News Shocks, Business Cycles, and the Disinflation Puzzle*

Hafedh Bouakez† Laurent Kemoe‡

March 2022

Abstract

We argue that key findings of the recent empirical literature on the effects of news about future technology — including their tendency to generate negative comovement of macroeconomic aggregates, and their puzzling disinflationary nature — are due to measurement errors in total factor productivity (TFP). Reduced-form innovations to TFP, which are typically identified as unanticipated technology shocks, are found to generate anomalous responses — notably, a persistent inflationary effect — that are inconsistent with the interpretation of these disturbances as supply shocks, thus hinting at the presence of an unpurged non-technological component in measured TFP. Such an impurity undermines existing identification schemes, which are based on the premise that measured TFP is entirely driven by surprise and news shocks to technology. In this paper, we estimate the macroeconomic effects of news shocks in the U.S. using an agnostic identification approach that is robust to measurement errors in TFP. Our strategy consists in identifying the surprise technology shock as the linear combination of reduced-form innovations that best accounts for TFP movements on impact, subject to the constraint that a positive realization of the shock yields a negative impulse response of inflation. We then identify the news shock as the linear combination of reduced-form innovations that is orthogonal to the surprise technology shock and that best predicts future movements in TFP at a long but finite horizon. We find no evidence of negative comovement conditional on a news shock, and the disinflation puzzle essentially vanishes under our identification strategy. Our results also indicate that news shocks have become an important driver of business-cycle fluctuations in recent years.

JEL classification: E20, C32, E32, O47
Key words: Identification, Measurement errors, News shocks, Sign restrictions, Technology, Total factor productivity.

---

*We thank the Editor, Sanjay K. Chugh, and two anonymous referees for very helpful comments and suggestions. We also thank our discussants Nadav Ben Zeev and Daniele Siena, as well as Paul Beaudry, Giacomo Candian, Ryan Chahrour, Patrick Fève, André Kurmann, Michel Normandin, and seminar participants at the Canadian Macro Study Group Meeting and the T2M Conference for useful feedback, and David Gutkovsky for excellent research assistance. Financial support from the HEC Montréal Foundation is gratefully acknowledged. The views expressed in this paper are those of the author(s) and do not necessarily represent the views of the IMF, its Executive Board, or IMF management.

†Department of Applied Economics, HEC Montréal, and CIREQ, 3000 chemin de la Côte-Sainte-Catherine, Montréal, Québec, Canada H3T 2A7. Tel.: 1-514-340-7003; Fax: 1-514-340-6469; E-mail: hafedh.bouakez@hec.ca.

‡International Monetary Fund, 700 19th St NW Washington, DC 20431 USA; E-mail: lkemoe2@imf.org.

1 Introduction

A long-standing and fundamental question in macroeconomics is: what causes business-cycle fluctuations? Following the seminal work of Beaudry & Portier (2006), interest has been rekindled in Pigou (1927)'s theory of business cycles, according to which changes and revisions in expectations about future fundamentals can give rise to boom-bust cycles. A number of empirical studies — based on vector autoregressions (VARs) — have therefore attempted to gauge the importance of news shocks about future productivity in generating the type of positive comovement of macroeconomic aggregates observed in the data, and, more generally, in explaining business cycles.[1]

Beaudry & Portier (2006) were the first to document using U.S. data that news shocks lead to positive comovement of consumption, hours worked, and investment, and account for the bulk of their variability at business-cycle frequencies. Beaudry & Lucke (2010) and Beaudry & Portier (2014) reach essentially the same conclusions.[2] These findings have been challenged, however, by some scholars who questioned the underlying identification strategies.[3] Using an alternative, more flexible, identification approach, Barsky & Sims (2011) find that good news about future technology tend to raise consumption but to decrease output, hours worked, and investment in the short run.[4] They also find that inflation declines sharply and persistently in response to a positive realization of the news shock; a result deemed puzzling in light of the standard New Keynesian model.[5] Though Barsky & Sims (2011) find that news shocks account for a significant fraction of output variability at business-cycle frequencies, they invoke the negative comovement to conclude that these shocks are unlikely to be a major driver of business cycles. These findings are confirmed by subsequent studies that propose alternative but related methodologies to Barsky & Sims' (e.g., Forni et al. (2014), Barsky et al. (2015), and Kurmann & Sims (2021)).

Existing empirical approaches to identify news shocks about future productivity are based on

[1]Two alternative approaches to evaluate the importance of news shocks exist in the literature. The first consists in estimating/calibrating dynamic stochastic general-equilibrium (DSGE) models that feature anticipated shocks to technology (e.g., Fujiwara et al. (2011), Karnizova (2012), Schmitt-Grohé & Uribe (2012), Khan & Tsoukalas (2012), Görtz & Tsoukalas (2017), and Görtz et al. (Forthcoming). The second uses directly observable news such as the standardization of new technologies (Baron & Schmidt (2014)), oil and gas discoveries (Arezki et al. (2017)), and changes in firms' stock-market valuation due to announcements of patent grants (Cascaldi-Garcia & Vukotić (2022)).

[2]Di Casola & Sichlimiris (2018) extend the methodology proposed by Beaudry & Portier (2006) and find that news shocks about future technology are inflationary.

[3]Beaudry & Portier (2006), Beaudry & Lucke (2010), and Beaudry & Portier (2014) estimate small-scale systems (two to five equations) in which news shocks are identified using a mix of short- and long-run restrictions. Kurmann & Mertens (2014) show that Beaudry & Portier (2006)'s identification scheme does not have a unique solution when applied to a Vector Error Correction Model (VECM) with more than two variables. This identification scheme is therefore uninformative about the effects of news shocks and their importance for business cycles. Kurmann & Mertens (2014) further point out that the validity of the identification strategy proposed by Beaudry & Lucke (2010) critically depends on the plausibility of zero restrictions for other non-news shocks necessary to identify news shocks. Finally, Forni et al. (2014) argue that small-scale VARs and VECMs do not contain enough information to recover anticipated technology shocks from observable variables, a problem commonly known as *non-fundamentalness*.

[4]Barsky & Sims (2011) identify the news shock as the shock that best explains future movements in total factor productivity not accounted for by its own innovation.

[5]See, for instance, Jinnai (2013), Barsky et al. (2015), and Kurmann & Otrok (2014).

the premise that total factor productivity (TFP) is entirely and exclusively driven by two orthogonal disturbances: unanticipated and news shocks, the latter generally affecting TFP with a lag. This assumption is consistent with the standard treatment of TFP in theoretical macroeconomic models. Hence, the above-mentioned studies invariably include a measure of TFP in the information set when attempting to identify technological news shocks from the data.

In this paper, we argue that the TFP measures typically utilized in the empirical literature contain important measurement errors that call into question the interpretation of measured TFP as a proxy for technology. This is despite the corrections aiming at purging TFP of its non-technological component by controlling for unobserved variations in labor and capital. Most importantly, we demonstrate that the negative comovement of macroeconomic aggregates and the disinflation puzzle documented in recent empirical studies are spurious and are just an artifact of using a polluted measure of technology.

We document the severity of measurement errors in the adjusted TFP measure constructed by Fernald (2014) — which is the most widely used TFP series — by examining the dynamic effects of an unanticipated technology shock, identified as the reduced-form innovation to TFP, as is done in *all* existing VAR-based studies on news shocks.[6] The most revealing symptom of the presence of measurement errors is that unanticipated technological improvements are found to be inflationary, an outcome that runs against the conventional interpretation of surprise technology shocks as supply shocks, and violates the prediction of standard theories of aggregate fluctuations. A favorable surprise technology shock is also found to have counter-intuitive effects on stock prices and consumer confidence, which are initially unresponsive to the shock but fall persistently in the subsequent periods. We interpret these anomalous responses as an indication that the TFP series used in the empirical literature is an uncleansed measure of technology. Using artificial data from a simple New Keynesian model, we illustrate how, for instance, failing to account for non-constant returns to scale when measuring TFP could lead one to conclude that a surprise technology shock identified as the reduced-form innovation to TFP is inflationary whereas the true shock is not. Since a correct identification of news shocks hinges on the surprise technology shocks being properly identified, measurement errors in TFP are likely to undermine existing identification approaches.

We then propose an agnostic identification strategy that is robust to the presence of measurement errors in TFP. Our methodology relaxes the assumption that only technological shocks affect measured TFP, and assumes instead that the latter can be perturbed by non-technological disturbances at any given horizon. To identify the surprise technology shock, we select the linear combination of reduced-form innovations that explains most of the forecast error variance of TFP at the one-quarter horizon, subject to the constraint that a positive realization of the shock yields

[6] The only exception is the study by Kurmann & Sims (2021), in which there is no attempt to identify surprise technology shocks.

a negative impulse response of inflation. Hence, by construction, our strategy avoids the inflation anomaly engendered by identification schemes that associate surprise technology shocks with reduced-form innovations to TFP. We then extract the news shock as the linear combination of reduced-form innovations that is orthogonal to the surprise technology shock, and that maximizes the contribution of the news shock to the forecast error variance of TFP at a long but finite horizon. The argument underlying this criterion, originally proposed by Francis et al. (2014) and commonly referred to as the Max Share, is that the contribution of non-technology shocks to movements in TFP is likely to be negligible at very low frequencies.

We take our agnostic approach to the data by estimating a seven-variable VAR similar to that considered by Barsky & Sims (2011), first using their original data set, which spans the period 1960Q1–2007Q3, and then using an updated sample that extends the data coverage to 2019Q4. We find that non-technology shocks account for nearly half of the forecast error variance of Fernald's TFP series at the one-quarter horizon. This observation confirms the existence of non-trivial measurement errors in measured TFP and raises skepticism about available estimates of the effects of news shocks. Our results also show that the estimated effects of unanticipated technology shocks are remarkably consistent both with the predictions of the medium-scale New Keynesian model of Smets & Wouters (2007) and with the empirical evidence based on identification via long-run restrictions. In addition to being disinflationary by construction, an unanticipated technological improvement leads to a persistent and hump-shaped increase in consumption and output, and to a short-run decline in hours worked. Moreover, the shock is found to have a delayed positive effect on stock prices and consumer confidence.

Turning to the effects of news shocks, we find no evidence of negative comovement between consumption, output, and hours worked using our methodology. In the sample ending in 2007Q3, a favorable news shock triggers an increase in consumption, but the short-run responses of output and hours worked are small and statistically indistinguishable from zero. In the updated sample, all three variables increase significantly and persistently in response to the shock. Importantly, this simultaneous increase — indicative of positive comovement — occurs even before TFP starts to rise, thus lending support to the view that aggregate fluctuations can be driven by expectations of higher productivity. Our results also indicate that the inflation response is mostly statistically insignificant in both samples. In other words, the disinflation puzzle essentially vanishes under our identification strategy. More generally, the effects of a news shock identified using our agnostic approach differ markedly from those identified based on Barsky & Sims' or the Max Share criterion along with the orthogonality requirement with respect to the reduced-form innovation to TFP.

Finally, variance-decomposition results and the historical decomposition of the time series of consumption, output, and hours strongly suggest that news shocks are unlikely to have been a major contributor to business-cycle fluctuations before 2007. In the extended sample, however, we find that news shocks account for roughly 30 to 50 percent of the forecast error variance of consumption, output, and hours worked at business-cycle frequencies, and that they explain a significant share of the decline in these quantities during the recent U.S. downturns, including the Great Recession. Together, these findings indicate that TFP news shocks have become an important source of business-cycle fluctuations in recent years, a conclusion that contradicts the verdict of the recent empirical literature that builds on Barsky & Sims' methodology (e.g., Forni et al. (2014), Barsky et al. (2015), and Kurmann & Sims (2021)). Interestingly, our results are consistent with those reported by Chahrour et al. (2020), who agnostically identify the shock that mainly drives the comovement between output and hours worked in the U.S. Their identified shock, which is found to account for a large fraction of aggregate fluctuations at business-cycle frequencies, turns out to be essentially orthogonal to adjusted TFP at short horizons while explaining most of its variability in the long run, thus bearing the interpretation of a TFP news shock.

The presumption that TFP is measured with error is of course not new; it has been discussed, for instance, in Christiano et al. (2004), Basu et al. (2006), and Fernald (2014). In a closely related paper, Kurmann & Sims (2021) also study the implications of measurement errors in TFP for the identification of news shocks. These authors, however, do not establish a link between the anomalous responses to a surprise technology shock and the existence of measurement errors in TFP. Instead, their suspicion of the presence of such errors is based on the observation that the effects of news shocks estimated using Barsky & Sims' methodology are highly sensitive to revisions in Fernald's adjusted TFP series, with the short-run response of hours worked ceasing to be negative when more recent vintages of TFP data are used. Kurmann & Sims (2021) propose an alternative identification strategy that is robust to revisions in TFP data, producing similar results to those originally estimated by Barsky & Sims (2011). Their approach relaxes the assumption that measured TFP does not react contemporaneously to news shocks, and instead identifies these shocks solely based on the Max Share criterion. A crucial implication of this strategy, however, is that the news shock is not orthogonalized with respect to the surprise technology shock. Because the latter is typically identified as the reduced-form innovation to TFP, imposing orthogonality with respect to this shock necessarily implies that the contemporaneous response of TFP to the news shock is nil,[7] which is precisely the restriction that Kurmann & Sims (2021) aim to relax. This in turn suggests that Kurmann & Sims' strategy is likely to confound surprise and anticipated technological shocks, as both shocks affect TFP in the short and in the long run, making it impossible — without further assumptions — to disentangle their respective contribution to the forecast error variance of TFP at any given horizon.

The rest of this paper is organized as follows. Section 2 discusses the symptoms of measurement errors in TFP. Section 3 presents our agnostic identification strategy. Section 4 discusses the results based on Barsky & Sims' original data and on an updated sample. Section 5 concludes.

[7] Assuming that TFP is ordered first in y_t, the impulse vector associated with the surprise technology shock has zeros everywhere except for the first element. For this impulse vector to be orthogonal to the one associated with the news shock, the latter must have zero as its first element.

2 The Inflation Anomaly and Other Symptoms of Measurement Errors in TFP

In this section, we illustrate the extent to which the effects of unanticipated technology shocks typically reported in the VAR-based "news" literature are inconsistent with the predictions of New Keynesian models and, more generally, standard theories of aggregate fluctuations. We view these inconsistencies as symptoms of the presence of measurement errors in the TFP series commonly used in the literature.

