

Differences in Quarterly Utilization-Adjusted TFP by Vintage, with an Application to News Shocks*

Eric Sims[†]

University of Notre Dame & NBER

First version: March 23, 2016

This version: May 24, 2016

Abstract

This paper documents large differences across vintages in the properties of the widely-used quarterly utilization-adjusted TFP series produced by [Fernald \(2014\)](#), who provides updated data each quarter on his website. The most recent vintage of the adjusted TFP series has correlations with earlier vintages of the series that are less than 0.6. Compared to earlier vintages, the most recent vintage of the adjusted TFP data is more weakly correlated with output and more strongly negatively correlated with hours worked. The most recent vintage of the adjusted TFP series is less predictable from other macroeconomic shocks than are earlier vintages. I revisit the empirical analysis from [Barsky and Sims \(2011\)](#), who use an earlier vintage of Fernald's adjusted TFP data to identify impulse responses to news shocks about future productivity in a structural VAR. The vintage of adjusted TFP data matters for their estimated impulse responses, and in some specifications the differences using the most recent vintage of the adjusted TFP data are qualitatively large in a way that is more favorable to theories of news-driven business cycles.

JEL Classification: E22, E23, E32, O47

Keywords: Total factor productivity, variable utilization, news shocks

*I am grateful to John Fernald for helpful conversations and for making some of the past vintages of data available to me. I also wish to thank Rudi Bachmann and Deokwoo Nam for helpful discussions concerning this project. Carlos Rondon Moreno provided capable research assistance.

[†]E-mail address: esims1@nd.edu.

1 Introduction

Identifying the shocks which drive business cycles is one of the central tasks of macroeconomic research. Ever since the contributions of [Kydland and Prescott \(1982\)](#) and [Long and Plosser \(1983\)](#), economists have studied and debated the role of productivity shocks in driving short run fluctuations. One of the principal difficulties in understanding the role of productivity shocks concerns the measurement of productivity in the data. While early approaches typically used growth accounting techniques to measure total factor productivity (TFP) as the residual component of aggregate output not explained by capital and labor, economists quickly realized that measured TFP may be a poor measure of aggregate productivity due to unobserved utilization of factors of production.¹

Partly in response to criticisms of traditional TFP series as a measure of productivity, [Basu, Fernald, and Kimball \(2006\)](#) use restrictions derived from economic theory and disaggregated industry-level data to construct an aggregate measure of TFP which is adjusted for unobserved factor utilization, as well as imperfect competition and sectoral heterogeneity, at an annual frequency. [Fernald \(2014\)](#) applies the basic analysis in [Basu et al. \(2006\)](#) to produce a measure of TFP which controls for unobserved utilization at a quarterly, rather than annual, frequency. He makes these data available for download on his personal website (<http://www.frbsf.org/economic-research/economists/john-fernalld/>). These data are updated each quarter. 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.² The quarterly, utilization-adjusted TFP series provided by Fernald has proven highly influential. As of mid-May of 2016, the working paper describing the construction of the adjusted TFP series has been cited 220 times on Google Scholar.

The purpose of this paper is to document that there are large and qualitatively important differences in the time series properties among different vintages of the adjusted TFP series provided on Fernald’s website. It is important that practitioners be aware of these differences and how they might affect substantive conclusions in applied work making use of these data. In the course of my own research, I have made use of three different vintages of Fernald’s

¹The 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.

²In particular, the [Basu et al. \(2006\)](#) methodology allows for non-constant returns to scale, markups of price over marginal cost, and observed factor shares not equaling output elasticities. These corrections are only available at an annual frequency and are therefore not part of Fernald’s quarterly series, which only adjusts TFP for variable utilization.

utilization-adjusted TFP series. These different vintages of data were created on December 1, 2007, September 16, 2011, and May 7, 2015. While none of the three vintages of adjusted TFP data are identical, the differences are most striking comparing the 2015 vintage of data to the two earlier vintages. In the text of [Fernald \(2014\)](#) and the accompanying notes on the Excel spreadsheet available for download, Fernald is completely transparent in noting how the construction of the series has varied over time. He does not, however, make a point to note that the time series properties of the adjusted TFP series have changed in important ways which might affect conclusions from empirical analyses using these data. My paper is intended to fill this gap.

Section 2 begins by presenting some simple descriptive statistics on the growth rates of the adjusted TFP series by vintage.³ While all three vintages have fairly similar means and standard deviations, the correlations of the 2015 vintage with either the 2007 or 2011 vintages are quite low. In particular, while the correlation between the 2007 and 2011 vintage series is 0.95, the correlations between the 2007 or 2011 vintages with the 2015 vintage are both smaller than 0.6. These correlations are not driven by behavior during particular subsamples. There are economically important differences in correlations of the different vintages of the adjusted TFP series with key macroeconomic variables. While the 2007 and 2011 vintages of adjusted TFP data are positively correlated with aggregate output and mildly negatively correlated with total hours worked, the 2015 vintage of adjusted TFP data is uncorrelated with output and strongly negatively correlated with total hours worked.

Changes in the construction of the quarterly utilization series used to adjust the conventional TFP series seem to be the principal driving force behind the differences in the time series properties of the utilization adjusted TFP series by vintage. The correlations of the utilization series used in the 2015 vintage with the utilization series used in the 2007 or 2011 vintages are only about 0.65. Constructing a synthetic 2015 vintage of adjusted TFP using earlier vintages of the utilization series results in an adjusted TFP series that is much more strongly correlated with either the 2007 or 2011 vintages of the data. [Fernald \(2014\)](#) notes that two changes were made to the utilization series in 2014. First, he switches to using updated utilization estimates from [Basu, Fernald, Fisher, and Kimball \(2013\)](#), which uses the same methodology as [Basu et al. \(2006\)](#) but on more recent data. Second, he uses new industry-level data to compute the aggregate utilization series. My own email exchanges with Fernald indicate that the switch to the [Basu et al. \(2013\)](#) methodology likely explains the disparate properties among newer and older vintages of the data – a vintage of the adjusted TFP series from immediately prior to the switch in measurement of utilization (December of

³All statistics are based on a fixed sample size of 1947q3-2007q3, the end data of which is constrained by data availability in the 2007 vintage of the data.

2013) looks very much like the 2007 and 2011 vintages, while a vintage from right after the switch (May 2014) has properties similar to the 2015 vintage of the data.⁴

Given the large differences in the properties of the adjusted TFP series by vintage, it is natural to ask which vintage of the data is “best.” Since most applied work making use of Fernald’s data takes the adjusted TFP series as a measure of exogenous productivity, I investigate whether the adjusted TFP series (by vintage) is predictable from macroeconomic shocks which are putatively unrelated to productivity. In particular, I estimate pairwise Granger causality tests making use of the defense news shock series from [Ramey \(2011\)](#), the tax shock series from [Romer and Romer \(2010\)](#), the monetary policy shock series produced in [Romer and Romer \(2004\)](#), and crude oil prices. One can reject the hypotheses that the defense news shock and monetary policy shock series do not Granger cause the 2007 and 2011 vintages of the adjusted TFP series at the five percent level or better. One cannot reject the hypotheses that any of the proposed shocks do not Granger cause the 2015 vintage of the adjusted TFP data at the five percent level. These results are indicative that the most recent vintage of the adjusted TFP series may represent an improvement over earlier vintages.

Having documented that different vintages of the utilization-adjusted TFP data might produce very different unconditional correlations with aggregate variables of interest, I then investigate whether the different vintages of adjusted TFP data might matter for conditional moments. Fernald’s utilization-adjusted TFP series has been used extensively in the empirical literature on news shocks about future productivity. In Section 3, I re-do the empirical analysis from [Barsky and Sims \(2011\)](#) using different vintages of adjusted TFP data. These authors estimate vector autoregression (VAR) models using the 2007 vintage of the adjusted TFP series and other aggregate variables. They identify a news shock as the structural shock uncorrelated with the reduced form innovation in adjusted TFP which maximally accounts for the variance share of adjusted TFP over a ten year forecast horizon.

