

CEP 18-05

Taking Stock of TFP News Shocks: The Inventory Comovement Puzzle

Christoph Görtz
University of Birmingham

Christopher Gunn
Carleton University

Thomas Lubik
Federal Reserve Bank of
Richmond

June 12th, 2018 (revised July 14th, 2018)

CARLETON ECONOMIC PAPERS



Department of Economics

1125 Colonel By Drive
Ottawa, Ontario, Canada
K1S 5B6

Taking Stock of TFP News Shocks: The Inventory Comovement Puzzle*

Christoph Görtz Christopher Gunn
University of Birmingham Carleton University

Thomas Lubik
Federal Reserve Bank of Richmond

July 2018

Abstract

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

Keywords: News shocks, Business cycles, Inventories.

JEL Classification: E2, E3.

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

1 Introduction

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

factor of production in the durables sector where they play a buffer stock role similar to fixed capital investment. While she does not explicitly discuss any empirical responses of inventories, she shows the inventory to sales ratio in durables sectors moves particularly strongly countercyclical in response to TFP news shocks, which is consistent with our findings for the whole economy.

The remainder of the paper is structured as follows. Sections 2.1 and 2.2 discuss the VAR identification strategy and the data used in the empirical analysis. Section 2.3 presents our main empirical findings. We corroborate these and study the transmission mechanism in sections 2.4 and 2.5. Sections 3 and 4 present a standard model with inventories and investigate the mechanisms underlying the comovement puzzle. Section 5 presents an extension of the standard model as an example of a model that addresses the inventory comovement puzzle. Section 6 concludes.

2 VAR Analysis

2.1 The VAR model

This section discusses the VAR model. The identification of a TFP news shock based on this model is discussed in the following section. Consider the following reduced form VAR(p) model,

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

where y_t is an $n \times 1$ vector of variables of interest, $A(L) = I + A_1L + A_2L^2 + \dots + A_pL^p$ is a lag polynomial, A_1, A_2, \dots, A_p are $n \times n$ matrices of coefficients and, finally, u_t is an error term with $n \times n$ covariance matrix Σ . Define a linear mapping between reduced form, u_t , and structural errors, ε_t ,

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

We can then write the structural moving average representation as

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

where $C(L) = A(L)B_0$, $\varepsilon_t = B_0^{-1}u_t$, and the matrix B_0 satisfies $B_0B_0' = \Sigma$. The B_0 matrix may 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 ($DD' = I$).

The h step ahead forecast error is,

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

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 selection vectors with one in the i -th position and zeros elsewhere. The e_j vectors pick out the j -th column of D , denoted by γ . $\tilde{B}_0\gamma$ is an $n \times 1$ vector corresponding to the j -th column of a possible orthogonalization and can be interpreted as an impulse response vector. In the following section, we discuss the estimation and identification methodology that yields an estimate for the TFP news shock from the VAR model.

2.2 Data and VAR estimation

We estimate the VAR using quarterly U.S. data for the period 1985:Q1–2015:Q2. This sample horizon is guided by the literature that documents differences in cross correlation patterns of several macro-aggregates in samples before and after the mid-1980s (e.g. Galí and Gambetti (2009)). Furthermore, McCarthy and Zakrajsek (2007) document significant changes in inventory dynamics occur in the mid-1980s due to improvements in inventory management. To identify the TFP news shock from the VAR model, we adopt the Max Share identification method (based on Francis et al. (2014)). It assumes that (i) the news shock does not move TFP on impact and (ii) maximizes the variance of TFP at a specific long but finite horizon (we set the horizon to 40 quarters). The time series included in the

VAR enter in levels, consistent with the treatment in the empirical VAR literature (see e.g. Barsky and Sims (2011), Beaudry and Portier (2004) and Beaudry and Portier (2014)). To estimate the VAR model we use three lags and a Minnesota prior. Confidence bands are computed by drawing from the posterior.

We consider two different measures for total inventories in the VAR. First, non-farm private inventories, which are defined as the physical volume of inventories owned by private non-farm business, valued at average prices of the period (the replacement costs of inventories). The second measure, business inventories, differs from the former as stockholdings are valued by the cost at acquisition of inventories that can differ from their price when sold. In the NIPAs, inventory profits and losses resulting from differences between acquisition and sales price are shown as adjustments to business income. Business inventories are only available from 1992Q1 which is why we reduce the sample horizon if these are included in the VAR. Output is defined as GDP and total hours as hours worked of all persons in the non-farm business sector. Investment is the sum of fixed investment and personal consumption expenditures for durable goods. Fixed investment is the component of gross private domestic investment that excludes changes in private inventories. Consumption is the sum of personal consumption expenditures for non-durable goods and services. All these time series are seasonally adjusted and in real per-capita terms (except for hours which are not deflated). Inflation is constructed using the GDP deflator. A measure of technology is key to identify the news shock. We follow the convention in the empirical literature and use the measure of utilization-adjusted TFP provided by Fernald (2014).⁷ We also use the Michigan consumer confidence indicator (E5Y) in our VAR system.⁸ The set of variables included in our VAR system is, apart from inventories, standard in the literature and considering the E5Y consumer confidence measure is a way to provide forward looking information that

⁷We use the 2015 vintage which contains updated corrections on utilization from industry data.

⁸The Michigan consumer confidence indicator summarizes responses to the following question: “Turning to economic conditions in the country as a whole, do you expect that over the next five years we will have mostly good times, or periods of widespread unemployment and depression, or what?” The variable is constructed as the percentage giving a favorable answer minus the percentage giving an unfavorable answer plus 100.

captures expectations.⁹

2.3 VAR results

Figure 1 shows impulse response functions to a TFP news shock from an eight-variable VAR. It is striking that all activity variables are increasing prior to a significant rise in TFP. The comovement between output, consumption, investment and hours over this post Great Moderation sample has been documented in existing work (e.g. Görtz et al. (2017)). The fact we add here is to document the increase in the stock of private non-farm inventories prior to a rise in TFP. The hump-shaped increase in the stock of inventories indicates that inventory investment is positive until about quarter 12, shortly before the higher productivity is actually realized. Additionally, we report a short-lived decline in inflation and an anticipation of the increase in TFP in the consumer confidence indicator E5Y, both consistent with findings in previous work. Barsky and Sims (2012) highlight that the inflation response is broadly consistent with the New Keynesian framework in which current inflation equals an expected present discounted value of future marginal costs. The significant increase in the E5Y is indicative of an increase in consumer confidence upon the arrival of news about higher TFP (see e.g. Barsky and Sims (2011) or Görtz and Tsoukalas (2018)). The TFP news shock is important for fluctuations in inventories and GDP as we find it to explain between 43-59% (44-66%) of the forecast error variance in inventories (GDP) over a horizon from 6-32 quarters. Details about the forecast error variance decomposition are provided in Appendix A.1.

Figure 2 shows that the rise in inventories prior to TFP is also robust when we use total business inventories as an alternative measure to private non-farm inventories. Evaluating the response of inventories to a TFP news shock also with this alternative measure is important as it is not a priori clear at which prices inventories should be measured. However, this measure is only available from 1992Q1 which restricts the sample for this VAR system.¹⁰ All

⁹See e.g. Barsky and Sims (2012). The S&P500 stock price index has also been considered for this purpose. Our results are robust to including the S&P500 instead of the E5Y.

¹⁰Note that data availability limits all VAR systems that include total business inventories or its subcom-

variables in Figure 2 show very similar qualitative responses to the ones in Figure 1, albeit the shorter sample results in somewhat wider confidence bands. Overall, this figure confirms the comovement of macroeconomic aggregates, including inventories, prior to the significant rise in TFP.

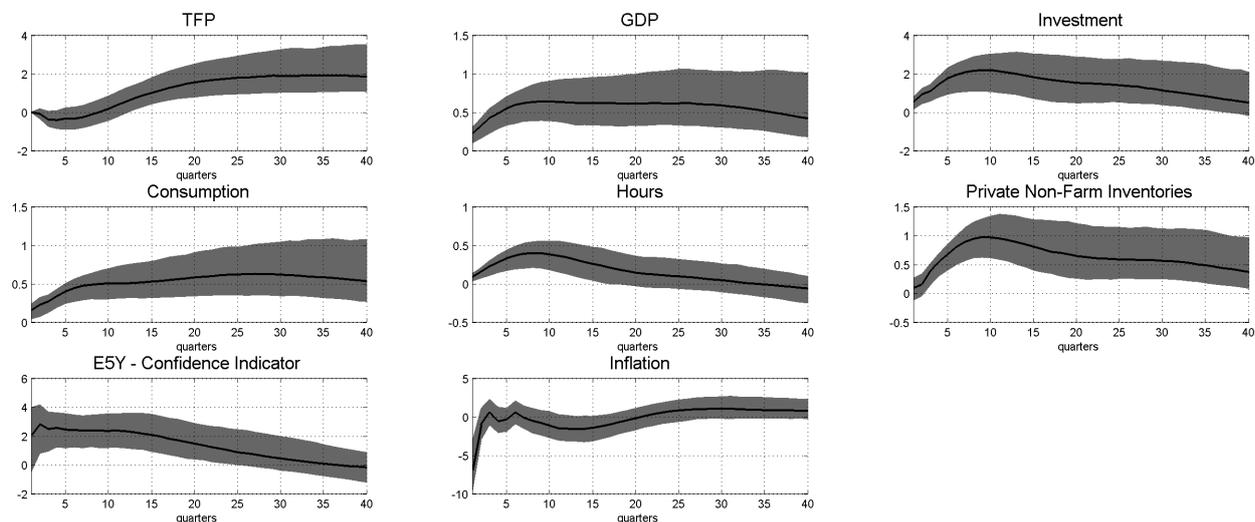


Figure 1: **IRF to TFP news shock – including Private Non-Farm Inventories.** Sample 1985Q1-2015Q2. The shaded gray areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

The vast majority of inventories are held in the manufacturing, wholesale and retail sectors (see e.g. Ramey and West (1999)). Figure 3 shows the responses of business inventories in each of these sectors to the (aggregate) TFP news shock when we alternate between including one of the three separate sectoral measures of inventory in our eight-variable VAR. It is evident that the expansion of the inventory stock is broad-based across the manufacturing, wholesale and retail sector. The two trade sectors almost entirely hold finished goods inventories (see e.g. Blinder and Maccini (1991)), while over our 1992Q2-2015Q2 sample the inventory stock held in the manufacturing sector is split across finished goods inventories (36%), work in process (30%) and input inventories in form of materials and supplies (34%).¹¹ Figure 4 shows the responses of inventory types in the manufacturing sector when

ponents to start in 1992Q1.

¹¹For the wholesale and retail sector, time series data that break the total stock down into inventory types is not available.

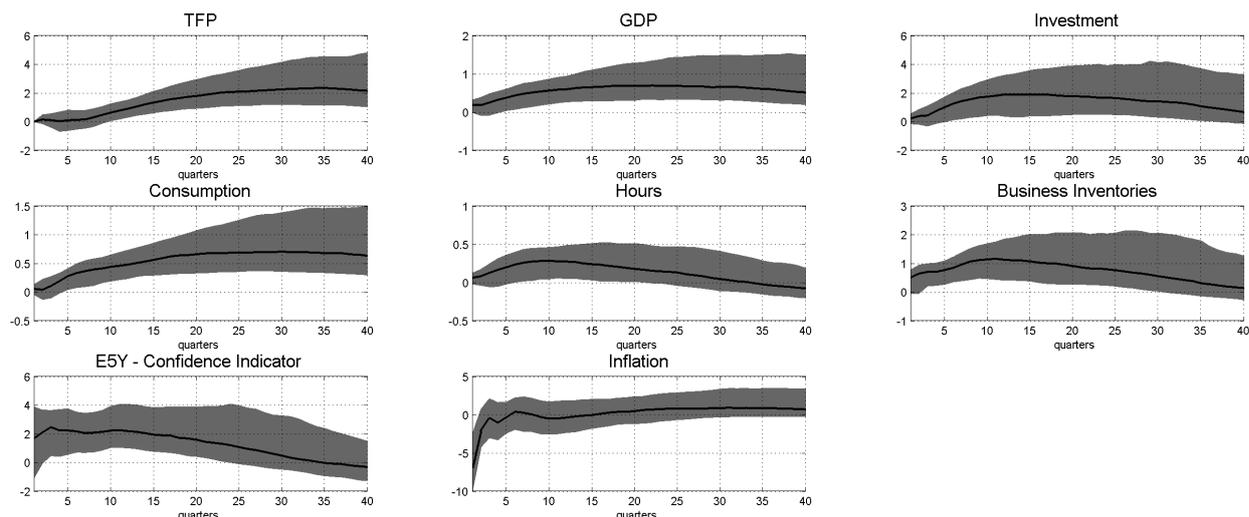


Figure 2: **IRF to TFP news shock – including Business Inventories.** Sample 1992Q1-2015Q2. The shaded gray areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

we include these one-by-one in our eight-variable VAR. In response to a TFP news shock, finished goods and input inventories in the manufacturing sector rise strongly before TFP increases significantly after about 10 quarters. The increase in work in process inventories is only marginally significant, but the median also shows the positive hump shaped behavior evident in the other subplots.¹² The strong response of inventories to news shocks is consistent with results in Kesavan et al. (2010) who find inventories to be a forward looking-variable closely linked to future expectations of economic conditions as they help to improve forecasts about sales.

2.4 The forces behind the inventory accumulation

The discussion above shows the increase in the inventory stock in anticipation of higher future TFP is a robust fact across different measures for total inventories, different sectors and different types of inventories. Two potentially counteracting effects may be at play

¹²The responses of the remaining seven variables in the VAR that we are not showing in Figures 3 and 4 are very similar to the ones reported in Figure 2 and are available upon request. We focus our discussion on sectoral data for business inventories as private non-farm inventories in the manufacturing, wholesale and retail sectors is available only from 1996Q4.

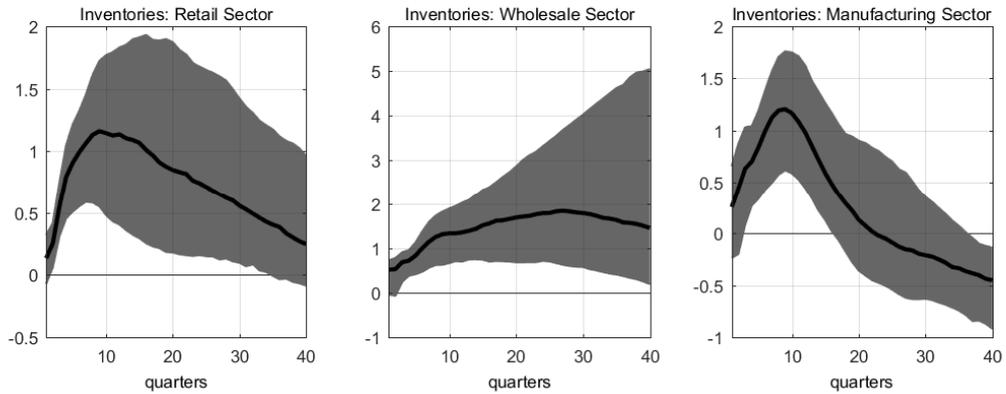


Figure 3: **IRF of business inventories by sector to TFP news shock.** Sample 1992Q1-2015Q2. Subplots result from eight variable VARs comprising TFP, GDP, consumption, investment, hours, inventory measure, inflation, E5Y. The inventory measures were included one-by-one in the VAR system. The shaded gray areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

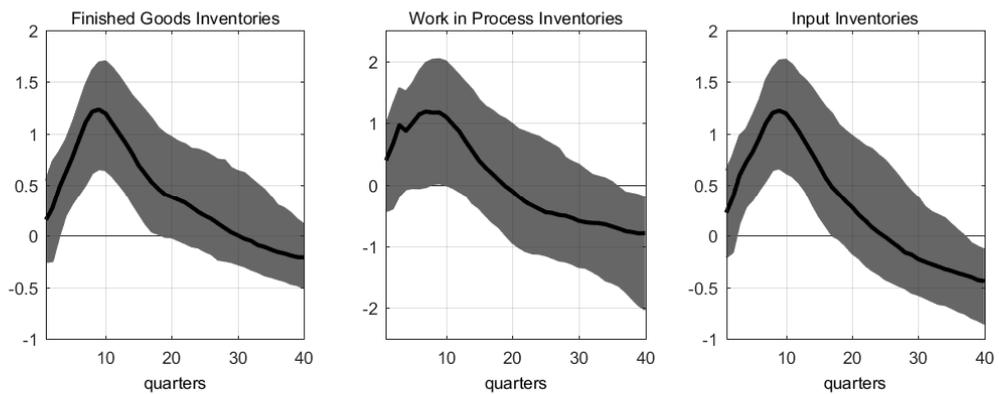


Figure 4: **IRF of business inventories in the manufacturing sector by inventory type to TFP news shock.** Sample 1992Q1-2015Q2. Subplots result from eight variable VARs comprising TFP, GDP, consumption, investment, hours, inventory measure, inflation, E5Y. The inventory measures were included one-by-one in the VAR system. The shaded gray areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

in response to news about higher future TFP. Firms may increase inventories in light of rising consumption and investment to ensure a relatively stable inventory to sales ratio and enhance demand. This demand effect, where inventories are productive for sales, has for example been documented in [Bils and Kahn \(2000\)](#). In addition to this positive demand effect, a negative substitution effect may provide incentives to run the current inventory stock down and increase stockholding in future when the higher productivity is actually realized. The fact that inventories rise in response to a TFP news shock points towards dominance of the positive demand effect over the negative substitution effect.

