# Patents, News, and Business Cycles*

Silvia Miranda-Agrippino*
Bank of England
CEPR and CfM(LSE)

Sinem Hacıoğlu Hoke†
Bank of England
CEPR and DAFM(KCL)

Kristina Bluwstein‡
Bank of England

First Version: August 22, 2018, Revised: July 21, 2022

**Abstract**

We exploit information in patent applications to construct an instrumental variable for the identification of technology news shocks that relaxes all the identifying assumptions traditionally used in the literature. The instrument recovers news shocks that have no effect on aggregate productivity in the short-run, but are a significant driver of its trend component. News shocks prompt a broad-based business cycle expansion in anticipation of the future increase in TFP, but only account for a modest share of fluctuations in macroeconomic aggregates at business cycle frequencies. The stock market prices-in news shocks on impact but consumer expectations, dragged down by labor market outcomes, take sensibly longer to adjust, consistent with the predictions of models of information frictions.

**Keywords:** Technology News Shocks; Business Cycle; SVAR-IV; Patents Applications; Information Frictions.

**JEL Classification:** E32, O33, O34, C36

*silvia.miranda-agrippino@bankofengland.co.uk Web: www.silviamirandaagrippino.com
†sinem.hacioglu@bankofengland.co.uk Web: www.sinemhaciogluhoke.com
‡kristina.bluwstein@bankofengland.co.uk Web: www.kristinabluwstein.com

* Previously "When Creativity Strikes: News Shocks and Business Cycle Fluctuations". We are grateful to Franck Portier and André Kurmann for insightful discussions and many useful suggestions on earlier versions of this paper. We are also grateful to Bob Barsky, Fabio Canova, Cristiano Cantore, Jordi Galí, Jean-Paul L'Huillier, Leonardo Melosi, Karel Mertens, Emanuel Mönch, Evi Pappa, Valerie Ramey, Morten Ravn, Ricardo Reis and Barbara Rossi for detailed comments. We also thank participants at the Advances in Applied Macro/Finance Conference, Bank of England, Catholic University of Milan, FRB New York, FRB Chicago, FRB San Francisco, FRB Dallas, ES European Winter Meeting, Northwestern, UCSD, LSE, UCL, UPF, Texas A&M for helpful discussions. Silvia Miranda-Agrippino gratefully acknowledges support and hospitality from Northwestern University where this research was partly conducted. The views expressed are those of the authors and do not reflect those of the Bank of England or any of its Committees.

# 1 Introduction

The idea that changes in agents' beliefs about the future may be an important driver of economic fluctuations has fascinated many scholars over the years. While the application to technology news is recent, and was revived following the seminal work of Beaudry and Portier (2004, 2006), the insight that expectations about future fundamentals could be a dominant source of economic fluctuations is a long-standing one in economics (e.g. Pigou, 1927). The news-driven business cycle hypothesis posits that economic fluctuations can arise because of changes in agents' expectations about future fundamentals, and absent any actual change in the fundamentals themselves. If the arrival of favorable news about future productivity can generate an economic boom, lower than expected realized productivity can set off a bust without any need for a change in productivity having effectively occurred. The plausibility of belief-driven business cycles is, however, still a hotly debated issue in the literature (see e.g. the extensive review in Ramey, 2016).[1]

In this paper we approach the topic from a different angle, and study the related question of how does the aggregate economy respond to shocks that raise expectations about future productivity growth. We provide an empirical answer in an information-rich VAR that includes many relevant aggregates, such as consumption, investment and labor inputs, as well as forward looking variables, such as asset prices and consumer expectations. The novelty in our approach is the identification of technology news shocks. We exploit information in patent applications to construct an instrumental variable (IV) for the shock that allows us to dispense from all the identifying assumptions traditionally used in the literature.[2]

[1]The empirical literature on technology news shocks is vast, and we review it when discussing our results in Section 4. At the poles of the debate are the advocates of the news-driven business cycle hypothesis, e.g. Beaudry and Portier (2006, 2014); Beaudry and Lucke (2010), and its opponents, e.g. Barsky and Sims (2011, 2009); Kurmann and Otrok (2013); Barsky et al. (2015); Kurmann and Sims (2021). Other contributions have highlighted the role played by different modeling assumptions and specifications, and by alternative data transformations (e.g. Christiano, Eichenbaum and Vigfusson, 2003; Francis and Ramey, 2009; Mertens and Ravn, 2011; Forni, Gambetti and Sala, 2014).

[2]Traditional identifications are motivated by economic theory, and typically combine zero restrictions on the impact response of TFP with assumptions about the drivers of productivity in the long-run. In Beaudry and Portier (2006) news shocks are orthogonal to current productivity, but are its sole driver in the long run (Galí, 1999; Francis and Ramey, 2005). Other works have relaxed this latter assumption and assumed news shocks maximize the forecast error variance of productivity at some long finite horizon (e.g. Francis, Owyang, Roush and DiCecio, 2014), or over a number of different horizons (e.g. Barsky and Sims, 2011).

The intuition behind our identification is simple: by their nature, patent applications embed a signal of potential future technological change. However, they may in turn be prompted by current economic booms and/or past news. To account for this endogeneity, when constructing the IV we control for expectations about the macro outlook that were formed prior to the application filing dates, and for other contemporaneous policy changes that could influence the decision of filing a patent either directly, or through their effects on other macro aggregates. Specifically, we recover the IV as the component of patent applications that is orthogonal to pre-existing macro beliefs as captured by the Survey of Professional Forecasters, to contemporaneous monetary and fiscal policy changes as summarized by narrative accounts, as well as to own lags.[3] We consider two different data sources for patent applications. Our baseline is the NBER 'USPTO Historical Patent Data Files' of Marco, Carley, Jackson and Myers (2015), that provides a comprehensive record of all patent applications, granted and not granted, filed at the U.S. Patents and Trademark Office (USPTO) since 1981 and aggregated at monthly frequency. The second source is Kogan, Papanikolaou, Seru and Stoffman (2017), that collects information on individual patents granted by the USPTO to large corporations between 1926 and 2010, including their application date, forward citations, and economic value they generate in the stock market. We use this latter source to study the robustness of our findings to weighting the patent applications in the construction of the IV.[4]

The exclusive rights granted to patent holders ensure that individuals and businesses have a set number of years to capitalize on their inventions, and act as a powerful incentive to engage in the patenting process. The length of time from the application to the grant date, and the eventual diffusion of the innovation within the economy can be in the

[3]To be clear, our strategy is in principle equivalent to identifying technology news shocks in a standard Cholesky triangularization as an innovation to patent applications in a VAR where the variables enter in the following order: (1) past (relative to the filing date of patent applications) expectations about current and future macro outcomes; other contemporaneous policy shocks; (2) patent applications; (3) TFP and other variables of interest. In practice, splitting the problem in two and constructing the instrument outside of the VAR grants us a number of advantages, including being able to accurately match the timing of the patent filings with that of the SPF forecasts, delivering an IV which can readily be used by other researchers, accounting for the presence of measurement error, and easily deal with different sample lengths.

[4]Kogan et al. (2017) use Google Patents to retrieve information on issued patents, and restrict their sample to only include patents granted to corporations whose returns are in the CRPS database. While covering a smaller cross-section of patent applications, this latter source is also useful to extend our IV to earlier years.

order of several years, depending on the type of patent and the characteristics of the industry sector.[5] Therefore, patent applications at any given time contain information about technological changes that may occur at some point in the future (see also e.g. Griliches, 1990; Lach, 1995; Hall and Trajtenberg, 2004). In other words, and importantly for our purpose, they represent an uncontroversial way to measure news about possible future technological progress, to a large extent regardless of whether such progress does indeed follow. Because patent applications are public, we can use the filing date as the first measurable time in which the news occurs, although it is clearly the case that the underlying idea, in the form of a private signal, predates it. Controlling for policy changes and for expectations about the macro outlook that precede the application filing is a necessary step to increase the likelihood that no other structural disturbances affect the US economy through the IV, except contemporaneous technology news. This is our identifying assumption.

Because of the minimal set of restrictions required for identification, our framework allows us to investigate whether news shocks generate the patterns that were assumed in earlier identification schemes. While it is not known ex ante whether technological innovation will effectively follow, the news we capture does eventually materialize on average, and aggregate TFP eventually rises. This allows us to label the recovered structural disturbance as news, as opposed to noise (see e.g. discussion in Chahrour and Jurado, 2018), overcoming the issues highlighted in Blanchard, L'Huillier and Lorenzoni (2013). Importantly, because innovations can in principle be released to the public under a 'patent-pending' status, our identification scheme does not warrant imposing orthogonality with respect to the current level of technology, which is a typical assumption in the news literature.[6] While such orthogonality condition is not imposed a priori, the IV recovers a shock that has essentially no effect on TFP either on impact, or in the years immediately afterwards. After this inertial initial reaction, aggregate TFP rises robustly, following the S-shaped pattern that is typical of the slow diffusion of technology (see e.g.

[5]From application filing to grant issuance the process takes two years on average. While not all applications result in granted patents, the share of successful applications can be substantial (up to 80%), with some heterogeneity across sectors (see Marco et al., 2015).

[6]In this respect, our identification is akin to Barsky et al. (2015); Kurmann and Sims (2021), who also relax the assumption of a zero impact response of TFP. Our approach is also robust to mismeasurements in commonly used empirical estimates of aggregate technology (see e.g. discussions in Fernald, 2014; Kurmann and Sims, 2021).

Rogers, 1962; Gort and Klepper, 1982). Similarly, albeit we impose no constraints on variance shares ex ante, the recovered shock explains only a modest fraction of the variation of TFP at frequencies higher or equal than those associated with standard business cycle durations, and is instead an important driver of its long-run/permanent component.

The arrival of positive news about future technology triggers a sustained and broad-based economic expansion. In the VAR output, consumption, investment and hours worked all rise to peak at the two-year horizon, and well before any material improvement in TFP is recorded. In this sense, the pattern of responses lends credit to a 'news-view' in the spirit of Beaudry and Portier (2006), whereby aggregate fluctuations arise in anticipation of changes in TFP. Indeed, the large asynchronicity in the timing of the estimated dynamic responses suggests that the aggregate effects of technology news that we unveil may be predominantly (if not entirely) driven by beliefs, rather than by future realized fundamentals. The expansion is not immediate. While consumption rises already upon realization of the shock, impact responses of output and investment are not significant at conventional levels. The impact response of the labor market can instead be best summarized as a short-lived leftward shift in labor demand, whereby both wages and aggregate hours fall briefly as the shock hits before increasing robustly (see also Basu, Fernald and Kimball, 2006; Barsky and Sims, 2011; Kurmann and Sims, 2021). The shock that we recover is, however, not a main driver of business cycles. At the relevant frequencies, it accounts for less than 10% of the variation in consumption, and for about 5% of the variation in hours, investment, and output. These findings echo results in Angeletos, Collard and Dellas (2020) that shows that shocks that account for the bulk of business cycle fluctuations are not those that are responsible for the long-run. This disconnect between what drives business cycles and long-run fluctuations is a result that we also confirm in our setting.

Finally, our results highlight important asymmetries in the way in which different agents within the economy respond to technology news shocks. On the one hand, the stock market is quick in pricing-in the news. On the other, consumers require sbustantially longer to upgrade their forecasts about the outlook. In apparent contrast with there being underlying positive news, but consistent with the immediate albeit short-lived deterioration in labor market conditions, consumers incorporate the positive signal only with a delay. These results point to a strong interaction between consumers' expectations and labor market dynamics, and the relevance of the latter in shaping the response of the former. More generally, they constitute additional evidence in support of the noisy information environment modelled in e.g. Woodford (2003); Sims (2003); Mackowiak and Wiederholt (2009), and documented in Coibion and Gorodnichenko (2012, 2015), for which news shocks represent the ideal case study.

Our work is closely related to a stream of studies that have relied on empirical measures of technological changes to identify technology news shocks. The first such study is Shea (1999). Here annual patent applications and R&D expenditures are used to estimate the effects of technology shocks on industry aggregates. Identification is achieved by ordering either measure last in a battery of small-scale VARs that also include labor inputs and productivity. Christiansen (2008) extends this study by using over a century of annual patent application data. The benchmark specification is a bivariate VAR with labor productivity and patents ordered first. Alexopoulos (2011) uses the number of book titles published in the field of technology to capture the time in which the novelty is commercialized. Responses of aggregate variables are estimated in a set of bivariate VARs with the publication index ordered last.[7] Our paper differs from these contributions in several ways. First, these studies address the fundamental endogeneity of empirical measures of technological changes only to the extent that it is captured in the reminder of variables included in the bi/tri-variate VARs. Other than relying on a richer VAR specification, in the construction of our instrument we explicitly control for the fact that the cyclical nature of patent applications may be influenced by current economic conditions, or indeed by past news. Second, and related, these studies have all implicitly assumed the empirical measure of technology being a near perfect measure of news shocks. In fact, their identifying assumptions amount to effectively retrieving the transmission coefficients by running a distributed lag regression (with some controls) of the variables on the patent data. In contrast, our identifying assumptions explicitly account for the possible presence of measurement error in the constructed instrument.

[7] More recently, Baron and Schmidt (2014) have used technology standards and a recursive identification to infer on the aggregate implications of anticipated technology shocks. In an international context, Arezki, Ramey and Sheng (2017) use giant oil discoveries as a directly observable measure of technology news shocks and estimate their effects in a dynamic panel distributed lag model.

Finally, these studies have all relied on annual data potentially overlooking important higher frequency variation which instead we exploit for the identification.

The structure of the paper is as follows. Section 2 introduces the external instrument and describes the patent data used for its construction. In Section 3 we lay out the identifying assumptions in our SVAR-IV and discuss the identification of technology news shocks using an illustrative 5-variable VAR. Section 4 contains our main results; here we we extend the analysis to an information-rich 16-variable VAR to explore the transmission mechanisms of technology news shocks more in detail. Section 5 concludes. Additional material is reported in the Appendix.

# 2 A Patent-Based IV for Technology News Shocks

## 2.1 Information in Patent Data

Our starting point for the analysis is the monthly flow of all new patent applications filed at the U.S. Patent and Trademark Office. The data are from the 'USPTO Historical Patent Data Files' compiled by Marco et al. (2015) as a follow up and extension of Hall et al. (2001). The dataset records the monthly stocks and flows of all publicly available applications and granted patents filed from January 1981 to December 2014. The stocks include pending applications and patents-in-force; flows include new applications, patent grants and abandonments.[8]

The patents in the dataset are classified as utility patents. Also known as patents for invention, these cover the creation of new or improved, and useful products, processes or machinery. We construct quarterly patent counts by summing up the monthly flows of all new patent applications within each quarter over the available sample. The left panel of Figure 1 plots the time series of quarterly patent applications aggregated at the industry level. In the figure, shaded areas denote NBER recession episodes, and we normalize 1981-I to be equal to 0 to highlight the different trends across different sectors. Patent applications have increased substantially over the past 40 years and, as visible from the chart, patents classified under 'computers and communications' have enjoyed a

[8] The dataset is available at `http://www.ustpo.gov/economics`.

Figure 1: Patent Applications & Aggregate Innovation

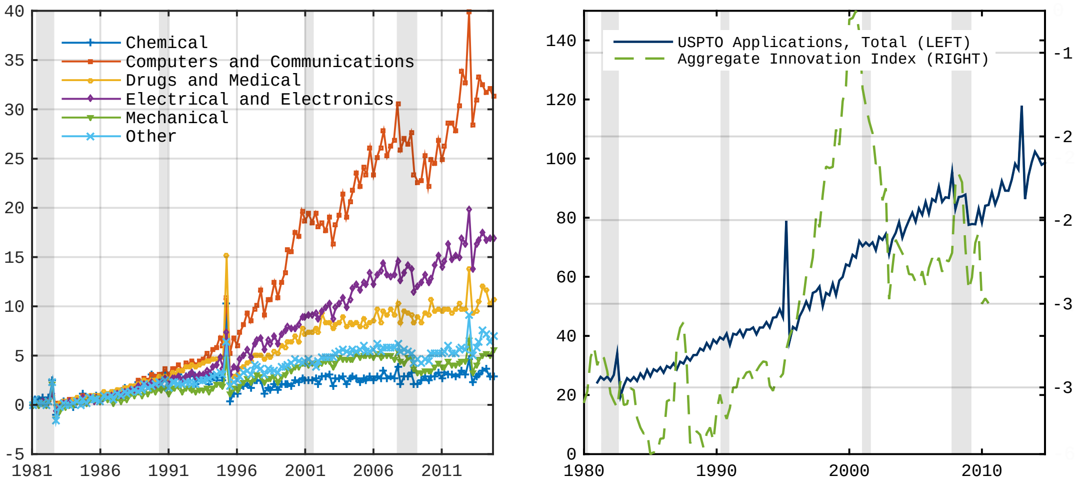


*Note*: [LEFT] Patent applications across all NBER categories. Quarterly figures obtained as sum of monthly readings, 1981-I=0. Thousands. Source: USTPO. [RIGHT] Total number of USTPO applications (sum across NBER categories, solid line), thousands, left axis. Kogan et al. (2017) aggregate innovation index, GDP weighted, log scale, USD, right axis. Shaded areas denote NBER recession episodes.

faster trend. Applications across all categories tend to slide after recessionary episodes, providing some preliminary evidence of their cyclical nature.

There have been three important regulatory changes in patenting in 1982, 1995, and 2013. All these regulations affected the number of applications when they came into effect, as shown by the spikes in Figure 1. However, since they were not legislated in response to considerations related to either current or anticipated economic conditions, they provide us with important exogenous variation which we exploit for the identification. Said differently, to the extent that each patent embeds a signal about potential future technological progress, the increase in applications induced by each piece of legislation is an exogenous (relative to macroeconomic conditions) increase in technology news, which is the focus of our identification.[9]

In 1982, the old Court for Customs and Patent Appeals was abolished, and a new Court of Appeals for the Federal Circuit was established; the new court provided more protection to patents' owners against infringement. In 1995, the U.S. implemented wide-ranging changes to patent law under the Agreement on Trade-Related Aspects of Intellectual Property Rights (TRIPS), as part of the Uruguay Round Agreements Act. The

[9]We explore the sensitivity of our results to these spikes in Section 3.

TRIPS agreement's main purpose was to harmonize patenting rules among all members of the World Intellectual Property Organization with the aim to contribute to the promotion of technological innovation and to the transfer and dissemination of technology.[10] One of the main changes introduced by the TRIPS Agreement was that of promoting transparency in patenting, and disincentivize strategic behaviour through stricter regulation.[11] This had two main effects. First, it shifted forward the timing of some applications, which resulted in the one-off increase highlighted in the chart. Second, it made applications more informative about future innovations (Encaoua, Guellec and Martínez, 2006). Finally, in March 2013, the U.S. implemented the rules dictated by the America Invents Act which further revised ownership rights.[12]

In the right panel of Figure 1 we compare the total number of USPTO applications (sum across industries in LHS chart, solid line) with the aggregate index of innovation proposed in Kogan et al. (2017). The index is a forward-looking measure of the private, economic value of innovations in the US, and constructed as the GDP-weighted sum of the economic value of all patents granted within each quarter.[13] We note that in the relevant sample, patent applications naturally lead the aggregate innovation index. Moreover, the large spikes in the number of applications tend to correspond to substantial future increases in aggregate innovation, particularly so after the TRIPS Agreement. We take this as an indication that the exogenous legislation-induced spikes in applications are informative about their 'innovation content', and thus contain important information for the purpose of identifying technology news shocks.[14]

[10] Article 7 ("Objectives") of the TRIPS Agreement states that the protection and enforcement of intellectual property rights should contribute to the promotion of technological innovation and to the transfer and dissemination of technology, to the mutual advantage of producers and users of technological knowledge and in a manner conducive to social and economic welfare, and to a balance of rights and obligations. Source: https://tinyurl.com/WTO-TRIPS-Technology-transfer.

[11] The change in legislation led to a significant reduction in the so-called submarine patents. These are patents whose issuance or publication is intentionally delayed for strategic purposes, and would often emerge decades later to prevent competitors from patenting on related topics. The TRIPS also modified patent terms which were set to 20 years from filing, and away from the previous practice of 17 years after issuance. For most industries this meant a reduction in the protection period. Source: https://www.wto.org/english/tratop_e/trips_e/inovationpolicytrips_e.htm.

[12] The new rules were designed to address the right to file a patent application, and switched the priority rule to the 'first-inventor-to-file', rather than the pre-existing 'first-to-invent'. Source: https://www.uspto.gov/sites/default/files/aia_implementation/20110916-pub-l112-29.pdf.

[13] The original index in Kogan et al. (2017) is annual. Using their data, we have reconstructed a quarterly version following the same procedure as in the original one.

