts, 2019). This challenge can be avoided by implementing alternative approaches of identifying topics in macroeconomic news. For instance, Fulop and Kocsis (2023) use regular expressions – a pattern-based technique to find matches of character combinations – to categorize news items into different topics.

4.3 Tone

The tone of (social) media content, sometimes also called tonality or valence, refers to the way an issue, event, or development is described. A few of the surveyed papers use keywords to capture the tone of macroeconomic news (e.g., Casey and Owen, 2013; Lott and Hassett, 2014; Xu et al., 2018). For instance, the keyword "recovery" could reflect good news, whereas the term "recession" could be indicative of bad news. Others categorize media stories into good and bad news based on content analyses by human coders (e.g., Alsem, 2008; Lei et al., 2015), which is arguably more accurate but also more costly.

Media tone is a concept that is closely related to media sentiment. A few studies use readily available measures of media sentiment obtained from analytics service providers (e.g., Alomari et al., 2021; Brandt and Gao, 2019). As discussed before, this approach is time-efficient, and it can be advantageous to use metrics relying on data as rich as those available to analytics service providers, but the access restrictions and lack of transparency of procedures used to create these metrics are not ideal.

Other authors create their own measures of sentiment by using dictionaries of positively and negatively connotated words. The procedure of checking for occurrences of words included in a dictionary is a natural-language processing technique that could be considered a sophisticated version of the manual, keyword-based approach of measuring tone. The methodological principle is the same, but the former approach is perhaps less arbitrary. Several studies listed in Table 1 create measures of sentiment of macroeconomic news by applying the bag-of-words approach and existing sentiment dictionaries (e.g., Ardia et al., 2019; Bannigidadmath and Narayan; 2021; Jiang and Hong, 2021). Some of these dictionaries are built for the general use of language, such as the Harvard IV dictionary (Stone, 1997) or the NRC Sentiment Lexicon (Mohammad & Turney, 2013). Others aim to take the context into account, such as the Loughran and McDonald (2011) financial dictionary.

However, existing dictionaries might still fail to accurately capture sentiment in a narrow research context, which is why some researchers create their own dictionaries. For instance, Bortoli et al. (2018) and Aprigliano et al. (2023) extract relevant terms from newspaper articles about the economy. Human coders then evaluate these terms in light of their implied sentiment. Once the annotations are stored in a dictionary, they can be used to quantify the sentiment of individual newspaper articles via the bag-of-words approach.

Consider the news headline “Economic recovery not in sight”. With the approaches mentioned above, the sentence would likely be rated as positive, due to the term “recovery”, which is listed as a positively connotated word in many dictionaries. The example illustrates that plain, dictionary-based sentiment analysis is prone to measurement error. For that reason, some studies implement more sophisticated approaches that account for negations (e.g., “not”) and other valence shifters in proximity to positively and negatively connotated words (Algaba et al., 2023; Aprigliano et al., 2023).

An alternative to dictionary-based analysis is sentiment classification based on supervised learning (Bowden et al., 2019; Rambaccussing and Kwiatkowski, 2020; Goshima et al., 2021). Human coders annotate a small subsample of a media dataset, for instance, whether a news headline or social media post is predominantly positive, negative, or neutral. These annotations are then used to train a machine-learning classifier, based on which ratings of sentiment can be obtained for the entire dataset (Grimmer at al., 2022). This approach has the potential to accurately capture subtle differences in sentiment even when the language is very context specific.

4.4 Computational methods

As Figure 1 shows, there is a clear methodological trend in the literature. The number of publications featuring “traditional analyses”, where media measures are constructed by counting keywords or through content analysis by human coders, has been stagnating. In contrast, there has been an increase in studies using computational approaches, especially tools from natural language processing and machine learning. This trend can be expected to continue, not least because of recent advances in generative artificial intelligence and learning technology.

A particularly promising tool are so-called transformer models, such as BERT (Bidirectional Encoder Representations from Transformers; see Devlin et al., 2018). These models use word embeddings, a technique that retains differential meanings of words when used in different contexts (e.g., “The board of directors will present their forecast” vs. “They will prepare their forecast on board the plane”). Transformer models can be combined with the methods discussed above, including sentiment analysis, topic modeling, and text classification. Transformer models are typically pre-trained on large text corpora that consist of billions of sentences, which is why these models tend to outperform other methods.

While the surveyed literature has not applied transformer models to create text-based measures of media coverage, other research in economics (e.g., Hansen et al., 2023) and the social sciences (Puranam et al., 2021; Chen et al., 2023; Licht, 2023) illustrates the potential of the technology. Online communities and companies specialized in machine-learning have developed open-access Python libraries, based on which researchers can apply transformer models without much programming knowledge. For example, the Transformers library provided by Hugging Face (Wolf et al., 2020) can be used to access numerous models that are pre-trained on diverse datasets. The library also facilitates the finetuning of models for the specific research context, based on which sophisticated measures of volume, topics, and tone can be created.

Another neglected but promising set of technologies for the study of macroeconomic news comes from the area of computer vision. The analysis of images has received limited attention in economics, likely because it is expensive and time-consuming when conducted by human coders. The absence of imagine analysis in macroeconomic news constitutes a major research gap, considering that images are often more influential and effective in conveying information than texts. Examples of applications of computer vision tools in economics include Ash et al. (2022) and Caprini (2023), who use image classification based on deep learning to identify visual stereotypes and partisan visual language, respectively, in media reports.

Much like natural language processing, computer vision tools have hugely benefited from the development of transformer models. These models allow the machine to acquire an understanding of the relationship between image and text. The principle is the same as in the case of text, except that human annotations of images are used to train and finetune these models. Common tasks include image captioning (where the machine produces a sentence describing what can be seen in the image), image classification (where images are categorized according to pre-defined criteria), and emotion recognition (where people’s facial expressions are categorized by the emotions they likely convey). These tools can be implemented via easy-to-use open-source libraries and pre-trained models in Python.