In the following subsections we will shed light on the transmission channel by providing evidence for the importance of the demand and substitution effects. [Jones and Tuzel \(2013\)](#) highlight an important role for the risk premium — measured e.g. by the implied cost of capital (ICC) or debt and equity cost of capital — as opportunity cost for holding inventories. Such cost of capital indicators reflect the risk of holding inventories, e.g. as a result of input inventories taking time to be transformed into final products, or finished goods inventories being subject to uncertainty about demand. A decline in the risk premium, and hence the opportunity cost of holding inventories, at the time the news about higher future TFP arrives, would be supportive of a strong demand effect and the cost of capital indicators can provide an indication about the strength of this effect. We discuss the responses of measures for debt and equity cost of capital constructed from aggregate data in [section 2.4.1](#) and the responses of ICC measures constructed from firm level data in [section 2.4.2](#). Behind the substitution effect is the assumption that the anticipation of higher future TFP implies a decline in future marginal cost and hence an incentive to substitute production from today into the future by drawing down the inventory stock. In [section 2.4.3](#) we consider the response of measures for marginal costs to a TFP news shock to evaluate whether there is evidence for such a negative substitution effect.

2.4.1 Debt and Equity Cost of Capital

[Jones and Tuzel \(2013\)](#) show that inventory investment is negatively related to debt and

equity cost of capital constructed from standard regressions of excess return on predictive variables. We construct these measures of risk premiums — the excess return on either portfolios of stocks or on bonds — in line with the methodology of Jones and Tuzel (2013). In particular, we use as dependent variable in these regressions either the return on the US stock market minus the one-month Treasury bill return (RMRF) or the return on corporate bonds minus the one-month Treasury bill return (RBRF). We include seven independent variables that have been found to be relevant in previous work as predictors for excess bond and stock returns and have also been used by Jones and Tuzel (2013). These are, the term spread (TERM), the default spread (DEF), the dividend yield (DP), the ratio of new orders to shipments of durable goods (NOS), and the consumption-wealth ratio (CAY) of Lettau and Ludvigson (2001) as well as the real return on a nominally riskless asset (RF) and the four-quarter moving average of this variable (RF4).¹³ All seven independent variables potentially enter with one lag and we select those predictors that minimize the Akaike Information Criterion (AIC). For the regression on excess stock market returns, RMRF, this criterion implies to include DP, which has a coefficient of 1.76*, and the intercept is -0.02. For the excess corporate bond return, RBRF, the regression includes TERM (3.5931***), RRF4 (1.1270***), DP (0.6617***), CAY (0.2527***) and the intercept (0.0433***) where coefficients are given in parenthesis.¹⁴ The fitted values of these regressions are the equity cost of capital and debt cost of capital, respectively.

We separately add the equity and debt cost of capital measures in an eight-variable VAR system as proxies for the opportunity costs of holding inventories. Figure 5 shows IRFs of

¹³The term spread is the difference between the 10 year and 3 months Treasury yields from the Federal Reserve’s H15 database. The default spread is Moody’s Seasoned Baa Corporate bond yield relative to the yield on a 10-Year Treasury constant maturity from FRED. The dividend yield is computed, using data from Robert Shiller’s website, as the quarterly average of past Standard & Poor’s (S&P) composite dividends divided by the end-of-quarter level of the S&P composite index. The ratio of new orders to shipments is provided by Jones and Tuzel (2013). The real return on a riskless asset is calculated as the one-month Treasury bill return from Kenneth French’s website minus the Consumer Price Index growth available from FRED. The market return and the one-month treasury bill is the Fama French market factor from Kenneth French’s website and for the bond return we employ Moody’s Seasoned Baa corporate bond yield.

¹⁴Significance at the 10% (1%) level is indicated by * (***) . Note that the time horizon for the regressions and all figures that include their fitted values is 1985Q1-2015Q1 due to the limited availability of NOS at the end of the sample.

selected variables of the two VAR systems in response to a TFP news shock. It is evident that both cost of capital measures decline significantly several years before TFP increases significantly. This countercyclical movement of excess returns indicates a decline in the opportunity costs for holding inventories and is hence consistent with the rise in inventories prior to movements in TFP driven by a strong demand effect.¹⁵ ¹⁶

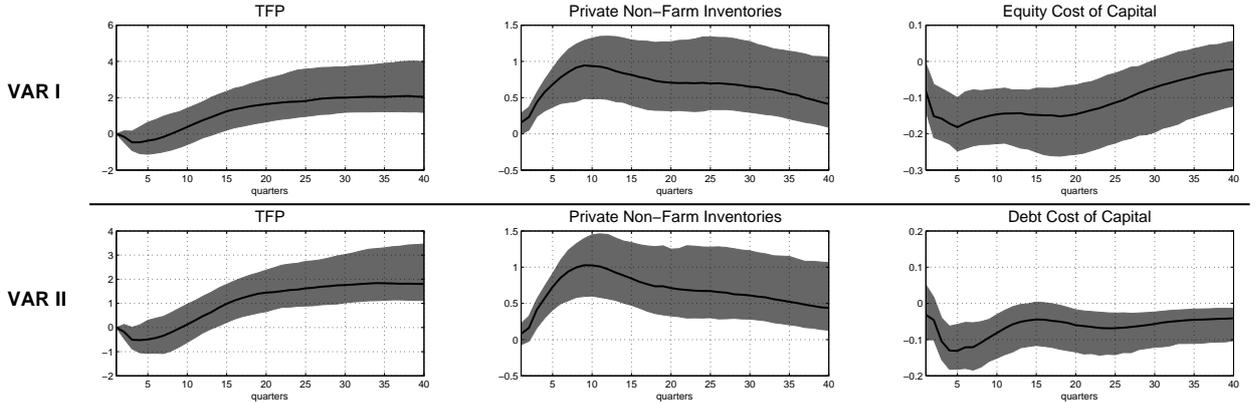


Figure 5: **IRF of Debt and Equity Cost of Capital measures to TFP news shock.** Sample 1985Q1-2015Q1. VAR I and II show selected IRFs from two separate eight-variable VAR systems. They differ in the use of either the equity or the debt cost of capital measure. The remaining variables in the VAR are TFP, GDP, consumption, hours, inventories, inflation, E5Y. The shaded gray areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

2.4.2 Implied Cost of Capital

Based on the evidence in Jones and Tuzel (2013), we also construct implied cost of capital measures from firm level data to proxy the opportunity costs of holding inventories. The implied cost of capital (ICC) is a firm’s internal rate of return that equates the present value of expected future cash flows with the current stock price; as such the ICC

¹⁵The responses of variables we do not show in Figure 5 in this section and Figure 6 in the next section are almost identical to the ones reported in Figure 1. Figures with responses for all variables are available on request.

¹⁶The unconditional correlation between HP filtered GDP and HP filtered ECC (DCC) is -0.43 (-0.58). This is much weaker than the unconditional correlation of HP filtered GDP and HP filtered activity variables: consumption (0.84), hours (0.87), investment (0.89) and inventories (0.75). The correlation of HP filtered private non-farm inventories with ECC (DCC) is -0.14 (-0.37). These correlations provide additional evidence that the strong decline of opportunity cost measures conditional on a TFP news shock is not purely based on an unconditional feature in the data.

is the market discount rate used to discount firm’s expected cash flows.¹⁷ To ensure our results are not driven by one specific method to construct ICCs, we employ four measures — according to Joseph R. Gordon (1997) (GORDON), Easton (2004) (MPEG), Ohlson and Juettner-Nauroth (2005) (OJ) and Gebhardt et al. (2001) (GLS) — which are widely used in the literature. These ICC measures can broadly be classified in three categories: OJ and MPEG are based on abnormal earnings growth models, GLS is based on the individual income valuation model, and GORDON is based on the Gordon growth model. These differ in assumptions regarding short- and long-term growth rates, their use of forecasted earnings, and the explicit forecast horizon. We use quarterly firm-level data from Compustat and CRSP of listed non-financial corporations to estimate expected earnings forecasts and subsequently construct the firm-level ICC measures. These procedures closely follows the methodologies summarized in Hou et al. (2012). We describe the details of the ICC construction in Appendix B. The quarterly firm level observations of a particular ICC measure are aggregated to a quarterly time series by taking the average per quarter. While the four measures are constructed from firm-level data that does not cover the whole economy, they can still provide useful corroborative evidence in addition to the aggregate cost of capital measures employed in the section above.

The resulting time series for the four ICCs are used one-by-one as an input to an eight-variable VAR system. Figure 6 shows that all measures decline significantly in response to a TFP news shock. Overall, and consistent with the cost of capital measures based on aggregate data in the section above, also these these measures constructed from firm-level data provide corroborative evidence for a decline in the opportunity costs of inventories, indicative of a strong positive demand effect driving up inventories in light of news about higher future TFP.

¹⁷Also Sarte et al. (2015) highlight the prominent role played by the discount rate for explaining inventory dynamics in a post Great Moderation sample.

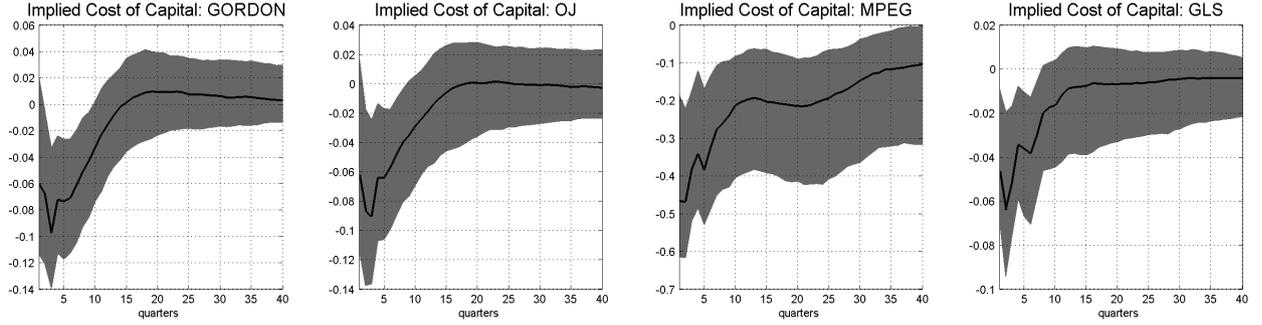


Figure 6: **IRF of Implied Cost of Capital measures to TFP news shock.** Sample 1985Q1-2015Q2. Subplots result from eight variable VARs comprising TFP, GDP, consumption, hours, inventories, ICC measure, inflation, E5Y. The ICC measures were included one-by-one in the VAR system. The shaded gray areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

2.4.3 Marginal Cost

While the sections above were concerned with the demand effect, in this section we evaluate whether there is evidence for a negative substitution effect. For this purpose we construct several measures for marginal costs considered in the literature.

Real marginal cost is given by

$$MC = \frac{W}{PF_h(k, h)}$$

where the real wage is given by W/P and $F_H(\cdot)$ is the marginal product of labor. Hence, any formulation of marginal cost will depend on assumptions on the production function. If the production function is Cobb-Douglas, log real marginal cost is proportional to the labor share

$$MC^{CD} = \ln(s)$$

where the labor share s is given by $Wh/(PF(k, h))$. Considering a generalization that allows

for the production function to be CES, the log real marginal cost is proportional to

$$MC^{CES} = \ln(s) - \left(\frac{1}{\sigma} - 1\right)[\ln(y) - \ln(Z \cdot h)],$$

where technology is denoted by Z , σ is the elasticity of substitution between capital and labor and y is output in terms of value added.¹⁸ The labor share can be measured by the BLS labor share in the private business sector (preferred measure by Nekarda and Ramey (2013)), or the nonfarm business version of this measure (preferred by Galí et al. (2007)). As a measure for technology, in line with Nekarda and Ramey (2013), we use the utilization adjusted TFP series provided by John Fernald and a baseline value of 0.5 for σ . We use nonfinancial corporate business gross value added as measure for output which is divided by population as well as the GDP deflator. Hours worked, h , is defined as outlined in the data section above.

We construct marginal cost measures based on the assumption of a Cobb-Douglas or CES production function. Each of these are constructed either using the labor share series preferred by Nekarda and Ramey (2013) or by Galí et al. (2007). In addition to these four measures, we construct two additional measures that aim to correct the labor share for overhead labor under the assumption of either a Cobb-Douglas or a CES production function. The labor share for these is constructed, in line with the methodology of Nekarda and Ramey (2013), by multiplying BLS data on employees, average weekly hours and average hourly wages (all of production and nonsupervisory employees in the private sector) and then dividing by current dollar output in private business.¹⁹

Figure 7 shows the responses of marginal cost measures when they are included one-by-one in our eight-variable VAR. The first subplot shows the response of a marginal cost

¹⁸If output is measured as gross output, one obtains the same expression as long as the production function is either (i) a generalized CES in which the elasticities of substitutions are equal across all inputs or (ii) a nested CES in which σ is the elasticity of substitution between the labor input and a composite of the other inputs (see Nekarda and Ramey (2013)). For this reason we use the value added measure for y .

¹⁹There is no time exact data on overhead labor and Nekarda and Ramey (2013) state that this correction can only be treated as an approximation. However, Rotemberg and Woodford (1999) argue that failing to account for overhead labor is an important reason for estimates of marginal costs to be biased.

measure based on the Cobb-Douglas production function and the labor share measure preferred by Nekarda and Ramey (2013) (CD: Nekarda-Ramey 1). Interestingly, marginal costs do not move in anticipation of news about higher future TFP but increase around the time when TFP rises significantly.²⁰ This doesn't provide any evidence for a strong negative substitution effect which would imply a decline of marginal costs at the time the productivity is realized. Under the assumption of a CES production function, also this marginal cost measure (CES: Nekarda-Ramey 1, shown in the second subplot) rises once the increase in is realized TFP. For the first few quarters after the arrival of the news we see a decline in this marginal cost measure, prior to the significant increase in TFP at quarter 15 (not shown). When using the labor share measure preferred by Galí et al. (2007) responses are qualitatively and quantitatively very similar. These are shown in Appendix A.2 where we provide further evidence on robustness of the exercises related to marginal cost measures. Subplots three and four in Figure 7 show responses of the marginal cost measures once accounting for overhead labor (Nekarda-Ramey 2). Under the assumption of either the Cobb-Douglas or the CES production function these measures do not move upon the arrival of news about higher future TFP and they increase after several quarters. While the CES based measure starts to decline from their peak only very slowly, after about 35 quarters, the Cobb-Douglas based measure declines somewhat earlier and even falls below zero after 33 quarters.

The differences in the responses of the marginal cost measures highlight that all these measures can only be approximations to marginal costs and the discussion about the most suitable construction is ongoing.²¹ Despite the differences in the responses, it is interesting to note though that none of the marginal cost measures indicates a significant decline of marginal costs upon the arrival of the TFP news shock or in the first quarters when the increase in TFP is realized around quarter 10. Only the Cobb-Douglas based measure that accounts for overhead labor (Nekarda-Ramey 2) falls below the zero line, but significantly

²⁰The behavior of the variables in the VAR that are not shown are very similar to the ones in Figure 1 where TFP increases significantly after about 12 quarters.

