Estimating Macroeconomic News and Surprise Shocks*

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

February 9, 2024

ABSTRACT

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. To help reconcile these differences, this paper examines the ability of the state-of-the-art VAR approach in the literature to identify responses to TFP news shocks and possibly surprise shocks in theory and practice. We find that this estimator tends to be strongly biased when applied to data generated from DSGE models with shock processes that match key TFP moments, 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. Applying the news-based identification methods 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 state-of-the-art method.

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

*We thank Filippo Ferroni and Andrew Fieldhouse for discussing our paper at the 2023 System Econometrics Meeting and the 2024 American Economics Association Meeting. We also thank Francesco Bianchi, Danilo Cascaldi-Garcia, Oliver Coibion, Sinem Hacıoğlu Hoke, and Todd Walker for helpful comments. The views 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.

†Research Department, Federal Reserve Bank of Dallas, 2200 N Pearl Street, Dallas, TX 75201, and CEPR (lkilian2019@gmail.com).

‡Research Department, Federal Reserve Bank of Dallas, 2200 N Pearl Street, Dallas, TX 75201 (michael.plante@dal.frb.org).

§Corresponding author: Research Department, 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 of the effects of shocks to total factor productivity (TFP) on macroeconomic aggregates, starting with Beaudry and Portier (2006). However, empirical evidence on the impact of TFP news shocks obtained from alternative identification strategies tends to be conflicting. To help reconcile these differences, we first use data simulated from dynamic stochastic general equilibrium (DSGE) models to compare the ability of the state-of-the-art approach in the literature to identify news shocks proposed by Kurmann and Sims (2021, henceforth, KS) to alternative approaches. We then employ the preferred estimator to reexamine the empirical importance of TFP news shocks for the U.S. economy. Our analysis shows that news shocks are slower to diffuse to TFP and have a smaller effect on real activity than implied by the state-of-the-art method.

The most common approach to identifying anticipated shocks to TFP ("news shocks") and unanticipated shocks to TFP ("surprise shocks") is to use the max share estimator popularized by Barsky and Sims (2011, henceforth, BS), who in turn built on Uhlig (2004) and Francis et al. (2014). This estimator selects parameters for the structural impact multiplier matrix of a vector autoregressive (VAR) model to maximize the sum of the forecast error variance shares of TFP over a ten-year horizon subject to the restriction that the news shock is orthogonal to current TFP. The latter assumption can be traced to Cochrane (1994) and Beaudry and Portier (2006) and has been central to most identification strategies seeking to recover TFP news shocks.[2]

As stressed by Barsky et al. (2015) and KS, however, the assumption that news shocks affect TFP only with a delay is hard to defend on *a priori* grounds. One reason is that new technologies may affect TFP immediately, even though their main effect on TFP takes many years due to the

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

[2] Variations of this max share approach have been widely used in applied work in a variety of economic contexts (e.g., Angeletos et al., 2020; Ben Zeev and Khan, 2015; Benhima and Cordonier, 2022; Bouakez and Kemoe, 2023; Carriero and Volpicella, 2022; Chen and Wemy, 2015; Fève and Guay, 2019; Forni et al., 2014; Francis and Kindberg-Hanlon, 2022; Görtz et al., 2022a,b; Levchenko and Pandalai-Nayar, 2020; Nam and Wang, 2015).

slow diffusion of those technologies. Another reason is that changes in measured TFP are difficult to distinguish from unobserved changes in factor utilization. As a result, TFP is likely to be mismeasured in empirical work, as documented by KS (also see Bouakez and Kemoe, 2023). These facts call into question any identification strategy involving restrictions on the short-run response of TFP.

In response to these concerns, KS proposed an alternative approach to estimating TFP news shocks that is conceptually similar to BS with two important differences. First, they allow news shocks to have contemporaneous effects on TFP. Second, they construct the shock that accounts for the maximum forecast error variance share at a given long horizon rather than maximizing the sum of the forecast error variance shares from the impact period up to that horizon. They interpret that shock as a news shock if it causes TFP to increase gradually, while causing TFP news indicators to jump on impact. KS show that their max share impulse response estimates are robust to revisions in the widely used measure of TFP developed by Fernald (2015), unlike the BS estimates.

It is widely believed that the KS estimator and, more generally, 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 the estimator. We begin by examining the accuracy of the KS estimator under ideal conditions in the absence of TFP measurement error using data generated from a conventional DSGE model. Our simulation results demonstrate that, even when virtually all variation in TFP at a 20-year horizon is explained by news shocks, the KS estimator may fail to recover the responses to news shocks, regardless of the sample size. This highlights that the accuracy of the KS estimator not only depends on the quantitative importance of news shocks at long horizons but also at shorter horizons. We then show that these results also hold when using the larger-scale DSGE model developed by KS, which allows the simulated TFP data to be contaminated by measurement error.

Our conclusion may seem at odds with simulation evidence reported by KS that their estimator comes somewhat close to the population responses to a news shock in very large samples. While we can replicate these findings under their parameterization of the TFP process, which implies that

surprise shocks are not an important determinant of TFP at any horizon, our evidence suggests that this parameterization is at odds with the data. When the parameters in our DSGE model are set to match key TFP moments in the data, impulse responses based on the KS estimator are strongly biased. This bias worsens in realistically small samples, and the estimator becomes highly variable. We trace this problem to the inability of the KS estimator to distinguish the effects of news and non-news shocks at short horizons.[3]

These results raise 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 KS estimator in the absence of TFP measurement error. We then compare the news variable VAR estimators to the KS estimator in the presence of TFP measurement error and show that these methods still substantially reduce the bias and RMSE of the responses to news shocks. The superior accuracy of these estimators is also robust to sizable measurement error in the news variable. This evidence indicates that estimators based on TFP news variables avoid the identification issues of the KS estimator.

[3]Similar results are obtained for the non-accumulated max share (NAMS) estimator recently proposed by Dieppe et al. (2021). We also show that when TFP is measured correctly, the same model may be estimated by maximizing the forecast error variance at short horizons to obtain an estimate of the surprise shock, from which the news shock may be derived. This alternative estimator is even less accurate than the original KS estimator.

[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).

Our third contribution is to illustrate the use of TFP news for identifying news shocks using a range of empirical measures of TFP news that have been used in other studies. We first show that two of these 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 KS VAR model, consistent with the findings of our simulation study.

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 output.[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. These results differ substantially from the state-of-the-art estimates implied by the KS approach. In the KS VAR model, news shocks not only quickly diffuse to TFP, but 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 KS estimator in the absence of TFP measurement error. In Section 4, we extend the analysis to the larger-scale KS 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 KS approach. Section 7 contains the concluding remarks.

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

2 IDENTIFICATION PROBLEM

This section first describes the KS method for identifying TFP news shocks. It then explains the conditions under which the KS method simulateneously 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 model 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 KURMANN-SIMS APPROACH 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 $j = 1, 2, 3$. $\gamma_{\ell,j}$ are the elements associated with an unnamed third shock. KS construct an estimate of the news shock based on

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

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 a 20-year horizon. KS stress the importance of validating the model estimate by showing that selected news indicators respond positively to the news shock in the short run.

One key difference between the BS estimator and the KS estimator is that BS restricted $\gamma_{n,1}$ and $\gamma_{\ell,1}$ to zero, which implies that $\gamma_{s,1} = 1$ given the orthogonality of Q. This follows from their assumption that TFP innovations are fully explained by the surprise shock. The KS estimator, in contrast, allows news shocks to affect TFP contemporaneously and hence leaves $\gamma_{n,1}$ unrestricted.

This leaves open the question of what to do about the additional restriction $\gamma_{\ell,1} = 0$ maintained in BS. Whether one imposes the added restriction $\gamma_{\ell,1} = 0$ or not does not affect the estimate of the news shock proposed by KS, but determines whether the surprise shock can be identified in this model. 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 $\gamma_{\ell,1} = 0$.

2.3 IDENTIFICATION CONDITIONS For concreteness, consider the same VAR model with three variables outlined above and assume TFP is affected only by news and surprise shocks ($\gamma_{\ell,1} = 0$).

The Q matrix is orthogonal if and only if $Q'Q = QQ' = I_3$. This 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

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

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

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

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

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

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

Following KS, 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 in the absence of TFP measurement error the KS 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 KS estimator. In particular, $\gamma_{s,1} = \pm\sqrt{1-\gamma_{n,1}^2}$ and $\gamma_{s,j} = -\gamma_{n,1}\gamma_{n,j}/\gamma_{s,1}$ for $j \in \{2,\ldots,K\}$.*

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 $K > 3$.

