

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

André Kurmann
Drexel University

Eric Sims
University of Notre Dame
& NBER

November 3, 2017

Abstract

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

JEL Classification: E22, E23, E32, O47

Keywords: Total factor productivity, variable utilization, news shocks

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

1 Introduction

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

In this paper, we critically revisit this assumption. Two conditions have to be met for the zero impact restriction on productivity to be satisfied. First, news about future productivity must not coincide with innovations in (true) productivity. Second, the empirical measure of productivity must not be confounded by current business cycle fluctuations. We show empirically and through model-based simulations that the identification of news shocks can be affected importantly if either condition is violated.² We then propose an alternative identification that is robust to these issues and apply it to U.S. data.

The starting point of our investigation is the quarterly utilization-adjusted series of total factor productivity (TFP) constructed by [Fernald \(2014\)](#) that has become the main measure of productivity in the news literature. Fernald frequently revises the adjusted TFP series based on new data and methodological refinements. We document that one of these revisions concerning the estimate of factor utilization significantly changes the cyclical behavior of Fernald's adjusted TFP series. We assess the consequences of this change in cyclical behavior for news shock identification by redoing the empirical analysis of [Barsky and Sims \(2011\)](#). Based on the 2007 vintage of Fernald's adjusted TFP series, as originally used by [Barsky and Sims \(2011\)](#), a positive news shock leads to a jump in consumption on impact but an initial decline in hours worked. As a result, the implied conditional correlation of consumption growth with hours growth is negative, leading [Barsky and Sims \(2011\)](#) to conclude that news shocks do not constitute a main driver of

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

²News shocks coincide with innovations in current productivity if, for example, the successful adoption of a new technology by a firm both raises current productivity and provides public information that other firms adopt the same technology in the future.

business cycles. Based on more recent vintages of Fernald’s adjusted TFP series published after the revision in utilization, in contrast, a positive news shock generally leads to an initial increase in consumption, hours, and other real aggregates, thereby affording a more favorable interpretation of the news-driven business cycle hypothesis laid out by [Beaudry and Portier \(2006\)](#).

The sensitivity of these results to a seemingly innocuous revision in Fernald’s adjusted TFP series suggests – as acknowledged by [Fernald \(2014\)](#) himself but otherwise mostly ignored by the literature – that measurement issues about productivity are a quantitatively important concern for news shock identification. Building on the insight of [Beaudry and Portier \(2006\)](#) that news shocks are related to innovations driving long-run variations in productivity, we consider an alternative identification that accounts for the maximum forecast error variance (FEV) share of adjusted TFP at a long but finite horizon. This identification is based on the “max-share” approach by [Francis et al. \(2013\)](#) – which in turn builds on earlier work by [Uhlig \(2003\)](#) – but differs in that we apply it to adjusted TFP instead of labor productivity and that we propose it as a possible news identification. While conceptually similar to [Barsky and Sims \(2011\)](#) and many close variants in the literature, the max-share identification does not impose orthogonality of the shock with current productivity and, more generally, does not rely on information about productivity fluctuations at short horizons. So long as mismeasurement of productivity is transient, the max-share identification should therefore be robust to measurement error, independent of whether innovations to expected future productivity have an immediate impact on (true) productivity or not. We verify this through several Monte Carlo simulations from a DSGE model in which the max-share identification outperforms the [Barsky and Sims \(2011\)](#) identification when adjusted TFP is subject to cyclical measurement error or when the shock moves true productivity on impact.

Of course, nothing guarantees that the max-share identification captures news shocks as opposed to other shocks driving future productivity. However, when applied to U.S. data, we find compelling evidence in favor of a news interpretation. First, the shock has no significant impact on adjusted TFP for several quarters but predicts sustained future productivity growth, accounting for 70 percent or more of TFP fluctuations at long forecast horizons. Second, the shock is associated with large impact responses of two novel indicators of innovation – an index of books published in the fields of technology by [Alexopoulos \(2011\)](#), and an index of technological standardization

by [Baron and Schmidt \(2015\)](#) – followed by a hump-shaped increase in R&D expenditures and a gradual decline in the relative price of investment goods. Third, the shock generates strong positive immediate reactions of forward-looking information variables and simultaneously leads to a decline in forecast uncertainty indices constructed by [Bachmann, Elstner, and Sims \(2013\)](#) and [Jurado, Ludvigson, and Ng \(2015\)](#). Taken together, these responses suggest that the max-share identification picks up current innovations in technology that trigger a subsequent rise in R&D and lead to higher productivity in the future, consistent with the observation that the diffusion of new technologies in many industries is slow because of costly experimentation and adoption.³ Furthermore, agents are aware of these innovations and update their expectations accordingly. The news interpretation therefore seems natural.

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

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

³See for example [Griliches \(1957\)](#), [Mansfield \(1961\)](#), [Mansfield \(1989\)](#), [Gort and Klepper \(1982\)](#), and [Rogers \(1995\)](#). Also see the discussion of these works in [Rotemberg \(2003\)](#).

⁴[Christiano et al. \(2004\)](#) base their argument on Granger causality tests with respect to hours. However, as pointed out by [Basu, Fernald, and Kimball \(2006\)](#), if technology is affected by news shocks, then hours can Granger-cause adjusted TFP even if technology is exogenous and adjusted TFP is a perfect measure of technology. The working paper by [Sims \(2016\)](#), which serves as a supporting document for the present paper, also runs Granger causality tests on different vintages of Fernald’s adjusted TFP series against different non-technology shock measures but finds the results to be inconclusive.

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

Second, the literature on long-run productivity shocks typically uses average labor productivity as the technology measure and is primarily concerned with the dynamics of hours worked in response to a shock.⁶ As such, this literature does not directly relate to the news literature and the idea that improvements in technology disseminate slowly and in a predictable manner. Indeed in many cases – including the max-share implementation by [Francis et al. \(2013\)](#) on which our identification is based – labor productivity jumps immediately to essentially its new permanent level.⁷ This would suggest that technology is a random walk process, which is very different from the results obtained in our paper. Moreover, as argued for example by [Uhlig \(2004\)](#), labor productivity is susceptible to being affected in the long-run by non-technology shocks that change the capital stock; e.g. permanent changes in taxation.

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

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

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

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

⁸The identification of [Barsky et al. \(2015\)](#) extracts the shock that accounts for *all* of the forecast revision of adjusted TFP at some long but finite horizon. We prefer the max-share approach because, as noted above, it does not impose that news shocks are the only source of predictable fluctuations at that particular horizon.

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

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

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

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

2 TFP and its Use in the News Literature

Following the lead of [Kydland and Prescott \(1982\)](#) and [Long and Plosser \(1983\)](#), the business cycle literature has typically measured productivity as the residual of aggregate output not accounted for by capital and labor inputs, commonly known as TFP. Economists quickly realized, however, that TFP may be a poor measure of productivity for a variety of reasons, most notably unobserved factor utilization.¹¹ In response to these concerns, [Basu, Fernald, and Kimball \(2006\)](#) apply restrictions derived from economic theory to industry-level data to construct an aggregate measure of TFP that takes into account sectoral heterogeneity, imperfect competition, compositional changes in the quality of labor and capital, and unobserved factor utilization. [Fernald \(2014\)](#) extends the analysis of [Basu et al. \(2006\)](#), which is carried out with annual data, to construct a quarterly measure of TFP. Because of the higher frequency, not all of the corrections in the original [Basu et al. \(2006\)](#) series are captured in the quarterly series, but perhaps the most important one – the correction for variable factor utilization – is. The resulting quarterly utilization-adjusted TFP series, which is available for download on John Fernald’s [website](#), has proven highly influential and has become the primary measure of productivity in the news literature.¹²

In what follows, we describe the construction of Fernald’s utilization-adjusted TFP. We then review how the literature uses this series to identify news shocks, and illustrate through the lens of a DSGE model how different measurement issues in the construction of the series may affect news shock 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

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

¹¹The idea that observed TFP fluctuations might be driven by endogenous responses 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 and conclude that much of the variation in TFP is not due to exogenous productivity shocks.

¹²As of November 2017, the working paper describing the construction of Fernald’s TFP series has been cited 363 times on [Google Scholar](#).

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 TFP, which is an empirical concept, and because the news literature 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

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

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

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

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 a quarterly estimate of TFP growth as

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

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.

For the utilization adjustment, 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 estimation is that cost-minimizing firms simultaneously vary inputs along all margins. This implies that variations in observed hours per worker can be used as a proxy for variations in 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. Finally, utilization-adjusted TFP growth is computed as

$$\Delta \ln TFP_t^U = \Delta \ln TFP_t - \Delta \ln U_t, \tag{5}$$

where $\Delta \ln U_t$ denotes the estimated growth in factor utilization.

2.2 Measurement issues for the news literature

Following Beaudry and Portier (2006), the news literature typically assumes that technology is driven by two exogenous components, one related to news shocks about expected future changes in fundamentals and the other capturing unanticipated or current shocks. The news shock is then identified by imposing that it affects technology only with a delay.¹³ While this zero impact restriction has intuitive appeal, it comes with potentially important caveats. As discussed in the Introduction, there is no *a priori* reason to think that news about growth-enhancing advances in technology are, despite their slow diffusion, completely unrelated to current productivity. Indeed, it seems equally intuitive to assume that market participants revise their expectations about future fundamentals only once there is evidence that at least some firms have successfully adopted the new

¹³As discussed further in Section 4, while this zero restriction is sufficient to identify news shocks in bivariate VARs, additional assumptions need to be imposed in VARs with more variables. Nevertheless, the zero restriction remains a key assumption in all of the empirical applications of which we are aware.

technology.¹⁴ Moreover, even if technology reacts to news shocks only with a delay, the zero impact restriction may still be violated in the data if the empirical measure of technology is confounded by measurement error. This is a distinct possibility for Fernald’s adjusted TFP series (as for any other series that one may consider in its place). First, the estimated utilization rate used to correct TFP may be an inaccurate measure of actual factor utilization. Second, as acknowledged 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).”

To illustrate the potential perils of identifying news shocks through the zero impact restriction when the underlying conditions are violated, we consider a medium scale New Keynesian DSGE model that is similar to Christiano, Eichenbaum, and Evans (2005), Smets and Wouters (2007) and Justiniano, Primiceri, and Tambalotti (2010). Here, we only discuss key parts of the model as they relate to the measurement of productivity. The full set of equilibrium conditions is provided in an online Appendix.

The aggregate production function of the model is

$$Y_t = A_t (Z_t K_t)^\alpha L_t^{1-\alpha} - FX_t, \tag{6}$$

where the different variables have the same interpretation as in (1); and $FX_t \geq 0$ is a fixed cost of production that increases with the economy’s trend growth factor X_t . Contrary to Fernald’s empirical approach, which allows for time-varying utilization of both capital and labor, the model features only capital utilization. This would be straightforward to change (e.g. by introducing labor hoarding) but would not affect the measurement issues we want to illustrate here.¹⁵

Following the news literature, technology is the sum of two components

¹⁴To our knowledge, the only other paper that discusses this possibility is Barsky, Basu, and Lee (2015) who write: “It is possible that news about future productivity arrives along with innovations in productivity today (page 233).”

¹⁵The model abstracts from several other features considered by Fernald, including time-variation in the quality of labor and capital, or sectoral heterogeneity in their utilization. Abstracting from these features does not invalidate the measurement issues illustrated by the model.

$$\ln A_t = \ln S_t + \ln \Gamma_t, \quad (7)$$

where S_t is the stationary component, governed by an AR(1) process

$$\ln S_t = \rho_S \ln S_{t-1} + \sigma_S \varepsilon_{S,t}, \quad (8)$$

and Γ_t is the permanent component, governed by an AR(1) process in the growth rate

$$\ln \Gamma_t - \ln \Gamma_{t-1} = (1 - \rho_\Gamma) \ln g + \rho_\Gamma (\ln \Gamma_{t-1} - \ln \Gamma_{t-2}) + \sigma_g \varepsilon_{g,t-q}, \quad (9)$$

with g denoting the steady state gross growth rate of technology; $0 < \rho_S < 1$, $0 < \rho_\Gamma < 1$; and $\varepsilon_{S,t}$, $\varepsilon_{g,t-q}$ independent and identically distributed. When $q > 0$, agents update expectations about the permanent component before the shock impacts technology, which is the identifying assumption in the news literature. However, even with $q = 0$, the process given in (9) embodies elements of slow diffusion as long as $\rho_\Gamma > 0$; i.e. a positive shock to the permanent component of technology today portends even larger increases in the level of technology in the future. Consistent with our argumentation in the Introduction and the empirical work below, we therefore refer to the shock to the permanent component as the news shock regardless of whether agents observe it in advance ($q > 0$) or not ($q = 0$).¹⁶

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

$$\Delta \ln TFP_t^M = \Delta \ln Y_t - \omega_{L,t} \Delta \ln L_t - (1 - \omega_{L,t}) \Delta \ln K_t, \quad (10)$$

where $\omega_{L,t}$ denotes the labor share. The growth rate of model adjusted TFP is then

$$\Delta \ln TFP_t^{M,U} = \Delta \ln TFP_t^M - (1 - \omega_{L,t}) \Delta \ln Z_t. \quad (11)$$

¹⁶Most of the literature associates news shocks with the permanent component of technology. As discussed in the Introduction and again in Section 5, this choice is motivated by the idea that growth-enhancing advances in technology diffuse slowly and are therefore anticipated by market participants, whereas the current component is taken to capture transitory unexpected changes in productive capacity. However, nothing about the illustration here would change if we associated news shocks with the stationary component of technology.

Even if Y_t , L_t , K_t , and Z_t are perfectly observed, there are nevertheless two incongruities between adjusted TFP, (11), and true technology, (7). First, if there is a fixed cost $F > 0$, then the production function is not constant returns to scale and the construction of TFP is misspecified. Second, since firms have market power and are subject to nominal price rigidities, $\omega_{L,t}$ and $1 - \omega_{L,t}$ do not in general correspond to the true factor elasticities $1 - \alpha$ and α of the model. Specifically, cost-minimization with respect to labor input implies

$$w_t L_t = (1 - \alpha) \mu_t^{-1} [Y_t - F X_t], \quad (12)$$

where μ_t^{-1} denotes the inverse price markup over marginal cost. If the fixed cost F is chosen to ensure zero profit along the balanced growth path, which is a standard assumption, then (12) becomes

$$\omega_{L,t} = \frac{w_t L_t}{Y_t} = (1 - \alpha) \mu_t^{-1} \mu^*, \quad (13)$$

with μ^* denoting the average markup. In this case, the labor share corresponds to the factor elasticity $1 - \alpha$ on average but fluctuates over time due to undesired fluctuations in the markup owing to price rigidity.¹⁷

In light of the empirical evidence presented in the next section, we introduce measurement error with respect to utilization of the form

$$u_t^{ob} = u_t + \mathbf{D} \mathbf{S}_t + \sigma_u \varepsilon_{u,t}, \quad (14)$$

where u_t^{ob} is the utilization rate observed by the econometrician; and $u_t \equiv (1 - \omega_{L,t}) \ln Z_t$ is the true utilization rate implied by the model. The difference between the two is measurement error composed of both a systematic component, $\mathbf{D} \mathbf{S}_t$, where \mathbf{S}_t denotes the state vector of the log-linearized model and \mathbf{D} is a conformable row vector of weights; and an idiosyncratic component $\sigma_u \varepsilon_{u,t}$, with $\varepsilon_{u,t}$ drawn from a standard normal distribution. While ad hoc, this formulation has the advantage that it can accommodate many different ways in which utilization may be mis-measured.