²¹For example Rotemberg and Woodford (1999) and Nekarda and Ramey (2013) discuss several shortcomings in the construction of marginal cost measures.

only after 33 quarters, a long time after the realization of higher TFP. Hence, none of the marginal cost measures indicates support for a strong negative substitution effect that would provide incentives to shift production into the future and draw down the inventory stock upon arrival of news about higher future TFP. Together with the evidence in the sections above on the presence of a strong demand effect the documented behavior of marginal cost is consistent with the increase in inventories in response to news about higher future TFP.²²

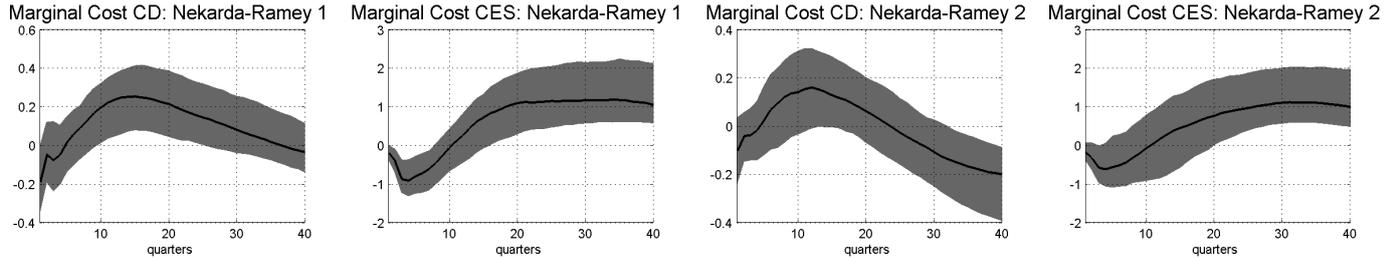


Figure 7: **IRF of marginal cost measures to TFP news shock.** Sample 1985Q1-2015Q2. Subplots result from eight variable VARs comprising TFP, GDP, consumption, hours, inventories, marginal cost measure, inflation, E5Y. The marginal cost measures were included one-by-one in the VAR system. The shaded gray areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

2.5 Corroborative evidence and guidance for models

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

Jaimovich and Rebelo (2009) state that for models to facilitate comovement of consumption and investment in response to news about future higher TFP a strong increase in utilization and/or hours worked is required. The VAR results in Figure 8 document a strong increase in capital utilization and the real wage in response to a TFP news shock. The positive hump-shaped response of the real wage is consistent with the increase in hours

²²Unconditionally, the correlations of HP filtered GDP with HP filtered measures for marginal cost indicate weak countercyclicality or acyclicity (for detailed correlations see Appendix A.2). This is in line with the pro-/acyclicity of markups reported in Nekarda and Ramey (2013). However, over a longer horizon our IRFs to a TFP news shock indicate procyclicality of marginal costs. This is in line with the findings in Rotemberg and Woodford (1999) that the labor share tends to rise late in expansions and fall late in recessions.

documented in Figure 1. Further, it is broadly consistent with a procyclical response of real marginal cost. In line with the evidence in the section above, it provides no evidence for a negative substitution effect that would be in support of drawing down inventories in response expected higher future productivity.

Figure 8 further shows that the inventory to sales ratio moves countercyclically in response to a TFP news shock. This is another indication that TFP news shocks are potentially important drivers of aggregate fluctuations as the unconditional countercyclicality of the inventory to sales ratio is a commonly accepted view in the literature (e.g. Blinder (1981)). In the sections below we will show that in our model this countercyclicality is a necessary condition for comovement of inventories with the other macroeconomic aggregates in a setup with flexible prices.²³

The following sections will discuss the implications of the empirically documented procyclical inventory movements in response to TFP news shocks for structural models.

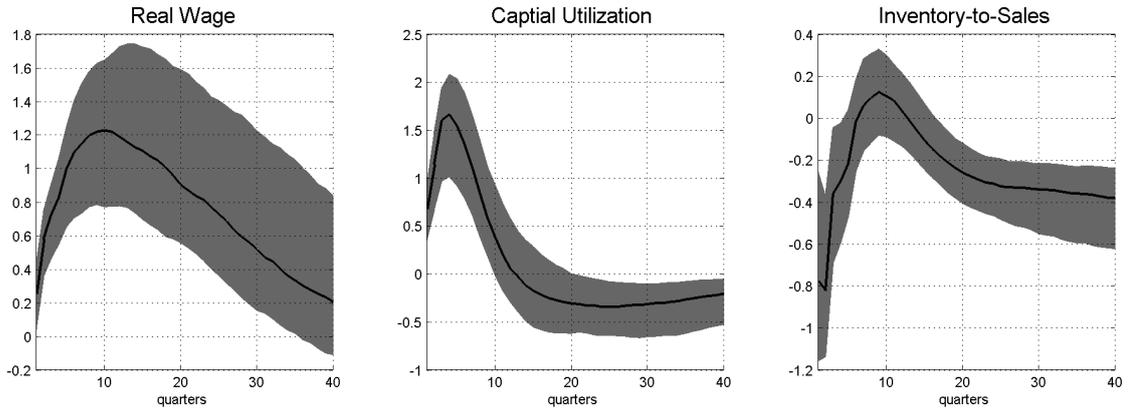


Figure 8: **IRF to TFP news shock.** Sample 1985Q1-2015Q2. Subplots result from seven-variable VARs comprising TFP, GDP, consumption, hours, inflation, E5Y and one of the plotted variables above at a time. The shaded gray areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

²³The inventory to sales ratio is the ratio of private non-farm inventories and final sales of domestic business as in Lubik and Teo (2012). Utilization is provided by Fernald (2014). The real wage is compensation of employees, nonfinancial corporate business, in real per capita terms. Responses of the other variables in the seven-variable VAR system are virtually identical to the ones in Figure 1 and are available upon request.

3 A standard business cycle model with inventories

We now present a relatively standard business cycle model supplemented with inventories and study the behaviour of inventories and other important macroeconomic variables in response to TFP news. The model is a New Keynesian framework with sticky prices and wages, augmented with inventories introduced as in Jung and Yun (2006) and Lubik and Teo (2012), in the spirit of Bils and Kahn (2000). Following Jaimovich and Rebelo (2009), we include the trio of a particular specification of preferences, investment adjustment costs and costly capacity utilization as a common and simple means for producing comovement of consumption, investment and hours-worked in response to news about TFP. Our New Keynesian framework nests the case of flexible prices and wages, and thus our standard model builds on the work of Crouzet and Oh (2016) who study the behaviour of inventories in response to TFP news in a flexible price and wage Jaimovich and Rebelo (2009) real business cycle model.

The model economy consists of a large number of identical infinitely-lived households, a competitive intermediate goods firm, a continuum of monopolistically competitive distributors indexed by $i \in [0, 1]$, a competitive final goods producer, a continuum of monopolistically competitive labor unions indexed by $j \in [0, 1]$, and a competitive employment agency. The intermediate goods firms owns its capital stock and produces a homogeneous good which it sells to distributors, who then differentiate the good into distributor-specific varieties which they sell to the final goods firm who aggregates the varieties into a final good. The final good may be used for consumption or investment. Distributors face frictions in setting prices for each distributor-specific variety and hold inventories of each variety. Monopolistic unions buy homogeneous labor from households, transform it into differentiated labor inputs, and sell them to the employment agency who aggregates the differentiated labor into a composite which it then sells to the intermediate goods producer. The unions face frictions in setting wages for each labor type. Since this particular decentralization of wage stickiness implies that consumption and hours are identical across households, henceforth we will refer to a

stand-in representative household.

3.1 Household, employment unions and employment agency

The stand-in household's lifetime utility is given by

$$E_0 \sum_{t=0}^{\infty} \beta^t \left\{ \frac{1}{1-\sigma} \left(c_t - \frac{n_t^{h^{1+1/\xi}}}{1+1/\xi} \right)^{1-\sigma} \right\}, \quad (6)$$

where c_t is consumption, n_t^h is hours supplied to the labor union, $0 < \beta < 1$ is the subjective discount factor, $\xi > 0$ is the Frisch labor supply elasticity and $\sigma > 0$ is the household's elasticity of intertemporal substitution. The particular functional form of preferences follows the one in Greenwood et al. (1988).²⁴ The household's period t budget constraint is given by

$$c_t + \frac{B_{t+1}^n}{P_t} = R_t^n \frac{B_t^n}{P_t} + \frac{W_t^h}{P_t} n_t^h + f_t, \quad (7)$$

where P_t is the price of the final good in terms of the nominal unit under the control of the central bank, B_t^n are nominal bonds earning a risk-free gross nominal rate of return R_t^n , W_t^h is the nominal wage rate paid by the labor union and f_t are any real profits flowing collectively from the various other entities in the model. The household chooses c_t , n_t^h , and B_{t+1}^n to maximize (6) subject to (7).

Our sticky-wage framework follows the decentralization of Schmitt-Grohé and Uribe (2006) and Smets and Wouters (2007) and so we discuss it only briefly, leaving the details to appendix C.1.1. Labor unions acquire homogenous labor n_t^h from the household at wage W_t^h , differentiate it into labor types n_{jt} , $j \in [0, 1]$, which they sell to the employment agency for wage W_{jt} . The employment agency acquires each j th intermediate labor type n_{jt} at wage W_{jt} from the labor unions, and combines the differentiated labor into a composite N_t , and then sells the composite labor to the intermediate goods producers for wage W_t .

²⁴These “no-income effect” preferences are a special case of the more general form used in Jaimovich and Rebelo (2009).

3.2 Intermediate Goods Firm

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

$$y_t = z_t n_t^\alpha \tilde{k}_t^{\theta_y}, \quad (8)$$

where z_t is total factor productivity that follows an exogenous stochastic process, and \tilde{k}_t is physical capital services. Capital services is defined by $\tilde{k}_t = u_t k_t$, where k_t is the firm's stock of physical capital and u_t is the utilization rate of that stock. The firm's capital stock evolves according to

$$k_{t+1} = (1 - \delta(u_t))k_t + i_t \left[1 - S \left(\frac{i_t}{i_{t-1}} \right) \right], \quad (9)$$

where $\delta(\cdot)$ is a depreciation function satisfying $\delta'(\cdot) > 0$, $\delta''(\cdot) > 0$ and $\delta(1) = \delta_k > 0$. $S(\cdot)$ is an investment adjustment cost function as in Christiano et al. (2005) with $S(1) = S'(1) = 0$ and $S''(1) = s'' > 0$.

Each period, the firm acquires n_t at wage W_t from the employment agency, i_t from the final goods producer at price P_t , and then sells its output y_t at nominal price $P_t \tau_t$ to the distributors. The above then implies the firm's real profits are given by $\Pi_t^y = \tau_t y_t - w_t n_t - i_t$. The firm's problem involves choosing k_{t+1} , i_t and n_t to maximize $E_0 \sum_{t=0}^{\infty} \frac{\beta^t \lambda_t}{\lambda_0} \Pi_t^y$ subject to equation (9).

Additionally, define the marginal cost of production, mc_t , for the intermediate goods firm as $mc_t = \frac{w_t}{MPN_t}$, where MPN_t is the marginal product of labor. From the intermediate goods firm's labor first-order condition, $w_t = \tau_t \alpha \frac{y_t}{n_t}$, it then follows that the output price τ_t is equal to the the marginal cost of production mc_t .

3.3 Final goods firm

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

$$s_t = \left[\int_0^1 \nu_{it}^{\frac{1}{\theta}} s_{it}^{\frac{\theta-1}{\theta}} di \right]^{\frac{\theta}{\theta-1}}, \quad \theta > 1, \quad (10)$$

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

$$\nu_{it} = \left(\frac{a_{it}}{a_t} \right)^\zeta, \quad \zeta > 0, \quad (11)$$

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

The firm acquires each i th goods variety at nominal price P_{it} from the distributors, and sells the final good at nominal price P_t where it may be used as a consumption or as an input into the production of investment goods. The firm maximizes the profit function $\Pi_t = P_t s_t - \int_0^1 P_{it} s_{it} di$ by choosing $s_{it} \forall i$, yielding a demand function for s_{it} for the j th variety,

$$s_{it} = \nu_{it} \left(\frac{P_{it}}{P_t} \right)^{-\theta} s_t, \quad (12)$$

and price index P_t ,

$$P_t = \left[\int_0^1 \nu_t(i) P_{it}^{1-\theta} dj \right]^{\frac{1}{1-\theta}}. \quad (13)$$

3.4 Distributors

Distributors acquire the homogenous good y_t from the intermediate goods firm y_t at real price τ_t , and then differentiate it into goods-variety y_{it} at zero cost, with a transformation rate of one unit of the homogeneous good to one unit of the differentiated good. Goods

available for sale are the sum of the differentiated output and depreciated previous period's inventories,

$$a_{it} = (1 - \delta_x)x_{it-1} + y_{it}, \quad (14)$$

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

$$x_{it} = a_{it} - s_{it}, \quad (15)$$

and δ_x is the period depreciation of the inventory stock. The distributors have market power over the sales of their differentiated varieties, and thus the i th distributor sets the price P_{it} for sales s_{it} of its variety, subject for to the demand curve for that variety. The distributors face frictions in setting their prices, and as in Lubik and Teo (2012), we assume that the i th distributor faces convex adjustments costs in the form $\frac{\kappa}{2} \left[\frac{P_{it+k}}{\pi_{t-1}^{\iota_p} \pi^{1-\iota_p} P_{it+k-1}} - 1 \right]^2 s_t$. Each period, the i th distributor then faces the problem of choosing P_{it} , s_{it} , y_{it} and a_{it} to maximize

$$E_t \sum_{k=0}^{\infty} \beta^k \frac{\lambda_{t+k}}{\lambda_t} \left\{ \frac{P_{it+k}}{P_{t+k}} s_{it+k} - \tau_t y_{t+k}(j) - \frac{\kappa}{2} \left[\frac{P_{it+k}}{\pi_{t-1}^{\iota_p} \pi^{1-\iota_p} P_{it+k-1}} - 1 \right]^2 s_t \right\}, \quad (16)$$

subject to the demand curve (12), the stock and inventory expression (14) and (15). Substituting in the demand curve (12) for s_{it} , and letting μ_t^a and μ_t^x be the multipliers on (14) and (15) respectively, the distributor's first-order conditions are given by

$$\tau_t = \mu_t^a \quad (17)$$

$$\mu_t^x = \beta E_t \frac{\lambda_{t+1}}{\lambda_t} \mu_{t+1}^a (1 - \delta_x) \quad (18)$$

$$\mu_t^a = \frac{P_{it}}{P_t} \zeta \frac{s_{it}}{a_{it}} + \mu_t^x \left[1 - \zeta \frac{s_{it}}{a_{it}} \right] \quad (19)$$

$$(1-\theta) \frac{s_{it}}{P_t} - \kappa \left[\frac{P_{it}}{\pi_{t-1}^{\iota_p} \pi^{1-\iota_p} P_{it-1}} - 1 \right] \frac{s_t}{\pi_{t-1}^{\iota_p} \pi^{1-\iota_p} P_{it-1}} + \beta E_t \frac{\lambda_{t+1}}{\lambda_t} \kappa \left[\frac{P_{it+1}}{\pi_t^{\iota_p} \pi^{1-\iota_p} P_{it}} - 1 \right] \frac{P_{it+1} s_{t+1}}{\pi_t^{\iota_p} \pi^{1-\iota_p} P_{it}^2} + \mu_t^x \theta \frac{s_{it}}{P_{it}} = 0, \quad (20)$$

where (17), (18), (19) and (20) describe the optimal choices of y_{it} , x_{it} , a_{it} and P_{it} .

Note from equation (14) that with beginning of period inventories predetermined, a

distributor can only further increase its stock of available goods for sale a_{it} in period t by acquiring additional output y_{it} , purchased at real price τ_t . Thus the cost of generating an additional unit of goods for sale is equal to the price of output (or marginal cost of output) τ_t , which from the intermediate goods firm's problem is also the marginal cost of production mc_t . At the optimum, equation (17) says that the cost of an additional unit of goods for sale τ_t is equal to the value of those goods for sale, μ_t^a ,

Next, note from the inventory definition (15) that for a given level of goods available for sales, a_{it} , any increase in sales s_{it} results in a reduction in inventory. Thus, the opportunity cost of sales for the distributor is equal to the value of foregone inventory, μ_t^x , which we can then interpret as the marginal cost of sales. The first-order condition (18) then says that the value of an additional unit of inventory today, μ_t^x , is the expected discounted value of the extra level of goods available for sale next period generated by that inventory, μ_{t+1}^a , whose value is in turn equal to the price of output next period, τ_{t+1} from (17). Thus, in a model with inventory, the marginal cost of sales is equal to the expected discounted value of next period's marginal cost of output, since increasing sales by drawing down inventories to forgo production today means that eventually the distributor will need to increase production in the future.

The first-order condition (19) says that the marginal value of extra goods for sale μ_t^a consists of the value of the extra sales generated by the additional a_{it} , plus the value of the additional inventory yield from the unsold portion of the additional a_{it} . Combining (17), (18) and (19) yields

$$\tau_t = \zeta \frac{s_{it}}{a_{it}} + \beta E_t \frac{\lambda_{t+1}}{\lambda_t} \tau_{t+1} (1 - \delta_x) \left[1 - \zeta \frac{s_{it}}{a_{it}} \right], \quad (21)$$

showing that the distributor chooses a_{it} such that this benefit is equal to the marginal cost of output τ_t . We will refer to (21) as the distributor's optimal stocking condition.

Finally, the first-order condition (20) describes the distributor's optimal choice of price p_{it} in terms of the marginal cost of sales μ_t^x and in response to the pricing frictions. The interpretation of this expression is standard, except for the presence of the marginal cost of

sales instead of the marginal cost of output as in a typical model without inventories. Indeed in standard models without inventories, the marginal cost of sales is equal to the marginal cost of output. Here however, the presence of inventories drives a wedge between the marginal cost of output and marginal cost of sales. Thus we can think of there being two additive markups: the markup between marginal cost of production and the marginal cost of sales, and the markup between the marginal cost of sales and the price. The distributor adjusts these two margins jointly through its joint decision of inventories and prices. The optimal stocking condition (19) describes the adjustment of the first markup through inventories; the optimal pricing condition (20) describes the adjustment of the second markup through price-setting.

