Revisions in Utilization-Adjusted TFP and Robust Identification of News Shocks*

André Kurmann
Drexel University

Eric Sims
University of Notre Dame
& NBER

August 1, 2019

Abstract

This paper documents large revisions in a widely-used series of utilization-adjusted total factor productivity (TFP) by Fernald (2014) and shows that these revisions can materially affect empirical results about the macroeconomic effects of news shocks. We trace these revisions to changes in estimated factor utilization that are evocative of cyclical measurement issues with productivity. We propose an alternative identification that is robust to these measurement issues including the revisions in Fernald's series. When applied to U.S. data, the shock predicts delayed productivity growth while simultaneously generating strong impact responses of novel indicators of technological innovation and forward-looking information variables. The shock does not lead to comovement in macroeconomic aggregates as typically associated with business cycles.

JEL Classification: E22, E23, E32, O47

Keywords: Total factor productivity, variable utilization, news shocks

*Kurmann: Drexel University, LeBow College of Business, School of Economics, 3220 Market Street, Philadelphia, PA 19104 (email: kurmann.andre@gmail.com); Sims: University of Notre Dame, 3060 Jenkins Nanovic Hall, Notre Dame, IN 46556 (esims1@nd.edu). First version: August 8, 2016. This paper combines the previous drafts by Sims ("Differences in Quarterly Utilization-Adjusted TFP by Vintage, with an Application to News Shocks", March 2016) and Kurmann and Otrok ("New Evidence on the Relationship between News Shocks and the Slope of the Term Structure", June 2016). We are grateful to John Fernald for helpful conversations and for generously sharing his code and data. We also thank Chris Otrok for earlier involvement in the project, Susanto Basu and Silvia Miranda-Agrippino for thoughtful discussions, as well as Rudi Bachmann, Deokwoo Nam, Yuriy Gorodnichenko, several anonymous referees, and many seminar participants for comments.

1 Introduction

Dating back to Pigou (1927), economists have argued that changes in expectations about future fundamentals are an important source of economic fluctuations. This view has reemerged recently in part due to an influential paper by Beaudry and Portier (2006) who report that news shocks about future productivity are closely related to innovations driving long-run variations in productivity and constitute one of the main drivers of business cycles. While the importance of news shocks for business cycle fluctuations remains debated, the main identifying restriction behind news shocks is almost universally accepted: productivity reacts to news shocks only with a delay.[1]

In this paper, we critically revisit the zero impact restriction. In spite of thoughtful attempts to control for unobserved input utilization, we argue that popular empirical measures of productivity are nevertheless likely to be confounded by business cycle fluctuations due to imperfect measurement of factor utilization. As a result, news shock identifications that rely on short-run restrictions, in particular the zero impact restriction, can produce misleading results. We then propose an alternative identification that is robust to cyclical mismeasurement of productivity and apply it to U.S. data.

The starting point of our investigation is the quarterly utilization-adjusted series of total factor productivity (TFP) constructed by Fernald (2014) that has become the main measure of productivity in the news literature. Fernald frequently revises the adjusted TFP series based on new data and methodological changes. We document that a switch in detrending methods in the estimation of utilization significantly changes the cyclical properties of this series. The sensitivity of adjusted TFP to a seemingly small change such as this suggests – as acknowledged by Fernald (2014) but otherwise mostly ignored by the literature – that measurement issues with productivity can be quantitatively important.

To assess the consequences of Fernald's revisions for news shock identification, we redo the estimation of Barsky and Sims (2011), which has emerged as one of the most popular identification approaches in the literature. Based on pre-revision vintages of Fernald's adjusted TFP series, a positive news shock leads to a jump in consumption on impact but an initial decline in hours worked. As a result, the implied conditional correlation of consumption growth with hours growth

[1] See Beaudry and Portier (2014) and Barsky, Basu, and Lee (2015) for excellent reviews of this literature.

is negative, leading Barsky and Sims (2011) to conclude that news shocks do not constitute a main driver of business cycles. Based on post-revision vintages constructed with the new estimate of utilization, in contrast, a positive news shock leads to a coincident increase in consumption, hours, and other real aggregates, thereby affording a news-driven interpretation of business cycles as proposed by Beaudry and Portier (2006).

To interpret these results and illustrate the consequences of productivity mismeasurement for news shock identification more generally, we consider a medium-scale New Keynesian business cycle model that allows for multiple sources of unobserved factor utilization. Under certain conditions, Fernald's estimate of utilization coincidences with factor utilization in the model and adjusted TFP provides an almost perfect measure of true productivity. But under alternative yet equally plausible conditions, Fernald's estimate of utilization and therefore productivity is confounded by substantial cyclical mismeasurement. We conduct Monte-Carlo simulations to study the quantitative significance of this mismeasurement. The main insight from these simulations is that identifications relying on short-run restrictions and in particular the zero impact restriction can be highly sensitive to differences between factor utilization in the model and its estimation by the econometrician. Since factor utilization is not observed directly in the data and different assumptions about factor utilization are difficult to test, a more fruitful approach consists instead of devising alternative identification restrictions that are robust to cyclical mismeasurement.

In the last part of the paper, we propose such an alternative identification. Building on the premise by Beaudry and Portier (2006) that news shocks capture information about slowly disseminating changes in technology and economic organization that drive long-run productivity, we extract the innovation that accounts for the maximum forecast error variance (FEV) share of adjusted TFP at a long but finite horizon.[2] This "max-share" approach, which builds on work by Uhlig (2003), has been used previously by Francis et al. (2013) to identify long-run technology shocks. We differ in that we apply it to adjusted TFP instead of labor productivity and that we propose it as a possible news identification. Conceptually, the max-share identification is also similar to Barsky and Sims (2011) and many close variants in the news literature, with the crucial difference, however, that it does not impose the zero impact restriction and, more generally, does

[2] The idea that new technologies diffuse slowly finds ample support in a large micro-empirical literature. See for example Griliches (1957), Mansfield (1961), Mansfield (1989), Gort and Klepper (1982), and Rogers (1995).

not rely on short-run fluctuations in productivity. The max-share identification should therefore be more robust to cyclical mismeasurement of productivity, and we verify this through Monte-Carlo simulations with our model.[3]

Of course, nothing guarantees that the max-share identification captures news shocks as opposed to other shocks driving future productivity. However, when applied to U.S. data, we find compelling evidence in favor of a news interpretation. The shock has no significant impact on adjusted TFP for several quarters but predicts sustained future productivity growth, accounting for 70 percent or more of TFP fluctuations at long forecast horizons. More importantly, the shock is associated with large impact responses of two novel indicators of innovation – an index of books published in the fields of technology by Alexopoulos (2011) and an index of technological standardization by Baron and Schmidt (2015) – followed by a hump-shaped increase in R&D expenditures and a gradual decline in the relative price of investment goods. Third, the shock generates strong positive immediate reactions of forward-looking information variables.

In terms of macroeconomic implications, the max-share identification implies very similar impulse responses as the ones originally reported in Barsky and Sims (2011), with the important difference that all the results are robust to the revisions in Fernald's adjusted TFP series. Consumption increases on impact of the shock and then gradually rises further to a new permanent level, while hours worked initially decline and later increase in a hump-shaped pattern before returning to the pre-shock level. The shock therefore implies a negative correlation between consumption growth and hours worked, which makes it an unlikely source of business cycle fluctuations. Nevertheless, the shock accounts for a large share of macroeconomic fluctuations at medium and longer horizons and generates sharp impact responses of inflation and asset prices.

Relation to the literature. The main lesson of the paper is that cyclical measurement issues can materially affect the identification of news shocks based on short-run restrictions. While measurement error occupies a central role in many fields of economics, it has generally taken a back seat in quantitative macroeconomics. A notable exception is Christiano et al. (2004) who argue, as we do, that adjusted TFP may be confounded by measurement error.[4] They then apply

[3]More generally, the max-share identification is also robust to whether whether innovations to expected future productivity have an immediate impact on (true) productivity or not. Such a situation occurs, for instance, when the successful adoption of a new technology by a firm raises current productivity and simultaneously provides public information that other firms adopt the same technology in the future.

[4]Chang and Li (2018) is another more general investigation about the sensitivity of recent research results to

the infinite-horizon strategy of Gali (1999) to identify long-run productivity shocks based on the assumption that measurement errors in adjusted TFP are transient. Our paper differs in important aspects from Christiano et al. (2004) and the related literature on long-run productivity shocks. First, the max-share identification proposed here does not impose that technology is the *only* source of long-run fluctuations in productivity and instead extracts the shock that accounts for the maximum FEV share of adjusted TFP at a long but finite horizon. The max-share approach therefore affords the possibility that other shocks (e.g. a surprise productivity shock) exert long-lasting effects on adjusted TFP and at the same time addresses the criticism that infinite-horizon restrictions imply potentially large biases in finite-order VARs.[5]

Second, the literature on long-run productivity shocks typically uses average labor productivity as the technology measure and is primarily concerned with the dynamics of hours worked in response to a shock.[6] As such, this literature does not directly relate to the news literature and the idea that improvements in technology disseminate slowly and in a predictable manner. Indeed in many cases – including the max-share implementation by Francis et al. (2013) on which our identification is based – labor productivity jumps immediately because hours fall on impact thus implying increased capital deepening.[7] This may lead to the inadvertent conclusion, often imposed in DSGE models, that technology follows a random walk process, which is very different from the results obtained in our paper.

Within the extensive VAR literature on news shocks, our paper is perhaps most closely related to the one by Barsky et al. (2015). They identify a news shock by imposing a longer-run restriction that is conceptually similar to the max-share approach proposed here but differs in potentially important details.[8] Using a pre-revision vintage of Fernald's adjusted TFP series, they find that

measurement error in gross domestic product.

[5]Bias-reduction in finite-order VARs is the main motivation of Francis et al. (2013) for the max-share identification. Also see Erceg et al. (2005); Christiano et al. (2006); and Chari et al. (2008) for important contributions in this respect. Another practical advantage of the max-share approach is that it can be implemented either with a VAR in levels that includes non-stationary variables, as we do, or a stationary VAR. In contrast, the infinite-horizon approach of Gali (1999) requires the VAR to be stationary, which implies that the researcher needs to take a stand on various cointegration restrictions that can affect the results in important ways.

[6]Aside from Christiano et al. (2004), one other exception is Chen and Wemy (2015) who, like us, apply the max-share approach to adjusted TFP. However, they do not investigate the robustness of the approach to revisions in adjusted TFP nor whether the resulting shock is a news shock.

[7]We confirm this result in our VAR specification. See the discussion in Section 5 for details.

[8]The identification of Barsky et al. (2015) extracts the shock that accounts for *all* of the forecast revision of adjusted TFP at some long but finite horizon subject to the zero impact restriction, although their results would be very similar if this restriction was not imposed. We prefer the max-share approach because, as noted above,

their news shock looks quite similar independent of whether they impose the zero impact restriction or not. Our results confirm their finding in the sense that the initial response of adjusted TFP to the proposed max-share shock is small and insignificantly different from zero for several quarters. Our contribution relative to the paper by Barsky et al. (2015) and the rest of the news literature is to document the large revisions in Fernald's utilization-adjusted TFP series and to show that these revisions can materially affect empirical conclusions about the effects of news shocks based on identifications in which short-run restrictions play an important role.[9] We propose the max-share approach as an alternative identification of news shocks and show that it is robust to measurement issues; and we go to considerable length to establish the news content of the extracted shock by relating it to measures of technological innovation and forward-looking information variables.

The idea that the slow dissemination of technology implies predictable long-run changes in productivity relates to a recent (non-news) literature on the macroeconomic effects of persistent productivity growth processes. Rotemberg (2003) discusses extensively the available evidence on the slow dissemination of technology and proposes a model in which random technological progress leads to stochastic variations in long-run output while deviations of output from trend are mostly driven by temporary shocks. As in our empirical investigation, he finds that slowly diffusing technical progress leads to a temporary drop in hours worked and economic activity.[10] Lindé (2009) incorporates autocorrelated shocks to the growth rate of productivity into an otherwise standard RBC model. Autocorrelated shocks to productivity growth have the flavor of news, and he shows that incorporating this feature can help reconcile the RBC model with empirical results on the effects of technology shocks on hours worked.

---

it does not impose that news shocks are the only source of predictable fluctuations at that particular horizon. Furthermore, our Monte-Carlo simulations reveal that imposing the zero impact restriction can have important consequences even if without this restriction, the impact response of adjusted TFP is close to zero.

[9] In contemporaneous work, Cascaldi-Garcia (2017) also points out that revisions in Fernald's adjusted TFP series affect the macroeconomic implications of news shocks based on the Barsky and Sims (2011) identification. The paper does not document the source of these revisions in detail, nor does the paper discuss why these revisions raise questions about the zero impact restriction imposed by the news literature. Instead, the paper is intended as a comment on Kurmann and Otrok (2013) to which Kurmann and Otrok (2017b) respond using the alternative identification approach proposed here.

[10] Other papers that document the slow diffusion of technology and build models of costly adoption are Comin and Gertler (2006), Comin and Hobijn (2010), or Comin, Gertler, and Santacreu (2009).

2 Revisions in Utilization-Adjusted TFP

Following the lead of Kydland and Prescott (1982) and Long and Plosser (1983), the business cycle literature has typically measured productivity as the residual of aggregate output not accounted for by capital and labor inputs, commonly known as TFP. Economists quickly realized, however, that TFP may be a poor proxy of technology for a variety of reasons, most notably changes in unobserved factor utilization. In response to these concerns, Basu et al. (2006) construct an aggregate measure of productivity that takes into account sectoral heterogeneity, imperfect competition, compositional changes in the quality of labor and capital, and unobserved factor utilization. Fernald (2014) extends the analysis of Basu et al. (2006), which is carried out with annual data, to construct a quarterly measure of TFP. Because of the higher frequency, not all of the corrections in the original Basu et al. (2006) series can be implemented, but perhaps the most important one – the adjustment for variable factor utilization – is.

In what follows, we briefly review the construction of Fernald's utilization-adjusted TFP series. We then document how seemingly small changes in the estimation of factor utilization lead to large revisions in utilization-adjusted TFP that materially affect its business cycle properties.

2.1 Fernald's utilization-adjusted TFP series

Fernald's series of utilization-adjusted TFP is based on the assumption that there exists an aggregate production function

$$Y_t = F(e_t L_t, z_t K_t, A_t), \tag{1}$$

where Y_t denotes output, L_t labor input (the product of average hours per worker, h_t, with employment, N_t), K_t capital input, e_t labor effort, z_t capital use, and A_t technology that should be understood broadly as a shifter of the production function.

