Estimating Macroeconomic News and Surprise Shocks*

Lutz Kilian† Michael D. Plante‡ Alexander W. Richter§

July 23, 2024

Abstract

The importance of understanding the economic effects of news and surprise shocks to TFP is widely recognized in the literature. A common VAR approach is to identify responses to TFP news shocks by maximizing the variance share of TFP over a long horizon. Under suitable conditions, this approach also implies an estimate of the surprise shock. We find that these TFP max share estimators tend to be strongly biased when applied to data generated from DSGE models with shock processes that match the TFP moments in the data, both in the presence of TFP measurement error and in its absence. Incorporating a measure of TFP news into the VAR model and adapting the identification strategy substantially reduces the bias and RMSE of the impulse response estimates, even when there is sizable measurement error in the news variable. When applying this method to the data, we find that news shocks are slower to diffuse to TFP and have a smaller effect on real activity than implied by the TFP max share method.

Keywords: Structural VAR; TFP; news; anticipated shocks; measurement error; max share; IV
JEL Classifications: C32, C51, C61, E32

*We thank Filippo Ferroni for discussing our paper at the 2023 System Econometrics Meeting and Andrew Fieldhouse for discussing our paper at the 2024 American Economics Association Meeting. We also thank Francesco Bianchi, Danilo Cascaldi-Garcia, Oliver Coibion, Sinem Hacıoğlu Hoke, Christian Matthes, and Todd Walker for helpful comments. The views expressed in this paper are our own and do not necessarily reflect the views of the Federal Reserve Bank of Dallas or the Federal Reserve System.

†Federal Reserve Bank of Dallas, 2200 N Pearl Street, Dallas, TX 75201, and CEPR (lkilian2019@gmail.com).
‡Federal Reserve Bank of Dallas, 2200 N Pearl Street, Dallas, TX 75201 (michael.plante@dal.frb.org).
§Federal Reserve Bank of Dallas, 2200 N Pearl Street, Dallas, TX 75201 (alex.richter@dal.frb.org).

1 INTRODUCTION

There is considerable interest in understanding the economic effects of shocks to expectations about future economic activity dating back to Pigou (1927).[1] Such news shocks have received particular attention in studies that explore the effects of shocks to total factor productivity (TFP) on macroeconomic aggregates, starting with Beaudry and Portier (2006).

A common approach to identifying an anticipated shock to TFP ("news shock") is to use the max share estimator popularized by Uhlig (2003, 2004), Barsky and Sims (2011), and Francis et al. (2014), among others. This estimator identifies the news shock by selecting parameters for the structural impact multiplier matrix of a vector autoregressive (VAR) model to maximize the forecast error variance shares of TFP over a long horizon. Under suitable conditions, this estimator also implies an estimate of the unanticipated shock to TFP ("surprise shock"). We will refer to this estimator as the "TFP max share" estimator. This class of estimators continues to be widely applied in empirical work, and studies using this estimator have given rise to theoretical work on news and surprise shocks (e.g., Bretscher et al., 2021; Chahrour and Jurado, 2018; Faccini and Melosi, 2022). Variations of this approach have also been applied in other economic contexts.[2]

Early applications of the TFP max share estimator imposed the restriction that the news shock is orthogonal to current TFP, which can be traced to Cochrane (1994) and Beaudry and Portier (2006). A further refinement of the TFP max share estimator was introduced by Kurmann and Sims (2021), who relaxed this exclusion restriction to allow for measurement error in TFP due to the fact that factor utilization is unobserved in the data and has to be estimated.[3] Their estimator also accounts for the fact that new technologies may affect TFP immediately, even though their effect on TFP may take many years to build up due to the slow diffusion of new technologies.

[1] See Beaudry and Portier (2014) for a review of the literature on news-driven business cycles.

[2] Recent examples include Forni et al. (2014), Barsky et al. (2015), Ben Zeev and Khan (2015), Chen and Wemy (2015), Nam and Wang (2015), Ben Zeev et al. (2017), Fève and Guay (2019), Angeletos et al. (2020), Levchenko and Pandalai-Nayar (2020), Cascaldi-Garcia and Galvao (2021), Dieppe et al. (2021), Kurmann and Sims (2021), Benhima and Cordonier (2022), Francis and Kindberg-Hanlon (2022), Görtz et al. (2022a), Görtz et al. (2022b), Bouakez and Kemoe (2023), Miyamoto et al. (2023), and Carriero and Volpicella (2024).

[3] Christiano et al. (2004) and Bouakez and Kemoe (2023) also discuss the ramifications of TFP measurement error.

It is widely believed that TFP max share estimators of news shocks work well as long as news shocks account for the bulk of the variation in TFP at long horizons. Our first contribution is to show that this condition is not sufficient to ensure the accuracy of this estimator. We begin by examining the accuracy of three variants of the TFP max share estimator in the ideal setting when there is no TFP measurement error. The data are simulated from a conventional dynamic stochastic general equilibrium (DSGE) model. Our simulation results demonstrate that, even when virtually all variation in TFP at a long horizon is explained by news shocks, the TFP max share estimator may fail to recover the responses to news shocks, regardless of the sample size. In practice, the accuracy of the estimator not only depends on the quantitative importance of news shocks at long horizons but also at short horizons. We then show that these results also hold when using a larger-scale DSGE model, which allows the simulated TFP data to be contaminated by measurement error as in Kurmann and Sims (2021). This evidence suggests that the TFP max share estimator is not well-identified and raises the question of what alternative methods are available to applied researchers.

Our second contribution is to show that adding a direct measure of TFP news to the VAR model and adapting the identification strategy, as suggested in some recent empirical studies, will substantially reduce the asymptotic bias.[4] We discuss two such identification strategies. One is based on maximizing the variance share of the news variable at a short horizon (as opposed to the variance share of TFP at a long horizon) and is new to the literature. The other treats news variable innovations as predetermined as in Alexopoulos (2011) and Cascaldi-Garcia and Vukotić (2022).

While we are not the first to employ direct measures of TFP news for identifying news shocks, we are the first to examine the ability of these estimators to recover the population responses from data generated by DSGE models. We first show that appropriately constructed estimators based on news variables have much lower bias and root mean squared error (RMSE) than the TFP max share estimator in the absence of TFP measurement error. We then evaluate the news-based estimators in

[4]Examples of studies employing measures of TFP news include Shea (1999), Christiansen (2008), Alexopoulos (2011), Baron and Schmidt (2019), Cascaldi-Garcia and Vukotić (2022), Miranda-Agrippino et al. (2022), and Fieldhouse and Mertens (2023).

the presence of TFP measurement error and show that these methods still substantially reduce the bias and RMSE of the responses to news shocks compared to the TFP max share estimator. While TFP news is not perfectly observed in the data, the superior accuracy of these estimators is robust to sizable measurement error in the news variable.

A fundamental concern with using the TFP max share estimator is that households are likely to account for TFP news in forming expectations about future TFP. This information typically has been omitted from the VAR information set in the literature. It is well-known that the identification of structural stocks in the VAR requires these information sets to be aligned (see, e.g., Hansen and Sargent, 1991; Leeper et al., 2013). We present evidence that this problem accounts for some, but not all of the asymptotic bias of the TFP max share estimator of the impulse responses. While we find that applying the TFP max share method to a VAR model that also includes a direct measure of TFP news alleviates this bias, the two news-based identification strategies are even less biased and have lower RMSE, especially in the presence of TFP measurement error. This evidence indicates that the bias of the TFP max share estimator cannot be explain by an information deficiency alone.

Our third contribution is to empirically illustrate the use of TFP news for identifying news shocks using a range of TFP news measures that have been used in the literature. We first show that two of these news measures generate plausible results in light of the underlying economic theory. Both yield impulse response estimates that are systematically and substantially different from the estimates generated by the TFP max share method, consistent with our simulation results.

We then reexamine the question of whether these shocks are an important driver of TFP and real activity. There are conflicting views in the literature about how quickly news shocks diffuse to TFP and about the extent to which they drive macroeconomic aggregates.[5] We find that news shocks are slow to diffuse to TFP, but have a more immediate effect on real activity, explaining 24% of the fluctuations in output at a five-year horizon. In the long-run, the share of the forecast error variance explained by news shocks is 24% for TFP and 36% for output. In contrast, the

[5] See, for example, Beaudry and Lucke (2010), Barsky and Sims (2011), Forni et al. (2014), Barsky et al. (2015), Fève and Guay (2019), Levchenko and Pandalai-Nayar (2020), Cascaldi-Garcia and Vukotić (2022), Görtz et al. (2022b), Miranda-Agrippino et al. (2022), and Bouakez and Kemoe (2023).

estimates based on the TFP max share estimator not only imply that news shocks quickly diffuse to TFP, but also that they explain 63% of the forecast error variance of output at a one-year horizon and almost 90% at horizons beyond five years.

The remainder of the paper is organized as follows. In Section 2, we review the estimation of news shocks obtained by maximizing the contribution of the news shock to the forecast error variance of TFP at long, but finite horizons, and derive the identification conditions for the surprise shock. In Section 3, we use data generated from a conventional DSGE model to examine the accuracy of the TFP max share estimator in the absence of TFP measurement error. In Section 4, we enlarge the DSGE model and allow for TFP measurement error. In Section 5, we use these DSGE models to examine the accuracy of two alternative identification strategies that accommodate TFP measurement error by including a direct measure of TFP news in the VAR model. In Section 6, we examine the empirical importance of news shocks in a range of VAR models based on alternative measures of TFP news and compare the results to those obtained using the TFP max share estimator. Section 7 contains the concluding remarks.

2 IDENTIFICATION PROBLEM

This section describes the TFP max share estimator for identifying news shocks. It also explains the conditions under which this method simultaneously identifies surprise TFP shocks.

2.1 NOTATION Consider a VAR model with K variables. Let $\mathbf{y}_t$ be a $K \times 1$ vector of variables. The reduced-form moving average representation of the VAR model is given by $\mathbf{y}_t = \Phi(L)\mathbf{u}_t$, where $\Phi(L) = I_K + \Phi_1 L + \Phi_2 L^2 + \cdots$, I_K is a K-dimensional identity matrix, L is a lag operator, and $\mathbf{u}_t$ is a $K \times 1$ vector of reduced-form shocks with variance-covariance matrix $\Sigma = E[\mathbf{u}_t\mathbf{u}_t']$.

Let $\mathbf{w}_t$ be a $K \times 1$ vector of structural shocks with $E[\mathbf{w}_t\mathbf{w}_t'] = I_K$. Under suitable normalizing assumptions, $\mathbf{u}_t = B_0^{-1}\mathbf{w}_t$, where the $K \times K$ structural impact multiplier matrix B_0^{-1} satisfies $B_0^{-1}(B_0^{-1})' = \Sigma$. The impact effect of shock j on variable i is given by the jth column and the ith row of B_0^{-1}. Let P denote the lower triangular Cholesky decomposition of Σ with the diagonal elements normalized to be positive, and let Q be a $K \times K$ orthogonal matrix. Since $Q'Q = QQ' = I_K$ and hence $(PQ)(PQ)' = PP' = \Sigma$, we can express the set of possible solutions for B_0^{-1} as PQ. Identification involves pinning down some or all columns of Q.

One way of proceeding is to observe that the h-step ahead forecast error is given by

$$\mathbf{y}_{t+h} - E_{t-1}\mathbf{y}_{t+h} = \sum_{\tau=0}^{h} \Phi_\tau PQ\mathbf{w}_{t+h-\tau},$$

where Φ_τ is the reduced-form matrix for the moving average coefficients, which may be constructed following Kilian and Lütkepohl (2017) with $\Phi_0 = I_K$. As a result, the share of the forecast error variance of variable i that is attributed to shock j at horizon h is given by

$$\Omega_{i,j}(h) = \frac{\sum_{\tau=0}^{h} \Phi_{i,\tau} P\gamma_j\gamma_{j'} P'\Phi_{i,\tau}'}{\sum_{\tau=0}^{h} \Phi_{i,\tau}\Sigma\Phi_{i,\tau}'},$$

where $\Phi_{i,\tau}$ is the ith row of the lag polynomial at lag τ and γ_j is the jth column of Q. A unique estimate of the impact effect of structural shock j may be obtained by choosing the values of γ_j to maximize $\Omega_{i,j}(h)$ for some horizon h (or its average over selected horizons).