[14] In a recent contribution Cascaldi-Garcia and Vukotić (2022) use the index of Kogan et al. (2017) to

When constructing our benchmark IV for technology news shocks, we use all patent applications submitted to the USPTO, including those that are ex-post not granted, and weigh them all equally (solid line in Figure 1, right panel). There are multiple reasons for this choice. First, we choose to work with patent applications rather than grants. Previous studies such as e.g. Christiansen (2008) have noted how most of the news content in patent applications may be exhausted by the time they are granted. One reason for it is that innovations can be disseminated under patent-pending status. Other anecdotal evidence reported in Kogan et al. (2017) suggests that "the market often had advance knowledge of which patent applications were filed, since firms often choose to publicize new products and the associated patent applications themselves." Thus, for the purpose of isolating technology news, applications are more likely to capture the effective time at which the news materializes.[15] Second, we choose to include in our set also patents that are ex-post not granted. This is primarily due to our data source supplying information on the total number of applications filed at the USPTO each month, with no information on which ones are ultimately successful. But it also makes sense from an identification perspective: at the time of the application, all patents arguably bear news. Third, it is possible, and indeed likely, that markets and applicants may attach to each patent an individual ex-ante probability of it being ex-post granted and/or more or less groundbreaking. This would be the optimal way to weigh the applications for the purpose of capturing news more accurately, but it is of course unfeasible. As a result, and in an attempt to account for all these aspects, as we detail in the next section we construct our baseline IV using all applications with equal weights.

There is a question of whether the IV can be ameliorated by weighting the patents differently. A common practice in the literature that uses patent data is to weigh them according to forward citation counts. That is, according to the number of citations that each patent receives in the future, which is typically regarded as a way to measure its scientific relevance. An alternative, proposed in Kogan et al. (2017), is to use weights that reflect the economic value that a patent generates in the stock market when it is granted. At the firm-patent level, the value of each patent is measured based on the three-day

identify technology news shocks as a follow up to our analysis in this paper.

[15] Christiansen (2008) also notes that grants tend to be significantly more cyclical than applications, and dependent on the intensity of labor and administrative cycles at the USPTO.

return that the patent owner's stock enjoys when the patent is granted. We discuss these options in detail in Appendix E, where we provide extensive robustness of our results to alternative weighting schemes. While results are generally robust to these alternatives, due to economic agents – including financial markets – not knowing at the application stage which patents will ex-post be granted, nor the expected citations or realized returns around the grant date, we are skeptical around the use of such weighting schemes for the purpose of constructing an instrument for technology news shocks, since they rest on information that was not available at the time in which the news materialized.

## 2.2 Instrument Construction

We recover an instrumental variable for the identification of technology news shocks as the component of patent applications that is orthogonal to pre-existing beliefs about the state of the economy, other contemporaneous policy shocks, and is unpredictable given its own history. Intuitively, we seek to remove endogenous variation in application filings that results from anticipation of economic conditions due to past news and other contemporaneous disturbances. This to increase the likelihood that the IV correlates with contemporaneous news shocks only, which is the required condition for correct identification.

Specifically, we propose as IV the residual of the following regression, estimated at quarterly frequency

$$pa_t = c + \gamma(L)pa_t + \sum_{h=1,4} \beta_h \mathbb{E}_t[x_{t+h}] + \sum_{j=0}^{2} \delta_j \eta_{t-j} + z_t. \tag{1}$$

In Eq. (1), $pa_t$ is the quarterly growth rate of all patent applications, i.e. $pa_t = 100 \times (\ln PA_t - \ln PA_{t-1})$, where $PA_t$ is the sum of all patent applications filed at the USPTO. $\gamma(L) = \sum_{j=1}^{4} \gamma_j L^j$, where $L$ is the lag operator, and $\mathbb{E}_t[x_{t+h}]$ is an $m \times 1$ vector of forecasts for the economic variables in $x_t$ that we take from the Survey of Professional Forecasters (SPF). The forecast horizon $h$ is equal to one and four quarters. The time index in $\mathbb{E}_t$ refers to the publication date of the survey. Because of the release schedule of the SPF, the information set conditional on which forecasts are made is in fact relative to the previous quarter; hence, the collection of forecasts in $\mathbb{E}_t[x_{t+h}]$ captures pre-existing beliefs about the macroeconomic outlook.[16] The vector $x_t$ includes the unemployment rate ($u_t$), inflation ($\pi_t$), and the growth rates of real non-residential fixed investments ($I_t$), and of real corporate profits net of taxes ($\Pi_t$).[17]

An important concern relates to the potential correlation of patent applications with other contemporaneous shocks, besides current technology news. If this were the case, the exclusion restrictions in our IV-based identification strategy would be violated. While there is no formal way to test for the exogeneity of the instrument, we address this concern by including in Eq. (1) further controls that capture monetary and fiscal policy changes up to the current quarter. The vector $\eta_t$ includes unexpected and anticipated exogenous tax changes as classified by Romer and Romer (2010) and Mertens and Ravn (2012), and the narrative series for monetary policy shocks of Romer and Romer (2004).[18] The rationale here is that monetary and tax policy, by affecting macro aggregates (especially investment) within the quarter, may have a direct effect on patent applications, and act as a confounding factor in the identification.

The regression results are presented in Table 1.[19] The table reports individual regression coefficients and robust standard errors in parentheses for five models. Eq. (1) corresponds to column (5) in the table. In columns (1) to (4) we consider subsets of controls for comparison. Due to the availability of the narrative tax series, the specifications in columns (4) and (5) are estimated over the sample 1981-I:2006-IV. Columns (1) to (3) use the full length of patent data (1981-I:2014-IV). At the bottom of the table, we report Wald test statistics for the joint significance of the controls (excluding own lags) in each regression. Patent applications exhibit a strong autocorrelation pattern.[20]

[16] SPF forecasts are published in the middle of the second month of each quarter. The information set of the respondents at the time of compiling the survey includes the advance report on the national income and product accounts of the Bureau of Economic Analysis, which is published at the end of the first month in each quarter, and contains advance releases for macroeconomic aggregates referring to the previous quarter. For further information see https://www.philadelphiafed.org/research-and-data/real-time-center/survey-of-professional-forecasters.

[17] SPF respondents forecast nominal corporate profits net of taxes. We construct a series for real corporate profits forecasts by deflating with the forecasts for the GDP deflator (our measure of inflation, see Section 4) at the relevant forecast horizons.

[18] We use an extension of the Romer and Romer (2004) series up to 2007. Controlling for the changes in tax policy follows from the intuition in Uhlig (2004) who noted that changes in capital income taxes would lead to permanent effects on labor productivity and hence be a confounding factor in the analysis of technology shocks. This intuition was further developed in Mertens and Ravn (2011).

[19] The instrument is plotted in Figure A.1 in the Appendix.

[20] The negative sign of the autoregressive coefficients, also noted in Adams et al. (1997), suggests the

## Table 1: Instrument Construction

| | (1) | (2) | (3) | (4) | (5) |
|---|---|---|---|---|---|
| *Own Lags* | | | | | |
| $pa_{t-1}$ | −0.849*** | −0.928*** | −0.901*** | −0.948*** | −0.952*** |
| | (0.10) | (0.11) | (0.10) | (0.09) | (0.08) |
| $pa_{t-2}$ | −0.480*** | −0.605*** | −0.574*** | −0.505*** | −0.548*** |
| | (0.10) | (0.11) | (0.11) | (0.12) | (0.11) |
| $pa_{t-3}$ | −0.273*** | −0.383*** | −0.365*** | −0.236** | −0.272** |
| | (0.09) | (0.08) | (0.08) | (0.11) | (0.11) |
| $pa_{t-4}$ | 0.002 | −0.061 | −0.056 | −0.012 | −0.033 |
| | (0.09) | (0.08) | (0.08) | (0.10) | (0.09) |
| *Pre-Existing Beliefs* | | | | | |
| $\mathbb{E}_t[u_{t+1}]$ | | −0.323 | | | 0.629 |
| | | (0.37) | | | (4.82) |
| $\mathbb{E}_t[\pi_{t+1}]$ | | 1.635** | | | 3.424* |
| | | (0.69) | | | (1.77) |
| $\mathbb{E}_t[I_{t+1}]$ | | 0.488** | | | 0.065 |
| | | (0.23) | | | (0.28) |
| $\mathbb{E}_t[\Pi_{t+1}]$ | | −0.137 | | | −0.221 |
| | | (0.23) | | | (0.34) |
| $\mathbb{E}_t[u_{t+4}]$ | | | −0.851* | | −1.513 |
| | | | (0.46) | | (5.57) |
| $\mathbb{E}_t[\pi_{t+4}]$ | | | 0.887 | | −2.979* |
| | | | (0.77) | | (1.57) |
| $\mathbb{E}_t[I_{t+4}]$ | | | 0.377 | | −0.101 |
| | | | (0.26) | | (0.40) |
| $\mathbb{E}_t[\Pi_{t+4}]$ | | | −0.673*** | | −0.224 |
| | | | (0.19) | | (0.27) |
| *Policy Shocks* | | | | | |
| $mpol_t$ | | | | −4.810** | −4.377** |
| | | | | (2.10) | (1.84) |
| $mpol_{t-1}$ | | | | 6.318 | 6.319 |
| | | | | (4.15) | (4.47) |
| $mpol_{t-2}$ | | | | 4.644** | 3.560* |
| | | | | (1.84) | (2.08) |
| $utax_t$ | | | | −0.902 | −1.979* |
| | | | | (0.89) | (1.14) |
| $utax_{t-1}$ | | | | 0.595 | −0.875 |
| | | | | (1.65) | (1.60) |
| $utax_{t-2}$ | | | | −0.884 | −2.976** |
| | | | | (0.67) | (1.47) |
| $atax_t$ | | | | 4.646 | 2.443 |
| | | | | (3.08) | (2.86) |
| $atax_{t-1}$ | | | | −1.645 | −3.332 |
| | | | | (1.45) | (2.02) |
| $atax_{t-2}$ | | | | −4.599 | −5.261 |
| | | | | (3.90) | (3.99) |
| intercept | 4.343*** | 0.977 | 7.610 | 5.027*** | 10.949* |
| | (0.80) | (2.86) | (5.02) | (0.85) | (6.33) |
| F-stat | 33.87 | 18.04 | 19.48 | 21.26 | 13.59 |
| | [0.000] | [0.000] | [0.000] | [0.000] | [0.000] |
| Adj-$R^2$ | 0.448 | 0.486 | 0.469 | 0.510 | 0.493 |
| N | 131 | 131 | 131 | 99 | 99 |
| *Wald Tests for Joint Significance of Controls* | | | | | |
| Quarter Ahead SPF | | 4.788 | | | |
| | | [0.001] | | | |
| Year Ahead SPF | | | 3.72 | | |
| | | | [0.007] | | |
| Policy Shocks | | | | 2.361 | |
| | | | | [0.020] | |
| SPF & Policy Shocks | | | | | 2.505 |
| | | | | | [0.003] |

*Notes:* Regression results based on Eq. (1). Dependent variable: $pa_t = 100 \times (\ln PA_t - \ln PA_{t-1})$. Robust standard errors in parentheses. SPF Forecasts are for the unemployment rate ($u_t$), inflation (GDP deflator, $\pi_t$), real non-residential investments ($I_t$), and real corporate profits net of taxes ($\Pi_t$). Policy controls include narrative monetary policy ($mpol_t$), narrative unanticipated ($utax_t$) and anticipated ($atax_t$) tax changes. The bottom panel reports Wald test statistics for the joint significance of the controls with associated p-values below in square brackets. *, **, *** denote statistical significance at 10, 5, and 1% respectively.

Moreover, pre-existing beliefs about the future as captured by the SPF forecasts contain information for patent applications beyond that included in own lags. This is consistent with patents being endogenous to the economic cycle, and, potentially, also related to past news embedded in the survey forecasts. Policy changes are also informative.

The procedure in Eq. (1) removes the autocorrelation and seasonal patterns in patent applications, and the dependence on pre-existing beliefs as captured by the SPF. Moreover, it ensures that the IV is also orthogonal to other contemporaneous policy shocks. The resulting IV is not forecastable also conditional on a wider set of predictors. Macro-financial factors extracted from large cross-sections and broader sets of forecasts that Granger-cause patent applications are uninformative for the IV.[21]

We argue that it is unlikely that structural disturbances other than current technology news may affect the US economy through $z_t$. This is our sole identifying assumption.

# 3 Identification of Technology News Shocks

In the news literature, it is common to think of the process for technology as a random walk with drift subject to two stochastic disturbances. A typical representation assumes technology to be the sum of a stationary and a permanent component, with news shocks affecting the latter (see e.g. Blanchard et al., 2013; Kurmann and Sims, 2021). Formally

$$\ln A_t = \ln S_t + \ln \Gamma_t \ , \tag{2}$$

where $S_t$ is the stationary component, assumed to follow an AR(1) process

$$\ln S_t = \phi_s \ln S_{t-1} + e_{\mathrm{A1},t} \ , \tag{3}$$

presence of seasonal patterns in patent applications data. It is likely that these may be the result of USPTO institutional features and characteristics of the patenting process itself. The inclusion of own lags in Eq. (1) removes dependency of the IV on its own past and ensures that the specific source of seasonality does not affect the identification.

[21]See Tables A.1 and A.2 for Granger-causality results on patent applications, and Tables A.3 and A.4 for the same on the IV.

and $\Gamma_t$ is the permanent component, characterized instead by the presence of a unit-root

$$\Delta\ln\Gamma_t = \Delta\ln A + \phi_\Gamma \Delta\ln\Gamma_{t-1} + e_{\mathrm{A2},t-k} \ . \tag{4}$$

In Eqs. (3) - (4) above $\Delta\ln A$ is the steady state growth rate of technology, the autoregressive coefficients $\phi_s$ and $\phi_\Gamma$ are in the interval $(0, 1)$, and $e_{\mathrm{A1},t}$ and $e_{\mathrm{A2},t-k}$ are zero-mean normally distributed i.i.d. processes with variance equal to $\sigma^2_{\mathrm{A1}}$ and $\sigma^2_{\mathrm{A2}}$ respectively. $A_t$ is typically understood as a shifter to the aggregate production function of the economy, and intended to capture a concept of technology related to the efficiency with which the factors of production are utilized, or the introduction of new processes altogether.

$e_{\mathrm{A2},t}$ is the news shock. The standard identifying assumption in the news literature is that agents learn about $e_{\mathrm{A2},t-k}$ before it hits the technology process, i.e. $k > 0$ (see e.g. Beaudry and Portier, 2006; Barsky and Sims, 2011, among many others). However, a number of more recent papers have argued that news shocks are also in principle compatible with $k = 0$, which would affect technology also on impact (see e.g. Barsky et al., 2015; Kurmann and Sims, 2021). This may happen because news about future productivity arrives along with an innovation in current technology, because innovations to current technology may signal significant improvements in the following years, or because technology slowly diffuses across sectors.

Allowing for $k = 0$ naturally makes the task of telling apart a news shock with effects also on current technology from an innovation in current technology ($e_{\mathrm{A1},t}$) a daunting one. In this respect, we rely on the information content of the instrument constructed in Section 2. As noted, while patent applications are most informative for news about possible future technological changes ($k > 0$), the fact that innovations can be distributed under a patent-pending status does not rule out the $k = 0$ case a priori. Hence, the use of the patent-based IV does not warrant imposing orthogonality with respect to the current level of technology. However, as we shall see in the reminder of this section, while no assumption on the impact response is made, the instrument recovers a shock that leads to an effectively muted response of TFP upon realization, while eliciting a strong and sustained response at further ahead horizons. This gives us confidence that the recovered shock has a large element of news embedded in it.

## 3.1 Identifying assumptions in our SVAR-IV

We use our patent-based IV to back out the dynamic causal effects of technology news shocks on a collection of macroeconomic and financial variables in a structural Vector Autoregression (SVAR-IV, Mertens and Ravn, 2013; Stock and Watson, 2012, 2018).

Let $y_t$ denote the $n$-dimensional vector of economic variables of interest, whose dynamics follow a VAR(p)

$$\Phi(L)y_t = u_t, \qquad u_t \sim \mathcal{WN}(0, \Sigma), \tag{5}$$

where $\Phi(L) \equiv \mathbb{I}_n - \sum_{j=1}^{p} \Phi_j L^j$, $L$ is the lag operator, $\Phi_j$ $j = 1, \ldots, p$ are conformable matrices of autoregressive coefficients, and $u_t$ is a white noise vector of zero-mean innovations, or one-step-ahead forecast errors, i.e. $u_t \equiv y_t - \text{Proj}(y_t | y_{t-1}, y_{t-2}, \ldots)$.

For the purpose of estimating the impulse response functions (IRFs) and error variance decompositions (EVDs) we require that the information in our VAR be sufficient to recover all the structural shocks. Specifically, that there exists an $n$-dimensional matrix $B_0$ such that

$$u_t = B_0 e_t, \tag{6}$$

where $e_t$ is a vector of $n$ structural disturbances, and $B_0$ collects the contemporaneous effects of $e_t$ on $y_t$. Given a suitable identification scheme, Eq. (6) guarantees that the structural disturbances can be recovered from the observables in the VAR. Full invertibility is not strictly required for IV-based identification of IRFs to a single shock of interest, as discussed in Miranda-Agrippino and Ricco (2018) and Plagborg-Møller and Wolf (2021). However, Forni et al. (2019) show that if Eq. (6) does not hold, then estimates of the forecast error variance contributions are distorted.

When agents anticipate future changes, as is the case with technology news shocks, non-fundamentalness is likely to arise (see e.g. Leeper et al., 2013). Intuitively, if the shock only has effect on future variables, current realizations are only informative about past shocks, and the mapping in Eq. (6) breaks down. In this context, a natural route towards the problem solution is to add information to the VAR, through variables that help revealing the state variables. This is the role of the stock price index in Beaudry and Portier (2006), or measures of consumers or business confidence as in Barsky and Sims (2012). In a similar vein, factors estimated from large cross-sections can be added to the VAR specification as in e.g. Giannone and Reichlin (2006); Forni and Gambetti (2011).[22]

Conditional on Eq. (6) holding, the conditions for identification in SVAR-IV are

$$\mathbb{E}[e_{\text{A2},t} z_t] = \rho, \qquad \rho \neq 0 \qquad (\textit{Relevance}) \tag{7}$$

$$\mathbb{E}[e_{i,t} z_t] = 0, \qquad \forall i \neq \text{A2} \qquad (\textit{Contemporaneous Exogeneity}), \tag{8}$$

where $z_t$ denotes the external instrument used for the identification of $e_{\text{A2},t}$. Under these conditions, the impact responses to $e_{\text{A2},t}$ of all variables in $y_t$ are consistently estimated (up to scale and sign) from the projection of the VAR innovations $\hat{u}_t$ on the instrument $z_t$ (Mertens and Ravn, 2013; Stock and Watson, 2012, 2018).

It is important to note that, by construction, the IV will correlate with technology news shocks insofar as these are captured by the patenting process, and may therefore leave other sources of variation in long-term productivity growth unaccounted for. Said differently, while all patent applications are an ex-ante measure of technology news, not all technology news are captured by patents. What is crucial for the identification is that no other structural disturbances affect the correlation between $\hat{u}_t$ and $z_t$ other than technology news.

## 3.2 Inspecting the Mechanism in an Illustrative VAR

In this section, we put our instrument to test in an illustrative 5-variable VAR and discuss the sensitivity of our results with respect to a number of perturbations. The variables included in the VAR are the quarterly estimates of TFP corrected for input utilization of Fernald (2014), output, consumption, total hours worked, and the Dow Jones Industrial Average as the stock market index. The variables are chosen as to encompass the sets used in the VARs of Beaudry and Portier (2006) and Barsky and Sims (2011). The variables enter the VAR in log levels, and are deflated and expressed in per-capita terms

[22] While non-fundamentalness is a theoretically binding constraint, empirically the VAR-based IRFs may still be accurate if the 'wedge' between the estimated and the true shocks is small (Sims, 2012). See also Beaudry and Portier (2014); Beaudry et al. (2019).

where appropriate. We use the GDP deflator to measure inflation and report a detailed description of the data and their construction in Table B.1 in the Appendix. The VAR is estimated with Bayesian techniques with 4 lags over the 60-year sample 1960-I:2019-IV. We refer to the sample used for the VAR estimation as the estimation sample, and the one used for the projection of the VAR residuals on the instrument as the identification sample respectively. Our identification sample equals the full length of $z_t$ (1982:I to 2006-IV).

For the estimation of the VAR, we use a standard Normal-Inverse Wishart prior centered around a random walk for each variable (Doan et al., 1983; Litterman, 1986; Kadiyala and Karlsson, 1997). The optimal priors' tightness is estimated as in Giannone et al. (2015). We present our empirical results in the form of impulse response functions at the mode of the posterior distribution of the parameters, and normalized such that the peak response of TFP equals 1%. The IRFs are identified with the two-step procedure of Mertens and Ravn (2013). Shaded areas correspond to 68% and 90% posterior credible sets.[23]

The IRFs are reported in Figure 2. A few elements stand out. First, while we have not imposed any restrictions on the effect of the shock on current TFP, the chart reveals that the shock recovered by the IV has essentially no effect on TFP neither on impact, nor in the following four to six years. TFP eventually rises robustly and remains elevated throughout, following a shape that resembles the S-shaped pattern that is typical of the slow diffusion of new technologies. A similarly shaped response is reported in Barsky et al. (2015) and Kurmann and Sims (2021) who identify technology news shocks based on the forecast error variance of TFP, and do not restrict the impact TFP response to zero. Second, output, consumption and hours worked all rise. Aggregate consumption increases robustly on impact, while the initial response of output and hours is more modest, albeit still positive. For all three variables, the rise is sudden, and the peak of the dynamic adjustment is reached long before any material increase in TFP materializes, within one or two years after the shock hits. Third, the stock market prices-in the news on impact, and remains elevated throughout.

