Estimating Macroeconomic News and Surprise Shocks*

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

June 12, 2025

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. 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 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, Todd Walker, Marco del Negro, and five referees 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). 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, such as Barsky and Sims (2011), 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 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 due to the slow diffusion of

[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), Forni 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), Chahrour et al. (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.

new technologies. A further refinement of the max share estimator, known as the NAMS estimator, was introduced in Dieppe et al. (2021).

It is widely believed that the TFP max share estimator of news shocks works well as long as news shocks account for the bulk of the variation in TFP at long horizons. For example, Kurmann and Sims (2021) express confidence that, generally, the max share identification performs well as long as news shocks account for a large part of the unpredictable variation in productivity 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. 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 responses to news shocks are not well-identified by the TFP max share estimator 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 will substantially reduce the asymptotic bias. We propose a novel identification strategy based on maximizing the variance share of the TFP news variable at a short horizon (as opposed to the variance share of TFP at a long horizon). This max share news estimator is more appealing than the alternative identification strategy of treating TFP news as an instrument because the impact effect of TFP news on TFP may be zero in population.

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 such an estimator to recover the population responses from data generated by DSGE models.[4] We first show that the max share news estimator tends to have

[4]Examples of empirical studies employing measures of TFP news include Shea (1999), Christiansen (2008), Alexopoulos (2011), Baron and Schmidt (2019), Cascaldi-Garcia and Vukotić (2022), and Miranda-Agrippino et al. (2024).

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 this news-based estimator in the presence of TFP measurement error and show that it still substantially reduces the bias and RMSE of the responses to news shocks compared to TFP max share estimators. While TFP news is not perfectly observed in the data, the superior accuracy of the max share news estimator is robust to sizable and possibly persistent measurement error in the news variable.

These improvements in accuracy can be traced to two key differences in the construction of the estimator. One is the reduced-form VAR specification. A fundamental concern with using the TFP max share estimator is that economic agents are likely to account for TFP news in forming expectations about future TFP. If this information is omitted from the VAR information set, the estimator cannot be expected to recover the structural stocks (see, e.g., Hansen and Sargent, 1991; Leeper et al., 2013). We present evidence that including a direct measure of TFP news in the VAR model substantially alleviates the large-sample bias of TFP max share estimators, but that the max share news estimator based on the same VAR model is even less biased and has even lower RMSE. Thus, the higher large-sample bias of the TFP max share estimator cannot be explained by an information deficiency alone. The other explanation is that TFP max share estimators tend to confound the news shock with non-news shocks even when the news shock is dominant, which creates asymptotic bias in the response estimator. Since the max share news estimator does not rely on longer-run restrictions, it dispenses with this concern.[5]

Our third contribution is to empirically illustrate the use of TFP news for identifying news shocks using 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

[5] Kurmann and Sims (2021) acknowledge that non-news shocks could be important at longer horizons, causing the accuracy of the TFP max share estimator to deteriorate, but consider that possibility unlikely. Dieppe et al. (2021) caution that some contamination of the news shock is inevitable. For a more in-depth analysis of the conditions required for valid TFP max share identification see Dou et al. (2025).

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.[6] 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 5-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 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 a 5-year horizon.

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. 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 alternative identification strategies based on VAR models that include a direct measure of TFP news. 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 reviews three TFP max share estimators for identifying news 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,

[6] 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. (2024), and Bouakez and Kemoe (2023).

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 ESTIMATORS

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 max share 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 target 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 range of horizons, 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 shocks with less persistent effects may contaminate the estimate 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 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 horizon H_n. 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 $H_n = 80$.[7]

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

As shown in Appendix E, in the absence of TFP measurement error the estimate of γ_n directly implies an estimate of γ_s without further identifying assumptions. We will make use of this result in Sections 3 and 4 when evaluating the KS max share and NAMS estimators in models not subject to TFP measurement error.

It should be noted that Kurmann and Sims (2021) and Dieppe et al. (2021) do not claim that their estimators are consistent. Rather they propose these estimators as an improvement over earlier TFP max share estimators with the understanding that in practice the estimated news shock could be contaminated by other shocks undermining their accuracy. How accurate these improved estimators are in practice must be evaluated by simulation.

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 $j \in [0,1]$ produces a differentiated good $y_t(j) = a_t k_{t-1}(j)^\alpha n_t(j)^{1-\alpha}$, where $k_{t-1}(j)$ and $n(j)$ 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. Agents anticipate the effects of this shock when forming expectations, consistent with the interpretation of a news shock. Allowing news shocks to contemporaneously affect TFP does not materially change our results.

Each intermediate firm chooses its inputs to minimize costs, $w_t n_t(j) + r_t^k k_{t-1}(j)$, 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 $y_t(j)$ units from each intermediate-goods firm to produce the final good, $y_t \equiv [\int_0^1 y_t(j)^{(\epsilon_p-1)/\epsilon_p} dj]^{\epsilon_p/(\epsilon_p-1)}$, where $\epsilon_p > 1$ measures the elasticity of substitution between intermediate goods. It then maximizes dividends to determine the demand function for good i, $y_t(j) = (p_t(j)/p_t)^{-\epsilon_p} y_t$, where $p_t = [\int_0^1 p_t(j)^{1-\epsilon_p} dj]^{1/(1-\epsilon_p)}$ is the aggregate price level.