The downside of transformer models is that they are computationally demanding (e.g., processing times tend to increase exponentially with the size of dataset). Hence, in most situations, researchers will likely have to resort to cloud computing. In addition, social scientists may find that those tools do not meet their requirements for rigor, replicability, and transparency. Computational methods often remain opaque, produce results that cannot be exactly replicated due to re-sampling and cross-validation techniques, or allow for error levels that social scientists may find unacceptable. The future will show how the different scholarly traditions in the social sciences and computer sciences can be reconciled.

5. Econometric approaches

The causal arrow between macroeconomic news and real economic conditions is theoretically ambiguous, due to the involvement of and interplay between multiple actors. As the directed acyclic graph in Figure 3 illustrates, media coverage could induce households, investors, and possibly firms to change their behavior (#1), which could lead to measurable changes in macroeconomic indicators (#2). For instance, households may restrict their consumption if they read news stories about increasing unemployment. However, media coverage also responds to changes in economic conditions (#3), as journalists and editors monitor and report about macroeconomic variables (#4). Journalists and editors may also respond to reader demand effects when they create news reports (#5). That is, reach and profits can be maximized by catering to the beliefs and preferences of the audience, as postulated by theories of endogenous news coverage (e.g., Gentzkow and Shapiro, 2006; 2010; Prat and Strömberg, 2013; Qian and Yanagizawa-Drott, 2017). The causal arrow between macroeconomic news coverage and real economic conditions could therefore point in either or both directions. In addition, there could be confounding factors that affect some or all of the above relationships, such as wars, catastrophic events, media capture by advertisers and political actors, or the release of information by governmental bodies. Hence, any empirical correlations between the variables in question are difficult to interpret, due to observational equivalence.

Methods investigating the temporal ordering of variables could be used to shed light on the causal arrow. 10 of the surveyed studies report results from Granger (1969) causality tests, checking whether one time series predicts another. While the limitations of this approach have been extensively discussed in the literature, it is useful to emphasize the role of expectations and anticipatory behavior when studying the relationship between macroeconomic news and real economic conditions. For example, media owners may hire those journalists and editors they expect to cater to the preferences of the audience. In turn, journalists and editors may adjust the volume, topics, and tone in anticipation of events that could be of interest to readers and viewers (e.g., reports about the expected inflation rate prior to the official release by the statistics agency). Consumers may decide to read or ignore the news depending on their expectations of the content, and they may adjust their economic behavior in line with beliefs about future developments. In short, time series of macroeconomic indicators, media coverage, and consumer/investor sentiment are likely "contaminated" with expectations and anticipatory behavior, which violates the assumptions required for the validity of Granger causality (e.g., Freeman, 1983).

An alternative approach to shed light on the causal arrow is to exploit exogenous variation in the variables of interest, as done by two of the surveyed studies: He (2017) investigates the impact of economic news on employment levels in US counties. To minimize the threat of reverse causality, the author isolates those news stories that could affect local conditions but are unlikely to be influenced by local circumstances (i.e., news stories about nationwide economic trends). Garz (2018) uses exogenous variation in unemployment news resulting from competing newsworthy events (i.e., disasters and terrorist attacks) to estimate media effects on household perceptions of the state of the economy.

The wider media literature in economics can provide inspiration for research designs that could be used to establish causal interpretations when studying macroeconomic news:

News pressure. Instrumental variables based on competing events have been particularly popular to isolate exogenous variation in news coverage; see Balles et al. (2023) for a review. Besides disasters and terrorist attacks, researchers have exploited major sports events (e.g., FIFA World Cup) and presidential inaugurations, elections, and other important political events (Durante and Zhuravskaya, 2018). The underlying idea is that those events soak up media attention and crowd out the type of news under investigation. This approach is particularly relevant for measures of the volume of macroeconomic news.

Technical disruptions. A related identification strategy involves variation in information flows due to outages. For instance, Müller and Schwarz (2021) use exogenous variation in social media content resulting from Internet outages to investigate the link between hate speech and hate crime. Variants of this approach could be applied in situations where high-frequency changes in economic conditions matter, and especially when online news and social media are in center of the analysis.

Round-number events. The description of macroeconomic developments in the media relies heavily on numbers and statistics. Situations in which an economic indicator crosses a round number or reaches a milestone often receive a lot of attention, due to left-digit bias. For example, Agarwal et al. (2022) show that attention to the stock market increases when the Dow Jones Index crosses a round number, such as 10,000 points, which in turn increases household demand for mortgages. This kind of identification strategy can be implemented by using a regression-discontinuity design, where the assignment variable is a macroeconomic indicator. Round-number events can be used to isolate variation in the volume and topics of coverage – but also the tone, as (social) media might react asymmetrically to threshold events (i.e., it could be good news if an index climbs past a round number and bad news if it falls below a round number).

Policy interventions. Researchers have previously examined situations where policy changes or other interventions cause changes in news content of some outlets while leaving others unaffected. For example, Piotroski et al. (2017) study state interventions resulting in changes of media ownership that lead to changes in political bias of treated newspapers; Djourelova (2023) uses changes in the language policy of the Associated Press to examine the implications for immigration-related news and public opinion; Garz and Szucs (2023) investigate responses of news outlets to a change in Facebook's content selection algorithm. Those interventions are usually investigated using difference-in-differences techniques. Depending on the policy change, all types of media measures – volume, topics, and tone – could be affected.

