Is there News in Inventories?*

Christoph Görtz
University of Birmingham

Christopher Gunn
Carleton University

Thomas Lubik
Federal Reserve Bank of Richmond

May 2019

Abstract

Inventories are an important, highly volatile and forward looking component of the business cycle, yet they have been largely neglected by the literature on TFP news shocks that argues such shocks are important drivers of macroeconomic fluctuations. We use a standard VAR identification to document a new fact: in response to TFP news, inventories move procyclically along with the other major macroeconomic aggregates. Our finding is not self-evident: conventional views would suggest news about higher future productivity provides incentives to run the current inventory stock down and increase stockholding in the future when productivity is high. We provide evidence that this substitution effect is dominated by a demand effect due to which firms increase inventories in response to sales in light of rising consumption and investment. Our empirical fact corroborates the view that TFP news shocks are important drivers of macroeconomic fluctuations. However, it imposes a challenge to existing theoretical frameworks as they fail to reproduce the procyclical inventory movements in response to TFP news shocks. We suggest this comovement puzzle can be solved through extending a standard framework with intangible capital and wage stickiness.

Keywords: News shocks, Business cycles, Inventories.
JEL Classification: E2, E3.

*We thank seminar and conference participants at the Canadian Economics Association Conference 2018 and the conference on Computing in Economics and Finance 2018. All remaining errors are our own. Görtz: University of Birmingham, Department of Economics, Email: c.g.gortz@bham.ac.uk. Gunn: Carleton University, Department of Economics, Email: chris.gunn@carleton.ca. Lubik: Research Department, Federal Reserve Bank of Richmond, P.O. Box 27622, Richmond, VA 23261, Email: thomas.lubik@rich.frb.org.

1 Introduction

Expectations about future total factor productivity (TFP) have been proposed as a potentially important source of aggregate fluctuations (Beaudry and Portier (2004), Beaudry and Portier (2014)). Substantial effort has been undertaken to understand how these, so called, 'TFP news shocks' can give rise to the empirically observed comovement of consumption, investment and hours worked in structural frameworks (e.g. Jaimovich and Rebelo (2009), Gunn and Johri (2011)) and whether these shocks play an important role once models are taken to the data (e.g. Schmitt-Grohe and Uribe (2012), Khan and Tsoukalas (2012), Görtz and Tsoukalas (2017)). Despite these important advances the literature on news shocks has largely neglected inventory investment — a margin that has long been recognized to play a large role in explaining aggregate fluctuations (e.g. Ramey and West (1999), Wen (2005)). Blinder and Maccini (1991) for example document that in a typical recession in the United States, the fall in inventory investment accounts for 87% of the decline in output; and Blinder (1981) states "to a large extent, business cycles are inventory fluctuations" (p. 500). While the literature tends to suggest news about shifts in future technology can indeed be a significant source of business cycles, to date we know very little about the relation of news shocks and movements in inventories. Does inventory investment co-move with consumption and fixed investment in response to TFP news shocks? Would this empirical finding support the importance assigned to news shocks as relevant drivers of aggregate fluctuations? Which structural frameworks can account for the empirically observed movements in inventories — do we need to rethink the existing ones? In this paper we make a step to answering these questions.

We document a TFP news shock identified from a vector autoregression (VAR) implies an increase in inventory investment along with the well documented expansion of output, consumption, investment and hours worked in a U.S. post-Great Moderation sample.[1] The expansion of inventories in response to a TFP news shock is a robust finding not only for the

[1] Our baseline identification scheme follows the approach in Francis et al. (2014). We discuss robustness to alternative identification approaches in section A.3.

whole economy, but also across the retail, wholesale and manufacturing sector as well as for finished goods, work in process and input inventories. It is a consensus in the literature that unconditionally inventory investment is procyclical (e.g. Ramey and West (1999)).[2] The consistency between the unconditional and conditional movements in inventories provides substantial support for the hypothesis that news shocks cannot be rejected as important drivers of business cycles.[3]

The documented expansion of the inventory stock in response to news about higher future TFP is not a priori self-evident. Conventional views about inventory behavior would suggest that on the one hand, such news would provide incentives to run the current inventory stock down and increase stockholdings in the future when the high productivity is realized. In addition to this negative substitution effect, one the other hand, the associated rise in sales of consumption and investment goods would create a demand effect that would lead to an incentive to increase inventories to avoid stockouts and enhance demand. To the extent that both these effects are present, our results suggest this negative substitution effect is dominated by the positive demand effect.

We investigate the transmission mechanism leading to the documented increase in inventories. Measures for the opportunity costs of holding inventories suggested by Jones and Tuzel (2013) point to the presence of a strong demand effect. In particular, we construct aggregate measures of debt and equity cost of capital and implied cost of capital measures from firm-level data. In response to a TFP news shock all measures decline significantly prior to the realization of higher TFP. This decline in the opportunity cost of holding inventories is supportive of the documented expansion in this margin. We further study the response of various measures of marginal cost to a TFP news shock. Declining marginal costs between the time the news about higher future TFP arrives and the actual realization of higher productivity is indicative for the presence of a negative substitution effect. However, once

[2]The correlation between HP-filtered GDP and inventory investment is 0.75 in our sample.

[3]Indeed, we find that the TFP news shock is important for fluctuations in key macroeconomic variables as it explains between 44-66% (43-59%) of the forecast error variance in GDP (inventories) over a horizon from 6-32 quarters.

introduced in our VAR system, none of our marginal cost measures shows such a decline in marginal costs that would point to a strong incentive to run down current inventories and build up stockholdings again once the higher productivity has been realized. Overall, we find evidence against a strong negative substitution effect, but in favor of a strong positive demand effect, which corroborates the increase in inventories we document in response to higher future TFP. Interestingly, this demand enhancing motive for holding more inventories in light of rising sales has received considerable support and is widely used in the theoretical literature following a seminal contribution by Bils and Kahn (2000).

Armed with these empirical results, we then ask whether a standard new-shock business cycle model supplemented with inventories can replicate these features of the data. We study the response to TFP news in a standard New Keynesian model that includes the trio of particular specification of preferences, investment adjustment costs and variable capital utilization.[4] The model is augmented with finished goods inventories that have a sales enhancing role as in Jung and Yun (2006), based on the stock-elastic demand model of Bils and Kahn (2000).[5] We show that our empirical evidence imposes two related challenges to this standard model. First, inventories respond countercyclically to TFP news. This holds for model versions with and without nominal rigidities. Second, the countercyclical response of inventories in turn suppresses the response of hours, and as a result dampens the response of utilization and output. This is not consistent with the narrative in the expectations driven business cycle literature of a strong boom in response to news about higher future productivity. We term these challenges the *inventory co-movement* and *path-of-hours* challenges, respectively. What is the basis of these two challenges? With respect to the first challenge, we show that the countercyclical movement in inventories results from a too-strong procyclical rise in marginal costs during the expansion in the standard model. The second challenge then follows from the first: since firms can satisfy any news-

[4]These model features are widely recognized in the news-literature as a simple means for producing comovement of consumption, investment and hours in response to a TFP news shock. As such, our model nests the frameworks of Jaimovich and Rebelo (2009) and Schmitt-Grohe and Uribe (2012).

[5]This mechansim received substantial empirical and theoretical support and is hence a widely used motive to give rise to inventory holdings, see e.g. Lubik and Teo (2012) and Jung and Yun (2013).

induced increase in sales by drawing down inventories, the demand for labor falls relative to a model without inventories, suppressing the response of hours, utilization and output relative to sales. As such, our empirical finding poses a new puzzle to the theoretical literature to develop frameworks that can account for the comovement and a strong expansion of inventories, output, consumption, investment and hours in response to TFP news shocks.

We take a first step in addressing this puzzle. We show that it is possible to generate an expansion of all macroeconomic aggregates, including inventories, with a simple variant of the standard model that assumes firms create productivity-increasing knowledge through a learning-by-doing producess. Following researchers such as Chang et al. (2002), Cooper and Johri (2002) and Gunn and Johri (2011) who have found such a mechanism helpful for allowing business cycle models to match other features of the data, we extend the standard model to include intangible capital as an additional input into production, and assume this *knowledge capital* accumulates through a learning-by-doing process involving labor. The mechanism then addresses the above challenges in an intuitive way. The arrival of news about an increase in TFP in the future raises the value of knowledge in the present, since firms can accumulate knowledge over time and enhance the impact of the rise in TFP in the future. Firms as a result increase their demand for labor prior to the arrival of TFP in order to accumulate knowledge, in the process driving up production levels and accumulating productivity-enhancing knowledge, limiting the rise in marginal costs through the boom and thereby increasing the incentive to accumulate inventories. Sticky wages are additionally helpful for limiting the initial rise in marginal costs while firms are first building up knowledge capital. We see this model as one example to resolve the comovement puzzle, but a rigorous investigation of data-generating mechanisms goes beyond the scope of this paper.

Our study is related to the large research agenda on the role of news shocks for aggregate fluctuations. The VAR methodology we employ to identify the empirical response to TFP news shocks has been widely used (e.g. Barsky and Sims (2011), Barsky and Sims (2012), Ben Zeev and Khan (2015)) and employed, amongst others, to document the comovement of macroeconomic aggregates (except inventories) over a post-Great Moderation sample (e.g.

in Görtz et al. (2017)). On the theoretical side, our paper links to a large strand of work that investigates ways of facilitating procyclical movements in consumption, investment and hours in response to TFP news shocks (e.g. Jaimovich and Rebelo (2009), Pavlov and Weder (2013)).

A large long-standing literature investigates the empirical relation of inventories with macroeconomic fluctuations and the implications of introducing inventories in theoretical frameworks to which we cannot do full justice here.[6] Bils and Kahn (2000) highlight the unconditionally limited role of intertemporal substitution for variations in inventories that is also documented in our work in the context of expectations about productivity.

To the best of our knowledge the only two papers that consider inventories in relation to TFP news shocks are the contributions by Crouzet and Oh (2016) and Vukotic (2016). Crouzet and Oh (2016) introduce inventories into existing models that had been successful in generating comovement of investment, consumption and hours in response to TFP news shocks. They provide a very valuable analysis that shows these extended models imply countercyclical movements of inventories under realistic calibrations. This evidence from theory is used to inform sign restrictions in a structural VAR to identify TFP news shocks. Given the unconditional procyclicality of inventory investment and the imposed negative sign restriction on this variable, Crouzet and Oh (2016) conclude TFP news shocks are of very limited importance for aggregate fluctuations. We approach the question on the relation between inventory movements and TFP news shocks the other way around. We use a standard and widely used VAR methodology to identify the response of inventory movements to a TFP news shock and let this empirical evidence inform our modelling choices.

Vukotic (2016) uses a TFP news shock identification similar to ours and documents the VAR responses of industries in the U.S. manufacturing sector. She finds the propagation of news shocks to be much stronger in the durables than in non-durables industries. Implications of a two-sector model can be aligned with this finding once inventories are introduced as factor of production in the durables sector where they play a buffer stock role similar to fixed

[6] Surveys are e.g. Blinder and Maccini (1991) and Ramey and West (1999).

capital investment. While she does not explicitly discuss any empirical responses of inventories, she shows the inventory to sales ratio in durables sectors moves particularly strongly countercyclical in response to TFP news shocks, which is consistent with our findings for the whole economy.

The remainder of the paper is structured as follows. Sections 2.1 and 2.2 discuss the VAR identification strategy and the data used in the empirical analysis. Section 2.3 presents our main empirical findings. We corroborate these in sections 2.4. Sections 3 presents a business cycle model with inventories model to study the response of inventory. Section 4 concludes.

2 VAR Analysis

2.1 The VAR model and news shock identification

Consider the following reduced form VAR(p) model,

$$y_t = A(L)u_t, \tag{1}$$

where y_t is an $n \times 1$ vector of variables of interest, $A(L) = I + A_1 L + A_2 L^2 + ... + A_p L^p$ is a lag polynomial, $A_1, A_2, ..., A_p$ are $n \times n$ matrices of coefficients and, finally, u_t is an error term with $n \times n$ covariance matrix Σ. Define a linear mapping between reduced form, u_t, and structural errors, ε_t,

$$u_t = B_0 \varepsilon_t, \tag{2}$$

We can then write the structural moving average representation as

$$y_t = C(L)\varepsilon_t, \tag{3}$$

where $C(L) = A(L)B_0$, $\varepsilon_t = B_0^{-1}u_t$, and the matrix B_0 satisfies $B_0 B_0^{'} = \Sigma$. The B_0 matrix may also be written as $B_0 = \tilde{B}_0 D$, where $\tilde{B}_0$ is any arbitrary orthogonalization of Σ and D is an orthonormal matrix ($DD' = I$).

The h step ahead forecast error is,

$$y_{t+h} - E_{t-1}y_{t+h} = \sum_{\tau=0}^{h} A_\tau \tilde{B}_0 D \varepsilon_{t+h-\tau}. \tag{4}$$

The share of the forecast error variance of variable i attributable to shock j at horizon h is then

$$V_{i,j}(h) = \frac{e_i' \Big(\sum_{\tau=0}^{h} A_\tau \tilde{B}_0 D e_j e_j' D' \tilde{B}_0' A_\tau' \Big) e_i}{e_i' \Big(\sum_{\tau=0}^{h} A_\tau \Sigma A_\tau' \Big) e_i} = \frac{\sum_{\tau=0}^{h} A_{i,\tau} \tilde{B}_0 \gamma \gamma' \tilde{B}_0' A_{i,\tau}'}{\sum_{\tau=0}^{h} A_{i,\tau} \Sigma A_{i,\tau}'}, \tag{5}$$

where e_i denotes selection vectors with one in the i-th position and zeros elsewhere. The e_j vectors pick out the j-th column of D, denoted by γ. $\tilde{B}_0\gamma$ is an $n \times 1$ vector corresponding to the j-th column of a possible orthogonalization and can be interpreted as an impulse response vector.

In the following, we discuss the methodology that identifies the TFP news shock from the VAR model. This so called Max Share methodology is consistent with Uhlig (2003) and based on Francis et al. (2014) who isolate unanticipated productivity shocks by maximizing the forecast error variance share of TFP at a long but finite horizon. At a long horizon h, all variations in TFP are either accounted for by anticipated or unanticipated shocks to this variable. Then we can write

$$V_{1,1}(h) + V_{1,2}(h) = 1, \tag{6}$$

where we assume TFP is ordered first in the VAR system and the unanticipated shock is indexed by 1 and the anticipated (news) shock by 2. The unanticipated shock is identified as the innovations to observed TFP and will be independent of the identification of the other $n-1$ structural shocks. Given the index for the unanticipated shock, the share of variance in TFP attributable to this shock at horizon h is summarized in $V_{1,1}(h)$. Following Barsky and Sims (2011) and Francis et al. (2014), choosing the elements of $\tilde{B}_0$ to make equation (6) hold as closely as possible is equivalent to choosing the impact matrix so that contributions to $V_{1,2}(h)$ are maximized. Hence, we choose the second column of the impact matrix to solve the following optimization problem:[7]

$$argmax\ V_{1,2}(h) = \frac{\sum_{\tau=0}^{h} A_{i,\tau}\tilde{B}_0\gamma\gamma'\tilde{B}_0' A_{i,\tau}'}{\sum_{\tau=0}^{h} A_{i,\tau}\Sigma A_{i,\tau}'},$$

$$\text{s.t. } \gamma'\gamma = 1, \quad \gamma(1,1) = 0, \quad \tilde{B}_0(1,j) = 0 \quad \forall j > 1.$$

In the above, we restrict γ to have unit length which ensures it is a column vector belonging to an orthonormal matrix. The second and third constraints impose that a news shock about TFP cannot affect TFP contemporaneously. To summarize, we identify the TFP news shock from the VAR model as the shock that (i) does not move TFP on impact and (ii) maximizes the share of variance explained in TFP at a long but finite horizon h.