Differentiating (1) with respect to time and further assuming constant returns to scale as well as price-taking by firms in perfectly competitive input and output markets, cost-minimization implies that technology growth can be expressed as

$$\frac{\dot{A}_t}{A_t} = \left(\frac{\dot{Y}_t}{Y_t} - \omega_{L,t}\frac{\dot{L}_t}{L_t} - \omega_{K,t}\frac{\dot{K}_t}{K_t}\right) - \left(\omega_{L,t}\frac{\dot{e}_t}{e_t} + \omega_{K,t}\frac{\dot{z}_t}{z_t}\right), \tag{2}$$

where $\omega_{L,t}$ denotes the cost share of labor and $\omega_{K,t}$ the cost share of capital, which under constant returns to scale equals $(1-\omega_{L,t})$. The term in the first parenthesis is typically referred to as TFP growth, and the term in the second parenthesis as the change in factor utilization.

Fernald constructs TFP growth from quarterly NIPA and BLS data as

$$\Delta \ln TFP_t = \Delta \ln Y_t - \omega_{L,t}\Delta \ln L_t - (1-\omega_{L,t})\Delta \ln K_t, \tag{3}$$

with output growth measured as the log change in the equally weighted average of real expenditures and income in the business sector; and labor and capital growth built up from quality-adjusted series of different labor and capital types. To adjust for variable labor effort and capital use, which are not directly observed in the data, Fernald follows Basu et al. (2006) and proxies the change in factor utilization by a weighted change in industry hours per worker; i.e.

$$\varDelta \ln \hat{u}_t = \sum_i \kappa_i \widehat{\beta}_i \varDelta \ln h^c_{it}, \tag{4}$$

where h^c_{it} denotes a measure of hours per worker in industry i discussed further below, κ_i the industry weights, and $\widehat{\beta}_i$ the industry-specific factors of proportionality estimated using demand-side shocks as instruments.[11] The idea behind this proxy is that industry capital stocks and employment are quasi-fixed (i.e. subject to adjustment costs) but hours per worker, labor effort and capital use can be adjusted costlessly. Under certain conditions – reviewed in detail in Section 4 – optimal firm behavior then implies that utilization (i.e. the weighted sum of labor effort and capital use) is proportional to hours per worker.

Given (3) and (4), utilization-adjusted TFP

$$\Delta \ln TFP^u_t = \Delta \ln TFP_t - \Delta \ln \hat{u}_t \tag{5}$$

[11] See Fernald (2014) for details on the data and instrumental variable estimation procedure.

provides an empirical estimate of aggregate technology growth as defined in (1). As explicitly acknowledged in Fernald (2014), *"...with markups, possibly heterogeneous across producers, of price above marginal cost, or with factor adjustment costs that lead the shadow cost of inputs to differ across firms...aggregate TFP and aggregate technology are not the same—even in the absence of variable factor utilization..."* Similarly, if the utilization proxy in (4) is incorrect, then this will also lead to mismeasurement. Despite these potential issues, adjusted TFP is an important benchmark and has become the primary measure of technology for business cycle macroeconomics.

2.2 Changes across vintages

Fernald regularly publishes revised estimates of adjusted TFP based on new data and methodological changes.[12] Table 1 reports key statistics for the vintages of December 2007, December 2013, May 2014, and May 2016, all over the same sample period 1947:3–2007:3.[13]

Table 1: Moments of Adjusted TFP Growth for Different Vintages

	$\Delta \ln TFP_t^{u,07}$	$\Delta \ln TFP_t^{u,13}$	$\Delta \ln TFP_t^{u,14}$	$\Delta \ln TFP_t^{u,16}$
Mean	1.49	1.41	1.42	1.42
Standard Deviation	3.41	3.30	3.79	3.46
Corr w/ $\Delta \ln TFP_t^{u,07}$	1.00	0.85	0.56	0.58
Corr w/ $\Delta \ln Y_t$	0.53	0.38	0.18	0.07
Corr w/ $\Delta \ln H_t$	-0.01	-0.06	-0.24	-0.35

Notes: $\Delta \ln TFP_t^{u,j}$ is the quarterly log change expressed in annualized percentage points of Fernald's adjusted TFP series for vintages $j = 07$, 13, 14 or 16; Y_t is real GDP; H_t is total hours worked in the non-farm business sector. All macroeconomic aggregates are from the NIPA tables and are expressed as quarterly log changes. The sample period for each of the statistics is 1947q3-2007q3.

The means and standard deviations are similar across vintages. However, there is a marked change in business cycle comovement from the 2014 vintage onward. The correlation coefficient between the 2007 vintage and post-2013 vintages of adjusted TFP growth are less than 0.6. Furthermore, while the 2007 and 2013 vintages of adjusted TFP growth are positively correlated with output growth and uncorrelated with total hours growth, the 2014 and 2016 vintages of adjusted TFP growth are uncorrelated with output growth and negatively correlated with total

[12] The different vintages of adjusted TFP as well as the underlying components are available on Fernald's website.

[13] The beginning and end of the sample is dictated by the availability of the December 2007 vintage. The results for other pre-2014 vintages are very similar to the 2007 and 2013 vintages, while the results for other post-2013 vintages are very similar to the 2014 and 2016 vintages.

hours growth. These changes in correlation across vintages of adjusted TFP occur for different subsamples and are not driven by a particular time period.

To assess what drives these changes, Figure 1 plots the log *levels* of the 2007 and 2016 vintages of adjusted TFP, unadjusted TFP, and estimated utilization.

Figure 1: Adjusted TFP, Unadjusted TFP and Utilization: 2007 vs. 2016 Vintages

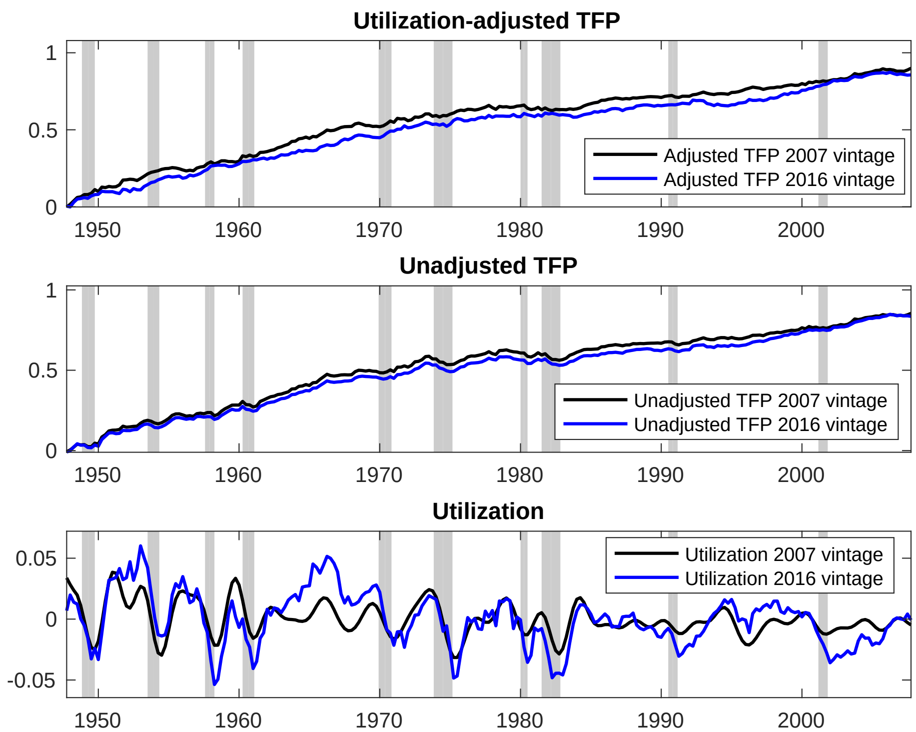

Notes: The figure plots the log levels of the 2007 and 2016 vintages of utilization-adjusted TFP (top panel), unadjusted TFP (middle panel) and estimated utilization (bottom panel). The 2007 vintages are depicted as black lines. The 2016 vintages are depicted as blue lines. The grey shaded bars show NBER recessions. The sample period for each of the graphs is 1947q3-2007q3.

As the top panel shows, the two vintages of adjusted TFP share roughly the same trend over the full sample although there are some differences over subsamples. As the middle and bottom panels show, only part of these differences are attributable to differences in non-adjusted TFP across vintages, and these differences occur mostly at low frequencies. A more important portion is due to differences in estimated utilization across vintages. While both vintages are stationary and display overall similar fluctuations, the 2007 vintage is substantially smoother and less persistent than the 2016 vintage, with the latter exhibiting particularly pronounced swings during the 1960s, the 1990s and the early 2000s.

Table 2 confirms this visual assessment. The business cycle properties of unadjusted TFP remain essentially unchanged across vintages. Variations in utilization, by contrast, are significantly larger for the 2014 and 2016 vintages and there is an important decline in correlation relative to 2007 and 2013 vintages.

Table 2: Moments of unadjusted TFP and Utilization Growth for Different Vintages

	$\Delta \ln TFP_t^{07}$	$\Delta \ln TFP_t^{13}$	$\Delta \ln TFP_t^{14}$	$\Delta \ln TFP_t^{16}$
Mean	1.42	1.37	1.37	1.39
Standard Deviation	3.75	3.55	3.55	3.55
Corr w/ $\Delta \ln TFP_t^{07}$	1.00	0.92	0.92	0.93
	$\Delta \ln \hat{u}_t^{07}$	$\Delta \ln \hat{u}_t^{13}$	$\Delta \ln \hat{u}_t^{14}$	$\Delta \ln \hat{u}_t^{16}$
Mean	-0.08	-0.04	-0.05	-0.03
Standard Deviation	2.34	2.94	3.75	3.76
Corr w/ $\Delta \ln \hat{u}_t^{07}$	1.00	0.94	0.58	0.65

Notes: This table shows descriptive statistics for the 2007, 2013, 2014 and 2016 vintages of the non-adjusted TFP and utilization series. The sample period for these statistics is fixed at 1947q3-2007q3.

This suggests that the large changes in business cycle properties of adjusted TFP are not due to revisions in non-adjusted TFP (i.e. due to data revisions or changes in NIPA methodology), but are instead driven primarily by revisions in utilization. We confirm this conjecture by combining the 2007 vintage of utilization with non-adjusted TFP from other vintages. Correlations for the resulting synthetic series of adjusted TFP are presented in the Appendix. The correlations of the 2007 adjusted TFP vintage with the synthetic 2014 and 2016 series are both 0.91, compared to 0.56 and 0.58 for the actual vintages. Hence, while revisions in utilization do not account for all of the changes in adjusted TFP, they explain the large majority.

2.3 Revisiting Fernald's estimation of utilization

What explains the changes in estimated utilization across vintages? Between December 2013 and May 2014, Fernald implemented two methodological changes. First, he switched from using estimates of industry weights and proportionality factors $\widehat{\beta}_i$ in (4) by Basu et al. (2006) to estimates from Basu et al. (2013), which are based on more recent data and a more detailed industry decomposition. Second, Fernald has to contend with the issue that hours per worker in many industries are trending over time. Up to the December 2013 vintage, Fernald follows Basu et al. (2006) and detrends industry hours per worker with the bandpass filter of Christiano and Fitzgerald

(2003) to isolate frequencies between 8 and 32 quarters. From the May 2014 vintage onward, Fernald instead detrends industry hours per worker with the bi-weight filter used in Stock and Watson (2012), which removes a much slower moving trend than the bandpass filter.

Using replication codes for the December 2013 and May 2014 vintages shared generously by Fernald, we assess the quantitative importance of the two changes. Table 3 reports the results.

Table 3: Changes in Fernald's utilization estimates

	$\Delta \ln \hat{u}_t^{13}$	$\Delta \ln \hat{u}_t^{13,BFFK}$	$\Delta \ln \hat{u}_t^{13,BW}$	$\Delta \ln \hat{u}_t^{13,BFFK\&BW}$	$\Delta \ln \hat{u}_t^{14}$
Standard Deviation	2.94	2.28	4.76	3.73	3.75
Corr w/ $\Delta \ln \hat{u}_t^{07}$	0.94	0.90	0.59	0.56	0.58

Notes: This table shows simulated utilization series based on the 2013 vintage data. See text for details. The sample period for all statistics is 1947q3-2007q3.

For comparison, the first and the last column replicate the business cycle properties of the actual December 2013 and May 2014 vintages of estimated utilization growth from Table 1. The second column, labelled $\Delta \ln \hat{u}_t^{13,BFFK}$, shows the effect of switching to the industry weights and proportionality factors from Basu et al. (2013). While this switch lowers the volatility of utilization somewhat, it leaves the correlation with the 2007 vintage essentially unchanged. As shown by the third column labelled $\Delta \ln \hat{u}_t^{13,BW}$, by contrast, changing the detrending method from bandpass filtering to bi-weight filtering leads to a substantial increase in the volatility of utilization growth and a concurrent decrease in the correlation with the 2007 vintage. Finally, as shown in the fourth column labelled $\Delta \ln \hat{u}_t^{13,BFFK\&BW}$, the two changes combined essentially replicate the 2014 vintage. The remaining small difference is due to data revisions.

The results make clear that the change in filtering of hours per worker is the main driver of the revisions in utilization growth. The bandpass filter not only removes slow moving trends but also high frequency fluctuations, resulting in a very smooth estimated utilization series. The bi-weight filter, in contrast, extracts only a slow-moving trend, implying an estimated utilization series that includes substantially more short-term fluctuations.

For the purpose of news identification that follows, the bottomline is that neither the bandpass filter nor the bi-weight filter – or any other statistical detrending method for that matter – are likely to capture secular trends in hours per worker appropriately. On the one hand, the bandpass filter may attribute too much of cyclical variations in hours per worker to the trend, thus underestimating the volatility of utilization. On the other hand, as shown in the Appendix, the bi-weight filter in the present context turns out to be almost equivalent to no filtering. The bi-weight filter may thus imply long-lasting swings in utilization that extend well beyond variations in true utilization. This means that for either filtering choice, utilization and therefore adjusted TFP are likely confounded by cyclical mismeasurement even if the conditions underlying the proportionality assumption in (4) are satisfied (a point to which we return in Section 4).

3 Implications for News Shocks Identification

Starting with Cochrane (1994), the modern macro literature has defined news shocks as information useful in predicting future fundamentals (often productivity) but unrelated to current and past fundamentals. As proposed by Beaudry and Portier (2006), this implies a zero impact restriction, which is that news shocks affect productivity only with a delay. This restriction is at the core of almost all news shocks identifications used to date.

