

SECTORAL MEDIA FOCUS AND AGGREGATE FLUCTUATIONS

RYAN CHAHROUR, KRISTOFFER NIMARK, AND STEFAN PITSCHNER

ABSTRACT. We formalize the editorial role of news media in a multi-sector economy and show that media can be an independent source of business cycle fluctuations, even when the information they report is accurate. Our approach tightly links agents' beliefs to real economic developments and allows for incomplete information without exogenous noise shocks. In the model, media monitor the economy, making state-dependent decisions on which subset of sectors to report. Accurate public reporting about sectoral developments that are newsworthy but unrepresentative, causes firms in all sectors to over- or underinvest in productive capacity. We construct historical time series of sectoral news coverage in the US and use them to calibrate a multi-sector model driven only by sectoral TFP shocks. Time-varying media focus generates demand-like aggregate fluctuations that are orthogonal to productivity. Presented with historical productivity shocks constructed by the BEA, the model reproduces the 2009 Great Recession.

1. INTRODUCTION

A fundamental question in macroeconomics regards the sources of aggregate fluctuations. Cochrane (1994) goes through a laundry list of plausible candidates, including technology, monetary policy, government spending, oil price and credit shocks, and argues that these types of shocks are either too small, or imply counterfactual correlations between different macro economic variables. He summarizes this state of affairs, writing

“It would be nice to point to recognizable events, of the type that is reported by newspapers, as the source of economic fluctuations, rather than to residuals from some equations.”

In this paper, we argue not only that aggregate fluctuations can be generated by the type of events that are reported by newspapers, but, in fact, that some events generate aggregate fluctuations *because* they are reported by newspapers. We propose a model in which accurate public reporting about sectoral developments that are unrepresentative of the economy as

Date: October 22, 2019. The authors thank George-Marios Angeletos, Xavier Gabaix, Alex Kohlhas, Ellen McGratten, Joerg Stoye, Mathieu Taschereau-Dumouchel, Nellie Zhao, seminar participants at the Cornell Macro Lunch, Bocconi University, Federal Reserve Bank of Boston, Boston College, Federal Reserve Bank of Cleveland, SED 2016, NBER Summer Institute 2019, NORMAC 2019, University of Florida, Georgetown University, Uppsala University and *EABCN Conference on New Approaches for Understanding Business Cycles* for useful comments and suggestions. Shilpy Agarwal and Linchen Zhang provided excellent research assistance.

Chahrour: Economics Department, Boston College. *e-mail:* ryan.chahrour@bc.edu *webpage:* www.chahrour.net. *Nimark:* Economics Department, Cornell University. *e-mail:* pkn8@cornell.edu *webpage:* www.kris-nimark.net. *Pitschner:* Department of Economics, Uppsala University. *e-mail:* pitschner@gmail.com, *webpage:* www.stefanpitschner.com.

a whole causes firms across all sectors to over- or underinvest in productive capacity. This creates the appearance of aggregate shocks that are orthogonal to productivity, even though the only source of exogenous variation are sector-specific productivity shocks.

A recent literature has demonstrated that, under certain conditions, production networks can lead firm- or sector-specific shocks to generate aggregate fluctuations, e.g. Horvath (1999), Carvalho (2010), Acemoglu, Carvalho, Ozdaglar and Tahbaz-Saleh (2012), Carvalho and Gabaix (2013), Baqaee and Fahri (2019) and Carvalho and Grassi (2019). Shocks to a single sector or firm propagate to other sectors or firms through the trade of intermediate inputs. Foerster, Sartre and Watson (2011) and Atalay (2017) quantify these channels, and their results suggest that sector-specific shocks can explain a substantial portion of observed aggregate output fluctuations. However, trade in intermediate inputs by itself does not induce strong enough correlation in production across sectors to account for all of the observed volatility of aggregate output.

We show that news media can serve as a powerful additional source of sectoral comovement. A basic premise of our argument is that individual firms do not have the resources to directly observe every sector in the production network. Instead, firms rely on news media to monitor the economy on their behalf and to report the most newsworthy developments. However, even accurate reports provide only a partial picture of the economy. Such partial information, in turn, may lead firms to over- or underestimate how much of their product other firms will demand. As in Angeletos and La’o (2010, 2013), a firm that is overly optimistic about demand for its output overinvests in productive capacity. If firms across different sectors receive the same partial information via news media, over- or underinvestment will be correlated across sectors. News media thus function as a coordination device for the economy, increasing the correlation of sectoral outputs beyond what would result from sectors’ trading relationships alone.

We embed state-dependent news reporting in a modified version of the multi-sector model of Acemoglu *et al* (2012). In our model, news media act as information intermediaries that relay information about the state of the economy to firms. We argue that there are two aspects of this role that are particularly relevant for understanding business cycles. First, news organizations monitor the economy by collecting and producing information about a large number of events. Second, they make editorial decisions about which events are sufficiently newsworthy to be reported.¹ We formalize these editorial decisions using *news selection functions*, first introduced in Nimark and Pitschner (2019). A news selection function is a mapping from the state of the world to a vector of reported outcomes. News selection functions provide a flexible way to model state-dependent editorial decisions, thereby capturing the changing focus of news coverage over time.

¹Our news selection functions represent what in the journalism and political science literature is referred to as the *gatekeeping process*. The former literature has studied where gatekeeping occurs, e.g. Shoemaker and Vos (2009) discusses whether the decision about what makes the news is made primarily at the news gathering (journalist) level, or at the news processing (copy-writing and editorial) level. The political science literature has focused mostly on how gatekeeping is affected by ideology and how it affects political opinion. Some of this literature has studied economic news directly, e.g. Soroka, Stecula and Wlezien (2015) who argue that news about future economic prospects affect public political opinion.

By determining what gets reported in which states of the world, a news selection function implicitly defines a notion of *newsworthiness*. A given notion of newsworthiness in turn implies a specific selection bias of events that end up in the news. For instance, if extreme events are considered newsworthy, they will be over-represented in media reports relative to the unconditional frequency with which they occur. The effects of selection bias increase with the number of sectors in the economy: With more potential events for news media to choose from, the reported outcome is likely to be more extreme.

State-dependent news selection also implies that what is reported is informative also about sectors that are not in the news. For instance, if more extreme developments are considered more newsworthy, a mundane news report allows agents to rule out extreme events in *every* sector, including in those sectors that were not reported on at all. Conditional beliefs about non-reported sectors thus depend on how newsworthy the reported outcomes are. Since firms in all sectors make the same inference, the effect of state-dependent news reporting on agents beliefs about non-reported sectors introduces additional correlation in firms' production decisions across sectors.

The effect that news media have on the economy depends on how newsworthiness is judged. To investigate empirically what these criteria may be in practice, we construct a measure of sectoral news coverage using articles from US newspapers. Using this new data set, we establish several facts. First, larger sectors receive more news coverage than smaller sectors. Second, after controlling for their size, some sectors receive a disproportionate amount of news coverage. Third, news coverage of individual sectors tends to increase when a sector experiences unusually large shocks.

We calibrate the model to match these features of the news data and the input-output structure of the US economy. For the production side, we choose parameters such that the model fits the data on intermediate input shares provided by the Bureau of Economic Analysis (BEA) aggregated to 29 sectors. In the calibrated model, state-dependent reporting decisions by news media contribute substantially to aggregate fluctuations. The variance of aggregate output is more than four times larger in the baseline model compared to the same model without news media and twice as large as in a model in which news media randomly chooses which sector to report on. Moreover, when we feed actual sectoral TFP shocks into the model, it predicts a severe recession in 2009, while a full information version of the same model does not.

Time-varying sectoral media focus generates fluctuations in aggregate output that are orthogonal to sectoral TFP. This is the case even though sectoral TFP shocks are the only exogenous source of variation in the model. Productivity in a sector has a bigger impact on aggregate output when the sector is in the news, compared to when it is not. This type of state dependence cannot be captured by a constant linear relationship between sectoral productivity and aggregate output. Researchers applying a Foerster, Sartre and Watson (2011) or Atalay (2017) -style filter to data generated from our calibrated model would conclude that the common component in GDP that is orthogonal to sectoral productivity accounts for about one third of the total variance of aggregate output.

In our model, firms choose their production capacity in anticipation of demand for their products. This mechanism is consistent with the evidence presented by Gennaioli, Ma and

Shleifer (2016), who show that firms' investment growth can be predicted by CFOs' expectations of sales growth, even after controlling for a plethora of other variables. Arif and Lee (2014) use information from firms' balance sheets to document that aggregate investment fluctuations are driven by firms' unduly optimistic expectations about future cash-flows that subsequently fail to materialize. Eisner (1978) and Greenwood and Hanson (2015) provide additional evidence that expectations about future sales drive investment decisions. Furthermore, Gennaioli *et al* (2016) document that expectation errors about sales growth are correlated across surveys and across different types of agents, suggesting that different agents may receive information from the same sources. In our model, news media provide the same partial information about the economy to firms in all sectors, thus providing a mechanism for why firms across different sectors make correlated prediction errors.

Our mechanism for translating changes in firms' beliefs into output decisions is similar to Angeletos and La'O (2013). In that paper, agents trade with randomly-matched trading partners and experience a sentiment shock that drives all firms to be optimistic about the production of their trading partner. In our paper, trading partners are fixed by the production structure, and news media reports on specific sectors drive optimism about production in other sectors. In both papers, firms produce more when they expect high demand for their product from other firms, i.e. when they expect more favorable terms-of-trade.

The idea that common but imperfect signals can generate demand-like disturbances is not new and was first formalized by Lorenzoni (2009). Both Nimark (2014) and Blanchard *et al* (2012) explore this idea empirically within fully specified structural models. Unlike in these earlier papers, however, consumers' expectations about future income plays no role in our model.

Chahrouh and Ulbricht (2019) develop a flexible empirical framework for quantifying the importance of information frictions for business cycles and argue that undue optimism or pessimism can explain up to 51% of the variation in output. Angeletos, Collard and Dellas (2018) develop a semi-structural method that is computationally simpler to implement than a fully structural approach. They find that sentiment shock can explain more than half of the variance of output, consumption and employment at business cycle frequencies. In a similar vein, but using a less structural approach, Angeletos, Collard and Dellas (2019) document that a single "main business cycle shock" appears to be driving most of the variation at business cycle frequencies of several aggregate variables. Time-varying sectoral media focus in our model generates aggregate fluctuations that share many of the properties of this shock, i.e. it generates positive co-movement between output, employment and consumption through fluctuations that are orthogonal to productivity.

In this paper we propose a new approach to model incomplete information. Instead of noisy signals about variables of common interest, firms in our model receive perfectly accurate information about some sectors in the economy. But because this information only provides a partial picture of the economy, firms do not have complete information about all developments that could potentially affect their production decisions. Like us, Tian (2019) studies the role of input-output linkages when firms may be overly optimistic, but that paper models common belief fluctuations driven by exogenous noise in a public signal.

One advantage of our approach to modeling incomplete information is that it avoids introducing exogenous informational shocks either at the firm or the aggregate level. This is more

than an aesthetic advantage: Given a specific news selection function, beliefs are completely determined by the cross-sectional profile of productivity shocks. The model thus tightly links agents' beliefs to the real economy, and it makes specific predictions about what realizations of sector-specific shocks should be associated with undue optimism or pessimism. Macroeconomic models with incomplete information have mostly used survey data on expectations to discipline agents' beliefs, or inferred these beliefs indirectly from agents' decisions, e.g. Melosi (2016), Blanchard, L'Huillier and Lorenzoni (2013), Nimark (2014) and Angeletos *et al* (2018). By explicitly modeling news media as information intermediaries, we can exploit our novel data on news coverage to discipline agents' beliefs.

There is a large literature that studies news media markets from the perspectives of industrial organization and political economy, but there are surprisingly few papers that have incorporated an explicit role for news media in macroeconomic models. Two important exceptions are Carroll (2003), who shows that news coverage can explain how inflation expectations spread through a population, and Veldkamp and Wolfers (2007). Like we do, Veldkamp and Wolfers argue that a common information source can explain why sectoral output is more correlated than sectoral productivity. In their model, information providers exist to exploit economies of scale in information dissemination. In equilibrium, information about aggregate shocks relevant for every sector is cheaper for firms to acquire than information about their own sector. Information consumption is therefore tilted towards aggregate shocks and away from sector specific shocks, implying that sectoral output is more correlated than sectoral productivity.

Blinder and Krueger (2004) and Curtin (2007) document that a majority of households get most of their economic news from either TV news shows or newspapers. The samples of these studies include periods during which the internet was still in its infancy, and one may reasonably ask how much news consumption patterns have changed due to the increasing importance and popularity of online information sources. Based on browser history data of 50,000 US households, Flaxman *et al* (2016) report that “*the vast majority of online news consumption is accounted for by individuals simply visiting the home pages of their favorite, typically mainstream, news outlets*”. Mainstream news outlets tend to cover the same news events online as in their print and broadcast editions, so the move of many news providers to an online format appears to be mostly a change in viewing technology rather than a change in the type of news content agents consume.

While there is relatively little theoretical work analyzing the role of news media in the macro economy, there exists a growing empirical literature that analyzes news based data sources and how they affect agents' expectations. Larsen, Thorsrud and Zhulanova (2019) document that news topics predict household inflation expectations, even after controlling for standard macro economic variables. They also document state-dependence in the degree to which households update their expectations that is consistent with news media being the driving force behind this pattern. Lamla, Lein and Sturm (2007) and Buchen (2014) both directly attempt to test Wolfers and Veldkamp's (2007) theory of sectoral co-movement using German news coverage data.

2. A MULTI-SECTOR ECONOMY

We study the role of state-dependent media focus in a simple multi-sector economy populated by two types of agents. A representative household decides how much labor to supply and how much to consume of each good. Firms decide how much labor and intermediate inputs to use in production. There are n sectors in the economy, and each sector consists of a continuum of firms that sell their goods in perfectly competitive markets. Sector $i \in \{1, 2, \dots, n\}$ is defined by how good i enters in the production function of other goods and by how goods produced by other sectors enter into the production function of a firm in sector i . The model structure is identical to that in Acemoglu *et al* (2012), with the exceptions that (i) aggregate labor supply is endogenous and (ii) firms choose labor inputs before demand for their product is known with certainty.

In the next sections, we embed news media in the model and then describe in detail how state-dependent reporting decisions determine what information is available to firms when they make their labor input decision. Here, we describe agents' preferences and the production structure of economy and discuss some of the properties of the model that are important for what follows.

2.1. Sectors and firms. A firm in sector i uses the Cobb-Douglas production function

$$Q_i = Z_i \left(\prod_j X_{ij}^{\alpha_{ij}} \right) L_i^{1-\alpha} \quad (2.1)$$

to produce good Q_i . Z_i is a sector-specific productivity shock, X_{ij} are intermediate inputs used by sector i and produced by sector j , and L_i is labor inputs used in sector i . The coefficients α_{ij} denote the share of good j used in the production of good i . The production function exhibits constant returns to scale so that $\sum_{j=1}^n \alpha_{ij} = \alpha$. The good Q_i produced by sector i can be used either for consumption C_i or as an intermediate input X_{ji} so that

$$C_i + \sum_j X_{ji} = Q_i. \quad (2.2)$$

Firms in sector i choose labor and intermediate inputs to maximize profits Π_i

$$\Pi_i = P_i Q_i - W L_i - \sum_j P_j X_{ij} \quad (2.3)$$

taking prices P_j of all goods as given.