Geographical particularities. Snyder and Strömberg (2010) exploit the level of congruence between newspaper markets and US congressional districts to investigate the impact of political news coverage on political accountability. For historical reasons, newspaper markets and congressional districts overlap a lot in some areas, while there is little overlap in others, which affects the incentives of newspapers to cover political issues for reasons that are exogenous to political accountability. Based on a similar idea, Couttenier et al. (2021) construct an instrumental variable capturing the spatial proximity between the location of an event and the headquarters of a media company. While they use this instrument to study effects of media coverage of immigrant crimes on voting, a modified version of this approach could be used to identify effects of macroeconomic news on relevant outcomes.

These examples are representative of research designs that could be readily applied – perhaps with some adaptations – in the context of macroeconomic news. The recommendation to exploit exogenous variation does not come without caveats though. There usually is a trade-off between methodological rigor and the economic significance of the results. Hence, it is advisable to verify that a research design not only supports causal interpretations but also provides insights that are important for theory/practice and reasonably generalizable.

6. Conclusion

The study of news about the economy has had a large impact in many fields in economics, including behavioral economics, finance, and macroeconomics. It has helped to refine theories of agenda-setting, expectations formation, investor sentiment, and others. While the empirical findings from the past 15 years are diverse and context-dependent, several key insights can be formulated: First, macroeconomic information is not always picked up by mass media and social media according to objective criteria. Instead, identical facts may receive different levels of attention and can be framed in different ways, depending on political circumstances and psychological factors. Second, it matters how much, about which issues, and in what manner information about the macroeconomy is disseminated. The surveyed literature clearly demonstrates that news spread by mass media and on social media includes information above and beyond of what can be measured by economic indicators. As a consequence, mediated information has the potential to affect households, investors, and firms regardless of the facts on the ground. A prominent example is negativity bias, which not only causes mass media and social media to respond asymmetrically to good and bad news but also has implications for downstream outcomes. This article closes with a summary of recommendations that scholars might find helpful for future research on the topic:

1) **Sources of media content data.** The surveyed literature has neglected data collection via APIs maintained by news outlets (e.g., New York Times), social media platforms (Facebook, Twitter), and open-source media collections (GDELT, Mediacloud). These APIs offer access to rich content data, based on which researchers can construct measures of macroeconomics news in terms volume, topics, and tone, as discussed in Section 3.
2) **Computational methods.** Research on macroeconomic news is increasingly using tools from natural language processing and machine learning to construct media measures. Future research on the topic could benefit from applying recently developed tools in computational linguistics (e.g., transformer models) and computer vision (e.g., image classification, emotion recognition). Thanks to advances in generative artificial intelligence, these tools has become both very powerful and easy to implement. Their application has the potential to open up new research possibilities when studying macroeconomic news, especially in the domain of image analysis, which has been completely neglected so far. Section 4 provides practical guidelines.
3) **Causal inference.** As discussed throughout this article, the causal arrow between (social) media and real economic conditions is theoretically ambiguous. Investigating the temporal order of variables faces limitations in resolving this ambiguity, as time series of the relevant variables (e.g., macroeconomic indicators, consumer perceptions, media coverage) are likely contaminated with expectations and anticipatory behavior of the relevant actors. It is therefore advisable to use or add designs that exploit exogenous variation. Section 5 provides various examples of those designs.

References

Adämmer, P., & Schüssler, R. A. (2020). Forecasting the Equity Premium: Mind the News!. *Review of Finance*, 24, 1313–1355.

AEA (2023). AEA Data Legality Policy and Explanations. URL: https://www.aeaweb.org/journals/data/data-legality-policy (last accessed: Nov 2, 2023).

Agarwal, S., Correia, S., Morais, B., & Song, C. (2022). Stock Market Milestones and Mortgage Demand: Evidence from US. Working Paper.

Agirdas, C. (2015). What Drives Media Bias? New Evidence from Recent Newspaper Closures. *Journal of Media Economics*, 28, 123–141.

Algaba, A., Borms, S., Boudt, K., & Verbeken, B. (2023). Daily news sentiment and monthly surveys: A mixed-frequency dynamic factor model for nowcasting consumer confidence. *International Journal of Forecasting*, 39, 266–278.

Alomari, M., Al rababa'a, A. R., El-Nader, G., & Alkhataybeh, A. (2021). Who's behind the wheel? The role of social and media news in driving the stock-bond correlation. *Review of Quantitative Finance and Accounting*, 57, 959–1007.

Alsem, K. J., Brakman, S., Hoogduin, L., & Kuper, G. (2008). The impact of newspapers on consumer confidence: Does spin bias exist? *Applied Economics*, 40, 531–539.

Anderson, S. P., Waldfogel, J., & Strömberg, D. (2015). *Handbook of Media Economics*. North Holland.

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

Aprigliano, V., Emiliozzi, S., Guaitoli, G., Luciani, A., Marcucci, J., & Monteforte, L. (2023). The power of text-based indicators in forecasting Italian economic activity. *International Journal of Forecasting*, 39, 791–808.

Ardia, D., Bluteau, K., & Boudt, K. (2019). Questioning the news about economic growth: Sparse forecasting using thousands of news-based sentiment values. *International Journal of Forecasting*, 35, 1370–1386.

Ash, E., Durante, R., Grebenshchikova, M., & Schwarz, C. (2022). Visual Representation and Stereotypes in News Media. SSRN Working Paper.

Ash, E., & Hansen, S. (2023). Text Algorithms in Economics. *Annual Review of Economics*, 15, 659–688.

Baker, M., & Wurgler, J. (2007). Investor Sentiment in the Stock Market. *Journal of Economic Perspectives*, 21, 129–151.

Baker, S. R., Bloom, N., & Davis, S. J. (2016). Measuring Economic Policy Uncertainty. *Quarterly Journal of Economics*, 131, 1593–1636.

Balles, P., Matter, U., & Stutzer, A. (2023). The Political Economics of Attention. In C. Björnskov and R.