A special case of this occurs under flexible prices, where the pricing condition (20) reduces to $(1 - \theta)\frac{s_{it}}{P_t} + \mu_t^x \theta \frac{s_{it}}{P_{it}} = 0$, implying

$$\frac{P_{it}}{P_t} = \frac{\theta}{\theta - 1} \mu_t^x, \quad (22)$$

such that the distributor sets its relative price as a constant markup over the marginal cost of sales. Even under this condition however, the distributor still dynamically adjusts the total markup between marginal cost of production and price by using inventories to adjust the markup between marginal cost of production and the marginal cost of sales.

3.5 Monetary policy, stochastic processes and equilibrium

The nominal interest rate, R_t^n , is set by the monetary authority according to a feedback rule,

$$\frac{R_t^n}{R^n} = \left(\frac{R_{t-1}^n}{R^n} \right)^{\rho_\pi} \left(\frac{\pi_t}{\pi} \right)^{(1-\rho_\pi)\phi_\pi},$$

where the parameters ρ_π and ϕ_π determine the interest rate's persistence and its response to deviations of inflation from target, respectively.

The productivity shock process follows

$$z_t = \rho z_{t-1} + \epsilon_{z,t},$$

where ρ determines the persistence and the innovation $\epsilon_{z,t}$ is the sum of a standard unanticipated shock, $\epsilon_{0,t}$, and an anticipated component, $\epsilon_{t-h,h}$, where news about $h > 0$ quarters ahead arrives at time $t - h$. The innovations are uncorrelated across time, *i.i.d* and $N(0, \sigma_0^2)$ and $N(0, \sigma_h^2)$, respectively.²⁵

We define a symmetric equilibrium for the model economy in the appendix. We solve the resulting non-linear system by taking a linear approximation around the steady state.

3.6 Calibration

In general, our calibration strategy involves using parameter ranges close to those in the news models of Jaimovich and Rebelo (2009) and Schmitt-Grohe and Uribe (2012), and the New Keynesian inventory model of Lubik and Teo (2012). Our analytical results below suggest that suppressing the rise of marginal costs relative to sales in an expansion will be critical to obtaining a procyclical response of inventories in response to news. To this end, we choose values of key parameters such as labor supply elasticity, elasticity of capacity utilization, and the Calvo sticky wage parameter intended to limit the rise of marginal costs and thus give the standard model the best possible chance of matching the data, while still using parameter values within the ranges in the literature.

Beginning with the household, we set the household's subjective discount factor β to 0.995 and elasticity of intertemporal substitution to σ to 1, as in Jaimovich and Rebelo (2009). We set the Frisch elasticity of labor supply ξ to 5, slightly larger than the 2.5 used in Jaimovich and Rebelo (2009). For goods production, we set the elasticity of output to current labor α to 0.64 as in Jaimovich and Rebelo (2009), and then following Schmitt-Grohe and Uribe (2012), we assume decreasing returns to scale of 10%, implying an elasticity of output

²⁵The way we introduce news shocks is standard in the literature, see for example in Schmitt-Grohe and Uribe (2012) and Khan and Tsoukalas (2012).

to capital θ_y of 0.26.²⁶ For the parameters related to physical capital, we set steady-state physical capital depreciation δ to 0.025, steady state utilization to $u_{ss} = 1$, the elasticity of marginal utilization $\delta_k''(1)/\delta_k'(1)$ to 0.05, slightly lower than the 0.25 used in Jaimovich and Rebelo (2009), and the investment adjustment cost parameter s'' to 2, slightly larger than the 1.3 used in Jaimovich and Rebelo (2009).

For the parameters related to inventory, we start by setting the inventory depreciation rate δ_x to 0.05, based on one of the two cases considered by Lubik and Teo (2012). We then set the goods aggregator curvature θ to 6.8 to yield a steady state goods markup of 10%, and then the taste shifter curvature ζ to 0.67 to yield a steady state sales-to-stock ratio of 0.55, as in Lubik and Teo (2012).

For the parameters related to the New Keynesian aspects of the model, we set the price and wage backwards indexation parameters ι_p and ι_w to 0.45, and the steady state wage markup mkp_w to 1.10. For the cases with sticky prices, we set the price adjustment cost parameter κ to 250 so that the coefficient κ_{mc} on the marginal cost term in the linearized New Keynesian Phillips Curve (NKPC) is equal to 0.025, consistent with the restricted GMM estimate of the NKPC in Lubik and Teo (2012). For the cases with sticky wages, we set the Calvo probability ζ_w to 0.85. For the monetary rule, we set the Taylor rule inertia ρ_π and inflation parameters to 0.8 and 1.5 respectively. Where we investigate flexible wages and/or prices, we set $\kappa = 0$ and/or $\zeta_w = 0$ accordingly.

Finally, we set the TFP process persistence ρ to 0.95. For all our simulations, we use a news shock horizon h of 8. Panel A of Table 1 summarizes all parameter values for the standard model.

²⁶Jaimovich and Rebelo (2009) discuss how decreasing returns to scale can be interpreted as a result of the presence of organizational (intangible) capital in production. In our present context, decreasing returns to scale has little impact on the response of inventory in the standard model. We nevertheless assume it here to hold constant the elasticity of output to current labour between our standard model and extended model, which will feature intangible capital.

Table 1: Summary of calibrated parameters

Description	Parameter	Value
Panel A: Standard model		
Subjective discount factor	β	0.99
Household elasticity of intertemporal substitution	σ	1
Frisch elasticity of labor supply	ξ	5
Labor elasticity in production	α	0.64
Capital elasticity in production	θ_y	0.26
Capital depreciation	δ_k	0.025
Depreciation elasticity of capacity utilization	$\delta_k''(1)/\delta_k'(1)$	0.05
Investment adjustment cost	s''	2
Inventory depreciation	δ_x	0.05
Goods aggregator curvature	θ	6.8
Taste shifter curvature	ζ	0.67
Calvo wage probability	ζ_w	0.85
Wage indexation	ι_w	0.45
Price adjustment cost	κ	250
Price indexation	ι_p	0.45
Steady state hours	n	0.2
Steady state capacity utilization	u	1
Steady state wage markup	mkp_w	1.1
TFP process persistence	ρ	0.95
TFP innovation standard deviation	σ_ϵ	1
Taylor rule inertia	ρ_π	0.8
Taylor rule inflation	ϕ_π	1.5
Panel B: Model with knowledge capital		
Knowledge capital elasticity in production	ϵ_y	0.15
Contribution of prior knowledge in its production	γ_h	0.8
Internalized labor elasticity in knowledge capital	ν_h	0.15
Externalized labor elasticity in knowledge capital	τ_h	0.15

4 Understanding the comovement puzzle

We now explore the behaviour of inventories in response to TFP news shocks in the standard model. We begin by deriving some analytic equilibrium expressions to gain insight into the model. We then explore the simulated response to TFP news in the calibrated version of the model.

4.1 Analytical framework

Combining the equilibrium versions of equations (15), (17) and (19) and then linearizing yields

$$\hat{x}_t = \hat{s}_t - \frac{\tau}{\Phi} \hat{\tau}_t + \frac{\tau - \zeta \frac{s}{a}}{\Phi} \hat{\mu}_t^x, \quad (23)$$

which we can decompose further, using equation (18), as

$$\hat{x}_t = \hat{s}_t - \frac{\tau}{\Phi} \hat{\tau}_t + \frac{\tau - \zeta \frac{s}{a}}{\Phi} \left\{ \hat{\tau}_{t+1} - \hat{R}_t^f \right\}, \quad (24)$$

where the real interest rate is given by $R_t^f = \frac{1}{\beta E_t \frac{\lambda_{t+1}}{\lambda_t}}$, $\Phi = (\tau - \mu_x) \frac{x}{a} > 0$, and where for any variable m_t , \hat{m}_t denotes the % deviations of variable m_t from the non-stochastic steady state. The equilibrium behaviour of inventories in the model is summarized in equation (24). Thus, we can frame any explanations about why inventories will or will not rise in response to news through the lens of this expression. The equation shows that there are potentially three different forces influencing the dynamics of inventories: the levels of sales, the path of the marginal cost of output, and the level of the real interest rate.

The first term on the right hand side of equation (24) represents a “demand” effect on inventories whereby they respond positively to contemporaneous increases in sales. We can think of the remaining terms as describing an “intertemporal substitution effect” comprised of current marginal costs and discounted future marginal costs, whereby firms use inventories to shift production through time in response to temporal variation in the cost of production.

All else equal, an expected decline in marginal costs over time puts downward pressure on inventory accumulation in the present as firms draw down inventories in order to shift production from the present to the relatively less costly future. For a given path of marginal costs, a rise in the real interest rate then puts downward pressure on inventories as firms discount the impact of future marginal costs more heavily.

Finally, knowing the relative size of the coefficients in equation (24) is important for determining the quantitative role these various influences at play. Using the steady state and parameter definitions, it follows that that $\tau > \tau - \zeta \frac{s}{a} > 0$ for all parameter values consistent with the parameter definitions, implying that there is more weight on the negative effect of current marginal costs than the positive effect of future marginal costs. Indeed, for our baseline calibration, $\frac{\tau}{\Phi} = 37.2$ and $\frac{\tau - \zeta \frac{s}{a}}{\Phi} = 22.1$, such that there is about 67% more weight on current marginal costs than future marginal costs.²⁷ Comparing the relative importance of the coefficients on sales and marginal costs is difficult here, because in general the percentage change of sales over the business cycle will differ from that of marginal costs. Looking at just the intertemporal substitution channel however, the above discussion of equation (24) suggests that the response of current marginal cost will have more influence on inventories than expected discounted future marginal costs.

4.1.1 Understanding the impact of TFP news

Equation (24) is helpful for quantifying the forces driving inventory dynamics in the model in general. To understand the particular impact of TFP news on inventories, we need to consider the general equilibrium effects of TFP news on the main components of this expression. In our particular decentralization, it is helpful to visualize the impact of TFP news on the market for output goods in $\tau_t - y_t$ space that exists in our particular

²⁷It follows from steady state definitions that the ratio of the coefficients on current and future marginal costs is given by $\frac{\frac{\tau}{\Phi}}{\frac{\tau - \zeta \frac{s}{a}}{\Phi}} = \frac{1}{1 - \theta(1 - \beta(1 - \delta_x))}$. This ratio is increasing in both the elasticity of substitution θ and the inventory depreciation rate δ_x . Decreasing θ much below our calibrated value however drives the markup into implausibly high values. For example, decreasing θ from 6.8 to 5 moves the ratio of marginal-costs-coefficients from 1.67 to 1.42, but also moves the markup from 10% to about 18%. There is more uncertainty in the data about δ_x . Decreasing this parameter to 0.025 from 0.05 reduces the marginal-costs-coefficient ratio from 1.67 to 1.31.

decentralization between the intermediate goods firm and the distributors. The intermediate goods firm sells into this market according to an upward-sloping output supply curve equal to its marginal cost of production curve given $w_t = \tau_t \alpha \frac{y_t}{n_t}$, and the distributors buy into the market according to a downward-sloping output demand curve formed by the combination of equilibrium versions of equations (14), (15) and (21). When TFP news arrives, as is well known in the literature, the wealth effect of the TFP news drives up the demand for consumption. In this particular context, this in turn drives up the demand for sales of distributors. Note that for a given value of discounted expected future marginal costs, the optimal stocking condition (21) establishes a link between the current marginal cost of output τ_t , and some optimal value of the sales-to-stock ratio $\frac{s_t}{a_t}$. Given the increase in demand for sales, for a given marginal cost τ_t , the optimal stocking condition suggests a_t must then rise to maintain the optimal sales to stock ratio $\frac{s_t}{a_t}$. With inventories fixed in the initial period, the only way distributors can increase stock available for sales a_t is by acquiring more output y_t in the market for output. Thus the news shock has the effect of shifting the distributors' demand for output. Additionally, this shift is either amplified or suppressed by the effect of any changes in the expected future marginal costs term in equation (19).

If output goods were supplied inelastically with a flat supply curve, then any increase in sales s_t would be met by an equal increase in inventory to maintain the fixed sales-to-stock ratio $\frac{s_t}{a_t} = \frac{s_t}{s_t + x_t} = \frac{1}{1 + \frac{x_t}{s_t}}$. Under an upward-sloping supply curve however, any equilibrium rise in τ_t then implies a higher optimal sales-to-stock ratio, which for a given level of sales s_t necessarily implies a reduction in the inventory-to-sales ratio $\frac{x_t}{s_t}$. While inventories could still rise under a drop in the inventory-to-sales ratio, if the rise in marginal costs is large enough, for the given rise in sales, inventories may actually need to decrease as it becomes more attractive for the distributors to draw down inventories in the present to avoid the high current production costs. Thus, as we can see from both this intuition and the linearized equations, whether inventory will rise or fall for a given increase in sales s_t depends on the magnitude of the rise in marginal costs relative to the increase in sales, as well as the

magnitude of any expected increases or decreases in expected marginal costs in the future. Additionally, it is evident from the above discussion that a countercyclical inventory-sales ratio — a feature consistent with our empirical evidence in Section 2.5 — is a necessary condition for co-movement.

4.1.2 Special cases of flexible wages or prices

Crouzet and Oh (2016) study the response of inventories to TFP news in a flexible price and wage real business cycle model built around the Jaimovich and Rebelo (2009) framework. Under the assumption of constant capacity utilization, they first derive an analytical expression showing that inventories do not rise in response to TFP news for economically significant parameter values. They then demonstrate that inventories fall in response to TFP news in simulated impulse response functions of the calibrated model with variable capacity utilization. Crouzet and Oh (2016) further highlight the generality of this result in an appendix where they show the analytical result holds across different ways of introducing inventories.²⁸ We now consider the special case of flexible wages and prices to show how our expression (24) nests their stock-elastic demand model case.

First, imposing flexible wages has no direct impact on the form of equation (24). All terms in the expression are the same with or without sticky wages, and thus sticky wages can only have an impact on the dynamics of inventory indirectly through its impacts on the response of the other variables in this expression, such as the impact of sticky wages on marginal costs.

Imposing flexible prices in contrast eliminates a key component of (24). Under flexible prices ($\kappa = 0$), the equilibrium version of the pricing expression (20) reduces to $\mu_t^x = \frac{\theta-1}{\theta}$, such that (23) is given by

$$\hat{x}_t = \hat{s}_t - \frac{\tau}{\Phi} \hat{\tau}_t. \quad (25)$$

²⁸This includes the stock-elastic demand model of Bils and Kahn (2000) and the stockout-avoidance model of Kryvtsov and Midrigan (2013).

Importantly, under flexible prices, only contemporaneous sales and marginal costs drive inventory dynamics; the intertemporal substitution effect that describes the role of time-variation in the path of marginal costs through time is not a factor. Said another way, it is time-variation in the level, and not the path of marginal costs that matters under flexible prices (independent of whether wages are flexible or sticky). While Crouzet and Oh (2016) derive an analytical expression for their stock-elastic demand case that includes a role for expected discounted growth in marginal costs, our analysis above shows that expected discounted growth in marginal costs are constant under flexible prices. Indeed we show Appendix C.1.3 that the expression of Crouzet and Oh (2016) is equivalent to our equation (25), and that equation (24) nests their case.

4.2 Response to news in the standard model

We now investigate the simulated responses to TFP news in various versions of the calibrated standard model. In all scenarios, we consider the effect of a news shock about a one standard deviation rise in TFP 8 periods in the future, which will eventually be realized as anticipated.

Response to news in a model version *without* inventories. Before looking at the response of the standard model with inventories, as a point of reference, we briefly consider the response of companion model *without* inventory under fully flexible wages and prices to TFP news. When we abstract from inventories, our model nests the type of frameworks in Jaimovich and Rebelo (2009) and Schmitt-Grohe and Uribe (2012) which are now widely used in the news-literature as they are successful in resembling the empirically observed comovement of consumption, investment and hours worked in response to news about future productivity. Figure 9 (dashed line) shows IRFs of the standard model with fully flexible wages and prices ($\kappa = 0$ and $\zeta_w \approx 0$), but *without inventories*, using the particular calibration of the non-inventory parameters in the associated model. In response to news about higher TFP in the future, this model can successfully generate comovement of output, consumption, investment and hours. As shown in Jaimovich and Rebelo (2009),

in response to news, investment adjustment costs induces a drop in the value of new capital and thereby a reduction in the cost of utilization of capital. This in turn leads to a rise in utilization and corresponding shift in labor demand. With the absence of the income-effect under this particular form of preferences, the increase in labor demand on its own leads to an increase in hours-worked in equilibrium.

Response to news *with* inventories under flexible prices and wages. Figure 9 (solid line) shows the response of the standard model *with* inventories under fully flexible wages and prices to TFP news ($\kappa = 0$ and $\zeta_w \approx 0$).

