When Creativity Strikes: News Shocks and Business Cycle Fluctuations☆

Silvia Miranda-Agrippino*
Northwestern University
Bank of England, CfM (LSE)

Sinem Hacıoğlu Hoke†
Bank of England
Data Analytics for Finance and Macro (KCL)

Kristina Bluwstein‡
Bank of England

First Version: August 22, 2018, Revised: May 29, 2019

Abstract

We propose a novel external instrument for the identification of technology news shocks based on patent applications. Technology diffuses slowly, and affects TFP in an S-shaped pattern. Responsible for a tenth of economic fluctuations at business cycle frequencies, the shock elicits a slow, positive response of quantities, and a sluggish contraction of prices followed by an endogenous easing of the monetary stance. The ensuing expansion substantially predates any material increase in TFP. The stock market prices in technology news on impact, but consumer expectations take longer to adjust, consistent with a New-Keynesian framework with nominal rigidities featuring informationally constrained agents.

Keywords: Technology News Shocks; Business Cycle; Identification with External Instruments; Patents Applications; Innovation.

JEL Classification: E32, O33, O34, C36

*Department of Economics, Northwestern University, 2211 Campus Drive, Evanston IL 60208, US. E-mail: silvia.miranda-agrippino@northwestern.edu Web: www.silviamirandaagrippino.com

†Monetary Analysis, Bank of England, Threadneedle Street, London EC2R 8AH, UK. E-mail: sinem.hacioglu@bankofengland.co.uk Web: www.sinemhaciogluhoke.com

‡Financial Stability, Strategy and Risk. Bank of England, Threadneedle Street, London EC2R 8AH, UK. E-mail: kristina.bluwstein@bankofengland.co.uk Web: www.kristinabluwstein.com/

☆ We are indebted to Franck Portier and André Kurmann for insightful discussions and many useful suggestions on earlier versions of this paper. We are grateful to Bob Barsky, Fabio Canova, Cristiano Cantore, Jordi Galí, Maria Luisa Mancusi, Michael McMahon, Karel Mertens, Emanuel Mönch, Evi Pappa, Valerie Ramey, Morten Ravn, Ricardo Reis and Barbara Rossi for detailed comments. We also thank participants at the Advances in Applied Macro/Finance Conference, Bank of England, Catholic University of Milan, Dallas Fed, ES European Winter Meeting, Northwestern, UCSD, and UPF for helpful discussions. We thank Luca Gambetti for sharing his code on the information sufficiency test. Silvia Miranda-Agrippino gratefully acknowledges support and hospitality from Northwestern University where this research was partly conducted. The views expressed are those of the authors and do not reflect those of the Bank of England or any of its Committees.

1 Introduction

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

In this paper, we set out to answer the following question: 'How does the aggregate economy react to a shock that raises expectations about future productivity growth?' We provide an empirical answer to this question in an information-rich quarterly VAR that incorporates many relevant aggregates, such as output, consumption, investment and labor inputs, as well as forward looking variables such as asset prices, interest rates, and consumer expectations. The novelty in our approach resides in the identification of technology news shocks. We construct an external instrument for identification by using the unforecastable component of all patent applications filed at the U.S. Patents and Trademark Office (USPTO) over the past forty years. The intuition behind our identification

[1]The empirical literature on technology news shocks is vast, and we review it when discussing our results in Sections 4 and 5. At the poles of the debate are the advocates of the news-driven business cycle hypothesis such as e.g. Beaudry and Portier (2006, 2014); Beaudry and Lucke (2010), and its opponents, such as e.g. Barsky and Sims (2011, 2009); Kurmann and Otrok (2013); Barsky et al. (2015); Kurmann and Sims (2017). In Beaudry and Portier (2006) news shocks are orthogonal to current productivity, but are the sole driver of TFP in the long run (e.g. Galí, 1999; Francis and Ramey, 2005). Other works have identified technology shocks as those maximizing the forecast error variance of productivity at some long finite horizon (e.g. Francis et al., 2014), or over a number of different horizons (e.g. Barsky and Sims, 2011). Other contributions have highlighted the differences arising from e.g. modeling variables in levels rather than in first differences, allowing for cointegrating relationships among variables (together with their number and their specification), accounting for low frequency structural breaks, accounting for other policy-related concomitant factors, and enriching the information set in the VAR. Examples include Christiano et al. (2003); Francis and Ramey (2009); Mertens and Ravn (2011); Forni et al. (2014).

is simple: Patent applications, by their very nature, are a potential promise of future technological change. However, they may themselves be the result of current economic booms, or of past news. We account for this endogeneity by controlling for expectations about the economy that were formed prior to the filing dates, other contemporaneous policy changes, and lagged applications.

Specifically, we recover the external instrument as the component of patent applications that is orthogonal to (i) its own lags; (ii) a selection of forecasts at different horizons intended to capture pre-existing expectations about macroeconomic developments that may influence the decision of filing a patent in the time unit, and which we take from the Survey of Professional Forecasters; and (iii) other contemporaneous unanticipated monetary and fiscal policy changes. The starting point for the construction of our external instrument are the monthly 'USPTO Historical Patent Data Files' assembled in Marco et al. (2015) that provide a comprehensive record of all publicly available applications and granted patents registered at the USPTO since 1981. To the best of our knowledge, the properties of these data have not been previously explored in empirical macroeconomics, or in the context of identifying technology news shocks.[2] Our instrument is associated with large increases in the aggregate measure of innovation of Kogan, Papanikolaou, Seru and Stoffman (2017). The index measures the expected economic importance of technological innovations, and correlates strongly and positively with forward citation counts, in turn a measure of their scientific value.

Contrary to the existing literature, our identification strategy allows us to dispense from potentially strong a priori assumptions related to the duration of the effects of news shocks, the long-run drivers of technology, or the length of time that is required for the news to affect the current level of technology. Because of the minimal set of restrictions required for identification, our framework allows us to investigate whether news shocks generate the type of behavior that was assumed in earlier identification schemes. In other words, what constituted an assumption in earlier studies becomes instead a result in our

[2] Earlier studies that have similarly employed patent applications to measure the effects of technology shocks (reviewed below) have typically relied on annual data. The use of patent data to measure technological advancements at industry level dates back at least to Lach (1995). Griliches (1990) provides a review of the uses of patent data as indicators for technological change in economic analysis. Hall and Trajtenberg (2004) use the annual NBER patent citations data file described in Hall et al. (2001) to show that granted patents, and their citations, can be used to measure evidence of General Purpose Technologies.

setting. Moreover, our approach is robust to mismeasurements in commonly used empirical estimates of aggregate technology (see e.g. discussions in Fernald, 2014; Kurmann and Sims, 2017). The identifying assumptions in our SVAR-IV (Mertens and Ravn, 2013; Stock and Watson, 2012, 2018) are that the instrument is informative about contemporaneous technology news, and that this is the only channel through which the instrument and the VAR innovations are related (Miranda-Agrippino and Ricco, 2018). Importantly, because innovations can in principle be released to the public under a 'patent-pending' status, our identification scheme does not warrant imposing orthogonality with respect to the current level of technology, which is a typical assumption in the news literature (see e.g. Beaudry and Portier, 2006; Barsky and Sims, 2011, among many others). In this respect, our identification is akin to Barsky et al. (2015); Kurmann and Sims (2017).

While such orthogonality condition is not imposed a priori, our external instrument recovers news shocks that have essentially no effect on TFP either on impact, or in the four years immediately afterwards. TFP then rises robustly following a persistent hump that reaches a peak after 6 to 7 years.[3] Both the shape and timing of the TFP response are consistent with the S-shaped pattern that is typical of the slow diffusion of new technologies documented, among many others, in Griliches (1957); Rogers (1962) and Gort and Klepper (1982).

By the time TFP materially departs from its initial level, all other variables in our VAR have reached the peak of their dynamic response to the shock. The arrival of news about future technological improvements triggers a sustained, albeit somewhat delayed, economic expansion: output, consumption, investment, hours worked and capacity utilization all rise to peak at the two-year horizon. Hence, the pattern of impulse response functions that we recover does lend credit to a 'news-view' in the spirit of what is described in e.g. Beaudry and Portier (2006). In recent influential work, Chahrour and Jurado (2018) have proven that any model with news has an observationally equivalent noise representation. Seen through this lens, news – that are realized on average –, confound the effects of 'pure beliefs' with those associated with changes in future fundamentals. Conversely, noise, orthogonal to fundamentals at all leads and lags, captures the

[3]The time that it takes for news to translate into meaningful changes in future TFP is sensibly longer than the two-year anticipation lag that is typically assumed in the literature (see e.g. Schmitt-Grohé and Uribe, 2012; Beaudry and Portier, 2014; Faccini and Melosi, 2018).

essence of 'pure beliefs'. The large asynchronicity in the timing of the estimated dynamic responses seems to suggest that the aggregate effects of technology news that we unveil may be predominantly (if not entirely) driven by beliefs. The shock that we recover, however, is not a main driver of economic fluctuations. At business cycle frequencies, about a tenth (on average) of aggregate fluctuations is accounted for by the estimated news shock; importantly, it accounts for at most a third of the variation of TFP in the very long run.

The pattern of dynamic responses that we recover is consistent with the predictions of New Keynesian models with nominal rigidities, particularly those where such frictions arise due to imperfect common knowledge (e.g. Mankiw and Reis, 2002; Woodford, 2003). After an inertial initial reaction, prices eventually decline. Conversely, real wages rise at medium horizons, but contract on impact. The monetary authority endogenously responds to the fall in (expected) inflation by lowering nominal rates on impact, and more than proportionally. Hence, real short-term rates decline at a time when the natural rate of interest, proportional to the expected growth rate of technology, is rising (see e.g. Christiano et al., 2010). This suboptimal response of the central bank can also be rationalized in terms of information rigidity: the central bank responds to its best forecast of current and future fundamentals, that may diverge from actual realizations (see e.g. discussion in Lorenzoni, 2011; Sims, 2012). The monetary easing also offers a potential amplification channel for news shocks that works through the compression of risk (term) premia (see also Crump et al., 2016). A noisy signal about future technological changes can also be responsible for agents overweighting current conditions when forming expectations about the future (see e.g. Coibion and Gorodnichenko, 2015). In this sense, the initial rise in consumers' expectations about future unemployment that we document is consistent with the initial deterioration in labor market conditions, reflected in the fall of both hours worked, and wages. In turn, this can help explain the initial negative reaction of consumers' expectations about current conditions, and about the expected business outlook five years hence. In this respect, our results suggest caution in interpreting innovations in consumer confidence indicators as a 'pure' measure of news (e.g. Cochrane, 1994; Barsky and Sims, 2012). In fact, when we compare responses to our news shock with those triggered by a positive contemporaneous TFP innovation, we find that consumer confidence jumps up on impact only in the latter case.

Our work is closely related to a stream of studies that have relied on empirical measures of technological changes in order to identify the effects of technology 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 contain labor inputs and productivity. Christiansen (2008) extends on the previous 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 construct a measure for technological changes intended to capture the time in which the novelty is effectively commercialized. Responses of aggregate variables are estimated in a set of bivariate VARs with the publication index ordered last. More recently, Baron and Schmidt (2014) have used technology standards and a recursive identification to infer on the aggregate implications of anticipated technology shocks. For what concerns the chronological placement in terms of anticipation lag, for each technological improvement, industry standardizations sit somewhere in between patent applications and the publication of the relevant title.[4] Our paper differs from these contributions in several ways. First, these studies address the fundamental endogeneity of empirical measures of technological changes only to the extent that it is captured in the reminder of variables included in the bi/tri-variate VARs. Other than relying on a richer VAR specification, in the construction of our instrument we recognize that the cyclical nature of patent applications may also be influenced by pre-existing expectations about the future, which we capture using an array of survey forecasts at different horizons, and by other contemporaneous policy changes. We argue that this is a crucial step for the correct identification of contemporaneous news as opposed to a convolution of current and past news, and current innovations to the technology process. 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

[4] In an international context, Arezki et al. (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.

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.

The structure of the paper is as follows. Section 2 introduces the external instrument that we design for the identification of technology news shocks, and describes the patent data that we use for its construction. Section 3 lays out the identification assumptions in our SVAR-IV. Section 4 collects the results, which we discuss in detail in Section 5 against the main transmission mechanisms proposed in the literature. Finally, Section 6 concludes. Additional material is reported in the Appendix.

2 A Patent-Based Instrument for News about Future Technological Changes

In the vast majority of industries, and particularly since the 20^{th} century, the introduction of technological innovations follows a relatively standardized process. Typically, before an invention – intended as either a brand new product or production process, as well as an amelioration to existing ones – is disclosed, the owner proceeds to file a patent application in order to protect her creation. The legal protection that is granted to patent holders ensures that an individual or business has a set number of years in which to capitalize on the invention. Hence, the incentive to protect new inventions through appropriate patent registrations is high. The length of time that elapses from the time in which a patent application is filed to when it is then granted, and the invention eventually diffuses 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 (see e.g. Lach, 1995; Hall and Trajtenberg, 2004). At the same time, although the bulk of informativeness of patent applications lies in the future, and patented products cannot be copied by others during the protected period, some inventions are released under a 'patent-pending' status. This initial diffusion of the invention spreads new knowledge to the public, some of whom may be able to improve upon that invention themselves. Thus, it is possible in principle that patent applications may also embed a signal for current technological changes.

The decision to file a patent application at any given time is, however, a fundamentally endogenous one. Hence, the direction of causality could run both ways, with patent applications being induced by current economic booms, and/or past favorable news. Separately, not all patent applications are ultimately granted, and therefore the signal about future productivity changes is necessarily only a partial one.

2.1 Information in Patent Data

We use the 'US Patents and Trademark Office (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, published and unpublished, and granted patents registered at the USPTO from January 1981 to December 2014. The stocks include pending applications and patents-in-force; flows include new applications, patent grants and abandonments.[5]

Our starting point for the analysis is the monthly flow of all new utility patent applications.[6] Utility patents, also known as 'patents for invention', cover the creation of new or improved, and useful, products, processes or machinery. We then construct quarterly variables by summing up the monthly flows within each quarter. The left panel of Figure 1 plots the time series of quarterly patent applications from 1981 to 2014, across all NBER categories (Marco et al., 2015). 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 categories. Patent applications have increased substantially over the past 40 years and, as visible from the chart, patents classified under 'computers and

[5]See also Hall et al. (2001). The dataset is available at `http://www.ustpo.gov/economics`.

[6]We discard information relative to both abandonments and patents granted. While granted patents can potentially provide a stronger signal about future technological changes, they tend to be significantly more cyclical than patent applications. Also, the production of the invention may already have started while the application was pending. Hence, most of the 'news content' in patent applications may be exhausted by the time it is granted. Moreover, as Christiansen (2008) discusses, the issuance highly depends on the intensity of labor and administrative cycles at the USPTO. Further details on patent data are in Appendix D.

