Is There News in Inventories?*

Christoph Görtz
University of Birmingham†

Christopher Gunn
Carleton University‡

Thomas A. Lubik
Federal Reserve Bank of Richmond§

June 2021

Abstract

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

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

JEL Classification: E2, E3.

---

*We are grateful to Paul Beaudry, Jean-Paul l'Hullier, Matteo Iacoviello, Alok Johri, Hashmat Khan, Andre Kurmann, Leonardo Melosi, Mathias Paustian, Franck Portier, Cedric Tille, and Mark Weder for useful comments and suggestions. We thank seminar and conference participants at the 2018 Canadian Economics Association Conference, the 2019 conference on Computing in Economics and Finance, the 7th Ghent University Workshop on Empirical Macroeconomics, the 2019 UVA-Richmond Fed Research Workshop, the 2019 Money, Macro and Finance Research Group Annual Conference, the 2019 AEA meeting, the 3rd University of Oxford NuCamp Conference, the College of William & Mary, the Deutsche Bundesbank, the Bank of Spain, the University of Sheffield, the University of Windsor, and Drexel University. The views expressed in this paper are those of the authors and not necessarily those of the Federal Reserve Bank of Richmond or the Federal Reserve System. Declarations of Interest: None.

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

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

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

1 Introduction

There is substantial evidence that expectations about future total factor productivity (TFP) are an important source of aggregate fluctuations (see Beaudry and Portier (2014), and references therein). Such TFP news shocks give rise to the observed comovement of aggregate quantities as identified in a large body of empirical work on the incidence and effects on news (e.g., Beaudry and Portier (2004)). Theoretical business cycle models can explain these findings under fairly general assumptions and modeling components (see Jaimovich and Rebelo (2009)) and imply substantial explanatory power of news shocks when taken to the data directly (e.g., Schmitt-Grohe and Uribe (2012); Görtz and Tsoukalas (2017)).

In this paper, we extend the news shock literature to account for inventories and show that they should take central stage in understanding the implications of news shocks. In the same vein, we argue that news shocks are an important component in understanding the behavior of inventory investment in addition to the standard mechanisms. Our paper uses inventories as a litmus test for the empirical relevance of TFP news shocks and we find these shocks are an important driver of aggregate fluctuations. In particular, we develop a new stylized fact and explain this fact in a general equilibrium model of inventory investment.

The news-shock literature has largely ignored inventory investment, which is a component of aggregate output and an adjustment margin to shocks that has long been recognized to play a large role in explaining aggregate fluctuations (see Ramey and West (1999); Wen (2005)). While inventory investment is only a small fraction of GDP, it plays an outsize role in contributing to the latter's volatility (see Blinder and Maccini (1991)). Aggregate inventories, in their dual role as input and output inventories, are also central to business cycle transmission via production networks (Iacoviello et al. (2011); Sarte et al. (2015)). Perhaps most importantly from our perspective is that inventories have a strategic role in buffering anticipated and unanticipated supply and demand disturbances. One might expect that news about such events would move inventories. Moreover, they are forward-looking in the sense that storage and acquisition requires planning. The forward-looking nature should make them responsive to news – which is precisely what we find.

Our paper makes two key contributions. First, we identify a new empirical fact in the inventory and news-shock literature. Using standard news-shock identification methodology for a structural vector autoregression (VAR) that includes inventories besides other quantity variables, we find that in response to anticipated news about higher future TFP, inventories rise on impact along with output, consumption, investment, and hours worked. This is a robust finding not only for the aggregate data, but also across the retail, wholesale and manufacturing sector as well as for finished goods, work-in-process, and input inventories. It is also robust across different approaches to identifying anticipated technology shocks. The consensus in the literature is that, unconditionally, inventory investment is procyclical (e.g., Ramey and West (1999)), whereby we identify a factor that induces conditional procyclicality.[1] Our findings therefore support the insight from the existing literature that news shocks are important drivers of business cycles.

Our second contribution is to identify the theoretical mechanism by which positive news about future TFP generates an expansion of all macroeconomic aggregates, including inventories, which is not a priori self-evident. In a conventional neoclassical framework with inventories, positive news about future TFP implies a wealth effect. The associated rise in sales of consumption and investment goods creates demand, which drives up inventories in order to avoid stockouts. However, the associated joint increase in sales and inventories can only be met through higher production. This implies rising marginal costs, which provides incentives for firms to partly satisfy higher demand by drawing down the inventory stock. This is reinforced by an intertemporal substitution effect, whereby positive news provides incentives to reduce current inventory stock, but build it up again in the future when high productivity is realized and marginal cost is lower.

We show that the standard news-shock model with inventories cannot explain our robust empirical finding that the news-driven demand effect dominates the substitution effect. By means of introducing general wedges into the standard model we isolate the components for labor supply and labor demand that are needed to replicate the empirical facts. We consider two potential mechanisms that operate on marginal costs, namely either sticky wages and prices, or knowledge capital. We find that the latter is qualitatively and quantitatively more successful, which we demonstrate by means of a calibration analysis and Bayesian estimation.

The core of our full model is the framework of Jaimovich and Rebelo (2009), which is closely related to Schmitt-Grohe and Uribe (2012). It includes the trio of particular specifications of pref-

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

erences, investment adjustment costs and variable capital utilization, which are features generally recognized in the news literature as needed for generating comovement of macroeconomic aggregates in response to a TFP news shock. We extend this model to include finished goods inventories based on the stock-elastic demand model of Bils and Kahn (2000). We then add knowledge capital, which can be interpreted as an intensive margin of hours worked, for instance, as the knowledge of how to best put to use an hour of work, based on earlier work by Chang et al. (2002a), Cooper and Johri (2002) and Gunn and Johri (2011).[2] We also impose a superstructure of nominal price and wage rigidities along the lines of Smets and Wouters (2007).

The accumulation of intangible knowledge through a learning-by-doing process involving labor addresses the shortcomings of the standard model in a straightforward manner. Firms acquire skill-enhancing knowledge through a learning-by-doing process from experience in production. The arrival of news about a future increase in TFP raises the value of knowledge in the present, inducing firms to increase their labor demand by varying markups in order to accumulate knowledge through experience. This has the effect of both contributing to the rise in hours worked, and thus production, and of suppressing the rise in the real wage during the initial boom. Consequently, the presence of knowledge capital limits the rise in marginal costs and increases the incentive to accumulate inventories. More succinctly, the accumulation of knowledge capital allows the news-shock-driven demand effect to dominate the substitution effect in production.

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

[2]This includes knowledge about operational processes, handling of machines and materials, and such. See Chang et al. (2002a) for an early application in a neoclassical business cycle model and d'Alessandro et al. (2019) for a recent application and further discussion.

sessment). In our theoretical modeling of inventories, we are guided by Bils and Kahn (2000), who highlight the unconditionally limited role of intertemporal substitution for variations in inventories that is also documented in our work in the context of expectations about productivity.

Our paper is most closely related to Crouzet and Oh (2016), who introduce inventories into a variant of the standard news-shock model of Jaimovich and Rebelo (2009), utilizing a reduced-form stockout-avoidance specification. They show that, while this setup can generate positive comovement of investment, consumption, and hours in response to stationary TFP news shocks, it fails to do so in the case of inventories. The countercyclical inventory movement is then used to inform sign restrictions in a structural VAR to identify TFP news shocks. Given the unconditional procyclicality of inventory investment and the imposed negative sign restriction on this variable, Crouzet and Oh (2016) come to the conclusion that such TFP news shocks are of limited importance for aggregate fluctuations. In contrast, we use a standard and widely used VAR methodology to identify first the response of inventory movements to news about the growth rate of TFP. The effects of these non-stationary shocks have been the focal point of the majority of the news literature, such as Barsky and Sims (2011) and Schmitt-Grohe and Uribe (2012). In response to these shocks, positive comovement of inventories emerges as a robust stylized fact that we then rationalize in an inventory model with a learning-by-doing propagation mechanism.

The remainder of the paper is structured as follows. Section 2 contains the main empirical results. Section 3 introduces the theoretical model used to rationalize the empirical findings. We trace out the required modeling elements and transmission mechanisms in general terms. We then identify potential specific candidates of which one is knowledge capital. In section 4, we estimate the theoretical model with knowledge capital using Bayesian methods. Section 5 concludes.

2 Inventories and news: Evidence from identified VARs

2.1 Data and estimation

We use quarterly U.S. data for the period 1983Q1-2018Q2.[3] Our main specification uses non-farm private inventories in the VAR. They are defined as the physical volume of inventories owned

[3]This choice is guided by the differences in cross-correlation patterns of several aggregate variables in samples before and after the mid-1980s (e.g., Galí and Gambetti (2009); Sarte et al. (2015)). In particular, McCarthy and Zakrajsek (2007) document that significant changes in inventory dynamics occur in the mid-1980s due to improvements in inventory management. In our robustness analysis, we document that our results generally hold for a longer sample.

by private non-farm businesses and are valued at average prices of the period, which captures the replacement costs of inventories.[4] Output is measured by GDP, and total hours as hours worked of all persons in the non-farm business sector. Investment is the sum of fixed investment and personal consumption expenditures for durable goods. Fixed investment is the component of gross private domestic investment that excludes changes in private inventories. Finally, consumption is defined as the sum of personal consumption expenditures for non-durable goods and services.

The time series are seasonally adjusted and expressed in real per-capita terms using total population, except for hours, which we do not deflate. In addition to the quantity aggregates, we also use a measure of inflation that we construct from the GDP deflator and a consumer confidence indicator that is based on the University of Michigan Consumer Sentiment Index.[5] This set of variables is standard in the literature, apart from inventories. The consumer confidence measure provides forward-looking information that potentially captures expectations or sentiment.[6]

Key to identifying the news shock in our baseline identification is a measure of observed technology. We follow the convention in the empirical literature and use the measure of utilization-adjusted TFP provided and regularly updated by Fernald (2014).[7] As a baseline, we identify TFP news shocks from the estimated VAR using the max-share method of Francis et al. (2014). This approach recovers the news shock by maximizing the variance of TFP at a specific long but finite horizon h, but does not move TFP on impact. The latter assumption implies that we impose a zero impact restriction on TFP conditional on the news shock. Following Francis et al. (2014) and the convention in the literature, we set the horizon h to 40 quarters. All variables enter in levels in line with the news shock VAR literature (e.g., Beaudry and Portier (2004); Barsky and Sims (2011)). We use Bayesian methods to estimate the VAR with three lags and a Minnesota prior. Confidence

[4]In a robustness exercise, we also consider business inventories as an alternative measure for stock holdings. This second measure differs in how the inventory stock is valued, namely by the cost at acquisition, which can be different from the replacement cost. In NIPA data, inventory profits and losses that derive from differences between acquisition and sales price are shown as adjustments to business income. Unfortunately, business inventories are available for only part of our sample (from 1992Q1). Apart from robustness considerations, the use of business inventories is appealing since this measure is available at a disaggregated level for different sectors and inventory types, which we subsequently use to evaluate robustness of our findings.

[5]This indicator, labeled E5Y, summarizes responses to the following question: "Turning to economic conditions in the country as a whole, do you expect that over the next five years we will have mostly good times, or periods of widespread unemployment and depression, or what?" The indicator is constructed as a diffusion index, namely as the percentage of respondents giving a favorable answer less the percentage giving an unfavorable answer plus 100.