Jong-A-Pin (eds): *Elgar Encyclopedia of Public Choice*. Cheltenham: Edward Elgar.

Bannigidadmath, D., & Narayan, P. K. (2021). Economic news and the cross-section of commodity futures returns. *Journal of Behavioral and Experimental Finance*, 31, 100540.

Barbaglia, L., Consoli, S., & Manzan, S. (2023). Forecasting with Economic News. *Journal of Business & Economic Statistics*, 41, 708–719.

Birz, G. (2017). Stale economic news, media and the stock market. *Journal of Economic Psychology*, 61, 87–102.

Birz, G., & Lott, J. R. (2011). The effect of macroeconomic news on stock returns: New evidence from newspaper coverage. *Journal of Banking & Finance*, 35, 2791–2800.

Blinder, A. S., & Krueger, A. B. (2004). What Does the Public Know about Economic Policy, and How Does It Know it? *Brookings Papers on Economic Activity*, 2004, 327–387.

Blood, D. J., & Phillips, P. C. B. (1997). Economic Headline News on the Agenda. In M. McCombs, D. L. Shaw, and D. Weaver: *Communication and Democracy*. Lawrence Earlbaum Associates, 97–113.

Bortoli, C., Combes, S., & Renault, T. (2018). Nowcasting GDP Growth by Reading Newspapers. *Economics and Statistics*, 505-506, 17–33.

Bowden, J., Kwiatkowski, A., & Rambaccussing, D. (2019). Economy through a lens: Distortions of policy coverage in UK national newspapers. *Journal of Comparative Economics*, 47, 881–906.

Brandt, M. W., & Gao, L. (2019). Macro fundamentals or geopolitical events? A textual analysis of news events for crude oil. *Journal of Empirical Finance*, 51, 64–94.

Calomiris, C. W., & Mamaysky, H. (2019). How news and its context drive risk and returns around the world. *Journal of Financial Economics*, 133, 299–336.

Campante, F., & Yanagizawa-Drott, D. (2018). Long-Range Growth: Economic Development in the Global Network of Air Links. *Quarterly Journal of Economics*, 133, 1395–1458.

Caprini, G. (2023). Visual Bias. Working Paper.

Carnazza, G. (2023). The Impact of the Social Mood on the Italian Sovereign Debt Market: A Twitter Perspective. *Italian Economic Journal*.

Casey, G. P., & Owen, A. L. (2013). Good News, Bad News, and Consumer Confidence. *Social Science Quarterly*, 94, 292–315.

Chen, B., Lukito, J., & Koo, G. H. (2023). Comparing the #StopTheSteal Movement across Multiple Platforms: Differentiating Discourse on Facebook, Twitter, and Parler. *Social Media + Society*, 9.

Corbet, S., Larkin, C., Lucey, B. M., Meegan, A., & Yarovaya, L. (2020). The impact of macroeconomic news on Bitcoin returns. *The European Journal of Finance*, 26, 1396–1416.

Couttenier, M., Hatte, S., Thoenig, M., & Vlachos, S. (2021). Anti-Muslim Voting and Media Coverage of Immigrant Crimes. *Review of Economics and Statistics*, forthcoming.

Daniel, V., & ter Steege, L. (2020). Inflation expectations and the recovery from the Great Depression in

Germany. *Explorations in Economic History*, 75, 101305.

DellaVigna, S., & Gentzkow, M. (2010). Persuasion: Empirical Evidence. *Annual Review of Economics*, 2, 643–669.

Devlin, J., Chang, M.-W., Lee, K., & Toutanova, K. (2018). BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. arXiv:1810.04805v2.

Djourelova, M. (2023). Persuasion through Slanted Language: Evidence from the Media Coverage of Immigration. *American Economic Review*, 113, 800–835.

Dräger, L. (2015). Inflation perceptions and expectations in Sweden – Are media reports the missing link? *Oxford Bulletin of Economics and Statistics*, 77, 681–700.

Durante, R., & Zhuravskaya, E. (2018). Attack When the World Is Not Watching? US News and the Israeli-Palestinian Conflict. *Journal of Political Economy*, 126, 1085–1133.

Dutta, S., & Thorson, J. (2019). Are hedge fund returns affected by media coverage of macroeconomic news? *Studies in Economics and Finance*, 36, 427–439.

Fisher, A., Martineau, C., & Sheng, J. (2022). Macroeconomic Attention and Announcement Risk Premia. *The Review of Financial Studies*, 35, 5057–5093.

Freeman, J. R. (1983). Granger Causality and the Times Series Analysis of Political Relationships. *American Journal of Political Science*, 27, 327–358.

Fulop, A., & Kocsis, Z. (2023). News indices on country fundamentals. *Journal of Banking & Finance*, 154, 106951.

Garz, M. (2013). Unemployment expectations, excessive pessimism, and news coverage. *Journal of Economic Psychology*, 34, 156–168.

Garz, M. (2014). Good news and bad news: Evidence of media bias in unemployment reports. *Public Choice*, 161, 499–515.

Garz, M. (2018). Effects of unemployment news on economic perceptions – Evidence from German Federal States. *Regional Science and Urban Economics*, 68, 172–190.

Garz, M., & Szucs, F. (2023). Algorithmic selection and supply of political news on Facebook. *Information Economics and Policy*, 62, 101020.

Gentzkow, M., Kelly, B., Taddy, M. (2019). Text as Data. *Journal of Economic Literature*, 57, 535–574.

Gentzkow, M., & Shapiro, J. M. (2006). Media Bias and Reputation. *Journal of Political Economy*, 114, 280–316.

Gentzkow, M., & Shapiro, J. M. (2010). What Drives Media Slant? Evidence from U.S. Daily Newspapers. *Econometrica*, 78, 35–71.