Using the 2007 or 2011 vintages of the adjusted TFP data, I identify impulse responses to a news shock that are very similar to those reported in [Barsky and Sims \(2011\)](#). In particular, output and hours decline immediately in response to good news about future productivity. Furthermore, the news shock accounts for a small fraction of the forecast error variance of output at horizons before adjusted TFP begins to change. The results are potentially different in an economically important way when using the 2015 vintage of adjusted TFP

⁴I am aware of one other paper which points out that Fernald’s utilization-adjusted TFP series differs by vintage and that the principal reason for these differences lies in the measurement of factor utilization. In particular, in a comment on [Kurmann and Otrok \(2013\)](#), [Cascardi-Garcia \(2016\)](#) notes that the adjusted TFP series differs by vintage and that this difference seems to be driven by changes in the measurement of utilization. This paper does not extensively document differences in the utilization-adjusted TFP series by vintage, and is written purely as a comment on [Kurmann and Otrok \(2013\)](#). At the original time of writing my paper, I was not aware of [Cascardi-Garcia \(2016\)](#).

data. In one of the VAR specifications in Barsky and Sims (2011), a positive news shock results in a large and immediate output expansion and accounts for a healthy share of the variance of output when the most recent vintage of adjusted TFP is used in the VAR. This result is more favorable to the news-driven business cycle hypothesis than what obtains using older vintages of the adjusted TFP data.

2 Data and Unconditional Moments

This section begins by briefly describing the construction of the utilization-adjusted total factor productivity series produced in Fernald (2014). I then proceed to present some summary statistics on the adjusted TFP series which vary by vintage of when the data were created, and show that different vintages of data display substantively different co-movements with a number of other aggregate series. I then provide some investigation into what accounts for the differences in the adjusted TFP series by vintage.

The theory underlying the construction of the adjusted TFP series is as follows. Suppose that there exists an aggregate production function of the form:

$$Y_t = A_t (z_t K_t)^{\alpha_t} (e_t L_t)^{1-\alpha_t} \quad (1)$$

Here Y_t is output, K_t is physical capital, and L_t is labor input. The variables z_t and e_t represent capital utilization and labor effort, respectively. The parameter α_t is a potentially time-varying capital share parameter. I have imposed constant returns to scale in the production function, so that the labor share is $1 - \alpha_t$. The variable A_t is an exogenous productivity variable.

A traditional measure of TFP is defined as log output less share-weighted log observable inputs (capital and labor). It is assumed that α_t can be measured from aggregate data on labor's share of income. Taking logs of (1) and first differencing:

$$\begin{aligned} \Delta \ln TFP_t &= \Delta \ln Y_t - \alpha_t \Delta \ln K_t - (1 - \alpha_t) \Delta \ln L_t \\ &= \Delta \ln A_t + \alpha_t \Delta \ln z_t + (1 - \alpha_t) \Delta \ln e_t \end{aligned} \quad (2)$$

From (2), it is straightforward to see that the traditional TFP series corresponds to the exogenous productivity variable in the model only if capital utilization and labor effort are constant.

Define u_t as total utilization, the share-weighted sum of capital utilization and labor effort:

$$\Delta \ln u_t = \alpha_t \Delta \ln z_t + (1 - \alpha_t) \Delta \ln e_t \quad (3)$$

If one could measure $\Delta \ln u_t$, one could define a new series, what I will label $\Delta \ln \widehat{A}_t$ and refer to as “adjusted TFP,” by subtracting the growth rate of the total utilization series from the growth rate of the traditional TFP series. To the extent to which the aggregate production function is properly specified, this series ought to measure the the growth rate of the exogenous productivity variable in (1):

$$\Delta \ln \widehat{A}_t = \Delta \ln TFP_t - \Delta \ln u_t \quad (4)$$

Based on the underlying assumptions about the aggregate production function, [Fernald \(2014\)](#) creates both a traditional TFP series and a utilization adjusted TFP series. The construction of the traditional TFP series is based on the following aggregate data. Output is measured as real output in the business sector. In earlier versions of the data, real output is measured from the expenditure side of the NIPA accounts. In more recent vintages of the data, he uses an average of the expenditure and income measures of business output. Observed labor input is measured as the product of hours worked and labor quality. In particular, let $L_t = H_t Q_t$, where H_t is hours worked and Q_t is labor quality. H_t is measured from the BLS series on hours worked in the business sector. The labor quality series is based on measurements in [Aaronson and Sullivan \(2001\)](#). The aggregate capital stock is based on weighted growth in fifteen disaggregated types of capital, where NIPA investment data are used to create perpetual inventory series. The share parameter α_t is constructed using interpolated annual NIPA data on payments to labor. These series are sufficient to construct a quarterly measure of $\Delta \ln TFP_t$.

To create a quarterly measure of adjusted TFP, $\Delta \ln \widehat{A}_t$, it is necessary to have a data series on $\Delta \ln u_t$. [Fernald \(2014\)](#) follows [Basu et al. \(2006\)](#) to construct an aggregate quarterly utilization series from disaggregated industry-level data. The basic idea behind this utilization series is that a cost-minimizing firm ought to want to vary inputs along all margins simultaneously. [Basu et al. \(2006\)](#) estimate a series of industry-level regressions where observed hours per worker is used as a proxy for unobserved input variation. The resulting industry utilization series are then aggregated to a quarterly frequency using average industry weights. Since March of 2014, the utilization series is constructed from updated utilization estimates from [Basu et al. \(2013\)](#), which extends the [Basu et al. \(2006\)](#) methodology to more recent data. Since August 2014, the utilization series have been based on updated hours per worker data at the industry level, which potentially affects the resulting utilization series.

I have data on adjusted TFP which were created on December 1, 2007, September 16,

2011, and May 7, 2015. I label the adjusted TFP series by vintage as $\Delta \ln \widehat{A}_t^{07}$, $\Delta \ln \widehat{A}_t^{11}$, and $\Delta \ln \widehat{A}_t^{15}$. Table 1 presents some summary statistics for these series. For these statistics, the sample period runs from 1947q3 through 2007q3, the last period available in the 2007 data. Data subsequent to 2007 are dropped from the 2011 and 2015 vintages for the purposes of comparison. The statistics are based on annualized percentage values. The means of each vintage of data are all close to 1.5 (i.e. the average annualized growth rate is about 1.5 percent). The mean value is highest for the 2007 series and lowest for the 2011 series. The standard deviations of the different vintages of adjusted TFP growth are in the neighborhood of 3.5, with the 2011 vintage data slightly more volatile than either the 2007 or 2015 vintages. The differences in means and volatilities are fairly small, and could be largely explained by data revisions near the end of the sample.

The Table also presents simple correlation coefficients for the different vintages of data on adjusted TFP growth. The correlation between the 2007 and 2011 vintages is very high at 0.95. The correlations between the 2007 and 2011 vintages with the 2015 vintage adjusted TFP growth are much smaller. In particular, the correlation between the 2007 and 2015 vintage data is 0.58, while the correlation between the the 2011 and 2015 vintages is 0.59. In other words, there are substantial differences in the time series properties of the early and late vintages of the adjusted TFP growth series.

The differences between the adjusted TFP series by vintage are not driven by particular parts of the sample. Table 2 presents correlations among different vintages of the adjusted TFP series over different samples. I consider the sample periods 1960q1-2007q3, 1960q1-1979q4, and 1984q1-2007q3. The first sample period is meant to exclude some large movements occurring in the immediate wake of World War II and some events surrounding the Korean War. The final sample, 1984q1-2007q3, focuses on the so-called “Great Moderation” period when the volatility of output and other aggregate variables fell. The 1960q1-1979q4 excludes data from the immediate post-war period as well as the Great Moderation period. The correlations among different vintages of the adjusted TFP data are very similar across different samples. In particular, the correlations between the 2007 and 2011 vintages with the 2015 vintages are always about 0.6 (or smaller), while the correlation between the 2007 and 2011 vintages are always 0.9 or higher.

2.1 Business Cycle Moments

Given the apparent large differences between the early (2007 and 2011) and late (2015) vintages of the adjusted TFP growth series, it is natural to ask how the different vintages of data co-move with different aggregate variables of interest. Tables 3 and 4 present correlations

of the different vintages of adjusted TFP data with several aggregate variables, the former for log first differences, and the latter for HP filtered (smoothing parameter 1600) log levels of the data. The aggregate variables considered include real GDP, consumption, and investment; inflation (defined as the log first difference of the GDP deflator); hours (as measured from both the non-farm business sector and the business sector); and real output (as measured both in the non-farm business sector and the business sector).⁵ The aggregate data series were downloaded in March of 2016. The sample period for all correlations is fixed across vintages of TFP data from 1947q3 to 2007q3.