Following Calvo (1983), a fraction, θ_p, of intermediate firms cannot choose their price in period t. Those firms index their price to steady-state inflation, so $p_t(j) = \bar{\pi} p_{t-1}(j)$. 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

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

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

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

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

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

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

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

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.

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 from 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 ($SD(\tilde{y}_t)$, $SD(\tilde{\imath}_t)$).[8] 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.[9]

[8] 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.

[9] 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 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.

3.3 SIMULATION EVIDENCE Since there are three structural shocks in the DSGE model, we fit a three-dimensional structural VAR model.[10] 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 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.

Given the concerns about the BS estimator raised in Kurmann and Sims (2021), Figure 1 focuses on the results for the KS max share estimator, but we provide summary statistics for all three TFP max share estimators below. 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. The responses to the surprise shock shown in the bottom row are also biased.

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 similar for the KS max share estimator and NAMS estimator, suggesting that there is little to choose between them. The BS estimator is somewhat more accurate. This result is expected given that there is no measurement error in TFP and the

[10] Adding more variables would render the error covariance matrix singular, invalidating the VAR analysis.

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

(a) News shock

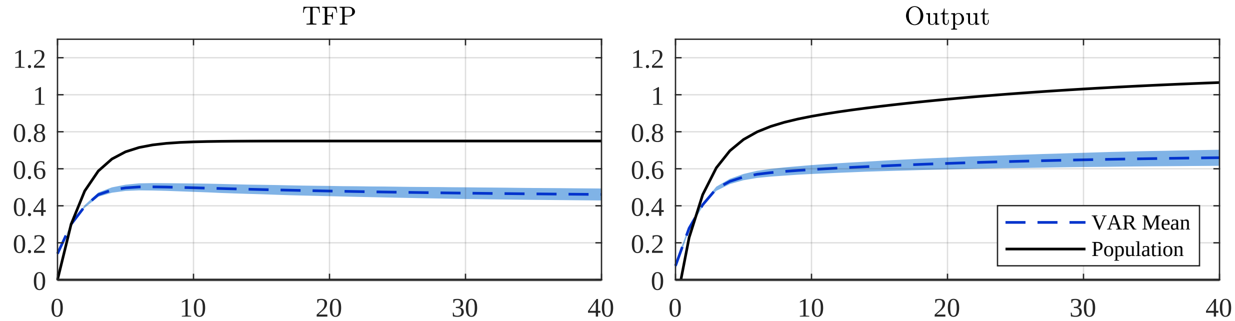

(b) Surprise shock

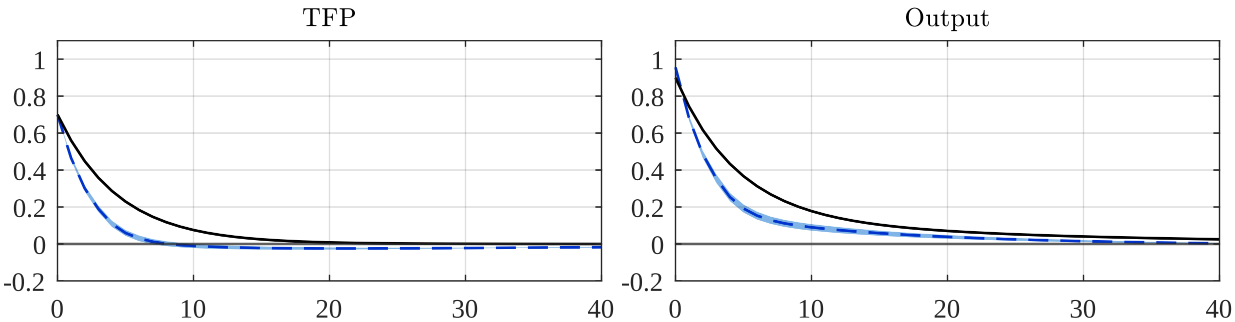

Notes: VAR(4) model with $T = 10{,}000$ and $\mathbf{y}_t = (a_t, y_t, i_t)'$. The responses are scaled so the estimated response of TFP matches the population value when the shock first takes effect.

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.[11]

There is a widespread belief that 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 account for only

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

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	9.9	1.9	12.8	2.0	26.7

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

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.

3% of the variation in TFP growth at horizon 80. Thus, even when surprise shocks (or other non-news shocks) are seemingly negligible in population at the horizons that matter for computing the estimator, TFP max share estimators tend to confound news and surprise shocks. If the TFP max share estimator does not work in this simple setting, there is no reason to expect it to work better in more general settings. In the next section, we show that this result does not rest on the assumption that the TFP measure is not contaminated by other structural shocks. Indeed, allowing for such contamination only makes the identification more challenging.

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 there are three shocks driving the measured TFP data. As a result, identifying surprise TFP shocks is not possible without further identifying restrictions. However, the news shock may 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 measurement error is recovered from the simulated model data by mimicking how Fernald (2014) construct the TFP variable. 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.

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

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

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

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

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

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).

Just like in the baseline model, we fit the shock processes to the data. The only difference is that we now infer the parameters of the process governing a_t by matching the moments for measured TFP in the model (TFP_t^u) to the observed TFP data. 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.

News shocks are the dominant determinant of the long-run variation in measured TFP in this larger-scale DSGE model. As illustrated in Table 5, news shocks explain almost all of the long-run variability of true TFP and 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.