2.2 Data and VAR estimation

We estimate the VAR using quarterly U.S. data for the period 1983:Q1–2018:Q2. This sample horizon is guided by the literature that documents differences in cross correlation patterns of several macro-aggregates in samples before and after the mid-1980s (e.g. Galí and Gambetti (2009)). Furthermore, McCarthy and Zakrajsek (2007) document significant changes in inventory dynamics occur in the mid-1980s due to improvements in inventory management. Several time series that we use in the following — e.g. total business inventories and its sectoral components — are also only available over (part of) a post-Great Moderation sample. However, data availability permitting, we show below for robustness that our main results hold also for a longer sample. To identify TFP news shock from the VAR model, we adopt the Max Share identification method outlined in section 2.1. Based on Francis et al. (2014) we set the horizon h to 40 quarters. The time series included in the VAR enter in levels, consistent with the treatment in the empirical VAR literature (see e.g. Barsky and Sims (2011), Beaudry and Portier (2004) and Beaudry and Portier (2014)). To estimate the VAR model we use three lags and a Minnesota prior. Confidence bands are computed by drawing from the posterior.

[7] The optimization problem is formulated in terms of choosing γ conditional on any arbitrary orthogonalization, $\tilde{B}_0$, to ensure the resulting identification belongs to the space of possible orthogonalizations of the reduced form.

We consider two different measures for total inventories in the VAR. First, non-farm private inventories, which are defined as the physical volume of inventories owned by private non-farm business, valued at average prices of the period (the replacement costs of inventories). The second measure, business inventories, differs from the former as stockholdings are valued by the cost at acquisition of inventories that can differ from their price when sold. In the NIPAs, inventory profits and losses resulting from differences between acquisition and sales price are shown as adjustments to business income. Business inventories are only available from 1992Q1 which is why we reduce the sample horizon if these are included in the VAR. Output is defined as GDP and total hours as hours worked of all persons in the non-farm business sector. Investment is the sum of fixed investment and personal consumption expenditures for durable goods. Fixed investment is the component of gross private domestic investment that excludes changes in private inventories. Consumption is the sum of personal consumption expenditures for non-durable goods and services. All these time series are seasonally adjusted and in real per-capita terms (except for hours which are not deflated). Inflation is constructed using the GDP deflator. A measure of technology is key to identify the news shock. We follow the convention in the empirical literature and use the measure of utilization-adjusted TFP provided by Fernald (2014).[8] We also use the Michigan consumer confidence indicator (E5Y) in our VAR system.[9] The set of variables included in our VAR system is, apart from inventories, standard in the literature and considering the E5Y consumer confidence measure is a way to provide forward looking information that captures expectations.[10]

[8] We use the 2018 vintage which contains updated corrections on utilization from industry data.

[9] The Michigan consumer confidence indicator summarizes responses to the following question: "Turning to economic conditions in the country as a whole, do you expect that over the next five years we will have mostly good times, or periods of widespread unemployment and depression, or what?" The variable is constructed as the percentage giving a favorable answer minus the percentage giving an unfavorable answer plus 100.

[10] See e.g. Barsky and Sims (2012). The S&P500 stock price index has also been considered for this purpose. Our results are robust to including the S&P500 instead of the E5Y.

2.3 VAR results

Figure 1 shows impulse response functions to a TFP news shock from an eight-variable VAR. It is striking that all activity variables are increasing prior to a significant rise in TFP. The comovement between output, consumption, investment and hours over this post Great Moderation sample has been documented in existing work (e.g. Görtz et al. (2017)). The fact we add here is to document the increase in the stock of private non-farm inventories prior to a rise in TFP. The hump-shaped increase in the stock of inventories indicates that inventory investment is positive until about quarter 12, shortly before the higher productivity is actually realized. Additionally, we report a short-lived decline in inflation and an anticipation of the increase in TFP in the consumer confidence indicator E5Y, both consistent with findings in previous work. Barsky and Sims (2012) highlight that the inflation response is broadly consistent with the New Keynesian framework in which current inflation equals an expected present discounted value of future marginal costs. The significant increase in the E5Y is indicative of an increase in consumer confidence upon the arrival of news about higher TFP (see e.g. Barsky and Sims (2011) or Görtz and Tsoukalas (2018)). We show in Appendix A.2 that the described IRF patterns are qualitatively, and to a large extent also quantitatively, very similar when the VAR is estimated over samples starting in 1948Q1 or 1960Q1. The TFP news shock is important for fluctuations in inventories and GDP as we find it to explain between 47-65% (47-71%) of the forecast error variance in inventories (GDP) over a horizon from 6-32 quarters.[11]

Figure 2 shows that the rise in inventories prior to TFP is also robust when we use total business inventories as an alternative measure to private non-farm inventories. Evaluating the response of inventories to a TFP news shock also with this alternative measure is important as it is not a priori clear at which prices inventories should be measured. However, this measure is only available from 1992Q1 which restricts the sample for this VAR system.[12] All variables in Figure 2 show very similar qualitative responses to the ones in Figure 1, albeit

[11]Details about the forecast error variance decomposition are provided in Appendix A.1.

[12]Note that data availability limits all VAR systems that include total business inventories or its subcomponents to start in 1992Q1.

the shorter sample results in somewhat wider confidence bands. Overall, this figure confirms the comovement of macroeconomic aggregates, including inventories, prior to the significant rise in TFP.

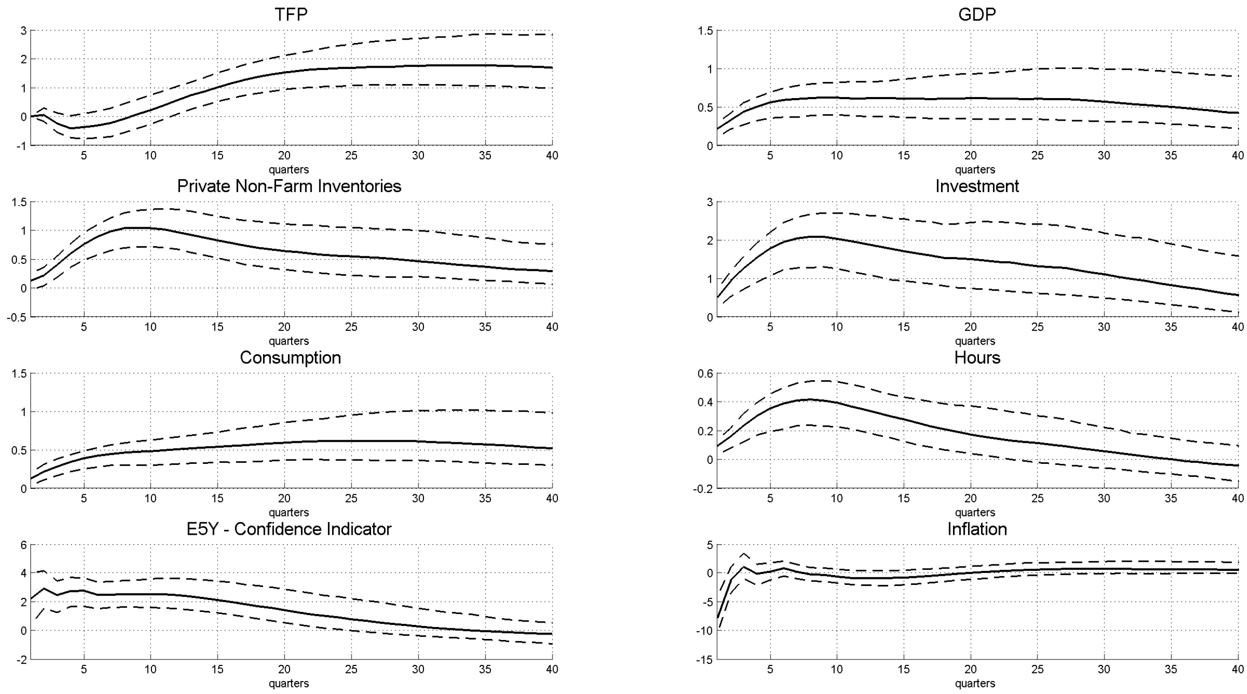

Figure 1: **IRF to TFP news shock – including Private Non-Farm Inventories.** Sample 1983Q1-2018Q2. The solid line is the median and the dashed lines are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

The vast majority of inventories are held in the manufacturing, wholesale and retail sectors (see e.g. Ramey and West (1999)). Figure 3 shows the responses of business inventories in each of these sectors to the (aggregate) TFP news shock when we alternate between including one of the three separate sectoral measures of inventory in our eight-variable VAR. It is evident that the expansion of the inventory stock is broad-based across the manufacturing, wholesale and retail sector. The two trade sectors almost entirely hold finished goods inventories (see e.g. Blinder and Maccini (1991)), while over our 1992Q2-2018Q2 sample the inventory stock held in the manufacturing sector is split across finished goods inventories (36%), work in process (30%) and input inventories in form of materials and supplies

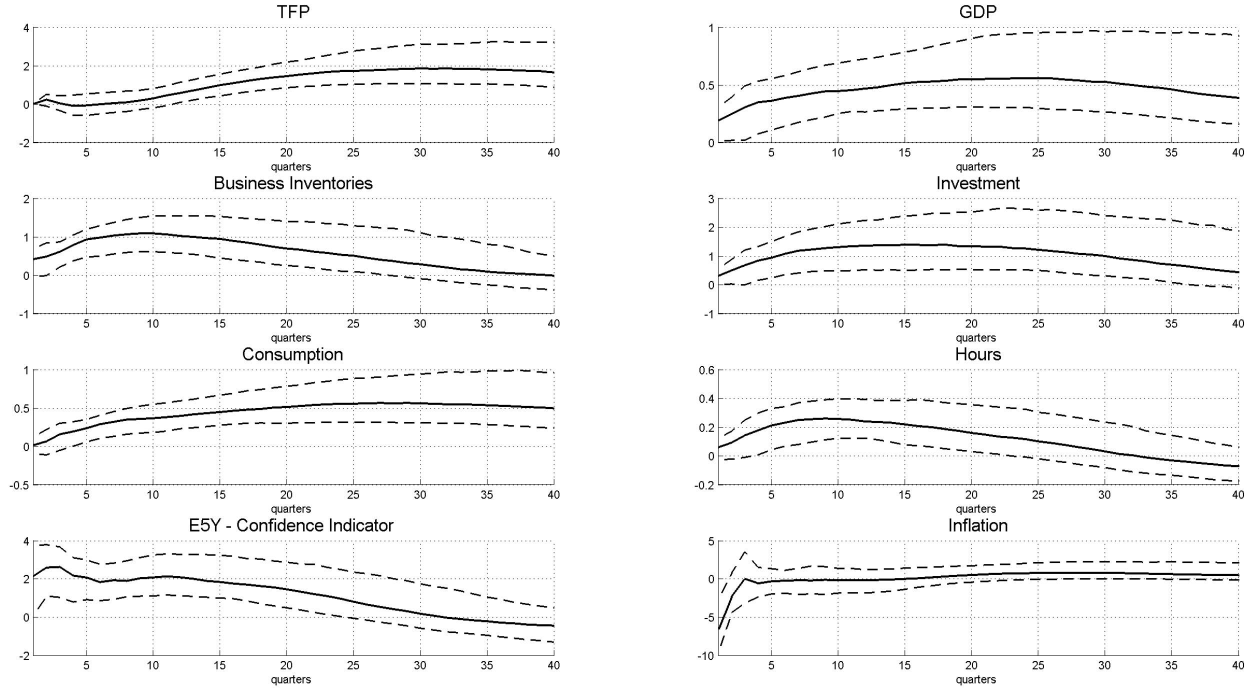

Figure 2: **IRF to TFP news shock – including Business Inventories.** Sample 1992Q1-2018Q2. The solid line is the median and the dashed lines are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

(34%).[13] Figure 4 shows the responses of inventory types in the manufacturing sector when we include these one-by-one in our eight-variable VAR. In response to a TFP news shock, finished goods and input inventories in the manufacturing sector rise strongly before TFP increases significantly after about 12 quarters.[14] The strong positive response of aggregate inventories, sectoral measures and different types of inventories, prior to the realization of anticipated higher productivity, is consistent with results in Kesavan et al. (2010) who find inventories to be a forward looking-variable closely linked to future expectations of economic conditions as they help to improve forecasts about sales. The new fact we document about the response of inventories to TFP news shocks is broad based across different aggregate measures, sectors and types of inventories.

[13]For the wholesale and retail sector, time series data that break the total stock down into inventory types is not available.

[14]The responses of the remaining seven variables in the VAR that we are not showing in Figures 3 and 4 are very similar to the ones reported in Figure 2 and are available upon request. We focus our discussion on sectoral data for business inventories as private non-farm inventories in the manufacturing, wholesale and retail sectors is available only from 1996Q4.

Figure 3: **IRF of business inventories by sector to TFP news shock.** Sample 1992Q1-2018Q2. Subplots result from eight variable VARs comprising TFP, GDP, consumption, investment, hours, inventory measure, inflation, E5Y. The inventory measures were included one-by-one in the VAR system. The solid line is the median and the dashed lines are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

Figure 4: **IRF of business inventories in the manufacturing sector by inventory type to TFP news shock.** Sample 1992Q1-2018Q2. Subplots result from eight variable VARs comprising TFP, GDP, consumption, investment, hours, inventory measure, inflation, E5Y. The inventory measures were included one-by-one in the VAR system. The solid line is the median and the dashed lines are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

2.4 Corroborative evidence and guidance for models

This section briefly discusses additional evidence that corroborates the results of the sections above and provides further guidance for modelling inventory behavior in response to news shocks.

Jaimovich and Rebelo (2009) state that for models to facilitate comovement of consumption and investment in response to news about future higher TFP a strong increase in utilization and/or hours worked is required. The VAR results in Figure 5 document a strong increase in capital utilization and the real wage in response to a TFP news shock. The positive hump-shaped response of the real wage is consistent with the increase in hours documented in Figure 1.

Figure 5 further shows that at long horizons the inventory to sales ratio moves countercyclically in response to a TFP news shock. This is another indication that TFP news shocks are potentially important drivers of aggregate fluctuations as the unconditional countercyclicality of the inventory to sales ratio is a commonly accepted view in the literature (e.g. Blinder (1981)). In the sections below we will show that in our model this countercyclicality is a necessary condition for comovement of inventories with the other macroeconomic aggregates in a setup with flexible prices. The literature on inventories often does not only consider the level of inventories but also the change in this variable which provides an indication about inventory investment (abstracting from depreciation). The fourth subplot in Figure 5 shows a positive response of inventory investment in light of a TFP news shock. It peaks at about five quarters before it declines towards zero. This pattern is broadly consistent with the response of the level of inventories documented in Figure 1.

The last subplot in Figure 5 shows that issued patents rise in response to a TFP news shock. This provides corroborative evidence for an important mechanism in the structural model discussed in the next section: the accumulation of knowledge capital in light of favorable news about future productivity.[15]

[15] The samples used for the subplots in figure 5 vary slightly due to data availability. The VARs including capital utilization, the inventory to sales ratio or the change in inventories, respectively are estimated over our baseline 1983Q1-2018Q2 sample. The sample for the VAR including the real wage is 1983Q1-2018Q1.