2.2 TFP MAX SHARE ESTIMATOR

For expository purposes, consider a stylized VAR model of the effects of shocks to TFP with $K = 3$. Without loss of generality, the TFP variable is ordered first. The orthogonal rotation matrix is given by

$$Q = \begin{pmatrix} \gamma_{s,1} & \gamma_{n,1} & \gamma_{\ell,1} \\ \gamma_{s,2} & \gamma_{n,2} & \gamma_{\ell,2} \\ \gamma_{s,3} & \gamma_{n,3} & \gamma_{\ell,3} \end{pmatrix}, \tag{1}$$

where $\gamma_{s,j}$ and $\gamma_{n,j}$ are elements associated with the impact of the surprise and news shock, respectively, on variable $j \in \{1, 2, 3\}$. $\gamma_{\ell,j}$ are the elements associated with an unnamed third shock.

We consider three variants of the TFP max share estimator. One variant proposed by Kurmann and Sims (2021), which we will refer to as the KS estimator, is based on

$$\gamma_n = \text{argmax } \Omega_{1,2}(H_n), \quad \Omega_{1,2}(H_n) \equiv \frac{\sum_{\tau=0}^{H_n} \Phi_{1,\tau} P\gamma_n\gamma_n' P'\Phi_{1,\tau}'}{\sum_{\tau=0}^{H_n} \Phi_{1,\tau}\Sigma\Phi_{1,\tau}'}, \tag{2}$$

subject to the restriction that $\gamma_n'\gamma_n = 1$, where $\gamma_n = (\gamma_{n,1}, \gamma_{n,2}, \gamma_{n,3})'$ denotes the second column of Q and H_n denotes the maximimum horizon.

An earlier variant of the TFP max share estimator proposed by Barsky and Sims (2011), which we will refer to as the BS estimator, imposes that news shocks do not affect TFP on impact and solves for the γ_n that maximizes the sum of the forecast error variance shares of TFP over a long horizon, rather than the variance share over a particular long horizon. Specially, the news shock estimate is based on

$$\gamma_n = \text{argmax} \sum_{h=0}^{H_n} \Omega_{1,2}(h), \quad \Omega_{1,2}(h) \equiv \frac{\sum_{\tau=0}^{h} \Phi_{1,\tau} P \gamma_n \gamma_n' P' \Phi_{1,\tau}'}{\sum_{\tau=0}^{h} \Phi_{1,\tau} \Sigma \Phi_{1,\tau}'}, \tag{3}$$

subject to $\gamma_{n,1} = 0$ and $\gamma_n'\gamma_n = 1$.

One concern with the BS estimator is that less persistent shocks may affect the estimates of the news shock, which prompted Kurmann and Sims (2021) to remove the cumulative sum from the objective function. Dieppe et al. (2021) go a step further and propose a third variant of the TFP max share estimator, referred to as the non-accumulated max share (NAMS) estimator. When adapted to our context, this estimator maximizes the squared TFP response to the news shock at H_n. In this case, the news shock estimate is based on

$$\gamma_n = \text{argmax} \frac{\Phi_{1,H_n} P \gamma_n \gamma_n' P' \Phi_{1,H_n}'}{\Phi_{1,H_n} \Sigma \Phi_{1,H_n}'}, \tag{4}$$

subject to the restriction that $\gamma_n'\gamma_n = 1$. Throughout the paper, we follow Kurmann and Sims (2021) in setting $H_n = 80$.[6]

2.3 Identification Conditions

As long as TFP innovations are fully explained by news and surprise shocks, as would be the case in the absence of TFP measurement error, it has to be the case that $\gamma_{\ell,1} = 0$. Whether one imposes this restriction does not affect the estimate of the news shock, but it determines whether the surprise shock can also be identified. To formalize this result, assume $\gamma_{\ell,1} = 0$, and note that the Q matrix is orthogonal if and only if $Q'Q = QQ' = I_3$. This

[6] Although Barsky and Sims (2011) set $H_n = 40$, raising H_n has little effect on the results for the BS estimator.

yields the restrictions

$$\begin{pmatrix} \gamma_{s,1} & \gamma_{s,2} & \gamma_{s,3} \\ \gamma_{n,1} & \gamma_{n,2} & \gamma_{n,3} \\ 0 & \gamma_{\ell,2} & \gamma_{\ell,3} \end{pmatrix} \begin{pmatrix} \gamma_{s,1} & \gamma_{n,1} & 0 \\ \gamma_{s,2} & \gamma_{n,2} & \gamma_{\ell,2} \\ \gamma_{s,3} & \gamma_{n,3} & \gamma_{\ell,3} \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}, \tag{R1}$$

$$\begin{pmatrix} \gamma_{s,1} & \gamma_{n,1} & 0 \\ \gamma_{s,2} & \gamma_{n,2} & \gamma_{\ell,2} \\ \gamma_{s,3} & \gamma_{n,3} & \gamma_{\ell,3} \end{pmatrix} \begin{pmatrix} \gamma_{s,1} & \gamma_{s,2} & \gamma_{s,3} \\ \gamma_{n,1} & \gamma_{n,2} & \gamma_{n,3} \\ 0 & \gamma_{\ell,2} & \gamma_{\ell,3} \end{pmatrix} = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix}. \tag{R2}$$

Restriction R1 implies

$$\gamma_{n,1}^2 + \gamma_{n,2}^2 + \gamma_{n,3}^2 = 1, \tag{R1-1}$$

$$\gamma_{s,2}\gamma_{\ell,2} + \gamma_{s,3}\gamma_{\ell,3} = 0, \tag{R1-2}$$

$$\gamma_{n,2}\gamma_{\ell,2} + \gamma_{n,3}\gamma_{\ell,3} = 0, \tag{R1-3}$$

$$\gamma_{s,1}^2 + \gamma_{s,2}^2 + \gamma_{s,3}^2 = 1, \tag{R1-4}$$

$$\gamma_{\ell,2}^2 + \gamma_{\ell,3}^2 = 1, \tag{R1-5}$$

$$\gamma_{s,1}\gamma_{n,1} + \gamma_{s,2}\gamma_{n,2} + \gamma_{s,3}\gamma_{n,3} = 0. \tag{R1-6}$$

Restriction R2 implies

$$\gamma_{s,1}^2 + \gamma_{n,1}^2 = 1, \tag{R2-1}$$

$$\gamma_{s,1}\gamma_{s,2} + \gamma_{n,1}\gamma_{n,2} = 0, \tag{R2-2}$$

$$\gamma_{s,1}\gamma_{s,3} + \gamma_{n,1}\gamma_{n,3} = 0, \tag{R2-3}$$

$$\gamma_{s,2}^2 + \gamma_{n,2}^2 + \gamma_{\ell,2}^2 = 1, \tag{R2-4}$$

$$\gamma_{s,3}^2 + \gamma_{n,3}^2 + \gamma_{\ell,3}^2 = 1, \tag{R2-5}$$

$$\gamma_{s,2}\gamma_{s,3} + \gamma_{n,2}\gamma_{n,3} + \gamma_{\ell,2}\gamma_{\ell,3} = 0. \tag{R2-6}$$

An estimate of γ_n is obtained by maximizing the forecast error variance share of the news shock subject to (R1-1). Given γ_n, (R2-1)-(R2-5) imply

$$\gamma_{s,1} = \pm\sqrt{1-\gamma_{n,1}^2}, \quad \gamma_{s,2} = -\frac{\gamma_{n,1}\gamma_{n,2}}{\gamma_{s,1}}, \quad \gamma_{s,3} = -\frac{\gamma_{n,1}\gamma_{n,3}}{\gamma_{s,1}},$$
$$\gamma_{\ell,2} = \pm\sqrt{1-\gamma_{s,2}^2-\gamma_{n,2}^2}, \quad \gamma_{\ell,3} = \pm\sqrt{1-\gamma_{s,3}^2-\gamma_{n,3}^2}.$$

Thus, for $K = 3$ the identifying restrictions uniquely identify all three structural response functions up to their sign. This means that all that is required to recover the news and surprise shocks is a normalizing assumption to the effect that the surprise shock has a positive impact effect on TFP and the news shock has a positive effect on TFP at H_n. For $K > 3$ only the news and surprise shocks are identified. This result means it is sufficient to compare the explanatory power of both TFP shocks without having to take a stand on the identification of the other $K-2$ structural shocks.

While our example is for $K = 3$, the following proposition shows that without TFP measurement error the TFP max share estimator of the news shock always identifies the surprise shock.

Proposition 1. *In the absence of TFP measurement error, γ_s will be uniquely identified for any given estimate of γ_n obtained using the TFP max share estimator. In particular, when TFP is ordered first in the VAR model, $\gamma_{s,1} = \pm\sqrt{1-\gamma_{n,1}^2}$ and $\gamma_{s,j} = -\gamma_{n,1}\gamma_{n,j}/\gamma_{s,1}$ for $j \in \{2,\dots,K\}$.*

The proof immediately follows from a generalization of the analysis for $K = 3$. Note that there are multiple solutions for Q, some of which will satisfy R1 and R2 and some of which may not. For $K = 3$, for example, there are 2^3 possible solutions. The validity of the estimator requires the existence of an orthogonal Q matrix. In Appendix B, we show that when solving for γ_n and γ_s, γ_ℓ can always be chosen such that Q is orthogonal. This result generalizes to $K > 3$.

Our analysis highlights that the TFP max share estimator will be able to deliver estimates of the surprise shock even when there is no restriction on $\gamma_{n,1}$, as long as there is no TFP measurement error. This allows us to shed light on the ability of the TFP max share estimator to recover the population responses to news and surprise shocks under ideal conditions without TFP measurement error. If the estimator does not work in this setting, it would not be expected to work in the more realistic setting when TFP has to be recovered from the observed data by removing the estimated factor utilization, creating TFP measurement error.

3 Accuracy of the TFP Max Share Estimator

3.1 Data Generating Process We begin our examination of the TFP max share estimator by focusing on the ideal setting where there is no TFP measurement error, so $\gamma_{\ell,1} = 0$. For this purpose, we simulate data from a conventional New Keynesian model (henceforth, the "baseline model"). The advantage of starting with a simple model is that it allows us to highlight that our results do not depend on any nonstandard model features. In Section 4, we will enlarge this model in order to examine the implications of TFP measurement error.

Households The representative household solves

$$J_t = \max_{c_t, n_t, b_t, i_i, k_t} \log c_t - \chi n_t^{1+\eta}/(1+\eta) + \beta E_t J_{t+1}$$

subject to

$$c_t + i_t + b_t = w_t n_t + r_t^k k_{t-1} + r_{t-1} b_{t-1}/\pi_t + d_t,$$

$$k_t = (1-\delta) k_{t-1} + \mu_t i_t,$$

where $\beta \in (0,1)$ is the subjective discount factor, $\chi > 0$ is a preference parameter, $1/\eta$ is the Frisch elasticity of labor supply, c_t is consumption, n_t is labor hours, b_t is the real value of a privately-issued one-period nominal bond, i_t is investment, k_t is the stock of capital that depreciates at rate δ, r_t^k is the real rental rate of capital, w_t is the real wage rate, d_t is real dividends rebated from intermediate goods firms, $\pi_t = p_t/p_{t-1}$ is the gross inflation rate, r_t is the gross nominal interest rate set by the central bank, and μ_t is an investment efficiency shock that evolves according to

$$\ln \mu_t = \rho_\mu \ln \mu_{t-1} + \sigma_\mu \varepsilon_{\mu,t}, \ -1 < \rho_\mu < 1, \ \varepsilon_{\mu,t} \sim \mathbb{N}(0,1).$$

The representative household's optimality conditions imply

$$w_t = \chi n_t^\eta c_t,$$
$$1/\mu_t = E_t\left[x_{t+1}\left(r_{t+1}^k + (1-\delta)/\mu_{t+1}\right)\right],$$
$$1 = E_t[x_{t+1} r_t/\pi_{t+1}],$$

where $x_{t+1} \equiv \beta c_t/c_{t+1}$ is the pricing kernel between periods t and $t+1$.