In what follows, we use the news identification approach by Barsky and Sims (2011) to quantify the implications of the revisions in adjusted TFP for news shocks, although the lessons learned are relevant for other identification approaches relying on the zero impact restriction as well. We focus on the Barsky-Sims approach because it does not require taking a stand on the nature of non-news shocks.[14] Furthermore, Monte-Carlo simulations show that the Barsky-Sims approach performs well in small samples (provided that productivity is measured correctly). As such, the Barsky-Sims approach has emerged as one of the most commonly used identifications in the literature.

3.1 Barsky-Sims identification

The Barsky-Sims identification consists of estimating a VAR and extracting the innovation that is orthogonal to Fernald's adjusted TFP series but maximally accounts for the FEV share of adjusted TFP over a ten year horizon. Since our alternative identification proposed in Section 5

[14] The zero impact restriction is sufficient to identify news shocks in bivariate VARs. In VARs with more than two variables, additional restrictions need to be imposed. Full identification approaches that do so by taking a stand on all structural shocks affecting the VAR are often subject to important robustness issues. See for example Kurmann and Mertens (2014) who show that the identification by Beaudry and Portier (2006) does not have a unique solution in their VAR systems with more than two variables; or Fisher (2010) who shows that the results by Beaudry and Lucke (2010) depend on the choice of cointegration restrictions imposed.

is conceptually very similar, we review the details here. Let $\mathbf{Y}_t$ be a $k \times 1$ random vector process of which the first variable is a measure of productivity (e.g. Fernald's utilization-adjusted TFP), and let the reduced form moving average representation of this process be given by $\mathbf{Y}_t = \mathbf{B}(\mathbf{L})\mathbf{u}_t$, where $\mathbf{u}_t$ is a $k \times 1$ vector of prediction errors with variance-covariance matrix $E(\mathbf{u}_t\mathbf{u}_t') = \mathbf{\Sigma}_u$, and $\mathbf{B}(L) = \mathbf{I} + \mathbf{B}_1 L + \mathbf{B}_2 L^2 + \ldots$ is a matrix lag polynomial.

Now assume that there exists a linear mapping between the prediction errors and the structural shocks, $\mathbf{u}_t = \mathbf{A}\epsilon_t$, where ϵ_t is a $k \times 1$ vector of structural shocks characterized by $E(\epsilon_t\epsilon_t') = \mathbf{I}$, and $\mathbf{A}$ is a $k \times k$ matrix satisfying $\mathbf{A}\mathbf{A}' = \mathbf{\Sigma}_u$. Given the symmetry of $\mathbf{\Sigma}_u$, there are a multitude of $\mathbf{A}$ consistent with $\mathbf{A}\mathbf{A}' = \mathbf{\Sigma}_u$. The Choleski decomposition of $\mathbf{\Sigma}_u$ is one potential solution. Denote this by $\widetilde{\mathbf{A}}$. The entire set of permissible values of $\mathbf{A}$ consistent with $\mathbf{A}\mathbf{A}' = \mathbf{\Sigma}_u$ is then described by $\widetilde{\mathbf{A}}\mathbf{Q}$, where $\mathbf{Q}$ is an orthonormal rotation matrix; and the structural moving average representation is $\mathbf{Y}_t = \mathbf{C}(\mathbf{L})\epsilon_t$, where $\mathbf{C}(\mathbf{L}) = \mathbf{B}(\mathbf{L})\widetilde{\mathbf{A}}\mathbf{Q}$.

The h step ahead forecast error of $\mathbf{Y}_t$ can be written as

$$\mathbf{Y}_{t+h} - E_{t-1}\mathbf{Y}_{t+h} = \sum_{l=0}^{h} \mathbf{B}_l \widetilde{\mathbf{A}}\mathbf{Q}\epsilon_{t+h-l}. \tag{6}$$

The FEV share of variable i attributable to shock j at horizon h is then

$$\mathbf{\Omega}_{i,j}(h) = \frac{\sum_{l=0}^{h} \mathbf{B}_{i,l}\widetilde{\mathbf{A}}\gamma\gamma'\widetilde{\mathbf{A}}'\mathbf{B}_{i,l}'}{\sum_{l=0}^{h} \mathbf{B}_{i,l}\mathbf{\Sigma}_u\mathbf{B}_{i,l}'}, \tag{7}$$

where $\mathbf{B}_{i,l}$ is the ith row of lag polynomial evaluated at $L = l$ and γ is the jth column of Q.

The news shock identification of Barsky and Sims (2011) consists of picking γ to maximize the sum of FEV shares of productivity (the first variable in the VAR) up to some truncation horizon H subject to the restriction that the shock is orthogonal to current productivity. Formally

$$\max_{\gamma} \quad \sum_{h=0}^{H} \Omega_{1,2}(h) \text{ s.t. } \gamma'\gamma = 1 \text{ and } \gamma(1,1) = 0, \tag{8}$$

where without loss of generality productivity is ordered first in $\mathbf{Y}_t$ and the news shock is defined as the second shock in ϵ_t. The first restriction ensures that γ belongs to an orthonormal matrix.

The second restriction imposes that the news shock affects productivity only with a delay.

3.2 Effect of revisions on news shock identification

We apply the Barsky-Sims identification to a four-variable VAR comprised of either the 2007 vintage or the 2016 vintage of Fernald's utilization-adjusted TFP series, real personal consumption expenditures per capita, total hours worked per capita in the non-farm business sector, and inflation as measured by the growth rate of the GDP price deflator.[15] Results for larger VARs that contain additional macro aggregates are similar. With the exception of the inflation rate, the variables enter the VAR in log levels. The VAR is estimated with four lags via Bayesian methods subject to a Minnesota prior.[16] Confidence bands are computed by drawing from the resulting posterior distribution. The sample period is fixed at 1960q1-2007q3.[17] As in Barsky and Sims (2011), the truncation horizon is set to $H = 40$.

Figure 2 presents impulse responses to a news shock using the Barsky and Sims (2011) news identification. Here and below, the solid black lines show the posterior median impulse responses implied by the posterior distribution of the VAR estimated with the 2016 vintage of adjusted TFP, and the gray bands are the corresponding 16 to 84 percent posterior coverage intervals. In turn, the red dash-dotted lines show the posterior median impulse responses implied by the posterior distribution of the VAR estimated with the 2007 vintage of adjusted TFP, and the red dashed lines are the corresponding 16 to 84 percent posterior coverage intervals.

Based on the 2007 vintage of adjusted TFP, the responses are very similar to those estimated by Barsky and Sims (2011). Adjusted TFP starts to increase the quarter after the shock; consumption jumps up while inflation falls significantly on impact; and hours worked initially decline, turning significantly positive only after about 12 quarters. As shown in the Appendix, these impulse responses imply that if the economy was buffeted solely by news shocks, the correlation between

[15] VARs based on any of the pre-2014 vintages of adjusted TFP produce impulse responses that are nearly identical to those based on the 2007 vintage, while VARs based on post-2014 vintages of adjusted TFP produce impulse responses that are very similar to those based on the 2016 vintage.

[16] The Minnesota prior assumes a random walk process for adjusted TFP and consumption, and a white noise process for hours worked and the inflation rate. Estimates are robust to assuming a random walk prior for hours worked and inflation as well.

[17] The beginning of the sample is chosen to facilitate comparison with Barsky and Sims (2011) and because some additional variables of interest studied below are unavailable prior to 1960. Furthermore, the omission of the immediate post-war data from the sample removes some large influences due the 1951 Treasury Accord and Korean War. The end date is the last available observation for the 2007 vintage of adjusted TFP data.

consumption growth and hours growth would be negative, whereas in the data this correlation is robustly positive.

Figure 2: Impulse Responses to Barsky-Sims News Shock, 2007 vs. 2016 Vintage of Adjusted TFP

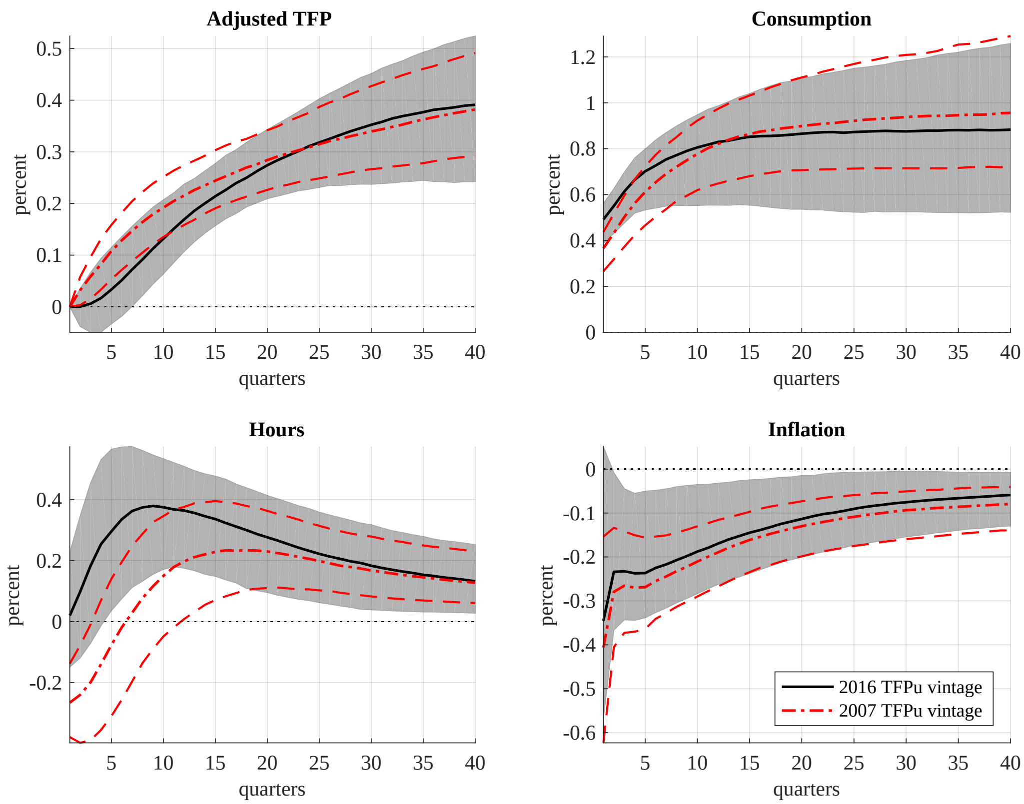

Notes: Solid black lines are the posterior median estimates from the VAR system estimated with the 2016 vintage of adjusted TFP. The gray bands correspond to the 16 to 84 percent posterior coverage intervals. The red dash-dotted lines are the posterior median estimates for the system estimated with the 2007 vintage of adjusted TFP. The red dashed lines correspond to the 16 to 84 percent posterior coverage intervals. The impulse responses are identified using the Barsky-Sims identification.

Based on the 2016 vintage of adjusted TFP, in contrast, the impulse responses look different in economically meaningful ways. Adjusted TFP reacts to the news shock only after several quarters while hours worked increase from the beginning (although not significantly so for the first few quarters), reaching peak response about 10 quarters earlier than based on the 2007 vintage. This difference in the response of hours worked implies that the correlation of consumption growth and hours growth conditional on news shocks is now significantly positive (see again the Appendix for details). Moreover, the deflationary impact of news shocks, which Barsky, Basu, and Lee (2015) cite as one of the most robust features of the data, is no longer statistically significant.

The difference in responses depending on the vintage of adjusted TFP used has important implications for the role of news shocks. Based on the 2007 vintage, the absence of comovement between hours and consumption leads Barsky and Sims (2011) to conclude that news shocks about future productivity are not a major source of business cycle fluctuations. Based on the 2016 vintage instead, the coincident increase in consumption and hours is consistent with the view espoused by Beaudry and Portier (2006) that news shocks have significant short-term demand effects and are a potentially important driver of business cycle fluctuations.

4 Interpreting the results through a DSGE Model

The results of the preceding sections illustrate that measurement issues about productivity can have important implications for news shock identifications that rely on short-run restrictions on productivity and in particular on the zero impact restriction. To interpret and better understand these results, we build a medium scale New Keynesian DSGE model and conduct different Monte Carlo simulations to address the following questions. Under what conditions does Fernald's adjusted TFP series appropriately measure technology? What are the consequences of different sources of productivity mismeasurement for news shock identification? Is Fernald's new bi-weight filtered estimate of utilization preferable to the previously used bandpass filtered estimate?

The DSGE model we use is based on Christiano et al. (2005), Smets and Wouters (2007) and Justiniano et al. (2010) but augmented with variable labor effort and hours worked so as to analyze the conditions under which the proportionality result between utilization and hours worked set forth in Basu et al. (2006) holds. The model abstracts from heterogeneity in production across industries and absence of an aggregate production function. While these considerations are potentially important, abstracting from them does not invalidate the measurement issues highlighted here.

4.1 Model

To save on space, we focus primarily on the components of the model that relate to the measurement of technology and utilization. A full description of the equilibrium conditions is provided in the Appendix.

The model is populated by intermediate goods producer, a representative final goods producer, a representative household, labor unions, a labor packer, and a monetary authority. Intermediate goods producers are indexed by $i \in [0\ 1]$ and produce output with

$$Y_t(i) = A_t \left(K_{s,t}(i)\right)^{\alpha} \left(L_{s,t}(i)\right)^{1-\alpha} - FX_t, \tag{9}$$

where A_t denotes exogenous technology (common across firms), $K_{s,t}(i)$ capital capital services, $L_{s,t}(i)$ labor services, and $FX_t \geq 0$ is a fixed cost that increases with the economy's trend X_t. The output is sold at price $P_t(i)$ to the final goods producer who uses it to produce final output Y_t with a CES technology with elasticity of substitution ϵ_p across the different intermediate goods. As in Calvo (1983), intermediate goods producers can reoptimize price $P_t(i)$ with fixed probability $1-\theta_p$ per period and otherwise adjust the price according to indexation rule $P_t(i) = P_{t-1}(i)\Pi_{t-1}^{\gamma_p}\Pi^{1-\gamma_p}$, where Π_{t-1} is lagged gross inflation, Π is steady state inflation, and $\gamma_p \in [0,1]$ is an indexation parameter.