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, 313–328.

Goshima, K., Ishijima, H., Shintani, M., & Yamamoto, H. (2021). Forecasting Japanese inflation with a news-based leading indicator of economic activities. *Studies in Nonlinear Dynamics & Econometrics*, 25, 111–133.

Granger, C. W. J. (1969). Investigating Causal Relations by Econometric Models and Cross-spectral Methods. *Econometrica*, 37, 424–438.

Grant, M.J., & Booth, A. (2009). A typology of reviews: An analysis of 14 review types and associated methodologies. *Health Information & Libraries Journal*, 26, 91–108.

Grimmer, J., Roberts, M. E., & Stewart, B. M. (2022). *Text as Data: A New Framework for Machine Learning and the Social Sciences*. Princeton University Press.

Hansen S., Lambert P.J., Bloom N., Davis S.J., Sadun, R., & Taska, B. (2023). Remote work across jobs, companies, and space. NBER Working Paper 31007.

Harrington, D. E. (1989). Economic News on Television: The Determinants of Coverage. *Public Opinion Quarterly*, 53, 17–40.

He, Z. (2017). The Contractionary Effect of Bad Economic News. *Journal of Money, Credit and Banking*, 49, 1339–1384.

Hetherington, M. J. (1996). The Media's Role in Forming Voters' National Economic Evaluations in 1992. *American Journal of Political Science*, 40, 372–395.

Hollanders, D., & Vliegenthart, R. (2011). The influence of negative newspaper coverage on consumer confidence: The Dutch case. *Journal of Economic Psychology*, 32, 367–373.

Iselin, D., & Siliverstovs, B. (2016). Using newspapers for tracking the business cycle: A comparative study for Germany and Switzerland. *Applied Economics*, 48, 1103–1118.

Jiang, Y., & Hong, Y. (2021). State Media, Institutional Environment, and Analyst Forecast Quality: Evidence from China. *Emerging Markets Finance and Trade*, 57, 3929–3943.

Kacperczyk, M., Van Nieuwerburgh, S., & Veldkamp, L. (2016). A Rational Theory of Mutual Funds' Attention Allocation. *Econometrica*, 84, 571–626.

Kahneman, D., & Tversky, A. (1979). Prospect Theory: An Analysis of Decision under Risk. *Econometrica*, 47, 263–291.

Krippendorff, K. (2004). *Content Analysis: An Introduction to Its Methodology* (second edition). Sage Publications.

Krotov, V., & Johnson, L. (2023). Big web data: Challenges related to data, technology, legality, and ethics. *Business Horizons*, 66, 481–491.

Lamla, M. J., & Lein, S. M. (2014). The role of media for consumers' inflation expectation formation. *Journal of Economic Behavior & Organization*, 106, 62–77.

Lamla, M. J., & Lein, S. M. (2015). Information Rigidities, Inflation Perceptions, And the Media: Lessons from The Euro Cash Changeover. *Economic Inquiry*, 53, 9–22.

Lamla, M. J., Lein, S. M., & Sturm, J.-E. (2020). Media reporting and business cycles: Empirical evidence based on news data. *Empirical Economics*, 59, 1085–1105.

Lamla, M. J., & Maag, T. (2012). The Role of Media for Inflation Forecast Disagreement of Households and Professional Forecasters. *Journal of Money, Credit and Banking*, 44, 1325–1350.

Larcinese, V., Puglisi, R., & Snyder, J. M. (2011). Partisan bias in economic news: Evidence on the agenda-setting behavior of U.S. newspapers. *Journal of Public Economics*, 95, 1178–1189.

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

Leetaru, K., & Schrodt, P. A. (2013). GDELT: Global Data on Events, Location and Tone, 1979-2012. International Studies Association meetings, San Francisco, April 2013.

Lei, C., Lu, Z., & Zhang, C. (2015). News on inflation and the epidemiology of inflation expectations in China. *Economic Systems*, 39, 644–653.

Licht, H. (2023). Cross-Lingual Classification of Political Texts Using Multilingual Sentence Embeddings. *Political Analysis*, forthcoming.

Lott, J. R., & Hassett, K. A. (2014). Is newspaper coverage of economic events politically biased? *Public Choice*, 160, 65–108.

Loughran, T., & McDonald, B. (2011). When is a liability not a liability? Textual analysis, dictionaries, and 10-Ks. *Journal of Finance*, 66, 35–65.

Luscombe, A., Dick, K., & Walby, K. (2022). Algorithmic thinking in the public interest: Navigating technical, legal, and ethical hurdles to web scraping in the social sciences. *Quality & Quantity*, 56, 1023–1044.

Manacorda, M., & Tesei, A. (2020). Liberation Technology: Mobile Phones and Political Mobilization in Africa. *Econometrica*, 88, 533–567.

Mankiw, N. G., & Reis, R. (2002). Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve. *Quarterly Journal of Economics*, 117, 1295–1328.

McCombs, M., & Shaw, D. L. (1972). The Agenda-Setting Function of Mass Media. *Public Opinion Quarterly*, 36, 176–187.

Mohammad, S. M., & Turney, P. D. (2013). Crowdsourcing a Word-Emotion Association Lexicon. *Computational Intelligence*, 29, 436–465.

Müller, K., & Schwarz, C. (2021). Fanning the Flames of Hate: Social Media and Hate Crime. *Journal of the European Economic Association*, 19, 2131–2167.

Piotroski, J. D., Wong, T. J., & Zhang, T. (2017). Political Bias in Corporate News: The Role of Conglomeration Reform in China. *Journal of Law and Economics*, 60, 173–207.

Prat, A., & Strömberg, D. (2013). The Political Economy of Mass Media. In D. Acemoglu, M. Arellano, and E. Dekel (eds.): *Advances in Economics and Econometrics*. Cambridge University Press, 135–187.