¹⁷In the absence of fixed costs, the production function is constant returns to scale, consistent with the assumption underlying the construction of TFP, while the labor share becomes $\omega_{L,t} = \frac{w_t L_t}{Y_t} = (1 - \alpha) \mu_t^{-1}$. Hence, the labor share differs from $1 - \alpha$ even on average. All the simulations below assume a positive fixed cost although we also experimented with zero fixed cost. The results remained largely unchanged.

We illustrate the extent to which the different misspecifications may affect identification by calibrating the model and computing the model-implied impulse responses of true technology and adjusted TFP to a news shock ε_g . The parameterization of the model is given in the online Appendix. The structural parameters are set in line with the estimates in [Justiniano, Primiceri, and Tambalotti \(2010\)](#). The parameters of the permanent component of technology in (9) are inspired by the empirical results obtained below although the exact values do not matter for the illustration.¹⁸

Figure 1 shows the impulse responses to a news shock for true technology (solid line) and for adjusted TFP (dotted line) under four alternative scenarios: (i) a “baseline” scenario according to which utilization is measured without error ($\mathbf{D} = \mathbf{O}$, $\sigma_u = 0$) and the news shock affects technology with a one period lag ($q = 1$); (ii) an “idiosyncratic measurement error” scenario ($\mathbf{D} = \mathbf{O}$, $\sigma_u \neq 0$, $q = 1$); (iii) a “systematic measurement error” scenario ($\mathbf{D} \neq \mathbf{O}$, $\sigma_u = 0$, $q = 1$); and (iv) a “contemporaneous news impact” scenario ($\mathbf{D} = \mathbf{O}$, $\sigma_u = 0$, $q = 0$). For the “systematic measurement error” case, we assume that all elements of \mathbf{D} are zero, with the exception of the element corresponding to the news shock itself being equal to -3 (note that the news shock is a state variable, given that it is anticipated in advance) and the combination of states corresponding to consumption equal to 1. This selection of coefficients is certainly arbitrary but motivated by the quantitative exercises performed in Section 6.

In the baseline scenario (upper left panel), there is virtually no discrepancy between the responses of adjusted TFP and true technology to a news shock. As we show later in Section 6, this result not only obtains conditional on a news shock but is more general: as long as utilization and all other inputs are measured correctly, adjusted TFP and true technology are almost collinear in spite of the fact that the presence of the fixed cost and time-varying markups invalidate Fernald’s assumptions in our model. This is an interesting result, as it suggests that inconsistencies between Fernald’s assumptions and common modeling assumptions about returns to scale and time-varying markups may not be quantitatively important.¹⁹

¹⁸Since we are only considering impulse responses to a news shock here, we do not need to take a stand about the parameterization of the stationary component of technology nor any of the other shock processes that may buffet the model economy.

¹⁹In our baseline calibration the price markup along the balanced growth path is 10%, in line with the estimates from [Basu et al. \(2006\)](#). This implies a relatively low fixed cost share of output when the fixed cost is chosen so that profits are zero along the balanced growth path. Our results are qualitatively similar even with much larger fixed cost shares (as in the [Smets and Wouters \(2007\)](#) estimate, which implies a fixed cost share of output

Similarly, in the idiosyncratic measurement error scenario (upper right panel), the responses of adjusted TFP and true technology lie virtually on top of each other. This is because the idiosyncratic measurement error is orthogonal to the news shock, making it irrelevant for the model-implied impulse responses. As we discuss in Section 6, however, idiosyncratic measurement error may make VAR identification more difficult.

The lower left panel shows the systematic measurement error scenario. Here we observe that the impulse responses of adjusted TFP and true technology do not align with one another. Adjusted TFP jumps up on impact of the news shock whereas true technology does not. This means that even though impact orthogonality between news and true technology is satisfied in the model, it is not satisfied with respect to adjusted TFP as measured by an econometrician.

The lower right panel shows the contemporaneous news impact scenario; i.e. when the news shock affects technology on impact ($q = 0$). Similarly to the first two scenarios considered in the upper row of the Figure, the responses of true technology and adjusted TFP lie virtually on top of one another. Differently than the other cases considered, however, true technology and hence adjusted TFP react on impact.

The impulse responses shown in Figure 1 illustrate that imposing the zero impact restriction on adjusted TFP can be problematic if utilization is mis-measured in a way that is systematically related to economic conditions, or if news affects technology immediately as opposed to with a lag. Of course, the extent to which these violations of the zero impact restrictions matter quantitatively for empirical applications remains a quantitative question. We examine this in detail in Sections 4 – 6 after documenting the revisions in Fernald’s adjusted TFP series.

3 Revisions to Utilization-Adjusted TFP

Fernald frequently revises his adjusted TFP series based on new data and methodological refinements. In this section, we analyze the consequences of these revisions. In particular, we document that a major change in business cycle properties of adjusted TFP occurred in March of 2014 when Fernald switched to using new estimates for the computation of unobserved factor utilization.

of 60 percent), albeit the incongruities between the data concept of adjusted TFP and true model technology are naturally somewhat larger with a larger fixed cost share.”

3.1 Time series properties of different vintages

The first part of Table 1 provides basic unconditional first and second moments of adjusted TFP growth for the December 2007, the December 2013, the May 2014, and the May 2016 vintages.²⁰ As in Fernald, TFP growth is computed as the quarterly log change, expressed in annualized percentage points. While the means and standard deviations of adjusted TFP growth are very similar across vintages, there is an important decline in comovement that occurs primarily from the 2014 vintage onward, with the correlation coefficient between the 2007 vintage and post-2013 vintages of adjusted TFP dropping to below 0.6. This decline in comovement across vintages of adjusted TFP holds for different subsamples and is therefore not driven by a change in business cycle behavior during a particular time period.²¹

As the second part of Table 1 shows, the decline in comovement across vintages of adjusted TFP growth also matters in significant ways for correlations with prominent macroeconomic aggregates. The 2007 and 2013 vintages of adjusted TFP growth are positively correlated with growth rates of output, consumption, and investment, and essentially uncorrelated with aggregate labor hours. In contrast, the 2014 and 2016 vintages of adjusted TFP growth are less positively correlated with the growth rates of output and consumption, uncorrelated with investment, and negatively correlated with hours.

3.2 What drives the changes across vintages?

Fernald’s data contains not only the adjusted TFP series but also the different aggregates used in calculating adjusted TFP. We use these variables to assess what drives the large differences in business cycle correlations across vintages of adjusted TFP.

Note from (4) and (5) that all variables but the estimate for utilization enter into the computation of non-adjusted TFP. As a first step, we therefore consider separately the time series behavior of non-adjusted TFP and utilization across different vintages. Figure 2 plots the log *levels* of the 2007 and 2016 vintages of adjusted TFP (top panel), non-adjusted TFP (middle

²⁰We also have data for the September 2011 vintage and the May 2015 vintage. The results for the 2011 vintage are very similar to the results for the 2007 and the 2013 vintage, while the results for the 2015 vintage are very similar to the results for the 2014 and the 2016 vintage.

²¹Very similar results obtain if we apply a Hodrick-Prescott filter to the log level of adjusted TFP instead of first-differencing.

panel), and utilization (bottom panel). The 2007 vintages are depicted as black lines while the 2016 vintages are shown as blue lines. The grey shaded bars show NBER-defined recessions.

As the top panel shows, the two vintages of the adjusted TFP series share roughly the same trend over the full sample but there are sizable differences over subsamples. The 2016 vintage grew more slowly during the first 25 years of the sample as well as from the mid-1980s to the mid-1990s. Concurrently, the 2016 vintage grew considerably faster from the late 1990s through the mid-2000s. Finally, whereas the 2016 vintage declined from 2005 onwards, the 2007 vintage showed an uptick near the end of the sample. As the middle panel shows, some of these differences in subsample trends are attributable to differences in non-adjusted TFP across vintages. But an equally if not more important part is driven by differences in utilization across vintages. As the bottom panel shows, while both the 2007 and the 2016 vintage of utilization are stationary and display overall similar fluctuations, there are sizable and persistent differences.²² In particular, the 2016 vintage displays substantially larger swings and is less smooth than the 2007 vintage.

To provide further insights, Table 2 reports the same summary statistics for the 2007, 2013, 2014 and 2016 vintages as in the first part of Table 1 but for non-adjusted TFP and utilization. Unadjusted TFP growth remains essentially unchanged across vintages in terms of overall mean, standard deviation, and comovement. Utilization growth, in contrast, becomes significantly more volatile starting with the 2014 vintage and there is an important decline in correlation with the 2007 vintage. This indicates that the large changes in business cycle properties of utilization-adjusted TFP across vintages are not due to data revisions in output, capital, labor, labor quality, or the labor share of income but rather are a direct consequence of the changes in volatility and business cycle properties of utilization that occurred primarily between the December 2013 and the May 2014 vintage.²³ This conclusion aligns with revision notes included in Fernald’s dataset. Indeed, in March 2014, Fernald switched from using industry utilization estimates by [Basu, Fernald, and](#)

²²By construction, Fernald’s utilization series evolves around a constant mean as the different industry-specific utilization rates are estimated based on bandpass-filtered (and therefore demeaned) data. The aggregate utilization rate does not have an exact zero mean because it is a weighted average of industry-specific rates. However, the actual means for both the 2007 and the 2016 vintages are close to zero (see Table 2). The figure reports the demeaned utilization rates to make the comparison of the two utilization vintages easier.

²³We should note that there are large differences in volatility and comovement of the labor quality series between the 2007 vintage and later vintages, including the 2011 vintage. These revisions to labor quality account for some of the differences in subsample trends of TFP across vintages, but not a significant amount. For example, the correlation between the 2011 vintage and the 2016 vintage of adjusted TFP growth is less than 0.6 despite the fact that the labor quality series for these two vintages is virtually identical.

Kimball (2006), which relied on data through 1995, to estimates from Basu, Fernald, Fisher, and Kimball (2013), which uses data through 2005.²⁴

To confirm that the change in the estimate of factor utilization after 2014 drives the change in time series behavior of adjusted TFP, we construct synthetic measures of adjusted TFP by combining the 2007 vintage of the utilization estimate with non-adjusted TFP series from other vintages. Correlations for the resulting synthetic measures are presented in Table 3. The correlation of the 2007 adjusted TFP vintage with the synthetic 2014 and 2016 vintages are both 0.91 (compared to 0.56 and 0.58 for the actual vintages of the adjusted TFP series). Hence, while the change in utilization after 2013 does not explain all of the divergence in the different vintages of adjusted TFP, it explains the large majority.

4 Implications for the Identification of News Shocks

To quantify the implications of the changes in business cycle behavior of adjusted TFP for the identification of news shocks, we redo the empirical analysis of Barsky and Sims (2011). They identify news shocks in a VAR context as the innovation that is orthogonal to Fernald’s adjusted TFP series but maximally accounts for the FEV share of adjusted TFP over a ten year horizon. We choose to focus on this identification rather than on one of the other identifications proposed in the literature because it performs well in small-sample Monte Carlo simulations (provided, of course, that the zero impact restriction is satisfied in the data generating process and technology is measured correctly) and because it is a partial identification approach that does not require taking a stand on the nature of non-news shocks.²⁵

4.1 Barsky-Sims identification

Since the alternative identification proposed below shares many of the elements with the identification by Barsky and Sims (2011), we review the details here. Let \mathbf{Y}_t be a $k \times 1$ random

²⁴An email exchange with Fernald confirms that the switch to the Basu, Fernald, Fisher, and Kimball (2013) estimates is the principal source for the change in the utilization series.

²⁵Indeed, full identification approaches are often subject to important robustness issues with respect to non-news shocks. See for example Kurmann and Mertens (2014) who show that the identification by Beaudry and Portier (2006) does not have a unique solution in their VAR systems with more than two variables; or Fisher (2010) who shows that the results by Beaudry and Lucke (2010) depend on the choice of cointegration restrictions imposed.

vector process of which the first variable is a measure of productivity (e.g. Fernald's utilization-adjusted TFP), and let the reduced form moving average representation of this process be given by $\mathbf{Y}_t = \mathbf{B}(\mathbf{L})\mathbf{u}_t$, where \mathbf{u}_t is a $k \times 1$ vector of prediction errors with variance-covariance matrix $E(\mathbf{u}_t\mathbf{u}_t') = \boldsymbol{\Sigma}_u$, and $\mathbf{B}(L) = \mathbf{I} + \mathbf{B}_1L + \mathbf{B}_2L^2 + \dots$ is a matrix lag polynomial. The coefficients of $\mathbf{B}(L)$ and $\boldsymbol{\Sigma}_u$ can be estimated with an unrestricted VAR.

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

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

$$\mathbf{Y}_{t+h} - E_{t-1}\mathbf{Y}_{t+h} = \sum_{l=0}^h \mathbf{B}_l\tilde{\mathbf{A}}\mathbf{Q}\epsilon_{t+h-l}. \quad (15)$$

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

$$\Omega_{i,j}(h) = \frac{\sum_{l=0}^h \mathbf{B}_{i,l}\tilde{\mathbf{A}}\gamma\gamma'\tilde{\mathbf{A}}'\mathbf{B}'_{i,l}}{\sum_{l=0}^h \mathbf{B}_{i,l}\boldsymbol{\Sigma}_u\mathbf{B}'_{i,l}}, \quad (16)$$

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

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

$$\begin{aligned} \max_{\gamma} \quad & \sum_{h=0}^H \Omega_{1,2}(h) \\ \text{s.t.} \quad & \gamma'\gamma = 1 \text{ and } \gamma(1,1) = 0, \end{aligned} \quad (17)$$

where without loss of generalization productivity is ordered first in \mathbf{Y}_t , and the news shock

is defined as the second shock in ϵ_t . The first restriction ensures that γ belongs to an orthonormal matrix. The second restriction imposes that the news shock has no immediate impact on productivity.

4.2 Results

We apply the Barsky-Sims identification to a four-variable VAR in Fernald’s utilization-adjusted TFP series (either the 2007 or 2016 vintage), real personal consumption expenditures per capita, total hours worked per capita in the non-farm business sector, and inflation as measured by the growth rate of the GDP price deflator.²⁶ Results for larger VARs that contain additional macro aggregates would be similar. With the exception of the inflation rate, the variables enter the VAR in log levels. The VAR is estimated with four lags via Bayesian methods subject to a Minnesota prior.²⁷ Confidence bands are computed by drawing from the resulting posterior distribution. The sample period is fixed at 1960q1-2007q3.²⁸ As in Barsky and Sims (2011), the truncation horizon is set to $H = 40$.