Following the news literature, log technology is the sum of two components $\ln A_t = \ln S_t + \ln \Gamma_t$, where S_t is the usual surprise component, governed by

$$\ln S_t = \rho_S \ln S_{t-1} + \sigma_S \varepsilon_{S,t}, \tag{10}$$

with $\varepsilon_{S,t}$ i.i.d. $(0,1)$, and Γ_t is a permanent component that evolves according to

$$\ln \Gamma_t - \ln \Gamma_{t-1} = (1-\rho_\Gamma)\ln g + \rho_\Gamma \left(\ln \Gamma_{t-1} - \ln \Gamma_{t-2}\right) + \sigma_g \varepsilon_{g,t-1}, \tag{11}$$

with $\varepsilon_{g,t-1}$ i.i.d. $(0,1)$. As is common in the news literature, this shock is assumed to occur before it starts to impact technology and agents update expectations about the permanent component accordingly. Moreover, since $\rho_\Gamma > 0$, this shock portends even larger increases in the level of technology in the future, consistent with the basic insight of Beaudry and Portier (2006) and the empirical work discussed in Section 5 that technological innovations leading to permanent changes in productivity diffuse slowly.

The representative household consists of a continuum of members, a fraction N_t of whom are working while a fraction $1 - N_t$ are not working. Employed members provide labor services $L_t = e_t h_t N_t$ to labor unions, where h_t denotes average hours worked and e_t is labor effort. Members of the household are randomly chosen to work, with the household head choosing the total fraction of workers, N_t. All workers supply the same hours and effort, and all members enjoy the same consumption regardless of whether they work or not (i.e. there is perfect intra-household insurance). The expected lifetime utility of the household is

$$E_0 \sum_{t=0}^{\infty} \beta^t \nu_t \left[\ln \left(C_t - bC_{t-1} \right) + \theta N_t (T - G(h_t, e_t)) + \theta (1 - N_t) T \right], \tag{12}$$

where ν_t denotes an intertemporal preference shock that evolves according to

$$\ln \nu_t = \rho_\nu \ln \nu_{t-1} + \sigma_\nu \varepsilon_{\nu,t},$$

with $\varepsilon_{\nu,t}$ i.i.d. $(0, 1)$; C_t denotes consumption; b the degree of habit formation; T the total time endowment; and $G(h_t, e_t)$ the effective time cost when working h_t hours at effort level e_t.

The household can save via investment in physical capital, I_t, or through one period nominal bonds, B_t, that pay gross nominal interest rate R_t. It receives lump sum transfers, D_t, from ownership in production firms and labor unions, R^k_t for each unit of capital services supplied, and W_t for each unit of labor services. The flow budget constraint is:

$$\begin{aligned} C_t + I_t + \frac{B_{t+1}}{P_t} \leq \frac{W_t}{P_t} L_t + (1 + R_t) \frac{B_t}{P_t} + D_t \\ + \frac{R^k_t}{P_t} z_t K_t - a(z_t) K_t - W_t N_t \Psi \left(\frac{N_t}{N_{t-1}} \right) - K_t J \left(\frac{I_t}{K_t} \right). \end{aligned} \tag{13}$$

where z_t is the intensity with which capital is utilized, $a(z_t)$ is a convex adjustment cost to utilizing capital, $\Psi(\cdot)$ is a convex cost of adjusting employment, and $J(\cdot)$ is a convex cost of adjusting investment.[18] Physical capital evolves according to

$$K_{t+1} = (1 - \delta) K_t + \mu_t I_t,$$

where μ_t denotes an exogenous shock to the marginal efficiency of investment (MEI) that evolves

[18] Adjustment costs to employment and capital are crucial here as without them, optimal hours, effort and capital use would be constant. See Burnside et al. (1993) or Basu et al. (2006) for details.

according to

$$\ln \mu_t = \ln \rho_\mu \mu_{t-1} + \sigma_\mu \varepsilon_{\mu,t},$$

with $\varepsilon_{\mu,t}$ i.i.d. $(0, 1)$.

Labor services are supplied in a competitive market to a continuum of labor unions $j \in [0\,1]$. Unions transform these inputs into differentiated types of intermediate labor and sell them to a labor packer at nominal wage $W_t(j)$. The labor packer combine the different unions' labor into final labor service $L_{s,t}$ via a CES technology with elasticity of substitution ϵ_w and hires it out to intermediate goods producers at nominal wage rate W_t^l. Unions are subject to Calvo style nominal wage rigidity. With fixed probability $1 - \theta_w$, they can reoptimize $W_t(j)$ and otherwise adjust the wage according to indexation rule $W_t(j) = W_{t-1}(j)\Pi_{t-1}^{\gamma_w}\Pi^{1-\gamma_w} g$.[19]

The monetary authority, finally, sets the nominal interest rate according to

$$\frac{R_t}{R} = \left(\frac{R_{t-1}}{R}\right)^{\rho_R} \left[\left(\frac{\Pi_t}{\Pi}\right)^{\phi_\pi} \left(\frac{Y_t/Y_{t-1}}{g_Y}\right)^{\phi_y}\right]^{1-\rho_R} \exp(\sigma_R \varepsilon_{R,t}),$$

where $\varepsilon_{R,t}$ i.i.d. $(0, 1)$ is a monetary policy shock.

To assess the conditions under which Fernald's adjusted TFP series accurately measures technology, we start with unadjusted TFP as defined in (3). Even if utilization was constant (or variations in utilization were perfectly corrected), the model would nevertheless imply two incongruities between adjusted TFP and technology. First, with fixed cost $F > 0$, the production function is not constant returns to scale, thus invalidating the assumption that the cost shares of labor and capital sum up to one (i.e. $\omega_{L,t} + \omega_{K,t} = 1$). Second, if intermediate goods firms have market power and are subject to nominal price rigidities, then $\omega_{L,t}$ and $1 - \omega_{L,t}$ do not in general correspond to the true factor elasticities $1 - \alpha$ and α of the model. Specifically, cost-minimization on part of intermediate goods firms with respect to labor services implies

$$w_t^l L_{s,t} = (1 - \alpha)\psi_t \left[Y_t - F X_t\right], \tag{14}$$

where $w_t^l = W_t^l / P_t$ is the real wage of the labor service composite hired by production firms, and

[19] The assumption that non-reoptimized wages are indexed to productivity growth g ensures that trend growth does not result in steady state wage dispersion.

ψ_t denotes the inverse of the average price markup over marginal cost across intermediate firms. If the fixed cost F is chosen to ensure zero profit along the balanced growth path, which is a standard assumption, then (14) becomes

$$\omega_{L,t} = \frac{w_t^l L_{s,t}}{Y_t} = (1-\alpha)\psi_t \psi^{-1}, \tag{15}$$

with ψ^{-1} denoting the steady state markup. In this case, the labor share corresponds to the factor elasticity $1-\alpha$ on average but fluctuates over time due to undesired fluctuations in the markup owing to price rigidity.[20] Hence, as foreshadowed by Fernald's quote from Section 2, only in the limiting case of no fixed costs and no markups is it the case that unadjusted TFP defined as in (3) correctly measures technology net of utilization.

Consider now factor utilization. We introduce an econometrician similar to Fernald who does not observe labor effort and capital use but instead proxies utilization with filtered hours per worker as in (4) except that there are no industry differences; i.e. $\Delta \ln \hat{u}_t = \widehat{\beta} \Delta \ln h_t^c$. Mismeasurement can come from three sources. First, true utilization in the model, $\ln u_t = \alpha \ln z_t + (1-\alpha)\ln e_t$, is generally not proportional to hours per worker. Optimal hours and effort supplied by workers results in

$$G_h(h_t, e_t) h_t = G_e(h_t, e_t) e_t, \tag{16}$$

which does imply proportionality between e_t and h_t, exactly as in Basu et al. (2006). For capital use, however, optimality implies

$$r_t^k = a'(z_t). \tag{17}$$

Since r_t^k is an equilibrium object determined by the capital-labor ratio, there is no time-invariant mapping between z_t and h_t. Hence, unless the elasticity of the marginal cost $a'(z_t)$ with respect to capital use is infinity so that optimal capital use is constant (the case we henceforth label as $\sigma_z = 0$), true utilization systematically differs from hours per worker.[21] The second source of

[20] In the absence of fixed costs, the production function is constant returns to scale, consistent with the assumption underlying the construction of TFP, while the labor share becomes $\omega_{L,t} = \frac{w_t^l L_{s,t}}{Y_t} = (1-\alpha)\psi_t$. Hence, the labor share differs from $1-\alpha$ even on average. All the simulations below assume a positive fixed cost although we also experimented with zero fixed cost. The results remain very similar.

[21] While the assumption that capital use results in a resource cost (or equivalently in higher depreciation of physical capital) is standard in the DSGE literature, an alternative view is that workers need to be compensated for undesirable shifts in order to operate capital more intensively. Specifically, assume that preferences for leisure take the form $(T - G(h_t, e_t)V(z_t))$, with the cost of capital use $V(z_t)$ interpreted as the additional disutility from

utilization mismeasurement is that the proportionality factor $\widehat{\beta}$ estimated by the econometrician is biased. Basu et al. (2006), respectively Basu et al. (2013), try to address this issue by using demand side instruments in their estimation. It remains an open question, however, to what extent these instruments truly satisfy the exogeneity conditions necessary for instrumental variable estimation. The third potential source of utilization mismeasurement, as already highlighted by the results in Section 2, is that the filtering of hours per worker prior to constructing the utilization proxy may be inappropriate.

4.2 Calibration

The calibration of the standard model parameters is based largely on the estimates in Justiniano et al. (2010) except that we impose a stronger degree of nominal wage rigidity so as to generate model impulse responses for total hours and inflation to a news shock that broadly resemble the ones obtained in the data.[22]

For the utility cost of work, we assume

$$G(h_t, e_t) = \kappa_0 + \frac{\kappa_1}{\kappa_2} h_t^{\kappa_2} + \frac{\kappa_3}{\kappa_4} e_t^{\kappa_4}. \tag{18}$$

Given the proportionality between e_t and h_t, this time cost can be expressed in terms of hours worked only; i.e. $\tilde{G}(h_t) = G(h_t, e_t(h_t))$. We set κ_2 and κ_4 so as to target a Frisch elasticity of the intensive margin labor supply of 1 and the relative volatility of effort to hours of 4. The former is a plausible middle ground in the empirical literature (see Keane and Rogerson (2012)); the latter is set so as to obtain a measurement error between true and observed utilization that moves inversely with hours worked (see equation (19) below for details). The remaining parameters of this function do not affect the linearized dynamics of the model and are set consistent with normalized steady state values for h, e, and $G(h, e)$.

The autoregressive parameters for exogenous processes take on standard values, and the standard deviations of shocks are chosen to generate an unconditional standard deviation of output

working shifts at undesirable times. As long as the labor market is frictionless, optimal behavior by workers and firms then also implies proportionality between z_t and h_t, and utilization comoves perfectly with hours per worker as proposed by Basu et al. (2006). The point of our model here is not to take a stand on whether this proportionality condition holds in the data but rather to illustrate the consequences when it does not hold.

[22]See Kurmann and Otrok (2017a) for a discussion.

growth of one percent. See the Appendix for more details. Consistent with Justiniano et al. (2010), the preference shock and the MEI shock are main drivers of macro fluctuations in the model, accounting together for 55% of unconditional variance of output growth variance and about 85% of the unconditional variance of total hours growth and consumption growth.

Finally, for the proportionality factor $\widehat{\beta}$, we either set it so as to match the variance ratio of true utilization to hours per worker in the model or to $\widehat{\beta} = 3$, which is approximately the variance ratio of Fernald's aggregate utilization estimate to aggregate hours per worker in the data. This is somewhat lower than the variance ratio of true utilization to hours per worker in the model and thus leads to utilization mismeasurement.

4.3 Monte-Carlo simulations

We simulate 100,000 periods of data from the model and assess the consequences of technology mismeasurement. Before doing so, it should be re-emphasized that since the model abstracts from several important features of Fernald's construction of adjusted TFP in the data, the simulations are primarily an illustration of the measurement issues that *can* arise rather than a full explanation of how Fernald's revisions give rise to the changes in business cycle properties of adjusted TFP that we observe in the data. Nevertheless, we think that these illustrations are quite informative.

Table 4 starts by reporting the unconditional correlations between true utilization and estimated utilization and between true technology and adjusted TFP under four different scenarios. The first column shows the case when utilization is measured correctly; i.e. capital use is constant ($\sigma_z = 0$) and $\widehat{\beta}$ is exactly correct so that the proportionality condition holds, and hours per worker are not filtered. The correlation between true utilization and estimated utilization is 1 by definition in this case, and adjusted TFP comoves very closely with true technology. This

Table 4: Model-implied mismeasurement of utilization and technology

	No mismeasurement of utilization	$\sigma_z > 0$, $\widehat{\beta} = 3$ hours unfiltered	$\sigma_z > 0$, $\widehat{\beta} = 3$ hours BP-filtered	$\sigma_z > 0$, $\widehat{\beta} = 3$ hours BW-filtered
$corr(\Delta u_t, \Delta \hat{u}_t)$	1.00	0.97	0.48	0.97
$corr(\Delta A_t, \Delta TFP^u_t)$	0.95	0.82	0.59	0.82

Notes: This table shows correlations of different variables implied by the solution of the medium scale DSGE model laid out in the text.

suggests that incongruities arising from time-varying markups and non-constant returns to scale by themselves do not matter quantitatively. The second column shows the case of variable capital use ($\sigma_z > 0$) and the proportionality factor $\widehat{\beta}$ set to 3, but no filtering of hours. Estimated utilization still comoves closely with true utilization and adjusted TFP remains strongly correlated with technology, albeit less so than when utilization is measured correctly.

The third and fourth columns keep variable capital use ($\sigma_z > 0$) and $\widehat{\beta} = 3$ but now also detrend hours per worker with either the bandpass filter or the bi-weight filter. Bandpass filtering clearly imparts substantial additional mismeasurement, with the correlations between true and estimated utilization and between technology and adjusted TFP dropping to around 0.5. In contrast, bi-weight filtering does not affect the comovement of measured utilization and adjusted TFP in any significant way. This is because the persistence of hours per worker in the model is low and thus, bi-weight filtering is almost equivalent to no filtering, which is very similar to what we obtain in the data (see the Appendix for details).

Next, we estimate the baseline four-variable VAR from above on the simulated data to illustrate the performance of the Barsky-Sims identification of news shocks under the different scenarios.[23] First, we consider the baseline scenario in which the proportionality condition holds (i.e. $\sigma_z = 0$; and the proportionality factor is correct). Figure 3 reports the results. Here and below, the solid black lines show the impulse responses to a news shock in the model; the dashed blue lines the VAR responses implied by the Barsky-Sims identification when hours per worker in the construction of utilization are bandpass filtered; and the dotted red lines the VAR responses implied by the Barsky-Sims identification when hours per worker are bi-weight filtered.