FIGURE 1: PATENT APPLICATIONS & AGGREGATE INNOVATION

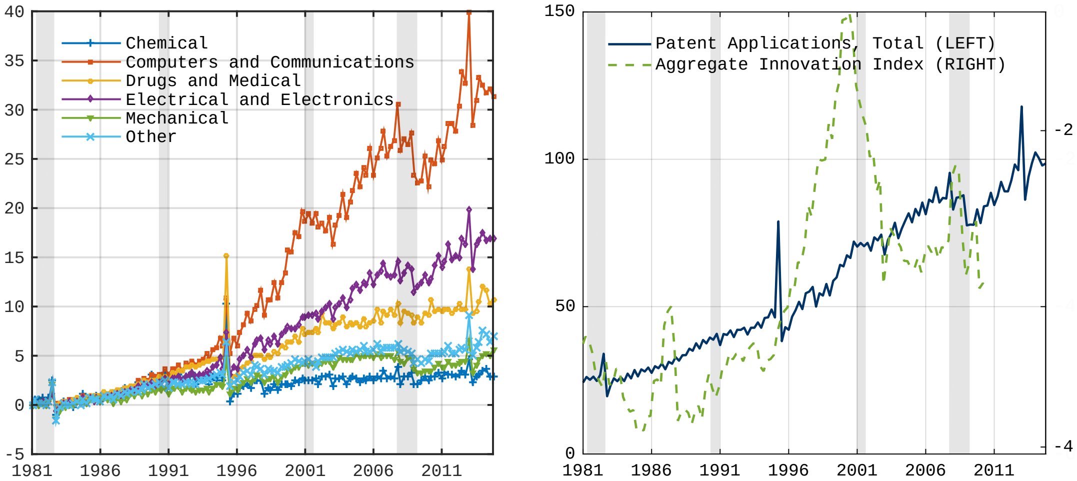

Note: [LEFT] Patent applications across all NBER categories. Quarterly figures obtained as sum of monthly readings, 1981-I=0. Thousands. [RIGHT] Total number of applications (sum across categories), thousands, left axis. Kogan et al. (2017) aggregate measure of economic value of innovations, GDP weighted,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

communications' have enjoyed a faster trend. Patent applications across all categories tend to slide after recessionary episodes, providing some preliminary evidence of their cyclical nature.

There have been three important regulatory changes in patenting in 1982, 1995, and 2013. All these regulations affected the number of applications when they came into effect, as shown by the spikes in Figure 1. In 1982, the old Court for Customs and Patent Appeals was abolished and a new Court of Appeals for the Federal Circuit was established; the new court provided more protection to the owners of patents against infringement. In 1995, the U.S. implemented the changes agreed upon in 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. The large impact on the number of patent filings was due to a change in patent terms; as of June 1995, patent terms were set to 20 years from filing, and away from the previous practice of 17 years after issuance. Finally, in March 2013, the U.S. implemented the rules under the America Invents Act. These sets of rules were designed to address the right to file a patent application, and implied that applications filed on or after March 2013 were to be governed under the new priority rule 'first-inventor-to-file', rather than the pre-existing 'first-to-invent'.[7] All these three regulatory changes led to an increase in applications prior to their implementation. However, since they were not legislated in response to considerations related to either current or anticipated economic conditions, they provide us with important sample variation which we exploit for the identification.

As an additional piece of evidence, the right panel of Figure 1 plots the total number of applications (sum across the NBER categories) against the aggregate index of innovation of Kogan et al. (2017) (dashed line). Using their data, we have constructed a quarterly equivalent of the index as a GDP-weighted sum of the economic value of all patents granted within each quarter. The data cover up to 2010-III. At the firm-patent level, the value of each patent is measured based on the change in the firm stock price in a three-day window that brackets the date in which the patent is granted to the firm. Because it is based on financial data, this is a forward looking measure of the private, economic value of innovations. Kogan et al. (2017) document that their measure is strongly positively correlated with forward citations. This in turn refers to the number of citations that the patent receives in the future, and is hence regarded as a proxy of the scientific value of the patented invention. We note that in the relevant sample, the large spikes in the number of applications tend to correspond to substantial increases in the innovation index, and this is particularly true in the nineties. Hence, while only a subset of the applications in our data are eventually granted, we confirm that the exogenous sample variation introduced by the changes in legislation is also informative about the overall 'innovation content'.

We investigate the endogeneity of quarterly patent applications in Table 1. Here we regress the quarterly growth rate of all patent applications on its first four lags, and on pre-existing expectations about the state of the economy at different forecast horizons taken from the Survey of Professional Forecasters (SPF). The vector of forecasts $\mathbb{E}_t[w_{t+h}]$ includes real output growth, the unemployment rate, inflation (GDP deflator), real federal government spending, real non-residential investments, and real corporate profits net of taxes.[8] The forecast horizon is expressed in quarters, such that $\mathbb{E}_t[w_t]$

[7]For a detailed description of the Leahy-Smith America Invents Act the reader is referred to `https://www.uspto.gov/sites/default/files/aia_implementation/20110916-pub-l112-29.pdf`

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

Table 1: Endogeneity of Patent Applications

	(1)	(2)	(3)	(4)
pa_{t-1}	−0.849	−0.941	−0.932	−0.894
	(0.095)	(0.091)	(0.094)	(0.087)
pa_{t-2}	−0.480	−0.638	−0.639	−0.592
	(0.105)	(0.099)	(0.099)	(0.098)
pa_{t-3}	−0.273	−0.439	−0.428	−0.389
	(0.085)	(0.081)	(0.077)	(0.068)
pa_{t-4}	0.002	−0.076	−0.079	−0.067
	(0.094)	(0.077)	(0.076)	(0.074)
$\mathbb{E}_t[w_t]$		5.59		
		[0.000]		
$\mathbb{E}_t[w_{t+1}]$			7.94	
			[0.000]	
$\mathbb{E}_t[w_{t+4}]$				4.19
				[0.001]
regulation dummy		✓	✓	✓
constant	✓	✓	✓	✓
Adj-R^2	0.465	0.811	0.810	0.787
N	131	131	131	131

Notes: Granger Causality. Dependent variable: $pa_t = 100 \times (\ln PA_t - \ln PA_{t-1})$. $\mathbb{E}_t[w_{t+h}]$ denotes SPF forecast for quarter $t+h$ published at t conditional on $t-1$. w_t contains real output growth, unemployment rate, inflation (GDP deflator), real federal government spending, real non-residential investments, and real corporate profits net of taxes. Top panel: robust standard errors in parentheses. Middle panel: Wald test statistics for joint significance of SPF forecasts and associated p-value in square brackets.

denotes SPF forecasts for the current quarter and the time index in $\mathbb{E}_t$ refers to the publication date of the survey. Because of the publication 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 $\mathbb{E}_t[w_{t+h}]$ captures pre-existing beliefs about the macroeconomic outlook.[9] Regressions include a constant and a regulatory dummy intended to capture the legal changes discussed above. We report standard errors for the autoregressive coefficients, and Wald test statistics for the joint significance of SPF forecasts at each horizon. Standard errors are HAC-corrected.

Patent applications exhibit a strong autocorrelation pattern. Moreover, as antici-

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

pated, there is evidence that current and expected economic conditions can influence the decision of filing patent applications in any given quarter. Patent applications also correlate with the first (lagged) factor extracted from the large collection of US macroeconomic and financial data assembled in McCracken and Ng (2015).[10] Typically, the first such factor is interpreted as a measure of economic activity. This too reinforces the evidence on the fundamentally cyclical nature of patent applications.

2.2 Instrument Construction

We recover an external instrument for the identification of technology news shocks as the component of quarterly utility patent applications that is orthogonal to agents' pre-existing beliefs about the state of the economy, and is unpredictable given its own history. Intuitively, we seek to remove endogenous variation in applications filings that results from anticipation of current or future economic conditions.

Specifically, we recover an instrument for identification of news shocks using the residuals of the following regression, estimated at quarterly frequency

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

pa_t is the quarterly growth rate of all utility patent applications in a given quarter t, i.e. $pa_t = 100 \times (\ln PA_t - \ln PA_{t-1})$. $\gamma(L) = \sum_{j=1}^{4} \gamma_j L^j$, and $\mathbb{E}_t[x_{t+h}]$ is an $m \times 1$ vector of forecasts compiled at t for the vector of economic variables x_{t+h}, where h is equal to one and four quarters. Again, we use median SPF forecasts conditional on the previous quarter to capture expectations about the state of the economy that pre-date the application filing. The vector x_t contains 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). $\forall t,\ x_t \subset w_t$ used in Table 1 in the previous subsection.

The procedure in Eq. (1) removes both the autocorrelation and the dependence on pre-existing beliefs about macroeconomic conditions as captured by the survey forecasts by construction. In Tables 2 and 3 we check for correlation of the recovered instrument both with other forecasts at different horizons (i.e. $\mathbb{E}_t[w_{t+h}]$), and with the same factors

[10] Results are reported in Table D.1 in the Appendix.

Table 2: Dependence of Instrument on Economic Forecasts

	$\mathbb{E}_t[w_t]$	$\mathbb{E}_t[w_{t+1}]$	$\mathbb{E}_t[w_{t+4}]$
Wald Test	0.270	0.850	0.290
p-value	0.949	0.531	0.882
Adj R^2	0.582	0.587	0.583
N	127	127	127

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

of Table D.1. In both cases, we do not find evidence against the null of no correlation (i.e. the null that the instrument is Granger caused by the variables in the tables).

Table 3: Dependence of Instrument on Lagged States

	F_1	F_2	F_3	F_4	F_5	F_6	F_7
Wald Test	0.880	1.160	0.290	0.810	1.040	0.190	0.290
p-value	0.481	0.330	0.882	0.521	0.389	0.945	0.885
Adj R^2	0.583	0.584	0.583	0.587	0.596	0.580	0.582
N	127	127	127	127	127	127	127

Notes: Dependent variable is the residual of Eq. (1). F_t are factors extracted from the quarterly dataset of McCracken and Ng (2015). Numbers reported are Wald test statistics for the joint significance of the first 4 lags of each factor. All the regressions include own 4 lags, regulation dummy and constant.

A final concern may relate to the potential correlation of patent application filings with other contemporaneous shocks occurring in the current quarter. In order to account for this, we augment Eq. (1) with a set of further controls intended to capture policy changes in the current quarter, as follows

$$pa_t = c + \gamma(L)pa_t + \sum_{h=1,4} \beta_h \mathbb{E}_t[x_{t+h}] + \sum_{j=0}^{2} \delta_j \eta_{t-j} + z_t. \qquad (1')$$

η_t in Eq. (1') includes unexpected and anticipated exogenous tax changes occurring in quarter t, as classified by Romer and Romer (2010) and Mertens and Ravn (2012), and the series of unanticipated changes to the intended Fed funds rate target of Romer and

Figure 2: Instrument for News Shocks

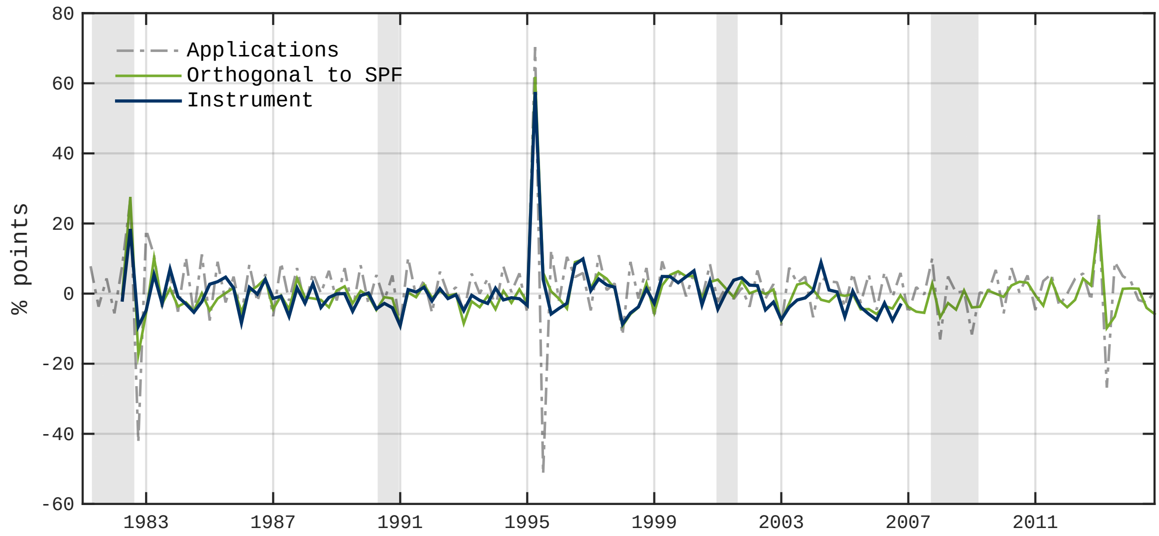

Note: Raw count of patent applications, quarterly growth rate (grey, dash-dotted line); residuals of Eq. (1), (green, solid); instrument for news shocks (blue, solid), residuals of Eq. (1'). Shaded areas denote NBER recession episodes.

Romer (2004).[11]

The variables pa_t and z_t are plotted in Figure 2. The grey dash-dotted line is the quarterly growth rate of patent applications pa_t. The green solid line shows the residuals of Eq. (1) where there is no control for current and lagged policy changes. The solid blue line depicts the residuals of Eq. (1'). Due to the availability of the narrative tax series, the latter is available only up to 2006-IV. We use this as our preferred instrument. Results obtained when not controlling for contemporaneous policy changes are largely equivalent to our benchmark and discussed in Section 4.

[11]We use a series of narrative changes in monetary policy extended 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).

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

$$\ln A_t = \Delta \ln A + \ln A_{t-1} + e_{\mathrm{A1},t} + e_{\mathrm{A2},t-k} \ , \tag{2}$$

where $\Delta \ln A$ is the steady state growth rate of technology, and $e_{\mathrm{A1},t}$ and $e_{\mathrm{A2},t-k}$ are zero-mean normally distributed i.i.d. processes with variance equal to σ^2_{A1} and σ^2_{A2} respectively. A_t is typically understood as a shifter to the aggregate production function of the economy, and intended to capture a concept of technology related to the efficiency with which the factors of production are utilized, or the introduction of new processes altogether.

$e_{\mathrm{A2},t}$ is the news shock.[12] The standard identifying assumption in the news literature is that agents learn about $e_{\mathrm{A2},t-k}$ before it hits the technology process, i.e. $k > 0$ (see e.g. Beaudry and Portier, 2006; Barsky and Sims, 2011, among many others). However, a number of more recent papers have argued that a news shock is also in principle compatible with $k = 0$, which would affect technology also on impact (see e.g. Barsky et al., 2015; Kurmann and Sims, 2017). 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. We remain agnostic, hence, empirically, we do not constrain news shocks to be orthogonal to the current level of technology. Allowing for $k = 0$ naturally makes the task of telling apart a news shock with effects on current technology from an innovation in current technology ($e_{\mathrm{A1},t}$) a daunting one. In this respect, we rely on the information content of the instrument constructed in Section 2. As discussed, while patent applications are most informative for 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 external instrument does not warrant imposing orthogonality with respect to the current level of technology. However, as we shall see in Section 4, while no assumption on the impact

[12] An alternative equivalent formalization 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, 2017).

responses is made, the instrument recovers a shock which leads to an effectively muted response of total factor productivity (TFP) upon realization, while eliciting a strong and sustained response at further ahead horizons. This gives us some confidence that the recovered shock has a large element of news embedded in it.