I focus first on correlations based on log first differences of data, as presented in Table 3. The 2007 and 2011 vintages of adjusted TFP growth are strongly positively correlated with real GDP growth, with correlations of 0.53 and 0.44, respectively. The 2015 vintage adjusted TFP growth rate is close to uncorrelated with GDP at 0.07. One sees very similar patterns when output is measured using real output in the non-farm business or business sectors – adjusted TFP is most positively correlated with output using the 2007 vintage adjusted TFP data, somewhat less positively correlated using the 2011 vintage data, and is close to uncorrelated with output in the 2015 vintage data. The correlation of adjusted TFP growth with consumption growth is small and positive and does not vary much by vintage of TFP data. The correlation between adjusted TFP growth and the growth rate of aggregate investment is positive across vintages of TFP data, but is substantially weaker with the 2015 vintage of adjusted TFP data (0.04) than with the 2007 (0.27) or 2011 vintages (0.19). Growth rates in adjusted TFP are mildly negatively correlated with inflation, with similar correlations across vintages. As with data on output, there are large differences in the correlations of the TFP growth rates by vintage with measures of aggregate hours growth. The 2007 vintage of adjusted TFP data is mildly negatively correlated with hours growth, as measured either using the non-farm business sector or the business sector. These correlations are slightly more negative using the 2011 vintage of adjusted TFP data, but are qualitatively similar to the correlations using the 2007 vintage data. The correlation of adjusted TFP growth with hours growth using the 2015 vintage of adjusted TFP data is substantially more negative (at about -0.4).

I next consider correlations between HP filtered log levels of the different vintages of adjusted TFP data with HP filtered log levels of different aggregate series. One observes much the same pattern as when focusing on correlations in growth rates. Adjusted TFP

⁵The real GDP series is the headline NIPA number. Consumption is defined as real consumption of non-durables and services. Investment is defined as the sum of durable consumption expenditures and private fixed investment. Both of these series are constructed by using nominal share-weighted real growth rates using individual price indexes. The inflation series is the log first difference of the GDP deflator. Hours and real output in both the non-farm business and business sectors are directly available from the BLS.

goes from mildly positively correlated with different measures of output in the 2007 vintage data (correlation close to 0.15), to weakly negatively correlated with output measures in the 2011 data (correlation of -0.05), to more strongly negatively correlated with different output measures in the 2015 vintage of adjusted TFP data (correlation of -0.2). Similar patterns are evident for the correlations with consumption and investment, where the correlations of adjusted TFP with these variables decrease going from early to later vintages of adjusted TFP data. There is little difference in the correlations of HP filtered adjusted TFP with inflation across vintages. As with output, there are large declines in the correlations between adjusted TFP and HP filtered hours across vintages. In the 2007 vintage data, adjusted TFP is weakly negatively correlated with either measure of hours (correlation of -0.1). In the 2015 vintage data, the correlation is much more negative (-0.4).

In summary, the large differences evident in the adjusted TFP series by vintage translate into substantive differences in correlations of the adjusted TFP series with different aggregate variables. While the correlations for the 2007 and 2011 vintages of the adjusted TFP series are qualitatively similar, there are important differences when using the 2015 vintage of the data. Adjusted TFP goes from procyclical to acyclical or countercyclical, and there is a large and quantitatively important decline in the correlations of the adjusted TFP series with hours worked when moving from 2007 or 2011 vintage data to the 2015 vintage data.

2.2 What Drives the Differences by Vintage?

Given the large differences in the 2015 vintage of adjusted TFP data from earlier vintages, a next natural question to ask is what drives these differences. [Fernald \(2014\)](#) makes available on his website not only the adjusted TFP series itself, but also the series for the variables that go into the construction of this series. These variables are the aggregate capital stock, K_t ; aggregate labor hours, H_t ; aggregate labor quality, Q_t ; and aggregate utilization, u_t .⁶

Table 5 presents descriptive statistics on the log first differences of the component series that go into the construction of the adjusted TFP series. These statistics are based on correlations of the data for the components series provided on Fernald’s website, and are computed over a fixed sample period of 1947q3-2007q3. The output series used in the construction of the 2007 and 2011 vintages of data are virtually identical, with a correlation of 0.996. The 2015 vintage of output data is less correlated with the 2007 or 2011 data, but the correlations are nevertheless still quite high at 0.96. These differences are likely

⁶On his website Fernald does not distinguish between capital utilization and labor effort, instead only reporting the total utilization series, u_t , defined in (3). Furthermore, there are some small differences in the variable α_t across vintages, but these are not important for differences in the final adjusted TFP series by vintage.

driven by the movement to measuring output as the average of expenditure and income based output, as opposed to only using expenditure-based measures in the 2007 and 2011 vintages of data. The correlations of the growth rates of the capital stock by vintage are all above 0.90. The 2011 and 2015 vintages of capital stock data appear almost identical. The correlations between the 2007 vintage of capital stock data and the 2011 and 2015 vintages are 0.94 and 0.91, respectively. There is virtually no difference in the hours series by vintage; the correlations between the growth rates of hours are all above 0.999 across the different vintages.

In summary, there are some small differences in the output and capital stock series by vintage, but these correlations are still fairly close to one, and it therefore seems unlikely that these differences could account for the large differences in the final adjusted TFP series by vintage. I next turn to a discussion of the labor quality series, Q_t , and the total utilization series, u_t . The labor quality series in the 2011 and 2015 vintages of data are virtually identical, with a correlation of 0.999. There is a large difference in the labor quality series used in the 2007 vintage compared to the other two vintages. Not only is the 2007 quality series about half as volatile as the 2011 and 2015 vintages (standard deviation of 0.4 compared to standard deviations of about 0.7 for the two later vintages), but the 2007 labor quality series is weakly correlated with either of the two later vintages of data (these correlations are both about 0.4). In spite of the large differences between the 2007 and 2015 labor quality series, it seems unlikely that different measures of the labor quality series are the primary driving force behind the differences in the final adjusted TFP series. This is because the correlation between growth rates of adjusted TFP in the 2011 and 2015 vintages of data is less than 0.6, in spite of the fact that the labor quality series are virtually identical.

The 2007 and 2011 vintages of the total utilization series are quite similar. In spite of the fact the 2011 vintage utilization series is a good bit more volatile than the 2007 series, the correlation between the utilization series across these two vintages is 0.96. In contrast, there are large differences between the 2015 vintage utilization data and either the 2007 or 2011 vintages. The correlation between the 2007 and 2015 vintage utilization series is 0.66, and the correlation between the 2011 and 2015 vintages is 0.64.

To investigate the possibility that differences in the utilization and/or labor quality series are driving the differences in the adjusted TFP series by vintage, I compute synthetic measures of adjusted TFP in 2011 and 2015 using 2007 vintage utilization and/or labor quality data. That is, I re-construct the 2011 or 2015 adjusted TFP series using component series from that same vintage, with the exception of utilization and/or labor quality, for which I use the 2007 vintage series. I then look at correlations between the synthetic 2011 or 2015 series with the actual 2007 vintage data. Table 6 presents these correlations.

The correlation between the synthetic 2011 adjusted TFP series using 2007 vintage utilization data and the actual 2007 vintage data is 0.98, slightly higher than the correlations between the actual series (0.95). The correlation between the synthetic 2015 adjusted TFP series using 2007 vintage utilization data and the actual 2007 vintage of data is 0.91. This is substantially higher than the correlation between the actual vintages of data, which is only 0.59. As the middle panel of the Table suggests, using the 2007 vintage labor quality series makes almost no difference for the correlations between the synthetic adjusted TFP series and the actual 2007 vintage of data. Finally, the last panel of the table constructs synthetic 2011 and 2015 vintages of the data using both the 2007 vintage utilization series as well as the 2007 vintage labor quality series. The correlations of these synthetic series with the actual 2007 vintage adjusted TFP data are slightly higher than when the synthetic series are constructed using just the 2007 vintage utilization series, but these differences are small.

The above analysis suggests that the principal factor driving the differences between the 2015 vintage of adjusted TFP data and the 2007 and 2011 vintages are differences in the utilization series. While there are large differences in the labor quality series across vintages, these differences do not seem very important for the final adjusted TFP series. These findings align with notes available on the data for download that two changes were made to the utilization series in 2014 – an update to using the [Basu et al. \(2013\)](#) methodology for producing the utilization series in March of 2014, and an update in August of 2014 making use of new data on hours per worker in the industry-level regressions underlying the final utilization series.