2.1 Unanticipated technology shocks: measurement...

In the VAR-based literature on news shocks, unanticipated technology shocks are usually identified as reduced-form innovations to TFP. Formally, let y_t be a $k \times 1$ vector of observables of length T, which includes TFP and which has the following moving-average (MA) representation:

$$y_t = B(L)u_t,$$

where u_t is a $k \times 1$ vector of statistical innovations, whose variance-covariance matrix is denoted by Σ. Let ϵ_t be a $k \times 1$ vector of structural innovations, including the unanticipated technology shock, whose variance-covariance matrix is normalized to I_k. If a linear mapping between the statistical innovations, u_t, and the structural shocks, ϵ_t, exists, then we can write

$$u_t = A\epsilon_t,$$

where the impact matrix, A, must be such that $AA' = \Sigma$. Assuming (without loss of generality) that TFP is ordered first in y_t and that the unanticipated technology shock is ordered first in ϵ_t, a Cholesky decomposition of Σ ensures that the surprise technology shock is proportional to the statistical innovation to TFP.

We use the strategy above to measure the effects of a surprise technology shock within a seven-variable VAR similar to that estimated by Barsky & Sims (2011). The vector of observables includes adjusted TFP, output, consumption, hours, inflation, stock prices, and consumer confidence, measured at a quarterly frequency. We start by using Barsky & Sims' original data, which span the period 1960Q1–2007Q3; we then update the sample by extending it to 2019Q4.[8]

The results are shown with solid black lines in Figure 1.[9] The (one-standard-error) confidence

[8]The series used in estimation are constructed as follows. Adjusted TFP is the quarterly series constructed by Fernald (2014), which controls for unobserved input variation. Output is measured by the log of real GDP in the non-farm business sector. Consumption is measured by the log of real personal spending on non-durables and services. Hours are measured by the log of total hours worked in the non-farm business sector. Output, consumption and hours are expressed in per capita terms by dividing them by the civilian, noninstitutional population, age 16 and over. Inflation is measured by the percentage change in the CPI for all urban consumers. Stock prices are measured by the log of the S&P index. Consumer confidence is retrieved from the Michigan Survey of Consumers.

[9]These results are based on a VAR with 3 lags. Alternative lag lengths yield similar results.

intervals around the estimated impulse responses are computed using the bias-corrected bootstrap procedure proposed by Kilian (1998). A surprise technology shock triggers a transitory increase in TFP, output, and consumption. In all three cases, the estimated response is rather monotonic and the variable reverts to its pre-shock level rather rapidly. In contrast, hours worked exhibit a relatively muted — and mostly statistically insignificant — response. The figure also shows that, in response to the identified surprise technology shock, inflation rises persistently and in a hump-shaped manner, with a peak occurring at around 10 quarters after the shock. Stock prices and consumer confidence, in contrast, are unresponsive on impact and eventually fall below their pre-shock levels for a prolonged period of time. Very similar results are reported by Forni et al. (2014), Barsky et al. (2015), and Fève & Guay (2019).

When we extend the sample to 2019Q4, we observe some differences with respect to the results discussed above, but only at short horizons. As Figure 2 shows, hours worked now fall while stock prices and consumer confidence rise during the first three quarters after the shock. At medium and long horizons, however, the results based on the extended sample are very similar to those based on the sample ending in 2007Q3. In particular, inflation continues to rise persistently and in a hump-shaped manner, while stock prices and consumer confidence continue to decline persistently, converging to their pre-shock levels from below.

2.2 ... and theory

How do the empirical findings discussed in the previous section compare with the predictions of New Keynesian theory of aggregate fluctuations? We answer this question by studying the effects of unanticipated technology shocks both within the simplest version of the New Keynesian model and the more realistic medium-scale version proposed by Smets & Wouters (2007). To do so, we assume that the log of TFP (in deviation from its mean), a_t, is governed by the following process:

$$a_t = \rho_a a_{t-1} + x_{t-1} + \epsilon_t^s, \tag{1}$$

$$x_t = \rho_x x_{t-1} + \epsilon_t^n, \tag{2}$$

where ϵ_t^s and ϵ_t^n are, respectively, the surprise and anticipated (or news) technology shocks, and $0 \leq \rho_a, \rho_x < 1$. Notice that ρ_x is irrelevant to the dynamic effects of the surprise shock and thus ρ_a and the size of the disturbance ϵ_t^s are the only parameters that one needs to calibrate to study those effects. We choose those two parameters such that the implied response of TFP to the surprise technology shock mimics as closely as possible the median response estimated from the data. The model-based responses of TFP, consumption, output, hours, and inflation are superimposed on their empirical counterparts in Figures 1 and 2.

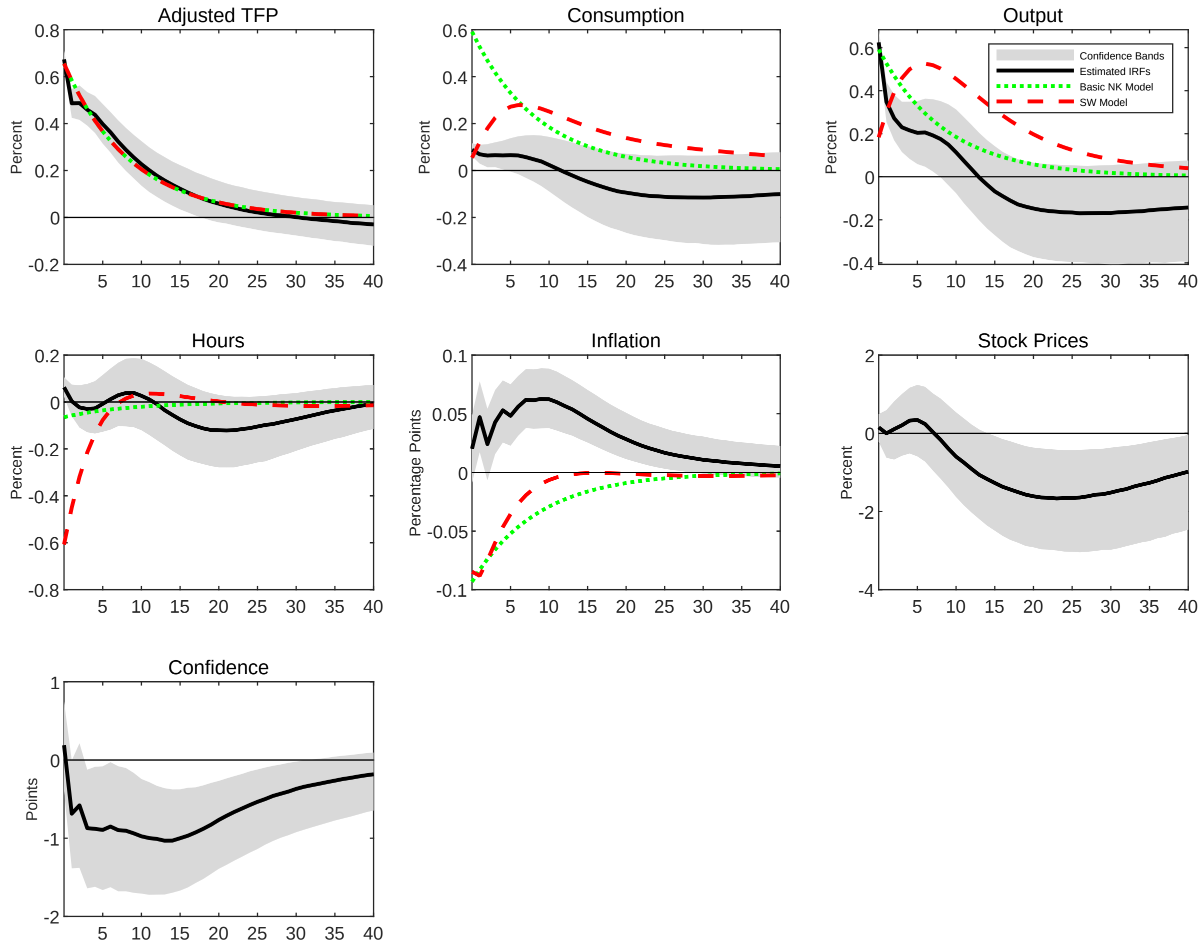

Figure 1: Impulse responses to a surprise technology shock. Sample: 1960Q1–2007Q3.

Notes: The figure shows the impulse responses to a surprise technology shock. The solid lines are the median impulse responses estimated based on the reduced-form innovation to TFP. The 68 percent confidence bands are the bias-corrected bootstrap confidence intervals computed using Kilian (1998)'s procedure with 2000 replications. The dotted lines are the impulse responses obtained from the standard New Keynesian model. The dashed lines are the impulse responses obtained from the Smets & Wouters (2007) model.

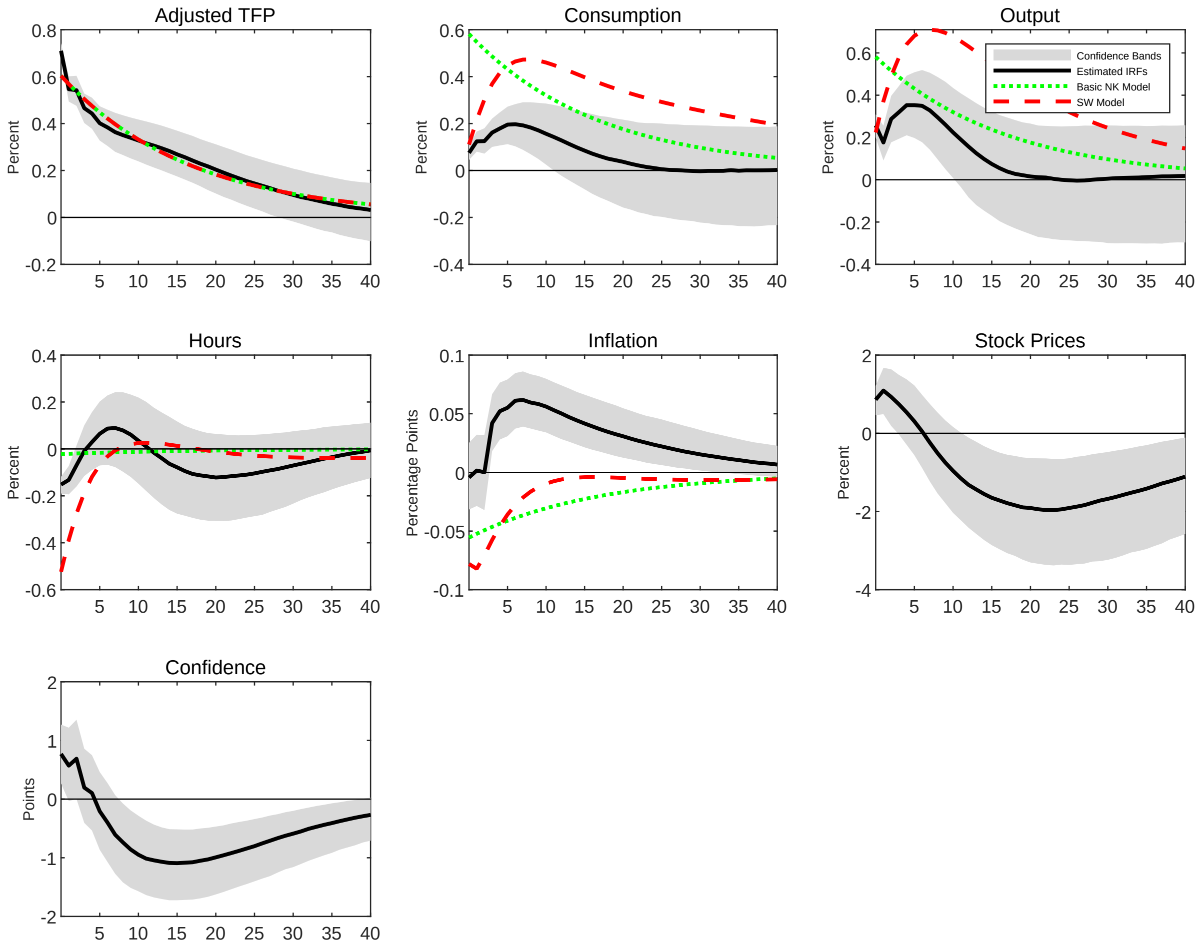

Figure 2: Impulse responses to a surprise technology shock. Sample: 1960Q1–2019Q4.

Notes: The figure shows the impulse responses to a surprise technology shock. The solid lines are the median impulse responses estimated based on the reduced-form innovation to TFP. The 68 percent confidence bands are the bias-corrected bootstrap confidence intervals computed using Kilian (1998)'s procedure with 2000 replications. The dotted lines are the impulse responses obtained from the standard New Keynesian model. The dashed lines are the impulse responses obtained from the Smets & Wouters (2007) model.

2.2.1 The basic New Keynesian model

Consider first the basic New Keynesian model, summarized by the following log-linearized equations (around a zero-inflation steady state):[10]

$$c_t = y_t, \tag{3}$$

$$y_t = a_t + \alpha n_t, \tag{4}$$

$$mc_t = \sigma c_t + \varphi n_t - a_t, \tag{5}$$

$$c_t = \mathbb{E}_t c_{t+1} - \sigma^{-1}(i_t - \mathbb{E}_t \pi_{t+1} - \ln \beta), \tag{6}$$

$$\pi_t = \beta \mathbb{E}_t \pi_{t+1} + \lambda mc_t, \tag{7}$$

$$i_t = \ln \beta + \phi_\pi \pi_t + \phi_y (y_t - y_t^f), \tag{8}$$

[10]This is essentially the model presented in Galí (2008).

where c_t is consumption, y_t is output, n_t is hours worked, mc_t is real marginal cost, π_t is the inflation rate, i_t is the nominal interest rate, and y_t^f is the flexible-price (or natural) level of output. All the variables are expressed as percentage deviations from their steady-state values except π_t and i_t, which are expressed in levels. The non-policy parameters are defined as follows: $\alpha > 0$ is labor intensity, $\sigma > 0$ is the inverse of the elasticity of intertemporal substitution, $\varphi > 0$ is the inverse of the Frisch elasticity of labor supply, $0 < \beta < 1$ is the discount factor, and $\lambda = \frac{(1-\theta)(1-\beta\theta)}{\theta} \frac{\alpha}{\alpha+(1-\alpha)\varepsilon} > 0$, with $0 < \theta < 1$ being the Calvo probability of not changing prices, and $\varepsilon > 1$ the elasticity of substitution between goods. The policy parameters $\phi_\pi, \phi_y > 0$ are the coefficients attached to, respectively, inflation and the output gap in the interest-rate rule.

