

PATENTS, NEWS, AND BUSINESS CYCLES[☆]

Silvia Miranda-Agrippino*
Federal Reserve Bank of New York,
CEPR

Sinem Hacıoğlu-Hoke[†]
Federal Reserve Board,
CEPR

Kristina Bluwstein[‡]
Bank of England

First Version: August 22, 2018, Revised: August 19, 2025

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. The shock prompts a broad-based expansion in anticipation of the future increase in TFP, with output, consumption, and investment all rising well before any material increase in TFP is recorded. Despite the positive conditional comovements, the shock only accounts for a modest share of fluctuations of macroeconomic aggregates at business cycle frequencies. Financial markets price-in news shocks on impact, while most of the macro aggregates respond with some delay.

Keywords: Technology News Shocks; Business Cycle; SVAR-IV; Patent Applications.

JEL Classification: E32, O33, O34, C36

*silvia.miranda-agrippino@ny.frb.org Web: www.silviamirandaagrippino.com

[†]sinem.haciogluhoke@frb.gov Web: www.sinemhaciogluhoke.com

[‡]kristina.bluwstein@bankofengland.co.uk Web: www.kristinabluwstein.com

* Previously “When Creativity Strikes: News Shocks and Business Cycle Fluctuations.” We thank the editor Kurt Mitman and the anonymous referees for valuable suggestions and comments. We are grateful to Franck Portier and André Kurmann for insightful discussions and many useful suggestions on earlier versions of this paper, and to John Fernald for his invaluable guidance in the construction of a monthly TFP series. Special thanks go to Richard Crump, Danial Lashkari and Ricardo Reis. We thank Bob Barsky, Fabio Canova, Cristiano Cantore, Jordi Galí, Jean-Paul L’Huillier, Leonardo Melosi, Karel Mertens, Emanuel Mönch, Evi Pappa, Valerie Ramey, Morten Ravn, Barbara Rossi and Mark Watson for detailed comments. We also thank seminar and conference participants at the Advances in Applied Macro/Finance Conference, Bank of England, Catholic University of Milan, Bocconi, FRB New York, FRB Chicago, FRB San Francisco, FRB Dallas, ES European Winter Meeting, Northwestern, UCSD, LSE, UCL, UPF, Texas A&M, and Oxford 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, the Board of Governors of the Federal Reserve System, the Federal Reserve Bank of New York or any other person associated with the Federal Reserve System.

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](#)).

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 enables us to dispense from all the identifying assumptions traditionally used in the literature.¹

The intuition behind our identification is simple: by their nature, patent applications embed information about potential future technological change (see also e.g. [Griliches, 1990](#); [Lach, 1995](#); [Hall and Trajtenberg, 2004](#)). At the same time, patent applications are cyclical, and may themselves be the result of current economic booms and/or past news.

¹Traditional identifications are motivated by economic theory and typically combine zero restrictions on the impact response of total factor productivity (TFP) with assumptions about its long-run drivers. 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 that news shocks maximize the forecast error variance of productivity either 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](#)).

To account for this endogeneity, we introduce explicit controls for expectations about the macroeconomic outlook that were prevalent at the time of the application filings, and for other policy changes that could influence the decision of filing a patent either directly or through their effect on other macro aggregates. Specifically, we recover an IV for technology news shocks as the component of patent applications that is orthogonal to pre-existing beliefs as captured by the Survey of Professional Forecasters (SPF), to contemporaneous and lagged monetary and fiscal policy changes as summarized by narrative accounts, as well as to own lags.²

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 between the application and the grant issuance, and the eventual diffusion of the innovation within the economy, can be in the order of several years, depending on the type of patent and the characteristics of the industry sector. Therefore, patent applications at any given time contain information about technological changes that may occur at some point in the future. 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, the filing date can be thought of as the first measurable time at 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 macroeconomic outlook that prevailed at the time of 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 sole identifying assumption.

Our main data source for patent applications are the NBER USPTO Historical Patent

²To 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 that can readily be used by other researchers, accounting for the presence of measurement error, and easily dealing with different sample lengths.

Data Files of [Marco, Carley, Jackson and Myers \(2015\)](#), that provide a comprehensive record of all patent applications—granted and not granted—filed at the U.S. Patent and Trademark Office (USPTO) since 1981, and aggregated at monthly frequency. We also discuss the appropriateness of weighting patent applications according to their scientific or economic value for the construction of the IV. For this we use data assembled in [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 generated in the stock market.

Because of the minimal set of restrictions required for identification, our framework enables us to investigate whether news shocks generate the patterns that were assumed in earlier identification schemes. Importantly, it allows us to dispense from assumptions about the long-run drivers of technology, as well as on the impact effects, such that assumptions that were made in earlier studies become instead results in our setting. While it is not known *ex ante* whether technological innovation will effectively follow, the news we capture does eventually materialize on average, and results in a persistent and gradual increase in aggregate TFP. 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 \(2013\)](#). Because innovations can in principle be released to the public under a patent-pending status, our identification scheme does not impose orthogonality with respect to the current level of technology, which is a typical assumption in the news literature.³ While this 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. [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

³In 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](#)).

its long-run/permanent component.

The empirical literature has long debated the potential for technology shocks to drive business cycle fluctuations.⁴ In particular, two critical aspects have animated the debate. First, whether technology shocks could generate the type of comovements in macroeconomic variables—particularly consumption and hours—that were typical of business cycles. Second, whether they accounted for a meaningful share of variation of economic aggregates at the relevant frequencies. We revisit these questions in light of our novel identification in an otherwise unrestricted VAR, and document four main patterns. First, macro aggregates react well in advance of any material increase in TFP, suggesting an important role for anticipatory effects. Second, the conditional comovements implied by our identified VAR are positive, and therefore enable technology shocks as a potential originator of business cycles. Third, most macro aggregates tend to respond to the shock with some delay, which cautions against placing too much weight on impact responses alone. Fourth, while an important driver of long-run dynamics, the recovered shock only explains a modest fraction of the variation of main macroeconomic aggregates at business cycle frequencies. Here it is important to note that while our identifying assumption rests on patent applications bearing news about future technological change, not all technological change necessarily goes through the patenting process, which in turn may leave some drivers of technology—and of business cycle volatility—unaccounted for.

Our results show that 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 within the first three years, 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,

⁴The empirical literature on technology news shocks is vast, and we review it when presenting our results in Section 5. 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 \(2009, 2011\)](#); [Kurmann and Otrok \(2013\)](#); [Barsky, Basu and Lee \(2015\)](#); [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](#)).

rather than by future realized fundamentals. The expansion is not immediate. While consumption rises somewhat already upon realization of the shock, the impact response of output and hours tends to be not significant at conventional levels. Investment also increases robustly. And so do real wages in the medium term. The shock triggers a significant response of the monetary authority that eases policy in anticipation of the expected decline in inflation. Lower borrowing rates and compressed risk premia appear as likely amplifiers of the short-term effects of the shock. We find that the identified shock generally accounts for less than 10% of the variation of main macro aggregates at business cycle frequencies, but it is an important driver of their long-run variation, a finding that echoes the results in [Angeletos, Collard and Dellas \(2020\)](#).