Figure 2 plots the responses of output and measured TFP for $T = 10{,}000$. We focus on the KS max share estimator as before. 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 6 shows the RMSE for the KS max share estimator, alongside the other TFP max share estimators considered in Section 3, in the presence of TFP measurement error. Once again, the performance of the KS max share and NAMS estimators is similar.[12] These results show that the

[12] Surprisingly, the BS estimator is more accurate than the other TFP max share estimators, but the BS estimates of 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.

Table 5: Forecast error variance decompositions for TFP in the larger-scale DSGE model

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

	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.

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

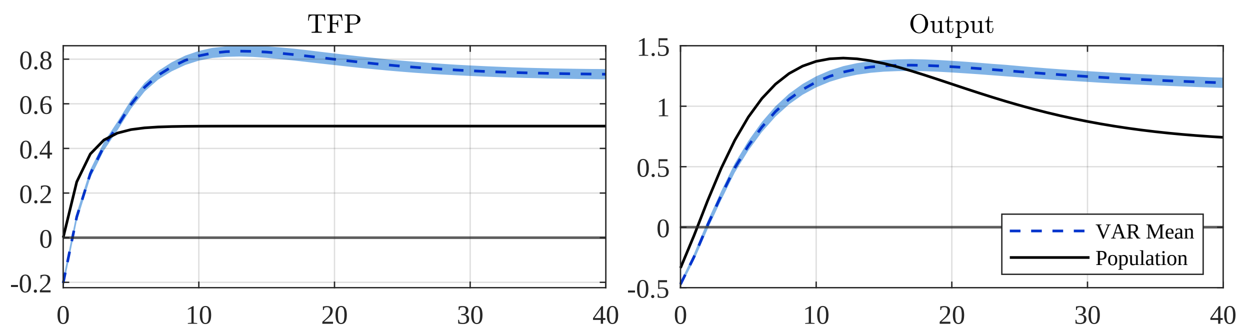

Notes: VAR(4) model with $T = 10{,}000$ and $\mathbf{y}_t = (\text{TFP}_t^u, y_t, i_t)'$.

Table 6: 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	10.1	10.8	21.3	42.2

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

concerns about the accuracy of the TFP max share estimator are robust to the presence of TFP measurement error.[13]

The strong bias of the TFP max share estimators in Figure 2 and their high RMSEs in Table 6 call into question the conventional wisdom that a sufficient condition for the validity of TFP max share estimators is that the news shock explains a large share of the long-run variation in measured TFP. Our evidence illustrates that neither the KS max share estimator nor the NAMS estimator are able to accommodate TFP measurement error. This 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 G, we show that this result is an artifact of their parameter choices for the TFP process. Under their parameterization, TFP growth is persistent, which is at odds with the data. Moreover, news shocks explain nearly all variation in TFP at all horizons, which greatly enhances the accuracy of the KS max share estimator because it effectively removes the identification challenge. When the parameters in the DSGE model are set to match the TFP moments in the data, in contrast, the impulse response estimator is strongly biased. This bias arises for any specification of the TFP news process that assigns a non-trivial role to non-news shocks in driving TFP at the relevant horizons, as illustrated in Appendix I.

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. TFP news data have been employed in a number of studies.[14] The premise of all these studies is that measures of TFP news should increase immediately as a positive news

[13]Similar conclusions also apply to max share estimators targeting labor productivity or output (see Appendix I).

[14]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 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 grants), Miranda-Agrippino et al. (2024, patent applications), and Fieldhouse and Mertens (2023, government R&D spending).

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 the merits of a novel identification strategy, 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 Our 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. We refer to this estimator as the "max share news" estimator. An obvious concern is that, in practice, the TFP news variable could be measured with error. We present simulation evidence below that the max share news estimator performs well with and without measurement error in the TFP news variable.[15]

An alternative approach to dealing with TFP news measurement error would have been 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. (2024). 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. However, as shown in Plagborg-Møller and Wolf (2021), the proxy VAR approach that uses the news variable as an external instrument is not valid when the shock of interest is non-invertible.[16]

Yet another approach would have been to treat the TFP news as an internal instrument, which

[15] When ordering the news variable and setting $H_n = 0$, the max share news estimator reduces to a block-recursive estimator as employed in Cascaldi-Garcia and Vukotić (2022), for example. The latter specification is similarly accurate in our simulations.

[16] 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 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 (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.

requires treating the news variable in the VAR as predetermined with respect to TFP (see, e.g., Plagborg-Møller and Wolf, 2021). While this helps address concerns about the shock of interest being noninvertible, neither the external instrument nor the internal instrument approach is appealing when estimating responses to news shocks. As discussed in Montiel Olea et al. (2021) and Plagborg-Møller and Wolf (2021), when constructing the IV estimator one needs to normalize the responses by scaling the response of interest by the impact response of the variable being instrumented. In our case, one would instrument the TFP residual by the TFP news instrument. As illustrated in Figure 1, one cannot rule out that the impact response of TFP to news is zero in population, resulting in responses that are infinite. This means that the IV estimator of the VAR responses is not well defined in general.

5.2 ACCURACY OF THE NEWS VARIABLE ESTIMATORS

We start by using the baseline DSGE model to examine the accuracy of 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.[17] 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 max share estimator. This timing mirrors the way observed TFP news has been used in applied work.[18] The choice of these variables is dictated by our interest in constructing the responses of TFP and output. All variables enter the VAR in logs and are directly observable in the DSGE model.