Model (3)–(8) can be solved analytically to determine the effects of a surprise technology shock. Assuming for simplicity that $\alpha = 1$ and setting $x_t = 0$ for all t, one can use the method of undetermined coefficients to show that

$$\pi_t = -\frac{\sigma\lambda\left(1+\varphi\right)\left(1-\rho_a\right)}{\Delta_a} a_t,$$

where $\Delta_a = \lambda(\sigma+\varphi)\left(\phi_\pi - \rho_a\right) + \left(1-\beta\rho_a\right)\left[\sigma\left(1-\rho_a\right) + \phi_y\right] > 0$.[11] Since the numerator in the expression above is positive, an unanticipated technological improvement will cause inflation to fall persistently as long as $\rho_a < 1$. This disinflationary effect reflects the persistent fall in real marginal cost or, equivalently, the negative output gap resulting from the shock.[12] This can be seen by noticing that

$$mc_t = -\frac{\sigma\left(1+\varphi\right)\left(1-\rho_a\right)\left(1-\beta\rho_a\right)}{\Delta_a} a_t.$$

The surprise technology shock has a positive effect on output (and thus consumption) but an ambiguous effect on hours worked. The solutions for these variables are given by

$$\begin{aligned} y_t &= \frac{\left(1+\varphi\right)\left[\lambda\left(\phi_\pi - \rho_a\right) + \left(\sigma+\varphi\right)^{-1}\left(1-\beta\rho_a\right)\phi_y\right]}{\Delta_a} a_t, \\ n_t &= \left\{\frac{\left(1+\varphi\right)\left[\lambda\left(\phi_\pi - \rho_a\right) + \left(\sigma+\varphi\right)^{-1}\left(1-\beta\rho_a\right)\phi_y\right]}{\Delta_a} - 1\right\} a_t. \end{aligned}$$

Under plausible parameter values, however, hours worked fall in response to a positive unanticipated technology shock. The responses depicted in Figures 1 and 2 (with green dotted lines) are obtained using the following standard parameterization of the model: $\sigma = \varphi = 1$, $\beta = 0.99$, $\theta = 0.75$, $\phi_\pi = 1.5$, and $\phi_y = 0.125$. Under these parameter values, a positive surprise shock to technology raises output and consumption and decreases hours worked and inflation.

[11] The necessary and sufficient condition for the existence of a unique linear rational expectations equilibrium is given by $\lambda(\sigma+\varphi)\left(\phi_\pi - 1\right) + \left(1-\beta\right)\phi_y > 0$. It is straightforward to see that this condition implies that $\Delta_a > 0$.

[12] By iterating equation (7) forward, inflation can be expressed as a discounted sum of current and expected future real marginal costs.

The dynamic responses implied by the model hardly match those estimated from the data, but the most striking discrepancy concerns the response of inflation, which has the opposite sign and a completely different shape relative to what is predicted by the VAR.

2.2.2 The Smets and Wouters (2007) model

Next, consider the medium-scale model developed by Smets & Wouters (2007). To conserve space, we only summarize the main features of the model and refer the reader to their paper for a more detailed description. The model features a representative household whose preferences exhibit habit formation in consumption. The final good is produced using an aggregator of intermediate goods characterized by a non-constant elasticity of substitution. Intermediate goods are produced using a technology that depends on TFP, labor, and capital, and that exhibits variable capital utilization and fixed costs. Capital accumulation is subject to investment adjustment costs. Both prices and wages are set in a staggered fashion à la Calvo, whereby the non-optimizing agents partially index their prices and wages to past inflation, thus giving rise to a New Keynesian Phillips curve that depends not only on current and expected future inflation but also on past inflation. Monetary policy follows an interest rate rule with a smoothing component. The model is estimated by Bayesian techniques using U.S. data over the period 1966Q1–2004Q4.

We use Smets and Wouters' posterior means for the structural parameters to generate the implied responses to an unanticipated positive technology shock, which are represented by the dashed red lines in Figures 1 and 2. Despite some quantitative differences, these responses are in line with the predictions of the basic New Keynesian model: output and consumption rise while hours worked and inflation fall in response to the shock. The persistent fall in inflation is again in stark contrast with the positive response obtained from the VAR.[13] Notice also that the VAR-based responses of output and consumption lack the persistent and hump-shaped pattern implied by the model.

2.3 Discussion

As we have just shown, reduced-form innovations to TFP are found to be inflationary, an outcome that runs against the conventional interpretation of technology shocks as supply shocks, and contradicts the prediction of standard theories of aggregate fluctuations. It is also at odds with the results reported by a number of empirical studies that rely on the long-run restriction approach proposed by Galí (1999) to identify exogenous technology shocks (e.g., Edge et al. (2003), Christiano et al. (2003), Fève & Guay (2010)). Moreover, the result that technology shocks have a delayed negative effect on stock prices and consumer confidence also appears hard to reconcile with the view that

[13]A persistent decline in inflation following a favorable surprise technology shock is also predicted by the New Keynesian models estimated by Ireland (2004) and Altig et al. (2011), though the inflation response is relatively small in magnitude in the latter case.

technology enhances efficiency and raises the productive capacity of the economy.

These observations cast serious doubt on the interpretation of reduced-form innovations to TFP as pure unanticipated technological improvements. The identified shocks appear to be contaminated by other non-technological disturbances that also affect measured TFP contemporaneously and whose effects are akin to those of a demand shock. Since a proper identification of news shocks about future productivity hinges on purging TFP of its non-technological component, the anomalous responses just discussed suggest that existing methodologies — albeit sound in theory — may still fail to correctly identify news shocks and their effects, due to measurement errors in TFP.

In the models discussed in Section 2.2, TFP is assumed to be exogenous to the state of the economy and, as such, is not expected to be affected by demand shocks — note that this is precisely the identifying assumption underlying the empirical literature on news shocks. TFP, however, is not readily observable in the data and must be inferred from production and input use, a task that poses a number of measurement challenges. First, some inputs may not be observable or measurable; second, input utilization varies in response to non-technology shocks; third, the production technology may have non-constant returns to scale; fourth, aggregating inputs across heterogeneous production sectors may introduce a bias. Failing to eliminate any of these potential sources of measurement errors may result in an incorrect measure of TFP and thus a poor proxy for technology. In their seminal paper, Basu et al. (2006) went a long way towards constructing a purified annual measure of technology by adjusting TFP for observed and unobserved input variations and non-constant returns to scale. The quarterly TFP series used in the empirical literature on news shocks was constructed by Fernald (2014) following Basu et al. (2006)'s methodology but without correction for non-constant returns to scale since the industry level data needed for this correction are only available at an annual frequency.

To get a sense of how this impacts the measurement of TFP, we plot in Figure 3 the series of annual TFP growth constructed by Fernald (2014) and Basu et al. (2006) for the period 1960–1996.[14] Although there is some similarity between the two series, their correlation is modest (0.60), suggesting that the constant-returns-to-scale assumption underlying the construction of the quarterly TFP series is counterfactual and is likely to be one of the culprits for the anomalous responses documented above.[15]

To further illustrate the importance of this assumption as a potential source of measurement errors, we estimate the effects of a surprise technology shock identified as the reduced form innovation to Basu et al. (2006)'s series using the same observable variables as in Section 2.1, measured annu-

[14] Basu et al. (2006)'s TFP series ends in 1996.

[15] There are a number of additional reasons why the two TFP series differ. These include data revisions — the output and investment data that Basu et al. (2006) used have been subject to multiple benchmark revisions — as well as differences in industry coverage and data sources. In addition, the series constructed by Fernald (2014) uses a measure of labor quality that contains significant high-frequency noise. These differences, however, are unlikely to cause Fernald's TFP series to be more endogenous to the state of the economy than the series constructed by Basu et al. (2006).

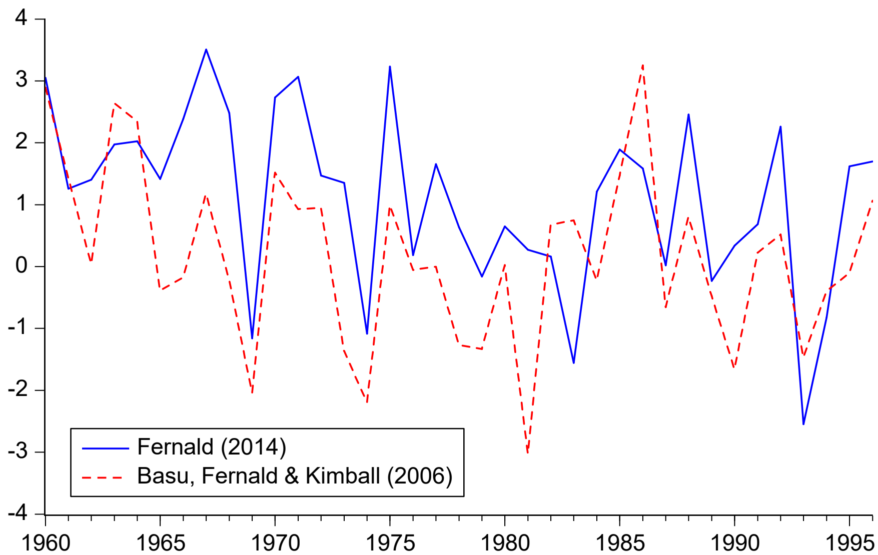

Figure 3: Annual TFP growth computed by Fernald (2014) and Basu et al. (2006).

ally. The estimated impulse responses and their confidence bands are shown in Figure 4, in which a period corresponds to a year.[16] The figure shows that, following a positive technology shock, output remains essentially unresponsive on impact but increases in a hump-shaped manner during the subsequent years, whereas hours worked fall significantly at the time of the shock. Inflation also falls sharply on impact, consistently with the expected disinflationary effect of a technological improvement, and in sharp contrast with the rise in inflation obtained using the quarterly TFP series. This observation suggests that Basu et al. (2006)'s TFP series is less polluted by non-technological factors than Fernald (2014)'s quarterly series.

In order to gain further insights into why failing to account for non-constant returns to scale can severely bias the estimation of the inflation response to surprise technology shocks, we append the New Keynesian model presented in Section 2.2.1 with a preference shock, and use it to generate artificial series under the assumptions of constant and increasing returns to scale, by setting α to 1 and 1.1, respectively. The additional shock captures changes in consumers' preference for safety and liquidity, and is assumed to follow an AR(1) process. The details of the experiment are described in detail in the Online Appendix. Now, assume that one measures TFP as $\tilde{a}_t \equiv y_t - n_t$. Under constant returns to scale, $\tilde{a}_t$ coincides with a_t, but the two series diverge under increasing returns to scale. Using the simulated series of $\tilde{a}_t$, y_t, and π_t, we estimate a three-equation VAR and identify the

[16]The results reported in Figure 4 are based on a VAR with one lag. We obtain very similar results when we include two lags. Because we are estimating a VAR with 7 variables using 36 annual observations, including more lags leaves too few degrees of freedom to obtain reliable estimates.

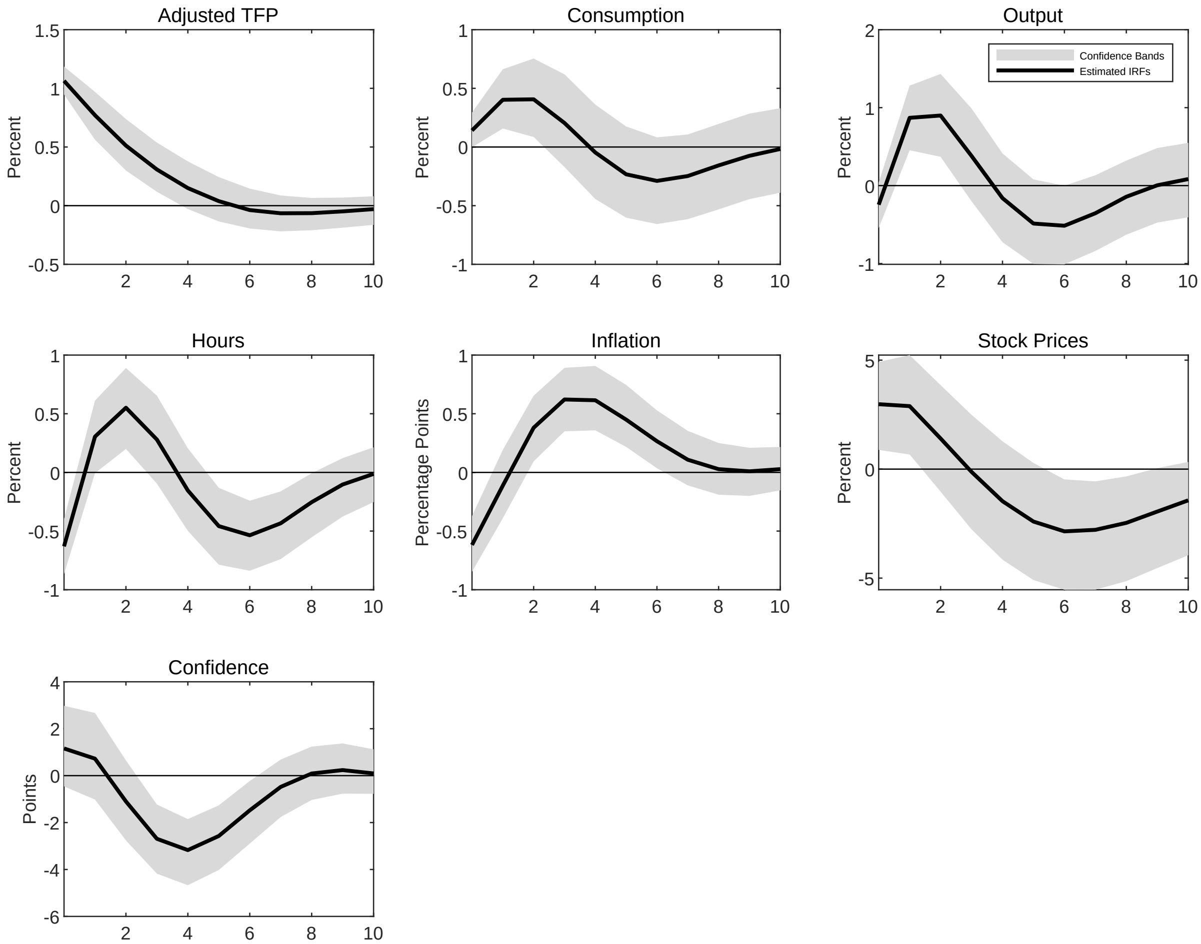

Figure 4: Impulse responses to a surprise technology shock based on Basu et al. (2006)'s annual TFP series.

Notes: The figure shows the impulse responses to a surprise technology shock. The solid lines are the median impulse responses estimated based on the reduced-form innovation to TFP. The 68 percent confidence bands are the bias-corrected bootstrap confidence intervals computed using Kilian (1998)'s procedure with 2000 replications.

surprise technology shock as the reduced-form innovation to measured TFP. The estimated impulse responses to this shock are depicted in Figure 5 along with the true responses (i.e., calculated from the model). The upper panels of the figure show that the identification procedure recovers the true shock and its disinflationary effect very accurately when there is no discrepancy between measured and actual TFP (i.e., when $\alpha = 1$). In contrast, when measured TFP fails to account for increasing returns to scale, the identified shock counterfactually leads to an increase in inflation.