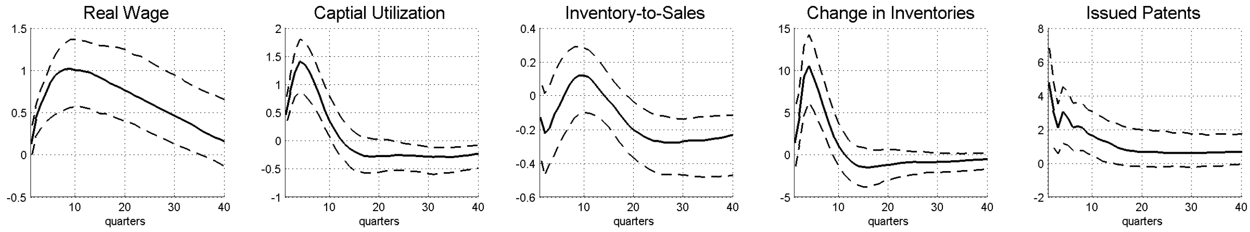

Figure 5: **IRF to TFP news shock.** Subplots result from eight-variable VARs comprising TFP, GDP, investment, consumption, hours, inflation, E5Y and one of the plotted variables above at a time. The solid line is the median and the dashed lines are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

3 A business cycle model with inventories

We now present a business cycle model supplemented with inventories and study the behaviour of inventories and other important macroeconomic variables in response to TFP news. The backbone of the model is based on the flexible price business cycle model of Schmitt-Grohe and Uribe (2012) that includes the particular specification of preferences, investment adjustment costs and costly capacity utilization following Jaimovich and Rebelo (2009) as the workhorse means within the literature for producing comovement of consumption, investment and hours-worked in response to news about TFP. We augment this core with inventories introduced as in Jung and Yun (2006) and Lubik and Teo (2012) whereby finished good inventories are "sales enhancing", based on the stock-elastic demand model of Bils and Kahn (2000). Finally, we add intangible capital, which we refer to as *knowledge capital*, as an additional input into production. Following Chang et al. (2002) and Cooper and Johri (2002), we assume knowledge capital evolves over time as a learning-by-doing process, whereby the household acquires new technological knowledge through its experiences in supplying labor to the production process. As we will see, such a supply-side process helps

The sample for the VAR including the patent data is 1983Q1-2014Q4. Responses of the other variables in the eight-variable VAR system are virtually identical to the ones in Figure 1 and are available upon request. The inventory to sales ratio is the ratio of private non-farm inventories and final sales of domestic business as in Lubik and Teo (2012). Utilization is provided by Fernald (2014) and consistent with our utilization adjusted measure for TFP. The real wage is compensation of employees, nonfinancial corporate business, in real per capita terms. The change in inventories is the change in private non-farm inventories. The data on issued patents is from the United States Patent and Trademark Office (USTPO).

suppress the rise in marginal costs during the demand-like expansion phase of the news boom in order to allow inventories to rise along with the other macroeconomic variables. Exploiting this property of learning-by-doing for inventories is novel within the literature, but the general feature of learning-by-doing as a supply-side side mechanism that enhances the dynamics of business cycles models is not new. While learning-by-doing as a modeling mechanism has had a long history in studying long-run issues such as growth, e.g. in Arrow (1962), work such as Chang et al. (2002), Cooper and Johri (2002) and Gunn and Johri (2011) examines the mechanism in terms of its propagation characteristics in response to various business cycle shocks (including TFP news shocks) at business cycle frequencies. This particular extension also has a distinct advantage in terms of its parsimony: the modification adds only an additional input into production and an accumulation equation for the learning, leaving the other elements of the model unaffected. Moreover, it nests the more standard model without intangible capital.

3.1 Model environment

We adopt a particular decentralization convenient for modeling finished goods inventories, separating the production side of the economy into distinct production, distribution and final goods aggregation phases. The model economy consists of a representative infinitely-lived household, a competitive intermediate goods-producing firm, a continuum of monopolistically competitive distributors indexed by $i \in [0, 1]$, and a competitive final goods producer. The intermediate goods firms owns its capital stock and produces a homogeneous good which it sells to distributors, who then differentiate the good into distributor-specific varieties which they sell to the final goods firm who aggregates the varieties into a final good. The final good may be used for consumption or investment. Following Chang et al. (2002), we assume that the household accumulates the knowledge capital, selling its *effective labor* to firms in the form of the product of knowledge capital and hours-worked.

3.1.1 Household and Government

The household's lifetime utility is defined over sequences of consumption C_t and hours-worked n_t given by

$$E_0 \sum_{t=0}^{\infty} \beta^t \Gamma_t \left\{ \frac{(C_t - bC_{t-1} - \psi n_t^{\xi} F_t)^{1-\sigma} - 1}{1-\sigma} \right\}, \tag{7}$$

where

$$F_t = (C_t - bC_{t-1})^{\gamma_f} F_{t-1}^{1-\gamma_f}, \tag{8}$$

and where Γ_t is a stationary exogenous stochastic *preference* shock process, $0 < \beta < 1$, $0 \leq b < 1$, $\psi > 0$, $\xi > 1$, $\sigma > 0$ and $0 < \gamma_f \leq 1$. See Schmitt-Grohe and Uribe (2012) for a discussion of these preferences, which are based on the those of Jaimovich and Rebelo (2009), with the addition of consumption habits. In general, as γ_f tends towards 0, the preferences approach the "no-income effect" preferences of Greenwood et al. (1988), yet still remain consistent with the balanced-growth path in a growing economy.

Letting knowledge capital h_t represent the household's state of technological knowledge (or skill level) based on past labor supplies, the household's knowledge impacts the effective units of labor supplied to firms, such that it earns wage w_t in exchange for effective labor $\tilde{n}_t = h_t n_t$. Following Chang et al. (2002), the household accumulates knowledge capital according to

$$h_{t+1} = h_t^{\gamma_h} n_t^{\nu_h}, \tag{9}$$

where $0 \leq \gamma_h < 1$, $\nu_h > 0$, such that it gains knowledge as it engages with the production process through supplying labor.[16] Note that equation (9) implies this form of knowledge capital is stationary even in a non-stationary economy, due to the stationarity of hours-worked, implying that the long-run growth path of output is determined by exogenous technological

[16] The log-linear specification of (9) used by Chang et al. (2002) is common in the literature, and similar to that of Cooper and Johri (2002). Other specifications that have been explored in the literature are also possible. See the discussion in Cooper and Johri (2002), as well as applications in McGrattan and Prescott (2010), Gunn (2015) and Hou and Johri (2018). One advantage of the specification in (9) is its analytical simplicity.

factors only. In this sense, this form of knowledge capital can be thought of a type of index which conditions on the effect of hours in production over the business cycle, as the household responds to fluctuations in the exogenous stochastic drivers of growth (such as news about permanent increases in TFP).

The household owns the stock of physical capital K_t, renting capital services $\tilde{K}_t = u_t K_t$ to the intermediate goods producers each period for rental rate r_t, where u_t is the utilization rate of the capital. The capital stock evolves according to

$$K_{t+1} = (1 - \delta(u_t))K_t + m_t I_t \left[1 - S\left(\frac{I_t}{I_{t-1}}\right)\right], \tag{10}$$

where $\delta(\cdot)$ is a depreciation function satisfying $\delta'(\cdot) > 0$, $\delta''(\cdot) > 0$ and $\delta(1) = \delta_k > 0$, and m_t is a stationary exogenous stochastic *marginal efficiency of investment* shock process. $S(\cdot)$ is an investment adjustment cost function as in Christiano et al. (2005) with $S(g^I) = S'(g^I) = 0$ and $S''(g^I) = s'' > 0$, where g^I is the steady state growth rate of investment.

The household's period t budget constraint is given by

$$C_t + \Upsilon_t I_t + T_t = w_t \tilde{n}_t + r_t u_t K_t + \Pi_t, \tag{11}$$

where Υ_t is non-stationary exogenous stochastic *investment-specific productivity* shock, T_t is lump sum taxes, and Π_t is collective profits flowing from firms. We assume that the *growth rate* of Υ_t, given by

$$g_t^{\Upsilon} = \frac{\Upsilon_t}{\Upsilon_{t-1}}, \tag{12}$$

follows a stationary stochastic process.

Government spending G_t is financed each period by the lump-sum taxes, such that $G_t = T_t$. We assume that government spending follows a process given by

$$G_t = \left(1 - \frac{1}{\varepsilon_t}\right) Y_t, \tag{13}$$

where ε_t is a stationary exogenous stochastic *government spending* shock.

The household chooses C_t, I_t, n_t, u_t and K_{t+1} to maximize equation (7) subject to (8), (9), (10) and (11). The household's first-order conditions are standard with the exception of the impact of knowledge capital. Highlighting just those first-order conditions with direct impacts of knowledge capital, the household's n_t and h_{t+1} first-order conditions are given by

$$\xi\psi\Gamma_t V_t^{-\sigma} n_t^{\xi-1} F_t = \lambda_t w_t h_t + \mu_t^h \nu_h \frac{h_{t+1}}{n_t}, \tag{14}$$

$$\mu_t^h = \beta E_t \left\{ \lambda_{t+1} w_{t+1} n_{t+1} + \mu_{t+1}^h \gamma_h \frac{h_{t+2}}{h_{t+1}} \right\}. \tag{15}$$

Note in equation (14) that the presence of knowledge capital adds an additional term into the household's hours-worked first order condition that drives a wedge in between the marginal utility of leisure and the marginal contribution of hours to earnings, and which serves to shift the household's labor supply. All else equal, a rise in the value of knowledge capital, μ_t^h, increases labor supply as the household attempts to increase its knowledge capital by engaging with production. Equation (15) then describes μ_t^h as a function of the expected discounted value of the marginal contribution of knowledge in wage earnings next period and the continuation value of that knowledge capital.

3.1.2 Intermediate Goods Firm

The competitive intermediate goods firm produces the homogeneous good Y_t according to the technology

$$Y_t = z_t \left(\Omega_t \tilde{n}_t\right)^{\alpha} \tilde{K}_t^{1-\alpha}, \tag{16}$$

where z_t is a stationary exogenous stochastic *productivity* shock process and Ω_t is a non-stationary exogenous stochastic *productivity* shock process. We assume that the *growth rate* of Ω_t, given by

$$g_t^{\Omega} = \frac{\Omega_t}{\Omega_{t-1}}, \tag{17}$$

which follows a stationary stochastic process.

Each period, the firm acquires effective labor, $\tilde{n}_t$, at wage w_t from the labor market, and capital services, $\tilde{K}_t$, at rental rate r_t from the capital services market, and then sells its output Y_t at real price τ_t to the distributors. The firm's problem involves choosing $\tilde{n}_t$ and $\tilde{K}_t$ to maximize profit $\Pi_t^y = \tau_t Y_t - w_t \tilde{n}_t - r_t \tilde{K}_t$.

Additionally, define the marginal cost of production, mc_t, for the intermediate goods firm as $\mathrm{mc}_t = \frac{w_t}{MP\tilde{N}_t}$, where $MP\tilde{N}_t$ is the marginal product of effective labor. From the intermediate goods firm's labor first-order condition, $w_t = \tau_t \alpha \frac{Y_t}{\tilde{n}_t}$, it then follows that the output price τ_t is equal to the the marginal cost of production mc_t.

3.1.3 Final goods firm

The competitive final goods firm produces goods for sale, S_t, by combining goods varieties S_{it}, $i \in [0, 1]$, according to the technology

$$S_t = \left[\int_0^1 \nu_{it}^{\frac{1}{\theta}} S_{it}^{\frac{\theta-1}{\theta}} di \right]^{\frac{\theta}{\theta-1}}, \quad \theta > 1, \tag{18}$$

where ν_{it} is a 'taste shifter' depending on the stock of goods available for sale a_{it} (taken as given by the final goods producer), defined as

$$\nu_{it} = \left(\frac{A_{it}}{A_t} \right)^{\zeta}, \quad \zeta > 0, \tag{19}$$

and where A_t is the economy-wide average stock of goods for sale, given by $A_t = \int_0^1 A_{it} di$. The parameters θ and ζ capture the elasticity of substitution between differentiated goods and the elasticity of demand with respect to the relative stock of goods, respectively.

The firm acquires each ith goods variety from the distributors at relative price p_{it} in terms of the final good, and sells the final good where it may be used for consumption or as an input into the production of investment goods. The firm maximizes the profit function $\Pi_t^s = p_{it} S_t - \int_0^1 p_{it} S_{it} di$ by choosing S_{it} $\forall i$, yielding a demand function for S_{it} for the ith variety,

$$S_{it} = \nu_{it} p_{it}^{-\theta} S_t. \tag{20}$$

3.1.4 Distributors

Distributors acquire the homogenous good Y_t from the intermediate goods firm at real price τ_t, and then differentiate it into goods-variety Y_{it} at zero cost, with a transformation rate of one unit of the homogeneous good to one unit of the differentiated good. Goods available for sale are the sum of the differentiated output and depreciated previous period's inventories,

$$A_{it} = (1 - \delta_x) X_{it-1} + Y_{it}, \tag{21}$$

where inventories, X_{it}, are the stock of goods remaining at the end of the period, given by

$$X_{it} = A_{it} - S_{it}, \tag{22}$$

and δ_x is the period depreciation of the inventory stock. The distributors have market power over the sales of their differentiated varieties, and thus the ith distributor sets the price p_{it} for sales S_{it} of its variety, subject for to the demand curve for that variety. Each period, the ith distributor then faces the problem of choosing p_{it}, S_{it}, Y_{it} and A_{it} to maximize

$$E_t \sum_{k=0}^{\infty} \beta^k \frac{\lambda_{t+k}}{\lambda_t} \left\{ p_{it+k} S_{it+k} - \tau_t Y_{t+k}(j) \right\}, \tag{23}$$

subject to the demand curve (20), the stock and inventory expression (21) and (22). Substituting in the demand curve (20) for S_{it}, and letting μ_t^a and μ_t^x be the multipliers on (21) and (22) respectively, the distributor's first-order conditions are given by

$$\tau_t = \mu_t^a \tag{24}$$

$$\mu_t^x = \beta E_t \frac{\lambda_{t+1}}{\lambda_t} \mu_{t+1}^a (1 - \delta_x) \tag{25}$$

$$\mu_t^a = p_{it} \zeta \frac{S_{it}}{A_{it}} + \mu_t^x \left[1 - \zeta \frac{S_{it}}{A_{it}}\right] \tag{26}$$

$$\frac{P_{it}}{P_t} = \frac{\theta}{\theta - 1} \mu_t^x, \tag{27}$$

where (24), (25), (26) and (27) describe the optimal choices of Y_{it}, X_{it}, A_{it} and $p_{it} = P_{it}/P_t$.

Note from equation (21) that with beginning of period inventories predetermined, a distributor can only further increase its stock of available goods for sale, A_{it}, in period t by acquiring additional output, Y_{it}, purchased at real price τ_t. Thus the cost of generating an additional unit of goods for sale is equal to the price of output (or marginal cost of output) τ_t, which from the intermediate goods firm's problem is also the marginal cost of production mc_t. At the optimum, equation (24) says that the cost of an additional unit of goods for sale τ_t is equal to the value of those goods for sale, μ_t^a,

Next, from the inventory definition (22), for a given level of goods available for sales, A_{it}, any increase in sales, S_{it}, results in a reduction in inventory. Thus, the opportunity cost of sales for the distributor is equal to the value of foregone inventory, μ_t^x, which we can then interpret as the marginal cost of sales. The first-order condition (25) then says that the value of an additional unit of inventory today, μ_t^x, is the expected discounted value of the extra level of goods available for sale next period generated by that inventory, μ_{t+1}^a, whose value is in turn equal to the price of output next period, τ_{t+1} from (24). Thus, in a model with inventory, the marginal cost of sales is equal to the expected discounted value of next period's marginal cost of output, since increasing sales by drawing down inventories to forgo production today means that eventually the distributor will need to increase production in the future.

The first-order condition (26) says that the marginal value of extra goods for sale μ_t^a

consists of the value of the extra sales generated by the additional goods available, A_{it}, plus the value of the additional inventory yield from the unsold portion of the additional goods available, A_{it}. Combining (24), (25) and (26) yields

$$\tau_t = \zeta \frac{S_{it}}{A_{it}} + \beta E_t \frac{\lambda_{t+1}}{\lambda_t} \tau_{t+1}(1-\delta_x)\Big[1 - \zeta \frac{S_{it}}{A_{it}}\Big], \tag{28}$$

showing that the distributor chooses A_{it} such that this benefit is equal to the marginal cost of output τ_t. We will refer to (28) as the distributor's optimal stocking condition.

Finally, the first-order condition (27) says that the distributor sets its relative price as a constant markup over the marginal cost of sales. In standard flexible price models with imperfect competition but *without* inventories, the marginal cost of sales is equal to the marginal cost of output, and thus (27) is the same as in the standard model. Yet the presence of inventories drives a wedge between the marginal cost of output and marginal cost of sales such that there is no longer a constant markup between relative price and the marginal cost of *output*. Thus overall we can think of there being two additive markups: the markup between marginal cost of output and the marginal cost of sales, and the markup between the marginal cost of sales and the price. The optimal stocking condition (26) describes the adjustment of the first markup through inventories; the optimal pricing condition (27) describes the adjustment of the second markup through price-setting. Under flexible prices, the latter markup is constant, but the former is not. Thus the total markup between marginal cost of output and price will vary dynamics as the distributors uses inventories to adjust the markup between marginal cost of production and the marginal cost of sales.