[6]See, for instance, Barsky and Sims (2012). An alternative measure of forward-looking information is the S&P 500 stock price index. Our results are robust to including the S&P 500 instead of the Michigan consumer confidence index which we document in the online appendix B.2.

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

bands are computed by drawing from the posterior. Since the VAR setup and our baseline news shock identification is standard in the literature, we refer the reader to appendix A for further details. We first report on the results from the baseline identification and then scrutinize our results against using alternative identification schemes proposed in the literature.

2.2 The empirical response of inventories to a TFP news shock

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

In contrast, all quantity variables significantly rise on impact and follow a hump-shaped pattern. Moreover, the peak response occurs before TFP hits its highest point. Positive comovement between output, consumption, investment, and hours over this post-Great Moderation sample in response to news has been documented before, for instance by Görtz et al. (2021). We add to these previously established stylized facts the behavior of private non-farm inventories. In response to a news shock, they rise somewhat on impact and continue to do so in a hump-shaped pattern until reaching a peak at about 10 quarters. The change in the stock of inventories, inventory investment, is negative afterwards, while its level never falls below the zero line, its starting point.[8] Importantly, the VAR results also reveal that the TFP news shock is a key driver for fluctuations in inventories and GDP as it explains between 47-65% and 47-71% of the respective forecast error variances over a horizon between 6-32 quarters.[9]

We consider a variety of additional specifications to assess the robustness of our findings. First, we show in appendix B.5 that the results are robust to alternative specifications for the news identification horizon h and also hold in a very small-scale VAR or if other variables are included in the VAR system. We also consider longer sample periods for the specification with non-farm private inventories, that is, samples starting in 1948Q1 and 1960Q1. These results are reported in

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

[9] The full set of results from the variance decomposition is reported in the online appendix B.1.

appendix B.2. We find that the impulse response patterns identified in our baseline specification carry over to the two longer samples qualitatively and to a large extent also quantitatively.[10]

2.3 Robustness: alternative news shock identification

While our baseline max-share identification is widely used in the literature, it crucially relies on the observed TFP series. The series we employ is arguably the best measure for TFP available, yet it is likely to suffer from a certain degree of measurement error. For this reason, we subject our empirical findings above to alternative identifications for news shocks recently suggested in the literature. The alternative identification approaches fall broadly into two categories. The first relies on Fernald's TFP series as an observable, but attempts to mitigate any effects of potential mis-measurement. The second does not rely on TFP, but uses patents to broadly capture news about future technology.

Kurmann and Sims (2019) argue that the TFP measure is likely to be confounded by business cycle fluctuations due to imperfect measurement of factor utilization. This is particularly problematic in light of the zero-impact restriction imposed in the baseline identification scheme. For this reason, Kurmann and Sims (2019) suggest to recover news shocks by maximising the forecast error variance of TFP at a long finite horizon, as in our baseline identification, but without imposing a zero-impact restriction on TFP. They argue that allowing TFP to jump freely on impact in response to the news shock, produces robust inference to cyclical measurement error in the construction of TFP. Figure 2 shows the impulse responses under the Kurmann-Sims identification. Over our considered time horizon, these responses are qualitatively and quantitatively very similar to the ones reported from our baseline. Importantly, both identification schemes suggest that inventories increase in anticipation of higher future TFP. Even without the impact restriction, TFP rises significantly only with a substantial delay.[11]

The second type of alternative identification schemes relies on patents and is independent of Fernald's productivity measure. We follow Cascaldi-Garcia and Vukotic (2020), who argue that

[10] A priori it is not obvious at which prices inventories should be measured. Appendix B.3 shows that our finding of a procyclical inventory response to TFP news shocks is robust to a specification with business inventories. Business inventories are measured at the cost at acquisition, which can be different from the replacement cost considered as a measure for private non-farm inventories. The availability of disaggregated data for business inventories allows us to verify the robustness of our results to inventories in different sectors (manufacturing, wholesale, retail) and of different types (input, work in process, and final goods inventories).

[11] Appendix B.4 shows that our baseline results are robust also to other, closely related, identification schemes proposed by Barsky and Sims (2011) and Forni et al. (2014).

patents include information about future TFP movements since firms engage in activities to take advantage of expected technological improvements or are the originators of such productivity advancements. The patent system is designed to reveal such news without the full set of improvements necessarily being in place. Following the methodology in Cascaldi-Garcia and Vukotic (2020) and Kogan et al. (2017) we construct a quarterly aggregate patent series from panel observations on patents associated with stock market listed firms in the CRSP database.[12]

We then follow Cascaldi-Garcia and Vukotic (2020) in using this series to identify responses to patent-based news shocks in a Bayesian VAR based on a simple Cholesky identification with the patent series ordered first. Figure 3 shows impulse responses to this patent-based news shock. They are qualitatively consistent with the responses in the baseline specification. TFP rises significantly only with a delay, even though there is no zero-impact restriction applied. Consistent with the findings in Cascaldi-Garcia and Vukotic (2020), activity variables as well as consumer confidence rise. We add to their findings by documenting a rise in inventories, which is consistent with the evidence based on the other news shock identification schemes considered above. These results are interesting on their own as our constructed time series include the most recent data. Due to data limitations, Cascaldi-Garcia and Vukotic (2020) only show responses for a time horizon up to 2010. We conclude that the consistency of all results in this section provides robust evidence for the rise in inventories in light of positive news about future technology.

2.4 The empirical evidence and structural models

We can summarize our findings at this point as follows. Evidence from an identified VAR shows that a news shock signalling higher future productivity leads to an increase and subsequent positive comovement of all aggregate variables we considered. The new fact that we document in our paper is that this pattern extends to the response of inventories and is broad-based across different news shock identification schemes. Why the behavior of inventories follows this pattern need not be obvious a priori. Conceivably, they could decline initially to satisfy higher demand instead of higher production. Moreover, higher TFP in the future reduces the cost of replenishing a drawn-down inventory stock. At the same time, firms may increase inventories to maintain a

[12]Kogan et al. (2017) compute the economic value of a patent based on a firm's stock-price reaction to observed news about a patent grant, controlling for factors that could move stock prices but are unrelated to the economic value of the patent. In particular, they aggregate value weighted patents by taking the sum of all patents issued in a particular quarter, scaled by aggregate output.

desired inventory-sales ratio, which counters this effect. It is along these margins that the success of a theoretical model to replicate the empirical findings rests.[13]

Jaimovich and Rebelo (2009) document the elements necessary in a theoretical model to facilitate comovement of consumption and investment in response to news about future higher TFP. Specifically, they show that a strong increase in utilization and hours worked are key components. Positive news stimulates consumption through a wealth and income effect. The latter is driven by increased hours worked to raise production in order to satisfy that demand. Similarly, investment increases to support the higher capital stock to take advantage of higher future TFP. This reasoning is corroborated in our baseline VAR corresponding to Figure 1, where we add additional variables one at a time. Selective impulse responses to a TFP news shock are reported in Figure 4.[14]

Figure 4 shows that the inventory-to-sales ratio moves countercyclically in response to a news shock. This is a key observation that informs our thinking about a theoretical model. Countercyclicality of the inventory-to-sales ratio is a necessary condition for comovement of inventories with the other macroeconomic aggregates. The literature on inventories often does not only consider their level but also their change, which provides an indication about inventory investment. The figure shows a positive response of inventory investment which is broadly consistent with the response of the level of inventories documented in Figure 1. Figure 4 also documents a strong increase in capital utilization. The positive hump-shaped response of the real wage is consistent with the increase in hours documented in Figure 1. It is also indicative of a hump-shaped increase in knowledge capital. In addition to the real wage, we consider two more variables that have been considered to understand the response of knowledge capital. Intellectual property products provide suggestive evidence for a possible channel of how news propagates and affects the production process. Figure 4 shows that intellectual property products rise in response to a news shock, com-

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

[14]The inventory-to-sales ratio is the ratio of private non-farm inventories and final sales of domestic business as in Lubik and Teo (2012). Utilization is provided by Fernald (2014) and consistent with our utilization-adjusted measure for TFP. The real wage is compensation of employees, non-financial corporate business, in real per-capita terms. The change in inventories is the change in private non-farm inventories. Issued patents are obtained from the US Patent and Trademark Office. The series for intellectual property products is real per-capita nonresidential intellectual property products available from the Bureau of Economic Analysis.

mensurate with the behavior of other variables considered so far. The same holds for the number of issued patents. This suggests that a central component of a news-driven business cycle model that is consistent with the empirical evidence could be the accumulation of knowledge, residing with households as human capital or embodied in physical capital. In the next section we build a theoretical model along the lines suggested by these findings.

3 Theoretical model

We now develop a business cycle model to rationalize the findings of the empirical analysis. Our baseline framework is the flexible wage and price model of Schmitt-Grohe and Uribe (2012) augmented by inventories. Their model uses the particular specification of preferences, investment adjustment costs and costly capacity utilization of Jaimovich and Rebelo (2009), which has become the workhorse framework in the news shock literature. We model inventories as in Lubik and Teo (2012), based on the stock-elastic demand model of Bils and Kahn (2000), where finished goods inventories are sales-enhancing.

3.1 Model description

The model economy consists of a large number of identical infinitely-lived households, a competitive intermediate goods-producing firm, a continuum of monopolistically competitive distributors, and a competitive final goods producer. The intermediate goods firm owns its capital stock and produces a homogeneous good that it sells to distributors. This good is then differentiated by the distributors into distributor-specific varieties that are sold to the final-goods firm. The varieties are aggregated into final output, which then becomes available for consumption or investment. We adopt this particular decentralization since it is convenient for modeling finished goods inventories by separating the production side of the economy into distinct production, distribution, and final goods aggregation phases. The model economy contains several stationary stochastic shock processes as well as non-stationary TFP and IST shocks. In addition to the TFP shocks, we include a suite of shocks that are standard in the literature to facilitate estimation later in the paper.

3.1.1 Intermediate goods firm

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

$$Y_t = F(N_t, \widetilde{K}_t; H, z_t, \Omega_t) = z_t \left(\Omega_t N_t\right)^{\alpha_n} \widetilde{K}_t^{\alpha_k} \left(\Omega_t H\right)^{1-\alpha_n-\alpha_k}, \tag{1}$$

where z_t is a stationary exogenous stochastic productivity process, Ω_t is a non-stationary exogenous stochastic productivity process, and H is a fixed factor that allows for decreasing-returns-to-scale to N_t and $\widetilde{K}_t$ as in Jaimovich and Rebelo (2009) and Schmitt-Grohe and Uribe (2012).[15] We assume that the growth rate of Ω_t, $g_t^{\Omega} = \Omega_t/\Omega_{t-1}$, is stationary.

In each period, the firm acquires labor N_t at wage w_t from the labor market, and capital services $\widetilde{K}_t$ at rental rate r_t from the capital services market. It then sells its output Y_t at real price τ_t to the distributors. The firm's profit maximization problem results in standard demand functions for labor and capital services, respectively: $w_t = \alpha_n \tau_t \frac{Y_t}{N_t}$ and $r_t = \alpha_k \tau_t \frac{Y_t}{\widetilde{K}_t}$. Additionally, we find it convenient to define the marginal cost of production for intermediate goods, $mc_t = \frac{w_t}{MPN_t} = \frac{w_t}{\alpha_n Y_t/N_t}$, where $MPN_t = F_{N_t}$ is the marginal product of labor. It then follows that the output price τ_t is equal to the marginal cost of production mc_t.