[23]Because our instrument is a residual generated regressor, OLS-based inference is asymptotically correct (Pagan, 1984).

Figure 2: Technology News Shocks in the 5-variable VAR

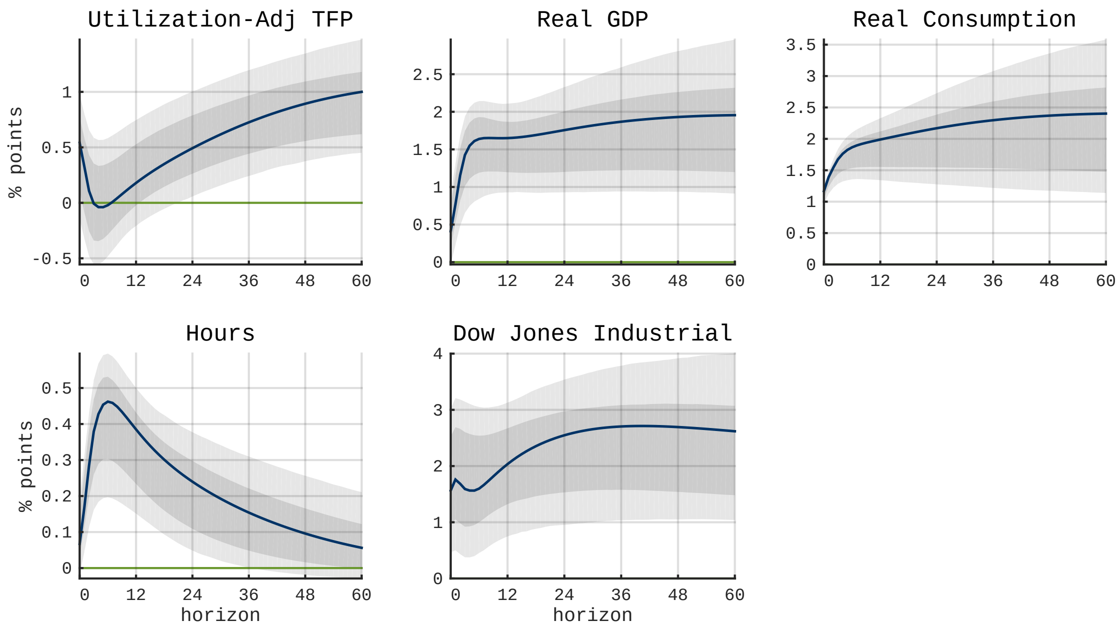


*Note*: Modal responses to a technology news shock identified with patent-based IV. Estimation sample 1960-I:2019-IV. Identification sample 1982-I:2006-IV. Shaded areas denote 68% and 90% posterior credible sets. Horizons in quarters.

The IRFs in Figure 2 are compatible with a 'news-driven' business cycle view in which macroeconomic aggregates react positively to positive news, and a business cycle expansion arises in anticipation of potential future technological improvements. Notwithstanding the minimal set of identifying restrictions, the pattern of IRFs recovered by our IV shares many similarities with those in prominent studies such as Beaudry and Portier (2006) and Barsky and Sims (2011), as we report in Figure D.1 in the Appendix. What is remarkable in this context is that the negligible impact response of TFP, the stock market pricing-in the news on impact, and, as we discuss below, the shock having maximum explanatory power for TFP at long horizons – assumed for identification in these earlier studies –, become instead results in our setting. The magnitude of the peak effects is also in line with previous literature (e.g. Barsky and Sims, 2011; Kurmann and Sims, 2021).

The identification is robust to removing the controls for other contemporaneous policy shocks, and to restricting the identification sample to start in 1995-III, which removes the regulation spikes and uses the portion of the sample where arguably patents were more

Figure 3: Shares of TFP explained variance in the 5-variable VAR

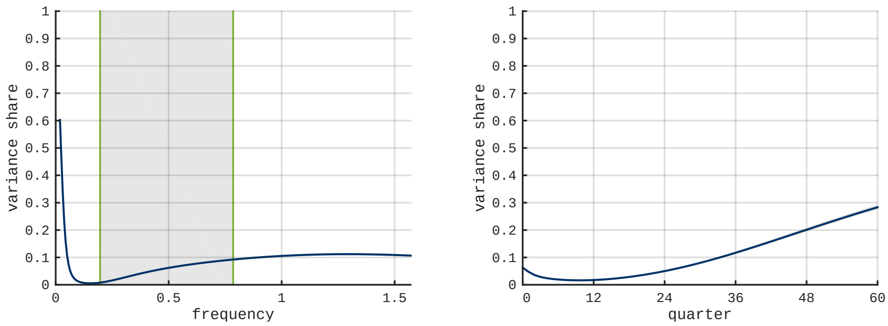


*Note*: Share of TFP error variance accounted for by technology news shock identified with patent-based IV. VAR(4) with standard macroeconomic priors. Estimation sample 1960-I : 2019-IV; Identification sample 1982-I : 2006-IV. In the left panel the shaded area delimits business cycle frequencies (between 8 and 32 quarters).

informative for future innovations.[24] In both these cases, for most of the variables the differences are minimal; for TFP, output and consumption the IRFs lie within the error bands of our baseline estimates for the most part. Some qualitative differences arise in the response of hours and the stock market, but do not alter our conclusions (see Figure D.2 in the Appendix).

The identification is also robust to only using ex-post granted patents in the construction of the IV, which corresponds to assigning a zero weight to patent applications that are eventually unsuccessful. And to alternative weighting schemes, as we discuss in detail in Appendix E. Using only ex-post granted patents to construct the IV yields somewhat stronger responses for hours and GDP. It is possible that ex-post granted patents may be embedding a somewhat stronger signal. Equally, the alternative dataset that we use for these robustness tests only including large firms may also have a bearing on the response of aggregate output and hours (see Figure E.2 in the Appendix).[25]

To complete the discussion, Figure 3 reports the share of TFP variance that is ac-

[24] See discussion in Section 2. The post-95 identification sample serves as a useful illustration, but it is based on a limited number of observations. To further evaluate the role of the TRIPS spike we have replaced the IV with a dummy variable that is equal to 1 in 1995-II, and zero otherwise. The TRIPS dummy recovers a different pattern of IRFs, suggesting that while important for the identification, the TRIPS spike is not entirely driving the results.

[25] The data used for this robustness exercise is from Kogan et al. (2017), that records information on individual patents granted by the USPTO to large US corporations for which a company match exists in the CRPS dataset. Among other things, for each patent the dataset reports information on application and grant dates, forward citations, and economic value. See Appendix E.

counted for by technology news shocks as identified by the IV.[26] Even if we have not imposed any such restriction ex ante, the shock recovered bu the IV is most explanatory for TFP at long horizons, and at very low frequencies. This is a pattern that we confirm also in the larger VAR of the next section, and that is consistent with the shock being a driver of the long-run component of aggregate productivity. The shares of aggregate fluctuations accounted for by the shock are however implausibly high in this illustrative VAR, reaching up to 80% for consumption and output. Two features the VAR are likely to account for such large variance shares. First, the 5-variable VAR is not informationally sufficient (see Forni and Gambetti, 2014); as noted in Section 3.1, this may introduce a bias in the forecast error variance decompositions (see Forni et al., 2019). Second, and related, the 5-variable VAR is likely not to be a plausible representation of the data generating process, since it is likely to omit other relevant variables. As is well known, this type of misspecification biases the estimation of the VAR coefficients, with the resulting distortions becoming more prevalent as the horizon increases. While this applies to both IRFs and error variance decompositions, the latter, calculated as ratios of potentially inconsistent quantities, are in practice especially sensitive to this type of misspecification (see Braun and Mittnik, 1993).

# 4 Technology News Shocks and Business Cycles

To study the propagation of technology news shocks to the broader economy we use a larger 16-variable VAR. The variables included cover real and nominal macroeconomic aggregates, financial markets, and expectations. This larger system allows us to characterize more carefully the role played by the different transmission channels, and the importance of these structural disturbances in the origination of economic fluctuations.

## 4.1 Dynamic Responses

As for the 5-variable VAR, we include 4 lags, and estimate the coefficients using standard Normal-Inverse Wishart priors over the sample 1960-I:2019-IV. With the exception of

[26] Variance decompositions for all variables are in Figures D.3 and D.4 in the Appendix. The algorithm is discussed in Appendix C.

Figure 4: Propagation of Technology News Shocks

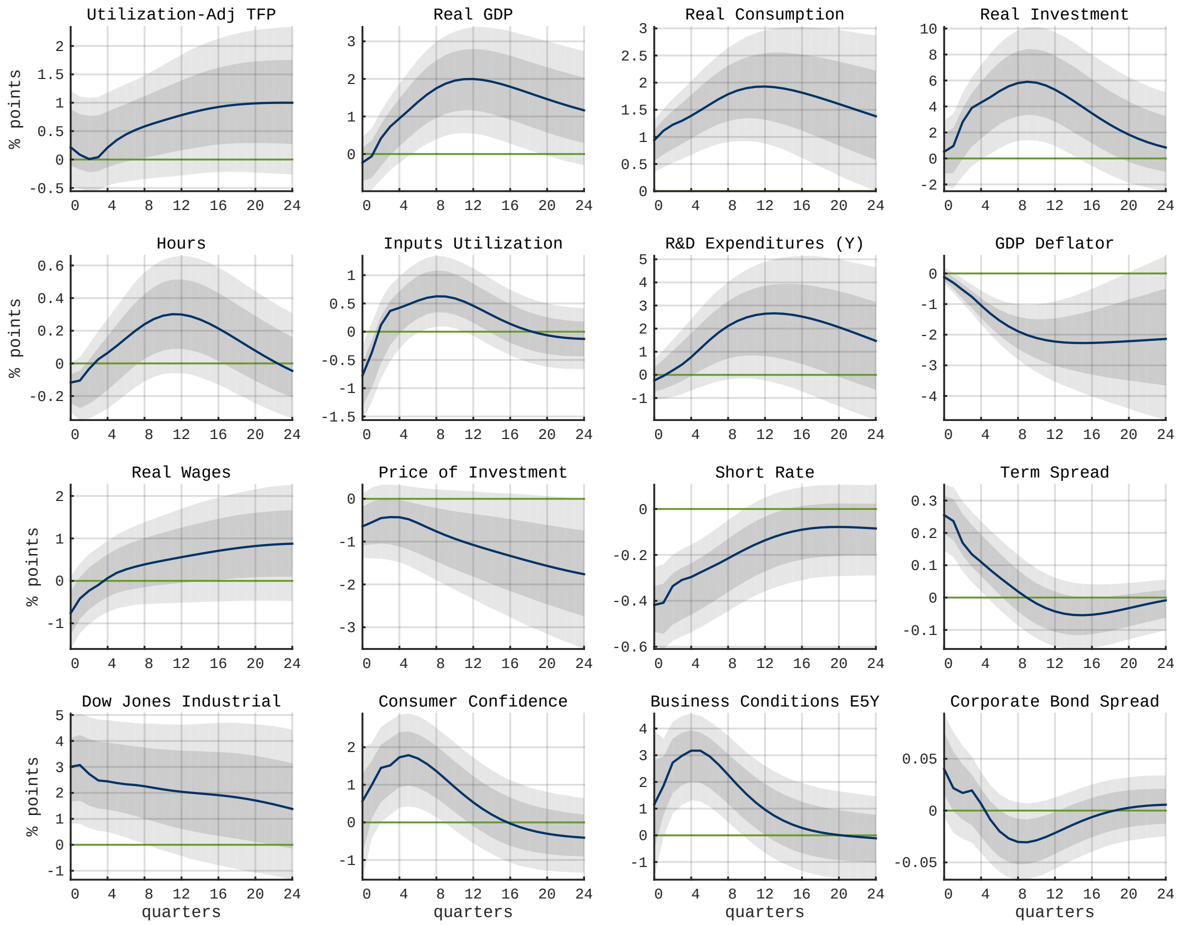


*Note*: Modal response to a technology news shock identified with patent-based external instrument. VAR(4). Estimation sample 1960-I:2019-IV. Identification sample 1982-I:2006-IV. Shaded areas denote 68% and 90% posterior credible sets.

interest rates and spreads, all the variables enter the specification in log levels, and are deflated and expressed in per-capita terms where appropriate. A complete description of the data and transformations is reported in Appendix B.

The IRFs to a positive technology news shock identified with the IV are reported in Figure 4. These are IRFs at the mode of the posterior distribution of the parameters, and are scaled such that the peak response of TFP equals 1% in annualized terms. Shaded areas correspond to 68% and 90% posterior credible sets. Robustness of our results is discussed below and the associated charts are reported in Appendix F.

**Productivity & Quantities** Most of the considerations made in the previous section carry through in the larger VAR. The initial response of TFP is muted, and the response becomes significant only years after the shock hits, albeit the estimates become less precise for this variable in the larger VAR. Consumption rises immediately, and remains elevated throughout. Output and investment do not respond on impact, and then rise persistently to reach a peak after about two years. R&D expenditures also rise eventually, pushed by the increase in investment and output. The magnitude of the responses is economically important. The output rise reaches 2 percentage points at peak, while investment increases by 6pp in annual terms. Total hours worked also rise robustly at the two year horizon, but fall marginally on impact. This impact negative response is significant but reabsorbed in the span of a few months. Inputs utilization also falls significantly on impact. This variable, distributed by Fernald (2014), combines estimates of both labor and capital utilization. Given the muted response of investment on impact, it is likely that the fall in utilization may be primarily driven by a fall in labor inputs.

Relative to the 5-variable VAR we note a few important differences. First, the impact response of hours worked and output become less positive. In particular, while the output response becomes zero, aggregate hours decline. Second, the IRFs become less persistent. This is due to nominal variables absorbing much of the persistence in the estimated VAR coefficients. Figure F.3 in the Appendix plots the response functions over 60 quarters for the baseline VAR and a VAR that excludes prices and wages. For most variables, the IRFs are equivalent at all horizons, but for TFP, output, consumption, investment, R&D and the stock market the exclusion of nominal variables makes the responses significantly more persistent, and more in line with the shock capturing permanent changes in productivity. Because prices and wages are important to study the propagation of technology news shocks, we have preferred to keep these variables in our baseline, even if this leads to a fall in persistence of the responses at long horizons. The correlation between the shocks estimated in the baseline VAR and the VAR that excludes nominal variables is equal to 0.97.

While the responses are somewhat delayed, also in the larger VAR they are consistent with positive technology news prompting a broad-based expansionary business cycle phase whereby all macroeconomic aggregates are significantly higher at the two-year mark, and long before any material increase in TFP is recorded. The sluggish response of R&D expenditures is also in line with this interpretation. In this sense, these results align with the 'news view' of Beaudry and Portier (2006); Beaudry and Lucke (2010) according to which the economy responds to current news in anticipation of potential future technological improvements. As we shall see in the reminder of the section, however, this broad-based expansion is not enough to make this type of shocks a main driver of business cycles, due to the modest share of fluctuations that it accounts for at the relevant frequencies. The initial contraction in total hours worked, and the short-lived deterioration of labor market conditions more generally, turn out to be an important element in understanding the response of consumers' expectations to technology news shocks, and we discuss it in greater detail below.

**Prices & Wages** In accordance with earlier studies, we find that technology news shocks are disinflationary (Jinnai, 2013; Kurmann and Otrok, 2017). Importantly, however, and consistent with nominal rigidities preventing an immediate adjustment, we find that the response of the price level is subdued initially, and only slowly builds up over time to reach a peak of about -2pp at the two year horizon in annualized terms. This translates into a response of inflation that is muted on impact, and followed by a negative hump-shape that reaches a peak of negative 20bps at the two year horizons, and reverts to zero thereafter. The muted impact response of inflation contrasts with findings in some earlier studies that document instead a sharp initial decline in prices (see e.g. Barsky and Sims, 2011; Barsky et al., 2015). Aggregate real wages fall marginally on impact to improve at longer horizons. Coupled with the response of aggregate prices, this points toward a short-lived decline in aggregate nominal wages. The response of the relative price of investment goods, that suffers a minor contraction on impact and keeps adjusting over time, indicates that the identified news shock makes investment goods progressively cheaper relative to consumption goods. Hence, the shock has some of the flavor of the investment-specific technological (IST) improvements of e.g. Fisher (2006) and Justiniano et al. (2010, 2011).[27]

[27] A similar response of the relative price of investments is reported in Kurmann and Sims (2021).

**Financial Markets & Consumers' Expectations** As in the 5-variable VAR, the stock market is quick in pricing-in positive news, and jumps up strongly on impact. The response of the stock market is stronger when the Dow Jones Industrial Average index is used compared to using broader indices such as the S&P 500. This is likely due to the DJIA including many of the heavy-weight information-technology companies, presumably those mostly affected by these types of shocks over the identification sample considered.

The disinflationary feature of the identified shock induces a significant endogenous response of the monetary authority, that responds more than proportionally to the decline in (expected) inflation. Due to the sample considered including the zero-lower-bound (ZLB) period, we use the one-year nominal interest rate as our measure for the short-term policy rate. The one-year rate falls by about 40 basis points on impact, which is almost twice the size of the peak decline of inflation. This implies that shorter maturity interest rates are likely to fall by more, and hence that short-term real rates fall following the shock. The slope of the yield curve, here measured as the spread between the 10-year and the 1-year Treasury rates, rises by about 25 bps on impact, mainly driven by changes at the short end, and implying a 15 bps fall in long term yields. The response of the yield curve is qualitatively similar to what documented in Kurmann and Otrok (2013), but the magnitudes in our case are significantly smaller. Comparing the responses of the short- and long-term rates, we note that the 1-year rate returns to trend relatively quickly, and is hence likely not to fully account for the impact fall in the 10-year Treasury yield. This implies that following a technology news shock risk premia decline.[28] In turn, this can act as an amplification mechanism for the propagation of news shocks. In contrast, the response of the BAA-AAA corporate bond spread is essentially flat. In Figure F.4 in the Appendix, we verify that neither the global financial crisis nor the ZLB sample drive or affect our results.

Finally, Figure 4 reports the responses of a consumer confidence indicator and a business confidence indicator reflecting expectations about economic conditions over a

[28] See Figure F.9 in the Appendix. This finding aligns with those in Crump et al. (2016). We use the VAR to decompose the response in the 10-year rate into its expectations and term-premium components by noting that, net of risk considerations, holding a 10-year bond should be equivalent to rolling 1-year bonds over 10 years. We calculate horizon $h$ term premium responses as the difference between the horizon $h$ response of the 10-year rate, and the average expected response of the 1-year rate at horizons $h, h+4, \ldots, h+36$.

horizon of 5 years, both taken from the Michigan Survey of Consumers. Interestingly, we find that while both measures robustly rise at medium horizons, they do not do so on impact. While impact responses are numerically positive, they are not significant at conventional levels. This result contrasts with previous findings in the literature (e.g. Barsky and Sims, 2011; Kurmann and Sims, 2021), but is consistent with consumers operating in a noisy information environment, and as a result overweighting the responses of current economic conditions when forming their expectations about the future. We return to this issue in greater detail below.

### 4.2 Variance Shares

Table 2 reports the average shares of explained variation over selected frequency intervals for all variables in our VAR. Specifically, the columns in Table 2 report the percentage share of variance accounted for by the identified shock in the short-run (average over frequencies corresponding to a period between 1 and 2 years), over the business cycle (between 2 and 8 years), and in the medium- and the long-run (between 8 and 25 years, and 50 and 60 years respectively).[29] Variance shares at all frequencies between 1 and 100 years are reported in Figure 5 for a selection of variables, and in Figure F.1 in the Appendix for the remainder of entries in our VAR. In the figure, the shaded areas highlight business cycle frequencies. The algorithm used for the decomposition builds on Altig et al. (2011) and is described in detail in Appendix C. The advantage of looking at variance decompositions in the frequency domain is that it allows us to separate among long, medium, and short-run fluctuations more clearly than a standard forecast error variance decomposition in the time domain.[30]

A few results are worth highlighting. First, similar to what found in the 5-variable VAR, the shock recovered by the IV is mostly explanatory for TFP in the very long run, where it accounts for about 20% of the overall variation. Conversely, the contribution

[29]Recall $\omega = 2\pi/t$, where $t$ denotes time and $\omega$ denotes the frequency. A period of 1 year (4 quarters) corresponds to $\omega \simeq 1.57$, while 100 years yield $\omega \simeq 0.02$. Business cycle frequencies, typically set between 8 and 32 quarters, correspond to frequencies between [0.2 0.8].

[30]Intuitively, even at relatively short forecast horizons, FEVDs in the time domain combine fluctuations at all frequencies. Because each horizon is a mixture of short, medium and long term components, evaluating the contribution of shocks at business cycle frequencies becomes more problematic. For comparison, time-based forecast error variance decompositions are reported in Figure F.2 in the Appendix.