We use our patent-based external instrument to back out the dynamic causal effects of technology news shocks on a collection of macroeconomic and financial variables in a structural Vector Autoregression (SVAR-IV, Mertens and Ravn, 2013; Stock and Watson, 2012, 2018). Let y_t denote the n-dimensional vector of economic variables of interest, whose dynamics follow a VAR(p)

$$\Phi(L) y_t = u_t, \qquad u_t \sim \mathcal{WN}(0, \Sigma), \tag{3}$$

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

For our purpose, we require that the information in the VAR be sufficient for the identification of $e_{\text{A2},t}$. Specifically, we assume that there exists a $1 \times n$ vector λ such that

$$e_{\text{A2},t} = \lambda u_t, \tag{4}$$

or, in other words, that there exists a suitable rotation of the VAR innovations that reveals the shock of interest $e_{\text{A2},t}$. Forni et al. (2019) and Miranda-Agrippino and Ricco (2018) show that conditional on a correct identification scheme and a VAR specification that correctly captures the dynamics of y_t, the estimated IRFs to the shock of interest converge to the 'true' ones, provided that the VAR is partially invertible in the shocks of interest, i.e. that Eq. (4) holds. Miranda-Agrippino and Ricco (2018) discuss in particular the conditions required for identification with external instruments in SVAR-IVs under partial invertibility. Let z_t denote the external instrument used for the identification of $e_{\text{A2},t}$, and let $^\perp$ denote the component of a process that is orthogonal to the space spanned by the lags of y_t up to $t-1$, such that $z_t^\perp = z_t - \text{Proj}(z_t | \mathcal{H}_{t-1}^{y})$. Recall here that results in Table 3 show that our patent-based instrument is uncorrelated with lagged state variables, and hence with lagged y_t. The required conditions are:

$$\mathbb{E}[e_{\text{A2},t} z_t^{\perp}] = \rho, \ \rho \neq 0 \qquad (\textit{Relevance}) \tag{5}$$

$$\mathbb{E}[e_{i,t} z_t^{\perp}] = 0, \ \forall i \neq \text{A2} \qquad (\textit{Contemporaneous Exogeneity}) \tag{6}$$

$$\mathbb{E}[e_{i,t+\tau} z_t^{\perp}] = 0, \ \forall i \neq \text{A2}, \tau \neq 0 : \mathbb{E}[e_{i,t+\tau} u_t'] \neq 0. \qquad (\textit{Limited Lead-Lag Exogeneity}) \tag{7}$$

Under these conditions, the impact responses to $e_{\text{A2},t}$ of all variables in y_t are consistently estimated (up to scale and sign) from the projection of the VAR innovations $\hat{u}_t$ onto the instrument z_t (Mertens and Ravn, 2013; Stock and Watson, 2012, 2018).[13] The first two conditions are the standard conditions for instruments' validity in IV identification. The third condition arises because of the dynamics, and essentially requires that the instrument and the VAR innovations are only related via contemporaneous realizations of the shock of interest, therefore allowing the instrument to be potentially contaminated by leads or lags of other shocks, so long as these are 'filtered out' by the VAR (i.e. those for which $\mathbb{E}[e_{i,t+\tau} u_t'] = 0$). Hence, with a potentially imperfect instrument –i.e. likely to fail the limited lead-lag exogeneity condition–, these conditions call for the use of information-rich VARs which make Eqs. (4) and (7) more plausible.[14] Furthermore, Forni et al. (2019) show that if the VAR is informationally sufficient for $e_{\text{A2},t}$ but not for the other shocks, then estimates of the forecast error variance contribution of $e_{\text{A2},t}$ are distorted.

4 Results: News Shocks and Business Cycle

We study the transmission and importance of technology news shocks in a 16-variable quarterly VAR that includes a rich and heterogeneous set of variables intended to both cover the relevant variables of interest, and capture possible anticipation of future events that is at the core of the transmission mechanism of news shocks. A complete description of our data is reported in Appendix A. Variables enter the VAR in log levels, with

[13] We compare our identification with the prominent ones in the literature proposed by Beaudry and Portier (2006) and Barsky and Sims (2011) in an illustrative VAR in Appendix C.

[14] Cascaldi-Garcia (2018) suggests the use of growth forecast revisions as instruments for news shocks. While theoretically appealing, the limited availability of sufficiently long-horizon forecasts makes survey-based forecast revisions problematic in relation to conditions (6) and (7).

the exception of interest rates and corporate spreads, and are deflated and expressed in per-capita terms where appropriate. We use the GDP deflator to measure inflation. The VAR is estimated with 4 lags and standard Normal-Inverse Wishart priors (Doan et al., 1983; Litterman, 1986; Kadiyala and Karlsson, 1997). The optimal priors' tightness is estimated as in Giannone et al. (2015). Minor perturbations to the number of lags included do not change the results.[15] The 16-variable VAR(4) is informationally sufficient.[16] We present our empirical results in the form of impulse response functions (IRFs) to a news shock identified using our patent-based external instrument and following the two-step procedure of Mertens and Ravn (2013). 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. For comparison, we report responses to a contemporaneous TFP innovation in Appendix E.

Our benchmark estimation sample is 1971-I:2016-IV, where the start date is constrained by the availability of the Nasdaq Composite stock market index, and by the quarterly series for capacity utilization.[17] Our preferred specification uses the patent-based external instrument that also controls for contemporaneous policy changes, which gives us an identification sample running from 1982:I to 2006-IV.[18] Robustness tests are discussed below and reported in Appendix F.

The IRFs to a positive technology news shock are reported in Figures 3 to 7. We discuss each in turn. These are IRFs at the mode of the posterior distribution of the parameters and are scaled such that the peak response of TFP equals 1%.[19] Shaded

[15]We 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 materially change our results. IRFs are not reported but available upon request.

[16]We use the test for informational sufficiency of Forni and Gambetti (2011) and do not find evidence of any of the lagged state variables Granger causing the VAR residuals. Quarterly factors are extracted from the McCracken and Ng (2015) quarterly FRED-MD dataset. The p-value associated with the null hypothesis of informational sufficiency is 0.98.

[17]We prefer to work with the Nasdaq index since this is more directly linked to developments in the industrial sector than the S&P 500. In fact, the latter also includes financial institutions including investment banks, and other entities such as insurance companies which can act as confounding elements, particularly in light of the financial crisis of late 2007-2008. We discuss results relative to the response of the S&P 500 below.

[18]Figure F.4 in the Appendix compares it with responses obtained without directly controlling for contemporaneous policy changes (i.e. the green line in Figure 2). Results are qualitatively the same, but estimated with a slightly larger degree of uncertainty. Error bands for both specifications are not reported for ease of readability, but available upon request.

[19]Median responses across the draws are virtually the same.

Figure 3: The Slow Diffusion of Technology

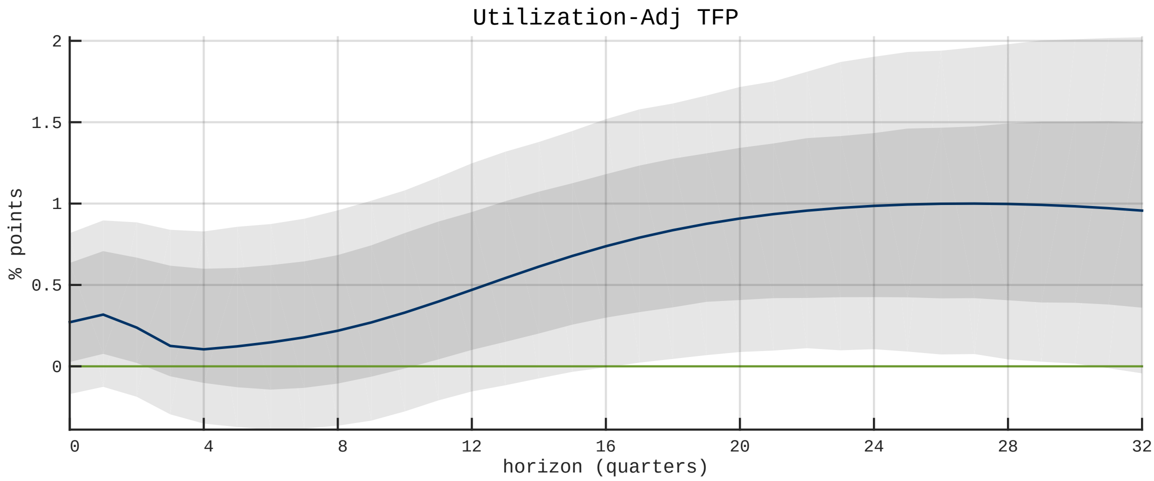

Note: Modal response of Utilization-Adjusted TFP to a technology news shock identified with patent-based external instrument. VAR(4). Estimation sample 1971-I : 2016-IV. Identification sample 1982-I : 2006-IV. Shaded areas denote 68% and 90% posterior coverage bands.

areas correspond to 68% and 90% posterior coverage bands.

Figure 3 plots the response of TFP to the identified technology news shock over a period of 32 quarters. We use the quarterly series of total factor productivity corrected for input utilization of Fernald (2014). TFP rises mildly on impact, then contracts slightly, and finally rises robustly following a persistent hump that reaches a peak between 6 and 7 years after the shock hits. The response is not significant for the first four years. The shape of the TFP response resembles the S-shaped pattern that is typical of the slow diffusion of new technologies documented, among others, in Griliches (1957); Mansfield (1961); Rogers (1962) and Gort and Klepper (1982). Technology diffuses slowly at first. This initial phase is then followed by a fast diffusion period that ends once the new technology has been fully absorbed, and diffusion reaches its maximum. A similarly shaped response is reported in Barsky et al. (2015) and Kurmann and Sims (2017). Both these papers identify technology news shocks based on the forecast error variance of TFP, and do not restrict the impact response of TFP to be zero.[20]

The responses of the variables related to economic activity are reported in Figure 4.

[20]Kurmann and Sims (2017) consider the case in which TFP measures true technology with an error that correlates with economic conditions. Assuming that the measurement error albeit systematic is nevertheless transient, identification based on the long-run forecast error variance of TFP avoids reliance on its short term fluctuations, and is thus robust to such mis-measurements.

Figure 4: Quantities

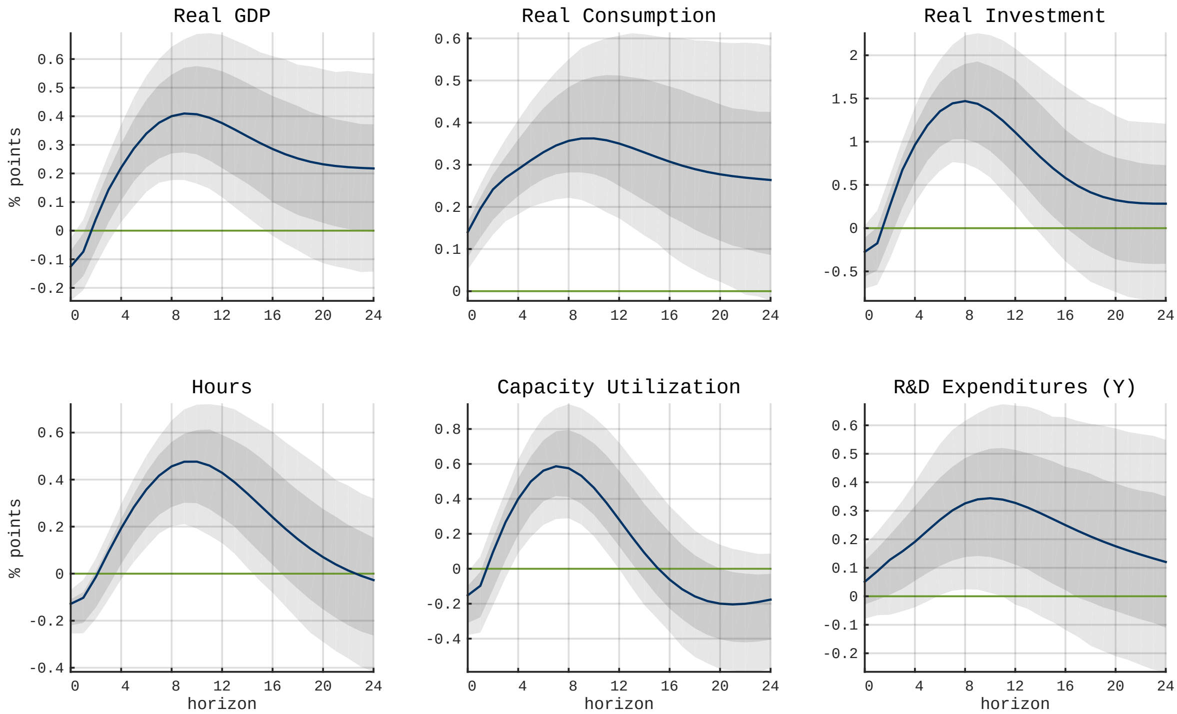

Note: Modal response of quantities to a technology news shock identified with patent-based external instrument. VAR(4). Estimation sample 1971-I : 2016-IV. Identification sample 1982-I : 2006-IV. Shaded areas denote 68% and 90% posterior coverage bands.

Consumption rises immediately following the shock, and remains elevated throughout. Output, investment, and capacity utilization stay mostly put on impact, and then rise persistently to reach a peak after about two years after the shock hits. Impact modal responses are negative, but only marginally significant at conventional levels, and fully reabsorbed in the span of two to three quarters. The magnitude of the responses is economically important. Output reaches almost half a percentage point at peak, while investment increases by 1.5%. The labor market witnesses similarly significant improvements at the two year horizon. Here, however, we note that the initial decline in labor inputs, albeit short-lived, is strongly significant, and more robust to changes to either the sample size or the VAR specification than the other negative impact responses of Figure 4. We explore the source of the contraction in total hours worked more in detail in the next section. R&D expenditures (as a component of output) do not seem to respond to the shock in significant ways. While modal reactions suggest R&D to be somewhat

Figure 5: Shares of Explained Variance

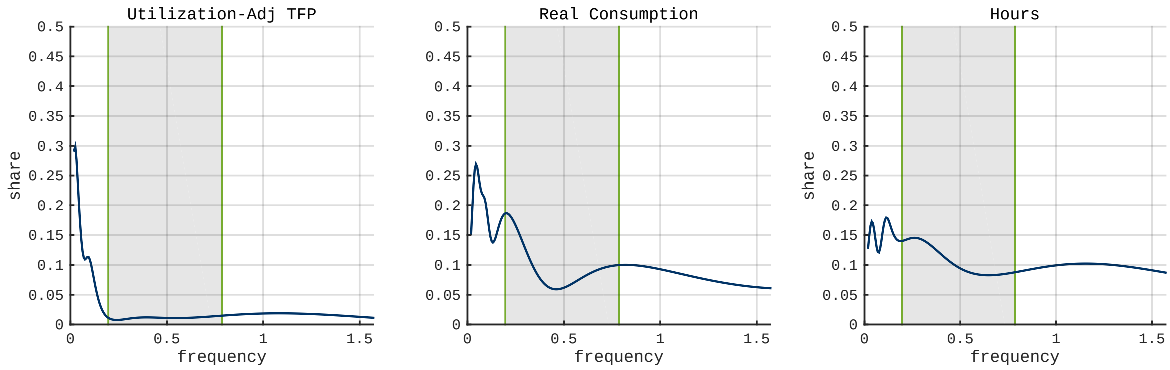

Note: Share of error variance accounted for by technology news shock identified with patent-based external instrument. VAR(4). Estimation sample 1971-I : 2016-IV. Identification sample 1982-I : 2006-IV. Shaded areas delimits business cycle frequencies (between 8 and 32 quarters). Frequencies on the x axis cover a period from 1 (highest) to 100 (lowest) years.

higher following the shock, the response is only significant at the two-year horizon. This could be entirely driven by the rise in output.