3.1.2 Final goods firm

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

$$S_t = \left[\int_0^1 v_{it}^{\frac{1}{\theta}} S_{it}^{\frac{\theta-1}{\theta}} di\right]^{\frac{\theta}{\theta-1}}, \qquad \text{with} \qquad v_{it} = \left(\frac{A_{it}}{A_t}\right)^{\zeta}, \qquad \text{and} \qquad \theta > 1,\ \zeta > 0.$$

where v_{it} is a taste shifter that depends on the stock of goods available for sale A_{it}. The latter is composed of current production and the stock of goods held in inventory.[16] We assume that v_{it} is taken as given by the final goods producer and A_t is the economy-wide average stock of goods for sale, given by $A_t = \int_0^1 A_{it} di$. The parameters θ and ζ capture, respectively, the elasticity of substitution between differentiated goods and the elasticity of demand with respect to the relative stock of goods.

The firm acquires each variety i from the distributors at relative price $p_{it} = P_{it}/P_t$, where $P_t =$

[15] These authors interpret the fixed factor H as land or organizational capital. A production function that is homogeneous-of-degree-1 in its inputs of labor, capital services and the fixed factor H introduces decreasing returns to scale to labor and capital services, thereby allowing for the possibility of a positive increase in the stock value of the firm in response to TFP news.

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

$\left[\int_0^1 v_{it} P_{it}^{1-\theta} di\right]^{\frac{1}{1-\theta}}$ is the aggregate price index. It sells the final good for use in consumption or as an input into the production of investment goods. The firm maximizes the profit function $\Pi_t^s = S_t - \int_0^1 \frac{P_{it}}{P_t} S_{it} di$ by choosing S_{it}, $\forall i$. This results in demand for S_{it} for the ith variety:

$$S_{it} = v_{it} p_{it}^{-\theta} S_t. \tag{2}$$

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

3.1.3 Distributors

We close the production side of the model by introducing inventories at the level of the distributors. We follow Bils and Kahn (2000) in modeling inventories as a mechanism that helps generate sales, while at the same time implying a target inventory-sales ratio that captures the idea of stock-out avoidance. Distributors acquire the homogeneous good Y_t from the intermediate goods firms at real price τ_t. They differentiate Y_t into goods variety Y_{it} at zero cost, with a transformation rate of one-to-one. Goods available for sale are the sum of the differentiated output and the previous period's inventories subject to depreciation:

$$A_{it} = (1 - \delta_x) X_{it-1} + Y_{it}, \tag{3}$$

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

$$X_{it} = A_{it} - S_{it}, \tag{4}$$

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

The distributors have market power over the sales of their differentiated varieties. The ith distributor sets price p_{it} for sales S_{it} of its variety subject to its demand curve (2). Each period, a distributor faces the problem of choosing p_{it}, S_{it}, Y_{it}, and A_{it} to maximize profits:

$$E_t \sum_{t=0}^{\infty} \beta^k \frac{\lambda_{t+k}}{\lambda_t} \left[\frac{P_{it+k}}{P_{t+k}} S_{it+k} - \tau_t Y_{it+k}\right],$$

subject to the demand curve (2), the law of motion for goods available for sale (3), and the definition of the inventory stock (4). Profit streams are evaluated at the household's marginal utility of wealth λ_t. Substituting the demand curve for S_{it}, and letting μ_t^a and μ_t^x be the multipliers on the two other constraints, we can then find a representative distributor's first-order conditions:

$$\tau_t = \mu_t^a, \quad (5)$$

$$\mu_t^x = (1-\delta_x)\beta E_t \frac{\lambda_{t+1}}{\lambda_t}\mu_{t+1}^a, \quad (6)$$

$$\mu_t^a = \zeta p_{it}\frac{S_{it}}{A_{it}} + \mu_t^x\left(1-\zeta\frac{S_{it}}{A_{it}}\right), \quad (7)$$

$$\frac{P_{it}}{P_t} = \frac{\theta}{\theta-1}\mu_t^x, \quad (8)$$

which are, respectively, the optimal choices of Y_{it}, X_{it}, A_{it}, and P_{it}. The optimality condition (5) implies that the cost of an additional unit of goods for sale, τ_t, is equal to the value of those goods for sale, namely μ_t^a. Since inventories at the beginning of a period are predetermined by the law of motion for A_{it}, a distributor can only further increase its stock of available goods for sale by acquiring additional output Y_{it}.

The optimality condition (6) relates the current value of an additional unit of inventory to the expected discounted value of the extra level of goods available for sale next period generated by holding inventory. Since any increase in sales results in a reduction in stock holdings, the opportunity cost of sales for the distributor is equal to the value of foregone inventory μ_t^x, which can be thought of as the marginal cost of a sale. The marginal cost of sales is thus equal to the expected discounted value of next period's marginal cost of output, since increasing sales by drawing down stock in order to forgo production today means that the distributor will need to increase production eventually in the future.

The optimality condition (7) connects the marginal value μ_t^a of a unit of goods available for sale to the value of the extra sales generated by the additional goods available plus the value of the additional inventory yield from the unsold portion of the additional goods. We can combine the marginal cost expressions (5)-(7) to derive:

$$\tau_t = \zeta\frac{P_{it}}{P_t}\frac{S_{it}}{A_{it}} + (1-\delta_x)\beta E_t\frac{\lambda_{t+1}}{\lambda_t}\tau_{t+1}\left(1-\zeta\frac{S_{it}}{A_{it}}\right). \quad (9)$$

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

Finally, the optimal pricing choice (8) sets the distributor's relative price as a constant markup over the marginal cost of sales as in a standard flexible price model with imperfect competition, but without inventories. The presence of inventories however drives a wedge between the marginal costs of output and of sales to the effect that there is no longer a constant markup between price and marginal costs of output, but one that varies with the value of foregone inventory μ_t^x.

3.1.4 Further model elements and model solution

The household and government side of the model economy are standard and follow Schmitt-Grohe and Uribe (2012). Further details and derivations are in appendix C.1.1. The non-stationary exogenous stochastic TFP process Ω_t, with growth rate g_t^{Ω} is given by:[17]

$$\ln\left(\frac{g_t^{\Omega}}{g^{\Omega}}\right) = \rho_{g^{\Omega}} \ln\left(\frac{g_{t-1}^{\Omega}}{g^{\Omega}}\right) + u_t^{g^{\Omega}}, \qquad \text{with} \qquad u_t^{g^{\Omega}} = \varepsilon_{g^{\Omega}t}^{0} + \varepsilon_{g^{\Omega}t-4}^{4} + \varepsilon_{g^{\Omega}t-8}^{8} + \varepsilon_{g^{\Omega}t-12}^{12},$$

where $\varepsilon_{g^{\Omega}t}^{0}$ is an unanticipated shock and $\varepsilon_{g^{\Omega}t-p}^{p}$ is a news shock that agents receive in period t about the innovation in time $t+p$. Model equilibrium, stationarization and solution method are standard and we discuss these in detail in appendix C.2.

3.2 Understanding inventory dynamics

We begin our model analysis by examining the response of inventories to TFP news in a calibrated version of the model introduced above. Our choice of parameter values is guided by the existing literature, where we maintain comparability with Jaimovich and Rebelo (2009) and Schmitt-Grohe and Uribe (2012) for the aspects of the news shock mechanism and Lubik and Teo (2012) for the inventory component. This calibration is detailed in Appendix C.3 as it is purely for illustrative purposes. Our main empirical analysis is in section 4, where we estimate a full New Keynesian version of the model.

Figure 5 reports the impulse responses of key model variables to news about a future permanent increase in TFP that will be realized in 8 quarters as anticipated. With the exception of consumption, all macroeconomic variables decline in response to the news. Moreover, after the initial drop, inventory declines rapidly over time until the actual realization of the TFP shock. Consequently, the response of the major variables in the model is at odds with our VAR-based empirical evidence. This finding is corroborated analytically in the following subsections. In addition, the figure also illustrates how incorporating inventories in an otherwise standard model can alter the dynamics of other model variables, despite a calibration close to that of Jaimovich and Rebelo (2009) designed to generate co-movement in consumption, investment and hours-worked

[17] We discuss details of the other shock process in section 4 where we estimate the model.

in response to news. Therefore, we now examine the key mechanisms of the model to understand the behavior and role of inventory holdings. We frame our discussion in terms of demand and supply schedules in the model economy's market for produced output Y_t with market-clearing price τ_t, which in the baseline model, is also the marginal cost of production.[18]

Output Demand. We derive the demand schedule from the optimal stocking condition for the distributors:

$$\tau_t = \frac{\zeta}{\theta}\frac{S_t}{A_t} + \frac{\theta - 1}{\theta} = \frac{\zeta/\theta}{1 + X_t/S_t} + \frac{\theta - 1}{\theta} = \tau(\chi_t), \tag{10}$$

where $\chi_t = \frac{X_t}{S_t}$, and $\tau'(\cdot) < 0$, and the inventory accumulation equation, formed by combining (3) and (4):

$$X_t = (1 - \delta_x) X_{t-1} + Y_t - S_t. \tag{11}$$

Equation (10) is the key equation governing inventory dynamics in the model. It implies that the distributor targets a sales-to-stock ratio $\frac{S_t}{A_t}$, or equivalently, an inventory-sales ratio, $\chi_t = \frac{X_t}{S_t}$, for a given level of marginal cost of output τ_t. All else equal, the distributor increases inventory holdings with a rise in sales, what may be labelled the demand channel. Similarly, inventory holdings are reduced with a rise in current marginal costs, what may be labelled the cost channel.[19] Equation (11) describes the law of motion of inventory accumulation and shows the two margins of adjustment: a given increase in sales S_t can be satisfied with either a decrease in inventory X_t, an increase in output Y_t, or some combination (which may involve both an increase in X_t along with Y_t). The optimality condition embedded $\tau(\chi_t)$ governs the trade-off between these two margins.

We now define $\chi(\tau_t) = \tau^{-1}(\chi_t)$, so that $\frac{X_t}{S_t} = \chi(\tau_t)$ expresses the optimal stocking condition that relates the inventory-sales ratio to a given level of marginal costs τ_t. Using this in the inventory accumulation equation (11) gives:

$$Y_t = (1 + \chi(\tau_t)) S_t - (1 - \delta_s) X_{t-1}, \tag{12}$$

which is downward-sloping in (Y_t, τ_t)-space. The optimal stocking condition combined with the

[18]Our analysis is focused on the news phase, which is the range of time defined from $t = 1$ when the news shock arrives, to the period $t + p - 1$, namely one period before TFP actually changes in period $t + p$. During the news phase, there are no changes in non-stationary TFP (and of course, no changes in any shock other than the considered TFP news shock). Appendix C.4 includes a detailed analytical and descriptive exposition.

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

inventory accumulation equation can thus be thought of as a demand curve for Y_t. All else equal, higher marginal cost implies a lower inventory-sales ratio, and thus lower demand for Y_t, as distributors seek to run down inventory stock. Similarly, an increase in sales shifts the curve outward and raises the demand for Y_t as the distributors seek to maintain their sales-inventory ratio by increasing their holdings.

