When Creativity Strikes: News Shocks and Business Cycle Fluctuations*

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

Sinem Hacıoğlu Hoke†
Bank of England
DAFM (KCL)

Kristina Bluwstein‡
Bank of England

First Version: August 22, 2018, Revised: July 16, 2020

Abstract

We exploit information in a new dataset of monthly patent applications to construct an instrumental variable for the identification of technology news shocks that relaxes all the identifying assumptions traditionally used in the literature. Our sole requirement is that no structural disturbances affect the US economy via our instrument except for contemporaneous technology news. The instrument recovers news shocks that have no appreciable effect on aggregate productivity on impact, but are a significant driver of its trend component. News shocks prompt a sustained business cycle expansion in anticipation of the future increase in TFP, and are responsible for a sizeable share of economic fluctuations at business cycle frequencies. The stock market prices-in news shocks on impact, but consumer expectations take sensibly longer to adjust, consistent with the predictions of models of information frictions.

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

JEL Classification: E32, O33, O34, C36

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

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

1 Introduction

The idea that changes in agents' beliefs about the future may be an important driver of economic fluctuations has fascinated many scholars over the years. While the application to technology news is 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 a related question: 'How does the aggregate economy react to a shock that raises expectations about future productivity growth?' We provide an empirical answer 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 exploit information in a novel dataset of monthly US patent applications to construct an instrumental variable that allows us to dispense from all the identifying assumptions traditionally used in the news literature.

[1] The empirical literature on technology news shocks is vast, and we review it when discussing our results in Section 4. At the poles of the debate are the advocates of the news-driven business cycle hypothesis 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).

The intuition behind our identification is simple: Patent applications, by their nature, are a promise of potential future technological change. However, they may themselves be prompted by current economic booms and/or past news. We account for this endogeneity by controlling for expectations about the economic outlook that were formed prior to the filing dates, and for other contemporaneous policy changes. 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 taken from the Survey of Professional Forecasters intended to capture pre-existing expectations about the outlook that may influence the decision of filing a patent; and (iii) other contemporaneous monetary and fiscal policy changes.[2] The starting point for the construction of our external instrument are the monthly 'USPTO Historical Patent Data Files' (Marco et al., 2015) that provide a comprehensive record of all publicly available patent applications filed at the U.S. Patents and Trademark Office (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.

The exclusive rights granted to patent holders ensure that individuals and businesses have a set number of years to capitalize on their inventions, and act as a powerful incentive to engage in the patenting process. The length of time that then elapses from the application to the grant date, and the eventual diffusion of the innovation within the economy can be in the order of several years, depending on the type of patent and the characteristics of the industry sector.[3] Therefore, patent applications at any given time contain information about technological changes that may occur at some point in the future (see e.g. Griliches, 1990; Lach, 1995; Hall and Trajtenberg, 2004). In other words, and importantly for our purpose, they represent an uncontroversial way to measure news

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

[3]From application filing to grant issuance, the process takes on average 2 years as documented by Marco et al. (2015). While not all applications result in granted patents, the share of successful applications is substantial (up to 80%), with some heterogeneity across sectors.

about possible future technological progress, to a large extent regardless of whether such progress does indeed follow. Because patent applications are public, we can use the application filing dates as the first measurable time in which the news occurs, although it is clearly the case that the underlying idea, in the form of a private signal, predates it. Controlling for policy changes and for expectations about the future that precede the application filing is a necessary step to increase the likelihood that no other structural disturbances affect the US economy through the instrumental variable (IV), except for contemporaneous news. This is the identifying assumption in our SVAR-IV (Mertens and Ravn, 2013; Stock and Watson, 2012, 2018).

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. While it is not known ex ante whether technological innovation will effectively follow, the news we capture does eventually materialize on average, which results in our IV being associated with large subsequent increases in indices of aggregate innovation (Kogan, Papanikolaou, Seru and Stoffman, 2017), and in aggregate TFP eventually rising following the recovered structural shock. This allows us to label the recovered structural disturbance as news, as opposed to noise (see e.g. discussion in Chahrour and Jurado, 2018), overcoming the issues highlighted in Blanchard, L'Huillier and Lorenzoni (2013). Importantly, because innovations can in principle be released to the public under a 'patent-pending' status, our identification scheme does not warrant imposing orthogonality with respect to the current level of technology, which is a typical assumption in the news literature (see e.g. Beaudry and Portier, 2006; Barsky and Sims, 2011, among many others).[4] While such orthogonality condition is not imposed a priori, our external instrument recovers a shock that has essentially no effect on TFP either on impact, or in the years immediately afterwards. After such inertial initial reaction, aggregate TFP rises robustly, following the S-shaped pattern that is typical of the slow diffusion of technology (see e.g. Rogers, 1962; Gort and Klepper, 1982). Similarly, albeit we impose no constraints on variance shares ex ante, the recovered shock explains

[4]In this respect, our identification is akin to Barsky et al. (2015); Kurmann and Sims (2017), who also relax the assumption of a zero impact response of TFP. Our approach is also robust to mismeasurements in commonly used empirical estimates of aggregate technology (see e.g. discussions in Fernald, 2014; Kurmann and Sims, 2017).

only a modest fraction of the variation of TFP at frequencies higher or equal than those associated with standard business cycle durations, and is instead an important driver of its long-run/permanent component.

When looking at the broader response of macroeconomic aggregates, we find that following positive news, consumption, investment, output, and hours worked all rise in anticipation of future technological improvement. In particular, by the time TFP materially departs from its initial level, all other variables in our VAR have reached the peak of their dynamic response. 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). Moreover, this large asynchronicity in the timing of the estimated dynamic responses suggests that the aggregate effects of technology news that we unveil may be predominantly (if not entirely) driven by beliefs, rather than by future realized fundamentals. The shock that we recover is responsible for a sizeable share of aggregate economic fluctuations. At business cycle frequencies, it accounts for about 15% of the variation in consumption and hours worked, and for 20% of that in the stock market. Moreover, it accounts for over a fifth of the high-frequency fluctuations in the unemployment rate (see also Faccini and Melosi, 2018, for corroborating evidence). These shares are economically meaningful, and particularly so given the size of our VAR.

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

Our work is closely related to a stream of studies that have relied on empirical measures of technological changes to identify technology news shocks. The first such study is Shea (1999). Here annual patent applications and R&D expenditures are used to estimate the effects of technology shocks on industry aggregates. Identification is achieved by ordering either measure last in a battery of small-scale VARs that also include labor inputs and productivity. Christiansen (2008) extends this study by using over a century of annual patent application data. The benchmark specification is a bivariate VAR with labor productivity and patents ordered first. Alexopoulos (2011) uses the number of book titles published in the field of technology to capture the time in which the novelty is commercialized. Responses of aggregate variables are estimated in a set of bivariate VARs with the publication index ordered last.[5] Our paper differs from these contributions in several ways. First, these studies address the fundamental endogeneity of empirical measures of technological changes only to the extent that it is captured in the reminder of variables included in the bi/tri-variate VARs. Other than relying on a richer VAR specification, in the construction of our instrument we explicitly control for the fact that the cyclical nature of patent applications may be influenced by current economic conditions or indeed by past news. Second, and related, these studies have all implicitly assumed the empirical measure of technology being a near perfect measure of news shocks. In fact, their identifying assumptions amount to effectively retrieving the transmission coefficients by running a distributed lag regression (with some controls) of the variables on the patent data. In contrast, our identifying assumptions explicitly account for the possible presence of measurement error in the constructed instrument. Finally, these studies have all relied on annual data potentially overlooking important higher frequency variation which instead we exploit for the identification.

The structure of the paper is as follows. Section 2 introduces the external instrument and describes the patent data used for its construction. In Section 3 we lay out the identifying assumptions in our SVAR-IV and discuss the identification of technology news shocks in the context of a core 5-variable VAR. In Section 4 we extend the analysis

[5]More recently, Baron and Schmidt (2014) have used technology standards and a recursive identification to infer on the aggregate implications of anticipated technology shocks. In an international context, Arezki 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.

to an information-rich 16-variable VAR to explore the transmission mechanisms more in detail. Section 5 concludes. Additional material is reported in the Appendix.

2 A Patent-Based IV for Technology News Shocks

2.1 Information in Patent Data

Our starting point for the analysis is the monthly flow of all new patent applications filed at the U.S. Patent and Trademark Office. The data are from the 'USPTO Historical Patent Data Files' compiled by Marco et al. (2015) as a follow up and extension of Hall et al. (2001). The dataset records the monthly stocks and flows of all publicly available applications and granted patents filed from January 1981 to December 2014. The stocks include pending applications and patents-in-force; flows include new applications, patent grants and abandonments.[6,7]

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

[6]The dataset is available at `http://www.ustpo.gov/economics`. We discard information relative to both abandonments and patents grants. Innovations can be released under patent-pending status, hence, most of the 'news content' in patent applications may be exhausted by the time it is granted. Moreover, grants tend to be significantly more cyclical than applications, and dependent on the intensity of labor and administrative cycles at the USPTO (Christiansen, 2008).

[7]The proportion of patent applications that eventually results in a grant being issued can vary substantially, both over time, and across industry sectors. Marco et al. (2015) provide an example for success rates of the 2002 cohort of patent applications: 57% in 'drugs & medical', to 81% in 'electrical & electronics' were eventually granted. This supports the intuition that patent applications give a strong signal about future technological changes.

Figure 1: Patent Applications & Aggregate Innovation

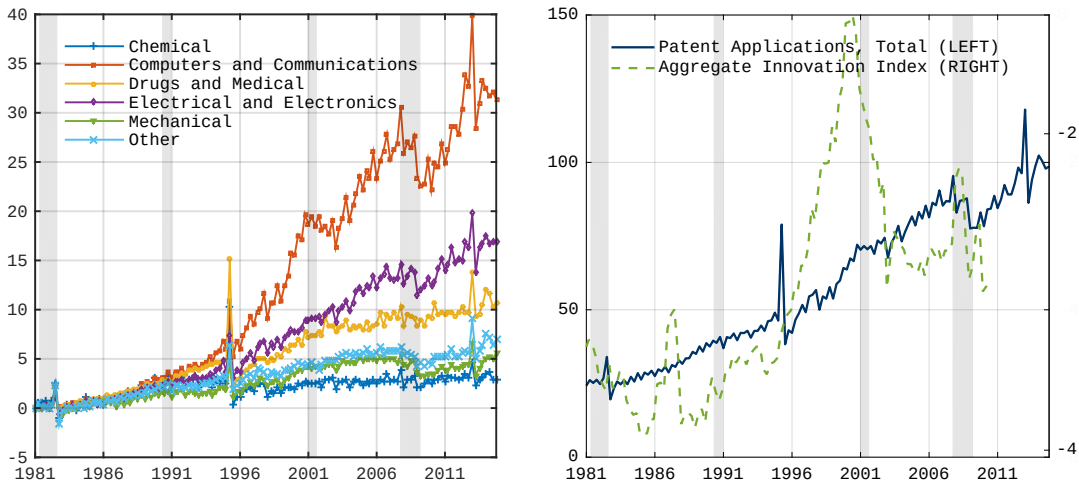

Note: [LEFT] Patent applications across all NBER categories. Quarterly figures obtained as sum of monthly readings, 1981-I=0. Thousands. [RIGHT] Total number of applications (sum across categories), thousands, left axis. Kogan et al. (2017) aggregate measure of economic value of innovations, GDP weighted,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

There have been three important regulatory changes in patenting in 1982, 1995, and 2013. All these regulations affected the number of applications when they came into effect, as shown by the spikes in Figure 1. However, since they were not legislated in response to considerations related to either current or anticipated economic conditions, they provide us with important exogenous variation which we exploit for the identification. Said differently, to the extent that each patent embeds a signal about potential future technological progress, the increase in applications induced by each piece of legislation is an exogenous (relative to macroeconomic conditions) increase in technology news, which is the focus of our identification. In Section 3 we explore the sensitivity of our results to these spikes.

