



Working Paper Series

Is There News in Inventories?

WP 20-03

Christoph Görtz
University of Birmingham

Christopher Gunn
Carleton University

Thomas A. Lubik
Federal Reserve Bank of Richmond

This paper can be downloaded without charge
from: <http://www.richmondfed.org/publications/>



Richmond • Baltimore • Charlotte

Is There News in Inventories?*

Christoph Görtz Christopher Gunn
University of Birmingham[†] Carleton University[‡]

Thomas A. Lubik
Federal Reserve Bank of Richmond[§]

May 2020

Abstract

We identify total factor productivity (TFP) news shocks using standard VAR methodology and document a new stylized fact: in response to news about future increases in TFP, inventories rise and comove positively with other major macroeconomic aggregates. We show that the standard theoretical model used to capture the effects of news shocks cannot replicate this fact when extended to include inventories. To explain the empirical inventory behavior, we therefore develop a framework that relies on the presence of knowledge capital accumulated through a learning-by-doing process. The desire to take advantage of higher future TFP through knowledge capital drives output and hours choices on the arrival of news and leads to inventory accumulation alongside the other macroeconomic variables. The broad-based comovement we document supports the view that news shocks are an important driver of aggregate fluctuations.

Keywords: News shocks, business cycles, inventories, knowledge capital, VAR.

JEL Classification: E2, E3.

*We are grateful to Paul Beaudry, Jean-Paul l'Hullier, Alok Johri, Hashmat Khan, Andre Kurmann, Mathias Paustian, Franck Portier, Cedric Tille, and Mark Weder for useful comments and suggestions. We thank seminar and conference participants at the 2018 Canadian Economics Association Conference, the 2019 conference on Computing in Economics and Finance, the 7th Ghent University Workshop on Empirical Macroeconomics, the 2019 UVA-Richmond Fed Research Workshop, the 2019 Money, Macro and Finance Research Group Annual Conference, the 2019 AEA meeting, the 3rd University of Oxford NuCamp Conference, the College of William & Mary, the Deutsche Bundesbank, the University of Sheffield, the University of Windsor, and Drexel University. The views expressed in this paper are those of the authors and not necessarily those of the Federal Reserve Bank of Richmond or the Federal Reserve System.

[†]Department of Economics. University House, Birmingham B15 2TT. United Kingdom. Tel.: +44 (0) 121 41 43279. Email: c.g.gortz@bham.ac.uk

[‡]Department of Economics. Loeb Building, 1125 Colonel By Drive. Ottawa, ON, K1S 5B6. Canada. Tel.: +1 613 520 2600x3748. Email: chris.gunn@carleton.ca.

[§]Research Department, P.O. Box 27622, Richmond, VA 23261. Tel.: +1-804-697-8246. Email: thomas.lubik@rich.frb.org.

1 Introduction

There is substantial evidence that expectations about future total factor productivity (TFP) are an important source of aggregate fluctuations (see Beaudry and Portier, 2014, and references therein). Such TFP news shocks give rise to the observed comovement of aggregate quantities as identified in a large body of empirical work on the incidence and effects on news (e.g., Beaudry and Portier, 2004). Theoretical business cycle models can explain these findings under reasonably general assumptions and modeling components (see Jaimovich and Rebelo, 2009) and imply substantial explanatory power of news shocks when taken to the data directly (e.g., Schmitt-Grohé and Uribe, 2012; Görtz and Tsoukalas, 2017). At the same time, the news-shock literature has largely ignored inventory investment – a component of aggregate output and an adjustment margin to shocks that has long been recognized to play a large role in explaining aggregate fluctuations (see Ramey and West, 1999; Wen, 2005).

In this paper, we argue that inventories should take central stage for understanding the implications of news shocks. In the same vein, we argue that news shocks are an important component in understanding the behavior of inventory investment in addition to the standard mechanisms. Our paper thereby provides further evidence that news is an important component of aggregate fluctuations and that it provides a litmus test by looking at inventories. In particular, we develop a new stylized fact and explain this fact in a general equilibrium model of inventory investment, where we introduce knowledge capital as a key new modeling element.

Our first contribution is the identification of a new fact for the inventory and news-shock literature. Using standard news-shock identification methodology¹ for a structural vector autoregression (VAR) that includes inventories besides other quantity variables, we find that in response to anticipated news about higher future TFP, inventories rise on impact along with output, consumption, investment, and hours worked. This is a robust finding not only for the aggregate data, but also across the retail, wholesale and manufacturing sector as well as for finished goods, work-in-process, and input inventories. Our findings support the insight from the existing literature that news shocks are important drivers of business cycles. Furthermore, the consensus in the literature is that, unconditionally, inventory investment is procyclical (e.g., Ramey and West, 1999), whereby we identify a

¹Our baseline identification scheme is an extension of the approach in Francis et al. (2014). We discuss robustness to alternative identification assumptions in the online appendix.

factor that induces conditional procyclicality.²

The observation that inventories rise in response to news about higher future TFP is not a priori self-evident. In a conventional neoclassical framework with inventories, positive news about future TFP implies a wealth effect. The associated rise in sales of consumption and investment goods creates a demand effect, which drives up inventories in order to avoid stockouts and enhance demand. However, the associated joint increase in sales and inventories can only be met through higher production. This implies rising marginal costs, which provides incentives for firms to partly satisfy higher demand by drawing down the inventory stock. This is reinforced by an intertemporal substitution effect, whereby positive news provides incentives to reduce current inventory stock, but build it up again in the future when high productivity is realized and marginal cost is lower. To the extent that both effects are present, our empirical results suggest the negative substitution effect is dominated by the positive demand effect.

Our second contribution is to identify a theoretical mechanism by which positive news about future TFP generates an expansion of all macroeconomic aggregates, including inventories. Specifically, we reconcile the empirical findings with the standard news-shock model with inventories by providing a role for intangible capital, which we refer to as knowledge capital³, based on earlier work by Chang et al. (2002), Cooper and Johri (2002) and Gunn and Johri (2011). The accumulation of intangible knowledge through a learning-by-doing process involving labor addresses the shortcomings of the standard model in a straightforward manner. Periods of accelerated technological change involve a reorganization of production as the economy prepares for the new technological environment, including the acquisition of new skills, machines, production processes, and materials. In a one-good neoclassical model where all these underlying changes are hidden, we argue that a simple mechanism whereby agents make investments in intangible knowledge to prepare for the future increase in TFP serves as a supply-side proxy for complex production reorganization.

Households acquire skill-enhancing knowledge through a learning-by-doing process from experience in production. The arrival of news about a future increase in TFP raises the value of knowledge in the present, inducing households to increase their labor supply in order

²We find that the TFP news shock explains between 47-71% and 47-65% of the forecast error variance in GDP and inventories, respectively, over a horizon from 6-32 quarters.

³Knowledge capital can be interpreted as an intensive margin of learned skills, for instance, as the knowledge of a worker how to best put to use an hour of work. This includes knowledge about operational processes, handling of machines and materials, and such. See Chang et al. (2002) for an early application in a neoclassical business cycle model and d'Alessandro et al. (2019) for a recent application and further discussion.

to accumulate knowledge through experience. This has the effect of both contributing to the rise in hours worked, and thus production, and of suppressing the rise in the real wage during the initial boom. Consequently, the presence of knowledge capital limits the rise in marginal costs and increases the incentive to accumulate inventories. More succinctly, the accumulation of knowledge capital allows the news-shock-driven demand effect to dominate the substitution effect in production.

The core of our model is the framework of Jaimovich and Rebelo (2009) which nests the model of Schmitt-Grohé and Uribe (2012). It includes the trio of particular specifications of preferences, investment adjustment costs and variable capital utilization, which are features generally recognized in the news literature as needed for generating comovement of macroeconomic aggregates in response to a TFP news shock. We extend this model to include finished goods inventories based on the stock-elastic demand model of Bils and Kahn (2000).⁴ The standard news-shock business cycle model supplemented with inventories cannot replicate the facts from our identified news-shock VAR, as inventories respond countercyclically to TFP news in the model. This behavior results from a too-strong procyclical rise in marginal costs during the expansion. In turn, this countercyclical response of inventories suppresses the positive response of hours and as a result dampens the response of utilization and output. Since firms can satisfy any news-induced increase in sales by drawing down inventories, the demand for labor falls, suppressing the response of hours, utilization, and output relative to sales.

Our findings contribute to the large literature on the role of news shocks as drivers of aggregate fluctuations. Considerable work has been done on studying mechanisms that generate procyclical movements in consumption, investment, and hours in response to TFP news shocks, e.g., Jaimovich and Rebelo (2009) and on studying their effects empirically in identified VARs and estimated DSGE models, for instance, Barsky and Sims (2011, 2012) and Schmitt-Grohé and Uribe (2012). The new aspect our paper adds to this literature is the focus on inventories, both in terms of their behavior in a VAR with news shocks and in developing a theoretical framework to study the empirical results. A large long-standing literature investigates the empirical relation of inventories with macroeconomic fluctuations and the implications of introducing inventories in theoretical frameworks (see Ramey and West, 1999, for a comprehensive survey and critical assessment). In our theoretical modeling of inventories, we are guided by Bils and Kahn (2000), who highlight the unconditionally limited role of intertemporal substitution for variations in inventories that

⁴This mechanism enjoys 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).

is also documented in our work in the context of expectations about productivity.

Our paper is most closely related to Crouzet and Oh (2016), who introduce inventories into a variant of the standard news-shock model of Jaimovich and Rebelo (2009), utilizing a reduced-form stockout-avoidance specification. They show that, while this setup can generate positive comovement of investment, consumption, and hours in response to TFP news shocks, it fails to do so in the case of inventories. The countercyclical inventory movement is then 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) come to the conclusion that TFP news shocks are of limited importance for aggregate fluctuations. In contrast, we use a standard and widely used VAR methodology to identify the response of inventory movements to TFP news first. We thereby establish positive comovement of inventories as a robust stylized fact that we then rationalize in an inventory model with a learning-by-doing propagation mechanism.

The remainder of the paper is structured as follows. Section 2 contains the main empirical results. We first discuss the identification strategy for news shocks and the data used in the VAR analysis, followed by a discussion of the baseline results. We corroborate these in an extensive robustness analysis. Section 3 introduces the theoretical model that we use to rationalize the empirical findings with a focus on inventory modeling and the role of knowledge capital. In section 4 we present the main quantitative results of the paper based on a calibration analysis, while section 5 contains a simulation study that reconciles the theoretical and empirical findings of the paper. Section 6 concludes.

2 Inventories and News: Evidence From an Identified VAR

This section presents our key empirical findings: the positive response of inventories to news shocks and the strong comovement with other macroeconomic aggregates. The results are based on an estimated structural VAR where we identify news shocks based on the so-called Max Share approach. We discuss the empirical model, the identification scheme and the data used in the estimation first, followed by the main empirical results and a wide-ranging robustness analysis.