Figure 3 presents impulse responses to a news shock using the Barsky and Sims (2011) news identification. Here and below, the solid black lines show the posterior median impulse responses implied by the posterior distribution of the VAR estimated with the 2016 vintage of adjusted TFP, and the gray bands are the corresponding 16 to 84 percent posterior coverage intervals. In turn, the red dash-dotted lines show the posterior median impulse responses implied by the posterior distribution of the VAR estimated with the 2007 vintage of adjusted TFP, and the red dashed lines are the corresponding 16 to 84 percent posterior coverage intervals.

Using the 2007 vintage of adjusted TFP, the responses are similar to those estimated by Barsky and Sims (2011). Adjusted TFP starts to increase within a couple of quarters of the

²⁶Though we focus on vintages of the adjusted TFP data from 2007 and 2016, we note that the results with the 2011 or 2013 vintages of the data result in impulse responses which are similar to those when using the 2007 vintage of adjusted TFP. In contrast, when using the 2014 (post-adjustment in the measurement of utilization) or 2015 vintages, the results are similar to when using the 2016 vintage.

²⁷The Minnesota prior assumes a random walk process for adjusted TFP and consumption, and a white noise process for hours worked and the inflation rate. Estimates are robust to assuming a random walk prior for hours worked and inflation as well.

²⁸The beginning of the sample is chosen to facilitate comparison with Barsky and Sims (2011) and because some additional variables of interest which we study later are unavailable prior to 1960. Furthermore, the omission of the immediate post-war data from the sample removes some large influences due the 1951 Treasury Accord and Korean War. The end date is the last available observation for the 2007 vintage of adjusted TFP data.

shock; consumption jumps up while inflation falls significantly on impact; and hours worked initially declines, turning positive only after several quarters. As shown in Table 4, these impulse responses imply that if the economy was buffeted solely by news shocks, the correlation between consumption growth and hours growth would be strongly negative (second row), whereas in the data it is robustly positive (first row).²⁹ This leads Barsky and Sims (2011) to conclude that news shocks about future productivity cannot be a major source of business cycle fluctuations and that qualitatively, at least, the impact decline of hours worked is consistent with the predictions of a relatively frictionless real business cycle model.

Using the 2016 vintage of adjusted TFP instead, the impulse responses look quite different in economically important ways. Adjusted TFP reacts to the news shock with a substantially longer delay while hours worked increase from the beginning (although not significantly so for the first two quarters) and reach a peak response only a few quarters after the shock. This difference in the response of hours worked implies that the correlation of consumption growth and hours growth conditional on news shocks is now significantly positive (third row of Table 4), making news shocks a potentially important source of business cycle fluctuations. Moreover, the deflationary impact of news shocks, which Barsky, Basu, and Lee (2015) cite as one of the most robust features of the data, is no longer statistically significant.

Overall, the results based on the 2016 vintage of adjusted TFP are different in economically important ways compared to the responses estimated with vintages of adjusted TFP that predate the March 2014 revision. This difference has important implications for how we interpret business cycle fluctuations. In particular, as Barsky and Sims (2011) argue, the impulse responses based on the 2007 vintage are at least qualitatively consistent with a basic RBC model in which news shocks have a delayed permanent effect on technology. By contrast, the impulse responses obtained with the 2016 vintage are more in line with the view that news shocks have significant short-term demand effects that are accommodated by the economy and thus lead to comovement in real macro aggregates as observed over the business cycle.

²⁹The different business cycle moments implied by the Barsky-Sims news shock are very precisely estimated. In particular, the 16-84 coverage interval for the correlation between consumption growth and hours growth is $[-0.57, -0.50]$ if the VAR is estimated with the 2007 vintage of adjusted TFP; and $[0.22, 0.37]$ if the VAR is estimated with the 2016 vintage of adjusted TFP. To save on space, we do not report these coverage bounds here.

5 An Alternative Identification of News Shocks

In this section, we consider an alternative identification of news shocks that provides similar impulse responses regardless of the vintage of adjusted TFP used. We first motivate the alternative identification and describe its implementation. Then we present results for the macroeconomic implications and the news shock interpretation of this alternative identification.

5.1 Motivation and implementation

A large 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 \(1989\)](#), for example, the time until half of potential adopters actually adopt a new technology varies between five and fifteen years, depending on technology. While the slow dissemination of new technologies and its implications for the modeling of productivity is discussed extensively by [Rotemberg \(2003\)](#) as well as [Comin and Gertler \(2006\)](#) and [Lindé \(2009\)](#) among others, much of the business cycle literature has modeled productivity as a jump process where innovations lead to an immediate change of productivity to a new level that is either permanent or highly persistent.

The central idea proposed here is that this slow dissemination should constitute news in as much as these changes are known in advance and therefore lead to predictable changes in productivity growth. This idea is consistent with the insight of [Beaudry and Portier \(2006\)](#) from their bivariate VAR systems with stock prices and a measure of productivity that news shocks identified through the zero impact restriction are closely related to the shocks driving long-run movements in productivity.³⁰ Our contribution here consists of exploring this insight further by extracting a long-run productivity shock in larger VAR systems, assessing its robustness to the above documented revisions in adjusted TFP, and using new information to interpret the extracted shock as a news shock.

Specifically, we estimate the same four-variable VAR as above as well as larger VARs and extract the shock that accounts for the maximum FEV share of adjusted TFP at a long but

³⁰Subsequently, the news literature has focused almost exclusively on news identifications centered around the zero impact restriction and what additional restrictions should be imposed to identify news shocks in larger VAR systems. See the discussion in the beginning of Section 4 for references.

finite horizon H . This shock extraction is based on the “max-share” approach by [Francis et al. \(2013\)](#) – which in turn builds on earlier work by [Uhlig \(2003\)](#) – but differs in that we apply it to adjusted TFP instead of labor productivity and that we propose it as a possible news identification. Formally, the shock is defined by the solution to the following optimization problem

$$\begin{aligned} \max_{\gamma} \quad & \Omega_{1,2}(H) \\ \text{s.t.} \quad & \gamma' \gamma = 1 \end{aligned} \tag{18}$$

where, as per equation (16), $\Omega_{1,2}(H)$ denotes the FEV share of adjusted TFP at horizon H accounted for by the second shock in shock vector ϵ_t ; and γ denotes a column vector belonging to orthonormal rotation matrix \mathbf{Q} of the Choleski decomposition of the reduced form variance covariance matrix. While conceptually similar to [Barsky and Sims \(2011\)](#), there are two important differences. First, we look for the shock that accounts for the maximum FEV share of adjusted TFP at a long horizon H instead of maximizing the sum of FEV shares from impact onward. Second, we drop the zero restriction and do not impose that the news shock is orthogonal with respect to the innovation in adjusted TFP (i.e. the first element of γ is not restricted to zero, which means that adjusted TFP is allowed to respond contemporaneously to the shock). As discussed in the Introduction, this max-share approach has the advantage that by focusing on a long forecast horizon only, it diminishes the potential bias imparted by measurement error, which is putatively transient in nature, on the extracted shock. Moreover, the approach avoids taking a stand on whether (true) technology reacts to the shock only with a lag or not. Monte-Carlo simulations based on a medium-scale DSGE model presented in Section 6 show that in the presence of transient measurement errors or contemporaneous responses of technology to news shocks, the proposed identification indeed outperforms the identification of [Barsky and Sims \(2011\)](#).

The max-share approach has two advantages relative to other long-run identification schemes employed in the VAR literature. First, as shown in [Francis et al. \(2013\)](#), by focusing on a finite horizon the max-share approach helps reduce small-sample biases in finite-order VARs that can have potentially important effects for infinite-horizon identifications of long-run shocks.³¹ Second,

³¹Bias-reduction is the main motivation of [Francis et al. \(2013\)](#) for the max-share identification. Also see [Erceg et al. \(2005\)](#); [Christiano et al. \(2006\)](#); and [Chari et al. \(2008\)](#) for important contributions in this respect. Another, practical advantage of the max-share approach is that it can be implemented either with a VAR in levels that

the max-share approach does not impose that technology is the *only* source of long-run fluctuations in productivity and instead extracts the shock that accounts for the maximum FEV share of adjusted TFP at a long but finite horizon. The max-share approach therefore affords the possibility that other shocks (e.g. a surprise productivity shock) exert long-lasting effects on adjusted TFP.

It is also instructive to contrast our use of the max-share approach to the specification by [Francis et al. \(2013\)](#). Aside from the fact that the primary objective here is to devise a more robust approach to identifying news shocks, our work differs from [Francis et al. \(2013\)](#) in that we use adjusted TFP instead of labor productivity as the target variable. One advantage of using adjusted TFP is that labor productivity may be affected in more important ways by business cycle fluctuations than adjusted TFP. Another advantage of using adjusted TFP as the target variable instead of labor productivity is that non-technology shocks such as persistent changes in capital taxes and worker composition may also affect labor productivity in the long-run but should leave long-run TFP unaffected (provided that Fernald's aggregate production function assumption and his measures of effective labor and capital are correct).³² This can lead to important differences in results. Indeed, [Francis et al. \(2013\)](#) report that in response to their max-share shock, labor productivity jumps immediately to what is essentially a new permanent level. We confirm this result in VAR estimations reported in the Appendix. This would suggest that technology should be modeled as a random walk process, which is very different from the results below where the response of adjusted TFP to the max-share shock is insignificant for several quarters before gradually increasing to a new permanent level, consistent with the empirical literature cited above that technology is slowly diffusing.

We close this description by briefly commenting on an important issue that has been brought up in the empirical literature on news shocks. [Fernandez-Villaverde et al. \(2007\)](#) document that many VAR systems may be non-invertible in the sense that it is impossible to recover structural

includes non-stationary variables, as we do, or a stationary VAR. In contrast, the infinite-horizon approach of [Gali \(1999\)](#) requires the VAR to be stationary, which implies that the researcher needs to take a stand on various cointegration restrictions that can affect the results in important ways.

³²See [Uhlig \(2004\)](#) or [Bocola et al. \(2014\)](#) for examples when non-technology shocks affect labor productivity in the long-run. Of course, non-technology shocks may affect adjusted TFP (as well as labor productivity) in the long-run if the discovery and adoption of new technologies arises endogenously as a function of the state of the business cycle. In this case, the proposed identification as well as the other existing identifications of technology shocks will confound news shocks with non-technology shocks. This point remains very much an open issue for the literature on technology shocks that we start to address below by examining the response of novel indicators of technological innovation to our extracted shock.

shocks from reduced form forecast errors. [Leeper et al. \(2013\)](#) argue that news shocks are likely to result in non-invertible VAR systems. Subsequent papers such as [Sims \(2012\)](#) and [Beaudry and Portier \(2014\)](#) have examined the consequences of non-invertibility through Monte Carlo simulations of DSGE models with news shocks. Both of these papers conclude that even if a VAR system is technically non-invertible, VAR methods may nevertheless perform quite well in recovering the impulse responses to a news shock. An important conclusion from these papers is that the inclusion of forward-looking variables as is done in particular in the larger VAR systems considered below helps to ameliorate potential invertibility issues.³³ Our Monte Carlo experiments in [Section 6](#) provide further support that VAR methods are adequately suited to identifying news shocks, subject to the caveats about measurement issues and zero impact restrictions highlighted above.

5.2 Results

We first apply the proposed max-share identification to the same four-variable VAR as in [Section 4](#). The horizon at which the FEV share of adjusted TFP is maximized is set to $H = 80$ quarters, although similar results would obtain for shorter horizons above 40 quarters. The estimated impulse responses are reported in [Figure 4](#). As before, the solid black lines show the median impulse responses when the VAR is estimated with the 2016 vintage of adjusted TFP while dash-dotted red lines show the median impulse responses when the 2007 vintage is used.

Although the max-share identification does not impose the zero impact restriction, adjusted TFP does not react significantly on impact of the shock and increases only gradually thereafter, with the estimated long run response of adjusted TFP two to three times larger than the impact response. In other words, the news shock predicts delayed but sustained future TFP growth.

In contrast with the results based on the [Barsky and Sims \(2011\)](#) news identification, there is very little difference in the impulse responses of the macroeconomic variables when estimating the VAR with the 2007 vintage of adjusted TFP instead of the 2016 vintage. In both cases, consumption jumps on impact and then gradually increases further to a permanently higher level; hours worked decline significantly on impact before turning positive after about five quarters; and inflation drops sharply and significantly on impact of the shock before gradually returning towards

³³See [Beaudry and Portier \(2014\)](#) or [Barsky et al. \(2015\)](#) for excellent discussions.

its initial level.

Overall, these results look very much in line with the original results reported in [Barsky and Sims \(2011\)](#). Indeed, as [Table 5](#) shows, the median correlation between consumption growth and hours growth implied by the max-share shock is robustly negative, contrary to what we observe in the data.³⁴ Similarly, consumption growth is strongly negatively correlated with inflation, whereas in the data the two are essentially uncorrelated.

5.3 Does the max-share identification capture news shocks?

The key question with the proposed alternative identification is whether it captures news about future productivity growth and to what extent the resulting shock can be considered as exogenous. To address this question we investigate whether the max-share shock has the key characteristics of what one would expect from a news shock about future productivity; i.e.

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

For the first question, we already know from the results with the four-variable VAR that the max-share shock leads to persistent and therefore predictable changes in future TFP growth. We now extend the analysis by considering an eight-variable VAR system that contains, on top of the four variables already included above, real gross domestic product (GDP) per capita, real private investment expenditures per capita, the real S&P500 index (deflated by the consumer price index) and the Federal Funds rate.³⁵ This choice of variables is motivated by the desire to learn about the effects of the max-share shock for other prominent macroeconomic aggregates and by the idea that including forward-looking information variables may help sharpen the results and address issues of non-fundamentalness discussed above. Indeed, as [Beaudry and Portier \(2006\)](#) argue, there is a large literature suggesting that stock prices reflect expectations about future economic conditions

³⁴Similar to the results reported above, the business cycle moments implied by the max-share shock are estimated very precisely.

³⁵The real S&P500 index is taken directly from Robert Shiller's [website](#). None of the results would change if the index was instead transformed into real terms with another deflator. The other variables are taken from the FRED database of the Federal Reserve Bank of St. Louis.

and should therefore be an important indicator of news. Similarly, the Federal Reserve with its hundreds of Ph.D. economists on staff should have superior forecasting abilities and thus, news might also be reflected in the Federal Funds rate, the main monetary policy instrument up until the recent financial crisis.

As before, the VAR is estimated with four lags for the 1960:1-2007:3 period subject to a Minnesota prior. Figure 5 displays the impulse responses to the max share shock, both when the VAR is estimated with the 2016 vintage of adjusted TFP (solid black lines) and when the VAR is estimated with the 2016 vintage of adjusted TFP (dash-dotted red lines).

