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

Andre Kurmann
Drexel University

Eric Sims
University of Notre Dame & NBER

August 8, 2016

Abstract

This paper documents that revisions to a widely-used quarterly utilization-adjusted TFP series produced by Fernald (2014) have significantly changed the cyclical behavior of the series. We show that these changes, which are primarily due to new estimates for the utilization adjustment, substantially affect the conclusions about the macroeconomic effects of news shocks as estimated for example by Barsky and Sims (2011). These differences in results cast doubt on the use of utilization-adjusted TFP as a measure of exogenous productivity and challenge the key identifying restriction of the news literature that utilization-adjusted TFP does not contemporaneously respond to news shocks. Building on the large empirical literature on the slow, S-shaped diffusion of new technology, we propose an alternative identification that does not rely on zero impact restrictions and instead accounts for most of unpredictable movements in TFP at long horizons. We show that this alternative identification is robust to revisions in Fernald's TFP series as well as a number of other measurement issues. When applied to U.S. data, we find that news about future productivity fail to generate comovement in the main macro aggregates and therefore do not constitute a main source of business cycle fluctuations. Nevertheless, news shocks account for a large part of of macroeconomic fluctuations at medium and longer horizons and generate sharp impact responses of prominent macro aggregates, including inflation and interest rates.

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, 710 Flanner Hall, Notre Dame, IN 46556 (esims1@nd.edu). We are grateful to John Fernald for helpful conversations and for making some of the past vintages of data available to us. We are grateful to Christopher Otrok for earlier involvement on this project. We also wish to thank Rudi Bachmann and Deokwoo Nam for helpful discussions concerning this project. Carlos Rondon Moreno provided capable research assistance.

1 Introduction

Dating back to Pigou (1927), economists have argued that news about future changes in fundamentals are an important source of economic fluctuations. This view has reemerged recently in large part due to a set of influential papers by Beaudry and Portier (2004) and Beaudry and Portier (2006), who report that news about future productivity improvements are closely related to shocks driving long-run variations in productivity and constitute one of the main drivers of business cycles. While the importance of news shocks for understanding business cycles is hotly debated in the literature,[1] the main identifying assumption behind news shocks is almost universally accepted: productivity does not react contemporaneously to a news shock.

In this paper, we critically revisit this assumption. Two conditions have to be met for the zero restriction on productivity to be satisfied. First, news about future productivity must affect productivity only with a lag. Second, the empirical measure of productivity must be exogenous to current business cycle conditions. Our focus is primarily on the second condition. We start by documenting large revisions in a utilization-adjusted series of total factor productivity (TFP) produced by Fernald (2014) that has served as one of the main measures of productivity in the news literature. These revisions affect the empirical results about the effects of news shocks in important ways. We argue that these revisions are indicative of pervasive measurement issues that make TFP as constructed by Fernald (2014) and others likely to be systematically affected by current business cycle conditions. But then, the assumption that news shocks have zero contemporaneous impact on TFP is not appropriate and renders results sensitive to even seemingly innocuous data revisions.[2]

In light of these considerations, we propose an alternative identification that does not rely on news shocks having zero contemporaneous impact on TFP. Building on the large empirical literature on the slow, S-shaped diffusion of new technology, we instead extract innovations that account for most of unpredictable movements in TFP at long horizons. These innovations are interpreted as news because they lead to persistent and therefore predictable growth in TFP over time while accounting for only a small fraction of TFP fluctuations at short forecast horizons. The identification is robust to measurement issues that affect the cyclical component of productivity, including the revisions to Fernald's utilization-adjusted TFP series. Moreover, the identification allows us to avoid taking a stand on whether productivity reacts only with a lag to news or not. When applied to U.S. data, we find that news about future productivity fail to generate comovement in the main real macro aggregates and therefore do not constitute a main source of business cycle fluctuations. Nevertheless, news about future productivity account for a large part of macroeconomic fluctuations at medium and longer horizons and generate sharp short-term fluctuations in prominent macroeconomic aggregates, including inflation and interest rates, that are important to understand.

Section 2 starts by reviewing the measurement of productivity in the data. Following the lead of Kydland and Prescott (1982) and Long and Plosser (1983), the modern business cycle literature has typically measured

[1] See for example Beuadry and Lucke (2010), Fisher (2010), Barsky and Sims (2011), Kurmann and Mertens (2014), Schmitt-Grohe and Uribe (2012), or Kahn and Tsoukalas (2012). Also see Barsky, Basu, and Lee (2015) or Beaudry and Portier (2014) for excellent reviews of this literature.

[2] As an example, suppose that news about future productivity lead to an immediate reaction in factor utilization. Then, if utilization is not measured appropriately, utilization-adjusted TFP reacts on impact of the shock even if underlying productivity does not.

productivity as the residual of aggregate output not accounted for by capital and labor inputs – commonly known as TFP – and modeled productivity shocks as a jump process consistent with the near random walk properties of TFP in the data. Economists quickly realized, however, that TFP may be a poor measure of productivity due to imperfect input and output markets, adjustment costs, and unobserved utilization of factors of production.[3] Partly in response to these criticisms, Basu, Fernald, and Kimball (2006) apply restrictions derived from economic theory to industry-level data to construct an aggregate measure of TFP that is adjusted for compositional changes in the quality of labor and capital, sectoral heterogeneity, imperfect competition as well as unobserved factor utilization. Fernald (2014) extends the analysis in Basu et al. (2006), which is carried out with annual data, to produce a quarterly measure of TFP. Because of the higher frequency, not all of the corrections in the original Basu et al. (2006) series are captured in the quarterly series, but perhaps the most important one – the correction for variable factor utilization – is. This quarterly, utilization-adjusted TFP series, which is freely available for download on Fernald's website, has proven highly influential and has become one of the most popular if not the main measure of productivity in the news literature.[4]

Fernald frequently updates his quarterly TFP series based on data revisions and methodological refinements.[5] In Section 3, we analyze the consequences of these updates. We document that while most updates affect the time series behavior only insignificantly, a switch in March 2014 from industry utilization estimates by Basu, Fernald, and Kimball (2006) to estimates from Basu, Fernald, Fisher, and Kimball (2013) had a large impact on aggregate utilization and therefore utilization-adjusted TFP. This is reflected in economically important differences in correlations of pre- and post-March 2014 vintages of adjusted TFP with key macroeconomic variables. For example, while pre-March 2014 vintages of adjusted TFP are positively correlated with aggregate output and mildly negatively correlated with total hours worked, post-March 2014 vintages are uncorrelated with output and strongly negatively correlated with total hours worked.[6]

The change in time series behavior across vintages of adjusted TFP has important consequences for the identification of news shocks. We illustrate this in Section 4 by redoing the empirical analysis of Barsky and Sims (2011) who identify news shocks in vector autoregression (VAR) models as the innovations that are orthogonal to Fernald's adjusted TFP seres but maximally account for the forecast error variance (FEV) share of adjusted TFP over a ten year horizon. Using the 2007 vintage as in Barsky and Sims (2011), we find results that are very similar to those reported in their paper. Consumption increases on impact of a positive

[3]The idea that observed TFP fluctuations might in fact be driven by an endogenous response of factor utilization to non-productivity shocks is mentioned by Summers (1986) in his early critique of real business cycle models. Burnside, Eichenbaum, and Rebelo (1993) construct a structural model with labor hoarding (another term for labor utilization) and conclude that much of the variation in TFP is not due to exogenous productivity shocks.

[4]As of August of 2016, the working paper describing the construction of the adjusted TFP series has been cited 236 times on Google Scholar.

[5]In the text and the notes accompanying the data, Fernald describes in detail how the construction of the series has varied over time. He does not, however, analyze the extent to which the different updates change the time series properties of the TFP series across vintages.

[6]In contemporaneous work, Cascaldi-Garcia (2016)also points out the differences in Fernald's utilization-adjusted TFP series across vintages and documents that these differences come from changes in measured factor utilization. The paper does not extensively document differences in the utilization-adjusted TFP series by vintage as we do, nor does the paper discuss why these changes may raise questions about the zero restriction employed in the literature. Instead, the paper is written purely as a comment on Kurmann and Otrok (2013) to which Kurmann and Otrok (2016) respond using the alternative identification approach proposed here.

news shock; output, investment, hours increase only with a delay; and inflation and nominal short-term rates drop markedly impact. Furthermore, the news shock accounts for a small fraction of the FEV of real aggregates at horizons before adjusted TFP begins to change. In contrast, using post-March 2014 vintages of the adjusted TFP series, we find that a positive news shock results in a large and immediate increase in hours and impact increases in output and investment; no longer leads to a significant drop in inflation and short-term interest rates; and accounts for a healthy share of the FEV at short horizons. These results are markedly different and would afford a more favorable interpretation of the news-driven business cycle hypothesis laid out by Beaudry and Portier (2006) than what obtains using older, pre-March 2014 vintages.

Given the large differences in the properties of the adjusted TFP series across vintages, it might seem natural to ask which vintage of the data is "best". However, we believe that at least for the purpose of identifying news shocks, a more fruitful approach is to consider alternative identifications of news shocks that do not rely on contemporaneous orthogonality restrictions and are therefore robust to mismeasurement of the cyclical component of TFP. The identification we propose in Section 5 is based on the observation that the type of news about productivity-enhancing changes in technology, factor input quality, and efficiency in allocation that we intend to capture is a major driver of long-run growth but exhibits slow, S-shaped patterns of dissemination. See for example Griliches (1957), Mansfield (1961), Mansfield (1989), Gort and Klepper (1982), or Rogers (1995). We implement this identification by extracting the shock that accounts for most of unpredictable movements in TFP at a long but finite forecast horizon without imposing a contemporaneous zero restriction. Through a host of Monte-Carlo simulations, we show that this alternative identification is robust to miseasurement of cyclical variables affecting TFP, including factor utilization, as well as an environment in which news shocks have an immediate effect on productivity. At the same time, nothing guarantees that the resulting shock captures news in the sense that that underlying changes in productivity are slowly disseminating. We find, however, that the shock anticipates predictable changes in TFP growth, and accounts for almost none of the fluctuations of TFP at short horizons but more than 70 percent of longer-run fluctuations. Hence, the news interpretation is natural. Further, we show that shock generates responses of macroeconomic aggregates that are indeed robust to the use of different vintages of Fernald's TFP series. In response to a positive shock, consumption jumps on impact before gradually increasing to a new permanent level. Investment and hours, by contrast, first decline modestly before starting to increase after a few quarters. The identified shock therefore does not generate the type of comovement that we typically associate with business cycle fluctuations. In turn, inflation and short-term interest rates drop markedly on impact of the shock before slowly returning to their initial values. Since the shock accounts for a substantial fraction of the variations in these variables, this reaction is important to understand.

The alternative identification we propose shares obvious parallels with work outside of the news literature on the drivers of long-term growth. In particular, Blanchard and Quah (1989), Shapiro and Watson (1988) and Gali (1999) among many others identify permanent technology shocks in VAR systems based on infinite-horizon restrictions. In response to criticism that infinite-horizon restrictions imply potentially large biases in finite order VARs, Francis, Owyang, Roush, and DiCecio (2013) propose to identify permanent technology

shocks at long but finite horizons. The difference to our work is that none of these papers directly target TFP and instead focus on labor productivity and other macroeconomic aggregates. As such these papers do not directly speak to the news literature and the idea that improvements in productivity disseminate slowly and in a predictable manner.[7] Our paper also relates to a recent literature on the macroeconomic effects of slowly disseminating technology due to costly adoption. The paper most closely related is Rotemberg (2003) who 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. 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 Utilization-adjusted TFP and its use in the news literature

This section begins by briefly describing the construction of the quarterly utilization-adjusted TFP series by Fernald (2014). We then review how the literature uses this series to identify news shocks, and how different measurement issues in the construction of the series raise potential challenges for this identification.