Lastly, and in an important departure from earlier studies, we test our results also in a novel monthly setting. For this purpose we construct monthly time series for TFP and Utilization-Adjusted TFP for the US economy. To our knowledge, ours are the first such estimates, and we make these data publicly available.

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 at which the novelty is commercialized. Responses of aggregate variables are estimated in a set of bivariate VARs with the publication index ordered last.⁵ 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 spec-

⁵More 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.

ification, in the construction of the 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. Finally, these studies have all relied on annual data potentially overlooking important higher frequency variation which instead we exploit for the identification. In a recent contribution, [Cascaldi-Garcia and Vukotić \(2022\)](#) use the innovation index of [Kogan, Papanikolaou, Seru and Stoffman \(2017\)](#) to identify technology news shocks. This index measures the dollar value that patents generate in the stock market once they are granted. Because patent grants post-date patent applications by possibly several years, and tend to depend on the intensity of labor and administrative cycles at the USPTO (see [Christiansen, 2008](#)), the innovation index may not necessarily be a good indicator of news.

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. In Section 4 we discuss how to translate our framework to a monthly setup. Section 5 contains our main results; here we extend the analysis to an information-rich 12-variable VAR to explore the transmission mechanisms of technology news shocks more in detail. A discussion of our results is reported in Section 6, and Section 7 concludes. Additional material is reported in the Online Appendix.

2 A Patent-Based IV for Technology News Shocks

2.1 Information in Patent Data

The starting point of our analysis is the monthly flow of all new patent applications filed at the US 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.⁶ In what follows, we operate at quarterly frequency for consistency with the existing literature, and due to the constraints imposed by data availability, particularly TFP. We discuss how one could apply our identification in a monthly VAR in Section 4.

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 figure, patents classified under Computers and Communications have enjoyed a faster growth. Applications across all categories tend to slide after recessionary episodes, providing some preliminary evidence of their cyclical nature.

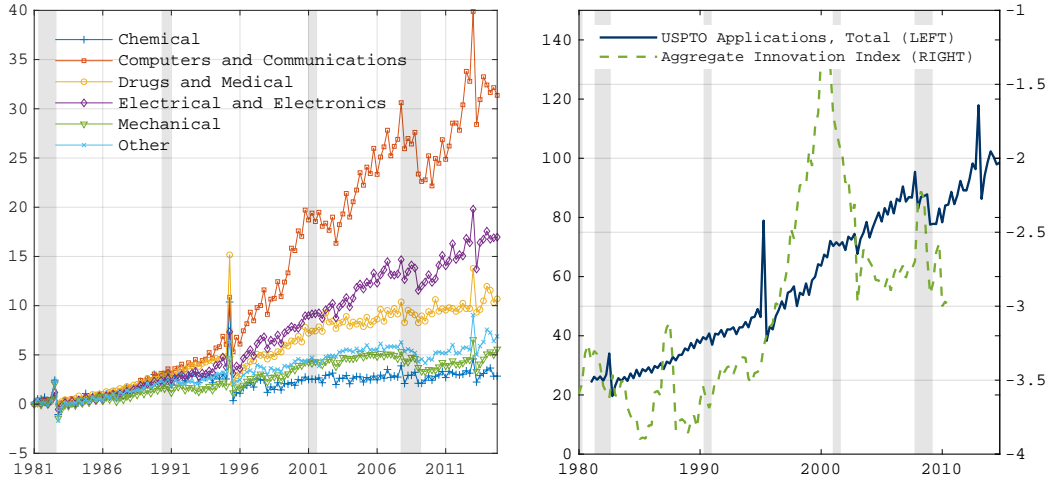
There have been three important regulatory changes in patenting in 1982, 1995, and 2013. All these regulations affected the number of applications when they came into effect, as shown by the spikes in the left panel of Figure 1. However, since they were not legislated in response to considerations related to either current or anticipated economic conditions, they provide us with important exogenous variation that we exploit for the identification. Said differently, to the extent that each patent embeds news about potential future technological progress, the increase in applications in anticipation of the upcoming regulatory changes represents an exogenous (relative to macroeconomic conditions) increase in technology news, which is the focus of our identification.⁷

In 1982, the old Court for Customs and Patent Appeals was abolished, and a new

⁶The dataset is available at <http://www.uspto.gov/economics>.

⁷We explore the sensitivity of our results to the regulation-induced spikes in the Appendix.

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.

Court of Appeals for the Federal Circuit was established. The new court provided more protection to patent 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 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.⁸ One of the main changes introduced by the TRIPS agreement was that of promoting transparency in patenting, and disincentivizing strategic behavior through stricter regulation.⁹ This had two main effects. First, it shifted forward the timing of some applications, which

⁸Article 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>.

⁹The 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 that 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/innovationpolicytrips_e.htm.

resulted in the one-off increase highlighted in Figure 1. 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.¹⁰

To provide a visual illustration of the link between patent applications and subsequent aggregate innovation, the right panel of Figure 1 compares the total number of USPTO applications (sum across industries in LHS chart, solid line) with the aggregate index of innovation of Kogan et al. (2017). The index is a forward-looking measure of the private, economic value of innovations in the U.S., and constructed as the GDP-weighted sum of the market value generated by patents granted within each quarter.¹¹ We note that, as expected, patent applications lead the aggregate innovation index. Moreover, the large spikes in the number of applications tend to correspond to substantial future increases in aggregate innovation, and particularly so after the TRIPS agreement. We take this as a preliminary indication that the exogenous legislation-induced increases in applications are informative about their innovation content, and thus contain important information for the purpose of identifying technology news shocks.

We construct the IV using all the patent applications submitted to the USPTO—including those that are ex-post not granted—and weighting 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 Christiansen (2008) have noted how most of the news content in patent applications may be exhausted by the time they are granted.¹² One reason 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

¹⁰The 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-1112-29.pdf.

¹¹The 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 index.

¹²From application filing to grant issuance the process takes about two years on average across industries. 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).

likely to capture the effective time at which the news materializes. Second, we choose to also include in our set 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, 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 return that the patent owner’s stock enjoys when the patent is granted. We discuss these options in detail in the Appendix. Here we note that, at the application stage, economic agents—including financial markets—do not know which patents will ex-post be granted, let alone their expected future citations or economic value. Therefore, we are skeptical about the use of these weighting schemes for the purpose of identifying technology news shocks, since they rest on information that was not available at the time at 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 beliefs about the state of the economy that are prevalent at the time of the application filings, to other contempo-

aneous 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 introduce three sets of controls. First, lagged patent applications to control for past shocks. Second, expectations about the macroeconomic outlook to control for other shocks, anticipated or otherwise, that are not captured in lagged patent applications. We align the timing of the survey forecasts such that the expectations reflect the most up-to-date predictions conditional on information available to the forecasters at the time of the patent filings. Finally, we include explicit controls for monetary and tax policy disturbances that may affect the decision of filing a patent either directly, or indirectly by affecting e.g. firms' investment plans.