Firms The production sector consists of a continuum of monopolistically competitive intermediate goods firms and a final goods firm. Intermediate firm $i \in [0,1]$ produces a differentiated good $y_t(i) = a_t k_{t-1}(i)^\alpha n_t(i)^{1-\alpha}$, where $k_{t-1}(i)$ and $n(i)$ are the capital and labor inputs. Following the literature, TFP (a_t) has a transitory component (s_t) and a permanent component (z_t) given by

$$\ln a_t = \ln s_t + \ln z_t,$$
$$\ln z_t = \ln g_t + \ln z_{t-1},$$
$$\ln s_t = \rho_s \ln s_{t-1} + \sigma_s \varepsilon_{s,t},\ -1 < \rho_s < 1,\ \varepsilon_s \sim \mathbb{N}(0,1),$$
$$\ln g_t = (1-\rho_g)\ln \bar{g} + \rho_g \ln g_{t-1} + \sigma_g \varepsilon_{g,t-1},\ -1 < \rho_g < 1,\ \varepsilon_g \sim \mathbb{N}(0,1),$$

where $\varepsilon_{g,t-1}$ is lagged so that the shock occurs one period before it affects TFP. Households anticipate the effects of this shock when forming expectations, consistent with the interpretation of a news shock. It can be shown that allowing news shocks to contemporaneously affect TFP would not materially change any of our results.

Each intermediate firm chooses its inputs to minimize costs, $w_t n_t(i) + r_t^k k_{t-1}(i)$, subject to the production function. After aggregating across intermediate firms, the optimality conditions imply

$$r_t^k = \alpha mc_t a_t k_{t-1}^{\alpha-1} n_t^{1-\alpha},$$
$$w_t = (1-\alpha) mc_t a_t k_{t-1}^{\alpha} n_t^{-\alpha},$$

where mc_t is the real marginal cost of producing an additional unit of output.

The final-goods firm purchases $y_t(i)$ units from each intermediate-goods firm to produce the final good, $y_t \equiv [\int_0^1 y_t(i)^{(\epsilon_p-1)/\epsilon_p} di]^{\epsilon_p/(\epsilon_p-1)}$, where $\epsilon_p > 1$ measures the elasticity of substitution between intermediate goods. It then maximizes dividends to determine the demand function for good i, $y_t(i) = (p_t(i)/p_t)^{-\epsilon_p} y_t$, where $p_t = [\int_0^1 p_t(i)^{1-\epsilon_p} di]^{1/(1-\epsilon_p)}$ is the aggregate price level.

Following Calvo (1983), a fraction, θ_p, of intermediate firms cannot choose their price in a given period. Those firms index their price to steady-state inflation, so $p_t(i) = \bar{\pi} p_{t-1}(i)$. A firm that can set its price at t chooses p_t^* to maximize $E_t \sum_{k=t}^{\infty} \theta_p^{k-t} x_{t,k} d_k^*$, where $x_{t,t} \equiv 1$, $x_{t,k} \equiv \prod_{j=t+1}^{k>t} x_j$, and $d_k^* = [(\bar{\pi}^{k-t} p_t^*/p_k)^{1-\epsilon_p} - mc_k(\bar{\pi}^{k-t} p_t^*/p_k)^{-\epsilon_p}] y_k$. Letting $p_{f,t} \equiv p_t^*/p_t$, optimality implies

$$p_{f,t} = \tfrac{\epsilon_p}{\epsilon_p - 1}(f_{1,t}/f_{2,t}),$$

$$f_{1,t} = mc_t y_t + \theta_p E_t[x_{t+1}(\pi_{t+1}/\bar{\pi})_p^{\epsilon} f_{1,t+1}],$$

$$f_{2,t} = y_t + \theta_p E_t[x_{t+1}(\pi_{t+1}/\bar{\pi})^{\epsilon_p - 1} f_{2,t+1}].$$

The aggregate price level, price dispersion ($\Delta_t^p \equiv \int_0^1 (p_t(i)/p_t)^{-\epsilon_p} di$), and the aggregate production function are given by

$$1 = (1-\theta_p) p_{f,t}^{1-\epsilon_p} + \theta_p(\pi_t/\bar{\pi})^{\epsilon_p - 1},$$

$$\Delta_t^p = (1-\theta_p) p_{f,t}^{-\epsilon_p} + \theta_p(\pi_t/\bar{\pi})_p^{\epsilon} \Delta_{t-1}^p,$$

$$\Delta_t^p y_t = a_t k_{t-1}^{\alpha} n_t^{1-\alpha}.$$

Equilibrium The central bank sets the nominal interest rate according to a Taylor rule given by

$$r_t = \bar{r}(\pi_t/\bar{\pi})^{\phi_\pi},$$

where ϕ_π controls the response to deviations of inflation from its steady-state level.

The aggregate resource constraint is given by

$$c_t + i_t = y_t.$$

Due to the permanent component of TFP, we detrend the model by dividing trended variables by $z_t^{1/(1-\alpha)}$. The detrended equilibrium system is provided in Appendix C. We solve the log-linearized model using Sims (2002) `gensys` algorithm.

Table 1: Data and model-implied moments from the baseline DSGE model

Moment	Data	Model	Moment	Data	Model
$SD(\tilde{a}_t)$	2.01	2.32	$SD(\tilde{\imath}_t)$	9.63	9.93
$SD(\Delta a_t)$	0.80	0.83	$AC(\tilde{a}_t)$	0.87	0.88
$SD(\tilde{y}_t)$	3.13	2.92	$AC(\Delta a_t)$	−0.09	0.04

Notes: A tilde denotes a detrended variable and Δ is a log change.

3.2 CALIBRATION Each period in the model is one quarter. The discount factor, $\beta = 0.995$, implies a 2% annual real interest rate. The Frisch elasticity of labor supply, $1/\eta = 0.5$, is set to the intensive margin estimate in Chetty et al. (2012). The steady-state inflation rate, $\bar{\pi} = 1.005$, is consistent with a 2% annual inflation target. The elasticity of substitution between goods, $\epsilon_p = 11$, the degree of price stickiness, $\theta_p = 0.75$, and the monetary response to inflation, $\phi_\pi = 1.5$, are set to the values used by Kurmann and Sims (2021). The capital depreciation rate, $\delta = 0.025$, matches the annual average rate on private fixed assets and durable goods over 1960 to 2019. The average growth rate of TFP, $\bar{g} = 1.0026$, and the income share of capital, $\alpha = 0.3343$, are based on the latest vintage of the Fernald TFP data.

Finally, we set the parameters of the TFP and marginal efficiency of investment (MEI) processes to match six moments in the data: the standard deviation and autocorrelation of TFP growth ($SD(\Delta a_t)$, $AC(\Delta a_t)$), the standard deviation and autocorrelation of detrended TFP ($SD(\tilde{a}_t)$, $AC(\tilde{a}_t)$), and the standard deviations of detrended output and investment ($SD(\tilde{y}_t)$, $SD(\tilde{\imath}_t)$).[7] This yields $\rho_g = 0.6$, $\rho_s = 0.8$, $\rho_\mu = 0.9$, $\sigma_g = 0.003$, $\sigma_s = 0.007$, and $\rho_\mu = 0.007$. Table 1 shows that these parameters imply a good model fit, suggesting that this model is a useful laboratory for evaluating the TFP max share identification strategy.

3.3 SIMULATION EVIDENCE Since there are three structural shocks in the DSGE model, we fit a three-dimensional structural VAR model. We work with a VAR model with intercept for $\mathbf{y}_t = (a_t, y_t, i_t)'$, given that investment has a strong connection with the MEI shock. All variables enter in logs. We first generate $1{,}000$ realizations of log-level data of length T for TFP, output, and

[7] We use the Hamilton (2018) filter with 4 lags and a delay of 8 quarters to detrend the data. Hodrick (2020) shows that this method is more accurate than a Hodrick and Prescott (1997) filter when log series are difference stationary.

Figure 1: KS estimator of responses based on the baseline DSGE model

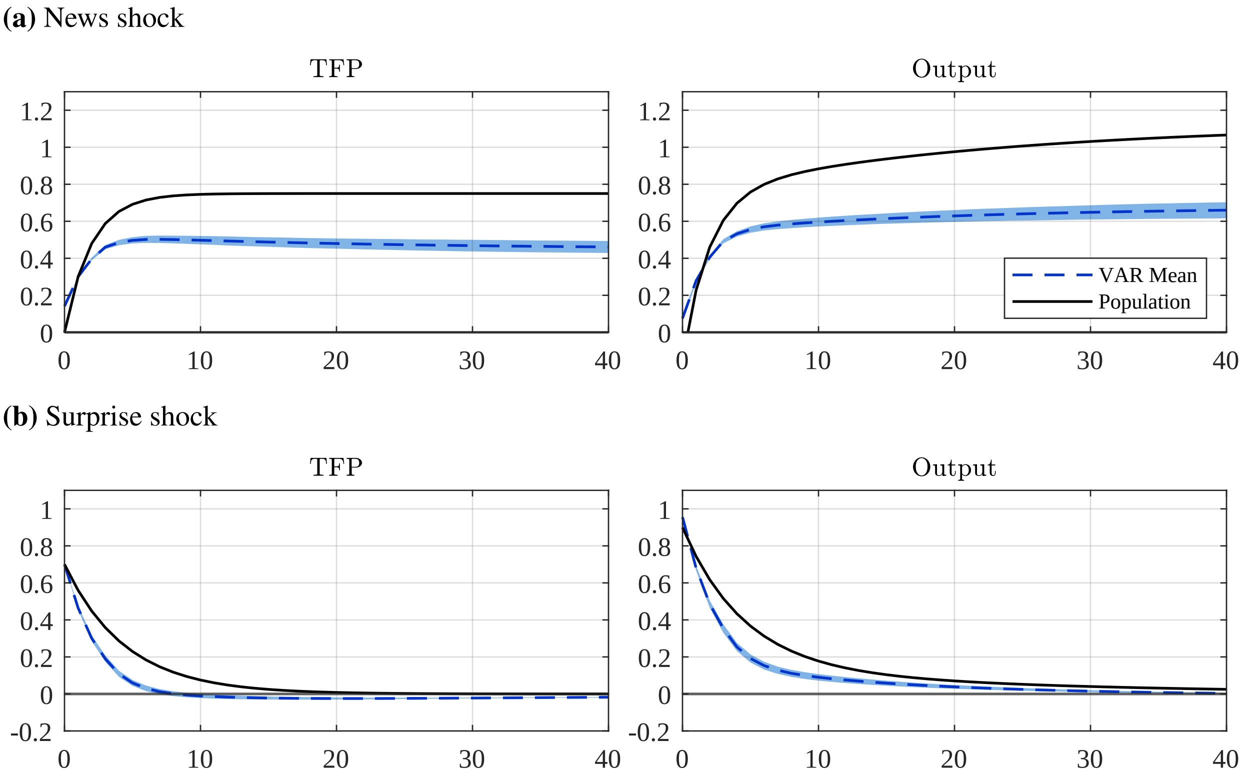

Notes: VAR(4) model with $T = 10{,}000$ and $\mathbf{y}_t = (a_t, y_t, i_t)'$. The responses are scaled so the estimated response of TFP matches the population value when the shock first takes effect.

investment by simulating the DSGE model, fit the VAR model on each of these data realizations, and construct the impulse responses. We then report the expected value of these responses, the underlying population response, and 68% quantiles of the distribution of the impulse response estimates. The distance between the expected value and the population value measures the bias of the estimator. The 68% quantiles provide a measure of the variability of the estimates.[8]

In the interest of space, Figure 1 focuses on the results for the KS estimator. We set $T = 10{,}000$ to approximate the large-sample properties of the estimator. The VAR lag order is set to 4. Our conclusions are robust to longer lags. The top row shows that the responses of TFP and output to a news shock are strongly biased downwards relative to the population responses. The responses to

[8]It can be shown that the sufficient condition for invertibility derived in Fernández-Villaverde et al. (2007) is met in both the baseline model and in the larger-scale DSGE model introduced in Section 4. This implies that both DSGE models have a VAR(∞) representation.