Table 2: Error Variance Decomposition

| | SHORT RUN [ 1 - 2 years ] | BUSINESS CYCLE [ 2 - 8 years ] | MEDIUM RUN [ 8 - 25 years ] | LONG RUN [ 50 - 60 years ] |
|---|---|---|---|---|
| Utilization-Adj TFP | 0.65 | 0.63 | 5.49 | 20.69 |
| Real GDP | 2.28 | 5.41 | 16.28 | 14.48 |
| Real Consumption | 11.68 | 8.51 | 25.53 | 13.40 |
| Real Investment | 1.37 | 5.01 | 16.33 | 19.37 |
| Hours | 3.61 | 3.35 | 8.96 | 12.73 |
| Inputs Utilization | 6.78 | 3.41 | 4.91 | 1.27 |
| R&D Expenditures (Y) | 0.46 | 6.37 | 8.69 | 1.99 |
| GDP Deflator | 0.69 | 14.94 | 11.72 | 3.30 |
| Real Wages | 3.06 | 2.24 | 4.70 | 12.41 |
| Price of Investment | 2.72 | 0.63 | 9.66 | 17.29 |
| Short Rate | 17.20 | 11.11 | 8.27 | 8.13 |
| Term Spread | 18.42 | 11.43 | 6.63 | 3.16 |
| Dow Jones | 7.36 | 5.49 | 10.20 | 4.35 |
| Consumer Confidence | 0.50 | 6.21 | 11.80 | 9.13 |
| Business Conditions E5Y | 0.66 | 9.74 | 14.26 | 8.48 |
| Corporate Bond Spread | 0.85 | 3.68 | 2.75 | 3.21 |

*Notes:* Average percentage share of variance accounted for by the identified technology news shock over different frequency intervals. Estimation sample 1960:I - 2019:IV. Identification sample 1982:I - 2006:IV.

of the shock to higher frequency fluctuations in productivity is negligible. Hence, while we have not imposed any such restriction ex ante, the recovered shock turns out to be mostly a driver of the trend component of TFP. Second, the shock is responsible for a small fraction of the fluctuations in both consumption and hours at business cycle frequencies, but it accounts for over a fifth of the variation in consumption, and about 10% of that in labor inputs in the long-run. Moreover, the shock explains about 15% of the variation in output and investment in the long-run. These numbers should be considered as a conservative estimate, due to the very long sample used in the VAR estimation relative to the length of the instrument. For example, when the VAR is estimated over the pre-crisis years only, the share of explained variation of consumption, hours and investment at business cycle frequencies rises to 10%. These shares are far from expressing the bulk of the business cycle variation in these variables, confirming the disconnect between drivers of business cycles and of long-run fluctuations discussed in Angeletos et al. (2020). Third, the shock explains around a 15% of the medium-run variance of the stock market, and

Figure 5: Shares of Explained Variance

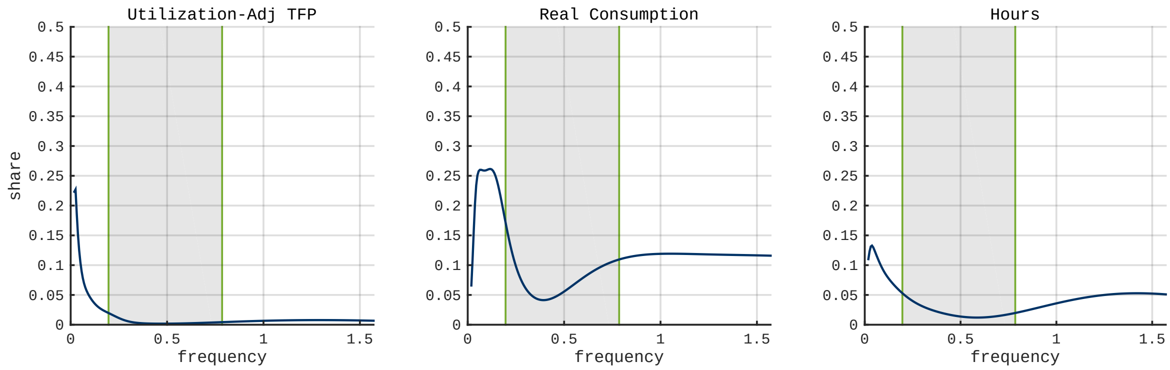


*Note*: Share of error variance accounted for by technology news shock identified with patent-based external instrument. VAR(4). Estimation sample 1960-I : 2019-IV. Identification sample 1982-I : 2006-IV. Shaded areas delimits business cycle frequencies (between 8 and 32 quarters). Frequencies on the $x$ axis cover a period from 1 (highest) to 100 (lowest) years.

is responsible for about 20% of variation in the yield curve in the short-term. A note of caution is in order. As discussed, the IV only captures technology news shocks insofar as these are captured by the patenting process, and may therefore leave other sources of variation in long-term productivity unaccounted for. As a result, caution should be used when comparing the shares of forecast error variance with those reported in other studies.

Finally, it is worth mentioning that the shock is a significant driver of the trend variation of the relative price of investments (20% at lowest frequencies). This variable is used in Justiniano et al. (2010, 2011) to disentangle IST shocks from neutral technology shocks. Our interpretation of this result is that the IV recovers technology news shocks that operate also through embodied technological change.[31]

## 4.3 Labor Market Response and Consumers' Expectations

According to the responses in Figure 4, the immediate reaction of the labor market to technology news shocks can be interpreted as a temporary leftward shift in the aggre-

[31]Whether this is the main channel through which the shock we identify operates remains however unclear. In a recent contribution Chen and Wemy (2015) show that IST shocks are an important driver of long-run movements in aggregate TFP, which is a useful complement to our findings. In fact, this paper shows that shocks that maximize the long-run FEV of TFP and those that maximize that of the relative price of investment are almost perfectly collinear. Due to our identification being fundamentally different, it is not clear that this interpretation can seamlessly applied in our context.

gate demand of labor, which results in a short-lived contraction of both hours worked and wages. This is compatible with e.g. firms switching to more capital-intensive technologies, automation, or with the skills of incumbent workers becoming obsolete as new technologies are introduced, consistent with what documented in e.g. Kogan, Papanikolaou, Schmidt and Seegmiller (2021). We do not take a stand on the microfoundation, but in this last section we take a closer look at the labor market response, and how it interacts with consumers' expectations. Contrary to what reported in previous studies, consumer expectations fail to fully adjust to the shock on impact, and the responses of indices of consumer confidence only become significantly positive after a few quarters.[32]

To explore this link further, in the VAR we replace total hours worked with the unemployment rate and the labor participation rate. We also add consumers' expectations about unemployment one year hence, again extracted from the Michigan Survey of Consumers; the survey asks respondents whether they expect unemployment over the next 12 months to be higher, lower or about the same as current. Figure 6 collects the responses (top panels) and variance shares (bottom panels) for these three variables, full IRFs are reported in Figure F.6 in the Appendix. While not based on an exact decomposition, the chart reveals that the variation in hours worked is unlikely to be accounted for by changes in labor participation rate, whose response is essentially flat at all horizons. Conversely, the unemployment rate rises upon realization of the shock, to revert at medium horizons. Moreover, the shock is responsible for about 15% of the variation in the unemployment rate in the short-run (see also Faccini and Melosi, 2018, for the role played by technology news on employment and its forecasts). Perhaps more interesting, however, is the response of consumers' expectations about future unemployment prospects. Consistent with the immediate rise in unemployment, and in apparent contrast with there being underlying positive news, the share of consumers that expect a higher unemployment rate going forward rises sharply, with the peak response realized well within the first year. We argue that the upward revision in consumers' expectations about unemployment can help to account for the muted impact response of the consumer confidence indicators.

The context of technology news shocks offers a natural environment in which different

[32] In some specifications, while still not significant at conventional levels, the impact response of indices of consumer confidence can be numerically negative. Thus, while the sign of the impact response is somewhat uncertain, the conclusion that the impact response is not different from zero is robust.

Figure 6: Unemployment and Unemployment Expectations

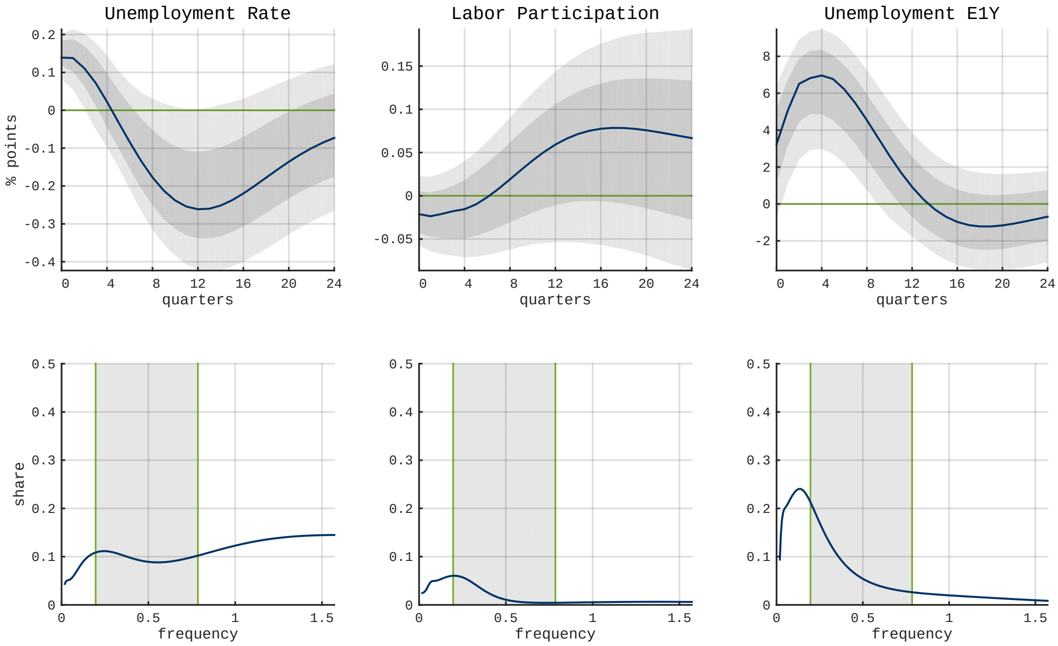


*Note*: Impulse response functions (top panels) and shares of explained variance (bottom panels) for the unemployment rate, the rate of labor participation, and the 1-year-ahead unemployment expectation. Survey forecasts are from the Michigan Survey of Consumers. VAR(4). Estimation sample 1960-I:2019-IV; Identification sample 1982-I:2006-IV.

agents in the economy are plausibly informed to different degrees. For example, it is plausible to postulate that market participants are more attentive, or more able to incorporate these types of news, relative to the average consumer. Here we do not attempt to speculate on the ultimate sources of such rigidities to information processing, but note that the IRFs to consumers' expectations about unemployment, and about current and expected business conditions fit nicely within the predictions of models of noisy information (e.g. Woodford, 2003; Sims, 2003; Mackowiak and Wiederholt, 2009). Consider the simple framework in which agents use a Kalman Filter to form expectations about the future. The lower the signal-to-noise ratio in the information they receive, the less the new information will be weighted-in in their expectations about the future, the more these expectations will be based on current realizations/past signals. Survey-based evidence reported in Coibion and Gorodnichenko (2012, 2015) suggests that this framework offers a plausible characterization of the process of expectation formation. News about future technological changes can be thought of as a quintessential signal extraction problem. Blanchard, L'Huillier and Lorenzoni (2013) in particular consider the case in which technology is driven by both temporary and permanent shocks (i.e. shocks that have long-lasting effects on the level of technology), and agents observe a noisy signal of the permanent component of technology. Agents are not able to disentangle news from noise. In their model the noisier the signal, the slower the consumption adjustment, the more likely that shocks to the permanent component result in an initial fall in employment.

We think of the initial rise in both actual and expected unemployment (Figure 6) as compatible with such noise-ridden environment, and with consumers overweighting the negative impact response of labor market variables to the shock. In turn, this can help explain the initial muted response of consumer confidence about both current and expected business conditions documented in Figure 4. In this respect, our results suggest caution in interpreting innovations in consumer confidence indicators as a 'pure' measure of news (Cochrane, 1994; Barsky and Sims, 2012). In fact, when we compare responses to our news shock with those elicited by a positive contemporaneous TFP innovation, we find that consumer confidence jumps up on impact only in the latter case (see Figure F.8 in the Appendix).[33] Finally, it is worth noting that in contrast to most papers in the existing empirical literature where the TFP rise is typically more sudden, the IV identifies technology news shocks that take years to materialize. Requiring that the average consumer is able to discern these shocks clearly and on impact may be somewhat unrealistic in this setting.

# 5 Conclusions

How does the aggregate economy react to a shock that raises expectations about future productivity growth? In this paper we have provided an answer to this question by proposing a novel patent-based instrumental variable for the identification of technology news shocks that allows us to dispense from all the traditional assumptions used in the empirical news literature. The IV is constructed as the component of patent applications that is orthogonal to pre-existing beliefs about the macro outlook, and to other

[33]While both indices of consumer confidence may embed some elements of forward-lookingness, in the time-series they track contemporaneous annual GDP growth particularly well. In this sense they may more appropriately be interpretable as coincident indicators, rather than forward-looking ones.

contemporaneous policy shocks. Our sole identifying assumption is no other structural disturbances affect the economy via the IV except for contemporaneous technology news.

The IV recovers technology news shocks that have essentially no impact on current productivity, but are a significant driver of its trend component. Positive news give rise to a broad-based business cycle expansion in anticipation of future technological improvements. The stock market prices-in news shocks on impact, and output, investment, consumption and hours all increase long before any material improvement in TFP is recorded. However, the shock only accounts for a modest share of economic fluctuations at business cycle frequencies, and is hence not a main driver of business cycles.

The immediate response of the labor market to technology news shocks as identified by our IV is best summarized as a leftward shift in aggregate labor demand. This is short-lived, and compatible with e.g. firms switching to more capital-intensive technologies, or with the skills of incumbent workers becoming obsolete as new technologies are introduced. But it is sufficient to drag down consumers' expectations that only incorporate the news with delay, and only as the outlook starts to improve. Models that can rationalize these dynamics embed news in frameworks in which, as is plausible, agents only observe a noisy signal about macro fundamentals, and are likely to overweigh current conditions when forming their expectations about the future.

Our paper is fundamentally empirical in nature, but our findings suggest that the heterogeneous degree to which expectations of firms, financial markets and consumers respond to news shocks plays an important role in their propagation, and offer new insights for the modelling of these types of disturbances.

# References


Adams, Kay, Douglas Kim, Frederick L. Joutz, Robert P. Trost, and Gus Mastrogianis (1997) "Modeling and forecasting U.S. Patent application filings," *Journal of Policy Modeling*, Vol. 19, No. 5, pp. 491–535, October.

Alexopoulos, Michelle (2011) "Read All about It!! What Happens Following a Technology Shock?," *American Economic Review*, Vol. 101, No. 4, pp. 1144–1179, June.

Altig, David, Lawrence Christiano, Martin Eichenbaum, and Jesper Linde (2011) "Firm-Specific Capital, Nominal Rigidities and the Business Cycle," *Review of Economic Dynamics*, Vol. 14, No. 2, pp. 225–247, April.

Angeletos, George-Marios, Fabrice Collard, and Harris Dellas (2020) "Business-Cycle Anatomy," *American Economic Review*, Vol. 110, No. 10, pp. 3030–70, October.

Arezki, Rabah, Valerie A. Ramey, and Liugang Sheng (2017) "News Shocks in Open Economies: Evidence from Giant Oil Discoveries," *The Quarterly Journal of Economics*, Vol. 132, No. 1, pp. 103–155.

Baron, J. and J. Schmidt (2014) "Technological Standardization, Endogenous Productivity and Transitory Dynamics," Working Papers 503, Banque de France.

Barsky, Robert B. and Eric R. Sims (2009) "News Shocks," Working Paper 15312, National Bureau of Economic Research.

——— (2011) "News shocks and business cycles," *Journal of Monetary Economics*, Vol. 58, No. 3, pp. 273–289.

——— (2012) "Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence," *American Economic Review*, Vol. 102, No. 4, pp. 1343–77, June.

Barsky, Robert B., Susanto Basu, and Keyoung Lee (2015) "Whither News Shocks?," *NBER Macroeconomics Annual*, Vol. 29, No. 1, pp. 225–264.

Basu, Susanto, John G. Fernald, and Miles S. Kimball (2006) "Are Technology Improvements Contractionary?" *American Economic Review*, Vol. 96, No. 5, pp. 1418–1448, December.

Beaudry, Paul and Bernd Lucke (2010) "Letting Different Views about Business Cycles Compete," *NBER Macroeconomics Annual*, Vol. 24, No. 1, pp. 413–456.

Beaudry, Paul and Franck Portier (2004) "An exploration into Pigou's theory of cycles," *Journal of Monetary Economics*, Vol. 51, No. 6, pp. 1183–1216, September.

——— (2006) "Stock Prices, News, and Economic Fluctuations," *American Economic Review*, Vol. 96, No. 4, pp. 1293–1307, September.

——— (2014) "News-Driven Business Cycles: Insights and Challenges," *Journal of Economic Literature*, Vol. 52, No. 4, pp. 993–1074, December.

Beaudry, Paul, Patrick FËve, Alain Guay, and Franck Portier (2019) "When is nonfundamentalness in SVARs a real problem?" *Review of Economic Dynamics*, Vol. 34, pp. 221–243.

Blanchard, Olivier J., Jean-Paul L'Huillier, and Guido Lorenzoni (2013) "News, Noise, and Fluctuations: An Empirical Exploration," *American Economic Review*, Vol. 103, No. 7, pp. 3045–3070, December.

Braun, Phillip A. and Stefan Mittnik (1993) "Misspecifications in vector autoregressions and their effects on impulse responses and variance decompositions," *Journal of Econometrics*, Vol. 59, No. 3, pp. 319–341, October.

Cascaldi-Garcia, Danilo and Marija Vukotić (2022) "Patent-Based News Shocks," *The Review of Economics and Statistics*, Vol. 104, No. 1, pp. 51–66, 01.

Chahrour, Ryan and Kyle Jurado (2018) "News or Noise? The Missing Link," *American Economic Review*, Vol. 108, No. 7, pp. 1702–36, July.

Chen, Kaiji and Edouard Wemy (2015) "Investment-specific technological changes: The source of long-run TFP fluctuations," *European Economic Review*, Vol. 80, pp. 230–252.

Christiano, Lawrence J., Martin Eichenbaum, and Robert Vigfusson (2003) "What Happens After a Technology Shock?," NBER Working Papers 9819, National Bureau of Economic Research, Inc.

Christiansen, Lone Engbo (2008) "Do Technology Shocks Lead to Productivity Slowdowns? Evidence from Patent Data," IMF Working Papers 08/24, International Monetary Fund.

Cochrane, John H. (1994) "Shocks," *Carnegie-Rochester Conference Series on Public Policy*, Vol. 41, No. 1, pp. 295–364, December.

Coibion, Olivier and Yuriy Gorodnichenko (2012) "What Can Survey Forecasts Tell Us about Information Rigidities?" *Journal of Political Economy*, Vol. 120, No. 1, pp. 116 – 159.

——— (2015) "Information Rigidity and the Expectations Formation Process: A Simple Framework and New Facts," *American Economic Review*, Vol. 105, No. 8, pp. 2644–78.

Crump, Richard K., Stefano Eusepi, and Emanuel Moench (2016) "The term structure of expectations and bond yields," Staff Reports, revised 2018 775, Federal Reserve Bank of New York.

Doan, Thomas, Robert B. Litterman, and Christopher A. Sims (1983) "Forecasting and Conditional Projection Using Realistic Prior Distributions," NBER Working Papers 1202, National Bureau of Economic Research, Inc.

Encaoua, David, Dominique Guellec, and Catalina Martínez (2006) "Patent systems for encouraging innovation: Lessons from economic analysis," *Research Policy*, Vol. 35, No. 9, pp. 1423 – 1440.

Faccini, Renato and Leonardo Melosi (2018) "The Role of News about TFP in U.S. Recessions and Booms," Working Paper Series WP-2018-6, Federal Reserve Bank of Chicago.

Fernald, John G. (2014) "A quarterly, utilization-adjusted series on total factor productivity," Working Paper Series 2012-19, Federal Reserve Bank of San Francisco.

Fisher, Jonas D. M. (2006) "The Dynamic Effects of Neutral and Investment-Specific Technology Shocks," *Journal of Political Economy*, Vol. 114, No. 3, pp. 413–451, June.

Forni, Mario and Luca Gambetti (2011) "Testing for Sufficient Information in Structural VARs," CEPR Discussion Papers 8209, C.E.P.R. Discussion Papers.

——— (2014) "Sufficient information in structural VARs," *Journal of Monetary Economics*, Vol. 66, No. C, pp. 124–136.

Forni, Mario, Luca Gambetti, and Luca Sala (2014) "No News in Business Cycles," *Economic Journal*, Vol. 124, No. 581, pp. 1168–1191, December.

——— (2019) "Structural VARs and noninvertible macroeconomic models," *Journal of Applied Econometrics*, Vol. 34, No. 2, pp. 221–246.