3.1.5 Exogenous Stochastic Processes

There are six stochastic processes in the model: Γ_t (preference), m_t (MEI), g_t^{Υ} (growth rate of permanent investment-specific productivity), ε_t (government spending), z_t (stationary productivity), and g_t^{Ω} (growth rate of non-stationary productivity). All the stochastic processes Ξ_t, where $\Xi = \Gamma_t, m_t, g_t^{\Upsilon}, \varepsilon_t, z_t, g_t^{\Omega}$, evolve according to the stationary process

$$\ln(\Xi_t/\Xi) = \rho_{\Xi} \ln(\Xi_{t-1}/\Xi) + u_{\Xi t}, \tag{29}$$

where $\rho_{\Xi} < 1$ and $u_{\Xi t}$ is the shock innovation. We potentially allow for news shocks to all stochastic processes with the exception of the preference shocks. For the processes with news shocks, the innovation $u_{\Xi t}$ contains both anticipated and unanticipated components, and we assume the news signals arrive with horizons of $4, 8$ and 12 periods. For the preference shock stochastic process, the innovation $u_{\Xi t}$ contains only an unanticipated component. As such, the innovations are given by

$$u_{\Xi t} = \begin{cases} \epsilon^0_{\Xi t} + \epsilon^4_{\Xi t-4} + \epsilon^8_{\Xi t-8} + \epsilon^{12}_{\Xi t-12}, & \Xi = \{m_t, g_t^{\Upsilon}, \varepsilon_t, z_t, g_t^{\Omega}\} \\ \epsilon^0_{\Xi t}, & \Xi = \Gamma_t, \end{cases} \tag{30}$$

where $\epsilon^0_{\Xi t}$ is a surprise shock, and for $p = 4, 8, 12$, $\epsilon^p_{\Xi t-p}$ is a news shock that agents receive in period $t - p$ about the innovation in t. All shocks are mean zero and uncorrelated over time and with each other. The news and surprise shocks have standard deviation σ^p_{Ξ} and σ^0_{Ξ} respectively.

3.1.6 Equilibrium

Defining $V_t = C_t - bC_{t-1} - \psi n_t^{\xi} F_t$ as the periodic utility function argument to ease notation, and letting μ_t^f, μ_t^h, μ_t^k and λ_t be the multipliers on (8), (9), (10) and (11) respectively, we define a symmetric competitive equilibrium as a set of stochastic processes $\{C_t, I_t, G_t, S_t, Y_t, n_t, u_t, F_t, h_t, K_t, X_t, A_t, w_t, r_t, \tau_t, \mu_t^f, \mu_t^h, \mu_t^k, \lambda_t\}_t^{\infty}$ satisfying

$$C_t + \Gamma_t I_t + G_t = S_t, \tag{31}$$

$$K_{t+1} = (1 - \delta(u_t))K_t + m_t I_t \left[1 - S\left(\frac{I_t}{I_{t-1}}\right)\right], \tag{32}$$

$$h_{t+1} = h_t^{\gamma_h} n_t^{\nu_h}, \tag{33}$$

$$F_t = (C_t - bC_{t-1})^{\gamma_f} F_{t-1}^{1-\gamma_f}, \tag{34}$$

$$G_t = \left(1 - \frac{1}{\varepsilon_t}\right) Y_t, \tag{35}$$

$$\Gamma_t V_t^{\sigma} + \mu_t^f \gamma_f \frac{F_t}{C_t - bC_{t-1}} - b\beta E_t \left\{ \Gamma_{t+1} V_{t+1}^{\sigma} + \mu_{t+1}^f \gamma_f \frac{F_{t+1}}{C_{t+1} - bC_t} \right\} = \lambda_t, \tag{36}$$

$$\xi \psi \Gamma_t V_t^{-\sigma} n_t^{\xi-1} F_t = \lambda_t w_t h_t + \mu_t^h \nu_h \frac{h_{t+1}}{n_t} \tag{37}$$

$$r_t = \frac{\mu_t^k}{\lambda_t} \delta'(u_t) \tag{38}$$

$$\Upsilon_t \lambda_t = \mu_t^k m_t \left\{ 1 - S\left(\frac{I_t}{I_{t-1}}\right) - S'\left(\frac{I_t}{I_{t-1}}\right) \frac{I_t}{I_{t-1}} \right\} + \beta E_t \mu_{t+1}^k m_{t+1} S'\left(\frac{I_{t+1}}{I_t}\right) \left(\frac{I_{t+1}}{I_t}\right)^2 \tag{39}$$

$$\mu_t^f = -\psi \Gamma_t V_t^{-\sigma} n_t^{\xi} + \beta(1 - \gamma_f) E_t \mu_{t+1}^f \frac{F_{t+1}}{F_t} \tag{40}$$

$$\mu_t^k = \beta E_t \left\{ \lambda_{t+1} r_{t+1} u_{t+1} + \mu_{t+1}^k [1 - \delta(u_{t+1})] \right\} \tag{41}$$

$$\mu_t^h = \beta E_t \left\{ \lambda_{t+1} w_{t+1} n_{t+1} + \mu_{t+1}^h \gamma_h \frac{h_{t+2}}{h_{t+1}} \right\} \tag{42}$$

$$Y_t = z_t \left(\Omega_t \tilde{n}_t\right)^{\alpha} \tilde{K}_t^{1-\alpha}, \tag{43}$$

$$w_t = \alpha \tau_t \frac{Y_t}{h_t n_t} \tag{44}$$

$$r_t = (1 - \alpha) \tau_t \frac{Y_t}{u_t K_t} \tag{45}$$

$$A_t = (1 - \delta_x) X_{t-1} + Y_t \tag{46}$$

$$X_t = A_t - S_t \tag{47}$$

$$\frac{\theta - 1}{\theta} = \beta(1 - \delta_x) E_t \frac{\lambda_{t+1}}{\lambda_t} \tau_{t+1} \tag{48}$$

$$\tau_t = \frac{\zeta}{\theta} \frac{S_t}{A_t} + \frac{\theta - 1}{\theta} \tag{49}$$

3.1.7 Stationarity and Solution Method

The endogenous model economy inherits stochastic trends from the two non-stationary stochastic processes for Υ_t and Ω_t. Our solution method focuses on isolating fluctuations about these stochastic trends. We divide non-stationary variables by their permanent component to yield a stationary version of the model, and then take a linear approximation of the dynamics about the steady state of the stationary system.

The stochastic trend components of output and capital are given by $X_t^y = \Upsilon_t^{\frac{\alpha-1}{\alpha}}\Omega_t$ and $X_t^k = \Upsilon^{\frac{-1}{\alpha}}\Omega_t$ accordingly. The stochastic trend components of all another non-stationary variables can then be expressed as some function of X_t^y and X_t^k. We provide the details of this transformation and resulting stationary equilibrium system in the Appendix.

3.2 Calibration

Our calibration strategy involves using parameter values close to those in the literature for the parameters which are common to the business cycle literature. We show robustness to the calibration of key parameters in Section 3.3.3.

Beginning with the household, we set the household's subjective discount factor β to 0.9957, implied by average inflation and the Federal Funds Rate over our sample, and the elasticity of intertemporal substitution, σ, to 1, as in Jaimovich and Rebelo (2009). We set ξ to 2, implying a Frisch elasticity of labor supply of 1 in the absence of habits in consumption. This places our Frisch elasticity in between the ranges of works such as Schmitt-Grohe and Uribe (2012), Christiano et al. (2014), Christiano et al. (2005), Justiniano et al. (2011), and Jaimovich and Rebelo (2009). For the consumption habits parameter, in general the business cycle literature tends to find non-zero values of b greater than 0.5. We set b equal to 0 to imply no habits as a baseline and explore robustness to this parameter later. Finally, we set γ_f, the preference parameter which determines the strength of the income effect, to 0.01 based on Schmitt-Grohe and Uribe (2012).

For goods production, we set the elasticity of output to current labor α to 0.64 as in Jaimovich and Rebelo (2009). For the parameters related to physical capital, we set steady-state physical capital depreciation δ to 0.025 and the elasticity of marginal utilization $\delta_k''(1)/\delta_k'(1)$ to 0.15. The literature finds a very wide range of values for this elasticity. For example, Smets and Wouters (2003) and Christiano et al. (2005) estimate 0.18 and 0.01

respectively; Schmitt-Grohe and Uribe (2012) estimate 0.34, and Smets and Wouters (2007) estimate 0.54. We choose a value of 0.15 within this range, close to the value of 0.25 used in Jaimovich and Rebelo (2009). The literature also finds a wide range of values for the investment adjustment cost parameter s''. Smets and Wouters (2007) estimate 5.7, Christiano et al. (2005) estimate 2.48, and Schmitt-Grohe and Uribe (2012) estimate 9.1. We set s'' to 5, within this range.

The parameters related to inventories are all based on Lubik and Teo (2012). The inventory depreciation rate, δ_x, is set to 0.05. The taste shifter curvature, ζ, is set to 0.67 to yield a steady state sales-to-stock ratio of 0.55, and the goods aggregator curvature, θ, is set to 6.8 to yield a steady state goods markup of 10%.

For the parameters related to intangible capital, we assume constant returns to scale in the knowledge accumulation equation, setting γ_h, the contribution of prior intangible capital in its own production, to 0.75, and ν_h, the elasticity of labor in intangible capital, to 0.25. These are within the ranges for both parameters in Chang et al. (2002), Cooper and Johri (2002) and Gunn and Johri (2011).

Finally, a number of steady state parameter values are implied by average values in the data, such as the steady state growth rates of GDP and the relative price of investment (RPI). The parameters ψ and are fixed to guarantee steady state utilization equal to unit and steady state hours equal to 0.2. Table 1 summarizes all parameter values used to calibrate the model.

Table 1: Summary of calibrated parameters

Description	Parameter	Value
Subjective discount factor	β	0.9957
Household elasticity of intertemporal substitution	σ	1
Determinant of Frisch elasticity of labor supply	ξ	2
Habit persistence in consumption	b	0
Wealth elasticity parameter	γ_f	0.01
Labor elasticity in production	α	0.64
Depreciation elasticity of capacity utilization	$\delta_k''(1)/\delta_k'(1)$	0.15
Capital depreciation	δ_k	0.025
Investment adjustment cost	s''	5
Inventory depreciation	δ_x	0.05
Goods aggregator curvature	θ	6.8
Taste shifter curvature	ζ	0.67
Contribution of prior intangible capital in its production	γ_h	0.75
Labor elasticity in intangible capital	ν_h	0.25
TFP growth process persistence	ρ_Ω	0.95
Steady state government spending over output	g/y	0.18
Steady state hours	n	0.2
Steady state capacity utilization	u	1
Steady state GDP growth rate (in %)	g^y	0.42545
Steady state RPI growth rate (in %)	g^{rpi}	-.58203

3.3 Model Results

3.3.1 The responses to TFP news shocks

We now investigate the impulse responses to the non-stationary TFP news shock. As an illustrative example, we consider the effect of a 12 quarter ahead news shock which will eventually be realized as anticipated. Figure 6 shows the response of the model economy to this shock. In response to news about a future increase in TFP, inventories rise over time along with the other macroeconomic quantities, in advance of the actual rise in TFP, consistent with our empirical VAR evidence.

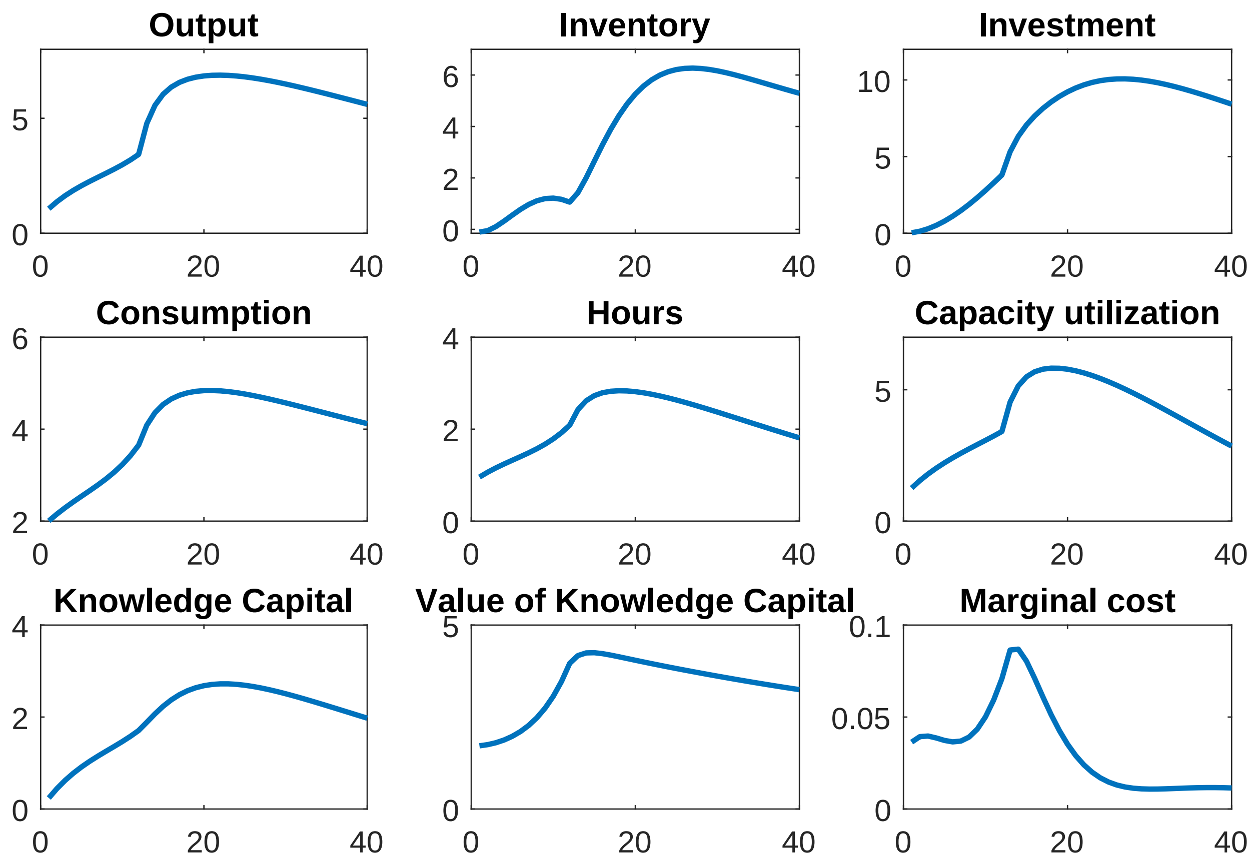

Figure 6: **IRF to 12-period out unit TFP news shock**

Before exploring how the knowledge capital mechanism produces procyclical inventory movements, it is helpful to first understand how the knowledge capital mechanism drives a boom in hours-worked and output. From the household's equilibrium n_t first-order condition (37), the value of additional knowledge capital today μ_t^h is a function of the additional future wage earnings that knowledge capital yields. When the news about future TFP arrives, the household knows that that wages will be high in the future when TFP eventually increases, thus increasing the marginal impact of extra knowledge in wage earnings in the future, driving up the value of knowledge capital μ_t^h in the present. This rise in μ_t^h in turn shifts the household's labour supply curve outwards as the household seeks to increase its knowledge by supplying additional labour, suppressing the rise in the real wage in the labour market, and contributing to an equilibrium rise in hours-worked and thus output. In other words, knowing that there will be technological change in the future, the household begins preparing for that change in the present, building up its knowledge in the present by engaging with production in order to respond optimally when the change comes in the future.[17]

3.3.2 Understanding the dynamics of inventory

Now turning to the response of inventory, the equilibrium distributor optimal stocking condition, given by

$$\tau_t = \frac{\zeta}{\theta}\frac{S_t}{A_t} + \frac{\theta - 1}{\theta} \tag{50}$$

is the key equation governing inventory dynamics in the model in general. In particular, it implies that the distributor targets a specific sales-to-stock ratio $\frac{S_t}{A_t}$ (or equivalently, a specific inventory-sales ratio, $\frac{X_t}{S_t}$, since $\frac{S_t}{A_t} = \frac{S_t}{S_t + X_t} = \frac{1}{1 + X_t/S_t}$), for a given for a given level of marginal costs τ_t. All else equal, the distributor will increase inventory along with a rise in sales (the "demand channel") and reduce inventory along with a rise in contemporaneous marginal costs (the "current cost channel")[18].