As in the four-variable VAR, the estimated responses match closely across the two vintages, confirming the robustness of the max-share approach to the revisions in Fernald's adjusted TFP series. Compared to the four-variable VAR, the reaction of adjusted TFP to the shock is more delayed and gradual, with an impact response that is not significantly different from zero for several quarters (more than 10 quarters in the case of the 2016 vintage).³⁶ This difference in results is primarily due to the inclusion of the real S&P500 index in the VAR, confirming the point of [Beaudry and Portier \(2006\)](#) that stock prices contain valuable information about market expectations of future economic conditions. The real S&P500 index itself reacts strongly on impact and then displays a hump-shaped response before gradually returning towards its pre-shock level. Investment and hours worked both decline initially while output rises slightly and consumption jumps up robustly on impact. Thereafter, output, consumption, and investment gradually increase towards a permanently higher level while hours worked increases in a hump-shaped manner before returning towards its original level before the shock. Inflation and the Federal Funds rate both decline significantly on impact and then remain persistently below their original values. The initial decline of inflation substantially exceeds the decline in the Federal Funds rate, implying that real short interest rates increase on impact of the shock. Hence, the shock triggers a contractionary monetary policy response despite the deflationary effect that the shock has on the economy.³⁷

³⁶It is interesting to compare this result with reaction of labor productivity in the same VAR specification except that adjusted TFP is replaced by labor productivity. The impulse responses are reported in the Appendix. As in [Francis et al. \(2013\)](#), we find that labor productivity jumps significantly on impact of the max-share shock and that the adjustment to the new permanently higher level thereafter is considerably smaller than for adjusted TFP. This difference in results is due primarily to capital deepening: since hours worked decline on impact of the shock, the capital-labor ratio increases, pushing the response of labor productivity above the response of adjusted TFP.

³⁷The initial drop in inflation would be even more pronounced if prices were measured with the CPI deflator or the PCE deflator instead of the GDP deflator.

The opposite-signed impact responses of consumption relative to hours and investment implies that the max-share shock generates negative business cycle comovement between these variables. This confirms the conclusion from the four-variable VAR that the shock is unlikely to be a main driver of business cycle dynamics. This does not mean, however, that the shock is unimportant for macroeconomic fluctuations more generally. Indeed, as Table 6 shows, while the shock accounts for only a small fraction of the FEV of real macroeconomic aggregates at short horizons (with consumption being the notable exception), the shock is the main driver of these variables at longer horizons with the exception of hours worked.³⁸ Indeed, at the 80 quarter horizon, the shock accounts for about three-fourths of unpredictable variations in adjusted TFP, GDP, consumption, and investment. Quite strikingly, the shock also accounts for almost half of unpredicted variations in the real S&P500 index and inflation at forecast horizons of 20 quarters and more, and about one-third of unpredictable variations in the Federal Funds rate at horizons of 40 quarters or more.

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

To answer the second and third question above, we re-estimate the eight-variable VAR with the Federal Funds rate replaced sequentially with different measures of technological innovation and forward-looking information variables. The rest of the VAR specification is kept unchanged except when we have to adapt the sample due to data availability, as described below. To save on space, we only report impulse responses for the variables that replace the Federal Funds rate. The seven other variables in the VAR, which are kept the same throughout the exercise, react very similarly to the max-share shock as reported above in Figure 5.

We first consider four different measures of technological innovation: the index of information and communications technology (ICT) standards by [Baron and Schmidt \(2015\)](#); the index of new technology manuals by [Alexopoulos \(2011\)](#); real R&D expenditures per capita from the NIPAs; and the inverse of the relative price of investment price from [Justiniano et al. \(2010\)](#). The index by [Baron and Schmidt \(2015\)](#) counts the number of new ICT industry standards per quarter

³⁸All of the results in Table 6 refer to median estimates from the VAR estimated with the 2016 vintage of adjusted TFP. The results are very similar for the VAR estimated with the 2007 vintage and are therefore omitted to save on space.

released by standard setting organizations (SSOs) in the U.S.³⁹ As [Baron and Schmidt \(2015\)](#) argue, standardization is an essential step in the introduction and adoption of new technologies. It precedes the implementation of new technologies but presumably provides an important signal about the commercial viability of an innovation and thus future growth opportunities. As such, standardization represents an ideal measure to assess the extent to which our max-share shock captures news. As in [Baron and Schmidt \(2015\)](#), we focus on ICT standards because ICT have constituted the dominant type of general purpose technology, although results are robust to using broader industry standards. [Alexopoulos \(2011\)](#)'s count of books published in the field of technology provides a complementary measure even though she develops her measure primarily to investigate the role of contemporaneous technology shocks.⁴⁰ As explained in her paper, new book titles in this area “... appear precisely when the innovation is first introduced to market, for the very good reason that the whole purpose of publications is to spread the word about the new product or process.” R&D expenditures and the relative investment price are common measures of the quality and/or efficiency of newly produced investment goods. If our max-share shock captures news about future productivity growth, then we would expect both of these measures to react gradually as new technologies are being implemented and start to affect productivity.⁴¹

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

Figure 6 reports the impulse responses to the max-share shock. Both the index of new ICT

³⁹SSOs are mostly private organizations that exist in many industries to establish voluntary and regulatory standards. Prominent examples include the electricity plug, the USB key, the WiFi communications protocol or quality standards (e.g. ISO). The standardization index by [Baron and Spulber \(2015\)](#) and [Baron and Schmidt \(2015\)](#) is based on information from the Searle Center database on technology standards and standard setting organizations. See their papers for details. We thank Justus Baron and Julia Schmidt for making their index available.

⁴⁰As emphasized above, the two are not necessarily distinct as news about future productivity growth may coincide with contemporaneous innovations to productivity. [Alexopoulos \(2011\)](#) also constructs different new book titles for different technology categories, including new titles for computer hardware and software, and telecommunications. The results presented below are robust to using these alternative measures.

⁴¹One could argue that as long as Fernald's TFP series appropriately controls for quality changes in the capital stock, news shocks derived from TFP should be unrelated to capital-embodied technological change. However, it is doubtful that the quality adjustments made to the different capital series that Fernald uses fully capture these quality changes. Moreover, as argued for example by [Chen and Wemy \(2015\)](#), there may be spillovers from capital-embodied technological change to neutral, general-purpose technology.

standards and the index of new technology manuals jump markedly on impact of the shock. The index of new ICT standards then declines back towards its pre-shock level while the new manuals measure remains permanently higher. The response of the ICT standards index is particularly striking and matches closely with the evidence reported in [Baron and Schmidt \(2015\)](#), who use a recursive identification approach based on zero impact restrictions. R&D expenditures and the (inverse of the) relative price of investment goods, in turn, increase only gradually after the shock, although this increase occurs at a considerably faster pace than for adjusted TFP, as reported in [Figure 5](#).

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

Next, we consider three forward-looking information variables that have been interpreted as capturing news: the spread between long-term (5-year) treasury bond yields and the Federal Funds rate as used in [Kurmann and Otrok \(2013\)](#); the Michigan Survey's 5-year ahead consumer confidence index as used in [Barsky and Sims \(2012\)](#); and the business confidence index from the Business Outlook Survey (BOS) conducted by the Federal Reserve Bank of Philadelphia as used in [Bachmann, Elstner, and Sims \(2013\)](#). [Figure 7](#) shows the impulse responses of these series to the max-share shock. For reference, we also include the impulse response of the real S&P500 index, which is part of the VAR used to generate these results. All of the indicators jump up sharply on impact of the news shock and then decline gradually back to their original level. These responses are highly significant and indicate that the identified max-share shock captures news about the future that is picked up not only by financial markets but also the Fed, consumers, and businesses.

We complete this investigation by considering an additional set of forward-looking information variables that, to our knowledge, have not formally been linked to news shocks: the excess stock market return by [Fama and French \(1993\)](#); the return on capital computed by [Gomme, Ravikumar, and Rupert \(2011\)](#); as well as two recent indicators of uncertainty, the forecast dispersion index constructed by [Bachmann, Elstner, and Sims \(2013\)](#), and the 12-month ahead macroeconomic uncertainty index constructed by [Jurado, Ludvigson, and Ng \(2015\)](#). The excess stock market

return is the value-weighted return on all CRSP firms incorporated in the U.S. and listed on the NYSE, AMEX, or NASDAQ minus the one month Treasury bill rate.⁴² The return on capital is computed from the investment first-order condition of a neoclassical business cycle model and is driven by fluctuations in the marginal product of capital and relative prices of investment goods. We view both of these returns as complementary measures of how the market reacts to news about the future value of capital. The uncertainty indicator by [Bachmann et al. \(2013\)](#) is sourced from the same BOS data as the business confidence index above and is constructed as the cross-sectional dispersion of survey participants' assessment of future business conditions. Lastly, the uncertainty indicator by [Jurado et al. \(2015\)](#) is constructed as the weighted mean of conditional volatilities of unforecastable components of future macroeconomic time series, with the forecasting part obtained from a diffusion index based on a large set of predictors. As opposed to other measures of uncertainty proposed in the literature, the two indicators are explicitly designed to capture uncertainty about the future.

Figure 8 shows the impulse responses of these series to the max-share shock. Similar to the news indicators above, both return variables jump up sharply on impact of the shock. This response is short-lived: after two quarters, the variables have essentially returned back to their pre-shock value. These impulse responses confirm that the max-share shock contains important news for capital markets. Interestingly, the two uncertainty indicators show a significant negative reaction on impact of the shock that, in the case of [Jurado et al. \(2015\)](#)'s measure, is quite persistent. In conjunction with the above results, this decline in uncertainty suggests that the news picked up by the max-share shock provides resolution of uncertainty about the productive potential of innovations.

The different results provide in our view compelling evidence that the max-share shock captures news about future productivity growth. The shock predicts delayed sustained future TFP growth, accounting for only a small fraction of TFP fluctuations at short forecast horizons but for 70 percent or more of TFP fluctuations at longer horizons. Perhaps more importantly, the shock is associated with large and persistent jumps in two novel measures of innovation, followed by a hump-shaped increase in R&D expenditures and a gradual decline in the relative price of investment goods; and the shock generates jumps in a wide variety of forward-looking informa-

⁴²The data are regularly updated and available on Kenneth French's [website](#).

tion variables. Taken together, these responses suggest that the max-share identification picks up technological innovation as opposed to other business cycle shocks or noise that endogenously lead to changes in productivity; and that market participants clearly update their forecasts about the economy. The news interpretation therefore seems natural.

6 Monte Carlo Simulations

This section returns to the New Keynesian DSGE model of Section 2 and uses it to further investigate the performance of the proposed max-share identification relative to the Barsky-Sims identification under different scenarios about measurement error and timing of news shocks. We simulate time series for adjusted TFP, consumption, inflation, and hours worked from the model and estimate a four-variable VAR from these series. We then compare the impulse responses to a news shock from the model with the impulse responses from the VAR under the Barsky-Sims and the max-share identifications.

6.1 Specification

The calibration of the different structural parameters is the same as in Section 2. To avoid stochastic singularity, we need at least as many shocks in the model as variables in the VAR. In addition to the unanticipated shock and the news shock from the technology processes in (8) and (9), we consider a shock to the marginal efficiency of investment (MEI) and a shock to the disutility from labor (which is isomorphic to a wage markup shock). The online appendix provides details. The shock magnitudes are parameterized so that the MEI shock accounts for the majority of short run fluctuations, as documented in [Justiniano, Primiceri, and Tambalotti \(2010\)](#).⁴³

Before proceeding, let us pause to note that we have not attempted to specify the mismeasurement of utilization to mimic the observed business cycle properties of Fernald's adjusted TFP series reported in Section 3. Given the flexibility of the model equation for observed utilization in (14), this would be relatively straightforward to accomplish. However, the model abstracts from many of the features that Fernald considers in his construction of the series and so, even if we

⁴³In particular, the MEI shock accounts for 50 percent, the news shock for 25 percent, and the labor supply and the unanticipated technology shocks for 12.5 percent each of the unconditional variance of output growth. The results that follow are similar under alternative parameterizations of the shock processes.

matched the unconditional business cycle properties of utilization, it is unclear what lessons one would derive from such an exercise. Instead, our Monte Carlo simulations simply seek to provide further illustration that mismeasurement of utilization potentially matters in important ways for the empirical identification of news shocks.

6.2 Business cycle moments

We start by reporting select unconditional business cycle moments for true technology and adjusted TFP as measured by an econometrician in the model under the four different scenarios considered in Section 2. Table 7 reports the results. The first column shows the unconditional correlations between true technology growth and adjusted TFP growth. The second and third columns show the correlations between true technology growth and hours growth and adjusted TFP growth and hours growth, both conditional on news shocks only (i.e. all other shocks are set to zero). These conditional correlations are what matters for the news shock identification.

Focusing on the first column, one observes that true technology growth and measured adjusted TFP growth are nearly perfectly collinear so long as utilization is measured accurately. This is true independent of whether the news shock impacts true technology immediately or with a lag and suggests, as noted in Section 2, that the incongruities between the model and Fernald’s assumptions concerning constant returns to scale and time-varying markups do not seem to matter quantitatively. The presence of either idiosyncratic or systematic measurement error in utilization, in contrast, lowers the correlation between true technology growth and adjusted TFP growth somewhat.⁴⁴

Moving to the second column, the contemporaneous correlation of true technology growth with hours growth conditional on news shocks is close to zero for the baseline scenario as well as for the two scenarios with measurement error (these correlations are exactly the same because measurement error only affects adjusted TFP but not true technology). When news shocks have a contemporaneous impact on technology, the correlation of true technology growth with hours growth becomes positive.

⁴⁴The extent of this reduction in correlation is modest, which obviously depends on the importance of the measurement error. One can make this correlation arbitrarily small by sufficiently increasing the variance of the measurement error.

As shown in the third column, finally, the correlations of adjusted TFP growth with hours growth conditional on news shocks are similar to the correlations for true technology growth reported in the second column, with the notable exception of the scenario with systematic measurement error for which the correlation becomes strongly positive. These results suggest, as we will see in what follows, that idiosyncratic measurement error ought not to affect identification of news shocks, since identification is conditional on the realization of a shock which induces movements in variables other than adjusted TFP (although idiosyncratic measurement error may make identification more difficult in finite samples). By contrast, systematic measurement error can result in significant biases in the empirical identification of news shocks since it has such a large impact on the comovement of adjusted TFP growth with news-driven fluctuations in spite of the fact that the reduction in the unconditional correlation between adjusted TFP and true technology in this specification is rather modest.

6.3 Impulse responses to a news shock

For our Monte-Carlo exercises, we simulate 100,000 periods of data of adjusted TFP, consumption, inflation, and hours worked from the model and estimate a VAR with twenty lags from these series. We then identify a news shock using the Barsky-Sims identification (labelled “BS”) and the proposed max-share identification (labelled “MS”). The point of basing our simulations on such a long sample is that we want to examine the asymptotic properties of the two identification procedures. The high VAR lag order, in turn, is intended to diminish the role of “lag-truncation bias” that arises when estimating data generated from linearized DSGE models with a finite-order VAR (although our results are qualitatively similar with a much smaller number of lags). Of course, one could also investigate the small-sample properties of our estimates [Francis, Owyang, Roush, and DiCecio \(2013\)](#) show that these properties are quite good for the max-share identification.