Similar to the data, consumption in the model jumps up on impact of the news shock and then gradually increases further to a new permanently higher level; total hours worked drop slightly on impact and then increase in a hump-shaped manner; and inflation falls sharply on impact and then returns to zero over the next ten quarters. When the proportionality condition holds and hours per worker are bi-weight filtered, the Barsky-Sims identification performs well in capturing the dynamics to a news shock. The fit is somewhat less close when hours per worker are bandpass filtered, with the responses of consumption, total hours, and inflation displaying

[23]The point of using such a long sample of simulated data is that we want to examine the asymptotic consequences of technology mismeasurement for news identification. Of course, could also investigate the small-sample properties of our estimates. We did so and found very similar results.

mild oscillatory behavior. This is not an issue of the Barsky-Sims identification per se but of bandpass filtering when constructing utilization, which appears to introduce artificial dynamics in the VAR.[24] Nevertheless, even with bandpass filtering, the fit with true model responses to a news shock remains good. This provides further confirmation that for reasonable markup variations as implied by our model, the differences between Fernald's construction of TFP and true TFP due to time-varying markups and non-constant returns to scale are quantitatively unimportant.

Figure 3: Simulated Responses to Barsky-Sims News Shock when Proportionality Holds

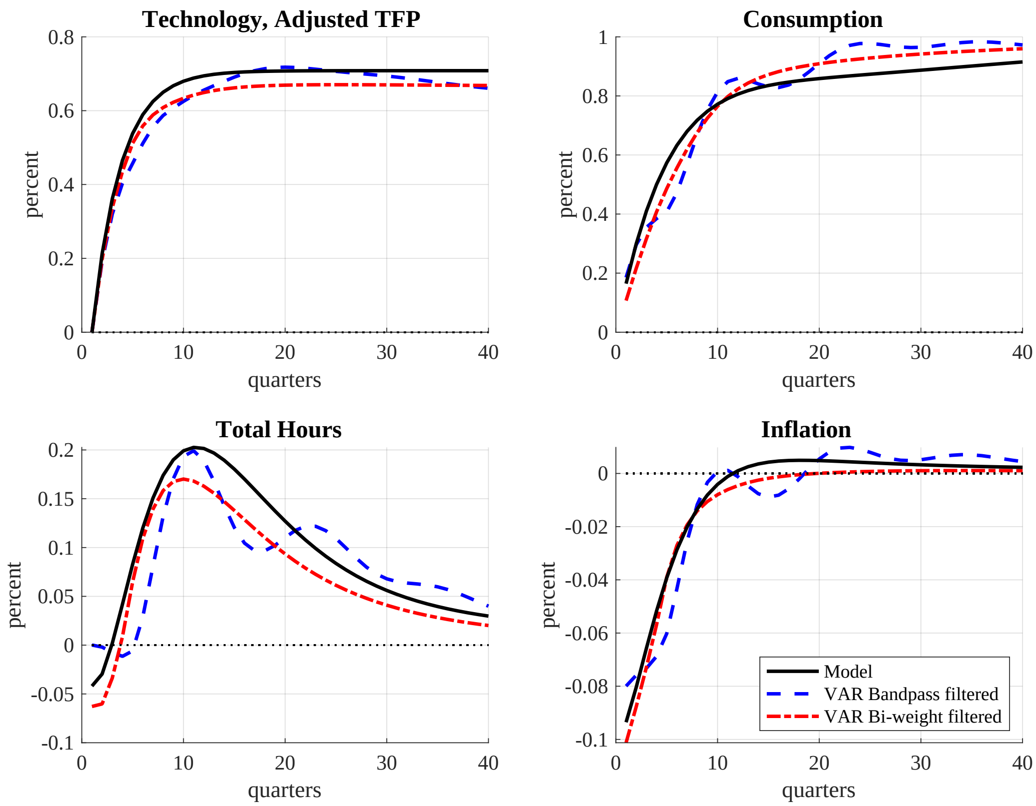

Notes: Solid lines are the true impulse responses to a news shock in the model. The dashed blue lines are the estimated responses using the Barsky-Sims identification based on the simulated data with bandpass filtered hours per worker in the construction of utilization. The dash-dotted red lines are the estimated responses using the Barsky-Sims identification based on the simulated data with bi-weight filtered hours per worker in the construction of utilization.

Second, we consider the case when the proportionality condition does not hold; i.e. capital use is non-constant ($\sigma_z > 0$) and $\widehat{\beta} = 3$. As shown in Figure 4, the VAR responses for consumption and inflation come again reasonably close to the ones implied by the model, independent of the filtering method for hours per worker in the construction of utilization. The response of total

[24] To avoid any confusion, total hours in the VAR are not filtered, only hours per worker in the construction of utilization are filtered.

hours in the VAR, however, new depends significantly on the filtering of utilization. Under bi-weight filtering, hours slightly increase on impact and remains above the model-implied response for about 10 quarters. In contrast, under bandpass-filtering, the total hours response is – aside from the initial period – negative for several quarters before increasing in line with what is implied by the model.

Figure 4: Simulated Responses to Barsky-Sims News Shock when Proportionality Fails to Hold

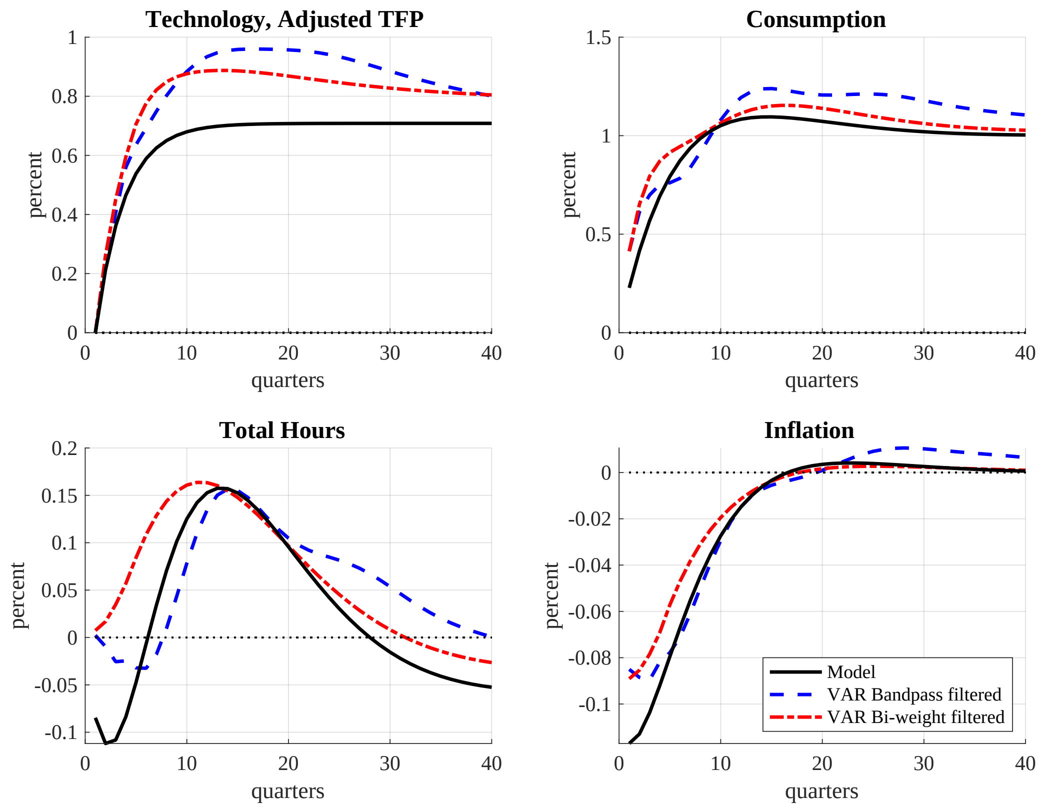

Notes: Solid lines are the true impulse responses to a news shock in the model. The dashed blue lines are the estimated responses using the Barsky-Sims identification based on the simulated data with bandpass filtered hours per worker in the construction of utilization. The dash-dotted red lines are the estimated responses using the Barsky-Sims identification based on the simulated data with bi-weight filtered hours per worker in the construction of utilization.

This difference in hours response depending on the filtering method is broadly similar to what we observe in Figure 2 for the 2007 (bandpass filtered) vintage and the 2016 (bi-weight filtered) vintage of adjusted TFP. This suggests that for the case when the proportionality condition between utilization and hours does not hold, bandpass filtering of hours per worker in the construction of utilization may actually be preferable to bi-weight filtering even though by itself, bandpass filtering introduces substantial mismeasurement. To understand this result, it is useful to express adjusted TFP growth as follows

$$\triangle \ln TFP_t^u = (\triangle \ln A_t - \triangle \epsilon_t^{TFP}) + (\triangle \ln u_t - \widehat{\beta} \triangle \ln h_t^c), \tag{19}$$

where $\triangle \epsilon_t^{TFP}$ is the difference between true TFP growth as defined in the model and TFP growth as defined in (3). Since $\triangle \epsilon_t^{TFP} \approx 0$ in our simulations, adjusted TFP moves either because of shocks to technology or because of non-technology shocks that imply $\triangle \ln u_t - \widehat{\beta} \triangle \ln h_t^c \neq 0$. According to our model calibration, preference shocks and MEI shocks both lead to sizable short-term fluctuations in $\triangle \ln u_t - \widehat{\beta} \triangle \ln h_t^c$ and thus adjusted TFP that comove with total hours; i.e. the short-run change in true utilization $\triangle \ln u_t$ in response to these shocks is larger than the short-run change in measured utilization $\widehat{\beta} \triangle \ln h_t^c$. Absent filtering of hours, the Barsky-Sims identification, by relying on short-term restrictions and in particular the zero impact restriction, picks up a combination of these shocks, confusing them as news shocks to technology. Since bi-weight filtering barely affects the cyclicality of hours per worker, this explains the positive VAR response of total hours in Figure 4. In comparison, band-pass filtering of hours per worker substantially alters the dynamic characteristics of $\triangle \ln u_t - \widehat{\beta} \triangle \ln h_t^c$, which in our case results in the Barsky-Sims identification picking up less of the combination of non-technology shocks and resulting in a more negative response of hours to the news shock, similar to what is implied by the model.

The bottomline of this discussion is that bandpass filtering hours per worker, despite inducing substantial mismeasurement of utilization, can in some cases help "smoothening out" departures from the proportionality assumption in Fernald's proxy of utilization. Bi-weight filtering, in contrast, does not attenuate such departures from proportionality and therefore leaves news shock identifications such as Barsky-Sims approach that focus on short-run restrictions in adjusted TFP more sensitive to utilization mismeasurement.

At the same time, it should be clear from this discussion that none of our simulation results are general. Indeed, for alternative model calibrations, bandpass filtering of utilization does not attenuate the effects of departures from proportionality and to the contrary, may in fact exacerbate it.[25] Hence, we conclude from these simulations that news shock identifications relying on short-

[25] In particular, as discussed above, VARs with bandpass filtered hours per worker in the construction of utilization have a tendency to induce oscillatory impulse responses.

run restrictions and in particular the zero impact assumption can be highly sensitive to cyclical mismeasurement of true technology. A more fruitful approach is instead to devise alternative identification restrictions that are robust to cyclical measurement issues. This is what we propose in the next section.

5 An Alternative Identification of News Shocks

The central idea behind our proposed alternative identification is that new productivity-enhancing technologies disseminate slowly and, if known to agents, constitute news about future productivity growth. As long as productivity in the long run is driven primarily by new technologies, an identification that accounts for most of productivity variations in the long-run should therefore capture news. At the same time, as long as this identification does not rely on short-run restrictions and in particular the zero impact restriction, it should be robust to cyclical measurement issues with productivity.

The idea that new technologies diffuse slowly finds ample support in a large micro-empirical literature; e.g. Griliches (1957), Mansfield (1961), Mansfield (1989), Gort and Klepper (1982) or Rogers (1995). According to Mansfield (1989), for example, the time until half of potential adopters actually adopt a new technology varies between five and fifteen years, depending on technology. While the slow dissemination of new technologies and its implications for the modeling of productivity is discussed extensively by Rotemberg (2003) as well as Comin and Gertler (2006) and Lindé (2009) among others, much of the business cycle literature has modeled productivity as a jump process where innovations lead to an immediate change of productivity to a new level that is either permanent or highly persistent. Yet, the assumption of slow dissemination is consistent with the basic insight of Beaudry and Portier (2006) from a bivariate VAR that news shocks identified through the zero impact restriction are closely related to the shocks driving long-run movements in productivity. Our contribution here consists of exploring this insight further by extracting a long-run productivity shock in larger VAR systems, assessing its robustness to the above documented revisions in adjusted TFP, and using additional information to interpret the extracted shock as a news shock.

5.1 Implementation and discussion

We implement our alternative news shock identification by estimating a VAR containing adjusted TFP and extracting the shock that accounts for the maximum FEV share of adjusted TFP at a long but finite horizon H; i.e.

$$\max_{\gamma} \quad \Omega_{1,2}(H) \quad \text{s.t. } \gamma'\gamma = 1, \tag{20}$$

where, as per equation (7), $\Omega_{1,2}(H)$ denotes the FEV share of adjusted TFP at horizon H accounted for by the second shock in shock vector ϵ_t; and γ denotes a column vector belonging to orthonormal rotation matrix $\mathbf{Q}$ of the Choleski decomposition of the reduced form variance covariance matrix. While conceptually similar to Barsky and Sims (2011), there are two important differences. First, we look for the shock that accounts for the maximum FEV share of adjusted TFP at a long horizon H instead of maximizing the sum of FEV shares from impact onward. Second, we drop the zero restriction (i.e. the first element of γ is not restricted to zero), which means that measured productivity is allowed to respond contemporaneously to the shock. By focusing on a long forecast horizon only, this max-share identification has the advantage that it reduces the potential bias imparted by cyclical mismeasurement of technology – especially coming from the mismeasurement of utilization. Moreover, the approach avoids taking a stand on whether (true) technology reacts to the shock only with a lag or not.[26]

Mechanically, the proposed max-share identification is the same as the technology shock identification of Francis et al. (2013) – which in turn builds on earlier work by Uhlig (2003) – but differs in that we propose it as an alternative identification of news shocks and that we apply it to adjusted TFP instead of labor productivity as the target variable. As we will discuss below, this latter difference is important. Because of variations in the ratio of capital to labor – capital deepening – labor productivity responds quite differently to the shock than adjusted TFP, thus making the news interpretation less obvious. Moreover, since capital deepening is endogenous,

[26]There is no *a priori* reason to think that news about growth-enhancing advances in technology are, despite their slow diffusion, completely unrelated to current productivity. Indeed, it seems equally intuitive to assume that market participants revise their expectations about future fundamentals only once there is evidence that at least some firms have successfully adopted the new technology. To our knowledge, the only other paper that discusses this possibility is Barsky, Basu, and Lee (2015) who write: *"It is possible that news about future productivity arrives along with innovations in productivity today (page 233)."*

labor productivity is affected even in the long-run by other non-technology shocks, potentially invalidating identification of technology shocks based on long-run restrictions.[27]

Compared to other long-run identification schemes employed in the VAR literature, the max-share approach has the advantage of focusing on a long but finite horizon. As Francis et al. (2013) show, this helps to reduce small-sample bias in VARs that, as discussed in the Introduction, can have potentially important effects for infinite-horizon identifications of long-run shocks. In addition, the max-share approach does not impose that technology is the *only* source of long-run fluctuations in productivity and instead affords the possibility that other shocks (e.g. a surprise productivity shock) exert at least some long-term effect on adjusted TFP.