2.1 VAR-Based Identification of News Shocks

We consider the following vector autoregression (VAR), which describes the joint evolution of an $n \times 1$ vector of variables y_t :

$$y_t = A(L)u_t. \quad (1)$$

$A(L) = I + A_1L + \dots + A_pL^p$ is a lag polynomial of order p over conformable coefficient matrices $\{A_p\}_{i=1}^p$. u_t is an error term with $n \times n$ covariance matrix Σ . We assume a linear mapping between the reduced form errors u_t and the structural errors ε_t :

$$u_t = B_0\varepsilon_t, \quad (2)$$

where B_0 is an identification matrix. We can then write the structural moving average representation of the VAR:

$$y_t = C(L)u_t, \quad (3)$$

where $C(L) = A(L)B_0$, $\varepsilon_t = B_0^{-1}u_t$, and the matrix B_0 satisfies $B_0B_0' = \Sigma$. B_0 can also be written as $B_0 = \tilde{B}_0D$, where \tilde{B}_0 is any arbitrary orthogonalization of Σ and D is an orthonormal matrix such that $DD' = I$.

Identification of news shocks in a structural VAR is based on the idea that information about future movements of a variable such as TFP, namely news, generally affects outcomes even before the shock is realized. At longer time horizons, however, it is likely that the dominant sources of movements in TFP are its own anticipated and unanticipated components. This idea can be utilized explicitly as an identifying assumption for news shocks. At the same time, a second assumption is needed to separate unanticipated shocks from news shocks to TFP. Consistent with Barsky and Sims (2011) and Forni et al. (2014), we impose a zero-impact restriction on TFP to recover the anticipated component based on the assumption that news does not affect TFP contemporaneously.

Mechanically, we identify the news shock by finding a rotation of the identification matrix \tilde{B}_0 , which maximizes the forecast error variance of the TFP series at some finite horizon. In this, we follow the Max Share approach of Francis et al. (2014). Specifically, the h -step ahead forecast error is given by:

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

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

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

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

At a long enough horizon h , variations in TFP are plausibly accounted for by anticipated or unanticipated shocks to this variable. We thus write as an identifying assumption that:

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

where we assume that TFP is ordered first in the VAR system and that the unanticipated and the anticipated (news) shocks are indexed by 1 and 2, respectively. We recover the unanticipated shock as the innovation to observed TFP. It is therefore independent of the identification of the other $n - 1$ structural shocks. The share of total TFP variance that can be attributed to this shock at horizon h is thus $V_{1,1}(h)$, while the remainder is due to news shocks.

The Max Share approach chooses the elements of \tilde{B}_0 to make this restriction on forecast error variance share hold as closely as possible. This is equivalent to choosing the impact matrix so that contributions to $V_{1,2}(h)$ are maximized. Consequently, we choose the second column of the impact matrix to solve the following optimization problem:⁵

$$\begin{aligned} \arg \max_{\gamma} V_{1,2}(h) &= \frac{\sum_{\tau=0}^h A_{i,\tau} \tilde{B}_0 \gamma \gamma' \tilde{B}_0' A_{i,\tau}'}{\sum_{\tau=0}^h A_{i,\tau} \Sigma A_{i,\tau}'}, \\ \text{s.t. } \gamma \gamma' &= 1, \gamma(1,1) = 0, \tilde{B}_0(1,j) = 0, \forall j > 1. \end{aligned} \quad (7)$$

We restrict γ to have unit length to be a column vector of an orthonormal matrix. The second and third constraints impose that a TFP news shock cannot affect TFP contemporaneously.⁶ We therefore identify a TFP news shock from the estimated VAR 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 Estimation

We use quarterly U.S. data for the period 1983Q1 – 2018Q2, which is guided by the observed differences in cross-correlation patterns of several macro-aggregates in samples before and

⁵The optimization problem is written in terms of choosing γ conditional on any arbitrary orthogonalization \tilde{B}_0 to guarantee that the resulting identification belongs to the space of possible orthogonalizations of the reduced form.

⁶Kurmann and Sims (2019) do not impose this exclusion restriction since they argue that TFP is mismeasured and that therefore anticipated and unanticipated movements are indistinguishable at short horizons. We show in the appendix that our results are robust to applying this and other identification methods used in the literature.

after the mid-1980s (e.g., Galí and Gambetti, 2009; Sarte et al., 2015). In particular, McCarthy and Zakrajsek (2007) document that significant changes in inventory dynamics occur in the mid-1980s due to improvements in inventory management. Moreover, several of the time series that we use in the analysis, such as total business inventories and its sectoral components, are only available over part of the post-Great Moderation sample. In our robustness analysis, we document that our results hold also for a longer sample, data availability permitting.

We consider two different measures of total inventories in the VAR. First, non-farm private inventories, which are defined as the physical volume of inventories owned by private non-farm businesses. These are valued at average prices of the period, which captures the replacement costs of inventories. Our second measure, business inventories, differs from the first in how the inventory stock is valued, namely by the cost at acquisition, which can be different from the replacement cost. In NIPA data, inventory profits and losses that derive from differences between acquisition and sales price are shown as adjustments to business income. Unfortunately, business inventories are available only from 1992Q1 on. We therefore reduce the sample horizon accordingly if they are included in the VAR.⁷

Output is measured by 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. Finally, consumption is defined as the sum of personal consumption expenditures for non-durable goods and services. The time series are seasonally adjusted and expressed in real per-capita terms using total population, except for hours, which we do not deflate. In addition to the quantity aggregates, we also use a measure of inflation that we construct from the GDP deflator and a consumer confidence indicator that is based on the University of Michigan Consumer Sentiment Index.⁸ This set of variables is standard in the literature, apart from inventories. The consumer confidence measure provides forward-looking information that potentially captures expectations or sentiment.⁹

⁷Apart from robustness considerations, the use of business inventories is appealing since this measure is available at a disaggregated level for different sectors and inventory types, which we subsequently also use in the VAR.

⁸This indicator, labeled E5Y, 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 indicator is constructed as a diffusion index, namely as the percentage of respondents giving a favorable answer less the percentage giving an unfavorable answer plus 100.

⁹See, for instance, Barsky and Sims (2012). An alternative measure to capture forward-looking information is the S&P 500 stock price index. Our results are robust to including the S&P 500 instead of the

Key to identifying the news shock is a measure of observed technology. We follow the convention in the empirical literature and use the measure of utilization-adjusted TFP provided and regularly updated by Fernald (2012).¹⁰ We identify TFP news shocks from the estimated VAR using the Max Share method outlined in the previous section. Following Francis et al. (2014) we set the horizon h to 40 quarters. All variables enter in levels in line with the news shock VAR literature (e.g., Beaudry and Portier, 2004; Barsky and Sims, 2011). We use Bayesian methods to estimate the VAR with three lags for a Minnesota prior. Confidence bands are computed by drawing from the posterior.

2.3 Results

Figure 1 shows impulse response functions to an identified TFP news shock in the specification with private non-farm inventories. What is striking is that all activity variables increase prior to a significant rise in TFP. In response to news about higher future productivity, TFP does not move significantly for the first 12 quarters. This pattern extends considerably beyond what is imposed by the zero impact assumption of no movements of TFP in the first period. The TFP response peaks toward the end of the horizon.

In contrast, all quantity variables significantly rise on impact and follow a hump-shaped pattern. Moreover, the peak response occurs considerably before TFP hits its highest point. Positive comovement between output, consumption, investment, and hours over this post-Great Moderation sample has been documented before, for instance by Görtz et al. (2017). We add to these previously established stylized facts the behavior of private non-farm inventories, which respond in a similar manner to a news shock: they rise somewhat on impact and continue to do so in a hump-shaped pattern until reaching a peak at about 10 quarters. The change in the stock of inventories, inventory investment, is negative afterwards, while its level never falls below the zero line, its starting point.¹¹

As a robustness exercise, we also consider longer sample periods for the specification with non-farm private inventories, namely samples starting in 1948Q1 and 1960Q1. These results are reported in the online appendix. We find that the impulse response patterns identified in our baseline specification carry over to the two longer samples qualitatively and to a large extent also quantitatively. Overall, across these different samples, the TFP

Michigan consumer confidence index which we document in the online appendix.

¹⁰We use the 2018 vintage, which contains updated corrections on utilization from industry data.

¹¹We also report a short-lived decline in inflation and an anticipation of the future increase in TFP in the consumer confidence indicator, both of which are consistent with previous findings. The significant increase in consumer confidence validates our news shock identification and confirms existing literature (e.g., Barsky and Sims, 2011).

news shock is important for fluctuations in inventories and GDP as it explains between 47-65% and 47-71% of the forecast error variance in inventories and GDP, respectively, over a horizon between 6-32 quarters.¹²

Figure 2 reports the impulse response functions of the specification with business inventories. By necessity, this sample is shorter as the inventory series and its subcomponents are only available since 1992Q1. We consider this alternative specification important as it is not a priori obvious at which prices inventories should be measured. The figure shows that the rise in inventories prior to TFP is robust when we use the business inventory series. All variables exhibit qualitative responses that are very similar to the baseline, although the shorter sample results in somewhat wider confidence bands. Overall, this specification confirms the comovement of macroeconomic aggregates, including inventories, in response to an anticipated TFP shock and prior to the rise in TFP itself.

In the next step, we study the effects of news shocks on inventories in the manufacturing, wholesale, and retail sectors, which comprise the overwhelming majority of inventory stocks. Figure 3 shows the responses of business inventories in each of these sectors to the aggregate TFP news shock. The VAR is estimated by including the sectoral inventories one by one instead of the aggregate inventory measure. The sectoral impulse responses exhibit almost identical hump-shaped patterns: a rise on impact towards a peak response around 10 quarters before declining gradually over the forecast horizon. These results support the finding from the aggregate baseline specification in that the expansion of the inventory stock and other variables is broad-based across sectors.

We also dig deeper into the composition of inventory holdings. The two trade sectors, wholesale and retail, hold almost entirely finished goods inventories, while the inventory stock in the manufacturing sector is split across finished goods inventories (36%), work in process (30%) and input inventories in the form of materials and supplies (34%) over the restricted 1992Q2–2018Q2 sample period for business inventories and their components. Figure 4 shows the responses of inventory types in the manufacturing sector when included one by one in the VAR.¹³ Finished goods and input inventories in the manufacturing sector rise strongly before the realization of anticipated higher productivity as in the baseline specification and all other variations considered above.

We can summarize our findings at this point as follows. Evidence from an identified VAR shows that a news shock about higher future productivity leads to an increase and

¹²The full set of results from the variance decomposition is reported in the online appendix.