Puranam, D., Kadiyali, V., & Narayan, V. (2021). The Impact of Increase in Minimum Wages on Consumer Perceptions of Service: A Transformer Model of Online Restaurant Reviews. *Marketing Science*, 40, 985–1004.

Qian, N., & Yanagizawa-Drott, D. (2017). Government Distortion in Independently Owned Media: Evidence from U.S. News Coverage of Human Rights. *Journal of the European Economic Association*, 15, 463–499.

Rambaccussing, D., & Kwiatkowski, A. (2020). Forecasting with news sentiment: Evidence with UK newspapers. *International Journal of Forecasting*, 36, 1501–1516.

Reed, M. (2016). A Study of Social Network Effects on the Stock Market. *Journal of Behavioral Finance*, 17, 342–351.

Roberts, H., Bhargava, R., Valiukas, L., Jen, D., Malik, M. M., Bishop, C. S., Ndulue, E. B., Dave, A., Clark, J., Etling, B., Faris, R., Shah, A., Rubinovitz, J., Hope, A., D'Ignazio, C., Bermejo, F., Benkler, Y., & Zuckerman, E. (2021). Media Cloud: Massive Open Source Collection of Global News on the Open Web. Proceedings of the International AAAI Conference on Web and Social Media, 15, 1034–1045.

Roberts M. E., Stewart, B. M., & Tingley, D. (2019). "stm: An R Package for Structural Topic Models." *Journal of Statistical Software*, 91, 1–40.

Rozin, P., & Royzman, E. B. (2001). Negativity Bias, Negativity Dominance, and Contagion. *Personality and Social Psychology Review*, 5, 296–320.

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

Snyder, J. M., & Strömberg, D. (2010). Press Coverage and Political Accountability. *Journal of Political Economy*, 118, 355–408.

Stone, P. J. (1997). Thematic text analysis: New agendas for analyzing text content. In C. W. Roberts (ed.): *Text analysis for the social sciences: Methods for drawing statistical inferences from texts and transcripts*. Routledge, 35–54.

Tausch, F., & Zumbuehl, M. (2018). Stability of risk attitudes and media coverage of economic news. *Journal of Economic Behavior & Organization*, 150, 295–310.

Thorsrud, L. A. (2020). Words are the New Numbers: A Newsy Coincident Index of the Business Cycle. *Journal of Business & Economic Statistics*, 38, 393–409.

Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., von Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Le Scao, T., Gugger, S., Drame, M., Lhoest, Q., & Rush, A. (2020). Transformers: State-of-the-Art Natural Language Processing. In Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations, 38–45. Association for Computational Linguistics.

Xu, Y., Liu, Z., Chen, J., & Salem, S. (2022). How official TV news affect public inflation expectations? Evidence from the Chinese national broadcaster China Central Television. *International Journal of Finance & Economics*, forthcoming.

Xu, Y., Liu, Z., & Ortiz, J. (2018). The relationship between media bias and inflation expectations in P.R. China. *Research in International Business and Finance*, 45, 402–412.

Figures and tables

Figure 1: Surveyed publications over time, by type of media analysis

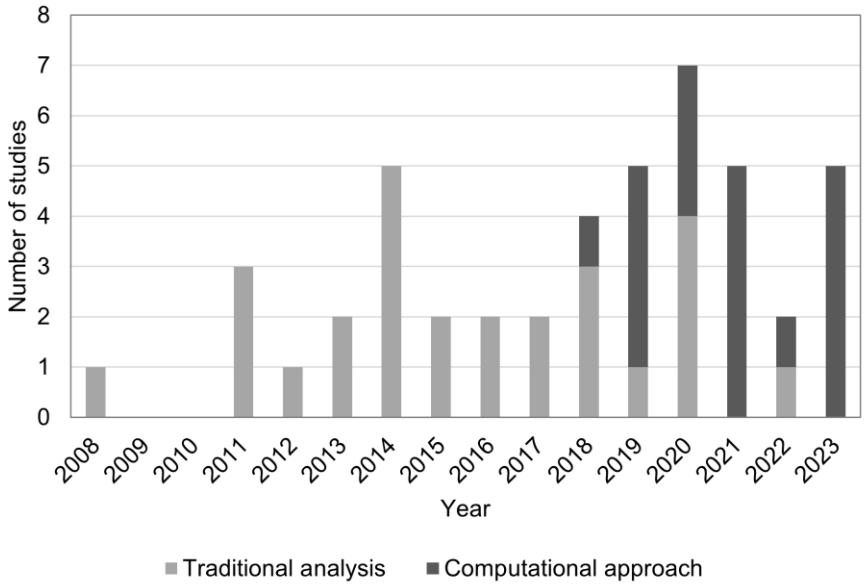

Notes: The figure shows yearly counts of studies meeting the selection criteria described in Section 2. *Traditional analysis* refers to studies using keyword-based measures of (social) media content or measures based on content analysis by human coders. *Computational approach* refers to studies employing tools from natural language processing, machine learning, or both.