Table 2: RMSE over 40 quarters based on the baseline DSGE model

	TFP Response		Output Response		
Estimator	News Shock	Surprise Shock	News Shock	Surprise Shock	Total
KS Max Share	10.2	2.3	13.1	2.4	28.0
BS Max Share	2.2	6.8	3.1	9.1	21.3
NAMS Max Share	9.9	1.9	12.8	2.0	26.7

Notes: VAR(4) model with $T = 10{,}000$ and $\mathbf{y}_t = (a_t, y_t, i_t)$.

the surprise shock shown in the second row are also biased downwards.

Table 2 reports the RMSE of the impulse responses for the alternative versions of the TFP max share estimator. The first four columns show the sum of the RMSEs over horizons 0 to 40 for output and TFP. The last column shows the sum of these entries across the four response functions. The RMSEs of the responses to a news shock are nearly identical for the KS and NAMS max share estimators, suggesting that there is little to choose between them. The BS estimator is more accurate. This result is expected given that there is no measurement error in TFP and the exclusion restriction in the VAR model aligns with the timing assumption of the news shock in the DSGE model. However, the improved accuracy of the BS estimator of the responses to the news shock comes at the expense of a higher RMSE for the responses to the surprise shock. In short, this evidence calls into question the ability of all three TFP max share estimators to recover the population responses, even asymptotically.[9]

There is a widespread belief the TFP max share estimator works well as long as news shocks account for the bulk of the variation in TFP at long horizons. Our results show that this condition in not sufficient. Table 3 provides the forecast error variance decomposition of TFP in the DSGE model. The large bias in Figure 1 arises even though surprise shocks only account for 3% of the variation in TFP growth at $h = 80$. This demonstrates that the accuracy of the TFP max share estimator hinges on the surprise shock playing a small role even at horizons much shorter than $h = 80$. Intuitively, when surprise shocks are nontrivial in population, the TFP max share

[9] It can be shown that imposing additional theoretically motivated sign and magnitude restrictions, as discussed in Francis and Kindberg-Hanlon (2022), does not help address these identification problems with the TFP max share estimator.

Table 3: Forecast error variance decompositions for TFP based on the baseline DSGE model

	Horizon				
Shock	4	8	20	40	80
News	37.0	66.4	87.3	93.8	96.9
Surprise	63.0	33.6	12.7	6.2	3.1
MEI	0.0	0.0	0.0	0.0	0.0

Notes: MEI denotes marginal efficiency of investment.

estimator tends to confound news and surprise shocks, as illustrated in Figure 1. Focusing on models without TFP measurement error makes this point especially clear.[10]

4 THE ROLE OF TFP MEASUREMENT ERROR

Our analysis so far has focused on DSGE models without unobserved changes in factor utilization. A key insight in Kurmann and Sims (2021) is that, in practice, one needs to be concerned about measurement error driving a wedge between measured and true TFP. In this section, we consider an environment where unobserved factor utilization causes TFP measurement error and discuss to what extent this changes our findings.

One key difference is that in this case the TFP innovation can no longer be written as a linear combination of news and surprise shocks, because their are three shocks driving the TFP data. As a result, identifying surprise TFP shocks is not possible without further identifying restrictions. However, the news shock can be identified as before. Our main finding in this section is that the presence of measurement error does not overturn the result that the TFP max share estimator is unable, in general, to recover the population responses to news shocks.

To illustrate this point, we evaluate the TFP max share estimator based on data simulated from the medium-scale DSGE model used by Kurmann and Sims (2021), allowing for TFP measurement error. The only difference in the model structure is that we turn off the preference and monetary policy shocks in the simulations, so the results are directly comparable to the baseline model.

[10] As discussed in Appendix E, one could alternatively estimate γ_n given an estimate of γ_s obtained by maximizing the TFP forecast error variance share at a short horizon. This alternative estimator performs very poorly, even in the absence of TFP measurement error.

Table 4: Data and model-implied moments from the larger-scale DSGE model

Moment	Data	Model	Moment	Data	Model
$SD(\tilde{a}_t)$	2.01	2.31	$SD(\tilde{\imath}_t)$	9.63	9.48
$SD(\Delta a_t)$	0.80	0.73	$AC(\tilde{a}_t)$	0.87	0.87
$SD(\tilde{y}_t)$	3.13	3.92	$AC(\Delta a_t)$	−0.09	0.01

Notes: A tilde denotes a detrended variable and Δ is a log change. In the data, a_t is Fernald utilization-adjusted TFP while in the model it is measured TFP (TFP_t^u).

We begin by briefly discussing how TFP is measured. The model allows for factor utilization, denoted by u_t, to vary over time due to changes in capital utilization and worker effort. The econometrician observes neither of these but does observe output (y_t), the capital stock (k_{t-1}), hours worked (h_t), and employment (n_t). The growth in (log) unadjusted TFP is

$$\Delta \ln \text{TFP}_t = y_t - (1-\omega_{\ell,t})\Delta \ln k_{t-1} - \omega_{\ell,t}(\Delta \ln h_t + \Delta \ln n_t),$$

where $\omega_{\ell,t}$ is the labor share. Changes in factor utilization ($\Delta \ln \hat{u}_t$) are assumed to be proportional to changes in detrended hours worked ($\Delta \ln \hat{h}_t$), so

$$\Delta \ln \hat{u}_t = \vartheta \Delta \ln \hat{h}_t,$$

where ϑ is a proportionality factor. We set $\vartheta = 3$ to match the value in Kurmann and Sims (2021). Hours worked are detrended using a biweight filter, consistent with the latest vintages of the Fernald TFP measure (see Fernald, 2015). The growth in utilization-adjusted TFP is given by

$$\Delta \ln \text{TFP}_t^u = \Delta \ln \text{TFP}_t - \Delta \ln \hat{u}_t.$$

In our simulations, we produce a series for the (log) level of utilization-adjusted TFP, $\ln \text{TFP}_t^u$, by cumulating the growth rates over time. This series represents measured TFP in the model.

Just like in the baseline model, we calibrate the shock processes. This exercise implies that $\rho_g = 0.5$, $\rho_s = 0.4$, $\rho_\mu = 0.95$, $\sigma_g = 0.0025$, $\sigma_s = 0.006$, and $\sigma_\mu = 0.004$. The other parameter values are set to those used in Kurmann and Sims (2021). Table 4 shows that the model fits the data reasonably well on along all dimensions.

Figure 2 plots the responses of output and measured TFP for $T = 10{,}000$. We focus on the KS estimator in the interest of space. There is a notable discrepancy between the response of measured TFP to a news shock and the population response of true TFP. This result is not surprising. With measurement error there is no reason to expect the VAR to recover this response, since the VAR is estimated with the mismeasured TFP variable and the population response is based on true TFP. What is more concerning is that there is strong bias in the output response, sometimes in the positive direction and sometimes in the negative direction.

Table 5 shows the RMSE for the KS estimator, alongside the other TFP max share estimators we considered in Section 3, in the presence of TFP measurement error. Once again, the performance of the KS and NAMS estimators is similar. The BS estimator is slightly more accurate than the other TFP max share estimators, but the responses to the news shock are not nearly as accurate as they are in the baseline model. This is because the presence of measurement error implies that the exclusion restriction underlying the BS estimator is no longer valid. These results show that the concerns about the TFP max share estimator are robust to the presence of TFP measurement error.[11]

Table 6a shows that in this application news shocks explain only 75% of the long-run variation in measured TFP, which is almost identical to the share Kurmann and Sims (2021) obtained when applying their estimator to actual data. Given the belief that news shocks must account for a large part of the unpredictable variation in TFP at long horizons in order for the TFP max share estimator to perform well, one might argue that the bias we documented is not surprising. While it remains true in our application that news shocks explain almost all of the long-run variability of true TFP, as shown in Table 6b, this does not necessarily mean that they are as important a determinant of measured TFP. Thus, VAR models based on measured TFP may not satisfy one of the maintained assumptions required for the TFP max share estimator to perform well, which provides another possible explanation of the bias.

[11] Our conclusion may seem at odds with simulation evidence reported in Kurmann and Sims (2021) that their estimator comes somewhat close to the population responses to a news shock in very large samples. In Appendix F, we show that this result is an artifact of their parameterization of the TFP process, which is at odds with the data. We show that under their parameterization news shocks explain nearly all variation in TFP at all horizons, which greatly enhances the accuracy of the KS estimator. When the parameters in the DSGE model are set to match the TFP moments in the data, in contrast, the impulse response estimates are strongly biased.

Figure 2: KS estimator of responses to a news shock based on the larger-scale DSGE model

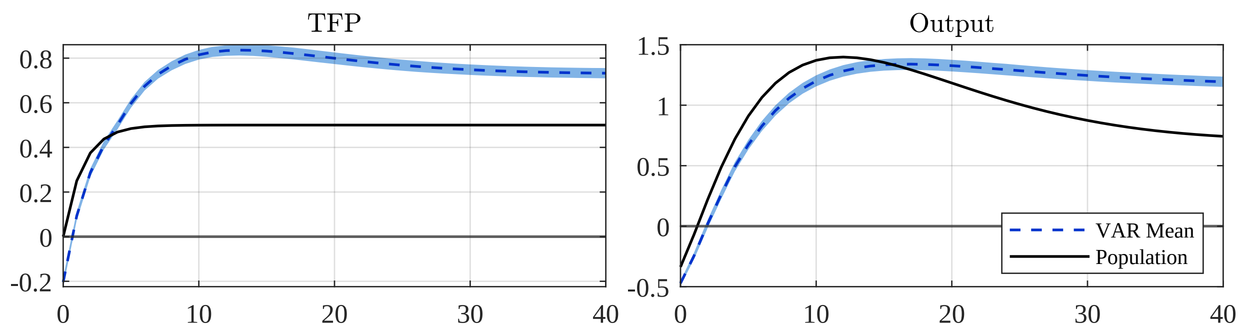

Notes: VAR(4) model with $T = 10{,}000$ and $\mathbf{y}_t = (\text{TFP}^u_t, y_t, i_t)'$.

Table 5: RMSE over 40 quarters based on the larger-scale DSGE model

Estimator	TFP Response	Output Response	Investment Response	Total
KS Max Share	10.3	10.4	19.3	40.0
BS Max Share	8.8	8.3	14.1	31.1
NAMS Max Share	10.1	10.8	21.3	42.2

Notes: VAR(4) model with $T = 10{,}000$ and $\mathbf{y}_t = (\text{TFP}^u_t, y_t, i_t)$.