Francis, Neville and Valerie A. Ramey (2005) "Is the technology-driven real business cycle hypothesis dead? Shocks and aggregate fluctuations revisited," *Journal of Monetary Economics*, Vol. 52, No. 8, pp. 1379–1399, November.

——— (2009) "Measures of per Capita Hours and Their Implications for the Technology-Hours Debate," *Journal of Money, Credit and Banking*, Vol. 41, No. 6, pp. 1071–1097, September.

Francis, Neville, Michael T. Owyang, Jennifer E. Roush, and Riccardo DiCecio (2014) "A Flexible Finite-Horizon Alternative to Long-Run Restrictions with an Application to Technology Shocks," *The Review of Economics and Statistics*, Vol. 96, No. 4, pp. 638–647.

Galí, Jordi (1999) "Technology, Employment, and the Business Cycle: Do Technology Shocks Explain Aggregate Fluctuations?" *American Economic Review*, Vol. 89, No. 1, pp. 249–271, March.

Giannone, Domenico and Lucrezia Reichlin (2006) "Does information help recovering structural shocks from past observations?," *Journal of the European Economic Association*, Vol. 4, No. 2-3, pp. 455–465, 04-05.

Giannone, Domenico, Michele Lenza, and Giorgio E. Primiceri (2015) "Prior Selection for Vector Autoregressions," *Review of Economics and Statistics*, Vol. 97, No. 2, pp. 436–451.

Gort, Michael and Steven Klepper (1982) "Time Paths in the Diffusion of Product Innovations," *The Economic Journal*, Vol. 92, No. 367, pp. 630–653.

Griliches, Zvi (1990) "Patent Statistics as Economic Indicators: A Survey," *Journal of Economic Literature*, Vol. 28, No. 4, pp. 1661–1707.

Hall, Bronwyn H. and Manuel Trajtenberg (2004) "Uncovering GPTS with Patent Data," Working Paper 10901, National Bureau of Economic Research.

Hall, Bronwyn H., Adam B. Jaffe, and Manuel Trajtenberg (2001) "The NBER Patent Citation Data File: Lessons, Insights and Methodological Tools," Working Paper 8498, National Bureau of Economic Research.

Jinnai, Ryo (2013) "News shocks and inflation," *Economics Letters*, Vol. 119, No. 2, pp. 176 – 179.

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti (2010) "Investment shocks and business cycles," *Journal of Monetary Economics*, Vol. 57, No. 2, pp. 132–145, March.

——— (2011) "Investment Shocks and the Relative Price of Investment," *Review of Economic Dynamics*, Vol. 14, No. 1, pp. 101–121, January.

Kadiyala, K Rao and Sune Karlsson (1997) "Numerical Methods for Estimation and Inference in Bayesian VAR-Models," *Journal of Applied Econometrics*, Vol. 12, No. 2, pp. 99–132, March-Apr.

Kogan, Leonid, Dimitris Papanikolaou, Lawrence D. W. Schmidt, and Bryan Seegmiller (2021) "Technology-Skill Complementarity and Labor Displacement: Evidence from Linking Two Centuries of Patents with Occupations," NBER Working Papers 29552, National Bureau of Economic Research, Inc.

Kogan, Leonid, Dimitris Papanikolaou, Amit Seru, and Noah Stoffman (2017) "Technological Innovation, Resource Allocation, and Growth," *The Quarterly Journal of Economics*, Vol. 132, No. 2, pp. 665–712, 03.

Kurmann, André and Christopher Otrok (2013) "News Shocks and the Slope of the Term Structure of Interest Rates," *American Economic Review*, Vol. 103, No. 6, pp. 2612–2632, October.

——— (2017) "News Shocks and Inflation: Lessons for New Keynesians." Unpublished.

Kurmann, André and Eric Sims (2021) "Revisions in Utilization-Adjusted TFP and Robust Identification of News Shocks," *The Review of Economics and Statistics*, Vol. 103, No. 2, pp. 216–235, 05.

Lach, Saul (1995) "Patents and productivity growth at the industry level: A first look," *Economics Letters*, Vol. 49, No. 1, pp. 101 – 108.

Leeper, Eric M., Todd B. Walker, and Shu-Chun Susan Yang (2013) "Fiscal Foresight and Information Flows," *Econometrica*, Vol. 81, No. 3, pp. 1115–1145.

Lerner, Josh and Amit Seru (2021) "The Use and Misuse of Patent Data: Issues for Finance and Beyond," *The Review of Financial Studies*, Vol. 35, No. 6, pp. 2667–2704, 07.

Litterman, Robert B (1986) "Forecasting with Bayesian Vector Autoregressions-Five Years of Experience," *Journal of Business & Economic Statistics*, Vol. 4, No. 1, pp. 25–38, January.

Mackowiak, Bartosz and Mirko Wiederholt (2009) "Optimal Sticky Prices under Rational Inattention," *American Economic Review*, Vol. 99, No. 3, pp. 769–803, June.

Marco, Alan C., Michael Carley, Steven Jackson, and Amanda Myers (2015) "The USPTO Historical Patent Data Files: Two Centuries of Innovation," USPTO Economic Working Papers 1, U.S. Patent and Trademark Office.

McCracken, Michael W. and Serena Ng (2016) "FRED-MD: A Monthly Database for Macroeconomic Research," *Journal of Business & Economic Statistics*, Vol. 34, No. 4, pp. 574–589.

Mertens, Karel and Morten O. Ravn (2011) "Technology-Hours Redux: Tax Changes and the Measurement of Technology Shocks," *NBER International Seminar on Macroeconomics*, Vol. 7, No. 1, pp. 41–76.

——— (2012) "Empirical Evidence on the Aggregate Effects of Anticipated and Unanticipated US Tax Policy Shocks," *American Economic Journal: Economic Policy*, Vol. 4, No. 2, pp. 145–181, May.

——— (2013) "The Dynamic Effects of Personal and Corporate Income Tax Changes in the United States," *American Economic Review*, Vol. 103, No. 4, pp. 1212–47, June.

Miranda-Agrippino, Silvia and Giovanni Ricco (2018) "Identification with External Instruments in Structural VARs under Partial Invertibility," Working Paper 24, OFCE.

Pagan, Adrian (1984) "Econometric Issues in the Analysis of Regressions with Generated Regressors," *International Economic Review*, Vol. 25, No. 1, pp. 221–247, February.

Pigou, A.C. (1927) *Industrial Fluctuations*: Macmillan and Company, limited.

Plagborg-Møller, Mikkel and Christian Wolf (2021) "Local Projections and VARs Estimate the Same Impulse Responses," *Econometrica*, Vol. 89, pp. 955–980, 01.

Ramey, Valerie A. (2016) "Macroeconomic Shocks and Their Propagation," in John B. Taylor and Harald Uhlig eds. *Handbook of Macroeconomics*, Vol. 2 of Handbook of Macroeconomics: Elsevier, Chap. 2, pp. 71 – 162.

Rogers, Everett M. (1962) *Diffusion of innovations*: The Free Press of Glencoe Division of The Macmillan Co., New York, 1st edition.

Romer, Christina D. and David H. Romer (2004) "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review*, Vol. 94, No. 4, pp. 1055–1084.

——— (2010) "The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks," *American Economic Review*, Vol. 100, No. 3, pp. 763–801, June.

Shea, John (1999) "What Do Technology Shocks Do?," in *NBER Macroeconomics Annual 1998, volume 13*: National Bureau of Economic Research, Inc, pp. 275–322.

Sims, Christopher A. (2003) "Implications of rational inattention," *Journal of Monetary Economics*, Vol. 50, No. 3, pp. 665 – 690. Swiss National Bank/Study Center Gerzensee Conference on Monetary Policy under Incomplete Information.

Sims, Eric R. (2012) "News, Non-Invertibility, and Structural VARs," Working Papers 013, University of Notre Dame, Department of Economics.

Stock, James H. and Mark W. Watson (2012) "Disentangling the Channels of the 2007-09 Recession," *Brookings Papers on Economic Activity*, Vol. 44, No. 1 (Spring, pp. 81–156.

——— (2018) "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments," *The Economic Journal*, Vol. 128, No. 610, pp. 917–948.

Uhlig, Harald (2004) "Do Technology Shocks Lead to a Fall in Total Hours Worked?," *Journal of the European Economic Association*, Vol. 2, No. 2-3, pp. 361–371, 04/05.

Woodford, Michael (2003) "Imperfect Common Knowledge and the Effects of Monetary Policy," in P. Aghion, R. Frydman, J. Stiglitz, and M. Woodford eds. *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund Phelps*: Princeton University Press.

# Appendix: Not For Publication

## A Additional Details on Instrument & Regression Tables

Figure A.1: Instrument for News Shocks

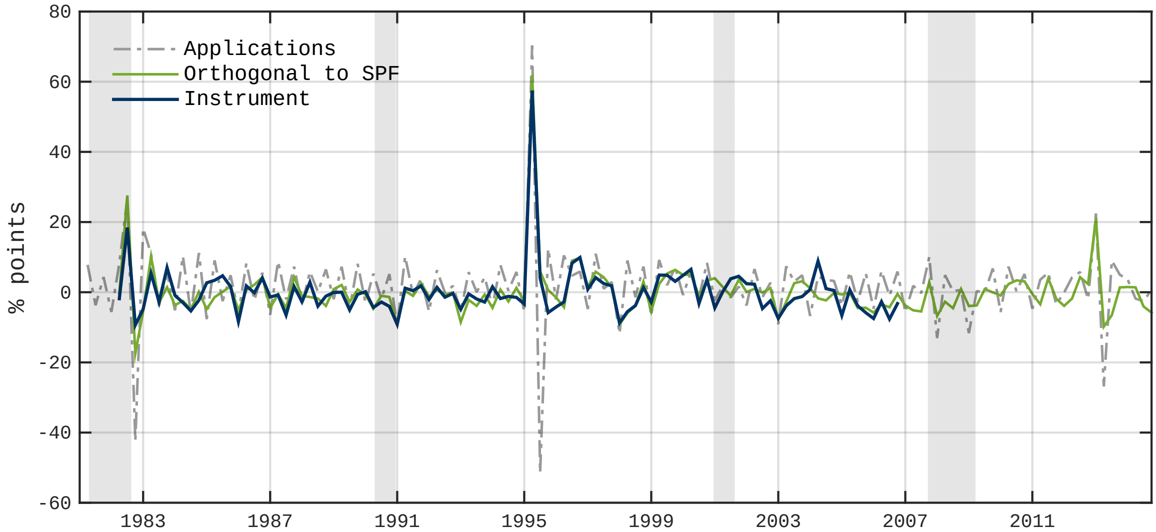


*Note*: Raw count of patent applications, quarterly growth rate (grey, dash-dotted line); instrument for news shocks (blue, solid), residuals of Eq. (1); residuals of Eq. (1) without policy controls, (green, solid). Shaded areas denote NBER recession episodes.

Table A.1: Dependence of Patent Applications on Pre-Existing Expectations

| | $\mathbb{E}_t[w_t]$ | $\mathbb{E}_t[w_{t+1}]$ | $\mathbb{E}_t[w_{t+4}]$ |
|---|---|---|---|
| Wald Test | 3.471 | 5.670 | 2.743 |
| p-value | 0.003 | 0.000 | 0.016 |
| Adj $R^2$ | 0.482 | 0.481 | 0.469 |
| N | 131 | 131 | 131 |

*Notes:* Dependent variable is the quarterly growth rate of patent applications. $\mathbb{E}_t[w_{t+h}]$ denotes SPF forecast for quarter $t+h$ published at $t$ conditional on $t-1$. $w_t$ includes real output growth, unemployment rate, inflation (GDP deflator), real federal government spending, real non-residential investments, and real corporate profits net of taxes. Numbers reported are Wald test statistics for joint significance of the SPF forecasts at each horizon. All the regressions include own 4 lags and constant.

Table A.2: Lagged Information in Patent Applications

| | $F_1$ | $F_2$ | $F_3$ | $F_4$ | $F_5$ | $F_6$ | $F_7$ |
|---|---|---|---|---|---|---|---|
| Wald Test | 6.901 | 0.475 | 0.365 | 1.548 | 1.160 | 1.284 | 0.582 |
| p-value | 0.000 | 0.754 | 0.834 | 0.193 | 0.332 | 0.280 | 0.676 |
| Adj $R^2$ | 0.504 | 0.436 | 0.432 | 0.480 | 0.459 | 0.459 | 0.439 |
| N | 131 | 131 | 131 | 131 | 131 | 131 | 131 |

*Notes:* Numbers reported are Wald test statistics for joint significance of the first 4 lags of each factor $F_t$. The factors are extracted from the quarterly dataset of McCracken and Ng (2016). The dependent variable is the quarterly growth rate of utility patent applications: $pa_t = 100(\ln PA_t - \ln PA_{t-1})$. All the regressions include own 4 lags and constant.

Table A.3: Dependence of Instrument on Pre-Existing Expectations

| | $\mathbb{E}_t[w_t]$ | $\mathbb{E}_t[w_{t+1}]$ | $\mathbb{E}_t[w_{t+4}]$ |
|---|---|---|---|
| Wald Test | 0.846 | 0.711 | 0.568 |
| p-value | 0.538 | 0.642 | 0.754 |
| Adj $R^2$ | −0.079 | −0.082 | −0.088 |
| N | 95 | 95 | 95 |

*Notes:* Dependent variable is the residual of Eq. (1). $\mathbb{E}_t[w_{t+h}]$ denotes SPF forecast for quarter $t+h$ published at $t$ conditional on $t-1$. $w_t$ includes real output growth, unemployment rate, inflation (GDP deflator), real federal government spending, real non-residential investments, and real corporate profits net of taxes. Numbers reported are Wald test statistics for joint significance of the SPF forecasts at each horizon. All the regressions include own 4 lags and constant.

Table A.4: Lagged Information in the Instrument

| | $F_1$ | $F_2$ | $F_3$ | $F_4$ | $F_5$ | $F_6$ | $F_7$ |
|---|---|---|---|---|---|---|---|
| Wald Test | 0.525 | 1.422 | 0.802 | 1.445 | 1.452 | 0.931 | 0.354 |
| p-value | 0.718 | 0.234 | 0.527 | 0.226 | 0.224 | 0.450 | 0.840 |
| Adj $R^2$ | −0.053 | −0.039 | −0.062 | −0.010 | −0.028 | −0.060 | −0.068 |
| N | 95 | 95 | 95 | 95 | 95 | 95 | 95 |

*Notes:* Numbers reported are Wald test statistics for joint significance of the first 4 lags of each factor $F_t$. The factors are extracted from the quarterly dataset of McCracken and Ng (2016). The dependent variable is the instrument (residuals of Eq. (1)). All the regressions include own 4 lags and constant.

## B Data in VAR

Table B.1 lists the variables included in the VAR. The construction of real consumption (RCONS), real investment (RINV), the relative price of investment (RPINV), and hours worked (HOURS) follows Justiniano et al. (2010, 2011); specifically,

$$RCON = 100 \times \ln\left(\frac{PCND + PCESV}{CNP16OV \times GDPDEF}\right)$$

$$RINV = 100 \times \ln\left(\frac{GPDI + PCDG}{CNP16OV \times GDPDEF}\right)$$

$$RPINV = 100 \times \ln\left(\frac{DDURRD3Q086SBEA + A006RD3Q086SBEA}{DNDGRD3Q086SBEA + DSERRD3Q086SBEA}\right)$$

$$HOURS = 100 \times \ln\left(\frac{HOANBS}{2080}\right),$$

where 2080 is the average numbers of hours worked in a year (i.e. 40 hours a week times 52 weeks). Consumption includes personal consumption expenditures in non-durable goods (PCND) and services (PCESV), whereas investment is constructed as the sum of private gross domestic investment (GPDI) and personal consumption expenditures in durable goods (PCDG). The relative price of investment goods is constructed as the ratio of the deflators of investment and consumption. Consistent with the definition above, these are constructed as the implicit price deflator for durable and investment, and the implicit price deflators for non-durable and services consumption respectively.

The level of Utilization-Adjusted TFP is obtained by cumulating the series of quarterly growth rates annualized of Fernald (2014). The short term rate and the yield curve slope are expressed in annualized terms. The yield curve slope (YCSLOPE) is constructed as the difference between the 10-year (DGS10) and 1-year (DGS1) Treasury constant-maturity rates. Variables are deflated using the GDP deflator, and transformed in per-capita terms by dividing for the trend in population (population variable: CNP16OV).

Table B.1: Variables Used

| Label | Variable Name | Source | FRED Codes | TREATMENT log | pc |
|---|---|---|---|---|---|
| TFPL | Utilization-Adj TFP | Fernald (2014)† | – | • | • |
| RGDP | Real GDP | FRED | GDPC1 | • | • |
| RCONS | Real Consumption | FRED | PCND; PCESV | • | • |
| RINV | Real Investment | FRED | GPDI; PCDG | • | • |
| RDGDP | R&D Expenditures (Y) | FRED | Y694RC1Q027SBEA | • | • |
| HOURS | Hours | FRED | HOANBS | • | • |
| UNRATE | Unemployment Rate | FRED | UNRATE | • | |
| LPR | Labor Force Participation Rate | FRED | CIVPART | • | |
| INPUTIL | Inputs Utilization | Fernald (2014)† | – | • | |
| GDPDEF | GDP Deflator | FRED | GDPDEF | • | |
| RPINV | Price of Investment | FRED | DDURRD3Q086SBEA; DNDGRD3Q086SBEA; DSERRD3Q086SBEA; A006RD3Q086SBEA | • | |
| RWAGE | Real Wages | FRED | COMPRNFB | • | |
| SHORTR | Short Rate | FRED | DGS1 | | |
| YCSLOPE | Term Spread | FRED | DGS1; DGS10 | | |
| SP500 | S&P 500 | DATASTREAM | – | • | |
| DJIA | Dow Jones Industrial Average | DATASTREAM | – | • | |
| CCONF | Consumer Confidence | UMICH | – | • | |
| BCE5Y | Business Conditions E5Y | UMICH | – | • | |
| UE1Y | Unemployment E1Y | UMICH | – | • | |
| CBSPREAD | Corporate Bond Spread | FRED | AAA; BAA | | |

*Notes:* Sources are: St Louis FRED Database (FRED); University of Michigan (UMICH) Survey of Consumers `https://data.sca.isr.umich.edu/charts.php`; † 2020 vintage of Fernald (2014) TFP series `https://www.frbsf.org/economic-research/indicators-data/total-factor-productivity-tfp/`. pc = per-capita.

# C Error Variance Decomposition

The content of this appendix extends on Altig et al. (2011). Let the Structural VAR be

$$B(L)y_t = B_0 e_t, \qquad e_t \sim \mathcal{WN}(0, \mathbb{I}_n), \tag{C.1}$$

where $B(L) \equiv \mathbb{I}_n - \sum_{j=1}^{p} B_j L^j$, $e_t$ are the structural shocks, and $B_0$ contains the contemporaneous transmission coefficients. Recall that under full invertibility

$$\Sigma = \mathbb{E}[u_t u_t'] = B_0 Q [e_t e_t'] Q' B_0' \tag{C.2}$$

for any orthogonal matrix Q. $u_t$ are the reduced-form VAR innovations. The external instrument of Section 3 allows identification of only one column $\mathrm{b}_0$ of $B_0$, which contains the impact effects of the identified technology news shock $e_{\text{A2},t}$ on $y_t$.

The spectral density of $y_t$ is

$$S_y(e^{-i\omega}) = [B(e^{-i\omega})]^{-1}\Sigma\,[B(e^{-i\omega})^{\top}]^{-1}, \tag{C.3}$$

where $i \equiv \sqrt{-1}$, we use $\omega$ to denote the frequency, and $B(e^{-i\omega})^{\top}$ is the conjugate transpose of $B(e^{-i\omega})$. Let $S_y^{\text{A2}}(e^{-i\omega})$ denote the spectral density of $y_t$ when only the technology news shock $e_{\text{A2},t}$ is activated. This is equal to

$$S_y^{\text{A2}}(e^{-i\omega}) = [B(e^{-i\omega})]^{-1}\mathrm{b}_0\sigma_{\text{A2}}\mathrm{b}_0'\,[B(e^{-i\omega})^{\top}]^{-1}. \tag{C.4}$$

$\sigma_{\text{A2}}$ is the variance of $e_{\text{A2},t}$ for which an estimator is given by $\sigma_{\text{A2}} = \left(\mathrm{b}_0'\Sigma^{-1}\mathrm{b}_0\right)^{-1}$ (see Stock and Watson, 2018). Hence, the share of variance due to $e_{\text{A2},t}$ at frequency $\omega$ can be calculated as

$$\gamma_{\text{A2}}(\omega) = \frac{\mathrm{diag}\left(S_y^{\text{A2}}(e^{-i\omega})\right)}{\mathrm{diag}\left(S_y(e^{-i\omega})\right)}, \tag{C.5}$$

where the ratio between the two vectors is calculated as the element-by-element division.