It becomes apparent from the figure that our VAR evidence imposes two challenges to the model. First, inventories move countercyclically which is at odds with our empirical evidence. Second, the countercyclical response of inventories adversely impact the dynamic paths of other key variables. In particular, the response of hours, utilization and output are nearly flat and close to the steady state for the first 8 periods, even though sales grows over time. Quantitatively, this is not consistent with the VAR evidence and the strong boom prior to the realization of higher productivity generally thought of in the news-literature. We discuss these two challenges in turn.²⁹

Countercyclical inventories. Recall that under flexible prices, equation (25) says that the response of inventories is governed by the response of current sales and current marginal costs. Given the relative size of the coefficients on sales versus marginal costs, marginal costs would either need to fall, or, rise by a small amount relative to sales in response to news for inventories to rise. Using the earlier intuition about the output goods market, the increase in sales demand s_t induced by TFP news shifts the distributor's demand curve outwards, and with no intertemporal substitution channel under flexible prices which could potentially

²⁹Tangential to the issue at hand, note that Figure 9 shows that adding inventories reduces the spike in the real interest rate in the period before the TFP shock hits. Such real interest rate spikes are well known in news-driven models, reflecting the high opportunity costs of consumption in the period before the TFP shock hits and increases the marginal product of capital. See Christiano et al. (2008) for a discussion. In the case here, adding inventories reduces this spike, because the incentive to accumulate inventories increases when the TFP shock hits and lowers marginal costs, increasing hours more so than in the model without inventory. Since with the non-separable preferences the marginal utility of consumption is increasing in hours worked, this rise in hours increases the marginal utility of consumption in the period of the TFP shock, thereby reducing the real interest rate in the period before.

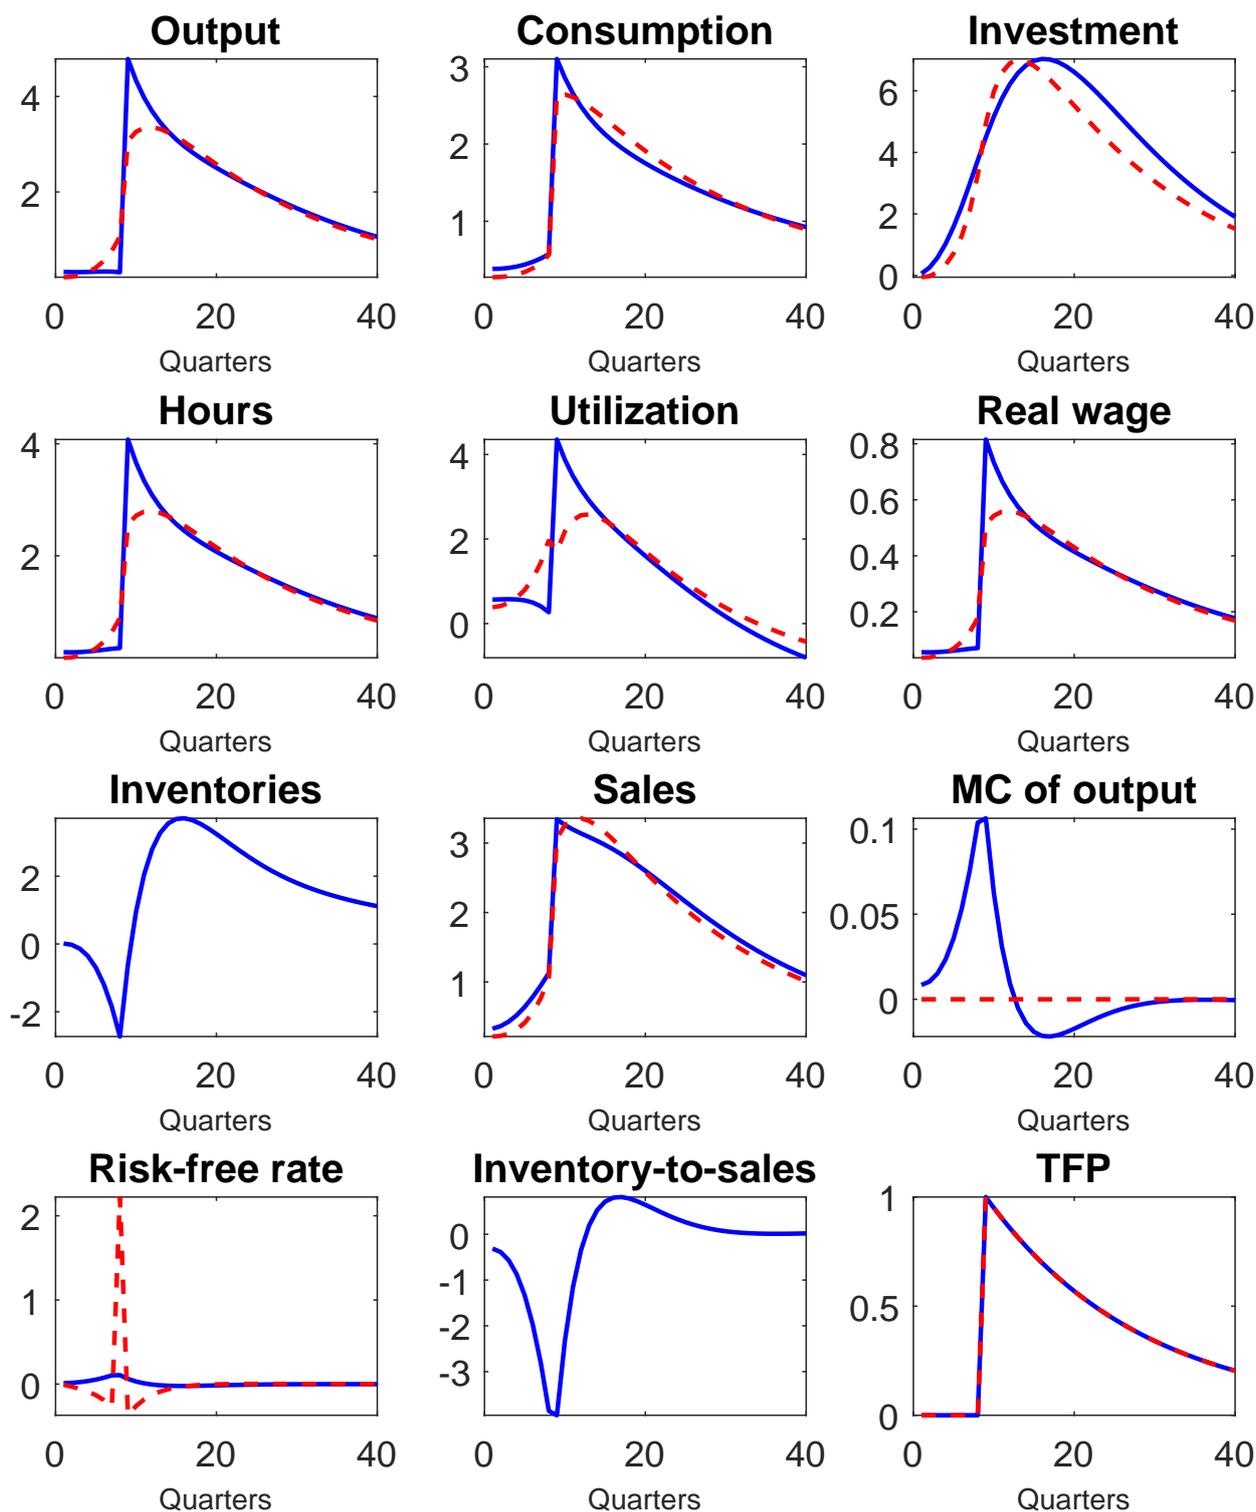


Figure 9: IRF to TFP news shock - Standard model with flexible wages and prices. Units are percentage deviations from steady state. Solid line: Standard model with flexible wages and prices; Dashed line: companion model *without* inventories.

temper or amplify this shift, the resulting equilibrium rise in marginal costs relative to sales is the only consideration for the response of inventory. Even though the effect of capacity utilization and a relatively elastic labor supply makes the intermediate goods firm's output supply curve relatively flat, the rise in marginal costs is still sufficient to drive to the optimal inventory-to-sales ratio enough that inventories fall over the first several periods.

Path of hours, utilization and output. Why does the countercyclical response of inventories adversely impact the response of the other macro variables relative to the model without inventories? To see this, combining the labor supply and demand conditions, equations (25), (14) and (15) yields

$$\left(1 + \frac{1}{\xi}\right)\hat{n}_t = A\hat{s}_t + B\hat{x}_t - (1 - \delta_x)\frac{x}{y}\hat{x}_{t-1}. \quad (26)$$

where $A = \frac{s}{y} + \frac{\Phi}{\tau} > 0$ and $B = \frac{x}{y} - \frac{\Phi}{\tau} > 0$ for $\frac{s}{a} < 1$.³⁰ Since x_{t-1} is predetermined, the response of hours varies directly with the relative movement of contemporaneous sales and inventories. As such, for a given increase in sales in response to TFP news, reductions in inventories will depress the response of hours. Intuitively, despite the increase in labor demand from the rise in utilization due to the investment adjustment costs per the Jaimovich-Rebelo mechanism, distributors can reduce their demand for produced goods (relative to the model without inventories), since they can meet some of the demand for sales by drawing down inventories, which in turn reduces the demand for labor and capacity utilization as inputs into production. The fall in inventories is thus intimately linked to the muted response of hours, which then leads to a muted response in output and utilization. As we can see from Figure 9, over the first several periods, sales rises and inventory falls, yet the response of hours, utilization and output are nearly flat. This result is consistent with the suppressing effect of inventories on hours.

Response to news under sticky prices. Figure 10 (dashed line) shows the response of the standard model to TFP news for the case of sticky prices and flexible wages. As is

³⁰The restriction $\frac{s}{a} = \frac{1}{1+x/s} < 1$ is supported by the data where the inventory to sales ratio, x/s , is reported to be consistently above unity.

clear from figure, while adding sticky prices enhances the overall response of the other macro variables, it does not change the countercyclical response of inventories. Indeed, introducing price stickiness actually causes inventories to drop further in the initial period.

With sticky prices (and flexible wages), the behaviour of inventories is governed by equation (24) such that the intertemporal substitution channel is present. Relative to the case with flexible prices, note that both sales and marginal costs rise much more in the initial period relative to the case with flexible prices and wages, as the TFP shock shifts the distributor's demand for output more so than in the flexible price economy, consistent with the countercyclical markup effect of the price-setters with sticky prices. The positive growth in marginal costs over the first eight or so periods implies a positive intertemporal substitution effect that puts upwards pressure on inventories in the earlier periods, but this is not enough to overcome the initial rise in marginal costs due to the larger rise in current marginal costs associated with the typical countercyclical markup channel operative under sticky prices.³¹ Thus, despite the presence of an intertemporal substitution channel (that puts upward pressure on inventories), the effect of current marginal costs on inventories is dominant.³²

Response to news under sticky wages. Figure 10 (dashed-dotted line) shows the response of the standard model to TFP news for the case of flexible prices and sticky wages. Sticky wages on their own similarly enhance the overall response of the other macro variables as with sticky prices on their own. Yet without sticky prices, the intertemporal substitution channel is shut down, and the behaviour of inventories is governed by equation (25) such that the response of inventory depends only on current sales and current marginal costs. Importantly, note that with sticky wages marginal costs rise much less than with sticky prices, despite a slightly larger initial increase in sales. Inventories as a result are actually able to rise above the steady state in the first period to the combined effect of the amplification of

³¹Recall from the discussion of equation (25) that the coefficient on current marginal costs is larger than that on future marginal costs).

³²It is possible to obtain a fall in marginal costs in the future relative to the present using an alternative calibration of the monetary policy rule. In this case, the substitution effect puts downward pressure on inventories, and thus inventories still fall.

sales and suppression of marginal costs. Nevertheless, beyond the initial period, inventories drop below steady state, moving countercyclically.

5 A possible solution to the comovement puzzle

Our discussion thus far has demonstrated the challenges involved in attempting to the address the comovement puzzle in the standard model.³³ Our theoretical analysis suggests that increasing the response of inventories requires both enhancing the response of sales and suppressing the reaction of marginal costs in the periods leading up to the shock. This suppression of marginal costs is also consistent with our conditional empirical evidence in sections 2.4.3 and 2.5, showing no large initial increases in marginal costs in response to news, as well as an only gradual rise in the real wage. It is also consistent with empirical evidence for example in Anderson et al. (2018) against significant movement in markups over the business cycle. In the following, we propose one way of extending the standard model that results in an alternated news shock transmission with a suppressed response of marginal costs and a enhanced response of sales which help to resolve the comovement problem.

5.1 Extended model with knowledge capital

We extend the standard model by modifying the homogeneous goods production technology to include an additional input to production in the form of intangible capital that we refer to as *knowledge capital*. Knowledge capital evolves over time as a learning-by-doing process, whereby the firm acquires new technological knowledge through its experiences in engaging labor in the production process. While learning-by-doing as a modeling mechanism has had a long history in studying long-run issues such as growth, e.g. in Arrow (1962), it also has a history in its role as a propagation mechanism in business cycle models, including

³³For flexible price versions of this model, this has also been discussed by Crouzet and Oh (2016), who also show the countercyclicality of inventories is independent of whether this margin is introduced in the Jaimovich and Rebelo (2009) type framework, a stock-elastic demand model as in Bils and Kahn (2000) or the stockout-avoidance model of Kryvtsov and Midrigan (2013). They show that in all cases the dynamics of inventories are governed by an expression that is nested in our equation (24).

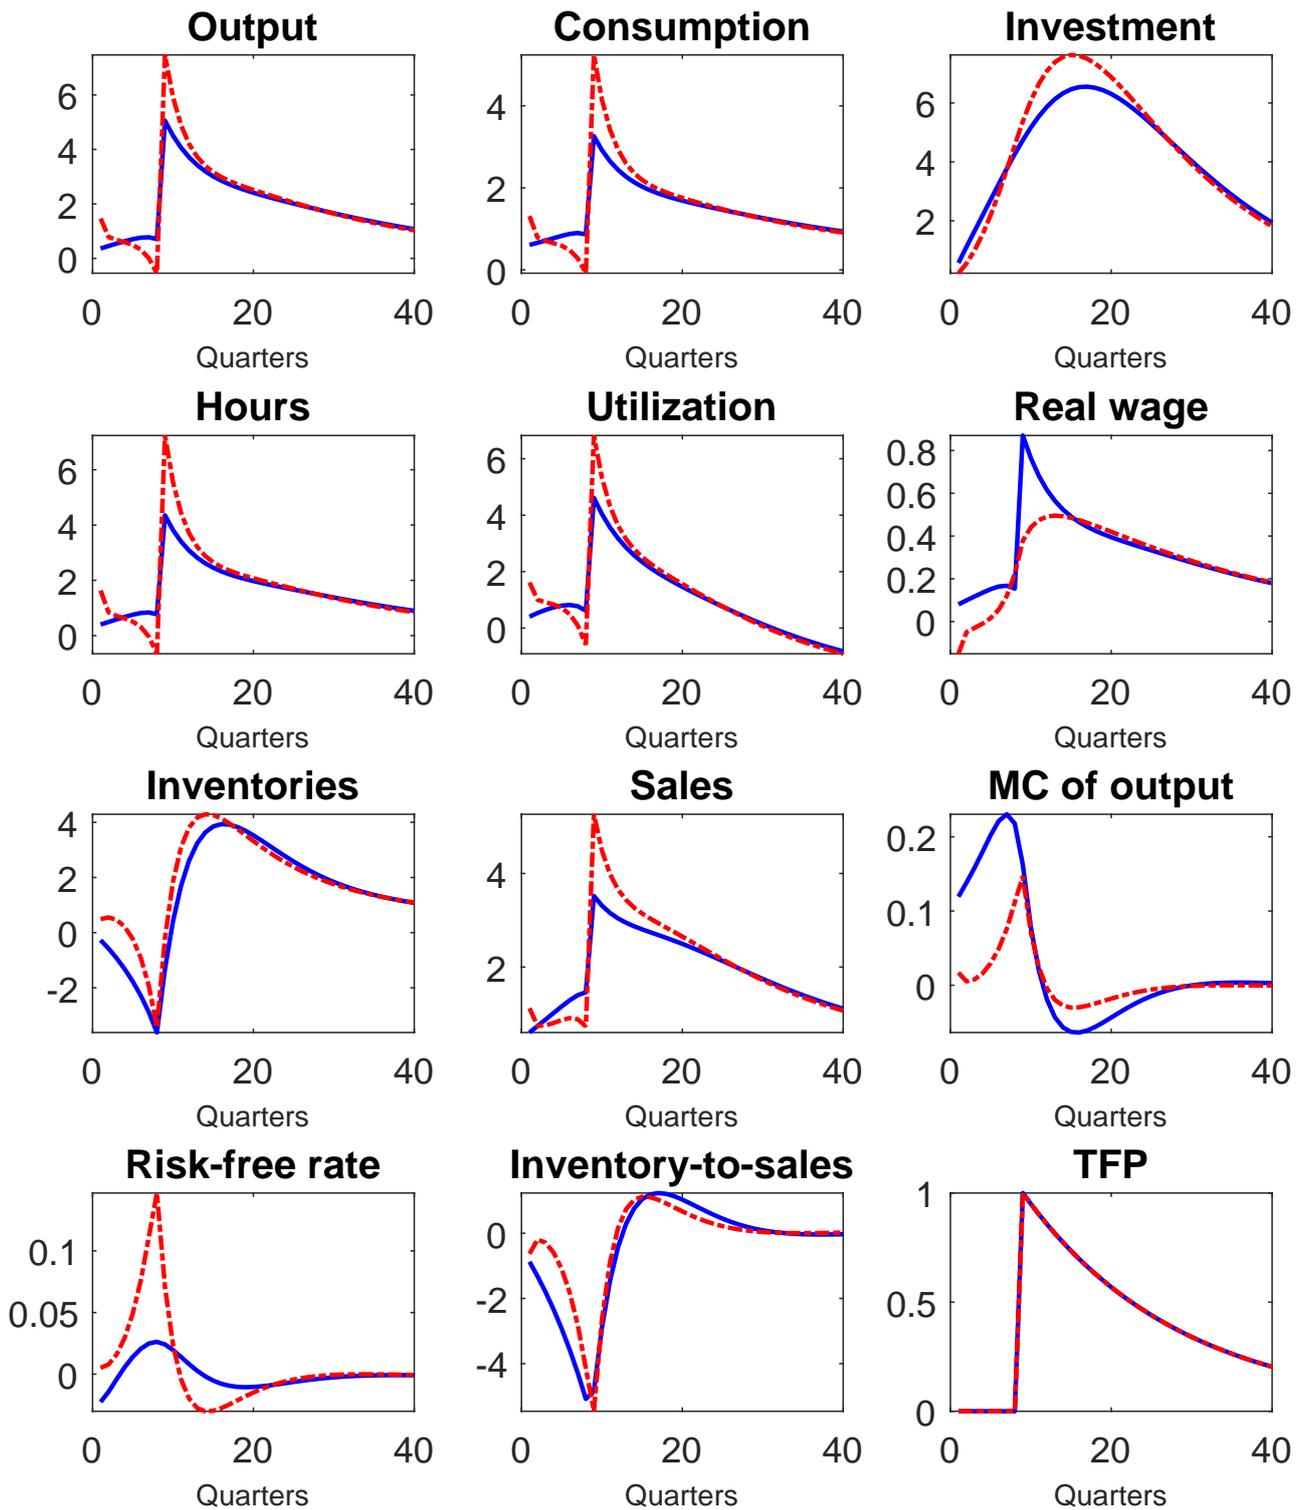


Figure 10: **IRF to TFP news shock - Standard model with nominal rigidities.** Units are percentage deviations from steady state. Solid line: Standard model with sticky prices and flexible wages; Dashed-dot line: Standard model with flexible prices and sticky wages.