While KS clearly show that it is important to account for TFP measurement error in empirical work, this result is important for two reasons. First, it implies that the KS estimator will be able to deliver estimates of the surprise shock if the accuracy of TFP measures can be improved. Second, and more importantly, it allows us to shed light on the ability of the KS estimator to recover the population responses to news and surprise shocks under ideal conditions without TFP measurement

error. If the KS estimator does not work in this ideal setting, clearly it would not be expected to work in less than ideal settings with TFP measurement error.

3 ASSESSING THE ACCURACY OF THE KS ESTIMATOR

3.1 DATA GENERATING PROCESS

We begin our examination of the KS estimator by focusing on the ideal setting where there is no measurement error, so $\gamma_{\ell,1} = 0$. To facilitate our analysis, we generate data from a conventional New Keynesian model (henceforth, the "baseline model"), which is a simplified version of the DSGE model used by KS. The key difference is that we abstract from time-varying factor utilization, which only matters when modeling TFP measurement error. The full KS model that includes TFP measurement error will be examined in Section 4.

Households The representative household solves the Bellman equation

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

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) that evolve according to

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

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.

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%

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

Moment	Data	Baseline Model	Baseline Model with KS TFP Parameterization
$SD(\tilde{a})$	2.01	2.31	1.45
$SD(\Delta a)$	0.80	0.83	0.30
$AC(\tilde{a})$	0.87	0.88	0.88
$AC(\Delta a)$	-0.09	0.04	0.67
$SD(\tilde{y})$	3.13	2.88	1.79
$SD(\tilde{\imath})$	9.63	9.62	7.20

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

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 in KS. 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 TFP data produced by Fernald.

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 ($SD(\Delta \ln a)$, $AC(\Delta \ln a)$), the standard deviation and autocorrelation of detrended TFP ($SD(\Delta \ln \tilde{a})$, $AC(\Delta \ln \tilde{a})$), and the standard deviations of detrended output and detrended investment ($SD(\ln \tilde{y})$, $SD(\ln \tilde{\imath})$).[6] This yields $\rho_g = 0.6$, $\rho_s = 0.8$, $\rho_\mu = 0.9$, $\sigma_g = 0.003$, $\sigma_s = 0.007$, and $\rho_\mu = 0.0075$. Table 1 shows that these parameters imply a good model fit, suggesting that this model is a useful laboratory for evaluating the KS identification strategy.

3.2 Simulation Evidence on the Accuracy of the KS Estimator

Since there are three structural shocks in the DSGE model, we fit a three-dimensional structural VAR model. We work with a VAR(4) 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, and the lag order matches that used by KS. We generate 1,000 realizations of log-level data of length T for TFP, output, and invest-

[6] 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 using a Hodrick and Prescott (1997) filter when log series are difference stationary.

Figure 1: KS max share identified impulse responses based on the baseline DSGE model

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

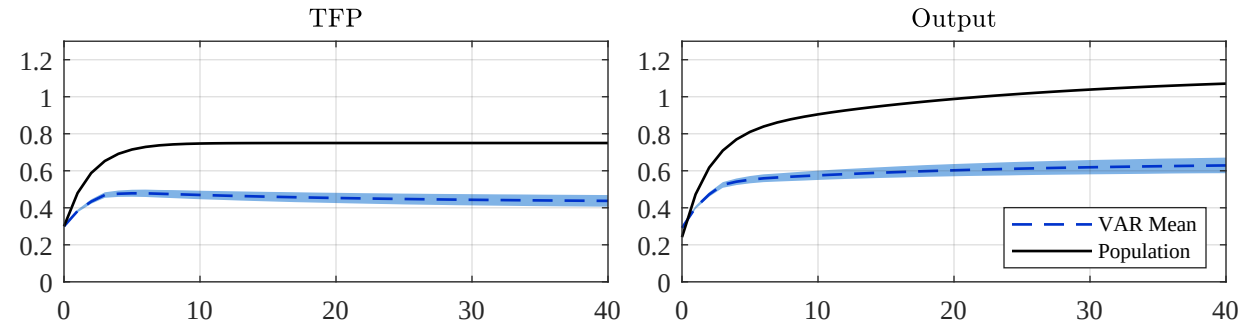

(b) Surprise shock, $T = 10{,}000$

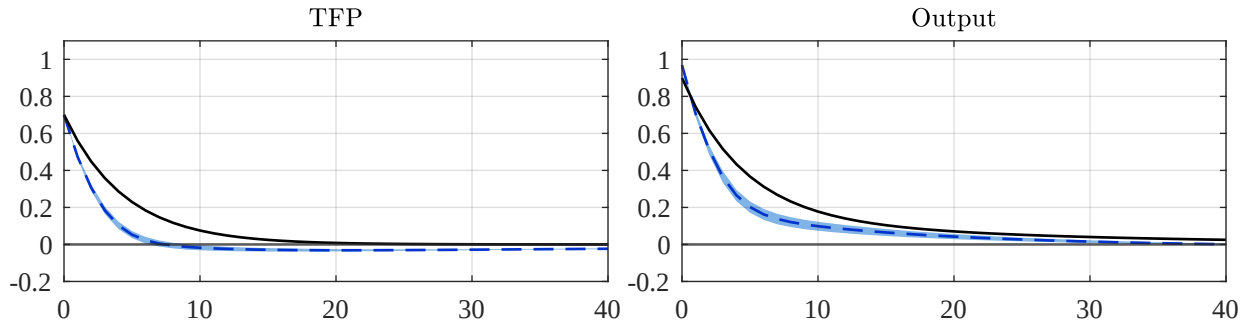

(c) News shock, $T = 240$

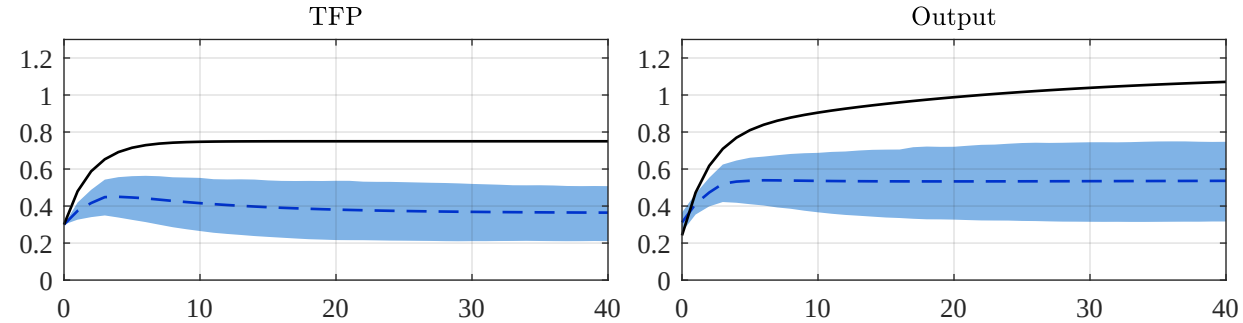

(d) Surprise shock, $T = 240$

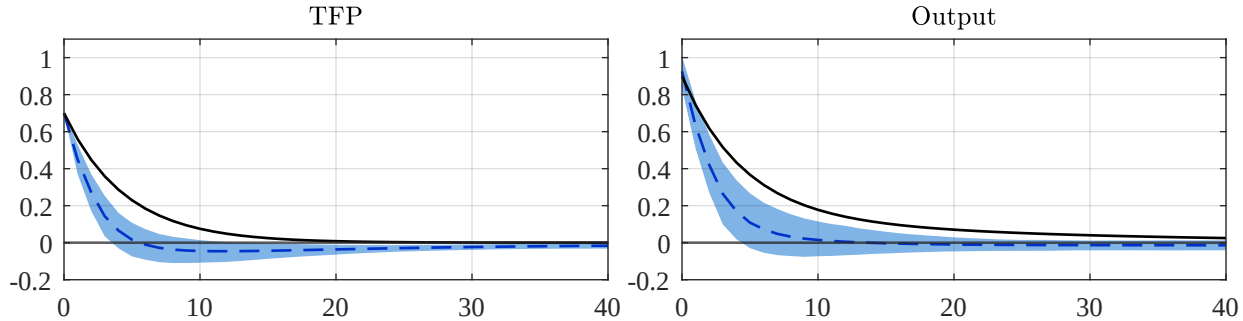

Notes: Based on a VAR(4) model for $\mathbf{y}_t = (a_t, y_t, i_t)'$. The responses have been scaled so the impact response of TFP matches the population value.

ment by simulating the DSGE model, fit the VAR model on each of these data realizations, and construct the impulse responses. Figure 1 reports the expected value of these responses, the underlying population response, and 68% quantiles of the distribution of the impulse response estimates, following KS. 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.[7]

It is useful to start with results for $T = 10{,}000$, which is the sample size KS used in their simulation study. The top row shows that in this case the responses of TFP and output to a news shock are strongly biased downwards relative to the population responses. The responses to the surprise shock shown in the second row are also biased downwards with the TFP response exhibiting the wrong sign at many horizons. This evidence calls into question the ability of the KS estimator to recover the population responses, even asymptotically. The concern here is not only that the signs of the responses and the shapes of the response functions may differ systematically from the population responses, but also the magnitude, which in turn may also affect the quantitative importance of the news shock in explaining the variability of the model data.