2.1 Fernald's utilization-adjusted TFP series

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

$$Y_t = F(E_t L_t, Z_t K_t, A_t), \tag{1}$$

where Y_t denotes output, L_t labor input, K_t capital input, E_t labor effort, Z_t capital utilization, and A_t technology. As discussed for example in Acemoglu (2009), A_t should be understood simply as a shifter of the production function that captures a *"...broad notion of technology, incorporating the effects of the organization of production and of markets on the efficiency with which the factors of production are utilized (page 28)."* We call A_t technology rather than productivity because we want to distinguish it from Total Factor Productivity (TFP), which is an empirical concept discussed below, and because the news literature typically treats A_t as exogenous to current business cycle conditions.

Differentiating with respect to time and omitting time subscripts to simplify notation, the contribution of technology to output growth can be expressed as

$$\frac{\dot{A}}{A} = \frac{\dot{Y}}{Y} - \varepsilon_E \frac{\dot{E}}{E} - \varepsilon_L \frac{\dot{L}}{L} - \varepsilon_Z \frac{\dot{Z}}{Z} - \varepsilon_K \frac{\dot{K}}{K}, \tag{2}$$

[7] Indeed, in many of these papers, labor productivity jumps immediately to essentially its new permanent level – a pattern that is very different from TFP as we document in our results.

where $\varepsilon_E \equiv F_E E/Y$ is defined as the elasticity of output with respect to labor effort and so forth for the other arguments of the production function, and the elasticity of output with respect to technology is normalized to one. Assuming further that input and output markets are perfectly competitive, that inputs can be instantaneously adjusted without cost, and that production is constant returns to scale, (2) can be rewritten as

$$\frac{\dot{A}}{A} = \left(\frac{\dot{Y}}{Y} - \omega_L \frac{\dot{L}}{L} - (1-\omega_L)\frac{\dot{K}}{K}\right) - \left(\omega_L \frac{\dot{E}}{E} + (1-\omega_L)\frac{\dot{Z}}{Z}\right), \tag{3}$$

where $\omega_L = \frac{WL}{PY}$ is the share of nominal labor payments, WL, in nominal output, PY. 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.

Based on (3), Fernald (2014) constructs quarterly estimates of both TFP growth and utilization-adjusted TFP growth. To construct TFP growth, he proceeds as follows. Output growth is measured as the log change in the equally weighted average of real expenditures and income in the business sector from the NIPAs. Labor input growth is measured as the sum of the log change in total hours worked in the business sector from the BLS and the log change in labor quality, which is based on worker skill estimates from wage regressions by Aaronson and Sullivan (2001) and BLS multifactor productivity data. Capital input growth is measured as the weighted log change of different capital stocks, with the weights determined by the relative income shares, and the different capital stocks computed from NIPA investment data using the perpetual inventory method. The labor share of income $\omega_{L,t}$ is constructed using interpolated annual NIPA data on payments to labor. The output elasticity of capital input, in turn, is set equal to $1 - \omega_{L,t}$. The quarterly estimate of TFP growth is computed as:

$$\Delta \ln TFP_t = \Delta \ln Y_t - \omega_{L,t} \Delta \ln L_t - (1-\omega_{L,t}) \Delta \ln K_t. \tag{4}$$

To construct utilization-adjusted TFP growth, Fernald (2014) follows Basu, Fernald, and Kimball (2006) and estimates an aggregate quarterly utilization series from disaggregated industry-level data. The basic idea behind this utilization series is that a cost-minimizing firm simultaneously varies inputs along all margins. This implies that variations in observed hours per worker can be used as a proxy for variations in the unobserved utilization rates, with the factor of proportionality estimated from industry-level data. The resulting industry utilization estimates are then aggregated to obtain a quarterly economy-wide utilization estimate using average industry weights, and utilization-adjusted TFP growth is computed as:

$$\Delta \ln TFP_t^U = \Delta \ln TFP_t - \Delta \ln U_t,$$

where $\Delta \ln U_t$ denotes the estimated change in factor utilization. Since March of 2014, Fernald (2014) computes the utilization series from updated utilization estimates by Basu, Fernald, Fisher, and Kimball (2013), which extends the Basu et al. (2006) methodology to more recent data. Since August 2014, the utilization series are based on updated industry level data of hours per worker.

2.2 Measurement Issues for the News Literature

Following Beaudry and Portier (2006), the news literature has generally assumed that technology is driven by two exogenous components, one capturing the slow diffusion of new technologies and the other capturing unanticipated or current shocks. In particular:

$$\ln A_t = d(L)\varepsilon_t^{news} + v(L)\varepsilon_t^{current}, \tag{5}$$

where ε_t^{news} and $\varepsilon_t^{current}$ are uncorrelated innovations, and $d(L)$ and $v(L)$ are lag polynomials governing the dynamic effects of the two innovations. Almost without exception, the key assumption that the news literature imposes on this process is $d(0) = 0$, identifying ε_t^{news} as the news shock in the sense that it affects technology only with a lag.

The assumption of $d(0) = 0$ naturally gives rise to a zero restriction that has been imposed in all empirical applications we are aware of. For example, Beaudry and Portier (2006) identify the news shock in their VARs as the shock that is orthogonal to a conventional measure of current TFP. Beuadry and Lucke (2010) and Barsky and Sims (2011) use a similar identification in VAR models using a utilization-corrected measure of TFP. In a fully-specified DSGE model, Schmitt-Grohe and Uribe (2012) assume that the news shock affects technology only with a lag.

As discussed in the Introduction, two important conditions have to be met for this zero restriction to be appropriate. News about future productivity affect productivity only with a lag; and productivity as measured in the data is unrelated to current business cycle conditions. However, there is no *a priori* reason to think that news shocks about future productivity should be completely unrelated to current productivity. Or as Barsky, Basu, and Lee (2015) put it: *"It is possible that news about future productivity arrives along with innovations in productivity today (page 233)*". Furthermore, if the TFP series used is an imperfect measure of technology, then TFP might be systematically related to current business cycle fluctuations, including fluctuations induced by news shocks. This could arrise for a number of reasons. First, the utilization series used to correct TFP series may be an inaccurate measure of actual factor utilization. Second, as emphasized by Fernald (2014) himself, *"...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...[s]imilarly,if observed factor shares do not equal output elasticities – as is the case with imperfect competition – then those effects will also show up in utilization-adjusted TFP growth (page 26)."*

This is not to say that utilization-adjusted TFP ceases to be an object of interest – it very much remains so even for this paper. Our point is simply that for a number of reasons, utilization-adjusted TFP may not be orthogonal to news shocks, which makes imposing the zero impact restriction to identify news shocks problematic. One could argue of course that these departures from orthogonality are small and do not matter quantitatively. In the next section, we show, however, that this is not generally the case. We illustrate that an otherwise unnoticed revision to Fernald (2014)'s factor utilization estimates result in large differences in

both the unconditional properties of adjusted TFP and large differences in estimated impulse responses to a news shock when imposing a zero impact effect of news.

3 Revisions to Utilization-Adjusted TFP

The utilization-adjusted measure of TFP produced by Fernald (2014) has been publicly available for several years. As we noted above, in March of 2014 he switched to a different methodology for measuring factor utilization. In a recent working paper, Sims (2016) documents that this revision in utilization is far from innocuous. It significantly changes the time series properties of the resulting adjusted TFP series and has potentially important implications for the empirical literature seeking to estimate the macroeconomic effects of news shocks. We briefly review the evidence here.

3.1 Data

We have vintages of Fernald's utilization-adjusted TFP series from the calendar years 2007, 2011, 2013, 2014, 2015, and 2016. For the purposes of the our analysis in this paper, we focus on the earliest and latest available vintages (i.e. 2007 and 2016). The 2011 and 2013 vintages are very similar to the 2007 vintage, while the 2014 and 2015 vintages are quite similar to the 2016 vintage.

Table 1 provides some basic unconditional moments of the adjusted TFP series for both the 2007 and 2016 vintages. We do so both for the log first differences of the series as well as for the HP cyclical component of the log levels, where we use a smoothing parameter of 1600.

Table 1: Moments for the 2007 and 2016 Vintages of the Adjusted TFP Series

	First Differences			HP Cyclical Component	
	$\Delta \ln TFP_t^{U,07}$	$\Delta \ln TFP_t^{U,16}$		$\ln TFP_t^{U,07}$	$\ln TFP_t^{U,16}$
Mean	0.0037	0.0036	Mean	-0.0000	-0.0000
Standard Deviation	0.0085	0.0087	Standard Deviation	0.0097	0.0100
Corr w/ $\Delta \ln TFP_t^{U,07}$	1.0000	0.5840	Corr w/ $\ln TFP_t^{U,07}$	1.0000	0.6637

Notes: This table shows descriptive statistics for the different vintages of adjusted TFP. The sample period for these statistics is fixed at 1947q3-2007q3. $\Delta \ln TFP_t^{U,j}$ is the log first difference of the adjusted TFP series, for vintages $j = 07$ or 16. $\ln TFP_t^{U,j}$ is the HP cyclical component of the log level (smoothing parameter 1600) for each vintage.

The means (of the log first differenced series) and standard deviations (both in log first differences and HP cyclical components) are very similar across vintages. In spite of these similarities, the two vintages of adjusted TFP are not strongly correlated with one another. In log first differences, the correlation between the two different vintages is 0.58. In terms of the HP cyclical component, the correlation is a bit higher at 0.66 but still far below 1. The differences in the two vintages of the adjusted TFP series are not driven by particular parts of the sample – the relatively weak correlation between the two vintages holds in various different subsamples.

Table 2 documents that the large apparent differences in the two vintages of the utilization adjusted TFP series matter in a significant way for some basic business cycle correlations. In log first differences, the 2007 vintage of adjusted TFP is strongly positively correlated with output, consumption, and investment, and essentially uncorrelated with the first difference of aggregate labor hours. The 2016 vintage of adjusted TFP, in contrast, is only weakly positively correlated with the log first differences of output and consumption, is uncorrelated with the log first difference of investment, and is strongly negatively correlated with the log first difference of aggregate labor hours. A similar pattern emerges when focusing instead on the HP cyclical component of the log levels of the different series. The correlations of the cyclical component of adjusted TFP with output, consumption, investment, and labor hours decrease significantly (in most cases going from positive to negative) when moving from the 2007 to 2016 vintages of the adjusted TFP series.

Table 2: Business Cycle Moments for the 2007 and 2016 Vintages of the Adjusted TFP Series

	First Differences		HP Cyclical Component	
	$\Delta \ln TFP_t^{U,07}$	$\Delta \ln TFP_t^{U,16}$	$\ln TFP_t^{U,07}$	$\ln TFP_t^{U,16}$
$\Delta \ln Y_t$	0.5280	0.0734		
$\Delta \ln C_t$	0.2599	0.1102		
$\Delta \ln I_t$	0.1987	-0.0173		
$\Delta \ln L_t$	-0.0103	-0.3542		
$\ln Y_t$			0.1519	-0.1867
$\ln C_t$			0.1521	-0.0365
$\ln I_t$			0.1458	-0.0619
$\ln L_t$			-0.1025	-0.3849

Notes: This table shows correlations between two vintages of the adjusted TFP series with several difference aggregate variables. Y_t is the headline number from the NIPA tables. C_t is personal consumption expenditures from the NIPA tables, and I_t is private fixed investment. Both personal consumption expenditures and private fixed investment are deflated by their own deflators. L_t is hours worked in the non-farm business sector. Δ denotes the first difference operator. The absence of the first difference operator denotes the HP cyclical component (smoothing parameter 1600). Statistics are generated over the sample period 1947q3-2007q3.

Sims (2016) documents that the large differences in the properties of the utilization-adjusted TFP series by vintage do not result because of data revisions to output, hours, or capital measures, but rather are a direct consequences of the switch in March 2014 to using a different approach to infer aggregate factor utilization. This is documented in Table 3. Whether measured in first differences or the HP cyclical component of the log level, the 2016 vintage of the utilization series is significantly more volatile than the 2007 vintage. In first differences, the correlation between the two vintages of the utilization series is only 0.65. It is somewhat higher, though still significantly less than 1, when looking at the HP cyclical component of the two vintages of the utilization series.