Formally, we recover the IV as the residuals of the following regression, estimated at quarterly frequency

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

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 number of patent applications filed at the USPTO each quarter. $\gamma(L) = \sum_{j=1}^4 \gamma_j L^j$, 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).¹³ $\mathbb{E}_t[x_{t+h}]$ captures the most up-to-date predictions that are prevalent at the time of the applications. 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.¹⁴

¹³Survey of Professional Forecasters data (Federal Reserve Bank of Philadelphia, 2023) are available at <https://www.philadelphiafed.org/surveys-and-data/real-time-data-research>.

¹⁴SPF 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

The vector x_t includes the unemployment rate (u_t), inflation (π_t), and the growth rates of real non-residential fixed investments (I_t), and of real corporate profits net of taxes (Π_t).¹⁵

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. Indeed, by affecting macro aggregates, and especially investment, monetary and tax policy may have a direct effect on patent applications, and act as a confounding factor in the identification. The vector η_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 \(2004\)](#).¹⁶

The regression results are presented in Table 1. 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.¹⁷ Moreover, pre-existing

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>.

¹⁵SPF 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 5) at the relevant forecast horizons.

¹⁶We 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 \(2011\)](#).

¹⁷The negative sign of the autoregressive coefficients, also noted in [Adams et al. \(1997\)](#), suggests the 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 the dependency of the IV on its own past and ensures that the specific source of

TABLE 1: INSTRUMENT CONSTRUCTION

	(1)	(2)	(3)	(4)	(5)
<i>Own Lags</i>					
pa_{t-1}	-0.849*** (0.10)	-0.928*** (0.11)	-0.900*** (0.10)	-0.985*** (0.10)	-0.977*** (0.10)
pa_{t-2}	-0.480*** (0.10)	-0.605*** (0.11)	-0.574*** (0.11)	-0.560*** (0.12)	-0.581*** (0.12)
pa_{t-3}	-0.273*** (0.09)	-0.384*** (0.08)	-0.366*** (0.08)	-0.292*** (0.11)	-0.303*** (0.11)
pa_{t-4}	0.002 (0.09)	-0.061 (0.08)	-0.056 (0.08)	-0.044 (0.10)	-0.048 (0.10)
<i>Pre-Existing Beliefs</i>					
$E_t[u_{t+1}]$		-0.326 (0.37)			0.574 (5.15)
$E_t[\pi_{t+1}]$		1.657** (0.68)			2.602 (1.99)
$E_t[I_{t+1}]$		0.485** (0.23)			0.072 (0.27)
$E_t[\Pi_{t+1}]$		-0.199 (0.34)			-0.279 (0.55)
$E_t[u_{t+4}]$			-0.845* (0.46)		-0.843 (5.70)
$E_t[\pi_{t+4}]$			0.969 (0.78)		-2.586 (1.62)
$E_t[I_{t+4}]$			0.385 (0.26)		-0.018 (0.39)
$E_t[\Pi_{t+4}]$			-0.982*** (0.28)		-0.307 (0.35)
<i>Policy Shocks</i>					
$mpol_t$				-4.022* (2.27)	-4.321* (2.33)
$mpol_{t-1}$				8.128* (4.81)	7.613 (5.54)
$mpol_{t-2}$				6.784** (2.70)	5.504* (3.16)
$utax_t$				-1.386 (0.94)	-1.886 (1.16)
$utax_{t-1}$				-1.241 (1.49)	-1.817 (1.65)
$utax_{t-2}$				-2.024** (0.86)	-3.209** (1.41)
$atax_t$				2.255 (2.60)	1.044 (2.78)
$atax_{t-1}$				-2.607 (1.65)	-3.681* (2.10)
$atax_{t-2}$				-4.273 (3.61)	-4.809 (3.61)
intercept	4.343*** (0.80)	0.814 (2.83)	6.639 (4.89)	5.100*** (0.88)	7.554 (5.85)
F-stat	33.87 [0.000]	17.99 [0.000]	19.38 [0.000]	17.88 [0.000]	11.32 [0.000]
Adj- R^2	0.448	0.486	0.469	0.535	0.503
N	131	131	131	99	99
<i>Wald Tests for Joint Significance of Controls</i>					
Quarter Ahead SPF		4.832 [0.001]			
Year Ahead SPF			3.727 [0.007]		
Policy Shocks				2.454 [0.016]	
SPF & Policy Shocks					2.634 [0.002]

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, π_t), real non-residential investments (I_t), and real corporate profits net of taxes (Π_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.

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, and particularly the contemporaneous ones, are also informative. Both shocks are normalized such that an increase corresponds to a tightening of policy. The table shows that it is typically the case that restrictive policies are associated with a decline in patent applications, a further indication of their cyclical nature.

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 orthogonal to other contemporaneous policy shocks. The IV is not forecastable also using a wider set of predictors. Macro-financial factors extracted from large cross-sections and broader sets of survey forecasts not included in Eq. (1) that Granger-cause patent applications are uninformative for the IV.¹⁸

We argue that it is unlikely that structural disturbances other than current technology news may affect the U.S. 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

$$\ln A_t = \ln S_t + \ln \Gamma_t , \quad (2)$$

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

$$\ln S_t = \phi_s \ln S_{t-1} + e_{A1,t} , \quad (3)$$

seasonality does not affect the identification.

¹⁸See Tables A.1 and A.2 for Granger-causality results on patent applications, and Tables A.4 and A.4 for the same on the IV.

and Γ_t is the permanent component, characterized instead by the presence of a unit-root

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

In Eqs. (3) - (4) above $\Delta \ln A$ is the steady state growth rate of technology, the autoregressive coefficients ϕ_s and ϕ_Γ are in the interval $(0, 1)$, and $e_{A1,t}$ and $e_{A2,t-k}$ are zero-mean normally distributed i.i.d. processes with variance equal to σ_{A1}^2 and σ_{A2}^2 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.

$e_{A2,t}$ is the news shock. The standard identifying assumption in the news literature is that agents learn about $e_{A2,t-k}$ 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 ($e_{A1,t}$) 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 the 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, 2018](#)).

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

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

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

For the purpose of estimating the impulse response functions (IRF) and forecast error variance decompositions (FEVD) 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

$$u_t = B_0 e_t, \quad (6)$$

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 [Plagborg-Møller and Wolf \(2021\)](#) and [Miranda-Agrippino and Ricco \(2023\)](#). However, [Forni, Gambetti and Sala \(2019\)](#) 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, Walker and Yang, 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 toward the problem solution is to add information to the VAR, through variables that help reveal the state variables. This is the role of the stock price index

in [Beaudry and Portier \(2006\)](#), or of measures of consumers and business confidence 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, Gambetti and Sala \(2014\)](#).¹⁹

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

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

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

where z_t denotes the external instrument used for the identification of $e_{A2,t}$. Under these conditions, the impact responses to $e_{A2,t}$ 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, 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 is 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.