2.2. The representative household. The representative household decides how much to work and how much to consume of each good. It solves the problem

$$\max_{X_1, \dots, X_n, L} C - \frac{L^{1+1/\nu}}{1 + 1/\nu} \quad (2.4)$$

subject to the budget constraint

$$C = W L + \Pi \quad (2.5)$$

where W is the wage, $L \equiv \sum_{i=1}^n L_i$ and $\Pi \equiv \sum_{i=1}^n \Pi_i$. The consumption bundle C is an equal-shares Cobb-Douglas aggregate of goods

$$C = \prod_i C_i^{\frac{1}{n}}. \quad (2.6)$$

where C_i denotes the amount of good i used for final consumption. We normalize the price of the aggregate consumption bundle C to 1.

2.3. Optimality conditions and timing of actions. To capture the notion that some production decisions are taken in anticipation of uncertain demand, firms choose the quantity of labor inputs before production takes place and before wages and prices are observed. In a second stage, firms choose how much intermediate inputs to use. The first stage of a firm's optimization problem is to solve

$$\max_{L_i} E \left[P_i Q_i - W L_i - \sum_j P_j X_{ij} \mid \Omega_i \right] \quad (2.7)$$

where Ω_i is the information set of a firm in sector i defined as

$$\Omega_i = \{Z_i, \mathbf{s}, \mathbf{r}\}. \quad (2.8)$$

A firm thus observes his own productivity as well as \mathbf{s} and \mathbf{r} , which summarize the information reported by news media. The vectors \mathbf{s} and \mathbf{r} are defined in the next section.

The optimal labor input decision equates expected marginal product of labor with the expected marginal cost, i.e. the real wage. A firm's equilibrium labor demand can thus be described as the labor share $(1 - \alpha)$ times the ratio of expected revenue and expected wage

$$L_i = (1 - \alpha) \frac{E [P_i Q_i \mid \Omega_i]}{E [W \mid \Omega_i]}. \quad (2.9)$$

After firms choose labor inputs, production takes place, sectors trade intermediate inputs and the household decides how much of each good to use for final consumption. From the Cobb-Douglas structure, equating marginal product with marginal cost of intermediate input X_{ij} implies that firms in sector i spend share α_{ij} on intermediate input good j

$$X_{ij} = \alpha_{ij} \frac{P_i Q_i}{P_j}. \quad (2.10)$$

Households supply labor until the marginal utility of consuming the real wage equals the marginal disutility of working

$$L^{\frac{1}{\nu}} = W \quad (2.11)$$

and spend a fraction $1/n$ of their income on each good i

$$P_i C_i = \frac{C}{n}. \quad (2.12)$$

2.4. Determinants of sectoral labor demand. State-dependent reporting affects output in the model via the expectations in the labor input decision described by (2.9). In equilibrium, demand for labor in sector i depends on the expected gross sales of sector i goods and the expected cost of labor, W . As in Angeletos and La'o (2010, 2013), labor inputs are strategic complements among firms. If other firms hire more labor, the demand for intermediate good i goes up, increasing its price. Firms in sector i , anticipating higher prices for their output, hire more labor themselves. However, this effect is partly offset by the fact that an increase in labor demand by other sectors increases the market wage. The strength of the second effect depends on the labor supply elasticity ν .

In the appendix, we show that the labor demand function in (2.9) can be expressed as a function of expected aggregate output C and wages W ,

$$L_i = (1 - \alpha) \lambda_i \frac{E[C | \Omega_i]}{E[W | \Omega_i]}, \quad (2.13)$$

where λ_i is the Domar weight of sector i . The Domar weight of a sector captures the importance of the sector as a supplier of intermediate goods to other sectors and is a function of the parameters of the production function (2.1) and the consumption good aggregator (2.6).² Hence, equation (2.13) implies a unit elasticity of labor demand with respect to expected consumption for all sectors, and an elasticity with respect to expected wages of minus one.

3. THE EDITORIAL ROLE OF NEWS MEDIA

In industrialized economies, firms are linked to each other through a complex network of trading relationships of intermediate goods. Shocks to a given sector propagate to other sectors through this network, and an individual firm's optimal production decisions partially depend on developments in other sectors. Given the complexity of a modern economy, arguably no individual firm has the resources to monitor every sector in the economy that could be relevant for its own production decision. Instead, many firms receive information about the economy via information intermediaries that monitor the economy and make state-dependent decisions about what to report. In this section, we describe how this editorial role of news media can be formalized within the multi-sector model presented above. This framework is based on the more abstract setting in Nimark and Pitschner (2019).

3.1. Formalizing state dependent reporting. The state of the economy is the n -dimensional vector of sector-specific productivity shocks $Z \in \mathcal{Z}_1 \times \mathcal{Z}_2 \times \dots \times \mathcal{Z}_n \equiv \mathcal{Z}$. News media monitor the state of the economy and make state dependent decisions about which elements of Z are most newsworthy. We formalize this monitoring and reporting behavior using *news selection functions*.

Definition 1. (News selection function) *A news selection function $\mathcal{S} : \mathcal{Z} \rightarrow (\mathbf{s}, \mathbf{r})$ is a mapping from n -dimensional states of the world $Z \in \mathcal{Z}$ into pairs (\mathbf{s}, \mathbf{r}) , where $\mathbf{s} \in \{0, 1\}^n$ is an n -dimensional indicator vector and $\mathbf{r} \in \mathbb{R}^r$ is an r -dimensional vector containing the elements Z_i of Z such that $s_i = 1$.*

²More precisely, the Domar weight λ_i is the log of the i^{th} element of the vector $\mathbf{\Lambda} \equiv (I - IO')^{-1} \frac{1}{n} \times \mathbf{1}$ where IO is the input-output matrix with element (i, j) equal to α_{ij} .

A news selection function \mathcal{S} thus associates a pair (\mathbf{s}, \mathbf{r}) with each state of the world $Z \in \mathcal{Z}$. The vector \mathbf{s} indicates which sectors are reported on. An element of \mathbf{s} equal to 1 indicates that the corresponding dimension of Z is reported, and a 0 indicates that the respective dimension is not reported. The vector \mathbf{r} contains the realized values of productivity in the reported sectors. For instance, $\mathbf{s}(Z) = (1, 0, \dots, 0)$ means that in state $Z = (Z_1, \dots, Z_n)$ only the first dimension is reported so that $\mathbf{r}(Z) = Z_1$. Similarly, $\mathbf{s}(\tilde{Z}) = (0, \dots, 0, 1, 1)$ means that in state $\tilde{Z} = (\tilde{Z}_1, \dots, \tilde{Z}_n)$ only the last two dimensions are reported so that $\mathbf{r}(\tilde{Z}) = (\tilde{Z}_{n-1}, \tilde{Z}_n)$. A news selection function \mathcal{S} thus assigns a 1 to element i of \mathbf{s} if the outcome Z_i is sufficiently newsworthy to be reported. Whether the element Z_i is reported or not generally depends on the entire state vector Z .

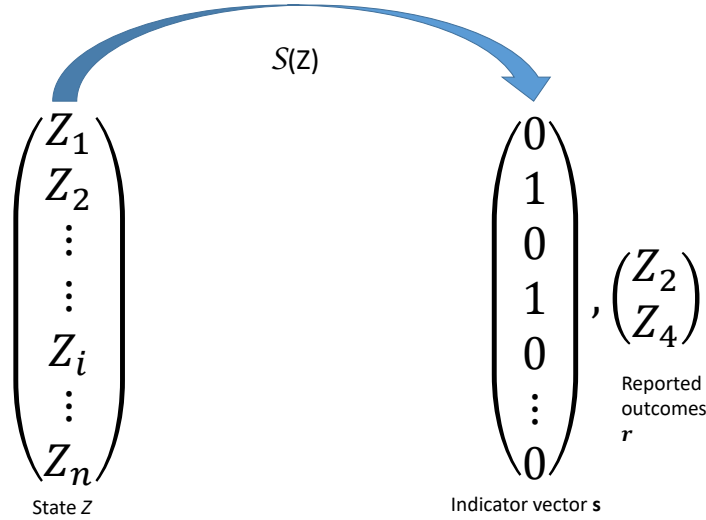


FIGURE 1. The news selection function \mathcal{S} reports productivity in sector 2 and sector 4 in state Z .

The dimension of \mathbf{r} (and the number of non-zero elements in \mathbf{s}) is r , so that all sector-specific shocks are reported if $r = n$. The elements in \mathbf{r} are reported accurately by the information provider. However, if $r < n$ only a subset of the sector-specific productivity shocks are reported. The vector \mathbf{r} then only provides a partial picture of the state of the economy.³

The mapping from realized states to reported sector-specific shocks is illustrated in Figure 1. There, the news selection function represented by \mathcal{S} reports Z_2 and Z_4 in state Z . An agent who receives reports from an information provider characterized by \mathcal{S} would then know the values of the productivity shocks in sector 2 and sector 4. This is the information contained in the vector \mathbf{r} . However, the agent would also know that the information provider chose to not to report about any of the other sectors. This information is contained in the

³Nimark and Pitschner (2019) show that agents who are constrained in terms of how many stories they can observe can achieve a lower posterior entropy by delegating the choice of what to observe to an organization or mechanism that can condition on the realized state.

indicator vector \mathbf{s} . To the extent that these reporting decisions are state-dependent, they will also reveal information about the unreported sectors, i.e. sectors 1, 3, 5, 6, ... n .

3.2. State-dependent reporting and beliefs. The firms in our model are Bayesian and understand the state-dependence of reporting decisions encoded in \mathcal{S} . A firm that observes \mathbf{r} and \mathbf{s} has the posterior beliefs $p(Z | \mathbf{r}, \mathbf{s})$. The posteriors are affected by the state-dependence of reporting decisions in two distinct ways. First, since some sectoral outcomes are considered more newsworthy than others, the distribution of reported sector-specific productivity shocks is different from the unconditional distribution of Z_i so that

$$p(Z_i | s_i = 1) \neq p(Z_i). \quad (3.1)$$

If not all outcomes of $Z_i \in \mathcal{Z}_i$ are equally newsworthy, the density $p(Z_i | s_i = 1)$ redistributes probability mass towards more newsworthy regions of the support of Z_i . Some types of outcomes may thus be over-represented in the news relative to their unconditional frequencies of occurring.

Second, state dependent reporting behavior implies that firms may also update their beliefs about non-reported sectors. Intuitively, this is because unreported sectoral outcomes that would have been reported had they occurred can be ruled out. More precisely, firms observing $s_j = 0$ can rule out any outcome, Z_j , that would have implied $s_j = 1$.

The selection bias introduced by news selection functions is related to, but conceptually distinct from the filtering biases that has been studied in the political economy literature. For instance, in the model of Stromberg (2004), media bias takes the form of giving more coverage to policy proposals that either affect larger groups of voters, or groups of voters that are more attractive from an advertising perspective. However, the editorial decision in that model is not state-dependent. Another form of filtering bias is proposed in Chan and Suen (2008). In their model, the state takes a continuous value in $(0, 1)$, but news media are restricted to reporting a binary signal. Like in our framework, news media thus provide a coarser signal than the true state of the world. However, in Chan and Suen (2008) news media do not make a decision about what events to report on.

More broadly, the political economy literature has mostly studied models in which reporting strategies relate to a single state variable, see for instance the survey by Gentzkow, Shapiro and Stone (2015). The news selection function framework presented here is more flexible and allows for the focus of what the news are about to depend on the state of the world. It can thus naturally capture the kind of crowding out effects generated by major news events that Eisensee and Stromberg (2007) as well as Nimark and Pitschner (2019) document empirically.

A key feature of our framework is that a news selection function classifies the sectoral outcomes in the state Z as either being newsworthy enough to be included in \mathbf{r} or not. The criteria used for this classification determine how the indicator vector \mathbf{s} depends on the state Z , and how the state-dependence of reporting decisions affects agents' beliefs. In the next section we discuss how three specific notions of newsworthiness, as encoded by different news selection functions, affect news selection biases and posterior beliefs.

4. THREE NOTIONS OF NEWSWORTHINESS

News media monitor the world and report those events that are considered most newsworthy. What kind of events get reported thus depends on the criteria used to judge how newsworthy an event is. In this section we study three different notions of newsworthiness and how the implied selection biases affect firms' beliefs. The three notions are (i) extreme (or unusual) outcomes are more newsworthy, (ii) negative outcomes are more newsworthy, and (iii) some sectors are inherently more newsworthy. The journalism literature has identified certain characteristics as contributing to the newsworthiness of an event, e.g. Shoemaker and Vos (2009) and Harcup and O'Neill (2016). The three criteria we consider here correspond to the subset of these that most naturally applies to economic news reporting decisions.

The notions of newsworthiness we study here are highly stylized, which helps us illustrate how the state-dependence of reporting decisions implied by each notion affect beliefs. In Section 5, we present empirical evidence on sectoral news coverage and discuss what makes sectoral developments more newsworthy in practice.

In this section, we assume that Z_i are distributed as independent log standard normals so that $z_i \equiv \log Z_i \sim N(0, 1) \forall i$ and $p(Z_j | Z_i) = p(Z_j) : j \neq i$. Neither of these assumptions are central to the mechanisms discussed here, but they help simplify the exposition. We relax the assumption of uncorrelated shocks when we solve and simulate the model.

4.1. Extreme outcomes are more newsworthy. The first notion of newsworthiness we study considers extreme or unusual events more newsworthy than more commonplace events. Shoemaker and Vos (2009) survey the literature that studies which criteria news organizations use to judge whether an event is newsworthy. They argue that one such criterion is *deviance*, which can be either normative, social or statistical. They define normative or social deviance as deviations from norms, laws and social status quos. Statistical deviance is defined as the degree to which an “*event is out of the ordinary or unusual*” and is the notion of newsworthiness that we study here. We formalize it as follows.

Definition 2. (Extreme outcomes more newsworthy) *The news selection function $\mathcal{S}_{|z|}$ treats more extreme outcomes as more newsworthy if for each pair i and j such that $s_i = 1$ and $s_j = 0$ we have that $|z_i| \geq |z_j|$.*

The news selection function $\mathcal{S}_{|z|}$ thus orders outcomes $z_i : i = 1, 2, \dots, n$ in terms of their absolute deviations from their means and reports the values of shocks that had the r largest such deviations. Given that the normal distribution is single-peaked and symmetric, this corresponds to reporting the r least probable outcomes. The news selection function $\mathcal{S}_{|z|}$ thus captures the notion that more unusual events are considered more newsworthy.⁴

The state dependence of reporting decision implied by $\mathcal{S}_{|z|}$ means that firms are more likely to observe extreme productivity outcomes. The next proposition proves this formally and shows that this selection effect grows with the number of sectors.

⁴The same notion of newsworthiness is also explored in Nimark (2014). The information structure in that paper did not specify the editorial decisions of news media explicitly. However, the set up there can be thought of as having been generated by a news selection function monitoring a high-dimensional state vector where one element is the common productivity shock and the other elements are unrelated to the macro economy.

Proposition 1. *For a given $r < n$, the variance of productivity shocks conditional on being reported $\text{var}(z_i | s_i = 1)$ is larger than the unconditional variance $\text{var}(z_i)$ and increasing in the number of sectors n .*

Proof. In the Appendix. □

To prove the first part of the proposition, we use that in every state of the world, the squared value of every reported productivity shock is larger than the squared value of every unreported shocks. The squared values of the reported shocks then state-wise dominates the squared values of the non-reported shocks, implying a higher expected squared value, i.e. a higher variance. To prove the second part, we use that adding dimensions to the state can only make the expected squared deviation of the r reported shocks larger.