Table 6: Forecast error variance decompositions for TFP in the larger-scale DSGE model

(a) Measured TFP ($\ln \text{TFP}^u_t$)

	Horizon				
Shock	4	8	20	40	80
News	33.3	46.3	47.7	62.3	75.2
Surprise	55.6	43.3	9.3	1.8	0.7
MEI	11.1	10.4	43.0	35.9	24.1

(b) True TFP ($\ln a_t$)

	Horizon				
Shock	4	8	20	40	80
News	47.9	75.7	91.0	95.6	97.8
Surprise	52.1	24.3	9.0	4.4	2.2
MEI	0.0	0.0	0.0	0.0	0.0

Notes: MEI is the marginal efficiency of investment.

5 Estimators Involving Measures of TFP News

The bias of the TFP max share estimator in large samples raises the question of whether there are alternative estimators that perform better. In this section, we consider identifying a TFP news shock by incorporating an observed measure of TFP news into the VAR model and adapting the identification strategy. Similar approaches have been employed in a number of studies. For example, Shea (1999) considers models that incorporate a measure of either government R&D spending or patent applications. Other examples include Christiansen (2008, patent applications), Alexopoulos (2011, new book titles in the fields of technology and computer science), Jinnai (2014, sector-specific productivity in the R&D sector), Baron and Schmidt (2019, counts of new information and communication technology standards), Cascaldi-Garcia and Vukotić (2022, patent applications), Miranda-Agrippino et al. (2022, patent applications), and Fieldhouse and Mertens (2023, government R&D spending). The premise of all these studies is that measures of TFP news should increase immediately as a positive news shock is realized, facilitating identification strategies based on short-run restrictions.

Despite the popularity of these identification strategies, there does not exist simulation evidence that quantifies the ability of these VAR models to recover news shocks (or for that matter surprise shocks) generated by DSGE models. In this section, we examine two such identification strategies, first in our baseline model abstracting from TFP measurement error and then in the larger-scale DSGE model, which allows the simulated TFP data to be contaminated by measurement error.

5.1 Identification Strategies Based on TFP News

One strategy is to identify the news shock as the shock that maximizes the forecast error variance contribution of the news variable at short horizons. We set $H_n = 4$, but our results are robust to smaller values for H_n. Another strategy is to treat the news measure as predetermined with respect to TFP, resulting in a block recursive VAR model with the news variable ordered first and TFP second. This approach is equivalent to using TFP news as an internal instrument (see, e.g., Plagborg-Møller and Wolf, 2021). We will refer to these estimators as the "max share news" and "Cholesky news" estimators.

An alternative approach to dealing with TFP news measurement error would be to use the news variable as an external instrument in a VAR model excluding the TFP news variable (e.g., Montiel Olea et al., 2021; Stock and Watson, 2018). This proxy VAR approach has been used, for example, by Cascaldi-Garcia and Vukotić (2022) and Miranda-Agrippino et al. (2022). Like the methods discussed in this section, the use of proxy VAR models allows the user to dispense with the assumption that news shocks do not affect TFP contemporaneously. As shown in Plagborg-Møller and Wolf (2021), the advantage of the strategy of ordering the instrument first in a block recursive VAR model is that it yields valid impulse response estimates even if the shock of interest is non-invertible. In contrast, the proxy VAR approach that uses the news variable as an external instrument is invalid in that case.[12]

An obvious concern is that, in practice, the TFP news variable could be measured with error. The instrumental variable approach allows consistent estimation of the impulse responses in that case. In contrast, there are no results in the literature about the robustness of the max share news approach to this form of measurement error, but we present simulation evidence below that both estimators perform well with and without measurement error in the TFP news variable.

5.2 ACCURACY OF THE NEWS VARIABLE ESTIMATORS

We start by using the baseline DSGE model to examine the accuracy of the news-based estimators in the absence of TFP measurement error. The news variable reflects the permanent component of TFP, z_t. Since the news shock is lagged by one period in the DSGE model, the TFP news variable only responds with a delay of one period. We therefore fit a VAR(4) model with intercept to $\mathbf{y}_t = (z_{t+1}, a_t, y_t)'$ for the news-based estimators and to $\mathbf{y_t} = (a_t, y_t, i_t)'$ for the KS estimator. This timing mirrors the way observed TFP

[12] Invertibility here refers to the ability to recover the structural shock of interest as a function of only current and past VAR model variables. When agents anticipate future changes, as is the premise in models with TFP news shocks, the maintained assumption that the VAR prediction errors are linearly related to the contemporaneous structural shocks fails whenever agents have more information than is contained in the reduced-form VAR model (for further discussion see, e.g., Hansen and Sargent, 1991; Kilian and Lütkepohl, 2017; Leeper et al., 2013). While this problem may be addressed by including additional variables in the reduced-form VAR model that capture the expected path of TFP, finding such variables is nontrivial. For example, stock price indices or measures of consumer sentiment, are not likely to be a good measures of expected TFP. The Cholesky approach avoids these complications.

Table 7: RMSE over 40 quarters based on the baseline DSGE model

	TFP Response		Output Response		
Estimator	News Shock	Surprise Shock	News Shock	Surprise Shock	Total
KS Max Share	10.2	2.3	13.1	2.4	28.0
Max Share News	1.4	0.8	1.9	1.0	5.0
Cholesky News	1.3	0.9	1.8	1.2	5.1
Alt KS Max Share	2.8	1.2	3.8	1.4	9.3

Notes: VAR(4) with $T = 10{,}000$, where $\mathbf{y}_t = (a_t, y_t, i_t)$ for the KS max share estimator and $\mathbf{y}_t = (z_{t+1}, a_t, y_t)$ for the max share news and Cholesky estimators. The Alt KS max share estimator is the KS estimator applied to $\mathbf{y}_t = (z_{t+1}, a_t, y_t)$.

news has been used in applied work.[13] The choice of these variables is dictated by our interest in constructing the responses of TFP and output.[14] All variables enter the VAR in logs and are directly observable in the DSGE model. Throughout this section, we compare the news-based estimators to the KS estimator. The corresponding results for the other TFP max share estimators have already been reported in Tables 2 and 5.[15]

Table 7 compares the RMSE of the impulse responses to news and surprise shocks across the estimators. There is strong evidence in favor of the news-based estimators. Both the max share news and Cholesky news estimators reduce the RMSE by 82% relative to the KS estimator. These improvements in accuracy are mainly due to RMSE reductions for the responses to news shocks.

To illustrate the improvement in accuracy, Figure 3 plots the responses of TFP and output to news and surprise shocks obtained using the max share news estimator. Both shocks appear properly identified by the max share approach with little bias in the mean estimates and small variability. Qualitatively similar results hold for the Cholesky news approach, as shown in Appendix G. These results suggest that identification strategies based on TFP news variables perform much

[13]It might seem more appropriate to compare the max share news and Cholesky news estimator with a two-variable VAR that includes only TFP (a_t) and output (y_t) but this would be inappropriate because the data-generating process has three unique shocks. Therefore, as before, we consider a three-variable VAR model that includes investment for the KS max share estimator, since it has a strong connection with the MEI shock.

[14]When $\mathbf{y}_t = (z_{t+1}, a_t, y_t)'$, the Q matrix is written as in (1), except that the zero restriction is in the second row of the matrix on $\gamma_{\ell,2}$. More generally, adding z_{t+1} as an additional variable in a VAR where a_t is ordered as the jth variable requires adding a column and row to Q and placing the zero restrictions in the jth row.

[15]It may seem that one could have replaced z_{t+1} in the VAR model with a_{t+1}, but it can be shown that the latter specification does not improve on the accuracy of the KS estimator.

Figure 3: Max share news estimator of responses based on the baseline DSGE model

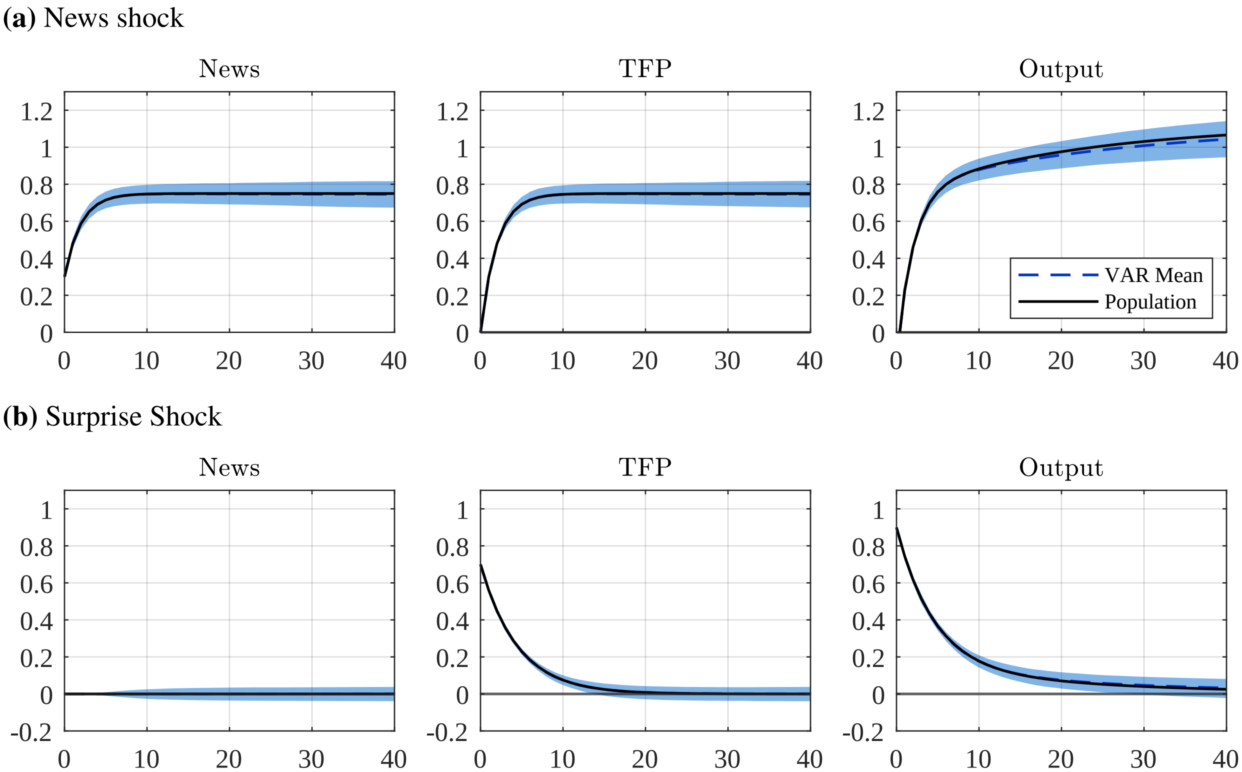

Notes: VAR(4) model with $T = 10{,}000$ and $\mathbf{y}_t = (z_{t+1}, a_t, y_t)'$. The responses are scaled so the estimated response of TFP matches the population value when the shock first takes effect.

better than the TFP max share estimator when both TFP and TFP news are accurately measured.

5.3 WHY DOES THE TFP MAX SHARE ESTIMATOR UNDER-PERFORM? It may seem puzzling that the TFP max share estimator performs so much worse than the news-based estimators. One potential reason is that the VAR models that these methods are applied to typically do not include a direct measure of TFP news. As noted by Hansen and Sargent (1991) and Leeper et al. (2013), when economic agents form expectations based on information not contained in the information set of the VAR model, standard approaches to identifying structural shocks tend to fail.

These insights suggest that the comparatively high RMSE of the TFP max share estimator may be explained by the absence of a forward-looking variable that captures TFP news in the VAR model. The last row in Table 7 explores this conjecture by applying the KS estimator to the VAR model that includes z_{t+1}. Our simulations show that indeed the accuracy of the KS estimator substantially improves when incorporating TFP news into the VAR model, but even in that case the KS estimator is less accurate than the news-based estimators.[16]

5.4 Impact of Measurement Error in TFP News In our analysis so far, we assumed that the econometrician perfectly observes the permanent component of TFP. However, the external measures of news used in empirical research are not perfectly correlated with the permanent component of TFP. To address this concern, we now allow the TFP news variable in the VAR model to be an imperfect measure of the permanent component of TFP news by introducing additive Gaussian measurement error, which is a standard approach in the econometrics literature (Plagborg-Møller and Wolf, 2022; Stock and Watson, 2018). Specifically, we replace z_{t+1} in the VAR model with $z^n_{t+1} = z_{t+1} + \sigma_n \epsilon^n_t$, where $\epsilon^n_t \sim \mathbb{N}(0, 1)$.