We first compare the max share news estimator to the KS max share and NAMS estimators.[19,20] Table 7 shows the RMSEs of the impulse responses to news and surprise shocks. As the first three

[17] In Appendix I, we show that our substantive findings are unaffected by the timing of the news shocks.

[18] It might seem more appropriate to compare the max share 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.

[19] 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.

[20] 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 or NAMS estimator.

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
NAMS	9.9	1.9	12.8	2.0	26.7
Max Share News	1.4	0.8	1.9	1.0	5.0
Alt KS Max Share	2.8	1.2	3.8	1.4	9.3
Alt NAMS	2.6	0.5	3.2	0.7	6.9

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 estimator. The Alt KS max share estimator uses the KS identification strategy and the max share news model variables. The Alt NAMS estimator uses the NAMS identification strategy and the max share news model variables.

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

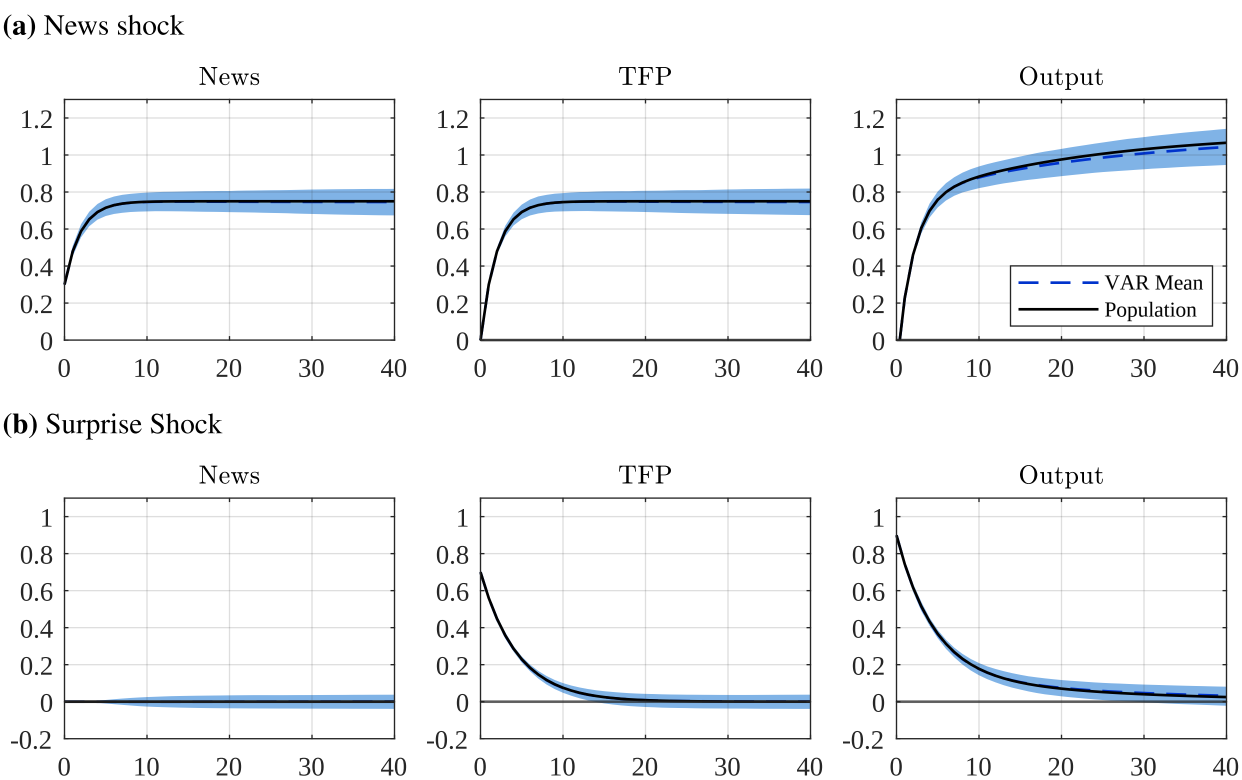

Notes: VAR(4) model with $T = 10{,}000$ and $\mathbf{y}_t = (z_{t+1}, a_t, y_t)'$. The responses are scaled so the estimated response of TFP matches the population value when the shock first takes effect.

rows show, the max share news estimator reduces the RMSE by 82% relative to the KS max share estimator and by 81% relative to the NAMS 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 news estimator with little bias in the mean estimates and small variability. These results suggest that identification strategies based on TFP news variables perform better than the TFP max share identification when both TFP and TFP news are accurately measured. Furthermore, the fact that the max share news estimator recovers the population responses illustrates that the reduced-form VAR specification is adequate for capturing the dynamics implied by the DSGE model.

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 max share news estimator. One reason is that the max share news estimator by construction is not subject to the confounding of the news shock. The other reason is that the VAR models underlying the TFP max share estimator do not include a direct measure of TFP news. This omission reflects standard practice in the literature. The inclusion of forward-looking variables such as TFP news is not mentioned as a prerequisite for TFP max share estimators to work in Barsky and Sims (2011), Kurmann and Sims (2021) and Dieppe et al. (2021). These studies evaluate the TFP max share estimator on data generated from DSGE models without including any news variables in the VAR model.[21] In this section, we demonstrate that including TFP news in the reduced-form VAR model is as essential for recovering news shocks as the news shock accounting for all of the variability in measured TFP at the horizons relevant for the construction of the estimator.