First, we consider the baseline scenario in which news shocks impact true technology with a lag and utilization is measured without error. [Figure 9](#) reports the results. Here and below, the solid lines show the impulse responses to a news shock in the model; the dotted red lines the responses implied by the BS identification; and the dashed blue lines the responses implied by the MS identification. In the model, consumption increases on impact, inflation falls, and hours worked increases in a hump-shaped manner. Though there are minor differences, both the BS and

MS identifications do a good job of recovering these impulse responses.

Second, we consider the scenario with idiosyncratic measurement error while keeping with the assumption that news shocks affect true technology with a delay. As Figure 10 reports, the different impulse responses are virtually identical to those shown in Figure 9. Once again, there is little difference between the performance of the BS and MS identifications – both do a good job at recovering the dynamic responses to a news shock in the model. This underscores the point made above that idiosyncratic measurement error ought not to affect identification in sufficiently large samples (such as the Monte Carlo exercise considered here). We should note that if idiosyncratic measurement error is sufficiently large, or sufficiently non-transient, the properties of the MS identification can be negatively affected relative to the BS identification, since some of what the MS identification will pick up is purely noise.

In both of these scenarios, the BS restriction that the news shock has no impact on adjusted TFP is satisfied in that true technology in the model only reacts to the news shock with a lag. Even though adjusted TFP is not identical to true technology because of non-constant factor shares and a fixed production cost, as documented above these differences are minor. In both cases, the BS and MS identifications perform similarly, which illustrates that one does not seem to lose accuracy by not imposing the zero impact restriction even if the restriction is approximately satisfied in the model.

Third, we consider the scenario with systematic measurement error. In this case, the zero impact restriction of the BS identification is satisfied with respect to true technology, but it is not consistent with the behavior of adjusted TFP due to the systematic mis-measurement of utilization. Figure 11 reports the results. The MS identification now clearly outperforms the BS identification. Although the MS procedure over-estimates the impact response of true technology on impact (since true technology only reacts with a lag), it closely matches the dynamic paths of adjusted TFP, consumption, inflation, and hours worked. The BS identification, in contrast, implies impulse responses for adjusted TFP and consumption that remain below their model counterparts, even at long horizons, produces the wrong sign for the impact response of inflation, and substantially over-estimates the impact response of hours.

Fourth and finally, we revert to assuming that utilization is measured without error but assume that the news shock impacts true technology in the model immediately rather than with a lag.

As Figure 12 shows, the MS identification again outperforms the BS identification. The BS identification generates the wrong sign of the impact response of hours, significantly underestimates the response of inflation, and underestimates the response of true technology and consumption at all horizons. The MS identification, in contrast, does a good job recovering the dynamic responses to the news shock.

In both scenarios (iii) and (iv), the BS assumption of impact orthogonality between the news shock and adjusted TFP is violated in the data generating process. This is what drives its poor relative performance relative to our MS identification. In scenarios (i) and (ii), where the BS assumption of impact orthogonality is consistent with the data generating process, the BS and MS identification perform about as well as one another. We conclude from these Monte Carlo exercises that there is little to gain from imposing impact orthogonality in the identification of the news shock, while the costs of doing so can be large.

7 Conclusion

An almost universally imposed assumption in the empirical literature is that news shocks impact 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 document large changes in the cyclical properties of the popular series of adjusted TFP produced by Fernald (2014). These changes materially affect empirical conclusions about the macroeconomic effects of news shock as identified by Barsky and Sims (2011). We therefore propose an alternative identification, based on the max-share approach by Francis et al. (2013), which does not impose the zero impact restriction. We show that our identification is robust to the revisions in Fernald's series, and performs well in Monte Carlo simulations under different assumptions about measurement of productivity and the data-generating process for technology. When applied to U.S. data, we find results that are consistent with a news interpretation: adjusted TFP increases gradually and with a significant lag whereas indicators of technological innovation and forward-looking information variables jump on impact. At the same time, the identified shock does not generate comovement in real macroeconomic aggregates and is therefore not a

main driver of business cycle fluctuations. This does not imply that the shock is unimportant for macroeconomics. The shock accounts for the majority of unpredictable fluctuations in real aggregates at medium- and long horizons and generates strong impact responses of inflation, the Federal Funds rate, asset prices as well as different measures of uncertainty. Investigating these results further and assessing the type of models that are consistent with these dynamics are important topics of future research.

References

- Aaronson, D. and D. Sullivan (2001). Growth in worker quality. *Economic Perspectives* 25(4), 53–74. [2.1](#)
- Acemoglu, D. (2009). *Introduction to Modern Economic Growth*. Princeton: Princeton University Press. [2.1](#)
- Alexopoulos, M. (2011, June). Read All about It!! What Happens Following a Technology Shock? *American Economic Review* 101(4), 1144–79. [1](#), [5.3](#), [40](#)
- Bachmann, R., S. Elstner, and E. R. Sims (2013, April). Uncertainty and Economic Activity: Evidence from Business Survey Data. *American Economic Journal: Macroeconomics* 5(2), 217–49. [1](#), [5.3](#)
- Baron, J. and J. Schmidt (2015). Technological Standardization, Endogenous Productivity and Transitory Dynamics. Working papers 503, Banque de France. [1](#), [5.3](#), [39](#)
- Baron, J. and D. F. Spulber (2015). Technology Standards and Standards Organizations: Introduction to the Searle Center Database. Technical report, Northwestern University Pritzker School of Law. [39](#)
- Barsky, R., S. Basu, and K. Lee (2015). Whither news shocks? In *NBER Macroeconomics Annual 2014, Volume 29*, NBER Chapters, pp. 225–264. National Bureau of Economic Research, Inc. [1](#), [8](#), [14](#), [4.2](#), [33](#)
- Barsky, R. and E. Sims (2011). News shocks and business cycles. *Journal of Monetary Economics* 58(3), 273–289. [1](#), [9](#), [4](#), [4.1](#), [4.1](#), [4.2](#), [28](#), [5.1](#), [5.2](#), [7](#)
- Barsky, R. B. and E. R. Sims (2012, June). Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence. *American Economic Review* 102(4), 1343–77. [5.3](#)
- Basu, S., J. Fernald, J. Fisher, and M. Kimball (2013). Sector-specific technical change. Technical report. [3.2](#), [24](#)

- Basu, S., J. G. Fernald, and M. S. Kimball (2006). Are technology improvements contractionary? *American Economic Review* 96(5), 1418–1448. [4](#), [2](#), [2.1](#), [19](#), [3.2](#)
- Beaudry, P. and B. Lucke (2010). Letting difference views about business cycles compete. In *NBER Macroeconomics Annual 209, Volume 23*, NBER Chapters, pp. 413–455. National Bureau of Economic Research, Inc. [25](#)
- Beaudry, P. and F. Portier (2006). Stock prices, news, and economic fluctuations. *American Economic Review* 96(4), 1293–1307. [1](#), [2.2](#), [25](#), [5.1](#), [5.3](#)
- Beaudry, P. and F. Portier (2014). News-driven business cycles: Insights and challenges. *Journal of Economic Literature* 52, 993–1074. [1](#), [5.1](#), [33](#)
- Bocola, L., I. Manovskii, and M. Hagedorn (2014). Identifying Neutral Technology Shocks. Technical report, Department of Economics, University of Pennsylvania. [32](#)
- Burnside, C., M. Eichenbaum, and S. Rebelo (1993). Labor hoarding and the business cycle. *Journal of Political Economy* 101(2), 245–273. [11](#)
- Cascaldi-Garcia, D. (2017). News shocks and the slope of the term structure of interest rates: Comment. *American Economic Review* 107(10), 3243–49. [9](#)
- Chari, V., P. J. Kehoe, and E. R. McGrattan (2008, November). Are structural VARs with long-run restrictions useful in developing business cycle theory? *Journal of Monetary Economics* 55(8), 1337–1352. [5](#), [31](#)
- Chen, K. and E. Wemy (2015). Investment-specific technology shocks: The source of long-run tfp fluctuations. *European Economic Review* 80, 230–252. [6](#), [41](#)
- 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. [2.2](#)
- Christiano, L. J., M. Eichenbaum, and R. Vigfusson (2004). The response of hours to a technology shock: Evidence based on direct measures of technology. *Journal of the European Economic Association* 2, 381–395. [1](#), [4](#), [6](#)

- Christiano, L. J., M. Eichenbaum, and R. Vigfusson (2006). Assessing structural VARs. *NBER Macroeconomics Annual* 21. 5, 31
- Comin, D. and M. Gertler (2006). Medium-term business cycles. *American Economic Review* 96(3), 523–551. 10, 5.1
- Comin, D., M. Gertler, and A. M. Santacreu (2009). Technology innovation and diffusion as sources of output and asset price fluctuations. Working Paper 15029, National Bureau of Economic Research. 10
- Comin, D. and B. Hobijn (2010). An exploration of technology diffusion. *American Economic Review* 100(5), 2031–2059. 10
- Erceg, C. J., L. Guerrieri, and C. Gust (2005). Can Long-Run Restrictions Identify Technology Shocks? *Journal of the European Economic Association* 3(6), 1237–1278. 5, 31
- Fama, E. and K. French (1993). Common risk factors in the returns on stocks and bonds. *Journal of Financial Economics* 33(1), 3–56. 5.3
- Fernald, J. (2014). A quarterly, utilization-adjusted series on total factor productivity. Federal Reserve Bank of San Francisco Working Paper Series 2012-19. (document), 1, 2, 2.1, 2.1, 2.2, 7
- Fernandez-Villaverde, J., J. F. R. Ramirez, T. J. Sargent, and M. W. Watson (2007). Abcs (and ds) of understanding vars. *American Economic Review* 97(3), 1021–1026. 5.1
- 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. 25
- Francis, N., M. Owyang, J. Roush, and R. DiCecio (2013). A flexible finite-horizon alternative to long-run restriction with an application to technology shocks. *Review of Economics and Statistics* 96(4), 638–648. 1, 5, 5.1, 5.1, 31, 36, 6.3, 7
- Gali, J. (1999). Technology, employment, and the business cycle: Do technology shocks explain aggregate fluctuations? *American Economic Review* 89(1), 249–271. 1, 5, 31

- Gomme, P., B. Ravikumar, and P. Rupert (2011, April). The Return to Capital and the Business Cycle. *Review of Economic Dynamics* 14(2), 262–278. [5.3](#)
- Gort, M. and S. Klepper (1982). Time paths in the diffusion of product innovations. *Economic Journal* 92(367), 630–653. [3](#), [5.1](#)
- Griliches, Z. (1957). Hybrid corn: An exploration in the economics of technological changes. *Econometrica* 25(4), 501–522. [3](#), [5.1](#)
- Jurado, K., S. Ludvigson, and S. Ng (2015). Measuring uncertainty. *American Economic Review* 105(3), 1177–1216. [1](#), [5.3](#)
- Justiniano, A., G. Primiceri, and A. Tambalotti (2010). Investment shocks and business cycles. *Journal of Monetary Economics* 57(2), 132–145. [2.2](#), [2.2](#), [5.3](#), [6.1](#)
- Kurmann, A. and E. Mertens (2014). Stock prices, news, and economic fluctuations: Comment. *American Economic Review* 104(4), 1439–1445. [25](#)
- Kurmann, A. and C. Otrok (2013). News shocks and the slope of the term structure of interest rates. *American Economic Review* 103(6), 2612–2632. [9](#), [5.3](#)
- Kurmann, A. and C. Otrok (2017). News shocks and the slope of the term structure of interest rates: Reply. *American Economic Review* 107(10), 3250–56. [9](#)
- Kydland, F. E. and E. C. Prescott (1982). Time to build and aggregate fluctuations. *Econometrica* 50(6), 1345–1370. [2](#)
- Leeper, E. M., T. B. Walker, and S.-C. S. Yang (2013). Fiscal foresight and information flows. *Econometrica* 81(3), 1115–1145. [5.1](#)
- Lindé, J. (2009). The effects of permanent technology shocks on hours: Can the rbc-model fit the var evidence? *Journal of Economic Dynamics and Control* 33, 597–613. [1](#), [5.1](#)
- Long, J. B. and C. Plosser (1983). Real business cycles. *Journal of Political Economy* 91(1), 39–69. [2](#)

- Mansfield, E. (1961). Technical change and the rate of imitation. *Econometrica* 29(4), 741–766. [3](#), [5.1](#)
- Mansfield, E. (1989). The diffusion of industrial robots in japan and the united states. *Research Policy* 18(4), 183–192. [3](#), [5.1](#)
- Pigou, A. C. (1927). *Industrial Fluctuations*. London: Macmillan. [1](#)
- Rogers, E. (1995). *Diffusion of Innovations*. New York: Free Press. [3](#), [5.1](#)
- Rotemberg, J. (2003). Stochastic technical progress, smooth trends, and nearly distinct business cycles. *American Economic Review* 93(5), 1543–1559. [3](#), [1](#), [5.1](#)
- Sims, E. (2012). News, non-invertibility, and structural vars. *Advances in Econometrics* 28, 81–136. [5.1](#)
- 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. [4](#)
- Smets, F. and R. Wouters (2007). Shocks and frictions in US business cycles: A bayesian DSGE approach. *American Economic Review* 97(3), 586–606. [2.2](#), [19](#)
- Summers, L. H. (1986). Some skeptical observations on real business cycle theory. *Federal Reserve Bank of Minneapolis Quarterly Review* Fall, 23–27. [11](#)
- Uhlig, H. (2003). What moves real GNP? Working paper, Humboldt University. [1](#), [5.1](#)
- Uhlig, H. (2004). Do technology shocks lead to a fall in total hours worked? *Journal of the European Economic Association* 2(2-3), 361–371. [1](#), [32](#)

Tables and Figures

Table 1: Moments of Adjusted TFP Growth for Different Vintages

	$\Delta \ln TFP_t^{U,07}$	$\Delta \ln TFP_t^{U,13}$	$\Delta \ln TFP_t^{U,14}$	$\Delta \ln TFP_t^{U,16}$
Mean	1.49	1.41	1.42	1.42
Standard Deviation	3.41	3.30	3.79	3.46
Corr w/ $\Delta \ln TFP_t^{U,07}$	1.00	0.85	0.56	0.58
Corr w/ $\Delta \ln Y_t$	0.53	0.38	0.18	0.07
Corr w/ $\Delta \ln C_t$	0.26	0.21	0.05	0.11
Corr w/ $\Delta \ln I_t$	0.20	0.13	-0.00	-0.02
Corr w/ $\Delta \ln L_t$	-0.01	-0.06	-0.24	-0.35

Notes: This table shows descriptive statistics for different vintages of adjusted TFP growth and correlations with growth rates of prominent macroeconomic aggregates. $\Delta \ln TFP_t^{U,j}$ is the log first difference of Fernald's adjusted TFP series for vintages $j = 07, 13, 14$ or 16 ; Y_t is headline real GDP; C_t is real personal consumption expenditures; I_t is real private fixed investment; and L_t is total hours worked in the non-farm business sector. Both personal consumption expenditures and private fixed investment are deflated by their own deflators. All macroeconomic aggregates are from the NIPA tables and are log first differenced. The sample period for each of the statistics is 1947q3-2007q3.