Chang et al. (2002), Cooper and Johri (2002) and Gunn and Johri (2011).

The extension has a distinct advantage in the present application in terms of its parsimony. Relative to the standard model, the modification only impacts the specification of the intermediate goods firm. All other aspects of the standard model remain unchanged, including the primary inventory mechanisms. Thus, in understanding what gives rise to positive co-movement of inventory in response to TFP news, we can still frame our analysis through the same inventory equation (24) and supporting framework for analysis that we used in the standard model.

5.2 Intermediate Goods Firm in extended model with knowledge capital

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

$$y_t = z_t n_t^\alpha \tilde{k}_t^{\theta_y} h_t^{\epsilon_y}, \quad (27)$$

where h_t is *knowledge capital*. The definition of all other variables is as in the intermediate goods problem in the standard model in section 3. The firm's stock of knowledge capital evolves according to

$$h_{t+1} = h_t^{\gamma_h} n_t^{\nu_h} N_t^{\tau_h}, \quad (28)$$

where n_t and N_t are the internalized and economy-wide (externalized) effects of labor in knowledge capital accordingly. The log-linear specification of (28) is similar to that of Chang et al. (2002) and Cooper and Johri (2002).³⁴

The firm's problem involves choosing k_{t+1} , h_{t+1} , i_t and n_t to maximize $E_0 \sum_{t=0}^{\infty} \frac{\beta^t \lambda_t}{\lambda_0} \Pi_t^y$ subject to its physical capital constraint (9) and its knowledge capital constraint (28). Rel-

³⁴Numerous other specifications of knowledge/intangible capital accumulation that have been explored in the literature are possible here. Variations include linear versus log-linear accumulation, as well as altering the assumption of by-product nature of learning-by-doing to one where firms must instead transfer resources out of production in order to produce intangible capital. See the discussion in Cooper and Johri (2002), as well as applications in McGrattan and Prescott (2010), Gunn (2015) and Hou and Johri (2018). The advantage of the specification in (28) is its analytical simplicity for our illustrative example.

ative to the standard model, the first-order condition with respect to n_t is modified and the first-order condition with respect to h_{t+1} is new. Defining q_t^h as the Lagrange multiplier on (28), these are given by

$$w_t = \tau_t \alpha \frac{y_t}{n_t} + q_t^h \nu_h \frac{h_{t+1}}{n_t} \quad (29)$$

$$q_t^h = \beta E_t \frac{\lambda_t}{\lambda_{t+1}} \left\{ \epsilon_y \tau_{t+1} \frac{x_{t+1}}{h_{t+1}} + q_{t+1}^h \gamma_h \frac{h_{t+1}}{h_t} \right\}. \quad (30)$$

Note in (29) that the presence of knowledge capital in the firm's technology adds an additional term into the firm's hours-worked first order condition that serves to shift labor demand. All else equal, a rise in the value of knowledge capital, q_t^h , increases labor demand as the firm attempts to increase its knowledge capital. Equation (30) then describes q_t^h as a function of the expected discounted value of the marginal product of that knowledge capital in production next period and the continuation value of that knowledge capital.

Additionally, as in the standard model, define the marginal cost of production, mc_t , for the intermediate goods firm as $mc_t = \frac{w_t}{MPN_t}$. From the firm's labor first-order condition (29), it then follows that

$$\tau_t = mc_t - q_t^h \frac{\nu_h}{\alpha} \frac{h_{t+1}}{y_t}, \quad (31)$$

such that the presence of knowledge capital drives a wedge between the output price τ_t (marginal cost of *output*) and the marginal cost of production mc_t . When the value of knowledge q_t^h is high, the firm increases hours-worked in order to increase knowledge, thereby in turn increasing output y_t for a given τ_t , and in effect, shifting the firm's output supply curve outwards.

5.3 Calibration

All parameters in our previous standard model are also present in our extended knowledge capital model, and thus we used the same values of these parameters, detailed in Panel A of Table 1. For the parameters unique to our knowledge capital model, we set ϵ_h , the knowledge capital elasticity in production, to 0.15 based on Gunn and Johri (2011), and consistent with

“learning rates” in the order of 20% measured in the empirical literature on learning-by-doing, as discussed by Cooper and Johri (2002). We then set γ_h , the contribution of prior knowledge in its own production, to 0.8, based on Gunn and Johri (2011), and then set ν_h , the elasticity of internalized labor in knowledge, to 0.15 in order to yield zero economic profits in steady state for the production firm. Finally, we impose that the internalized and externalized elasticities of labor in knowledge capital are equal, yielding $\tau_h = \nu_h = 0.8$. Panel B of Table 1 summarizes the parameter values unique to the knowledge capital model.

5.4 Response to news in extended model with knowledge capital

We now investigate the simulated response to TFP news in the extended model.

Response to news with sticky wages and flexible prices. Figure 11 shows the response of the knowledge capital model *with inventory* with sticky wages and flexible prices to TFP news. In contrast to the standard model, inventories now increase in response to news. Moreover, the response of the main macroeconomic variables like hours-worked, consumption and sales is enhanced relative to the standard model. Consistent with our VAR results, the real wage increases gradually, and the initial rise in the marginal cost of output τ_t is muted relative to the overall expansion.

To understand how this knowledge capital mechanism produces procyclical inventory movements, it is helpful to first understand how the mechanism drives a boom in hours-worked and output. Since knowledge capital is a complement in production to labor and physical capital, when TFP news arrives, the firm knows that in the future, the direct positive impact of the TFP shock on labor and capital will also increase the return on knowledge in the future, increasing the value of knowledge in the future. By virtue of being an accumulated stock, this in turn raises the value of knowledge capital in the present, evidenced by the immediate increase in the shadow-price of knowledge in the present, leading in turn to an outward shift in labour demand. In other words, knowing that there will be technological change in the future, the firm begins preparing for that change in the present, building up its knowledge in the present by engaging labour in order to respond optimally when the change

comes in the future.³⁵

How then does this lead to an increase in inventory? Since the firm has an incentive to immediately increase its hours worked in response to news, it is foregoing profits in the present in anticipation of more in the future, and therefore is willing to sell a greater quantity of homogenous goods to the distributors for a given price τ_t . In effect, in the market for output goods in τ_t - y_t space, this shifts the intermediate goods firm's output supply curve to the right, putting downward pressure on the price of output τ_t . All else equal, the distributors purchasing these goods respond to this drop in τ_t by increasing their purchases of these homogenous goods. In tandem with this increase supply however is the increase in demand for sales due to the wealth effect of the TFP shock on consumption that was present in the standard model also. The net effect then is both an outward shift in the supply and demand curves in the market for output. This has the benefit of reducing the equilibrium price τ_t associated with a given rise in sales, thereby increasing the chances that inventory can rise along with sales in maintain the optimal sales to stock ratio. For a given rise in sales, the extent to which output increases in equilibrium depends on the relative steepness of the supply curve, and the relative magnitude of the supply shift. The equilibrium increase in output relative to the increase in sales then in turn determines whether inventories will rise or fall.

Note that in the initial period upon the receipt of news, since the stock's of k_t and h_t are predetermined, the firm still actually faces the same mechanism driving the marginal cost of production as the firm without knowledge capital in the standard model. The difference now however is that the firm is willing to sell this production at a lower price than it would otherwise in the standard model. In periods $t + 1$ and beyond however, knowledge accumulates at a faster pace, increasing the stock of knowledge in production, directly lowering the firm's marginal costs of production.

³⁵See Gunn and Johri (2011) for additional discussion.

5.5 Robustness

Table 2 shows the robustness of the knowledge capital model's inventory comovement properties to changes in several key parameters. For each parameter in question, the table shows the upper and lower bounds of values that generate comovement on impact of output, consumption, investment, hours and inventories. In general, the most sensitive parameters for the response of inventories on impact for the baseline calibration are the depreciation of the elasticity of capital utilization and the Calvo wage stickiness parameter. For higher values of the former and lower values of the latter, current marginal costs rise more initially, and inventories fall below steady state on impact. Importantly, this finding about the importance of the role of these two parameters in suppressing marginal costs is consistent with findings elsewhere in the DSGE literature for models without inventory. For example, Christiano et al. (2017) discuss the importance of variable capacity utilization and sticky wages for ensuring marginal costs are relatively a-cyclical, a feature that helps the models address various features of the data such as the response of inflation to demand shocks. There are variants of the current calibration that allow comovement of inventories under a larger range of sticky wages, and indeed, under fully flexible wages and prices or various combinations of nominal rigidities. We view our exercise as being illustrative however, and leave more complete explorations of parameter spaces preferred by the data to future work.

Table 2: Parameter range for comovement of all macroeconomic aggregates

	lower bound	baseline calibration	upper bound
Investment adjustment cost	1.01	2	24.85
Depreciation elasticity of capital utilization	0	0.05	0.85
Wage Calvo probability	0.67	0.85	0.92
Wage indexation	0	0.45	1
Frisch elasticity of labour supply	0.57	5	5.92*

Table entries show the upper and lower bounds of parameter values that generate comovement on impact of output, consumption, investment, hours and inventories. We vary one parameter value at a time and keep the others fixed at baseline values. * indicates that Blanchard-Kahn conditions are violated above the upper range.

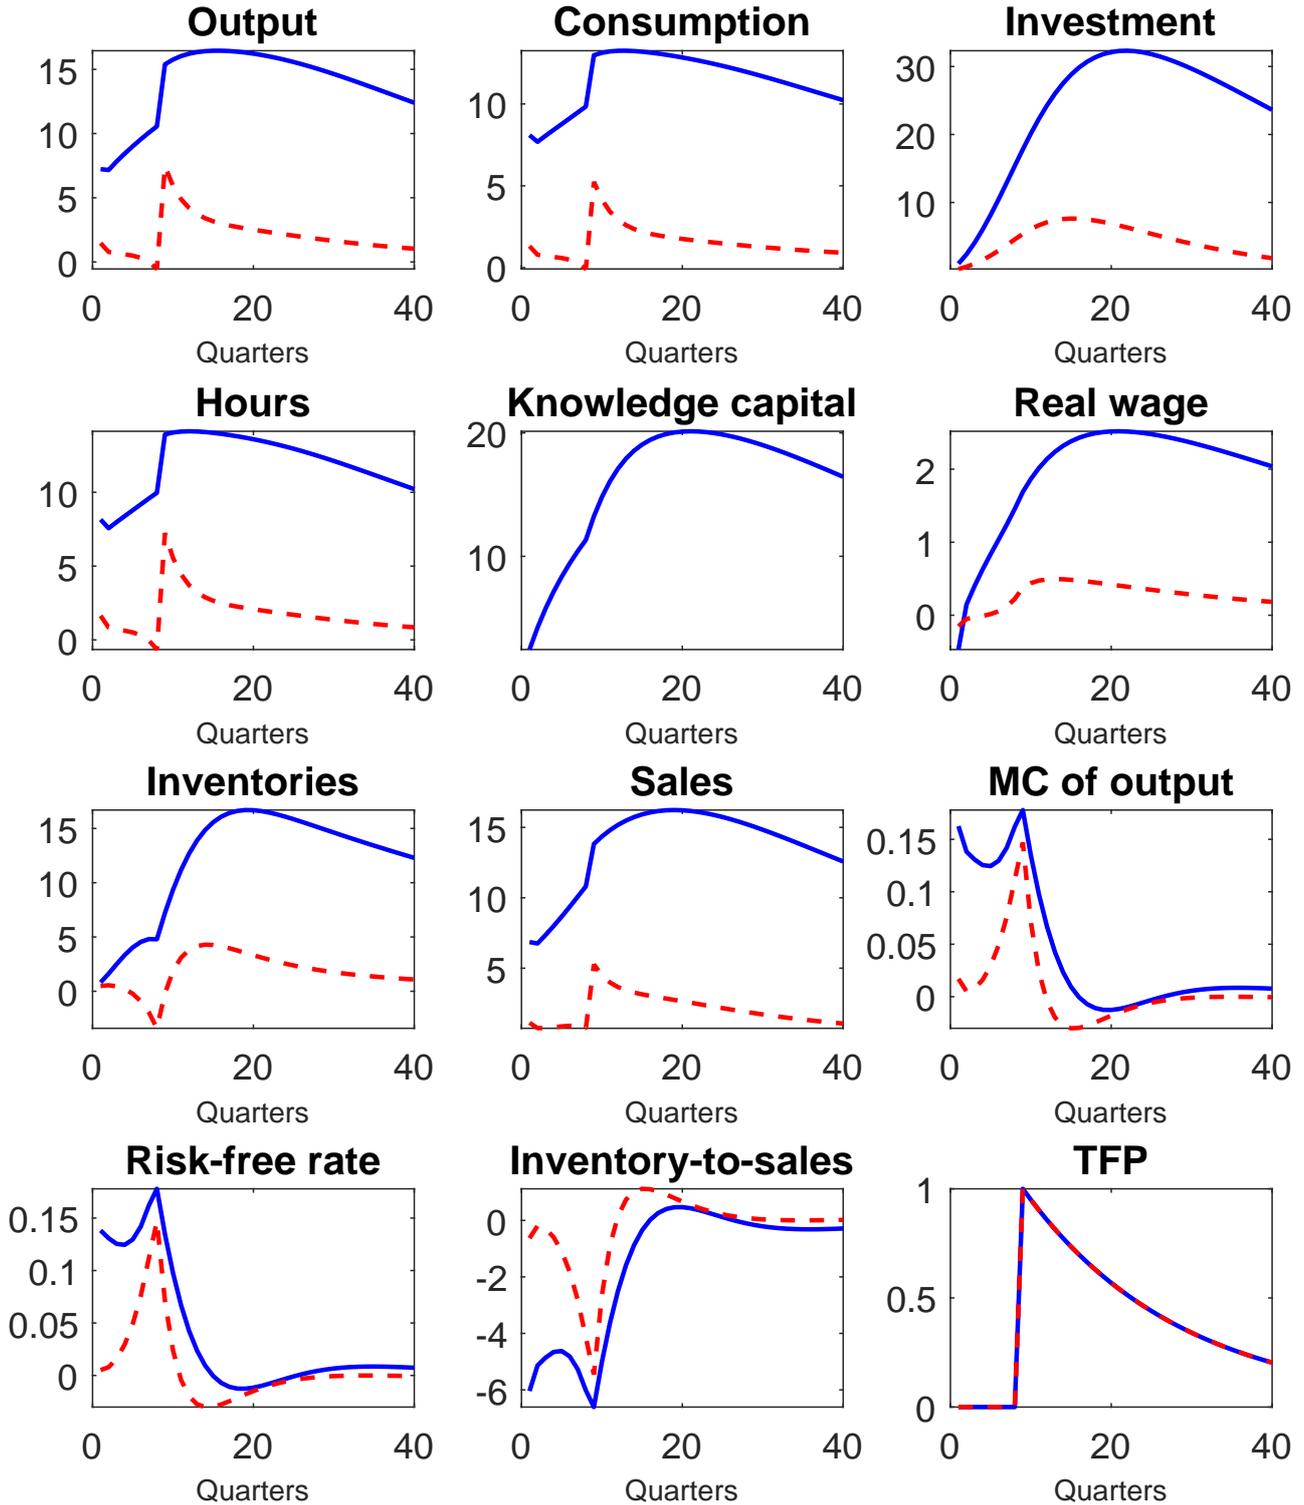


Figure 11: **IRF to TFP news shock - Knowledge capital model with flexible prices and sticky wages.** Units are percentage deviations from steady state. Solid line: Knowledge capital model with flexible prices and sticky wages; Dashed line: Standard model with flexible prices and sticky wages.