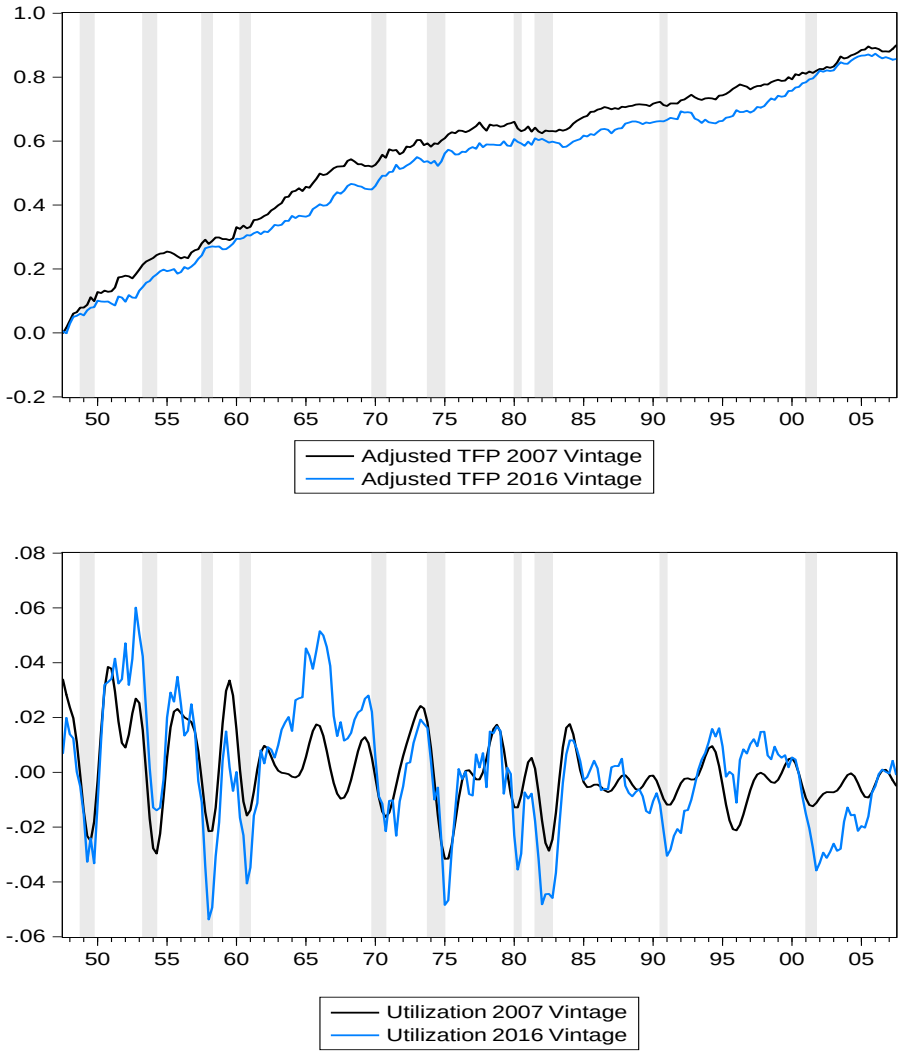
Table 3: Moments of the 2007 and 2016 Vintages of the Utilization Series

	First Differences			HP Cyclical Component	
	$\Delta \ln u_t^{07}$	$\Delta \ln u_t^{16}$		$\ln u_t^{07}$	$\ln u_t^{16}$
Mean	-0.0002	-0.0001	Mean	0.0000	0.0000
Standard Deviation	0.0058	0.0094	Standard Deviation	0.0120	0.0147
Corr w/ $\Delta \ln u_t^{07}$	1.0000	0.6537	Corr w/ $\ln u_t^{07}$	1.0000	0.8815

Notes: This table shows descriptive statistics for the different vintages of the utilization series. The sample period for these statistics is fixed at 1947q3-2007q3.

To visualize the differences according to vintage, Figure 1 plots both the 2007 and 2016 vintages of the adjusted TFP series (upper panel) and the utilization series (lower panel). The 2007 vintages are depicted with the black lines, while the 2016 vintages are shown with blue lines. While the two vintages of the adjusted TFP series share roughly the same trend, some differences are apparent. For example, the 2016 vintage of adjusted TFP grew much faster than the 2007 vintage from the late 1990s through the mid-2000s. Whereas the 2016 vintage was decline from 2005 onwards, the 2007 vintage showed an uptick near the end of the sample. Even stronger differences are apparent for the utilization series. While each vintage is clearly procyclical, the 2016 vintage is more volatile and significantly less smooth than the 2007 vintage.

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



Notes: This figure plots the log levels of both the 2007 and 2016 vintages of the utilization-adjusted TFP series (black lines) and the utilization series (blue lines). The upper panel plots the log level of the adjusted TFP series, while the lower panel plots the (demeaned) log level of the utilization series. Shaded gray bars are recessions as identified by the NBER.

4 Implications for the Empirical Identification of News Shocks

In an influential paper, Barsky and Sims (2011) propose an identification procedure to identify news shocks. This identification procedure is a modification of similar approaches used to identify conventional technology shocks in Uhlig (2004) and Francis, Owyang, Roush, and DiCecio (2013). Barsky and Sims (2011) show that their procedure performs well in a Monte Carlo experiment, and it is not plagued by the lack of identification in the Beaudry and Portier (2006) approach, an issue which was pointed out in Kurmann and Mertens (2014).

Suppose that one has a vector of observed time series, $\mathbf{Y}_t$, the first variable of which is a measure of utilization-adjusted TFP. Let $\mathbf{Y}_t$ be dimension $k \times 1$ with T total observations. One can form a reduced form moving average representation of the this vector of variables by estimating an unrestricted VAR in levels.

The reduced form moving average representation is given by:

$$\mathbf{Y}_t = \mathbf{B}(\mathbf{L})\mathbf{u}_t, \tag{6}$$

where $\mathbf{u}_t$ is a $k \times 1$ vector of innovations in $\mathbf{B}(L) = \mathbf{I} + \mathbf{C}_1 L + \mathbf{C}_2 L^2 + \ldots$ is a matrix lag polynomial. It is assume that there exists a linear mapping between reduced form innovations and structural shocks, ϵ_t:

$$\mathbf{u}_t = \mathbf{A}\epsilon_t, \tag{7}$$

where $\mathbf{A}$ is a $k \times k$ matrix. A structural moving average representation is:

$$\mathbf{Y}_t = \mathbf{C}(\mathbf{L})\epsilon_t, \tag{8}$$

where $\mathbf{C}(\mathbf{L}) = \mathbf{B}(\mathbf{L})\mathbf{A}$ and $\epsilon_t = \mathbf{A}^{-1}\mathbf{u}_t$. Normalizing the variance-covariance matrix of structural shocks to be an identity matrix (i.e. $E\left(\epsilon_t \epsilon_t'\right) = \mathbf{I}$), it must be the case that:

$$\mathbf{A}\mathbf{A}' = \mathbf{\Sigma}_u, \tag{9}$$

where $\mathbf{\Sigma}_u$ is the variance-covariance matrix of reduced form innovations. $\mathbf{C}(\mathbf{L})$ and $\mathbf{\Sigma}_u$ can be recovered from the data using conventional estimation techniques, but there are a multitude of $\mathbf{A}$ consistent with (9). 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 (9) is given by $\widetilde{\mathbf{A}}\mathbf{D}$, where $\mathbf{D}$ is an orthonormal rotation vector.

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

$$\mathbf{Y}_{t+h} - E_{t-1}\mathbf{Y}_{t+h} = \sum_{l=0}^{h} \mathbf{B}_l \widetilde{\mathbf{A}}\mathbf{D}\epsilon_{t+h-l}. \tag{10}$$

The share of the forecast error variance of variable i attributable to shock j at horizon h is then:

$$\mathbf{\Omega}_{i,j}(h) = \frac{\sum_{l=0}^{h} \mathbf{B}_{i,l}\widetilde{\mathbf{A}}\gamma\gamma'\widetilde{\mathbf{A}}'\mathbf{B}_{i,l}'}{\sum_{l=0}^{h} \mathbf{B}_{i,l}\mathbf{\Sigma}_u\mathbf{B}_{i,l}'}, \tag{11}$$

where $\mathbf{B}_{i,l}$ is the ith row of lag polynomial evaluated at $L = l$ and γ is the jth column of $\mathbf{D}$. Barsky and Sims (2011) propose identifying the news shock by picking γ to maximize the sum of variance shares due to the news shock from horizons up to H, which is a truncation horizon. Formally, the problem is:

$$\max_{\gamma} \quad \sum_{h=0}^{H} \Omega_{1,2}(h) \tag{12}$$

$$\text{s.t.}$$

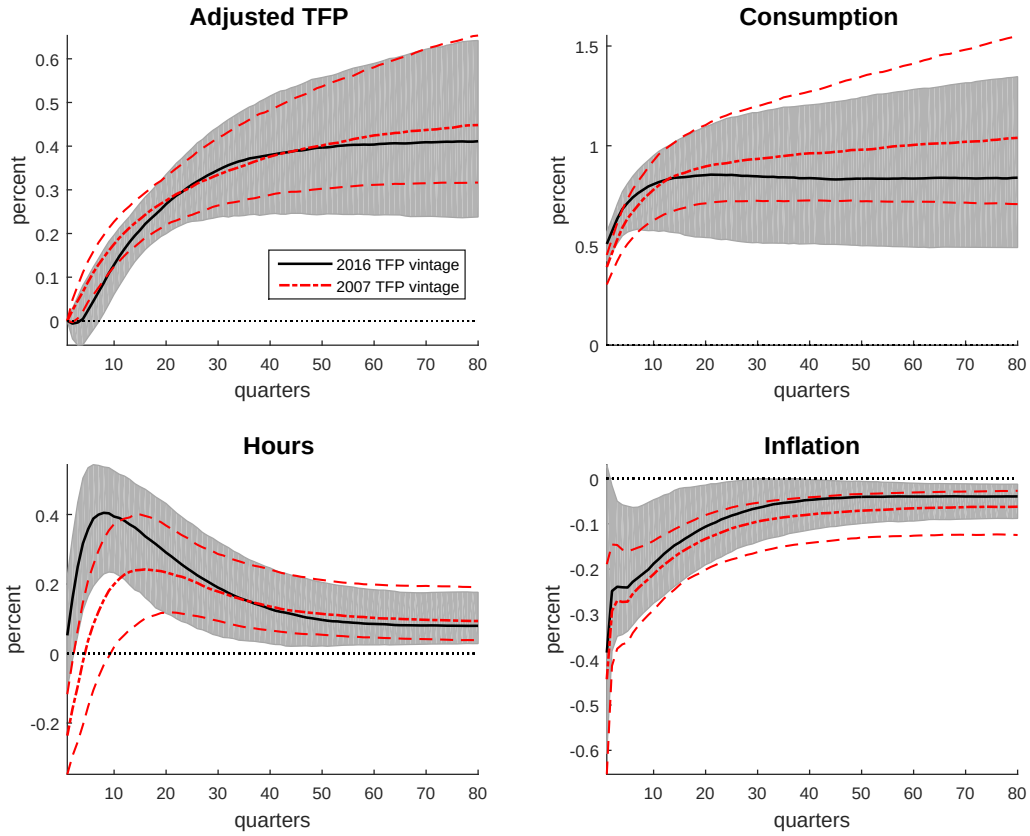
$$\gamma'\gamma = 1$$
$$\gamma(1,1) = 0 \tag{13}$$

As noted earlier, adjusted TFP occupies the first element of $\mathbf{Y}_t$, and the news shock is specified as the second structural shock in the system. The first restriction ensures that γ belongs to an orthonormal matrix, while the second restriction imposes that the news shock have no immediate impact on adjusted TFP. The assumption that the news shock is orthogonal with respect to current adjusted TFP is consistent with how the theoretical literature has thought about news shocks, but requires that the empirical measure of adjusted TFP accurately measure true technology. Otherwise, the impact orthogonality restriction is likely to be violated.