In 1982, the old Court for Customs and Patent Appeals was abolished, and a new Court of Appeals for the Federal Circuit was established; the new court provided more protection to patents' owners against infringement. In 1995, the U.S. implemented wide-ranging changes to patent law under the Agreement on Trade-Related Aspects of Intellectual Property Rights (TRIPS), as part of the Uruguay Round Agreements Act. The TRIPS agreement's main purpose was to harmonize patenting rules among all members of the World Intellectual Property Organization with the aim to contribute to the pro-

motion of technological innovation and to the transfer and dissemination of technology.[8] One of the main changes introduced by the TRIPS Agreement was that of promoting transparency in patenting, and disincentivize strategic behaviour through stricter regulation.[9] This had two main effects. First, it shifted forward the timing of some applications, which resulted in the one-off increase highlighted in the chart. Second, it made applications more informative about future innovations (Encaoua et al., 2006). Finally, in March 2013, the U.S. implemented the rules dictated by the America Invents Act which further revised ownership rights.[10]

In the right panel of Figure 1 we plot the total number of applications (sum across industries) against the aggregate innovation index of Kogan et al. (2017) (dashed line). Using their data, we have constructed a quarterly version of their index as the 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's stock price in a three-day window that brackets the date in which the patent is granted to the firm. Kogan et al. (2017) document that their measure is strongly positively correlated with forward citations, which 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. Because it is based on financial data, this index is a forward looking measure of the private, economic value of innovations. Because it is based on patent grants, we expect it to lag the time series of patent applications. We note that in the relevant sample, the large spikes in the number of applications tend to correspond to substantial subsequent increases in the innovation index, and this is particularly true after the TRIPS Agreement. We take this as preliminary indication

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

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

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

that the exogenous sample variation introduced by the changes in legislation that gives rise to an increase in news is also informative about their future 'innovation content'.[11]

2.2 Instrument Construction

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

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

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

In Eq. (1), pa_t is the quarterly growth rate of patent applications, i.e. $pa_t = 100 \times (\ln PA_t - \ln PA_{t-1})$, $\gamma(L) = \sum_{j=1}^{4} \gamma_j L^j$ where L is the lag operator, and $\mathbb{E}_t[x_{t+h}]$ is an $m \times 1$ vector of forecasts for the economic variables in x_t that we take from the Survey of Professional Forecasters (SPF). The forecast horizon h is equal to one and four quarters. The time index in $\mathbb{E}_t$ refers to the publication date of the survey. Because of the release schedule of the SPF, the information set conditional on which forecasts are made is in fact relative to the previous quarter; hence, the collection of forecasts in $\mathbb{E}_t[x_{t+h}]$ captures pre-existing beliefs about the macroeconomic outlook.[12] The vector x_t includes the unemployment rate (u_t), inflation (π_t), and the growth rates of real non-residential fixed investments

[11] Building on this evidence, in a recent contribution Cascaldi-Garcia and Vukotic (2019) use the index of Kogan et al. (2017) to identify technology news shocks as a follow up to our analysis in this paper.

[12] 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.

(I_t), and of real corporate profits net of taxes (Π_t).[13]

An important concern relates to the potential correlation of patent applications with other contemporaneous shocks, besides current technology news. If this were the case, the exclusion restrictions in our IV-based identification strategy would be violated. While there is no formal way to test for the exogeneity of the instrument, we address this concern by including in Eq. (1) further controls that capture monetary and fiscal policy changes up to the current quarter. The vector η_t includes unexpected and anticipated exogenous tax changes as classified by Romer and Romer (2010) and Mertens and Ravn (2012), and the narrative series for monetary policy shocks of Romer and Romer (2004).[14] The rationale here is that monetary and tax policy, by affecting macro aggregates (especially investment) within the quarter, may have a direct effect on patent applications, and act as a confounding factor in the identification.

The regression results are presented in Table 1. The table reports individual regression coefficients and robust standard errors in parentheses for five models. Eq. (1) corresponds to column (5) in the table. In columns (1) to (4) we consider subsets of controls for comparison. Due to the availability of the narrative tax series, the specifications in columns (4) and (5) are estimated over the sample 1981-I:2006-IV. Columns (1) to (3) use the full length of patent data (1981-I:2014-IV). At the bottom of the table, we report Wald test statistics for the joint significance of the controls (excluding own lags) in each regression. Patent applications exhibit a strong autocorrelation pattern.[15] Moreover, pre-existing beliefs about the future as captured by the SPF forecasts contain information for patent applications beyond that included in own lags. This is consistent with patents being endogenous to the economic cycle, and, potentially, also related to past news embedded in the survey forecasts. Policy changes are also informative.

[13]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.

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

[15]Because we use the instrument to identify technology news shocks in VARs of order 4, the inclusion of lags in Eq. (1) is not strictly necessary for our application. Removing the dependence of the instrument on its past however makes it ready to use also in other applications such as VARs with different lag orders, or local projections. Moreover, it controls for the seasonal pattern in patents data which the sign of the autocorrelation coefficients in Table 1 points towards.

Table 1: Instrument Construction

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

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

The procedure in Eq. (1) removes both the autocorrelation and the dependence on pre-existing beliefs as captured by the SPF, and ensures the orthogonality of the IV to other contemporaneous policy shocks. The resulting instrument is plotted in Figure A.1 in the Appendix. In Tables A.3 and A.4 also in the Appendix we check for correlation of the recovered instrument both with a broader set of forecasts, and with lagged macro-financial factors extracted from the large set assembled in McCracken and Ng (2015). In both cases, we do not find evidence against the null that the instrument is not Granger-caused by these variables. On the other hand, Tables A.1 and A.2 show that these variable do Granger cause patent applications. We argue that it is unlikely that structural disturbances other than current technology news may affect the US economy through z_t. This is our sole identifying assumption.

3 Identification of Technology News Shocks

In the news literature, it is common to think of the process for technology as a random walk with drift subject to two stochastic disturbances. A typical representation assumes technology to be the sum of a stationary and a permanent component, with news shocks affecting the latter (see e.g. Blanchard et al., 2013; Kurmann and Sims, 2017). Formally

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

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

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

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

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

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

$e_{A2,t}$ is the news shock. The standard identifying assumption in the news literature is that agents learn about $e_{A2,t-k}$ before it hits the technology process, i.e. $k > 0$ (see e.g. Beaudry and Portier, 2006; Barsky and Sims, 2011, among many others). However, a number of more recent papers have argued that news shocks are also in principle compatible with $k = 0$, which would affect technology also on impact (see e.g. Barsky et al., 2015; Kurmann and Sims, 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.

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

3.1 Identifying assumptions in our SVAR-IV

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

Let y_t denote the n-dimensional vector of economic variables of interest, whose dy-

namics follow a VAR(p)

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

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

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

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

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

When agents anticipate future changes, as is the case with technology news shocks, non-fundamentalness is likely to arise (see e.g. Leeper et al., 2013). Intuitively, if the shock only has effect on future variables, current realizations are only informative about past shocks, and the mapping in Eq. (6) breaks down. In this context, a natural route towards the problem solution is to add information to the VAR, through variables that help revealing the state variables.[16,17] This is the role of e.g. the stock price index

[16] In the context of technology news shocks, the issue arises because, due to anticipation, news shocks also become state variables that agents need to keep track of when solving their equilibrium problem. However, these being unobservable, they cannot be conditioned upon, and the problem essentially becomes one of missing information: the observables are insufficient to reveal the true states.

[17] While the issue of non-fundamentalness is a theoretically binding constraint for the usefulness of empirical VARs, Sims (2012) shows that, empirically, it should not be thought of as an 'either/or' problem. Even with non-invertibility, the 'wedge' between the shocks estimated in a structural VAR and the theoretical ones may be small enough that VAR-based inference may still deliver accurate results, in the form of impulse response functions to the identified shocks. This point is further discussed in Beaudry and Portier (2014); Beaudry et al. (2015, 2016).

in Beaudry and Portier (2006), or measures of consumers or business confidence as in Barsky and Sims (2012). In a similar vein, factors estimated from large cross-sections can be added to the VAR specification as in e.g. Giannone and Reichlin (2006); Forni and Gambetti (2011).

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

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

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

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

3.2 Technology News Shocks in a Core 5-variable VAR

In this section, we put our instrument to test in a core 5-variable VAR and discuss the sensitivity of our results with respect to a number of perturbations. The variables included in the VAR are the quarterly estimate of TFP corrected for input utilization of Fernald (2014), output, consumption, total hours worked, and a stock market index. The variables are chosen as to encompass the sets used in the VARs of Beaudry and Portier (2006) and Barsky and Sims (2011). The variables enter the VAR in log levels, and are deflated and expressed in per-capita terms where appropriate. We use the GDP deflator to measure inflation and report a detailed description of the data and their construction in Table B.1 in the Appendix. The VAR is estimated with Bayesian techniques with 4 lags over the sample 1971-I:2016-IV, where the starting date is constrained by the availability of the Nasdaq Composite stock market index.[18] We refer to the sample used for the VAR estimation as the estimation sample, and the one used for the projection of the VAR residuals on the instrument as the identification sample respectively. Our identification sample equals the full length of z_t (1982:I to 2006-IV).

[18] We discuss results relative to variations in the estimation sample and of the response of the S&P 500 in the next section.

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

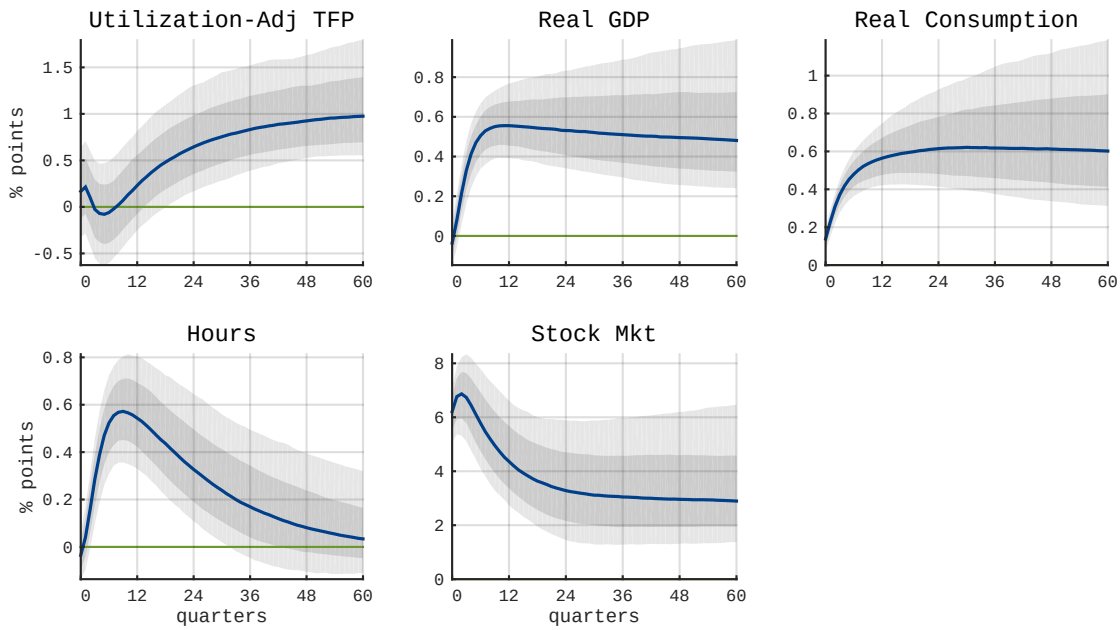

Note: Modal responses to a technology news shock identified with patent-based IV. Estimation sample 1971-I:2016-IV. Identification sample 1982-I:2006-IV. Shaded areas denote 68% and 90% posterior credible sets.