Hence, the identified technology news shock can induce comovements among variables that are typical of business cycle fluctuations over medium horizons, but does not seem to do so on impact. In this respect, our findings align with some of the results in e.g. Francis and Ramey (2005); Basu et al. (2006) and Barsky and Sims (2011), although the responses in Figure 4 (and with the notable exception of hours worked) point towards a muted initial response of real activity, rather than a fully recessionary episode. The timing of the responses in Figure 4 does lend credit to a 'news view' in the spirit of what is described in e.g. Beaudry and Portier (2006); Beaudry and Lucke (2010), to the extent that the movements in the quantity variables substantially anticipate the actual increase in TFP. Hence, there seems to be evidence in favor of news triggering business cycle-type fluctuations before any significant change in technology is effectively realized.

The shock, however, is not the main driver of fluctuations in economic variables at business cycle frequencies. Figure 5 plots the share of variance of TFP, consumption, and hours that is accounted for by the identified technology news shock at all frequencies between 1 (highest frequency) and 100 (lowest frequency) years.[21] The algorithm used

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

for the decomposition is described in detail in Appendix B.[22] The advantage of looking at variance decompositions in the frequency domain is that it allows us to separate long, medium, and short-run fluctuations in the variables in our VAR; importantly, this allows us to isolate the contribution of news shocks to business cycle fluctuations more clearly than with a standard forecast error variance decomposition in the time domain.[23] The identified shock explains at most a third of the variation of TFP in the very long run (100 years). Table 4 reports the shares of explained variation at selected frequency intervals for all variables. The recovered news shock is responsible for virtually none of the variation in TFP either in the short-run (i.e. area under the curve in rightmost section of the left panel of Figure 5, corresponding to a period of 1 to 2 years), or at business cycle frequencies (2 to 8 years), and accounts for about 10% of its variation in the long-run (8 to 25 years, see Table 4). At the same time, it is responsible for about 15% of the fluctuations in both consumption and hours at business cycle frequencies, and accounts for over a fifth of the variation in consumption, and about 15% of that in labor inputs in the long-run. These shares are sizeable and economically relevant, but far from capturing the bulk of variation in these variables.

The responses of prices are reported in Figure 6. Similarly to what is found in Barsky and Sims (2011); Kurmann and Otrok (2013) and Barsky et al. (2015), we find that technology news shocks are disinflationary. Importantly, however, while these authors document a sudden and persistent drop, we unveil a rather sluggish response of prices upon realization of the shock. The GDP deflator contracts only marginally on impact, but keeps sliding over the subsequent quarters, reaching a peak response of about -0.3% at the two year horizon, consistent with a sluggish adjustment of prices over time. A similarly sluggish adjustment is characteristic of the relative price of investment goods, that suffers a minor contraction already on impact, but keeps adjusting over time. The

[22]The algorithm builds on Altig et al. (2011). We discuss the contribution of the news shock to fluctuations in the remaining variables in our VAR at the end of this section (see Table 4). Variance decompositions for all variables at all frequencies between 1 and 100 years are in Figure B.1 in the Appendix.

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

FIGURE 6: PRICES & WAGES

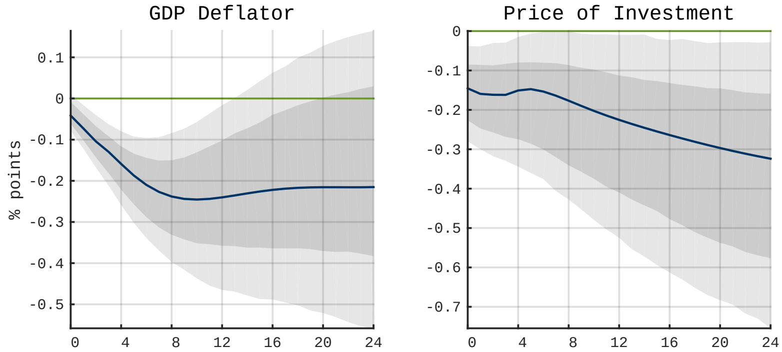

Note: Modal response of price variables to a technology news shock identified with patent-based external instrument. VAR(4). Estimation sample 1971-I : 2016-IV. Identification sample 1982-I : 2006-IV. Shaded areas denote 68% and 90% posterior coverage bands.

response indicates that the identified news shock makes investment goods cheaper relative to consumption goods. Hence, the shock has some of the flavor of the investment-specific technology improvements of e.g. Fisher (2006) and Justiniano et al. (2010, 2011).[24] Figure 6 also reports the response of real wages. We find that wages significantly contract on impact, to increase at longer horizons. We discuss the response of real wages in greater detail in the next section.

Lastly, we collect responses of asset prices and measures of consumers' expectations in Figure 7. The stock market prices in the news shock strongly and significantly on impact – the Nasdaq index jumps up by 5% upon realization of the shock.[25] The strong response of the stock market is more notable when the Nasdaq is used, due to the index composition being heavily weighted towards information-technology companies. These are presumably those mostly affected by these types of shocks over the identification sample considered (1982-I:2006-IV). Figure F.2 in the appendix compares IRFs in our benchmark sample with those obtained when estimating the VAR from 1962-I, and substituting the Nasdaq Composite with the S&P 500 (same identification sample). The response of the S&P is positive on impact, but the magnitude is about a third of that of the Nasdaq.[26]

[24] See also Ben Zeev and Khan (2015).

[25] Bretscher et al. (2019) use a New Keynesian DSGE model to study the implications of news shocks for asset pricing, and find that macroeconomic risk factors that derive from agents' accounting of news also help price the cross-section of expected returns.

[26] In this case we drop the capacity utilization variable which is unavailable prior to the 1970s, and substitute the Nasdaq with the S&P 500. The start date coincides with the availability of daily data for interest rates (DGS1 and DGS10) that enter the VAR in quarterly averages. We note that in this case

Figure 7: Expectations & Financial Markets

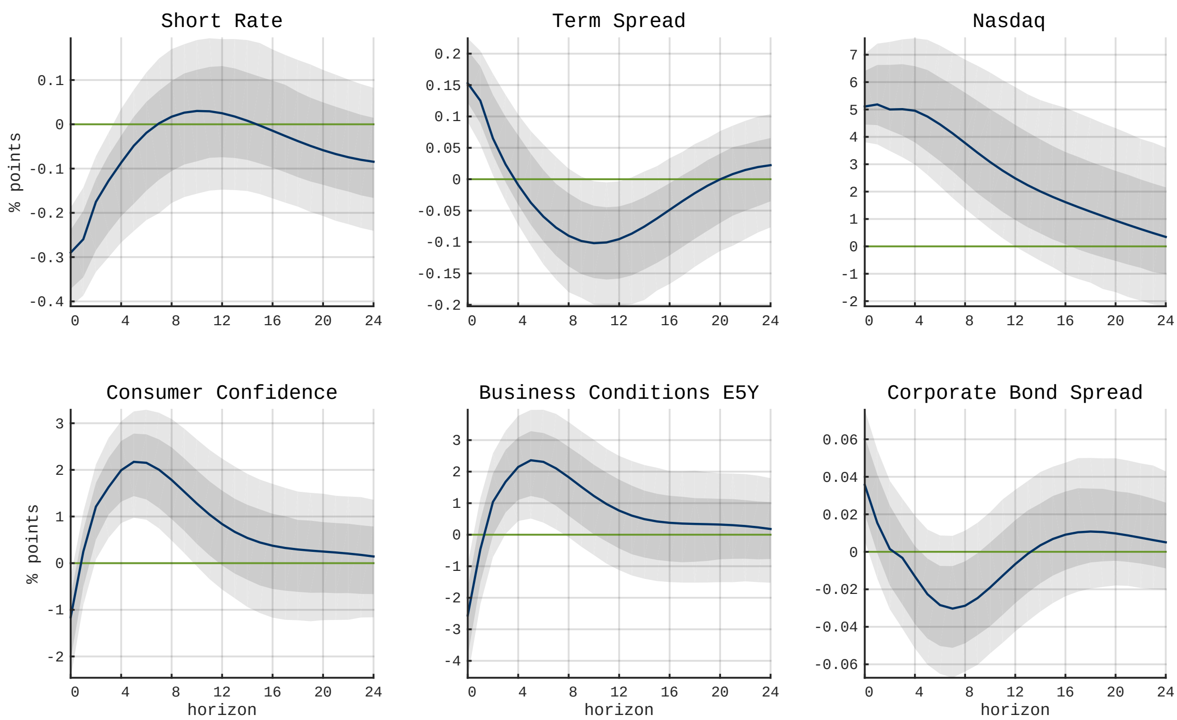

Note: Modal response of consumers' expectations and financial markets to a technology news shock identified with patent-based external instrument. VAR(4). Estimation sample 1971-I : 2016-IV. Identification sample 1982-I : 2006-IV. Shaded areas denote 68% and 90% posterior coverage bands.

Consistent with a world in which companies only slowly adjust to the introduction of new technologies, the BAA-AAA corporate bond spread slightly increases on impact, to improve at medium horizons. This response is however not particularly significant.

The significant disinflationary characteristic of the identified news shock induces an endogenous strong response of the monetary authority, that responds more than proportionally to the decline in inflation. Due to the sample considered including the zero-lower-bound (ZLB) period, we use the one year nominal interest rate as our measure for the short term policy rate. In Figure F.1 we verify that neither the global financial crisis nor the ZLB sample drive or affect our results. The one year rate falls by about 30 basis points on impact, which is roughly the same magnitude as the peak decline of prices (see Figure 6). This implies that shorter maturity interest rates are likely to fall by more, and hence that short-term real interest rates fall following the shock. The slope of the

the magnitude of the peak responses of both prices and interest rates is larger.

TABLE 4: ERROR VARIANCE DECOMPOSITION

		SHORT RUN [4 - 8 quarters]	BUSINESS CYCLE [8 - 32 quarters]	LONG RUN [32 - 100 quarters]
TFPL	Utilization-Adj TFP	1.68	1.03	9.57
RGDP	Real GDP	8.39	10.68	14.04
RCONS	Real Consumption	8.66	13.28	20.39
RINV	Real Investment	6.64	11.78	10.94
RDGDP	R&D Expenditures (Y)	0.56	3.85	7.27
HOURS	Hours	9.58	13.39	14.32
CAPUTIL	Capacity Utilization	5.55	10.59	13.65
GDPDEF	GDP Deflator	2.23	7.24	12.87
RPINV	Price of Investment	3.46	2.33	6.02
RWAGE	Real Wages	8.40	4.59	11.89
SHORTR	Short Rate	15.39	9.87	1.88
YCSLOPE	Term Spread	12.93	10.56	5.38
EQY2	Nasdaq	22.57	20.46	21.11
CCONF	Consumer Confidence	7.34	11.88	12.66
BCE5Y	Business Conditions E5Y	8.47	8.04	8.77
CBSPREAD	Corporate Bond Spread	2.05	4.58	1.55

Notes: Share of error variance accounted for by the identified technology news shock over different frequency intervals. Numbers are percentage points.

yield curve, here measured as the spread between the 10-year and the 1-year Treasury rates, rises by about 15 bps on impact, mainly driven by changes at the short end, and implying a 15 bps fall in long term yields. Similar types of impact responses are reported in Kurmann and Otrok (2013), where the identified news shock is also responsible for most of the unexplained variation in the slope of the term structure of interest rate. We do not find this to be the case. Table 4 shows that the shock is most explanatory over the short-run, where it can account for about 15% of movements in the term structure, but it captures very little variation in interest rates in the long run. The impact response of the short term rate also contrasts with findings in Kurmann and Sims (2017), where the response of the monetary authority is mildly contractionary.[27]

Finally, Figure 7 reports responses of a consumer confidence indicator and a business confidence indicator reflecting expectations about economic conditions over a horizon of

[27]For a broader discussion on the role played by different vintages of TFP data on the response of the term structure slope to technology news shocks see Cascaldi-Garcia (2017); Kurmann and Otrok (2017).

5 years, both taken from the Michigan Survey of Consumers. Interestingly, we find that while both measures of confidence robustly rise at medium horizons, they do not do so on impact. In fact, the responses tend to be negative upon realization of the shock. This finding is consistent with agents overweighting the responses of current economic conditions discussed in Figure 4 when forming their expectations about the future, and echoes the implications of models in which agents are subject to strong informational rigidities. We return to this issue in greater detail in the next section.

5 Discussion of the Results: the Propagation of News about Future Technology

Equipped with the empirical results reported in Section 4, in this section we try to shed some light on the likely transmission mechanisms by evaluating our findings against the different models proposed in the literature.

Total Factor Productivity As noted, the impulse response function of TFP to the identified news shock supports the hypothesis of slow diffusion of technology over time (see e.g. Rotemberg, 2003, and references therein).[28] While there is evidence of some (non-significant) positive spillover to current TFP, productivity does not materially move away from zero before the first 4 years after the shock hits. Hence, the effect of a news shock on current TFP is estimated to be effectively zero, even if we have not imposed such restriction ex ante. Moreover, by the time the TFP response becomes positive, and perhaps with the exception of real wages and the relative price of investment, all the other variables in the VAR have reached the peak of their dynamic adjustment. This large asynchronicity in the timing of the responses favors the hypothesis that macro aggregates can in fact move as a result of a change in expectations about future productivity growth, and before the change in aggregate technology materializes. The ensuing business cycle

[28] Among others, Rogers (1962); David (1990), and Hall (2006), have rejected the RBC view and have produced evidence that suggests a slow S-shaped diffusion of technology. While the implications of the slow diffusion for the modeling of productivity is discussed extensively in e.g. Rotemberg (2003); Comin and Gertler (2006); Lindé (2009), much of the business cycle literature has modeled productivity as a jump process. Other papers that build models of costly adoption of new technologies that are consistent with a slow diffusion pattern are e.g. Comin et al. (2009) and Comin and Hobijn (2010).

expansion is not estimated to be immediate, and we return to this point below. Hence, while it appears that business-cycle types of comovements can in fact be triggered, here we note that the relatively small share of explained variance that is accounted for by the identified news shock at business cycle frequencies casts substantial doubts on it being a main driver of economic fluctuations.[29] Importantly, this also holds true for the TFP process itself. Our estimates suggest that technology news account at best for a third of the variance of TFP at very low frequencies (see Figure 5).