An email exchange with Fernald confirms that the switch to the [Basu et al. \(2013\)](#) methodology for computing the utilization series is the principal driving force for changes in the adjusted TFP series by vintage. He provided me with vintages of the data both immediately prior (December 2013) and immediately after (May 2014) the switch to using [Basu et al. \(2013\)](#), which occurred in March of 2014. I label these vintages of adjusted TFP and utilization as $\Delta \ln \widehat{A}_t^{13}$, $\Delta \ln \widehat{A}_t^{14}$, $\Delta \ln u_t^{13}$, and $\Delta \ln u_t^{14}$. Table 7 shows correlations of these two vintages of the data with the other vintages which I previously had in my possession. The December 2013 vintage of adjusted TFP data is strongly positively correlated with the 2007 and 2011 vintages of the data. It is weakly correlated with the vintage from five months later, with a correlation of only 0.63. One observes a similar pattern for the utilization data. The December 2013 vintage of the utilization data is very strongly correlated with the 2007 and 2011 vintages, with correlation coefficients of 0.94 and 0.97, respectively. The correlation drops dramatically after the switch the [Basu et al. \(2013\)](#) methodology, with the correlation between the December 2013 vintage utilization series and the May 2014 vintage only 0.59.

It is interesting to examine how Fernald’s reported utilization series compares to another

popular measure of utilization, the Federal Reserve Board’s capacity utilization series. This series is available at a monthly frequency beginning in 1967. I convert it to a quarterly frequency by averaging observations within quarter, and then compute correlations of this series with Fernald’s reported utilization series by vintages over the period 1967q1 through 2007q3. The correlation between the 2007 vintage utilization data and the FRB’s series is 0.65. The correlation between the 2011 vintage utilization series and the FRB capacity utilization series is virtually identical. The 2015 vintage utilization data is much more weakly correlated with the FRB capacity utilization series, however. In particular, the correlation between the 2015 vintage utilization series and the FRB capacity utilization series is 0.48.

2.3 Which Vintage is Best?

Given the large and important differences in the time series properties of the different vintages of the adjusted TFP series, a natural question is which vintage is the “best”? “Best” is put in quotation marks because it is not clear which criterion one ought to use to determine whether a vintage of the data has desirable features. In his email communications with me, Fernald does not take a stand on which vintage of his data is the best, though he notes that the more recent vintages based on the [Basu et al. \(2013\)](#) methodology might be more desirable because they are based on more recent data.

Applied researchers have typically used Fernald’s adjusted TFP series as a measure of exogenous productivity shocks. A natural test of whether the series can in fact be interpreted as exogenous productivity is to examine whether it is predictable from variables which are putatively unrelated to exogenous productivity. To implement such a test, I run a series of pairwise Granger causality tests with different vintages of the adjusted TFP series and measures of exogenous macroeconomic shocks which have been identified through various means. These tests amount to estimating regressions of the form:

$$\Delta \ln \widehat{A}_t^s = \alpha + \sum_{j=1}^p \rho_j \Delta \ln \widehat{A}_{t-j}^s + \sum_{j=1}^p \gamma_j X_{t-j} + \epsilon_t, \quad s = \{07, 11, 15\} \quad (5)$$

In (5), the growth rate of adjusted TFP (where s indexes the vintage) is projected on p of its own lags and p lags of another variable, which I denote by X_t . The Granger causality test is a test of the hypothesis that the γ_j coefficients are jointly equal to zero. For variables on which to project the adjusted TFP series, I consider macroeconomic shocks identified in the empirical literature. In particular, I use the defense news shock series from [Ramey \(2011\)](#), the exogenous tax shocks from [Romer and Romer \(2010\)](#), and the exogenous monetary policy shocks from [Romer and Romer \(2004\)](#). These series were downloaded from the websites of the authors. The [Ramey \(2011\)](#) and [Romer and Romer \(2010\)](#) series go back to 1947q3, the

first observation for which I have adjusted TFP data. The [Romer and Romer \(2004\)](#) series only spans the period 1969q1-1996q4.⁷ For the [Romer and Romer \(2010\)](#) tax shock series, I use the exogenous tax change series based on the present value of tax changes. As another variable, I also use the growth rate of the price of crude oil as a metric for oil supply shocks.⁸ For these pairwise Granger causality tests, I use the longest available sample given the data in each test. I use four lags for all tests.

Table 8 presents F statistics and p values for the hypotheses that different exogenous shocks do not Granger cause the growth rate of adjusted TFP. For the 2007 and 2011 vintages of the adjusted TFP data, one can reject the hypotheses that the [Ramey \(2011\)](#) defense news shock series and the [Romer and Romer \(2004\)](#) monetary policy shock series do not Granger cause adjusted TFP at the five percent significance level or better. One cannot reject the hypotheses that the [Romer and Romer \(2010\)](#) tax shock series and the crude oil price do not Granger cause adjusted TFP growth for either the 2007 or 2011 vintages. When using the 2015 vintage of the adjusted TFP data, the hypotheses that any of the four variables under consideration do not Granger cause adjusted TFP cannot be rejected at the 5 percent significant level. The series with the most predictive value for adjusted TFP is the defense news shock series from [Ramey \(2011\)](#), for which the p value from the Granger causality test is 0.09. To the extent to which one takes the different series included in these tests as shocks unrelated to exogenous productivity, the results in Table 8 indicate that the most recent vintage of the adjusted TFP series represents an improvement over earlier vintages. It is far less predictable from these shocks than either the 2007 or 2011 vintages.

Although not the principal focus of the present paper, it is also of interest to investigate the extent to which the adjusted TFP series helps to forecast the exogenous shock series. If these series are in fact exogenous shocks, they ought to be unpredictable. Table 9 presents F statistics and p values from Granger causality tests that different vintages of the adjusted TFP series do not Granger cause these shocks. The 2007 and 2011 vintages of the adjusted TFP series have important predictive value for the [Ramey \(2011\)](#) defense news shocks, [Romer and Romer \(2010\)](#) tax shocks, and [Romer and Romer \(2004\)](#) monetary shocks. The p values for the tests that these vintages of the adjusted TFP series do not Granger cause these shocks are all 0.05 or lower. Neither the 2007 nor 2011 vintages of the adjusted TFP series Granger cause crude oil prices. These results cast some doubt that these different shocks identified in the literature can in fact be considered shocks – if they were shocks, they ought to be unpredictable. Interestingly, the results are somewhat different when using the 2015 vintage of

⁷The provided series is monthly. I convert to a quarterly frequency by averaging observations within a quarter.

⁸In particular, I use the BLS index for domestic crude oil production, series IS WPU0561.

the adjusted TFP series. While one can reject the hypothesis that it does not Granger cause the [Romer and Romer \(2004\)](#) monetary policy shock series, one cannot reject the hypotheses that it does not Granger cause the [Ramey \(2011\)](#) or [Romer and Romer \(2010\)](#) shocks.

3 Application to News Shocks

Section 2 documents significant differences in the adjusted TFP series by vintage provided by [Fernald \(2014\)](#). It also shows that these differences have substantive effects on unconditional correlations of the adjusted TFP series by vintage with several aggregate variables. This section considers how different vintages of the adjusted TFP data might affect conditional moments. In particular, I examine how different vintages of the adjusted TFP data might affect conclusions about the macroeconomic effects of “news shocks” about future productivity. The utilization adjusted TFP data provided by [Fernald \(2014\)](#) have been extensively used in this literature.⁹

I follow the empirical strategy outlined in [Barsky and Sims \(2011\)](#). They seek to identify the macroeconomic effects of news shocks about future productivity, where the news shock is defined as a shock to future productivity which is observed by agents in advance. In their empirical analysis they make use of the 2007 vintage of adjusted TFP data provided by [Fernald \(2014\)](#). Their substantive conclusion is that while news shocks are an important driver of output movements in the medium run, in the short run a favorable news shock causes output and hours worked to decline, while consumption rises. That is, not only does a positive news shock lead to a contraction in economic activity, it induces negative co-movement between consumption and output. This co-movement makes news shocks a poor candidate source of business cycle fluctuations. These responses are qualitatively consistent with the implications of flexible price, real business cycle models.

The interested reader is referred to [Barsky and Sims \(2011\)](#) for technical details of the empirical approach. The basic identification strategy can be summarized briefly as follows. First, a reduced form finite-order vector autoregression (VAR) is estimated on a set of data which includes a measure of adjusted TFP. A surprise productivity shock is identified as the reduced form innovation in adjusted TFP (i.e. this is equivalent to ordering adjusted TFP first in a Choleski decomposition). The news shock is then identified as the linear combination of reduced form innovations in the VAR, restricted to be orthogonal to the surprise adjusted TFP innovation, which maximally explains the remaining variance share of adjusted TFP

⁹In addition [Barsky and Sims \(2011\)](#), other prominent examples of papers which seek to identify news shocks using Fernald’s quarterly adjusted TFP series include [Kurmann and Otrok \(2013\)](#), [Beaudry and Portier \(2014\)](#), [Kurmann and Mertens \(2014\)](#), [Barsky, Basu, and Lee \(2015\)](#), [Levchenko and Pandalai-Nayar \(2015\)](#), and [Ben-Zeev and Kahn \(2015\)](#).

over some specified forecast horizon. The approach can be applied to a VAR in levels, a VAR in first differences with imposed cointegrating relationships, or to a vector error correction model (VECM) where cointegrating relationships are formally estimated.