Output Supply. The supply schedule in the market for output is derived from the labor market equilibrium condition and the production technology. For ease of exposition, we abstract from the income effect in the utility function ($\gamma_j \approx 0$) and assume no habits in consumption ($b = 0$). This results in:

$$\tau_t = \psi \frac{\xi}{\alpha_n} Q_t^{-\frac{\xi}{\alpha_n}} Y_t^{\frac{\xi}{\alpha_n}-1}, \tag{13}$$

where $Q_t = z_t \Omega_t^{1-\alpha_k} \left(\tilde{K}_t\right)^{\alpha_k}$, and $\frac{\partial \tau_t}{\partial Y_t} > 0$ for $\xi > \alpha_n$, so that the curve is upward-sloping for reasonably elastic labor supply.

Response to TFP News. The supply and demand schedules for output Y_t at marginal cost τ_t are depicted in Figure 6. Arrival of positive news about future TFP implies a wealth effect that drives up current demand for consumption. In our inventory framework, this also raises the demand for sales of distributors, which shifts their output demand curve (equation (12)) outward from D to D' in Figure 6 as agents increase their demand for newly produced goods. The shift in demand puts upward pressure on τ_t, which would imply a lower inventory-sales ratio via the optimal stocking condition. We can see from equation (12) that for a given rise in sales the extent of the rise in marginal cost determines whether inventories rise or fall. If the rise in marginal costs is large, inventories must fall in order to reduce the inventory-to-sales ratio enough for equation (12) to still hold as it becomes more attractive for distributors to draw down stock in the present in order to avoid the high current production costs. On the other hand, if the rise in marginal costs is small, inventories can still rise along with increasing sales as long as the rise is proportionally less than sales such that the inventory-to-sales ratio still falls and (12) holds. In fact, as long as marginal costs increase, a countercyclical inventory-sales ratio, which is consistent with our empirical evidence in Section 2.4, is a necessary condition for positive comovement of inventories with other aggregate quantities.

Inventory Comovement. We now build on the previous discussion to characterize conditions under which inventory responds procyclically. We combine (10) and (11) to eliminate sales S_t:

$$\left(1+\frac{1}{\chi(\tau_t)}\right)X_t = (1-\delta_x)X_{t-1} + Y_t, \tag{14}$$

such that the output demand equation reads:

$$\tau_t = Q^d(Y_t; X_t, X_{t-1}). \tag{15}$$

Similarly, we use the capital market equilibrium conditions to eliminate capacity utilization from the supply schedule (where q_t^k is the price of capital):

$$\tau_t = Q^s(Y_t; q_t^k, K_t). \tag{16}$$

We can then use equations (15) and (16) to characterize the dynamics of X_t relative to Y_t for given values of q_t^k and K_t. To gain additional insight, we focus on the linear approximation of the detrended equivalents of these equations around the steady state. We are interested in the conditions under which inventory co-moves with output. As such, we wish to isolate the conditions under which $\hat{x}_t > 0$ for $\hat{y}_t > 0$, where "hats" denote percent deviations from the detrended stationary steady state. Linearizing (15) and (16) and imposing $\hat{x}_t > 0$ for $\hat{y}_t > 0$ yields the inventory comvement condition (see appendix C.4 for the detailed derivations):

$$\left(\frac{\left(\frac{\xi}{\alpha_n}-1\right)-\theta_u}{1+\theta_u} - \frac{y}{s}\frac{1}{\varepsilon_x}\right)\hat{y}_t - \frac{\theta_u}{1+\theta_u}\varepsilon_u\hat{k}_t + \theta_u\hat{q}_t^k - \frac{x}{s}\frac{1}{\varepsilon_x}\frac{(1-\delta_x)}{g^y}\hat{x}_{t-1} < 0, \tag{17}$$

where $\hat{y}_t > 0$, $\varepsilon_x = |\frac{\chi'(\tau)}{\chi(\tau)}\tau|$ and $\theta_u = \frac{\xi}{\alpha_n}\frac{\alpha_k}{1+\varepsilon_u}$. This inequality describes the equilibrium response consistent with $\hat{x}_t > 0$ for $\hat{y}_t > 0$ in the market for output, conditional on the general equilibrium response of $\hat{q}_k^k$, $\hat{K}_t$ and $\hat{x}_{t-1}$. As such, the sign of the expression on the left-hand is a function of both the sign of the coefficients, as well as the sign and magnitude of the particular general equilibrium response of $\hat{y}_t$, $\hat{k}_t$, $\hat{q}_k^k$, and $\hat{x}_{t-1}$.

We provide a detailed discussion of the co-movement condition (17) in appendix C.4, where we derive analytic conditions for inventory co-movement to hold. We summarize these results as follows. In the initial period $t = 1$ when news arrives, $\hat{k}_t = 0$ and $\hat{x}_{t-1} = 0$. Satisfying the equation (17) for $\hat{y}_t > 0$ thus depends only on the sign of the coefficient on $\hat{y}_t$ and the sign and magnitude of $\hat{q}_t^k$. The coefficient on $\hat{y}_t$ measures the relative slope of the output demand and supply schedules and is positive for all realistic values of the pertinent parameters. Initial inventory comovement then rests on the response of $\hat{q}_t^k$. As is well known in the literature, with the flow-form of investment adjustment costs used in the model, $\hat{q}_t^k$ does respond negatively to news of a future rise in TFP. However, it is not enough to satisfy condition (17) on its own on impact. Consequently, inventories fall for all relevant parameter values.

During the transition period $t = 2$ to $t + p - 1$, a rise in $\hat{k}_t$ and $\hat{x}_{t-1}$ or a fall in q_t^k can potentially shift the output supply curve enough to relax condition (17). Yet if $\hat{x}_{t-1} < 0$ as it is here on impact, the $\hat{x}_{t-1}$ terms actually works in the wrong direction, making the condition more difficult to satisfy. Additionally, assuming an expansion where output growth is positive for several periods such that $\hat{y}_{t+1} > \hat{y}_t$, the positive coefficient on $\hat{y}_t$ in (17) means that any factors that shift the output supply curve have to shift it to overcome the increase in $\hat{y}_t$ over time. While movements in $\hat{k}_t$ and q_t^k offer the potential to shift the output supply curve over time, our simulations suggest that these factors are not enough, and that their combined effect is overwhelmed by the rise in $\hat{y}_t$.

We conclude that the baseline model is likely not consistent with inventory comovement. Specifically, the respective slopes of the output supply and demand curves do not on their own satisfy the inventory comovement condition during the news-period. However, our analysis points to the endogenous response of factors that shift either of these curves on impact and in subsequent periods. Investment adjustment costs is a possibility, yet our simulations suggest that variation in q_t^k on its own is unable to satisfy the comovement condition.

3.3 Uncovering the missing elements: a wedges approach

We now re-examine the inventory dynamics of the baseline model to understand the potential missing elements that would otherwise allow inventory to respond procyclically. The analysis in the previous section points towards missing endogenous shifters in the output supply curve. We study this aspect by introducing wedges into the model in the spirit of Chari et al. (2007). Such wedges can be interpreted as endogenous equilibrium objects that represent deviations of some other candidate model in equilibrium from the baseline model.

The intermediate goods firm produces output according to the production technology (1). Consider an alternative model, where the production technology is now given by

$$Y_t = \phi_t^e F(N_t, \widetilde{K}_t; H, z_t, \Omega_t) = \phi_t^e z_t (\Omega_t N_t)^{\alpha_n} \widetilde{K}_t^{\alpha_k} (\Omega_t H)^{1-\alpha_n-\alpha_k},$$

where ϕ_t^e is an *efficiency wedge*. The firm's optimal labor demand in the baseline model is given by $\frac{w_t}{F_{nt}} = \tau_t$, where $F_{nt} = MPN_t$, while in the alternative model this same condition is:

$$\frac{w_t}{\phi_t^e F_{Nt}} = \frac{\tau_t}{\phi_t^{ld}}, \tag{18}$$

where $\phi_t^e F_{Nt} = MPN_t$, and where ϕ_t^{ld} is a labor demand wedge. Consequently, time variation in ϕ_t^{ld} serves as an additional source of shifts in labor demand relative to the baseline model.

We note that the labor demand wedge ϕ_t^{ld} affects the optimality condition but not the production technology directly, whereas the efficiency wedge ϕ_t^e enters into both. ϕ_t^{ld} can thus be interpreted as a type of markup, such that a decrease is associated with an increase in labor demand. On the other hand, an increase in the efficiency wedge ϕ_t^e raises both labor demand and goods production. Given our earlier definition of marginal cost of production as $mc_t = w_t/MPN_t$, we can alternatively write equation (18) as $\phi_t^{ld} = \frac{\tau_t}{mc_t}$, which highlights the interpretation of the labor demand wedge as a markup of the price of output over marginal cost of production.

Turning to the households, the labor first-order condition in the baseline model is $MRS_t = w_t$. We introduce a *labor supply wedge* ϕ_t^{ls} operating in an alternative model, which implies the labor supply condition:

$$MRS_t = \frac{w_t}{\phi_t^{ls}},$$

All else equal, time-variation in ϕ_t^{ls} serves as an additional source of shifts in labor supply relative to the baseline model. As with the labor demand wedge, ϕ_t^{ld} can be interpreted as a markup, such that a reduction in ϕ_t^{ld} is associated with an increase in labor supply. Labor market equilibrium then results in the expression

$$MRS_t = \Phi_t \tau_t F_{Nt}, \tag{19}$$

where $\Phi_t = \frac{\phi_t^e}{\phi_t^l}$ is the overall labor wedge, and $\phi_t^l = \phi_t^{ls}\phi_t^{ld}$ is the (combined) labor markup wedge.

We can now incorporate the wedges into the demand and supply schedules for output. This implies the following modified output supply curve:

$$\tau_t = \psi \frac{\xi}{\alpha_n} \Phi_t^{-1} Q_t^{-\frac{\xi}{\alpha_n}} Y_t^{\frac{\xi}{\alpha_n}-1}.$$

Since $\frac{\partial \tau_t}{\partial \Phi_t} < 0$, the output supply curve is shifted outwards by a reduction in the labor supply wedge ϕ_t^{ls}, a reduction in the labor demand wedge ϕ_t^{ld}, or an increase in the efficiency wedge ϕ_t^e. This limits the rise in τ_t for any given increase in sales associated with news and thereby reduces the required decline in the inventory-sales ratio from the distributor's optimal stocking equation (10). Consequently, such changes in the respective wedges increase the possibility that inventories rise along with sales.

Similarly, we can extend the linearized co-movement conditions $\hat{x}_t > 0$ for $\hat{y}_t > 0$ to incorporate the wedges. This yields:

$$\left(\frac{\left(\frac{\xi}{\alpha_n} - 1\right) - \theta_u}{1+\theta_u} - \frac{y}{s}\frac{1}{\varepsilon_x} \right) \hat{y}_t - \frac{\theta_u}{1+\theta_u}\varepsilon_u \hat{k}_t + \theta_u \hat{q}_t^k - \frac{x}{s}\frac{1}{\varepsilon_x}\frac{(1-\delta_x)}{g^y} x_{t-1} - \frac{1+\frac{\xi}{\alpha_n}}{1+\theta_u}\hat{\phi}_t^e + \frac{\theta_u}{1+\theta_u}\hat{\phi}_t^l < 0. \tag{20}$$

where $\hat{y}_t > 0$.