The intuition for this result can be understood by noticing that measured TFP can be expressed as $\tilde{a}_t = \frac{1}{\alpha} a_t + \frac{\alpha-1}{\alpha} y_t$. Since the preference shock — which is a demand disturbance — moves y_t and π_t in the same direction, it gives rise to positive comovement of $\tilde{a}_t$ and π_t as long as $\alpha > 1$ (see

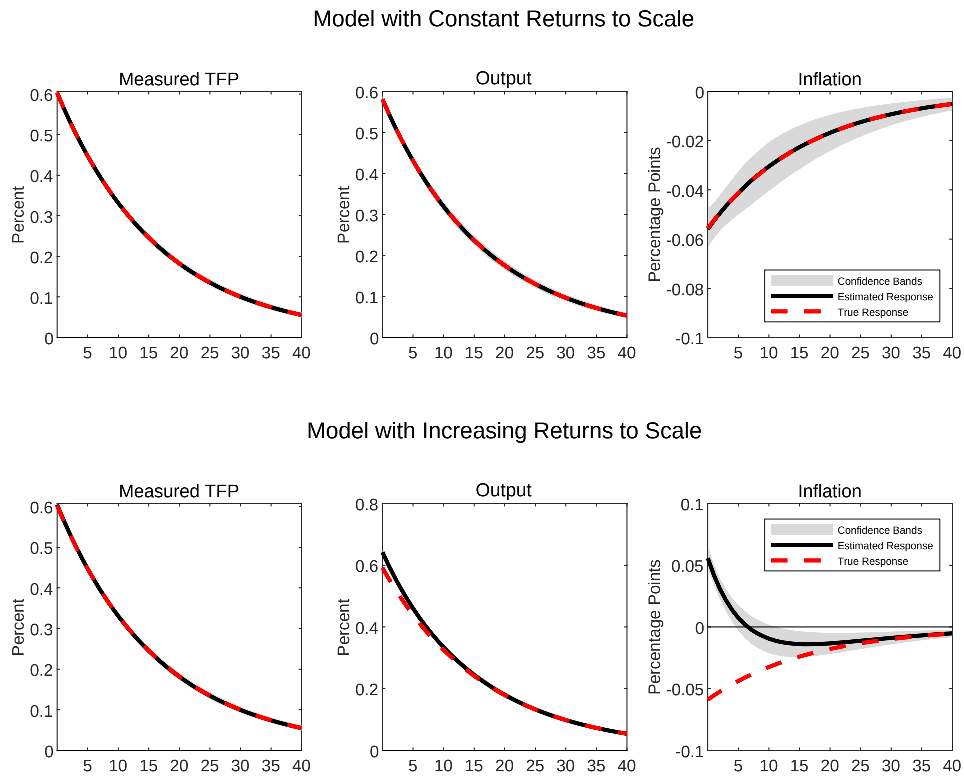

Figure 5: Impulse responses to a surprise technology shock based on artificial data.

Notes: The figure shows the impulse responses to a surprise technology shock estimated within a three-equation VAR using artificial data generated from the basic New Keynesian model. The series used in estimation are those of measured TFP, output, and inflation, each including 20000 observations. The solid lines are the median impulse responses to a reduced-form innovation to TFP. The 68 percent confidence bands are based on 2000 draws. TFP is always measured by assuming constant returns to scale. The upper panels report the results when the data-geenrating model features constant returns to scale. The lower panels report the results when the data-generating model features increasing returns to scale. The dashed lines are the theoretical (true) responses.

Figure A.3 in the Online Appendix). To the extent that the variance of this shock is sufficiently large, the reduced-form innovation to $\tilde{a}_t$ can generate a positive response of inflation.

We close this section by noting that although Basu et al. (2006)'s purified TFP series allows one to recover the disinflationary effect of surprise technology shocks in the short run, it still generates some anomalies that are hard to reconcile with conventional wisdom about the effects of those shocks. For instance, the initial decline in inflation followed by a protracted episode (of several years) during which inflation is above average. Moreover, while stock prices initially rise in response to a positive technology shock, they decline persistently during the subsequent years. Likewise, the shock triggers a delayed fall in consumer confidence that persists for a prolonged period of time. These responses cast doubt on the interpretation of the shock as a pure technological disturbance.

In sum, despite the colossal work carried out by Basu et al. (2006) and Fernald (2014) to construct a cleansed measure of technology, it is probably unrealistic to believe that the corrected TFP series is purged of all its non-technological factors, which in turn suggests that TFP-based measures of technology shocks are most likely contaminated by measurement errors. This conclusion motivates the agnostic approach that we describe in the next section.

3 An Agnostic Identification Approach

3.1 Idea

The maintained assumption underlying the empirical identification of news shocks about future productivity is that measured TFP is exclusively driven by surprise and anticipated technology shocks, the latter affecting TFP only with a lag. The common approach to identify the news shock is then to select the linear combination of reduced-form innovations that best explains (or forecasts) future movement in TFP while being orthogonal to the surprise technology shock. This strategy will correctly identify news shocks only to the extent that surprise technology shocks are the only disturbances that affect measured TFP contemporaneously, which, as we just argued above, seems highly unlikely.

We propose an alternative empirical strategy based on the assumption that measured TFP is affected by two types of disturbances: technological and non-technological shocks. The latter capture measurement errors due to the imperfect observability of inputs and their utilization rates, to the potential misspecification of the production function, and to aggregation bias. From this perspective, it may be inappropriate to characterize these shocks as structural, given that they do not bear a clear economic interpretation. However, this is not a concern for our methodology since we need not identify these shocks; we simply allow them to affect measured TFP contemporaneously and at any future horizon, just as surprise technology shocks.

To identify the surprise technology shock, we adopt the following agnostic strategy. We select the linear combination of reduced-form innovations that accounts for most of the forecast error variance of TFP at the one-quarter horizon, subject to the constraint that a positive realization of the shock yields a negative impulse response of inflation, consistently with the prediction of New Keynesian theory.[17] Hence, by construction, our strategy avoids the inflation anomaly engendered by identification schemes that associate surprise technology shocks with TFP innovations. We then identify the news shock as the linear combination of reduced-form innovations that is orthogonal to the surprise technology shock and that maximizes the contribution of the news shock to the forecast error variance of TFP at a long but finite horizon, H.[18] The latter criterion, initially

[17]The requirement that the surprise technology shock explains most of the impact variation in TFP may, to some extent, seem arbitrary. However, any identifying assumption needs minimal ground, which in this case presupposes that non-technological shocks do not account for the largest fraction of the one-quarter-ahead forecast error variance of TFP. This assumption is weaker than the one commonly used in the literature, which attributes 100 percent of this variance to technology shocks and none to non-technology shocks. Leaving the impact response of TFP completely unrestricted implies that our strategy would pick any disinflationary shock, even if it did not affect TFP at all. This would be an alternative source of contamination of the identified surprise technology shock.

[18]In related work, Nam & Wang (2019) also rely on sign restrictions to identify shocks that reflect changes in expectations, and to measure their importance for aggregate fluctuations. However, our focus and methodology differ

proposed by Francis et al. (2014) and commonly referred to as the Max Share, differs from the one used by Barsky & Sims (2011), which involves maximizing the contribution of the news shocks to the forecast error variance of TFP over all horizons up to a finite truncation horizon. Barsky & Sims' approach has been criticized on the ground that it may confound shocks that have either permanent or temporary effects on TFP, and has been shown to be quite sensitive to the truncation horizon (see Beaudry et al. (2011)). Since our approach allows for the presence of non-technology shocks, whose effects on measured TFP are likely to be much more important at short horizons than at more distant ones, the Max Share criterion seems more appropriated than the one proposed by Barsky & Sims (2011) to identify the news shock.

3.2 Implementation

Let $\tilde{A}$ denote the Cholesky decomposition of Σ, and assume again that TFP is ordered first in y_t. Any impact matrix $A = \tilde{A}D$, where D is an orthonormal matrix, also satisfies the requirement $AA' = \Sigma$. Let γ_j denote the jth column of D, $\epsilon_{1,t}$ denote the surprise technology shock (ordered first in ϵ_t), and $\epsilon_{2,t}$ denote the news shock (ordered second in ϵ_t).

We identify the surprise technology shock by selecting the orthonormal matrix D that satisfies the following two requirements: (i) the shock accounts for the largest fraction of the one-quarter-ahead forecast error variance of TFP, and (ii) positive realizations of the shock are non-inflationary. Note that because the impulse vector to this shock is $\tilde{A}\gamma_1$ (the first column of $\tilde{A}D$), we only need to characterize γ_1.

The h-step-ahead forecast error of vector y_t is

$$y_{t+h} - \mathbb{E}_t y_{t+h} = \sum_{\tau=0}^{h-1} B_\tau \tilde{A} D \epsilon_{t+h-\tau},$$

with $B_0 = I_k$. Denoting by $\Omega_{i,j}(h)$ the share of the forecast error variance of variable i attributable to structural shock j at horizon h ($h = 1, 2...$), this quantity is given by

$$\Omega_{i,j}(h) \equiv \frac{e_i' \left(\sum_{\tau=0}^{h-1} B_\tau \tilde{A} D e_j e_j' D' \tilde{A} B_\tau' \right) e_i}{e_i' \left(\sum_{\tau=0}^{h-1} B_\tau \Sigma B_\tau' \right) e_i} = \frac{\sum_{\tau=0}^{h-1} B_{i,\tau} \tilde{A} \gamma_j \gamma_j' \tilde{A} B_{i,\tau}'}{\sum_{\tau=0}^{h-1} B_{i,\tau} \Sigma B_{i,\tau}'},$$

where

$$B_{i,\tau} = e_i' B_\tau, \qquad \gamma_j = D e_j,$$

and e_i is a selection vector with 1 in the ith position and zero elsewhere.

from theirs. Nam & Wang (2019) are interested in identifying optimism shocks, which they define as non-technological non-inflationary disturbances that raise consumption and stock prices. To capture the non-inflationary effect of these shocks, they impose a non-negative response of the real interest rate. We instead leave the sign of the responses to our identified news shock unrestricted but impose a sign restriction on the response of inflation to a surprise technology shock, to which the news shock is orthogonal.

Denote by $r_{\pi,\gamma_1}(h)$ the impulse response of inflation to the impulse vector $\tilde{A}\gamma_1$ at horizon h. Our strategy to identify the surprise technology shock therefore amounts to selecting the vector γ_1 that solves the following maximization problem:

$$\max_{\{\gamma_1\}} \Omega_{1,1}(1) \equiv \frac{B_{1,0}\tilde{A}\gamma_1\gamma_1^{'}\tilde{A}B_{1,0}^{'}}{B_{1,0}\Sigma B_{1,0}^{'}}$$

s.t.

$$\gamma_1^{'}\gamma_1 = 1, \qquad r_{\pi,\gamma_1}(h) \leq 0 \text{ for } h = 1, ..., J.$$

The first constraint ensures that γ_1 is a column vector of an orthonormal matrix, while the second constraint restricts the dynamic response of inflation to be non-negative during the first J quarters after the shock. We choose J to be sufficiently long to ensure that the (median) inflation response at horizons $J+1$, $J+2$, ... is either negative, or positive but statistically insignificant. In practice, we set $J = 10$ quarters.[19]

It is worth emphasizing that, if it exists, the solution to the maximization problem above is unique. From this perspective, our strategy differs from the pure sign-restriction approach (e.g., Uhlig (2005) and Arias et al. (2018)), which would potentially yield a range of impulse vectors that are compatible with the desired sign restriction(s). Instead, our approach selects among that range the impulse vector that maximizes the contribution of the surprise technology shock to the one-quarter-ahead forecast error variance of TFP.[20]

Once the surprise technology shock, $\epsilon_{1,t}$, is identified, we identify the news shock, $\epsilon_{2,t}$, as the linear combination of reduced-form residuals that is orthogonal to $\epsilon_{1,t}$ and that explains the largest fraction of the forecast error variance of TFP at a long but finite horizon, H. This consists in selecting the vector γ_2 that solves

$$\max_{\{\gamma_2\}} \Omega_{1,2}(H) \equiv \frac{\sum_{\tau=0}^{H-1} B_{1,\tau}\tilde{A}\gamma_2\gamma_2^{'}\tilde{A}B_{1,\tau}^{'}}{\sum_{\tau=0}^{H-1} B_{1,\tau}\Sigma B_{1,\tau}^{'}}$$

s.t.

$$\gamma_2(1) = 0, \qquad \gamma_2^{'}\gamma_1 = 0, \qquad \gamma_2^{'}\gamma_2 = 1.$$

The first constraint ensures that the news shock does not affect TFP contemporaneously;[21] the

[19]In the Online Appendix, we show the estimated impulse responses when J is set to 1 and 4 quarters.

[20]In theory, our procedure may confound the surprise technology shock with a disinflationary non-technology supply shock that explains a substantial fraction of the forecast error variance of inflation on impact — a plausible candidate would be a negative energy-price shock. This however, requires the contribution of the non-technology supply shock to the volatility of inflation to be inordinately large. As we state below, we find our identified surprise technology shock to be essentially uncorrelated with an oil-price shock.

[21]If input utilization or returns to scale are imperfectly measured, then the news shock may affect measured TFP even on impact, as in the model described in Section 2.3. As can be seen in Figure 6 below, however, the initial response of measured TFP predicted by the model is virtually nil. In the Online Appendix, we show that our results are remarkably robust to relaxing the zero-impact restriction on the response of TFP to the news shock.

second constraint ensures that the news shock is orthogonal to ϵ_1; and the third constraint ensures that γ_2 is a column vector of an orthonormal matrix. In practice, we choose $H = 80$ quarters.

3.3 Simulation

Before applying our identification approach to actual data, we illustrate its usefulness in the controlled environment described in Section 2.3. Using the simulated series of measured TFP, output and inflation, we check whether, and to what extent, our agnostic procedure successfully recovers the true effects of a technological news shock when the data-generating model features increasing returns to scale whereas TFP is (incorrectly) measured under the assumption of constant returns to scale. The results are reported in Figure 6. For comparison, we also report the results based on two alternative identification strategies: the one proposed by Barsky & Sims (2011) (to which we henceforth refer as BS) and one that relies on the Max Share criterion to identify the news shock while leaving unrestricted the effects of the surprise technology shock to which the news shock is orthogonalized (to which we henceforth refer as MS). Like BS, MS identifies the surprise technology shock as the reduced-form innovation to measured TFP.