The share of variance due to $e_{\text{A2},t}$ over a range of frequencies is calculated using the following formula for the variance

$$\frac{1}{2\pi}\int_{-\pi}^{\pi} S_y(e^{-i\omega})d\omega = \lim_{N\to\infty}\frac{1}{N}\sum_{k=-N/2+1}^{N/2} S_y(e^{-i\omega_k}), \tag{C.6}$$

where $\omega_k = 2\pi k/N,\ k = -N/2,\dots,N/2$.

Recall that the spectrum is symmetric around zero. Let the object of interest be the share of variance explained by $e_{\text{A2},t}$ at business cycle frequencies. These are typically between 2 and 8 years which, with quarterly data, correspond to a period between 8 and 32 quarters. Recall the mapping between frequency and period $\omega = 2\pi/t$. Business cycle frequencies are then in the range $\left[2\pi\underline{k}/N\ 2\pi\bar{k}/N\right]$, where $\underline{k} = N/32$ and $\bar{k} = N/8$. It follows that the share of fluctuations in $y_t$ that is accounted for by $e_{\text{A2},t}$ at business cycle frequencies is equal to

$$\frac{\sum_{k=\underline{k}}^{\bar{k}}\mathrm{diag}\left(S_y^{\text{A2}}(e^{-i\omega})\right)}{\sum_{k=\underline{k}}^{\bar{k}}\mathrm{diag}\left(S_y(e^{-i\omega})\right)}. \tag{C.7}$$

# D Robustness & Additional Charts: 5-Variable VAR

Figure D.1 compares the IRFs retrieved by our baseline patent-based instrument with the identification schemes of Beaudry and Portier (2006), denoted 'EQY/LR', and of Barsky and Sims (2011), denoted 'Max-FEV', in the same VAR. All responses are scaled such that the peak response of TFP is equal to 1% across all identification schemes. Beaudry and Portier (2006) identify technology news shocks as an innovation to the stock market index that is orthogonal to the current level of TFP. Barsky and Sims (2011) identify news shock as being orthogonal to current TFP, and maximizing the forecast error variance of TFP at all horizons between 0 and 40 quarters.

Figure D.2 compares the IRFs obtained in the benchmark case with IV constructed without controlling for contemporaneous policy shocks (i.e. setting $\delta = 0$ in Eq. (1).

Figure D.3 plots the share of variance that is due to $e_{\text{A2},t}$ for all the variables included in the 5-variable VAR at all frequencies between 1 (highest frequency) and 100 (lowest frequency) years. Grey areas highlight business cycle frequencies.

Figure D.4 reports for comparison the share of forecast error variance accounted for by the identified shocks in the time domain (i.e. across forecast horizons).

Figure D.1: Different Identifications in 5-variable VAR

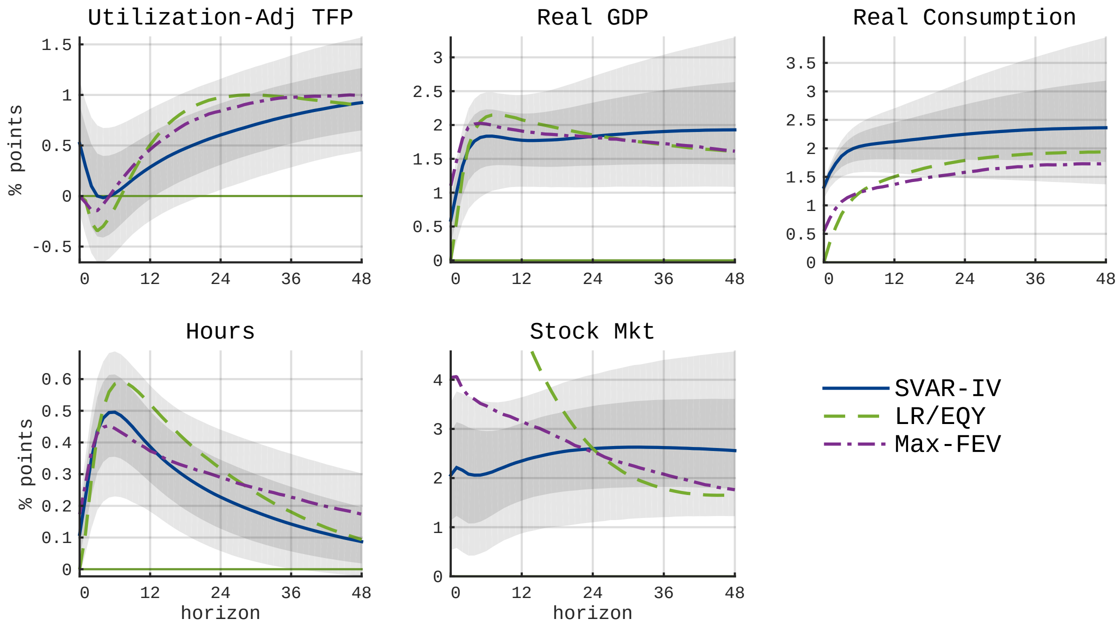


*Note*: Modal response to a technology news shock identified with (1) patent-based IV (SVAR-IV, blue), (2) long-run restrictions (LR/EQY, green dashed), and (3) maximum forecast error variance share (Max-FEV, purple dash-dotted). Estimation sample 1960-I:2019-IV. Identification sample 1982-I:2006-IV. Shaded areas denote 68% and 90% posterior credible sets for the SVAR-IV.

Figure D.2: Sensitivity to Other Policy Shocks & Spikes in IV

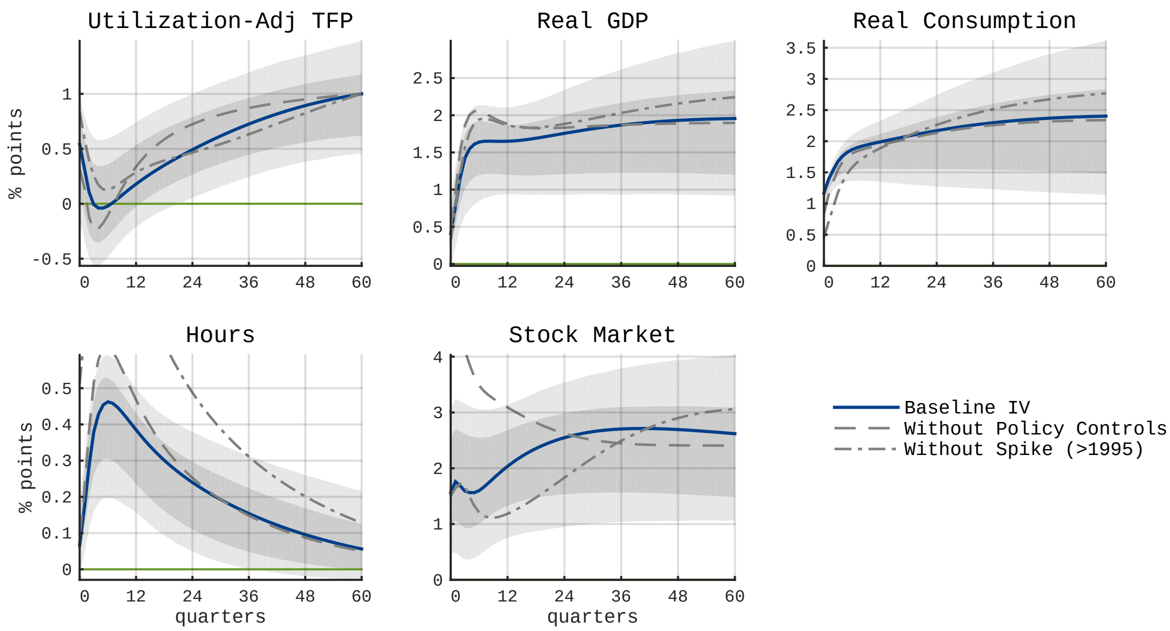


*Note*: Modal responses. Estimation sample 1960-I:2019-IV. Identification samples are: 1982-I:2006-IV with the baseline IV (solid lines); 1982-I:2014-IV for the IV that does not control for policy shocks (dashed lines); 1996-III:2006-IV for the IV that excluded the regulation spikes (dash-dotted lines). Shaded areas denote 68% and 90% posterior credible sets for the baseline IV.

Figure D.3: Error Variance Decomposition: Frequency, Small VAR

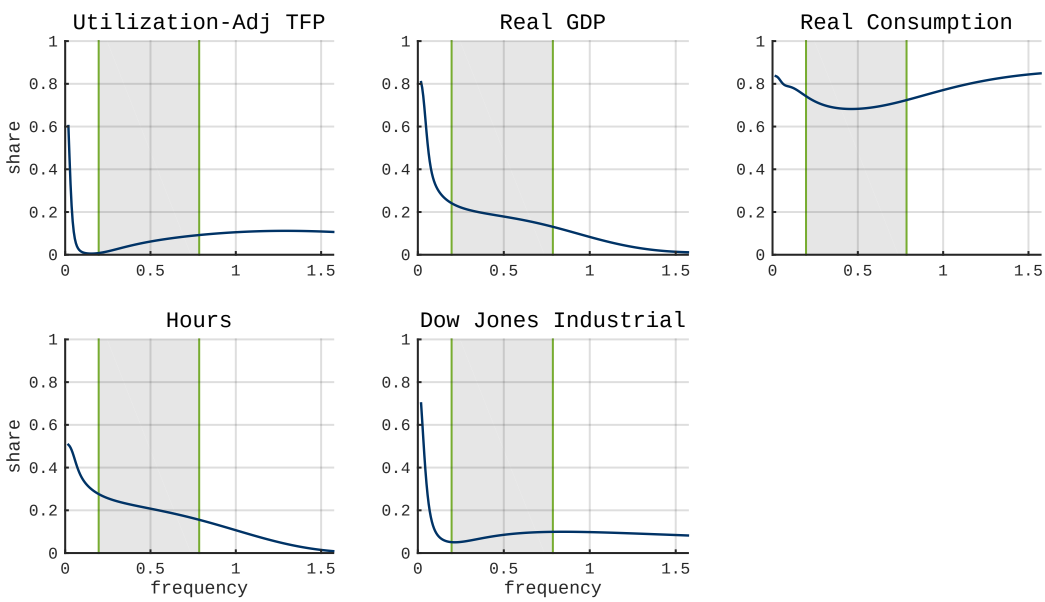


*Note*: Share of error variance accounted for by technology news shock identified with patent-based external instrument. VAR(4) with standard macroeconomic priors. Estimation sample 1960-I:2019-IV; Identification sample 1982-I:2006-IV. Shaded areas delimits business cycle frequencies (between 8 and 32 quarters).

Figure D.4: Forecast Error Variance Decomposition: Time, Small VAR

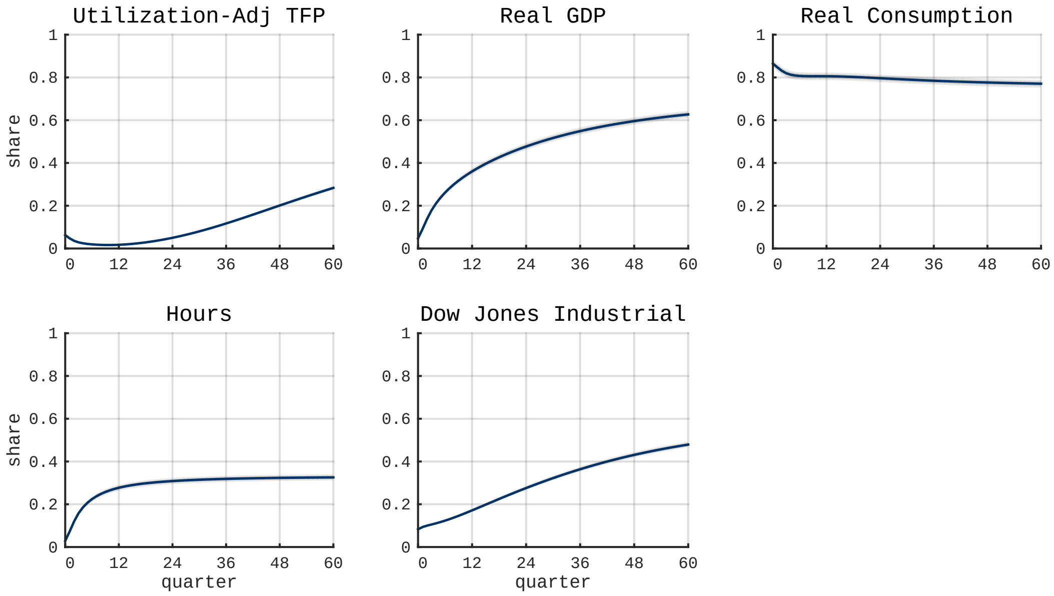


*Note*: Share of forecast error variance accounted for by technology news shock identified with patent-based external instrument. VAR(4). Estimation sample 1960-I:2019-IV; Identification sample 1982-I:2006-IV.

# E Alternative Patent Data Source

Kogan, Papanikolaou, Seru and Stoffman (2017), KPSS henceforth, assemble a dataset of patents granted by the USPTO to large US firms from 1926 to 2010. For each granted patent for which a company match exists in the CRPS database, KPSS collect information on the patent number, application, grant and publication dates, CRPS identifier of the patent owner, technology class and subclass, number of forward citations, and estimated economic value that the patent generates in the stock market in nominal US dollars. The latter being computed using the company's returns in a three-day window that brackets the grant date.[34]

Relative to the USPTO dataset, the KPSS set covers a smaller cross-section. However, the availability of citation counts and economic value of each patent allows us to address the extent to which our IV can be ameliorated by weighting the patents.

In order to retain consistency with the USPTO data and with our main intuition, we align the patents in the KPSS set according to their application date. The resulting patent application series is plotted in Figure E.1 against our original source. In the figure, the solid line is the same as in the right panel of Figure 1, and corresponds to the total number of applications filed at the USPTO. The dashed line is obtained by ordering the granted patents in the KPSS set according to their application date. The time lag between the application and the grant date makes the application series constructed using the KPSS data mechanically drop to near zero in the latest part of the sample (i.e. applications filed towards the end of the sample are granted much later, beyond the 2010 cut-off date in the KPSS dataset). This phenomenon – known as truncation bias – is immediately apparent in the figure. As extensively discussed in Lerner and Seru (2021), this type of bias is present more dramatically in recent years, and is not uniformly distributed across technology classes, industries, and regions. In order to partially account for it, in what follows we only use data in the KPSS set up to the end of 2002, which coincides with the time when the trends in applications in the USPTO and KPSS datasets start to visibly and artificially diverge.

It is also worth noting that because our original data source includes information on

[34] For a detailed description of the construction of the KPSS dataset see https://mitsloan.mit.edu/shared/ods/documents?PublicationDocumentID=5894.

Figure E.1: Patent Applications Data: USPTO vs KPSS

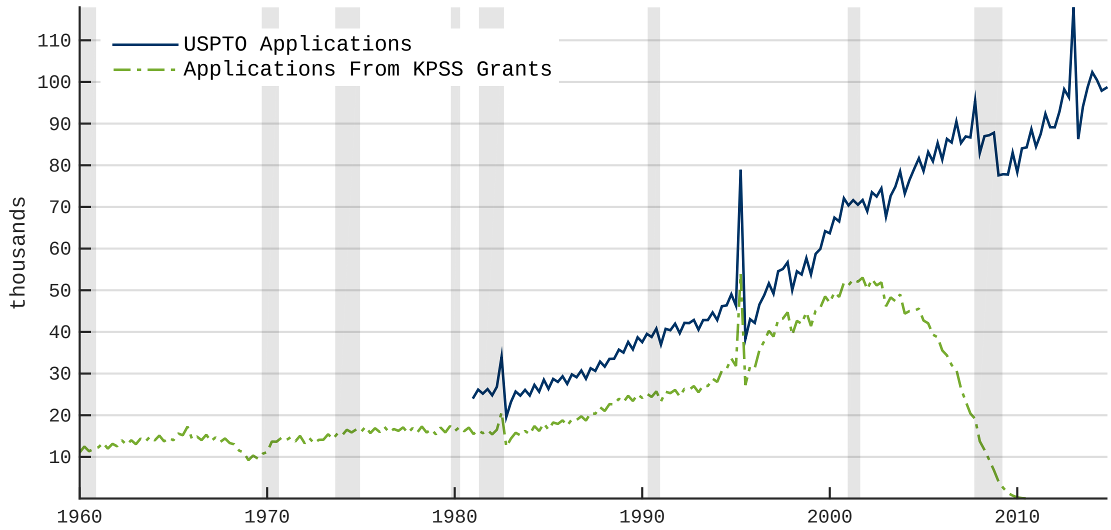


*Note*: Patent applications. Solid line, all patent applications filed at USPTO, source Marco et al. (2015). Dashed line, patent applications from granted patents in Kogan et al. (2017). Thousands.

the universe of patent applications, including those that are ex post not granted, it is naturally higher than the KPSS one. However, it is reassuring to verify that over the overlapping years, the two series share many similarities, including the large TRIPS spike. This is confirmed in Table E.1, which reports the coefficients of the instrument regression – Eq. (1) in the paper – using the two alternative sources. While in the KPSS case the estimates are slightly less precise due to the smaller number of data-points used, the picture that emerges is by and large equivalent. The regressions start in 1981 when the full set of SPF become available, but we end the sample at the end of 2002 for the KPSS data to partially account for the truncation bias.

Figure E.2 compares the impulse responses with the baseline IV based on USPTO data, with those obtained when using the KPSS source instead. Results are robust to the use of this alternative data source. As discussed in Section 3, some qualitative differences emerge in the response of output and hours, potentially due to the signal in the KPSS series being somewhat stronger since it is only based on applications of patents that are ex-post granted, or due to the fact that the KPSS set only includes large US firms. The use of the KPSS data can be thought of as one way of weighting the patents applications such that those that are ex-post not granted are assigned a zero weight, while all ex-

Table E.1: Instrument Construction, Alternative Data Sources

| | USPTO | KPSS |
|---|---|---|
| *Own Lags* | | |
| $pa_{t-1}$ | −0.952 | -0.893 |
| | (0.08) | (0.08) |
| $pa_{t-2}$ | −0.548 | -0.399 |
| | (0.11) | (0.13) |
| $pa_{t-3}$ | −0.272 | -0.128 |
| | (0.11) | (0.14) |
| $pa_{t-4}$ | −0.033 | 0.04 |
| | (0.09) | (0.13) |
| *Pre-Existing Beliefs* | | |
| $\mathbb{E}_t[u_{t+1}]$ | 0.629 | -0.093 |
| | (4.82) | (6.21) |
| $\mathbb{E}_t[\pi_{t+1}]$ | 3.424 | 3.029 |
| | (1.77) | (2.06) |
| $\mathbb{E}_t[I_{t+1}]$ | 0.065 | -0.01 |
| | (0.28) | (0.32) |
| $\mathbb{E}_t[\Pi_{t+1}]$ | −0.221 | -0.181 |
| | (0.34) | (0.45) |
| $\mathbb{E}_t[u_{t+4}]$ | −1.513 | -0.243 |
| | (5.57) | (7.15) |
| $\mathbb{E}_t[\pi_{t+4}]$ | −2.979 | -3.947 |
| | (1.57) | (2.16) |
| $\mathbb{E}_t[I_{t+4}]$ | −0.101 | -0.094 |
| | (0.40) | (0.47) |
| $\mathbb{E}_t[\Pi_{t+4}]$ | −0.224 | -0.438 |
| | (0.27) | (0.45) |
| *Policy Shocks* | | |
| $mpol_t$ | −4.377 | -3.118 |
| | (1.84) | (1.93) |
| $mpol_{t-1}$ | 6.319 | 8.676 |
| | (4.47) | (5.67) |
| $mpol_{t-2}$ | 3.560 | 5.179 |
| | (2.08) | (2.84) |
| $utax_t$ | −1.979 | -6.42 |
| | (1.14) | (4.15) |
| $utax_{t-1}$ | −0.875 | -4.492 |
| | (1.60) | (3.60) |
| $utax_{t-2}$ | −2.976 | -3.643 |
| | (1.47) | (2.25) |
| $atax_t$ | 2.443 | 5.395 |
| | (2.86) | (4.64) |
| $atax_{t-1}$ | −3.332 | -4.256 |
| | (2.02) | (2.16) |
| $atax_{t-2}$ | −5.261 | -7.983 |
| | (3.99) | (5.97) |
| intercept | 10.949 | 12.644 |
| | (6.33) | (8.19) |
| F-stat | 13.59 | 20.845 |
| | [0.000] | [0.000] |
| Adj-$R^2$ | 0.493 | 0.467 |
| N | 99 | 83 |
| *Wald Tests for Joint Significance of Controls* | | |
| SPF & Policy Shocks | 2.505 | 1.744 |
| | [0.003] | [0.059] |

*Notes:* Regression results based on Eq. (1). Dependent variable: $pa_t = 100 \times (\ln PA_t - \ln PA_{t-1})$. Robust standard errors in parentheses. SPF Forecasts are for the unemployment rate ($u_t$), inflation (GDP deflator, $\pi_t$), real non-residential investments ($I_t$), and real corporate profits net of taxes ($\Pi_t$). Policy controls include narrative monetary policy ($mpol_t$), narrative unanticipated ($utax_t$) and anticipated ($atax_t$) tax changes. The bottom panel reports Wald test statistics for the joint significance of the controls with associated p-values below in square brackets. USPTO sample: 1981-2006, KPSS sample: 1981:2002.