The wedges framework highlights the margins required to satisfy the comovement condition through either increases in the efficiency wedge $\hat{\phi}_t^e$ or decreases in the labour supply and demand markup wedges through $\hat{\phi}_t^l$. While there are potentially many different models that could yield movement in these wedges, we can isolate two general characterizations of the required movement in the wedges relative to the baseline model. First, a wedge should respond on impact in order to prevent an initial drop in inventory. Second, the combined effect of the wedges should grow over time in order to match the positive growth in $\hat{y}_t$ through the expansion and allow inventory to rise along with $\hat{y}_t$.

3.4 Two potential candidates

We consider two candidate models for generating movement in the labor wedges discussed above. The first model uses nominal rigidities; while the second model is based on a specific type of a real rigidity. We discuss each in turn, analyzing their impact on inventory dynamics relative to the baseline model.

3.4.1 Nominal rigidities: sticky Wages and prices

Our first candidate model uses sticky wages and prices to generate endogenous movement in the labor wedges. These are natural candidates to examine in our context since they operate by ultimately altering markups in the labor market. We introduce sticky prices as in Lubik and Teo (2012), whereby we assume that distributors face convex adjustments costs in setting prices. The sticky-wage component follows the decentralization of Schmitt-Grohe and Uribe (2012) and Smets and Wouters (2007). Finally, we close the model with a standard monetary policy nominal interest rate rule. Since these extensions to the baseline model are relatively standard, we discuss them only briefly, leaving the details to appendix D.

Labor Supply and Output Demand Wedges. The sticky-wage framework results in a time-varying markup μ_t^w between the wage w_t paid by the intermediate goods firm and the wage w_t^h paid to the household, such that:

$$\mu_t^w = \frac{w_t}{w_t^h}.$$

The dynamics of μ_t^w is captured by a wage Phillips curve. In the context of our wedges framework in the labor market, the presence of sticky wages corresponds to $\phi_t^{ls} = \mu_t^w$, $\phi_t^{ld} = 1$ and $\phi_t^e = 1$.

The sticky-price framework results in an additional wedge in the output demand side of the model. Unlike in the flexible price version, where the markup between the marginal cost of sales and price is constant, the distributor's pricing condition under sticky prices implies that this markup is time-varying. This means that the value of forgone inventory, μ_t^x, which we previously interpreted as the marginal cost of sales, is no longer constant. As such, this introduces μ_t^x as a time-varying wedge into the firm's optimal stocking equation:

$$\tau_t = \zeta p_{it} \frac{S_{it}}{A_{it}} + \mu_t^x \left(1 - \zeta \frac{S_{it}}{A_{it}}\right). \tag{21}$$

Solving for $\chi_t = \frac{X_t}{S_t}$ yields:

$$\chi_t = \zeta \frac{1-\mu_t^x}{\tau_t - \mu_t^x} - 1 = \chi(\tau_t, \mu_t^x),$$

where $\chi_\tau(t) = \frac{\partial \chi(\tau_t, \mu_t^x)}{\partial \tau_t} < 0$ and $\chi_{\mu^x}(t) = \frac{\partial \chi(\tau_t, \mu_t^x)}{\partial \mu_t^x} < 0$. μ_t^x is equal to the expected discounted value of future marginal costs, $\mu_t^x = (1-\delta_x)\beta E_t \frac{\lambda_{t+1}}{\lambda_t}\tau_{t+1}$. The derivative $\chi_{\mu^x}(t)$ represents an intertemporal substitution effect on the inventory decision: all else equal, if marginal costs are expected to be lower in the future relative to the present, it is optimal to defer inventory accumulation and run down inventory levels today. Compared to the baseline model where we identified a demand channel and a cost channel to the inventory decision, we can now think about a current and expected future cost channel in addition to the demand channel as key transmission mechanisms.

Introducing sticky prices adds an additional term to the comovement condition, which is now given by the following expression in the presence of wedges:

$$\left(\frac{\left(\frac{\xi}{\alpha_n} - 1\right) - \theta_u}{1+\theta_u} - \frac{y}{s}\frac{1}{\varepsilon_x} \right) \hat{y}_t - \frac{\theta_u}{1+\theta_u}\varepsilon_u \hat{k}_t + \theta_u \hat{q}_t^k - \frac{x}{s}\frac{1}{\varepsilon_x}\frac{(1-\delta_x)}{g^y}x_{t-1} - \frac{1+\frac{\xi}{\alpha_n}}{1+\theta_u}\hat{\phi}_t^e + \frac{\theta_u}{1+\theta_u}\hat{\phi}_t^l - \mu^x \hat{\mu}_t^x < 0, \tag{22}$$

for $\hat{y}_t > 0$. If expected discounted future marginal costs are low relative to today (for instance, due to the effect of a future expected increase in TFP), distributors have an incentive to run down inventories in the present. We note that this makes the comovement condition potentially more difficult to satisfy.[20]

Response to TFP News. Figure 7 reports the impulse responses of key model variables to news about a future permanent increase in TFP that will be realized in 8 quarters as anticipated.[21] In contrast to the results discussed in section 3.2 for the baseline model, consumption, investment, hours, utilization and output now rise on impact and then grow in subsequent periods. Inventories increase slightly on impact, however, it falls thereafter as output booms and only rises over the following periods.

From the perspective of our wedges analysis through the lens of our co-movement condition (22), sticky wages cause a drop in the labour supply wedge ϕ_t^{ls} on impact. This shifts the output supply curve outward and contains the initial rise in output price τ_t, thereby allowing inventories to increase along with hours and output. In the following periods, however, the rise in Y_t drives up marginal costs, making condition (22) more difficult to satisfy without further endogenous shifts in output demand or supply. In fact, the gradual adjustment of nominal wages over time means that wage markups rise back towards their steady-state levels. As a consequence, the effect of the labor supply wedge ϕ_t^{ls} diminishes through the expansion.

We therefore conclude that the sticky wage and price model only achieves one of the two requirements for wedges that we discussed earlier. While sticky wages produce a drop in the labor wedges on impact, there is no further sustained decline in either the labor or efficiency wedges over the ensuing periods to overcome the rise in marginal costs from the rise in output. Thus, inventories fall over time while the rest of the economy booms.

3.4.2 Learning-by-doing model

Our second candidate model uses real rigidities to generate endogenous movement in the labor wedges. Specifically, we allow for time-variation in the production input H of the baseline model. One interpretation of this input is as a type of intangible capital that we refer to as knowledge capital. Following Chang et al. (2002a) and Cooper and Johri (2002), we assume that this input evolves as an internalized learning-by-doing process to capture the idea that agents acquire new

[20] We emphasize that the additional $\hat{\mu}_t^x$ term in (22) is due to sticky prices, not sticky wages. In a version of the model with sticky wages but flexible prices, the distributor's pricing condition implies that the markup between marginal cost of sales and price is constant, as in the baseline model and thus the additional $\hat{\mu}_t^x$ term would drop out of (22).

[21] We detail the values of the additional parameters unique to the sticky wage and price model in the Appendix D.

technological knowledge through their experiences in engaging labor in the production process.[22]

Introducing Knowledge Capital in the Baseline Model. We assume that the acquired technological knowledge resides with the firm. This has the distinct advantage that relative to the baseline model the modification only impacts the specification of the intermediate goods firm. The respective firm now produces the homogeneous good Y_t using the technology:

$$Y_t = z_t \left(\Omega_t N_t\right)^{\alpha_n} \widetilde{K}_t^{\alpha_k} \left(\Omega_t H_t\right)^{1-\alpha_n-\alpha_k},$$

where the stock of time-varying knowledge capital H_t evolves according to:

$$H_{t+1} = (1-\delta_h)H_t + H_t^{\gamma_h} N_t^{1-\gamma_h}, \qquad \text{where} \qquad 0 \leq \delta_h \leq 1, \quad 0 \leq \gamma_h < 1, \quad \nu_h > 0. \tag{23}$$

The knowledge capital accumulation (23) nests a log-linear specification for $\delta_h = 1$ common in the literature such as in Chang et al. (2002a), Cooper and Johri (2002) and d'Alessandro et al. (2019), but also allows for a more general linear formulation for $0 < \delta_h < 1$.[23]

The intermediate goods firm's optimization problem now involves choosing N_t, $\widetilde{K}_t$ and H_{t+1} to maximize $E_0 \sum_{t=0}^{\infty} \frac{\beta^t \lambda_t}{\lambda_0} \Pi_t^y$ subject to the production function and knowledge capital accumulation equation, where $\Pi_t^Y = \tau_t Y_t - w_t N_t - r_t \widetilde{K}_t$. Relative to the baseline 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 (23), these are given by, respectively:

$$w_t = \tau_t \alpha \frac{Y_t}{N_t} + q_t^h (1-\gamma_h) \frac{H_t^{\gamma_h} N_t^{1-\gamma_h}}{N_t}, \tag{24}$$

$$q_t^h = \beta E_t \frac{\lambda_t}{\lambda_{t+1}} \left\{ (1-\alpha_n-\alpha_h)\tau_{t+1} \frac{Y_{t+1}}{H_{t+1}} + q_{t+1}^h \left(1 - \delta_h + \gamma_h \frac{H_{t+1}^{\gamma_h} N_{t+1}^{1-\gamma_h}}{H_t} \right) \right\}. \tag{25}$$

The presence of internalized knowledge capital in the firm's technology adds an additional term into the firm's hours-worked first order condition (24) that shifts labor demand. A rise in the value of knowledge capital, q_t^h, increases labor demand as the firm attempts to increase H_t. Then

[22] The idea of learning-by-doing, and in particular skill-accumulation through work experience, has a long history in labor economics, where empirical researchers have found a significant effect of past work effort on current wage earnings. Learning-by-doing also plays a key role in growth, e.g., Arrow (1962). The general aspect of learning-by-doing as a supply-side mechanism that enhances the dynamics of business cycle models is, of course, not new. Both Chang et al. (2002a) and Cooper and Johri (2002) study the propagation properties of learning-by-doing in the context of business cycle models. Since then various researchers have exploited these properties to help business cycle models better fit various features of the data. This includes Gunn and Johri (2011), who show how learning-by-doing can yield comovement of consumption, investment, hours worked, and stock prices in response to TFP news. More recently, d'Alessandro et al. (2019) extend a standard New Keynesian model with learning-by-doing to account for the response of various macroeconomic aggregates to a government spending shock.

[23] In specification (23), knowledge capital is stationary on the balanced growth path due to the stationarity of hours worked. This implies that the long-run growth path of output is determined by exogenous technological factors only. This form of knowledge capital can be thought of as an index that conditions on the effect of hours in production over the business cycle as the firm responds to fluctuations in the exogenous stochastic drivers of growth.

(25) 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.