¹³The responses of the other variables in the VAR are very similar to the ones reported in Figure 2 and are available upon request.

subsequent positive comovement of all aggregate variables considered. The new fact that we document is that this pattern extends to the response of inventories and is broad-based across different aggregate measures, sectors, and types of inventories. Why the behavior of inventories follows this pattern is a priori not obvious. Conceivably, they could decline initially to satisfy higher demand instead of higher production. Moreover, higher TFP in the future reduces the cost of replenishing a drawn down inventory stock. At the same time, firms may increase inventories to maintain a desired inventory-sales ratio, which counters this effect. It is along these margins that the success of a theoretical model to replicate the empirical findings rests.¹⁴

Jaimovich and Rebelo (2009) document the necessary model elements to facilitate comovement of consumption and investment in response to news about future higher TFP. Specifically, they show that a strong increase in utilization and hours worked are key components. Positive news stimulates consumption through a wealth and income effect. The latter is driven by increased hours worked to raise production in order to satisfy that demand. Similarly, investment increases to support the higher capital stock to take advantage of higher future TFP. This reasoning is corroborated in our structural VAR, where we add additional variables one at a time. Selective impulse responses to a TFP news shock are reported in Figure 5.¹⁵

We find a strong increase in capital utilization which turns negative after about four years once a sufficiently larger capital stock is in place. The positive hump-shaped response of the real wage is consistent with the increase in hours documented in Figure 1. The pattern of the real wage is also indicative of a hump-shaped increase in knowledge capital. Figure 5 further shows that the inventory-to-sales ratio moves countercyclically in response to a news shock. This is a key observation that informs our thinking about a theoretical model. Countercyclicity of the inventory-to-sales ratio is a necessary condition for comovement

¹⁴Görtz et al. (2019) 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. We also study the response of various measures of marginal cost to a TFP news shock. However, none of these measures shows a decline in marginal costs that would point to a strong incentive to run down current inventories and build up stocks again once the higher productivity is realized. Overall, we find evidence against a strong negative substitution effect, but support for a strong positive demand effect. This finding serves further to motivate a demand-enhancing motive for holding more inventories in line with [Bils and Kahn \(2000\)](#).

¹⁵The 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 \(2012\)](#) and consistent with our utilization-adjusted measure for TFP. The real wage is compensation of employees, non-financial corporate business, in real per-capita terms. The change in inventories is the change in private non-farm inventories. The series for intellectual property products is real per-capita nonresidential intellectual property products available from the Bureau of Economic Analysis.

of inventories with the other macroeconomic aggregates. The literature on inventories often does not only consider their level but also their change, which provides an indication about inventory investment. The fourth subplot in Figure 5 shows a positive response of inventory investment in light of a TFP news shock. It peaks at about four quarters before it declines towards zero. This pattern is broadly consistent with the response of the level of inventories documented in Figure 1.

Finally, we include intellectual property products in the VAR to provide suggestive evidence for a possible channel of how news propagates and affects the production process.¹⁶ The third subplot in the figure shows that intellectual property products rise in response to a news shock, commensurate with the behavior of other variables considered so far. This suggests that a key component of a news-driven business cycle model that is consistent with the empirical evidence is the accumulation of knowledge, residing with households as human capital or embodied in physical capital. In the next section we build a theoretical model along these lines.

3 A News Shock Driven Business Cycle Model with Inventories

We now develop a business cycle model that rationalizes the findings of the empirical analysis. The core of the model is the framework of Jaimovich and Rebelo (2009), which includes a particular specification of preferences, investment adjustment costs and costly capacity utilization. This model has become the workhorse model in the news shock literature designed to capture comovement of consumption, investment and hours-worked in response to news about TFP. We augment this model with two additional elements. First, we introduce inventories as in Lubik and Teo (2012), based on the stock-elastic demand model of Bils and Kahn (2000), where finished goods inventories are sales-enhancing.¹⁷ Second, we add intangible capital as an additional input into production. We think of this input as capturing knowledge that evolves over time as a learning-by-doing process. Following Chang et al. (2002) and Cooper and Johri (2002), we assume that households acquire new technological knowledge through their experiences in supplying labor to the production process. This aspect of the model is key to capturing the behavior of inventories to news shocks that we

¹⁶We are not aware of any direct and readily available empirical measure of knowledge capital as interpreted in this paper. We thus provide an indication by capturing some of the effects with this proxy.

¹⁷Our framework thereby abstracts from materials or input inventories that are unquestionably important but constitute the smaller part of total inventories in the data.

see in the data.¹⁸

3.1 Model Environment

The model economy consists of a representative infinitely lived household, a competitive intermediate goods-producing firm, a continuum of monopolistically competitive distributors, and a competitive final goods producer. The intermediate goods firm owns its capital stock and produces a homogeneous good that it sells to distributors. This good is then differentiated by the distributors into distributor-specific varieties that are sold to the final-goods firm. The varieties are aggregated into final output, which then becomes available for consumption or investment. We adopt this particular decentralization since it is convenient for modeling finished goods inventories by separating the production side of the economy into distinct production, distribution, and final goods aggregation phases. Following Chang et al. (2002), we assume that the household accumulates knowledge capital and supplies effective labor to firms as the product of knowledge capital and hours worked. The model economy contains several stochastic shock processes. We include a suite of other shocks in addition to the TFP shocks to facilitate estimation and simulation later in the paper.

3.1.1 Household and Government

The household's lifetime utility is defined over sequences of consumption C_t and hours worked N_t :

$$E_0 \sum_{t=0}^{\infty} \beta^t \Gamma_t \frac{\left(C_t - \psi N_t^\xi F_t\right)^{1-\sigma} - 1}{1-\sigma}, \quad (8)$$

where:

$$F_t = C_t^{\gamma_f} F_{t-1}^{1-\gamma_f} \quad (9)$$

is a preference component that makes consumption and labor non-time-separable and is consistent with the balanced-growth path in a growing economy. This preference structure is based on Jaimovich and Rebelo (2009) and nests the no-income effect structure of

¹⁸The idea of learning-by-doing, and in particular skill-accumulation through work experience, has a long history in labor economics, where empirical researchers have found a significant effect of past work effort on current wage earnings. Both Chang et al. (2002) and Cooper and Johri (2002) study the propagation properties of learning-by-doing in the context of business cycle models. Since then various researchers have exploited these properties to help business cycle models better fit various features of the data. This includes Gunn and Johri (2011), who show how learning-by-doing can yield comovement of consumption, investment, hours worked, and stock prices in response to TFP news. More recently, d'Allesandro et al. (2019) extend a standard New Keynesian model with learning-by-doing to account for the response of various macroeconomic aggregates to a government spending shock.

Greenwood et al. (1988) in the limit as γ_f tends toward zero. Γ_t is a stationary stochastic preference shock process, and $0 < \beta < 1$, $\psi > 0$, $\xi > 1$, $\sigma > 0$, and $0 < \gamma_f \leq 1$.

The household owns the stock of physical capital K_t . Each period, it rents capital services $\tilde{K}_t = u_t K_t$ to the intermediate goods producers at a rental rate r_t , whereby u_t is the utilization rate of the capital. The capital stock evolves according to:

$$K_{t+1} = [1 - \delta(u_t)] K_t + m_t I_t [1 - S(I_t/I_{t-1})], \quad (10)$$

where $\delta(\cdot)$ is a depreciation function that satisfies $\delta'(\cdot) > 0$, $\delta''(\cdot) > 0$ and $\delta(1) = \delta_k$, with $0 < \delta_k < 1$. m_t is a stationary exogenous stochastic process and captures the marginal efficiency of investment. $S(\cdot)$ 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.

We assume that the household accumulates knowledge capital H_t according to:

$$H_{t+1} = H_t^{\gamma_h} N_t^{\nu_h}, \quad (11)$$

where $0 \leq \gamma_h < 1$, and $\nu_h > 0$.¹⁹ It represents the household's state of technological knowledge (or skill level) based on past labor supplies in the vein of the learning-by-doing framework of Chang et al. (2002). The household gains knowledge as it engages with the production process through supplying labor.²⁰ The household's skill level directly affects the effective units of labor supplied to the firms, $\tilde{N}_t = H_t N_t$, for which it earns the wage w_t . This element is the key mechanism that explains the inventory response to a news shock in our framework. It helps suppress the rise in marginal costs during the demand-driven expansion phase of the news boom. This effect of learning-by-doing on inventories is novel within the literature.²¹ Moreover, this particular extension also has a distinct advantage in terms of its parsimony: it adds only an additional input into production and an accumulation equation, while leaving the other elements of the model unaffected. In addition, it nests the more standard model without intangible capital.

¹⁹The log-linear specification used by Chang et al. (2002) and d'Alessandro et al. (2019) is common in the literature.

²⁰In this specification, knowledge capital is stationary due to the stationarity of hours-worked even in a growing economy. This implies that the long-run growth path of output is determined by exogenous technological factors only. This form of knowledge capital can be thought of as an 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.

²¹The general aspect of learning-by-doing as a supply-side mechanism that enhances the dynamics of business cycle models is, of course, not new. While learning-by-doing has a long history in studying long-run issues such as growth, e.g. in Arrow (1962), more recent 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).

The household's budget constraint is given by:

$$C_t + \Upsilon_t I_t + T_t = w_t \tilde{N}_t + r_t u_t K_t + \Pi_t, \quad (12)$$

where Υ_t is a non-stationary exogenous stochastic investment-specific productivity process, T_t denotes lump-sum taxes, and Π_t captures collective profits flowing from firms. We assume that the growth rate of Υ_t , namely $g_t^\Upsilon = \Upsilon_t / \Upsilon_{t-1}$, is stationary. Revenues from taxation go directly to government spending G_t , where we assume that the budget is always balanced such that $G_t = T_t$. Furthermore, government spending follows the process $G_t = \left(1 - \frac{1}{\varepsilon_t}\right) Y_t$, where ε_t is a stationary stochastic government spending shock.

The household chooses sequences of C_t , I_t , N_t , u_t , K_{t+1} and H_{t+1} to maximize intertemporal utility subject to the constraints above. In the following, we only highlight those optimality conditions with a direct impact of knowledge capital, namely optimal choices for N_t and H_{t+1} , since the remainder are standard.²² Respectively, we have:

$$\xi \psi \Gamma_t F_t V_t^{-\sigma} N_t^{\xi-1} = \lambda_t w_t H_t + \nu_h \mu_t^h \frac{H_{t+1}}{N_t}, \quad (13)$$

$$\mu_t^h = \beta E_t \left(\lambda_{t+1} w_{t+1} N_{t+1} + \gamma_h \mu_{t+1}^h \frac{H_{t+2}}{H_{t+1}} \right), \quad (14)$$

where F_t is the utility component defined above and $V_t = C_t - \psi N_t^\xi F_t$ is the periodic utility function to ease notation; λ_t and μ_t^h are the multipliers on the household's budget constraint and the law of motion for knowledge capital.

The presence of knowledge capital adds an additional term into the household's optimality condition for supplying labor, equation (13). This drives a wedge between the marginal utility of leisure and the marginal contribution of hours to earnings, which serves as a shift to the labor supply. All else equal, a rise in the value of knowledge capital μ_t^h increases labor supply as the household desires to increase its knowledge by engaging in production. The optimality condition for H_t , equation (14), then describes the marginal value of knowledge as a function of the expected discounted value of its marginal contribution to wage earnings next period and the continuation value of that knowledge capital. The intertemporal accumulation of knowledge capital makes it worthwhile to increase labor on the arrival of news despite the potential presence of a wealth effect that dominates standard models. Once knowledge capital is in place, the returns are higher than they otherwise would be in the face of higher future productivity. We now turn to the production side of the model to develop this link.

²²We list the full set of optimality conditions in the online appendix.

3.1.2 Intermediate Goods Firm

The competitive intermediate goods firm produces the homogeneous good Y_t using the technology:

$$Y_t = \left(\Omega_t \tilde{N}_t \right)^\alpha \tilde{K}_t^{1-\alpha}, \quad (15)$$

where Ω_t is a non-stationary exogenous stochastic productivity process. We assume that the growth rate of Ω_t , namely $g_t^\Omega = \Omega_t/\Omega_{t-1}$, is stationary. In 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. It then sells its output Y_t at real price τ_t to the distributors.