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

The IRFs are reported in Figure 2. A few elements stand out. First, while we have not imposed any restrictions on the effect of the shock on current TFP, the chart reveals that the shock recovered by our instrument has essentially no effect on TFP neither on impact, nor in the first few years immediately following. TFP eventually rises robustly and remains elevated throughout. The shape of the TFP response resembles the S-shaped pattern that is typical of the slow diffusion of new technologies (see e.g. Rogers, 1962; Gort and Klepper, 1982). 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.[19] Second, output, consumption and hours worked all rise. Aggregate consumption increases already on impact, while the initial response of output and hours is more muted. For all three variables, the rise is sudden, and the peak of the dynamic adjustment is reached long before any material increase in TFP materializes. Third, the stock market prices-in the news on impact, and only slowly reverts back.

The IRFs in Figure 2 are compatible with a 'news-driven' business cycle view in which macroeconomic aggregates react positively to positive news, and a business cycle expansion arises in anticipation of potential future technological improvements. Notwithstanding the minimal set of identifying restrictions, the pattern of IRFs recovered by our IV shares many similarities with those in prominent studies such as Beaudry and Portier (2006) and Barsky and Sims (2011), as we report in Figure D.1 in the Appendix.[20] What is remarkable in this context is that the negligible impact response of TFP, the stock market pricing-in the news on impact, or, as we discuss below, the shock having maximum explanatory power for TFP at long horizons – assumed for identification in these earlier studies –, become instead results in our setting.

Before turning to the variance shares, in Figure 3 we evaluate the sensitivity of these results to the definition of our IV, and to the 1995-TRIPS spike. Each subplot in the figure reports three lines. The solid lines trace out the responses identified with our IV. These are the same as in Figure 2. The dashed lines are obtained with an IV that only controls for the lags in patent applications and for prior beliefs (i.e. we set all the δ coefficients to zero in Eq. 1), and allows us to evaluate to what extent the potential contamination by other policy shocks is an issue. The identification sample in this case is also longer (1982-I:2014-IV). Finally, the dash-dotted lines are obtained with the baseline IV but

[19] 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.

[20] Beaudry and Portier (2006) identify technology news shocks as an innovation to the stock market index that is orthogonal to the current level of TFP. 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. Kurmann and Mertens (2014) document that Beaudry and Portier (2006)'s identification does not have a unique solution when more variables are added into their model. Barsky and Sims (2011) identify news shock as being orthogonal to current TFP, and maximizing the forecast error variance of TFP at all horizons between 0 and 40 quarters.

Figure 3: Sensitivity to Other Policy Shocks & Spikes in IV

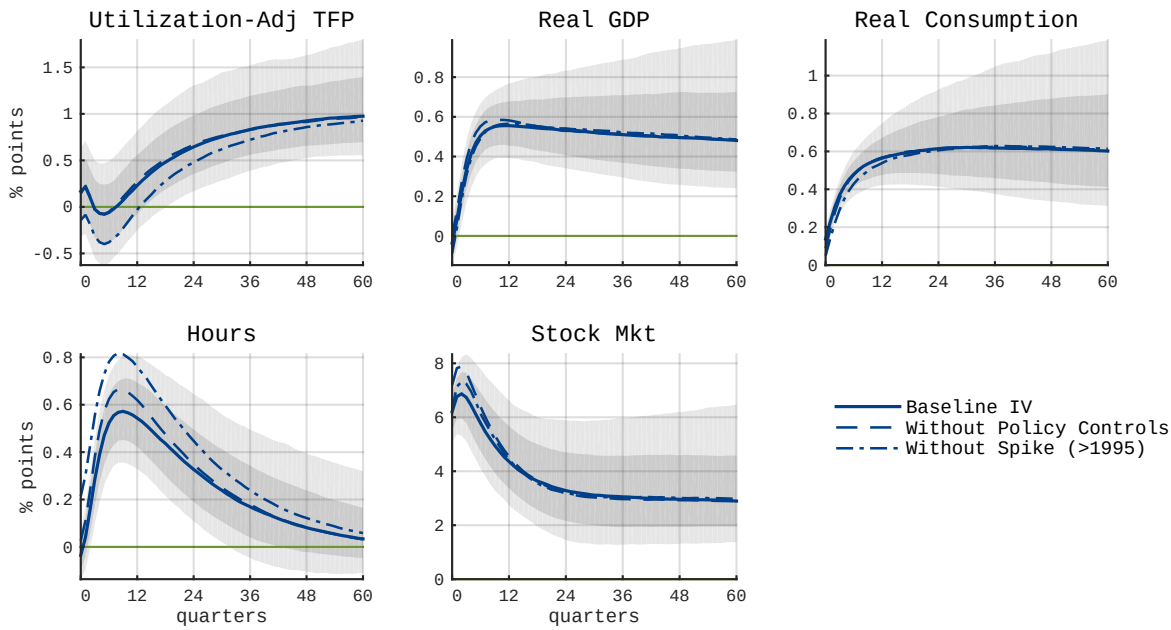

Note: Modal responses. Estimation sample 1971-I:2016-IV. Identification samples are: 1982-I:2006-IV with the baseline IV (solid lines); 1982-I:2014-IV for the IV that does not control for policy shocks (dashed lines); 1996-III:2006-IV for the IV that excluded the regulation spikes (dash-dotted lines). Shaded areas denote 68% and 90% posterior credible sets for the baseline IV.

on a restricted identification sample that starts in 1995-III. This effectively removes the regulation spikes and uses the portion of the sample where arguably patents were more informative for future innovations, as noted in Section 2. We note that for all the variables in the VAR the differences are minimal.[21]

To complete the discussion, Figures D.2 and D.3 in the Appendix report the share of variance that is accounted for by the shock identified with our IV in the 5-variable VAR. In Figure D.2 the decomposition is performed across different frequencies, with those corresponding to standard business cycle lengths highlighted by grey areas, while Figure D.3 reports the standard forecast error variance decomposition across horizons. Two main results stand out. The first, that we also confirm in the larger VAR of the next section, is that even if we have not imposed any such restriction ex ante, the shock

[21] Using the post-95 identification sample serves as a useful illustration in this instance, but because of the very small number of observations on which it is based, we would not want to use it as our benchmark. To further evaluate the role played by the TRIPS spike in the IV we have evaluated IRFs identified with an IV defined as a dummy variable that is equal to 1 in 1995-II, and zero otherwise. The TRIPS dummy does not recover the same IRFs as our baseline IV, which suggests that while important for the identification, the TRIPS spike is not entirely driving the results.

recovered by the patent-based IV is most explanatory for TFP at long horizons and at very low frequencies. This is consistent with it being a driver of the long-run component of aggregate productivity. The second is that in the core 5-variable VAR, the shock explains an implausibly high share of the variation in aggregate macro variables, which reaches about 80% in the case of consumption and output. Two features of the core VAR are likely to account for such large variance shares. First, the residuals of this VAR are Granger-caused by factors extracted from large cross-sections (see Forni and Gambetti, 2014). This may introduce a bias, as discussed in Forni et al. (2019). Second, the 5-variable VAR has a unit-root. This is quite visible in the strong persistence of the IRFs reported both in this section and in Appendix D, and introduces a further distortion in the computation of the variance shares, which are essentially a function of powers of the autoregressive coefficients, particularly at long horizons.

In the next section we evaluate the transmission of technology news shocks, both in terms of IRFs and EVDs, in a large VAR for which these limitations do not apply.

4 The Broader Propagation of Technology News

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

4.1 Dynamic Responses

As for the 5-variable VAR, we include 4 lags, and estimate the coefficients using standard Normal-Inverse Wishart priors over the sample 1971-I:2016-IV.[22] With the exception of interest rates and spreads, all the variables enter the specification in log levels, and are deflated and expressed in per-capita terms where appropriate. A complete description

[22] 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.

Figure 4: Propagation of Technology News Shocks

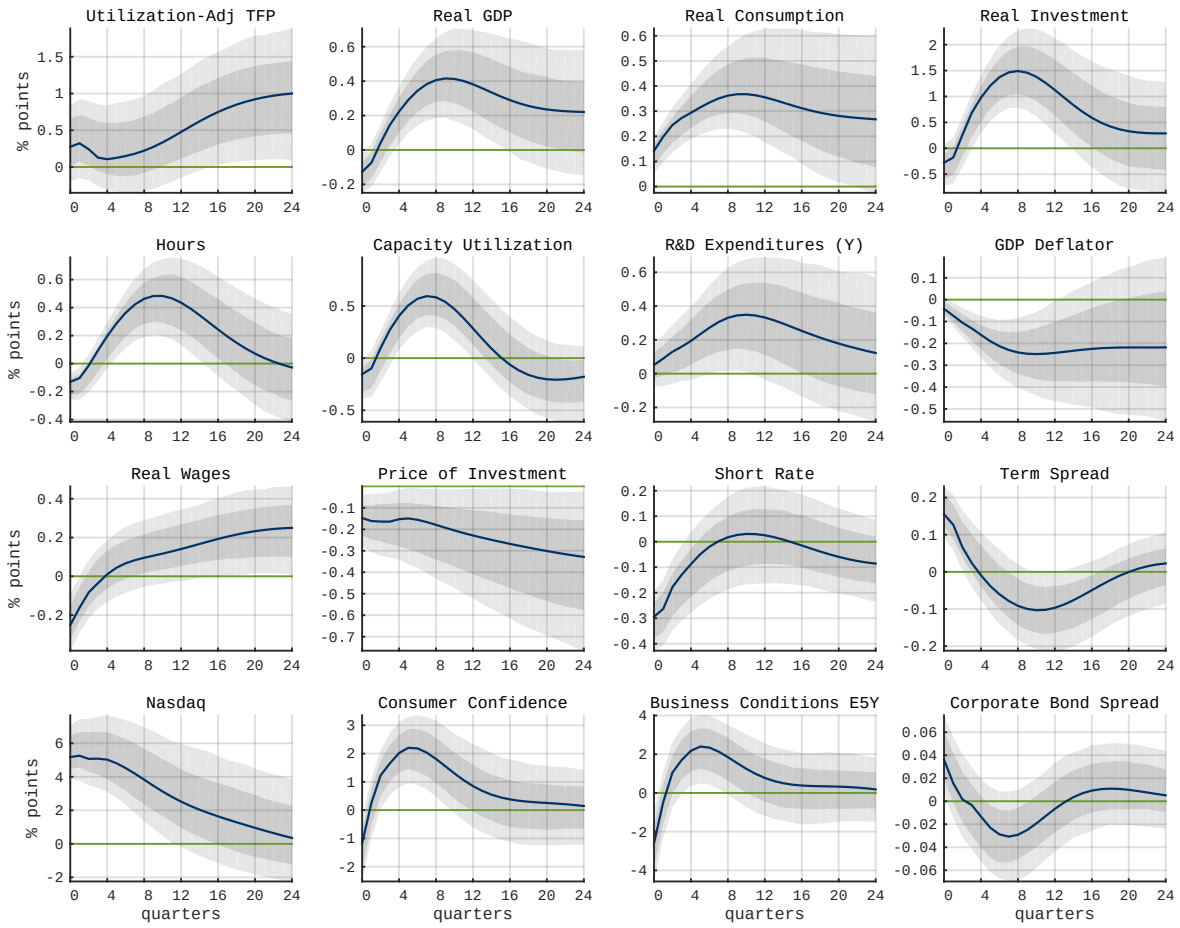

Note: Modal response to a technology news shock identified with patent-based external instrument. VAR(4). Estimation sample 1971-I:2016-IV. Identification sample 1982-I:2006-IV. Shaded areas denote 68% and 90% posterior credible sets.

of the data and transformations is reported in Appendix B. The 16-variable VAR(4) is informationally sufficient.[23]

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

[23]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.