Figure E.2: Technology News Shocks: USPTO vs KPSS application data

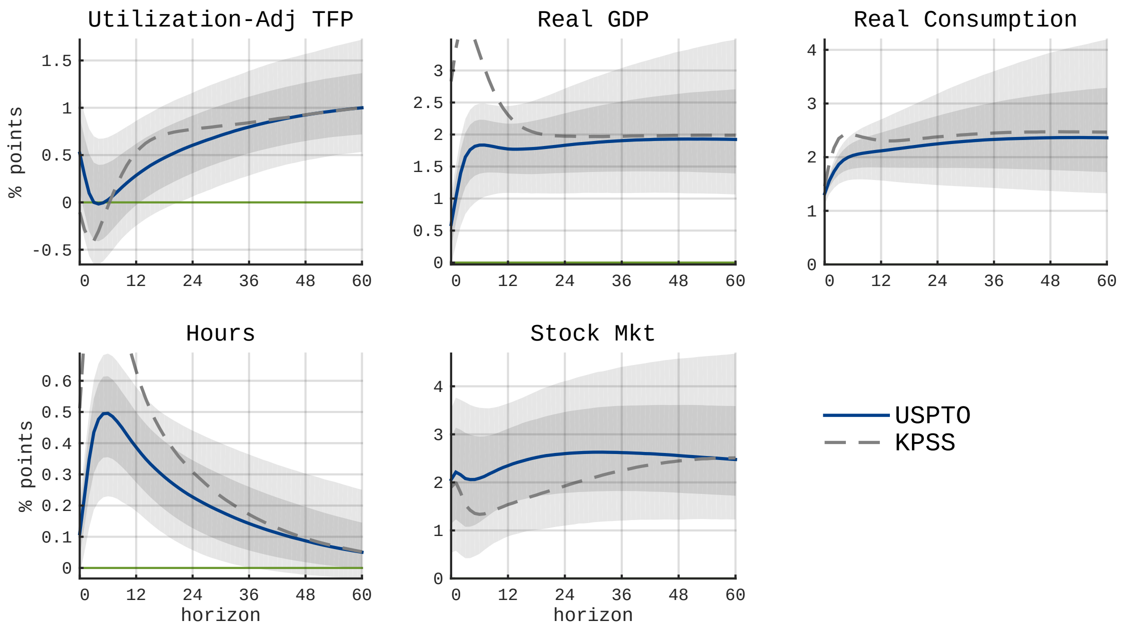


*Note*: Modal responses. Estimation sample 1960-I:2019-IV. Identification samples are: 1982-I:2006-IV with the baseline IV (solid lines); 1982-I:2002-IV for KPSS-based IV (dashed lines). Shaded areas denote 68% and 90% posterior credible sets for the baseline IV.

post granted patents are assigned equal weights. It is unclear whether this is a desirable approach in our context, since also patents that are ultimately not granted may contain an element of news that this weighting scheme disregards by construction. However, provided that our main results are robust to the change in source, the longer history in the KPSS set allows to potentially extend the IV backwards, provided that suitable proxies for pre-existing beliefs can be collected for these earlier years.[35]

Restricting the attention to the ex-post granted patents only, the KPSS dataset allows us to also explore alternative weighting schemes based on either forward citation counts, or the estimated economic value generated by the patent. Figure E.3 plots the raw number of patent applications in the KPSS data (solid line in both subplots) against the alternatives weighted either by citation (dashed line, left panel), or economic value (dashed line, right panel).

Forward citation counts record the number of citations that a patent receives in the

[35]The SPF started recording forecasts for corporate profits only from 1981. Unsurprisingly, this variable turns out to be particularly important when used as a control in the construction of the instrument.

Figure E.3: KPSS Data: Weighting Alternatives

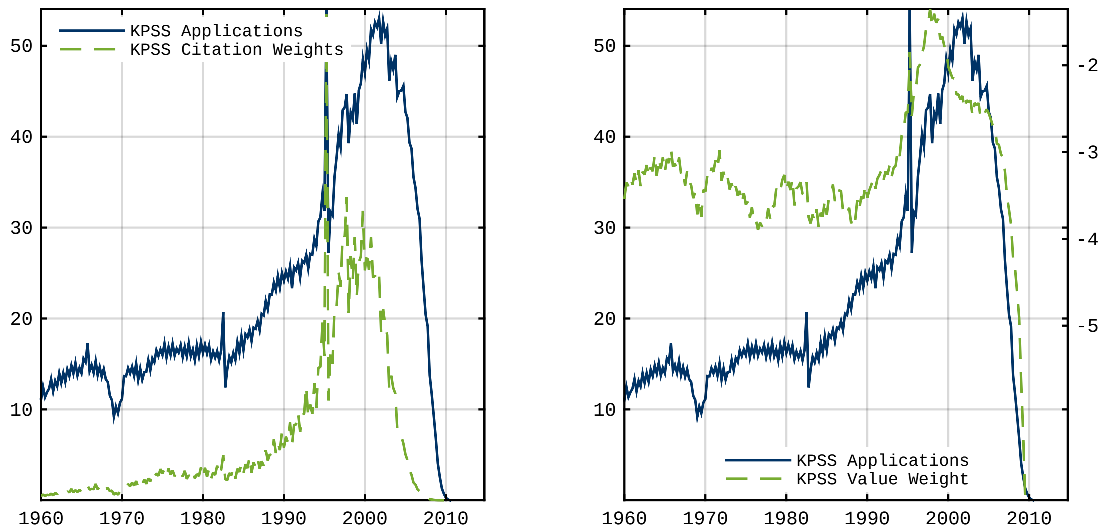


*Note*: In both figures the solid line is the number of applications in the KPSS dataset in each quarter (thousands). Dashed lines are for patent applications weighted by their forward citation count (thousands, left panel) and their economic value as measured by the firm's stock market reaction on the issue date (USD, right panel). The dollar value of innovation is deflated to 1982 million dollars using the CPI as in line with KPSS.

future. As noted in Lerner and Seru (2021), citation weights aggravate the truncation bias. Intuitively, patents are unlikely to be cited before being issued, and the number of citations is also not likely to pick up immediately after the issue date. This is clearly visible in the left panel of Figure E.3, where the citation-weighted applications artificially peak towards the end of the nineties. This additional truncation bias is also not uniformly distributed across technology classes. A further complication with citation-based weights is that the number of citations a patent receives can only increase over time. In turn, this implies that more recent patents are mechanically less cited, and thus assigned a smaller weight regardless of their intrinsic innovation content. Taking from Lerner and Seru (2021), *"the time lag between the filing of a patent application and its subsequent grant results in a mechanical tail-off in patent grants toward the end of the sample. Moreover, it may be a decade or longer after a patent is filed before one can get a good sense of how influential it is from citations. While it is possible to adjust the number of patent grants and number of patent citations received in early years based on historical patterns – and thus project the total number of patents or number of citations likely to be ultimately*

Figure E.4: Baseline vs Value Weights

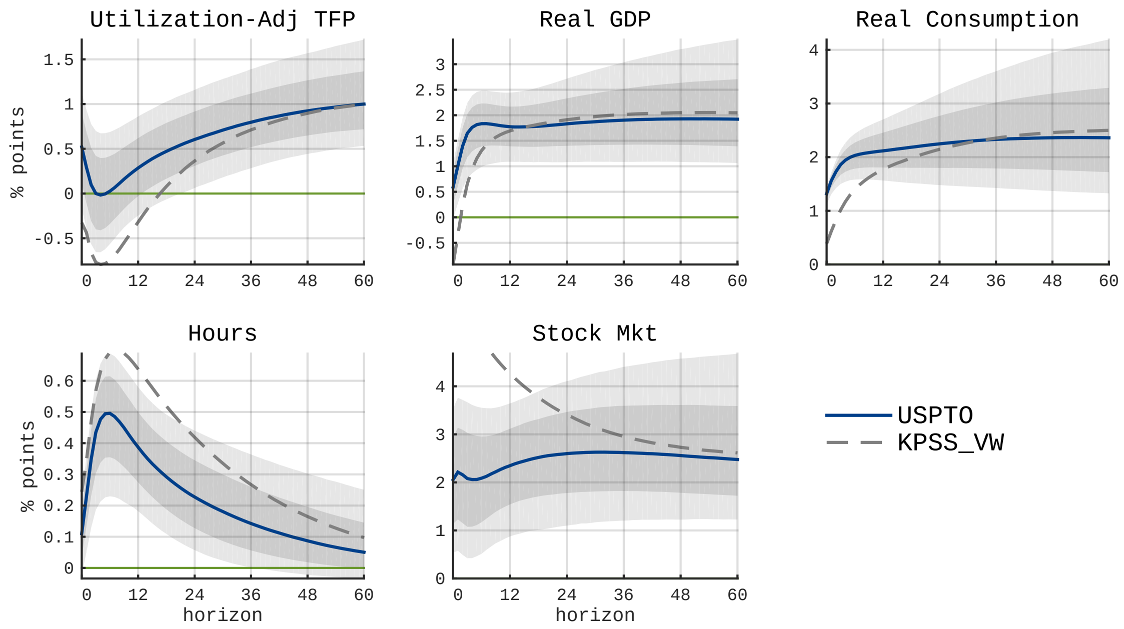


*Note*: Modal responses. Estimation sample 1960-I:2019-IV. Identification samples are: 1982-I:2006-IV with the baseline IV (solid lines); 1982-I:2002-IV for KPSS-based value-weighted IV (dashed lines). Shaded areas denote 68% and 90% posterior credible sets for the baseline IV.

*received – these estimates can be quite imprecise and potentially biased."* Based on these considerations, we do not explore the construction of the IV based on citation-weighted applications.

KPSS introduce an alternative weight that is based on the estimated economic value that a patent generates in the stock market. This is calculated based on the return that the patent owner's stock enjoys when the patent is granted. The three-day event window over which the return is calculated goes from the day before to the day after the grant date, and controls are used for competing events that fall within the measurement window (see Kogan et al., 2017, for details). Similar to the forward citations, this measure of economic value is obviously not known at the time the application is filed, which can create issues when using this weighting pattern to capture technological news at the application stage. However, to the extent that the value is computed over a fixed three-day window, and is hence not changing over time, this weighting scheme resolves some of the issues that are instead intrinsic to the citation-based weights. Truncation however remains a concern. To partially account for it, as in the case above we discard observations from 2002 onward when constructing the IV using the value-weighted patents. Figure E.4 plots the responses against our baseline. Both sets of IRFs are normalized to yield a peak response of TFP of 1pp.

The IRFs are broadly similar in the medium run, but some important differences emerge. The value-weighted IV recovers a shock that leads to a muted response of TFP on impact (also at 68% level) but to a subsequent significant decline of TFP in the first two years, after which TFP slowly rises. The initial fall in TFP is likely to account at least in part for the short-lived but significant impact fall in output, and the more muted initial response of consumption. It is also worth noting that value-weighting the patent applications data changes the time-series properties of the series quite dramatically (see Figure E.3); it is therefore not entirely surprising that the IRFs in this case are somewhat different.

In all, due to agents – including financial markets – not knowing at the application stage which patents will ex-post be granted, nor the expected realized return around the grant date, we are skeptical around the use of such weighting scheme for the purpose of constructing an instrument for technology news shocks, since it rests on information that was not available to economic agents at the time in which the news materialized.

# F Robustness & Additional Charts: Large VAR

Figure F.1 plots the share of variance that is due to $e_{A2,t}$ for all the variables included in the large VAR at all frequencies between 1 (highest frequency) and 100 (lowest frequency) years. Grey areas highlight business cycle frequencies. Table 2 in Section 4 reports the share of variance due to $e_{A2,t}$ over three different ranges of frequencies for the same variables. Figure F.2 reports for comparison the share of forecast error variance accounted for by the identified shocks in the two VARs.

All the IRFs reported in Figures F.3 to F.8 are scaled such that the peak response of utilization-adjusted TFP equals 1%.

Figure F.3 reports IRFs over 60 quarters for the baseline VAR with 16 variables, and a VAR that excludes prices and wages. Estimation and identification sample as in baseline.

Figure F.4 reports IRFs estimated over a sample that excludes the 2008 financial crisis (estimation sample 1960-I : 2007-IV).

Figure F.5 reports IRFs identified with an IV that does not control for contemporaneous policy shocks.

Figure F.6 reports IRFs for a VAR that includes households expectations about unemployment a year ahead and total hours worked are replaced by the unemployment rate and the labor participation rate. Estimation and identification samples as in baseline.

Figure F.7 reports IRFs for a VAR that includes GDP inflation instead of the GDP deflator in (log) levels. Estimation and identification samples as in baseline.

Figure F.8 reports impact responses for a selection of the variables in our VAR to a contemporaneous TFP innovation that raises TFP on impact by 1%, and obtained with a standard Cholesky factorization with TFP ordered first.

## Figure F.1: Error Variance Decomposition: Frequency

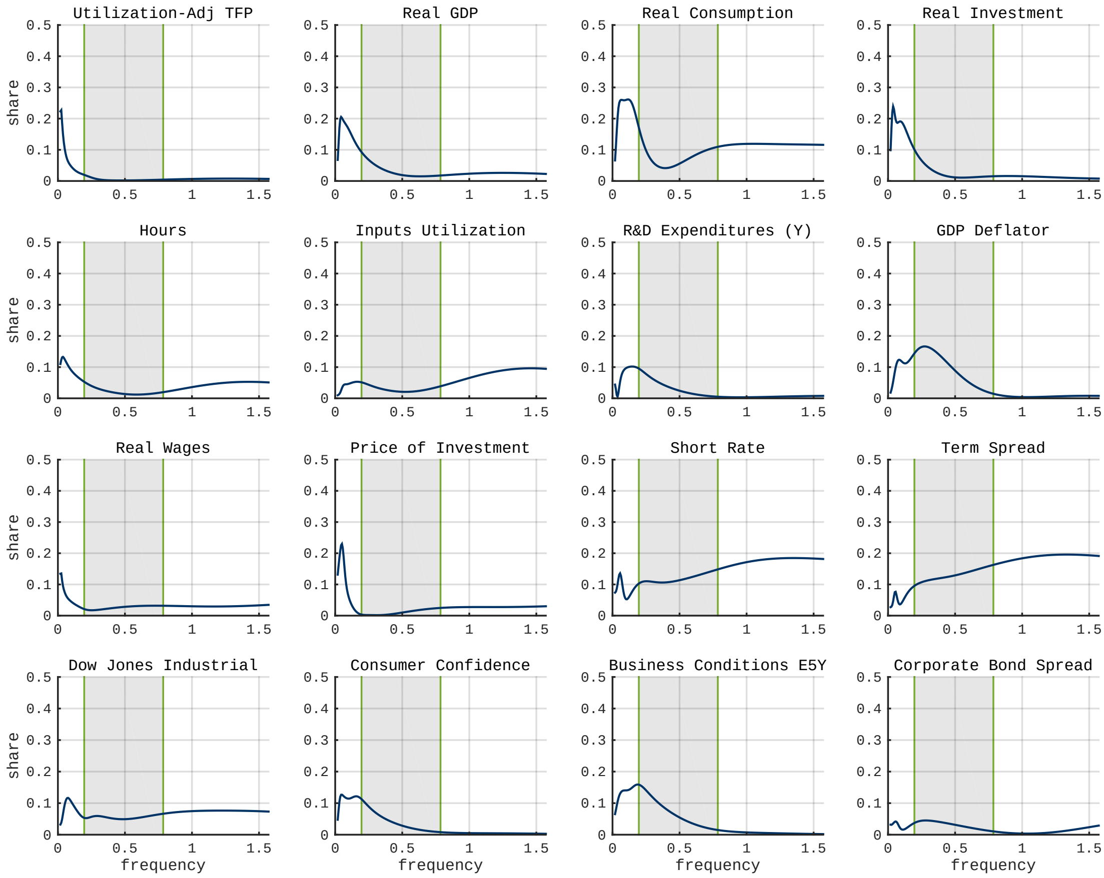


*Note*: Share of error variance accounted for by technology news shock identified with patent-based external instrument. VAR(4) with standard macroeconomic priors. Estimation sample 1960-I:2019-IV; Identification sample 1982-I:2006-IV. Shaded areas delimits business cycle frequencies (between 8 and 32 quarters).

Figure F.2: Forecast Error Variance Decomposition: Time

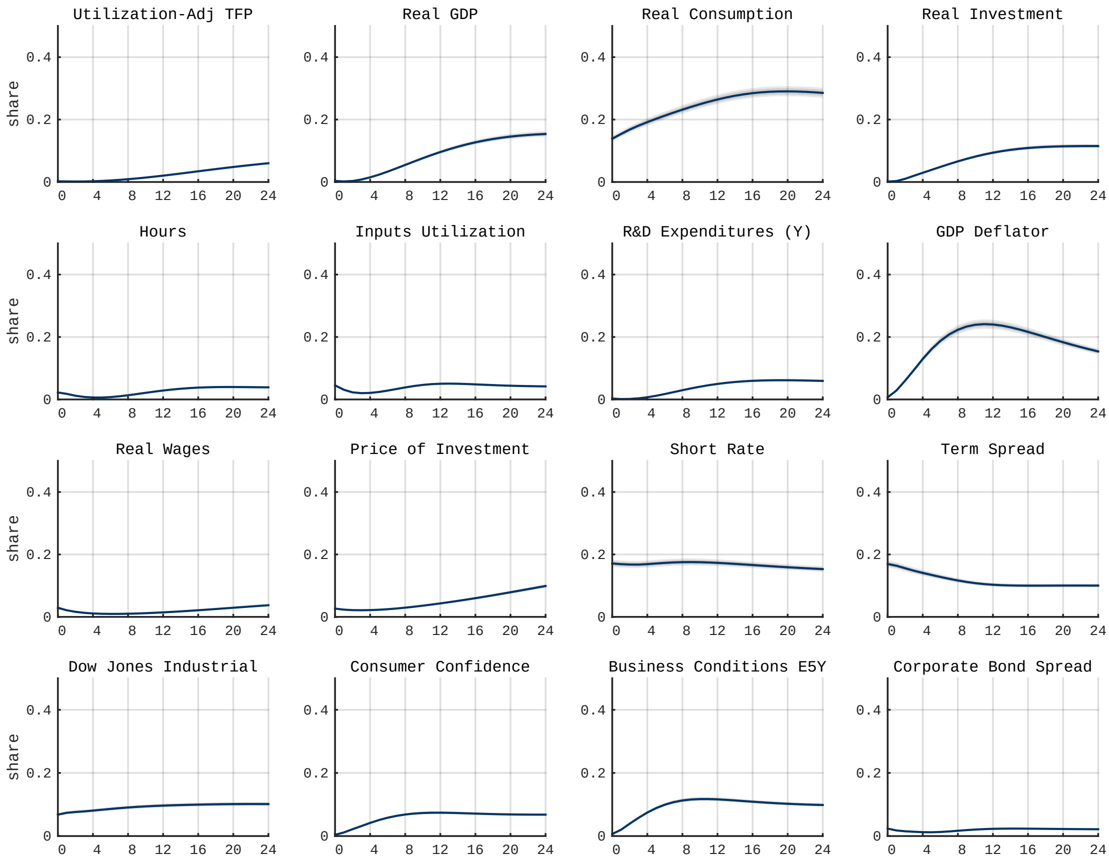


*Note*: Share of forecast error variance accounted for by technology news shock identified with patent-based external instrument. VAR(4). Estimation 1960-I:2019-IV; Identification 1982-I:2006-IV.