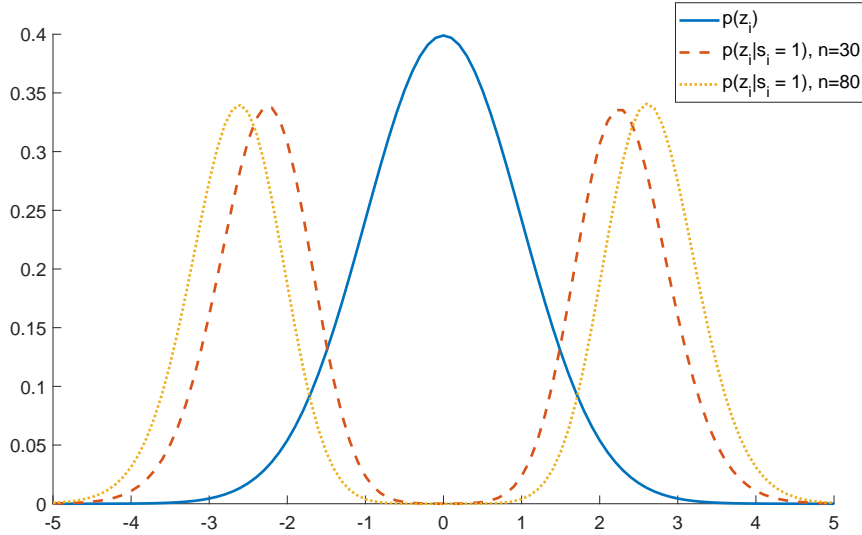


FIGURE 2. The distribution of Z_i conditional on $s_i = 1$ for $n = 30$ and $n = 80$ implied by the news selection function $\mathcal{S}_{|z|}$.

Figure 2 illustrates the selection bias implied by $\mathcal{S}_{|z|}$. It shows the distribution of z_i conditional on $s_i = 1$ for $n = 30$ and $n = 80$ when news media reports a single sector (i.e. $r = 1$). For comparison, we also plot the unconditional distribution of z_i . The distribution $p(z_i | s_i = 1)$ depends on the number of sectors in the economy. With a larger number of sectors, the most extreme outcome is likely to be more extreme. This consequence of having a larger number of sectors is illustrated in Figure 2, where the distribution $p(z_i | s_i = 1)$ associated with $n = 80$ has more mass further from zero than the distribution that arises when $n = 30$.

The state-dependence of the news selection function thus affects what kind of events are more likely to be reported. This state-dependence also affects how firms update their beliefs about non-reported sectors, as summarized by the following proposition.

Proposition 2. *The conditional variance of unreported productivity shocks $\text{var}(z_j | \mathbf{r}, \mathbf{s}, s_j = 0)$ is increasing in the minimum value of the reported productivity shocks $\min\{|z_i| : s_i = 1\}$.*

Proof. In the Appendix. □

Proposition 2 implies that firms update their beliefs about the unreported sector shocks $\{z_j : s_j = 0\}$ when they observe the values of the reported sector shocks in \mathbf{r} , even if shocks are independent across sectors. The logic is as follows. If only the most extreme productivity outcomes are reported, any non-reported outcome must be less extreme than the least extreme among the reported outcomes. The conditional distribution of the unreported sector shocks are thus symmetrically truncated normal distributions where the truncation points are $-\min\{|z_i| : s_i = 1\}$ and $\min\{|z_i| : s_i = 1\}$. The proposition then follows from the fact that the variance of a symmetric truncated normal is increasing in the distance of the truncation points from the mean. In Figure 3, the shaded blue areas indicate the regions of the support of the unconditional distribution of z_j that have zero posterior probability conditional on $s_j = 0$ and $\min\{|z_i| : s_i = 1\} = 1.5$.

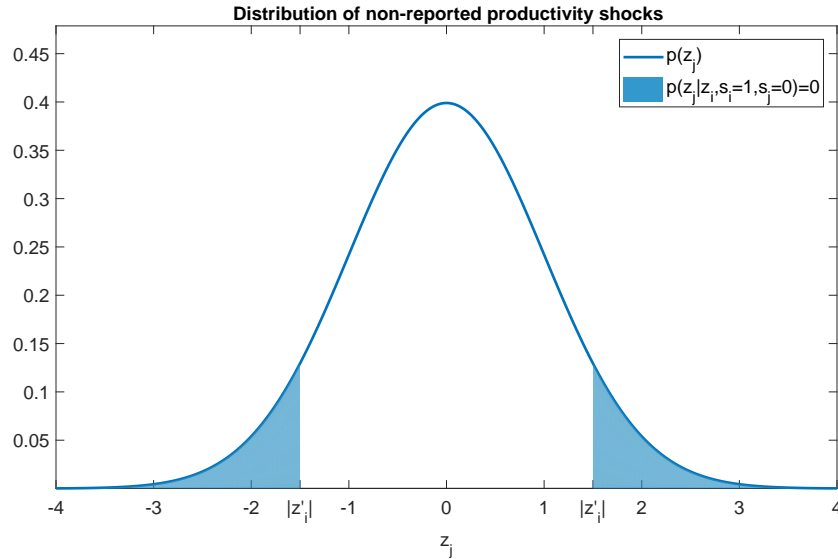


FIGURE 3. The distribution of Z_j conditional on $s_j = 0$ and $\min\{|Z_i| : s_i = 1\}$ implied by the news selection function $\mathcal{S}_{|z|}$

Since firms can rule out outcomes more extreme than $\min\{|z_i| : s_i = 1\}$ for unreported sectors, their conditional uncertainty rises when more extreme events are reported. When something extreme (e.g. a financial crisis) occurs, it is always reported. Major, but less extreme events may then be crowded out of the news coverage and go unreported. However, if something mundane is in fact reported, firms can infer that whatever has occurred in the non-reported sectors must be even more mundane. In such cases, they can rule out large portions of the tail in the distributions of the non-reported sectors. State-dependent reporting decision that treat extreme outcomes as more newsworthy thus generates time-varying conditional uncertainty about productivity in non-reported sectors.

4.2. Negative outcomes are more newsworthy. Another notion of newsworthiness that is potentially relevant is that negative events may be considered more newsworthy than

positive ones. That negative economic news are indeed considered more newsworthy by news organizations is shown by Harrington (1989), who documents that network television news overemphasize bad economic news. Similarly, Soroka (2012) documents that bad news about unemployment, inflation and interest rates are more likely to be reported by the *New York Times* than good news about the same variables. In a recent survey of the news values literature, Harcup and O’Neill (2016) lists *bad news* as one characteristic that makes an event more newsworthy.

To formalize the notion that negative outcomes are considered more newsworthy, we can define a news selection function \mathcal{S}_- that orders the newsworthiness of sectoral outcomes according to their relative position in \mathbb{R} .

Definition 3. (Negative outcomes more newsworthy) *More negative outcomes are considered more newsworthy according to the news selection function \mathcal{S}_- for any pair $i, j \in \{1, 2, \dots, n\}$ such that $s_i = 1$ and $s_j = 0$ we have that $z_i \leq z_j$.*

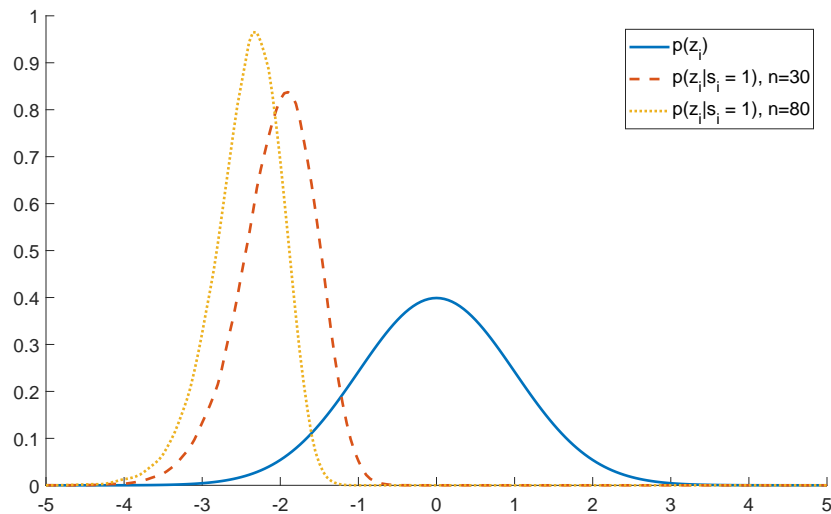


FIGURE 4. The distribution of Z_i conditional on $s_i = 1$ for $n = 30$ and $n = 80$ implied by the news selection function \mathcal{S}_- .

The news selection function \mathcal{S}_- thus reports the r lowest elements in z . The state dependence of \mathcal{S}_- affects the conditional mean of both reported and unreported outcomes.

Proposition 3. *The mean of reported productivity shocks $E(z_i | s_i = 1)$ is lower than the unconditional mean of productivity shocks and decreasing in the number of sectors n .*

Proof. In the Appendix. □

The proof uses that the values of reported sector shocks are lower than the unreported sector shocks in all states of the world, and that the weighted conditional means of reported and unreported shocks must equal the unconditional mean. The selection bias underlying Proposition 3 is illustrated in Figure 4. There, we plot the unconditional distribution of z_i

together with the distributions of the same variable conditional on being reported for $n = 30$ and $n = 80$. Both the conditional mean and variance are decreasing in the number of sectors n . With a larger number of sectors, the most negative outcome is more likely to be far out in the left tail of the distribution, but the dispersion around that mean is also decreasing.

Again, the selection bias introduced by \mathcal{S}_- affects the conditional distributions of unreported sector shocks.

Proposition 4. *The expected value of non-reported productivity shocks $E(z_j | \mathbf{r}, \mathbf{s}, s_j = 0)$ is increasing in the maximum value of the reported productivity shocks $\max\{z_i : s_i = 1\}$.*

Proof. In the Appendix. □

Since all non-reported sector shocks must be (weakly) more positive than the reported shocks, the conditional distribution of a non-reported shock is a left-truncated normal. The proposition then follows from observing that the truncation point is given by $\max\{z_i : s_i = 1\}$ and because the mean of a left truncated distribution is increasing in the truncation point. This is illustrated in Figure 5. If the most negative outcomes are reported, no unreported outcome can be smaller than the largest reported outcome. This means that realizations to the left of $\max\{z_i : s_i = 1\}$ in the support of the unreported shocks z_j can be ruled out. In the figure, this region is shaded in blue.

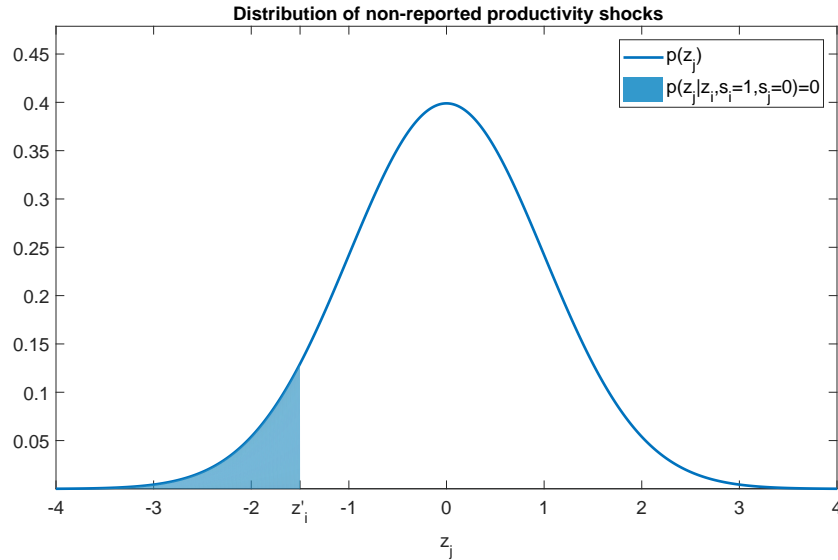


FIGURE 5. The distribution of a non-reported productivity shock $E(z_j | z^s, s, s_j = 0)$ implied by the news selection function \mathcal{S}_- .

4.3. Unconditionally more newsworthy sectors. The framework also allows for modeling some sectors as being inherently more newsworthy regardless of the realized state. For instance, some sectors may receive more news coverage because they are larger than others, or because they have trading relationships with a large number of other sectors. That this type of considerations may make a sector more newsworthy corresponds to what Harcup and

O’Neill (2016) refer to as *magnitude*. In their terminology, magnitude describe the number of people affected by an event, and large magnitude events have been documented as being considered more newsworthy.

We define a sector as being inherently more newsworthy than another sector as follows.

Definition 4. (Unconditionally more newsworthy sectors) *Sector i is unconditionally more newsworthy than sector j if for each pair i and j whenever $z_i = z_j$ and $s_i \neq s_j$ we have that $s_i = 1$ and $s_j = 0$.*

Definition 4 does not specify a unique news selection function, since it only specifies whether sector i or j is reported when $z_i = z_j$. To construct a complete ordering of the newsworthiness of different outcomes, the criteria in Definition 4 needs to be combined with some additional criteria. For instance, a news selection function may always report z_i instead of z_j regardless of the state. Another possibility is that deviance or negativity determines newsworthiness, but that the newsworthiness of sectoral developments are also weighted based on the inherent relative newsworthiness of different sectors. Combining sector-specific weights with the previously discussed criteria can be done as follows.

Definition 5. (Weighted news selection functions) *For an n -dimensional vector ω with typical element $\omega_i \in \mathbb{R}_+$, the weighted composite news selection functions $\mathcal{S}_{|\omega|}$ and $\mathcal{S}_{-\omega}$ are constructed by defining their corresponding indicator vectors as $\mathbf{s}_{|\omega|} = \mathbf{s}_{||}(\omega \circ z)$ and $\mathbf{s}_{-\omega} = \mathbf{s}_{-}(\omega \circ [z - \max(z)])$, where \circ denotes the n -dimensional Hadarmard (i.e. element-wise) product.*

The weights ω_i and ω_j in the definition regulates the relative newsworthiness of sectors. The larger ω_i is relative to ω_j , the more likely is sector i to be reported instead of sector j . Developments in an more newsworthy sector will therefore ceteris paribus have a bigger impact on aggregate economy than a less newsworthy sector. In the next section, we present empirical evidence on sectoral news coverage in US newspapers and analyze what makes it more likely that a sector ends up in the news. When we calibrate the model, the vector of weights ω is used to match model moments of news coverage to the corresponding moments in the news coverage data.

5. SECTORAL COVERAGE IN U.S. NEWSPAPERS

How news reporting affects the economy depends on what kind of events are considered most newsworthy. In this section we present and analyze empirical measures of US sectoral news coverage that we will use below to calibrate the model. For our baseline measure, we take a company-sector matching approach where we first identify company names in news articles and assign each company to a sector. For each sector, we then compute the fraction of the total news coverage that refers to companies in that specific sector. This approach allows us to compute a news coverage measure using sectoral definitions that are consistent with those of the BEA’s production accounts. We can thus calibrate the production side and the news media side of the model using consistent sector definitions, and we can study how sectoral news coverage responds to sectoral developments as measured by the BEA. However, we also present some evidence using news coverage data based on articles that make explicit references to sectors or industries as a unit, rather than individual companies.

Our data is from Dow Jones Factiva. We use news articles from six major US outlets that covers the period from 1988 to 2018. The outlets in our sample are the Wall Street Journal, the New York Times, USA Today, the Boston Globe, the Charleston Gazette and the Atlanta Journal Constitution. The first three of these are the largest US newspapers by circulation, and all six have consistent coverage by Factiva. Importantly, these six newspapers are the ones for which Factiva provides the entity tags that we use to match newspaper articles to company names and their respective sectors.