The upper and middle panels of Figure 6 show that both output and inflation fall during several quarters in response to a favorable news shock. These findings are at odds with the theoretical predictions, which indicate that output is initially unresponsive to the shock, before rising in the subsequent periods, whereas the inflation response is positive but muted at all horizons. In contrast, the responses obtained using our agnostic approach align remarkably well with those predictions. In particular, the inflation response is tiny and statistically indistinguishable from zero.

This experiment clearly illustrates how measurement errors in TFP can severely bias the estimated effects of anticipated technology shocks. A correct identification of these shocks hinges on the surprise technology shocks being properly identified, which is unlikely to be the achieved using the standard approach in the presence of measurement errors in TFP (see Section 2.3). On the other hand, simulation results suggest that our methodology is robust to the mis-measurement of TFP. Below, we use it to estimate the effects of news shocks based on actual data.

4 Results

We apply our agnostic identification strategy to the same seven-variable VAR estimated by Barsky & Sims (2011). We consider two data sets: the one originally used by these authors, which spans the period 1960Q1–2007Q3, and an updated data set that extends the data coverage through 2019Q4. For each of these samples, we discuss the impulse responses to a surprise and an anticipated technology shock, the contribution of news shocks to the forecast error variance of macroeconomic aggregates, and the historical decomposition of consumption, output, and hours worked. In the process, we contrast our findings with those obtained using BS and MS.

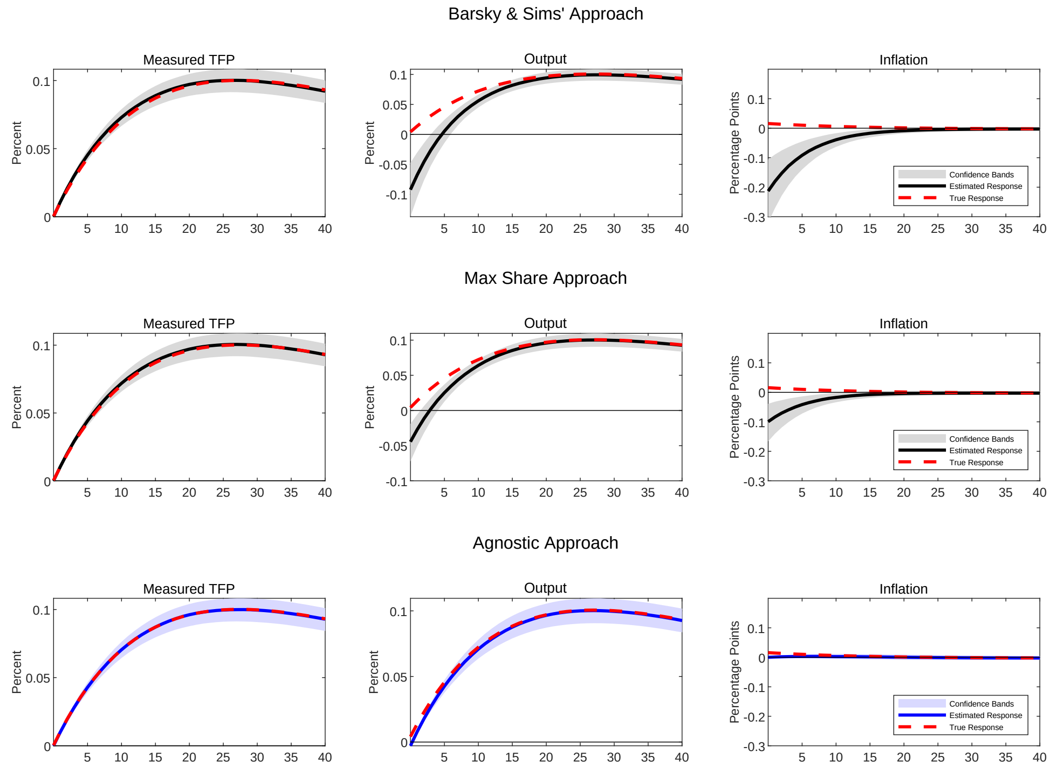

Figure 6: Impulse responses to a news shock based on artificial data and mis-measured TFP.

Notes: The figure shows the impulse responses to a news shock estimated within a three-equation VAR using artificial data generated from the basic New Keynesian model. TFP is measured by assuming constant returns to scale, but the data-generating model features increasing returns to scale. The series used in estimation are those of measured TFP, output, and inflation, each including 20000 observations. The solid lines are the median impulse responses to a new shock. The 68 percent confidence bands are based on 2000 draws. The top panels report the results based on Barsky & Sims (2011)' approach. The middle panels report the results based on the Max Share approach. The bottom panels report the results based on the agnostic approach. The dashed lines are the theoretical (true) responses.

4.1 Impulse responses

We start by discussing the estimated impulse responses to a surprise and an anticipated technology shock. To gauge these responses from the standpoint of New Keynesian theory, we compare them with those implied by the Smets & Wouters (2007) model. To do so, we again assume that TFP is described by process (1)–(2) and calibrate the parameters ρ_a and ρ_x and the size of the disturbances ϵ_t^s and ϵ_t^n so as to replicate as closely as possible the estimated (median) response of TFP to the surprise and the news shock. The confidence intervals around the estimated impulse responses are computed using Kilian (1998)'s bias-corrected bootstrap procedure.

4.1.1 Surprise shock

Our estimated impulse responses to a surprise technology shock are reported in the right panels of Figures 7 and 8 for the samples ending in 2007Q3 (i.e., used by Barsky & Sims (2011)) and 2019Q4, respectively. The left panels of each figure show the responses based on the reduced-form innovation to TFP (as in the empirical literature on news shocks).

Consider first the results based on the sample ending in 2007Q3. TFP increases on impact and remains persistently higher than its pre-shock level, a pattern that contrasts with the rapid return obtained when surprise technology shocks are identified as TFP innovations (shown in the upper left panel of Figure 7). Consumption and output also increase persistently and in a hump-shaped fashion. The estimated responses are remarkably similar to those implied by the Smets & Wouters (2007) model (particularly for consumption), and sharply contrast with the small, transitory, and rather monotonic pattern obtained from the identification scheme associating the shock with the TFP innovation.

Hours worked initially fall, then increase in a hump-shaped manner, converging to their pre-shock level from above. This pattern is consistent — at least qualitatively — with the prediction of the Smets & Wouters (2007) model, and differs from the muted reaction shown in the corresponding left panel. The result that unanticipated technological improvement has a contractionary effect on employment in the short run has been documented in several studies using different empirical approaches.[22]

Our estimated response for inflation is, by construction, restricted to be negative for the first ten quarters after the shock, as indicated by the shaded red area. Beyond that horizon, the inflation response becomes negligible and statistically insignificant. Interestingly, although our identification strategy does not impose a precise numerical value for the inflation response, the estimated response is strikingly similar to that implied by the Smets & Wouters (2007) model. The latter lies within the estimated confidence band at almost any given horizon.

Our identified surprise technology shock raises stock prices and consumer confidence. Stock prices are initially unresponsive but increase significantly and persistently at long horizons. The increase in consumer confidence is more transitory and is only statistically significant between the sixth and eighth quarters after the shock. These responses are at variance with the persistent decline in stock prices and consumer confidence shown in the left panels of Figure 7.

The results based on the extended sample, reported in Figure 8, convey similar messages. The shock has a long-lasting effect on TFP, consumption, and output. Hours worked fall significantly during the year following the shock, but their response is now statistically insignificant during the subsequent horizons. The inflation response to the surprise technology shock is negative by construction during the first ten quarters, and is virtually nil afterward. Stock prices rise but their response is statistically significant only at long horizons, while consumer confidence increases in the short run before returning to its pre-shock level.

In sum, these findings show that identifying surprise technology shocks as those explaining the largest fraction of the forecast error variance of TFP at the one-quarter horizon subject to a sign restriction on the response of inflation produces impulse responses that are more consistent

[22] See Galí & Rabanal (2005) for a survey.

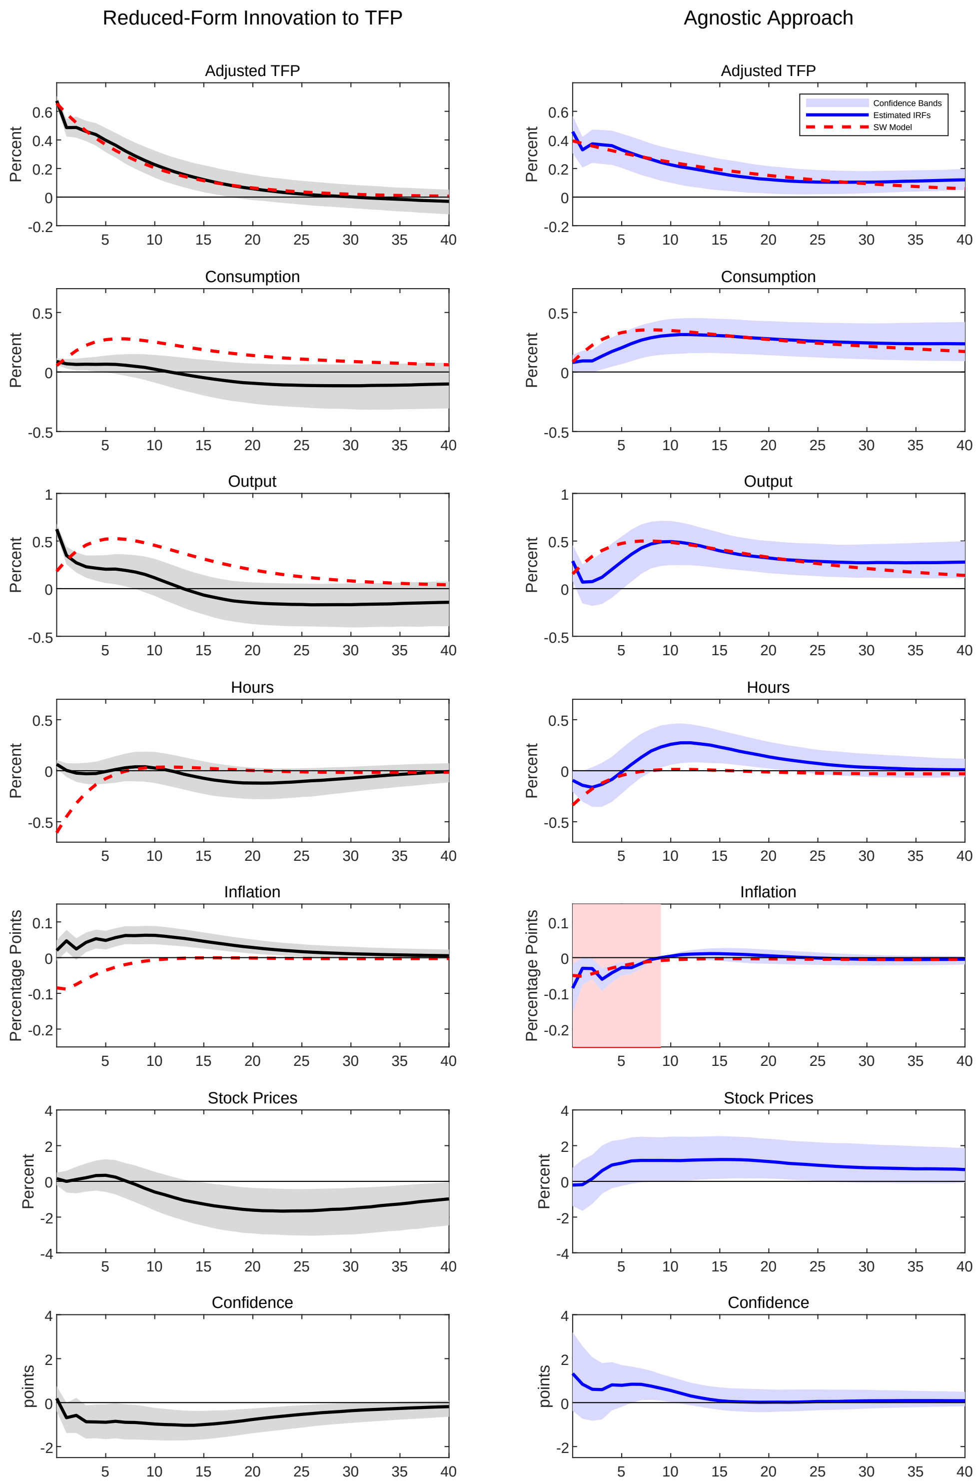

Figure 7: Impulse responses to a surprise technology shock. Sample: 1960Q1–2007Q3.

Notes: The figure shows the impulse responses to a surprise technology shock. The solid lines are the median impulse responses estimated based on the reduced-form innovation to TFP (left panels) and on the agnostic approach (right panels). The 68 percent confidence bands are the bias-corrected bootstrap confidence intervals computed using Kilian (1998)'s procedure with 2000 replications. The shaded red area indicates the horizons at which the inflation response is constrained to be negative. The dashed lines are the impulse responses obtained from the Smets & Wouters (2007) model.