As noted by Hansen and Sargent (1991) and Leeper et al. (2013), when economic agents form

[21] While some empirical studies based on TFP max share estimators include forward-looking variables in their VAR models, typically these variables are not TFP news, but variables such as stock returns or consumer sentiment. The only empirical application of a TFP max share estimator to include indicators of TFP news that we are aware of is in Kurmann and Sims (2021, p. 232), but their objective is to illustrate that these indicators respond to news shocks as expected, not to use them for the identification of the news shock.

expectations based on information not contained in the information set of the VAR model, standard approaches to identifying structural shocks fail. This insight suggests that the comparatively high RMSE of the TFP max share estimator may be explained in part by the absence of a forward-looking variable that captures TFP news in the VAR model. In the last two rows of Table 7, we explore this conjecture by applying the KS max share estimator to the VAR model that includes z_{t+1}, labeled as the Alt KS max share estimator. We also report results for the corresponding Alt NAMS estimator. Our simulations show that indeed the accuracy of the KS max share and NAMS estimators substantially improves when incorporating TFP news into the VAR model. However, the Alt TFP max share estimators are less accurate than the max share news estimator based on the same reduced-form VAR model because they remain subject to the contamination of the news shock. The max share news estimator has 46% lower RMSE than the Alt KS max share estimator and 28% lower RMSE than the Alt NAMS estimator.[22]

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, next we 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 $z^n_{t+1} = z_{t+1} + \epsilon^n_t$, where $\epsilon^n_t = \rho^n \epsilon^n_{t-1} + \sigma^n \upsilon^n_t$ with $\upsilon^n \sim \mathbb{N}(0, 1)$.

Table 8 focuses on the responses to news shocks only. The left side shows results for the base-

[22] 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 is close to one 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. Unlike in Forni et al. (2019), the VAR models of interest in our paper are specified in log levels rather than on detrended data. Applying their procedure to the residuals from a VAR model in log levels suggests that even the R^2 for the VAR model not augmented with TFP news is close to one, indicating that the model is informationally sufficient. It is unclear, however, whether the Forni et al. (2019) procedure remains valid in our setting, so this result must be discounted. Instead, we assess the informational sufficiency of the VAR models more directly. Our results in Table 7 (and similarly in Table 8 for the model with TFP measurement error) show clear evidence of an informational deficiency when not including TFP news and that including TFP news remedies that deficiency.

Table 8: RMSE over 40 quarters with and without measurement error (ME) in the news variable

	Baseline Model			Larger-Scale Model			
Estimator	TFP	Output	Total	TFP	Output	Invest	Total
KS Max Share	10.2	13.1	23.2	10.3	10.4	19.3	40.0
NAMS	9.9	12.8	22.7	10.1	10.8	21.3	42.2
	(A) No ME ($\rho_n = 0$, $\sigma_n = 0$)						
Max Share News	1.4	1.9	3.3	6.4	2.8	7.6	16.8
Alt KS Max Share	2.8	3.8	6.7	7.0	4.9	13.9	25.7
Alt NAMS	2.6	3.2	5.8	6.9	4.5	12.9	24.3
	(B) 20% iid ME ($\rho_n = 0$, $\sigma_n = 0.2\sigma_g$)						
Max Share News	1.6	2.4	4.0	6.4	3.0	8.1	17.4
Alt KS Max Share	3.4	4.5	7.9	7.1	4.9	13.7	25.7
Alt NAMS	2.6	3.4	6.0	6.9	4.7	13.1	24.7
	(C) 50% iid ME ($\rho_n = 0$, $\sigma_n = 0.5\sigma_g$)						
Max Share News	3.6	4.9	8.5	6.4	3.6	9.1	19.1
Alt KS Max Share	5.1	6.9	12.0	7.4	5.3	13.8	26.5
Alt NAMS	3.9	5.3	9.2	7.1	5.1	13.2	25.4
	(D) 20% Persistent ME ($\rho_n = 0.5$, $\sigma_n = 0.2\sigma_g$)						
Max Share News	1.5	2.2	3.7	6.4	3.0	8.1	17.5
Alt KS Max Share	3.2	4.3	7.6	7.1	4.9	13.8	25.8
Alt NAMS	2.6	3.3	5.9	6.9	4.7	13.1	24.8
	(E) 50% Persistent ME ($\rho_n = 0.5$, $\sigma_n = 0.5\sigma_g$)						
Max Share News	3.3	4.6	7.9	6.3	3.7	9.3	19.3
Alt KS Max Share	4.9	6.6	11.4	7.4	5.4	14.0	26.8
Alt NAMS	3.6	5.0	8.6	7.1	5.2	13.4	25.8
	(F) 20% Highly Persistent ME ($\rho_n = 0.9$, $\sigma_n = 0.2\sigma_g$)						
Max Share News	1.5	2.1	3.5	6.4	3.0	8.1	17.5
Alt KS Max Share	3.2	4.3	7.6	7.1	5.0	13.8	25.8
Alt NAMS	2.6	3.3	5.9	6.9	4.8	13.2	24.9
	(G) 50% Highly Persistent ME ($\rho_n = 0.9$, $\sigma_n = 0.5\sigma_g$)						
Max Share News	2.6	3.6	6.2	6.3	3.9	9.5	19.6
Alt KS Max Share	5.2	6.9	12.1	7.5	5.6	14.2	27.3
Alt NAMS	3.9	5.3	9.2	7.2	5.5	14.0	26.7