5.2 Results

We apply the proposed max-share identification to the same four-variable VAR as in Section 3. The horizon at which the FEV share of adjusted TFP is maximized is set to $H = 80$ quarters, although similar results would obtain for other long horizons. The estimated impulses responses are reported in Figure 5.

In contrast with the results based on the Barsky and Sims (2011) news identification, there is very little difference in the impulse responses between the VAR estimated with the 2007 vintage of adjusted TFP and the VAR estimated with the 2016 vintage. In both cases, consumption jumps on impact and then gradually increases further to a permanently higher level; hours worked decline significantly on impact before turning positive after about five quarters; and inflation drops sharply and significantly on impact of the shock before gradually returning towards its initial level. The only discernible difference is the short-run response of adjusted TFP, which should not be surprising given their difference in cyclical properties. For both vintages, adjusted TFP jumps on impact, although insignificantly so. The 2007 vintage then increases gradually whereas the 2016

[27]In particular, persistent changes in capital taxes and worker composition are likely to affect labor productivity even in the long-run but should leave long-run TFP unaffected (provided that Fernald's aggregate production function assumption and his measures of effective labor and capital are correct). See Uhlig (2004) or Bocola et al. (2014) for examples. Of course, non-technology shocks may affect adjusted TFP (as well as labor productivity) in the long-run if the discovery and adoption of new technologies arises endogenously as a function of the state of the business cycle. In this case, the proposed identification as well as the other existing identifications of technology shocks will confound news shocks with non-technology shocks. This point remains an unresolved issue in the literature that we start to address below by examining the response of novel indicators of technological innovation to our extracted shock.

vintage temporarily declines and remains insignificant for more than 10 quarters. Both vintages, however, increase gradually at longer horizons and end up two to three times higher than their impact responses. In other words, the max-share shock predicts delayed but sustained future productivity growth.

Figure 5: Impulse Responses of Four-Variable VAR to Max-Share Shock

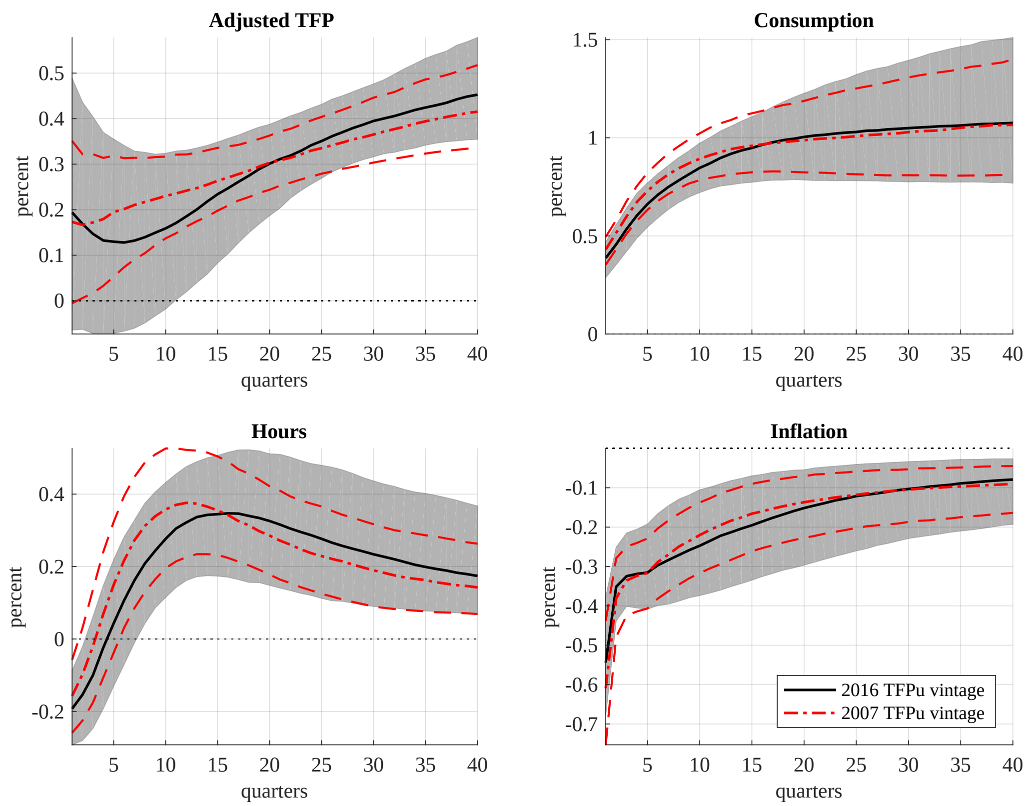

Notes: Solid black lines are the posterior median estimates from the VAR system estimated with the 2016 vintage of adjusted TFP. The grey bands correspond to the 16 to 84 percent posterior coverage intervals. The red dash-dotted lines are the posterior median estimates for the system estimated with the 2007 vintage of adjusted TFP. The red dashed lines correspond to the 16 to 84 percent posterior coverage intervals. The shock is identified using the max-share identification, which does not impose the zero impact restriction with respect to adjusted TFP and instead just maximizes the FEV share of adjusted TFP at a 80 quarter horizon.

Aside from the adjusted TFP response, these results look close to the original results reported in Barsky and Sims (2011) based on the 2007 vintage of adjusted TFP. Indeed, as shown in the Appendix, the median correlation between consumption growth and hours growth implied by the max-share shock is robustly negative, contrary to what we observe unconditionally in the data.[28]

As also shown in the Appendix, the responses for consumption, hours and inflation are robust to replacing adjusted TFP with either unadjusted TFP or labor productivity. The only difference is

[28]Similar to the results reported above, the business cycle moments implied by the max-share shock are estimated very precisely.

that the response of these alternative target variables to the max-share shock is quite different. In particular, consistent with results in Francis et al. (2013), labor productivity jumps up considerably more on impact and then increases further to a permanently higher level. This is due to an additional capital deepening effect and may lead to the inadvertent conclusion that technology should be modeled as a random walk process – as is quite frequently assumed in the business cycle literature – which is very different from our finding that the response of adjusted TFP to the max-share shock is insignificant on impact before gradually increasing to a new permanent level that is substantially higher, consistent with the empirical literature cited above that technology is slowly diffusing.

5.3 Does the max-share identification capture news shocks?

As discussed above, the news shock interpretation of the proposed alternative identification rests on several important questions, in particular:

1. Does the max-share shock lead to delayed predictable changes in future TFP?
2. Is the max-share shock correlated with measures of technological innovation?
3. Does the max-share shock generate sizable responses in forward-looking news indicators?

For the first question, we already know from the results with the four-variable VAR that the max-share shock leads to persistent and therefore predictable changes in future TFP growth. We now extend the analysis by considering an eight-variable VAR system that contains, in addition to the four variables already included above, real gross domestic product (GDP) per capita, real private investment expenditures per capita, the real S&P500 index (deflated by the consumer price index) and the Federal Funds rate.[29] This choice of variables is motivated by the desire to learn about the effects of the max-share shock for other prominent macroeconomic aggregates and by the idea that including forward-looking information variables may help sharpen the results and address issues of non-fundamentalness (e.g. Leeper et al., 2013). Indeed, as Beaudry and Portier (2006) argue, there is a large literature suggesting that stock prices reflect expectations about future economic

[29]The real S&P500 index is taken directly from Robert Shiller's website. None of the results would change if the index was instead transformed into real terms with another deflator. The other variables are taken from the FRED database of the Federal Reserve Bank of St. Louis.

conditions and should therefore be an important indicator of news. Similarly, the Federal Reserve with its large staff of economists may have superior forecasting abilities and thus, news could also be reflected in the Federal Funds rate, the main monetary policy instrument up until the recent financial crisis. As before, the VAR is estimated with four lags for the 1960:1-2007:3 period subject to a Minnesota prior.

Figure 6: Impulse Responses of Eight-Variable VAR to Max-Share Shock

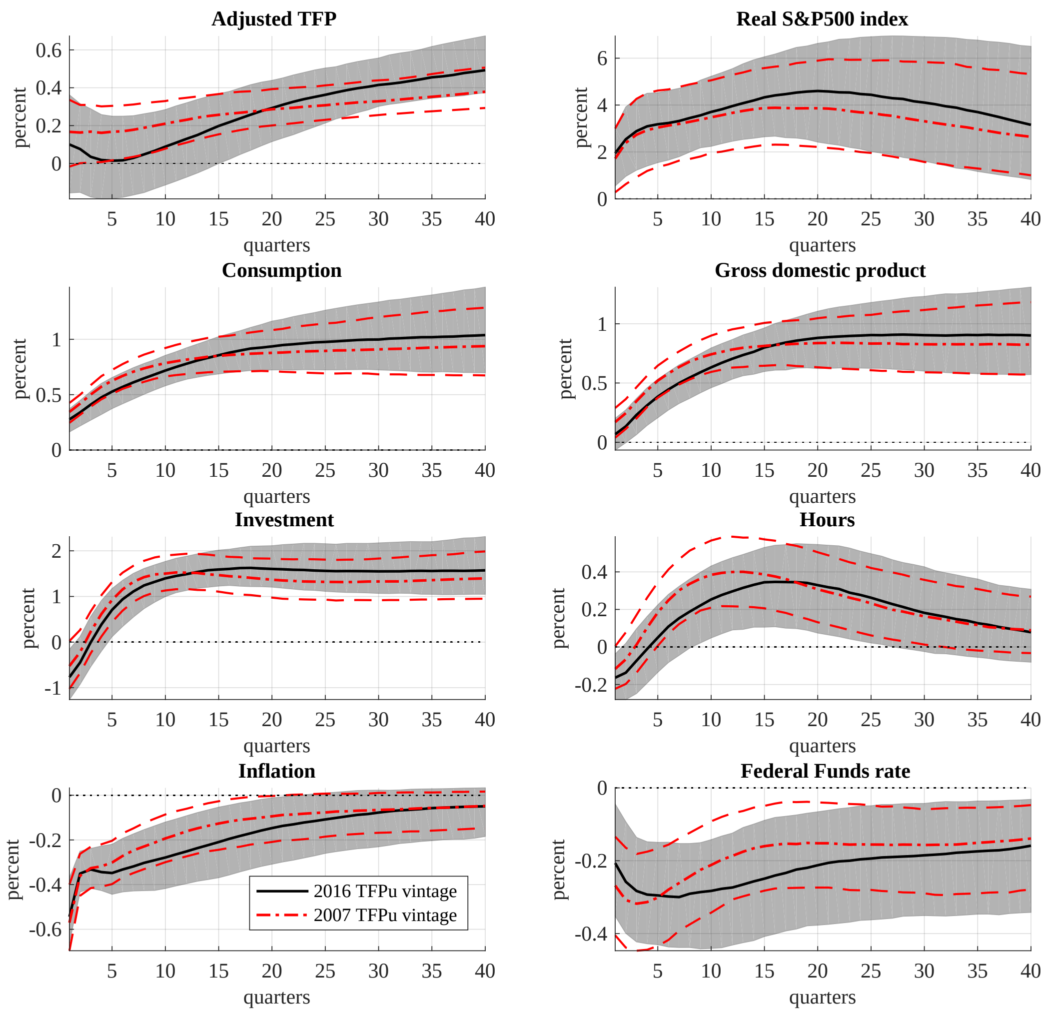

Notes: Solid black lines are the posterior median estimates from the VAR system estimated with the 2016 vintage of adjusted TFP. The grey bands correspond to the 16 to 84 percent posterior coverage intervals. The red dash-dotted lines are the posterior median estimates for the system estimated with the 2007 vintage of adjusted TFP. The red dashed lines correspond to the 16 to 84 percent posterior coverage intervals. The impulse responses are identified using the max-share identification.

Figure 6 displays the impulse responses. The estimated responses again match closely across the two vintages, confirming the robustness of the max-share approach to the revisions in Fernald's adjusted TFP series. Compared to the four-variable VAR, the reaction of adjusted TFP to the shock is more delayed and gradual, with an impact response for the 2016 vintage that starts closer to zero. This difference in results is primarily due to the inclusion of the real S&P500 index in the VAR, confirming the point of Beaudry and Portier (2006) that stock prices contain valuable information about market expectations of future economic conditions.

The real S&P500 index itself reacts strongly on impact of the shock and then displays a mild hump-shaped response that is quite persistent. Investment and total hours worked both decline initially while output rises slightly and consumption jumps up robustly on impact. Thereafter, output, consumption, and investment gradually increase towards a permanently higher level while total hours worked responds in a hump-shaped manner similar to its response in the four-variable VAR. Inflation and the Federal Funds rate both decline significantly on impact and then remain persistently below their original values. The initial decline of inflation substantially exceeds the decline in the Federal Funds rate, implying that real short interest rates increase on impact of the shock. Hence, the shock triggers a contractionary monetary policy response despite the deflationary effect that the shock has on the economy.