3.1 A Four Variable VAR

[Barsky and Sims \(2011\)](#) begin by estimating a four variable VAR. The variables included in the VAR are the 2007 vintage of adjusted TFP, real consumption of non-durables and services (the construction of which is explained in Footnote 5), real output in the non-farm business sector, and total hours worked in the non-farm business sector. The consumption, output, and hours series are expressed in per capita units by dividing by the civilian non-institutionalized population aged sixteen and over. The VAR is estimated in log-levels of the variables, with four lags. The sample period is 1960q1-2007q3. The news shock is identified by maximizing its contribution to the variance share of adjusted TFP over a forty quarter horizon.

I re-estimate this four variable VAR using the aforementioned three different vintages of adjusted TFP data – the 2007 vintage as in [Barsky and Sims \(2011\)](#), as well as the 2011 and 2015 vintages. The same consumption, output, and hours data are used in each VAR. Impulse responses to the identified news shock are shown in Figure 1. Solid lines depict the responses estimated using the 2007 vintage adjusted TFP data, dashed lines correspond to responses estimated from the 2011 vintage of data, and dotted lines are from estimation with the 2015 vintage adjusted TFP data.

The impulse responses estimated from the 2007 vintage data are virtually identical to the responses reported in [Barsky and Sims \(2011\)](#). Adjusted TFP increases rather quickly and continues to grow for an extended period of time. Consumption jumps up on impact, while output and hours decline. The impact decline in hours is particularly strong. Subsequent to impact, output, consumption, and hours all grow with adjusted TFP. The output response quickly turns positive, while it takes about four quarters for the hours response to turn positive. After turning positive, hours worked grows for a while before reverting back to its pre-shock value. In the long run, consumption and output are about 0.5 percent higher, while adjusted TFP is higher by about 0.3 percent. These long run responses are roughly consistent with balanced growth implications with a labor’s share parameter of two-thirds.

The impulse responses estimated with the 2011 vintage of adjusted TFP are very similar to those estimated from the 2007 vintage data. Relative to the 2007 vintage data, the adjusted TFP response to the news shock is a bit slower, consumption increases by less on impact, hours worked fall substantially on impact but by a little less than with the 2007 data, and output actually rises slightly on impact. Nevertheless, these differences are quantitatively

small, and qualitatively the impulse responses look quite similar to those estimated with the 2007 vintage data.

There are large and substantive differences with the impulse responses estimated using the 2015 vintage of adjusted TFP data relative to either the 2007 or 2011 vintages. Most notably, output is estimated to rise significantly on impact, with an increase of about 0.5 percent, close to its estimated long run response. Hours worked essentially does not respond on impact, before turning significantly positive after about four quarters. These responses appear much more favorable to the news-driven business cycle story than the responses estimated using either the 2007 or 2011 vintages of adjusted TFP data.

The large differences in the estimated impulse response functions are also evident in a forecast error variance decomposition. Table 10 presents the forecast error variance decomposition of output, hours, and consumption for the three different vintages of the adjusted TFP data under consideration. The numbers in the table show the percentage of the forecast error variance share of each variable attributable to the identified news shock at different forecast horizons ($h = 0$ (impact), $h = 4$, $h = 8$, and $h = 16$). Using either the 2007 or 2011 vintages of adjusted TFP data, the news shock explains essentially none of the variance share of output on impact and a large share of the hours variance. At longer forecast horizons, the news shock accounts for somewhat more of the output variance but substantially less of the variance of hours worked. The long run variance share of output and consumption accounted for by the news shock is between forty and fifty percent, in line with the estimates in Barsky and Sims (2011). The variance decomposition using the 2015 vintage of data is quite different. The news shock accounts for more than a quarter of the impact variance share of output and virtually none of the impact variance share of hours worked. The variance share due to the news shock grows for output as the forecast horizon is longer, and the news shock accounts for a larger share of medium horizon fluctuations in hours worked than with either the 2007 or 2011 data.

Barsky and Sims (2011) also show impulse responses to a surprise productivity shock, identified as the reduced form innovation in adjusted TFP ordered first. Figure 2 plots the responses of adjusted TFP and hours worked to the surprise productivity shock identified using the three different vintages of adjusted TFP. Regardless of vintage of data, the surprise productivity shock is persistent but clearly transitory. With the 2007 or 2011 vintages of adjusted TFP data, hours worked increases (albeit not by much) in response to a positive surprise productivity shock. This is in line with the results in Barsky and Sims (2011). Using the 2015 data, in contrast, hours worked declines after a positive productivity shock, both on impact and even more so at longer forecast horizons. These differences are important when thinking about mapping the results back into a model. Hours worked increasing following a

positive productivity shock is consistent with a flexible price, real business cycle model, while hours worked declining when productivity increases is potentially consistent with a sticky price, New Keynesian model.

3.2 A Seven Variable VAR

Barsky and Sims (2011) also identify a news shock in a seven variable VAR. In addition to the adjusted TFP series, output, consumption, and hours worked, their seven variable VAR also includes a measure of inflation, a measure of consumer confidence, and a measure of stock prices.¹⁰ There are tradeoffs at play when considering how many variables to include in a VAR. On the one hand, including extraneous variables reduces statistical precision and could potentially induce spurious results. On the other hand, it is important that the variables included in the VAR span the information set of agents in an economy. News shocks are known to potentially induce non-invertibility into a VAR model, which means that structural shocks cannot be recovered from reduced form innovations. Sims (2012) shows that non-invertibility is essentially a problem of missing information. Adding more variables to a VAR, particularly forward-looking variables like stock prices or consumer confidence, can solve or alleviate an invertibility problem.

Figure 3 plots impulse responses to identified news shocks in the seven variable VAR system. As in the four variable system, the VARs are estimated in log levels with four lags. The solid lines show responses using the 2007 vintage of adjusted TFP data, the dashed lines the 2011 vintage of adjusted TFP data, and the dotted lines the 2015 vintage. The estimated responses using the 2007 or 2011 vintage are very similar, but larger differences emerge when considering the 2015 vintage of adjusted TFP data. Regardless of vintage, hours worked declines on impact after a favorable news shock, though the impact decline is significantly smaller with the 2015 vintage of adjusted TFP data. Output declines on impact when using either the 2007 or 2011 vintages of data, but rises slightly on impact with the most recent adjusted TFP data. The impact rise in output with the 2015 vintage of adjusted TFP data is much smaller, however, when compared to the results in the four variable VAR. Consumption, consumer confidence, and stock prices all rise on impact regardless of vintage. Inflation falls

¹⁰The measure of consumer confidence is a qualitative index of five year ahead business condition expectations available from the Michigan Survey of Consumers. The inflation measure is the log first difference of the GDP deflator. The real stock price series is the nominal S&P 500 index, aggregated to a quarterly frequency by averaging over months, deflated by the GDP deflator. The stock price series is downloaded from Robert Shiller's website: <http://www.econ.yale.edu/~shiller/data.htm>. The definitions of inflation and the real stock price differ slightly from those used in Barsky and Sims (2011). They measure inflation via the percentage change in the CPI, and deflate the nominal stock price via the CPI. Furthermore, they use end of quarter observations to transform the underlying monthly series to a quarterly frequency, whereas here I take within-quarter averages.

significantly on impact when using the 2007 or 2011 vintages of adjusted TFP data and is persistently negative. With the 2015 data, in contrast, the impact response of inflation is actually slightly positive, though the response quickly turns negative.

Figure 4 plots impulse responses of adjusted TFP and hours worked to the surprise productivity shock identified as the reduced form innovation in the seven variable VAR. Similarly to the four variable VAR, the impact response of hours is positive using either the 2007 or 2011 vintages of adjusted TFP. In contrast, the impact response of hours is negative when the VAR is estimated using the 2015 vintage of adjusted TFP data.

In summary, the vintage of adjusted TFP data from [Fernald \(2014\)](#) is potentially important when drawing inference about the macroeconomic consequences of news shocks about future productivity. Using the most recent vintage of data, in a four variable VAR the news shock is estimated to lead to a large increase in output in advance of changes in productivity, potentially leading credence to the news-driven business cycle hypothesis. The results in the seven variable VAR are qualitatively similar though the impact increase in output is much smaller than in the four variable VAR. In either VAR, hours worked decline on impact, rather than rise as in earlier vintages of the data, to a surprise productivity shock.