Two approaches are available when identification is attained using instrumental variables. One is to include the IV among the endogenous variables in the VAR, and retrieve the IRFs using a standard Cholesky triangularization with the IV ordered first. Alternatively, one can follow the two-step procedure of [Mertens and Ravn \(2013\)](#) which entails estimating a VAR on the observables, and then regressing the VAR residuals on the IV. We choose this second option. Under invertibility and exogeneity of the IV, the two approaches yield the same results in population (see [Plagborg-Møller and Wolf, 2021](#); [Miranda-Agrippino and Ricco, 2023](#)). In empirical samples, however, one critical advantage of the two-step procedure is that it allows us to operate on samples of different

¹⁹While 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. \(2019\)](#).

length, which is particularly useful in our case due to the short life span of our IV. In particular, it allows us to estimate the VAR dynamics on sufficiently long samples—which in turn allows us to entertain meaningful discussions around the effects of the shock in the long run—while at the same time using the full length of the IV to estimate the impact effects.²⁰

3.2 Inspecting the Mechanism in an Illustrative VAR

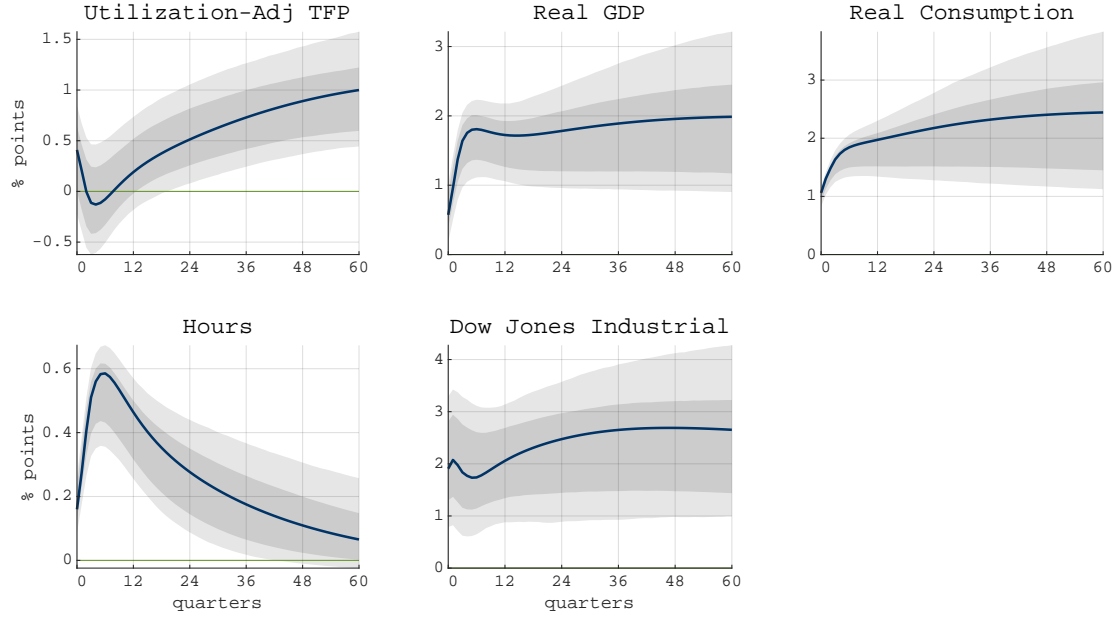
In this section, we put our instrument to test in an illustrative 5-variable VAR and discuss the sensitivity of our results with respect to a number of perturbations. The variables included in the VAR are the quarterly estimates of TFP corrected for input utilization of [Fernald \(2014\)](#), output, consumption, total hours worked, and the Dow Jones Industrial Average as the stock market index.²¹ These 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 using the GDP deflator and expressed in per-capita terms, where appropriate. We report a detailed description of the data and their construction 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. The 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 and estimate the optimal priors' tightness as in [Giannone, Lenza and Primiceri \(2015\)](#). 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%. Recall that the identification procedure leaves the full shape of the IRFs unrestricted, including the impact effects. Shaded areas correspond to 68% and

²⁰We report a more in-depth discussion on this point and results obtained using the alternative approach in Appendix E.

²¹The TFP series is obtained from [Fernald \(2021\)](#). Macroeconomic and financial series are obtained from FRED ([Federal Reserve Bank of St. Louis \(2023\)](#), <https://fred.stlouisfed.org>) and Refinitiv Datastream. A full list of variables used in this section and throughout the paper, and sources is included in the replication README.

FIGURE 2: TECHNOLOGY NEWS SHOCKS: 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. Horizon in quarters. First-stage F-stat=1.488.

90% posterior credible sets.²²

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 shock recovered by the IV has essentially no effect on TFP neither on impact, nor in the following three to five 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.²³

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

²²Because the instrument is a residual generated regressor, OLS-based inference is asymptotically correct (Pagan, 1984).

²³A 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.

TABLE 2: CONDITIONAL CORRELATIONS: 5-VARIABLE VAR

	$h = 1$	$h = 4$	$h = 12$	$h = 40$
Real GDP	0.989**	0.985**	0.995**	0.997**
Hours	0.988**	0.979**	0.990**	0.889**

Notes: Conditional correlations between consumption, output and hours implied by the identified VAR at selected horizons. Estimation sample 1960:I - 2019:IV. Identification sample 1982:I - 2006:IV. *, ** denote statistical significance at 68% and 90% levels respectively. Horizons in quarters.


on impact, and remains elevated throughout.²⁴ Broadly, the shock induces an immediate and strong economic expansion in anticipation of the rise of TFP. This is confirmed by the results in Table 2, where we report the implied conditional correlations of consumption with the main real activity aggregates at some selected horizons, calculated following Galí (1999).

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 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 downplaying or altogether removing the TRIPS observation (see Appendix). Removing the explicit controls for other policy shocks leads to responses for TFP, output and consumption that lie within the error bands of the 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. Similarly, the IRFs lie comfortably within

²⁴We use the residuals of the stock market for the first-stage, due to it being the variable most likely to respond to the shock on impact in a significant way. The first-stage F-statistic is equal to 1.488. For the larger VAR in Section 5 it is equal to 5.99. While these may appear somewhat weaker than desirable, it should be stressed that the first-stage regression involves stock market returns which are notoriously very volatile and difficult to predict. The fairly sizable first-stage F-stats in this case are therefore more encouraging than otherwise.

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).

the estimated error bands when we disregard the large observations corresponding to the implementation of the TRIPS agreement. Intuitively, this affects the precision of the estimates, but does not alter the broad picture.

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—mindful of the caveats highlighted in Section 2—also to alternative weighting schemes, as we discuss in detail in the Appendix. 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 exercises only including listed firms may also have a bearing on the response of aggregate output and hours.

To complete the discussion, Figure 3 reports the share of TFP variance that is accounted for by technology news shocks as identified by the IV.²⁵ Even if we have not imposed any such restriction ex ante, the shock recovered by the IV is most explanatory for TFP at long horizons, and at very low frequencies. This is consistent with the identified shock being a driver of the long-run component of aggregate productivity.²⁶

²⁵Variance decompositions for all variables as well as the algorithm are reported in the Appendix.