Knowledge Capital and Labor Wedges. We can write equation (24) as:

$$\frac{\tau_t}{w_t/(\alpha_n \frac{Y_t}{N_t})} = \frac{\tau_t}{mc_t} = 1 - q_t^h(1-\gamma_h)\left(\frac{H_t^{\gamma_h} N_t^{1-\gamma_h}}{w_t N_t}\right).$$

Given our definition of the labor demand wedge $\phi_t^{ld} = \frac{\tau_t}{mc_t}$ it then follows that this wedge in the learning-by-doing model is given by:

$$\phi_t^{ld} = 1 - q_t^h(1-\gamma_h)\left(\frac{H_t^{\gamma_h} N_t^{1-\gamma_h}}{w_t N_t}\right). \tag{26}$$

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$_t$ that acts like a markup. When the value of knowledge q_t^h is high, the firm increases hours-worked in order to increase knowledge, thereby decreasing the markup. Similarly, we can derive a modified efficiency wedge:

$$\frac{Y_t}{z_t\left(\Omega_t N_t\right)^{\alpha_n} \widetilde{K}_t^{\alpha_k}\left(\Omega_t H\right)^{1-\alpha_n-\alpha_k}} = \frac{1}{H} H_t^{1-\alpha_n-\alpha_k}.$$

By virtue of H_t being predetermined in production, the efficiency wedge does not move on impact. Rather, it grows over time as the firm accumulates knowledge, shifting the firm's marginal product of labor.

Overall, the learning-by-doing specification results in two wedges: a labor demand wedge ϕ_t^{ld} which moves on impact with the arrival of TFP news as the firm seeks to ramp-up production and reduce its markup; and an efficiency wedge ϕ_t^e, which reflects the gradual increase of knowledge in the production function, putting downward pressure on the marginal cost of production.

Response to TFP News. Figure 8 reports the impulse responses of the learning-by-doing specification to the same 8-quarter ahead TFP news shock as considered before. Notably, inventories now rise on impact and then increase in the ensuing periods along with the other major macroeconomic variables. We can again understand this response through the perspective of our wedges analysis and the co-movement condition (20) for flexible wages and prices.

The value of an incremental unit of knowledge, q_t^h, depends on the additional future profits that it returns for the firm (see the firm's h_{t+1} first-order condition, (25)). When news of higher future TFP arrives, the firm anticipates that output and profits will be higher in the future relative to today. This increases the marginal product of knowledge in the future in a manner that is complementary to the effect of higher TFP and physical capital. The rise in q_t^h shifts the firm's labor demand outwards as it seeks to increase its knowledge by using additional labor (see the firm's first-order condition (24)). In effect, the rise in the value of knowledge causes the firm to increase hours and to lower the markup between the output price τ_t and the marginal cost of production, mc_t, which reduces the labor demand wedge ϕ_t^{ld}. This shifts the output supply curve outward on impact, which limit the rise in τ_t and allows inventories to increase along with hours and output.

As the firm accumulates additional knowledge capital in subsequent periods, the efficiency wedge gradually rises. This offsets the rise in marginal costs over time on account of growing output demand that shifts the output supply curve increasingly outwards. Consequently, the increase in τ_t over time is limited, which in turn allows inventories to rise along with the other macroeconomic variables. This efficiency wedge effect thereby allows the co-movement condition (20) to be satisfied in the following periods after impact with increasingly higher levels of output.

Overall, the baseline model with knowledge capital achieves both requirements for wedges that are needed to facilitate the rise in inventories: the fast-moving labor demand wedge ϕ_t^{ld} that falls on impact of the news shock, and the sustained rise in the efficiency wedge ϕ_t^{ld} over the following periods, which is needed to overcome the rise in marginal costs from sustained growth in output demand.

4 Bayesian estimation

The analysis above shows that in a standard news shock model with inventories, a knowledge capital wedge can generate a positive inventory response alongside an expansion in all other macroeconomic aggregates in response to a TFP news shock. We also show that while nominal rigidities are not enough on their own, they help with the model's qualitative performance. We now go a step beyond this analysis and assess the performance of a full version of the model more formally. The specification features both knowledge capital and sticky wages and prices and it allows the TFP news shocks to compete with other disturbances found relevant in the literature.

We estimate the model using Bayesian methods. The specification of the shock processes, the treatment of observables, and prior choice is standard and close to related studies such as Smets and Wouters (2007) or Schmitt-Grohe and Uribe (2012). Further details on calibrated parameters and prior distributions can be found in appendix F. We estimate the model over the horizon 1983:Q1 - 2018:Q2, which is the same as in the VAR analysis. We use eight observables: output, consumption, investment, inventories, hours worked, wages, the nominal interest rate and the inflation rate. These are the seven observables of Smets and Wouters (2007) plus inventories.

We consider eight stochastic processes: a shock to the level of stationary TFP (z_t), a shock to the growth rate of non-stationary TFP (g_t^{Ω}), a shock to the growth rate of non-stationony IST (g_t^{Υ}), a marginal efficiency of investment (MEI) shock (m_t), a preference shock (Γ_t), a government spending shock (ε_t), a wage markup shock (ν_t^w), a price markup shock (ν_t^p) and a monetary policy shock η_t. Each exogenous disturbance is expressed in log-deviations from its mean as an AR(1) process, whose stochastic innovation is uncorrelated with other shocks, has zero mean, and is normally distributed. In addition to the unanticipated innovations to the above shocks, the model allows for anticipation effects for the stationary and non-stationary TFP processes as well as the non-stationary IST processes. Our treatment of anticipated and unanticipated components is standard and in line with the literature. Further details about the shock processes are provided in appendix E.2.1. For the processes with anticipated components we include four, eight and twelve quarter ahead innovations. The prior means assumed for the news components imply that the sum of the variance of news components is, evaluated at prior means, at most one half of the variance of the corresponding unanticipated component.

The full set of posterior estimates is reported in appendix F. Broadly speaking, the posterior parameter means are in line with those found in the literature. The estimated model features a highly elastic labor supply, a weak wealth effect (via Greenwood et al. (1988) preferences) and a typical degree of habit formation. There is a high degree of capital adjustment costs, while nominal adjustments costs (wage and price adjustment and indexation parameters) are reduced relative to the prior and smaller than in comparable New Keynesian settings. This indicates that much of the persistence arises from real rigidities, which is also borne out by the estimates of the shock parameters. In terms of model fit, we compare the New Keynesian model with knowledge capital to a version without this feature. The knowledge capital version scores considerably higher on account of the (log) marginal data density (-1303.6 vs -1318.5). While there is an implicit penalty for model complexity, the model with knowledge capital easily overcomes it.

In Figure 9, we report the impulse response functions at the estimated median value for all parameters to a news shock, specified as the arrival of news on an anticipated and realized increase in permanent TFP 8 periods out. From this figure it is evident that the estimated model generates responses to an anticipated TFP shock that are qualitatively consistent with the empirical responses reported in section 2 and those in the illustrative discussion in Section 3.4.2: all macroeconomic aggregates, including inventories, rise in light of news about higher future TFP, fuelled by a strong rise in the accumulation of knowledge capital.[24] These results provide evidence in favor of the news shock view of aggregate fluctuations since anticipated technology shocks can in principle replicate the unconditional comovement of output, investment, consumption, hours and inventories observed over the business cycle. In fact, we find this notion confirmed by the results of a variance decomposition exercise summarized in Table 1. TFP news shocks explain 30.0% of the fluctuations in output growth and also a relatively similar share in the variations of the other macroeconomic aggregates.

5 Conclusion

Our paper makes two contributions to the literatures on news shocks and inventory dynamics. First, based on standard VAR identification, we establish robust empirical evidence that an anticipated future rise in TFP raises inventory holdings in the present and induces positive comovement with other macroeconomic aggregates. Our evidence corroborates the view that TFP news shocks are important drivers of macroeconomic fluctuations. Moreover, it provides an additional dimension along which standard inventory frameworks can be evaluated as to their empirical viability. This is where our second contribution lies.

We show that the standard theoretical model used in the news shock literature, augmented with a standard inventory framework, cannot explain procyclical inventory movements in response to TFP news shocks. We discuss conditions that allow for a procyclical inventory response and employ a general wedges approach to show analytically on which margin and in which direction the wedges have to operate. This analysis suggest two potential frameworks, nominal rigidities in form of sticky wages and prices and a real rigidity in form of an additional factor of production, namely knowledge accumulated via learning-by-doing in production. We show that knowledge capital is the more likely candidate needed to capture the behavior of inventories. An estimated New Keynesian model with knowledge capital generates an expansion in all macroeconomic aggregates,

[24] We also investigate the model's ability to capture the typical behavior in response to other shocks, e.g. to unanticipated TFP shocks. We report selected impulse responses in appendix F.1. Overall, the responses are very much consistent with what is found in the literature.

including inventories, in response to news about higher future TFP.

References

Arrow, K. (1962). The economic implications of learning by doing. *Review of Economic Studies*, 3(29):155–173.

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

Barsky, R. B. and Sims, E. R. (2012). Information, animal spirits, and the meaning of innovations in consumer confidence. *American Economic Review*, 102(4):1343–77.

Beaudry, P. and Portier, F. (2004). An exploration into Pigou's theory of cycles. *Journal of Monetary Economics*, 51(6):1183–1216.

Beaudry, P. and Portier, F. (2014). News driven business cycles: Insights and challenges. *Journal of Economic Literature*, 52(4):993–1074.

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

Blinder, A. S. and Maccini, L. J. (1991). Taking stock: A critical assessment of recent research on inventories. *Journal of Economic Perspectives*, 5(1):73–96.

Cascaldi-Garcia, D. and Vukotic, M. (2020). Patent-based news shocks. *Review of Economics and Statistics*, forthcoming.

Chang, Y., Gomes, J. F., and Schorfheide, F. (2002a). Learning-by-Doing as a Propagation Mechanism. *American Economic Review*, 92(5):1498–1520.

Chang, Y., Gomes, J. F., and Schorfheide, F. (2002b). Learning-by-doing as a propagation mechanism. *American Economic Review*, 92(5):1498–1520.

Chari, V. V., Kehoe, P. J., and McGrattan, E. R. (2007). Business cycle accounting. *Econometrica*, 75(3):781–836.

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

Crouzet, N. and Oh, H. (2016). What do inventories tell us about news-driven business cycles? *Journal of Monetary Economics*, 79:49–66.

d'Alessandro, A., Fella, G., and Melosi, L. (2019). Fiscal Stimulus with Learning-By-Doing. *International Economic Review*, 60(3):1413–1432.

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

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

Francis, N., Owyang, M., Roush, J., and DiCecio, R. (2014). A flexible finite-horizon alternative to long-run restrictions with an application to technology shocks. *Review of Economics and Statistics*, 96:638–647.

Galí, J. and Gambetti, L. (2009). On the sources of the great moderation. *The American Economic Journal: Macroeconomics*, 1:26–57.

Görtz, C., Gunn, C., and Lubik, T. A. (2019). What Drives Inventory Accumulation? News on Rates of Return and Marginal Costs. Working Paper 19-18, Federal Reserve Bank of Richmond.

Görtz, C. and Tsoukalas, J. (2017). News and financial intermediation in aggregate fluctuations. *Review of Economics and Statistics*, 99(3):514–530.

Görtz, C., Tsoukalas, J., and Zanetti, F. (2021). News Shocks under Financial Frictions. *American Economic Journal: Macroeconomics*, forthcoming.

Greenwood, J., Hercowitz, Z., and Huffman, G. (1988). Investment, capacity utilization, and the real business cycle. *The American Economic Review*, 78:402–217.

Gunn, C. and Johri, A. (2011). News and knowledge capital. *Review of Economic Dynamics*, 14(1):92–101.

Iacoviello, M., Schiantarelli, F., and Schuh, S. (2011). Input And Output Inventories In General Equilibrium. *International Economic Review*, 52(4):1179–1213.

Jaimovich, N. and Rebelo, S. (2009). Can news about the future drive the business cycle? *American Economic Review*, 99(4):1097–1118.