4 Conclusion

This paper documents quantitatively large and qualitatively meaningful differences across vintages in the properties of the quarterly utilization-adjusted TFP series developed by [Fernald \(2014\)](#). The most recent vintage of the adjusted TFP series is much less procyclical, and more negatively correlated with labor market indicators, than earlier vintages of the series. The differences across vintage seem to be driven primarily by changes in 2014 in the way that factor utilization series is computed.

The quarterly utilization-adjusted TFP series provided on Fernald's website has proven highly influential in applied work. There are many papers which have made use of these data whose empirical results may differ in important ways when using the most recent vintage of the data. As an illustration of this, I revisit the literature on the effects of news shocks about future productivity. Fernald's data have been used extensively in this literature. I re-estimate both the four and seven variable VARs in [Barsky and Sims \(2011\)](#) using three different vintages of the adjusted TFP series. In the seven variable VAR, their results are qualitatively robust to using the most recent vintage of the adjusted TFP data, though with the most recent vintage output and hours fall less in response to favorable news than with older vintages of the adjusted TFP data. In their four variable VAR, the results with the most recent vintage of the adjusted TFP data are qualitatively different in way that is much more favorable to the news-driven business cycle hypothesis.

References

- AARONSON, D. AND D. SULLIVAN (2001): “Growth in Worker Quality,” *Economic Perspectives*, 25, 53–74.
- BARSKY, R., S. BASU, AND K. LEE (2015): “Whither News Shocks?” in *NBER Macroeconomics Annual 2014, Volume 29*, National Bureau of Economic Research, Inc, NBER Chapters, 225–264.
- BARSKY, R. AND E. SIMS (2011): “News Shocks and Business Cycles,” *Journal of Monetary Economics*, 58, 273–289.
- BASU, S., J. FERNALD, J. FISHER, AND M. KIMBALL (2013): “Sector-Specific Technical Change,” Tech. rep.
- BASU, S., J. G. FERNALD, AND M. S. KIMBALL (2006): “Are Technology Improvements Contractionary?” *American Economic Review*, 96, 1418–1448.
- BEAUDRY, P. AND F. PORTIER (2014): “News-Driven Business Cycles: Insights and Challenges,” *Journal of Economic Literature*, 52, 993–1074.
- BEN-ZEEV, N. AND H. KAHN (2015): “Investment-Specific News Shocks and US Business Cycles,” *Journal of Money, Credit, and Banking*, 47, 1443–1464.
- BURNSIDE, C., M. EICHENBAUM, AND S. REBELO (1993): “Labor Hoarding and the Business Cycle,” *Journal of Political Economy*, 101, 245–273.
- CASCALDI-GARCIA, D. (2016): “News Shocks and the Slope of the Term Structure of Interest Rates: Comment,” Tech. rep., Warwick Business School.
- FERNALD, J. (2014): “A Quarterly, Utilization-Adjusted Series on Total Factor Productivity,” Federal Reserve Bank of San Francisco Working Paper Series 2012-19.
- KURMANN, A. AND K. MERTENS (2014): “Stock Prices, News, and Economic Fluctuations: Comment,” *American Economic Review*, 104, 1439–1445.
- KURMANN, A. AND C. OTROK (2013): “News Shocks and the Slope of the Term Structure of Interest Rates,” *American Economic Review*, 103, 2612–2632.
- KYDLAND, F. E. AND E. C. PRESCOTT (1982): “Time to Build and Aggregate Fluctuations,” *Econometrica*, 50, 1345–1370.

- LEVCHENKO, A. AND N. PANDALAI-NAYAR (2015): “TFP, News, and “Sentiments:” The International Transmission of Business Cycles,” NBER Working Papers 21010, National Bureau of Economic Research.
- LONG, J. B. AND C. PLOSSER (1983): “Real Business Cycles,” *Journal of Political Economy*, 91, 39–69.
- RAMEY, V. A. (2011): “Identifying Government Spending Shocks: It’s All in the Timing,” *Quarterly Journal of Economics*, 126, 1–50.
- ROMER, C. D. AND D. H. ROMER (2004): “A New Measure of Monetary Shocks: Derivation and Implications,” *American Economic Review*, 94, 1055–1084.
- (2010): “The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks,” *American Economic Review*, 100, 763–801.
- SIMS, E. (2012): “News, Non-Invertibility, and Structural VARs,” *Advances in Econometrics*, 28, 81–136.
- SUMMERS, L. H. (1986): “Some Skeptical Observations on Real Business Cycle Theory,” *Federal Reserve Bank of Minneapolis Quarterly Review*, Fall, 23–27.

Table 1: Descriptive Statistics – Adjusted TFP Series by Vintage

	$\Delta \ln \widehat{A}_t^{07}$	$\Delta \ln \widehat{A}_t^{11}$	$\Delta \ln \widehat{A}_t^{15}$
Mean	1.4929	1.3694	1.4008
Standard Deviation	3.4075	3.6226	3.4140
Corr w/ $\Delta \ln \widehat{A}_t^{07}$	1.0000	0.9502	0.5859
Corr w/ $\Delta \ln \widehat{A}_t^{11}$	0.9502	1.0000	0.5947

Notes: This table shows descriptive statistics for the different vintages of adjusted TFP growth. The sample period for these statistics is fixed at 1947q3-2007q3. The underlying series are expressed in annualized percentage points.

Table 2: Descriptive Statistics – Adjusted TFP Series by Vintage, Different Sample Periods

	1960q1-2007q3	1960q-1979q4	1984q1-2007q3
Corr b/w $\Delta \ln \widehat{A}_t^{07}$, $\Delta \ln \widehat{A}_t^{11}$	0.952	0.976	0.896
Corr b/w $\Delta \ln \widehat{A}_t^{07}$, $\Delta \ln \widehat{A}_t^{15}$	0.586	0.580	0.569
Corr b/w $\Delta \ln \widehat{A}_t^{11}$, $\Delta \ln \widehat{A}_t^{15}$	0.596	0.543	0.636

Notes: This table presents correlations among different vintages of the adjusted TFP series, where the sample period is restricted to periods indicated in columns.

Table 3: Correlations with Aggregate Variables – Adjusted TFP Series by Vintage

	$\Delta \ln \widehat{A}_t^{07}$	$\Delta \ln \widehat{A}_t^{11}$	$\Delta \ln \widehat{A}_t^{15}$
$\Delta \ln Y_t$	0.528	0.436	0.069
$\Delta \ln Y_t^{nfb}$	0.493	0.405	0.023
$\Delta \ln Y_t^b$	0.554	0.463	0.109
$\Delta \ln C_t$	0.139	0.125	0.105
$\Delta \ln I_t$	0.268	0.192	0.038
π_t	-0.186	-0.157	-0.165
$\Delta \ln H_t^{nfb}$	-0.010	-0.065	-0.362
$\Delta \ln H_t^b$	-0.067	-0.112	-0.403

Notes: This table shows correlations of the growth rates of the adjusted TFP series with the log first differences of several different aggregate variables (with the exception of the inflation rate, π_t , which appears in levels). Y_t is the headline real GDP series from the NIPA tables. Y_t^{nfb} and Y_t^b are the real output series in the non-farm business and business sectors, respectively. C_t is real consumption of non-durables and services, constructed using nominal share weights to sum the growth rates of real nondurable and real services consumption. I_t is real investment, defined as the sum of expenditure on durable goods and private fixed investment. The real series is constructed in the same way as the consumption series. π_t is the log first difference of the GDP price deflator. H_t is hours worked in the non-farm business sector, and H_t^b is hours worked in the business sector. These correlations are computed over the sample period 1947q3-2007q3.

Table 4: Correlations with Aggregate Variables – Adjusted TFP Series by Vintage, HP Filtered Data

	$\ln \widehat{A}_t^{07}$	$\ln \widehat{A}_t^{11}$	$\ln \widehat{A}_t^{15}$
$\ln Y_t$	0.152	-0.056	-0.203
$\ln Y_t^{nfb}$	0.165	-0.055	-0.189
$\ln Y_t^b$	0.201	-0.012	-0.141
$\ln C_t$	0.092	-0.072	-0.103
$\ln I_t$	0.174	0.011	-0.038
π_t	-0.269	-0.320	-0.236
$\ln H_t^{nfb}$	-0.103	-0.253	-0.405
$\ln H_t^b$	-0.112	-0.254	-0.406

Notes: This table is similar to Table 3, but uses HP filtered log levels of the variables (with smoothing parameter 1600).