While our results for $T = 10{,}000$ are informative about the asymptotic properties of the estimator, they do not speak to the properties of the KS estimator for sample sizes encountered in applied work. Therefore, we also examine the performance of the KS estimator for $T = 240$ (60 years of quarterly data), which is a reasonably long estimation period in practice. The bottom two rows of Figure 1 show that the bias of the impulse response estimator is exacerbated, while the variability of the estimator increases substantially. This is particularly true for the estimator of the responses to news shocks.[8]

3.3 Comparison with KS Simulation Evidence

Contrary to our findings, KS report having some success identifying the news shock in a Monte Carlo exercise with $T = 10{,}000$ based on a larger-scale DSGE model. The key difference is not the inclusion of additional DSGE model

[7] 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 KS DSGE model introduced in Section 4. This implies that both DSGE models have a VAR(∞) representation.

[8] 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.

Table 2: Forecast error variance decompositions for TFP based on the baseline DSGE model

	Baseline Model			Baseline Model with KS TFP Parameterization		
h	News Shock	Surprise Shock	MEI Shock	News Shock	Surprise Shock	MEI Shock
1	15.5	84.5	0.0	96.2	3.8	0.0
8	70.5	29.5	0.0	99.7	0.3	0.0
20	87.9	12.1	0.0	99.9	0.1	0.0
40	93.9	6.1	0.0	99.9	0.1	0.0
80	97.0	3.0	0.0	100.0	0.0	0.0

Notes: h is the horizon of the variance decomposition and MEI denotes marginal efficiency of investment.

features or that KS allow for TFP measurement error, but that they use a different parameterization for the TFP process ($\rho_g = 0.7$, $\rho_s = 0.9$, $\sigma_g = 0.002125$ and $\sigma_s = 0.000425$). KS note that their TFP parameterization is based on standard values in the literature. However, most DSGE models feature either a stationary or a permanent TFP shock process. When a model features both processes—as is the case in their model—standard values from models with only one process can lead to TFP moments that are at odds with actual data. The most notable difference from our calibration is that the standard deviation of their surprise shock is only about 6% of our baseline value. The last column of Table 1 reports the implied model moments when using their parameterization of the TFP process in our model. This shows the KS specification is clearly at odds with the data.

The unrealistically high persistence of the KS TFP growth process (0.67 compared to about zero in the data) is important for understanding their findings because it drives the forecast error variance decomposition of TFP in the DSGE model, as shown in Table 2. Under our baseline calibration, the news shock plays an important role only at longer horizons. Under the KS parameterization, the news shock explains most of the variance at all horizons, which effectively eliminates the surprise shock and makes it easier for the KS procedure to identify the news shock. This explains the comparatively high accuracy of the KS estimator in their simulation analysis.

Our analysis suggests that the KS method will not work unless the role of the surprise shock is negligible at all horizons, rendering the identification of the news shock trivial. This conclusion differs materially from the widespread belief that the KS estimator and, more generally, max share

estimators of TFP news shocks work well as long as news shocks account for the bulk of the variation in TFP at long horizons. As Table 2 shows, even when surprise shocks only account for 3% of the variation in TFP growth at $h = 80$, the KS estimator fails to identify the response to the news shock. In other words, the success of the KS 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 estimator proposed by KS tends to confound news and surprise shocks, as illustrated in Figure 1. Focusing on models without TFP measurement error makes this point especially clear. Since one does not know how important surprise shocks are in applied work, this result cautions against the use of the KS estimator. After all, the objective of KS is to quantify the importance of news shocks, so a procedure that only works when news shocks are known *a priori* to explain nearly all of the variation in TFP at all horizons is of limited use in applied work. In the latter case, one could dispense with the max share VAR approach and estimate the news shock directly from the TFP data.

3.4 ALTERNATIVE ESTIMATORS

This section briefly considers two alternatives to the KS estimator. We examine whether either of them can help reduce the bias and RMSE of the KS estimator.

Non-Accumulated Max-Share Estimator As noted earlier, a key difference between the BS and KS estimators is that KS maximize the forecast error variance of the news shock at H_n rather than the sum of the forecast error variances over all periods up to H_n. Dieppe et al. (2021) go a step further and propose a NAMS estimator that, when adapted to our context, maximizes the square of the TFP response to the news shock at H_n. Specifically, the estimate of the news shock is based on

$$\gamma_n = \text{argmax} \; \frac{\Phi_{1,\tau} P \gamma_n \gamma_n' P' \Phi_{1,\tau}'}{\Phi_{1,\tau} \Sigma \Phi_{1,\tau}'}, \tag{3}$$

subject to the restriction that $\gamma_n' \gamma_n = 1$. The belief is that this estimator can help reduce bias of the KS estimator by further reducing the influence of the surprise shock on γ_n. Table 3, however, shows that the sum of the RMSE under the KS and NAMS max share estimators are nearly identical, suggesting that there is little to choose between these estimates. Similar results hold for the KS

Table 3: Sum of the RMSE over 40 quarters for each estimator based on the baseline DSGE model

	TFP Response		Output Response		
Estimator	News Shock	Surprise Shock	News Shock	Surprise Shock	Total
KS Max Share	11.3	2.6	14.5	2.3	30.7
NAMS Max Share	11.0	2.0	14.1	1.9	28.9

Notes: The VAR includes $\mathbf{y}_t = (a_t, y_t, i_t)$ for both max share estimators.

model, regardless of whether TFP is measured with error.

Surprise Shock Max Share Estimator A direct implication of our analysis in Section 2.2 is that we can either estimate γ_s given an estimate of γ_n obtained by maximizing the TFP forecast error variance share at a long horizon or, alternatively, we can estimate γ_n given an estimate of γ_s obtained by maximizing the TFP forecast error variance share at a short horizon. In other words, the estimator of γ_n is not unique. This raises the question of which estimator should be used when there is no TFP measurement error. A surprise shock max share estimator can be defined as

$$\gamma_s = \text{argmax } \Omega_{1,1}(H_s) = \text{argmax } \frac{\sum_{\tau=0}^{H_s} \Phi_{1,\tau} P \gamma_s \gamma_s' P' \Phi_{1,\tau}'}{\sum_{\tau=0}^{H_s} \Phi_{1,\tau} \Sigma \Phi_{1,\tau}'}, \tag{4}$$

subject to the restriction that $\gamma_s'\gamma_s = 1$ and that the responses of selected variables to the surprise shock match patterns that would be expected of a surprise shock, where H_s is set to a one-year horizon and $\gamma_s = (\gamma_{s,1}, \gamma_{s,2}, \gamma_{s,3})'$ denotes the first column in the orthogonal rotation matrix Q.

Figure 2 shows that not only is the alternative estimator much more biased than the original estimator in large samples, but it also tends to generate impulse responses that are increasing when the population response is declining and that are declining when the population response is increasing. In fact, responses to these surprise shocks look much like one would expect responses to a news shock to look like. Moreover, the responses to the news shock are of the opposite sign of the population responses. Thus, this alternative estimator should not be used in applied work.