²⁶The variance shares tend to be exceptionally high for consumption and output, reaching up to 80%. This is likely due to the 5-variable VAR not being informationally sufficient (see Forni and Gambetti, 2014) which, as noted in Section 3.1, may introduce a bias in the forecast error variance decompositions

4 Technology News Shocks in a Monthly VAR

In this section, we discuss how one could apply our identification strategy in a monthly VAR. Virtually all the existing empirical literature on technology news shocks relies on estimated quarterly models. One main reason for this choice is that this is the highest frequency at which benchmark estimates of TFP and utilization-adjusted TFP are available (Fernald, 2014).

The higher sampling frequency of patent data, however, raises the question of whether we may be discarding relevant variation due to time aggregation. In what follows, we discuss how to construct the IV using monthly data, and test our results in a monthly VAR. For this purpose, we also construct monthly estimates of TFP and of utilization-adjusted TFP for the U.S. economy from 1966:01 to 2023:12. We report all details in the Appendix.²⁷

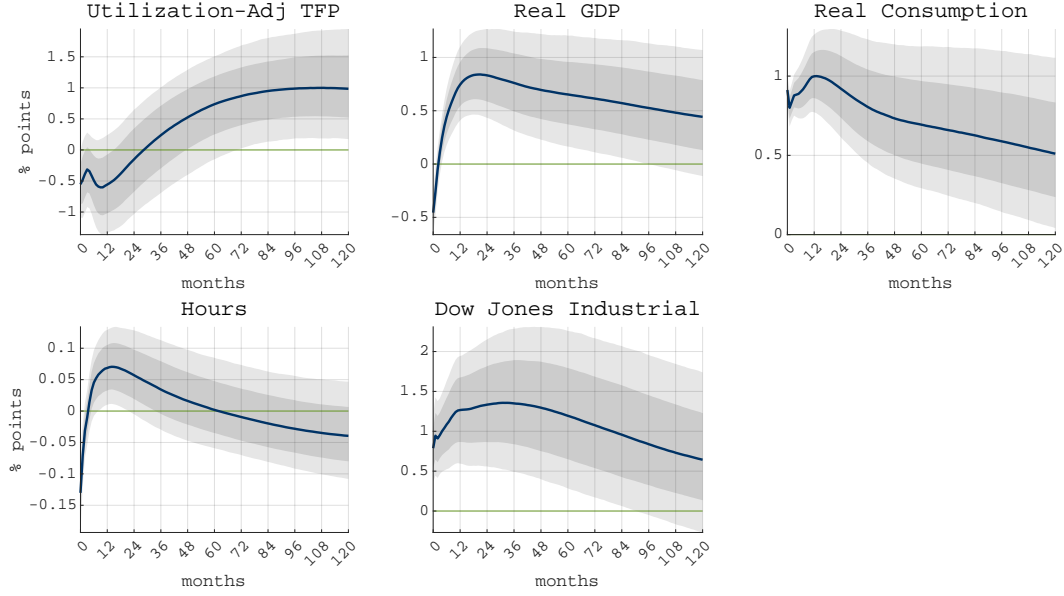
To construct an IV at monthly frequency we adapt the specification in Eq. (1) accordingly. While it was not always feasible to find an exact match for the entries used in the quarterly specification, we have attempted to preserve the nature of the exercise as much as possible. The main ingredients needed for the IV are patent applications (pa_t), forecasts for the macro outlook that capture up-to-date predictions prevalent at the time of the application filings ($\mathbb{E}_t[x_{t+h}]$), and policy controls (η_t).

Patent applications data are already available at monthly frequency from our source (Marco et al., 2015). The SPF forecasts that we use in our benchmark specification are distributed quarterly, which requires switching to a different survey. One possibility, and the one we have adopted, is to use the monthly Blue Chip forecasts. Blue Chip forecasts are published once a month and collect predictions about an array of different indicators at different quarterly horizons. However, the SPF and Blue Chip do not forecast the same set of variables, such that a match was only possible for the unemployment rate and inflation. In the quarterly specification we also included forecasts for non-residential investment and real corporate profits. These only become available in the Blue Chip Economic Indicators in 1993, therefore we have substituted them with the forecast for

(see Forni et al., 2019).

²⁷We thank John Fernald for his invaluable guidance and assistance in this endeavour. We make these data publicly available.

FIGURE 4: TECHNOLOGY NEWS SHOCKS IN A MONTHLY VAR



Note: Response to a technology news shock identified with patent-based external instrument. Monthly specification. VAR(12) with standard macroeconomic priors. Estimation sample January 1966 to December 2019. Identification sample January 1982 to December 2014, Shaded areas denote 68% and 90% posterior credible sets.

GDP growth in an attempt to encompass both. Similar to the quarterly case, we have included forecasts for the next quarter and next year in the monthly version.

Policy controls in the quarterly benchmark include narrative shocks for both monetary and tax policy. Monetary policy shocks are technically available at FOMC announcement frequency. For the monthly specification we have re-estimated and extended the series of [Romer and Romer \(2004\)](#) at monthly frequency. However, we were not able to switch to a monthly series for tax shocks.

The monthly IV is then estimated as the residual of the following regression

$$pa_t = c + \gamma(L)pa_t + \sum_{h=3,12} \beta_h \mathbb{E}_t[x_{t+h}] + \delta\eta_t + z_t, \quad (9)$$

where now the time indices t and h refer to months. Accordingly, pa_t denotes the monthly growth rate of patent applications, and $\gamma(L) = \sum_{j=1}^{12} \gamma_j L^j$. $\mathbb{E}_t[x_{t+h}]$ are the Blue Chip forecasts one quarter and one year ahead. And η_t is the monetary policy control.

We test the monthly IV in an illustrative 5-variable VAR that mimics the composition

of the VAR in Section 3. The VAR includes our new monthly series of utilization-adjusted TFP, a monthly series for GDP constructed as in Arias et al. (2019), real personal consumption expenditures, hours worked, and the stock market index. The VAR is estimated with 12 lags over the sample 1966-2019, and identified using the monthly IV. The shock is normalized to yield a peak response of TFP of 1 percentage point. Figure 4 reports the results.

Despite all the caveats associated with the construction of monthly versions of the IV and of utilization-adjusted TFP, results are remarkably in line with what discussed in the quarterly specification.