Figure 2: Main theories and corresponding empirical concepts

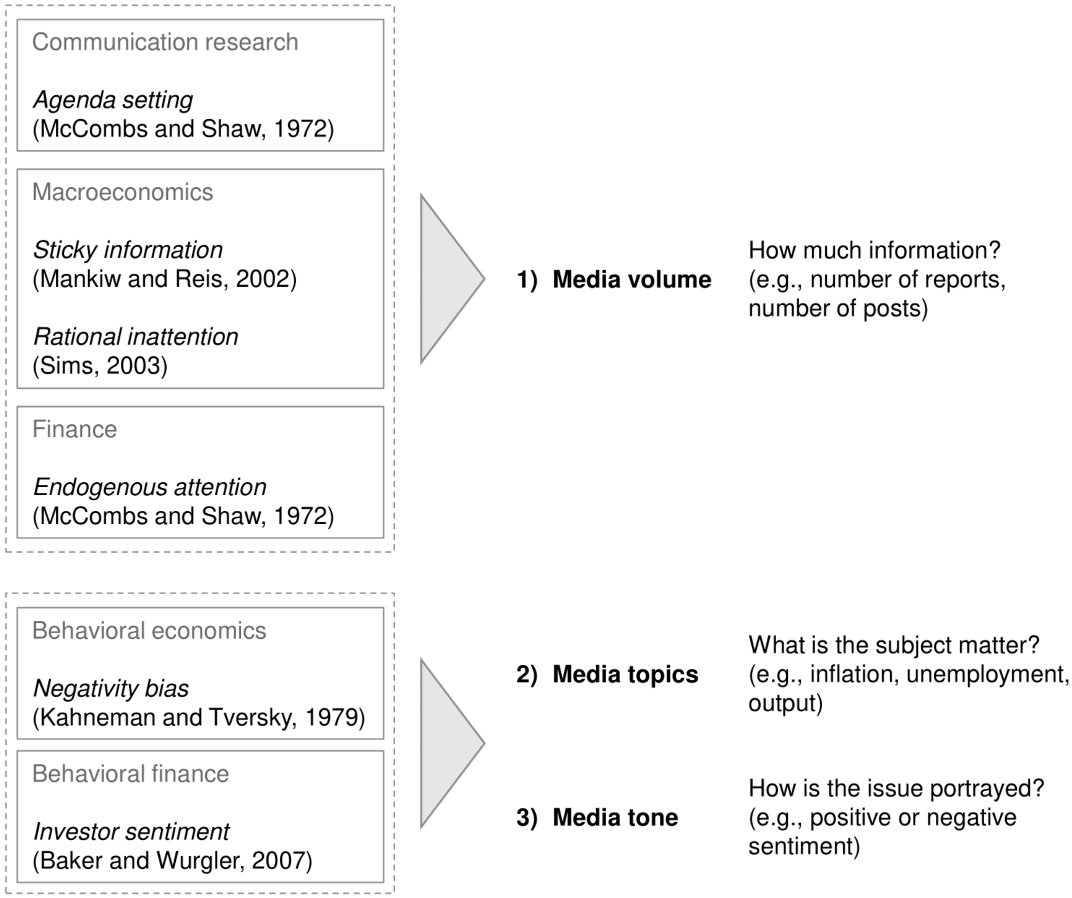

Figure 3: Main links between news and economic conditions

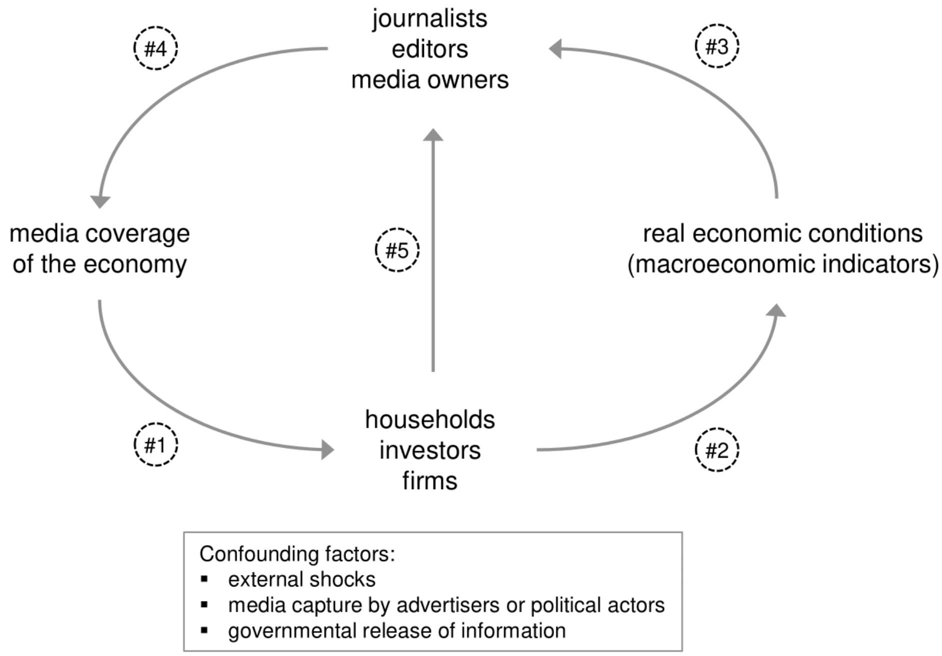

Notes: The figure omits social media for the sake of simplicity.

Table 1: Overview of quantitative studies of macroeconomic news

Reference	Topic	Media type(s)	Content theme(s)	Media data			Main econo-metric method	Causal inference
				Main source	Analysis	Measure(s)		
Adämmer and Schüssler (2020)	asset prices	newspapers	multiple or general	news archive	NLP	topics	time series	no
Agirdas (2015)	determinants of coverage	newspapers	unemployment	news archive	keyword-based counts	volume	panel data	no
Algaba et al. (2023)	consumer confidence	multiple	multiple or general	news archive	NLP	tone	time series	no
Alomari et al. (2021)	asset prices	multiple	multiple or general	analytics service provider	third-party analysis	tone	time series	no
Alsem et al. (2008)	consumer confidence	newspapers	multiple or general	news archive	content analysis (human)	tone	time series	no
Angelico et al. (2022)	inflation expectations	social media	inflation	public data online	NLP + ML	topics, tone	time series	no
Aprigliano et al. (2023)	macroeconomic forecasting	newspapers	multiple or general	news archive	NLP	tone	time series	no
Ardia et al. (2019)	GDP forecasting	newspapers	multiple or general	news archive	NLP	tone	time series	no
Bannigidadmath and Narayan (2021)	asset prices	newspapers	multiple or general	news archive	NLP	tone	time series	no
Barbaglia et al. (2023)	macroeconomic forecasting	multiple	multiple or general	analytics service provider	NLP	tone	time series	no