To understand the particular impact of TFP news on inventories, we need to consider the general equilibrium effects of TFP news on this expression. In our particular decentralization, one way to think of this equation is as a demand curve for produced output Y_t, in the market

[17]Note that the well-known Jaimovich and Rebelo (2009) mechanism is also in operation here in addition to the above knowledge capital mechanism: under the particular form of investment adjustment costs, the shadow value of capital drops today on account of the value of increasing investment today to lower future adjustment costs. This in turn leads to an outward shift in labour demand by the intermediate goods firm as the firm increases capacity utilization, whose cost depends inversely on the value of capital. Gunn and Johri (2011) show that the knowledge capital mechanism on its own is sufficient to promote co-movement of consumption, investment and hours-worked in the absence of the Jaimovich and Rebelo (2009) mechanism under certain scenarios. In the present application, the low-income effect preferences and variable capacity utilization elements of the Jaimovich and Rebelo (2009) mechanism help to enhance the boom, and as well, the variable capacity utilization helps suppress the rise in marginal costs. Unlike with the Jaimovich and Rebelo (2009) own its own, the particular form of investment adjustments costs here do not play a major role in the model results.

[18]Note that in the above stocking condition, the constant term $\frac{\theta-1}{\theta}$ represents the expected value of future marginal costs, since in equilibrium, $\frac{\theta-1}{\theta} = \beta(1-\delta_x)E_t\frac{\lambda_{t+1}}{\lambda_t}\tau_{t+1}$. In effect, the distributor considers the level of marginal costs today compared to expected future marginal costs in the future (the "intertemporal substitution" channel) when adjusting inventory, but since the later is constant, only dynamic variation in the former will impact variation in inventory. The constancy of expected future marginal costs is an artifact of flexible prices in the current model. Under sticky prices, expected future marginal costs would become dynamic as the distributor optimally varies its markup between price and the marginal cost of sales in response to shocks. We discuss the case of sticky prices in the companion paper Gortz, Gunn and Lubik (2019b).

for Y_t with market-clearing price τ_t (where we recall that τ_t is also equivalent to the marginal cost of production). Indeed, using (46), we visualize the optimal stocking condition

$$\tau_t = \frac{\frac{\zeta}{\theta} S_t}{(1-\delta_x)X_{t-1} + Y_t} + \frac{\theta - 1}{\theta}, \tag{51}$$

as a downward-sloping demand curve for Y_t in a diagram with τ_t on the vertical axis and Y_t on the horizontal axis, such that all else equal, increases in τ_t imply a lower optimal inventory-sales ratio, and thus lower demand for produced output Y_t, as distributors seek to run down inventories. For a given τ_t, an increase in sales increases the demand for produced output Y_t as the distributors seek to maintain their sales-inventory ratio by increasing inventory, shifting this demand curve to the right.

On the other side of the market, we can combine the household's labour supply, the intermediate firm's labor demand, and the production technology to form a output supply curve as a function of τ_t. Assuming no consumption habits ($b = 0$), no-income effect ($\gamma_f = 0$) and fixed the preference shock Γ_t at unity for simplicity, this curve is given by

$$\tau_t = \frac{1}{Y_t}\left[\frac{\xi}{\alpha}\psi Q_t^{-\frac{\xi}{\alpha}} Y_t^{\frac{\xi}{\alpha}} - \frac{\nu_h}{\alpha}\phi_t^h\right] \tag{52}$$

where $\frac{\partial \tau_t}{\partial Y_t} > 0$ for $\xi > \alpha$ such that the curve is upward-sloping, $Q_t = z_t(\Omega_t h_t)^{\alpha}\tilde{K}_t^{1-\alpha}$, and $\phi_t^h = \frac{\mu_t^h}{\lambda_t} h_{t+1}$ is given by

$$\phi_t^h = \beta E_t \frac{\lambda_{t+1}}{\lambda_t}\left\{\alpha\tau_{t+1}Y_{t+1} + \gamma_h \phi_{t+1}^h\right\}. \tag{53}$$

Note in (52) that a rise in the value of knowledge μ_t^h increases Y_t for a given level of τ_t, shifting the output supply curve to the right. In effect, when the value of knowledge rises, the household shifts its labour supply curve outwards in order to acquire more knowledge, lowering the real wage for a given level of hours. Since the intermediate firm's marginal cost given by $\mathrm{mc}_t = \frac{w_t}{MP\tilde{N}_t}$, this implies a reduction in marginal cost for a given level of hours, such that the firm's output supply curve shifts outwards in the market for produced output.

We are now in a position to understand the response of inventory to TFP news using the impact of the news on the supply and demand in the market for produced output Y_t. When TFP news arrives, as is well known in the literature, the wealth effect of the TFP news drives up the demand for consumption. In this particular context, this in turn drives up the demand for sales of distributors. This increase in demand for sales then shifts the distributors' output demand curve rightwards as they increase their demand for new produced goods in order to maintain their sales-to-inventory ratio for a given level of marginal costs. In the market for Y_t, for a given position of the output supply curve, the shift in demand puts upwards pressure on τ_t, implying a necessarily lower level of the inventory-to-sales ratio through (51). While inventory could still rise under a drop in the inventory-to-sales ratio, if the rise in marginal costs is large enough, for the given rise in sales, inventories may actually need to decrease as it becomes more attractive for the distributors to draw down inventories in the present to avoid the high current production costs. Indeed, whether inventory will rise or fall for a given increase in sales S_t depends on the magnitude of the rise in marginal costs relative to the increase in sales. For a given position of the output supply curve, the rise in τ_t will be a function of the slope of the supply curve.[19] Yet with the presence of knowledge capital, the rise in the value of knowledge upon receipt of the news suppresses the rise in τ_t as it shifts the output supply curve outwards, allowing for a smaller drop in the inventory-sales ratio and increasing the chance the inventory can rise along with sales. Note however that as long as marginal costs τ_t rise, a countercyclical inventory-sales ratio — a feature consistent with our empirical evidence in Section 2.4 — is a necessary condition for co-movement of inventory.

We can quantitatively explore the role of knowledge capital by examining the impulse response to the same shock except with the knowledge capital mechanism shut-down, represented as the nested case $\gamma_h \approx 1$ and $\nu_h \approx 0$. Under this case, the is no change to the output

[19]Note the dependence of the slope of the output supply curve and thus magnitude of the rise in marginal costs on the Frisch elasticity parameter ξ. Note also that the slope is decreasing in Q_t, such that contemporaneous increases in capacity utilization u_t flatten the slope, suppressing the rise in marginal costs. Moreover, dynamic increases in the predetermined stocks h_t and K_t over time flatten the curve over time, suppressing the rise in marginal costs over time through the boom.

demand curve (51), but the output supply curve reduces down to

$$\tau_t = \frac{\xi}{\alpha}\psi Q_t^{-\frac{\xi}{\alpha}} Y_t^{\frac{\xi}{\alpha}-1}. \tag{54}$$

Figure 7 shows the response of the model economy with knowledge capital to the same shock as in Figure 6. In response to news about a future increase in TFP, inventories now fall over time in advance of the rise in TFP. Without the shift in the output supply curve from knowledge capital, marginal costs rise too much, leading to too-large a fall in the inventory-sales ratio and thus a fall in inventory. Additionally, the presence of inventory actually negatively impacts the co-movement of other macroeconomic variables such as hours, output and investment. Why should the presence of inventory adversely impact these variables? Despite the increase in labor demand from the Jaimovich-Rebelo mechanism, distributors can reduce their demand for produced goods (relative to the model without inventories), since they can meet some of the demand for sales by drawing down inventories, which in turn reduces the demand for labor and capacity utilization as inputs into production. The fall in inventories is thus intimately linked to the muted response of hours, which then leads to a muted response in output and utilization and other quantities.

3.3.3 Robustness

We now explore the sensitivity of the results to variation in the calibration of key parameters. In general, the response of inventory to TFP news is most sensitive to those parameters that impact the response of marginal costs. In particular, as discussed earlier, ξ, which parameterizes the Frisch elasticity, $\delta_k''(1)/\delta_k'(1)$ the elasticity of marginal utilization, and ν, which paramaterizes the contribution of labour in knowledge capital, accumulation all affect of the output supply curve and thus drive the response of marginal costs.

Figures 10, 9 and 10 show the response of the model economy to news about a rise in TFP 12 periods in the future over ranges of ξ, $\delta_k''(1)/\delta_k'(1)$ and ν respectively. Figures 11 and 12 show the response of the model economy with the knowledge capital marginal shut

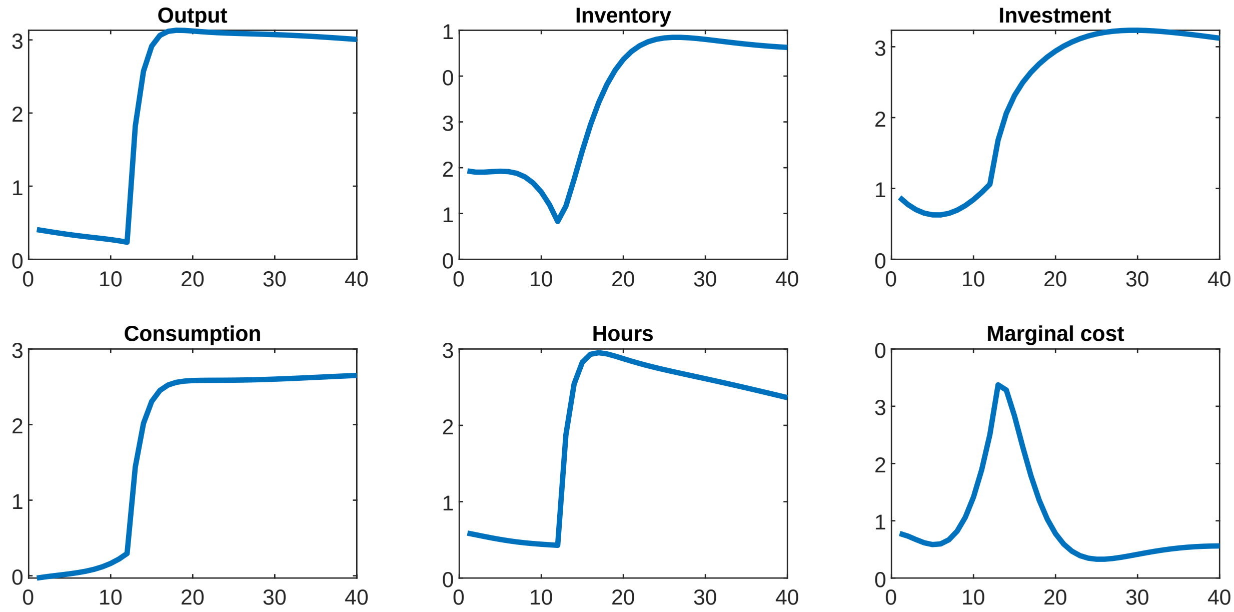

Figure 7: **IRF to 12-period out unit TFP news shock - Model *without* knowledge capital**

down to the same shock, over ranges of ξ, $\delta_k''(1)/\delta_k'(1)$.

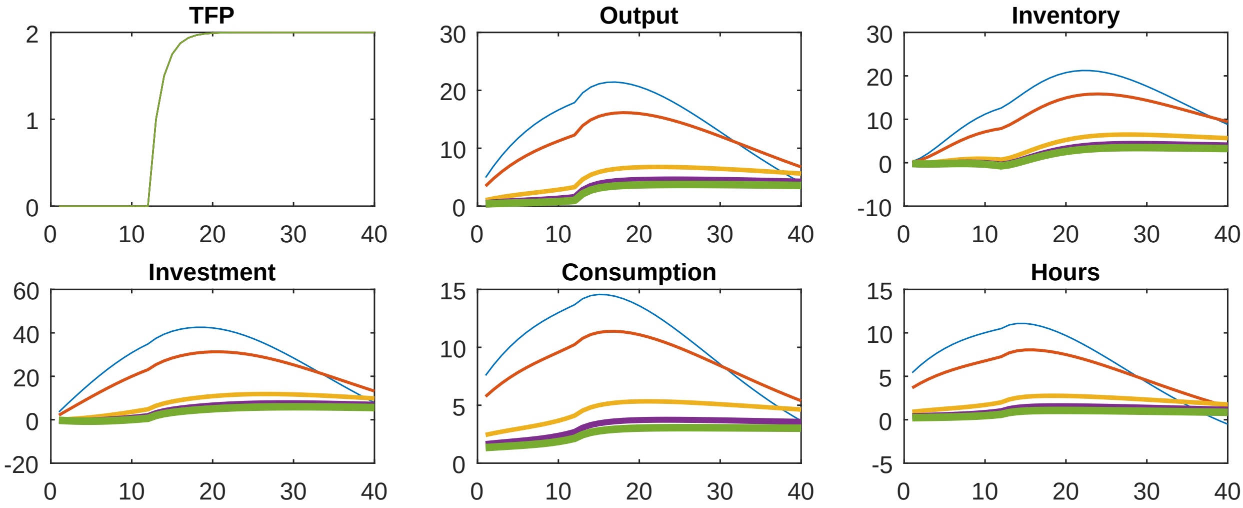

Figure 8: **IRF sensitivity to $\xi = \{1.4, 1.5, 2, 2.5, 3\}$ (thin to thick lines)**

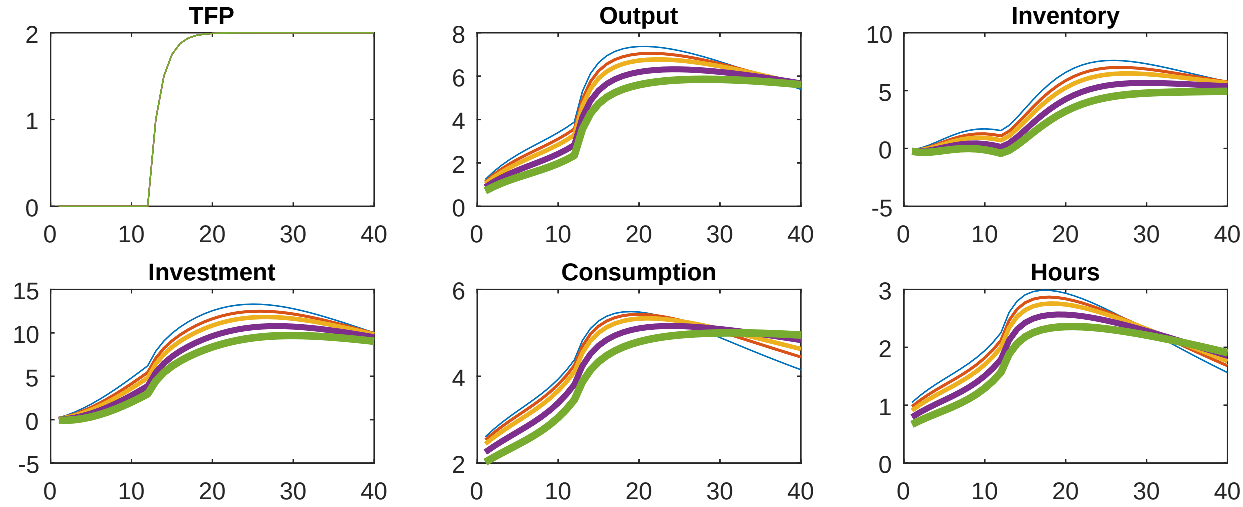

Figure 9: **IRF sensitivity to $\delta_k''(1)/\delta_k'(1) = \{0.05, 0.1, 0.15, 0.25, 0.4\}$ (thin to thick lines)**