Quantities: Output, Consumption, Investment, and Hours In our VAR output, investment, consumption and hours worked are all significantly higher a few quarters after the shock hits, with peak effects realized in the span of two years. On impact, consumption rises strongly, hours decline, and although modal responses are negative, investment and output do not meaningfully move away from zero before they start increasing. Capacity utilization also rises after staying still on impact. These types of responses are hard to rationalize under the standard neoclassical real business cycle (RBC) paradigm. The rise in consumption is understood to be the result of a wealth effect: expectations of future higher productivity raise expectations about future income, which in turn induce households to smooth consumption towards higher current levels. The same wealth effect also increases the desire for leisure, while higher expected future productivity redirects the capital stock away from investment and towards consumption until the higher productivity level is realized. Hence, consumption, labor effort and investment must in this case move in opposite directions (Barro and King, 1984; Cochrane, 1994). Moreover, in the classical RBC setting, a fixed labor demand implies that the fall in hours worked must come from a shift in the labor supply curve, which in turn requires an increase in wages. This too contrasts with our findings: real wages significantly contract upon realization of the shock, and only slowly increase over time.

We interpret the delayed business cycle expansion that is triggered by the news shock as indicative of the presence of potentially different sources of inertia that delay the

[29]Similar conclusions have been reached in a DSGE framework in e.g. Fujiwara et al. (2011); Schmitt-Grohé and Uribe (2012) and Khan and Tsoukalas (2012). Sims (2016) argues that earlier empirical works may be confounding current and past news shocks, hence implying a potentially systematic overstatement of the relative importance of news shocks.

adjustments. In fact, the responses of quantities documented here is consistent with New Keynesian models with nominal rigidities that influence the setting of prices, wages, or both (e.g. Barsky and Sims, 2009; Christiano et al., 2010; Barsky et al., 2015), and with RBC models augmented with real rigidities such as e.g. habit formation in consumption, and adjustment costs associated with changes in either the stock of capital or the rate of investment, and equipped with a system of preferences that allows to fine-tune the wealth elasticity of labor supply (e.g. Jaimovich and Rebelo, 2009; Schmitt-Grohé and Uribe, 2012). A weakened short-run wealth effect can in fact induce a right shift in the labor supply. At the same time, the presence of adjustment costs and variable capital utilization can induce positive shifts in labor demand if the price of capital decreases as a consequence of the shock. However, while these types of mechanisms can account for the positive comovements, they cannot reproduce other important effects, such as e.g. the increase in asset prices. In fact, these models predict that asset prices will move with the cost of capital, and will hence decrease (see e.g. Christiano et al., 2010).[30]

Prices: Inflation and Wages New Keynesian models with nominal rigidities, including those where such frictions arise endogenously due to imperfect common knowledge (e.g. Mankiw and Reis, 2002; Woodford, 2003) seem to offer a more varied array of mechanisms through which our findings can be rationalized. This is because they allow the dynamics to be dominated by the demand side, at least in the short-run (see discussion in e.g. Lorenzoni, 2009, 2011). In the VAR, the shock triggers a sudden and marked contraction of real wages followed by a slow, but significant deflationary episode. Prices drop mildly on impact, and continue to slide over time to reach a peak contraction two years after the shock. Real wages eventually increase; the time taken for the wage inflationary pressure to materialize goes from 8 to 16 quarters depending on the chosen significance level. The deflationary effect of news shocks is a robust finding, and has been documented in Christiano et al. (2010); Jinnai (2013); Kurmann and Otrok (2014) and Barsky et al. (2015) among others. However, contrary to findings in e.g. Barsky and Sims (2009, 2012) and Kurmann and Otrok (2014), we find that the bulk of the drop in inflation is not realized on impact. Rather, and consistent with nominal rigidities pre-

[30]Business cycle comovements are also reproduced in standard RBC frameworks augmented with dispersed information (see e.g. Angeletos and La'O, 2010).

venting an immediate impact adjustment, the response of prices is subdued initially, and only slowly builds up over time. Christiano et al. (2010) and Barsky and Sims (2011) interpret the fall in prices as a manifestation of the forward-looking nature of inflation in the New Keynesian model, where current inflation is a function of both current and future expected marginal costs. As also discussed in Barsky et al. (2015), expected future productivity improvements lower expected real marginal costs. If real wages do not rise too sharply, the expectation that marginal costs will be lower in the future creates downward pressure on current inflation. Whether this happens in practice depends on the persistence of the news process, the monetary policy rule, and the potency of nominal rigidities. For a given news process, and leaving temporarily aside the role of the monetary authority, the fall in inflation following the news shocks can be obtained under two different specifications of nominal frictions: a case of pure sticky prices as in e.g. Calvo (1983), and one in which prices are flexible, but wages are staggered like in e.g. Erceg et al. (2000). Christiano et al. (2010) show that while both scenarios give rise to a deflation, the range of parameters across which this happens in a sticky wage environment is larger (see also Barsky and Sims, 2009; Jinnai, 2013).

Monetary Policy, the Natural Rate of Interest, and the Term Premium Expectations that productivity will be higher in the future, but that do not change the level of current technology, give rise to an inefficient rise in current spending, primarily driven by the desire to increase current consumption. In order to keep spending anchored to the current (unchanged) level of technology, the natural rate of interest, proportional to the expected growth rate of technology, rises sharply. Consider now a central bank that sets the nominal interest rate as a function of expected inflation. This is the situation analyzed in detail in Christiano et al. (2010). Expectations that inflation will be lower in the future lead the central bank to lower the nominal interest rate precisely when the natural rate is increasing, thus creating an amplification mechanism for the propagation of the news shock.[31] The dynamic responses from our VAR abide by this narrative. As discussed, the one year rate moves by roughly the same amount as the deflator at peak, implying an even larger drop of shorter maturity interest rates. Hence, in our empiri-

[31] See also discussion in Sims (2012).

Figure 8: Long Rate Response

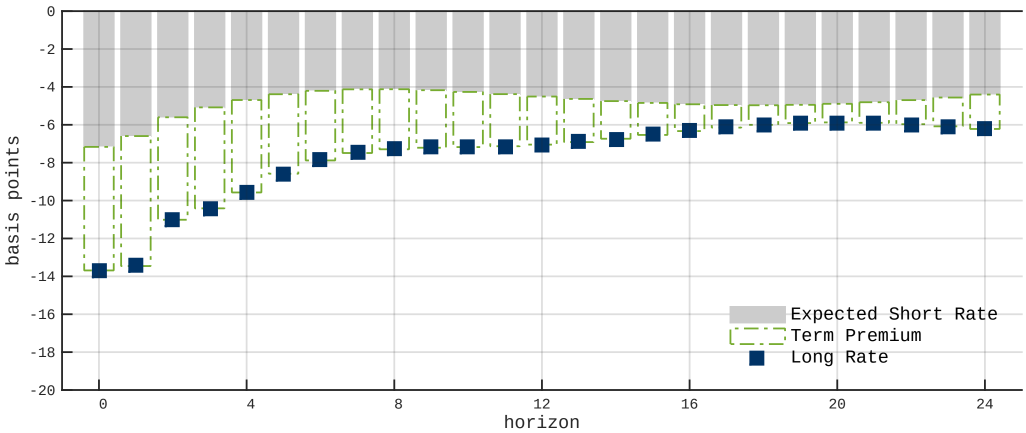

Note: Implied modal responses of the 10-year Treasury yield and VAR-based expectation and term premium components. VAR(4). Estimation sample 1971-I:2016-IV; Identification sample 1981-1:2006-IV.

cal setting the monetary authority responds to the news shock by aggressively reacting to the fall in (expected) inflation. The suboptimal response of the central bank can be rationalized in terms of information rigidities: the monetary authority may have to calibrate its response based on its best forecasts for (current and future) technology, which may diverge from the realized values (see e.g. Lorenzoni, 2011).[32] Finally, comparing the responses of the short and long term rates, we note that the 1-year rate returns to trend relatively quickly, and is hence likely not to fully account for the impact fall in the 10-year Treasury yield. This implies that following the news shock term premia decline.

We confirm this intuition in Figure 8. Here we plot the responses of the long term rate implied by Figure 7, and use the VAR to decompose it into its expectation and term premium components.[33] About 3/5 of the impact decline in the long term interest rate is estimated to be due to a fall in term premia; and the response dies out relatively slowly.

[32]As a partial solution to this issue, Christiano et al. (2010) suggest introducing variables that help to proxy for the natural rate, such as e.g. credit growth, in the reaction function.

[33]The 10-year yield can be decomposed into the expected 1-year rate over 10 years, plus a term premium ζ_t. If t denotes quarters,

$$y_t^{(10)} = \mathbb{E}_t\left[\frac{1}{10}\sum_{\tau=1}^{10} y_{t+4\times(\tau-1)}^{(1)}\right] + \zeta_t^{(10)}. \tag{8}$$

Net of risk considerations, holding a 10-year bond should be equivalent to rolling 1-year bonds over 10 years. We calculate horizon h term premium responses as the difference between the horizon h response of the 10-year rate, and the average expected response of the 1-year rate at horizons $h, h+4, \ldots, h+36$.

These findings align with those in Crump et al. (2016).

Consumer Expectations We lastly turn to analyze the response of consumer expectations. As discussed, upon realization of the shock both the index of consumer confidence and the component of the Michigan Survey of Consumers that relates to business conditions expected 5 years hence decline. The decline is short-lived, and both indicators robustly rise above trend within a year after the shock hits. While the negative responses are only marginally significant at conventional levels, they are nevertheless somewhat puzzling. We offer an interpretation for this finding that builds on the presence of information rigidities. In fact, a potential explanation for this result is that agents only imperfectly observe future fundamentals, and overweigh current economic conditions when forming their expectations about the future when the signal-to-noise ratio is low.

In a comprehensive study, Coibion and Gorodnichenko (2012, 2015) analyze survey forecasts of consumers, firms, professional forecasters and central banks, and find that economic agents face strong information rigidities, irrespective of their type. The empirical regularities unveiled in these works describe frictions to information processing that seem to be more coherent with frameworks in which agents continuously update their information set, but only receive noisy signals about the state of the economy (noisy information, Woodford, 2003; Sims, 2003; Mackowiak and Wiederholt, 2009), as opposed to alternatives in which the update itself is sluggish (sticky information, Mankiw and Reis, 2002). In the noisy information environment, agents never fully observe the true states, and form expectations about fundamentals via a signal extraction problem. Hence, at any given time agents' forecasts are a combination of existing beliefs and new information received, with relative weights determined by the degree of information rigidity (i.e. noise in the signal). Coibion and Gorodnichenko (2015) estimate that new information receives less than half the weight it would otherwise have under full-information. News about future technological changes can be thought of as a quintessential signal extraction problem (see also Chahrour and Jurado, 2018). Blanchard et al. (2013) consider the case in which technology is driven by both temporary and permanent shocks (i.e. shocks that have long-lasting effects on the level of technology), and agents observe a noisy signal of the permanent component of technology. Agents are not able to disentangle news from

Figure 9: Role of Unemployment and Unemployment Expectations

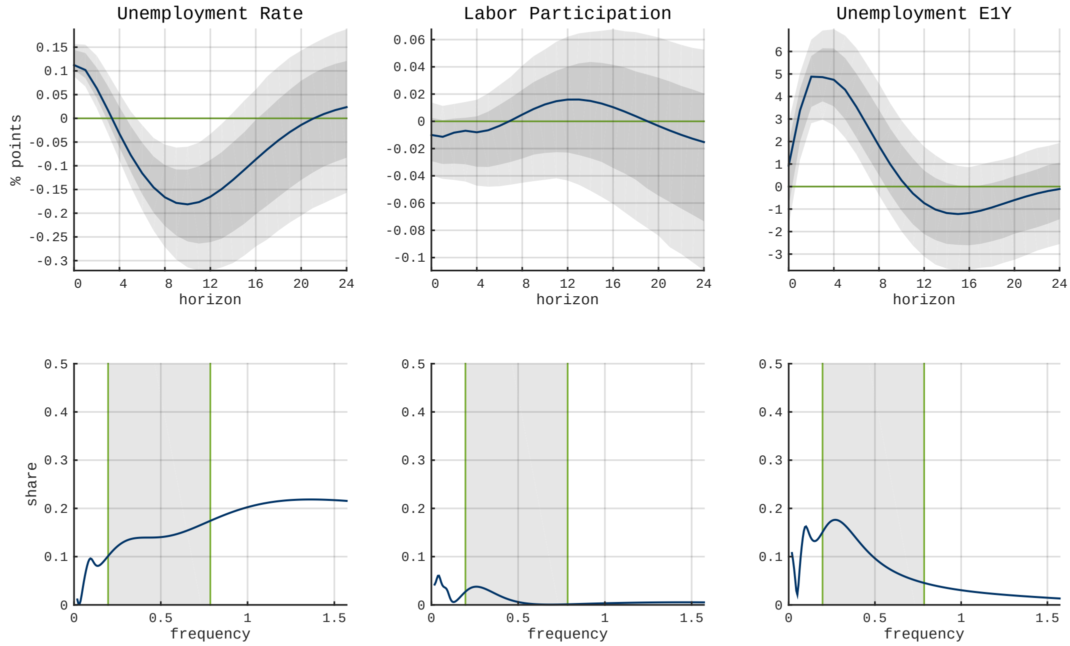

Note: Impulse response functions (top panels) and shares of explained variance (bottom panels) for the unemployment rate, the rate of labor participation, and the 1-year-ahead unemployment expectation. Survey forecasts are from the Michigan Survey of Consumers. VAR(4). Estimation sample 1971-I:2016-IV; Identification sample 1981-1:2006-IV.

noise; moreover, the noisier the signal, the slower the consumption adjustment, the more likely that shocks to the permanent component result in an initial fall in employment.

To offer some additional insights, in Figure 9 we look more in detail at the response of the labor market, and report IRFs and variance shares for the unemployment rate, the rate of labor participation, and consumers' expectations about one year ahead unemployment, again extracted from the Michigan Survey of Consumers.[34] Following the positive news shock, the unemployment rate rises on impact. Given the muted response of labor participation at all horizons, it seems to be the case that the initial fall in total hours (Figure 4) is essentially the result of an increase in the unemployment rate. Consistently, consumers' expectations about unemployment rise, and do so very significantly. The peak is realized well within the first year, and the shock explains a non trivial fraction of variation of unemployment forecasts at business cycle frequencies and

[34]We augment the VAR of Section 4 with these three variables and remove total hours worked. All other details of the VAR specification stay the same. Full IRFs are in Figure F.3.

about 20% of the short-run fluctuations in the unemployment rate (see also Faccini and Melosi, 2018, for the role played by technology news on employment and its forecasts). We think of the rise in expected unemployment as compatible with such noise-ridden environment, and with agents (consumers) overweighting the negative impact response of labor market variables to the shock. In turn, this can help explain the initial fall in consumer confidence about both current and expected economic conditions. Barsky and Sims (2009, 2012) use innovations in consumer confidence to infer on the effect of news shocks, arguing that measures of confidence aggregate information about future income that is otherwise unavailable in current consumption data, an intuition first offered in Cochrane (1994). The responses in Figure 7 suggest that confidence 'innovations' may in fact be anticipated. Indeed, Figure E.1 in Appendix E shows that consumer confidence robustly rises on impact only following a positive contemporaneous TFP innovation.