While there is no way of knowing the extent of measurement error in TFP news, Table 8 shows that our news-based estimators improve accuracy even if the news variable is measured with substantial error. For example, with 50% measurement error, expressed as a percentage of the standard deviation of the true news shock ($\sigma_n = 0.5\sigma_g$), both the max share news and Cholesky news estimators are still far more accurate than the KS estimator or related TFP max share estimators. Similarly accurate responses to news shocks are obtained even when allowing the measurement error to be serially correlated.

5.5 News-Based Estimators in the Presence of TFP Measurement Error An important question is whether the max share news and Cholesky news estimators can also reduce impulse response bias in the presence of TFP measurement error. We explore this question by

[16] In related work, Forni et al. (2019) recommend establishing the informational sufficiency of the reduced-form model for impulse response analysis by showing that the R^2 in a regression of the structural news shock in the DSGE model on the reduced-form VAR errors implied by the detrended DSGE model data is close to one. When applying this procedure, we find that the R^2 for the VAR(4) model augmented with the news variable is indeed essentially one, while that of the original VAR(4) model is far from one. Unlike in Forni et al. (2019), however, the VAR models of interest in our paper are specified in log levels rather than on detrended data. It is unclear whether this procedure remains valid in this case. Applying the procedure to the residuals from a log level VAR model suggests that even the R^2 for the original model is close to one. Thus, it is unclear what to make of these diagnostics. Our results in Table 7 (and similarly in Table 9 below for the model with TFP measurement error) show clear evidence of an informational deficiency in the original model, but they also show that the KS estimator cannot recover the true news shocks even when using the VAR model augmented by the news variable. This evidence indicates that the problems of the TFP max share estimator run deeper than merely an informational deficiency of the reduced-form VAR model.

Table 8: RMSE over 40 quarters based on the baseline DSGE model

	TFP Response		Output Response		
Estimator	News Shock	Surprise Shock	News Shock	Surprise Shock	Total
KS Max Share	10.2	2.3	13.1	2.4	28.0
Max Share News ($\sigma_n = 0$)	1.4	0.8	1.9	1.0	5.0
Max Share News ($\sigma_n = 0.2\sigma_g$)	1.6	0.8	2.4	1.0	5.8
Max Share News ($\sigma_n = 0.5\sigma_g$)	3.6	0.8	4.9	1.0	10.2
Cholesky News ($\sigma_n = 0$)	1.3	0.9	1.8	1.2	5.1
Cholesky News ($\sigma_n = 0.2\sigma_g$)	1.7	0.9	2.4	1.2	6.3
Cholesky News ($\sigma_n = 0.5\sigma_g$)	3.8	1.0	4.7	1.2	10.6

Notes: VAR(4) model with $T = 10{,}000$, where $\mathbf{y}_t = (a_t, y_t, i_t)'$ for the KS max share estimator and $\mathbf{y}_t = (z^n_{t+1}, a_t, y_t)'$ for the max share news and Cholesky news estimators.

fitting VAR(4) models to data for $\mathbf{y}_t = (z_{t+1}, \mathrm{TFP}^u_t, y_t, i_t)$ simulated from the larger-scale DSGE model for $T = 10{,}000$, where TFP^u_t denotes measured TFP. As Figure 4 shows, even in this case there is only modest bias in the responses of output and investment. The Cholesky news estimator produces similar results, as shown in Appendix G.

Table 9a quantifies the improvement in accuracy. It reports the RMSE of the impulse responses associated with each estimator. Adding the TFP news variable and adapting the identification procedure reduces the RMSE by 58% relative to the KS estimator. A similar reduction in accuracy also occurs when compared against the Alt KS max share estimator that includes the TFP news variable in the VAR model. Thus, the max share news and Cholesky news estimators improve accuracy, both in the presence of TFP measurement error and in its absence.

5.6 ACCURACY IN SMALL SAMPLES While our results for $T = 10{,}000$ indicate that the news-based estimators are much more accurate than the TFP max share estimator in long samples, they do not speak to the properties of the news-based estimators in sample sizes encountered in applied work. Therefore, we also examine the performance of the news-based estimators with $T = 320$ (80 years of quarterly data), which reflects the data available in the post-World War II period.

Table 9b shows the RMSE for the various estimators. Both the max share news and Cholesky news estimators lead to a 30% improvement in accuracy over the KS estimator and a 25% im-

Figure 4: Max share news estimator of responses based on the larger-scale DSGE model

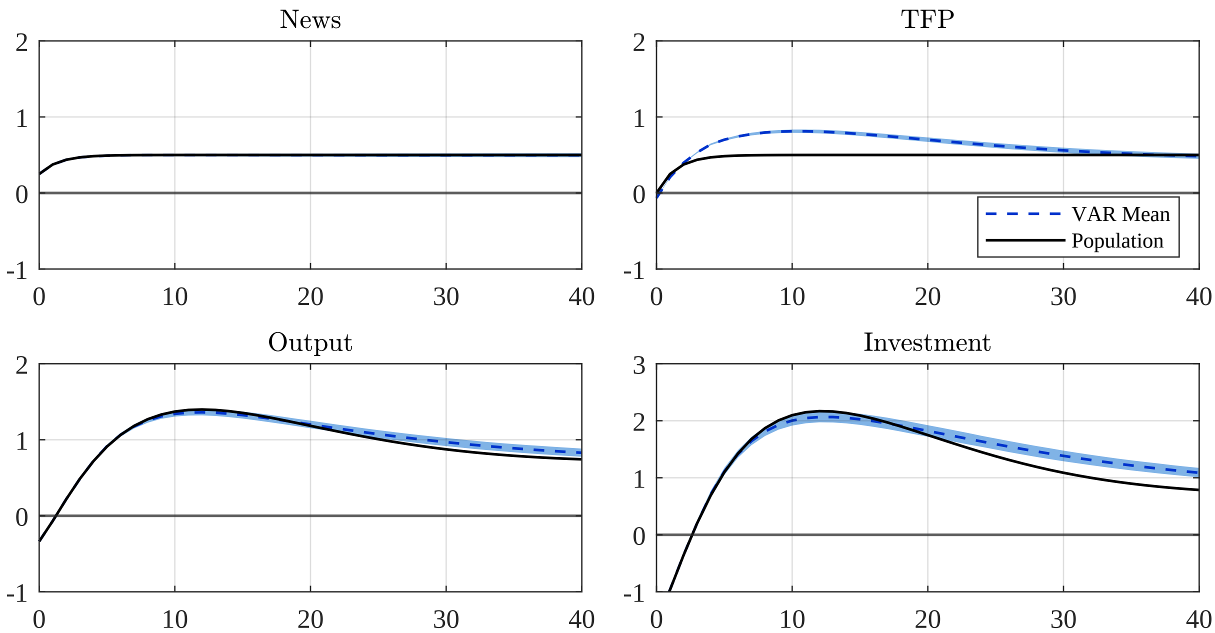

Notes: VAR(4) model with $T = 10{,}000$ and $\mathbf{y}_t = (z_{t+1}, \text{TFP}_t^u, y_t, i_t)'$.

Table 9: RMSE over 40 quarters based on the larger-scale DSGE model

(a) $T = 10{,}000$

Estimator	TFP Response	Output Response	Investment Response	Total
KS Max Share	10.3	10.4	19.3	40.0
Max Share News	6.4	2.8	7.6	16.8
Cholesky News	6.4	2.7	7.5	16.6
Alt KS Max Share	7.0	4.9	13.9	25.7

(b) $T = 320$

Estimator	TFP Response	Output Response	Investment Response	Total
KS Max Share	10.3	14.3	29.7	54.3
Max Share News	8.1	10.8	19.0	37.9
Cholesky News	8.0	10.8	18.6	37.4
Alt KS Max Share	8.1	13.2	28.2	49.4

Notes: VAR(4) model, where $\mathbf{y}_t = (\text{TFP}_t^u, y_t, i_t)$ for the KS max share estimator and $\mathbf{y}_t = (z_{t+1}, \text{TFP}_t^u, y_t, i_t)$ for the max share news and Cholesky news estimators. The Alt KS max share estimator uses the KS identification strategy and the max share news model variables.

provement over the Alt KS estimator. These results suggest that the benefits of the news-based estimators extend to realistic sample sizes and go beyond just aligning information sets.

6 EMPIRICAL FINDINGS

Our simulation evidence suggests that incorporating a measure of TFP news into the VAR model and adapting the identification strategy may improve the identification of the news shock. In practice, however, this approach will only be as good as the underlying measure of TFP news. Thus, we consider a range of VAR models that include one of four news variables: (1) **R&D**: real R&D expenditures, building on related work by Shea (1999) and Christiansen (2008); (2) **ICT**: the new information and communications technologies standards index introduced in Baron and Schmidt (2019); (3) **CGV**: the patent series used in Cascaldi-Garcia and Vukotić (2022); and (4) **MAHB**: the exogenous patent-innovation series in Miranda-Agrippino et al. (2022), which is based on quarterly total patent applications from the USPTO historical patent data file in Marco et al. (2015).[17]

For each series, we estimate a 9-variable VAR(4) model that includes one of the four news variables in addition to the 8 variables from the empirical VAR model used in Kurmann and Sims (2021).[18] Specifically, the model includes a measure of TFP news, utilization-adjusted TFP, per capita output, consumption, investment, and hours worked, the inflation rate, the real S&P 500 index, and the federal funds rate. The data sources are provided in Appendix A. The sample for each VAR varies due to differences in the availability of the news variables. We identify the

[17] Baron and Schmidt (2019) treat technological standardization as a prerequisite for new technologies to be implemented and show that shocks to the ICT series cause increases in TFP, output, and investment over medium-run horizons. Cascaldi-Garcia and Vukotić (2022) use a quarterly version of the patent series introduced by Kogan et al. (2017). This series weights patents by their value, measured as the response of each company's stock price due to news about the patent grant. The USPTO series is monthly and provides a record of all patent applications filed at the U.S. Patents and Trademark Office (USPTO) since 1981. The exogenous patent series in Miranda-Agrippino et al. (2022) is the residual from regressing the quarterly growth rate of the USPTO series on lags of itself and a set of control variables that can include SPF forecasts and exogenous policy shocks. To provide the longest sample possible, we consider the regression where the control variables exclude exogenous policy shocks. Miranda-Agrippino et al. (2022) note that their identification is robust to excluding these policy shocks.

[18] Cascaldi-Garcia and Vukotić (2022) use the same variables in their VAR model, except they also include a measure of consumer sentiment. Our results are robust to including this additional variable. There are also some differences in the data sources. Most notably, they use output from the nonfarm business sector, instead of real GDP. When we use this alternative definition of output, the impulse responses to a news shock are closer to what they report.

Table 10: Empirical results from VAR models including alternative measures of TFP news

	R&D	ICT	CGV	MAHB
Cholesky-identified VAR model				
Max response at longer horizons	Y	Y	N	N
Positive long-run responses	Y	Y	N	N
Max share identified VAR model				
Max response at longer horizons	Y	Y	N	N
Positive long-run responses	Y	Y	N	N

Notes: The Cholesky model is identified with the news variable ordered first and TFP second. The max share model maximizes the variance contribution of the news variable at $H_n = 4$.

structural shocks based on the max share news and Cholesky news models introduced in Section 5.