Productivity & Quantities Most of the considerations made in the previous section carry through in the larger VAR. The initial response of TFP is muted, and the response becomes significant only years after the shock hits. Conversely, consumption rises immediately, 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. 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% in quarterly space. Total hours worked also rise robustly at the two year horizon. Relative to the 5-variable VAR we note a few differences. First, the IRFs are generally less persistent. This is due to the unit-root in the small VAR being absorbed by 'low-frequency variables' such as capacity utilization.[24] Second, the impact responses of hours worked and output become significantly negative, although the latter is not a robust finding (see Appendix E). The immediate fall in output is also noted in Francis and Ramey (2005); Basu et al. (2006) and Barsky and Sims (2011).

While the responses are somewhat delayed, also in the larger VAR they are broadly consistent with positive technology news prompting an expansionary business cycle phase whereby all macroeconomic aggregates are significantly higher at the two-year mark, and long before any material increase in TFP is recorded. The sluggish response of R&D expenditures (as a component of output) is also in line with this interpretation. In this sense, these results align with the 'news view' of Beaudry and Portier (2006); Beaudry and Lucke (2010) according to which the economy responds to current news in anticipation of potential future technological improvements. The initial significant reduction in total hours worked, and the deterioration of labor market conditions more generally, turn out to be a crucial element in shaping consumers' expectations in reaction to news shocks, and we discuss it in greater detail below.

Prices & Wages In accordance with earlier studies, we find that technology news shocks are disinflationary (Jinnai, 2013; Kurmann and Otrok, 2014). Importantly, how-

[24] We use the quarterly series for capacity utilization distributed by FRED; substituting with the series compiled by Fernald (2014) yields the same results.

ever, and consistent with nominal rigidities preventing an immediate adjustment, we find that the response of prices is subdued initially, and only slowly builds up over time to reach a peak of about -0.3% at the two year horizon. This contrasts with the substantial impact contractions in earlier studies (see e.g. Barsky and Sims, 2011; Barsky et al., 2015). Aggregate real wages fall on impact to improve at longer horizons. Coupled with the response of aggregate prices, this points toward a short-lived decline in aggregate nominal wages. The response of the relative price of investment goods, that suffers a minor contraction on impact and keeps adjusting over time, indicates that the identified news shock makes investment goods progressively cheaper relative to consumption goods. Hence, the shock has some of the flavor of the investment-specific technological (IST) improvements of e.g. Fisher (2006) and Justiniano et al. (2010, 2011).

Financial Markets & Consumers' Expectations As in the 5-variable VAR, the stock market is quick in pricing in positive news, and jumps up strongly on impact. The response of the stock market is stronger when the Nasdaq is used compared to using more general indices such as the S&P 500. This is likely due to the Nasdaq composition being heavily skewed toward information-technology companies, presumably those mostly affected by these types of shocks over the identification sample considered (1982-I:2006-IV). That said, the overall picture does not change if we substitute in the S&P 500.[25]

The disinflationary feature of the identified shock induces a strong endogenous response of the monetary authority, that responds more than proportionally to the decline in (expected) inflation. Due to the sample considered including the zero-lower-bound (ZLB) period, we use the one-year nominal interest rate as our measure for the short-term policy rate. The one-year rate falls by about 30 basis points on impact, which is roughly the same magnitude as the peak decline of prices. This implies that shorter maturity interest rates are likely to fall by more, and hence that short-term real rates fall following the shock. The slope of the yield curve, here measured as the spread between the 10-year and the 1-year Treasury rates, rises by about 15 bps on impact, mainly driven

[25]See Figure E.4 in the Appendix. The S&P 500 is available over a longer sample, which allows us to extend back the estimation to 1962-I, date at which daily data for interest rates (DGS1 and DGS10) that enter the VAR in quarterly averages become available. We note that over this sample the magnitude of the peak responses of both prices and interest rates is larger.

by changes at the short end, and implying a 15 bps fall in long term yields. 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.[26] In turn, this can act as an amplification mechanism for the propagation of the news shock. In contrast, the response of the BAA-AAA corporate bond spread is essentially flat. In Figure E.3 in the Appendix, we verify that neither the global financial crisis nor the ZLB sample drive or affect our results.

Finally, Figure 4 reports the responses of a consumer confidence indicator and a business confidence indicator reflecting expectations about economic conditions over a horizon of 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 is consistent with consumers overweighting the responses of current economic conditions 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 below.

4.2 Variance Shares

Table 2 reports the shares of explained variation at selected frequency intervals for all variables in our VAR. Specifically, the columns in Table 2 report the share of variance (in percentage points) that is accounted for by the identified shock in the short-run (frequencies corresponding to a period between 1 and 8 quarters), over the business cycle (between 8 and 32 quarters), and in the long run (between 32 and 100 quarters).[27] Variance shares at all frequencies between 1 and 100 years are reported in Figure 5

[26] See Figure E.8 in the Appendix. This finding aligns with those in Crump et al. (2016). We use the VAR to decompose the response in the 10-year rate into its expectations and term-premium components by noting that, net of risk considerations, holding a 10-year bond should be equivalent to rolling 1-year bonds over 10 years. We calculate horizon h term premium responses as the difference between the horizon h response of the 10-year rate, and the average expected response of the 1-year rate at horizons $h, h+4, \ldots, h+36$.

[27] 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].

Table 2: 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.

for a selection of variables, and in Figure E.1 in the Appendix for the remainder of entries in our VAR. In the figure, the shaded areas highlight business cycle frequencies. The algorithm used for the decomposition builds on Altig et al. (2011) and is described in detail in Appendix C. The advantage of looking at variance decompositions in the frequency domain is that it allows us to separate among long, medium, and short-run fluctuations more clearly than a standard forecast error variance decomposition in the time domain.[28]

A few results are worth highlighting. First, similar to what found in the 5-variable VAR, the shock recovered by our IV is mostly explanatory for TFP in the very long run, where it accounts for about a third of the overall variation (Figure 5). Conversely, the contribution of the shock to high-frequency fluctuations in productivity is negligible.

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

Figure 5: Shares of Explained Variance

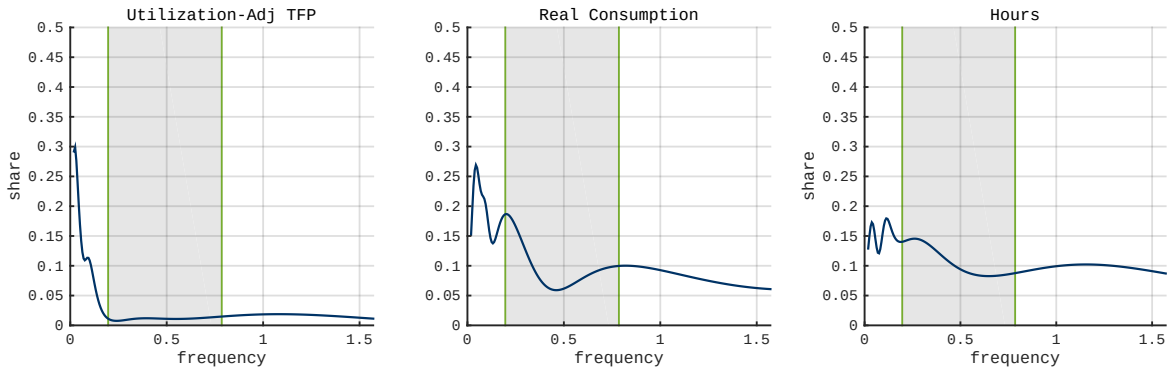

Note: Share of error variance accounted for by technology news shock identified with patent-based external instrument. VAR(4). Estimation sample 1971-I : 2016-IV. Identification sample 1982-I : 2006-IV. Shaded areas delimits business cycle frequencies (between 8 and 32 quarters). Frequencies on the x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

Hence, while we have not imposed any such restriction ex ante, the recovered shock turns out to be mostly a driver of the trend component of TFP. Second, the shock is responsible for 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. Moreover, the shock explains about 10% of the variation in output, investment, and capacity utilization over the business cycle. These shares are far from expressing the bulk of variation in these variables, but are sizeable and economically relevant, particularly in light of the relatively large size of our VAR. Hence, the recovered shock is an important source of economic fluctuations. Third, the shock explains around a fifth of the variation in the stock market at all frequencies, and is responsible for a non-trivial share of variation in short-term interest rates in the short-run (15%), and over the business cycle (10%).[29] Interestingly, the shock only accounts for about a tenth of the variation in the slope of the term structure, which contrasts with findings in Kurmann and Otrok (2013). Finally, it is worth mentioning that the shock is a significant driver of the trend variation of the relative price of investments (30% at lowest frequencies, see Figure E.1). This variable is used in Justiniano et al. (2010, 2011) to disentangle IST shocks from neutral technology shocks.

[29] 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 help price the cross-section of expected returns.

4.3 A Closer Look at the Role of the Labor Market

According to the responses in Figure 4, the immediate reaction of the labor market to technology news shocks is essentially a temporary leftward shift in the aggregate demand of labor, which results in a short-lived contraction of both hours worked and wages. In this last section, we take a closer look at the labor market response, and how it interacts with consumers' expectations.

In the VAR, we replace total hours worked with its components – the unemployment rate and the labor participation rate –, and add consumers' expectations about unemployment one year hence, again extracted from the Michigan Survey of Consumers. The question asks respondents whether they expect unemployment over the next 12 months to be higher, lower or about the same as current. All other details of the VAR specification are the same. Figure 6 collects the responses (top panels) and variance shares (bottom panels) for these three variables, full IRFs are reported in Figure E.6 in the Appendix. The chart reveals that the variation in hours worked is all accounted for by changes in the unemployment rate, while labor participation is essentially unresponsive. The unemployment rate rises on impact, to then decrease significantly at medium horizons. Notably, the shock is responsible for over 20% of the short-run fluctuations in the unemployment rate, a substantial share (see also Faccini and Melosi, 2018, for the role played by technology news on employment and its forecasts). Perhaps more interesting, however, is the response of consumers' expectations. Consistent with the rise in unemployment, and in apparent contrast with there being underlying positive news, consumers expect the unemployment rate to rise sharply, with the peak response realized well within the first year.