Figure F.3: IRFs Persistence

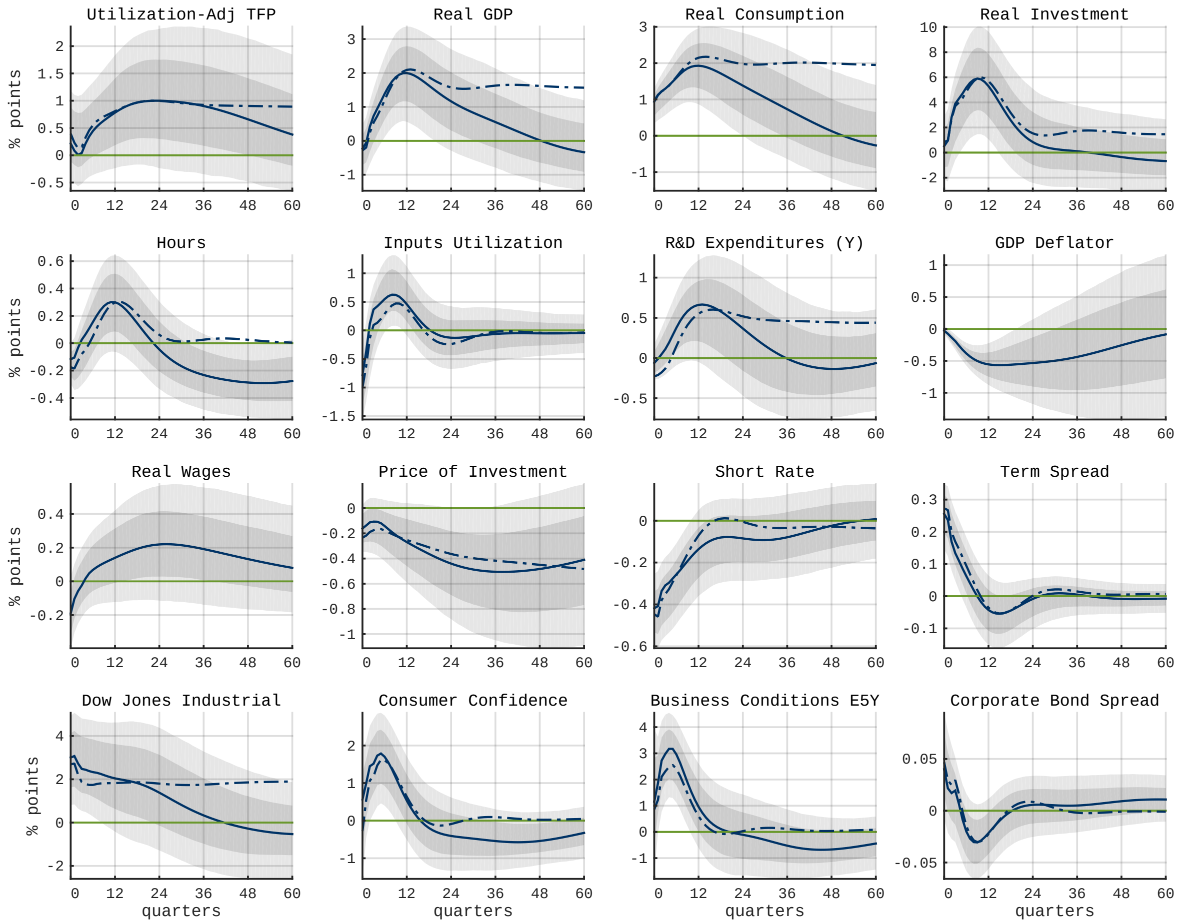


*Note*: Response of all variables to a technology news shock identified with patents-based external instrument. VAR(4) with standard macroeconomic priors. Estimation sample 1960-I:2019-IV; Identification sample 1982-I:2006-IV. Shaded areas denote 68% and 90% posterior credible sets. Solid lines: full system; Dash-dotted lines: VAR excludes prices and wages.

Figure F.4: IRFs Pre-Crisis Sample

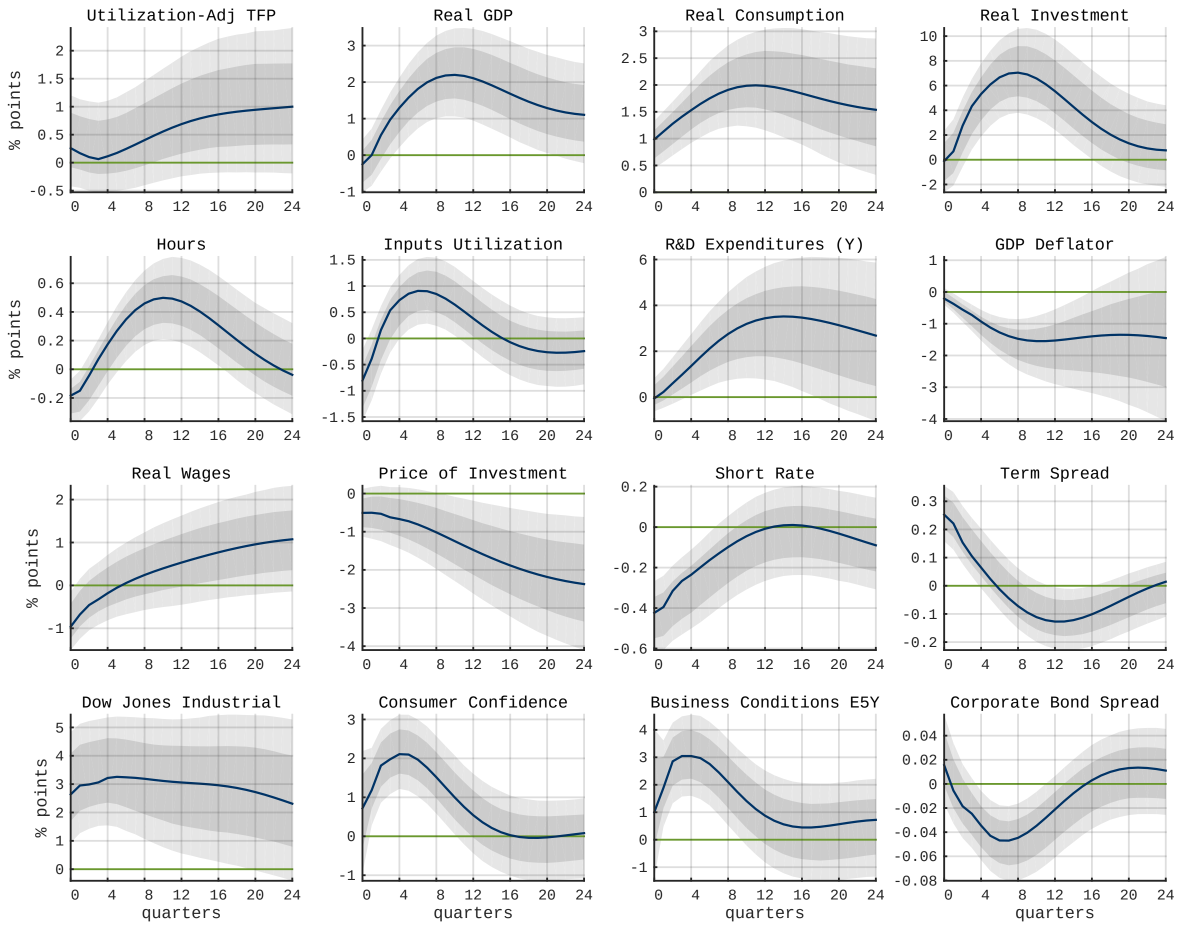


*Note*: Response of all variables to a technology news shock identified with patents-based external instrument. VAR(4) with standard macroeconomic priors. Estimation sample 1960-I:2007-IV; Identification sample 1982-I:2006-IV. Shaded areas denote 68% and 90% posterior credible sets.

Figure F.5: IRFs IV without Policy Controls

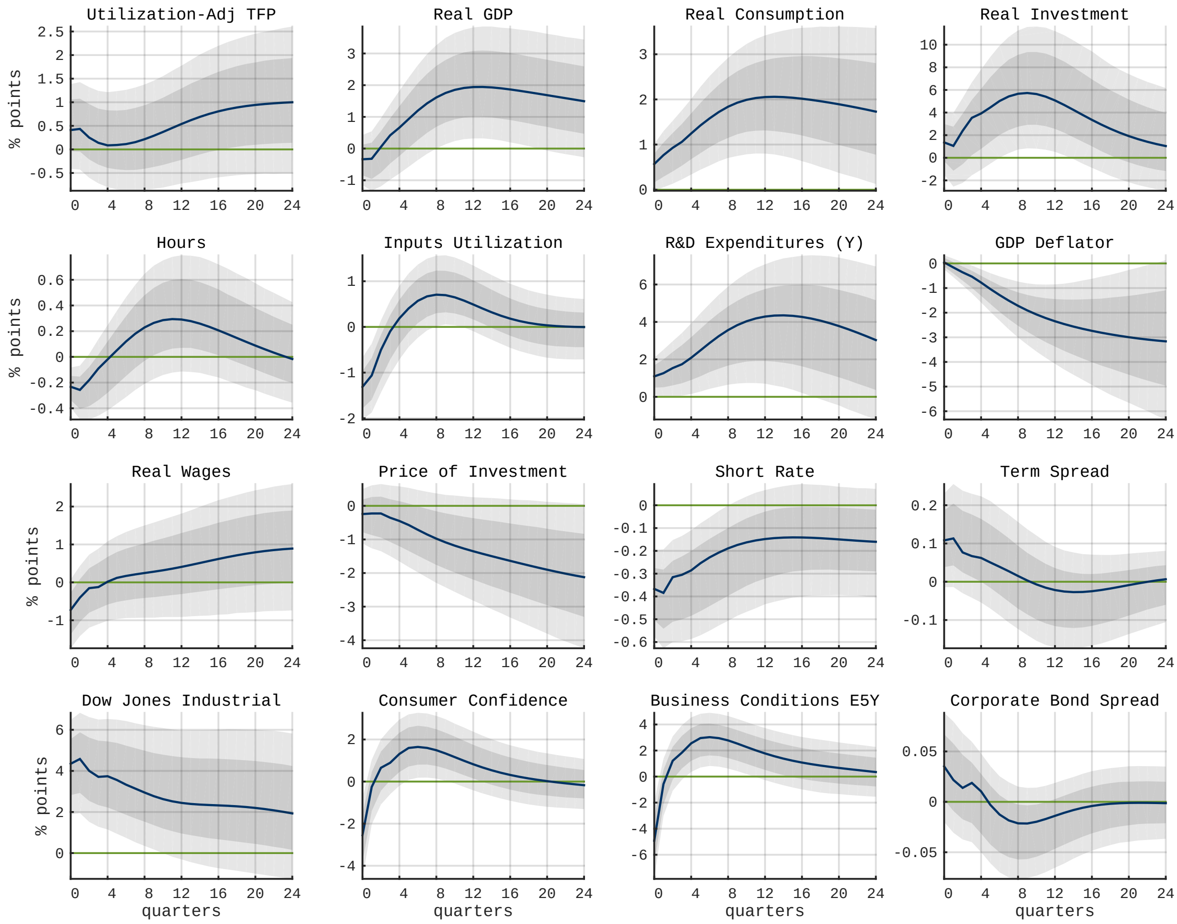


*Note*: Response of all variables to a technology news shock identified with patent-based external instrument that does not control for contemporaneous policy shocks. VAR(4) with standard macroeconomic priors. Estimation sample 1960-I:2019-IV; Identification sample 1982-I:2006-IV. Shaded areas denote 68% and 90% posterior credible sets.

Figure F.6: IRFs with Unemployment Expectations

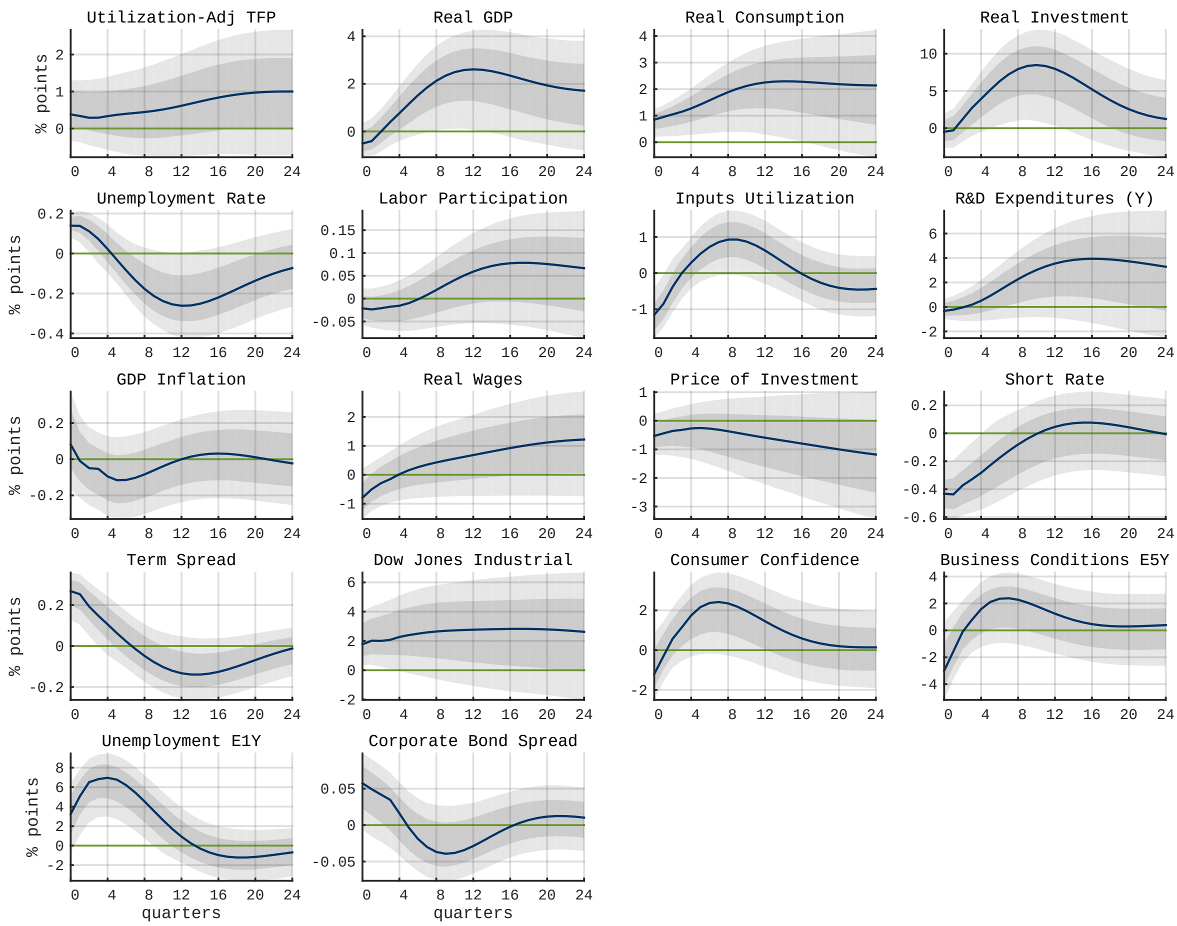


*Note*: Response of all variables to a technology news shock identified with patent-based external instrument. VAR(4) with standard macroeconomic priors. Estimation sample 1960-I:2019-IV; Identification sample 1982-I:2006-IV. Shaded areas denote 68% and 90% posterior credible sets.

Figure F.7: IRFs with GDP inflation

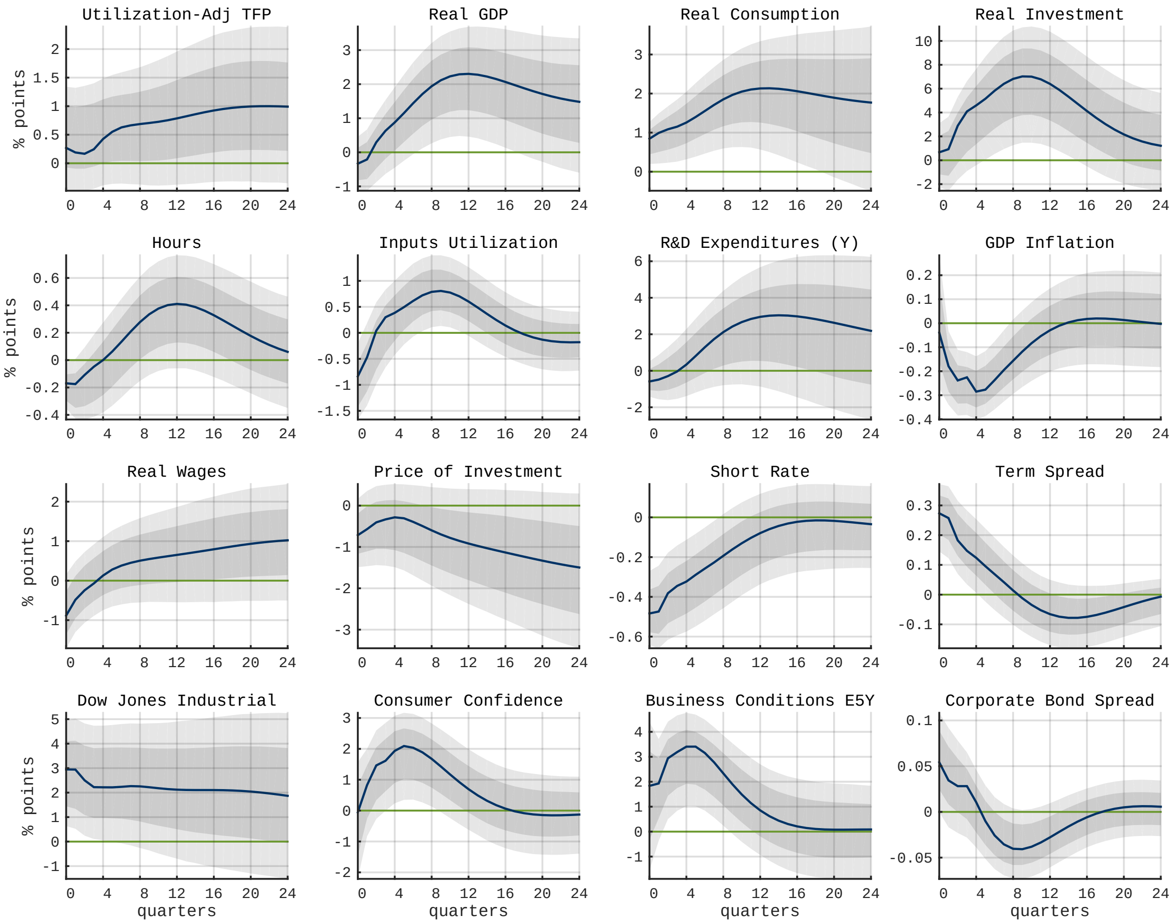


*Note*: Response of all variables to a technology news shock identified with patent-based external instrument that does not control for contemporaneous policy shocks. VAR(4) with standard macroeconomic priors. Estimation sample 1960-I:2019-IV; Identification sample 1982-I:2006-IV. Shaded areas denote 68% and 90% posterior credible sets.

## Figure F.8: Impact Responses to a Contemporaneous TFP Innovation

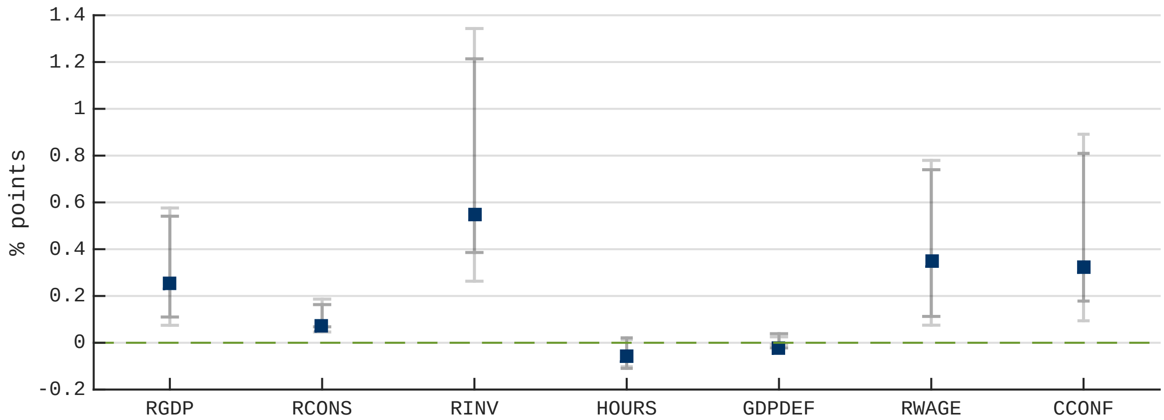


*Note*: Impact responses of selected variables to a TFP innovation that increases Utilization-Adjusted TFP by 1%. VAR(4). Estimation sample 1960-I:2019-IV. Grey bars delimit 68% and 90% posterior credible sets.

## Figure F.9: Long Rate Response

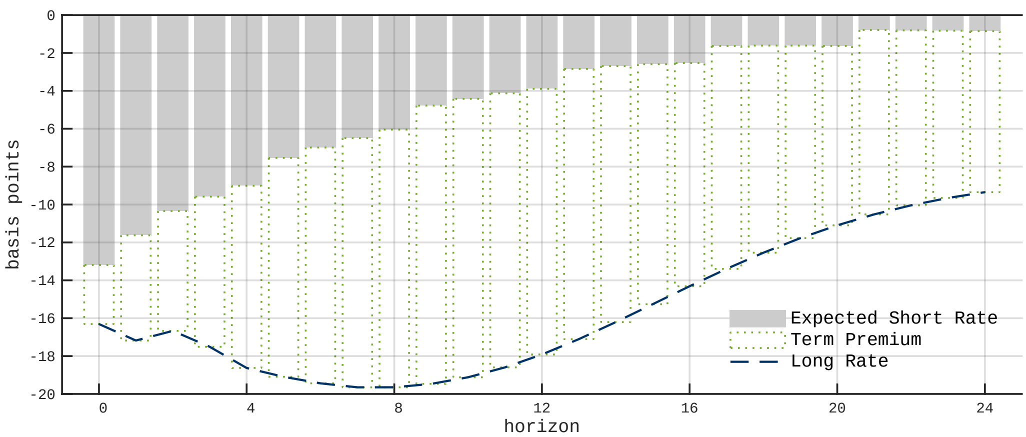


*Note*: Implied modal responses of the 10-year Treasury yield and VAR-based expectation and term premium components. VAR(4). Estimation sample 1960-I:2019-IV; Identification sample 1982-I:2006-IV.