Reference	Topic	Media type(s)	Content theme(s)	Media data			Main econo-metric method	Causal inference
				Main source	Analysis	Measure(s)		
Birz (2017)	asset prices	newspapers	unemployment	news archive	content analysis (semi-automated)	volume, tone	time series	no
Birz and Lott (2011)	asset prices	newspapers	multiple or general	news archive	content analysis (semi-automated)	volume, tone	time series	no
Bortoli et al. (2018)	GDP forecasting	newspapers	multiple or general	public data online	NLP + ML	tone	time series	no
Bowden et al. (2019)	consumer confidence	newspapers	multiple or general	news archive	NLP + ML	tone	time series	no
Brandt and Gao (2019)	asset prices	multiple	multiple or general	analytics service provider	third-party analysis	tone	panel data	no
Calomiris and Mamaysky (2019)	asset prices	multiple	multiple or general	analytics service provider	third-party analysis	tone	panel data	no
Carnazza (2023)	asset prices	social media	multiple or general	public data online	NLP	tone	time series	Granger
Casey and Owen (2013)	determinants of coverage	newspapers	multiple or general	news archive	keyword-based counts	tone	time series	no
Corbet et al. (2020)	asset prices	newspapers	multiple or general	news archive	content analysis (human)	tone	time series	no
Daniel and ter Steege (2020)	inflation expectations	newspapers	inflation	news archive	keyword-based counts	volume	time series	no

Reference	Topic	Media type(s)	Content theme(s)	Media data			Main econo-metric method	Causal inference
				Main source	Analysis	Measure(s)		
Dräger (2014)	inflation perceptions	multiple	inflation	content analyses provider	content analysis (human)	volume, tone	time series	Granger
Dutta and Thorson (2019)	asset prices	newspapers	unemployment	news archive	content analysis (semi-automated)	volume, tone	time series	no
Fisher et al. (2022)	asset prices	newspapers	multiple or general	news archive	keyword-based counts	volume	time series	no
Fulop and Kocsis (2023)	macroeconomic forecasting	multiple	multiple or general	news archive	NLP + ML	topics, tone	panel data	no
Garz (2013)	unemployment expectations	multiple	multiple or general	content analyses provider	content analysis (human)	volume, tone	time series	no
Garz (2014)	determinants of coverage	multiple	unemployment	content analyses provider	content analysis (human)	volume, tone	panel data	no
Garz (2018)	consumer confidence	newspapers	unemployment	news archive	keyword-based counts	volume	panel data	exogenous variation
Goshima et al. (2021)	inflation forecasting	newspapers	multiple or general	public data online	NLP + ML	tone	time series	no
He (2017)	employment	newspapers	multiple or general	news archive	keyword-based counts	volume	panel data	exogenous variation
Hollanders and Vliegenthart (2011)	consumer confidence	newspapers	multiple or general	news archive	keyword-based counts	volume	time series	Granger

Reference	Topic	Media type(s)	Content theme(s)	Media data			Main econo-metric method	Causal inference
				Main source	Analysis	Measure(s)		
Iselin and Siliverstovs (2016)	GDP forecasting	newspapers	multiple or general	news archive	keyword-based counts	volume	time series	no
Jiang and Hong (2021)	macroeconomic forecasting	online news	multiple or general	public data online	NLP	tone	panel data	no
Lamla and Lein (2014)	inflation expectations	multiple	Inflation	content analyses provider	content analysis (human)	volume, tone	time series	Granger
Lamla and Lein (2015)	inflation perceptions	multiple	Inflation	content analyses provider	content analysis (human)	volume, tone	time series	Granger
Lamla and Maag (2012)	inflation forecasting	multiple	Inflation	content analyses provider	content analysis (human)	volume, tone	time series	no
Lamla et al. (2020)	business expectations	multiple	multiple or general	content analyses provider	content analysis (human)	volume, tone	time series	Granger
Larcinese et al. (2011)	determinants of coverage	newspapers	multiple or general	news archive	keyword-based counts	volume	panel data	no
Larsen et al. (2021)	inflation expectations	multiple	multiple or general	news archive	NLP + ML	topics	time series	no
Lei et al. (2015)	inflation expectations	newspapers	inflation	news archive	content analysis (human)	volume, tone	time series	Granger

Reference	Topic	Media type(s)	Content theme(s)	Media data			Main econo-metric method	Causal inference
				Main source	Analysis	Measure(s)		
Lott and Hassett (2014)	determinants of coverage	newspapers	multiple or general	news archive	content analysis (semi-automated)	volume, tone	time series	no
Rambaccussing and Kwiatkowski (2020)	macroeconomic forecasting	newspapers	multiple or general	news archive	NLP + ML	tone	time series	no
Reed (2016)	asset prices	social media	multiple or general	public data online	keyword-based counts	volume	time series	Granger
Tausch and Zumbuehl (2018)	willing to take risks	multiple	multiple or general	content analyses provider	content analysis (human)	tone	panel data	no
Thorsrud (2020)	GDP forecasting	newspapers	multiple or general	news archive	NLP	topics	time series	no
Xu et al. (2018)	inflation expectations	newspapers	inflation	news archive	keyword-based counts	volume, tone	time series	Granger
Xu et al. (2022)	inflation expectations	television	inflation	news archive	keyword-based counts	volume	time series	Granger

Notes: *NLP* refers to natural language processing, while *ML* signifies machine learning.