Notes: VAR(4) model with $T = 10{,}000$. Baseline model: $\mathbf{y}_t = (a_t, y_t, i_t)'$ for the KS max share estimator and $\mathbf{y}_t = (z^n_{t+1}, a_t, y_t)'$ for the max share news estimator. Larger-scale model: $\mathbf{y}_t = (\mathrm{TFP}^u_t, y_t, i_t)'$ for the KS max share estimator and $\mathbf{y}_t = (z^n_{t+1}, \mathrm{TFP}^u_t, y_t, i_t)'$ for the max share news estimator. The Alt KS max share estimator uses the KS identification strategy and the max share news model variables. The Alt NAMS estimator uses the NAMS identification strategy and the max share news model variables.

line model. Panel A mirrors the results in the absence of TFP measurement error in Table 7. While there is no way of knowing the extent of measurement error in TFP news, Panels B and C show that the max share news estimator remains more accurate 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 ($\sigma_n = 0.5\sigma_g$), the max share news estimator is still 63% more accurate than the KS max share estimator.

While these results are promising, in practice one would expect TFP news to be systematically mismeasured, inducing persistent deviations between true and measured TFP news. Panels D-G show that similarly accurate responses to news shocks are obtained even when allowing the measurement error to be serially correlated, which mimics a situation in which TFP news are systematically mismeasured. The max share news estimator reduces the RMSE by between 66% and 85% relative to the KS max share estimator. Largely the same results hold for the NAMS estimator. The max share news estimator also remains more accurate than the Alt KS max share estimator by between 10% and 16% and more accurate than the Alt NAMS estimator by between 7% and 13%, depending on the specification of the measurement error.

5.5 NEWS-BASED ESTIMATORS IN THE PRESENCE OF TFP MEASUREMENT ERROR

An important question is whether the max share news estimator can also reduce the impulse response bias in the presence of TFP measurement error. We explore this question by fitting VAR(4) models to data for $\mathbf{y}_t = (z_{t+1}, \mathrm{TFP}^u_t, y_t, i_t)$ simulated from the larger-scale DSGE model for $T = 10{,}000$, where TFP^u_t denotes measured TFP.[23] As Figure 4 shows, even when TFP is mismeasured there is only modest bias in the responses of output and investment. The right side of Table 8, quantifies the improvement in accuracy relative to the other estimators. Panel A focuses on the case of no measurement error in TFP news. In this case, the max share news estimator reduces the RMSE by 58% relative to the KS max share estimator and 60% relative to the NAMS estimator. A smaller but still substantial reduction in accuracy of 35% occurs compared to the Alt KS max share esti-

[23] As shown in Appendix I, similar results would be obtained when augmenting the DSGE model with two more structural shocks and augmenting the dimension of the approximating VAR model accordingly.

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

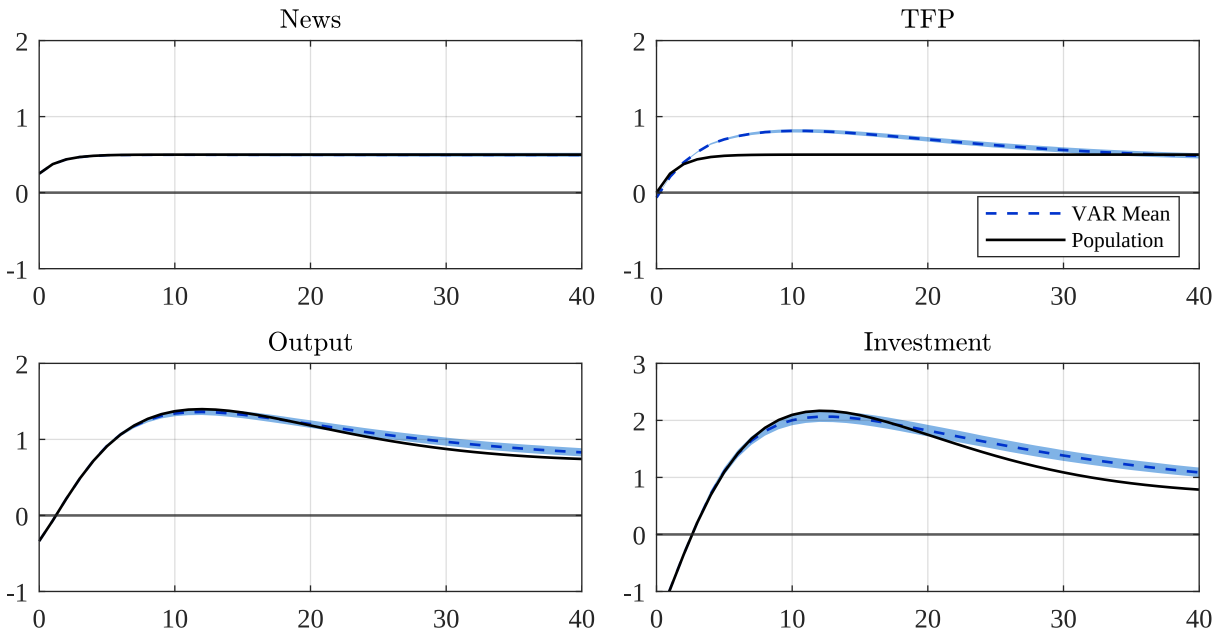

Notes: VAR(4) model with $T = 10{,}000$ and $\mathbf{y}_t = (z_{t+1}, \mathrm{TFP}^u_t, y_t, i_t)'$.

mator that includes the TFP news variable in the VAR model. The corresponding gain relative to the Alt NAMS estimator is 31%. In short, the max share news estimator is systematically more accurate. Panels B-G show the corresponding results when allowing both TFP and TFP news to be mismeasured. The improvements in accuracy from using the max share news estimator remain substantial in all cases relative to the KS max share and NAMS estimators. The gains are between 28% and 32% relative to the Alt KS max share estimator and between 25% and 30% relative to the Alt NAMS estimator.