Table 5: Fraction of FEV Explained by Max-Share Shock

	Forecast horizon (quarters)			
	4	20	40	80
Adjusted TFP (2016)	0.07	0.16	0.49	0.77
Gross domestic product	0.08	0.60	0.77	0.83
Consumption	0.34	0.74	0.86	0.88
Investment	0.05	0.41	0.60	0.72
Hours	0.04	0.18	0.23	0.25
Real S&P500 index	0.20	0.45	0.52	0.49
Inflation	0.38	0.47	0.46	0.45
Federal Funds rate	0.13	0.26	0.32	0.33

Notes: The sample period for each of the statistics is 1960q1-2007q3. The model statistics pertain to medians from the posterior distribution of each data series implied by the max-share shock. All results are rounded to two digits after the decimal point.

The opposite-signed impact responses of consumption relative to hours and investment implies that the max-share shock generates negative business cycle comovement between these variables. This confirms the conclusion from the four-variable VAR that the shock is unlikely to be a main driver of business cycle dynamics. This does not mean, however, that the shock is unimportant for macroeconomic fluctuations more generally. Indeed, as Table 5 shows, while the shock accounts for only a small fraction of the FEV of real macroeconomic aggregates at short horizons (with consumption being the notable exception), the shock is the main driver of these variables at longer horizons with the exception of hours worked.[30] Indeed, at the 80 quarter horizon, the shock accounts for about three-fourths of unpredictable variations in adjusted TFP, GDP, consumption, and investment. Quite strikingly, the shock also accounts for almost half of unpredicted variations in the real S&P500 index and inflation at forecast horizons of 20 quarters and more, and about one-third of unpredictable variations in the Federal Funds rate at horizons of 40 quarters or more.

The results in Figure 6 and Table 5 indicate that the max-share shock predicts a delayed but sustained increase in future productivity, accounting for almost none of the fluctuations in adjusted TFP at short horizons but three-fourths of fluctuations at long horizons. Stock market participants, consumers, firms, and the Federal Reserve immediately react to the shock.

To answer the second and third question above, we re-estimate the eight-variable VAR with the Federal Funds rate replaced sequentially with different measures of technological innovation and forward-looking information variables. The rest of the VAR specification is kept unchanged except when we have to adapt the sample due to data availability, as described below. To save on space, we only report impulse responses for the variables that replace the Federal Funds rate. The seven other variables in the VAR, which are kept the same throughout the exercise, react very similarly to the max-share shock as reported above in Figure 6.

We first consider four different measures of technological innovation: the index of information and communications technology (ICT) standards by Baron and Schmidt (2015); the index of new technology manuals by Alexopoulos (2011); real R&D expenditures per capita from the NIPAs; and the inverse of the relative price of investment price from Justiniano et al. (2010). The index by Baron and Schmidt (2015) counts the number of new ICT industry standards per quarter released by standard setting organizations (SSOs) in the U.S.[31] As Baron and Schmidt (2015) argue, standardization is an essential step in the introduction and adoption of new technologies.

[30] All of the results in Table 5 refer to median estimates from the VAR estimated with the 2016 vintage of adjusted TFP. The results are very similar for the VAR estimated with the 2007 vintage and are therefore omitted to save on space.

[31] SSOs are mostly private organizations that exist in many industries to establish voluntary and regulatory standards. Prominent examples include the electricity plug, the USB key, the WiFi communications protocol or quality standards (e.g. ISO). Also see *"The Joy of Standards"* by Andrew Russell and Lee Vinsel in the New York Times. The standardization index by Baron and Spulber (2015) and Baron and Schmidt (2015) is based on information from the Searle Center database on technology standards and standard setting organizations. See their papers for details. We thank Justus Baron and Julia Schmidt for making their index available.

It precedes the implementation of new technologies but presumably provides an important signal about the commercial viability of an innovation and thus future growth opportunities. As such, standardization represents an ideal measure to assess the extent to which our max-share shock captures news. As in Baron and Schmidt (2015), we focus on ICT standards because ICT have constituted the dominant type of general purpose technology, although results are robust to using broader industry standards. Alexopoulos (2011)'s count of books published in the field of technology provides a complementary measure even though she develops her measure primarily to investigate the role of contemporaneous technology shocks.[32] As explained in her paper, new book titles in this area *"... appear precisely when the innovation is first introduced to market, for the very good reason that the whole purpose of publications is to spread the word about the new product or process."* R&D expenditures and the relative investment price are common measures of the quality and/or efficiency of newly produced investment goods. If our max-share shock captures news about future productivity growth, then we would expect both of these measures to react gradually as new technologies are being implemented and start to affect productivity.[33]

Alexopoulos' book measure is only available at an annual frequency and stops in 1995. We therefore estimate a smaller, annual VAR for this case, containing adjusted TFP, consumption, inflation and Alexopoulos' book measure. For all the other variables, the impulse responses are estimated with the above described VAR based on quarterly data for the 1960:3-2007:3 sample.

Figure 7 reports the impulse responses. Both the index of new ICT standards and the index of new technology manuals jump markedly on impact of the shock. The index of new ICT standards then declines back towards its pre-shock level while the new manuals measure remains permanently higher. The response of the ICT standards index is particularly striking and matches closely with the evidence reported in Baron and Schmidt (2015), who use a recursive identification approach based on zero impact restrictions. R&D expenditures and the (inverse of the) relative price of

[32] As emphasized above, the two are not necessarily distinct as news about future productivity growth may coincide with contemporaneous innovations to productivity. Alexopoulos (2011) also constructs different new book titles for different technology categories, including new titles for computer hardware and software, and telecommunications. The results presented below are robust to using these alternative measures.

[33] One could argue that as long as Fernald's TFP series appropriately controls for quality changes in the capital stock, news shocks derived from TFP should be unrelated to capital-embodied technological change. However, it is doubtful that the the quality adjustments made to the different capital series that Fernald uses fully capture these quality changes. Moreover, as argued for example by Chen and Wemy (2015), there may be spillovers from capital-embodied technological change to neutral, general-purpose technology.

investment goods, in turn, increase only gradually after the shock, although this increase occurs at a considerably faster pace than for adjusted TFP, as reported in Figure 6.

Figure 7: Impulse Responses of Innovation Measures to Max-Share Shock

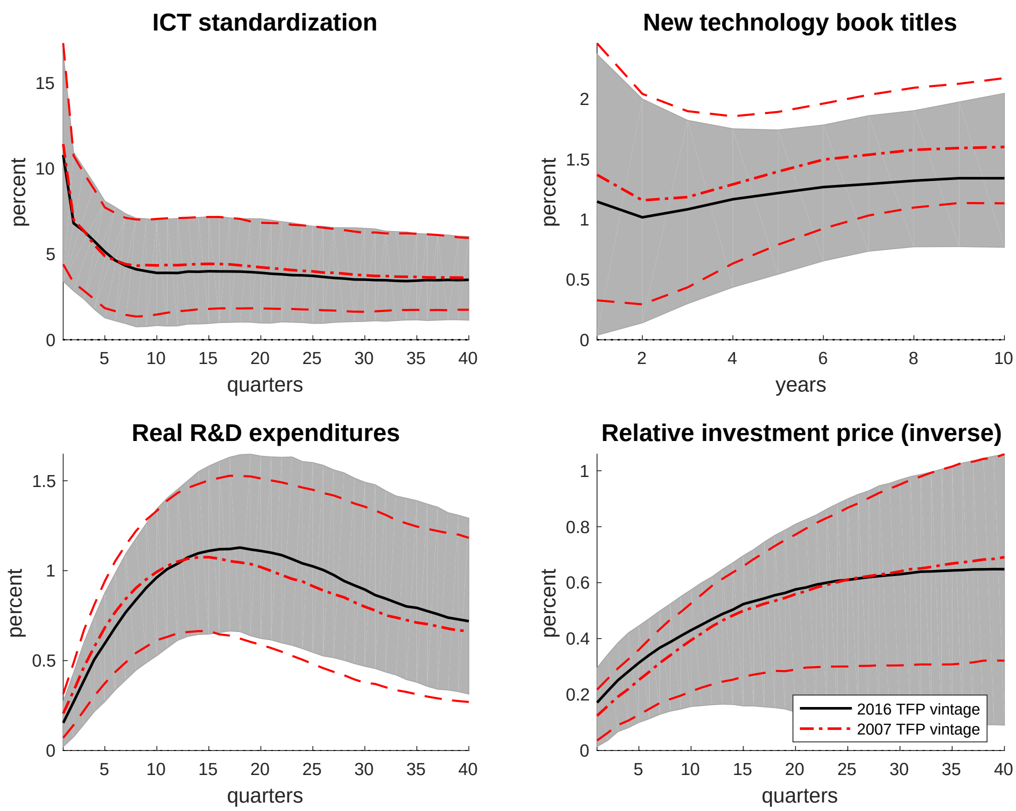

Notes: Solid black lines are the posterior median estimates from the VAR system estimated with the 2016 vintage of adjusted TFP. The gray bands correspond to the 16 to 84 percent posterior coverage intervals. The red dash-dotted lines are the posterior median estimates for the system estimated with the 2007 vintage of adjusted TFP. The red dashed lines correspond to the 16 to 84 percent posterior coverage intervals. The impulse responses are identified using the max-share identification.

Taken together, the impulse responses indicate that the max-share shock picks up the introduction of new technologies to markets instead of other shocks that endogenously lead to more R&D activity and eventually more innovation and higher productivity. Otherwise, one would expect ICT standards and new technology book titles to respond not with an initial jump but only gradually and with a delay relative to R&D expenditures.

Next, we consider three forward-looking information variables that have been interpreted as capturing news: the spread between long-term (5-year) treasury bond yields and the Federal Funds rate as used in Kurmann and Otrok (2013); the Michigan Survey's 5-year ahead consumer confidence index as used in Barsky and Sims (2012); and the business confidence index from the Business Outlook Survey (BOS) conducted by the Federal Reserve Bank of Philadelphia as used in Bachmann et al. (2013). Figure 8 shows the impulse responses of these series. For reference, we also include the impulse response of the real S&P500 index, which is part of the VAR used to generate these results. All of the indicators jump up sharply on impact of the news shock and then decline gradually back to their original level. These responses are highly significant and indicate that the identified max-share shock captures news about the future that is picked up not only by financial markets but also the Fed, consumers, and businesses.[34]

Figure 8: Impulse Responses of News Indicators to Max-Share Shock

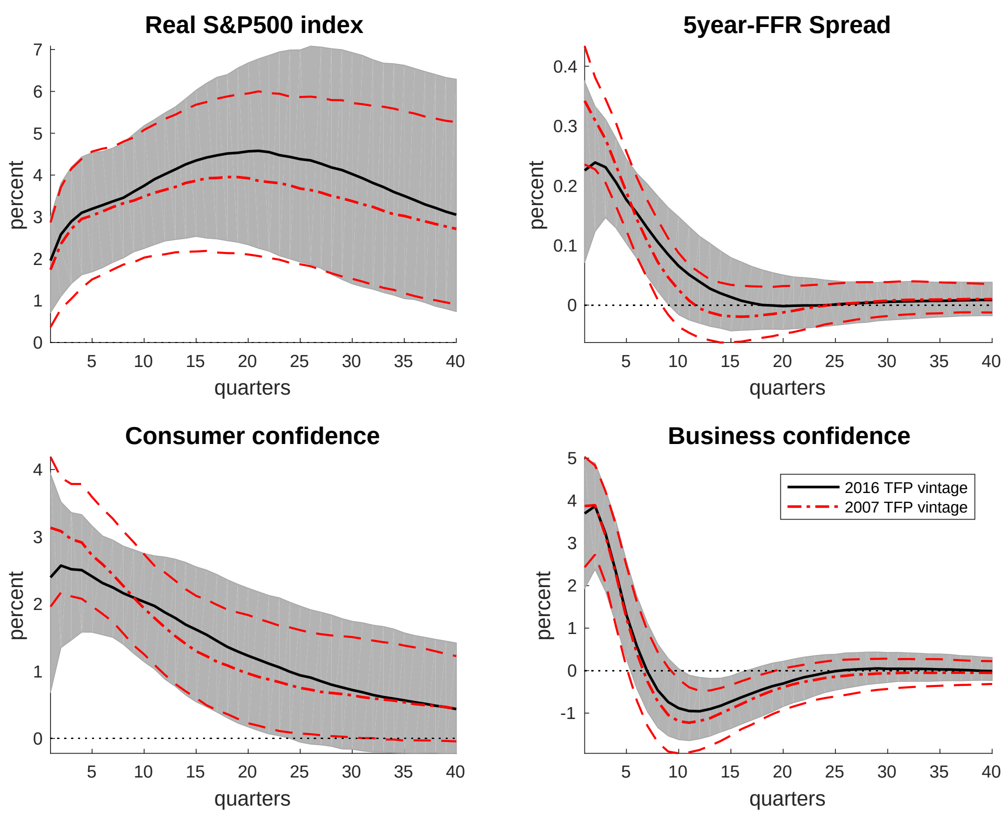

Notes: Solid black lines are the posterior median estimates from the VAR system estimated with the 2016 vintage of adjusted TFP. The gray bands correspond to the 16 to 84 percent posterior coverage intervals. The red dash-dotted lines are the posterior median estimates for the system estimated with the 2007 vintage of adjusted TFP. The red dashed lines correspond to the 16 to 84 percent posterior coverage intervals. The impulse responses are identified using the max-share identification.

The results provide compelling evidence that the max-share shock captures news about future productivity growth. The shock predicts delayed sustained future TFP growth, accounting for

[34]In previous versions of the paper, we also reported that the max-share shock leads to strong positive impact responses of capital returns, providing further evidence that the identified shock contains important information about the future that market participants know about. Likewise, the max share shock leads to a strong negative impact responses of different measures of uncertainty, suggesting that the news picked up by the max-share shock provides resolution of uncertainty about the productive potential of innovations. Details of these results are available upon request.

only a small fraction of TFP fluctuations at short forecast horizons but for 70 percent or more of TFP fluctuations at longer horizons. Perhaps more importantly, the shock is associated with large and persistent jumps in two novel measures of innovation, followed by a hump-shaped increase in R&D expenditures and a gradual decline in the relative price of investment goods; and the shock generates jumps in a wide variety of forward-looking information variables. Taken together, these responses suggest that the max-share identification picks up technological innovation as opposed to other business cycle shocks or noise that endogenously lead to changes in productivity; and that market participants clearly update their forecasts about the economy. The news interpretation therefore seems natural.

5.4 Monte-Carlo simulations

As a final check, we perform the same Monte-Carlo simulations as above to assess whether the max-share identification captures more robustly the model responses to a news shock than the Barsky-Sims identification. To save on space, we consider directly the situation where the proportionality condition for utilization does not hold; i.e. capital use is variable ($\sigma_z > 0$) and the factor of proportionality $\widehat{\beta} = 3$ is different from variance ratio of true utilization to hours per worker in the model. Results for the case when the proportionality conditions holds are reported in the Appendix and match the model responses as well as the ones obtained with the Barsky-Sims identification.