6 Conclusions

'How does the aggregate economy react to a shock that raises expectations about future productivity growth?' In this paper we have provided an answer to this question by introducing a novel external instrument for the identification of technology news shocks based on patent counts. Importantly, by controlling for expectations about current and future macroeconomic developments formed prior to the patent filings, as well as for other contemporaneous policy changes, we were able to account for the endogeneity in patent applications, and isolate contemporaneous news. We have evaluated the effects of news shocks on an array of macro aggregates, financial market data, and expectations. Our results are consistent with the predictions of New-Keynesian models with nominal rigidities, particularly those that arise endogenously due to noisy-information environments.

Our main conclusions are as follows. (i) Our IRFs support a 'news view' whereby an economic expansion is realized in anticipation of future technological improvements. The identified news shock has no effect on TFP during the first four years, while all other variables in our VAR reach the peak of their dynamic adjustment within two years. This suggests that the shift in economic aggregates is likely to be predominantly driven by a change in beliefs. (ii) While economically relevant, the shock is not a main driver of business cycles: it accounts, on average, for about a tenth of aggregate fluctuations at business cycle frequencies, and for about a third of the variation of TFP in the very long run. (iii) While the stock market prices-in technology news strongly on impact, consumers expectations take longer to adjust. In fact, the complexity of the signal extraction problem leads consumers to overweigh the initial deterioration in labor market conditions when forming expectations about the future, leading to an initial fall in confidence. Similarly, the central bank reacts to the positive news by easing the monetary stance in response to the fall in expected inflation, potentially acting as an amplifier of its effects through the compression of premia.

The reactions of consumers, market participants and the central bank to the identified news shock seem to point towards a substantial degree of heterogeneity in their expectation formation process. All concur to highlight the role that dispersed information about the future may have in shaping the response of different types of agents to such types of disturbances. We leave further investigation of these relevant issues to future research.

References

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 and Jennifer La'O (2010) "Noisy Business Cycles," *NBER Macroeconomics Annual*, Vol. 24, No. 1, pp. 319–378.

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

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

Barro, Robert and Robert G. King (1984) "Time-Separable Preferences and Intertemporal-Substitution Models of Business Cycles," *The Quarterly Journal of Economics*, Vol. 99, No. 4, pp. 817–839.

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

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

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

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

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

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

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

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

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

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.

Bretscher, Lorenzo, Aytek Malkhozov, and Andrea Tamoni (2019) "News Shocks and Asset Prices," March. SSRN 2367196.

Calvo, Guillermo A. (1983) "Staggered prices in a utility-maximizing framework," *Journal of Monetary Economics*, Vol. 12, No. 3, pp. 383 – 398.

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.

Carley, Michael, Deepak Hedge, and Alan C. Marco (2015) "What is the Probability of Receiving a U.S. Patent?" *Yale Journal of Law and Technology*, Vol. 17, No. 1, pp. 201–223.

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.

——— (2018) "Forecast Revisions as Instruments for News Shocks." FRB, Unpublished.

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

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.

Christiano, Lawrence J., Cosmin Ilut, Roberto Motto, and Massimo Rostagno (2010) "Monetary policy and stock market booms," *Proceedings - Economic Policy Symposium - Jackson Hole*, pp. 85–145.

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

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

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

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

Comin, Diego and Mark Gertler (2006) "Medium-Term Business Cycles," *American Economic Review*, Vol. 96, No. 3, pp. 523–551, June.

Comin, Diego and Bart Hobijn (2010) "An Exploration of Technology Diffusion," *American Economic Review*, Vol. 100, No. 5, pp. 2031–2059, December.

Comin, Diego A., Mark Gertler, and Ana Maria Santacreu (2009) "Technology Innovation and Diffusion as Sources of Output and Asset Price Fluctuations," NBER Working Papers 15029, National Bureau of Economic Research, Inc.

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.

David, Paul A. (1990) "The Dynamo and the Computer: An Historical Perspective on the Modern Productivity Paradox," *The American Economic Review*, Vol. 80, No. 2, pp. 355–361.

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

Erceg, Christopher J., Dale W. Henderson, and Andrew T. Levin (2000) "Optimal monetary policy with staggered wage and price contracts," *Journal of Monetary Economics*, Vol. 46, No. 2, pp. 281–313, October.

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

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

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

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.

Fujiwara, Ippei, Yasuo Hirose, and Mototsugu Shintani (2011) "Can News Be a Major Source of Aggregate Fluctuations? A Bayesian DSGE Approach," *Journal of Money, Credit and Banking*, Vol. 43, No. 1, pp. 1–29.

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, 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 (1957) "Hybrid Corn: An Exploration in the Economics of Technological Change," *Econometrica*, Vol. 25, No. 4, pp. 501–522.

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

Hall, Bronwyn H. (2006) "Innovation and Diffusion," in Jan Fagerberg and David C. Mowery eds. *The Oxford Handbook of Innovation*: Oxford University Press.

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.

Jaimovich, Nir and Sergio Rebelo (2009) "Can News about the Future Drive the Business Cycle?," *American Economic Review*, Vol. 99, No. 4, pp. 1097–1118, September.

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

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

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

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

Khan, Hashmat and John Tsoukalas (2012) "The Quantitative Importance of News Shocks in Estimated DSGE Models," *Journal of Money, Credit and Banking*, Vol. 44, No. 8, pp. 1535–1561, December.

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.

——— (2014) "News Shocks and Inflation: Lessons for New Keynesians." Unpublished.

——— (2017) "News Shocks and the Slope of the Term Structure of Interest Rates: Reply," *American Economic Review*, Vol. 107, No. 10, pp. 3250–56, October.

Kurmann, André and Eric R. Sims (2017) "Revisions in Utilization-Adjusted TFP and Robust Identification of News Shocks," Working Paper 23142, National Bureau of Economic Research.

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

Lindé, Jesper (2009) "The effects of permanent technology shocks on hours: Can the RBC-model fit the VAR evidence?" *Journal of Economic Dynamics and Control*, Vol. 33, No. 3, pp. 597 – 613.

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

Lorenzoni, Guido (2009) "A Theory of Demand Shocks," *American Economic Review*, Vol. 99, No. 5, pp. 2050–84, December.

——— (2011) "News and Aggregate Demand Shocks," *Annual Review of Economics*, Vol. 3, No. 1, pp. 537–557, September.

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

Mankiw, N. Gregory and Ricardo Reis (2002) "Sticky Information Versus Sticky Prices: A Proposal To Replace The New Keynesian Phillips Curve," *The Quarterly Journal of Economics*, Vol. 117, No. 4, pp. 1295–1328, November.

Mansfield, Edwin (1961) "Technical Change and the Rate of Imitation," *Econometrica*, Vol. 29, No. 4, pp. 741–766.

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

McCracken, Michael W. and Serena Ng (2015) "FRED-MD: A Monthly Database for Macroeconomic Research," Working Papers 2015-12, Federal Reserve Bank of St. Louis.

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

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

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

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

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

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.

Rotemberg, Julio J. (2003) "Stochastic Technical Progress, Smooth Trends, and Nearly Distinct Business Cycles," *American Economic Review*, Vol. 93, No. 5, pp. 1543–1559, December.

Schmitt-Grohé, Stephanie and Martín Uribe (2012) "What's News in Business Cycles," *Econometrica*, Vol. 80, No. 6, pp. 2733–2764, November.

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

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

Sims, Eric R. (2012) "Taylor rules and technology shocks," *Economics Letters*, Vol. 116, No. 1, pp. 92–95.

——— (2016) "Whats news in News? A cautionary note on using a variance decomposition to assess the quantitative importance of news shocks," *Journal of Economic Dynamics and Control*, Vol. 73, No. C, pp. 41–60.

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

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

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

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

Appendix: Not For Publication

A Data in VAR

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

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

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

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

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

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

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

Table A.1: Variables Used

Label	Variable Name	Source	FRED Codes	Treatment: log	Treatment: pc
TFPL	Utilization-Adj TFP	Fernald (2014)†	–	•	•
RGDP	Real GDP	FRED	GDPC1	•	•
RCONS	Real Consumption	FRED	PCND; PCESV	•	•
RINV	Real Investment	FRED	GPDI; PCDG	•	•
RDGDP	R&D Expenditures (Y)	FRED	Y694RC1Q027SBEA	•	•
HOURS	Hours	FRED	HOANBS	•	•
UNRATE	Unemployment Rate	FRED	UNRATE	•	
LPR	Labor Force Participation Rate	FRED	CIVPART	•	
CAPUTIL	Capacity Utilization	FRED	TCU	•	
GDPDEF	GDP Deflator	FRED	GDPDEF	•	
RPINV	Price of Investment	FRED	DDURRD3Q086SBEA; DNDGRD3Q086SBEA; DSERRD3Q086SBEA; A006RD3Q086SBEA	•	
RWAGE	Real Wages	FRED	COMPRNFB	•	
SHORTR	Short Rate	FRED	DGS1		
YCSLOPE	Term Spread	FRED	DGS1; DGS10		
EQY	Equity Index	FRED*	SP500	•	
EQY2	Nasdaq	FRED	NASDAQCOM	•	
CCONF	Consumer Confidence	UMICH	–	•	
BCE5Y	Business Conditions E5Y	UMICH	–	•	
UE1Y	Unemployment E1Y	UMICH	–	•	
CBSPREAD	Corporate Bond Spread	FRED	AAA; BAA		

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

B Error Variance Decomposition

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

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

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

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

for any orthogonal matrix Q. u_t are the reduced-form VAR innovations. The external instrument of Section 3 allows identification of only one column b_0 of B_0, which contains

Figure B.1: Error Variance Decomposition: Frequency

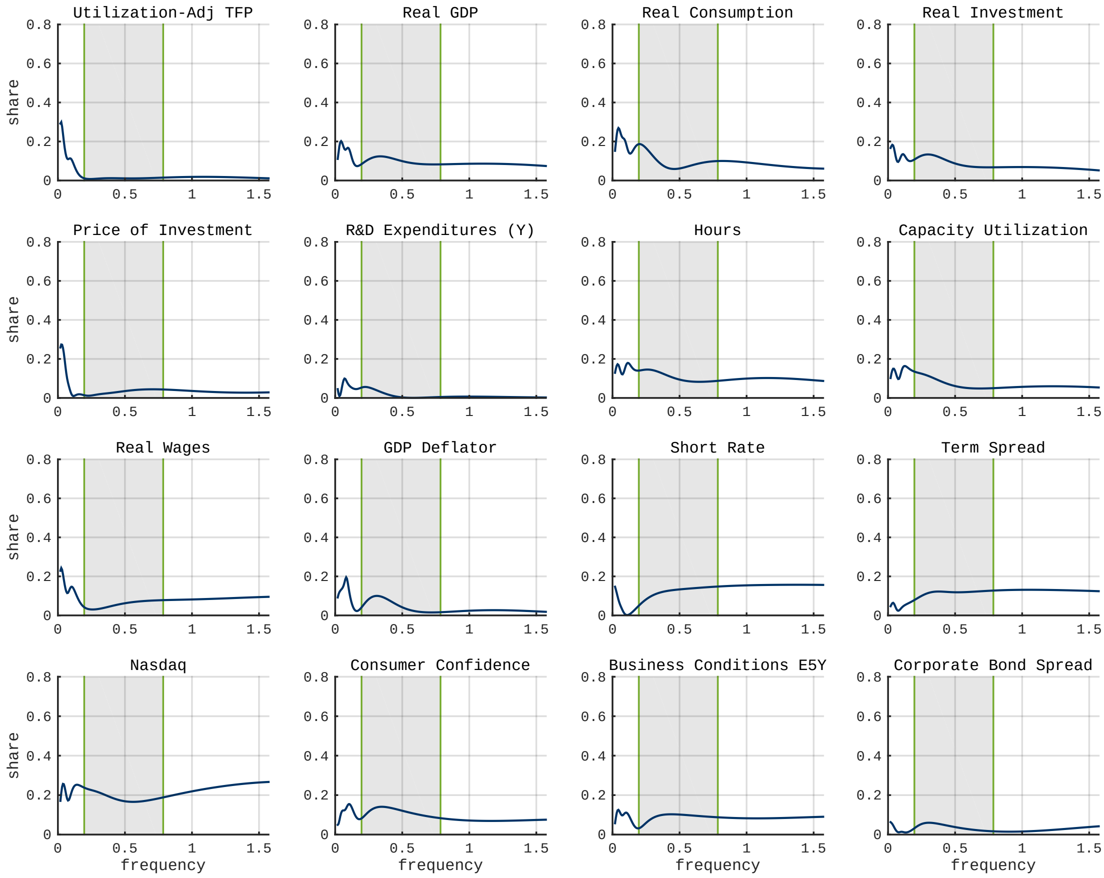

Note: Share of error variance accounted for by technology news shock identified with patent-based external instrument. VAR(4) with standard macroeconomic priors. Estimation sample 1971-I : 2016-IV; Identification sample 1982-I : 2006-IV. Shaded areas delimits business cycle frequencies (between 8 and 32 quarters).

the impact effects of the identified technology news shock $e_{\text{A2},t}$ on y_t.

The spectral density of y_t is

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

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

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

Figure B.2: Forecast Error Variance Decomposition: Time

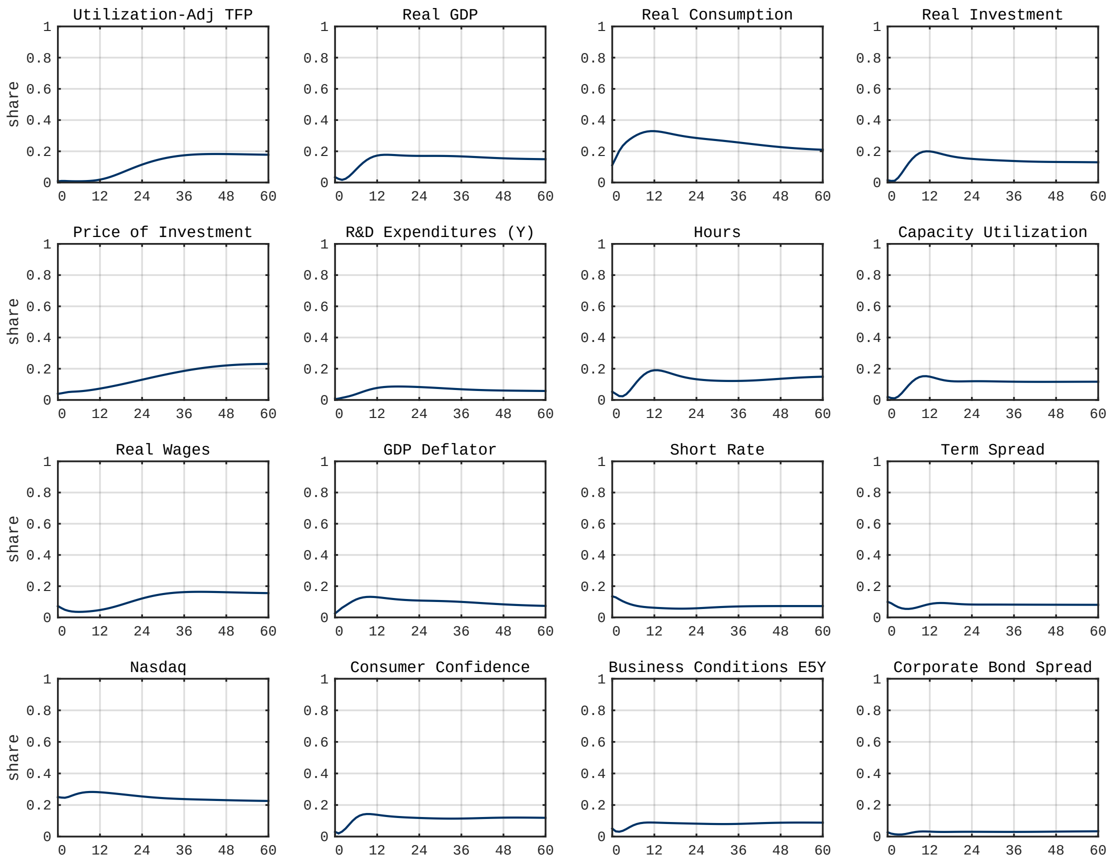

Note: Share of forecast error variance accounted for by technology news shock identified with patent-based external instrument. VAR(4). Estimation 1971-I : 2016-IV; Identification 1982-I : 2006-IV.