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 estimators 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 $T = 240$ (60 years of quarterly data), which reflects a more typical length of TFP news series in practice.

Table 9 shows the RMSEs for the various estimators. The max share news estimator yields a

Table 9: RMSE over 40 quarters based on the larger-scale DSGE model for $T = 240$

Estimator	TFP Response	Output Response	Investment Response	Total
KS Max Share	10.4	15.4	32.0	57.7
NAMS	10.2	16.8	36.1	63.1
Max Share News	8.8	13.4	22.9	45.0
Alt KS Max Share	8.8	15.4	32.0	56.3
Alt NAMS	9.0	15.7	31.4	56.1

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 estimator. The Alt KS max share estimator uses the KS identification strategy and the max share news model variables. The Alt NAMS estimator uses the NAMS identification strategy and the max share news model variables.

22% improvement in accuracy over the KS max share estimator and a 20% improvement over the Alt KS max share estimator. Relative to the NAMS and Alt NAMS estimator, the gains are 29% and 20%, respectively. These results suggest that the benefits of the max share news estimator extend to realistic sample sizes and go beyond just aligning information sets. In other words, the superior accuracy of the max share news estimator for $T = 240$ cannot be replicated by traditional TFP max share estimators, even when providing them with the same information set.

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 VAR models that include one of three 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); and (3) **CGV**: the patent grant series used in Cascaldi-Garcia and Vukotić (2022).[24]

[24] 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.

For each series, we estimate a 9-variable VAR(4) model that includes one of the three news variables in addition to the 8 variables from the empirical VAR model used in Kurmann and Sims (2021). 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 model extends from 1960Q1 to 2010Q4. We identify the structural shocks based on the max share news estimator 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 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. We find that only the R&D and ICT models satisfy the two criteria (see Appendix H).[25]

Figure 5 shows the responses in the ICT model. Similar results hold when using the R&D model.[26] The news shock increases TFP, output, consumption, and investment in the long run, and the peak effects occur after 12 quarters. TFP initially falls, consistent with the view that new technologies are inconsistent with the incumbent technology (Baron and Schmidt, 2019). Hours respond positively, so there is positive comovement between real GDP, consumption, and investment. Inflation declines, consistent with the interpretation of the news shock as a positive

[25]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.

[26]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 5: Comparison of max share news and TFP max share identified impulse responses