Figure 2 presents impulse responses to an identified news shock using the Barsky and Sims (2011) identification strategy. The variables included in the VAR are Fernald's utilization-adjusted TFP series (either the 2007 or 2016 vintage), real personal consumption expenditures, total hours worked per capita in the non-farm business sector, and inflation as measured by the growth rate of the GDP price deflator. With the exception of the inflation rate, the variables enter the VAR in log levels. Estimation is conduced with four lags via Bayesian methods subject to a Minnesota prior, with confidence bands computed by drawing from the resulting posterior distribution.[8] The sample period is fixed at 1960q1-2007q3, which is the last period of available data for the 2007 vintage of adjusted TFP. The truncation horizon is set to $H = 40$.

[8]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.

Figure 2: Empirical Responses, BS Identification, 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 BS identification.

In Figure 2, the 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. Using the 2007 vintage of adjusted TFP, the responses are similar to those estimated by Barsky and Sims (2011). Adjusted TFP slowly grows, consumption jumps up on impact, inflation falls significantly, and hours worked initially declines for several periods before turning positive after several quarters. Because of the impact decline in hours, Barsky and Sims (2011) conclude that news shocks cannot be a major driving force behind the business cycle. Qualitatively, at least, the impact decline in hours worked is consistent with the predictions of a relatively frictionless real business cycle model.

While the point estimates of the responses of adjusted TFP, consumption, and inflation are fairly similar across the two different vintages, the response of hours is quite different in an economically meaningful way

when the VAR is estimated with the 2016 vintage of adjusted TFP data. Hours worked now increases on impact and reaches a peak positive response only a few quarters after the shock. This difference relative to the response of hours with the 2007 vintage of adjusted TFP is economically meaningful because it (i) suggests that news shock might generate positive business cycle co-movement, and therefore could be an important contributor to business cycles, and (ii) suggests a different class of structural economic models which might be consistent with this pattern of responses. Furthermore, the impact decline in inflation is statistically insignificant when using the 2016 vintage of the adjusted TFP series. Barsky, Basu, and Lee (2015) cite the deflationary impacts of a news shock as one of the most robust features of the data. This is no longer true with the 2016 vintage of the adjusted TFP series.

5 An Alternative Identification of News Shocks

An apparently innocuous change in the construction of the utilization series significantly alters the unconditional moments of the Fernald's utilization-adjusted TFP series and leads to substantial differences in the estimated response to a news shock as identified by Barsky and Sims (2011). In this section, we propose an alternative identification of news shocks that turns out to provide consistent estimates of the responses to a news shock regardless of the vintage of adjusted TFP. We first motivate the alternative identification strategy and describe its implementation. Then, we present Monte Carlo simulations based on a medium-scale DSGE model to show that the proposed identification is more more robust to a variety of misspecifications than is the Barsky and Sims (2011) identification. Finally, we present results for the same data and VAR specification as above.

5.1 Motivation and implementation of alternative identification

An extensive empirical literature documents that new technologies diffuse slowly in an S-shaped pattern. See for example Griliches (1957), Mansfield (1961), Mansfield (1989), Gort and Klepper (1982) or Rogers (1995). According to Mansfield (1989), the time until half of potential adopters actually adopt a new technology varies between 5 and 15 years, depending on technology. This strongly contrasts with most of the business cycle literature, including much of the literature on news shocks, which typically models technology as a jump process where innovations lead to an immediate change of productivity to a new level that is either permanent or highly persistent. Exceptions to this practice are Rotemberg (2003)), Comin and Gertler (2006)), Comin and Hobijn (2010) or Comin, Gertler, and Santacreu (2009) who explicitly focus on the slow adoption of new technologies.

The alternative identification we propose takes the insights from this literature one step further. We posit – and verify later in our empirical application – that the slow dissemination of productivity-enhancing changes in technology, factor input quality and efficiency in allocation constitutes news because it leads to persistent and therefore predictable changes in productivity growth. We implement the identification of this news shock using similar VAR analytics as Barsky and Sims (2011) but with two important exceptions. The

first is that we identify the news shock by maximizing its contribution to the variance share of adjusted TFP at one long horizon, as opposed to maximizing the sum of variance shares from 0 onward. In practice, this difference turns out to be relatively unimportant. Second, we drop the zero restriction that the news shock has no effect on adjusted TFP on impact. Or, put differently, in contrast to the majority of work on news shocks to date, we do not impose that the news shock is orthogonal with respect to the innovation in adjusted TFP.

Formally, our identification of a news shock is given by the solution to the following optimization problem:

$$\max_{\gamma} \quad \Omega_{1,2}(H) \tag{14}$$

s.t.

$$\gamma'\gamma = 1, \tag{15}$$

Where γ is a column belonging to an orthonormal rotation matrix of the Choleski factor of the reduced form. It is restricted to have unit length, but we do not restrict it to have a zero in its first entry, meaning that we allow the news shock to impact adjusted TFP immediately. Also, we maximize the variance share at one horizon, H, instead of over many horizons.

As discussed in the Introduction, our alternative identification approach has two potentially important advantages over the existing practice of identifying news shocks based on the zero impact restriction that TFP does not contemporaneously react to news shocks. First, the identification allows us to avoid taking a stand on whether productivity reacts only with a lag to news or not. Second, if adjusted TFP is for any reason an imperfect measure of cyclical variations in technology, then imposing a zero impact restriction could induce potentially significant bias into the estimated response, as suggested by the results above based on different vintages of Fernald's utilization-adjusted TFP series.

Our identification strategy is essentially the same as the "max share" identification proposed in Francis, Owyang, Roush, and DiCecio (2013). Our work differs in that we include a measure of adjusted TFP, rather than average labor productivity, in our VAR system. One advantage of using a measure of TFP rather than labor productivity is that non-technology shocks may affect labor productivity in the long run, as argued by Uhlig (2004). Moreover, labor productivity is presumably affected in more important ways by current business cycle conditions than adjusted TFP. This can lead to important differences in the impact response to a shock. Indeed, the VAR results by Francis, Owyang, Roush, and DiCecio (2013)indicate that in response to their long-run shock, labor productivity jumps immediately to what is essentially as new permanent level. This would suggest that productivity may indeed be modeled as a jump process as discussed above. Our results below show that if instead the identification is applied to TFP, the results suggest that technology is slowly diffusing, consistent with the empirical literature.

In the next subsection, we provide some Monte Carlo results which speak to the desirable properties of our proposed identification. In particular, we show that it likely performs better than Barsky and Sims (2011). We then turn to the data. When using our identification strategy in place of Barsky and Sims (2011), we find that the differences arising from using different vintages of adjusted TFP largely disappear. While it appears that technology diffuses slowly, we consistently find that hours worked decline on impact.

5.2 Monte Carlo Results

For our Monte Carlo experiments, we consider a fairly standard medium scale New Keynesian DSGE model. The model is very similar to Christiano, Eichenbaum, and Evans (2005), Smets and Wouters (2007), and Justiniano, Primiceri, and Tambalotti (2010). It features monopolistic competition in both goods and labor markets, price and wage stickiness, and a number of other real frictions. We discuss only the key features of the model in the main body of the paper. The full set of equilibrium conditions in the model is given in Appendix A.

Abstracting from trend growth and price dispersion, the aggregate production function in the model is:

$$Y_t = A_t \left(u_t K_t\right)^{\alpha} L_t^{1-\alpha} - F, \tag{16}$$

where $F \geq 0$ is a fixed cost of production, K_t is physical capital, u_t is capital utilization, L_t is labor input, and Y_t is output. Although Fernald's empirical work allows for time-varying utilization of both capital and labor, in our model there is only utilization of capital. A_t is aggregate technology. We assume that log technology is the sum of a stationary and a permanent component:

$$\ln A_t = \ln S_t + \ln \Gamma_t, \tag{17}$$

where S_t is the stationary component of technology and Γ_t is the permanent component. The stationary component obeys an AR(1) process, and the permanent component an AR(1) process in the growth rate:

$$\ln S_t = \rho_S \ln S_{t-1} + s_S \varepsilon_{S,t}, \tag{18}$$

$$\ln \Gamma_t - \ln \Gamma_{t-1} = (1-\rho_\Gamma) \ln g + \rho_\Gamma \left(\ln \Gamma_{t-1} - \ln \Gamma_{t-2}\right) + s_g \varepsilon_{g,t-q}. \tag{19}$$

In (18)-(19), the autoregressive parameters are restricted to lie between 0 and 1. The ϵ are innovatins drawn from standard normal distributions, and the s parameters are standard deviations. g is the steady state gross growth rate of technology. The innovation for the permanent component of technology is observed by agents $q \geq 0$ periods prior to impacting the level of aggregate technology. In the news literature it is typically assumed that $q > 0$. However, even with $q = 0$, the process given in (19) embodies some elements of slow diffusion if $\rho_\Gamma > 0$. In particular, if $\rho_\Gamma > 0$, then a positive shock to the permanent component of technology today portends even larger increases in the level of technology in the future. In practice, we will assume that $\rho_\Gamma > 0$, so we will refer to a shock to the permanent component of technology as a news shock

regardless of whether agents observe it in advance ($q > 0$) or not ($q = 0$).

We measure adjusted TFP in the model exactly as Fernald does in the data. In particular, the growth rate of model TFP is defined as:

$$\Delta \ln TFP_t^M = \Delta \ln Y_t - \omega_{L,t} \Delta \ln L_t - (1 - \omega_{L,t}) \Delta \ln K_t. \tag{20}$$

The growth rate of model adjusted TFP is then:

$$\Delta \ln TFP_t^{M,U} = \Delta \ln TFP_t^M - (1 - \omega_{L,t}) \Delta \ln u_t \tag{21}$$

Even if Y_t, L_t, K_t, and u_t are perfectly observed, there are nevertheless two potential incongruities between (21) and true technology in the model, (17). First, if the production function features a fixed cost then the construction of TFP is mis-specified. Second, $\omega_{L,t}$ and $1 - \omega_{L,t}$ will in general not correspond to the true factor elasticities in the model. Firms have market-power in setting their prices, though they are not freely able to adjust price each period due to a Calvo friction. Nevertheless, each period firms will choose inputs to minimize cost. Cost-minimization implies the following:

$$w_t L_t = (1 - \alpha) \mu_t^{-1} \left[Y_t - F \right]. \tag{22}$$

In (22), μ_t^{-1} is the inverse price markup over marginal cost. Because of monopoly power in price-setting, the average markup, μ^*, will be greater than 1. Though firms desire a constant markup, due to price stickiness the markup will be time-varying. Suppose that there is no fixed cost, i.e. $F = 0$. Then (22) can be written:

$$\omega_{L,t} = \frac{w_t L_t}{Y_t} = (1 - \alpha) \mu_t^{-1}. \tag{23}$$

While the the construction of TFP is not mis-specified with respect to the true production function when there is no fixed cost, average labor share will not correspond to $1 - \alpha$ due to the markup of price over marginal cost. Furthermore, measured labor share will vary across time due to undesired fluctuations in the markup owing to price rigidity. Suppose instead that the model features a fixed cost, and that the size of this fixed cost is chosen to ensure zero profit in steady state, which is a standard assumption used in the literature. This requires that $F = Y^*(\mu^* - 1)$. In this case, (23) can be written:

$$\omega_{L,t} = \frac{w_t L_t}{Y_t} = (1 - \alpha) \mu_t^{-1} \mu^*. \tag{24}$$

In this case, the average measured labor share will correspond to the factor elasticity $1 - \alpha$. But measure labor share will still fluctuate over time due to undesired fluctuations in the markup owing to price rigidity. Furthermore, will on average the measured labor share will be correct, with the fixed cost the construction of TFP is now mis-specified with respect to the true production function.