The tags assigned by Factiva to any given news article are names of entities that may or may not be US companies.⁵ Our sample contains 1,178,716 such tags that correspond to 5,175 unique entities. To construct measures of sectoral news coverage from this data, we query Factiva for the NAICS code of each entity as well as its primary location. We also perform a name-based fuzzy-match to Compustat (which we verify manually) to obtain additional information on sector affiliation and country.⁶ Finally, for the 200 most frequent entity tags for which neither Factiva nor Compustat contain sector and country information, we obtain it manually via web searches.

We keep all entities that represent US companies, and that we are able to assign to a specific sector via their NAICS codes. We consider a company US-based if (i) Factiva lists its primary location as the US or (ii) Compustat lists both its postal address and its country of incorporation as the US or (iii) a web search yields that a company has substantial business activities in the US. Together, Factiva and Compustat allow us to identify 3,640 companies and their respective NAICS codes. These companies account for approximately 80% of the total number of entity tags in the sample. In addition, the 200 manually classified entities yield another 106 US companies and increase the fraction covered to approximately 88% percent of all tags. Finally, for the remaining 12% of tags we find that they primarily capture organizations that are not US companies (e.g. sports teams or political institutions).

Based on the extracted NAICS code for each US company in our data, we group news coverage into 29 different sectors that approximately correspond to the definitions in Atalay (2017).⁷ The sector labels are listed in Table 1. Most of the labels are self-explanatory, with perhaps the exception of *F.I.R.E.* which denotes the *Finance, insurance and real estate* sectors. We then measure sectoral news coverage as the number of times US companies belonging to a given sector are mentioned in the news articles in our data set. This approach establishes a correspondence between the news coverage of the sectors that uses sector definitions that are consistent with the sector definitions used by the BEA to construct sectoral output accounts.

For our baseline measure of sectoral news coverage, we use only articles from the Wall Street Journal and the New York Times. The reason for this is that it is only for these two news papers that Factiva entity tags are consistently available from 1988 onwards. In

⁵For instance, the European Union and ISIS are identified by Factiva as entity names in articles, but are neither companies nor US based.

⁶The fuzzy match is based on the Levenshtein distance (Levenshtein 1966) between the name tag provided by Factiva and the company name as it appears in Compustat.

⁷We exclude the government sector from our analysis, since news coverage of government entities is dominated by news that are unrelated to the economy, such as supreme court decisions and political debates. Our sector classification is described in more detail in the Data Appendix.

TABLE 1. Sector Labels

Sector	Sector Name	Sector	Sector Name
1	agriculture & forestry	16	primary metals
2	mining	17	fabric. metal products
3	oil & gas extraction	18	non-electrical machinery
4	construction	19	electrical machinery
5	food & kindred products	20	motor vehicles
6	textile mill products	21	other transportation equipment
7	apparel & leather	22	instruments
8	lumber	23	misc. manufacturing
9	furniture & fixtures	24	transportation & warehousing
10	paper & allied products	25	communications
11	printing & publishing	26	electric & gas utilities
12	chemicals	27	wholesale & retail
13	petroleum refining	28	F.I.R.E.
14	rubber & plastics	29	other services
15	non-metallic minerals		

Notes: See Data Appendix for sector definitions.

addition, the baseline specification excludes companies that are only mentioned by one of the two outlets in any given quarter. This filter excludes more minor events and thus brings the data closer to the notion of public news reports as implemented in the model. Below, we also discuss the implications of using two alternative measures. The first does not impose the filter that both the Wall Street Journal and the New York Times must report about a company. The second alternative measure includes coverage also from USA Today, the Boston Globe, the Charleston Gazette and the Atlanta Journal Constitution, but starts only in 1997, the first year for which entity tags are available for these additional outlets.

5.1. Sample averages of sectoral news coverage. One of the most salient facts in the data is the degree to which the average amount of news coverage received varies across sectors. Figure 6 plots the sectoral shares of total news coverage against their contributions to gross output, together with a 45 degree line. The sample correlation between sectoral news coverage and the sectoral shares of gross output is 0.75. Larger sectors thus tend to receive more coverage than smaller ones. The most widely featured sector in our news data is *Finance, insurance and real estate* followed by *Communications, Other services, Motor vehicles* and *Instruments*.

Sectors that are above the 45 degree line in Figure 6 are over-represented in the news relative to their economic size. More specifically, the *Communications, Motor vehicles, Instruments* and *Printing & publishing* sectors all receive substantially more news coverage than their economic size alone would indicate. That these sectors are over-represented in the news relative to their economic size accords well with a casual reading of recent history. One of the major developments over this period was the rise of the Bay Area Tech industry. *Communications* include mobile phone and cable tv companies such as AT&T, Verizon and Comcast, but also newer companies such as Facebook, Ebay, Netflix and Twitter. The

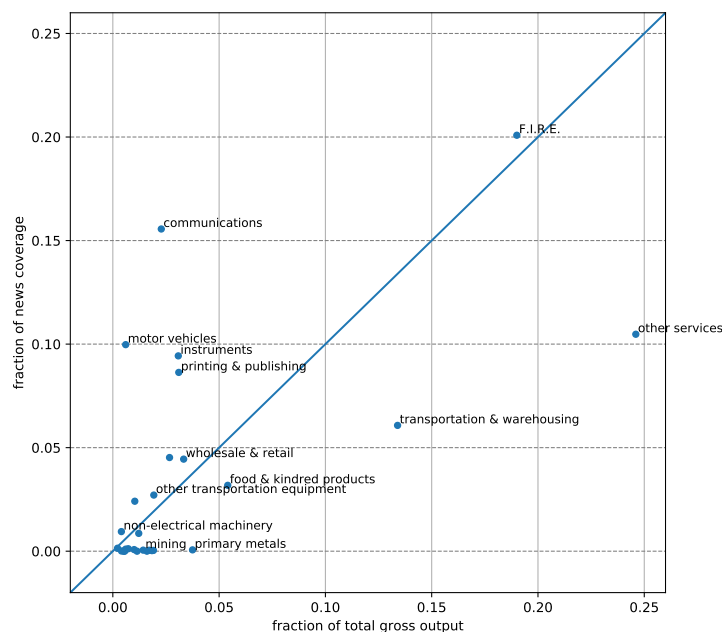


FIGURE 6. Sectoral news coverage and contribution to GDP. The horizontal axis measure the sample average fraction of aggregate value added by a sector, the vertical axis measures the sample average of the fraction of news coverage a sector received.

three most frequently mentioned companies in *Instruments* are Apple, Intel and Hewlett-Packard and news coverage of *Printing & publishing* is completely dominated by articles about Microsoft and (Google’s parent company) Alphabet. Another major economic story over the sample period was the financial crisis and the resulting bail-out of the Detroit based auto industry. The most-frequently mentioned companies in the 10 sectors that receive the most coverage overall are reported in Figure 7.

Another finding is that about half of the sectors receive approximately zero news coverage. This is illustrated in Figure 8, where we plot the cumulative sum of the sectoral shares of news coverage together with the sectoral shares of output. While the 10 most reported on sectors together receive more than 90% of the total news coverage, the 15 least reported on sectors together receive less than 1% of the news coverage. This asymmetry is not as strong in terms of shares of gross output. The 15 smallest sectors produce about 10% of gross output. There are also some large sectors that are substantially under-represented in the news. For instance, the sector *Other services*, which includes companies as varied as IBM and Walt Disney, produces almost a quarter of GDP, but receives only about 10% of the news coverage. (IBM and Walt Disney are the most reported on companies within this sector, so these companies are not necessarily themselves under-represented in the news.)

Overall, the sample averages remain largely unchanged when we use the two alternative measures. *Finance, insurance and real estate* receives somewhat smaller share of the news coverage when use all six newspapers, as both the Wall St Journal and the New York Times tend to cover these industries more than the other newspapers. Imposing the filter that both the Wall St Journal and the New York Times must mention a company in a given quarter,

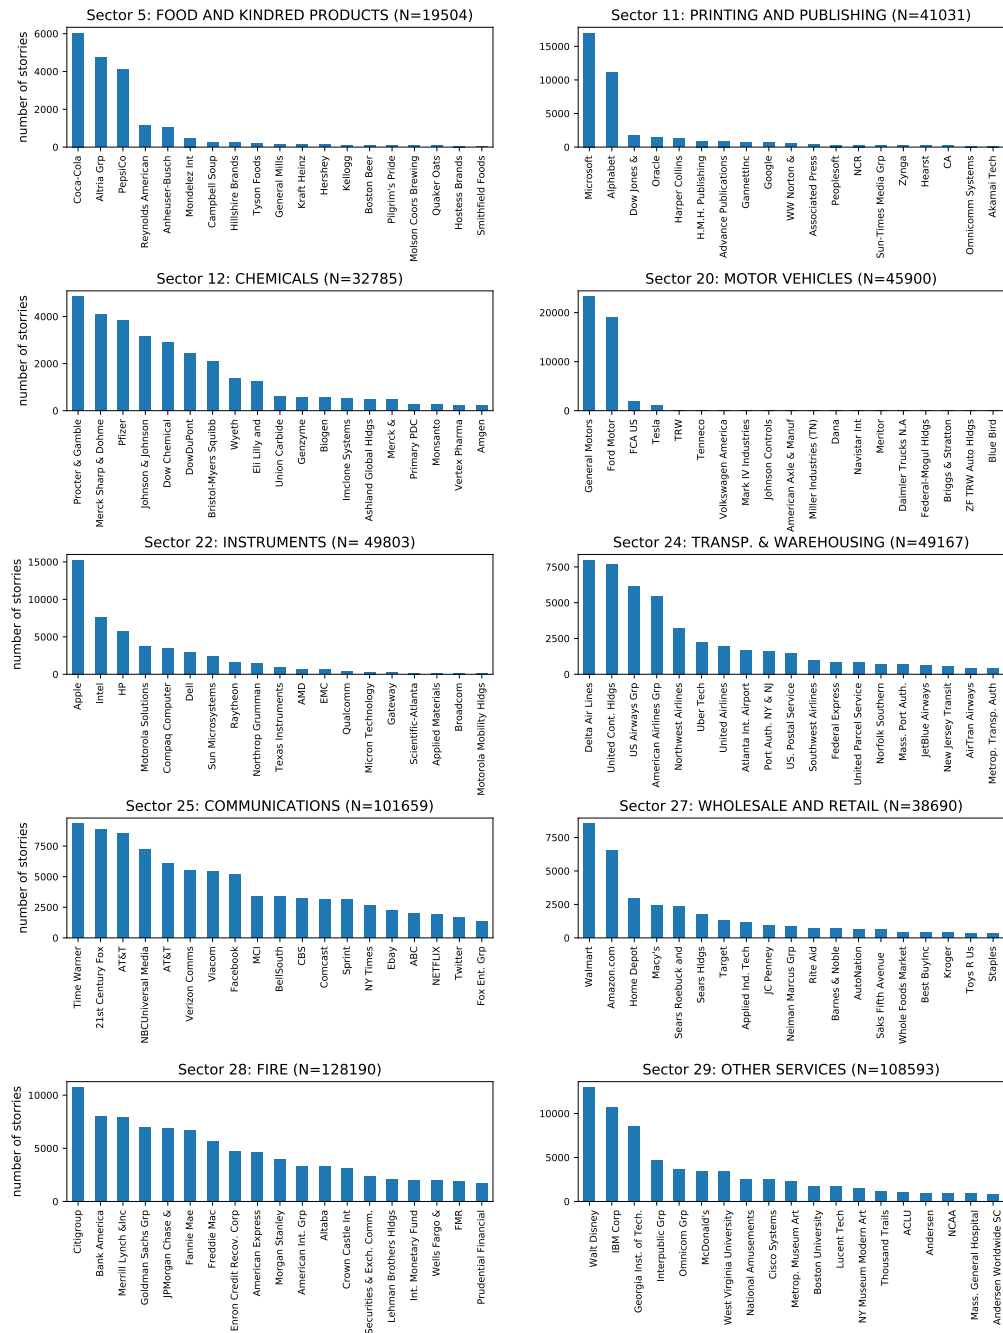


FIGURE 7. Most frequently mentioned company names for the 10 sectors that received the most coverage over the sample.

increases the fraction of news coverage received by the *Motor vehicle* industry, relative to the baseline.

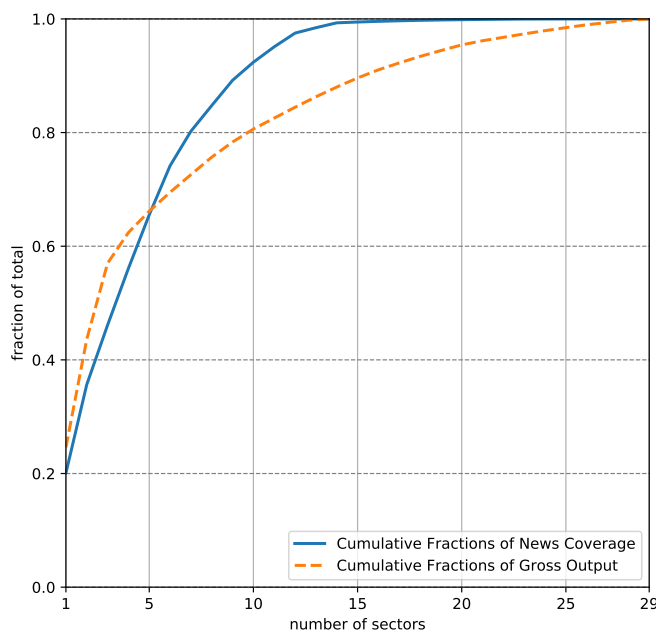


FIGURE 8. Cumulative sum of the sectoral shares of news coverage and gross output.

5.2. State dependence of sectoral news coverage. In addition to its variation across sectors, news focus also varies substantially over time. This is illustrated in Figure 9 where we plot the time series of sectoral news coverage (expressed as fractions of total coverage) for the 10 sectors that receive the most news coverage on average over the sample period. The figure also illustrates that for most sectors and most time periods, the three alternative measures result in broadly similar time series.

The largest changes in news coverage occur during the financial crisis in 2008 and 2009. In this period, news coverage of the *Finance, insurance and real estate* sector increased from a pre-crisis average of around 20% to more than 50%. News coverage of the *Motor vehicle* sector increased from around 10% to more than 20%. Together, these two sectors accounted for about three quarters of all news coverage in 2009. Other sectors that normally receive a substantial fraction of the news coverage naturally received a smaller share in this period. Both the *Printing and publishing* sector and the *Communications* sector saw their fraction of news coverage fall by approximately half during the crisis.

There are less dramatic movements of sectoral news coverage that are also likely to be driven by sectoral developments. The tech sectors discussed above experienced an increasing trend in news coverage in the 1990s and a sustained high level of news coverage in the decade since the financial crisis. The *Printing and publishing* sector, which includes Microsoft and Alphabet, saw a sharp and short-lived spike in news coverage during the dot-com boom of the late 1990s. We can also see that the *Transportation and warehousing* sector experienced a sharp spike in news coverage in 2016 - 2017. This is mostly driven Uber, which while classified as a transportation company, may also be considered part of the tech industry.