The firm's profit maximization problem involves choosing \tilde{N}_t and \tilde{K}_t to maximize $\Pi_t^Y = \tau_t Y_t - w_t \tilde{N}_t - r_t \tilde{K}_t$ subject to the production function. This results in standard demand functions for labor and capital services, respectively: $w_t = \alpha \tau_t \frac{Y_t}{\tilde{N}_t}$ and $r_t = (1 - \alpha) \tau_t \frac{Y_t}{\tilde{K}_t}$. Additionally, we find it convenient to define the marginal cost of production for intermediate goods as $mc_t = \frac{w_t}{MP\tilde{N}_t}$, where $MP\tilde{N}_t$ is the marginal product of effective labor. It then follows from the intermediate goods firm's first-order condition 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 varieties S_{it} , $i \in [0, 1]$ according to the technology:

$$S_t = \left[\int_0^1 \nu_{it}^{\frac{1}{\theta}} S_{it}^{\frac{\theta-1}{\theta}} di \right]^{\frac{\theta}{\theta-1}}, \quad \theta > 1, \quad (16)$$

where ν_{it} is a taste shifter that depends on the stock of goods available for sale A_{it} . The latter is composed of current production and the stock of goods held in inventory.²³ We assume that ν_{it} is taken as given by the final goods producer:

$$\nu_{it} = \left(\frac{A_{it}}{A_t} \right)^\zeta, \quad \zeta > 0. \quad (17)$$

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, respectively, the elasticity of substitution between differentiated goods and the elasticity of demand with respect to the relative stock of goods.

²³This structure follows Bilal and Kahn (2000) and is standard in modeling demand for goods drawn from inventories. It also supports a convenient decentralization of production.

The firm acquires each variety i from the distributors at relative price $p_{it} = P_{it}/P_t$, where $P_t = \int_0^1 P_{it}^{\frac{\theta}{\theta-1}} di$ is the aggregate price index. It sells the final good for use in 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 $S_{it}, \forall i$. This results in a demand function for S_{it} for the i th variety:

$$S_{it} = \nu_{it}p_{it}^{-\theta}S_t. \quad (18)$$

An increase in ν_{it} shifts the demand for variety i outwards. This preference shift is influenced by the availability of goods for sale of variety i , relative to aggregate sales, which thereby provides an incentive for firms to maintain inventory to drive customer demand and avoid stockouts.

3.1.4 Distributors

We now close the production side of the model by introducing inventories at the level of the distributors. In the nomenclature of the literature, these are finished goods or output inventories that are ready for sale. Intuitively, they can be thought of as warehouses attached to retail establishments. We follow [Bils and Kahn \(2000\)](#) in modeling inventories as a mechanism that helps generate sales, while at the same time implying a target inventory-sales ratio that captures the idea of stockout avoidance. In addition, this modeling framework creates a wedge between the marginal cost of producing finished goods and the marginal cost of generating a sale, which can come either from inventory stock or new production. It is this margin along which the substitution effect of inventories in response to news shocks operates.

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

$$A_{it} = (1 - \delta_x) X_{it-1} + Y_{it}, \quad (19)$$

where the stock of inventories X_{it} are the goods remaining at the end of the period:

$$X_{it} = A_{it} - S_{it}, \quad (20)$$

and $0 < \delta_x < 1$ is the rate of depreciation of the inventory stock.

The distributors have market power over the sales of their differentiated varieties. The i th distributor sets price p_{it} for sales S_{it} of its variety subject to its demand curve (18).

Each period, a distributor faces the problem of choosing p_{it} , S_{it} , Y_{it} , and A_{it} to maximize profits:

$$E_t \sum_{k=0}^{\infty} \beta^k \frac{\lambda_{t+k}}{\lambda_t} [p_{it+k} S_{it+k} - \tau_t Y_{it+k}], \quad (21)$$

subject to the demand curve (18), the law of motion for goods available for sale (19), and the definition of the inventory stock (20). Profit streams are evaluated at the household's marginal utility of wealth λ_t . Substituting the demand curve for S_{it} , and letting μ_t^a and μ_t^x be the multipliers on the two other constraints, we can then find a representative distributor's first-order conditions:

$$\tau_t = \mu_t^a, \quad (22)$$

$$\mu_t^x = (1 - \delta_x) \beta E_t \frac{\lambda_{t+1}}{\lambda_t} \mu_{t+1}^a, \quad (23)$$

$$\mu_t^a = \zeta p_{it} \frac{S_{it}}{A_{it}} + \mu_t^x \left(1 - \zeta \frac{S_{it}}{A_{it}} \right), \quad (24)$$

$$p_{it} = \frac{\theta}{\theta - 1} \mu_t^x, \quad (25)$$

which are, respectively, the optimal choices of Y_{it} , X_{it} , A_{it} , and P_{it} .

The distributor's optimality conditions allow us to connect the varying marginal costs of production, sales, and inventory holdings in an intuitive manner. The law of motion for A_{it} , equation (19), implies that inventories at the beginning of a period are predetermined. A distributor can only further increase its stock of available goods for sale by acquiring additional output Y_{it} , which has to be purchased at price τ_t . Therefore, the cost of generating an additional unit of A_{it} is equal to the price of output, that is, its marginal cost of production mc_t , as derived from the intermediate firm's profit maximization problem. At the optimum, equation (22) implies that the cost of an additional unit of goods for sale τ_t is equal to the value of those goods for sale, namely μ_t^a .

The inventory definition (20) implies that for a given level of goods available for sale, any increase in sales results in a reduction in stock holdings. The opportunity cost of sales for the distributor is equal to the value of foregone inventory μ_t^x , which can be thought of as the marginal cost of a sale. The optimality condition (23) relates the current value of an additional unit of inventory to the expected discounted value of the extra level of goods available for sale next period generated by holding inventory. This, in turn, equals the price of future output. We can therefore conclude that in this model of inventory holdings the marginal cost of sales is equal to the expected discounted value of next period's marginal cost of output. Increasing sales by drawing down stock in order to forgo production today means that the distributor will need to increase production eventually in the future.

The optimality condition (24) connects the marginal value μ_t^a of a unit of goods available for sale to the value of the extra sales generated by the additional goods available plus the value of the additional inventory yield from the unsold portion of the additional goods. We can combine the marginal cost expressions to derive:

$$\tau_t = \zeta p_{it} \frac{S_{it}}{A_{it}} + (1 - \delta_x) \beta E_t \frac{\lambda_{t+1}}{\lambda_t} \tau_{t+1} \left(1 - \zeta \frac{S_{it}}{A_{it}} \right). \quad (26)$$

This equation implies that the distributor chooses A_{it} , such that the benefit of accumulating goods for sale, either via purchasing new production or stocking inventory, is equal to the marginal cost of output τ_t . We will refer to this equation as the distributor's optimal stocking condition.

Finally, the optimal pricing choice (25) sets the distributor's 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. It follows that the pricing condition is the same as in the standard model. However, the presence of inventories drives a wedge between the marginal costs of output and of sales to the effect that there is no longer a constant markup but one that varies with the value of foregone inventory μ_t^x . Essentially, the optimality condition combines two types of markups: those between marginal costs of output and of sales, and the markup between the marginal cost of sales and price.

The optimal stocking condition (26) describes the adjustment of the first markup through inventories; the optimal pricing condition (25) describes the adjustment of the second markup through price-setting. With flexible prices the latter markup is constant, but the former is not. The total markup between marginal cost of output and price varies as the distributors use inventories to adjust the markup between marginal cost of production and the marginal cost of sales. We can thus combine the distributor's optimality conditions into equations at the aggregate level that reflect the trade-offs faced by inventory accumulation in terms of the various marginal cost concepts:

$$\frac{\theta - 1}{\theta} = (1 - \delta_x) \beta E_t \frac{\lambda_{t+1}}{\lambda_t} \tau_{t+1}, \quad (27)$$

$$\tau_t = \frac{\zeta S_t}{\theta A_t} + \frac{\theta - 1}{\theta}, \quad (28)$$

where we have imposed symmetry on the monopolistic agents' actions. We now turn to a discussion of the stochastic driving processes and calibration of the model before presenting the results from a first quantitative evaluation of the theoretical model.

3.2 Model Solution and Calibration

The model economy contains five stochastic processes: a preference shock Γ_t , a shock to the marginal efficiency of investment m_t , a shock to the growth rate of permanent investment-specific productivity g_t^Υ , a government spending shock ε_t , and a shock to the growth rate of non-stationary productivity g_t^Ω . We assume that these stochastic processes follow individually stationary first-order processes and are mutually uncorrelated. We allow for news shocks to all stochastic processes with the exception of the preference shock. We thus assume that the innovation u_t^j , $j \in \{m_t, g_t^\Upsilon, \varepsilon_t, g_t^\Omega\}$, in a shock process contains both anticipated and unanticipated components. Moreover, news signals arrive with horizons of 4, 8 and 12 quarters as is standard in the literature. The innovations are thus given by:

$$u_t^j = \begin{cases} \epsilon_{jt}^0 + \epsilon_{jt-4}^4 + \epsilon_{jt-8}^8 + \epsilon_{jt-12}^{12}, & j = \{m_t, g_t^\Upsilon, \varepsilon_t, g_t^\Omega\} \\ \epsilon_{jt}^0, & j = \Gamma_t \end{cases}, \quad (29)$$

where ϵ_{jt}^0 is an unanticipated shock, whereas for $h = 4, 8, 12$, ϵ_{jt-h}^p is a news shock that agents receive in period $t-h$ about the innovation in time t . All innovations are mean zero and uncorrelated over time and with each other.

The model economy contains two non-stationary stochastic processes, for productivity Ω_t and for investment-specific productivity Υ_t . In order to find a stationary solution for the model, we express the variables in terms of deviations from their respective stochastic trends. Specifically, we divide non-stationary variables by their permanent component to yield a stationary version of the model. The resulting equation system is then linearized around the steady state of the stationary system and solved using standard methods for linear rational expectations models. The original levels of the trending variables can be recovered by adding the respective stochastic trends back in. The stochastic trend components of output and capital are given by $X_t^y = \Upsilon_t^{-\frac{1-\alpha}{\alpha}} \Omega_t$ and $X_t^k = \Upsilon_t^{-\frac{1}{\alpha}} \Omega_t$, respectively. The stochastic trends 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 show the resulting stationary equilibrium system in the online appendix.

We report the baseline calibration in Table 1. Our choice of parameter values is guided by the existing literature where we strive to maintain comparability with Jaimovich and Rebelo (2009) and Schmitt-Grohé and Uribe (2012) for the aspects of the news shock mechanism, Lubik and Teo (2012) for the inventory component, and Chang et al. (2002) and Gunn and Johri (2011) in terms of knowledge capital. We conduct a robustness analysis for the key parameters underlying our mechanism in section 4.3.