Figure 2: Impulse responses from two max share estimators based on the baseline DSGE model

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

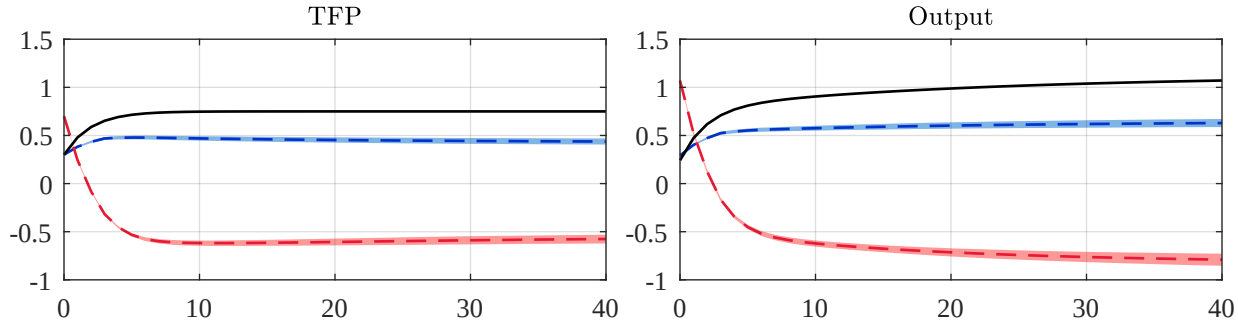

(b) Surprise shock, $T = 10{,}000$

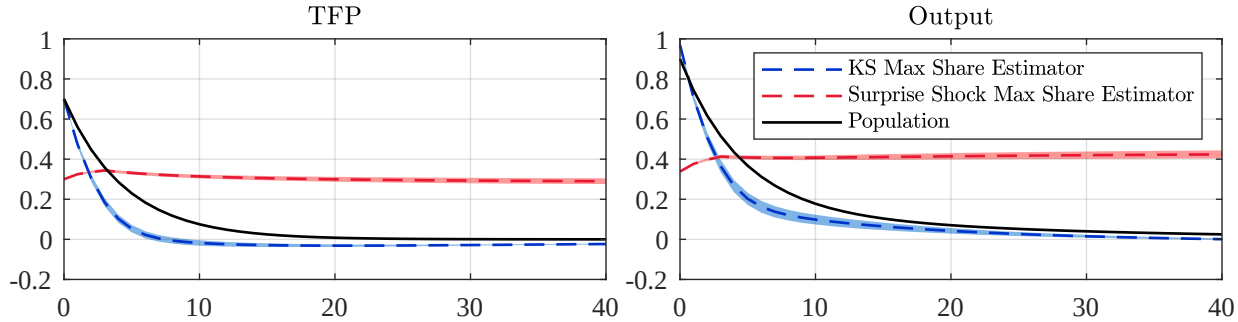

Notes: Based on a VAR(4) model for $\mathbf{y}_t = (a_t, y_t, i_t)'$.

4 THE ROLE OF TFP MEASUREMENT ERROR

Our analysis so far has focused on DSGE models without unobserved changes in factor utilization. Of course, a key point of KS is that one needs to be concerned about measurement error driving a wedge between measured and true TFP. In this section, we consider an environment with TFP measurement error and discuss to what extent this changes our findings.

One key difference is that the TFP innovation can no longer be written as a linear combination of news and surprise shocks, so 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 invalidate the result that the KS estimator is unable in general to recover the population responses to news shocks. The estimator works reasonably well under the KS TFP parameterization, but the impulse responses have large

bias when the TFP parameters are calibrated so that measured TFP matches key moments in the data.

To illustrate this point, we evaluate the KS estimator based on data generated from the larger-scale DSGE model employed by KS, allowing for TFP measurement error.[9] There are only two differences, which are made so the results are directly comparable to the baseline model. One is that we turn off the preference and monetary policy shocks in the simulations. The other is that the shock in the TFP growth process is contemporaneous, whereas in the KS model it is lagged by one period. Our specification imposes that TFP news has a contemporaneous effect on TFP, consistent with the specification of the empirical model in KS. Appendix E shows that the timing of the shock has little impact on the accuracy of the KS estimator.

We begin by briefly discussing how TFP is measured. Our approach mirrors KS. Their model allows for factor utilization, denoted by u, to vary over time due to changes in capital utilization and worker effort. The econometrician observes neither of these but does observe output (y), the capital stock (k), hours worked (h) and employment (n). The growth in (log) unadjusted TFP is

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

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

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

where $\hat{\beta}$ is a proportionality factor. Following KS, we set $\hat{\beta} = 3$. Hours worked are detrended using a biweight filter, consistent with the latest vintages of the Fernald TFP measure (see Fernald, 2015). The growth in utilitzation-adjusted TFP is given by

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

[9]The parameter choices KS used in their replication code differ slightly from those stated in the appendix of their article. We rely on the parameter values in their code. We also corrected a few errors in their derivation of the equilibrium system. The corrected equations can be found in Appendix D. These corrections do not change the main point in KS that TFP measurement error is important, but they affect some of the quantitative results.

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

Moment	Data	KS Model No Measurement Error (1)	KS Model Measurement Error (2)	KS Model with Calibrated Shock Processes Measurement Error (3)
$SD(\tilde{a})$	2.01	1.44	2.60	2.31
$SD(\Delta a)$	0.80	0.30	0.55	0.73
$AC(\tilde{a})$	0.87	0.88	0.88	0.87
$AC(\Delta a)$	-0.09	0.67	0.52	0.02
$SD(\tilde{y})$	3.13	5.22	5.22	3.99
$SD(\tilde{\imath})$	9.63	11.79	11.79	9.36

Notes: A tilde denotes a detrended variable and Δ is a log change. In the data, a is Fernald utilization-adjusted TFP. In the model, a is true TFP when there is no TFP measurement error and utilization-adjusted TFP (TFP^u) when there is measurement error.

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

In the previous section, we highlighted that the moments of the TFP process in KS are at odds with the data. This continues to be the case even when we allow for TFP measurement error and work with measured TFP rather than true TFP in the VAR model. Column (1) in Table 4 confirms that the KS model implies similar TFP moments as the baseline model in the absence of measurement error. Column (2) shows that the fit of the model does not substantially improve when incorporating measurement error. While the 0.52 autocorrelation of the growth rate of measured TFP is lower than what is reported in Column (1), it remains well above the data. Given this poor fit, we calibrate the shock processes to fit the data moments shown in Table 4. This exercise implies that $\rho_g = 0.5$, $\rho_s = 0.45$, $\rho_\mu = 0.95$, $\sigma_g = 0.0025$, $\sigma_s = 0.0065$, $\sigma_\mu = 0.004$. Column (3) shows that using the calibrated parameters improves the fit along all dimensions, particularly for the autocorrelation of TFP growth, which drops to 0.02.[10]

As discussed in Section 3, the KS estimator may be biased, even asymptotically, when the variance contribution of the surprise TFP shock at short horizons is non-trivial. Table 5a examines

[10] Alternatively, one could have estimated the entire model, but then all of the parameters would have differed from those in KS, raising the question of whether these differences are driving our results. We instead kept the DSGE model as close as possible to the original KS model while fitting key macro moments that are central to the performance of the max share estimator.

Table 5: Forecast error variance decompositions for TFP in the KS DSGE model

(a) Measured TFP ($\ln \mathrm{TFP}^u$)

	KS Model Measurement Error			KS Model with Calibrated Shock Processes Measurement Error		
h	News Shock	Surprise Shock	MEI Shock	News Shock	Surprise Shock	MEI Shock
1	71.5	1.7	26.7	15.5	81.7	2.8
8	83.9	0.4	15.7	50.6	41.3	8.0
20	61.4	0.3	38.3	53.2	10.4	36.4
40	86.1	0.2	13.7	64.8	2.3	32.9
80	94.8	0.1	5.1	76.0	1.0	23.0

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

	KS Model Measurement Error			KS Model with Calibrated Shock Processes Measurement Error		
h	News Shock	Surprise Shock	MEI Shock	News Shock	Surprise Shock	MEI Shock
1	96.2	3.8	0.0	12.9	87.1	0.0
8	99.7	0.3	0.0	75.0	25.0	0.0
20	99.9	0.1	0.0	89.6	10.4	0.0
40	99.9	0.1	0.0	94.8	5.2	0.0
80	100.0	0.0	0.0	97.4	2.6	0.0

Notes: h is the horizon of the variance decomposition and MEI is the marginal efficiency of investment.

whether the same problem arises in the larger-scale KS model. The left side shows the forecast error variance decomposition for measured TFP under the KS parameterization of the shock processes, while the right side shows the corresponding results when the shock processes are calibrated to the data. Under the KS parameterization, the news shock plays a substantial role at all horizons, so one would expect the KS estimator to perform well. In contrast, under the calibrated parameterization, the relative importance of news shocks is diminished at all horizons, while that of surprise and MEI shocks increases. This would lead one to expect much lower accuracy from the KS estimator than under the KS parameterization, much like in the baseline model.