Kogan, L., Papanikolaou, D., Seru, A., and Stoffman, N. (2017). Technological Innovation, Resource Allocation, and Growth. *The Quarterly Journal of Economics*, 132(2):665–712.

Kurmann, A. and Sims, E. (2019). Revisions in utilization-adjusted tfp and robust identification of news shocks. *Review of Economics and Statistics*, forthcoming.

Lubik, T. A. and Teo, W. L. (2012). Inventories, inflation dynamics and the new keynesian phillips curve. *European Economic Review*, 56(3):327–346.

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

Ramey, V. A. and West, K. D. (1999). Inventories. In Taylor, J. B. and Woodford, M., editors, *Handbook of Macroeconomics*, volume 1 of *Handbook of Macroeconomics*, chapter 13, pages 863–923. Elsevier.

Sarte, P.-D., Schwartzman, F., and Lubik, T. A. (2015). What inventory behavior tells us about how business cycles have changed. *Journal of Monetary Economics*, 76:264 – 283.

Schmitt-Grohe, S. and Uribe, M. (2011). Business cycles with a common trend in neutral and investment-specific productivity. *Review of Economic Dynamics*, 14:122–135.

Schmitt-Grohe, S. and Uribe, M. (2012). What's news in business cycles? *Econometrica*, 80(6):2733–2764.

Smets, F. and Wouters, R. (2007). Shocks and frictions in US business cycles: A Bayesian DSGE approach. *American Economic Review*, 97(3):586–606.

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

Tables and Figures

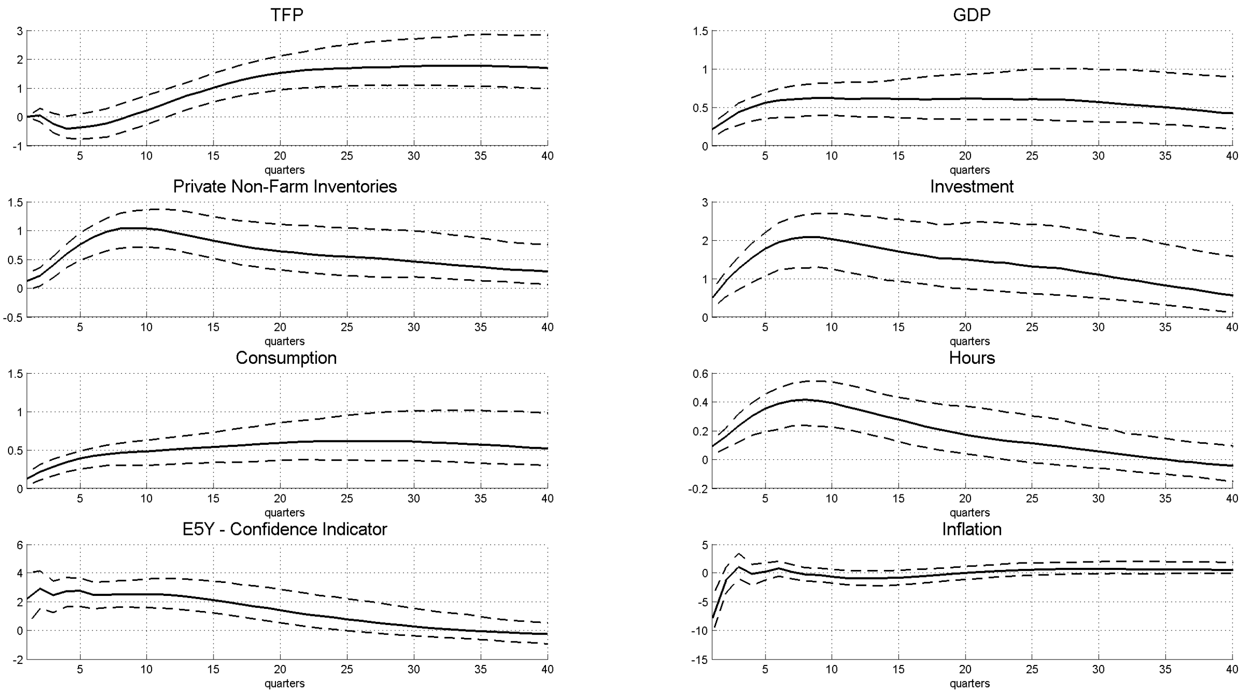

Figure 1: **IRF to TFP news shock – including Private Non-Farm Inventories.** Sample 1983Q1-2018Q2. The solid line is the median and the dashed lines are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

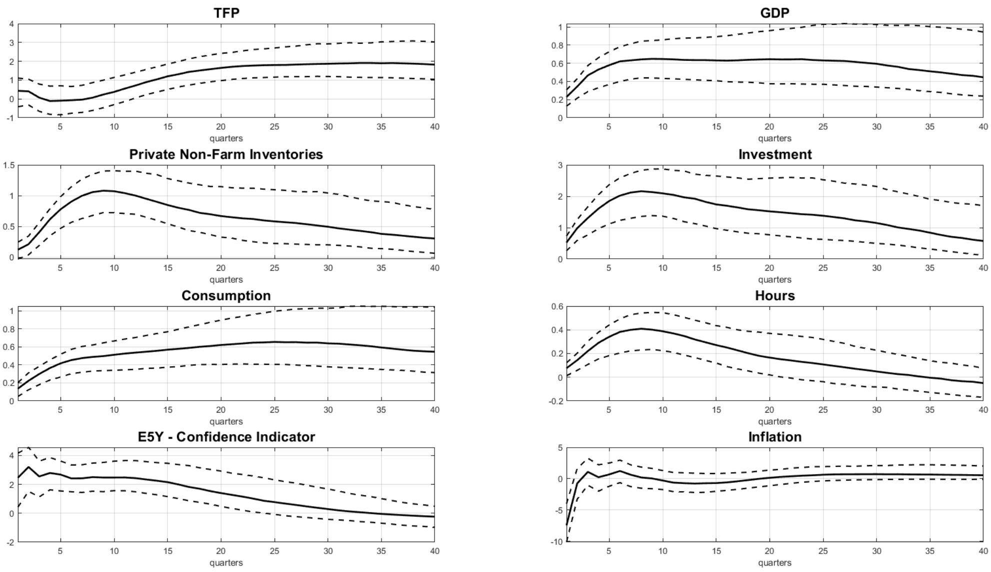

Figure 2: **IRF to TFP news shock. Kurmann-Sims identification.** Sample 1983Q1-2018Q2. The black solid line is the median response. The shaded dashed lines are the corresponding 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

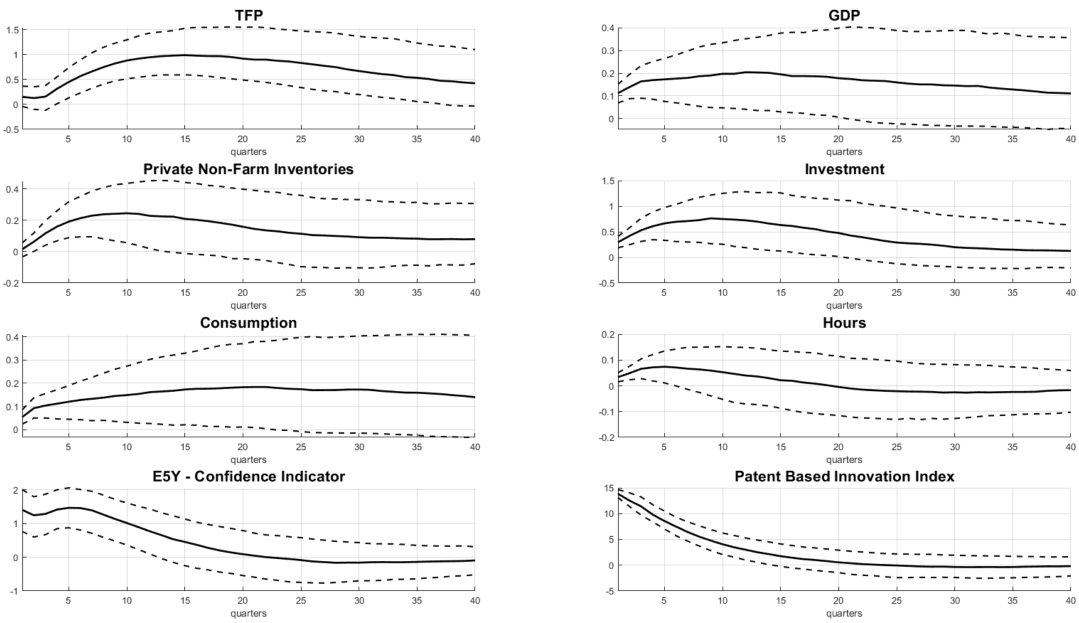

Figure 3: **IRF to patent based innovation shock.** The solid line is the median and the dashed lines are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

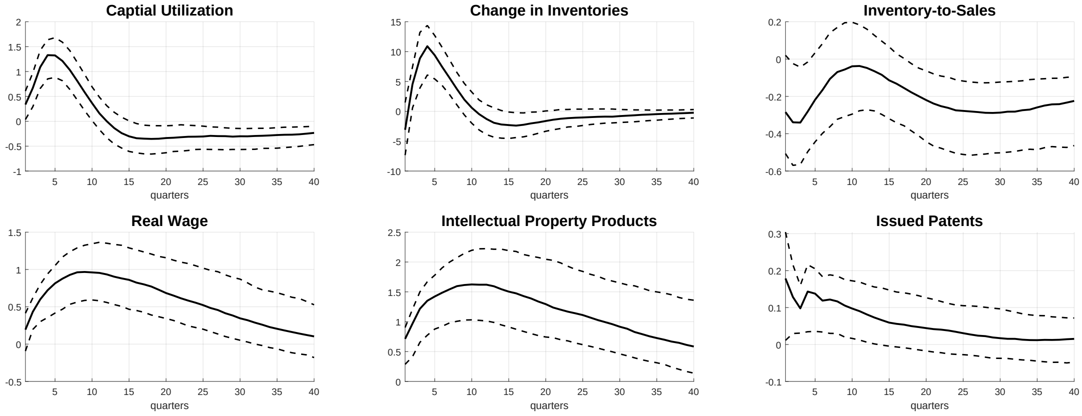

Figure 4: **IRF to TFP news shock. Max Share identification.** Subplots result from VARs comprising TFP, GDP, investment, hours, inflation and one of the plotted variables above at a time. The solid line is the median and the dashed lines are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