There are two natural criteria for judging whether the news shocks have been properly identified. These criteria are suggested by the population responses in the DSGE models used in Section 3 and Section 4, and by many other business cycle models. First, while the identification does not constrain the short-run response of TFP and output to a news shock, its effect on TFP and output should peak at horizons longer than 12 quarters. This criterion allows for weakly increasing as well as hump-shaped response functions. A peak at horizons shorter than 12 quarters would clearly be incompatible with the notion that the impact of news is largest at long horizons. Second, the news shock should have positive effects in the long run on TFP and output.

It is understood that sampling error and small-sample biases may cause slight violations of these criteria. Our focus in this section is on strong violations of these criteria. Table 10 summarizes the results based on a maximum horizon of 40 quarters. The full set of impulse responses is provided in Appendix G. We find that only the R&D and ICT models satisfy the two criteria.[19] This is true regardless of whether we use the max share news or Cholesky news identification method.

Figure 5 shows that the responses of other key variables also look reasonable for the ICT model.

[19]The estimates for the MAHB specification differ somewhat from those reported in Miranda-Agrippino et al. (2022). One likely reason is that they combine their estimate of the impact of news with a longer sample for the reduced-form model than the instrument is available for. A key difference between the MAHB instrument and the other measures of TFP news is that Miranda-Agrippino et al. (2022) purge their instrument of all dynamics. We also produced results based on an instrument that does not control for lags of the patent series. This did not affect the results reported in Table 10.

Figure 5: Comparison of max share news and TFP max share identified impulse responses

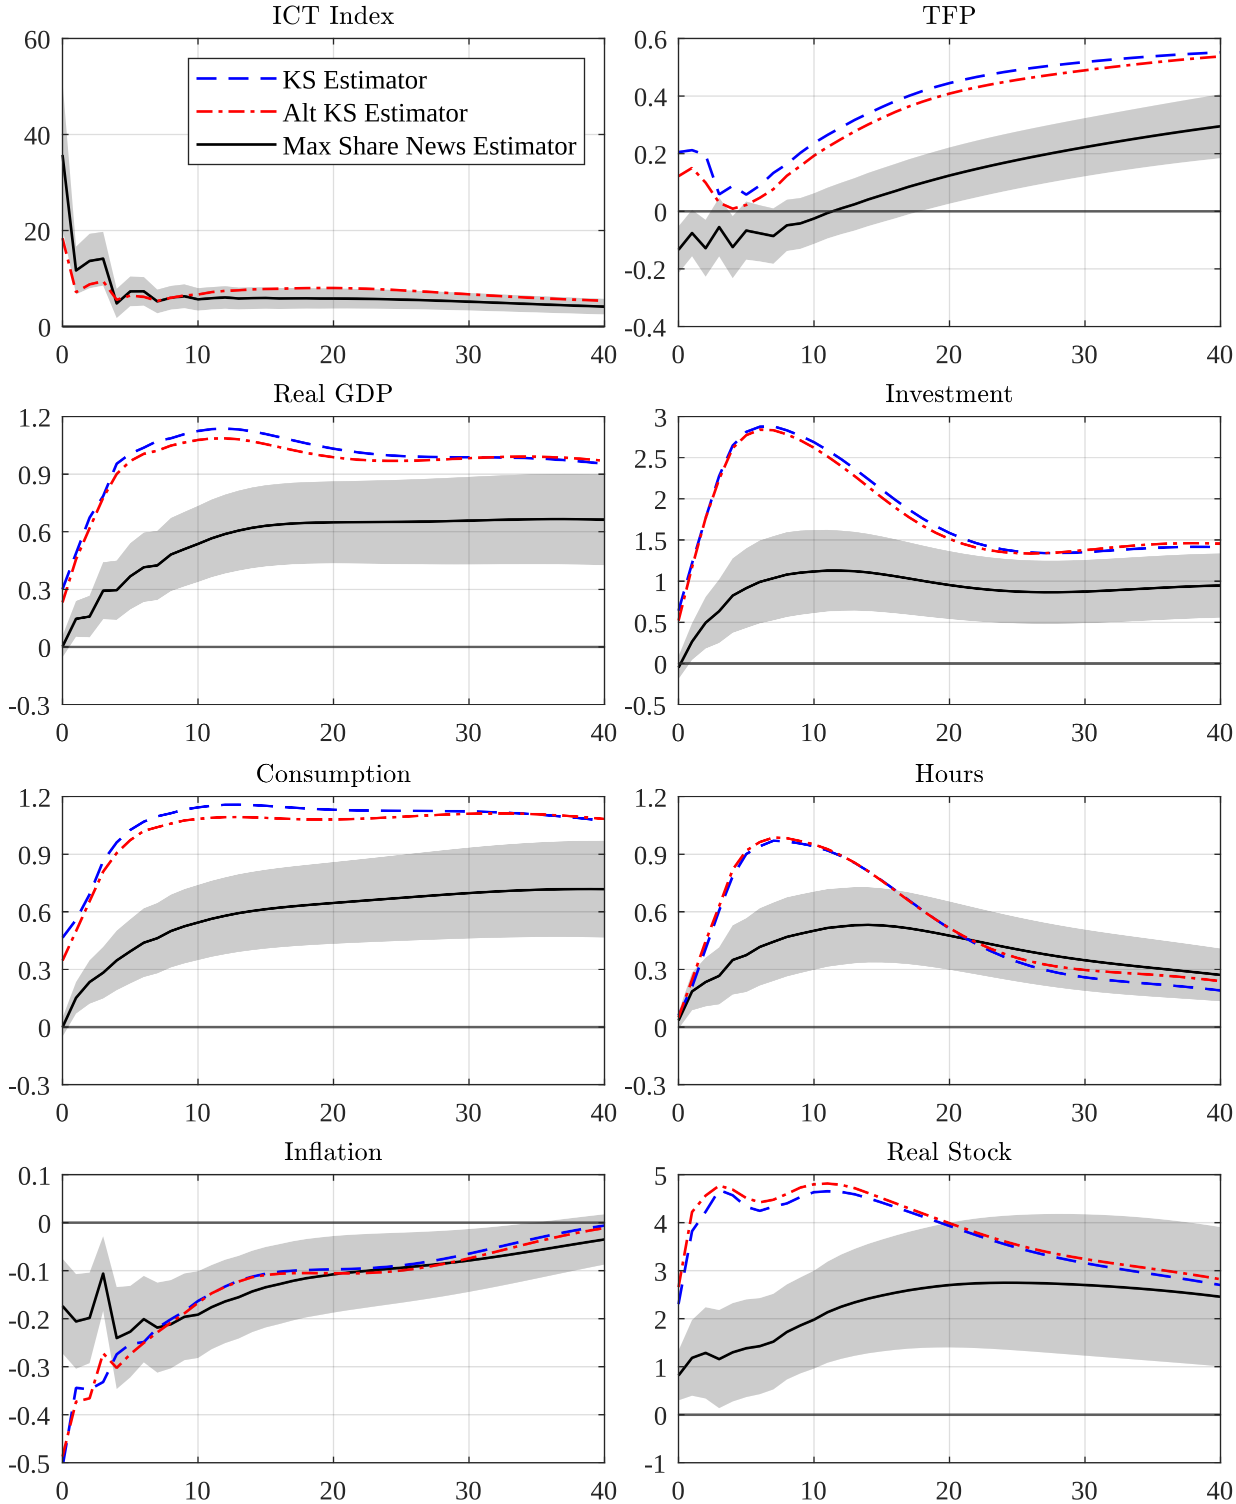

Notes: VAR(4) models estimated on identical samples from 1960-2014. Shaded regions represent 1-standard deviation error bands computed by residual-based bootstrap for the max share news estimator. Responses are in percent deviations from the baseline. The inflation response is annualized.

Similar results hold when using the R&D model.[20] The news shock increases both consumption and investment in the long run, and the peak effects occur after 12 quarters. Hours increase in the first 12 quarters, so there is positive comovement between real GDP, consumption, and investment. Inflation declines in the short-run, consistent with the interpretation of the news shock as a positive supply shock and declining real marginal costs in a New Keynesian model. Finally, the real S&P 500 index increases on impact and over the long-run, reflecting positive expectations of future economic conditions. The latter finding is consistent with the results in Beaudry and Portier (2006).

We plot these response estimates next to the estimates from the original 8-variable VAR model reported in Kurmann and Sims (2021) using the same estimation period, providing an apples-to-apples comparison between the two estimators. There are systematic and substantial differences between the two sets of response estimates, consistent with the bias documented in our simulation study. We also report the response estimates for the Alt KS max share estimator. The responses are almost identical to the 8-variable VAR model without ICT news, once again highlighting that the differences between the max share news estimator and the TFP max share estimator cannot be simply explained by an information deficiency of the original VAR.

The differences in the impulse responses translate to large differences in the forecast error variance decompositions for most variables. A forecast error variance decomposition helps assess whether news shocks are an important driver of TFP and real activity. There is no consensus on this question in the literature. Some studies find that news shocks diffuse to TFP quickly (e.g., Barsky et al., 2015; Barsky and Sims, 2011), while others find that it can take many years (e.g., Beaudry and Lucke, 2010; Cascaldi-Garcia and Vukotić, 2022; Fève and Guay, 2019; Forni et al., 2014; Levchenko and Pandalai-Nayar, 2020; Miranda-Agrippino et al., 2022). Similarly, some studies find that news shocks are the dominant driver of real activity in the medium run (e.g., Beaudry and Lucke, 2010; Fève and Guay, 2019; Forni et al., 2014), while others find that news shock play

[20] We also explored the orthogonalized nondefense R&D shock series of Fieldhouse and Mertens (2023) that was designed to mitigate the potential endogeneity of the R&D series provided by the Bureau of Economic Analysis. When we replace the federal funds and inflation rates in the VAR model with their government R&D capital and cumulated nondefense R&D appropriations series, the identified shock satisfies our criteria for a news shock and yields estimates similar to the ICT and R&D specifications.

Table 11: Forecast error variance decompositions based on actual data

	Max Share News Estimator				KS Max Share Estimator			
	4	20	40	80	4	20	40	80
TFP	2.6	2.1	9.9	24.3	6.1	25.5	55.7	71.8
Output	6.1	24.1	31.7	35.9	62.5	87.9	87.0	86.1
Consumption	9.1	24.4	31.4	35.9	81.9	94.0	93.3	90.2
Investment	4.0	12.9	18.4	24.3	48.8	71.1	74.8	76.8
Hours	6.9	21.2	25.4	24.8	29.0	59.8	52.1	49.2
Real Stock	2.9	10.5	14.3	15.4	34.7	49.2	38.3	34.0
Fed Funds	0.2	0.3	2.2	3.0	3.8	2.2	6.3	7.0
Inflation	10.4	14.8	14.5	14.0	38.8	26.5	23.3	22.0

Notes: Max share news estimates based on the ICT news variable. Qualitatively similar results hold for the R&D news variable. Columns denote the horizon of the forecast error variance decomposition. Results based on estimates from 1960-2014.

a smaller role (e.g., Cascaldi-Garcia and Vukotić, 2022; Levchenko and Pandalai-Nayar, 2020; Miranda-Agrippino et al., 2022).

Table 11 shows that under the KS identification method news shocks diffuse relatively quickly, explaining 56% of the fluctuations in TFP after just ten years. News shocks also explain the vast majority of the forecast error variance in real activity, even at relatively short horizons, leaving little room for other shocks. In contrast, the news shocks recovered by the max share news estimator are much slower to diffuse to TFP and explain a much smaller share of the fluctuations in real activity. These estimates suggest that news shocks play an important role, but one that is much smaller than suggested by the TFP max share estimator.