The context of technology news shocks offers a natural environment in which different agents in the economy are plausibly informed to different degrees. For example, it is plausible to postulate that market participants are more attentive, or more able to incorporate these types of news, than the average consumer. Here we do not attempt to speculate on the ultimate sources of such rigidities to information processing, but note that the IRFs to consumers' expectations about unemployment, and about current and expected business conditions fit nicely within the predictions of models of imperfect in-

Figure 6: Unemployment and Unemployment Expectations

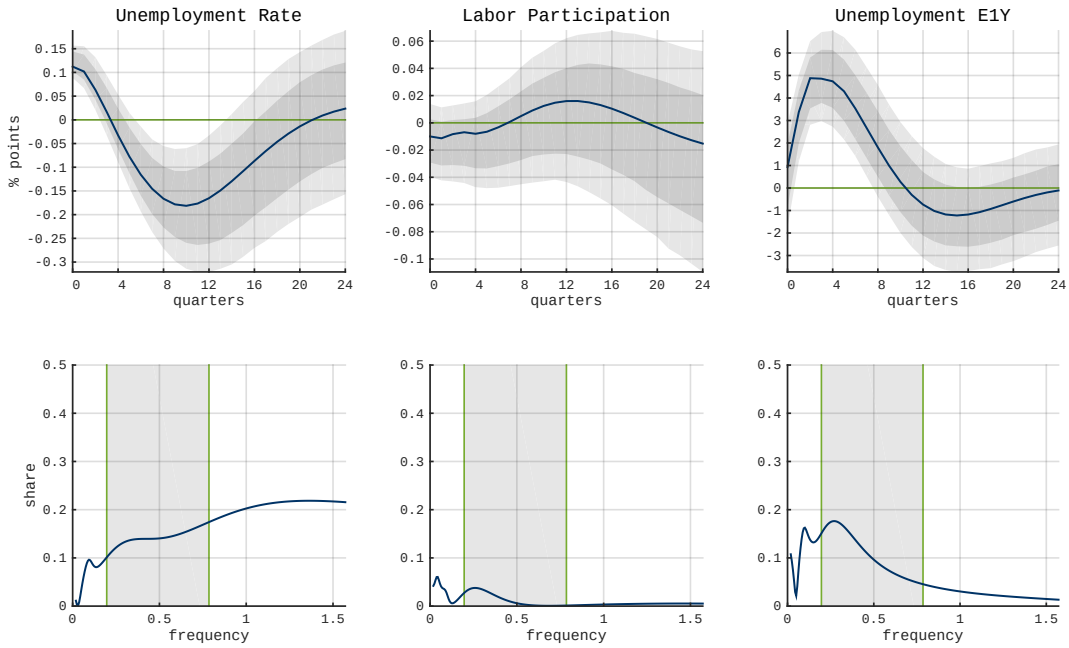

Note: Impulse response functions (top panels) and shares of explained variance (bottom panels) for the unemployment rate, the rate of labor participation, and the 1-year-ahead unemployment expectation. Survey forecasts are from the Michigan Survey of Consumers. VAR(4). Estimation sample 1971-I:2016-IV;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

formation (e.g. Woodford, 2003; Sims, 2003; Mackowiak and Wiederholt, 2009). Consider the simple framework in which agents use a Kalman Filter to form expectations about the future. The lower the signal-to-noise ratio in the information they receive, the less the new information will be weighted-in in their expectations about the future, the more these expectations will be based on current realizations/past signals (see Coibion and Gorodnichenko, 2012, 2015). News about future technological changes can be thought of as a quintessential signal extraction problem. Blanchard et al. (2013) in particular consider the case in which technology is driven by both temporary and permanent shocks (i.e. shocks that have long-lasting effects on the level of technology), and agents observe a noisy signal of the permanent component of technology. Agents are not able to disentangle news from noise. In their model the noisier the signal, the slower the consumption adjustment, the more likely that shocks to the permanent component result in an initial fall in employment.

We think of the initial rise in both actual and expected unemployment (Figure 6) as

compatible with such noise-ridden environment, and with 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 business conditions documented in Figure 4.

5 Conclusions

'How does the aggregate economy react to a shock that raises expectations about future productivity growth?' In this paper we have provided an answer to this question by proposing a novel patent-based instrumental variable that allows us to dispense from all the traditional assumptions used in the news literature.

Our IV recovers structural technology news shocks that have essentially no impact on current productivity, but are a significant driver of its trend component, and are responsible for a significant share of economic fluctuations at business cycle frequencies. We confirm many of the standard channels: positive news give rise to a sustained business cycle expansion in anticipation of future technological improvements. The stock market booms while term premia fall, acting as potentially important amplification channels. But we also unveil interesting new dynamics that suggest a central role for expectations, labor market dynamics, and their interaction. The immediate response of the labor market to technology news shocks as identified by our IV is best summarized as a leftward shift in aggregate labor demand. This is rationalized in models that embed news in frameworks in which, as is plausible, agents only observe a noisy signal about macro fundamentals. Consistently, we document that consumers' expectations only sluggishly adjust to the positive technology news, being initially dragged down by the deterioration in labor market conditions.

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

References

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.

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

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

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

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

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

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

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

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

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

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

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

Beaudry, Paul, Patrick Fève, Alain Guay, and Franck Portier (2015) "When is Nonfundamentalness in VARs a Real Problem? An Application to News Shocks," NBER Working Papers 21466, National Bureau of Economic Research, Inc.

——— (2016) "When is Nonfundamentalness in SVARs A Real Problem?," TSE Working Papers 16-738, Toulouse School of Economics (TSE).

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.

Canova, Fabio, David Lopez-Salido, and Claudio Michelacci (2009) "The effects of technology shocks on hours and output: a robustness analysis," *Journal of Applied Econometrics*, Vol. 25, No. 5, pp. 755–773.

Cascaldi-Garcia, Danilo and Marija Vukotic (2019) "Patent-Based News Shocks," The Warwick Economics Research Paper Series (TWERPS) 1225, University of Warwick, Department of Economics.

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.

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

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

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

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

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

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

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

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

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.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Kogan, Leonid, Dimitris Papanikolaou, 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 Elmar Mertens (2014) "Stock Prices, News, and Economic Fluctuations: Comment," *American Economic Review*, Vol. 104, No. 4, pp. 1439–1445, April.

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.

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.

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

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

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

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

McCracken, Michael W. and Serena Ng (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.

Plagborg-Møller, Mikkel and Christian K. Wolf (2018) "Local Projections and VARs Estimate the Same Impulse Responses," working paper, Princeton University.

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

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

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

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

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

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

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

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

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

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

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

Appendix: Not For Publication

A Additional Details on Instrument & Regression Tables

Figure A.1 plots our instrument. The grey dash-dotted line is the quarterly growth rate of patent applications pa_t. The green solid line are the residuals of Eq. (1) where there is no control for other contemporaneous policy shocks. The blue solid line is our baseline instrument, i.e. residuals of Eq. (1). Regression results are reported in the main section in Table 1.

Figure A.1: Instrument for News Shocks

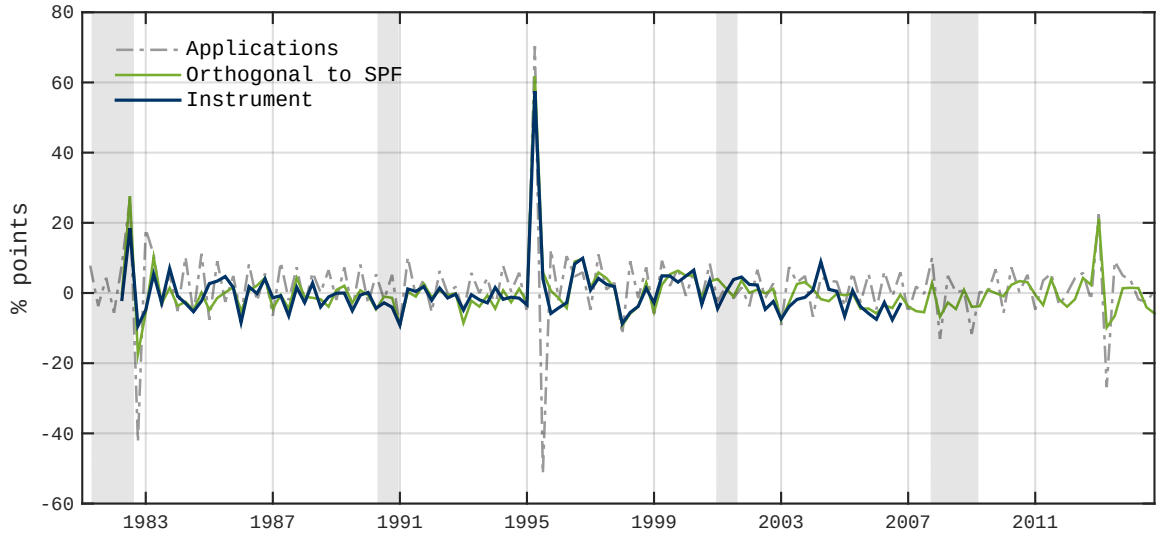

Note: Raw count of patent applications, quarterly growth rate (grey, dash-dotted line); instrument for news shocks (blue, solid), residuals of Eq. (1); residuals of Eq. (1) without policy controls, (green, solid). Shaded areas denote NBER recession episodes.

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

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

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

Table A.2: Lagged Information in Patent Applications

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

Notes: Numbers reported are Wald test statistics for joint significance of the first 4 lags of each factor F_t. The factors are extracted from the quarterly dataset of McCracken and Ng (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 and constant.

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

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

Notes: Dependent variable is the residual of Eq. (1). $\mathbb{E}_t[w_{t+h}]$ denotes SPF forecast for quarter $t+h$ published at t conditional on $t-1$. w_t 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 and constant.

Table A.4: Lagged Information in the Instrument

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

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

B Data in VAR

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

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

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

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

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

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

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

Table B.1: Variables Used

Label	Variable Name	Source	FRED Codes	TREATMENT log	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.

C Error Variance Decomposition

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

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

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

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

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

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

The spectral density of y_t is

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

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

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

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

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

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

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

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

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

Recall that the spectrum is symmetric around zero. Let the object of interest be the share of variance explained by $e_{A2,t}$ 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 $e_{A2,t}$ at business cycle frequencies is equal to

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

D Robustness & Additional Charts: 5-Variable VAR

Figure D.1 compares the IRFs retrieved by our baseline patent-based instrument with the identification schemes of Beaudry and Portier (2006), denoted 'EQY/LR', and of Barsky and Sims (2011), denoted 'Max-FEV', in the same VAR. All responses are scaled such that the peak response of TFP is equal to 1% across all identification schemes. Figure D.2 plots the share of variance that is due to $e_{A2,t}$ for all the variables included in the 5-variable VAR at all frequencies between 1 (highest frequency) and 100 (lowest frequency) years. Grey areas highlight business cycle frequencies. Figure D.3 reports for comparison the share of forecast error variance accounted for by the identified shocks in the two VARs.