We set the household's discount factor β to 0.9957, which is implied by the real interest

rate computed from average inflation and the federal funds rate over our sample period. The elasticity of intertemporal substitution is as in Jaimovich and Rebelo (2009), $\sigma = 1$. The disutility of working parameter ξ is set to 2, which implies a unit Frisch elasticity of labor supply. This choice places us between the ranges found in Christiano et al. (2005), Jaimovich and Rebelo (2009), and Schmitt-Grohé and Uribe (2012). Finally, we set γ_f , the preference parameter that determines the strength of the income effect, to 0.01 based on Schmitt-Grohé and Uribe (2012).

On the firm side, we set the elasticity parameter in the production function to $\alpha = 0.64$ as in Jaimovich and Rebelo (2009). For the parameters related to physical capital, we fix steady-state physical capital depreciation at $\delta = 0.025$ and the elasticity of marginal utilization $\delta_k''(1)/\delta_k'(1) = 0.15$. There is a wide range of values for this elasticity to be found in the literature. For example, Christiano et al. (2005) find estimates of 0.01, while Schmitt-Grohé and Uribe (2012) have 0.34, and Smets and Wouters (2007) report 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). In our robustness analysis we find that our results are essentially invariant to a wide range of these values. Similarly, the literature also finds a wide range of values for the investment adjustment cost parameter s'' . Smets and Wouters (2007) estimate it to be 5.7, Christiano et al. (2005) find 2.48, and Schmitt-Grohé and Uribe (2012) 9.1. We choose the middle ground in this range and set $s'' = 5$.

The parameters related to inventories are based on the empirical estimates in Lubik and Teo (2012). The inventory depreciation rate δ_x is set to 0.05. The taste shifter curvature ζ is chosen as 0.67 to yield a steady-state sales-to-stock ratio of 0.55, as in Lubik and Teo (2012). The goods aggregator curvature parameter θ is set to 6.8, which results in a steady-state goods markup of 10%. We assume constant returns to scale in the knowledge accumulation equation, setting $\gamma_h = 0.75$, the contribution of prior knowledge capital in its own production, which implies $\nu_h = 0.25$.

Finally, a number of steady-state parameter values are implied by average values in the data, such as the (quarterly) steady-state growth rates of GDP g^y and the relative price of investment (RPI) g^{RPI} , which we find to be 0.43 and -0.58 , respectively (for further discussion and derivation see the online appendix). We also set the steady-state government-spending ratio to output to $g/y = 0.18$ following Smets and Wouters (2007) and target a level of hours in steady state of 0.2, while steady-state capacity utilization is targeted at one. We choose the persistence parameters of the TFP shock process $\rho_\Omega = 0.95$ for the calibration analysis alone. The variances and persistence parameters of all shocks

are later estimated using likelihood-based methods.

4 Model Results

Our analysis focuses on the model’s behavior when subjected to TFP news shocks. We aim to identify the modeling elements needed to understand the empirical facts we uncovered in section 2. We first study the dynamic responses of the key variables to TFP news shocks. In the next step, we disentangle the contribution of the modeling components in generating these outcomes. Finally, we assess the robustness of our baseline calibration to alternative parameter choices. We leave it to section 5 to contrast the simulation findings from the theoretical model with the empirical VAR more formally.

4.1 Response to News Shocks

We first investigate the response of our model economy to a non-stationary TFP news shock, which corresponds conceptually to the identified shock in the empirical VAR analysis. In Figure 6 we report the impulse responses of key model variables to current news about a future increase in TFP that will be realized in 12 quarters as anticipated (solid blue lines). The actual behavior of TFP is depicted in the bottom right hand corner of the panel. When the shock materializes, TFP rises quickly to its new long-run level since the level of TFP is a random walk with drift. All other variables either rise on arrival of the news or increase steadily. Notably, output increases on impact on account of a strong hours and capacity utilization response.

In addition, inventories increase on impact and continue rising through the boom before the actual increase in TFP. Over the adjustment period, until the actual TFP rise occurs, the expansion is supported by a rise in investment and thus capital. With higher TFP in place in period 12, activity continues to expand and eventually overshoots after around 5-6 years when the wealth and income effects take hold. Figure 6 shows that in response to news about a future increase in TFP, inventories rise over time, which is the central finding from the VAR results and substantiated by our theoretical model. Before we demonstrate how the knowledge capital mechanism produces procyclical inventory movements we find it helpful to first discuss how this channel drives an expansion in hours and output.

The value of an additional unit of knowledge capital today, μ_t^h , depends on the additional future wage earnings that knowledge capital yields (see the household’s first-order condition for labor, equation (13)). When news about higher future TFP arrives, the household anticipates that wages will be high in the future relative to today as TFP eventually increases.

This raises the marginal value of having additional knowledge in terms of higher wage earnings in the future and drives up the current value of knowledge capital μ_t^h in a manner that is complementary to the effect of higher TFP and physical capital. The rise in μ_t^h shifts the household's labor supply curve outwards as it seeks to increase its knowledge by supplying additional labor. This, in turn, suppresses the real wage rise, which contributes to an increase in hours and thereby output. In that sense, the mechanism behind the knowledge capital channel is akin to the physical capital channel, whereby the household builds up its knowledge base to take advantage of higher productivity in the future.²⁴

4.2 Understanding Inventory Dynamics

We now turn to a discussion of the behavior of inventories in our model and show how the introduction of knowledge capital into a standard news shock framework is the key element for understanding the comovement we uncovered in the empirical section. The exposition centers on the optimal stocking condition from the distributor's first-order conditions:

$$\tau_t = \frac{\zeta S_t}{\theta A_t} + \frac{\theta - 1}{\theta} = \frac{\zeta}{\theta} \frac{1}{1 + X_t/S_t} + \frac{\theta - 1}{\theta}, \quad (30)$$

which governs inventory dynamics in the model. 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 level of marginal costs τ_t . All else equal, the distributor increases inventory holdings along with a rise in sales, what may be labeled the demand channel, and reduces it along with a rise in current marginal costs, the cost channel.²⁵

We now consider the effects of a TFP news shock on the joint dynamics of inventories and their determinants. We find it convenient to frame the discussion in terms of demand and supply schedules in the market for produced output Y_t with market-clearing price τ_t ,

²⁴We note that the mechanism and crucial modeling elements identified by Jaimovich and Rebelo (2009) are in operation here in addition to the new knowledge capital mechanism. In the former, given the particular form of investment adjustment costs, the shadow value of capital declines today on account of the value of increasing investment today so as to lower future adjustment costs. This, in turn, leads to a reduction in the cost of capacity utilization and as a result an outward shift in labor demand by the intermediate goods firm, whose cost depends inversely on the value of capital, as it increases capacity utilization.

Gunn and Johri (2011) show that the knowledge capital mechanism on its own is sufficient to induce comovement of consumption, investment, and hours in the absence of the Jaimovich and Rebelo (2009) mechanism. In our framework, the low-income effect preferences and variable capacity utilization of Jaimovich and Rebelo (2009) help to enhance the boom, while variable capacity utilization helps suppress the rise in marginal costs.

²⁵The constant term $\frac{\theta-1}{\theta}$ represents the expected value of future marginal costs since $\frac{\theta-1}{\theta} = \beta(1 - \delta_x) E_t \frac{\lambda_{t+1}}{\lambda_t} \tau_{t+1}$. When adjusting inventory holdings, the distributor considers the level of marginal costs today relative to expected future marginal costs, which can be described as an intertemporal substitution channel. Since the former is constant, only variation in the latter impacts inventory. The constancy of expected future marginal costs is an artifact of flexible prices in the current model.

which is also the marginal cost of production. The optimal stocking condition above can thus be thought of as a demand curve for Y_t . We can rewrite it as:

$$\tau_t = \frac{\zeta}{\theta} \frac{S_t}{(1 - \delta_x) X_{t-1} + Y_t} + \frac{\theta - 1}{\theta}, \quad (31)$$

which is downward-sloping in (Y_t, τ_t) -space. All else equal, higher τ_t implies a lower inventory-sales ratio, and thus lower demand for Y_t , as distributors seek to run down inventory stock. Similarly, an increase in sales shifts the curve outward and raises the demand for Y_t as the distributors seek to maintain their sales-inventory ratio by increasing their holdings.

We can combine the household's labor supply conditions, the intermediate firm's labor demand, and the production technology to derive a supply curve for output as a function of τ_t . Abstracting from the (small) income effect ($\gamma_f = 0$) for ease of exposition and normalization of the preference shock Γ_t to unity, this results in:

$$\tau_t = \frac{1}{Y_t} \left(\psi \frac{\xi}{\alpha} Q_t^{-\frac{\xi}{\alpha}} Y_t^{\frac{\xi}{\alpha}} - \frac{\nu_h \mu_t^h}{\alpha \lambda_t} H_{t+1} \right), \quad (32)$$

where $\frac{\partial \tau_t}{\partial Y_t} > 0$ for $\xi > \alpha$, so that the curve is upward-sloping for reasonably elastic labor supply, all else equal. We note that $Q_t = \Omega_t^\alpha \tilde{K}_t^{1-\alpha}$, and $\frac{\mu_t^h}{\lambda_t} H_{t+1} = \beta E_t \frac{\lambda_{t+1}}{\lambda_t} (\alpha \tau_{t+1} Y_{t+1} + \gamma_h \phi_{t+1}^h)$. A rise in the value of knowledge capital μ_t^h shifts the output supply curve outward as the household increases its labor supply in order to acquire more knowledge. This lowers the real wage for a given level of hours and implies a reduction in marginal cost for a given level of output. We depict the supply and demand curves for output Y_t at marginal cost τ_t in Figure 7.

We can now study the response of inventories to TFP news using their impact on supply and demand in the market for produced output. Arrival of positive news about future TFP movements implies a wealth effect that drives up current demand for consumption. In our inventory framework, this also raises the demand for sales of distributors, which shifts their output demand curve (31) outward from D to D' in Figure 7 as they increase their demand for newly produced goods. That is, at given marginal costs, output needs to be higher to be consistent with the TFP news-driven increase in demand as captured by the shifter S_t in this representation.

Alternatively, given supply, the shift in demand puts upward pressure on τ_t , which would imply a lower inventory-sales ratio via the optimal stocking condition. We can see from equation (31) that for a given rise in sales, the extent of the rise in marginal cost determines whether inventories rise or fall. If the rise in marginal costs is large, inventories

must fall in order to reduce the inventory-to-sales ratio enough for equation (31) to still hold, as it becomes more attractive for distributors to draw down stock in the present in order to avoid the high current production costs. On the other hand, if the rise in marginal costs is small, inventories can still rise along with increasing sales, as long as the rise is proportionally less than sales such that the inventory-to-sales ratio still falls and (31) holds.

Therefore, whether inventories rise or fall for a given increase in sales depends on the magnitude of the increase in marginal costs relative to sales. This is determined by the slope of the supply curve and the labor supply elasticity parameter ξ specifically. The slope is decreasing in Q_t , and thereby \tilde{K}_t , such that contemporaneous increases in capacity utilization flatten it, which helps suppress the rise in marginal costs. In the presence of knowledge capital, the rise in its value μ_t^h on arrival of the news mitigates the rise in τ_t as it shifts the output supply curve outward, from S to S' in Figure 7. This allows for a smaller drop in the inventory-sales ratio and thereby supports an increase in inventories along with sales. However, as long as marginal costs increase, a countercyclical inventory-sales ratio, which is consistent with our empirical evidence in section 2.4, is a necessary condition for positive comovement of inventories with other aggregate quantities.