Before turning to Monte Carlo experiments to assess the adequacy of our VAR identification procedure, we first investigate the extent to which these mis-specifications matter for Fernald's adjusted TFP series. The parameterization of the model is given in Table 1 of the Appendix. In addition to the two technology shocks, the model features two other shocks. The first is a shock to the marginal efficiency of investment. This demand shock is parameterized to account for the majority of short run fluctuations in output, as in Justiniano et al. (2010). The second is a shock to the disutility from labor, and is isomorphic to the wage markup shock in Smets and Wouters (2007).

Table 4 presents correlations of the log first differences of true productivity in the model, $\Delta \ln A_t$, with Fernald's adjusted TFP series, $\Delta \ln TFP_t^{M,U}$, and a conventional measure of TFP, $\Delta \ln TFP_t^{M}$. The upper panel considers our baseline parameterization. In spite of the obvious measurement issues in the model related to Fernald's approach, the adjusted TFP series is almost perfectly correlated with true technology whether a fixed cost is included in the model or not. Measured TFP is less positively correlated with true technology. This is primarily due to fluctuations in utilization. Conventional TFP and adjusted TFP are positively correlated with one another, with a correlation coefficient of about 0.9. This is close to the correlation between the 2007 vintages of adjusted and conventional TFP in Fernald's data, which is 0.8. The correlation between the adjusted and conventional TFP series in the 2016 vintage is much lower, however, at 0.4.

Table 4: Correlations Between Different Productivity Measures

	Corr $\Delta \ln A_t$, $\Delta \ln TFP_t^{M,U}$	Corr $\Delta \ln A_t$, $\Delta \ln TFP_t^{M}$	Corr $\Delta \ln TFP_t^{M,U}$, $\Delta \ln TFP_t^{M}$
		Baseline	
No fixed cost	0.9995	0.8850	0.8846
Fixed cost	0.9942	0.8867	0.9213
		Low utilization cost, $\delta_2 = 0.00025$	
	Corr $\Delta \ln A_t$, $\Delta \ln TFP_t^{M,U}$	Corr $\Delta \ln A_t$, $\Delta \ln TFP_t^{M}$	Corr $\Delta \ln TFP_t^{M,U}$, $\Delta \ln TFP_t^{M}$
No fixed cost	0.9984	0.7711	0.7422
Fixed cost	0.9939	0.7775	0.8260
		Very sticky prices, $\theta_p = 0.95$	
	Corr $\Delta \ln A_t$, $\Delta \ln TFP_t^{M,U}$	Corr $\Delta \ln A_t$, $\Delta \ln TFP_t^{M}$	Corr $\Delta \ln TFP_t^{M,U}$, $\Delta \ln TFP_t^{M}$
No fixed cost	0.9996	0.8209	0.8228
Fixed cost	0.9879	0.8162	0.8813
		Very sticky wages, $\theta_w = 0.90$	
	Corr $\Delta \ln A_t$, $\Delta \ln TFP_t^{M,U}$	Corr $\Delta \ln A_t$, $\Delta \ln TFP_t^{M}$	Corr $\Delta \ln TFP_t^{M,U}$, $\Delta \ln TFP_t^{M}$
No fixed cost	0.9997	0.8532	0.8573
Fixed cost	0.9892	0.8492	0.9061

Notes: This table shows correlations of different measures of productivity from the solution of the medium scale DSGE model laid out in the text. $\Delta \ln A_t$ refers to the log first difference of true technology in the model, $\Delta \ln TFP_t^{M,U}$ refers to the log first difference of utilization-adjusted TFP as measured in the model, and $\Delta \ln TFP_t^{M}$ refers to the log first difference of TFP as measured in the model.

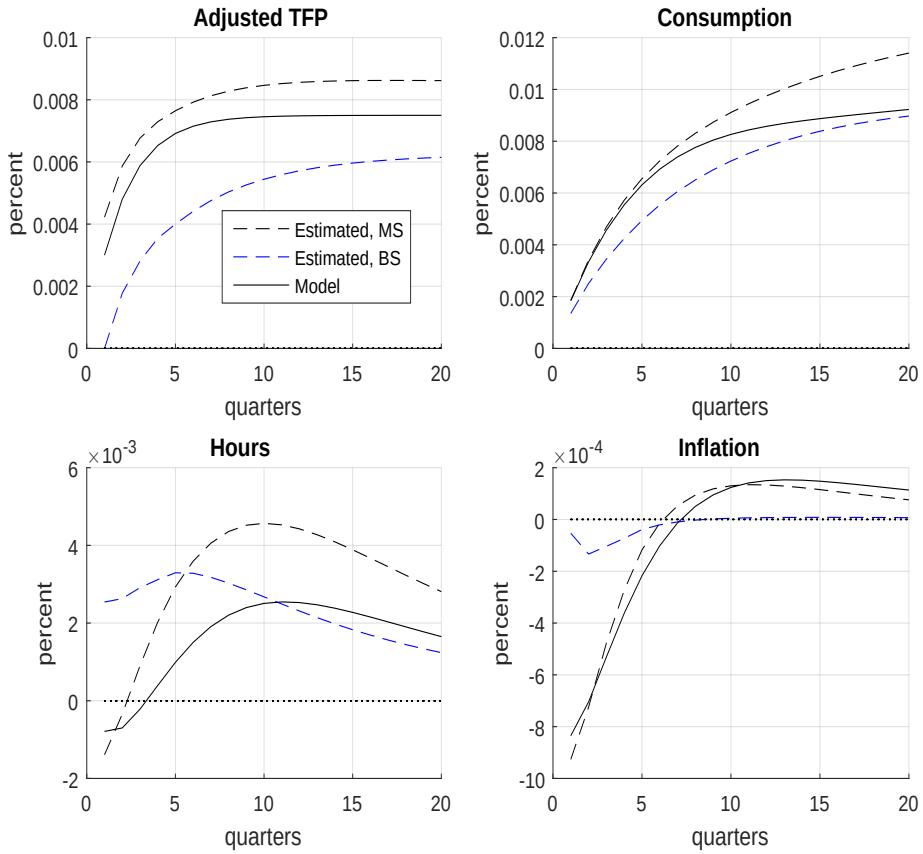
Our results hold qualitatively in several alternative parameterizations of the model – adjusted TFP is close to perfectly correlated with true technology regardless of parameter values and regardless of whether the model features a fixed cost of production or not. When utilization is less costly, the correlation between conventional TFP and true technology is lower. This is also true when prices or wages are stickier, which arises because factor utilization reacts more to non-technology shocks, driving a larger wedge. Nevertheless, our Monte Carlo experiments suggest that, in spite of the fact the assumptions underlying its construction are violated in a canonical macro model, Fernald's adjusted TFP series ought to be very close to true technology in the model, provided utilization and other factors are measured appropriately.

We next turn to Monte Carlo experiments designed to assess the adequacy of our proposed identification procedure, particular as it relates to Barsky and Sims (2011). Our experiments are conducted as follows. We simulate 10,000 periods of data from the model using the parameter values listed in Table 1. For these simulation, we assume that the model features a fixed cost of production, but our results are similar whether there is a fixed cost or not. From the simulation, we recover simulated time series of adjusted TFP, consumption, inflation, and hours worked. These are the same series we use in the empirical VAR from the previous section. On this one long data set, we estimate a VAR with four lags. We then identify a news shock using either our identification procedure (which we label "MS" for max-share) or Barsky and Sims

(2011), which we label "BS." In the figures below, solid lines correspond to the true impulse response to a news shock in the model. Dashed black lines are the estimated response from our max-share identification, while dashed blue lines are the responses under the Barsky and Sims (2011) identification. While the dashed lines are estimates of the response of adjusted TFP as measured from the model, the solid line in the panels labeled "Adjusted TFP" corresponds to true technology from the model.

We first consider a case in which $q = 0$, but $\rho_\Gamma = 0.6$. This means that the news shock portends future increases in technology which do not immediately materialize, but that the news shock does affect technology contemporaneously. Model and estimated response from model-generated data are shown in Figure 3. In the model, consumption increases, while hours and inflation decline on impact. The max-share identification captures these features well. The Barsky and Sims (2011) identification, which imposes orthogonality between the news shock and current adjusted TFP, performs considerably worse. This should not be particularly surprising, as the assumption of impact orthogonality is inconsistent with the structure of the model. The Barsky and Sims (2011) identification gets the sign of the impact response of hours incorrect and delivers little or no response of inflation. While the max-share identification is not perfect, qualitatively it generates the correct pattern of responses and performs much better than Barsky and Sims (2011).

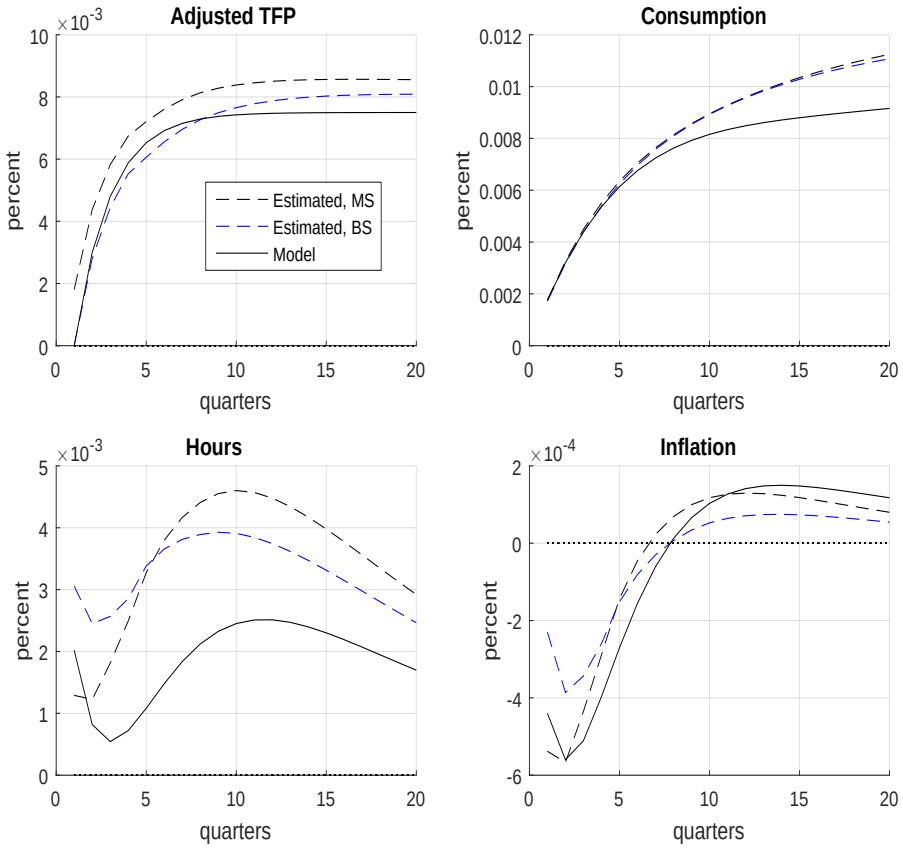
Figure 3: Model and Estimated Responses to a News Shock



Notes: Solid lines are the true impulse responses to an estimated growth shock in the model. The dashed black lines are the estimated responses using the max share identification at a 80 quarter horizon, without imposing impact orthogonality. This is labeled "MS" for "Max Share." The dashed blue lines are the estimated responses using the BS identification with a 40 quarter truncation horizon, imposing impact orthogonality between the identified shock an adjusted TFP. This is labeled "BS".