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

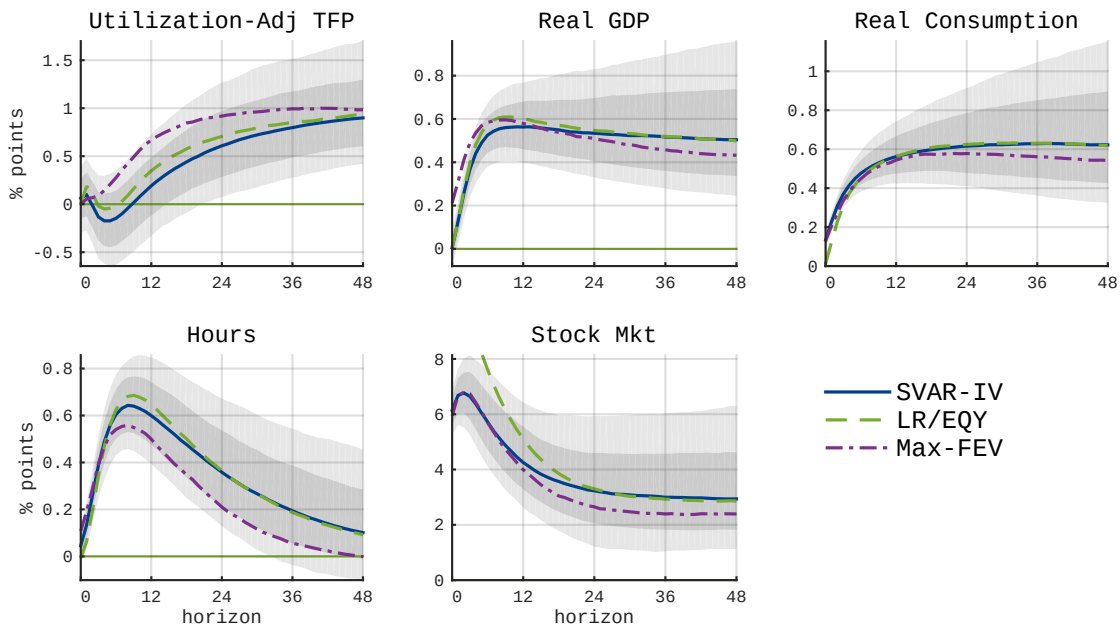

Note: Modal response to a technology news shock identified with (1) patent-based external instrument (SVAR-IV in blue), (2) long-run restrictions (LR/EQY in green dashed), and (3) maximum forecast error variance share (Max-FEV in purple dotted). Estimation sample 1971-I : 2016-IV. Identification sample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

Figure D.2: Error Variance Decomposition: Frequency, Small VAR

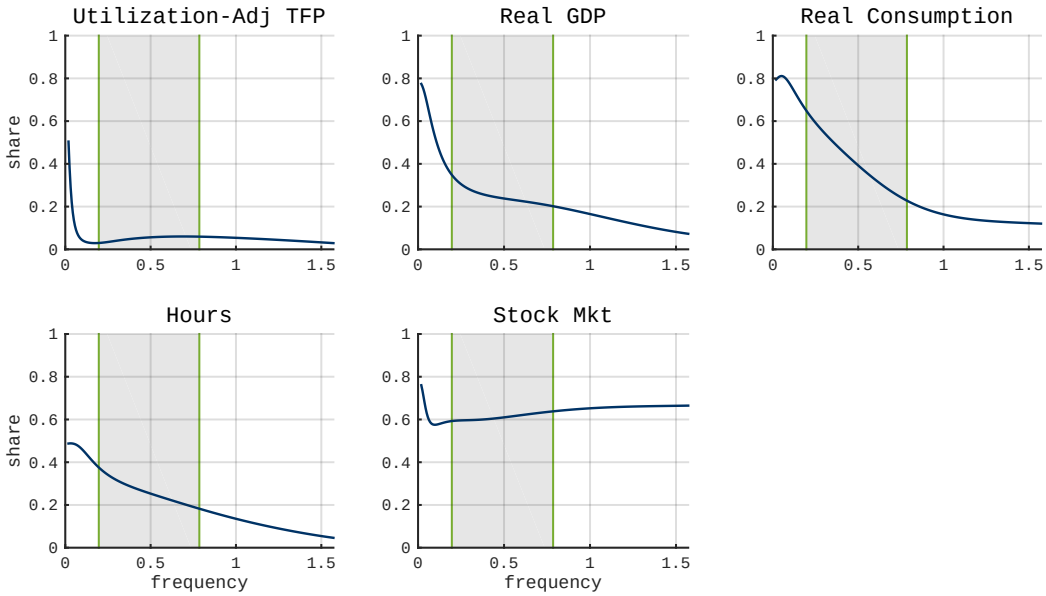

Note: Share of error variance accounted for by technology news shock identified with patent-based external instrument. VAR(4) with standard macroeconomic priors. Estimation sample 1971-I : 2016-IV; Identification sample 1982-I : 2006-IV. Shaded areas delimits business cycle frequencies (between 8 and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

Figure D.3: Forecast Error Variance Decomposition: Time, Small VAR

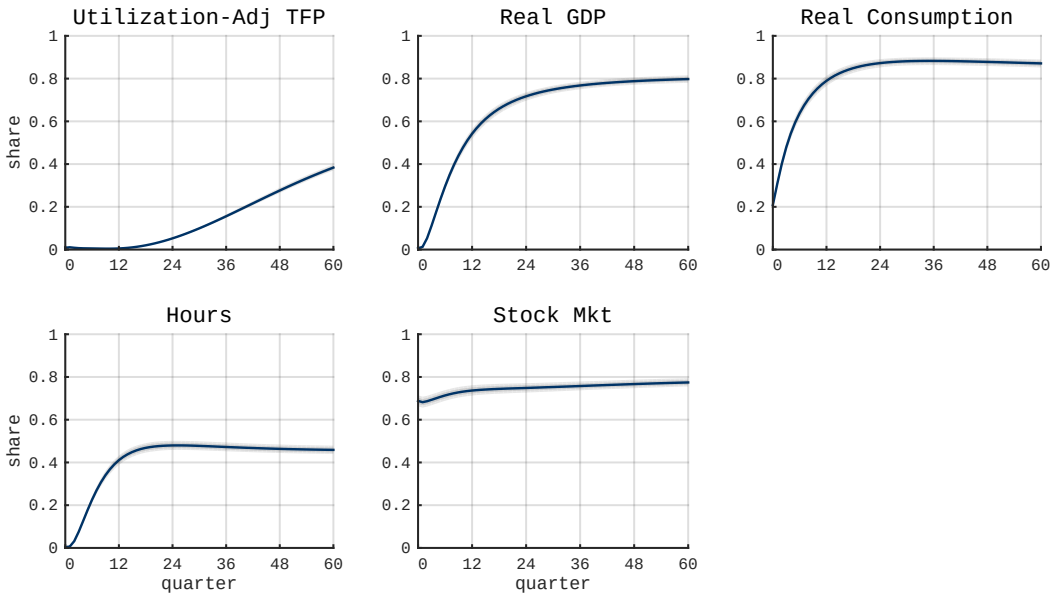

Note: Share of forecast error variance accounted for by technology news shock identified with patent-based external instrument. VAR(4). Estimation 1971-I : 2016-IV; Identification 1982-I : 2006-IV.

E Robustness & Additional Charts: Large VAR

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

All the IRFs reported in Figures E.3 to E.7 are scaled such that the peak response of utilization-adjusted TFP equals 1%.

Figure E.3 compares baseline IRFs with those obtained in a VAR that is estimated over a sample that excludes the 2008 financial crisis (estimation sample 1971-I : 2007-IV) using the same instrument (residuals of Eq. 1).

Figure E.4 compares baseline IRFs with those obtained in a VAR that is estimated over a sample that starts in 1962 (estimation sample 1962-I : 2016-IV) using the same instrument. In the VAR with the longer sample the stock market index is the S&P 500 and capacity utilization is not included due to it being available only since 1971.

Figure E.5 compares IRFs recovered by the baseline IV and the IV that does not control for contemporaneous policy shocks (SPF orthogonal).

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

Figure E.7 reports impact responses for a selection of the variables in our VAR to a contemporaneous TFP innovation that raises TFP on impact by 1%, and obtained with a standard Cholesky factorization with TFP ordered first.

Figure E.1: Error Variance Decomposition: Frequency

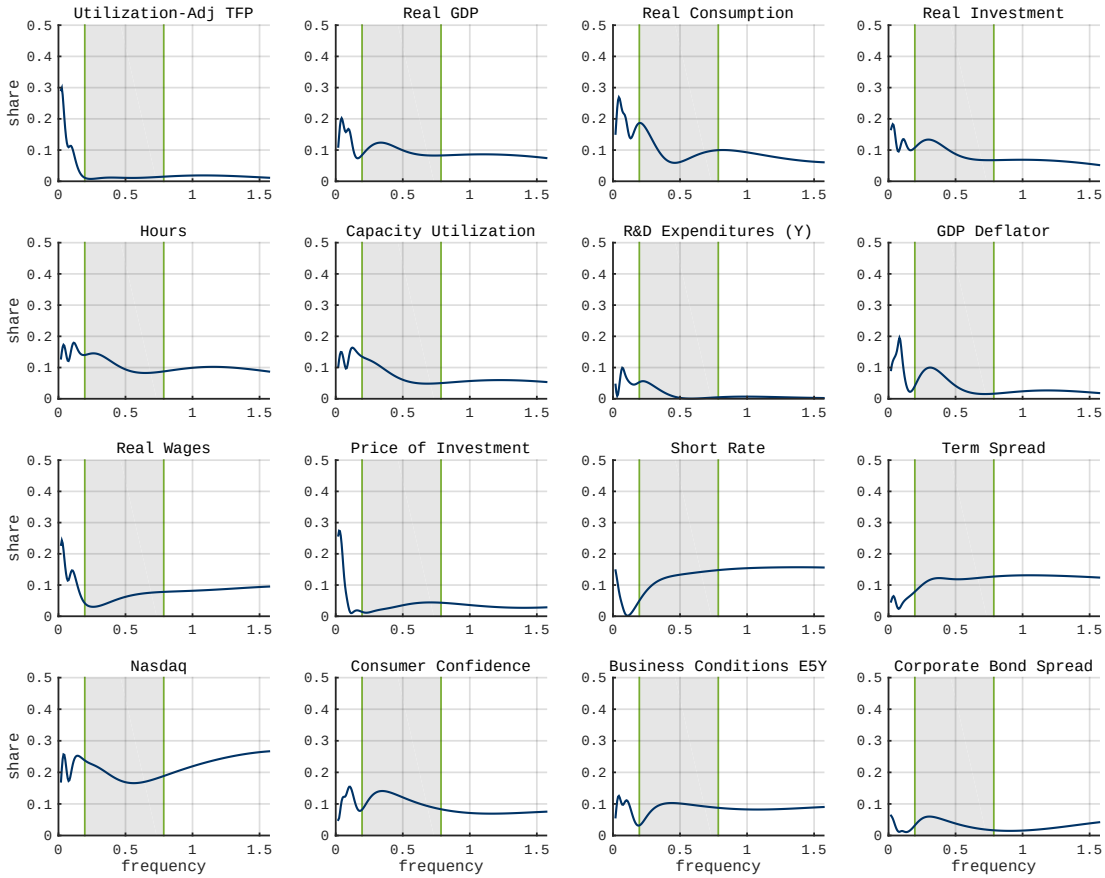

Note: Share of error variance accounted for by technology news shock identified with patent-based external instrument. VAR(4) with standard macroeconomic priors. Estimation sample 1971-I : 2016-IV; Identification sample 1982-I : 2006-IV. Shaded areas delimits business cycle frequencies (between 8 and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