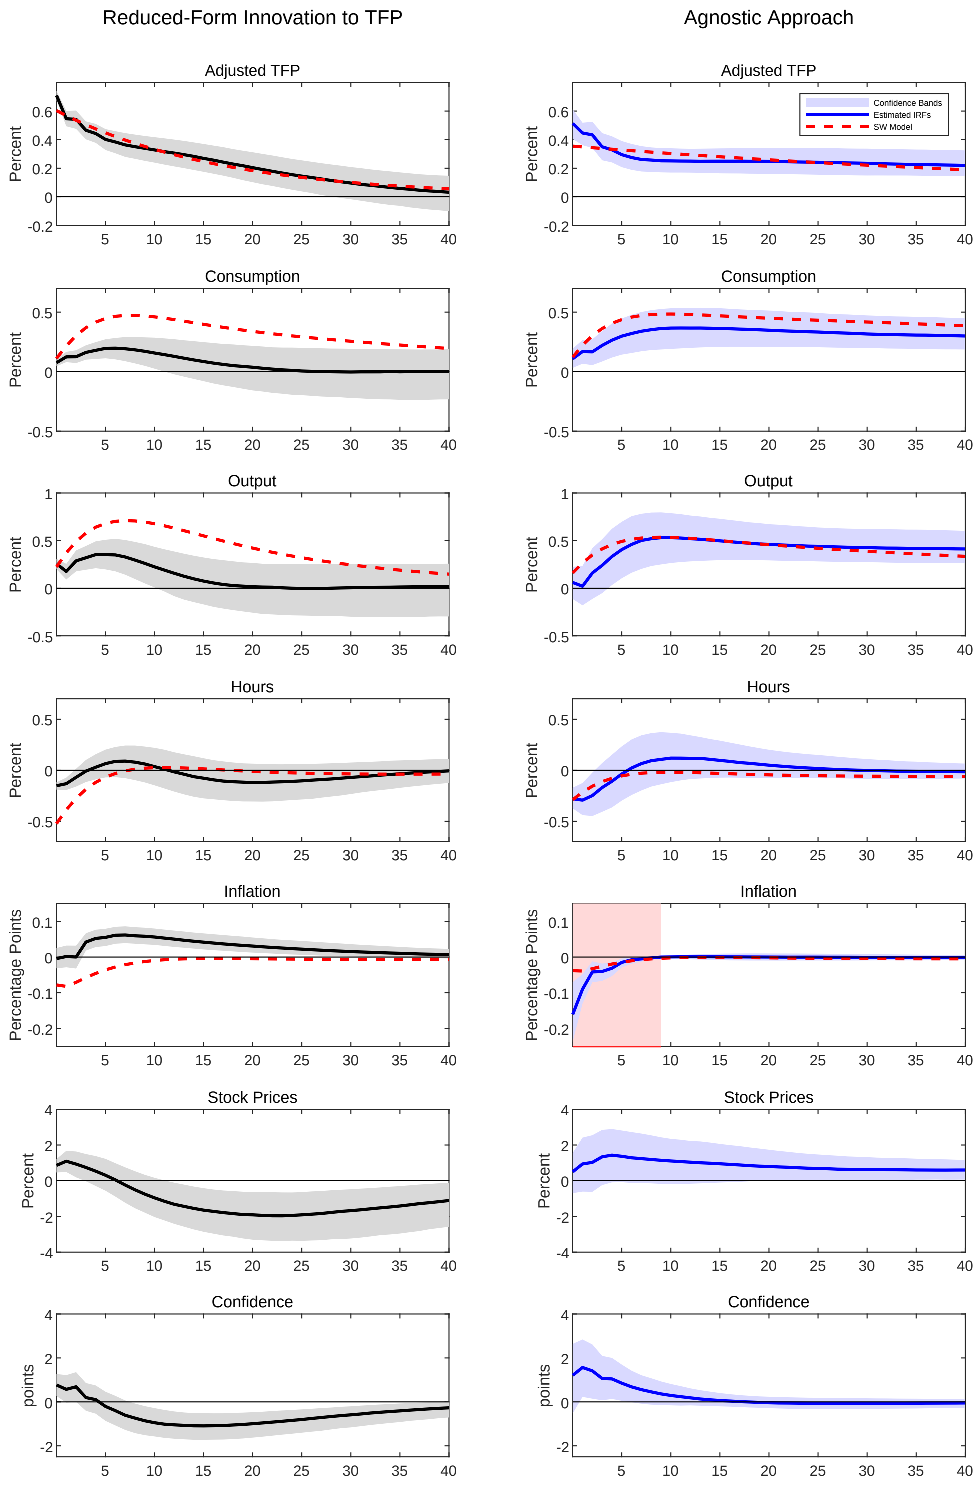

Figure 8: Impulse responses to a surprise technology shock. Sample: 1960Q1–2019Q4.

Notes: The figure shows the impulse responses to a surprise technology shock. The solid lines are the median impulse responses estimated based on the reduced-form innovation to TFP (left panels) and on the agnostic approach (right panels). The 68 percent confidence bands are the bias-corrected bootstrap confidence intervals computed using Kilian (1998)'s procedure with 2000 replications. The shaded red area indicates the horizons at which the inflation response is constrained to be negative. The dashed lines are the impulse responses obtained from the Smets & Wouters (2007) model.

with conventional wisdom and better grounded in theory than those obtained by using reduced-form innovations to TFP as a measure of surprise technology shocks. Interestingly, our estimated responses mimic remarkably well those implied by the Smets & Wouters (2007) model. The latter mostly lie within the confidence bands of the VAR-based responses.[23]

4.1.2 News shock

The estimated responses to a news shock are illustrated in the right panels of Figures 9 and 10 for the short and extended samples, respectively. Those based on BS and MS are shown in the left and middle panels of each figure, respectively.

Figure 9 shows that the response of TFP estimated using our agnostic strategy is similar in shape but smaller in magnitude than those implied by the alternative approaches. In line with the theoretical prediction, a favorable news shock about future productivity raises consumption significantly and persistently, regardless of the identification strategy. An important implication of both BS and MS is that output and hours worked initially decline in response to a positive realization of the news shock (see the third and fourth panels in the left and middle columns of Figure 9), an outcome that violates the predictions of the Smets & Wouters (2007) model. Both variables then rise persistently during the subsequent quarters, although the rise in hours is mostly statistically insignificant under BS. A similar pattern for hours is reported by Forni et al. (2014), Barsky et al. (2015), and Kurmann & Sims (2021).[24] The short-run contractionary effect of the news shock on aggregate output and hours worked no longer occurs, however, when we use our agnostic empirical methodology, as the output response is now statistically insignificant during the first two quarters after the shock, and that of hours worked is statistically indistinguishable from zero at any given horizon. In other words, we find no evidence of negative comovement between macroeconomic aggregates conditional on our identified news shock.[25]

As regards the response of inflation, both BS and MS imply that a favorable news shock about future technology decreases inflation sharply and persistently. This disinflationary effect, also

[23] As an additional check that the shock extracted using our methodology is technological in nature and does not capture other supply-side factors such as oil-price shocks, we compute the correlation between our identified shock and the innovation estimated by projecting changes in the price of crude oil (measured by the West Texas Intermediate spot price) on their past values, where the number of lags is selected based on the Schwarz criterion. The median correlation (across 2000 replications) is −0.057 in the sample ending in 2007Q3 and −0.044 in the sample ending in 2019Q4, indicating that the two shocks are essentially orthogonal to each other.

[24] Forni et al. (2014)'s approach is based on an estimated factor-augmented VAR in which the news shock is identified as the shock that best anticipates TFP at the 60-quarter horizon while being orthogonal to the reduced-form innovation in TFP. Barsky et al. (2015) identify the news shock as the innovation in the expectation of TFP at a fixed horizon in the future (20 quarters). Kurmann & Sims (2021) rely on the Max Share method (with $H = 80$) but without imposing the orthogonality of the news shock with respect to current TFP.

[25] On the other hand, even when estimated using our agnostic strategy, the responses of output and, especially, hours worked deviate markedly — in terms of magnitude — from those implied by the Smets & Wouters (2007) model. This discrepancy is attenuated but is still pronounced when we consider alternative preferences that mitigate the wealth effect of news shocks on households' labor supply (i.e., Jaimovich & Rebelo (2009)'s preferences). Bridging the gap between theory and empirical evidence along this dimension is an interesting challenge for future research.

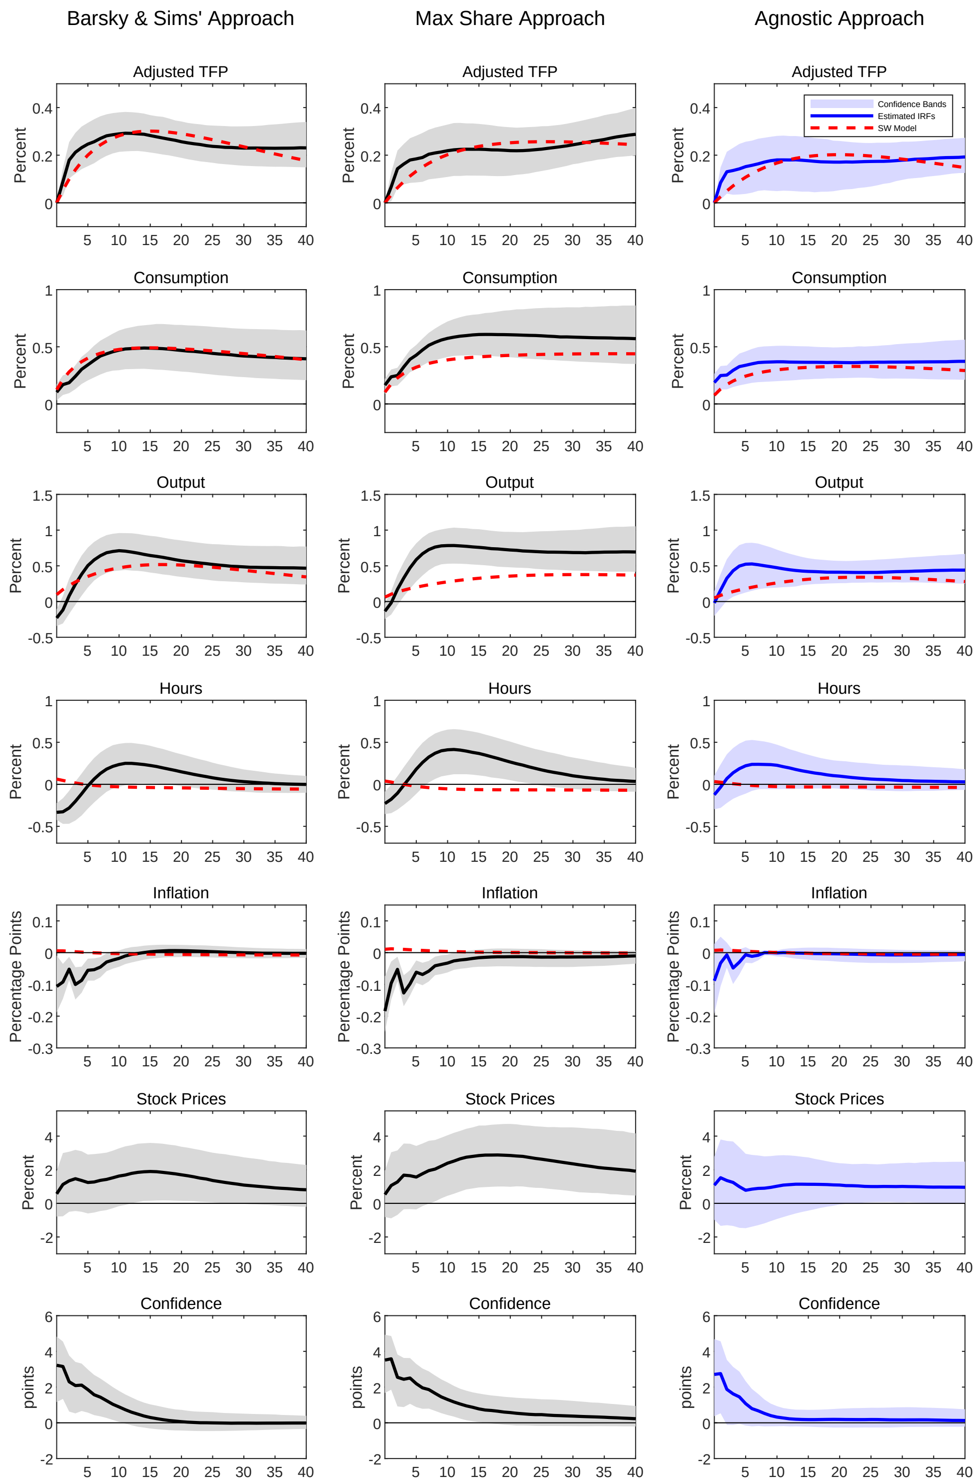

Figure 9: Impulse responses to a news shock. Sample: 1960Q1–2007Q3.

Notes: The figure shows the impulse responses to a news shock. The solid lines are the median impulse responses estimated based on Barsky and Sims' approach (left panels) and on the agnostic approach (right panels). The 68 percent confidence bands are the bias-corrected bootstrap confidence intervals computed using Kilian (1998)'s procedure with 2000 replications. The dashed lines are the impulse responses obtained from the Smets & Wouters (2007) model.

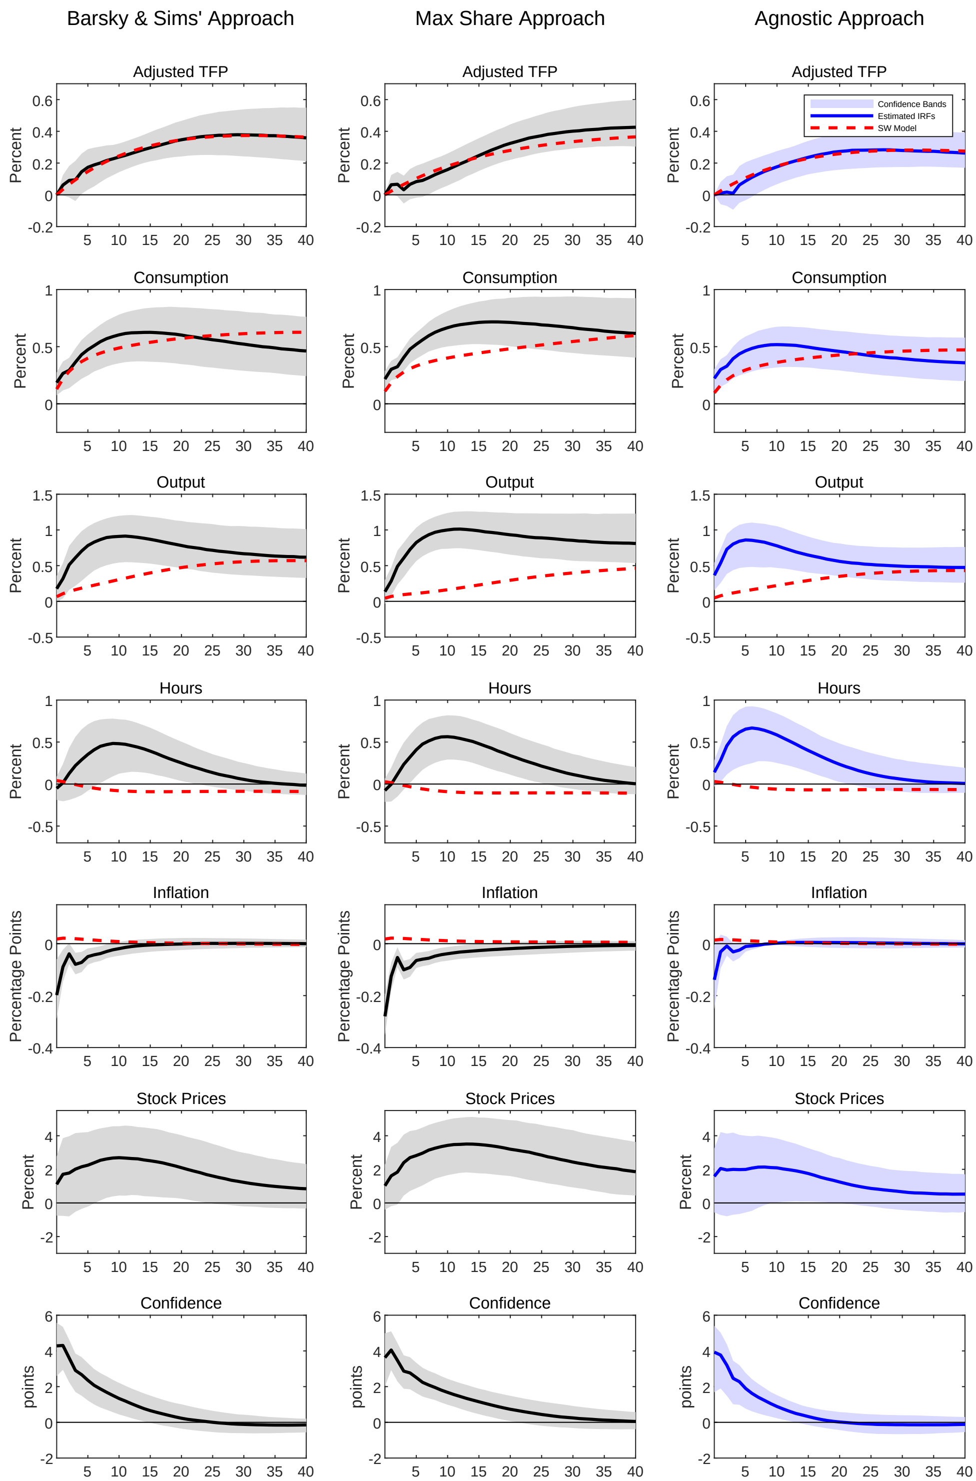

Figure 10: Impulse responses to a news shock. Sample: 1960Q1–2019Q4.