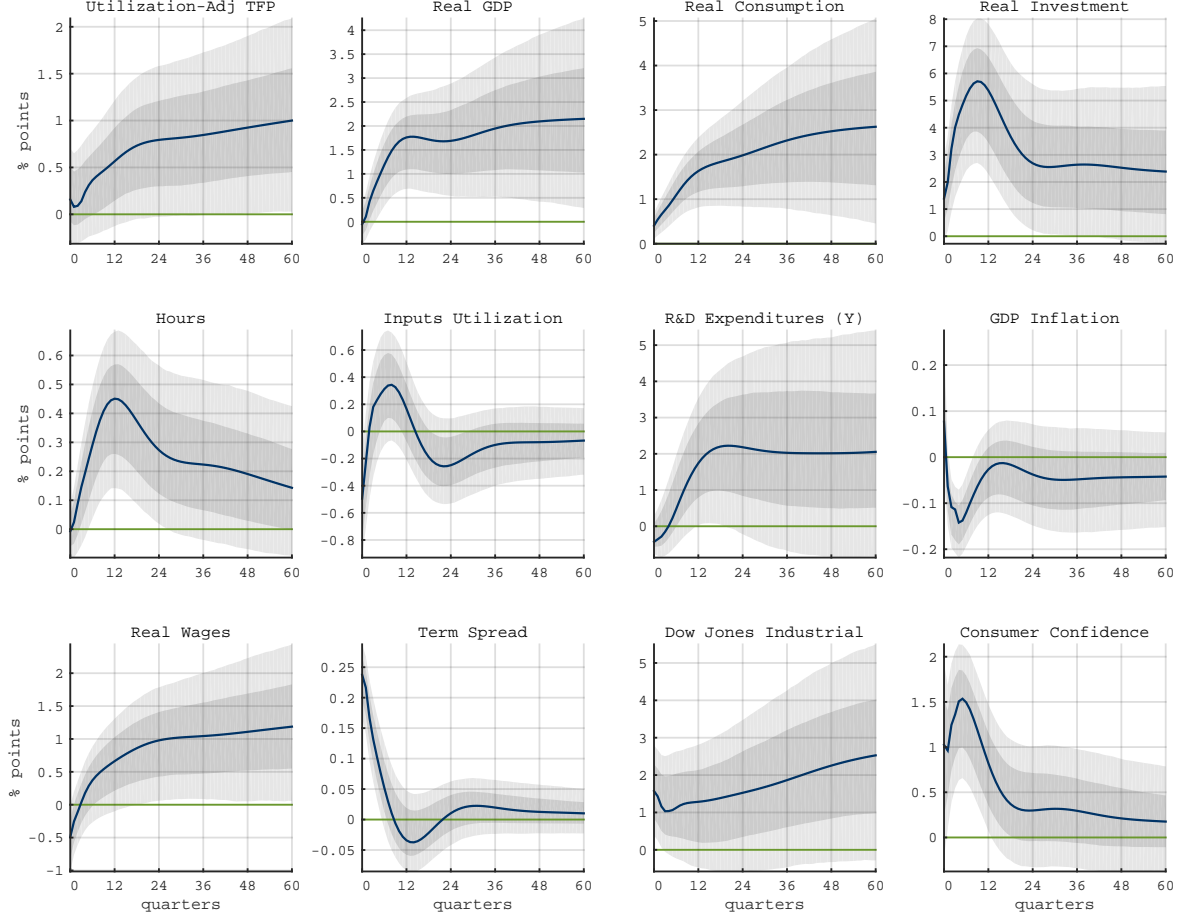
5 Technology News Shocks and Business Cycles

To study the propagation of technology news shocks to the broader economy we use a larger 12-variable VAR. The variables included cover real macroeconomic aggregates, financial markets, and expectations, and encompass the main indicators that feature in the theoretical literature on technology news shocks. This larger system enables us to characterize more carefully the response of the aggregate economy, and the importance of these structural disturbances in originating economic fluctuations. We offer a more in-depth discussion of our results in the next section.

In addition to the variables analyzed in the previous section, the VAR includes real investment, inputs utilization, R&D expenditures, the inflation rate and real wages, the term spread, and an index of consumer confidence taken from the Michigan Survey of Consumers. With the exception of inflation and the term spread, 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 the Appendix. The main features of the estimation are the same as in the previous section.²⁸ The IRFs are reported in Figure 5 and scaled such that the peak TFP response is equal to 1 percentage point. We discuss the robustness of our results below and report the associated charts in the Appendix.

²⁸We address concerns in e.g. Canova et al. (2009) and Fève et al. (2009) by re-estimating our baseline VAR with 12 lags. The richer parametrization substantially increases the computational burden but does not change our results. IRFs are reported in the Appendix.

FIGURE 5: 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. First-stage F-stat=5.99.

Most of the considerations made in the previous section carry through in the larger VAR. Albeit less precisely estimated, the response of TFP retains the main features discussed earlier. Namely, an initial muted response followed by a slow and persistent rise that becomes significant only years after the shock hits. Conversely, all other macro aggregates respond more swiftly, and tend to peak within the first three years. Both output and hours do not respond on impact, and are in distinctly positive territory thereafter. But while the response of hours tends to revert over time, output remains elevated throughout. Investment displays a similarly shaped response. While positive, the initial reaction is only marginally significant at conventional levels. The magnitude

TABLE 3: CONDITIONAL CORRELATIONS

	$h = 1$	$h = 4$	$h = 12$	$h = 40$
Real GDP	0.450	0.896*	0.985**	0.993**
Hours	0.597	0.896*	0.982**	0.916**

Notes: Conditional correlations between consumption, GDP, and hours implied by the identified VAR at selected horizons. Estimation sample 1960:I - 2019:IV. Identification sample 1982:I - 2006:IV. *, ** denote statistical significance at 68% and 90% levels respectively. Horizons in quarters.

of the responses is economically important. Output rises by almost 2 ppt at peak, while investment increases by about 6 ppt in annual terms. Consumption retains the positive and significant impact response observed earlier, although the magnitude of the initial adjustment is significantly more modest at less than half a percentage point. We return on the response of consumption in the discussion of our results in the next section. Inputs utilization—the same variable used to correct TFP—drops modestly on impact to increase a few quarters afterward. R&D expenditures also increase with delay, presumably as a result of the increase in both investment and output.

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 long before any material increase in TFP is recorded. We quantify the extent of the comovement in Table 3, where we report the conditional correlation between consumption, output, and hours worked at selected horizons.²⁹ We note that while the delayed response of output and hours makes the short-horizon correlations not significant, the correlations are generally large and positive at all horizons, which makes the shock a plausible enabler of business cycles. This aligns with findings in e.g. [Beaudry and Portier \(2006\)](#); [Christiano et al. \(2003\)](#) but contrasts with e.g. [Barsky and Sims \(2011\)](#). Although the latter identification scheme and associated comovements are shown to be sensitive to the TFP vintage used (see [Cascaldi-Garcia, 2017](#); [Kurmann and Sims, 2021](#)).

The identified shock is mildly deflationary. While the initial response is not significant, inflation falls within the first year following a negative hump-shape that reaches a trough

²⁹The full set of correlations is reported in the Appendix.

of about negative 14 bps 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 a sharp initial decline instead (see e.g. [Barsky and Sims, 2011](#); [Barsky et al., 2015](#)). Aggregate real wages fall marginally on impact to improve robustly at longer horizons.

Financial variables respond strongly and on impact. The stock market is quick in pricing-in positive news, and remains elevated throughout, although the response becomes less precisely estimated relative to the 5-variable VAR. Using broader stock market indices such as the S&P 500 makes the estimated response more uncertain. 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 slope of the yield curve, here measured as the spread between the 10-year and the 1-year Treasury rates, rises by about 20 bps on impact. The response of the yield curve is qualitatively similar to what is documented in [Kurmann and Otrok \(2013\)](#), but the magnitudes in our case are smaller. We return to the response of the yield curve and the likely monetary policy response to the shock in the next section. Finally, consumer confidence rises robustly at medium horizons, but the impact response is only marginally significant at conventional levels. We verify that neither the global financial crisis nor the ZLB sample drive or affect our results (see Appendix).