Table 2: Moments of Non-adjusted TFP and Utilization Growth for Different Vintages

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

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

Table 3: Moments of the Synthetic Measures of Adjusted TFP Growth for Different Vintages

	$\Delta \ln TFP_t^{U,07,07}$	$\Delta \ln TFP_t^{U,13,07}$	$\Delta \ln TFP_t^{U,14,07}$	$\Delta \ln TFP_t^{U,16,07}$
Mean	1.49	1.45	1.44	1.46
Standard Deviation	3.41	3.16	3.16	3.16
Corr w/ $\Delta \ln TFP_t^{U,07,07}$	1.00	0.91	0.91	0.91
Corr w/ $\Delta \ln TFP_t^{U,13,07}$		1.00	0.99	0.99
Corr w/ $\Delta \ln TFP_t^{U,14,07}$			1.00	0.99
Corr w/ $\Delta \ln TFP_t^{U,16,07}$				1.00

Notes: This table shows descriptive statistics for synthetic measures of adjusted TFP for the 2007, 2013, 2014, and 2016 vintages, constructed by combining the 2007 vintage of utilization and combining it with non-adjusted TFP from vintage $j = 2007, 2013, 2014,$ or 2015.

Table 4: Unconditional Business Cycle Moments Implied by Barsky-Sims Shock

	$sdev(\Delta \ln C_t)$	$sdev(\Delta \ln H_t)$	$sdev(\pi_t)$	$corr(\Delta \ln C_t, \Delta \ln H_t)$	$corr(\Delta \ln C_t, \pi_t)$
Data	0.68	0.79	1.13	0.42	-0.09
Barsky-Sims shock (TFP 2007)	0.36	0.30	0.39	-0.56	-0.73
Barsky-Sims shock (TFP 2016)	0.44	0.18	0.27	0.27	-0.75

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

Table 5: Unconditional Business Cycle Moments Implied by Max-Share Shock

	$sdev(\Delta \ln C_t)$	$sdev(\Delta \ln H_t)$	$sdev(\pi_t)$	$corr(\Delta \ln C_t, \Delta \ln H_t)$	$corr(\Delta \ln C_t, \pi_t)$
Data	0.68	0.79	1.13	0.42	-0.09
Max-share shock (TFP 2007)	0.43	0.25	0.60	-0.23	-0.73
Max-share shock (TFP 2016)	0.39	0.26	0.53	-0.34	-0.71

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

Table 6: Fraction of FEV Explained by Max-Share Shock

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

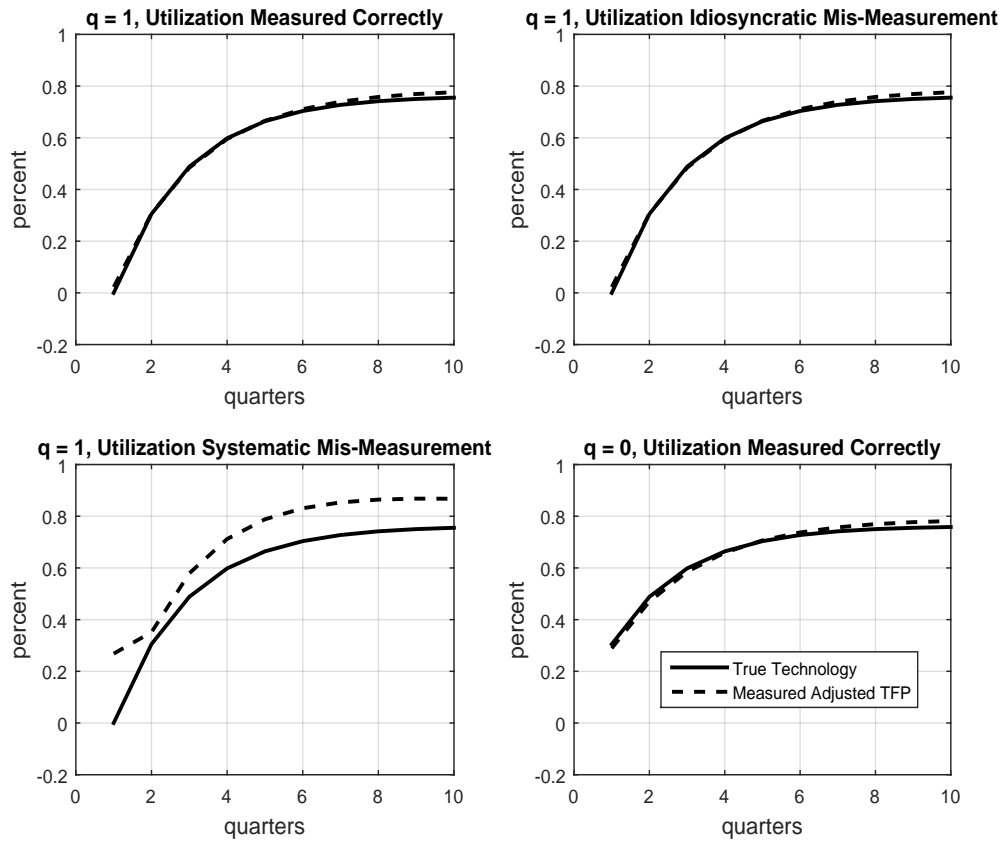
Notes: The sample period for each of the statistics is 1960q1-2007q3. The model statistics pertain to medians from the posterior distribution of each data series implied by the max-share shock. All results are rounded to two digits after the decimal point.

Table 7: Model-implied business cycle correlations

	$corr(\Delta A_t, \Delta TFP_t^{M,U})$	$corr(\Delta A_t, \Delta H_t news)$	$corr(\Delta TFP_t^{M,U}, \Delta H_t news)$
Baseline ($\mathbf{D} = \mathbf{O}, \sigma_u = 0, q = 1$)	0.99	-0.08	-0.01
Idiosyncratic measurement error ($\mathbf{D} = \mathbf{O}, \sigma_u = 0.01, q = 1$)	0.85	-0.08	-0.01
Systematic measurement error ($\mathbf{D} \neq \mathbf{O}, \sigma_u = 0, q = 1$)	0.89	-0.08	0.66
Contemporaneous news impact ($\mathbf{D} = \mathbf{O}, \sigma_u = 0, q = 0$)	0.99	0.24	0.30

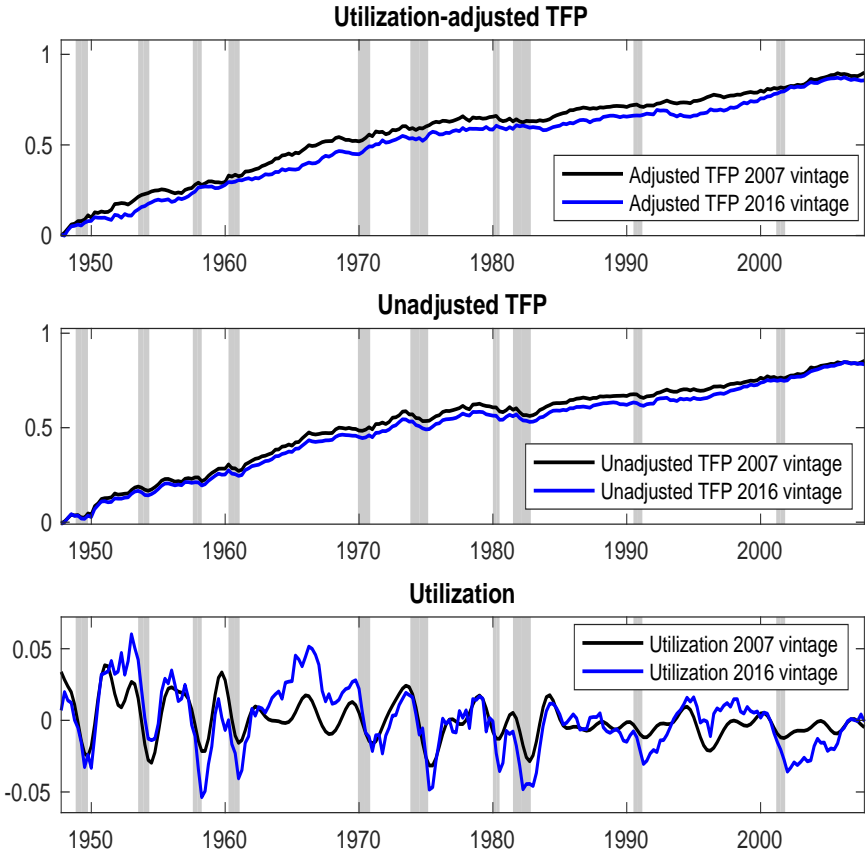
Notes: This table shows business correlations of different variables implied by the solution of the medium scale DSGE model laid out in the text. $\Delta \ln A_t$ refers to the log first difference of true technology; $\Delta \ln TFP_t^{M,U}$ refers to the log first difference of utilization-adjusted TFP as measured in the model; and ΔH_t refers to the log first difference of hours worked. The first column shows the unconditional correlation of $\Delta \ln A_t$ and $\Delta \ln TFP_t^{M,U}$ while the second and third column show the correlation of $\Delta \ln A_t$, respectively $\Delta \ln TFP_t^{M,U}$ with ΔH_t conditional on news shocks only.

Figure 1: Model Responses of True Technology and Adjusted TFP to a News Shock



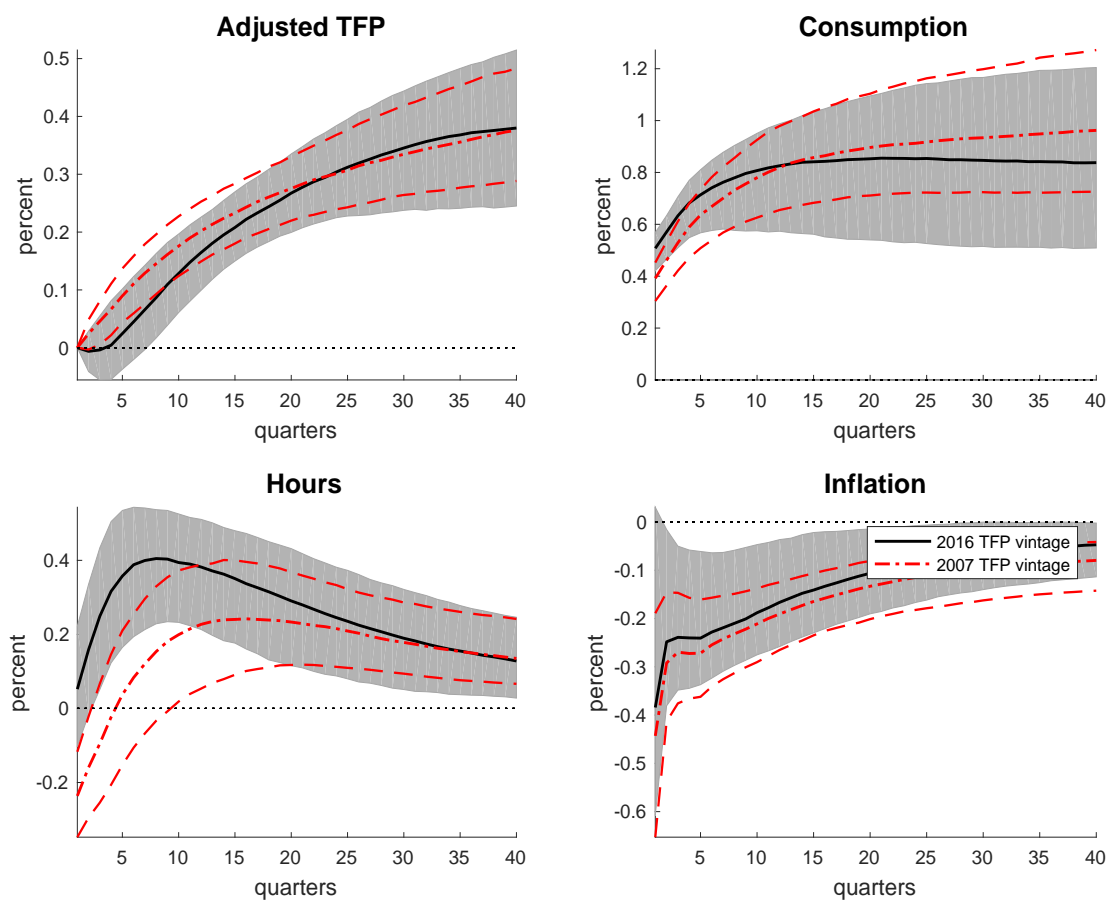
Notes: This figure plots model impulse responses of true technology (solid lines) and adjusted TFP (dashed lines) to a news shock for the four scenarios described in the text.