Notes: The figure shows the impulse responses to a news shock. The solid lines are the median impulse responses estimated based on Barsky and Sims' approach (left panels) and on the agnostic approach (right panels). The 68 percent confidence bands are the bias-corrected bootstrap confidence intervals computed using Kilian (1998)'s procedure with 2000 replications. The dashed lines are the impulse responses obtained from the Smets & Wouters (2007) model.

documented by Forni et al. (2014), Barsky et al. (2015), Fève & Guay (2019), and Kurmann & Sims (2021), is puzzling in light of New Keynesian theory, as pointed out by Barsky & Sims (2009), Jinnai (2013), and Kurmann & Otrok (2014). In the context of the basic New Keynesian model presented in Section 2.2.1, and assuming again that $\alpha = 1$, it is possible to show (using the method of undetermined coefficients) that the initial response of inflation to a news shock is given by

$$\frac{d\pi_t}{d\epsilon_t^n} = \frac{\sigma\lambda(1+\varphi)\left[\lambda(\sigma+\varphi)(\phi_\pi - 1) + (1-\beta)\phi_y - \beta\sigma(1-\rho_a)(1-\rho_x)\right]}{\Delta_a \Delta_x},$$

where $\Delta_x = \lambda(\sigma+\varphi)(\phi_\pi - \rho_x) + (1-\beta\rho_x)\left[\sigma(1-\rho_x) + \phi_y\right] > 0$. While the sign of the expression above is, in principle, ambiguous, it typically tends to be positive under sufficiently high values of ρ_a and ρ_x and a plausible calibration of the remaining parameters. Using the estimated values of ρ_a and ρ_x and the calibration discussed in Section 2.2.1, the basic New Keynesian model predicts a positive response of inflation to a favorable TFP news shock. The Smets & Wouters (2007) model also implies that inflation rises temporarily after a positive news shock but the response is tiny and essentially indistinguishable from 0 at any given horizon.[26]

This disinflation puzzle has prompted some researchers to suggest modifications to the prototype New Keynesian model so as to reconcile its predictions with existing empirical evidence.[27] Contrasting with this evidence, however, our results indicate that the inflation response to a news shock is rather muted and statistically insignificant at all horizons, consistently with the theoretical prediction. In other words, the disinflation puzzle vanishes under our agnostic identification strategy. The disinflationary effect documented in earlier studies appears to be caused by the mis-identification of anticipated technology shocks, due to measurement errors in TFP.[28]

Turning to the responses based on the sample ending in 2019Q4, the left column of Figure 10 shows that one of Barsky & Sims' main results, namely the contractionary effect of an anticipated technology shock on output and hours, disappears when we apply their identification strategy to the updated sample. A similar observation also holds under MS. In particular, both approaches predict that hours worked exhibit a statistically insignificant response during several quarters after the shock, before starting to increase. The rest of the responses are consistent with those based on the shorter sample. In particular, inflation falls significantly and persistently in response to a favorable news shock.

The results based on our agnostic identification strategy show important differences with respect to those implied by the alternative approaches as well as those based on the shorter sample. First,

[26] Görtz et al. (Forthcoming) show that the sign of the inflation response is more likely to be positive if the improvement in TFP is anticipated several quarters before it occurs.

[27] See, for instance, Jinnai (2013), Barsky et al. (2015), and Kurmann & Otrok (2014).

[28] Interestingly, using a different empirical methodology that identifies technological news shocks based on changes in firms' stock-market valuation due to announcements of patent grants, Cascaldi-Garcia & Vukotić (2022) also find no evidence of a disinflation puzzle or negative comovement between macroeconomic aggregates conditional on a news shock.

TFP exhibits an inertial response to the shock, starting to increase in a statistically significant manner only after about two years. This slowly diffusing process contrasts with the rapid increase in TFP estimated based on the shorter sample and using BS. Second, consumption, output, and hours worked increase significantly and persistently in response to the news shock. This simultaneous increase in macroeconomic aggregates — indicative of positive comovement — occurs well before TFP starts to rise; a result that corroborates Beaudry & Portier (2006)'s original findings. Interestingly, our estimated response for TFP, output, consumption, hours worked, and stock prices are extremely similar to those reported by Chahrour et al. (2020), who agnostically identify the shock that mainly drives the covariance between output and hours worked in the U.S.[29] Third, inflation falls in response to the shock but its response exhibits very little persistence and is (barely) statistically significant only on impact. In other words, the disinflation puzzle appears to be much less acute under our identification strategy.

To summarize, the decline in hours worked and inflation following a favorable news shock, typically documented in the literature, appear to be an artifact of using a polluted measure of technology, which impairs existing identification approaches. In fact, our results suggest that the news shocks identified using these approaches are in fact largely picking up the effects of *correctly identified* surprise shocks.[30]

4.2 Variance decomposition

Before evaluating the contribution of news shocks to the variability of macroeconomic aggregates, it is worth discussing the relative importance of the identified surprise technology shocks in explaining TFP. The results are reported in Table 1.[31] By construction, when surprise technology shocks are identified as reduced-form innovations to TFP, they explain *all* of the one-quarter-ahead forecast error variance of TFP (recall that the news shock does not affect TFP contemporaneously). Under our agnostic strategy, however, this need not be the case. In both sample periods, our identified surprise technology shocks account for roughly half of the one-quarter ahead forecast error variance

[29] Chahrour et al. (2020) use data that span the period 1966Q1–2018Q4.

[30] Kurmann & Sims (2021) point out that the results based on BS are highly sensitive to revisions in Fernald's adjusted TFP series. Using the same dataset as Barsky & Sims (2011), but replacing the 2007 vintage of TFP data with the 2016 vintage, they find, based on a four-variable VAR, that the response of hours worked to a favorable news shock is statistically insignificant during the first two quarters and positive thereafter. They interpret the absence of a negative response of hours as an indication of the presence of measurement errors. In an earlier version of this paper, we checked whether our agnostic identification approach is robust to revisions in measured TFP. To do so, we compared the impulse responses to a news shock based on the 2007 and 2016 vintages of TFP. The remaining series are those originally used by Barsky & Sims (2011). The estimated responses were found to be very similar across the two vintages, both in terms of shape and magnitude. In particular, we found no evidence of a negative response of hours worked or inflation. These findings confirm that our identification strategy is robust to revisions in TFP data, and substantiate the claim that news shocks are unlikely to give rise to negative comovement of aggregate variables and to disinflation.

[31] The contribution of the surprise shock to the forecast error variance of the remaining variables is reported in the Online Appendix.

of TFP, thus implying that non-technological shocks (potentially reflecting measurement errors) account for the remaining half, which in turn raises a serious objection against the interpretation of the estimated TFP series as a purified measure of technology.

Table 1: Share of Forecast Error Variance of TFP Attributed to Surprise Technology Shocks.

	Horizon					
	$h = 1$	$h = 4$	$h = 8$	$h = 16$	$h = 24$	$h = 40$
	Sample: 1960Q1–2007Q3					
Reduced-Form Innovation to TFP	1.000	0.876	0.782	0.631	0.532	0.407
Agnostic Approach	0.512	0.516	0.529	0.487	0.432	0.352
	Sample: 1960Q1–2019Q4					
Reduced-Form Innovation to TFP	1.000	0.886	0.804	0.708	0.589	0.401
Agnostic Approach	0.541	0.550	0.491	0.460	0.440	0.376

Note: The table reports the median fraction (across 2000 bootstrap replications) of the h-step ahead forecast error variance of TFP due to surprise technology shocks identified as reduced-form innovations to TFP, and using our agnostic approach.

Table 2 shows the contribution of news shocks to the h-step-ahead forecast error variance of the series used in estimation for both sample periods. The table also reports the results implied by BS and MS. In the sample ending in 2007Q3, our identified news shocks explain less than 5 percent of the conditional variance of TFP at the one-year horizon and roughly 25 percent at the ten-year horizon. They account for more than 40 percent of the forecast error variance of consumption, but barely 13 percent of the forecast error variance of output at the one-year horizon. The contribution of news shocks to output variability rises steadily with the forecasting horizon, reaching 28 percent at the ten-year horizon. For hours worked, inflation, stock prices, and consumer confidence, the share of the forecast error variance attributed to news shocks never exceeds 18 percent at any given horizon. Compared with the results based on BS and MS, we generally find a smaller contribution of the news shock to aggregate fluctuations at business-cycle frequencies.

Turning to the variance decomposition based on the extended sample, one of the striking differences with respect to the results based on the shorter sample and on BS and MS is that news shocks account for a relatively large fraction of the forecast error variance of output and hours worked at short horizons. At the one-year horizon, this fraction amounts to 41 percent for output and 25 percent for hours. In contrast, the corresponding numbers are 20 and 7 percent under BS, and 18 and 6 percent under MS. At business-cycle frequencies, the contribution of news shocks to the variability of consumption, output, and hours worked ranges roughly between 30 and 50 percent. Our agnostic approach continues to attribute a relatively small role to news shocks in accounting for the forecast error variance of inflation and stock prices in the updated sample, but these shocks appear to have become more important in explaining the variability of consumer confidence at all forecasting horizons.

Table 2: Share of Forecast Error Variance Attributed to News Shocks.

	1960Q1-2007Q3						1960Q1-2019Q4					
Horizon	$h = 1$	$h = 4$	$h = 8$	$h = 16$	$h = 24$	$h = 40$	$h = 1$	$h = 4$	$h = 8$	$h = 16$	$h = 24$	$h = 40$
Barsky & Sims' Approach												
TFP	0.000	0.068	0.159	0.311	0.397	0.460	0.000	0.022	0.073	0.195	0.324	0.479
Consumption	0.086	0.212	0.372	0.496	0.516	0.488	0.192	0.333	0.470	0.599	0.632	0.620
Output	0.081	0.076	0.204	0.393	0.438	0.436	0.053	0.200	0.354	0.522	0.564	0.573
Hours	0.424	0.174	0.129	0.161	0.165	0.160	0.040	0.072	0.126	0.219	0.236	0.230
Inflation	0.105	0.175	0.198	0.182	0.176	0.173	0.225	0.230	0.233	0.216	0.211	0.210
Stock Prices	0.040	0.069	0.082	0.116	0.138	0.142	0.079	0.109	0.149	0.206	0.218	0.220
Confidence	0.211	0.231	0.243	0.234	0.224	0.215	0.345	0.446	0.464	0.437	0.406	0.381
Max Share Approach												
TFP	0.000	0.043	0.095	0.186	0.260	0.399	0.000	0.012	0.024	0.096	0.233	0.473
Consumption	0.224	0.401	0.613	0.773	0.814	0.826	0.265	0.417	0.575	0.729	0.793	0.830
Output	0.030	0.071	0.268	0.506	0.601	0.679	0.035	0.182	0.380	0.599	0.678	0.742
Hours	0.201	0.075	0.123	0.245	0.283	0.282	0.045	0.058	0.131	0.270	0.304	0.300
Inflation	0.310	0.322	0.340	0.313	0.306	0.319	0.447	0.424	0.416	0.390	0.379	0.371
Stock Prices	0.038	0.075	0.107	0.195	0.271	0.342	0.053	0.096	0.178	0.298	0.352	0.382
Confidence	0.255	0.293	0.324	0.331	0.329	0.335	0.246	0.379	0.428	0.440	0.428	0.414
Agnostic Approach												
TFP	0.000	0.038	0.072	0.139	0.186	0.254	0.000	0.010	0.030	0.109	0.196	0.283
Consumption	0.278	0.399	0.423	0.358	0.338	0.324	0.280	0.425	0.484	0.486	0.454	0.413
Output	0.023	0.132	0.238	0.263	0.268	0.282	0.203	0.409	0.493	0.485	0.444	0.401
Hours	0.117	0.099	0.127	0.132	0.132	0.142	0.092	0.250	0.341	0.388	0.375	0.356
Inflation	0.097	0.102	0.090	0.081	0.085	0.099	0.120	0.102	0.092	0.085	0.087	0.098
Stock Prices	0.097	0.112	0.110	0.103	0.102	0.115	0.122	0.134	0.147	0.157	0.148	0.136
Confidence	0.156	0.183	0.171	0.155	0.154	0.160	0.293	0.364	0.359	0.321	0.294	0.284

Note: The table reports the median fraction (across 2000 bootstrap replications) of the h-step ahead forecast error variance of each variable due to news shocks identified using Barsky & Sims' approach (top panel), the Max Share approach (middle panel), and our agnostic approach (bottom panel).

4.3 Historical decomposition

In order to further investigate the importance of news shocks in accounting for business-cycle fluctuations, we simulate the time paths of consumption, output, and hours worked from the estimated VAR assuming that the news shocks are the only stochastic disturbances driving the data. The median simulations (across 2000 bootstrap replications) and the actual series are depicted in Figure 11, where the series are expressed in growth rates.