Figure 3 plots the responses of output and measured TFP under alternative parameterizations of the KS DSGE model for $T = 10{,}000$. The top panel shows the results in the absence of mea-

Figure 3: KS max share identified responses to news shocks based on the KS DSGE model

(a) KS model without measurement error, $T = 10{,}000$

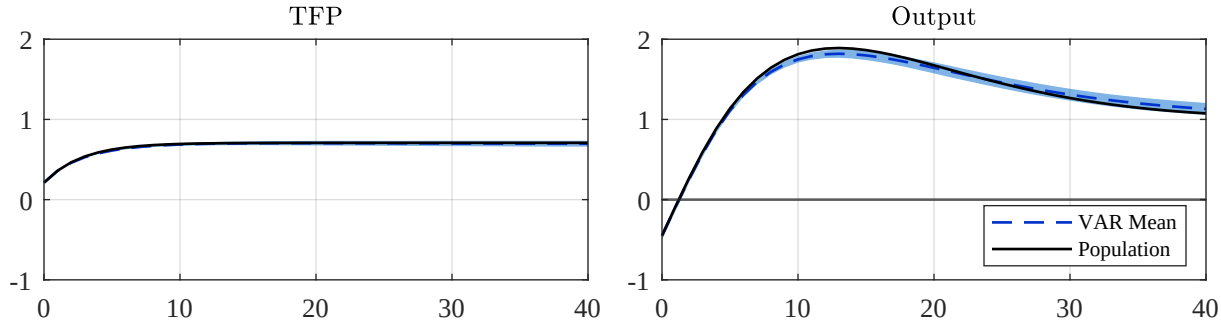

(b) KS model with measurement error, $T = 10{,}000$

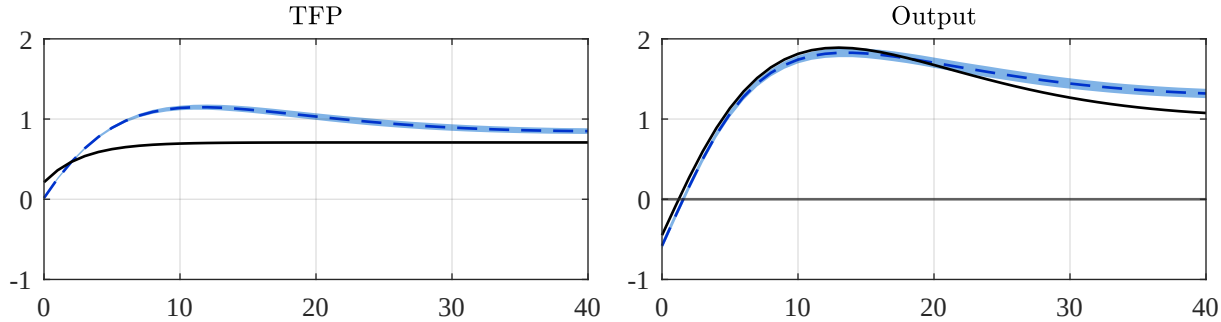

(c) KS model with calibrated shock processes and measurement error, $T = 10{,}000$

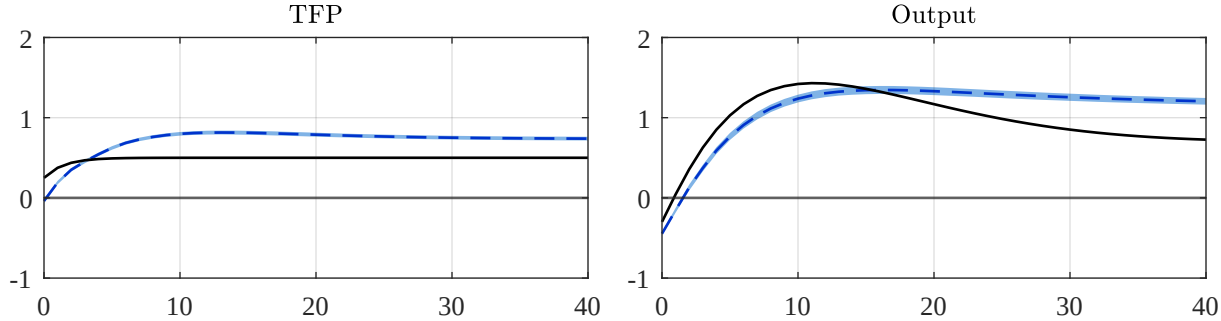

Notes: Based on a VAR(4) model for $\mathbf{y}_t = (a_t, y_t, i_t)'$ when there is no measurement error and $\mathbf{y}_t = (\mathrm{TFP}_t^u, y_t, i_t)'$ where there is measurement error. The responses from models with measurement error are not scaled given that the estimates of the TFP impact responses tend to be close to zero.

surement error. As expected, the KS estimator does an excellent job at recovering the population responses when the sample size is sufficiently large. The middle panel shows the results with measurement error under the KS TFP parameterization. While the bias of the output response is slightly larger, the fit remains quite good. There is a notable discrepancy between the response of

measured TFP to a news shock and the population response of true TFP, which is also apparent in the simulation evidence reported in KS. However, this result is not surprising. With measurement error there is no reason to expect the VAR to recover the TFP response, since the VAR is estimated with measured TFP and the population response is based on true TFP. The bottom panel reports results for the same model after calibrating the shock processes. Not only do the discrepancies between the TFP responses remain, but there now is strong bias in the output responses, sometimes in the positive direction and sometimes in the negative direction. As expected, the variance of the estimator increases when $T = 240$, but the impulse response patterns are similar.[11]

One might argue that this result is not surprising given the belief that news shocks must account for a large part of the unpredictable variation in measured TFP at long horizons for the KS estimator to perform well. Of course, we have no way of judging how important news shocks are for measured TFP. All we know is that news shocks are expected to explain most of the long-run variation in true TFP. As shown in Table 5b, it remains true under our calibration that news shocks explain 97% of the long-run variability of true TFP, but this does not mean that they are as important a determinant of measured TFP. Table 5a shows that news shocks explain only 76% of the long-run variation in measured TFP, which is almost identical to the share KS obtained when applying their estimator to actual data. This example illustrates that DSGE models that are consistent with the data may not satisfy one of the maintained assumptions required for the KS estimator to perform well, which provides another possible explanation of the bias in the KS estimator.

5 ALTERNATIVE ESTIMATORS INVOLVING DIRECT MEASURES OF TFP NEWS

The bias of the KS estimator even 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 recent studies. For example,

[11] Throughout the paper, we followed KS in setting the lag order of the VAR to 4. For $T = 10{,}000$, the accuracy of the KS estimator can be improved by increasing the lag order up to a certain point. This, however, comes at the cost of substantially increasing the variability of the estimator for $T = 240$.

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 KS 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).

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 and like the KS approach, 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 in-

strument 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 SIMULATION EVIDENCE ON THE ACCURACY OF THE NEWS VARIABLE ESTIMATORS

We start by using the baseline DSGE model to compare the accuracy of the news-based estimators to that of the KS estimator in the absence of TFP measurement error. We fit a VAR(4) model with intercept to $\mathbf{y}_t = (z_t, a_t, y_t)'$ for the max share news estimators and to $\mathbf{y_t} = (a_t, y_t, i_t)'$ for the KS estimator.[13] The variables enter the VAR in logs and are directly observable in the DSGE model. The choice of these variables is dictated by our interest in constructing the responses of TFP and output.[14]

Table 6 compares the RSME of the impulse responses to surprise and news shocks across these three estimators. The first four columns show the sum of the RMSEs over horizons 0 through 40 for output and TFP. The last column shows the sum of these entries across the four response

[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 related discussion see, e.g., Leeper et al., 2013; Mertens and Ravn, 2014). This renders the VAR model nonfundamental and distorts the impulse response estimates (see, e.g., Kilian and Lütkepohl, 2017). 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.

[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 $\mathbf{y}_t = (z_t, a_t, y_t)'$, 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 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.

Table 6: Sum of the RMSE over 40 quarters for each estimator based on the baseline DSGE model

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

	TFP Response		Output Response		
Estimator	News Shock	Surprise Shock	News Shock	Surprise Shock	Total
KS Max Share	11.3	2.6	14.5	2.3	30.7
Max Share News	1.4	0.8	1.9	1.1	5.2
Cholesky News	1.2	0.9	1.7	1.2	5.1
Alt KS Max Share	3.6	1.4	4.9	1.7	11.6

(b) $T = 240$

	TFP Response		Output Response		
Estimator	News Shock	Surprise Shock	News Shock	Surprise Shock	Total
KS Max Share	15.5	3.5	19.2	5.2	43.3
Max Share News	11.3	4.7	14.4	6.4	36.8
Cholesky News	8.8	5.5	11.3	7.3	33.0
Alt KS Max Share	15.1	2.7	19.4	4.5	41.6

Notes: The VAR includes $\mathbf{y}_t = (a_t, y_t, i_t)$ for the KS max share estimator and $\mathbf{y}_t = (z_t, a_t, y_t)$ for the max share news and Cholesky estimators. The Alt KS Max Share estimator uses the KS identification strategy and the Max Share News model variables.

functions. There is compelling evidence that the max share news estimator has substantially lower RMSE than the KS estimator not only for $T = 10{,}000$ (83% reduction in the RMSE), but in realistically small samples (15% reduction in the RMSE). The RMSE reductions obtained from the Cholesky news estimator is even larger with 85% for $T = 10{,}000$ and 29% for $T = 240$. These improvements in accuracy are mainly due to RMSE reductions for the responses to news shocks.