Next, we continue to assume that $\rho_\Gamma = 0.6$, but now set $q = 1$. This means that agents in the economy observe the news shock one period in advance of it impacting technology. In this situation, impact orthogonality between the news shock and technology is satisfied in the model. Figure 4 plots true model and estimated resposne from both our max-share identification and the BS identification. Qualitatively, while the BS identification performs better than when $q = 0$, it is not obviously better than the max-share identification which does not impose impact orthogonality. The max-share identification does a better job at estimating the impact response of hours, and better reproduces the dynamics of inflation to a news shock than does the BS identification.

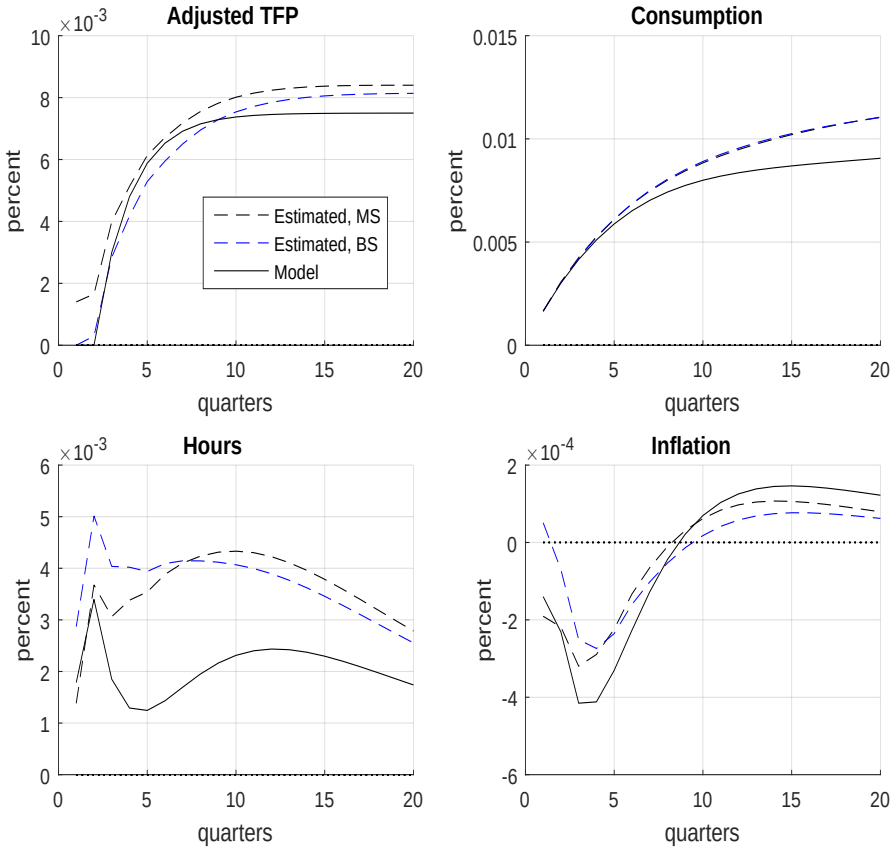
Figure 4: Model and Estimated Responses to a News Shock with a One Period Lag



Notes: Solid lines are the true impulse responses to an estimated growth shock which affects technology with a one period lag in the model. The dashed black lines are the estimated responses using the max share identification at a 80 quarter horizon, without imposing impact orthogonality. This is labeled "MS" for "Max Share." The dashed blue lines are the estimated responses using the BS identification with a 40 quarter truncation horizon, imposing impact orthogonality between the identified shock an adjusted TFP. This is labeled "BS".

Figure 5 shows model and estimated responses from model-generated data when $q = 2$, so that agents observe a news shock two quarters prior to it impacting technology. Qualitatively, the patterns are similar to what is shown in Figure 4. The max-share identification better captures the dynamics of inflation after a news shock, and provides a better estimate of the impact response of hours, than does the BS identification.

Figure 5: Model and Estimated Responses to a News Shock with a Two Period Lag



Notes: Solid lines are the true impulse responses to an estimated growth shock which affects technology with a two period lag in the model. The dashed black lines are the estimated responses using the max share identification at a 80 quarter horizon, without imposing impact orthogonality. This is labeled "MS" for "Max Share." The dashed blue lines are the estimated responses using the BS identification with a 40 quarter truncation horizon, imposing impact orthogonality between the identified shock an adjusted TFP. This is labeled "BS".

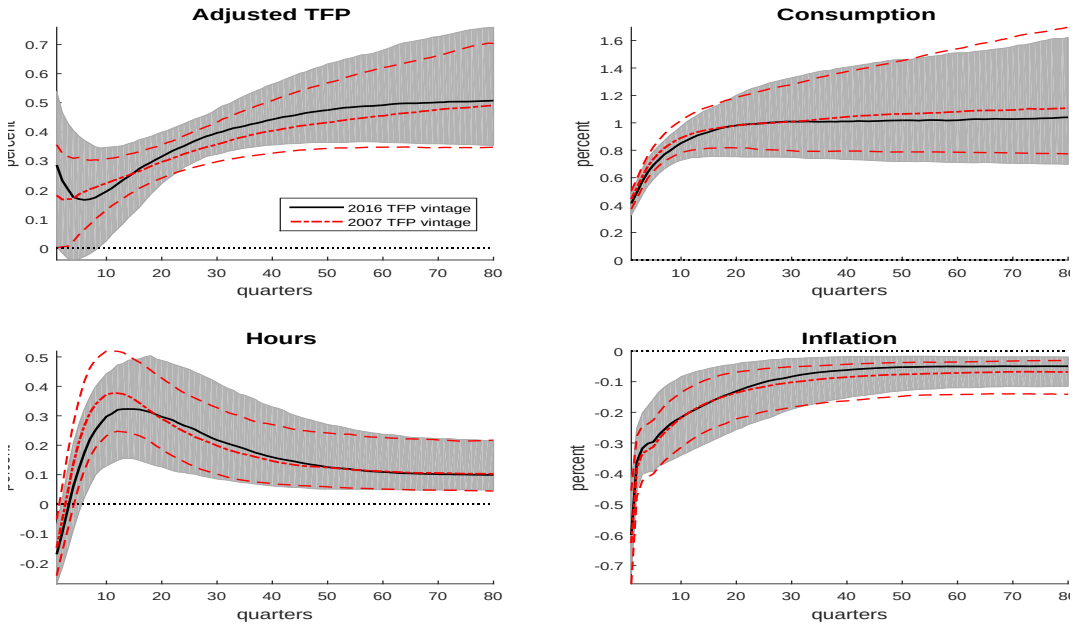
We conclude from these Monte Carlo exercises that our proposed identification performs well at qualitatively recovering the impulse responses to a news shock. It performs significantly better than the BS identification if impact orthogonality between news and technology is not satisfied in the model. Even if this restriction is satisfied in the model, our identification continues to perform well, and many respects is better than the BS identification. These Monte Carlo exercises suggest that there is little to gain from imposing impact orthogonality in the identification of the news shock, while the costs of doing so are potentially if this assumption is not satisfied in the model.

All of our Monte Carlo experiments assume that utilization is observed perfectly when constructing TFP or adjusted TFP in the model. Given the large differences in the Fernald's utilization series by vintage, this assumption may not be innocuous. In ongoing work, we are examining the suitability of our approach to identifying news shocks, compared to Barsky and Sims (2011), when the observed measure of adjusted TFP is mismeasured in the model.

5.3 Empirical Results

Having established the good properties of our proposed empirical identification, we next implement this identification on actual data. The data used in this empirical exercise are the same as in Section 3. The estimated responses for VAR systems estimate with both the 2007 and 2016 vintages of the adjusted TFP series are shown in Figure 6. Solid lines are responses when the 2016 vintage of adjusted TFP is included in the VAR, while dash-dotted red lines are responses when using the 2007 vintage of adjusted TFP.

Figure 6: Empirical Responses, Alternative Identification, 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 shock is identified using our alternative identification, which does not impose impact orthogonality between the news shock and adjusted TFP. The shock is identified as the structural shock which maximizes the variance share of adjusted TFP at a 80 quarter horizon.

There are several notable features evident in Figure 6. First, although our identification does not impose impact orthogonality, the estimated response of adjusted TFP to a news shock is slow and protracted, with the the estimated long run response of adjusted TFP two to three times larger than the impact response. This is true regardless of the vintage of adjusted TFP. Thus, a "news" interpretation of the response to the

identified shock seems natural.

Second, there is very little difference in the responses when estimating the VAR using the 2007 vintage compared to the 2016 vintage of the adjusted TFP series, which is very different than when using the Barsky and Sims (2011) identification. Hours worked declines on impact before turning positive after several quarters. Inflation falls on impact. The impact decline in inflation is statistically significant when estimate the VAR on either vintage of adjusted TFP data, unlike the Barsky and Sims (2011) identification.

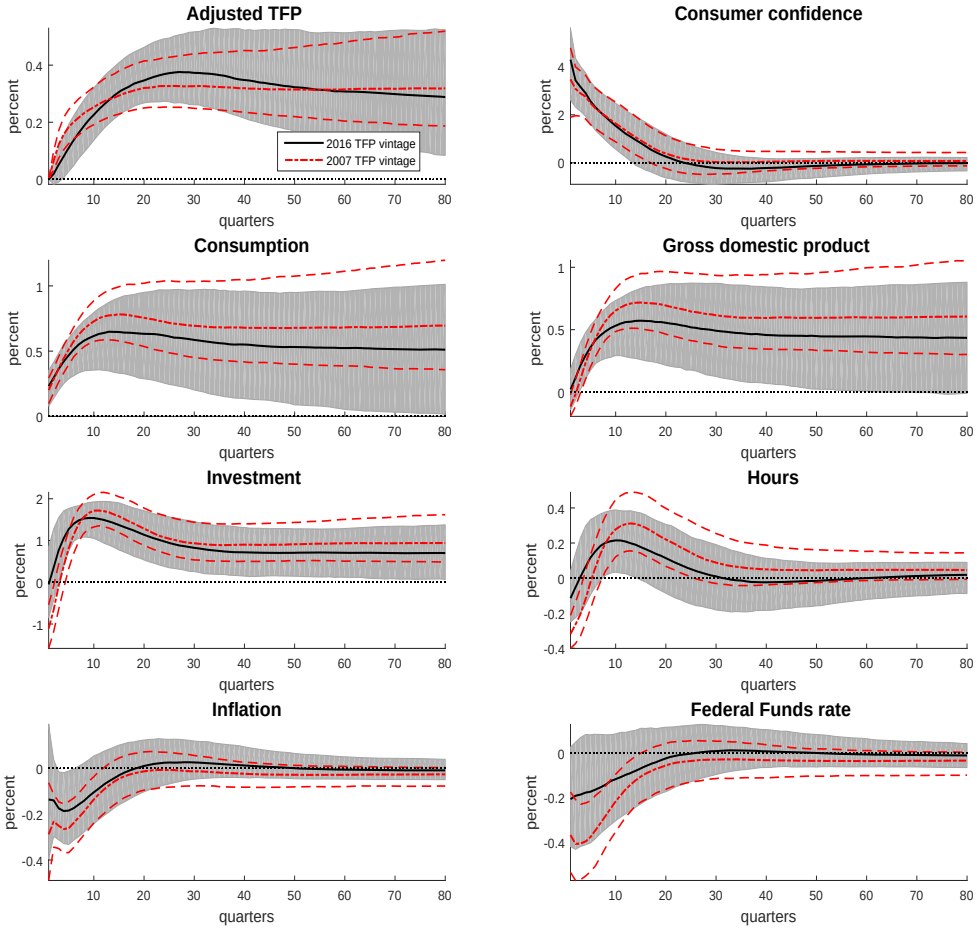
Third, the estimated responses are quite similar to the theoretical response to a news shock (when $q = 0$) in our medium scale model, which can be seen by comparing Figure 6 to the solid lines in Figure 3. This in spite of the fact that the parameterization of the model is taken from the literature, rather than being chosen to force the model produce impulse responses similar to what we observe in the data. In both model and data, adjusted TFP jumps up slightly on impact and continues to grow for several quarters thereafter. Consumption follows a similar path. Hours worked declines on impact before turning positive, and inflation falls significantly. Qualitatively, these responses are similar to those estimated in Barsky and Sims (2011), even with the most recent vintage of the adjusted TFP data.