We assess the role of knowledge capital in determining the strength of this mechanism by imposing $\gamma_h = 0$ and $\nu_h = 0$; that is, we shut down the knowledge capital mechanism in the model. This leaves the demand curve unaffected, while the output supply curve reduces to $\tau_t = \frac{\xi}{\alpha} \psi Q_t^{-\frac{\xi}{\alpha}} Y_t^{\frac{\xi}{\alpha}-1}$. The red lines in Figure 6 are the responses of the model economy without knowledge capital to a TFP news shock as discussed above. In response to this shock, inventories now fall over time in advance of the anticipated rise in TFP. Without the shift in the supply curve due to the presence of knowledge capital, marginal costs rise too much, which leads to a larger fall in the inventory-sales ratio and thus a fall in inventories. In Figure 8, we show the responses of sales, inventories, and marginal cost for the specifications with (blue lines) and without (red lines) knowledge capital. The response of inventories depends on the relative response of sales and marginal cost via equation (31). In order to assess the strength of this mechanism, we scale the responses in the model without knowledge capital to have the same peak sales response as in the full model.²⁶

It is notable that sales and inventories rise more in the model with knowledge capital, while marginal costs are at first below those for the version without knowledge capital. Moreover, they reach a lower peak later when the TFP shock is realized. TFP news raises

²⁶Scaling to the same impact response for marginal costs delivers the same result. Since the presence of knowledge capital engenders considerably more propagation (see Figure 6), we scale the responses to focus on and isolate the respective demand and substitution patterns.

sales demand and thereby demand for new production to increase the inventory stock. This effectively shifts output demand along an upward-sloping output supply curve, while the supply curve also shifts right on account of the presence of knowledge capital, but not enough to make marginal cost fall (see Figure 7). This effect is not unlike diminishing returns to labor in the absence of any shift in productivity, which as a result drives up marginal cost in a standard neoclassical production model.

In the model without knowledge capital, we see the reverse of this pattern. Without any rightward shift in the output supply curve, it becomes too costly to satisfy sales demand with new production. The firm therefore runs down its inventory. Consequently, there is less of an increase in demand for new output, that is, less of a rightward shift in new output demand, and as a result, less of a rise in marginal cost. When the TFP shock arrives in period 12, the knowledge capital model has accumulated a stock of this component, which ultimately drives down marginal cost. Moreover, agents still have an incentive to keep increasing their labor supply because the shock is persistent and the value of knowledge remains high on account of its continued benefit in the future. The key to explaining the inventory response to news is therefore the behavior of labor supply engendered by the incentive effects of knowledge capital.

Additionally, the presence of inventories in the Jaimovich-Rebelo model (without knowledge capital) impacts the comovement of other macroeconomic variables such as hours, output, and investment negatively. Despite the increase in labor demand via the standard Jaimovich-Rebelo channels, distributors can reduce their demand for produced goods when compared to the model without inventories. This is possible since they can meet sales demand 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. In addition, investment falls initially until higher TFP is realized, which suggests that, at least in our baseline calibration, the Jaimovich-Rebelo result breaks down in the presence of inventories. However, comovement of investment is restored in our model with knowledge capital.²⁷

Finally, this discussion highlights similarities and differences between our approach and Crouzet and Oh (2016). Consistent with our discussion above, they derive a fairly general condition to show that under realistic calibrations inventories fall in an otherwise standard

²⁷This finding of a fall in physical capital investment in the Jaimovich-Rebelo model without knowledge capital is not general over the entire plausible parameter space. At best, however, the investment response to TFP news is very much muted in the presence of inventories.

Jaimovich-Rebelo model. They demonstrate that this general condition nests the stock-elastic demand model as well as a specification with an explicit stockout avoidance motive. However, they focus on stationary TFP shocks, while we consider the empirically more relevant non-stationary case. Therefore, Crouzet and Oh (2016) derive their identifying restrictions from a different specification so that their empirical results capture responses to a shock that is not directly comparable to the non-stationary TFP shock considered in our analysis above.

4.3 Robustness

We now assess the sensitivity of our central finding to variations in some key parameters. The results are reported in Figures 9-13 which consider robustness to the labor supply parameter ξ , the elasticity of marginal utilization $\delta_k''(1)/\delta_k'(1)$, and the share of labor in knowledge capital ν_h in the specification with and without knowledge capital, whereby we maintain the assumption of constant returns to scale in the accumulation of knowledge capital. All three parameters affect the output supply curve directly and thus drive the response of marginal costs, which we argue above is the key component of the inventory mechanism. As before, we consider a news shock about an anticipated rise in TFP 12 quarters ahead. We report the impulse responses for a wide range of parameter variations in the same graph.

Figures 9-11 show the responses for our benchmark specification. The model appears sensitive to the labor supply elasticity. Changes appear large for somewhat lower values than in the benchmark calibration case of $\xi = 2$, but the positive comovement pattern remains robust. A less elastic labor supply makes the responses less volatile, as can be expected, but is not sufficient for overturning the positive investment response. In contrast, in the corresponding Figure 12 without knowledge capital, inventory declines over the time horizon until the TFP shock materializes for all variations of the labor supply elasticity, while the other aggregate variables increase. This pattern therefore lends strong support to the centrality of the knowledge capital channel in driving positive inventory comovement.

Figures 10 and 13 contain the dynamic responses for variations of the utilization parameter, which has no significant impact on comovement patterns. Finally, Figure 11 reports variations to the labor elasticity in intangible capital, where the baseline calibrated value is $\nu_h = 0.25$. While there is some variation in the extent of the response, the identified comovement patterns remain robust. With a larger value of ν_h , the pattern strengthens, while for lower values it weakens but is largely unchanged. It is only for $\nu_h = 0.1$ that the

inventory response can turn negative over the anticipation horizon.

We conclude that our key finding from the benchmark calibration is invariant to these parameter robustness checks. What explains the across-the-board positive comovement to anticipated TFP shocks is the presence of a knowledge capital channel, which stimulates production on arrival of the news and tends to negate the strong intertemporal substitution effect via marginal costs. Elastic labor supply supports our mechanism, as it does for Jaimovich and Rebelo (2009), but it is not sufficient.

5 Confronting the DSGE Model with the Empirical VAR Evidence

We establish in a structural VAR framework that a positive news shock induces strong positive comovement of aggregate quantities, especially of inventories. This new fact proves to be difficult to explain in standard theoretical models of news shocks, such as Jaimovich and Rebelo (2009) and Crouzet and Oh (2016). We have demonstrated in the section above that a standard model with knowledge capital can generate a positive inventory response alongside an expansion in all other macroeconomic aggregates. We now go a step beyond this analysis and assess the model’s performance somewhat more formally. Specifically, we now allow news to arrive at multiple horizons and let the TFP news shocks compete with other disturbances that have been found relevant in the literature.

We estimate the model using Bayesian techniques, where we retain the structural parameter values of the baseline calibration and only 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 model dynamics as detailed in section 3.2. Our setup of shock processes, treatment of observables, and prior specifications is standard and close to related studies such as Schmitt-Grohé and Uribe (2012) or Khan and Tsoukalas (2012). We estimate the model over the horizon 1983:Q1 - 2018:Q2, which is the same as in the VAR analysis, using GDP, consumption, investment, hours worked, and inventories as observables. Details on the estimation are provided in the online appendix.

Once the model is estimated, we perform a Monte Carlo experiment. We generate 500 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, consistent with the sample length in the empirical VAR analysis. We then compare the empirical responses from the VAR model with the responses estimated

on the artificial data samples under identical VAR specifications.

In order to facilitate comparison between empirical and model-implied TFP news shocks in the VAR, we construct the productivity series based on model variables as:

$$TFP_t = \frac{Y_t}{N_t^\alpha (u_t K_t)^{1-\alpha}} = (\Omega_t H_t)^\alpha. \quad (33)$$

This specification corresponds to the empirical measure for productivity as in Fernald (2012). The latter is adjusted for capital utilization, but given the lack of a precise measure for knowledge capital, it cannot fully account for the fact that variations in this variable impinge on TFP movements and thereby on the identification of news shocks. In the on-line appendix, we provide additional evidence that our empirical findings on the positive, news-driven comovement of all macroeconomic aggregates, including inventories, in section 2 are robust to a potential contamination of productivity by knowledge capital.²⁸

Figure 14 shows impulse response functions at the posterior median (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 areas) from the Monte Carlo experiment. The dynamic responses from the VAR on simulated data are qualitatively in line with the responses from the empirical VAR. Crucially, inventories rise on impact in response to the TFP news shock as do output, investment, consumption, and hours worked. Quantitatively, the empirical and model-implied responses are close as posterior bands overlap for the large majority of periods. Given that the DSGE estimation includes a much larger 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.²⁹

Overall, we find that the responses are qualitatively consistent between the actual and simulated samples. We regard this as strongly suggestive evidence that our framework with

²⁸Fernald’s productivity measure is widely used in the literature and is, despite potential measurement error, arguably the most comprehensive aggregate measure for US productivity. The robustness findings in the appendix address the following concerns as to the use of this measure. First, knowledge capital is a state variable, so that the zero-impact restriction in the VAR is not affected by including this variable in the productivity measure. We show that a Max Share identification without zero-impact restriction delivers almost identical results to our baseline responses. Second, we consider a news shock identification based on patents, suggested by Cascaldi-Garcia and Vucotic (2019) that is independent of Fernald’s productivity measure. Consistent with the results in section 2, we show in the appendix that a news shock under this alternative identification delivers broad comovement of all macroeconomic aggregates and a delayed response of TFP.

²⁹Conceptually, the news shock identification in the VAR and DSGE methodologies is very different, which arguably underlies the observation that the respective responses are quantitatively different. Specifically, the VAR identifies the shock based on the TFP series. In the DSGE model the whole spectrum of auto- and cross-correlations of all observables are used to identify the shock.

knowledge capital can reproduce our new empirical fact, namely that inventories comove alongside the other macroeconomic variables in response to TFP news shocks. This is notwithstanding that our parsimonious framework, which eases the discussion of propagation mechanisms, limits 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.³⁰

6 Conclusion

Our paper makes two contributions to the literatures on news shocks and inventory dynamics. First, we establish empirically that a news shock in terms of an anticipated rise in TFP in the future raises inventory holdings in the present and induces positive comovement with other macroeconomic aggregates. Based on standard VAR identification, this fact is robust across many dimensions, such as sectors, types of inventories, and alternative identification schemes for news shocks. Our empirical finding corroborates the view that TFP news shocks are important drivers of macroeconomic fluctuations. We also consider this an important result as it provides a dimension along which standard inventory frameworks can be evaluated as to their empirical viability. This is where our second contribution lies.