6 Conclusion

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

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

References

- Anderson, Eric, Sergio Rebelo, and Arlene Wong**, “Markups Across Space and Time,” Working Paper 24434, National Bureau of Economic Research March 2018.
- Arrow, K.J.**, “The Economic Implications of Learning by Doing,” *Review of Economic Studies*, 1962, 3 (29), 155–173.
- Barsky, Robert B. and Eric R. Sims**, “News shocks and business cycles,” *Journal of Monetary Economics*, 2011, 58 (3), 273–289.
- **and** –, “Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence,” *American Economic Review*, June 2012, 102 (4), 1343–77.
- Beaudry, Paul and Franck Portier**, “An exploration into Pigou’s theory of cycles,” *Journal of Monetary Economics*, September 2004, 51 (6), 1183–1216.
- **and Frank Portier**, “News Driven Business Cycles: Insights and Challenges,” *Journal of Economic Literature*, 2014, 52 (4), 993–1074.
- Ben Zeev and Hashmat Khan**, “Investment Specific News Shocks and U.S. Business Cycles,” *Journal of Money, Credit and Banking*, December 2015, 47 (8), 1443–1464.
- Bils, Mark and James A Kahn**, “What inventory behavior tells us about business cycles,” *American Economic Review*, 2000, 90 (3), 458–481.
- Blinder, A. S.**, “Retail Inventory Behavior and Business Fluctuations,” *Brookings Papers on Economic Activity*, 1981, 12 (1981-2), 443–520.
- **and L. J. Maccini**, “Taking Stock: A Critical Assessment of Recent Research on Inventories,” *Journal of Economic Perspectives*, Winter 1991, 5 (1), 73–96.
- Chang, Yongsung, Joao F. Gomes, and Frank Schorfheide**, “Learning-by-Doing as a Propagation Mechanism,” *American Economic Review*, December 2002, 92 (5), 1498–1520.

- Chirinko, Robert S.**, “ σ : The long and short of it,” *Journal of Macroeconomics*, 2008, 30 (2), 671 – 686.
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans**, “Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy,” *Journal of Political Economy*, February 2005, 113 (1), 1–45.
- Christiano, Lawrence J., Martin S Eichenbaum, and Mathias Trabandt**, “On DSGE Models,” 2017.
- Christiano, Lawrence, Roberto Motto, and Massimo Rostagno**, “Monetary policy and stock market boom-bust cycles,” Working Paper Series 955, European Central Bank October 2008.
- Cooper, Russell and Alok Johri**, “Learning-by-doing and aggregate fluctuations,” *Journal of Monetary Economics*, 2002, 49 (8), 1539–1566.
- Crouzet, Nicolas and Hyunseung Oh**, “What do inventories tell us about news-driven business cycles?,” *Journal of Monetary Economics*, 2016, 79, 49–66.
- Easton, Peter D.**, “PE Ratios, PEG Ratios, and Estimating the Implied Expected Rate of Return on Equity Capital,” *The Accounting Review*, 2004, 79 (1), 73–95.
- Erceg, Christopher J, Dale W Henderson, and Andrew T Levin**, “Optimal monetary policy with staggered wage and price contracts,” *Journal of monetary Economics*, 2000, 46 (2), 281–313.
- Fernald, J.**, “A quarterly, utilization–adjusted series on total factor productivity,” *Working Paper*, 2014, (2012-19).
- Forni, M., L. Gambetti, and L. Sala**, “No news in business cycles,” *Economic Journal*, 2014, 124, 1168–1191.

- Francis, N., M. Owyang, J. Roush, and R. DiCecio**, “A Flexible Finite-Horizon Alternative to Long-run Restrictions with an Application to Technology Shocks,” *Review of Economics and Statistics*, 2014, *96*, 638–647.
- Galí, J. and L. Gambetti**, “On the sources of the Great Moderation,” *The American Economic Journal: Macroeconomics*, 2009, *1*, 26–57.
- Galí, Jordi, Mark Gertler, and J. David López-Salido**, “Markups, Gaps, and the Welfare Costs of Business Fluctuations,” *The Review of Economics and Statistics*, 2007, *89* (1), 44–59.
- Gebhardt, William R., Charles M. C. Lee, and Bhaskaran Swaminathan**, “Toward an Implied Cost of Capital,” *Journal of Accounting Research*, 2001, *39* (1), 135–176.
- Gordon, Myron J. Gordon Joseph R.**, “The Finite Horizon Expected Return Model,” *Financial Analysts Journal*, 1997, *53* (3), 52–61.
- Görtz, C. and J. Tsoukalas**, “News and Financial Intermediation in Aggregate Fluctuations,” *Review of Economics and Statistics*, 2017, *99* (3), 514–530.
- Görtz, Christoph and John Tsoukalas**, “Sectoral TFP News Shocks,” *Economics Letters*, 2018, *forthcoming*.
- , – , and **Francesco Zanetti**, “News Shocks under Financial Frictions,” Technical Report 2017.
- Greenwood, J., Z. Hercowitz, and G. Huffman**, “Investment, Capacity Utilization, and the Real Business Cycle,” *The American Economic Review*, 1988, *78*, 402–217.
- Gunn, Christopher and Alok Johri**, “News and Knowledge Capital,” *Review of Economic Dynamics*, 2011, *14* (1), 92–101.
- Gunn, Christopher M.**, “Animal spirits as an engine of boom-busts and throttle of productivity growth,” *Journal of Economic Dynamics and Control*, 2015, *57*, 24–53.

- Hou, Keqiang and Alok Johri**, “Intangible capital, the labor wedge and the volatility of corporate profits,” *Review of Economic Dynamics*, 2018, *29*, 216 – 234.
- Hou, Kewei, Mathijs A. van Dijk, and Yinglei Zhang**, “The implied cost of capital: A new approach,” *Journal of Accounting and Economics*, 2012, *53* (3), 504 – 526.
- Hribar, Paul and Daniel W. Collins**, “Errors in Estimating Accruals: Implications for Empirical Research,” *Journal of Accounting Research*, 2002, *40* (1), 105–134.
- Jaimovich, Nir and Sergio Rebelo**, “Can News about the Future Drive the Business Cycle?,” *American Economic Review*, 2009, *99* (4), 1097–1118.
- Jones, Christopher S. and Selale Tuzel**, “New Orders and Asset Prices,” *Review of Financial Studies*, 2013, *26* (1), 115–157.
- Jung, Yongseung and Tack Yun**, “Monetary policy shocks, inventory dynamics, and price-setting behavior,” *Federal Reserve Bank of San Francisco Working Paper Series*, 2006, *2006-02*.
- Jung, YongSeung and Tack Yun**, “Inventory investment and the empirical Phillips curve,” *Journal of Money, Credit and Banking*, 2013, *45* (1), 201–231.
- Kesavan, Saravanan, Vishal Gaur, and Ananth Raman**, “Do Inventory and Gross Margin Data Improve Sales Forecasts for U.S. Public Retailers?,” *Management Science*, 2010, *56* (9), 1519–1533.
- Khan, Hashmat and John Tsoukalas**, “The Quantitative Importance of News Shocks in Estimated DSGE Models,” *Journal of Money, Credit and Banking*, December 2012, *44* (8), 1535–1561.
- Kryvtsov, Oleksiy and Virgiliu Midrigan**, “Inventories, Markups, and Real Rigidities in Menu Cost Models,” *The Review of Economic Studies*, 2013, *80* (1), 249–276.

- Kurmann, Andre and Eric Sims**, “Revisions in Utilization-Adjusted TFP and Robust Identification of News Shocks,” mimeo 2016.
- Lettau, Martin and Sydney Ludvigson**, “Consumption, Aggregate Wealth, and Expected Stock Returns,” *The Journal of Finance*, 2001, *56* (3), 815–849.
- Lubik, Thomas A and Wing Leong Teo**, “Inventories, inflation dynamics and the New Keynesian Phillips curve,” *European Economic Review*, 2012, *56* (3), 327–346.
- McCarthy, Jonathan and Egon Zakrajsek**, “Inventory Dynamics and Business Cycles: What Has Changed?,” *Journal of Money, Credit and Banking*, 03 2007, *39* (2-3), 591–613.
- McGrattan, Ellen R and Edward C Prescott**, “Unmeasured Investment and the Puzzling US Boom in the 1990s,” *American Economic Journal: Macroeconomics*, 2010, *2* (4), 88–123.
- Nekarda, Christopher J. and Valerie A. Ramey**, “The Cyclical Behavior of the Price-Cost Markup,” NBER Working Papers 19099, National Bureau of Economic Research, Inc June 2013.
- Ohlson, James A. and Beate E. Juettner-Nauroth**, “Expected EPS and EPS Growth as Determinants of Value,” *Review of Accounting Studies*, 2005, *10* (2-3), 349–365.
- Pavlov, Oscar and Mark Weder**, “Countercyclical Markups and News-Driven Business Cycles,” *Review of Economic Dynamics*, April 2013, *16* (2), 371–382.
- Ramey, V. A. and K. D. West**, “Inventories,” in J. B. Taylor and M. Woodford, eds., *Handbook of Macroeconomics*, Vol. 1 of *Handbook of Macroeconomics*, Elsevier, December 1999, chapter 13, pp. 863–923.
- Rotemberg, Julio J. and Michael Woodford**, “The cyclical behavior of prices and costs,” in J. B. Taylor and M. Woodford, eds., *Handbook of Macroeconomics*, Vol. 1 of *Handbook of Macroeconomics*, Elsevier, 1999, chapter 16, pp. 1051–1135.

- Sarte, Pierre-Daniel, Felipe Schwartzman, and Thomas A. Lubik**, “What inventory behavior tells us about how business cycles have changed,” *Journal of Monetary Economics*, 2015, *76*, 264 – 283.
- Schmitt-Grohe, Stephanie and Martin Uribe**, “Optimal inflation stabilization in a medium-scale macroeconomic model,” *NBER Working Papers*, 2005, (11854).
- Schmitt-Grohé, Stephanie and Martin Uribe**, “Comparing two variants of Calvo-type wage stickiness,” Technical Report, National Bureau of Economic Research 2006.
- Schmitt-Grohe, Stephanie and Martin Uribe**, “What’s news in business cycles?,” *Econometrica*, November 2012, *80* (6), 2733–2764.
- Smets, Frank and Rafael Wouters**, “Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach,” *American Economic Review*, 2007, *97* (3), 586–606.
- Vukotic, Marija**, “Sectoral Effects of News Shocks,” mimeo, University of Warwick January 2016.
- Wen, Y.**, “Understanding the inventory cycle,” *Journal of Monetary Economics*, November 2005, *52* (8), 1533–1555.

7 Appendix

A Additional VAR evidence

A.1 Forecast Error Variance Decomposition

Figure 12 displays the variance shares explained by the TFP news shock.

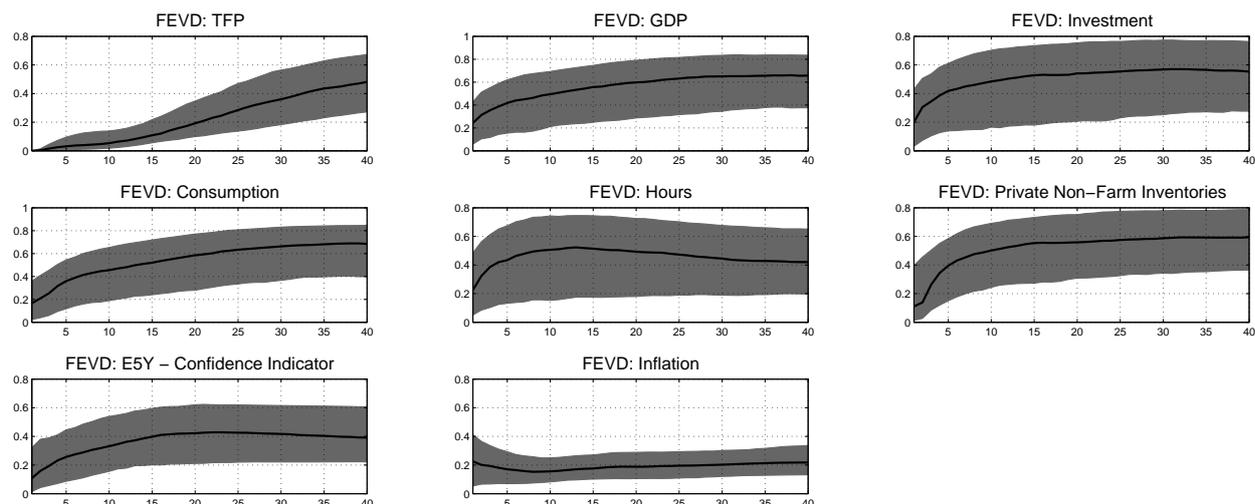


Figure 12: Forecast error variance decomposition (FEVD) of variables to the TFP news shock (solid line – median). Sample 1985Q1-2015Q2. The gray shaded areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters.

A.2 Additional evidence related to marginal cost measures

This section provides additional evidence to the discussion of the response of marginal cost measures in Section 2.4.3.

Figure 13 shows the response of two marginal cost measures to a TFP news shock when they are included one-by-one in an eight-variable VAR. The two marginal cost measures in the figure are constructed using the preferred measure for the labor share by Galí et al. (2007), the BLS labor share in the non-farm business sector, and are either based on the CES (CES: Gali et al.) or Cobb-Douglas (CD: Gali et al.) production function. Qualitatively and quantitatively the responses of these two marginal cost measures to a TFP news shock are

very similar to the responses shown in Figure 7 when using the labor share measure preferred by Nekarda and Ramey (2013) (CES: Nekarda-Ramey 1, CD: Nekarda-Ramey 1). In line with the discussion in the main body, neither of the two marginal cost measures in Figure 13 provides evidence for a strong negative substitution effect through a fall in marginal costs. This is consistent with the rise in inventories we report in response to a TFP news shock driven by a positive demand effect dominating the negative substitution effect.

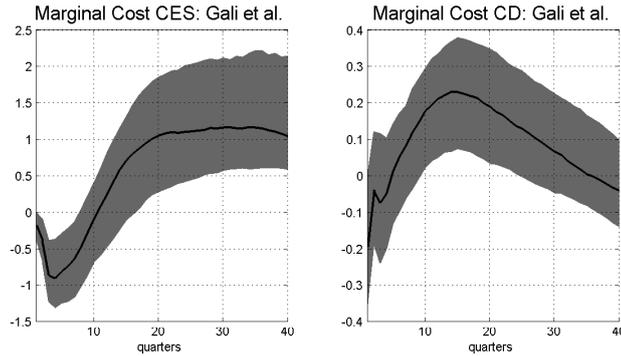


Figure 13: **IRF of marginal cost measures to TFP news shock.** Sample 1985Q1-2015Q2. Subplots result from eight variable VARs comprising TFP, GDP, consumption, hours, inventories, marginal cost measure, inflation, E5Y. The marginal cost measures were included one-by-one in the VAR system. The shaded gray areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

Table 3 shows the unconditional correlations of HP-filtered GDP with all our considered measures for marginal costs. Marginal costs are acyclical or mildly countercyclical which is in line with the evidence in Nekarda and Ramey (2013) who report that markups are acyclical or mildly procyclical. In addition to the abbreviations explained in the paragraph above, CD: Nekarda-Ramey 2 and CES: Nekarda-Ramey 2 refer to the marginal cost measures constructed considering a measure for overhead labor (as suggested by Nekarda and Ramey (2013)) under the assumption of either a Cobb-Douglas or a CES production function. Details of the construction are outlined in Section 2.4.3 in the main body.

The results shown in Figures 7 and 13 are robust to variations of the elasticity of substitution between capital and labor, σ , in the construction of the marginal cost measures. Chirinko (2008) concludes that the literature estimates σ , in the range of 0.4 and 0.6. (our

Table 3: Correlations of marginal cost measures with GDP

CES: Nekarda-Ramey 1	-0.31
CES: Gali et al.	-0.30
CD: Nekarda-Ramey 1	-0.06
CD: Gali et al.	-0.04
CD: Nekarda-Ramey 2	-0.21
CES: Nekarda-Ramey 2	-0.38

Notes. Time series are HP(1600)-filtered.
Sample is 1985Q1-2015Q2.

baseline calibration is 0.5). Robustness using these two values yields very similar responses of all marginal cost measures to a TFP news shock. Qualitatively they are virtually unchanged. Detailed results are available upon request.