Figure E.2: Forecast Error Variance Decomposition: Time

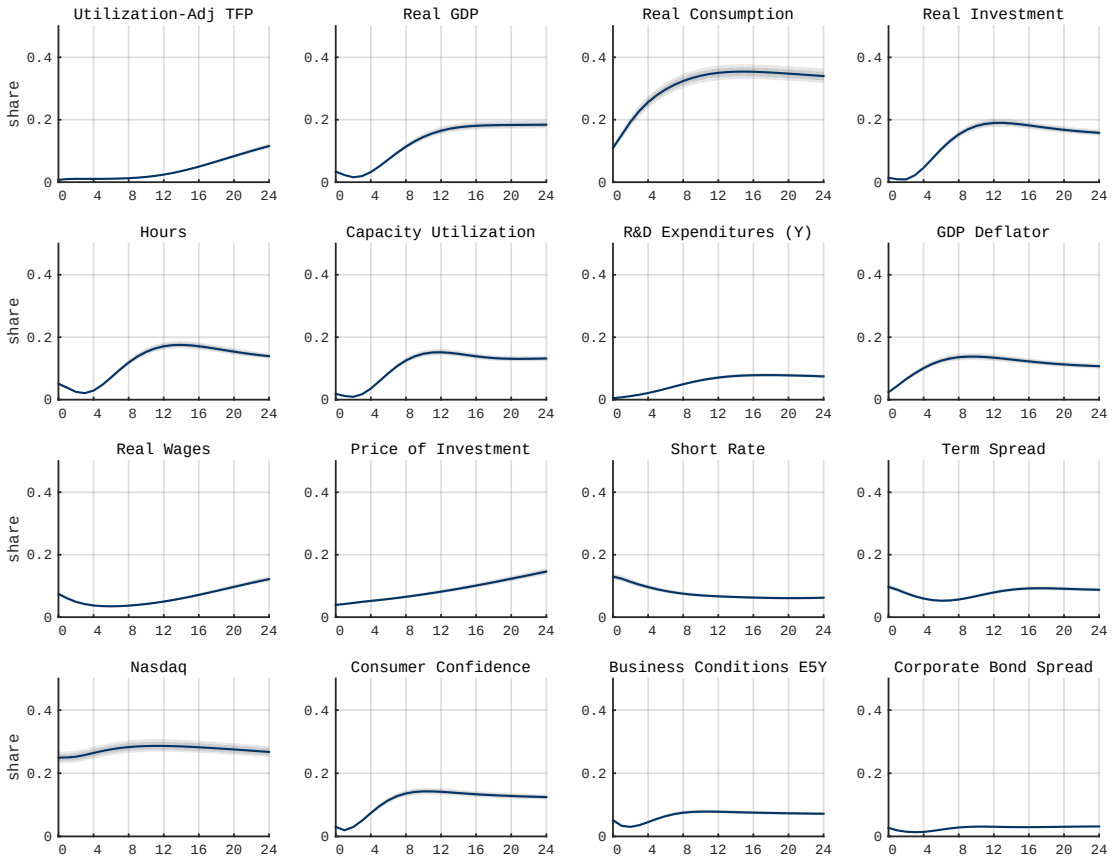

Note: Share of forecast error variance accounted for by technology news shock identified with patent-based external instrument. VAR(4). Estimation 1971-I : 2016-IV; Identification 1982-I : 2006-IV.

Figure E.3: IRFs Full vs Pre-Crisis Sample

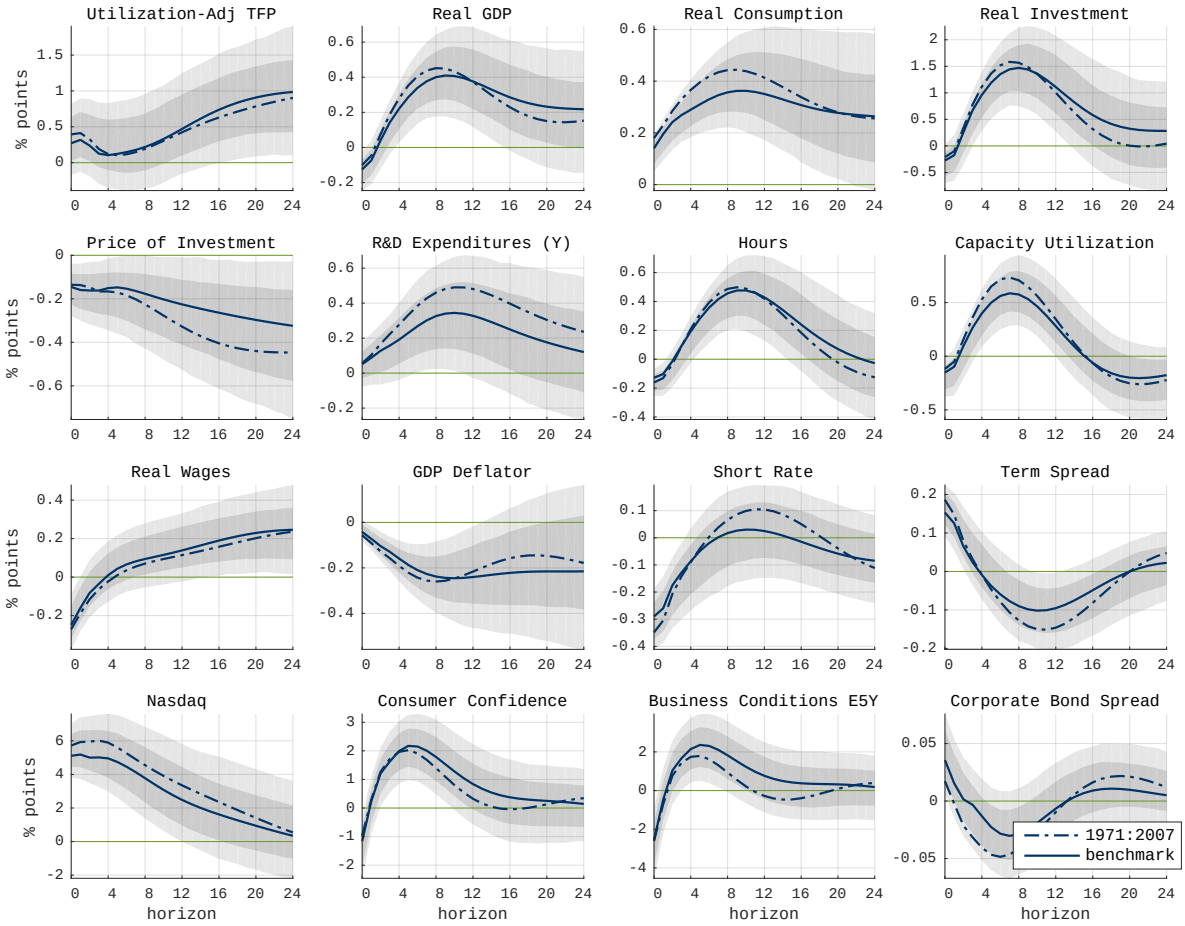

Note: Response of all variables to a technology news shock identified with patents-based external instrument. VAR(4) with standard macroeconomic priors. Solid lines: Estimation sample 1971-I : 2016-IV; Identification sample 1982-I : 2006-IV. Dash-dotted lines: Estimation sample: Estimation sample 1971-I : 2007-IV; Identification sample 1982-I : 2006-IV. Shaded areas denote 68% and 90% posterior credible sets.

Figure E.4: IRFs Longer Sample

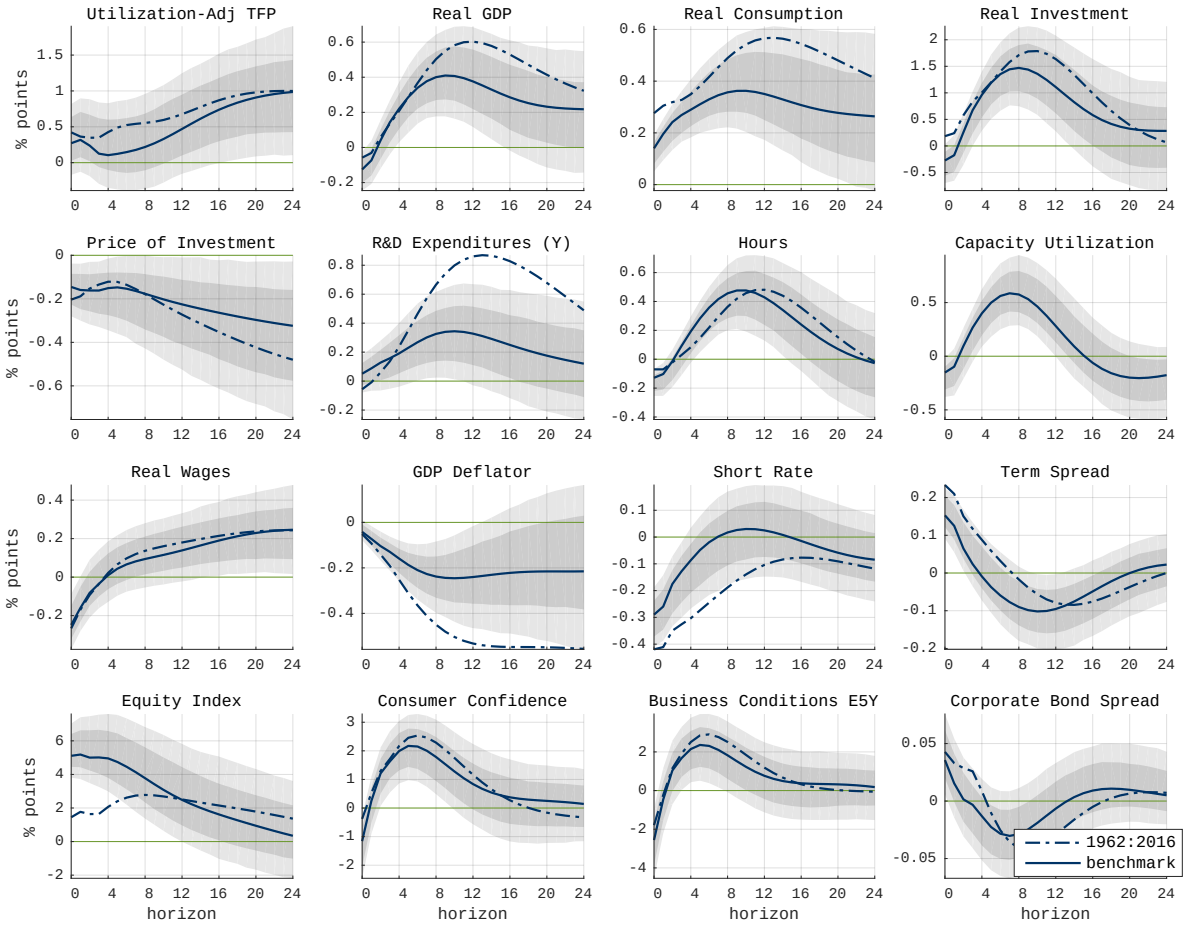

Note: Response of all variables to a technology news shock identified with patent-based external instrument. VAR(4) with standard macroeconomic priors. Solid Lines: Estimation sample 1971-I : 2016-IV; Identification sample 1982-I : 2006-IV. Dash-dotted lines: Estimation sample 1962-I : 2016-IV; Identification 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 credible sets.