One potential explanation for the lower explanatory power of news shocks in the ICT model is that, in practice, any one proxy for TFP news is likely to capture only a subset of all such news. This concern, however, is alleviated by additional simulation evidence that the max share news estimator tends to be a nearly unbiased estimator of the forecast error variance decomposition, even when the observed TFP news variable fails to capture all of the variation in TFP news. This is true, for example, when measured TFP news explains only half of the variability of the latent TFP news variable. Qualitatively similar results also hold for the Cholesky news estimator.

7 CONCLUSION

The importance of understanding the economic effects of TFP news and surprise shocks is widely recognized in the literature, but the empirical evidence obtained from alternative identification strategies tends to be conflicting. A common VAR approach is to identify responses to TFP news shocks by maximizing the variance share of TFP over a long horizon. Under suitable conditions, this approach also implies an estimate of the surprise shock. We find that these TFP max share estimators tend to be strongly biased when applied to data generated from DSGE models with shock processes that match the TFP moments in the data, both in the presence of TFP measurement error and in its absence. This occurs even in settings when news shocks explain almost all of the long-run variation in TFP. We further show that this bias cannot be simply explained by an information deficiency of the VAR.

This evidence raises the question of how to proceed in applied work. We showed that including measures of TFP news in VAR models and adapting the identification strategy substantially reduces the bias and RMSE of the impulse responses, regardless of whether TFP is measured with error, and even when there is substantial measurement error in the TFP news variable. We then reported empirical estimates of the responses to news shocks for a range of TFP news measures. Two of these specifications appeared economically plausible in light of the underlying theory. Our estimates suggest that news shocks are slower to diffuse to TFP and have a smaller effect on real activity than implied by the TFP max share method.

REFERENCES

ALEXOPOULOS, M. (2011): "Read All about It!! What Happens Following a Technology Shock?" *American Economic Review*, 101, 1144–1179.

ANGELETOS, G.-M., F. COLLARD, AND H. DELLAS (2020): "Business-Cycle Anatomy," *American Economic Review*, 110, 3030–3070.

BARON, J. AND J. SCHMIDT (2019): "Technological Standardization, Endogenous Productivity and Transitory Dynamics," Manuscript, Northwestern University.

BARSKY, R. B., S. BASU, AND K. LEE (2015): "Whither News Shocks?" *NBER Macroeconomics Annual*, 29, 225–264.

BARSKY, R. B. AND E. R. SIMS (2011): "News Shocks and Business Cycles," *Journal of Monetary Economics*, 58, 273–289.

BEAUDRY, P. AND B. LUCKE (2010): "Letting Different Views about Business Cycles Compete," in *NBER Macroeconomics Annual 2009, Volume 24*, ed. by D. Acemoglu, K. Rogoff, and M. Woodford, University of Chicago Press, NBER Chapters, 413–455.

BEAUDRY, P. AND F. PORTIER (2006): "Stock Prices, News, and Economic Fluctuations," *American Economic Review*, 96, 1293–1307.

——— (2014): "News-Driven Business Cycles: Insights and Challenges," *Journal of Economic Literature*, 52, 993–1074.

BEN ZEEV, N. AND H. KHAN (2015): "Investment-Specific News Shocks and U.S. Business Cycles," *Journal of Money, Credit and Banking*, 47, 1443–1464.

BEN ZEEV, N., E. PAPPA, AND A. VICONDOA (2017): "Emerging Economies Business Cycles: The Role of Commodity Terms of Trade News," *Journal of International Economics*, 108, 368–376.

BENHIMA, K. AND R. CORDONIER (2022): "News, Sentiment and Capital Flows," *Journal of International Economics*, 137.

BOUAKEZ, H. AND L. KEMOE (2023): "News Shocks, Business Cycles, and the Disinflation Puzzle," *Journal of Money, Credit and Banking*, 55, 2115–2151.

BRETSCHER, L., A. MALKHOZOV, AND A. TAMONI (2021): "Expectations and Aggregate Risk," *Journal of Monetary Economics*, 123, 91–108.

CALVO, G. A. (1983): "Staggered Prices in a Utility-Maximizing Framework," *Journal of Monetary Economics*, 12, 383–398.

CARRIERO, A. AND A. VOLPICELLA (2024): "Max Share Identification of Multiple Shocks: An Application to Uncertainty and Financial Conditions," *Journal of Business & Economic Statistics*, forthcoming.

CASCALDI-GARCIA, D. AND A. B. GALVAO (2021): "News and Uncertainty Shocks," *Journal of Money, Credit and Banking*, 53, 779–811.

CASCALDI-GARCIA, D. AND M. VUKOTIĆ (2022): "Patent-Based News Shocks," *Review of Economics and Statistics*, 104, 51–66.

CHAHROUR, R. AND K. JURADO (2018): "News or Noise? The Missing Link," *American Economic Review*, 108, 1702–1736.

CHEN, K. AND E. WEMY (2015): "Investment-Specific Technological Changes: The Source of Long-Run TFP Fluctuations," *European Economic Review*, 80, 230–252.

CHETTY, R., A. GUREN, D. MANOLI, AND A. WEBER (2012): "Does Indivisible Labor Explain the Difference between Micro and Macro Elasticities? A Meta-Analysis of Extensive Margin Elasticities," in *NBER Macroeconomics Annual 2012, Volume 27*, ed. by D. Acemoglu, J. Parker, and M. Woodford, MIT Press, Cambridge, 1–56.

CHRISTIANO, L. J., M. EICHENBAUM, AND R. VIGFUSSON (2004): "The Response of Hours to a Technology Shock: Evidence Based on Direct Measures of Technology," *Journal of the European Economic Association*, 2, 381–395.

CHRISTIANSEN, L. E. (2008): "Do Technology Shocks Lead to Productivity Slowdowns? Evidence from Patent Data," IMF Working Paper 2008/024.

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

DIEPPE, A., N. FRANCIS, AND G. KINDBERG-HANLON (2021): "The Identification of Dominant Macroeconomic Drivers: Coping with Confounding Shocks," European Central Bank Working Paper 2534.

FACCINI, R. AND L. MELOSI (2022): "Pigouvian Cycles," *American Economic Journal: Macroeconomics*, 14, 281–318.

FERNALD, J. G. (2015): "Productivity and Potential Output before, during, and after the Great Recession," in *NBER Macroeconomics Annual 2014*, ed. by J. A. Parker and M. Woodford, University of Chicago Press, vol. 29, 1–51.

FERNÁNDEZ-VILLAVERDE, J., J. F. RUBIO-RAMÍREZ, T. J. SARGENT, AND M. W. WATSON (2007): "ABCs (and Ds) of Understanding VARs," *American Economic Review*, 97, 1021–1026.

FÈVE, P. AND A. GUAY (2019): "Sentiments in SVARs," *Economic Journal*, 129, 877–896.

FIELDHOUSE, A. J. AND K. MERTENS (2023): "The Returns to Government R&D: Evidence from U.S. Appropriations Shocks," Federal Reserve Bank of Dallas Working Paper 2305.

FORNI, M., L. GAMBETTI, AND L. SALA (2014): "No News in Business Cycles," *Economic Journal*, 124, 1168–1191.

——— (2019): "Structural VARs and Noninvertible Macroeconomic Models," *Journal of Applied Econometrics*, 34, 221–246.

FRANCIS, N. AND G. KINDBERG-HANLON (2022): "Signing Out Confounding Shocks in Variance-Maximizing Identification Methods," *AEA Papers and Proceedings*, 112, 476–80.

FRANCIS, N., M. T. OWYANG, J. E. ROUSH, AND R. DICECIO (2014): "A Flexible Finite-Horizon Alternative to Long-Run Restrictions with an Application to Technology Shocks," *Review of Economics and Statistics*, 96, 638–647.

GÖRTZ, C., C. GUNN, AND T. A. LUBIK (2022a): "Is There News in Inventories?" *Journal of Monetary Economics*, 126, 87–104.

GÖRTZ, C., J. D. TSOUKALAS, AND F. ZANETTI (2022b): "News Shocks under Financial Frictions," *American Economic Journal: Macroeconomics*, 14, 210–243.

HAMILTON, J. D. (2018): "Why You Should Never Use the Hodrick-Prescott Filter," *Review of Economics and Statistics*, 100, 831–843.

HANSEN, L. P. AND T. J. SARGENT (1991): "Two Difficulties in Interpreting Vector Autoregressions," in *Rational Expectations Econometrics*, ed. by L. P. Hansen and T. J. Sargent, Boulder, CO: Westview Press, 77–119.

HODRICK, R. J. (2020): "An Exploration of Trend-Cycle Decomposition Methodologies in Simulated Data," NBER Working Paper 26750.

HODRICK, R. J. AND E. C. PRESCOTT (1997): "Postwar U.S. Business Cycles: An Empirical Investigation," *Journal of Money, Credit and Banking*, 29, 1–16.

JINNAI, R. (2014): "R&D Shocks and News Shocks," *Journal of Money, Credit and Banking*, 46, 1457–1478.

KILIAN, L. AND H. LÜTKEPOHL (2017): *Structural Vector Autoregressive Analysis*, Cambridge University Press, New York.

KOGAN, L., D. PAPANIKOLAOU, A. SERU, AND N. STOFFMAN (2017): "Technological Innovation, Resource Allocation, and Growth," *Quarterly Journal of Economics*, 132, 665–712.

KURMANN, A. AND E. SIMS (2021): "Revisions in Utilization-Adjusted TFP and Robust Identification of News Shocks," *Review of Economics and Statistics*, 103, 216–235.

LEEPER, E. M., T. B. WALKER, AND S.-C. S. YANG (2013): "Fiscal Foresight and Information Flows," *Econometrica*, 81, 1115–1145.

LEVCHENKO, A. A. AND N. PANDALAI-NAYAR (2020): "TFP, News, and "Sentiments": the International Transmission of Business Cycles," *Journal of the European Economic Association*, 18, 302–341.

MARCO, A., M. CARLEY, S. JACKSON, AND A. F. MYERS (2015): "The USPTO Historical Patent Data Files Two Centuries of Innovation," USPTO Economic Working Paper No. 2015-1.

MIRANDA-AGRIPPINO, S., S. H. HOKE, AND K. BLUWSTEIN (2022): "Patents, News, and Business Cycles," Manuscript, Bank of England.

MIYAMOTO, W., T. L. NGUYEN, AND H. OH (2023): "In Search of Dominant Drivers of the Real Exchange Rate," *The Review of Economics and Statistics*, forthcoming.

MONTIEL OLEA, J. L., J. H. STOCK, AND M. W. WATSON (2021): "Inference in Structural Vector Autoregressions identified with an external instrument," *Journal of Econometrics*, 225, 74–87.

NAM, D. AND J. WANG (2015): "The Effects of Surprise and Anticipated Technology Changes on International Relative Prices and Trade," *Journal of International Economics*, 97, 162–177.

PIGOU, A. C. (1927): *Industrial Fluctuations*, London: Macmilan.

PLAGBORG-MØLLER, M. AND C. K. WOLF (2021): "Local Projections and VARs Estimate the Same Impulse Responses," *Econometrica*, 89, 955–980.

——— (2022): "Instrumental Variable Identification of Dynamic Variance Decompositions," *Journal of Political Economy*, 130, 2164–2202.

SHEA, J. (1999): "What Do Technology Shocks Do?" in *NBER Macroeconomics Annual 1998, volume 13*, ed. by B. S. Bernanke and J. Rotemberg, MIT Press, 275–322.

SIMS, C. A. (2002): "Solving Linear Rational Expectations Models," *Computational Economics*, 20, 1–20.

STOCK, J. H. AND M. W. WATSON (2018): "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments," *Economic Journal*, 128, 917–948.

UHLIG, H. (2003): "What moves real GNP?" Working paper, Humboldt University.

——— (2004): "Do Technology Shocks Lead to a Fall in Total Hours Worked?" *Journal of the European Economic Association*, 2, 361–371.