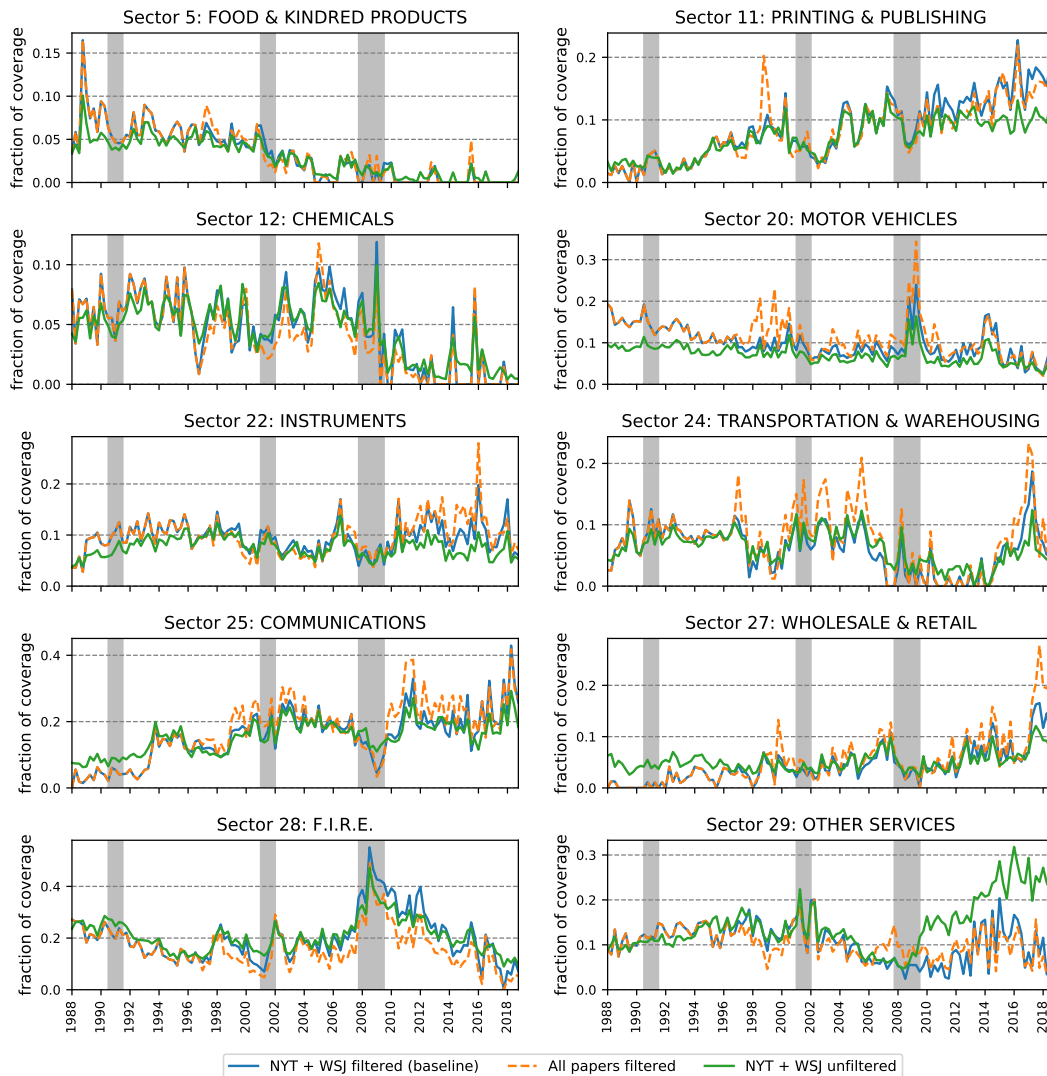


FIGURE 9. Sectoral news coverage over time for the 10 sectors that received the most coverage over the sample. Vertical axis measures fraction of news coverage the sector received.

The mirror image of the increase in news coverage of the tech sector in the last decade is also visible in Figure 9. Traditional sectors such as *Food and kindred products* and *Chemicals*, which both received substantial coverage throughout the 1990s, now receive a very small fraction of the total news coverage.

To investigate more formally if news coverage of a given sector is correlated with economic developments in that sector, we regress sectoral news coverage on observable economic outcomes in the same sectors. Table 2 displays the results of regressing news focus on the log differences in sectoral gross output, TFP and hours worked. If a bad economic outcome in a sector is considered newsworthy, this should manifest itself as negative coefficients on these variables. We also include the absolute values of the same three variables. If extreme

outcomes, either good or bad, in a sector are considered newsworthy, then this would result in a positive coefficient on these variables. The sample is annual and covers the period from 1988 to 2016, with annual news focus calculated as the simple average of the quarterly news focus in any given year. The table contains the result of these regressions for the ten sectors that receive the most coverage on average, i.e. the subset of sectors that typically receive at least some attention by the media.

The results show that news focus is systematically related to economic variables at the sector level. As an example, consider the *Finance, insurance and real estate* sector. For this sector, news focus is negatively correlated with changes in gross output and positively with the corresponding absolute value. Changes in output are positively correlated with news coverage in the *Instruments* sector, which includes computer hardware companies. We also find evidence for a relationship between sectoral news focus and productivity. For the *Transportation and warehousing* sector, we observe a highly significant relationship between the absolute values of changes in TFP, and for the *Motor vehicle* sector we find that news coverage was on average higher in times where TFP decreased. In general, many of the correlations we observe are significant, and the total number of such relationships is higher than what we would expect to see by random chance. The results are thus consistent with news media making state dependent reporting decisions, with both very positive and very negative sectoral outcomes being considered newsworthy.

5.3. Direct Media References to Specific Sectors. The news measures shown so far reflect how frequently companies affiliated with specific sectors are mentioned in the articles in our database. Thus, a sector whose companies are discussed frequently, is considered to receive a large amount of coverage. As discussed above, an important advantage of these measures is that the sector definitions used are consistent with those applied by the BEA when constructing sectoral output accounts.

An alternative approach to quantifying sectoral news focus is to consider articles that contain direct references to specific sectors. For example, a newspaper may directly refer to the conditions in the “auto sector” or the “motor vehicle industry”. Because such industry definitions primarily reflect how news editors tend to partition the economy in their reporting, they may not coincide exactly with the sector definitions used by economists and the BEA. Therefore, we would expect them to differ somewhat from the NAICS-based measures shown above. Nevertheless, explicit references to specific sectors may be an important part of sectoral news coverage and, to the extent possible, it is worth comparing sectoral news coverage measured using this approach to our baseline measure as a robustness check.

To identify and measure direct references to specific sectors consists, we first systematically search for word-pairs (2-grams and 3-grams) that contain the terms “industry” or “sector”. This allows us to construct a comprehensive list of expressions newspapers commonly use for direct references to sectors. Second, we group these expressions into meaningful categories and then quantify how frequently they occur. We find that the resulting sector definitions are not entirely consistent with their NAICS-based counterparts, but a number of sectors with frequent coverage are closely related and thus allow for a direct comparison. For example, we find that newspapers commonly make direct references to the auto sector, the food and tobacco industry, and the financial sector.

TABLE 2. Sectoral News Coverage and Observable Sector Properties

Sector	Statistic	Const.	Δl	Δy	Δz	$ \Delta l $	$ \Delta y $	$ \Delta z $
FOOD & KINDRED PRODUCTS	coeff	0.04***	0.21	1.06	-0.08	-1.34**	-0.71	0.76
	t-stat	3.82	0.65	1.35	-0.2	-2.64	-0.72	1.4
PRINTING & PUBLISHING	coeff	0.1***	-0.25	-0.34	0.27	-0.92	-0.02	0.33
	t-stat	3.96	-0.57	-0.5	0.92	-1.31	-0.04	0.75
CHEMICALS	coeff	0.04**	-0.09	0.03	0.46	0.24	-0.09	0.55
	t-stat	2.66	-0.31	0.17	1.49	0.44	-0.52	1.12
MOTOR VEHICLES	coeff	0.1***	-0.02	0.1	-0.77**	-0.23	0.17	0.18
	t-stat	10.79	-0.09	0.65	-2.21	-1.25	1.18	0.49
INSTRUMENTS	coeff	0.11***	-0.11	0.19*	-0.26	-0.06	0.05	-0.21
	t-stat	14.85	-0.69	2.01	-0.14	-0.28	0.52	-0.11
TRANSP. & WAREHOUSING	coeff	0.05**	0.24	-0.2	0.59	0.47	-0.51	1.17
	t-stat	2.62	0.48	-0.46	0.98	0.9	-1.64	1.29
COMMUNICATIONS	coeff	0.07*	-1.38*	1.97	-0.02	-0.62	0.03	0.61
	t-stat	1.78	-1.94	0.83	-0.03	-0.94	0.01	0.83
WHOLESALE & RETAIL	coeff	0.07***	1.25**	-0.72**	0.56	-1.87**	0.55**	-1.06
	t-stat	5.47	2.27	-2.28	0.98	-2.7	2.12	-1.53
F.I.R.E.	coeff	0.26***	0.35	-4.58***	1.9*	-0.67	3.3***	-2.26
	t-stat	6.56	0.48	-4.14	1.94	-0.72	3.09	-1.01
OTHER SERVICES	coeff	0.09***	-0.45	1.71	-2.06	-0.6	-0.31	0.09
	t-stat	3.47	-0.19	1.04	-1.19	-0.25	-0.18	0.04

Notes: The table shows results of multivariate regressions at the sector level. The sectors are those that are in the top ten in terms of average news coverage over the sample period. The independent variable is the fraction of news coverage received by a given sector. The dependent variables are a constant, log differences in labor, log differences in output, log differences in productivity as well as the corresponding absolute values. *, ** and *** denote statistical significance at the 10 percent, 5 percent and 1 percent levels, respectively.

Figure 10 shows the time series for these three sectors. Panel A displays the sector classification based on NAICS codes of the companies they contain. The corresponding sector coverage as measured using our vocabulary for direct references is shown in Panel B.⁸ Our principal finding here is that, while the behavior is not exactly identical, the key features are consistent. In terms of the ordering, the financial sector receives the most coverage under both definitions, followed by the auto sector and the food/tobacco industry, respectively. In

⁸The terms we use to identify each of these sectors are shown in the Data Appendix, together with time series of the remaining ones.

terms of the time-series behavior, we also observe clear similarities. While the food and tobacco sector receives relatively stable coverage over the sample period, both the auto sector and the financial industry are mentioned significantly more in the context of the 2008-2009 crisis.⁹

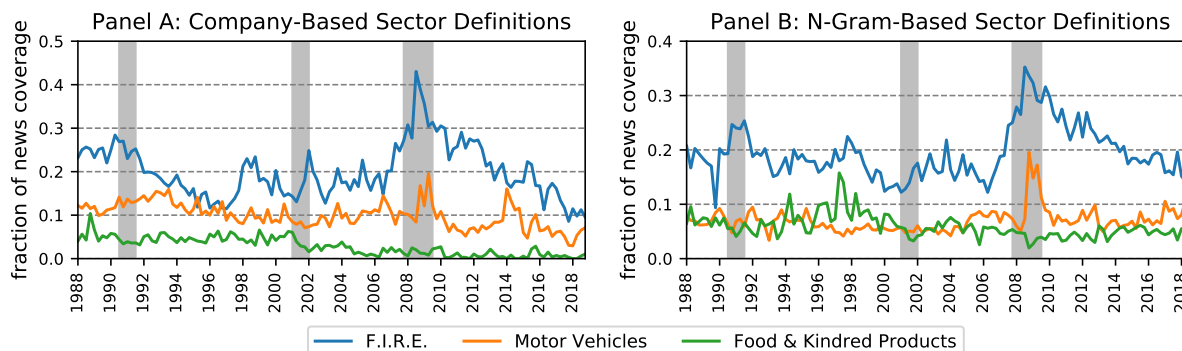


FIGURE 10. The Figure shows fractions of news coverage received by three different sectors. Sectors in Panel A are defined in terms of NAICS codes. Panel B reflects direct references to sectors.

6. AGGREGATE FLUCTUATIONS AND STATE-DEPENDENT REPORTING

The empirical evidence presented above shows that sectoral news coverage reflects the size of sectors and that it responds to sectoral developments. In this section we analyze the implications of these systematic reporting decisions for aggregate fluctuations. We calibrate the model to match the main features of sectoral news coverage and the input-output structure of the US economy.

6.1. Calibrating the production structure. The number of sectors is set to $n = 29$, with sectoral definitions listed in Table I. We use the BEA input-output accounts to calibrate the input shares parameters α_{ij} so that the relative importance of different intermediate inputs for each sector matches that in the data. We impose two simplifying assumptions: (i) A common labor share $(1 - \alpha) = 0.377$ and (ii) constant returns to scale so that $\sum_j \alpha_{ij} = \alpha$ across all sectors. The labor share is consistent with the average share of labor in gross output, after removing payments to capital. The aggregate labor elasticity parameter ν is set to 1. The log of sectoral productivity shocks are normally distributed white noise processes. The sectoral covariances are consistent with linearly detrended annual sectoral TFP data constructed from the BEA multi-factor productivity tables for the years 1987-2016.¹⁰ The average correlation of linearly detrended sectoral TFP in the sample is 0.08.

⁹The Data Appendix contains more details on how the n -gram based measures were constructed.

¹⁰The BEA tables in turn are based on the KLEMS accounting approach of Jorgensen, Ho and Samuels (2012).

6.2. Calibrating the news selection function. To calibrate the news selection function we need to specify (i) what makes a sector newsworthy and (ii) how many sectors news media report about in each period. In the baseline model, we use the weighted composite news selection function $\mathcal{S}_{|\omega|}$ that reports the largest weighted deviations of sectoral productivity shocks with $r = 1$ so that news media report on one sector in each period. The sector weights in the vector ω are chosen such that the average fraction of news coverage received by each sector in the model matches that in the news coverage data from Section 5. The news selection function also captures that sectors are more likely to be in the news when they experience large shocks.¹¹

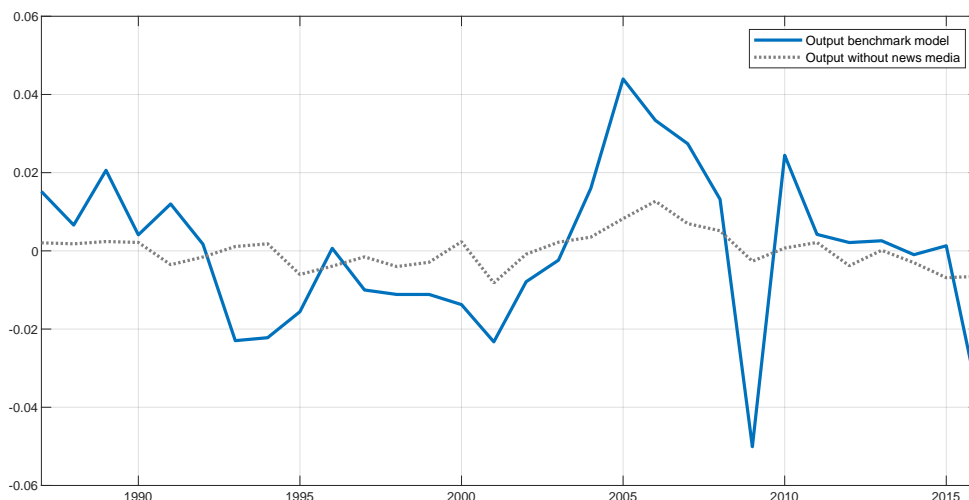


FIGURE 11. Output fluctuations around steady state with and without news media.

6.3. Output fluctuations with and without news media. Below, we present model simulations where we feed in historical sectoral productivity shocks from the years 1987 - 2016. This allows us to discuss specific historical episodes, using the realized cross-section of productivity to illustrate the mechanism through which news media affects aggregate outcomes. To analyze how news media affect aggregate fluctuations on average, we also compare unconditional population moments under different assumptions about what information is available to firms.

In the model, aggregate real value added output is equal to final consumption C as defined by (2.6). To quantify the importance of news media for aggregate fluctuations in C , we start by computing the log of aggregate output in the baseline model conditional on the historical productivity shocks. The deviation of the log of aggregate output from its mean is plotted in Figure 11 (blue solid line). We can compare this to output in the model without news media (grey dashed line).