Table 5: Descriptive Statistics – Components of Adjusted TFP

	$\Delta \ln Y_t^{07}$	$\Delta \ln Y_t^{11}$	$\Delta \ln Y_t^{15}$
Mean	3.5153	3.5329	3.5857
Standard Deviation	4.8642	4.8718	4.6848
Corr w/ $\Delta \ln Y_t^{07}$	1.0000	0.9961	0.9612
Corr w/ $\Delta \ln Y_t^{11}$	0.9961	1.0000	0.9640
	$\Delta \ln K_t^{07}$	$\Delta \ln K_t^{11}$	$\Delta \ln K_t^{15}$
Mean	3.5395	3.6897	3.8861
Standard Deviation	1.1851	1.1857	1.1031
Corr w/ $\Delta \ln K_t^{07}$	1.0000	0.9388	0.9107
Corr w/ $\Delta \ln K_t^{11}$	0.9388	1.0000	0.9864
	$\Delta \ln H_t^{07}$	$\Delta \ln H_t^{11}$	$\Delta \ln H_t^{15}$
Mean	1.0865	1.0845	1.0936
Standard Deviation	3.5478	3.5461	3.5459
Corr w/ $\Delta \ln L_t^{07}$	1.0000	0.9995	0.9995
Corr w/ $\Delta \ln L_t^{11}$	0.9995	1.0000	0.9997
	$\Delta \ln Q_t^{07}$	$\Delta \ln Q_t^{11}$	$\Delta \ln Q_t^{15}$
Mean	0.3575	0.3531	0.3528
Standard Deviation	0.3713	0.7111	0.7049
Corr w/ $\Delta \ln Q_t^{07}$	1.0000	0.3627	0.3670
Corr w/ $\Delta \ln Q_t^{11}$	0.3627	1.0000	0.9994
	$\Delta \ln u_t^{07}$	$\Delta \ln u_t^{11}$	$\Delta \ln u_t^{15}$
Mean	-0.0750	-0.0020	-0.0350
Standard Deviation	2.3388	2.9489	3.7119
Corr w/ $\Delta \ln u_t^{07}$	1.0000	0.9607	0.6611
Corr w/ $\Delta \ln u_t^{11}$	0.9607	1.0000	0.6443

Notes: This table presents descriptive statistics on the different vintages of the underlying series used in the construction of the adjusted TFP series. K_t is the capital stock series, H_t is hours worked, Q_t is the labor quality series, and u_t is the total utilization series. The sample period used to compute these statistics is fixed across vintages at 1947q3-2007q3.

Table 6: Descriptive Statistics – Synthetic Adjusted TFP Series

	Correlation with $\Delta \ln \widehat{A}_t^{07}$
Constructed w/ 2007 vintage utilization data	
$\Delta \ln \widehat{A}_t^{11}$	0.983
$\Delta \ln \widehat{A}_t^{15}$	0.907
Constructed w/ 2007 vintage labor quality data	
$\Delta \ln \widehat{A}_t^{11}$	0.956
$\Delta \ln \widehat{A}_t^{15}$	0.585
Constructed w/ 2007 vintage utilization and labor quality data	
$\Delta \ln \widehat{A}_t^{11}$	0.992
$\Delta \ln \widehat{A}_t^{15}$	0.916

Notes: This table presents correlations of “synthetic” vintages of the 2011 and 2015 vintages of adjusted TFP growth with the actual 2007 vintage adjusted TFP series. The synthetic vintages are constructed using the same year’s vintage component series, with the exception of the utilization and/or labor quality series, which are measured according to the 2007 vintage.

Table 7: Vintages Before and After the Switch to [Basu, Fernald, Fisher, and Kimball \(2013\)](#) Methodology of Measuring Utilization

	$\ln \widehat{A}_t^{13}$	$\ln \widehat{A}_t^{14}$	$\Delta \ln u_t^{13}$	$\Delta \ln u_t^{14}$
$\Delta \ln \widehat{A}_t^{14}$	0.625	1.000		
$\Delta \ln \widehat{A}_t^{07}$	0.852	0.562		
$\Delta \ln \widehat{A}_t^{11}$	0.910	0.557		
$\Delta \ln \widehat{A}_t^{15}$	0.641	0.860		
$\Delta \ln u_t^{14}$			0.593	1.000
$\Delta \ln u_t^{07}$			0.938	0.579
$\Delta \ln u_t^{11}$			0.971	0.572
$\Delta \ln u_t^{15}$			0.656	0.866

Notes: This table shows correlations of the December 2013 and May 2014 vintages of the adjusted TFP and utilization data with other vintages of the same data. These vintages are immediately before and after the switch to the [Basu et al. \(2013\)](#) methodology for measuring factor utilization, which occurred in March 2014.

Table 8: Granger Causality Results
 Tests of Hypothesis Exogenous Shocks do not Granger Cause Adjusted TFP

Vintage	Exogenous Shock	F statistic	p value
$\Delta \ln \widehat{A}_t^{07}$	Ramey (2011) Defense News	3.5057	0.0084
	Romer and Romer (2010) Tax Shocks	0.2254	0.9240
	Romer and Romer (2004) Monetary Shocks	2.5596	0.0432
	Crude Oil Price	1.0347	0.3901
$\Delta \ln \widehat{A}_t^{11}$	Ramey (2011) Defense News	3.0335	0.0183
	Romer and Romer (2010) Tax Shocks	0.2235	0.9251
	Romer and Romer (2004) Monetary Shocks	2.4732	0.0493
	Crude Oil Price	1.3824	0.2408
$\Delta \ln \widehat{A}_t^{15}$	Ramey (2011) Defense News	2.0110	0.0938
	Romer and Romer (2010) Tax Shocks	1.0300	0.3925
	Romer and Romer (2004) Monetary Shocks	0.8524	0.4954
	Crude Oil Price	0.6405	0.6271

Notes: This table presents results from pairwise Granger causality tests of the hypotheses that different exogenous shocks do not Granger cause the growth rate of adjusted TFP (by vintage). These causality tests are conducted using four lags and the longest available sample period. The sample periods for the Ramey (2011), Romer and Romer (2010), and crude oil price tests run from 1947q3-2007q3 (237 observations). The sample period for the Romer and Romer (2004) test runs from 1969q3-1996q4 (108 observations).

Table 9: Granger Causality Results
 Tests of Hypothesis Adjusted TFP does not Granger Cause Exogenous Shocks

Vintage	Exogenous Shock	F statistic	p value
$\Delta \ln \widehat{A}_t^{07}$	Ramey (2011) Defense News	3.6871	0.0062
	Romer and Romer (2010) Tax Shocks	2.4389	0.0478
	Romer and Romer (2004) Monetary Shocks	3.1236	0.0182
	Crude Oil Price	0.7043	0.5897
$\Delta \ln \widehat{A}_t^{11}$	Ramey (2011) Defense News	2.9834	0.0199
	Romer and Romer (2010) Tax Shocks	2.3697	0.0534
	Romer and Romer (2004) Monetary Shocks	3.3105	0.0137
	Crude Oil Price	0.9343	0.4448
$\Delta \ln \widehat{A}_t^{15}$	Ramey (2011) Defense News	1.1042	0.3554
	Romer and Romer (2010) Tax Shocks	1.6759	0.1564
	Romer and Romer (2004) Monetary Shocks	3.7952	0.0065
	Crude Oil Price	1.3570	0.2498

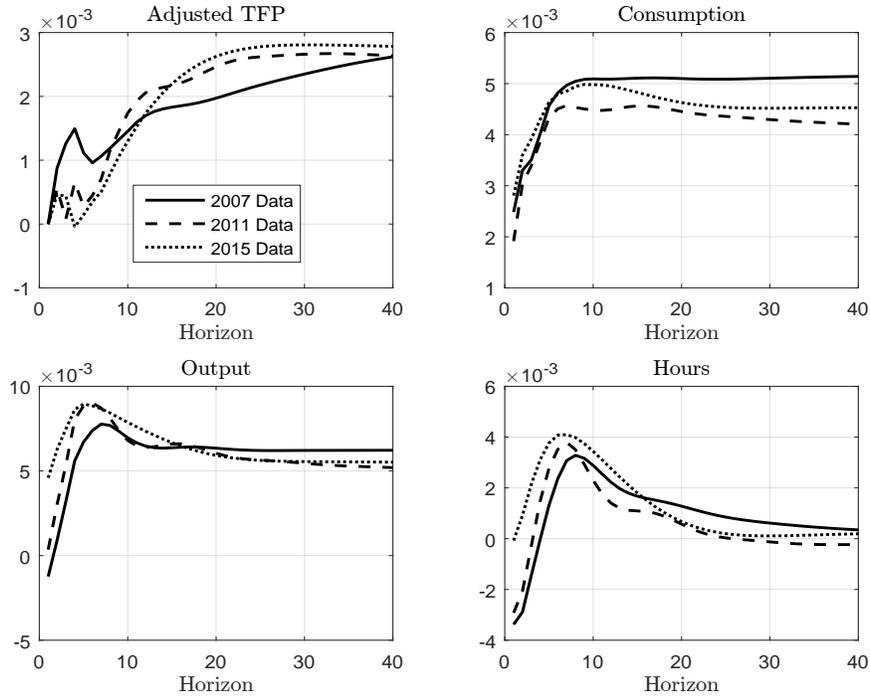
Notes: This table presents results from pairwise Granger causality tests of the hypotheses that the growth rate of adjusted TFP (by vintage) does not Granger cause different exogenous shocks. These causality tests are conducted using four lags and the longest available sample period. The sample periods for the Ramey (2011), Romer and Romer (2010), and crude oil price tests run from 1947q3-2007q3 (237 observations). The sample period for the Romer and Romer (2004) test runs from 1969q3-1996q4 (108 observations).