In the sample ending in 2007Q3 (upper panels), the correlation between the actual and simulated series is high for consumption (0.71) but fairly low for output and hours worked (0.22 and 0.12, respectively). News shocks appear to have played a very limited role in explaining post-war U.S. recessions, especially the 1969–1970, 1973–1975, and 1981–1982 recessions. Together with the variance-decomposition results discussed above, these observations lead us to conclude that news shocks are unlikely to have been a major driver of business-cycle fluctuations during the period 1960–2007. While this conclusion corroborates that reached by Barsky & Sims (2011), our argument for making such a claim differs from theirs. Indeed, Barsky & Sims (2011) base their conclusion on the fact that consumption co-moves negatively with output and hours worked in response to a news shock, a result that, as we have shown, is largely driven by measurement errors in TFP, just as the disinflationary effect of the shock. Instead, our conclusion is founded on the observation that news shocks explain only a modest fraction of the variability of output and hours worked at business-cycle frequencies.

The correlations between the actual and simulated series are larger in the extended sample (lower panels): 0.71, 0.32, and 0.35 for the growth rates of consumption, output, and hours worked, respectively. However, because these correlations are computed using different estimates across the two samples, the larger figures obtained in the extended sample need not reflect an increase in the contribution of news shocks to business-cycle fluctuations in recent years. To determine whether this is indeed the case, we compute the correlation between the actual and the simulated series in increasingly long sub-periods starting from 1960 (1960–1980 and 1960–2000, and 1960–2019) using the same estimates. The results are shown in Table 3. They indicate that the actual and simulated growth rates of output and, especially, hours worked have become more strongly correlated over time. A visual inspection of Figure 11 reveals that news shocks indeed account for a significant share of the decline in consumption, output, and hours worked during the recent U.S. downturns, including the Great Recession. These observations, along with the impulse-response and variance-decomposition results discussed above, indicate that news shocks have become an important driver of business-cycle fluctuations in recent years.

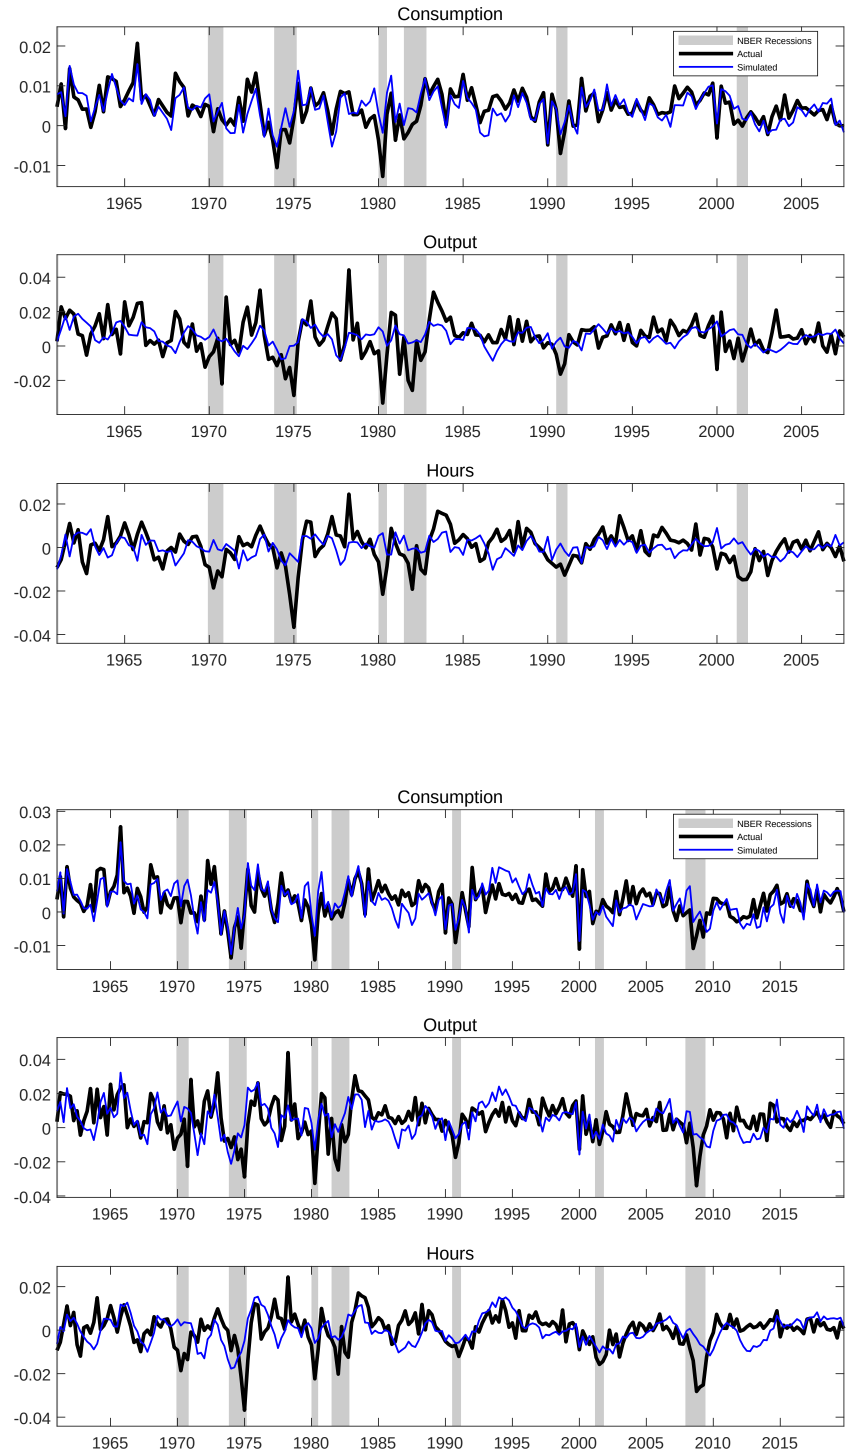

Figure 11: Historical decomposition.

Notes: The figure shows the actual series (thick black lines) and the ones simulated from the VAR assuming that news shocks are the only stochastic disturbances (thin blue lines). The simulated series are the median across 2000 bootstrap replications. The upper panels report the results for the 1960Q1–2007Q3 period while the lower panels report the results for the 1960Q1–2019Q4 period. The shaded areas indicate the dates of the U.S. recessions identified by the NBER.

Table 3: Correlations between the Actual and the Simulated (News Based) Series in Different Sub-Periods. Sample: 1960Q1–2019Q4.

	1960–1980	1960–2000	1960–2019
$\Delta \ln C_t$	0.761	0.714	0.710
$\Delta \ln Y_t$	0.276	0.308	0.325
$\Delta \ln N_t$	0.244	0.309	0.352

Notes: The table reports the correlations between the actual series and the series simulated under the assumption that news shocks are the only stochastic disturbances (medians across 2000 bootstrap replications). The variables C_t, Y_t, and N_t denote, respectively, consumption, output, and hours worked.

5 Conclusion

Much of the recent VAR-based evidence on the effects of news shocks about future productivity casts doubt on the plausibility and importance of TFP-news-driven business cycles, as these shocks are found to generate negative comovement between consumption and hours worked. Another robust finding of this literature is that favorable news shocks tend to be associated with a sharp and persistent decline in inflation.

In this paper, we have shown that these conclusions are spurious and are largely due to the presence of measurement errors in TFP. We have documented the severity of these errors by examining the effects of unanticipated technology shocks, usually identified as reduced-form innovations to TFP. We found these effects to be inconsistent with the interpretation of unanticipated technological disturbances as supply shocks. We have then proposed an agnostic identification strategy that is robust to measurement errors, successfully isolating the technological component of TFP. We found no evidence of negative comovement between consumption and hours worked conditional on a news shock, and the disinflation puzzle essentially disappears under our identification strategy. Importantly, we found that news shocks have become a major source of business-cycle fluctuations in recent years, consistently with Beaudry & Portier (2006)'s original view.

News about TFP, however, are clearly not the only factor that can cause changes in agents' expectations. Some recent studies have empirically examined the importance of changes in expectations caused by factors unrelated to TFP, such as news about investment-specific technology (e.g., Ben Zeev & Khan (2015)) or sentiments (e.g., Levchenko & Pandalai-Nayar (2015), Nam & Wang (2019), and Fève & Guay (2019)). The identification of these shocks, however, usually relies on the prior identification of TFP surprise and/or news shocks, which implies that the empirical approaches developed in this strand of the literature are also likely to be plagued by measurement errors in TFP. From this perspective, the empirical strategy developed in this paper can also help shed light on the relative importance of non-TFP news shocks for aggregate fluctuations.

References

Altig, D., Christiano, L., Eichenbaum, M. & Linde, J. (2011), 'Firm-Specific Capital, Nominal Rigidities and the Business Cycle', *Review of Economic Dynamics* **14**(2), 225–247.

Arezki, R., Ramey, V. A. & Sheng, L. (2017), 'News Shocks in Open Economies: Evidence from Giant Oil Discoveries', *The Quarterly Journal of Economics* **132**(1), 103–155.

Arias, J., Rubio-Ramirez, J. & Waggoner, D. (2018), 'Inference Based on SVARs Identified with Sign and Zero Restrictions: Theory and Applications', *Econometrica* **86**(2), 685–720.

Baron, J. & Schmidt, J. (2014), Technological Standardization, Endogenous Productivity and Transitory Dynamics, Technical report.

Barsky, R. B., Basu, S. & Lee, K. (2015), 'whither news shocks?', *in* J. A. Parker & M. Woodford, eds, 'NBER Macroeconomics Annual 2014', Vol. 29, University of Chicago Press, pp. 225–264.

Barsky, R. B. & Sims, E. R. (2009), 'News shocks', *NBER Working Paper Series* (15312).

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

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

Beaudry, P. & Lucke, B. (2010), 'letting different views about business cycles compete', *in* K. R. Daron Acemoglu & M. Woodford, eds, 'NBER Macroeconomics Annual 2009', Vol. 24, University of Chicago Press, pp. 413–455.

Beaudry, P., Nam, D. & Wang, J. (2011), 'Do mood swings drive business cycles and is it rational?', Technical report.

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

Beaudry, P. & Portier, F. (2014), 'News-driven business cycles: Insights and challenges', *Journal of Economic Literature* **52**(4), 993–1074.

Ben Zeev, N. & Khan, H. (2015), 'Investment-Specific News Shocks and U.S. Business Cycles', *Journal of Money, Credit and Banking* **47**(7), 1443–1464.

Cascaldi-Garcia, D. & Vukotić, M. (2022), ''Patent-based news shocks'', *The Review of Economics and Statistics* **104**(1), 51–66.

Chahrour, R., Chugh, S. & Potter, T. (2020), Anticipated Productivity and the Labor Market, Boston College Working Papers in Economics 992, Boston College Department of Economics.

Christiano, L. J., Eichenbaum, M. & Vigfusson, R. (2003), 'What Happens After a Technology Shock?', NBER Working Papers 9819, National Bureau of Economic Research, Inc.

Christiano, L. J., Eichenbaum, M. & Vigfusson, R. (2004), 'The response of hours to a technology shock: Evidence based on direct measures of technology', *Journal of the European Economic Association* **2**(2-3), 381–395.

Di Casola, P. & Sichlimiris, S. (2018), Towards Technology-News-Driven Business Cycles, Working Paper Series 360, Sveriges Riksbank (Central Bank of Sweden).

Edge, R. M., Laubach, T. & Williams, J. C. (2003), ''the responses of wages and prices to technology shocks", *Finance and Economics Discussion Series* (2003-65).

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

Fève, P. & Guay, A. (2010), 'Identification of technology shocks in structural vars', *The Economic Journal* **120**(549), 1284–1318.

Fève, P. & Guay, A. (2019), 'Sentiments in SVARs', *Economic Journal* **129**(618), 877–896.

Forni, M., Gambetti, L. & Sala, L. (2014), 'No news in business cycles', *The Economic Journal* **124**(581), 1168–1191.

Francis, N., Owyang, M. T., Roush, J. E. & DiCecio, R. (2014), 'A Flexible Finite-Horizon Alternative to Long-Run Restrictions with an Application to Technology Shocks', *The Review of Economics and Statistics* **96**(4), 638–647.

Fujiwara, I., Hirose, Y. & Shintani, m. (2011), 'Can news be a major source of aggregate fluctuations? a bayesian dsge approach', *Journal of Money, Credit and Banking* **43**(1), 1–29.

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

Galí, J. (2008), *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications*, Princeton University Press.

Galí, J. & Rabanal, P. (2005), 'technology shocks and aggregate fluctuations: How well does the real business cycle model fit postwar u.s. data?', *in* M. Gertler & K. Rogoff, eds, 'NBER Macroeconomics Annual 2004', Vol. 19, University of Chicago Press, pp. 225–317.

Görtz, C. & Tsoukalas, J. D. (2017), 'News and Financial Intermediation in Aggregate Fluctuations', *The Review of Economics and Statistics* **99**(3), 514–530.

Görtz, C., Tsoukalas, J. D. & Zanetti, F. (Forthcoming), 'News Shocks under Financial Frictions', *American Economic Journal: Macroeconomics* .

Ireland, P. N. (2004), 'Technology Shocks in the New Keynesian Model', *The Review of Economics and Statistics* **86**(4), 923–936.

Jaimovich, N. & Rebelo, S. (2009), 'Can news about the future drive the business cycle?', *American Economic Review* **99**(4), 1097–1118.

Jinnai, R. (2013), 'News shocks and inflation', *Economics Letters* **119**(2), 176 – 179.

Karnizova, L. (2012), 'News Shocks, Productivity and the U.S. Investment Boom-Bust Cycle', Technical report.

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

Kilian, L. (1998), 'Small sample confidence intervals for impulse response functions', *Review of Economics and Statistics* **80**(2), 218–230.

Kurmann, A. & Mertens, E. (2014), 'Stock Prices, News, and Economic Fluctuations: Comment', *American Economic Review* **104**(4), 1439–1445.

Kurmann, A. & Otrok, C. (2014), 'News shocks and inflation: Lessons for new keynesians', *The Review of Economics and Statistics* .

Kurmann, A. & Sims, E. (2021), 'Revisions in Utilization-Adjusted TFP and Robust Identification of News Shocks', *The Review of Economics and Statistics* **103**(2), 216–235.

Levchenko, A. A. & Pandalai-Nayar, N. (2015), 'TFP, News, and "Sentiments": The International Transmission of Business Cycles', NBER Working Papers 21010, National Bureau of Economic Research, Inc.

Nam, D. & Wang, J. (2019), 'Mood Swings and Business Cycles: Evidence from Sign Restrictions', *Journal of Money, Credit and Banking* **51**(6), 1623–1649.

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

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

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

Uhlig, H. (2005), 'What are the effects of monetary policy on output? Results from an agnostic identification procedure', *Journal of Monetary Economics* **52**(2), 381–419.