Figure 2: Adjusted TFP, Non-adjusted TFP and Utilization: 2007 vs. 2016 Vintages



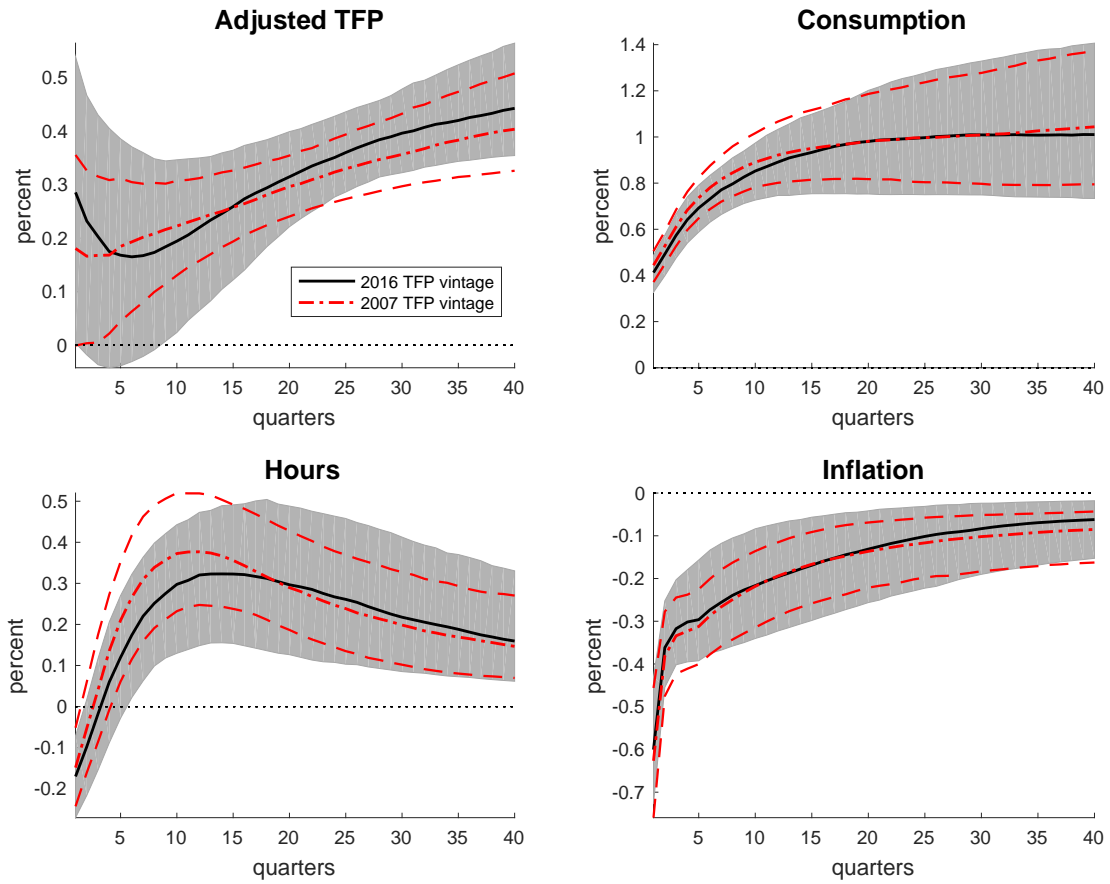
Notes: This figure plots the log levels of both the 2007 and 2016 vintages of the utilization-adjusted TFP series (top panel), the unadjusted TFP series (middle panel) and the utilization series (bottom panel). The 2007 vintages are depicted as black lines. The 2016 vintages are depicted as blue lines.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

Figure 3: Empirical Responses, BS Identification, 2007 vs. 2016 Vintage of Adjusted TFP



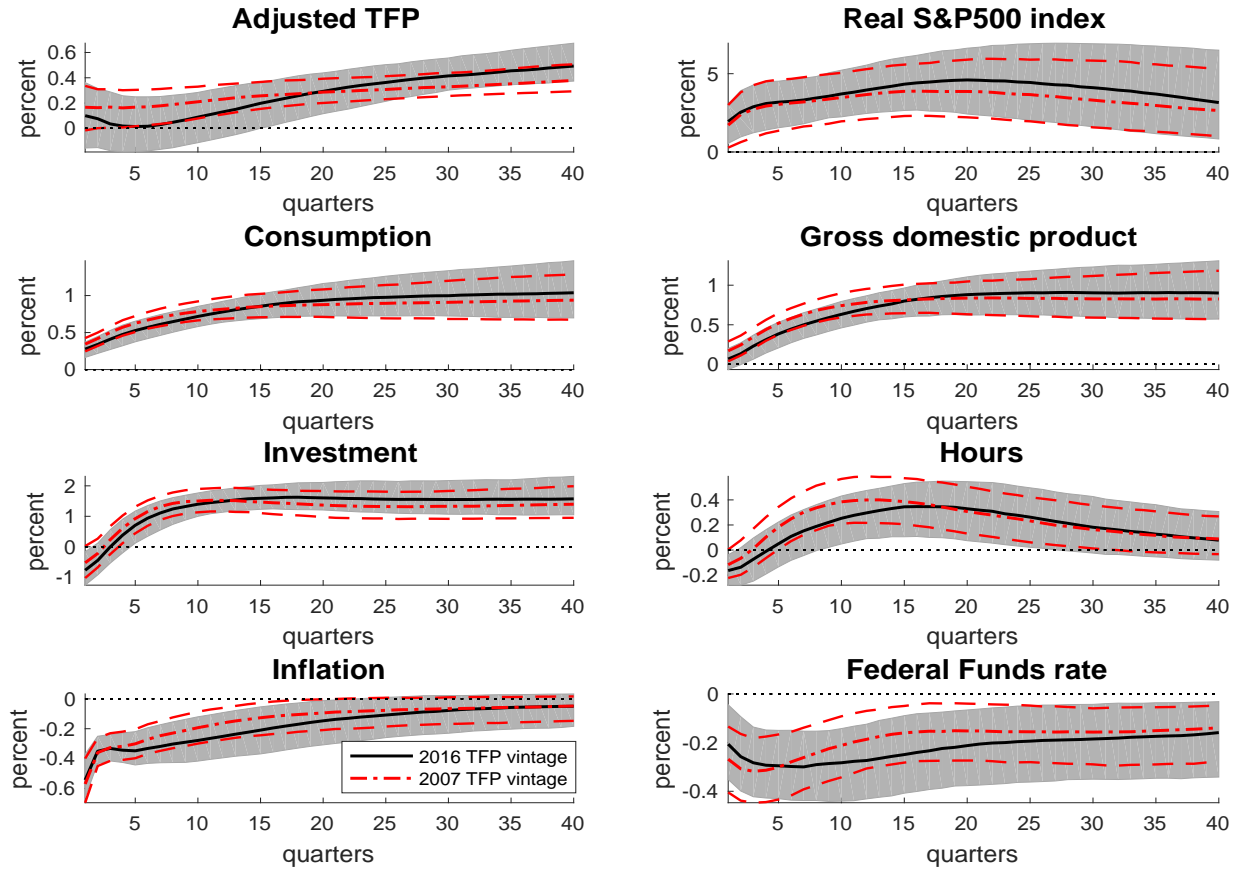
Notes: Solid black lines are the posterior median estimates from the VAR system estimated with the 2016 vintage of adjusted TFP. The gray bands correspond to the 16 to 84 percent posterior coverage intervals. The red dash-dotted lines are the posterior median estimates for the system estimated with the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

Figure 4: Impulse Responses of Four-Variable VAR to Max-Share Shock



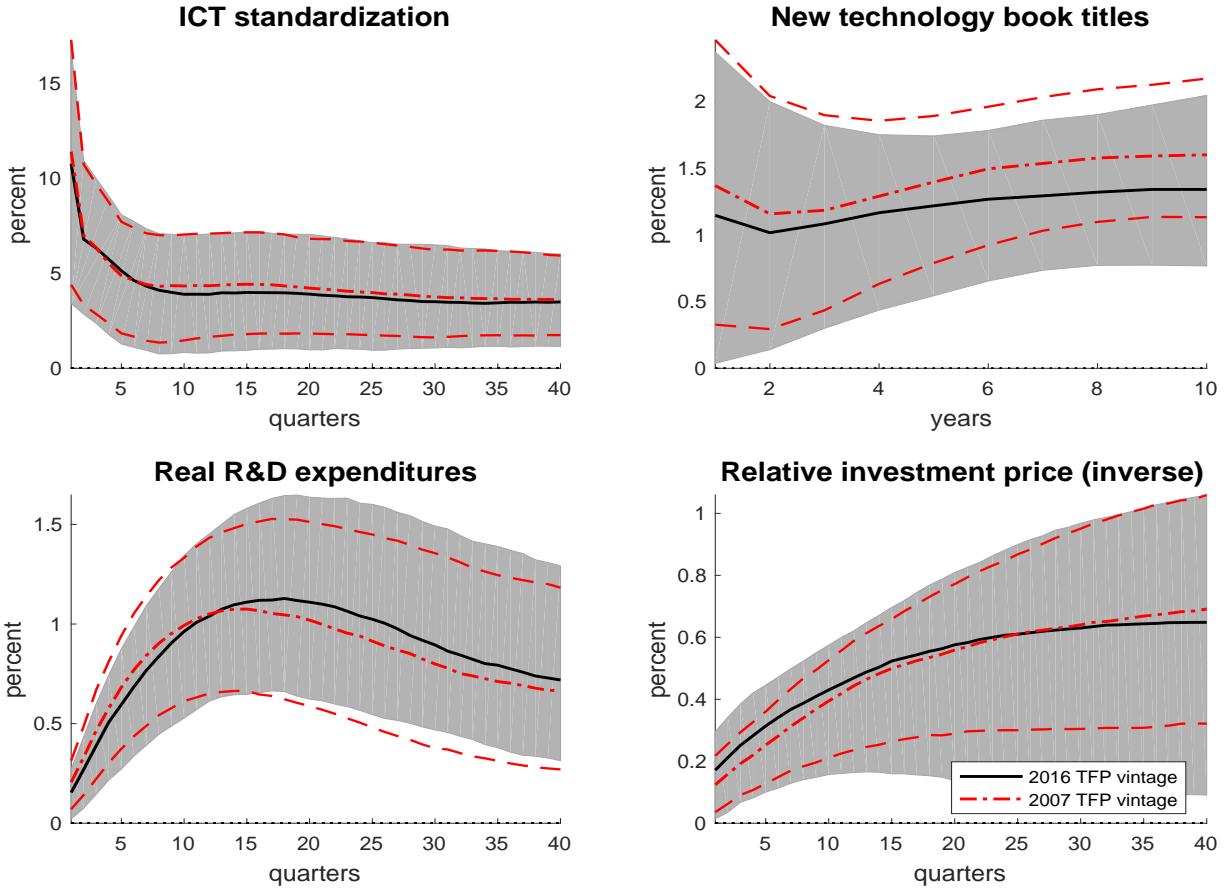
Notes: Solid black lines are the posterior median estimates from the VAR system estimated with the 2016 vintage of adjusted TFP. The grey bands correspond to the 16 to 84 percent posterior coverage intervals. The red dash-dotted lines are the posterior median estimates for the system estimated with the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

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



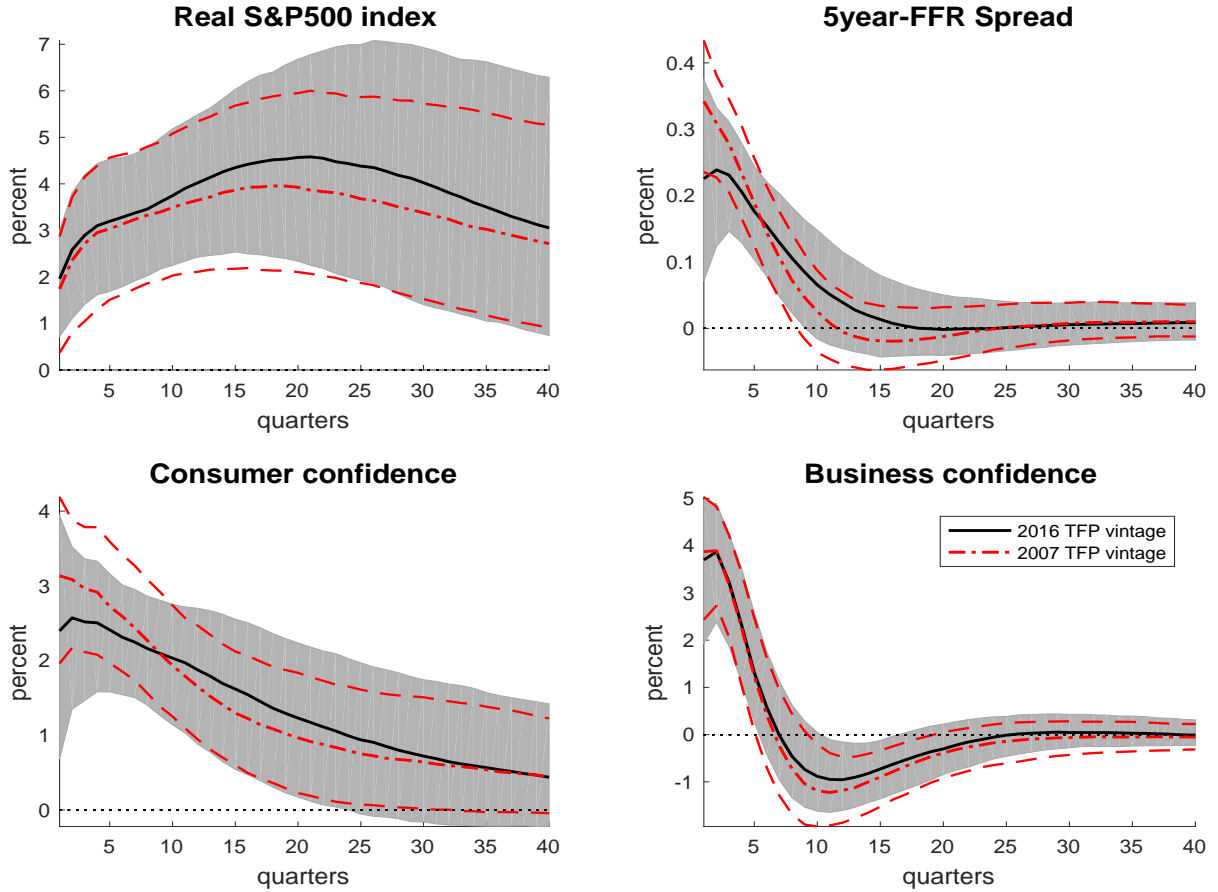
Notes: Solid black lines are the posterior median estimates from the VAR system estimated with the 2016 vintage of adjusted TFP. The grey bands correspond to the 16 to 84 percent posterior coverage intervals. The red dash-dotted lines are the posterior median estimates for the system estimated with the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

Figure 6: Impulse Responses of Innovation Measures to Max-Share Shock



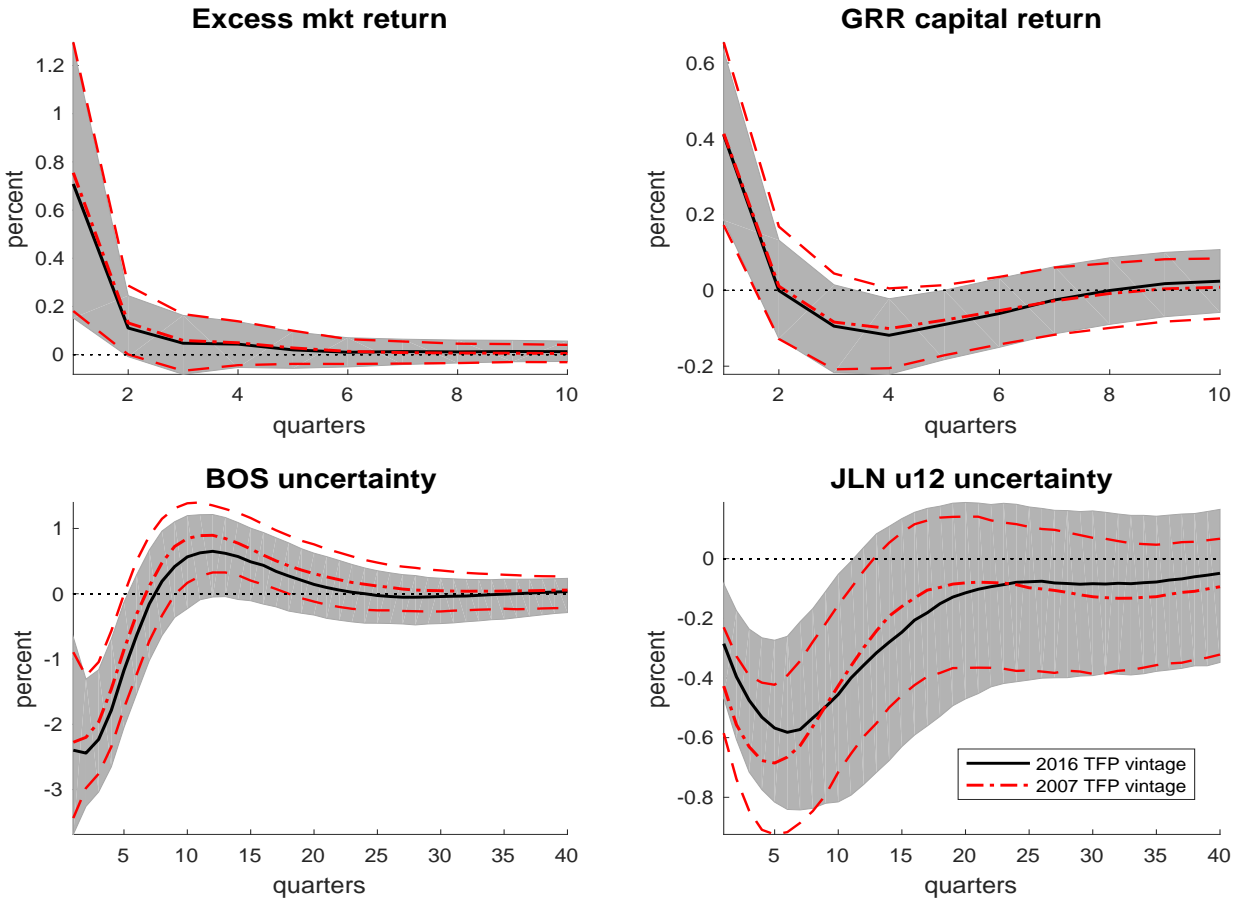
Notes: Solid black lines are the posterior median estimates from the VAR system estimated with the 2016 vintage of adjusted TFP. The gray bands correspond to the 16 to 84 percent posterior coverage intervals. The red dash-dotted lines are the posterior median estimates for the system estimated with the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