¹¹The weights in ω are thus a function both of the average fraction of news coverage a sector receives and of the standard deviation of sectoral productivity shocks.

Output fluctuations are visibly larger in the baseline model relative to the model without news media. The population moments of the calibrated model also show that news media reporting contributes substantially to output volatility. The standard deviation of aggregate output is 2.3% when firms have access to reports by news media, but only 0.5% when they do not. News media affect output fluctuations not only by providing more information that individual firms respond to, but also by increasing coordination of labor input decisions across sectors. The average correlation of sectoral output in the baseline model is 0.82 compared to 0.10 in the model without news media.

The period of the Great Recession provides a particularly clear example of how news reporting changes the aggregate consequences of sectoral shocks. The baseline model predicts a severe recession in 2009, with aggregate output 5 percent below steady state. However, in the model without news reporting, output barely falls below its steady state level in the same period. Aggregate output in both models is conditional on the same sequence of productivity shocks, so this difference must be driven by differences in firms' beliefs, which themselves are entirely determined by the cross-section of sectoral productivity. This cross-section is illustrated in the left panel of Figure 12.

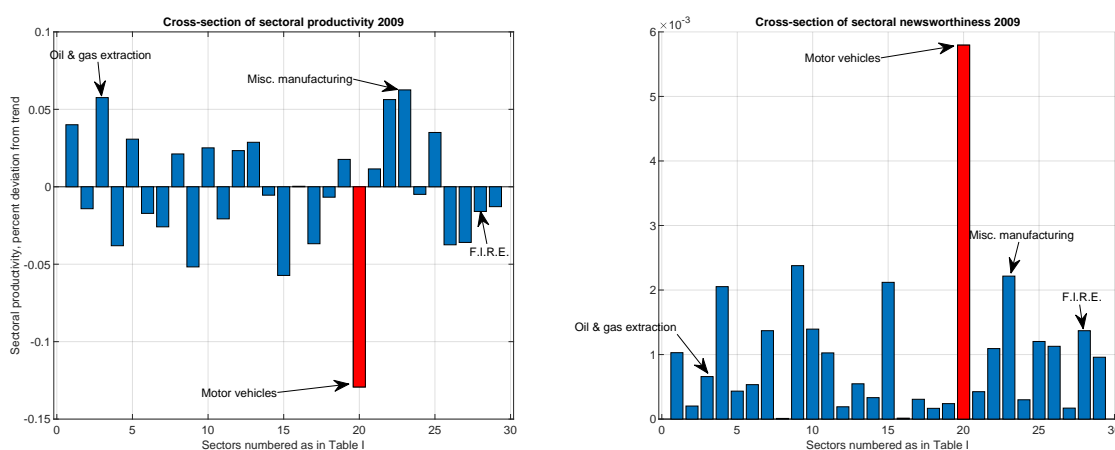


FIGURE 12. The left panel illustrates the cross-sectional profile of sector-specific (log) productivity z in 2009. The right panel illustrates the cross-sectional newsworthiness of sectoral productivity $|\omega \circ z|$.

The average sectoral productivity in 2009 is only slightly negative, but as shown in the figure, the *Motor vehicles* sector experienced a very large negative productivity shock in that year (red bar). This shock is also what was reported on by news media in the model. Other sectors, such as *Oil and gas extraction* and *Miscellaneous manufacturing* experienced substantial positive productivity shocks in the same period. However, these were not reported by the news media. The sector that news media did report on, and that firms across all sectors therefore knew about, experienced a large negative shock. Firms across all sectors therefore hired less labor than they would have, had they observed only their own productivity.

Moreover, the effect of this common pessimism is amplified by the strategic complementarity embedded in the labor demand function (2.13), as firms anticipate lower labor demand, and therefore lower demand for their output, in other sectors. Hence, media reporting drives an aggregate contraction both directly because it affects all firms' information and, indirectly, because what is reported is common knowledge among firms.

Figure 12 also illustrates the relative newsworthiness of the different sectors in 2009 according to the calibrated news selection function. The right panel of the figure shows the absolute values of the cross-section productivity shocks weighted by ω . It is clear that not only is the sectoral productivity shock hitting the *Motor vehicle* industry the largest in absolute terms, it is also by far the most newsworthy. The right panel also illustrates a limitation of the simple model where news media report on a single sector, and where newsworthiness is based only on productivity outcomes. We know from the data that *Finance, insurance and real estate* actually received more news coverage than *Motor vehicles* in 2009. However, in the model, the finance sector is not the most newsworthy sector in that period.

The model's predictions for the 2009 episode thus highlight both one of its strengths and a dimension in which it is too simple. The mechanism is strong enough to replicate the depth of the Great Recession, but it is somewhat unsatisfying to have a model that does so without any special role for the financial sector. However, given that the mechanism in the model relies on news media reporting unrepresentatively bad news in 2009, a richer model that implied additional reporting on the financial sector in that period would likely generate a recession at least as deep as in the baseline model.

6.4. Output fluctuations in baseline and full information model. One reason that the baseline model generates a large recession in 2009 while the model without news media does not is that firms in every sector know about the bad news coming out of the motor vehicle sector. In the model without news media, only firms in the motor vehicle sector are aware of this. If firms could observe productivity in every sector, they would also all know about the motor vehicle sector. However, since the cross-section of productivity in 2009 was not particularly bad overall, the full information model does not generate a strong recession in that period. This is illustrated in Figure 13. The full information model generates only a mild recession in 2009, with output about 1.5% below average. The reason the partial information model generates a strong recession and the full information does not is that the sector shock reported by news media in 2009 is unrepresentative of the shocks the economy experienced as a whole.¹²

In Figure 13 we also plot actual (demeaned) output growth for the same period. Both models are simple and highly stylized, and we should not expect either of them to fit actual data closely. The correlation between actual output growth and output in the full information model is 0.24 and improves only marginally in the baseline partial information model.

6.5. Time varying media focus and common non-productivity shocks. Atalay (2017) uses a multi-sector model that, unlike our model, includes capital as a production factor and allows for a richer specification of consumption and production elasticities. Using a filter

¹²We set the labor elasticity parameter ν equal to 0.65 in the full information model. The baseline model and the full information model then has the same unconditional standard deviation of aggregate output.

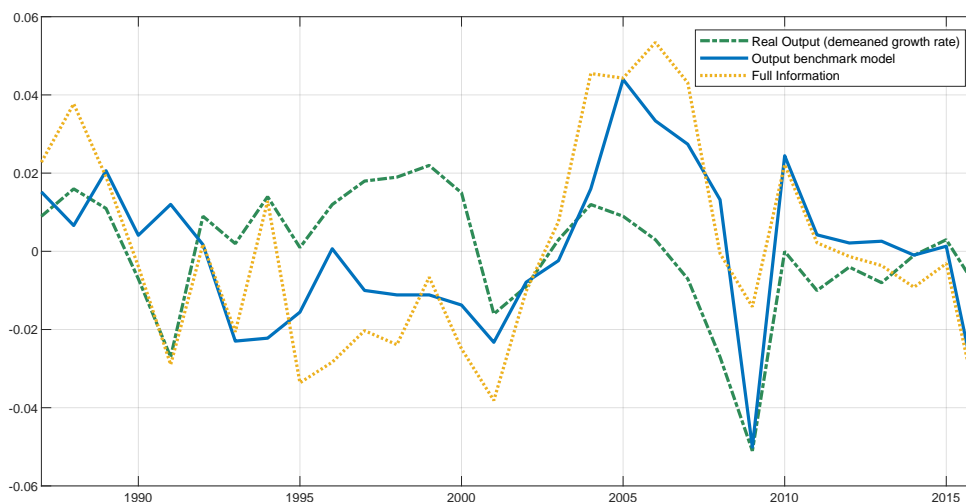


FIGURE 13. Output fluctuations around steady state in baseline model (blue solid line), full information model (yellow dotted line) and historical demeaned GDP growth (green dash-dot line).

implied by a log-linearized equation from his model, he estimates that for realistic values of elasticities of substitution, sectoral productivity shocks explain approximately 80% of the variance of aggregate output. The remaining variance is attributed to common non-productivity shocks.

In our model, sectoral productivity is the only source of exogenous variation, but the relationship between sectoral productivities and output is strongly non-linear: When a sector is in the news, productivity in that sector has a bigger impact on aggregate output than it does when that sector is not in the news. Since the impact of a sectoral shock depends on whether it is reported, a researcher applying a filter that imposes a constant (log-) linear relationship between sectoral productivity and output would conclude that sectoral productivity cannot explain sectoral or aggregate output.

Imposing linearity on data generated by our model would lead a researcher to conclude that output is driven in part by aggregate (non-productivity) shocks. To quantify how much of aggregate output fluctuations that incorrectly would be attributed to common non-productivity shocks by a linear filter, we first generate a long sample (50,000 periods) from our baseline model. We then run the regression

$$c_t = \gamma + \sum_{i=1}^n \beta_i z_{i,t} + \epsilon_t. \tag{6.1}$$

on the generated sample, where c_t is the simulated time-series of the log of aggregate value added. The fitted value of c_t from this regression is the linear projection of log output on the log of sectoral productivity shocks and thus represents the best possible fit that can be achieved by any linear model.

The regression coefficients from (6.1) imply a decomposition of model generated aggregate output into one component that can be predicted by a linear combination of sectoral productivity shocks and a residual aggregate component. We then use the regression coefficients computed from the simulated sample to decompose the path of aggregate output predicted by the baseline model conditional on the actual productivity shocks from the BEA. In figure 14 we plot model output along with the fitted values and the residuals from this decomposition. The residual, which by construction is orthogonal to any linear combination of sectoral productivity, accounts for a substantial fraction of the variation in output in the figure. The deep recession in 2009 is almost entirely explained by the residual, which accounts for about four percentage points, or 80%, of the deviation of output below its mean. In terms of population moments, the residual accounts for one third of the variation in aggregate output. An observer studying output data generated by our model through the lens of a linear filter would conclude that aggregate non-productivity shocks are needed to explain the observed fluctuations in aggregate output. The model can thus qualitatively account for the finding of Atalay (2017) without any additional shocks.

This exercise also demonstrates that time-varying sectoral media focus produces demand-driven business cycle fluctuations that share qualitative properties with the *Main Business Cycle* (MBC) shock identified by Angeletos, Collard and Dellas (2019). They find that a shock that is orthogonal to productivity, but increases output, employment and consumption is responsible for a large fraction of business cycle fluctuations.

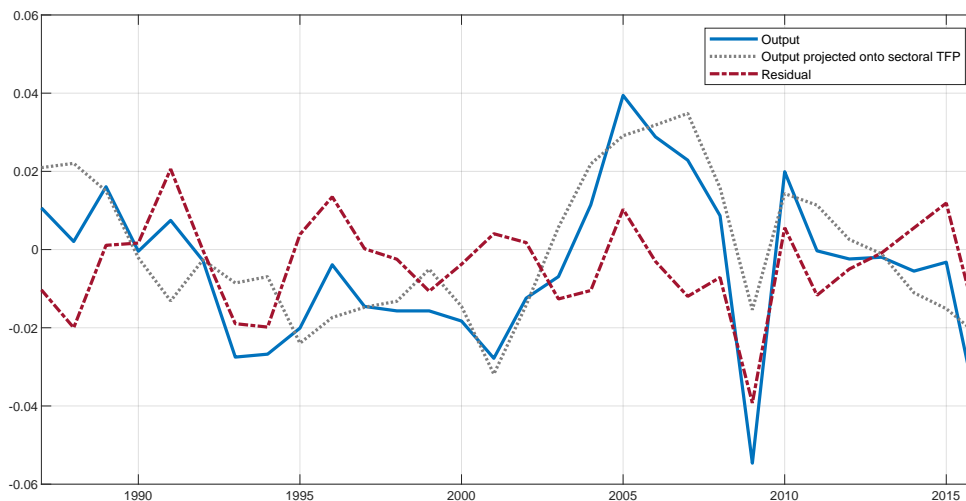


FIGURE 14. Decomposing fluctuations in log of aggregate output. The dotted grey line is the projection of c_t onto the sectoral productivity shocks. The residual is the component of aggregate output that cannot be expressed as a linear function of the cross-section of productivity shocks.

6.6. Selection bias and inference from state-dependent reporting decisions. State-dependent reporting decisions affect aggregate output through two distinct channels. First,

the selection bias towards more extreme shocks increases the standard deviation of firms' labor input decisions. Second, as shown in Section 2, the state dependence of reporting decisions allows firms to make inference not only about those shocks that are reported by news media, but also about those shocks that news media chose not to report.

To quantify the importance of the state-dependent news reporting in the model, we solve the model under the assumption that news media randomly choose which sector to report on. The population standard deviation of output in this version of the model is 1.3%, or about one half of that in the baseline model. The news selection function in the baseline version also weighs larger sectors more when evaluating newsworthiness. However, the effect on output of this systematic bias towards reporting on larger sectors is relatively small. The standard deviation of output in the model when sectors simply report the productivity shock with the largest (unweighted) absolute deviation from its mean is only marginally lower than in the baseline model.

We also compute how much output would change if firms did not take into account the state dependence of reporting decisions when forming beliefs about non-reported sectors. The effect of time-variation in conditional beliefs on output through this channel is very small, accounting for less than 0.1 percentage point of the standard deviation of output. This is due to both the weak sectoral correlations in the model and the absence of strong non-linearities in the labor input decision.

6.7. Model simplicity and model-free evidence. Above, we proposed a novel mechanism through which news media may affect the economy. In order to keep the model tractable, we have made several simplifying assumptions and it is worthwhile discussing some of them in more detail. One such assumption is that the news selection function of media is exogenous. In reality, the market for news is likely to be responsive to incentives just like any other market. One may therefore reasonably ask whether our analysis is robust to abstracting from any endogenous decisions by news suppliers and news consumers. The news data that we use to calibrate the news selection function in the model are the observed equilibrium outcome in the market for news. For our purposes, what matters is the joint distribution of economic outcomes and reporting decisions. Arguably, none of the exercises we perform imply that an endogenous change in this joint distribution is likely. There is thus little reason to believe that our results are sensitive to this assumption.

In practice, the news media market includes a diverse set of outlets with different focus and intended audiences. Again, for reasons of tractability, we treat news media as monolithic in the model. One implication of this is that anything the news media reports in the model is common knowledge among firms. This means that the news reports in the model correspond more closely to news reports about major events, which as shown in Nimark and Pitschner (2019) are more likely to result in homogenous news coverage across different outlets, that in turn are more likely to be closer to common knowledge.

While simplicity brings the benefits of tractability and transparency, it also imposes a lot of structure. To address concerns that our results are unduly driven by the assumptions we have made to keep the model tractable, we end this section by presenting some evidence that, while derived without direct reference to our model, is nevertheless consistent with its main mechanism.

6.7.1. *News coverage and sectoral correlations.* If the mechanism in the model is relevant in reality, we would expect that output in sectors that are over-represented in the news are more strongly correlated with aggregate output relative to less-reported sectors of the same size. We therefore first run a regression of average sectoral news coverage on Domar weights

$$\frac{1}{T} \sum_{t=1}^T f_{i,t} = \gamma_f + \beta_f \lambda_i + \varepsilon_i^f \quad (6.2)$$

where $f_{i,t}$ denotes the fraction of news coverage received by sector i in period t . A positive residual ε_i^f implies that the sector is over-represented in the news relative to its economic size. (The same information is contained in Figure 6.) We then run the regression

$$\rho_{i,y} = \gamma_\rho + \beta_\rho \lambda_i + \varepsilon_i^\rho \quad (6.3)$$

where $\rho_{i,y}$ denotes the correlation between gross output growth in sector i and aggregate output growth. A positive residual ε_i^ρ indicates that sector i is more strongly correlated with aggregate output than what would be implied by its economic size alone.