Figure E.5: IRFs: Benchmark and SPF Orthogonal Instruments

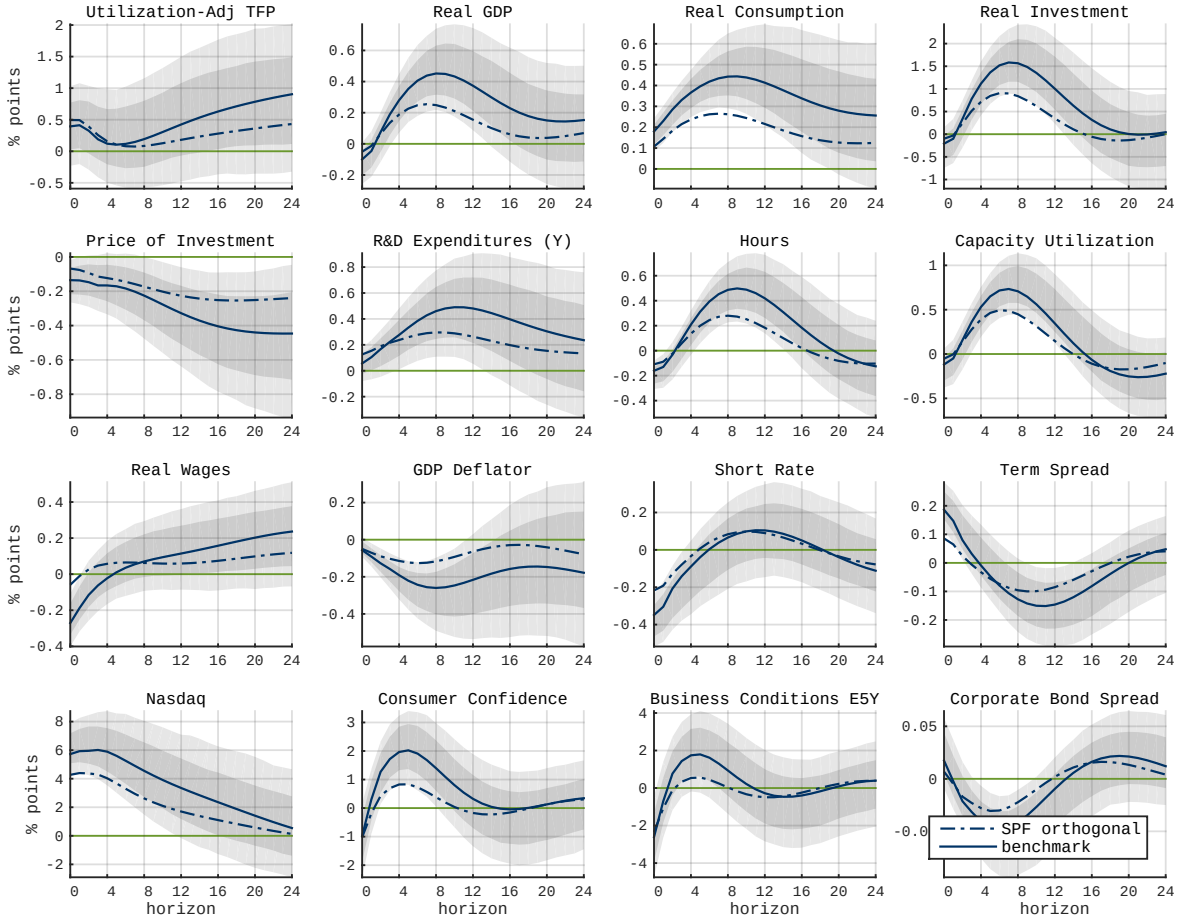

Note: Response of all variables to a technology news shock identified with patent-based external instrument. VAR(4) with standard macroeconomic priors. Solid Lines: Estimation sample 1971-I : 2007-IV; Identification sample 1982-I : 2006-IV. The instrument controls for contemporaneous policy shocks. Dash-dotted Lines: Estimation sample 1971-I : 2007-IV; Identification sample 1982-I : 2007-IV. The instrument does not control for contemporaneous policy shocks. Shaded areas denote 68% and 90% posterior credible sets.

Figure E.6: IRFs with Unemployment Expectations

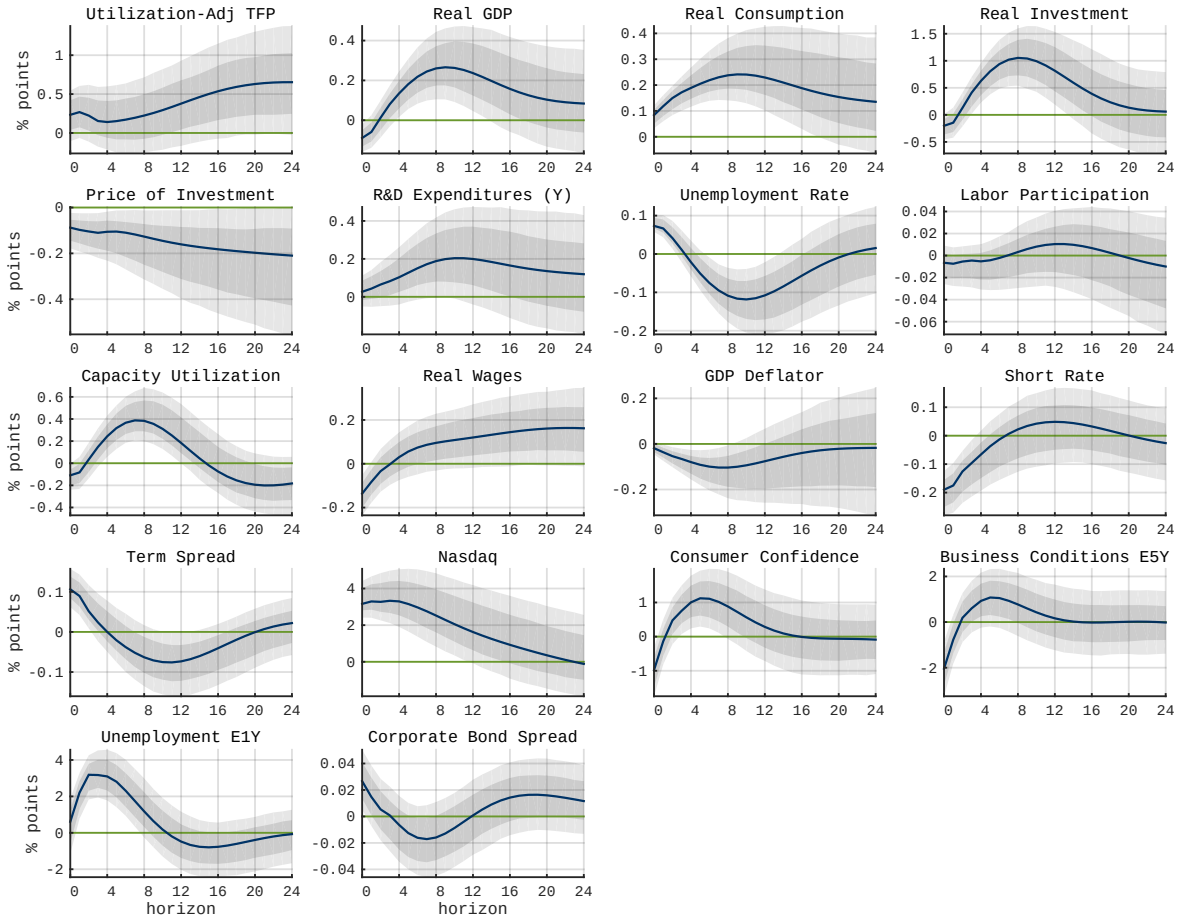

Note: Response of all variables to a technology news shock identified with patent-based external instrument. VAR(4) with standard macroeconomic priors. Instrument controls for contemporaneous policy changes. Estimation sample 1971-I : 2016-IV; Identification sample 1982-I : 2006-IV. Shaded areas denote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

Figure E.7: Impact Responses to a Contemporaneous TFP Innovation

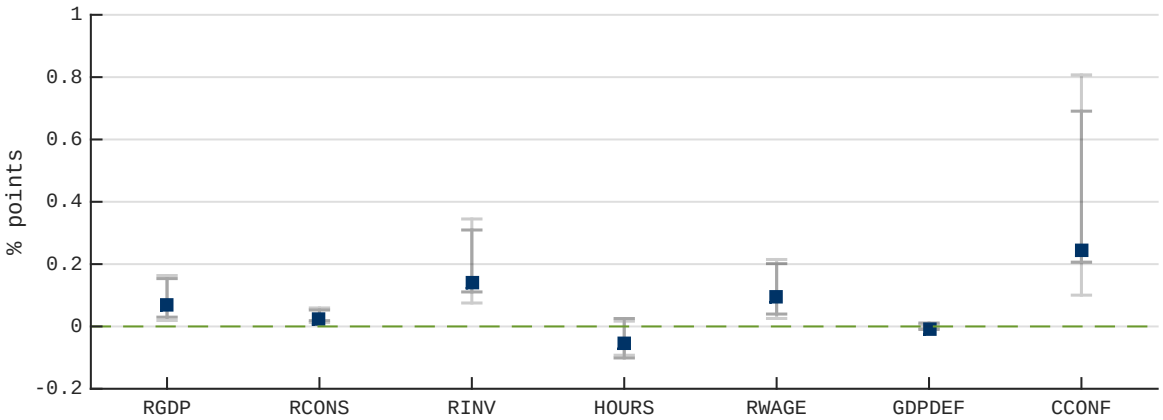

Note: Impact responses of selected variables to a TFP innovation that increases Utilization-Adjusted TFP by 1%. VAR(4). Estimation sample 1971-I:2016-IV. Grey bars delimit 68% and 90% posterior credible sets.

Figure E.8: Long Rate Response

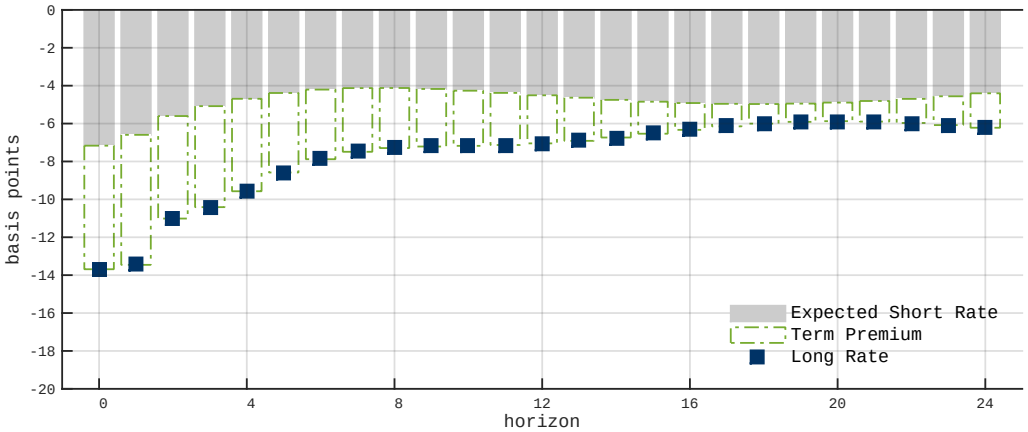

Note: Implied modal responses of the 10-year Treasury yield and VAR-based expectation and term premium components. VAR(4). Estimation sample 1971-I:2016-IV; Identification sample 1981-1:2006-IV.