Figure 7: Impulse Responses of News Indicators to Max-Share Shock



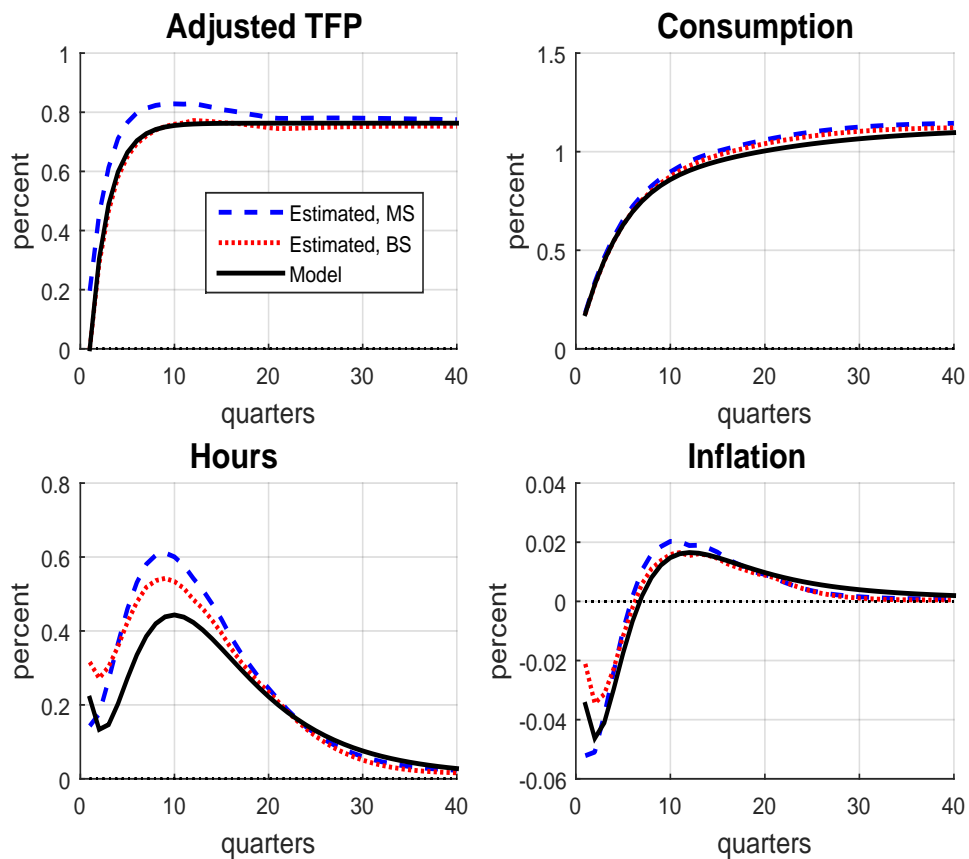
Notes: Solid black lines are the posterior median estimates from the VAR system estimated with the 2016 vintage of adjusted TFP. The gray bands correspond to the 16 to 84 percent posterior coverage intervals. The red dash-dotted lines are the posterior median estimates for the system estimated with the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

Figure 8: Impulse Responses of Asset Returns and Uncertainty Measures to Max-Share Shock



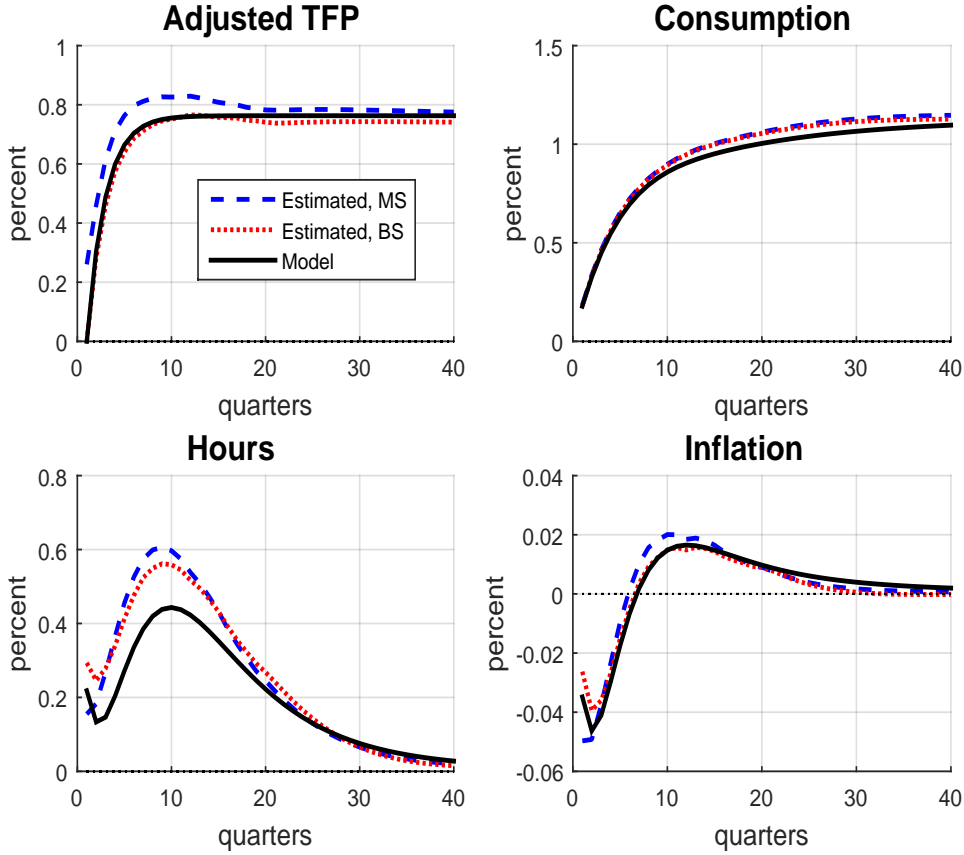
Notes: Solid black lines are the posterior median estimates from the VAR system estimated with the 2016 vintage of adjusted TFP. The gray bands correspond to the 16 to 84 percent posterior coverage intervals. The red dash-dotted lines are the posterior median estimates for the system estimated with the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

Figure 9: Model and Estimated Responses to a News Shock; No Measurement Error and $q = 1$



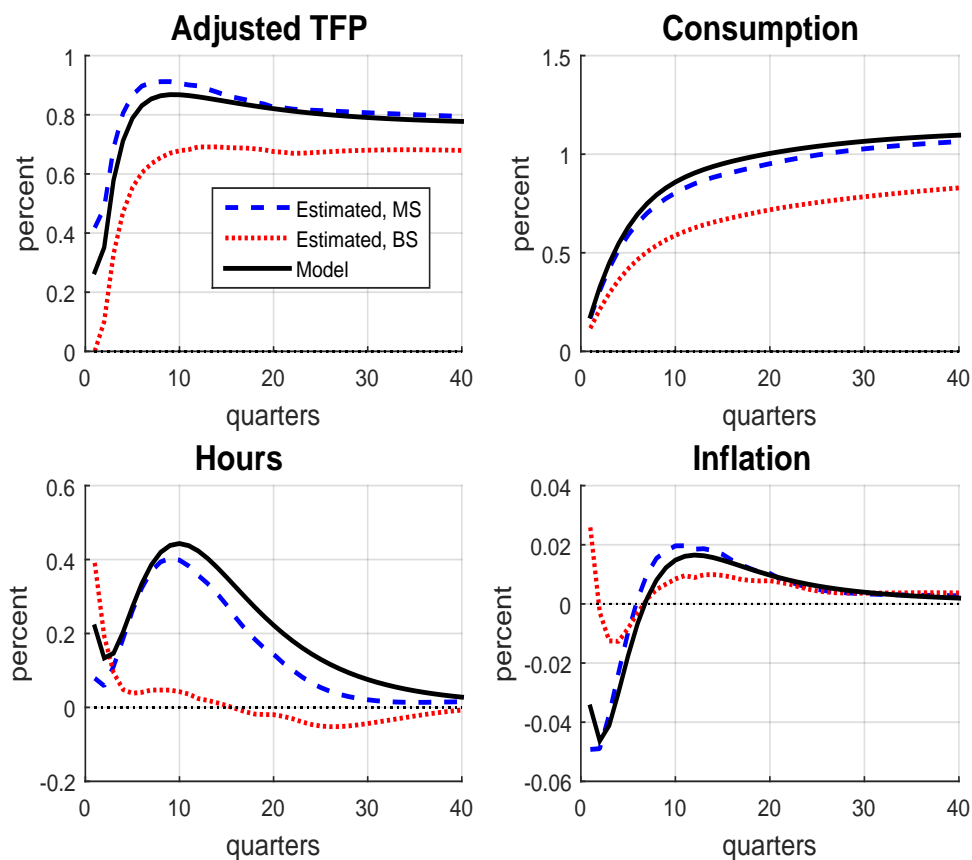
Notes: Solid lines are the true impulse responses to an estimated news shock in the model. The dashed black lines are the estimated responses using the max share identification at a 80 quarter horizon, without imposing impact orthogonality. This is labeled “MS” for “Max Share.” The dashed blue lines are the estimated responses using the BS identification with a 40 quarter truncation horizon, imposing impact orthogonality between the identified shock and adjusted TFP. This is labeled “BS”. In the model generating the data, we assume that $q = 1$, so that there is a one period lag between when agents observe news and when it impacts true productivity in the model.

Figure 10: Model and Estimated Responses to a News Shock; Idiosyncratic Measurement Error and $q = 1$



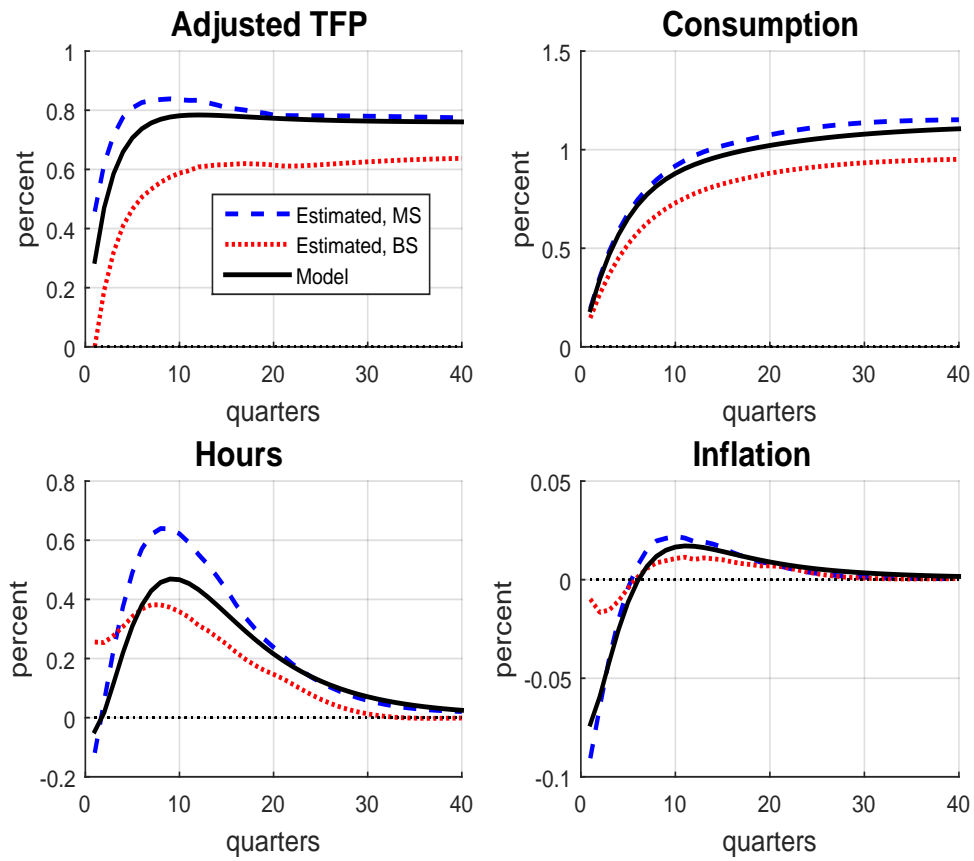
Notes: Solid lines are the true impulse responses to an estimated news shock in the model. The dashed black lines are the estimated responses using the max share identification at a 80 quarter horizon, without imposing impact orthogonality. This is labeled “MS” for “Max Share.” The dashed blue lines are the estimated responses using the BS identification with a 40 quarter truncation horizon, imposing impact orthogonality between the identified shock and adjusted TFP. This is labeled “BS”. In the model generating the data, we assume that $q = 1$, so that there is a one period lag between when agents observe news and when it impacts true productivity in the model.

Figure 11: Model and Estimated Responses to a News Shock; Systematic Measurement Error and $q = 1$



Notes: Solid lines are the true impulse responses to an estimated news 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 and adjusted TFP. This is labeled “BS”. In the model generating the data, we assume that $q = 1$, so that there is a one period lag between when agents observe news and when it impacts true productivity in the model.

Figure 12: Model and Estimated Responses to a News Shock; No Measurement Error and $q = 0$



Notes: Solid lines are the true impulse responses to an estimated news 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 and adjusted TFP. This is labeled “BS”. In the model generating the data, we assume that $q = 0$, so that there is a one period lag between when agents observe news and when it impacts true productivity in the model.