We show that the standard theoretical framework used in the news shock literature cannot explain procyclical inventory movements in response to TFP news shocks. We argue that an additional mechanism, namely the accumulation of knowledge capital, is needed to capture the behavior of inventories. This mechanism addresses two shortcomings of previous frameworks. First, they fail to reproduce the procyclical inventory movements in response to TFP news shocks due to a strong substitution effect that moves production into the future. Second, introducing inventories in standard frameworks implies an intertemporal labor choice that makes even comovement of consumption, investment, and hours much harder to achieve. Knowledge capital provides an incentive for firms and workers to engage in production today to accumulate the know-how needed for taking full advantage of higher TFP in the future. This leads to inventory accumulation in the present.

Even though inventories are strongly procyclical unconditionally, conditional on TFP news shocks, our empirical finding is not a priori self-evident. Conventional views would suggest two potential counteracting effects on inventories in response to news. A negative substitution effect provides incentives to reduce the current inventory stock and increase

³⁰For 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).

stockholding in the future when the higher productivity is actually realized. We provide evidence in Görtz et al. (2019) 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. This finding is based on firm-level data and supports the insights from the aggregate data in the current paper. In addition, our theoretical insights provide a new transmission channel for news shocks to the literature. A rigorous investigation of the data-generating mechanism, including multiple sectors and the use of input as well as finished goods inventories, goes beyond the scope of this paper and is left for future research.

References

- [1] Arrow, Kenneth J. (1962): “The economic implications of learning by doing”. *Review of Economic Studies*, 29(3), pp. 155-173.
- [2] Barsky, Robert B., and Eric R. Sims (2011): “News shocks and business cycles”. *Journal of Monetary Economics*, 58(3), pp. 273-289.
- [3] Barsky, Robert B., and Eric R. Sims (2012): “Information, animal spirits, and the meaning of innovations in consumer confidence”. *American Economic Review*, 102(4), pp. 1343-77.
- [4] Beaudry, Paul, and Franck Portier (2004): “An exploration into Pigou’s theory of cycles”. *Journal of Monetary Economics*, 51(6), pp. 1183-1216.
- [5] Beaudry, Paul, and Franck Portier (2014): “News driven business cycles: Insights and challenges”. *Journal of Economic Literature*, 52(4), pp. 993-1074.
- [6] Bils, Mark, and James A. Kahn (2000): “What inventory behavior tells us about business cycles”. *American Economic Review*, 90(3), pp. 458-481.
- [7] Cascaldi-Garcia, Danilo, and Marija Vokotić (2019): “Patent-based news shocks”. Forthcoming, *Review of Economics and Statistics*.
- [8] Chang, Yongsung, Joao F. Gomes, and Frank Schorfheide (2002): “Learning-by-doing as a propagation mechanism”. *American Economic Review*, 92(5), pp. 1498-1520.
- [9] Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans (2005): “Nominal rigidities and the dynamic effects of a shock to monetary policy”. *Journal of Political Economy*, 113(1), pp. 1-45.

- [10] Cooper, Russell, and Alok Johri (2002): “Learning-by-doing and aggregate fluctuations”. *Journal of Monetary Economics*, 49(8), pp. 1539-1566.
- [11] Crouzet, Nicolas, and Hyunseung Oh (2016): “What do inventories tell us about news-driven business cycles?” *Journal of Monetary Economics*, 79, pp. 49-66.
- [12] d’Alessandro, Antonello, Giulio Fella, and Leonardo Melosi (2019): “Fiscal stimulus with learning by doing”. *International Economic Review*, 60(3), pp. 1413-1432.
- [13] Fernald, John (2012): “A quarterly, utilization adjusted series on total factor productivity”. Federal Reserve Bank of San Francisco Working Paper Series 2012-19.
- [14] Forni, Mario, Luca Gambetti, and Luca Sala (2014): “No news in business cycles”. *Economic Journal*, 124, pp. 1168-1191.
- [15] Francis, Neville, Michael Owyang, Jennifer Roush, and Riccardo DiCeccio (2014): “A flexible finite-horizon alternative to long-run restrictions with an application to technology shocks”. *Review of Economics and Statistics*, 96, pp. 638-647.
- [16] Galí, Jordi, and Luca Gambetti (2009): “On the sources of the Great Moderation”. *American Economic Journal: Macroeconomics*, 1, pp. 26-57.
- [17] Görtz, Christoph, and John Tsoukalas (2017): “News and financial intermediation in aggregate fluctuations”. *Review of Economics and Statistics*, 99(3), pp. 514-530.
- [18] Görtz, Christoph, Christopher Gunn, and Thomas A. Lubik (2019): “What drives inventory accumulation? News on rates of return and marginal costs”. Federal Reserve Bank of Richmond Working Paper No. 19-18.
- [19] Görtz, Christoph, John Tsoukalas, and Francesco Zanetti (2017): “News shocks under financial frictions”. Technical Report.
- [20] Greenwood, Jeremy, Zvi Hercowitz, and Gregory Huffman (1988): “Investment, capacity utilization, and the Real Business Cycle”. *American Economic Review*, 78, pp. 402-217.
- [21] Gunn, Christopher M., and Alok Johri (2011): “News and knowledge capital”. *Review of Economic Dynamics*, 14(1), pp. 92-101.
- [22] Jaimovich, Nir, and Sergio Rebelo (2009): “Can news about the future drive the business cycle?” *American Economic Review*, 99(4), pp. 1097-1118.

- [23] Jung, YongSeung, and Tack Yun (2013): “Inventory investment and the empirical Phillips curve”. *Journal of Money, Credit and Banking*, 45(1), pp. 201-231.
- [24] Khan, Hashmat, and John Tsoukalas (2012): “The quantitative importance of news shocks in estimated DSGE models”. *Journal of Money, Credit and Banking*, 44(8), pp. 1535-1561.
- [25] Kurmann, Andre, and Eric Sims (2019): “Revisions in utilization-adjusted TFP and robust identification of news shocks”. Forthcoming, *Review of Economics and Statistics*.
- [26] Lubik, Thomas A., and Wing Leong Teo (2012): “Inventories, inflation dynamics and the New Keynesian Phillips curve”. *European Economic Review*, 56(3), pp. 327-346.
- [27] McCarthy, Jonathan, and Egon Zakrajsek (2007): “Inventory dynamics and business cycles: What has changed?” *Journal of Money, Credit and Banking*, 39(2-3), pp. 591-613.
- [28] Ramey, Valerie A., and Kenneth D. West (1999): “Inventories”. In: John B. Taylor and Michael Woodford (eds.): *Handbook of Macroeconomics*, Vol. 1, Chapter 13, pp. 863-923. Elsevier.
- [29] Sarte, Pierre-Daniel G., Felipe F. Schwartzman, and Thomas A. Lubik (2015): “What inventory behavior tells us about how business cycles have changed”. *Journal of Monetary Economics*, 76, pp. 264-283.
- [30] Schmitt-Grohé, Stephanie and Martín Uribe (2012): “What’s news in business cycles?” *Econometrica*, 80(6), pp. 2733-2764.
- [31] Smets, Frank, and Rafael Wouters (2007): “Shocks and frictions in US business cycles: A Bayesian DSGE approach”. *American Economic Review*, 97(3), pp. 586-606.
- [32] Wen, Yi (2005): “Understanding the inventory cycle”. *Journal of Monetary Economics*, 52(8), pp. 1533-1555.

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
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	0.50
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




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.




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.

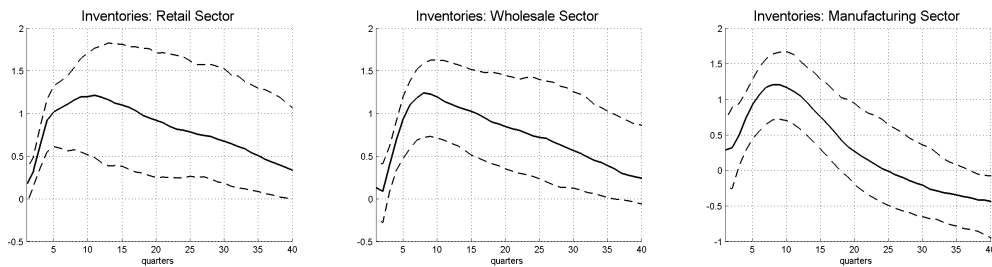


Figure 3: **IRF of business inventories by sector to TFP news shock.** Sample 1992Q1-2018Q2. Subplots result from eight variable VARs comprising TFP, GDP, consumption, investment, hours, inventory measure, inflation, E5Y. The inventory measures were included one-by-one in the VAR system. The solid line 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.

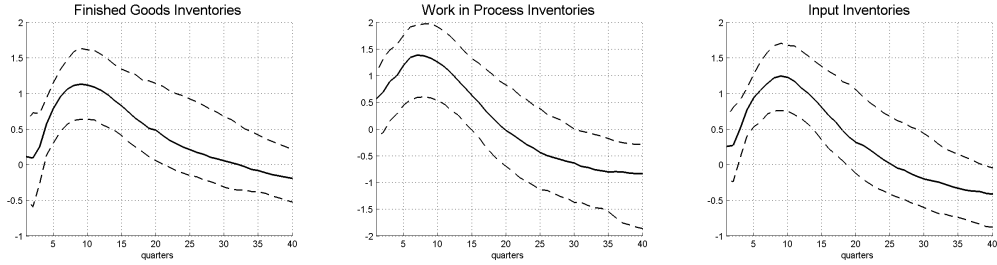


Figure 4: **IRF of business inventories in the manufacturing sector by inventory type to TFP news shock.** Sample 1992Q1-2018Q2. Subplots result from eight variable VARs comprising TFP, GDP, consumption, investment, hours, inventory measure, inflation, E5Y. The inventory measures were included one-by-one 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.




Figure 5: **IRF to TFP news shock.** Subplots result from VARs comprising TFP, GDP, investment, hours, inflation 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.

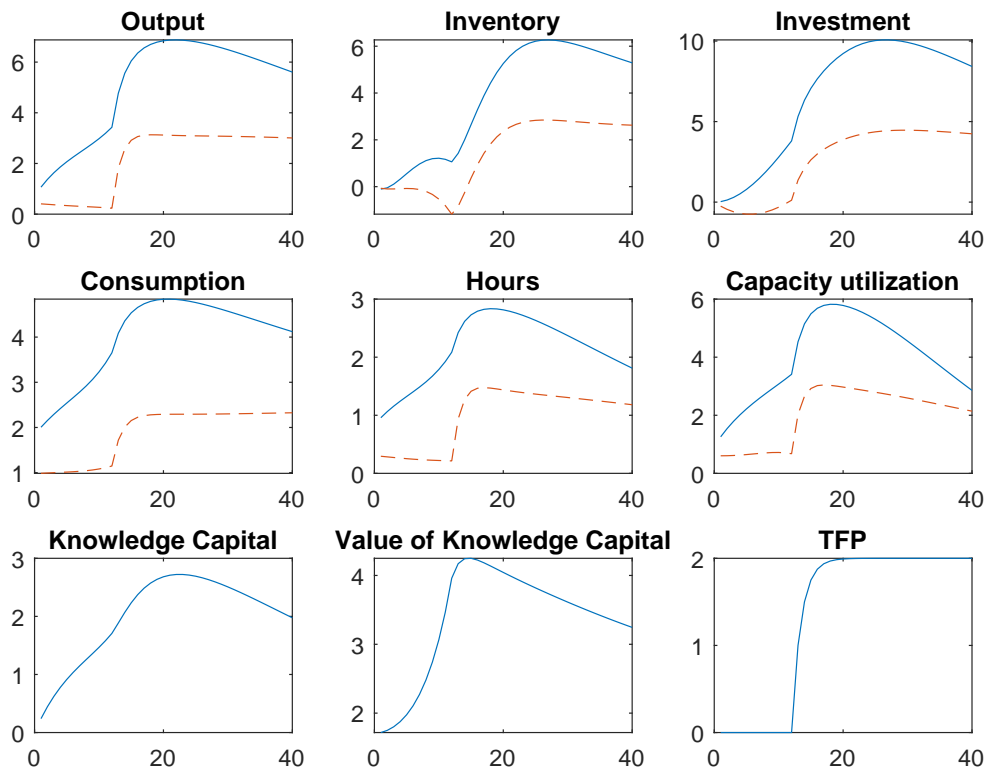


Figure 6: **IRF to 12 period ahead unit TFP news shock.** Baseline model (solid-blue) and model without knowledge capital (dashed-red).