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

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

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

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

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

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

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

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

Figure B.1 plots the share of variance that is due to $e_{\text{A2},t}$ for all the variables included in our benchmark VAR at all frequencies between 1 (highest frequency) and 100 (lowest frequency) years. Grey areas highlight business cycle frequencies. Table 4 in Section 4 reports the share of variance due to $e_{\text{A2},t}$ over three different ranges of frequencies. Figure B.2 reports for comparison the share of forecast error variance accounted for by the identified shocks; forecast horizons, a mixture of high, medium and low frequencies, make isolating business cycle fluctuations more problematic.

C Comparison with Standard Identifications: 5-Variable VAR

Figure C.1 compares responses to a news shock identified with our patent-based instrument with two prominent identification schemes in the literature.

FIGURE C.1: DIFFERENT IDENTIFICATIONS IN 5-VARIABLE VAR

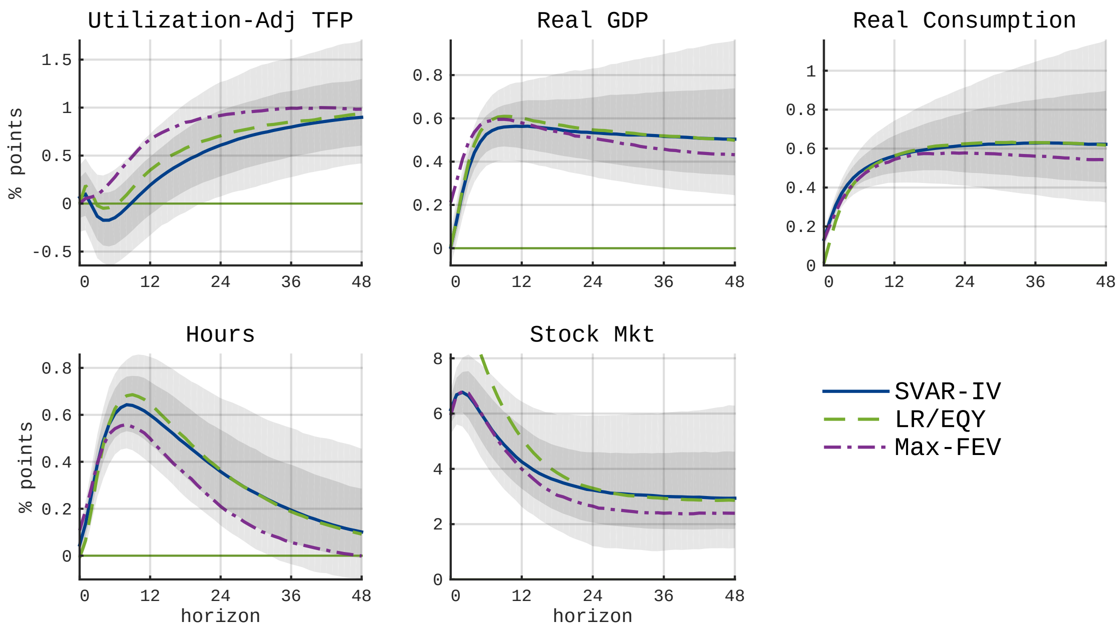

Note: Modal response to a technology news shock identified with (1) patent-based external instrument (SVAR-IV in blue), (2) long-run restrictions (LR/EQY in green dashed), and (3) maximum forecast error variance share (Max-FEV in purple dotted). Estimation sample 1971-I : 2016-IV. Identification sample 1982-I : 2006-IV. Shaded areas denote 68% and 90% posterior coverage bands for the SVAR-IV.

The first one is the one proposed in Beaudry and Portier (2006), denoted 'EQY/LR', and implemented as an innovation to the stock market index that is orthogonal to the current level of TFP. Beaudry and Portier (2006) show that, at least in their bivariate VAR, this is equivalent to identifying the news shock as being orthogonal to current TFP, but responsible for its long run variance. The second identification strategy is the one proposed in Barsky and Sims (2011), denoted 'Max-FEV'. Here the news shock is identified as being orthogonal to current TFP, and the one that maximizes the forecast error variance of TFP at all horizons between 0 and 40 quarters. All responses are scaled such that the peak response of TFP is equal to 1% across all identification schemes.

The vector of endogenous variables includes output, consumption, investment, total hours worked, and the stock market index. The variables are chosen as to encompass the sets used in the original VARs in both Beaudry and Portier (2006) and Barsky and Sims (2011).

D Additional Details on Patent Data

Figure D.1: Allowance Rates

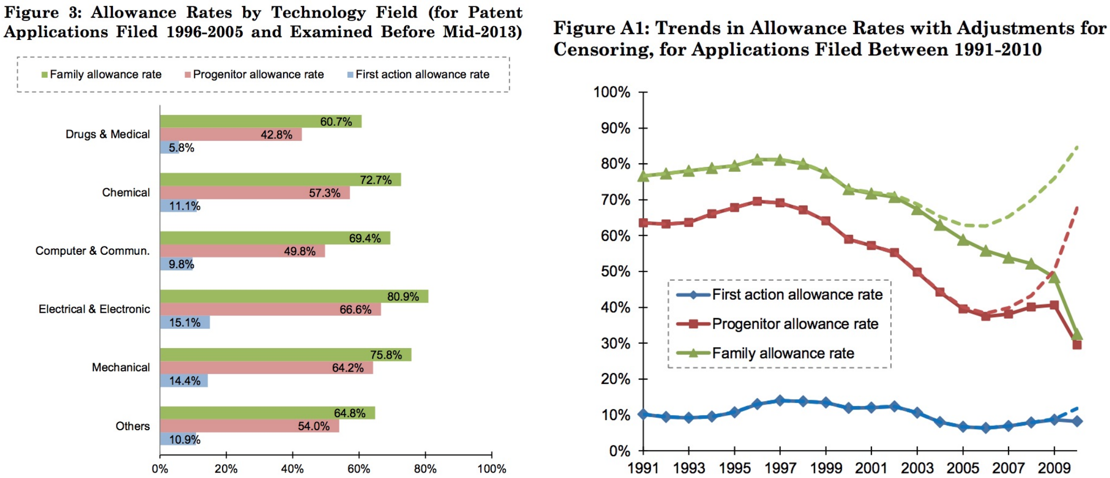

Note: [LEFT]: Allowance rates (proportion of patents granted out of submitted applications) across the six NBER patent-technology fields for applications filed at the USPTO for the first time between 1996 and 2005 and examined before mid2013. (i) *first-action allowance rate*: proportion of progenitor applications allowed without further examination; (ii) *progenitor allowance rate* (or simply, allowance rate): proportion of progenitor applications allowed without any continuation procedure, and (iii) *family allowance rate*: proportion of progenitor applications that produce at least one patent, including the outcomes of continuation applications that emerge from progenitor applications. [RIGHT]: Trends in allowance rates from 1991 to 2010, all categories. *Source*: Carley et al. (2015).

Table D.1: Lagged Information in Patent Applications

	F_1	F_2	F_3	F_4	F_5	F_6	F_7
Wald Test	9.215	1.252	1.835	0.642	1.437	0.256	0.209
p-value	0.000	0.293	0.126	0.634	0.226	0.905	0.933
Adj R^2	0.790	0.732	0.738	0.732	0.736	0.727	0.726
N	131	131	131	131	131	131	131

Notes: Numbers reported are Wald test statistics for joint significance of the first 4 lags of each factor F_t. The factors are extracted from the quarterly dataset of McCracken and Ng (2015). The dependent variable is the quarterly growth rate of utility patent applications: $pa_t = 100(\ln PA_t - \ln PA_{t-1})$. All the regressions include own 4 lags, regulation dummy and constant.

FIGURE D.2: CUMULATIVE DISPOSAL PROPORTION BY NBER CATEGORY

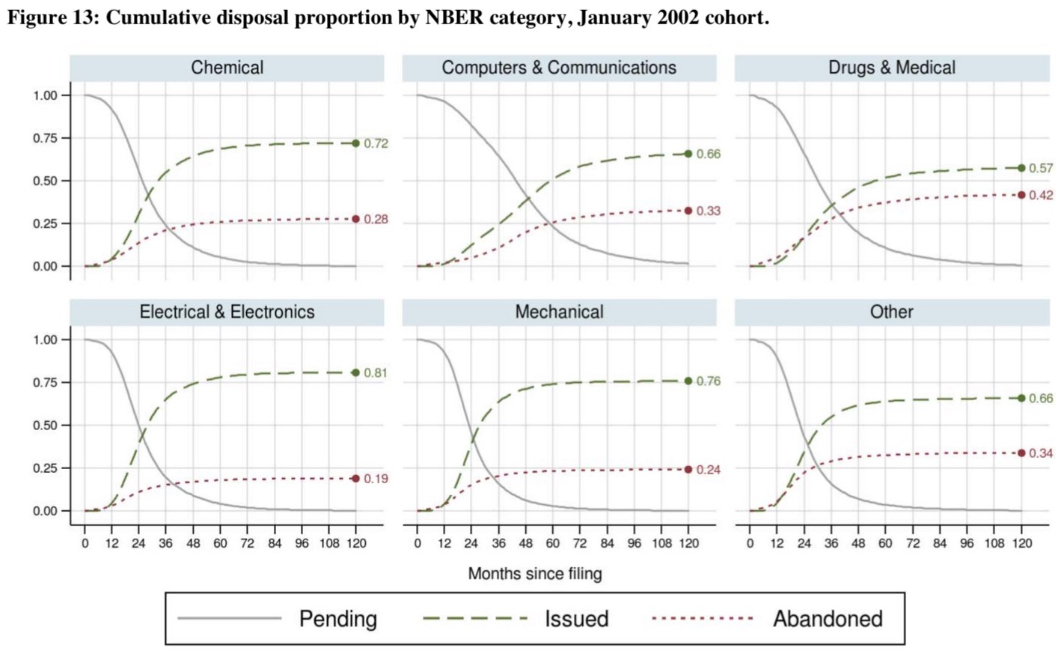

Source: Marco et al. (2015).

E Contemporaneous TFP Innovations

Figure E.1 reports impact responses for a selection of the variables in our VAR to a contemporaneous TFP innovation that raises TFP on impact by 1%, and obtained with a standard Cholesky factorization with TFP ordered first. The VAR is the same one used in Section 4. This identification scheme assumes that TFP is exogenous and only driven by technology shocks, and that the quarterly series of Fernald (2014) measures true technology without systematic error. Both these assumptions have been questioned in the literature (e.g. in Kurmann and Sims, 2017). This is however inconsequential; here we rely on standard Cholesky ordering only to highlight the differences between the impact effects of the two types of shocks. Full IRFs are not reported for space considerations, but available upon request.

The pattern of impact responses in Figure E.1 is fundamentally different from those elicited by the news shock. A positive contemporaneous TFP innovation significantly moves up output, consumption and investment on impact, while the response of hours is muted. Real wages increase robustly, and so do consumer expectations. Finally, there

Figure E.1: Impact Responses to a Contemporaneous TFP Innovation

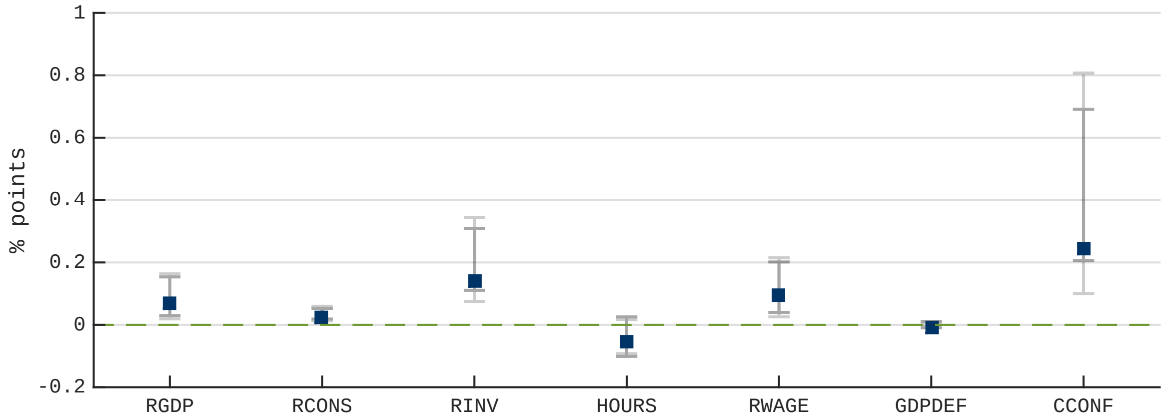

Note: Impact responses of selected variables to a TFP innovation that increases Utilization-Adjusted TFP by 1%. VAR(4). Estimation sample 1971-I:2016-IV. Grey bars delimit 68% and 90% posterior coverage bands.

seems to be no appreciable impact reaction of the price level.

F Additional Material

The impulse response functions reported in this Appendix are all scaled such that the peak response of utilization adjusted TFP equals to 1%.