A.3 Robustness to alternative VAR news identification

The results in the main body of the paper are generated using the Max-share method proposed by Francis et al. (2014). This section reports VAR findings using three alternative approaches. First, the identification scheme in Barsky and Sims (2011) that recovers the news shock by maximizing the variance of TFP over the horizons zero to 40 quarters, and the restriction that the news shock does not move TFP on impact. Second, the identification scheme in Kurmann and Sims (2016), that recovers the news shock by maximizing the FEV of TFP at a very long horizon (60 quarters) without however imposing the zero impact restriction on TFP conditional on the news shock.³⁶ Third, the Forni et al. (2014) long-run identification scheme which is similar in spirit to the Max Share method and has been used in an application with news shocks. The latter method identifies the news shock by imposing the zero impact restriction on TFP, and seeks to maximise the impact of the news shock on TFP in the long run.

Figure 14 shows, the median responses between the Max share method and the methods

³⁶These authors argue that allowing TFP to jump freely on impact, conditional on a news shock, produces robust inference to cyclical measurement error in the construction of TFP.

proposed by Barsky and Sims (2011) and Forni et al. (2014) are virtually indistinguishable. Figure 15 also shows responses based on the methodology proposed by Kurmann and Sims (2016) are qualitatively and quantitatively very similar to the ones based on the Max-share method. Importantly, all methods suggest inventories increase in anticipation of higher future TFP.

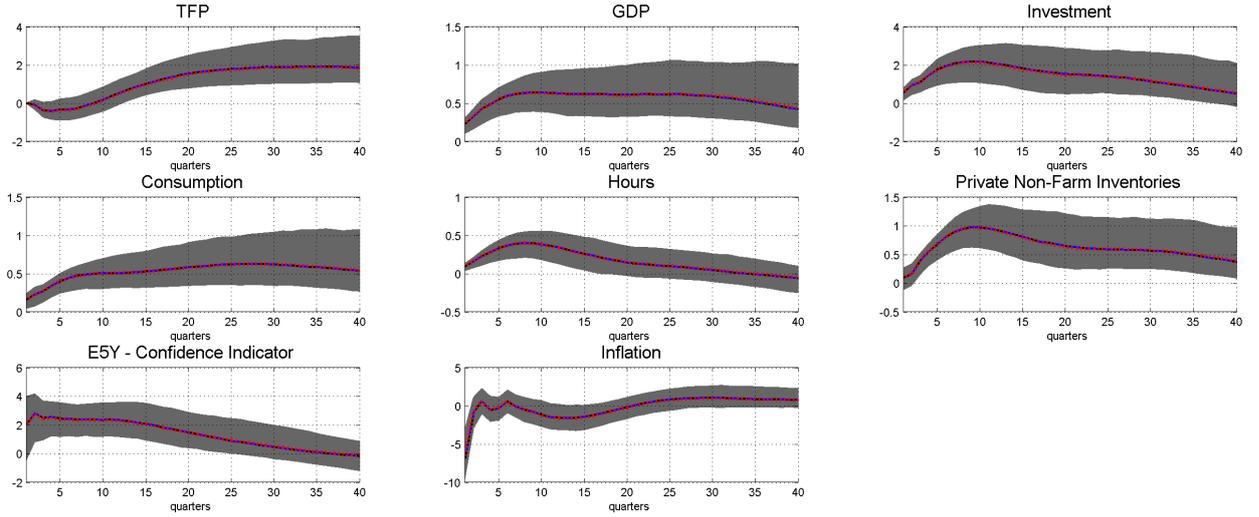


Figure 14: **IRF to TFP news shock**. Sample 1985Q1-2015Q2. The black solid line is the median response identified using the Max-share method. The shaded gray areas are the corresponding 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The blue dashed (red dotted) line is the median response identified using the Barsky and Sims (2011) (Forni et al. (2014)) methodology. The units of the vertical axes are percentage deviations.

B Implied Cost of Capital Construction

The construction of firm-level ICCs using data from Compustat-CRSP requires a measure for earnings forecasts. Based on Hou et al. (2012), we generate earnings forecasts by estimating the following pooled cross-sectional regression for each quarter from 1985Q1, using the the previous ten years of data

$$E_{i,t+\tau} = \beta_0 + \beta_1 A_{i,t} + \beta_2 D_{i,t} + \beta_3 DD_{i,t} + \beta_4 E_{i,t} + \beta_5 NegE_{i,t} + \beta_6 AC_{i,t} + \varepsilon_{i,t+\tau}.$$

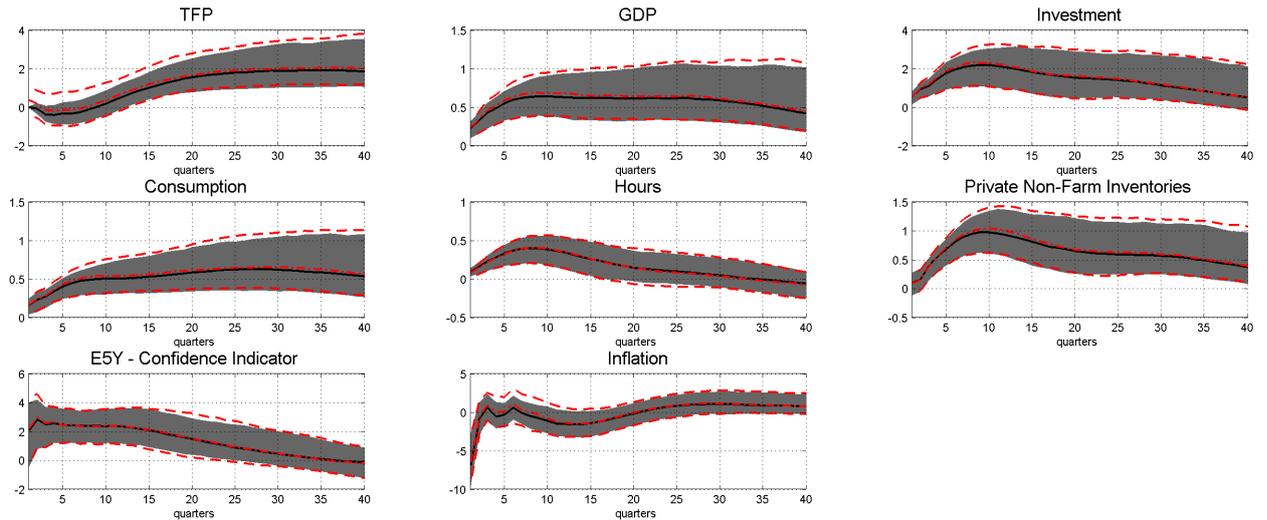


Figure 15: **IRF to TFP news shock**. Sample 1985Q1-2015Q2. The black solid (red dash-dotted) line is the median response identified using the Max-share (Kurmann and Sims (2016)) method. The shaded gray areas (dashed red lines) are the corresponding 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

Here, $E_{i,t+\tau}$ denotes earnings of firm i at time $t+\tau$. Earnings in Compustat is Income Before Extraordinary Items (IBQ). $NegE_{i,t}$ is a dummy variable that equals one for firms with negative earnings and zero otherwise. $A_{i,t}$ is Total Assets (ATQ). $D_{i,t}$ is dividend payments (DVTQ) and $DD_{i,t}$ is the associated a dummy variable that equals one for dividend payers and zero otherwise. $AC_{i,t}$ is accruals. These are calculated in our dataset as change in Current Assets (ACTQ) minus change in Current Liabilities (LCTQ) minus change in Cash and Short-Term Investments (CHEQ) plus change in Debt in Current Liabilities (DLCQ) minus Depreciation and Amortization (DPQ) according to Hribar and Collins (2002).

We closely follow the methodology outlined in Hou et al. (2012), Appendix A, to generate the four different ICC measures, which is why we do not repeat the detailed equations for the ICC measures here.³⁷ Note that for this purpose we merge the Compustat data with information from CRSP on market equity (MVAL) defined as the product of Number of Shares Outstanding (CSHO) and the Stock Price at the end of the quarter (PRCC). We further use the 1-Year Treasury Constant Maturity Rate as risk free rate. Prior to computing

³⁷Hou et al. (2012) show these ICC measures based on earnings forecasts outperform the measures when constructed using IBES analyst earnings forecast data.

earnings forecasts and ICC measures we have applied the cleaning procedures outlined in the section below to the Compustat-CRSP data. We follow the convention in the literature and set any ICC estimates below zero to missing. We further set the top one percentile of all firm-time observations for a particular ICC measure to missing prior to aggregating the firm observations by taking averages over each quarter.

Cleaning procedures applied to Compustat-CRSP data:

We select the sample by making the following adjustments to the data retrieved from Compustat-CRSP:

- We delete all regulated, quasi-public or financial firms (primary SIC classification is between 4900-4999 and 6000-6999).
- We delete firms reported earnings in a currency other than USD.
- We account for the effects of mergers and acquisitions by deleting all observations including of firms with (i) acquisitions (ACQ) exceeding 15% of total assets (ATQ), or (ii) sales growth exceeding 50% in any year due to a merger.
- We drop companies with all values for total assets (AT) or investment in plant, property and equipment (CAPX) missing or zero. We drop missing observations for CAPX if they are at the beginning or end of a company's reported data. If CAPX is missing in the middle of a company's reported data we drop the entire company.
- We drop firms with less than three quarters of data.
- We apply the following filters to key variables:
 - We replace missing values of DPQ with zero.
 - We set negative values of CHEQ, DLCQ, DPQ and DVPQ to missing.
 - We set values smaller or equal to zero of ACTQ, LCTQ, ATQ and MVAL to missing.

- We winsorise IBQ at the top and bottom percentile.
- We winsorise ATQ, ACTQ, LCTQ, CHEQ, DLCQ, DPQ, DVPQ and MVAL at the top percentile.
- ATQ, ACTQ, LCTQ, CHEQ, MVAL, DLCQ, IBQ and DPQ are deflated applying the Gross Domestic Product: Implicit Price Deflator. DPQ is deflated applying the Gross Private Domestic Fixed Investment: Nonresidential Implicit Price Deflator.

C Additional model detail

C.1 Standard model

C.1.1 Employment unions and employment agency

Typical sticky-wage frameworks based on Erceg et al. (2000) model households as monopoly suppliers of differentiated labor, introducing equilibrium heterogeneity across households in hours-worked, which can be problematic when household preferences are non-separable in consumption and leisure as we consider in this model. Instead, we follow the decentralization of Schmitt-Grohe and Uribe (2005) and Smets and Wouters (2007) whereby a monopolistic union buys homogenous labor from households, transforms it into a differentiated labor inputs $q \in [0, 1]$, and sells it to the employment agency who aggregates the differentiated labor into a composite which it then sells to intermediate goods producers.³⁸ This particular decentralization of wage stickiness implies that consumption and hours are identical across households.

Labor unions acquire homogenous labor n_t^h from the household at wage W_t^h , differentiate it into labor types n_{jt} , $j \in [0, 1]$, and then sell the differentiated labor it to the employment agency for wage W_{jt} . The unions have market power, and can thus choose the wage for each labor type subject to the labor demand curve for that labor type. The unions face

³⁸Schmitt-Grohe and Uribe (2005) show this decentralization yields a wage Phillips curve that is identical to that from the Erceg et al. (2000) model, up to a log-linear approximation.

Calvo frictions in setting their wages, such that each period they can re-optimize wages with probability $1 - \zeta_w$. A union that is unable to re-optimize wages re-sets it according to the indexation rule $W_{jt} = W_{jt-1} \pi_{t-1}^{\nu_w} \pi^{1-\nu_w}$, $0 \leq \nu_w \leq 1$, where $\pi_t = P_t/P_{t-1}$ and π is its steady state, and where $0 \leq \nu_w \leq 1$. A union that can re-optimize its wage in period t chooses its wage W_{jt}^* to maximize

$$E_t \sum_{s=0}^{\infty} \zeta_w^s \beta^s \frac{\lambda_{t+s} P_t}{\lambda_t P_{t+1}} [W_{jt}^* (\prod_{k=0}^s \pi_{t+k-1}^{\nu_w} \pi^{1-\nu_w}) - W_{t+s}^h] n_{jt+s},$$

subject to the demand curve for n_{jt} .

The employment agency acquires each j th intermediate labor type n_{jt} , $j \in [0, 1]$, at wage W_{jt} from the labor unions, and combines the differentiated labor into a composite n_t according to

$$n_t = \left[\int_0^1 n_{jt}^{\nu_w} dj \right]^{\frac{1}{\nu_w}}, \quad 0 < \nu_w \leq 1.$$

The agency sells the composite labor to the intermediate goods producers for wage W_t . The agency chooses $n_{jt} \forall j$ to maximize profits $W_t n_t - \int_0^1 W_{jt} n_{jt} dj$, yielding a demand function n_{jt} for the j th labor type,

$$n_{jt} = \left[\frac{W_{jt}}{W_t} \right]^{\frac{1}{\nu_w-1}} n_t,$$

and wage index W_t , given respectively by

$$W_t = \left[\int_0^1 W_{jt}^{\nu_w/(\nu_w-1)} dj \right]^{\frac{(\nu_w-1)}{\nu_w}}.$$

C.1.2 Equilibrium

In a symmetric equilibrium, $y_{it} = y_t^*$, $a_{it} = a_t^*$, $x_{it} = x_t^*$, $P_{it} = P_t^*$ and $s_{it} = s_t^* \forall i$, and $W_{jt}^* = W_t^*$, $n_{jt} = n_t^* \forall j$. It then follows that $y_t = \int_0^1 y_t^* di = y_t^*$, $a_t = \int_0^1 a_t^* di = a_t^*$, $x_t = \int_0^1 x_t^* di = x_t^*$, $n_t^h = \int_0^1 n_t^* dj = n_t^*$. Integrating over the taste shifter then yields

$$\int_0^1 \nu_{it} di = \int_0^1 \left(\frac{a_{it}}{a_t} \right)^\zeta dj = \frac{1}{a_t^\zeta} \int_0^1 a_{it}^\zeta di = 1,$$

and hence

$$P_t = \left[\int_0^1 \nu_{it} (P_t^*)^{1-\theta} di \right]^{\frac{1}{1-\theta}} = P_t^*$$

and

$$s_t = \left[\int_0^1 \nu_{it}^{\frac{1}{\theta}} s_t^* \frac{\theta-1}{\theta} di \right]^{\frac{\theta}{\theta-1}} = s_t^*.$$

Similarly,

$$W_t = \left[\int_0^1 W_t^{*\nu_w / (\nu_w - 1)} dj \right]^{\frac{(\nu_w - 1)}{\nu_w}} = W_t^*,$$

and

$$n_t = \left[\int_0^1 n_t^{*\nu_w} dj \right]^{\frac{1}{\nu_w}} = n_t^*.$$

Additionally, define $\text{mkp}_t^w = \frac{W_t}{W_t^h}$ to eliminate the household wage W_t^h , and define $b_t^n = \frac{B_t^n}{P_t}$.

An equilibrium consists of contingent sequences of $c_t, n_t, i_t, k_{t+1}, b_{t+1}^n, u_t, s_t, x_t, y_t, a_t, \text{mkp}_t^w, \tau_t, \lambda_t$ and prices w_t, R_t^n, π_t such that the household solves its problem, all firms solves their problems, and markets clear.

C.1.3 Relation to the results of Crouzet and Oh (2016)

In their appendix, Crouzet and Oh (2016) consider a stock-elastic demand model of inventory with flexible prices and wages. They define the expected discounted growth in marginal costs of production as $\Gamma_t = \frac{\beta E_t \frac{\lambda_{t+1}}{\lambda_t} \tau_{t+1} (1-\delta_n)}{\tau_t} = \frac{\mu_t^x}{\tau_t}$. Combining this with our equilibrium versions of (15) and (19) and then linearizing yields

$$\hat{x}_t = \hat{s}_t - \frac{\zeta \gamma}{\Phi} \hat{\mu}_t^x + \frac{\tau}{\Phi} \hat{\Gamma}_t. \quad (32)$$

This expression is similar to that of Crouzet and Oh (2016), with the addition of the second term on the right hand side, which arises in our model due to the impact of sticky prices. Under flexible prices, μ_t^x is constant, so substituting $\hat{\mu}_t^x = 0$ into (32) yields

$$\hat{x}_t = \hat{s}_t + \frac{\tau}{\Phi} \hat{\Gamma}_t, \quad (33)$$

as in Crouzet and Oh (2016). Note however that despite the presence of the expected discounted growth in marginal costs term $\hat{\Gamma}_t$ in (33), with $\mu_t^x = \beta E_t \frac{\lambda_{t+1}}{\lambda_t} mc_{t+1} (1 - \delta_n)$ constant, fluctuations in Γ_t are actually only due to movements in current marginal costs mc_t , since $\Gamma_t = \frac{\mu^x}{mc_t} = \frac{\theta/(\theta-1)}{mc_t}$. We can thus write (33) as $\hat{x}_t = \hat{s}_t - \frac{\tau}{\Phi} \hat{mc}_t$ as in equation (25) in the main text, such that only variation in sales and current marginal costs drive the response of inventory accumulation under flexible prices.