6 Results in larger VAR systems

A natural extension is to examine how robust these results are to higher dimensional VAR systems. To that end, we estimate a larger VAR system with eight variables. In addition to the four variables in our baseline VAR, we include real GDP, real investment, the Federal Funds Rate, and a survey measure of consumer confidence. This is similar to the larger dimensional system in Barsky and Sims (2011).

To illustrate how the different vintage of adjusted TFP might matter, we first consider the BS identification, choosing the news shock to the be shock orthogonal to adjusted TFP on impact which maximizes the sum of variance shares to adjusted TFP from impact to 40 quarters. Like in the four variable VAR, there are important differences between the estimated response depending on which vintage of adjusted TFP data one uses. With the 2007 vintage of data, output, investment, and hours all decline on impact in response to a favorable news shock. Inflation falls significantly, as does the Federal Funds Rate. With the 2016 vintage of data, output and investment increase on impact, and the impact decline in hours is significantly smaller than with the 2007 vintage of adjusted TFP. The impact declines in inflation and the Federal Funds Rate are smaller with the 2016 vintage of data compared to the 2007 vintage, and cease to be statistically significant.

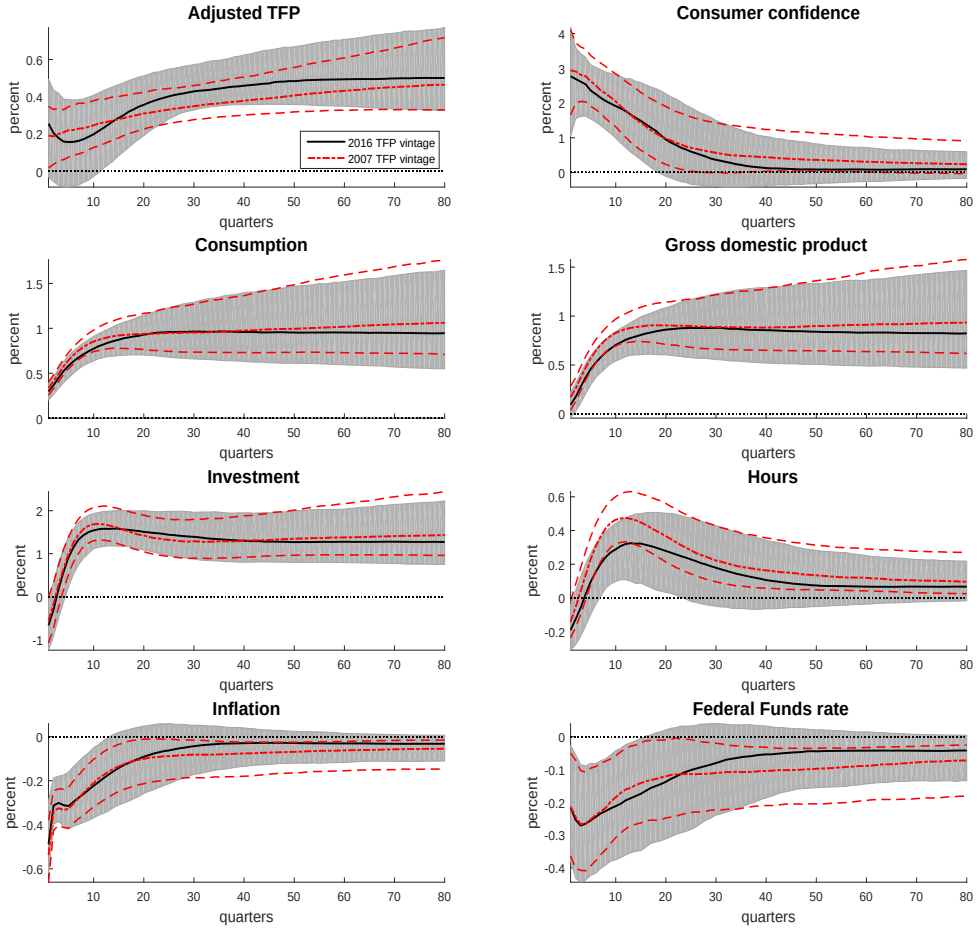
Figure 7: Empirical Responses, BS Identification, Eight Variable VAR, 2007 vs. 2016 Vintage of Adjusted TFP



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 BS identification.

Figure 8 repeats this exercise, but uses our alternative identification, which does not impose impact orthgonality between the news shock and adjusted TFP and which maximizes the variance share due to news at a 80 quarter horizon. As in Figure 6, and unlike what is shown in Figure 7, the estimated responses are much more similar for the two different vintages of the adjusted TFP series. Investment and hours decline on impact, while output rises slightly (albeit insignificantly). Furthermore, inflation and the Funds rate decline significantly on impact with both vintages.

Figure 8: Empirical Responses, Alternative Identification, Eight Variable VAR, 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 BS identification.

Our proposed identification of news shock, which differs primarily from Barsky and Sims (2011) in not imposing impact orthogonality between the news shock and adjusted TFP, delivers consistent, robust results regardless of the vintage of the adjusted TFP series. This is true in both a small and larger dimensional VAR system. We find that hours worked and investment decline on impact, as do inflation and the Federal Funds rate. In contrast, consumption and consumer confidence increase. As such, news shocks are unlikely to be a major driver of business cycle dynamics. Furthermore, our estimated responses are broadly consistent with the theoretical responses to a news shock in a conventionally specified medium scale DSGE model.

7 Conclusion

An almost universally imposed assumption in the empirical literature on news shocks is that a news shock impacts productivity only with a delay. This assumption is not necessarily consistent with a broader view of slow technology adoption, and may be violated empirically if one's measure of productivity does not perfectly align with true technology.

In this paper, we have proposed an alternative identification of news shocks which does not impose that a news shock have no impact effect on productivity. We have shown that our procedure performs well in a Monte Carlo experiment, and is more robust to misspecifications with respect to the data generating process than an identification assumption which imposes impact orthogonality between a news shock and productivity. Applied to the data, we find that the estimated response of a measure of utilization-adjusted TFP to our shock is consistent with a news interpretation – measured TFP increases slowly and with a significant lag. Hours worked declines on impact, while consumption rises. Inflation and the Fed Funds rate significantly decline in response to good news. These results are robust to using different vintages of Fernald (2014)'s popular utilization-adjusted TFP series, which is not the case with the Barsky and Sims (2011) identification approach which imposes impact orthogonality between news and TFP.

References

Aaronson, D. and D. Sullivan (2001). Growth in worker quality. *Economic Perspectives 25*(4), 53–74.

Acemoglu, D. (2009). *Introduction to Modern Economic Growth*. Princeton: Princeton University Press.

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.

Barsky, R. and E. Sims (2011). News shocks and business cycles. *Journal of Monetary Economics 58*(3), 273–289.

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

Basu, S., J. G. Fernald, and M. S. Kimball (2006). Are technology improvements contractionary? *American Economic Review 96*(5), 1418–1448.

Beaudry, P. and F. Portier (2004). An exploration into pigou's theory of cycles. *Journal of Monetary Economics 51*(6), 1183–1216.

Beaudry, P. and F. Portier (2006). Stock prices, news, and economic fluctuations. *American Economic Review 96*(4), 1293–1307.

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

Beuadry, 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.

Blanchard, O. and D. Quah (1989). The dynamic effects of aggregate demand and supply disturbances. *American Economic Review 79*(4), 132–145.

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

Cascaldi-Garcia, D. (2016). News shocks and the slope of the term structure of interest rates: Comment. Technical report, Warwick Business School.

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.

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

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.

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

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

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.

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.

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

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

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

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

Kahn, H. and J. Tsoukalas (2012). The quantitative importance of news shocks in estimated dsge models. *Journal of Money, Credit, and Banking 44*(8), 1535–1561.

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

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.

Kurmann, A. and C. Otrok (2016). New evidence on the relationship between news shocks and the slope of the term structure. Technical report.

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

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

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

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

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

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

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

Schmitt-Grohe, S. and M. Uribe (2012). What's news in business cycles. *Econometrica 80*(6), 2733–2764.

Shapiro, M. and M. Watson (1988). Sources of business cycle fluctuations. In *NBER Macroeconomics Annual 1988, Volume 23*, NBER Chapters, pp. 111–156. National Bureau of Economic Research, Inc.

Sims, E. (2016). Differences in quarterly utilization-adjusted tfp by vintage, with an application to news shocks. Working Paper 22154, National Bureau of Economic Research.

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

Summers, L. H. (1986). Some skeptical observations on real business cycle theory. *Federal Reserve Bank of Minneapolis Quarterly Review Fall*, 23–27.

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.

A A Medium Scale DSGE Model

For the Monte Carlo experiments considered in the paper, we use as a laboratory a conventionally specified medium scale DSGE model. The model features sticky prices and wages, capital accumulation, habit formation in consumption, an investment adjustment cost, variable capital utilization, and a central bank which implements monetary policy according to a Taylor rule. The model is very similar to Christiano, Eichenbaum, and Evans (2005), Smets and Wouters (2007), and Justiniano, Primiceri, and Tambalotti (2010). As such, we only list the full set of equilibrium conditions here, rather than fully laying out the decision problems of each type of agent in the model.

The full set of equilibrium conditions are listed below. A brief discussion of each expression follows.

$$\lambda_t = (C_t - bC_{t-1})^{-1} - \beta b E_t \left(C_{t+1} - bC_t\right)^{-1} \tag{A.1}$$

$$\lambda_t = \beta(1+i_t)E_t\lambda_{t+1}(1+\pi_{t+1})^{-1} \tag{A.2}$$

$$\lambda_t R_t = \xi_t(\delta_1 + \delta_2(u_t - 1)) \tag{A.3}$$

$$\xi_t = \beta E_t \left[\lambda_{t+1} R_{t+1} u_{t+1} + (1 - \delta(u_{t+1}))\xi_{t+1}\right] \tag{A.4}$$

$$\lambda_t = \xi_t Z_t \left[1 - \frac{\kappa}{2}\left(\frac{I_t}{I_{t-1}} - g_I\right) - \kappa\left(\frac{I_t}{I_{t-1}} - g_I\right)\frac{I_t}{I_{t-1}}\right] + \beta E_t \xi_{t+1} Z_{t+1} \kappa \left(\frac{I_{t+1}}{I_t} - g_I\right)\left(\frac{I_{t+1}}{I_t}\right)^2 \tag{A.5}$$

$$w_t^{\#} = \frac{\epsilon_w}{\epsilon_w - 1}\frac{f_{1,t}}{f_{2,t}} \tag{A.6}$$

$$f_{1,t} = \varphi_t \psi \left(\frac{w_t}{w_t^{\#}}\right)^{\epsilon_w(1+\chi)} L_t^{1+\chi} + \beta\theta_w E_t \left(\frac{w_{t+1}^{\#}}{w_t^{\#}}\right)^{\epsilon_w(1+\chi)} \left(\frac{(1+\pi_t)^{\zeta_w}}{1+\pi_{t+1}}\right)^{-\epsilon_w(1+\chi)} f_{1,t+1} \tag{A.7}$$

$$f_{2,t} = \lambda_t \left(\frac{w_t}{w_t^{\#}}\right)^{\epsilon_w} L_t + \beta\theta_w E_t \left(\frac{w_{t+1}^{\#}}{w_t^{\#}}\right)^{\epsilon_w} \left(\frac{(1+\pi_t)^{\zeta_w}}{1+\pi_{t+1}}\right)^{1-\epsilon_w} f_{2,t+1} \tag{A.8}$$