The correlation of the residuals from the two regressions is 0.21, suggesting that sectors that are over-represented in the news are indeed also more strongly correlated with aggregate output, as predicted by our model.

6.7.2. *News coverage-weighted productivity and aggregate output.* Our model predicts that productivity in a given sector has a bigger impact on aggregate output when the sector in question is in the news. We therefore compute two news-weighted aggregate productivity series as follows

$$z_t^f \equiv \sum_{i=1}^n f_{i,t} z_{i,t}, \quad z_t^{f,\lambda} \equiv \sum_{i=1}^n \left(f_{i,t} - \frac{\lambda_i}{\sum_{i=1}^n \lambda_i} \right) z_{i,t}. \quad (6.4)$$

The measure z_t^f simply weighs sectoral productivity in period t by the fraction of news coverage a sector received in that period. The correlation between z_t^f and aggregate output growth is 0.46. Given the strong correlation between news coverage and the size of a sector, this positive correlation is unsurprising. The second measure, $z_t^{f,\lambda}$ therefore weighs sectoral productivity by the fraction of news coverage in period t that is not simply a reflection of the sector's size. The correlation between the second measure $z_t^{f,\lambda}$ and aggregate output growth is 0.26, again consistent with the idea that a sector has a bigger impact on aggregate output when it is widely reported on by news media.

In the analysis above, we argued that time-varying media focus could generate output fluctuations that a linear filter would attribute to common non-productivity shocks. We therefore also compute the correlation between the weighted productivity measure $z_t^{f,\lambda}$ and the time series of common shocks from the Atalay (2017) filter. The correlation between $z_t^{f,\lambda}$ and the common shocks produced by Atalay's high elasticity calibration is 0.32. The high-elasticity calibration attributes more of the variance to common shocks than Atalay's low elasticity calibration. The correlation of $z_t^{f,\lambda}$ with the common shocks extracted using the latter calibration is also positive, but substantially lower at 0.08.

7. CONCLUSIONS

In this paper we have demonstrated that time varying sectoral media focus can generate aggregate fluctuations that are orthogonal to productivity, even in a model where the only source of exogenous variation are sectoral TFP shocks. That aggregate output fluctuations are partially orthogonal to productivity has been well-documented at least since the early 1990s, e.g. Hall (1993), Blanchard (1993) and Cochrane (1994). However, no consensus has emerged regarding the causes of the non-productivity related fluctuations. Inspired by Lucas' (1977) statement that "business cycles are all alike", Angeletos, Collard and Dellas (2019) use flexible VAR methods to document the properties of what they label the *Main Business Cycle* (MBC) shock. This shock appears to be responsible for most of the business cycle variation in several key macroeconomic variables.

While our model is too stylized to account for all of the dynamics associated with MBC shocks, many of our findings are consistent with them. Like that shock, time varying sectoral media focus generates fluctuations that are orthogonal to aggregate productivity and positively correlated with output, consumption and employment. Angeletos *et al* (2019) further argue that the facts they document are consistent with fluctuations in firms' beliefs about the demand for their products. We have provided a theory that can explain why the demand expectations of firms across different sectors move together. Discussing financial markets, Shiller (2001) writes that "*Significant market events generally occur only if there is similar thinking among large groups of people, and the news media are essential vehicles for the spread of ideas.*" We thus argue here that the same logic applies to macroeconomic fluctuations.

In this paper we have also proposed a conceptually new approach to model incomplete information. Firms in our model receive accurate but partial information from news media, and what media report depends deterministically on the cross-section of productivity shocks. By constructing a novel data set of sectoral news coverage, we are able to discipline the reporting decisions of news media in the model. This approach avoids introducing exogenous noise shocks and provides a tight link between beliefs, developments in the real economy, and observable patterns in news coverage.

REFERENCES

- [1] Acemoglu, Daron, Vasco M. Carvalho, Asuman Ozdaglar, and Alireza Tahbaz-Salehi, 2012. The network origins of aggregate fluctuations. *Econometrica* 80, no. 5 (2012): 1977-2016.
- [2] Acemoglu, Daron, Ufuk Akcigit, and William Kerr, 2015. Networks and the macroeconomy: An empirical exploration. No. w21344. National Bureau of Economic Research.
- [3] Angeletos, G.M. and La'o, J., 2010. Noisy business cycles. *NBER Macroeconomics Annual*, 24(1), pp.319-378.
- [4] Angeletos, George-Marios, and Jennifer La'O, 2013. Sentiments. *Econometrica* 81, no. 2 : 739-779.
- [5] Angeletos, G.M., Collard, F. and Dellas, H., 2018. Quantifying confidence. *Econometrica*, 86(5), pp.1689-1726.
- [6] Angeletos, G.M., Collard, F. and Dellas, H., 2019. Business Cycle Anatomy. mimeo, MIT.
- [7] Arif, S. and Lee, C.M., 2014. Aggregate investment and investor sentiment. *The Review of Financial Studies*, 27(11), pp.3241-3279.
- [8] Atalay, E., 2017. How important are sectoral shocks?. *American Economic Journal: Macroeconomics*, 9(4), pp.254-80.

- [9] Baqaee, David, and Emmanuel Farhi. 2019. The Macroeconomic Impact of Microeconomic Shocks: Beyond Hulten's Theorem.' *Econometrica* 87 (4): 1155-1203.
- [10] Blanchard, O., 1993. Consumption and the Recession of 1990-1991. *The American Economic Review*, 83(2), pp.270-274.
- [11] Blanchard, O.J., L'Huillier, J.P. and Lorenzoni, G., 2013. News, noise, and fluctuations: An empirical exploration. *American Economic Review*, 103(7), pp.3045-70.
- [12] Blinder, A.S. and Krueger, A.B., 2004. What Does the Public Know about Economic Policy, and How Does It Know It?. *Brookings Papers on Economic Activity*, pp.327-397.
- [13] Buchen, Teresa (2014). News Media, Common Information, and Sectoral Comovement, *Beiträge zur Jahrestagung des Vereins für Socialpolitik 2014: Evidenzbasierte Wirtschaftspolitik - Session: Applied Microeconometrics*, No. C21-V3
- [14] Carroll, C.D., 2003. Macroeconomic expectations of households and professional forecasters. *the Quarterly Journal of economics*, 118(1), pp.269-298.
- [15] Carvalho, V.M., 2010. Aggregate Fluctuations and the Network Structure of Intersectoral Trade.
- [16] Carvalho, Vasco, and Xavier Gabaix, 2013. The great diversification and its undoing. *The American Economic Review* 103, no. 5: 1697-1727.
- [17] Carvalho, Vasco M., and Basile Grassi. 2019. Large Firm Dynamics and the Business Cycle. *American Economic Review*, 109 (4): 1375-1425.
- [18] Chahrour, R. and Ulbricht, R., 2019. Information-driven business cycles: A primal approach. Available at SSRN 3135375.
- [19] Chan, J. and Suen, W., 2008. A spatial theory of news consumption and electoral competition. *The Review of Economic Studies*, 75(3), pp.699-728.
- [20] Cochrane, John H. 1994. Shocks. *Carnegie-Rochester Conference series on public policy*, vol. 41, pp. 295-364..
- [21] Curtin, R., 2007. What US Consumers Know About Economic Conditions. *Statistics, Knowledge and Policy* 2007, p.153.
- [22] Eisensee, T. and Strömberg, D., 2007. News droughts, news floods, and US disaster relief. *The Quarterly Journal of Economics*, 122(2), pp.693-728.
- [23] Eisner, Robert., 1978. Factors in business investment. NBER Books.
- [24] Flaxman, S., Goel, S. and Rao, J.M., 2016. Filter bubbles, echo chambers, and online news consumption, *Public Opinion Quarterly*, 80(S1), pp.298-320.
- [25] Foerster, A.T., Sarte, P.D.G. and Watson, M.W., 2011. Sectoral versus aggregate shocks: A structural factor analysis of industrial production. *Journal of Political Economy*, 119(1), pp.1-38.
- [26] Gennaioli, N., Y. Ma, and A. Shleifer, 2015. Expectations and investment." In *NBER Macroeconomics Annual 2015*, Volume 30. University of Chicago Press, 2015.
- [27] Gentzkow, M., Shapiro, J.M. and Stone, D.F., 2015. Media bias in the marketplace: Theory. In *Handbook of media economics* (Vol. 1, pp. 623-645). North-Holland.
- [28] Greenwood, Robin, and Samuel G. Hanson, 2015. Waves in Ship Prices and Investment", *Quarterly Journal of Economics*, 130(1), pp.55-109.
- [29] Hall, R.E., 1993. Macro Theory and the Recession of 1990-1991. *The American Economic Review*, 83(2), pp.275-279.
- [30] Harcup, T. and O'Neill, D., 2017. What is news? News values revisited (again). *Journalism studies*, 18(12), pp.1470-1488.
- [31] Harrington, D.E., 1989. Economic news on television: The determinants of coverage. *Public Opinion Quarterly*, 53(1), pp.17-40.
- [32] Horvath, M., 2000. Sectoral shocks and aggregate fluctuations. *Journal of Monetary Economics*, 45(1), pp.69-106.
- [33] Jorgenson, D.W., Ho, M.S. and Samuels, J., 2012, August. A prototype industry-level production account for the United States, 1947-2010. In *Second World KLEMS Conference*, Harvard University (Vol. 9).

- [34] Lamla, Michael J., Lein, Sarah M., Sturm, Jan-Egbert, 2007, News and sectoral comovement, KOF working papers Konjunkturforschungsstelle, Eidgenossische Technische Hochschule Zurich, No. 183,
- [35] Larsen, V.H., Thorsrud, L.A. and J. Zhulanova, 2019, News-driven inflation expectations and information rigidities, working paper, Bank of Norway.
- [36] Levenshtein, V.I., 1966, February. Binary codes capable of correcting deletions, insertions, and reversals. In Soviet physics doklady (Vol. 10, No. 8, pp. 707-710).
- [37] Lorenzoni, Guido, 2009. A Theory of Demand Shocks, *American Economic Review*, American Economic Association, vol. 99(5), pages 2050-84, December.
- [38] Melosi, L., 2016. Signalling effects of monetary policy. *The Review of Economic Studies*, 84(2), pp.853-884.
- [39] Nimark, K.P., 2014,. Man-bites-dog business cycles, *American Economic Review* 104, pp2320–2367.
- [40] Nimark, K.P. and Pitschner, S., 2019. News Media and Delegated Information Choice. *Journal of Economic Theory*.
- [41] Shoemaker, P.J. and Vos, T.P., 2009. Gatekeeping theory. Routledge.
- [42] Soroka, S.N., 2012. The gatekeeping function: Distributions of information in media and the real world, *Journal of Politics*, 74(2), pp.514-528.
- [43] Strömberg, D., 2004. Mass media competition, political competition, and public policy. *The Review of Economic Studies*, 71(1), pp.265-284.
- [44] Tian, Can, 2019, Input-Output Linkages in Pigouvian Industrial Fluctuations, *working paper* University of North Carolina.
- [45] Veldkamp, L. and Wolfers, J., 2007. Aggregate shocks or aggregate information? Costly information and business cycle comovement. *Journal of Monetary Economics*, 54, pp.37-55.

APPENDIX A. SOLVING THE MODEL

In this appendix we describe how to solve the model. We first present the primitives of the model and the optimality conditions of firms and households. We then derive the expressions that use to solve the model numerically and describe an algorithm for doing so. Throughout, we use the notation conventions that lower case letters denote the log of the corresponding uppercase letter. Bold letters and symbols denote vectors.

A.1. Model primitives. Households maximize the utility function

$$\max_{x_1, \dots, x_n, L_i} C - \frac{L^{1+1/\nu}}{1 + 1/\nu} \tag{A.1}$$

where L is labor supply and C is the final consumption good. The final good C is a Cobb-Douglas aggregate of sector-specific goods C_i given by

$$C = \prod_i C_i^{\frac{1}{n}}. \tag{A.2}$$

Sector i produces quantity Q_i of good i using the Cobb-Douglas production function

$$Q_i = Z_i \left(\prod_j X_{ij}^{\alpha_{ij}} \right) L_i^{1-\alpha} \tag{A.3}$$

where Z_i is a sector-specific productivity shock. Total output in sector i can be used either for the final consumption C_i or as an intermediate input X_{ij} in sector j so that

$$C_i + \sum_j X_{ji} = Q_i. \tag{A.4}$$

Sector specific labor demand L_i adds up to total labor demand L , i.e.

$$\sum_i L_i = L. \quad (\text{A.5})$$

Households spend the income they receive from working and from owning the firms so that

$$C = WL + \Pi. \quad (\text{A.6})$$

Under full information, profits Π are zero of course. When firms face information frictions, however, informational errors may lead Π to be non-zero.

A.2. Optimality conditions. Households supply labor until marginal disutility of working equals marginal utility of consuming wage

$$W = L^{\frac{1}{\nu}}. \quad (\text{A.7})$$

The intermediate goods are combined into the final consumption good using Cobb-Douglas aggregator (A.2). The optimal expenditure on good i , holding total expenditure PC fixed, is then given by

$$P_i C_i = \frac{1}{n} PC \quad (\text{A.8})$$

We normalize the price of the aggregate good to 1 and use (A.2) to replace C to get

$$P_i = \frac{1}{n} \frac{1}{C_i} \prod_j C_j^{\frac{1}{n}}. \quad (\text{A.9})$$

Labor markets are competitive, so households earn the same wage in every sector. Since firms choose labor before observing all prices, firm choose labor inputs so that expected marginal cost equals expected marginal product

$$E[W \mid \Omega_i] = (1 - \alpha) \frac{E[P_i Q_i \mid \Omega_i]}{L_i}. \quad (\text{A.10})$$

Marginal product of intermediate input j equals its marginal cost so that

$$P_j = \frac{\alpha_{ij}}{X_{ij}} P_i Q_i \quad (\text{A.11})$$

holds in equilibrium.

A.3. Solving for L_i as function of \mathbf{L} and \mathbf{Z} . The only decision taken under incomplete information is a firm's decision of how much labor to employ. To solve the model, we need to be able to express that choice as a function of a firm's expectations about the exogenous sector-specific productivity shocks Z_i and the labor input choices of firms in other sectors.