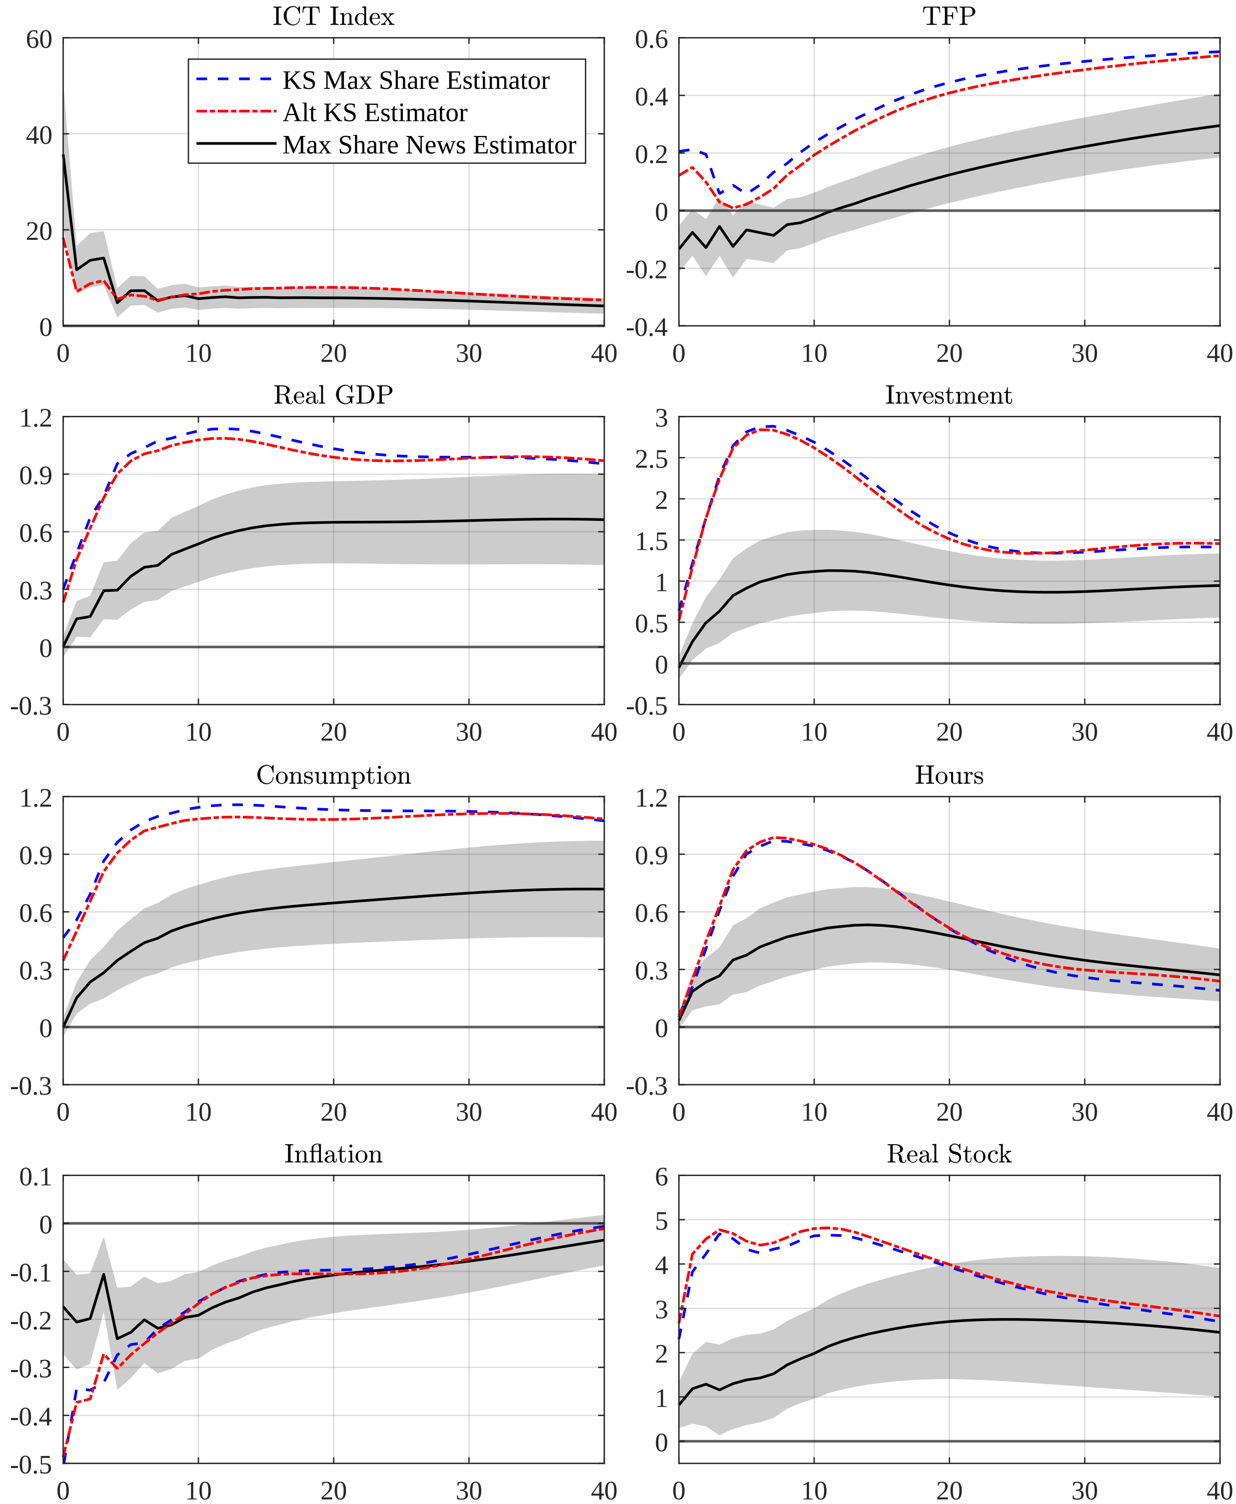

Notes: VAR(4) models estimated on identical samples from 1960-2014. Shaded regions represent 1-standard deviation error bands computed by residual-based bootstrap for the max share news estimator. Responses are in percent deviations from the baseline. The inflation response is annualized.

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 model.

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., 2024). 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., 2024).

Table 10 shows that, according to the KS max share estimator, 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

Table 10: 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 for 1960-2014.

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.

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 even in large samples 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 when news shocks explain almost all of the long-run variation in TFP. We uncovered and documented the sources of this bias.

Our evidence raises the question of how to proceed in applied work. We showed that including measures of TFP news in the VAR model and adopting a novel max share news estimator 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 reported empirical estimates of the responses to news shocks based on this alternative estimator for three 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., S. K. CHUGH, AND T. POTTER (2023): "Anticipated Productivity and the Labor Market," *Quantitative Economics*, 14, 897–934.

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.

DOU, L., P. HO, AND T. A. LUBIK (2025): "Max-Share Misidentification," Working Paper, Federal Reserve Bank of Richmond.

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

FERNALD, J. G. (2014): "A Quarterly, Utilization-Adjusted Series on Total Factor Productivity," Federal Reserve Bank of San Francisco Working Paper 2012-19.

——— (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, M. LIPPI, AND L. SALA (2017): "Noisy News in Business Cycles," *American Economic Journal: Macroeconomics*, 9, 122–152.

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.

MIRANDA-AGRIPPINO, S., S. H. HOKE, AND K. BLUWSTEIN (2024): "Patents, News, and Business Cycles," Bank of England Staff Working Paper 788.

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