$$R_t = \alpha m c_t A_t \widetilde{K}_t^{\alpha-1} N_t^{1-\alpha} \tag{A.9}$$

$$w_t = (1-\alpha) m c_t A_t \widetilde{K}_t^{\alpha} N_t^{-\alpha} \tag{A.10}$$

$$\frac{1+\pi_t^{\#}}{1+\pi_t} = \frac{\epsilon_p}{\epsilon_p - 1}\frac{x_{1,t}}{x_{2,t}} \tag{A.11}$$

$$x_{1,t} = \lambda_t m c_t Y_t + \beta\theta_p E_t (1+\pi_t)^{-\zeta_p\epsilon_p}(1+\pi_{t+1})^{\epsilon_p} x_{1,t+1} \tag{A.12}$$

$$x_{2,t} = \lambda_t Y_t + \beta\theta_p (1+\pi_t)^{\zeta_p(1-\epsilon_p)} E_t (1+\pi_{t+1})^{\epsilon_p - 1} x_{2,t+1} \tag{A.13}$$

$$Y_t v_t^p = A_t \widetilde{K}_t^{\alpha} L_t^{1-\alpha} - X_t F \tag{A.14}$$

$$v_t^p = (1+\pi_t)^{\epsilon_p}\left((1-\theta_p)(1+\pi_t^{\#})^{-\epsilon_p} + \theta_p (1+\pi_{t-1})^{-\epsilon_p\zeta_p} v_{t-1}^p\right) \tag{A.15}$$

$$K_{t+1} = Z_t \left[1 - \frac{\kappa}{2}\left(\frac{I_t}{I_{t-1}} - g_I\right)^2\right] I_t + (1 - \delta(u_t))K_t \tag{A.16}$$

$$\delta(u_t) = \delta_0 + \delta_1(u_t - 1) + \frac{\delta_2}{2}(u_t - 1)^2 \tag{A.17}$$

$$Y_t = C_t + I_t \tag{A.18}$$

$$\widetilde{K}_t = u_t K_t \tag{A.19}$$

$$(1+\pi_t)^{1-\epsilon_p} = (1-\theta_p)(1+\pi_t^{\#})^{1-\epsilon_p} + \theta_p(1+\pi_{t-1})^{\zeta_p(1-\epsilon_p)} \tag{A.20}$$

$$w_t^{1-\epsilon_w} = (1-\theta_w)w_t^{\#,1-\epsilon_w} + \theta_w\left(\frac{(1+\pi_{t-1})^{\zeta_w}}{1+\pi_t}w_{t-1}\right)^{1-\epsilon_w} \tag{A.21}$$

$$i_t = (1-\rho_i)i + \rho_i i_{t-1} + (1-\rho_i)\left[\phi_\pi(\pi_t - \pi_t^*) + \phi_y(\ln Y_t - \ln Y_{t-1} - \ln g_Y)\right] \tag{A.22}$$

$$\ln Z_t = \rho_Z \ln Z_{t-1} + s_Z \varepsilon_{z,t} \tag{A.23}$$

$$\ln \varphi_t = \rho_\varphi \ln \varphi_{t-1} + s_\varphi \varepsilon_{\varphi,t} \tag{A.24}$$

$$\ln S_t = \rho_S \ln S_{t-1} + s_S \varepsilon_{S,t} \tag{A.25}$$

$$\ln \Gamma_t - \ln \Gamma_{t-1} = (1-\rho_\Gamma)\ln g + \rho_\Gamma(\ln \Gamma_{t-1} - \ln \Gamma_{t-2}) + s_g \varepsilon_{g,t-q} \tag{A.26}$$

$$\mu_t = mc_t^{-1} \tag{A.27}$$

$$A_t = S_t \Gamma_t \tag{A.28}$$

In these equations λ_t is the Lagrange multiplier on the flow budget constraint of a household and ξ_t is the Lagrange multiplier on the capital accumulation equation. C_t denotes consumption, Y_t output, I_t investment, K_t physical capital, and L_t aggregate labor input. w_t is the aggregate real wage, and R_t is the real rental rate on capital services. u_t denotes capital utilization, with $\widetilde{K}_t = u_t K_t$ denoting capital services. π_t is the net inflation rate and i_t is the net nominal interest rate. v_t^p is a measure of price dispersion across firms. $w_t^{\#}$ is the reset real wage for a household given the opportunity to adjust its wage in a given period, while $\pi_t^{\#}$ is the reset inflation rate for firms given the opportunity to adjust their price. $f_{1,t}$ and $f_{2,t}$ are auxiliary variables related to optimal wage-setting, and $x_{1,t}$ and $x_{2,t}$ are auxiliary variables related to price-setting. mc_t is real marginal cost, the inverse of which is the price markup, μ_t. S_t is a stationary technology shock, while Γ_t is a non-stationary technology. A_t is technology, which is the product of these two terms. Z_t is a shock to the marginal efficiency of investment and φ_t is a shock to the disutility from labor. X_t is a trend factor, to be discussed below. g_I is the steady state gross growth rate of investment, and g_Y is the steady state gross growth rate of output.

(A.1) defines λ_t, the shadow value on the flow budget constraint facing a household. The parameter b measures internal habit formation and β is a discount factor. (A.2) is the Euler equation for bonds, which prices the nominal interest rate, i_t. (A.3) is the first order condition for capital utilization. The cost of capital utilization is faster depreciation of capital, which is governed by (A.17). δ_0 is the steady state depreciation rate, while δ_1 and δ_2 are parameters governing the linear and quadratic terms of the utilization adjustment

cost function. The optimality condition for the choice of future capital is given by (A.4), while the first order condition for investment is given by (A.5). An investment adjustment cost is captured by the parameter κ. Each period, there is a $1-\theta_w$ probability that a household can adjust its wage. Optimal wage-setting for updating households is characterized by (A.6)-(A.8). ϵ_w is a parameter governing how much market power a household has in setting its wage, χ is the inverse Frisch labor supply elasticity, and ψ is a fixed parameter governing the disutility from labor. ζ_w is measures how much non-updated wages may be indexed to lagged inflation.

Cost-minimization by firms gives rise to factor demand curves for capital and labor in (A.9)-(A.10). α is the exponent on capital services in the production function, with $1-\alpha$ the exponent on labor. Because firms face the same factor prices, they all have the same real marginal cost and hire capital and labor in the same ratio. Each period, firms face a $1-\theta_p$ probability of being able to adjust their price. Optimal price-setting for updating firms is characterized by (A.11)-(A.13). ϵ_p measures the extent of monopoly power in price-setting. ζ_p is a parameter measuring how much non-updated prices are indexed to lagged inflation. The aggregate production function is given by (A.14). F is a fixed cost of production, scaled by X_t, which measures the economy's trend growth. v_t^p is a measure of price dispersion, the evolution of which is given by (A.15).

Physical capital accumulates according to the law of motion given in (A.16). The aggregate resource constraint is given by (A.18). The evolution of inflation and the aggregate real wage are governed by (A.20) and (A.21), respectively. Monetary policy is governed by a Taylor rule, (A.22), which features interest rate smoothing governed by the parameter of rho_i, a reaction of inflation, ϕ_π to deviations from a long run target, π^*, and a reaction to deviations of output growth from its steady state level, ϕ_y. The exogenous processes for the marginal efficiency of investment shock, the labor supply preference shock, and the stationary technology shock are given by (A.23)-(A.25). Each follows a stationary AR(1) process with steady state levels normalized to unity. The permanent productivity process is a stationary AR(1) in the growth rate and is given by (A.26). g denotes the steady state growth rate. The innovation is dated $t-s$, for $q \geq 0$. $q=0$ means that the technology improvement materializes immediately. $q>0$ means that agents observe the shock before it impacts productivity. The price markup is defined as the inverse of real marginal cost in (A.27), and composite technology is the composite of the stationary and non-stationary terms, given in (A.28).

Many of the variables in the model inherit the stochastic trend from Γ_t. It is straightforward to show that the stochastic trend factor is $X_t = \Gamma_t^{\frac{1}{1-\alpha}}$. Re-scaling trending variables by this factor renders the model stationary, and permits solution of the model using standard techniques. We solve the model via linearization about the non-stochastic steady state in the re-scaled variables.

The parameters of the model are set to values listed in Table 1 below:

Table 1: Calibrated Parameters

Parameter	Value	Description
β	0.99	Discount factor
α	1/3	Capital's share
δ_0	0.025	Steady state depreciation
δ_1	$u^* = 1$	Utilization linear term
δ_2	0.025	Utilization squared term
ψ	$L^* = 1/3$	Labor disutility
π^*	0	Steady state inflation
ϕ_π	1.5	TR inflation
ϕ_y	0.3	TR output growth
ρ_i	0.85	TR smoothing
ϵ_p	11	Elasticity sub goods
ϵ_w	11	Elasticity sub labor
χ	1	Inverse Frisch
F	$\Pi^* = 0$	Fixed cost
θ_p	0.75	Calvo prices
θ_w	0.75	Calvo wages
ζ_p	0.00	Price indexation
ζ_w	0.00	Wage indexation
b	0.8	Habit formation
κ	4	Investment adjustment cost
g	1.0025	SS growth of productivity
ρ_Γ	0.6	AR productivity growth
ρ_A	0.85	AR surprise productivity
ρ_Z	0.80	AR investment
ρ_φ	0.90	AR labor supply
s_g	0.003	SD growth shock
s_A	0.005	SD surprise productivity
s_Z	0.03	SD investment shock
s_φ	0.05	SD labor supply shock

These parameter values are largely drawn from the literature and are therefore quite standard. The frequency of time is a quarter. The discount factor is set to 0.99 and the parameter α in the production function is 1/3. The steady state depreciation rate is set to 0.025. The linear parameter in the utilization adjustment cost function is set to be consistent with a normalization of steady state utilization to unity. The coefficient on the squared term in the utilization cost term is set to $\delta_2 = 0.025$. The disutility of labor parameter is chosen to be consistent with steady state hours worked of 1/3. The parameters ϵ_p and ϵ_w are chosen to be consistent with steady state price and wage markups of 10 percent. The Calvo price and wage parameters are both set to 0.75, implying that the average duration between price or wage changes is four quarters. We assume no backward indexation of prices or wages to lagged inflation. The habit formation parameter is set to $b = 0.8$, and the investment adjustment cost parameter is set to $\kappa = 4$, both standard

values in the literature. We assume that there is no trend inflation, so $\pi^* = 0$. The Taylor rule features significant interest smoothing and strong responses to both inflation and output growth. The steady state gross growth rate of productivity is set to 1.0025, which implies output growth of 1.5 percent per year in steady state. The fixed cost of production, F, is chosen so that profits are zero in steady state. We have also experimented with having no fixed cost in the model.

We have less to guide us in terms of the parameterization of the shock processes. Our model features four stochastic shocks, which coincides with the number of variables included in our estimated VARs. Our parameterization of the marginal efficiency of investment shock is based on Justiniano et al. (2010). Our parameterizations of the shock processes for the stationary technology shock and the labor supply preference shock are fairly conventional. As a benchmark, we assume that $s = 0$, so that the permanent productivity shock impacts the level of productivity in the period in which agents observe the shock.

When we solve the model with these parameter values, the model generates plausible moments. The standard deviation of output growth comes out to 0.0078, which is close to its value in the data over the entire post-war sample. Consumption is about half as volatile as output and investment is three times more volatile than output. In terms of an unconditional variance decomposition, the investment shock accounts for about 65 percent of the unconditional variance of output growth. The next most important shock is the productivity growth shock, which accounts for about 25 percent of the unconditional variance of output growth. The labor supply shock explains 7 percent of the variance of output growth and the stationary productivity shock accounts for about 3 percent of the variance of output growth. With our parameterization, the stationary technology shock accounts for roughly two-thirds of the variance of the growth rate of technology, whereas the permanent shock accounts for about one-third of the variance of technology growth.