Figure F.1: IRFs Full vs Pre-Crisis Sample

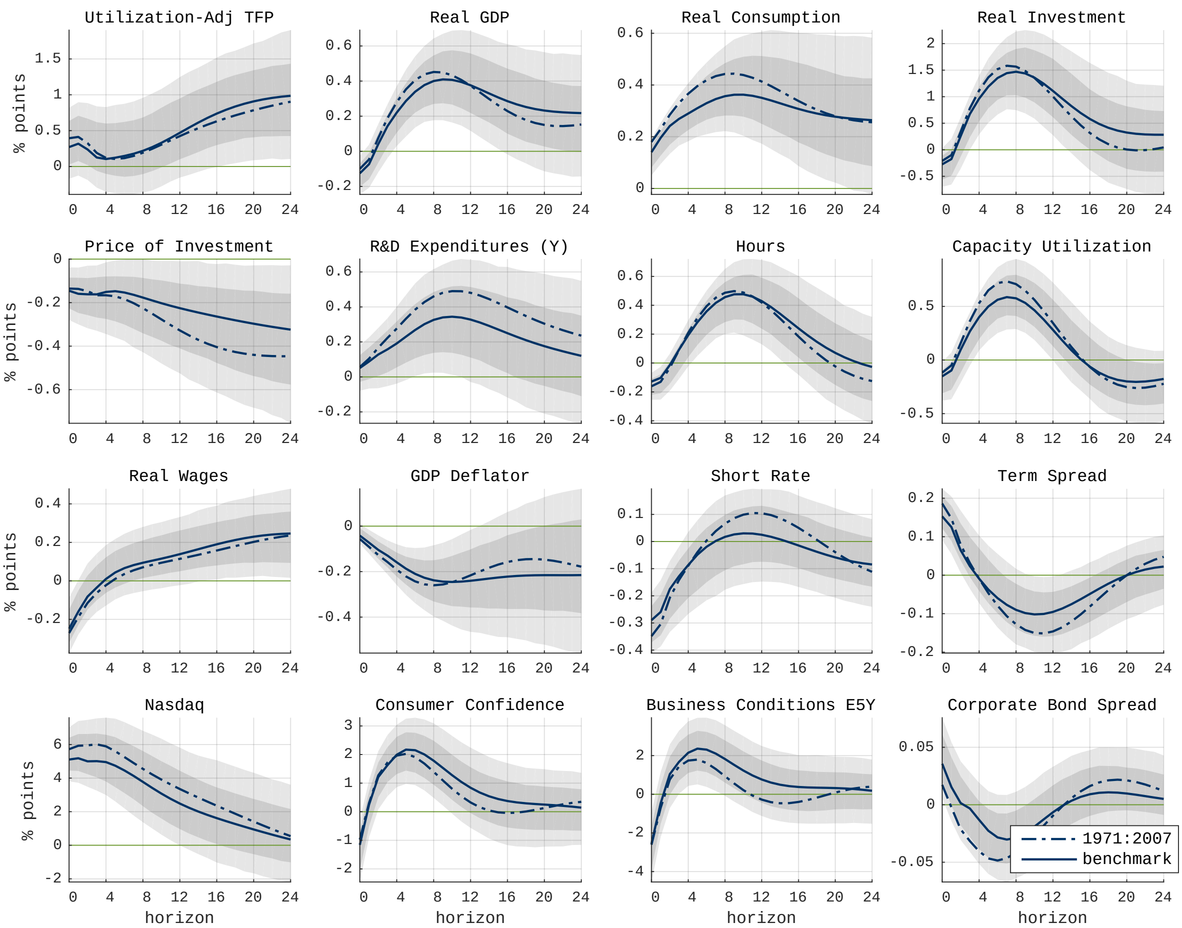

Note: Response of all variables to a technology news shock identified with patents-based external instrument. VAR(4) with standard macroeconomic priors. Estimation sample 1971-I : 2007-IV; Identification sample 1982-I : 2006-IV. Solid Lines: Instrument also controls for contemporaneous policy changes, benchmark. Dash-Dotted Lines: Instrument controls for SPF forecasts and lagged pa_t. Shaded areas denote 68% and 90% posterior coverage bands.

FIGURE F.2: IRFs LONGER SAMPLE

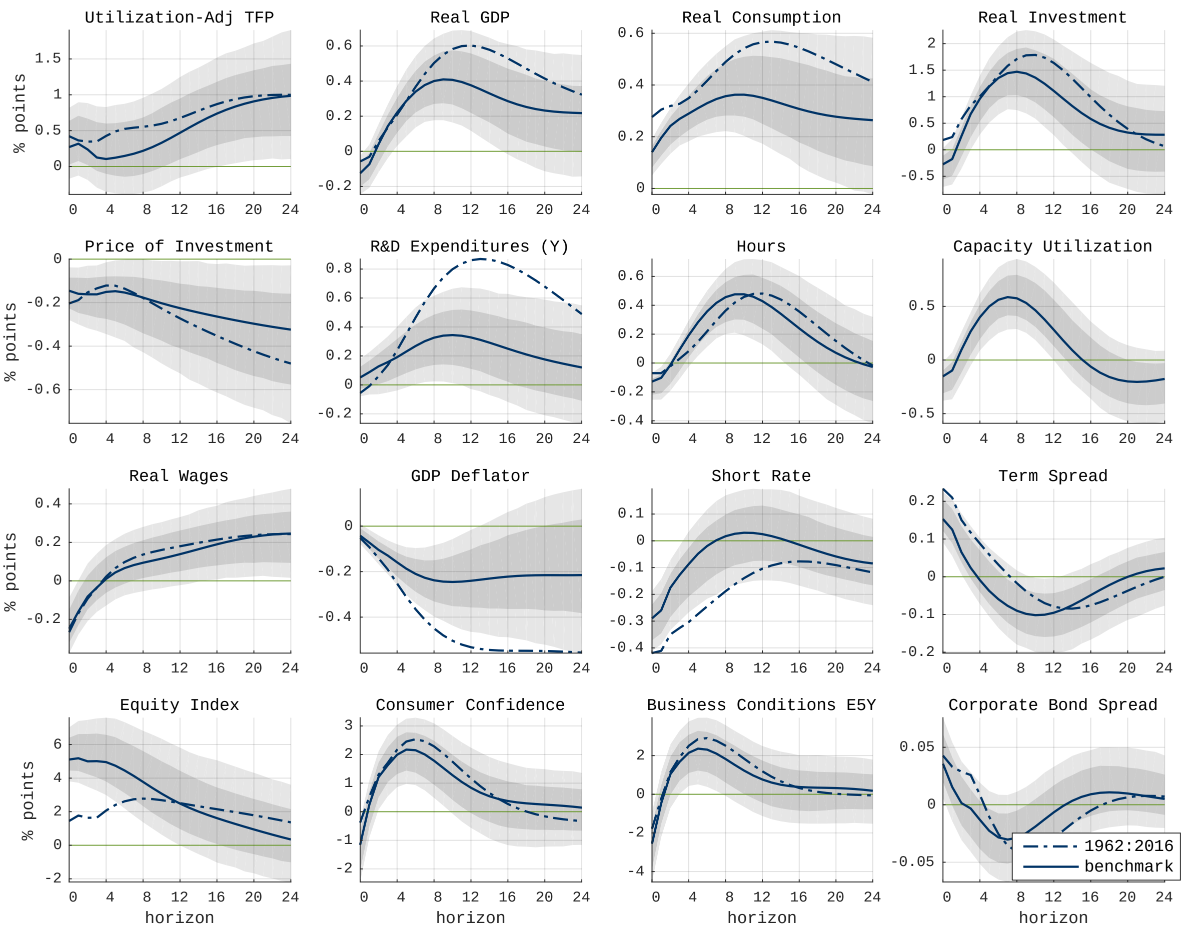

Note: Response of all variables to a technology news shock identified with patent-based external instrument. VAR(4) with standard macroeconomic priors. Solid Lines = Estimation sample 1971-I : 2016-IV; Identification sample 1982-I : 2006-IV. Dash-dotted Lines: Estimation sample 1962-I : 2016-IV; Identification sample 1982-I : 2006-IV. The equity index on the longer sample is the S&P 500 shown in the Nasdaq sub-plot as a dashed-dotted line. Shaded areas denote 68% and 90% posterior coverage bands.

Figure F.3: IRFs with Unemployment Expectations

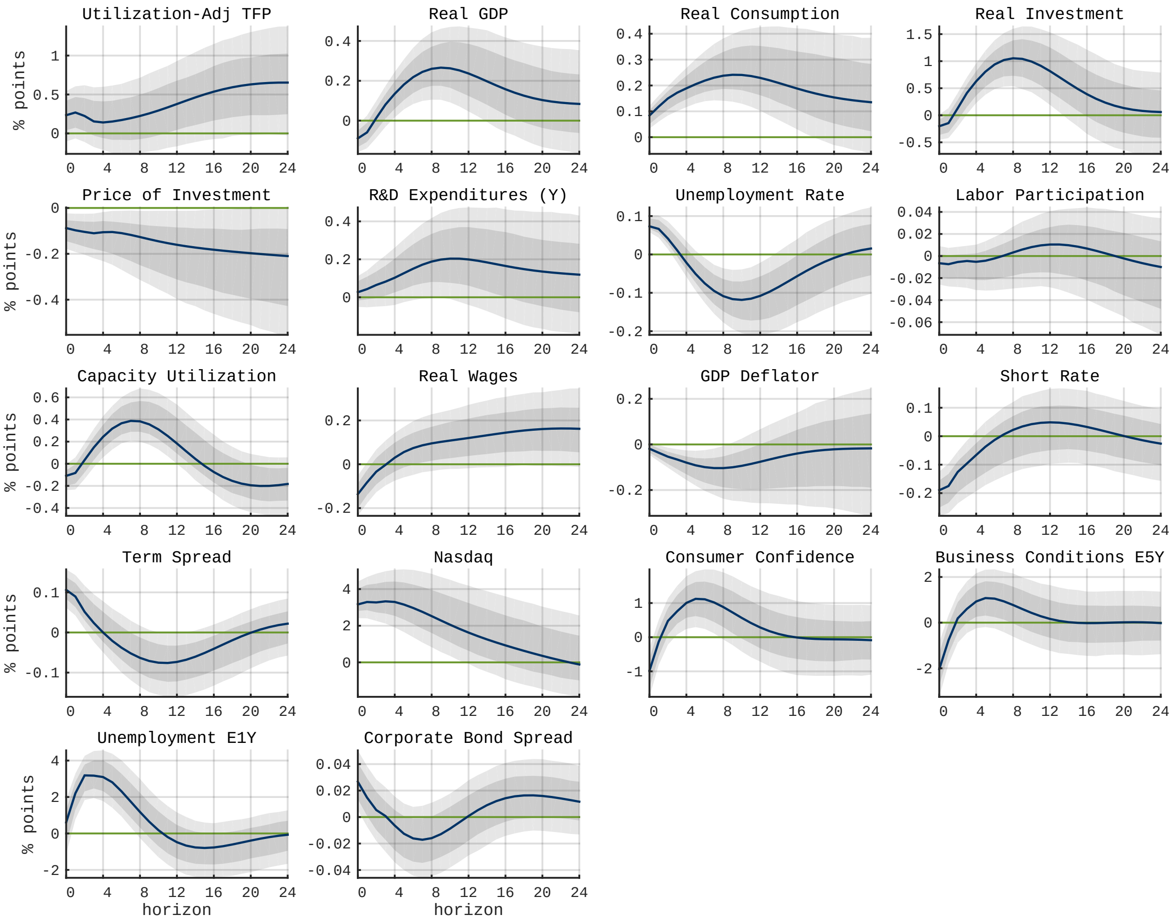

Note: Response of all variables to a technology news shock identified with patent-based external instrument. VAR(4) with standard macroeconomic priors. Instrument controls for contemporaneous policy changes. Estimation sample 1971-I : 2016-IV; Identification sample 1982-I : 2006-IV. Shaded areas denote 68% and 90% posterior coverage bands.

Figure F.4: IRFs Pre-Crisis Sample: Instruments

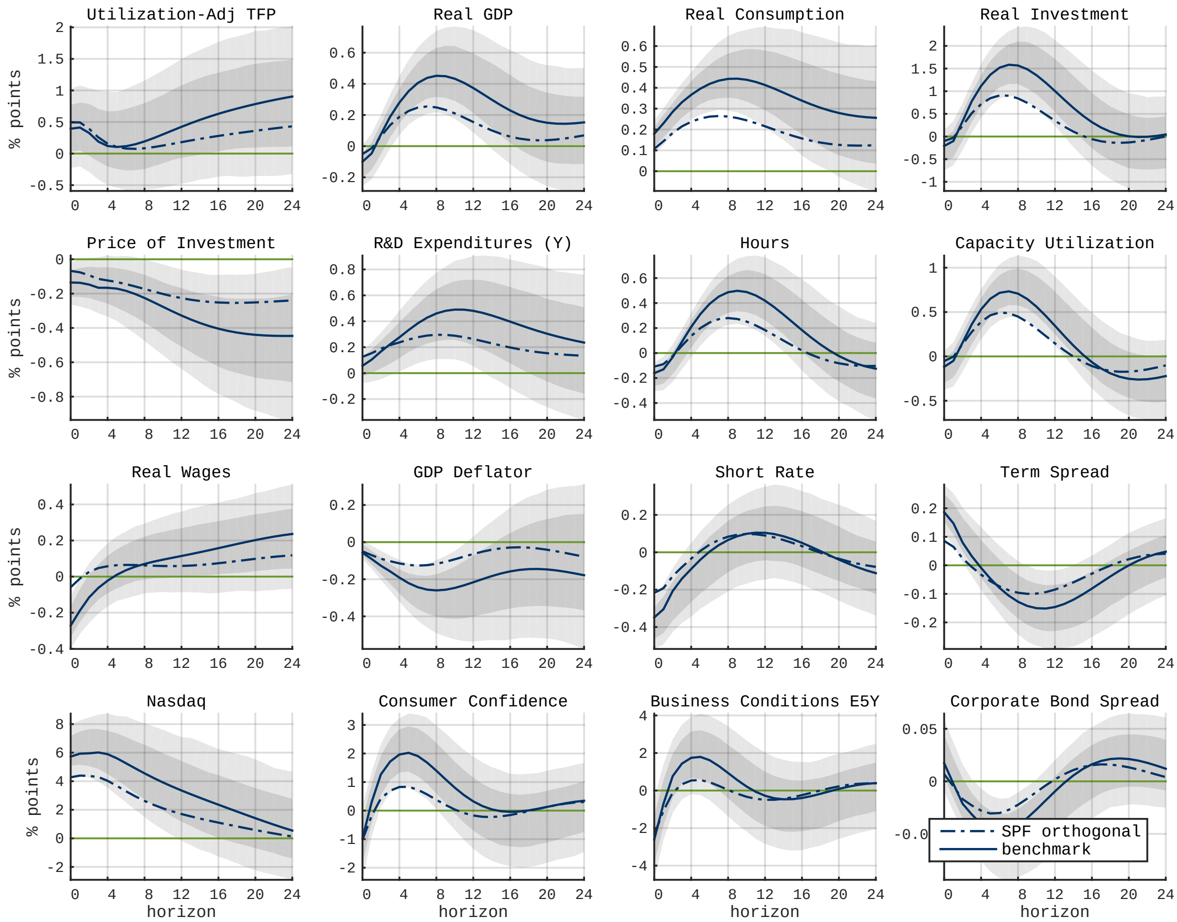

Note: Response of all variables to a technology news shock identified with patent-based external instrument. VAR(4) with standard macroeconomic priors. Solid Lines = Estimation sample 1971-I : 2007-IV; Identification sample 1982-I : 2006-IV. Dash-dotted Lines: Estimation sample 1971-I : 2007-IV; Identification sample 1982-I : 2007-IV. Shaded areas denote 68% and 90% posterior coverage bands.