Figure 9 reports the results of this simulation. The max-share identification clearly outperforms the Barsky-Sims identification (compare to Figure 4), closely matching the impulse responses of not only consumption but also total hours and inflation, regardless of whether utilization is constructed with bandpass-filtered or bi-weight filtered hours per worker. In particular, for both cases, the max-share identification implies a drop in total hours on impact followed by a hump-shaped increase after about 10 quarters that matches the model response of total hours. This further confirms the robustness of the max-share identification approach to different measures of utilization.

Figure 9: Simulated Responses to Max-Share Shock when Proportionality Fails to Hold

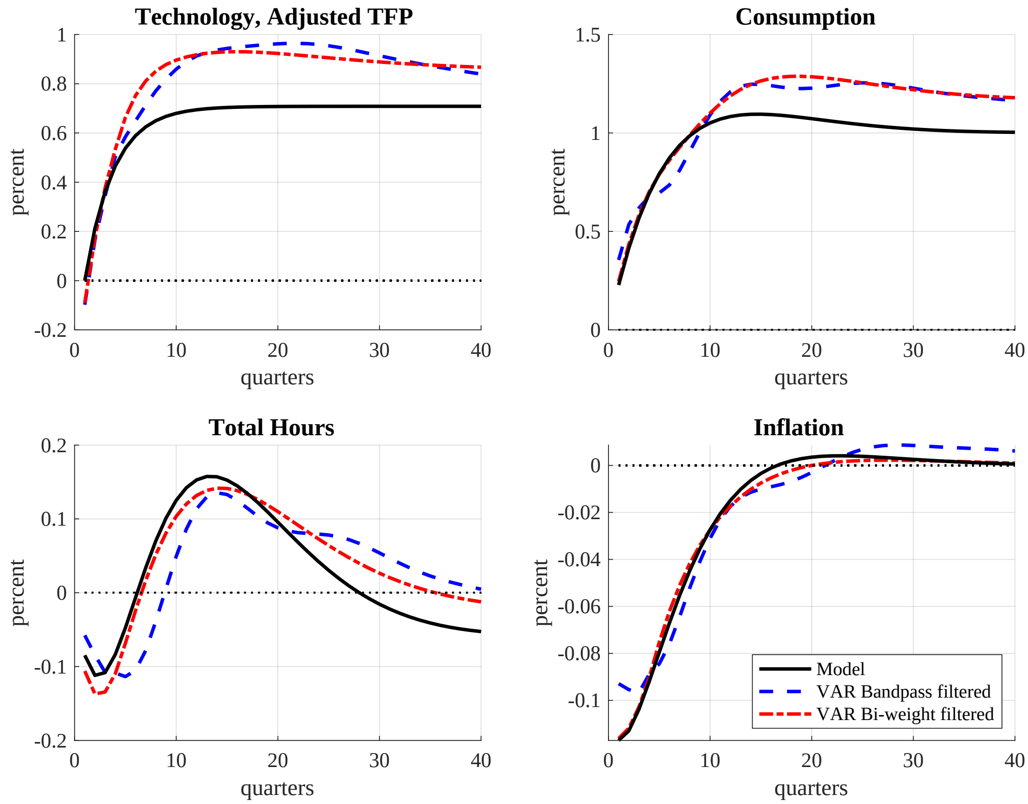

Notes: Solid lines are the true impulse responses to a news shock in the model. The dashed blue lines are the estimated responses using the max-share identification based on the simulated data with bandpass filtered hours per worker in the construction of utilization. The dash-dotted red lines are the estimated responses using the max-share identification based on the simulated data with bi-weight filtered hours per worker in the construction of utilization.

The Monte-Carlo simulation also allows us to assess the robustness of the max-share identification to alternative data generating scenarios. Generally, the max-share identification performs well as long as news shocks account for a large part of the unpredictable variation in adjusted TFP at long horizons (or whatever measure of productivity that one may choose). This performance gradually deteriorates as the importance of other shocks for long-run movements in adjusted TFP is increased, either because these shocks directly impact neutral technology or because of measurement issues. Nevertheless, as argued above, since long-term changes in productivity are typically slow to diffuse, the assumption that surprise (unanticipated) changes in productivity are important at long horizons seems unlikely. With regards to other shocks that impact adjusted TFP due to measurement error, this is of course a possibility, although one that is true of any identification of neutral technology shocks.

6 Conclusion

An almost universally imposed restriction in the news literature is that news shocks impact productivity only with a delay. This restriction may be violated if empirical series of productivity systematically mismeasure true technology.

In this paper, we document large revisions in one of the most popular measures of productivity, adjusted TFP by Fernald (2014), and show that these revisions are due to a switch in filtering of hours per worker in the estimation of factor utilization. These changes are evocative of cyclical mismeasurement and materially affect empirical conclusions about the macroeconomic effects of news shock as identified by Barsky and Sims (2011). We therefore propose an alternative identification, based on the max-share approach by Francis et al. (2013), which does not rely on short-run restrictions, in particular the zero impact restriction. We show that our identification is robust to the revisions in Fernald's series, and performs well in Monte Carlo simulations under different assumptions about cyclical mismeasurement of productivity. When applied to U.S. data, we find results that are consistent with a news interpretation: adjusted TFP increases gradually and with a significant lag whereas indicators of technological innovation and forward-looking information variables jump on impact. At the same time, the identified shock does not generate comovement in real macroeconomic aggregates and is therefore not a main driver of business cycle fluctuations. This does not imply that the shock is unimportant for macroeconomics. The shock accounts for the majority of unpredictable fluctuations in real aggregates at medium- and long horizons and generates strong impact responses of inflation, the Federal Funds rate, and asset prices. Investigating these results further and assessing the type of models that are consistent with these dynamics are important topics of future research.

References

Alexopoulos, M. (2011, June). Read All about It!! What Happens Following a Technology Shock? *American Economic Review 101*(4), 1144–79. 1, 5.3, 32

Bachmann, R., S. Elstner, and E. R. Sims (2013, April). Uncertainty and Economic Activity: Evidence from Business Survey Data. *American Economic Journal: Macroeconomics 5*(2), 217–49. 5.3

Baron, J. and J. Schmidt (2015). Technological Standardization, Endogenous Productivity and Transitory Dynamics. Working papers 503, Banque de France. 1, 5.3, 31

Baron, J. and D. F. Spulber (2015). Technology Standards and Standards Organizations: Introduction to the Searle Center Database. Technical report, Northwestern University Pritzker School of Law. 31

Barsky, R., S. Basu, and K. Lee (2015). Whither news shocks? In *NBER Macroeconomics Annual 2014, Volume 29*, NBER Chapters, pp. 225–264. National Bureau of Economic Research, Inc. 1, 1, 8, 3.2, 26

Barsky, R. and E. Sims (2011). News shocks and business cycles. *Journal of Monetary Economics 58*(3), 273–289. 1, 9, 3, 3.1, 3.2, 17, 3.2, 5.1, 5.2, 5.2, 6

Barsky, R. B. and E. R. Sims (2012, June). Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence. *American Economic Review 102*(4), 1343–77. 5.3

Basu, S., J. Fernald, J. Fisher, and M. Kimball (2013). Sector-specific technical change. Technical report. 2.3, 2.3, 4.1

Basu, S., J. G. Fernald, and M. S. Kimball (2006). Are technology improvements contractionary? *American Economic Review 96*(5), 1418–1448. 2, 2.1, 2.3, 4, 18, 4.1, 4.1, 21

Beaudry, P. and B. Lucke (2010). Letting difference views about business cycles compete. In *NBER Macroeconomics Annual 209, Volume 23*, NBER Chapters, pp. 413–455. National Bureau of Economic Research, Inc. 14

Beaudry, P. and F. Portier (2006). Stock prices, news, and economic fluctuations. *American Economic Review 96*(4), 1293–1307. 1, 3, 14, 3.2, 4.1, 5, 5.3, 5.3

Beaudry, P. and F. Portier (2014). News-driven business cycles: Insights and challenges. *Journal of Economic Literature 52*, 993–1074. 1

Bocola, L., I. Manovskii, and M. Hagedorn (2014). Identifying Neutral Technology Shocks. Technical report, Department of Economics, University of Pennsylvania. 27

Burnside, C., M. Eichenbaum, and S. Rebelo (1993). Labor hoarding and the business cycle. *Journal of Political Economy 101*(2), 245–273. 18

Cascaldi-Garcia, D. (2017). News shocks and the slope of the term structure of interest rates: Comment. *American Economic Review 107*(10), 3243–49. 9

Chang, A. C. and P. Li (2018, July). Measurement error in macroeconomic data and economics research: Data revisions, gross domestic product, and gross domestic income. *Economic Inquiry 56*(3), 1846–1869. 4

Chari, V., P. J. Kehoe, and E. R. McGrattan (2008, November). Are structural VARs with long-run restrictions useful in developing business cycle theory? *Journal of Monetary Economics 55*(8), 1337–1352. 5

Chen, K. and E. Wemy (2015). Investment-specific technology shocks: The source of long-run tfp fluctuations. *European Economic Review 80*, 230–252. 6, 33

Christiano, L. J., M. Eichenbaum, and C. L. Evans (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy 113*(1), 1–45. 4

Christiano, L. J., M. Eichenbaum, and R. Vigfusson (2004). The response of hours to a technology shock: Evidence based on direct measures of technology. *Journal of the European Economic Association 2*, 381–395. 1, 6

Christiano, L. J., M. Eichenbaum, and R. Vigfusson (2006). Assessing structural VARs. *NBER Macreoconomics Annual 21*. 5

Christiano, L. J. and T. J. Fitzgerald (2003). The band pass filter. *International Economic Review 44*(2), 435–465. 2.3

Cochrane, J. H. (1994). Shocks. *Carnegie-Rochester Conference Series on Public Policy 41*, 295–364. 3

Comin, D. and M. Gertler (2006). Medium-term business cycles. *American Economic Review 96*(3), 523–551. 10, 5

Comin, D., M. Gertler, and A. M. Santacreu (2009). Technology innovation and diffusion as sources of output and asset price fluctuations. Working Paper 15029, National Bureau of Economic Research. 10

Comin, D. and B. Hobijn (2010). An exploration of technology diffusion. *American Economic Review 100*(5), 2031–2059. 10

Erceg, C. J., L. Guerrieri, and C. Gust (2005). Can Long-Run Restrictions Identify Technology Shocks? *Journal of the European Economic Association 3*(6), 1237–1278. 5

Fernald, J. (2014). A quarterly, utilization-adjusted series on total factor productivity. Federal Reserve Bank of San Francisco Working Paper Series 2012-19. (document), 1, 2, 11, 2.1, 6

Fisher, J. (2010). Comment on: Letting difference views about business cycles compete. In *NBER Macroeconomics Annual 2009, Volume 23*, NBER Chapters, pp. 457–474. National Bureau of Economic Research, Inc. 14

Francis, N., M. Owyang, J. Roush, and R. DiCecio (2013). A flexible finite-horizon alternative to long-run restriction with an appication to technology shocks. *Review of Economics and Statistics 96*(4), 638–648. 1, 5, 5.1, 5.2, 6

Gali, J. (1999). Technology, employment, and the business cycle: Do technology shocks explain aggregate fluctuations? *American Economic Review 89*(1), 249–271. 1, 5

Gort, M. and S. Klepper (1982). Time paths in the diffusion of product innovations. *Economic Journal 92*(367), 630–653. 2, 5

Griliches, Z. (1957). Hybrid corn: An exploration in the economics of technological changes. *Econometrica 25*(4), 501–522. 2, 5

Justiniano, A., G. Primiceri, and A. Tambalotti (2010). Investment shocks and business cycles. *Journal of Monetary Economics 57*(2), 132–145. 4, 4.2, 4.2, 5.3

Keane, M. and R. Rogerson (2012). Micro and macro labor supply elasticities: A reassessment of conventional wisdom. *Journal of Economic Literature 50*(2), 464–76. 4.2

Kurmann, A. and E. Mertens (2014). Stock prices, news, and economic fluctuations: Comment. *American Economic Review 104*(4), 1439–1445. 14

Kurmann, A. and C. Otrok (2013). News shocks and the slope of the term structure of interest rates. *American Economic Review 103*(6), 2612–2632. 9, 5.3

Kurmann, A. and C. Otrok (2017a). News shocks and inflation: Lessons for new keynesians. 22

Kurmann, A. and C. Otrok (2017b). News shocks and the slope of the term structure of interest rates: Reply. *American Economic Review 107*(10), 3250–56. 9

Kydland, F. E. and E. C. Prescott (1982). Time to build and aggregate fluctuations. *Econometrica 50*(6), 1345–1370. 2

Leeper, E. M., T. B. Walker, and S.-C. S. Yang (2013). Fiscal foresight and information flows. *Econometrica 81*(3), 1115–1145. 5.3

Lindé, J. (2009). The effects of permanent technology shocks on hours: Can the rbc-model fit the var evidence? *Journal of Economic Dynamics and Control 33*, 597–613. 1, 5

Long, J. B. and C. Plosser (1983). Real business cycles. *Journal of Political Economy 91*(1), 39–69. 2

Mansfield, E. (1961). Technical change and the rate of imitation. *Econometrica 29*(4), 741–766. 2, 5

Mansfield, E. (1989). The diffusion of industrial robots in japan and the united states. *Research Policy 18*(4), 183–192. 2, 5

Pigou, A. C. (1927). *Industrial Fluctuations*. London: Macmillan. 1

Rogers, E. (1995). *Diffusion of Innovations*. New York: Free Press. 2, 5

Rotemberg, J. (2003). Stochastic technical progress, smooth trends, and nearly distinct business cycles. *American Economic Review 93*(5), 1543–1559. 1, 5

Smets, F. and R. Wouters (2007). Shocks and frictions in US business cycles: A bayesian DSGE approach. *American Economic Review 97*(3), 586–606. 4

Stock, J. H. and M. W. Watson (2012). Disentangling the channels of the 2007-2009 recession. *Brookings Papers on Economic Activity* (Spring), 81–141. 2.3

Uhlig, H. (2003). What moves real GNP? Working paper, Humboldt University. 1, 5.1

Uhlig, H. (2004). Do technology shocks lead to a fall in total hours worked? *Journal of the European Economic Association 2*(2-3), 361–371. 27