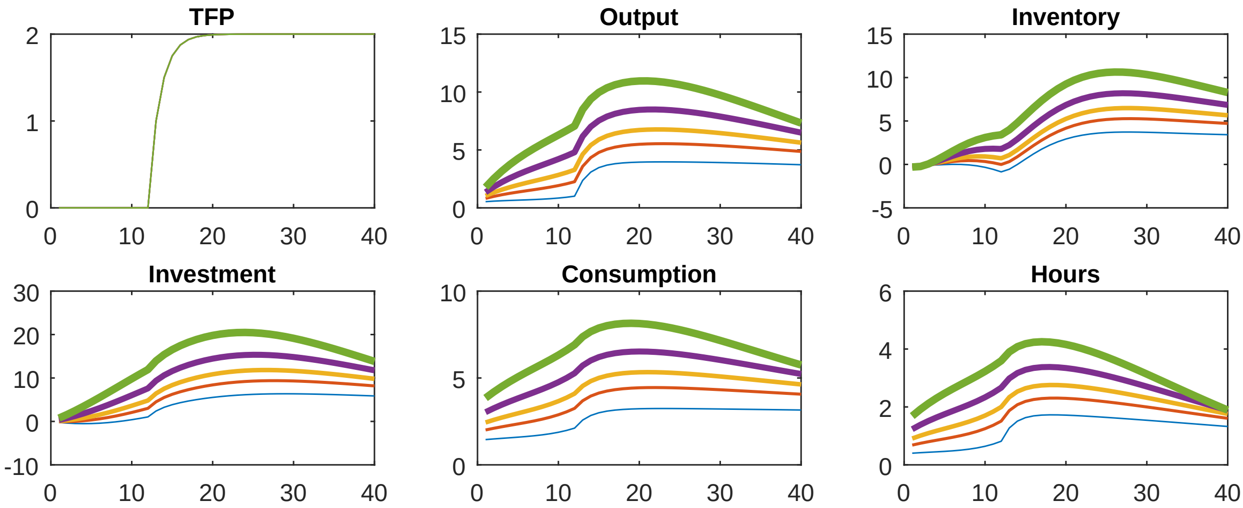

Figure 10: **IRF sensitivity to $\nu = \{0.1, 0.2, 0.25, 0.3, 0.35\}$ (thin to thick lines)**

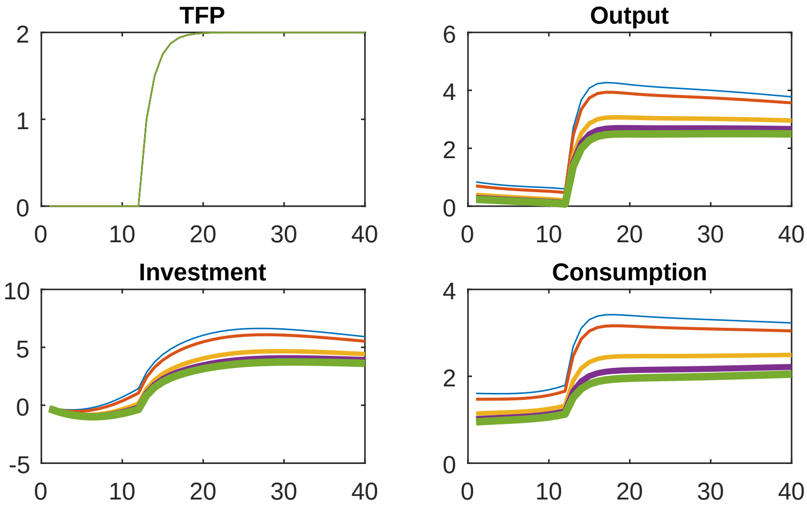

Figure 11: **IRF sensitivity to $\xi = \{1.4, 1.5, 2, 2.5, 3\}$ (thin to thick lines) - Model *without* knowledge capital**

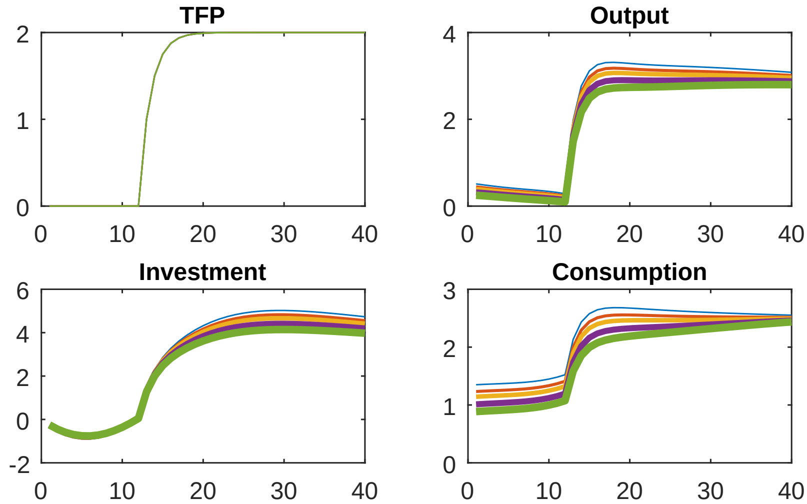

Figure 12: **IRF sensitivity to $\delta_k''(1)/\delta_k'(1) = \{0.05, 0.1, 0.15, 0.25, 0.4\}$ (thin to thick lines) - Model *without* knowledge capital**

3.4 Confronting the DSGE model with the empirical VAR evidence

In the section above, we have shown that our baseline model with knowledge capital can generate a positive inventory response — alongside an expansion in all other macroeconomic aggregates — in light of positive news about future TFP. This has been exemplified showing IRFs to a 12 quarter ahead TFP news shock. We want to evaluate, how well anticipated TFP news shocks, at a variety of horizons, can generate the comovement patterns of inventories and the other macroeconomic aggregates established from the empirical VAR in Section 2.3.

For this purpose, we retain the calibrated parameters values used above and apply Bayesian econometric techniques to estimate parameters related to the model's shock processes. We allow for four, eight and twelve quarter ahead news shocks to the growth rate of TFP. These TFP news shocks compete with a number of other anticipated and unanticipated shocks in explaining the auto- and cross-correlation spectrum of macroeconomic aggregates.[20] Estimating the shock processes allows us to perform a Monte Carlo experiment. We generate 1,000 samples of artificial data from the DSGE model by drawing parameter values from the posterior distribution. For each sample, we construct the level of the model generated time series for 142 periods – this series length is consistent with the time horizon used in Section 2.3 on the empirical VAR evidence. We then compare the empirical responses from the VAR model with those estimated with identical VAR specifications on the artificial data samples.

Figure **??** shows median IRFs (thick blue line) and 16% and 84% posterior bands (dashed blue lines) from the empirical VAR model, as well as the median (thin black line) and posterior bands (gray shaded area) from the Monte Carlo experiment. The dynamic responses from the VAR on artificial data are qualitatively in line with the responses from the empirical VAR. Importantly, inventories rise on impact in response to the TFP news shock — as do output, investment, consumption and hours worked. Also quantitatively the empirical and model-implied VAR responses are close as variables' posterior bands overlap for the vast majority of periods. A noticeable exception is the response of TFP where the empirical IRFs rise much stronger in the long run than the counterpart based on artificial data. We note however that the observable set used to estimate the DSGE model does not include TFP which is in line with the vast majority of studies that use DSGE models to provide inference on technology shocks.[21] Given that the DSGE estimation includes a much larger

[20]We estimate the model over the horizon 1983:Q1–2018:Q2 (same as the VAR evidence in the previous sections) using GDP, consumption, investment, hours worked and inventories as observables. The model includes stationary and non-stationary TFP shocks, non-stationary IST shocks, a marginal efficiency of investment (MEI) shock, a government spending shock and a preference shock. All shocks except the preference shock include in addition to the surprise innovations also four, eight and twelve quarter ahead news components. Our setup of shock processes, treatment of observables and prior specifications is standard and close to related studies such as Schmitt-Grohe and Uribe (2012) or Khan and Tsoukalas (2012). Details about the shock processes and the estimation are provided in Appendix B.

[21]We do not use the utilization adjusted TFP measure in the DSGE estimation since it lacks corrections for

number of anticipated and unanticipated innovations than the six-variable VAR, any comparison between the two methodologies to identify TFP news shocks has its limitations.[22] For our purposes the qualitative consistency across IRFs provides useful evidence that our baseline model is able to reproduce the new empirical fact that inventories co-move alongside the other macroeconomic variables in response to TFP news shocks. We see the qualitative similarity across responses as an additional success of our model. Particularly since our parsimonious framework – which eases the discussion of propagation mechanisms in the sections above – may limit the quantitative consistency between empirical and model implied VAR responses due to the omission of transmission mechanisms that have been found important in the literature on estimated DSGE models.[23]

In the above exercise the model implied time series for TFP, that was subsequently used to identify the news shock using the VAR methodology, is based on $TFP_t = z_t \Omega_t^{\alpha}$ according to the production function (16). While TFP_t is the productivity process based on our model, the intangible capital term in our production function would show up in the empirical construction of John Fernald's productivity series as part of the Solow residual since the underlying production function to his computations is a composite of labor and capital services only. For this reason, we extract an alternative expression for productivity from our model

imperfect competition and potential mark-up variation as well as factor reallocation that are only available with annual data. Thus, as emphasized also in Kurmann and Sims (2016), short-run movements in quarterly TFP series may potentially reflect non-technology factors and therefore a noisy measure of the true underlying technological process. This is problematic since the DSGE estimation would force the model-implied TFP to exactly replicate the imperfect measure of TFP. Our VAR identification is much less prone to be affected by the noise as it relies on a long-run restriction. While the zero-impact restriction might in principle be affected, we emphasize robustness of our VAR results in Appendix A.3 by using a VAR identification recently proposed by Kurmann and Sims (2016) that solely relies on long-run restrictions.

[22] Note also that there are considerable differences in news shock identification between the two methodologies. The VAR-identification relies exclusively on TFP as observable to identify the news shock. In contrast the Bayesian methodology uses the whole auto- and cross-correlation spectrum of all observables to identify TFP news shocks, however the set of observables typically does not include TFP.

[23] For a discussion on the importance of nominal rigidities and financial frictions in estimated models with anticipated technology shocks, see for example Görtz and Tsoukalas (2017).

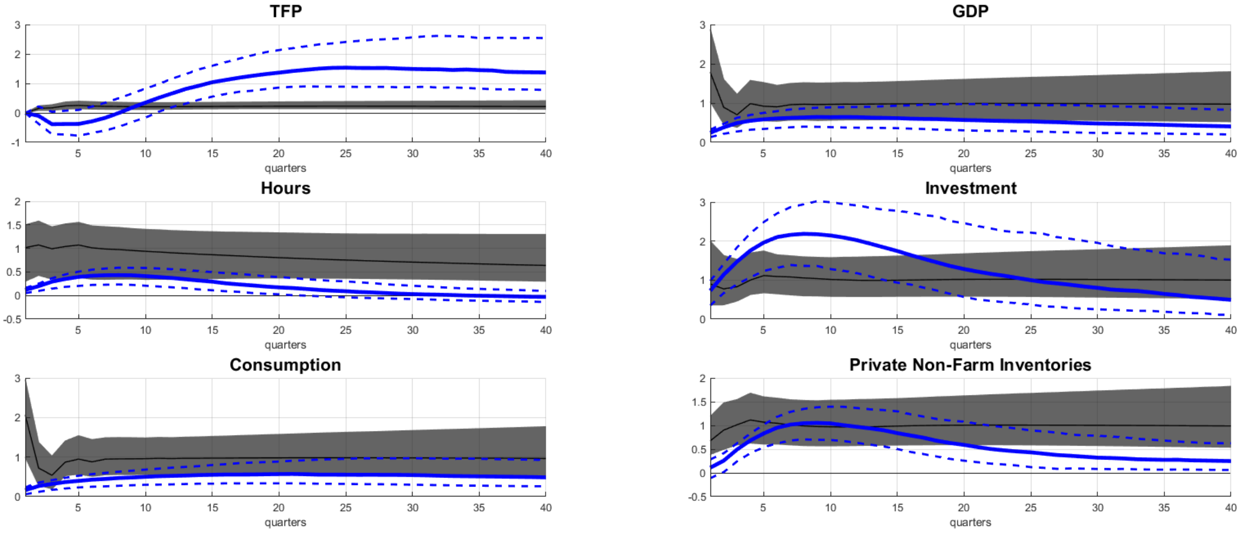

Figure 13: **TFP news shock.** The blue solid (dashed) line is the median (16% and 84% posterior band) response to a TFP news shock from a six-variable VAR. The solid black line (gray shaded areas) is the median (16% and 84% posterior band) response to a TFP news shock estimated from a VAR on 1,000 samples generated from the DSGE model. Units of the vertical axes are percentage deviations.

$$TFP_t^{alt} = \frac{Y_t}{n_t^{\alpha}(u_t K_t)^{1-\alpha}} = z_t \left(\Omega_t h_t\right)^{\alpha},$$

that corresponds to the construction of Fernald's empirical measure. Figure 14 shows median responses (blue thick line) and posterior bands (blue dashed lines) from the empirical VAR model as well as the median (thin black line) and posterior bands (gray shaded areas) from the Monte Carlo experiment. In this case the news shock on artificial data is identified based on our alternative productivity measure, TFP_t^{alt}. IRFs are in line with the ones in Figure **??**: qualitatively the responses on artificial data are consistent with the empirically observed comovement of inventories with the other macroeconomic aggregates — and also qualitatively it is striking that posterior bands overlap for the majority of periods. Additionally, it is noticeable though that using the alternative measure for productivity the medium- and long-run responses based on empirical and artificial data are now also qualitatively very closely aligned.[24]

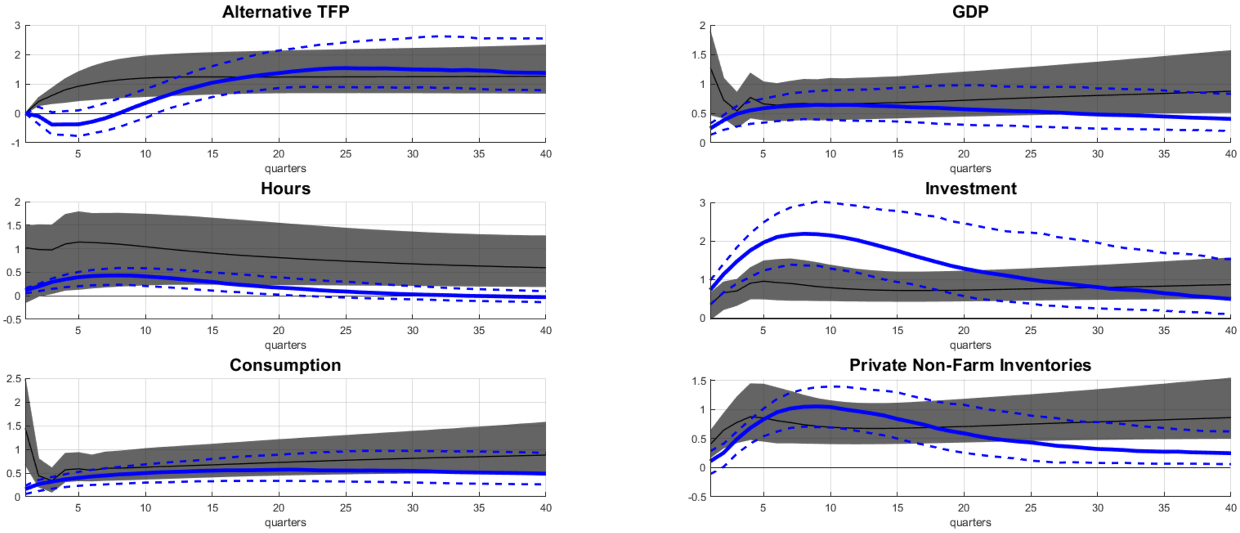

Figure 14: **TFP news shock.** The blue solid (dashed) line is the median (16% and 84% posterior band) response to a TFP news shock from a six-variable VAR. The solid black line (gray shaded areas) is the median (16% and 84% posterior band) response to a TFP news shock estimated from a VAR on 1,000 samples generated from the DSGE model. Units of the vertical axes are percentage deviations.