Figure 7: Supply and Demand curves for output, Y_t , and marginal cost τ_t .




Figure 8: **IRF to 12 period ahead TFP news shock.** Baseline model (solid blue) and model without knowledge capital (dashed red). The responses of the model without knowledge capital are scaled so that the maximum impact of sales corresponds to the one in the baseline model.

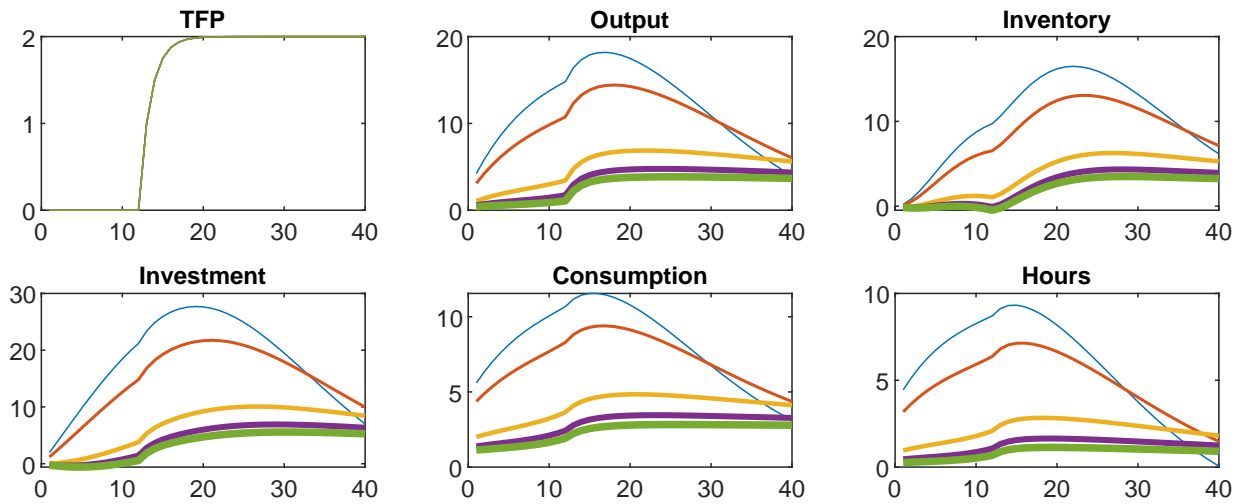


Figure 9: **IRF sensitivity for 12 period ahead TFP shock.** Baseline model. $\xi = \{1.4, 1.5, 2, 2.5, 3\}$ (thin to thick lines).

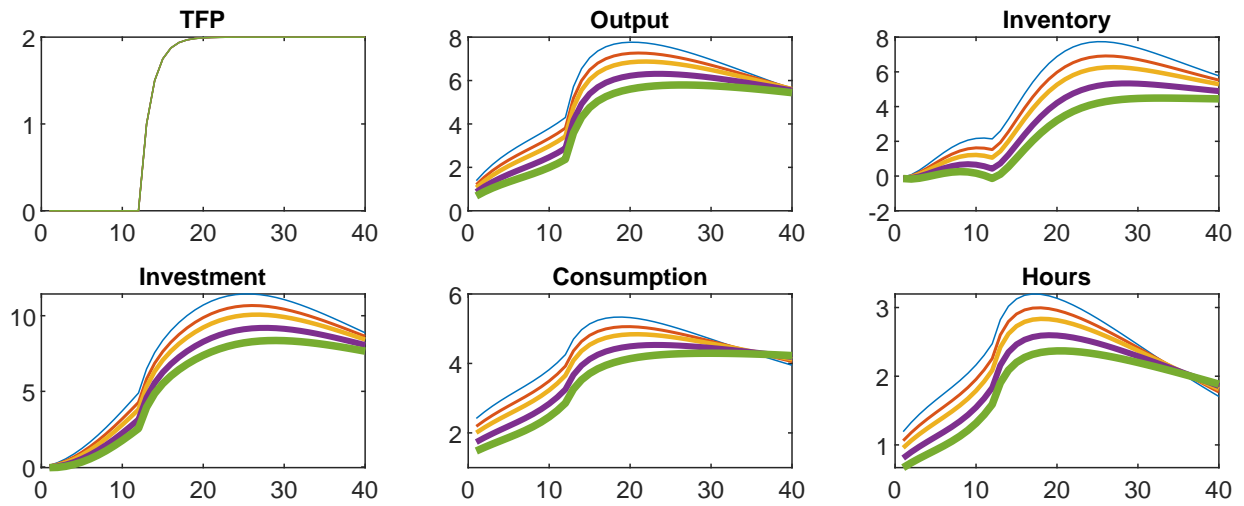


Figure 10: **IRF sensitivity for 12 period ahead TFP shock.** Baseline model. $\delta_k''(1)/\delta_k'(1) = \{0.05, 0.1, 0.15, 0.25, 0.4\}$ (thin to thick lines).




Figure 11: **IRF sensitivity for 12 period ahead TFP shock.** Baseline model. $\nu = \{0.1, 0.2, 0.25, 0.3, 0.35\}$ (thin to thick lines).




Figure 12: **IRF sensitivity for 12 period ahead TFP shock.** Model *without* knowledge capital. $\xi = \{1.4, 1.5, 2, 2.5, 3\}$ (thin to thick lines).

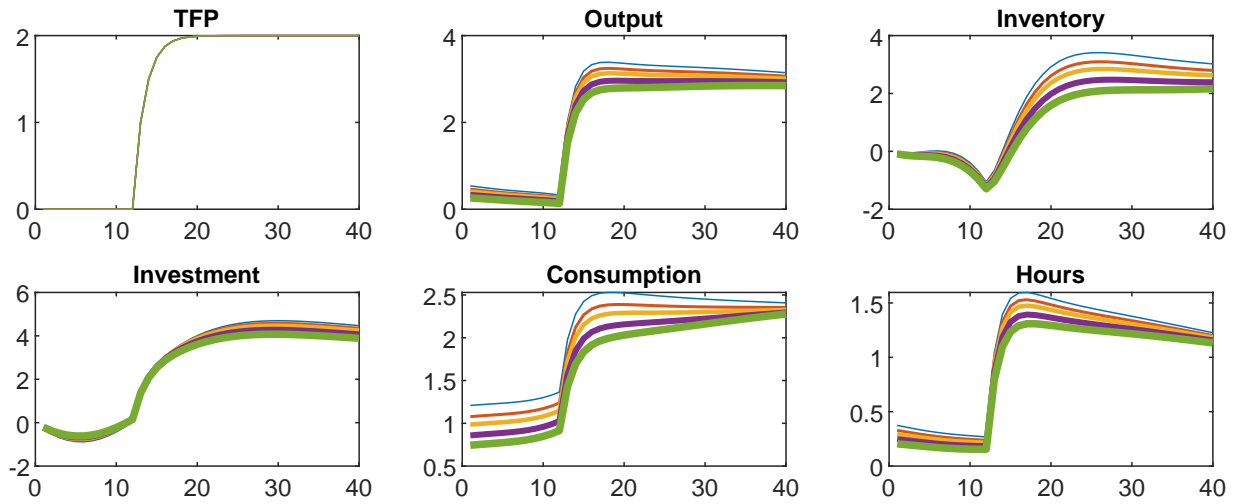


Figure 13: **IRF sensitivity for 12 period ahead TFP shock.** Model *without* knowledge capital. $\delta_k''(1)/\delta_k'(1) = \{0.05, 0.1, 0.15, 0.25, 0.4\}$ (thin to thick lines).

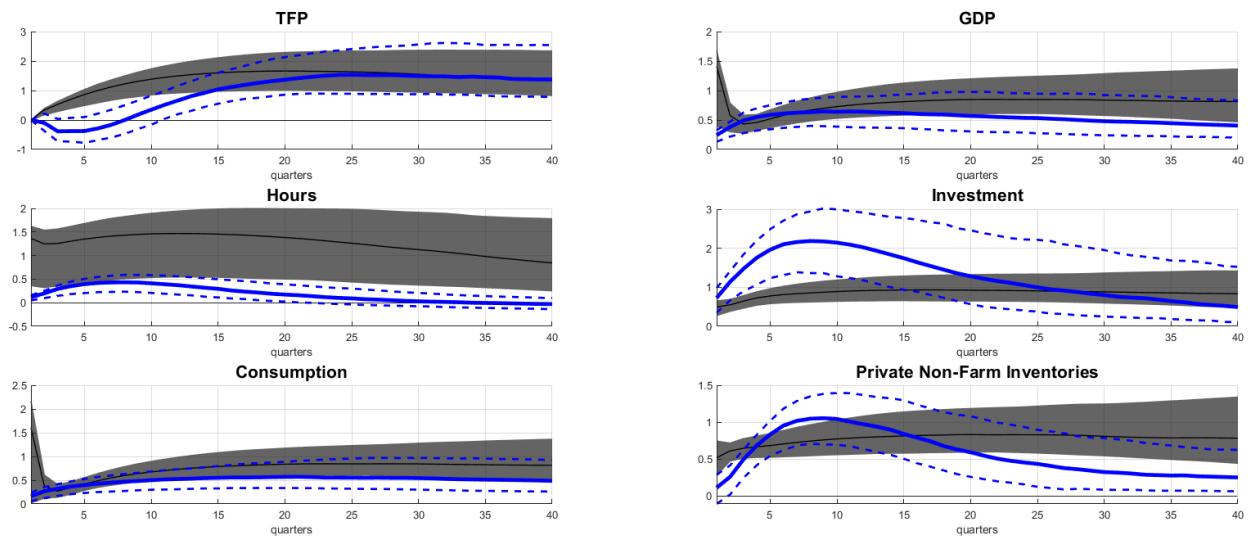


Figure 14: **TFP news shock.** The blue solid (blue dashed) line is the median (16% and 84% posterior band) response to a TFP news shock from a six-variable VAR. The solid black line (gray shaded areas) is the median (16% and 84% posterior band) response to a TFP news shock estimated from a VAR on 500 samples generated from the DSGE model. Units of the vertical axes are percentage deviations.