Start by substituting in the optimal demand for intermediate inputs X_{ij} into the production function (A.3) using (A.11) to get

$$Q_i = Z_i \left(\prod_j \left(\alpha_{ij} \frac{P_i Q_i}{P_j} \right)^{\alpha_{ij}} \right) L_i^{1-\alpha}. \quad (\text{A.12})$$

Use that $\sum_{j=1}^n \alpha_{ij} = \alpha$ to compute $\prod_j (P_i Q_i)^{\alpha_{ij}} = (P_i Q_i)^\alpha$ and move this term outside the product in the parenthesis, so that

$$Q_i = Z_i (P_i Q_i)^\alpha \left(\prod_j \left(\frac{\alpha_{ij}}{P_j} \right)^{\alpha_{ij}} \right) L_i^{1-\alpha}. \quad (\text{A.13})$$

Divide both sides by Q_i^α

$$Q_i^{1-\alpha} = Z_i P_i^\alpha \left(\prod_j \left(\frac{\alpha_{ij}}{P_j} \right)^{\alpha_{ij}} \right) L_i^{1-\alpha} \quad (\text{A.14})$$

and multiply by $P_i^{1-\alpha}$

$$(P_i Q_i)^{1-\alpha} = Z_i P_i \left(\prod_j \left(\frac{\alpha_{ij}}{P_j} \right)^{\alpha_{ij}} \right) L_i^{1-\alpha}. \quad (\text{A.15})$$

Define gross sales V_i as

$$V_i \equiv P_i Q_i, \quad (\text{A.16})$$

take logs of both sides of (A.15)

$$(1 - \alpha) v_i = z_i + p_i + (1 - \alpha) l_i + \sum_j \alpha_{ij} (\log(\alpha_{ij}) - p_j). \quad (\text{A.17})$$

and rearrange the resulting expression to get

$$(1 - \alpha) (v_i - l_i) - z_i - \sum_j \alpha_{ij} \log(\alpha_{ij}) = p_i - \sum_j \alpha_{ij} p_j. \quad (\text{A.18})$$

Define the input-output matrix IO so that the typical i^{th} row and j^{th} element is α_{ij} . We can then write the r.h.s. of (A.18) as

$$p_i - \sum_j \alpha_{ij} p_j = p_i - IO_i \mathbf{p} \quad (\text{A.19})$$

where IO_i is the i^{th} row of IO and use that (A.18) has to hold for every sector i so that

$$[(1 - \alpha) (\mathbf{v} - \mathbf{l}) - \mathbf{z} - \boldsymbol{\tau}] = (I - IO) \mathbf{p} \quad (\text{A.20})$$

where $\boldsymbol{\tau}$ is a vector with the i^{th} element given by $\sum_j \alpha_{ij} \log(\alpha_{ij})$.

Now, use that

$$P_i Q_i = \sum_j \alpha_{ji} P_j Q_j + P_i C_i \quad (\text{A.21})$$

and

$$P_i C_i = \frac{1}{n} C \quad (\text{A.22})$$

together with definition (A.16) of V_i to get

$$V_i = \sum_j \alpha_{ji} V_j + \frac{1}{n} C \quad (\text{A.23})$$

Since this has to hold for each i we get

$$V = (I - IO')^{-1} \frac{C}{n} \times \mathbf{1} \quad (\text{A.24})$$

Define the vector Λ as

$$\Lambda = (I - IO')^{-1} \frac{1}{n} \times \mathbf{1} \quad (\text{A.25})$$

and $\lambda = \log(\Lambda)$ with typical element λ_i . We then have

$$\mathbf{p} = (I - IO)^{-1} [(1 - \alpha)(\lambda + c \times \mathbf{1} - \mathbf{1}) - \mathbf{z} - \tau] \quad (\text{A.26})$$

or equivalently

$$\mathbf{p} = (I - IO)^{-1} [(1 - \alpha)(\lambda - \mathbf{1}) - \mathbf{z} - \tau] + (I - IO)^{-1} [(1 - \alpha)c \times \mathbf{1}] \quad (\text{A.27})$$

so that

$$P_i = \exp(\Gamma_i [(1 - \alpha)(\lambda - \mathbf{1}) - \mathbf{z} - \tau] + \Gamma_i (1 - \alpha)c \times \mathbf{1}) \quad (\text{A.28})$$

where Γ_i is the i^{th} row of the Leontief inverse $(I - IO)^{-1}$.

A.4. Final expressions. To solve the model we need to compute the optimal labor demand

$$L_i = (1 - \alpha) \frac{E[P_i Q_i | \Omega_i]}{E[W | \Omega_i]} \quad (\text{A.29})$$

as a function of expected labor inputs and productivity in every sector $L_j : j \in \{1, 2, \dots, n\}$, \mathbf{z} . To that end, first use that

$$\begin{aligned} P_i Q_i &= \lambda_i C \\ W &= \left(\sum L_i \right)^{\frac{1}{\nu}}. \end{aligned}$$

We then need to find an expression of C as a function of L_i and \mathbf{z} . Combining and rearranging the following three equations

$$C = \prod_i C_i^{\frac{1}{n}} \quad (\text{A.30})$$

$$C_i = C \frac{1}{n} \frac{1}{P_i} \quad (\text{A.31})$$

$$P_i = \exp(\Gamma_i [(1 - \alpha)(\lambda - \mathbf{1}) - \mathbf{z} - \tau] + \Gamma_i (1 - \alpha)c \times \mathbf{1}) \quad (\text{A.32})$$

gives

$$C_i = \frac{1}{n} \frac{C^{1 - (1 - \alpha)\Gamma_i \times \mathbf{1}}}{\exp(\Gamma_i [(1 - \alpha)(\lambda - \mathbf{1}) - \mathbf{z} - \tau])}. \quad (\text{A.33})$$

Substitute into (A.30) to get

$$C = \prod_i \left(\frac{1}{n} \frac{C^{1-(1-\alpha)\Gamma_i \times \mathbf{1}}}{\exp(\Gamma_i [(1-\alpha)(\lambda-1) - \mathbf{z} - \tau])} \right)^{\frac{1}{n}} \quad (\text{A.34})$$

$$= \prod_i \left(\frac{1}{n^{\frac{1}{n}}} \frac{C^{(1-(1-\alpha)\Gamma_i \times \mathbf{1})\frac{1}{n}}}{\exp(\Gamma_i [(1-\alpha)(\lambda-1) - \mathbf{z} - \tau])^{\frac{1}{n}}} \right) \quad (\text{A.35})$$

$$= C^{(1-\frac{1}{n}(1-\alpha)\Sigma_i \Sigma_j \Gamma_{ij})} \prod_i \left(\frac{1}{n^{\frac{1}{n}}} \frac{1}{\exp(\Gamma_i [(1-\alpha)(\lambda-1) - \mathbf{z} - \tau])^{\frac{1}{n}}} \right) \quad (\text{A.36})$$

so that

$$C = \prod_i \left(\frac{1}{n^{\frac{1}{n}} \exp(\Gamma_i [(1-\alpha)(\lambda-1) - \mathbf{z} - \tau])^{\frac{1}{n}}} \right)^{\frac{n}{(1-\alpha)\Sigma_i \Sigma_j \Gamma_{ij}}}. \quad (\text{A.37})$$

We then have the desired expression

$$L_i = (1-\alpha) \lambda_i \frac{E \left[\prod_i \left(\frac{1}{n^{\frac{1}{n}} \exp(\Gamma_i [(1-\alpha)(\lambda-1) - \mathbf{z} - \tau])^{\frac{1}{n}}} \right)^{\frac{n}{(1-\alpha)\Sigma_i \Sigma_j \Gamma_{ij}}} \mid \Omega_i \right]}{E_i \left[(\sum L_i)^{\frac{1}{\nu}} \mid \Omega_i \right]}. \quad (\text{A.38})$$

A.5. Numerical solution algorithm. We solve the model by evaluating the conditional expectation in (A.38) using a simulation based MCMC method. The simulation is initialized by solving the model under full information for T draws from the process for sectoral productivities.

The algorithm is described by the following steps.

- (1) Take S draws from the distribution of the vector of sectoral productivity shocks z .
- (2) For the first T draws of z find the full information equilibrium vector \mathbf{L} .
- (3) For draw $T+s$
 - (a) Apply the news selection function \mathcal{S} to find the vectors $\mathbf{r}(\mathbf{z}_{T+s}), \mathbf{s}(\mathbf{z}_{T+s})$.
 - (b) Compute $L_{i,T+s}$ for every i by evaluating the conditions expectation (A.38) where $\Omega_i = \{z_{i,T+s}, \mathbf{r}(\mathbf{z}_{T+s}), \mathbf{s}(\mathbf{z}_{T+s})\}$.
- (4) Repeat steps 2-4 until convergence.

The conditional expectation in Step 4 is computed by first identifying the set of indices in the chain up to draw $T+s$ such that $\mathbf{s}(\mathbf{z}_{T+s}) = \mathbf{s}(\mathbf{z}_{T+t}) : t < s$, i.e the set of draws for which the news selection function chose the same sector to report on as in draw $T+s$. Within this set, find the K draws that minimizes the distance $(z_{i,T+s} - z_{i,T+t})^2 + (z_{j,T+s} - z_{j,T+t})^2$. Over these K draws compute the average of expression (A.38) over the vectors $\mathbf{L}_{T+s}, \mathbf{Z}_{T+s}$.

As the simulated time series grows, the distribution fills out and becomes dense so that the distance in Step 2 shrinks. The algorithm thus use a discretized state space, but the bin-size shrinks over time. For the simulations used in the paper, we set $T = 1000$, $S = 200000$ and $K = 10$. Increasing S or K increase computational time, but does not increase precision in

a meaningful way. For the graphs in the paper, we added the BEA productivity shocks to the end of the random draws of \mathbf{z} .

APPENDIX B. PROOFS OF PROPOSITIONS

B.1. Proof of Proposition 1. The proposition states that for a given $r < n$, the variance of productivity shocks conditional on being reported $\text{var}(z_i | s_i = 1)$ is larger than the unconditional variance $\text{var}(z_i)$ and increasing in n .

Proof. We start by proving that $\text{var}(z_i | s_i = 1) > \text{var}(z_i)$. Define the variable $x_i \equiv z_i^2$. Since $E(z_i) = 0$, $E(x_i) = \text{var}(z_i)$. Denote the k^{th} order statistic of $\{x_1, x_2, \dots, x_n\}$ as $x_{(k)}$ so that

$$x_{(1)} \equiv \min \{x_1, x_2, \dots, x_n\} \quad (\text{B.1})$$

$$x_{(2)} \equiv \min \{x_1, x_2, \dots, x_n\} - x_{(1)} \quad (\text{B.2})$$

\vdots

$$x_{(k)} \equiv \min \{x_1, x_2, \dots, x_n\} - \{x_{(1)}, x_{(2)}, \dots, x_{(k-1)}\} \quad (\text{B.3})$$

Note that $s_i = 1$ implies that

$$x_i \in \{x_{(n)}, x_{(n-1)}, \dots, x_{(n-r+1)}\}. \quad (\text{B.4})$$

Since $x_{(k)} \geq x_{(k-j)}$ for any $j > 0$, $x_{(k)}$ first order dominates $x_{(k-j)}$, and hence

$$E(x_{(k)}) \geq E(x_{(k-j)}) \quad (\text{B.5})$$

so that

$$\text{var}(z_i | s_i = 1) \geq \text{var}(z_j | s_j = 0). \quad (\text{B.6})$$

Combining (B.6) with the fact that

$$\text{var}(z_i) = p(s_i = 1)\text{var}(z_i | s_i = 1) + p(s_i = 0)\text{var}(z_j | s_j = 0) \quad (\text{B.7})$$

gives the desired result

$$\text{var}(z_i) = \text{var}(z_i | s_i = 1) - p(s_i = 1)[\text{var}(z_i | s_i = 1) - \text{var}(z_j | s_j = 0)] \quad (\text{B.8})$$

$$< \text{var}(z_i | s_i = 1). \quad (\text{B.9})$$

To prove the second part of the proposition, we also need to show that $\text{var}(z_i | s_i = 1)$ is increasing in n . Using the same notation as above, consider $n = l$, so that the squared values of the reported sectors is the set $\{x_{(l)}, x_{(l-1)}, \dots, x_{(l-r+1)}\}$. Now, consider adding one dimension to the state so that $n = l + 1$. If $x_{l+1} > x_{(l-r+1)}$ we have that the elements in the random vector $(x_{(l+1)}, x_{(l)}, \dots, x_{(l-r+2)})$ is larger or equal to the corresponding elements in the vector $(x_{(l)}, x_{(l-1)}, \dots, x_{(l-r+1)})$. The former vector thus first order dominates the latter vector, implying that $E(x_{(l+1)}, x_{(l)}, \dots, x_{(l-r+2)}) > E(x_{(l)}, x_{(l-1)}, \dots, x_{(l-r+1)})$. The desired result then follows from the fact that the definition of x_i implies that $E(x_i) = \text{var}(z_i)$. \square

B.2. Proof of Proposition 2. The proposition states that the conditional variance of unreported productivity shocks $\text{var}(z_j \mid \mathbf{s}, \mathbf{r}, s_j = 0)$ is increasing in the minimum value of the reported productivity shocks $\min\{|z_i| : s_i = 1\}$.

Proof. The news selection function $\mathcal{S}_{|z|}$ implies that

$$p(|Z_j| > \min\{|Z_i| : s_i = 1\} \mid s_j = 0) = 0.$$

The distribution $p(Z_j \mid \mathbf{r}, \mathbf{s}, s_j = 0)$ is therefore a truncated normal with density function

$$p(Z_j \mid \mathbf{r}, \mathbf{s}, s_j = 0) = \begin{cases} 0 & \text{if } Z_j < -\min\{|Z_i| : s_i = 1\} \\ \frac{\phi(Z_j)}{\Phi(\min\{|Z_i| : s_i = 1\}) - \Phi(-\min\{|Z_i| : s_i = 1\})} & \cdot \\ 0 & \text{if } Z_j > \min\{|Z_i| : s_i = 1\} \end{cases}$$

where ϕ and Φ are the pdf and cdf of the unconditional distribution of Z_j .

A zero mean symmetric two-sided truncated normal distribution with truncation points $-a$ and a is a mean preserving spread of any zero mean symmetric two-sided truncated normal distribution with truncation points $-b$ and b such that $a > b$. The proposition then follows from that a mean preserving spread increases the variance of a distribution. \square

B.3. Proof of Proposition 3. The proposition states that the mean of reported productivity shocks $E(z_i \mid s_i = 1)$ is lower than the unconditional mean of productivity shocks and decreasing in the number of sectors n .

Proof. Denote the k^{th} order statistic of the vector z as $z_{(k)}$. Then $z_{(k)}$ state-wise dominates $z_{(k-j)}$ for any $j > 0$ so that $E(z_{(k)}) < E(z_{(k-j)})$. To prove the first part of the proposition, note that the elements in \mathbf{r} consists of the first r order statistics of z , so that

$$E(z_i \mid s_i = 1) < E(z_j \mid s_j = 0). \quad (\text{B.10})$$

The result then follows from that

$$E(z_i) = p(s_i = 1)E(z_i \mid s_i = 1) + p(s_j = 0)E(z_j \mid s_j = 0) \quad (\text{B.11})$$

$$= 0 \quad (\text{B.12})$$

To prove the second part of the proposition, set the dimension of the state $n = l$. The reported sector shocks in the set $\{z_{(1)}, z_{(2)}, \dots, z_{(r)}\}$ then consists of the first r order statistics of an l dimensional vector. Now, consider adding one dimension to the state so that $n = l + 1$. If $z_{l+1} < z_{(r)}$ the value of one element in the random vector $(z_{(1)}, z_{(2)}, \dots, z_{(r)})$ is then replaced by a smaller value. If $z_{l+1} > z_{(r)}$ the vector is unchanged. The values of the first r order statistics drawn from l sectors thus first order stochastically dominates the first r order statistics drawn from $l + 1$ sectors. The conditional mean $E(z_i \mid s_i = 1)$ is thus decreasing in the number of sectors n . \square

B.4. Proof of Proposition 4. The proposition states that the expected value of non-reported productivity shocks $E(z_j | \mathbf{s}, \mathbf{r}, s_j = 0)$ is increasing in the maximum value of the reported productivity shocks $\max\{z_i : s_i = 1\}$.

Proof. The news selection function \mathcal{S}_- implies that

$$p(Z_j < \max\{Z_i : s_i = 1\} | s_j = 0) = 0.$$

The conditional distribution of a non-reported sector shocks is thus normal but left-truncated at $\max\{Z_i : s_i = 1\}$ with expected value given by

$$E(Z_j | s_j = 0, \max\{Z_i : s_i = 1\}) = \frac{\phi(\max\{Z_i : s_i = 1\})}{1 - \Phi(\max\{Z_i : s_i = 1\})}$$

which is increasing in $\max\{Z_i : s_i = 1\}$. □