4 Conclusion

In this paper we use standard VAR identification to document a new empirical fact: in response to TFP news, inventories move procyclically along with the other major macroeconomic aggregates. This fact is robust across many dimensions such as sectors and types of inventories. Even though unconditionally inventories are strongly procyclical, conditional on TFP news shocks our finding is not a priori self-evident. Conventional views would suggest two potential counteracting effects on inventories in response to news about higher future productivity. A negative substitution effect provides incentives to run the current inventory stock down and increase stockholding in the future when the higher productivity is actually

[24]The alternative productivity measure TFP_t^{alt} is a function of knowledge capital, amongst technology shocks. Our VAR methodology imposes a zero-impact restriction on TFP to identify the news shock. Note that this methodology can still identify a TFP news shock based on the alternative productivity measure: news in period t about future productivity in the model would not move knowledge capital in the same period as the latter is a state variable and hence predetermined.

realized. We provide evidence that this substitution effect is dominated by a demand effect due to which firms increase inventories in response to sales in light of rising consumption and investment.

Our empirical finding corroborates the view that TFP news shocks are important drivers of macroeconomic fluctuations. However, we show this finding imposes two challenges to existing theoretical frameworks used in the news-literature: First, they fail to reproduce the procyclical inventory movements in response to TFP news shocks due to a strong negative substitution effect. Second, introducing inventories in standard frameworks implies an intertemporal labor choice that makes even comovement of consumption, investment and hours much harder to achieve. Our empirical findings impose this new comovement puzzle to the theoretical literature. A rigorous investigation of data-generating mechanisms goes beyond the scope of this paper and is left for future research. However, we suggest one way to solve the comovement puzzle by extending a standard framework with intangible capital and sticky wages.

References

Arrow, K.J., "The Economic Implications of Learning by Doing," *Review of Economic Studies*, 1962, *3* (29), 155–173.

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

_ **and** _ , "Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence," *American Economic Review*, June 2012, *102* (4), 1343–77.

Beaudry, Paul and Franck Portier, "An exploration into Pigou's theory of cycles," *Journal of Monetary Economics*, September 2004, *51* (6), 1183–1216.

_ **and Frank Portier**, "News Driven Business Cycles: Insights and Challenges," *Journal of Economic Literature*, 2014, *52* (4), 993–1074.

Ben Zeev and Hashmat Khan, "Investment Specific News Shocks and U.S. Business Cycles," *Journal of Money, Credit and Banking*, December 2015, *47* (8), 1443–1464.

Bils, Mark and James A Kahn, "What inventory behavior tells us about business cycles," *American Economic Review*, 2000, *90* (3), 458–481.

Blinder, A. S., "Retail Inventory Behavior and Business Fluctuations," *Brookings Papers on Economic Activity*, 1981, *12* (1981-2), 443–520.

_ **and L. J. Maccini**, "Taking Stock: A Critical Assessment of Recent Research on Inventories," *Journal of Economic Perspectives*, Winter 1991, *5* (1), 73–96.

Chang, Yongsung, Joao F. Gomes, and Frank Schorfheide, "Learning-by-Doing as a Propagation Mechanism," *American Economic Review*, December 2002, *92* (5), 1498–1520.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans, "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy," *Journal of Political Economy*, February 2005, *113* (1), 1–45.

Christiano, Lawrence, Roberto Motto, and Massimo Rostagno, "Risk Shocks," *American Economic Review*, 2014, *104* (1), 27–65.

Cooper, Russell and Alok Johri, "Learning-by-doing and aggregate fluctuations," *Journal of Monetary Economics*, 2002, *49* (8), 1539–1566.

Crouzet, Nicolas and Hyunseung Oh, "What do inventories tell us about news-driven business cycles?," *Journal of Monetary Economics*, 2016, *79*, 49–66.

Fernald, J., "A quarterly, utilization–adjusted series on total factor productivity," *Working Paper*, 2014, (2012-19).

Forni, M., L. Gambetti, and L. Sala, "No news in business cycles," *Economic Journal*, 2014, *124*, 1168–1191.

Francis, N., M. Owyang, J. Roush, and R. DiCecio, "A Flexible Finite-Horizon Alternative to Long-run Restrictions with an Application to Technology Shocks," *Review of Economics and Statistics*, 2014, *96*, 638–647.

Galí, J. and L. Gambetti, "On the sources of the Great Moderation," *The American Economic Journal: Macroeconomics*, 2009, *1*, 26–57.

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

Görtz, Christoph and John Tsoukalas, "Sectoral TFP News Shocks," *Economics Letters*, 2018, *forthcoming.*

_ , _ , and Francesco Zanetti, "News Shocks under Financial Frictions," Technical Report 2017.

Greenwood, J., Z. Hercowitz, and G. Huffman, "Investment, Capacity Utilization, and the Real Business Cycle," *The American Economic Review*, 1988, *78*, 402–217.

Gunn, Christopher and Alok Johri, "News and Knowledge Capital," *Review of Economic Dynamics*, 2011, *14* (1), 92–101.

Gunn, Christopher M, "Animal spirits as an engine of boom-busts and throttle of productivity growth," *Journal of Economic Dynamics and Control*, 2015, *57*, 24–53.

Hou, Keqiang and Alok Johri, "Intangible capital, the labor wedge and the volatility of corporate profits," *Review of Economic Dynamics*, 2018, *29*, 216 – 234.

Jaimovich, Nir and Sergio Rebelo, "Can News about the Future Drive the Business Cycle?," *American Economic Review*, 2009, *99* (4), 1097–1118.

Jones, Christopher S. and Selale Tuzel, "New Orders and Asset Prices," *Review of Financial Studies*, 2013, *26* (1), 115–157.

Jung, Yongseung and Tack Yun, "Monetary policy shocks, inventory dynamics, and price-setting behavior," *Federal Reserve Bank of San Francisco Working Paper Series*, 2006, *2006-02.*

Jung, YongSeung and Tack Yun, "Inventory investment and the empirical Phillips curve," *Journal of Money, Credit and Banking*, 2013, *45* (1), 201–231.

Justiniano, Alejandro, Giorgio E. Primiceri, and Andrea Tambalotti, "Investment shocks and the relative price of investment," *Review of Economic Dynamics*, 2011, *14* (1), 101 – 121.

Kesavan, Saravanan, Vishal Gaur, and Ananth Raman, "Do Inventory and Gross Margin Data Improve Sales Forecasts for U.S. Public Retailers?," *Management Science*, 2010, *56* (9), 1519–1533.

Khan, Hashmat and John Tsoukalas, "The Quantitative Importance of News Shocks in Estimated DSGE Models," *Journal of Money, Credit and Banking*, December 2012, *44* (8), 1535–1561.

Kurmann, Andre and Eric Sims, "Revisions in Utilization-Adjusted TFP and Robust Identification of News Shocks," mimeo 2016.

Lubik, Thomas A and Wing Leong Teo, "Inventories, inflation dynamics and the New Keynesian Phillips curve," *European Economic Review*, 2012, *56* (3), 327–346.

McCarthy, Jonathan and Egon Zakrajsek, "Inventory Dynamics and Business Cycles: What Has Changed?," *Journal of Money, Credit and Banking*, 03 2007, *39* (2-3), 591–613.

McGrattan, Ellen R and Edward C Prescott, "Unmeasured Investment and the Puzzling US Boom in the 1990s," *American Economic Journal: Macroeconomics*, 2010, *2* (4), 88–123.

Pavlov, Oscar and Mark Weder, "Countercyclical Markups and News-Driven Business Cycles," *Review of Economic Dynamics*, April 2013, *16* (2), 371–382.

Ramey, V. A. and K. D. West, "Inventories," in J. B. Taylor and M. Woodford, eds., *Handbook of Macroeconomics*, Vol. 1 of *Handbook of Macroeconomics*, Elsevier, December 1999, chapter 13, pp. 863–923.

Schmitt-Grohe, Stephanie and Martin Uribe, "What's news in business cycles?," *Econometrica*, November 2012, *80* (6), 2733–2764.

Smets, F. and R. Wouters, "An estimated dynamic stochastic general equilibrium model of the Euro Area," *Journal of the European Economic Association*, 2003, *1*, 1123–75.

Smets, Frank and Rafael Wouters, "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach," *American Economic Review*, 2007, *97* (3), 586–606.

Uhlig, Harald, "What moves real GNP?," Technical Report, Humboldt University Mimeo May 2003.

Vukotic, Marija, "Sectoral Effects of News Shocks," mimeo, University of Warwick January 2016.

Wen, Y., "Understanding the inventory cycle," *Journal of Monetary Economics*, November 2005, *52* (8), 1533–1555.

5 Appendix

A Additional VAR evidence

A.1 Forecast Error Variance Decomposition

Figure 15 displays the variance shares explained by the TFP news shock.

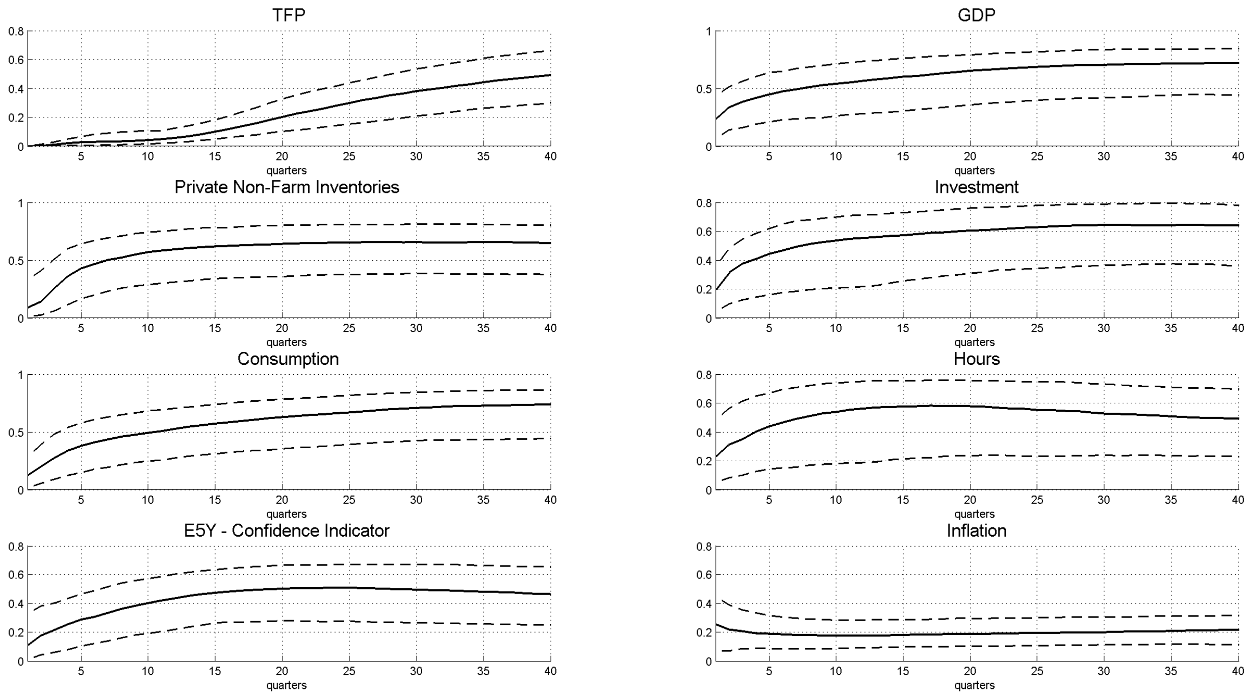

Figure 15: Forecast error variance decomposition (FEVD) of variables to the TFP news shock. Sample 1983Q1-2018Q2. The solid line is the median and the dashed lines are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters.

A.2 Response of inventories over a longer sample.

Changes in the behavior of inventories coinciding with the onset of the Great Moderation have been widely documented in the literature (e.g. McCarthy and Zakrajsek (2007)). This debate, in addition to data availability issues highlighted in the main body, motivates our focus on a post post Great Moderation sample. However, it is interesting to evaluate whether the rise of inventories we document in anticipation of higher future TFP is present also when considering longer samples. Figure 16 shows this is indeed the case for a 1960Q1-2018Q2 sample. This figure reports strong comovement of all macroeconomic aggregates, including inventories, several quarters before TFP increases significantly. This sample is restricted by the availability of the E5Y. If we substitute this variable by the S&P500 we can consider a 1948Q1-2018Q2 sample. Figure 17 shows that also IRFs based on this sample are qualitatively and largely also quantitatively very similar to the results based on our 1983Q1-2018Q2 baseline sample and the 1960Q1-2018Q2 sample. Overall, the fact that inventories rise in response to a TFP news shock is very robust if our baseline sample is extended.

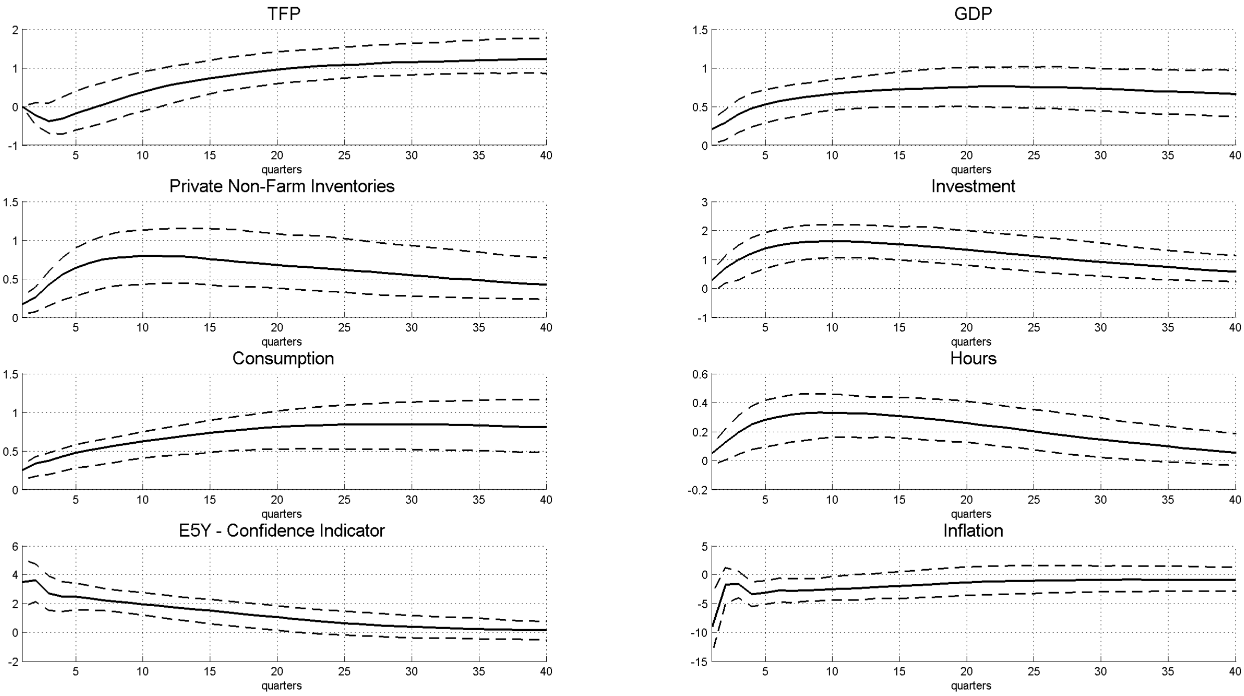

Figure 16: **IRF to TFP news shock. Sample 1960Q1-2018Q2.** The solid line is the median and the dashed lines are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