These results suggest that identification strategies based on TFP news variables perform much better than the KS estimator when both TFP and TFP news are accurately measured. For illustrative purposes, Figure 4 plots the responses of TFP and output to news and surprise shocks obtained using the max share news estimator. The top two rows show the results for $T = 10{,}000$. Both shocks appear properly identified by the max share approach with very little bias in the mean estimates and small variance. The bottom two rows show the corresponding results for $T = 240$. There is some bias in the responses in this case, but there is also an increase in the variability of the

Figure 4: Max share news identified impulse responses based on the baseline DSGE model

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

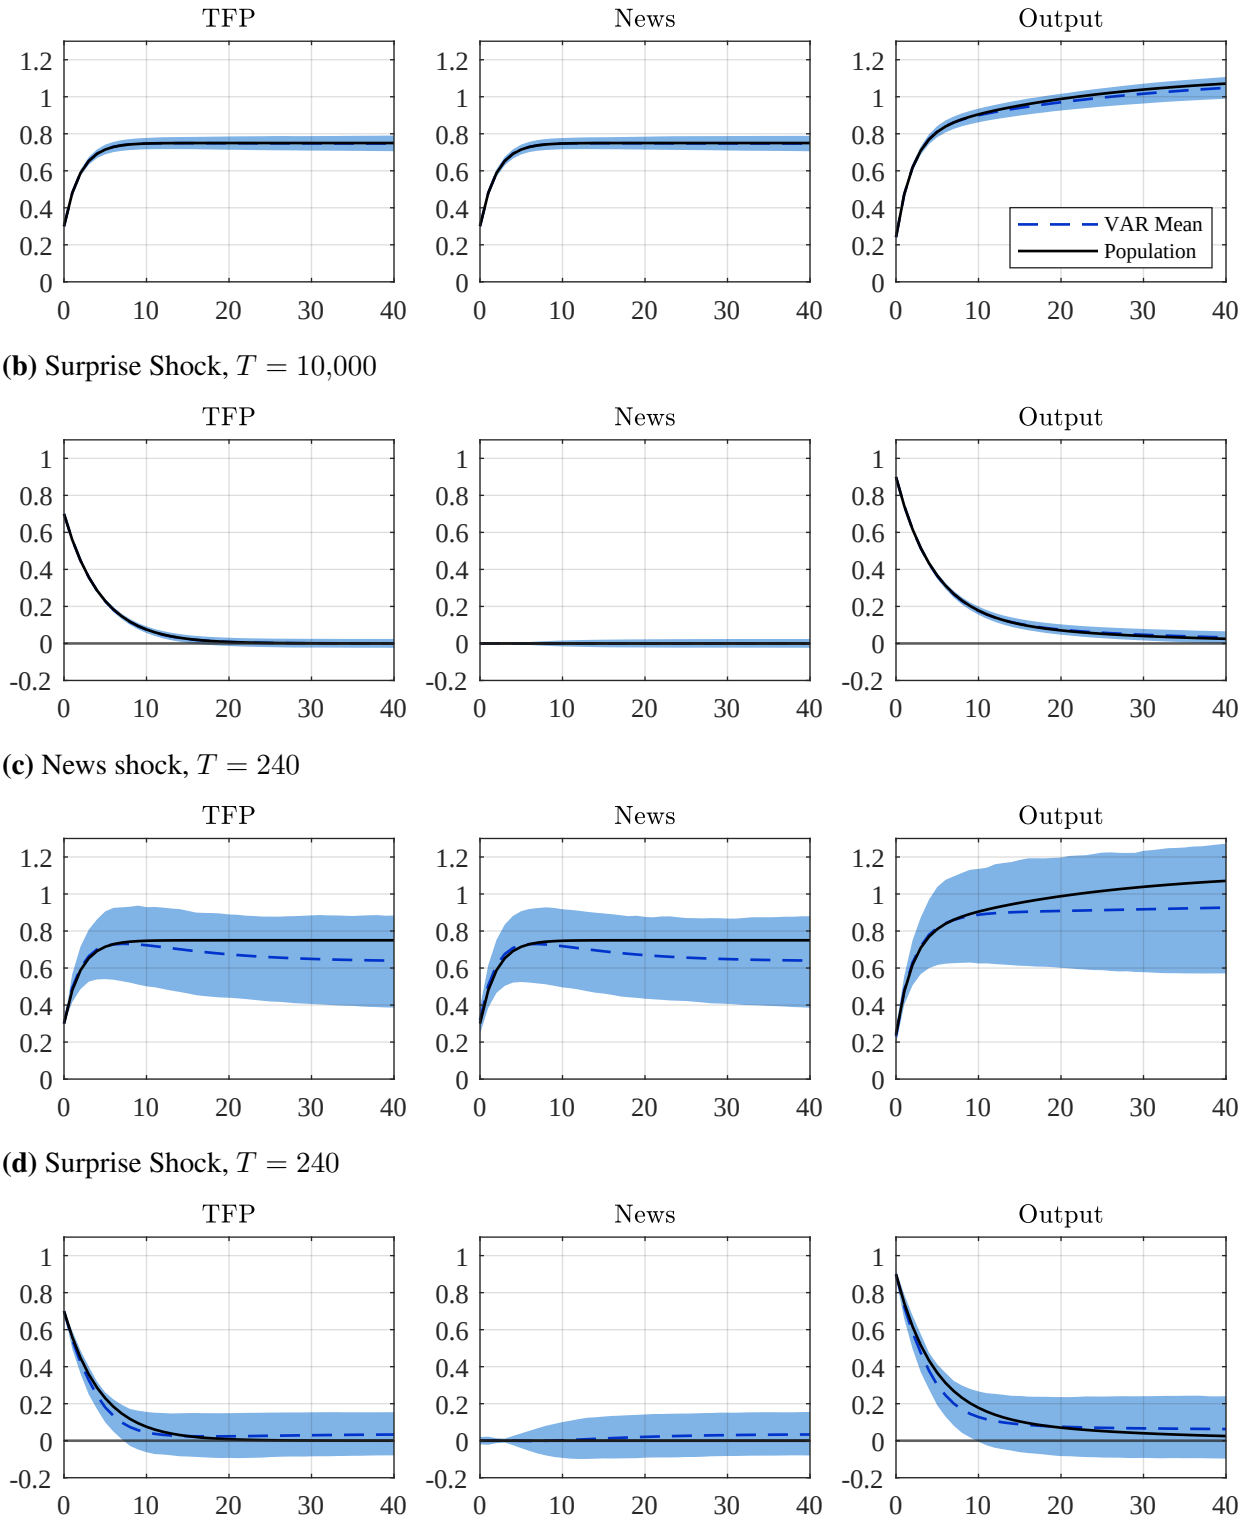

(b) Surprise Shock, $T = 10{,}000$

(c) News shock, $T = 240$

(d) Surprise Shock, $T = 240$

Notes: Based on a VAR(4) model for $\mathbf{y}_t = (z_t, a_t, y_t)'$. The response estimates have been scaled, so the impact response of TFP matches the population value.

estimator, as measured by the 68% quantiles of the distribution of VAR estimates. Thus, one would not necessarily expect the max share news estimator to be reliable in small samples. Qualitatively similar results hold for the Cholesky news approach, as shown in Appendix E.

5.3 WHY DOES THE KS MAX SHARE ESTIMATOR UNDER-PERFORM?

It may seem puzzling that the KS max share estimator performs so much worse in Table 6 than the news-based estimators. One potential reason is that the KS VAR model does not include a direct measure of TFP news. As noted by Leeper et al. (2013) and Mertens and Ravn (2014), for example, 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. This insight suggests that the comparatively high RMSE of the KS estimator may be explained by the absence of a forward-looking variable that captures TFP news in their VAR model.[15] The last row in the table explores this conjecture by applying the KS identification strategy to the max share news VAR model that includes z_t. Our simulations show that indeed the accuracy of the KS max share estimator improves when incorporating TFP news in the VAR model, but even in that case the KS estimator remains less accurate than the news-based estimators.

5.4 IMPACT OF MEASUREMENT ERROR IN TFP NEWS

In our analysis above, 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 allow the TFP news variable in the VAR 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 in the VAR model with $z_t^n = z_t + \sigma_n \epsilon_t^n$, where $\epsilon_t^n \overset{iid}{\sim} \mathbb{N}(0,1)$.