Using an instrumental variable that is based on patent applications rather than grants is important to unveil the full extent of the anticipatory effects. Comparing our results with those in [Cascaidi-Garcia and Vukotić \(2022\)](#) reveals that the latter choice leads to responses that are mostly significant on impact, and very short-lived, suggesting that patent grants may be capturing some non-trivial cyclical variation as well, as also noted in [Christiansen \(2008\)](#).

The set of response functions is compatible with the identified shock being an originator of business-cycle type of fluctuations. But whether it can be thought of as a meaningful driver of business cycles ultimately rests on the share of aggregate fluctuations that it can account for.

Table 4 reports the average shares of explained variation over selected frequency intervals for all variables in our VAR. Each column reports the percentage share of variance

TABLE 4: 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	0.33	0.42	3.76	11.88
Real GDP	1.52	6.61	13.13	32.85
Real Consumption	4.74	6.78	20.16	34.27
Real Investment	1.44	8.70	26.37	34.77
Hours	1.17	6.04	15.73	29.80
Inputs Utilization	5.54	3.82	4.98	4.05
R&D Expenditures	1.92	6.78	6.16	9.14
GDP Inflation	1.98	10.28	2.30	3.54
Real Wages	3.34	3.67	6.64	19.10
Term Spread	33.62	23.24	10.65	7.42
Dow Jones	4.67	2.00	1.43	14.32
Consumer Confidence	1.42	9.45	16.7	20.60

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.

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). The algorithm used for the decomposition builds on [Altig et al. \(2011\)](#) and is described in detail in the Appendix. 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, and similar to what we found in the 5-variable VAR, the shock recovered by the IV is most explanatory for TFP in the very long run. Conversely, the contribution of the shock to higher frequency fluctuations in productivity is negligible. This is consistent with the identified shock being mostly a driver of the trend component of TFP. Second, the shock is responsible for a relatively

³⁰Intuitively, 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, frequency-based and time-based forecast error variance decompositions are reported in the Appendix.

small fraction of the fluctuations in main business cycle aggregates at business cycle frequencies, but it accounts for over a third of the variation in consumption, investment, output and hours in the very long-run. This apparent disconnect between drivers of business cycles and of long-run fluctuations echoes findings in [Angeletos et al. \(2020\)](#). Third, the shock explains around 15% of the long-run variance of the stock market, and is responsible for over a third of the variation of the yield curve in the short-term, which points in the direction of [Kurmann and Otrok \(2013\)](#). 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 productivity unaccounted for. As a result, caution should be used when comparing the shares of forecast error variance with those reported in other studies.


6 Discussion of the Results

In this section we take stock of our results and use them as guide to interpret the features of the identified shock, and how it may diffuse through the economy. In this context, it is important to bear in mind that the aggregate IRFs that we report are likely to result from a combination of multiple and distinct effects that jointly determine how households, firms, financial markets, and the central bank respond to the shock, and that the empirical nature of our exercise does not allow to disentangle. In what follows, we make use of additional variables to aid with the interpretation, and leave a more formal model-based characterization for future research.³¹

As noted, and consistent with patent applications marking the early stages of the innovation process, the IV recovers a shock that improves long-term productivity significantly, but has no noticeable bite on TFP in the short-run. One interesting question is what type of technological change is the IV likely to be picking up. To this purpose, recall that our identification strategy centers on the signal embedded in so-called utility patents. These patents encompass advancements in products, machinery, and processes. In turn, advancements are intended as improvements of existing technologies as well

³¹We study the response of these additional variables by separately including them in the VAR. Full IRFs are reported in the Appendix.

FIGURE 6: PRICE OF INVESTMENT. UNEMPLOYMENT & CONSUMER EXPECTATIONS



Note: Response of selected variables separately included in the VAR. 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.

as the creation of new technologies altogether. This definition makes it likely that the identified shock may combine elements of both embodied and disembodied technological change. Some evidence in this direction is provided by the response of the relative price of investment (Figure 6 panel a) which tends to contract persistently over time, indicating that the shock may have some of the flavor of the investment-specific technological improvements of e.g. Fisher (2006).³²

News about this future (and potentially investment-intensive) productivity improvement is released in advance—and channeled by the IV as per our identifying assumption—which opens up the door for the economy to adjust and react in anticipation to it. Our results show that output, consumption, investment and hours all expand a few quarters

³²See also Justiniano et al. (2010, 2011); Ben Zeev and Khan (2015). However, whether this is the main channel through which the shock operates remains unclear. 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 be seamlessly applied in our context.

after the shock hits.³³ The large asynchronicity between the speed of adjustment of these macro aggregates relative to the improvement in TFP is consistent with such anticipatory effects being active and playing a potentially important role.

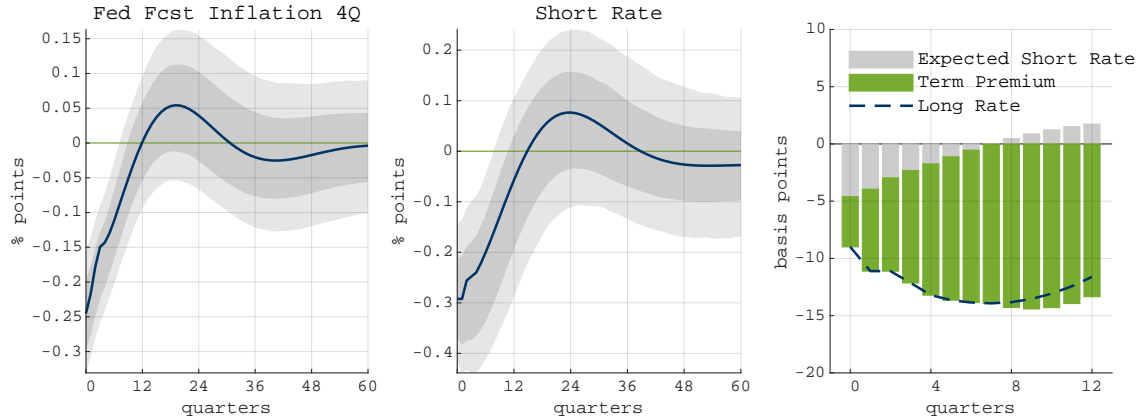
Anticipatory effects are also typically advocated to make sense of the systematic increase in consumption, which is a fixture of virtually all empirical studies. To shed more light on the reaction of households behavior, Figure 6 reports the response of the unemployment rate (panel b) as well as of consumers' expectations about unemployment and business conditions over a one- and five-year horizon respectively (panels c and d). Taken together, these responses paint a rather nuanced picture. As noted, consumer confidence tends to improve robustly shortly after the shock hits, even though the impact rise is only marginally significant. Very interestingly, short-term expectations of unemployment rise sharply upon realization of the shock, to quickly revert thereafter. The survey asks respondents whether they expect unemployment over the next twelve months to be higher, lower, or about the same as current, and returns the balance of responses as an indicator. Therefore, the IRF in the figure is to be interpreted as an increase in the share of respondents that expect unemployment to rise. While to different degrees, these two sets of responses seem to suggest that the perception of the short-term effects of technology news may be potentially unfavorable, or at least not unequivocally benign. This initial reaction however dissipates over a relatively short horizon. And, consistently, expectations about the medium-term outlook rise significantly (panel d).³⁴

How these expectations may interact with the remainder of the variables to concur to determine the response of aggregate consumption is a question that is best addressed in the context of a model. But, based on our results, we posit that there may be at least two elements at play. On the one hand, the aggregate responses may mask compositional effects and heterogeneity across workers. Consider for example the case in which firms switch to more capital-intensive technologies, or reconfigure towards automation, or introduce technologies that render the skills of some incumbent workers obsolete (e.g. Kogan et al., 2021). These cases can plausibly lead to expectations of unemployment to

³³Note that, differently from theoretical models, the VAR is unrestricted, and does not impose an aggregate resource constraint; as such, some discrepancies may arise when comparing the impact responses of output, consumption and investment.