A.3 Robustness to alternative VAR news identification

The results in the main body of the paper are generated using the Max-share method proposed by Francis et al. (2014). This method is widely used in the literature and implies the news shock is identified as the shock that (i) does not move TFP on impact and (ii)

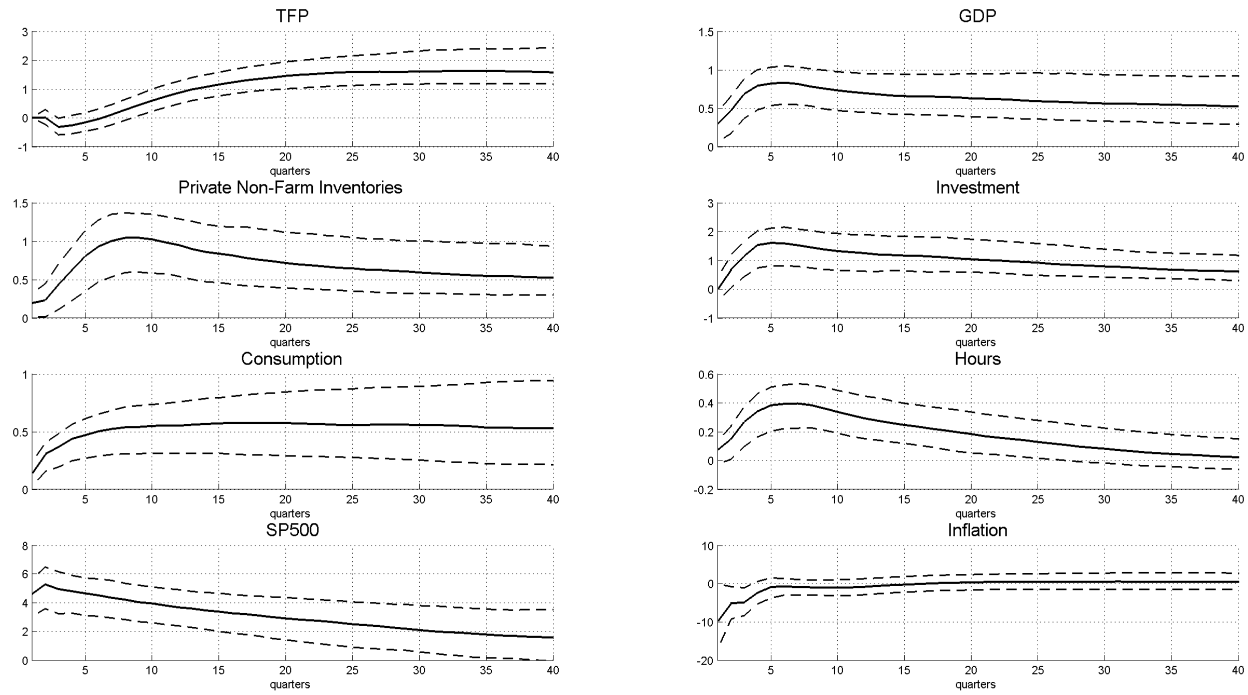

Figure 17: **IRF to TFP news shock. Sample 1948Q1-2018Q2.** The solid line is the median and the dashed lines are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

maximizes the variance of TFP at the 40 quarter horizon. This section shows robustness of findings using three closely related alternative approaches. First, the identification scheme suggested in Barsky and Sims (2011) that recovers the news shock by maximizing the variance of TFP over horizons from zero to 40 quarters, and the restriction that the news shock does not move TFP on impact. Second, the Forni et al. (2014) long-run identification scheme which is similar in spirit to the Max Share method. This method identifies the news shock by imposing the zero impact restriction on TFP and seeks to maximise the impact of the news shock on TFP in the long run. Third, the identification scheme in Kurmann and Sims (2016), that recovers the news shock by maximizing the forecast error variance of TFP at a long horizon without however imposing the zero impact restriction on TFP conditional on the news shock.[25]

Figure 18 provides a comparision between the median responses based on the Max share

[25] Kurmann and Sims (2016) argue that allowing TFP to jump freely on impact, conditional on a news shock, produces robust inference to cyclical measurement error in the construction of TFP.

method and the methods proposed by Barsky and Sims (2011) and Forni et al. (2014). The median responses of the Max Share methodology and the Forni et al. (2014) methodology are virtually indistinguishable and also the median based on the Barsky and Sims (2011) methodology is very similar. Figure 19 also shows that responses based on the methodology proposed by Kurmann and Sims (2016) are qualitatively and quantitatively very similar to the ones based on the Max-share method. Importantly, all methods suggest inventories increase in anticipation of higher future TFP.

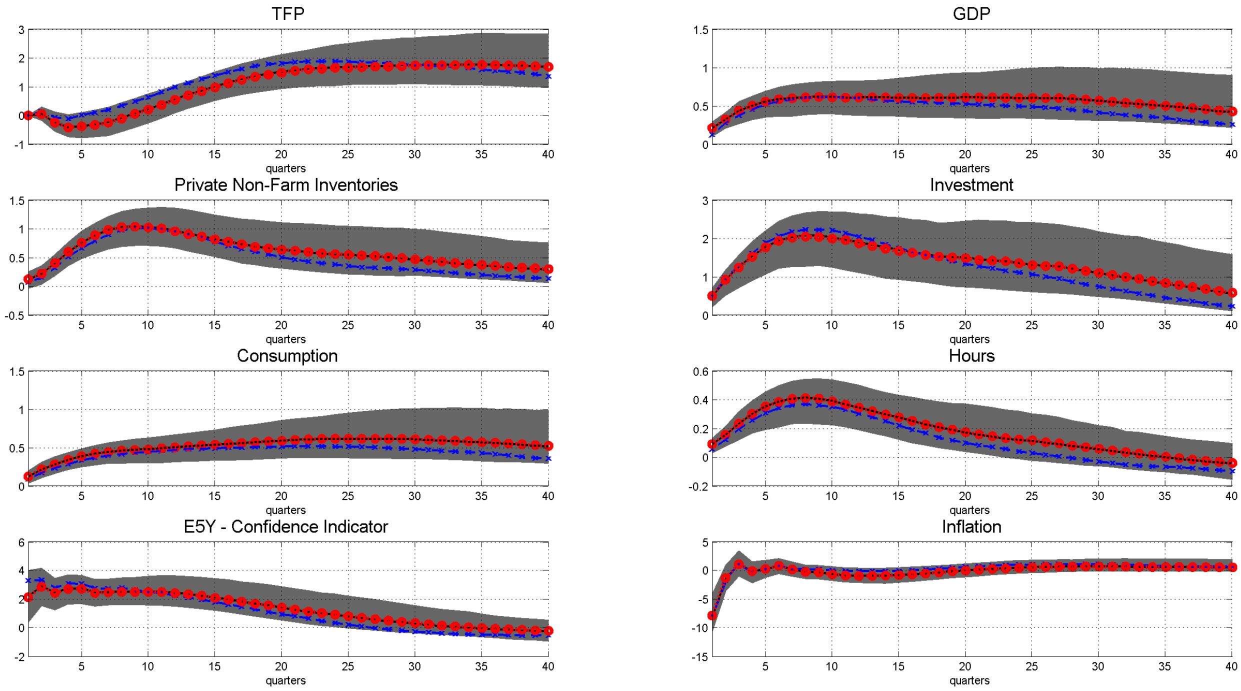

Figure 18: **IRF to TFP news shock.** Sample 1983Q1-2018Q2. The black solid line is the median response identified using the Max-share method. The shaded gray areas are the corresponding 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The blue line with crosses (red line with circles) is the median response identified using the Barsky and Sims (2011) (Forni et al. (2014)) methodology. The units of the vertical axes are percentage deviations.

B Shock processes and Bayesian Estimation

To estimate the model, we include the following exogenous disturbances: a shock to the growth rate of TFP (a_t), a shock to the level of TFP (z_t), a shock to the growth

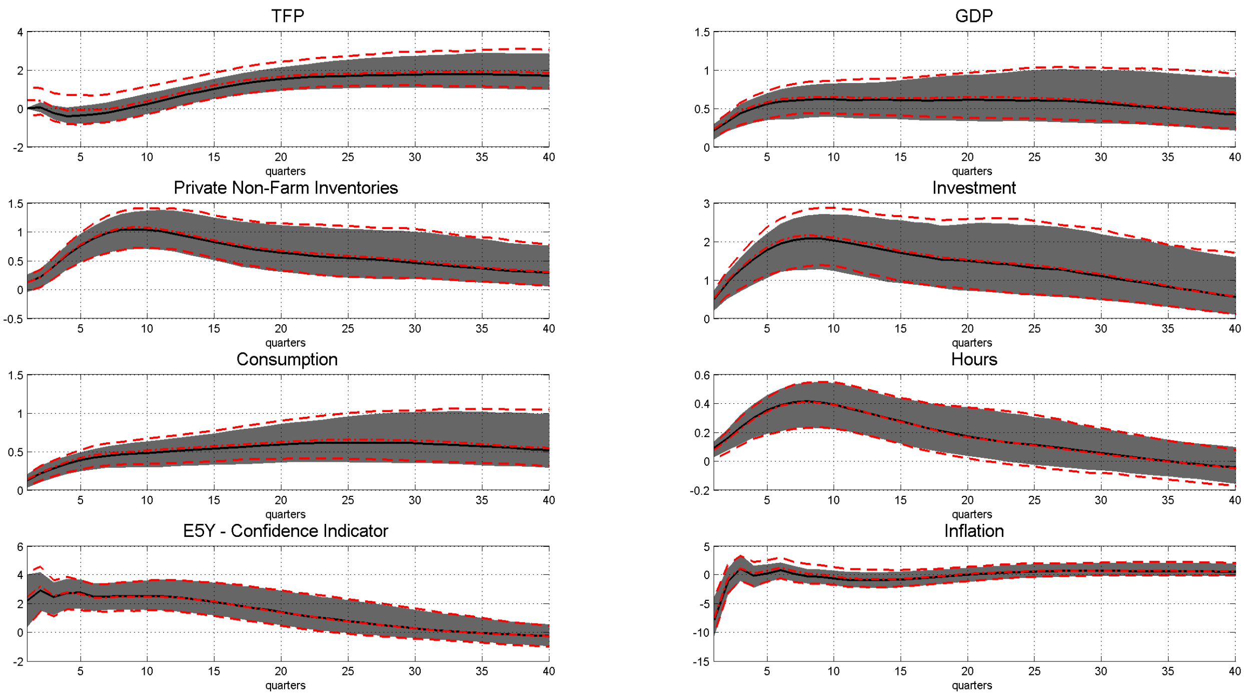

Figure 19: **IRF to TFP news shock.** Sample 1983Q1-2018Q2. The black solid (red dash-dotted) line is the median response identified using the Max-share (Kurmann and Sims (2016)) method. The shaded gray areas (dashed red lines) are the corresponding 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

rate of IST (v_t), a marginal efficiency of investment (MEI) shock (μ_t), a preference shock (b_t) and a government spending shock (g_t). Each exogenous disturbance is expressed in log-deviations from the steady state as a first-order autoregressive (AR(1)) process whose stochastic innovation is uncorrelated with other shocks, has a zero-mean, and is normally distributed. In addition to the unanticipated innovations to the above shocks, the model allows for anticipation effects. In particular, all shock processes (with the exception of the preference shock) include four, eight and twelve quarter ahead innovations. Our treatment of anticipated and unanticipated components is standard and in line with the literature.[26]

We estimate the model over the horizon 1983:Q1–2018:Q2 (same as the VAR evidence in the previous sections) using GDP, consumption, investment, inventories and hours worked as observables. The variables are expressed in real (except hours worked), per-capita terms as outlined in Section 2.2 and GDP, consumption, investment and inventories enter the vector of observables in first-differences. We demean the data prior to estimation.[27]

Model parameters shown in Table 1 are calibrated and we estimate the shocks' persistence and standard deviations so that we can use the model as a data generating process. The prior distributions conform to the assumptions in Schmitt-Grohe and Uribe (2012) or Khan and Tsoukalas (2012). The prior means assumed for the news components are in line with these studies and imply that the sum of the variance of news components is, evaluated at prior means, at most one half of the variance of the corresponding unanticipated component. Table 2 provides an overview about prior and posterior distributions.

[26] For example Schmitt-Grohe and Uribe (2012) also include news components in the processes for government spending shocks and stationary as well as non-stationary neutral and investment specific technology shocks. News shocks also arrive at the four, eight and twelve quarter horizons for example in Görtz et al. (2017).

[27] Removing sample means from the data prevents the possibility that counterfactual implications of the model for the low frequencies may distort inference on business cycle dynamics (see e.g. Christiano et al. (2014)).

Table 2: Prior and Posterior Distributions

Parameter	Description	Prior Distribution			Posterior Distribution		
		Distribution	Mean	Std. dev.	Mean	10%	90%
Shocks: Persistence							
ρ_z	TFP level	Beta	0.5	0.2	0.7362	0.562	0.9602
ρ_b	Preference	Beta	0.5	0.2	0.5101	0.5021	0.5164
ρ_μ	Marginal efficiency of investment	Beta	0.5	0.2	0.9997	0.9996	0.9999
ρ_g	Government spending	Beta	0.5	0.2	0.9234	0.8976	0.9485
ρ_a	TFP growth	Beta	0.5	0.2	0.3985	0.3275	0.4883
ρ_v	IST growth	Beta	0.5	0.2	0.3266	0.0076	0.6483
Shocks: Volatilities							
σ_z	TFP level	Inv Gamma	0.5	2*	0.2808	0.1577	0.4714
σ_z^4	TFP level. 4Q ahead news	Inv Gamma	0.289	2*	0.1309	0.0739	0.1765
σ_z^8	TFP level. 8Q ahead news	Inv Gamma	0.289	2*	0.1018	0.0704	0.1282
σ_z^{12}	TFP level. 12Q ahead news	Inv Gamma	0.289	2*	0.1201	0.0832	0.1632
σ_b	Preference	Inv Gamma	0.5	2*	0.3242	0.1411	0.4899
σ_μ	Marginal efficiency of investment	Inv Gamma	0.5	2*	0.383	0.2177	0.6068
σ_μ^4	MEI. 4Q ahead news	Inv Gamma	0.289	2*	0.1491	0.0837	0.2128
σ_μ^8	MEI. 8Q ahead news	Inv Gamma	0.289	2*	0.2013	0.0726	0.3268
σ_μ^{12}	MEI. 12Q ahead news	Inv Gamma	0.289	2*	1.8434	1.7244	1.9677
σ_g	Government spending	Inv Gamma	0.5	2*	0.5833	0.2013	0.9444
σ_g^4	Gov. spending. 4Q ahead news	Inv Gamma	0.289	2*	1.7283	0.6408	2.8799
σ_g^8	Gov. spending. 8Q ahead news	Inv Gamma	0.289	2*	0.1679	0.0705	0.2862
σ_g^{12}	Gov. spending. 12Q ahead news	Inv Gamma	0.289	2*	2.526	1.6557	3.381
σ_g	TFP growth	Inv Gamma	0.5	2*	0.4606	0.1828	0.7886
σ_a^4	TFP growth. 4Q ahead news	Inv Gamma	0.289	2*	0.1059	0.0549	0.1459
σ_a^8	TFP growth. 8Q ahead news	Inv Gamma	0.289	2*	0.2034	0.0958	0.3047
σ_a^{12}	TFP growth. 12Q ahead news	Inv Gamma	0.289	2*	0.5899	0.4879	0.714
σ_v	IST growth	Inv Gamma	0.5	2*	1.02	0.2276	1.8342
σ_v^4	IST growth. 4Q ahead news	Inv Gamma	0.289	2*	0.5701	0.3949	0.7065
σ_v^8	IST growth. 8Q ahead news	Inv Gamma	0.289	2*	0.3432	0.0906	0.5539
σ_v^{12}	IST growth. 12Q ahead news	Inv Gamma	0.289	2*	0.4953	0.1959	0.8646

Notes. The posterior distribution of parameters is evaluated numerically using the random walk Metropolis-Hastings algorithm. We simulate the posterior using a sample of 500,000 draws and discard the first 100,000 of the draws.