While there is no way of knowing the extent of measurement error in TFP news, Table 7 shows that our news-based estimators improve accuracy even if the news variable is measured

[15] While KS include such forward looking variables in some of their empirical work, they do not include them in their simulation study.

Table 7: Sum of the RMSE over 40 quarters for each estimator based on the baseline DSGE model

	TFP Response		Output Response		
Estimator	News Shock	Surprise Shock	News Shock	Surprise Shock	Total
KS Max Share	11.3	2.6	14.5	2.3	30.7
Max Share News ($\sigma_n = 0$)	1.4	0.8	1.9	1.1	5.2
Max Share News ($\sigma_n = 0.2\sigma_g$)	1.6	0.8	2.4	1.1	6.0
Max Share News ($\sigma_n = 0.5\sigma_g$)	3.6	0.9	5.0	1.3	10.8
Cholesky News ($\sigma_n = 0$)	1.2	0.9	1.7	1.2	5.1
Cholesky News ($\sigma_n = 0.2\sigma_g$)	1.3	1.1	1.6	1.5	5.5
Cholesky News ($\sigma_n = 0.5\sigma_g$)	1.7	2.9	2.3	3.7	10.6

Notes: Results based on $T = 10{,}000$. The VAR includes $\mathbf{y}_t = (a_t, y_t, i_t)'$ for the KS max share estimator and $\mathbf{y}_t = (z_t^n, a_t, y_t)'$ for the max share news and Cholesky news estimators.

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 more accurate than the original KS estimator with $T = 10{,}000$. Similarly accurate estimates of the responses to news shocks are obtained when allowing the measurement error to be serially correlated.

5.5 THE 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 fitting VAR(4) models to data for $\mathbf{y}_t = (z_t, a_t, y_t, i_t)$ simulated from the KS model for $T = 10{,}000$, either under their parameterization or under our alternative parameterization discussed in Section 4.

The top panel of Figure 5 plots the impulse responses under the KS parameterization for the max share news estimator. There is a very good fit, except for the response of measured TFP for the reasons discussed in Section 4. The bottom panel plots the impulse responses under our calibrated shock processes. Even in that case, there is only modest bias in the responses of output and investment. The Cholesky news estimator produces similar results, as shown in Appendix E.

Table 8 quantifies the improvement in accuracy. It reports the RMSE of the impulse response estimator associated with each method under our calibration. As in the case of no measurement

Figure 5: Max-share news identified responses to news shocks based on the KS DSGE model

(a) KS model with measurement error, $T = 10{,}000$

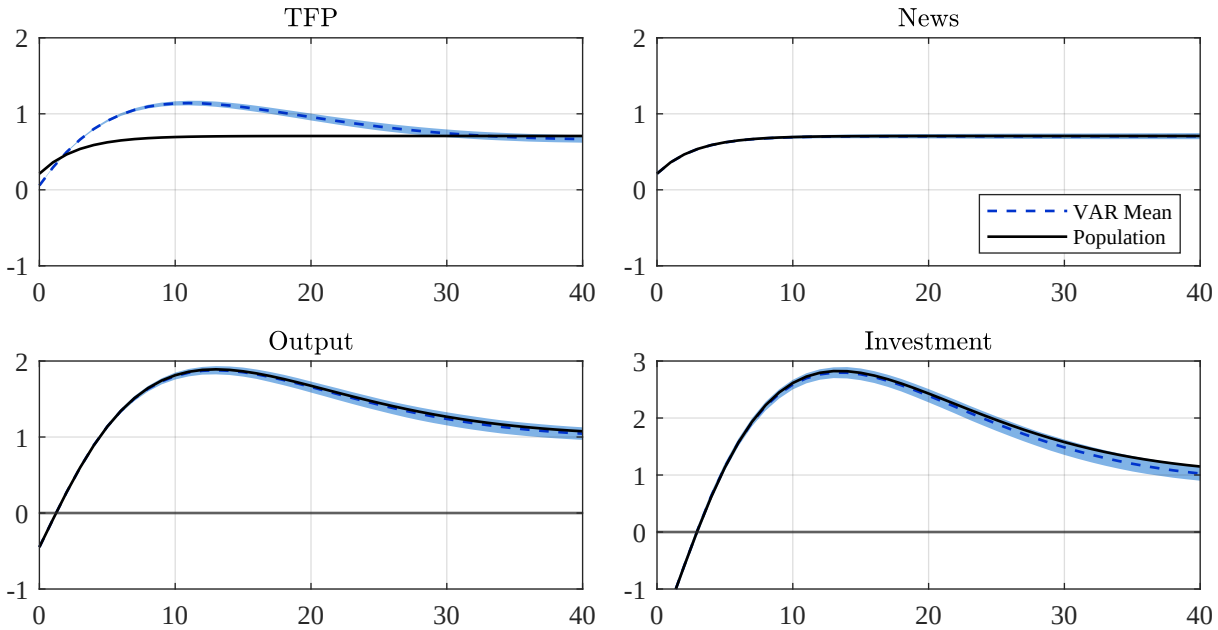

(b) KS model with measurement error and our calibrated shock processes, $T = 10{,}000$

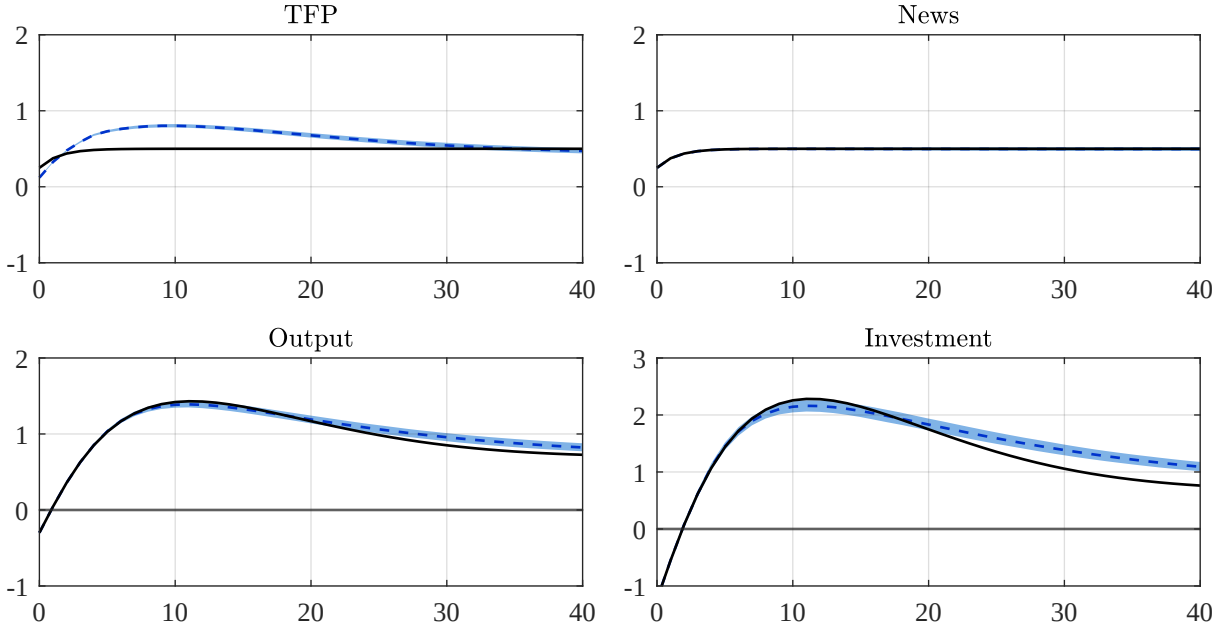

Notes: Based on a VAR(4) model for $\mathbf{y}_t = (z_t, \mathrm{TFP}_t^u, y_t, i_t)'$. The responses are not scaled given that the impact responses of TFP are very close to zero.

Table 8: Sum of the RMSE over 40 quarters for each estimator based on the KS DSGE model

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

Estimator	TFP	Output	Investment	Total
KS Max Share	10.2	11.3	20.9	42.4
Max Share News	6.1	3.0	8.3	17.4
Cholesky News	6.1	3.0	8.2	17.3
Alt KS Max Share	6.6	5.1	14.3	26.0

(b) $T = 240$

Estimator	TFP	Output	Investment	Total
KS Max Share	10.4	16.3	34.1	60.9
Max Share News	9.0	14.0	23.8	46.9
Cholesky News	8.6	13.5	22.7	44.9
Alt KS Max Share	8.7	16.0	33.7	58.4

Notes: Results based on the KS model with calibrated shock processes. The VAR includes $\mathbf{y}_t = (\mathrm{TFP}_t^u, y_t, i_t)$ for the KS max share estimator and $\mathbf{y}_t = (z_t, \mathrm{TFP}_t^u, y_t, i_t)$ 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.

error, adding the TFP news variable and adapting the identification procedure produces substantial improvements in the accuracy of the impulse response estimator relative to the KS approach, even when there is substantial measurement error in the TFP news variable (see Appendix E). For example, the RMSE is reduced by 59% for $T = 10{,}000$ and by between 23% and 26%, depending on the method, for $T = 240$. Thus, the max share news and Cholesky news estimators improve accuracy, both in the presence of TFP measurement error and in its absence. Qualitatively similar results hold even when applying the KS estimator to the max share news model that includes z_t, as discussed in Section 5.3.