Table 10: Variance Share due to News Shock, Four Variable VAR

2007 Vintage Adjusted TFP			
Horizon	Output	Hours	Consumption
0	0.02	0.39	0.36
4	0.14	0.06	0.39
8	0.26	0.07	0.43
16	0.36	0.08	0.47
2011 Vintage Adjusted TFP			
Horizon	Output	Hours	Consumption
0	0.00	0.31	0.22
4	0.30	0.06	0.35
8	0.38	0.08	0.36
16	0.43	0.07	0.39
2015 Vintage Adjusted TFP			
Horizon	Output	Hours	Consumption
0	0.26	0.00	0.46
4	0.47	0.08	0.51
8	0.58	0.14	0.54
16	0.60	0.14	0.54

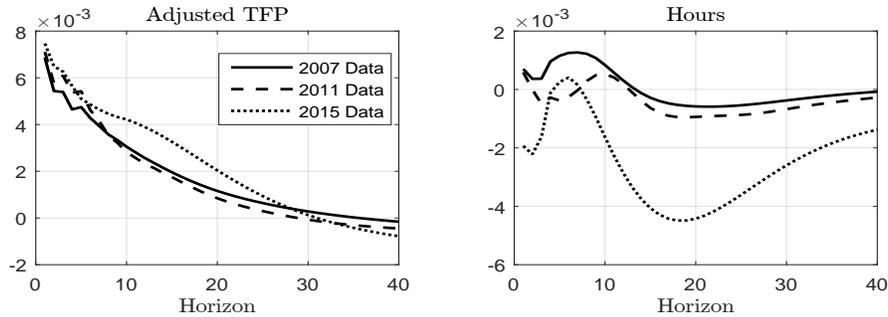
Notes: This table presents a forecast error variance decomposition of output, consumption, and hours based on the four variable VARs estimated in Section 3.1. The left column corresponds to the forecast horizon, where $h = 0$ is impact. The numbers in the table are the fraction of the forecast error variance of each variable attributable to the identified news shock.

Figure 1: IRFs to News Shock: Four Variable VAR



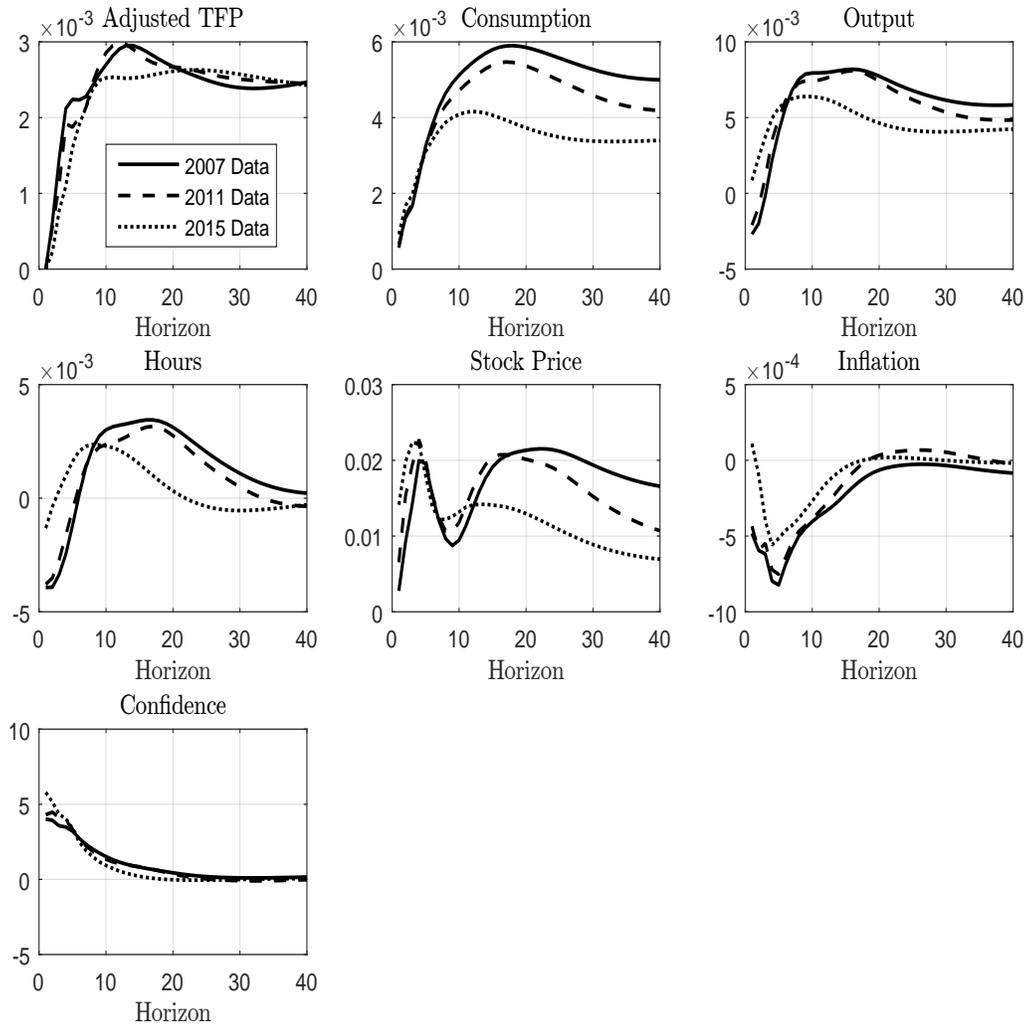
Notes: this figure plots impulse responses of adjusted TFP, consumption, output, and hours to the news shock identified from the four variable VAR described in Section 3.1. The solid lines are the impulse responses obtained using the 2007 vintage of adjusted TFP data, the dashed lines are the responses when using the 2011 vintage of adjusted TFP data, and the dotted lines are from a VAR with the 2015 vintage of adjusted TFP data.

Figure 2: IRFs to Surprise Productivity Shock: Four Variable VAR



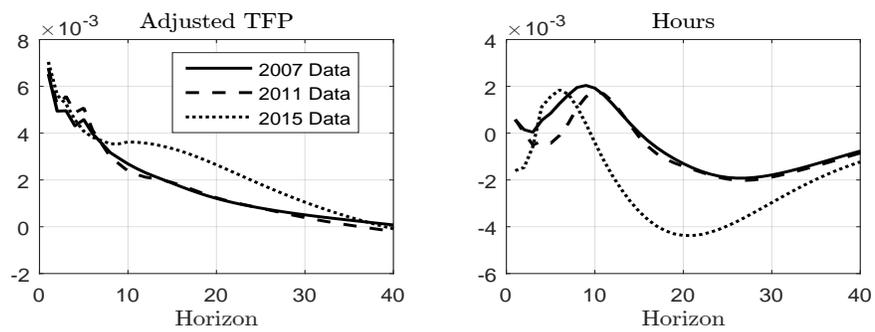
Notes: this table plots impulse responses of adjusted TFP and hours worked to the surprise productivity shock identified from the four variable VAR described in Section 3.1.

Figure 3: IRFs to News Shock: Seven Variable VAR



Notes: this table shows impulse responses of adjusted TFP, consumption, output, hours, inflation, stock prices, and consumer confidence to a news shock as identified from the seven variable VAR described in Section 3.2.

Figure 4: IRFs to Surprise Productivity Shock: Seven Variable VAR



Notes: this figure plots impulse responses of adjusted TFP and hours worked to the surprise productivity shock identified from the seven variable VAR described in Section 3.2.