³⁴Barsky and Sims (2012) argue that this variable in particular is likely to embed news about future productivity. See also Cochrane (1994).

FIGURE 7: MONETARY POLICY RESPONSE



Note: Response of selected variables separately included in the VAR. VAR(4) with standard macroeconomic priors. Estimation sample 1975-I:2018-IV; Identification sample 1982-I:2006-IV. Shaded areas denote 68% and 90% posterior credible sets.

increase in the short-term. And workers that are negatively impacted may reasonably reduce their consumption. However, there is no a priori reason to believe that this should apply in equal measure to all workers, or that indeed this should be thought of as the representative or predominant case. In the VAR the impact response of aggregate hours is muted, but the unemployment rate rises on impact (Figure 6 panel b), suggesting that adjustments along both the intensive and extensive margins may be at play. On the other hand, there may be meaningful heterogeneity across the income distribution. While aggregate wages decline mildly on impact, the stock market rises significantly. Depending on the relative distribution of labor income and financial wealth, it is plausible that the combination of responses may leave some segments of the population significantly better off.

A final point refers to the possible amplification that may result from the endogenous response of the monetary authority to the shock (see also Kurmann and Otrok, 2013). Figure 7 reports the response of the short-term interest rate, of the Federal Reserve's expectation of inflation a year hence, which we take from the official Greenbook/Tealbook publication, and of the decomposition of the response of the 10-year rate into its expectation and term premia components, as implied by our VAR.³⁵ Due to the sample

³⁵Note that the availability of Greenbook forecasts for inflation restricts the VAR sample to 1975-2018. 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

considered including the zero-lower-bound period, we use the one-year nominal interest rate to measure the short-term policy rate.

The one-year rate falls by about 30 bps on impact, which roughly matches the size of the decline in expected 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. Recall also that the slope of the yield curve—the spread between the 10-year and the 1-year Treasury rates—rises by about 20 bps on impact. Together with the short-term interest rate response, this implies an impact decline of long-term yields of about 10 bps. We further note that the 1-year rate returns to pre-shock levels relatively quickly, and is hence likely not to fully account for the impact fall in the 10-year Treasury yield. In turn, this implies that following a technology news shock the term premium declines. Indeed, the decomposition of Figure 7 shows that the term premium remains compressed for an extended period of time, which aligns with findings in [Crump, Eusepi and Moench \(2016\)](#). In addition to anticipatory effects, the fall in borrowing costs, coupled with compressed risk premia, may act as a further powerful amplifier for the propagation of news shocks.

7 Conclusions

How does the aggregate economy react to a shock that raises expectations about future productivity growth? In this paper, we have provided an empirical answer to this question using a novel patent-based instrumental variable for the identification of technology news shocks that enables 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 contemporaneous policy shocks. Our sole identifying assumption is that 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. Our results reveal four main patterns. First, macro aggregates react well in advance of any material increase in TFP, suggesting an important role for anticipatory effects. Second, the conditional of the 10-year rate, and the average expected response of the 1-year rate at horizons $h, h + 4, \dots, h + 36$.

comovements implied by our identified VAR are positive, and therefore enable technology shocks as a potential originator of business cycles. Third, most macro aggregates tend to respond to the shock with some delay. Fourth, while an important driver of long-run dynamics, the recovered shock only explains a relatively modest fraction of the variation of main macroeconomic aggregates at business cycle frequencies.

We further document a nuanced response of consumers' expectations in response to the shock, and that the central bank tends to respond to the shock by easing policy. Lower borrowing rates and compressed term premia appear as likely amplifiers of the short-term effects of news shocks.

In our analysis we have focused on the aggregate effects of technology news shocks, and therefore looked at the aggregate signal embedded in the universe of USPTO patent applications. In practice, however, different sectors may behave differently. Equally, patents in certain industries may carry a stronger or more pervasive news signal than in others. This opens up interesting questions concerning heterogeneous effects at sectoral and even at the firm level. We leave these important questions for future research.

Data Availability Statement The data and code underlying this research are available on Zenodo at <https://doi.org/10.5281/zenodo.15298755>.

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.
- Arias, Jonas E., Dario Caldara, and Juan F. Rubio-Ramírez (2019) “The systematic component of monetary policy in SVARs: An agnostic identification procedure,” *Journal of Monetary Economics*, Vol. 101, pp. 1–13.
- 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.
- 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 Fève, Alain Guay, and Franck Portier (2019) “When is nonfundamentality in SVARs a real problem?” *Review of Economic Dynamics*, Vol. 34, pp. 221–243.

- Ben Zeev, Nadav and Hashmat Khan (2015) “Investment-Specific News Shocks and U.S. Business Cycles,” *Journal of Money, Credit and Banking*, Vol. 47, No. 7, pp. 1443–1464.
- 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.
- Canova, Fabio, David Lopez-Salido, and Claudio Michelacci (2009) “The effects of technology shocks on hours and output: a robustness analysis,” *Journal of Applied Econometrics*, Vol. 25, No. 5, pp. 755–773.
- Cascaldi-Garcia, Danilo (2017) “News Shocks and the Slope of the Term Structure of Interest Rates: Comment,” *American Economic Review*, Vol. 107, No. 10, pp. 3243–49, 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.
- 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.
- 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.
- Federal Reserve Bank of Philadelphia (2023) “Survey of Professional Forecasters.”
- Federal Reserve Bank of St. Louis (2023) “FRED - Federal Reserve Economic Data.”
- Fernald, John G. (2014) “A quarterly, utilization-adjusted series on total factor productivity,” Working Paper Series 2012-19, Federal Reserve Bank of San Francisco.
- Fernald, John (2021) “A quarterly, utilization-adjusted series on total factor productivity.” 2021/12 Vintage.
- Fève, Patrick, Julien Matheron, and Jean-Guillaume Sahuc (2009) “On the dynamic implications of news shocks,” *Economics Letters*, Vol. 102, No. 2, pp. 96 – 98.

- 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 (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.
- 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.

- 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.
- 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.
- 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.
- 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 (2023) “Identification with External Instruments in Structural VARs,” *Journal of Monetary Economics*, Vol. 135, pp. 1–19.
- 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, 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 (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.