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. In this

section, we therefore 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).[16]

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 by KS.[17] The sample for each VAR varies due to differences in the availability of the news variables. We identify the 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

[16]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.

[17]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 9: Empirical results from VAR models including alternative measures of TFP news.

Criterion	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 $H_n = 4$.

these criteria. For example, a response may turn marginally negative a longer horizons. Our focus in this section is on strong violations of these criteria. Table 9 summarizes the results based on a maximum horizon of 40 quarters. The full set of impulse responses is provided in Appendix E. We find that only the R&D and ICT models satisfy the two criteria. This is true regardless of whether we use the max share news or Cholesky news identification method.[18]

Figure 6 shows that the responses of other key variables also look reasonable for the ICT model. Similar results hold when using the R&D model.[19] 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

[18]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. One 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 does not affect the results reported in Table 9.

[19]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.

Figure 6: Comparison of max share news and KS max share identified impulse responses

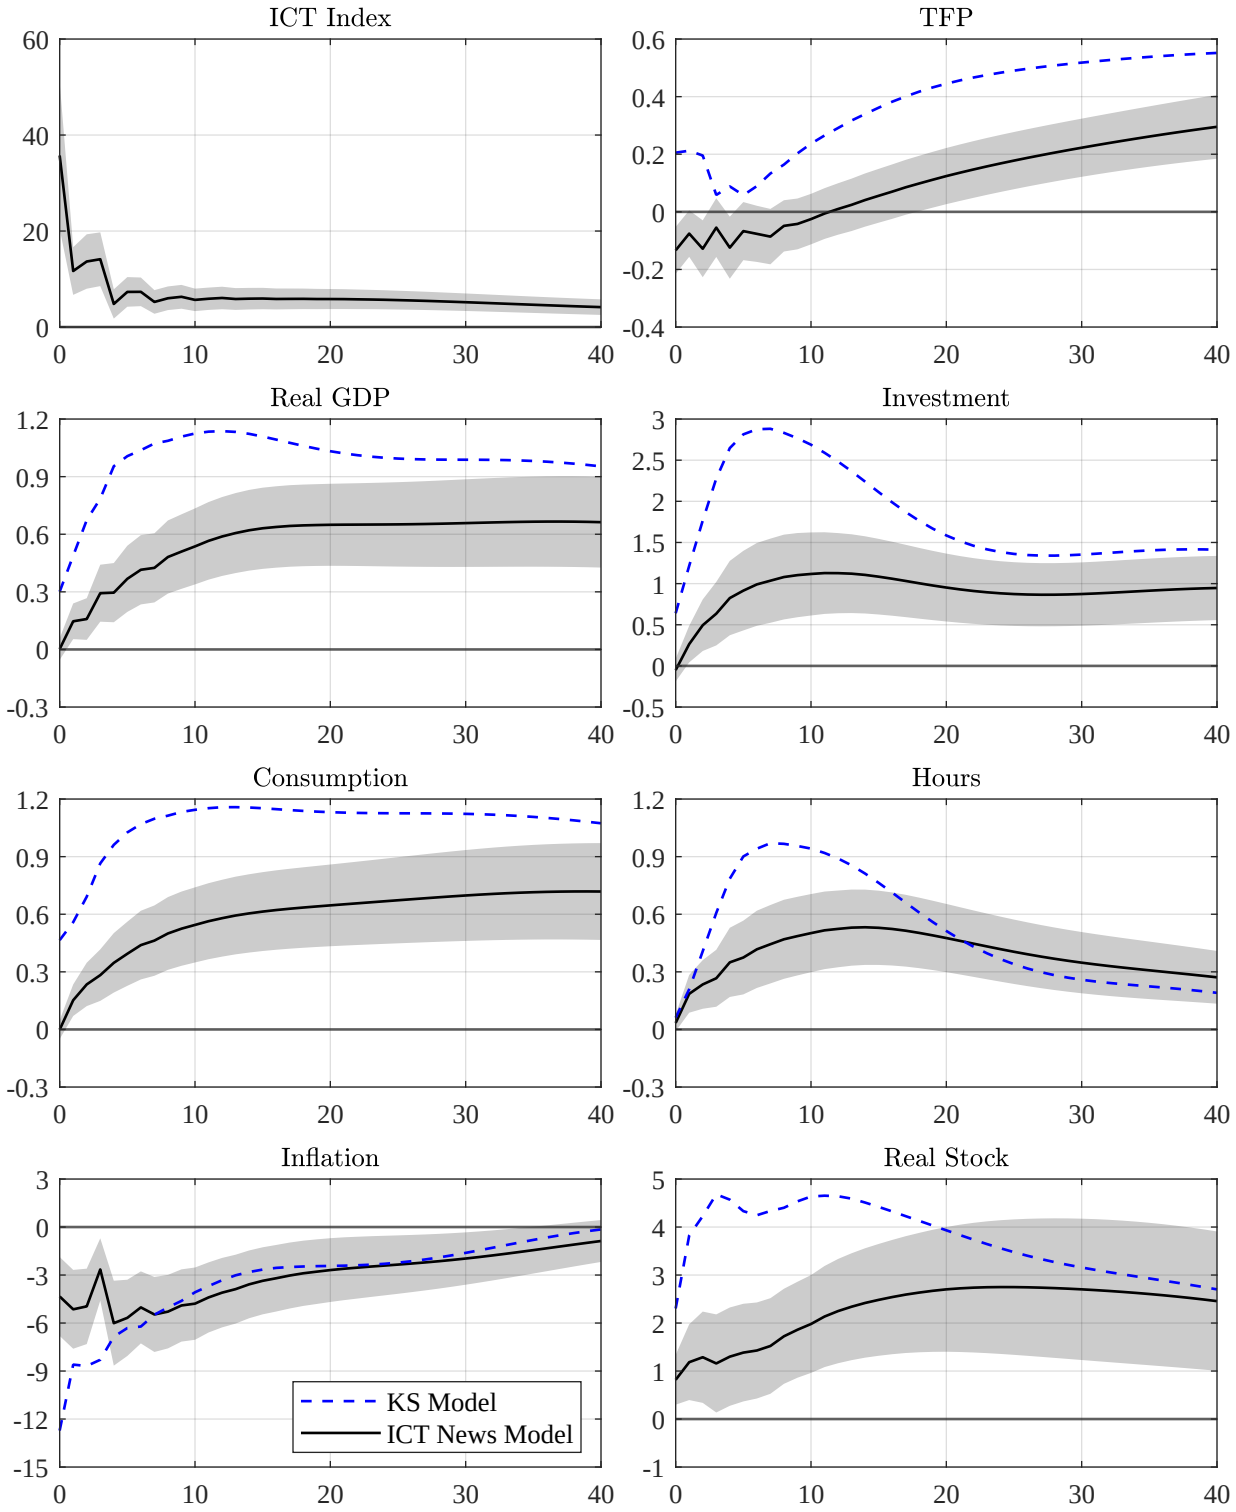

Notes: VARs estimated on identical samples from 1960-2014. Shaded regions represent 1-standard deviation error bands computed by bootstrap for the ICT news model.

Portier (2006). We plot these response estimates next to the estimates from the original 8-variable KS VAR model 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.

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

Table 10 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 state-of-the-art method in the literature.

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

Table 10: Forecast error variance decompositions

	Max Share News				KS Max Share			
	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.

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. The state of the art in this literature is the max share identification strategy recently proposed by KS, which allows news shocks to affect TFP contemporaneously. Using data generated from a range of DSGE models that are calibrated to match key data moments, we showed that the KS estimator is strongly biased, 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.

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 state-of-the-art 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.

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.

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

CARRIERO, A. AND A. VOLPICELLA (2022): "Generalizing the Max Share Identification to Multiple Shocks Identification: An Application to Uncertainty," University of Surrey Discussion Paper 03/22.

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

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.

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.

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.

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.

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.

MERTENS, K. AND M. O. RAVN (2014): "A Reconciliation of SVAR and Narrative Estimates of Tax Multipliers," *Journal of Monetary Economics*, 68, 1–19.

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

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