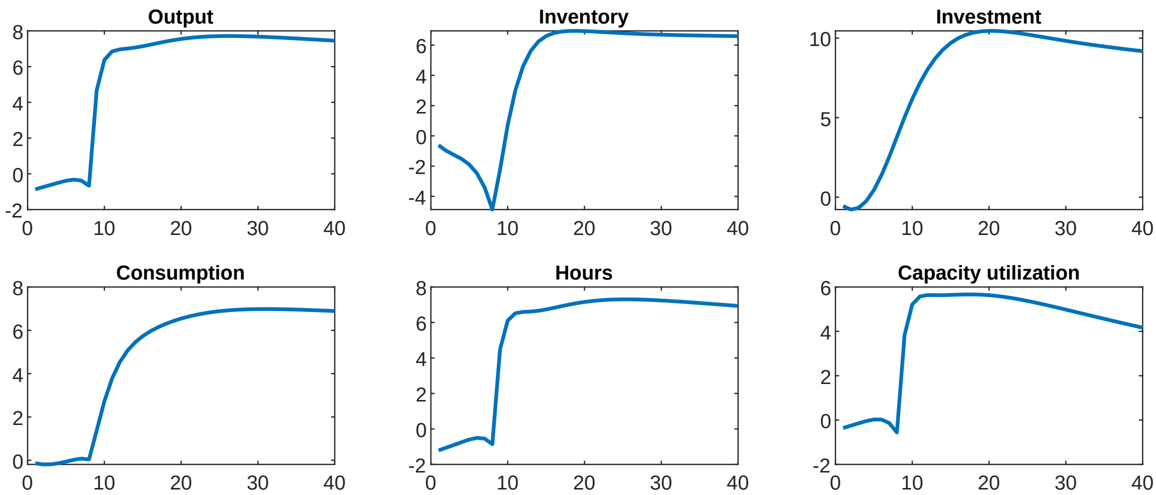

Figure 5: **IRF to 8-period out non-stationary TFP news shock: baseline model**

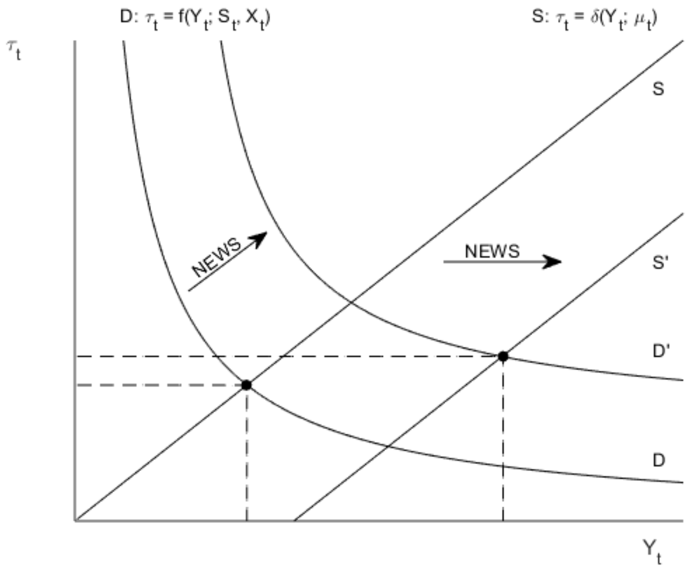

Figure 6: **Supply and Demand curves for Output, Y_t, and marginal cost, τ_t.**

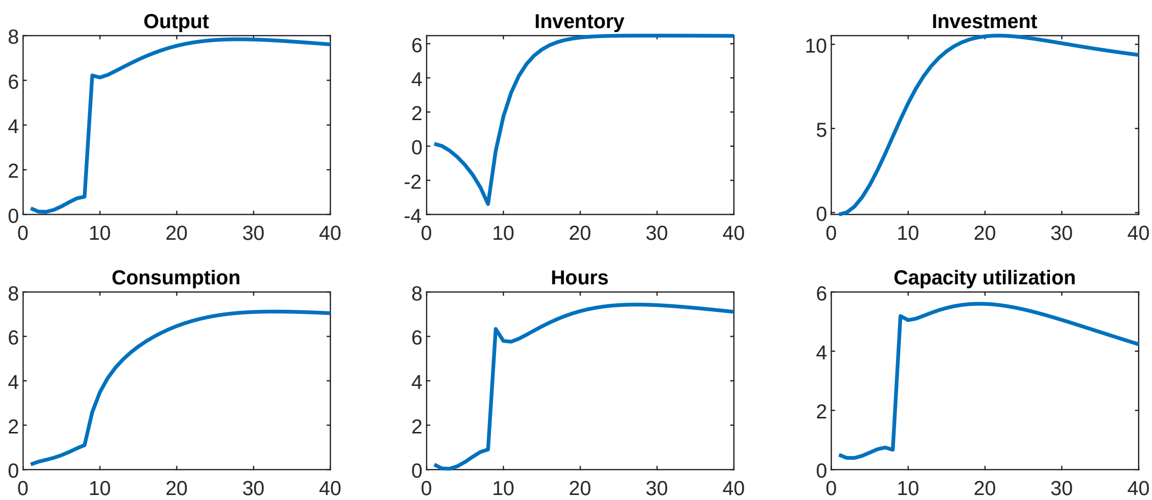

Figure 7: **IRF to 8-period out non-stationary TFP news shock:** ***Sticky wage and price model***

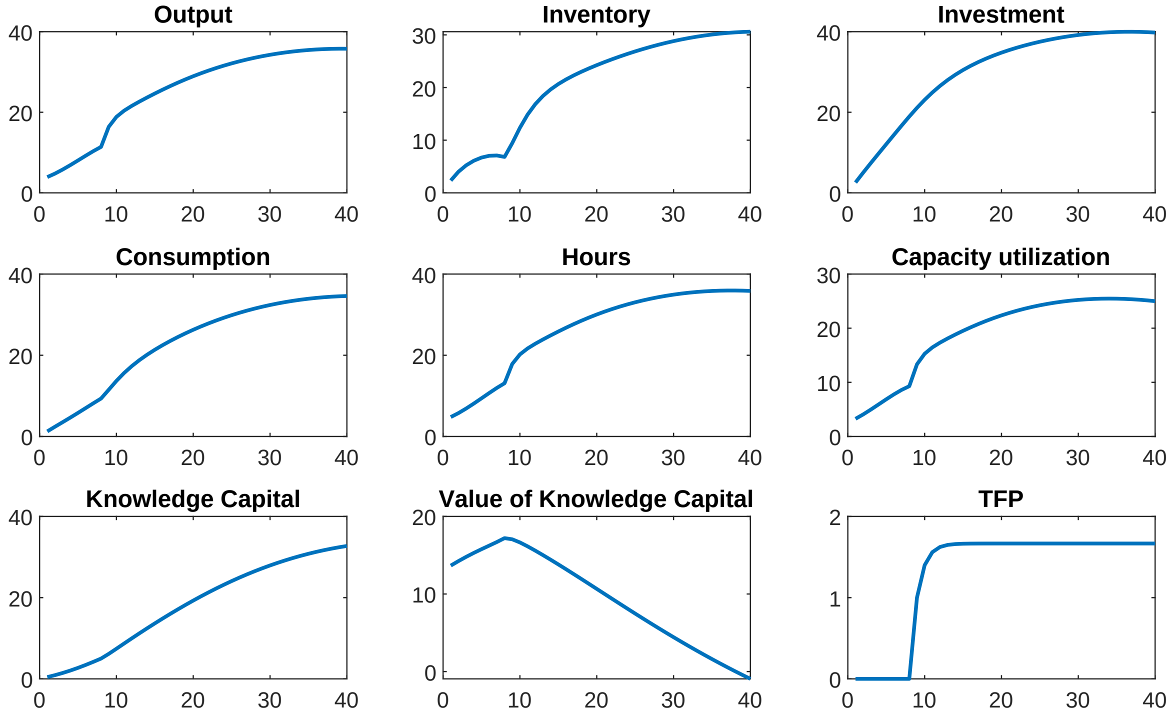

Figure 8: **IRF to 8-period out non-stationary TFP news shock:** ***Learning-by-doing model***

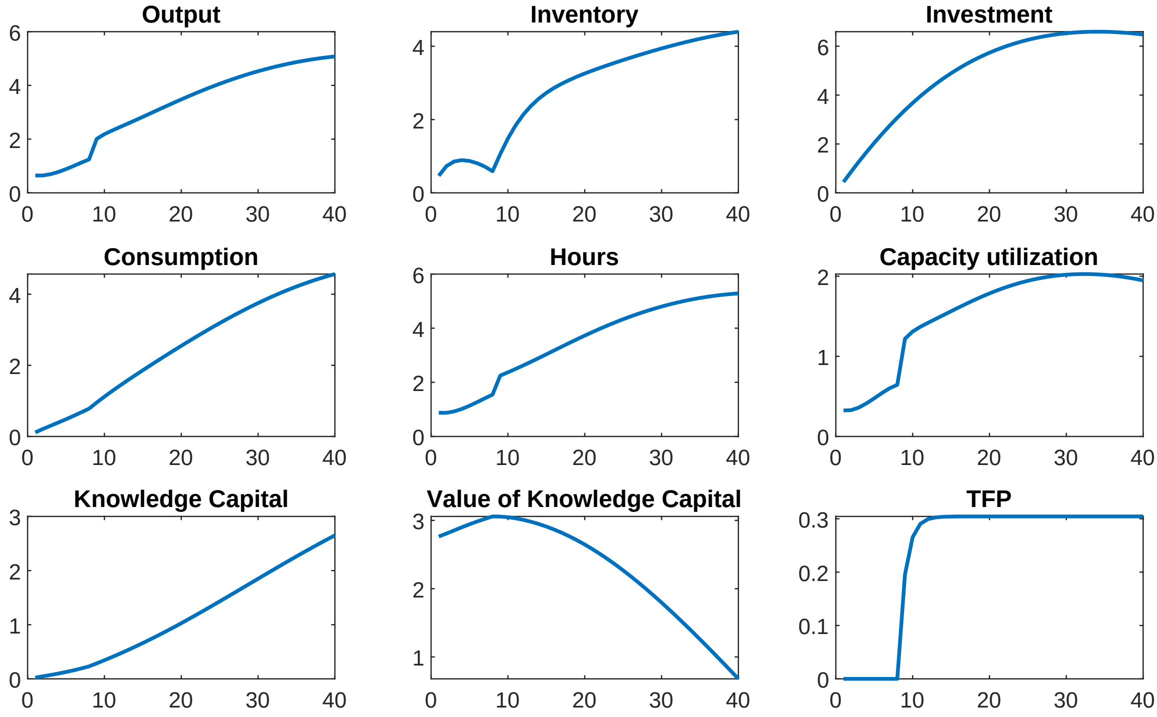

Figure 9: **IRF to 8-period out non-stationary TFP news shock: Estimated model (*Learning-by-doing + sticky wages and prices*)**

Table 1: Variance Decomposition: *NK Model with Knowledge Capital*

	Perm. TFP	*Perm. TFP News*	*All other shocks*
Output Growth	27.6	30.3	42.0
Investment Growth	26.4	29.7	43.9
Consumption Growth	30.5	35.6	34.0
Inventory Growth	25.4	26.8	47.8
Hours	32.1	37.6	30.3
Wage Growth	14.5	15.8	69.7
Interest Rate	26.0	29.3	44.6
Inflation	20.6	20.9	58.5

Online Appendix (not for publication)

A Details on the VAR model

This appendix provides details on the VAR model, the baseline news shock identification and prior specifications.

A.1 VAR-Based Identification of News Shocks

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

$$y_t = A(L)u_t. \tag{27}$$

$A(L) = I + A_1 L + ... + A_p L^p$ is a lag polynomial of order p over conformable coefficient matrices $\left\{A_p\right\}_{i=1}^{p}$. u_t is an error term with $n \times n$ covariance matrix Σ. We assume a linear mapping between the reduced form errors u_t and the structural errors ε_t:

$$u_t = B_0 \varepsilon_t, \tag{28}$$

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

$$y_t = C(L)\varepsilon_t, \tag{29}$$

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

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

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

$$y_{t+h} - E_{t-1}y_{t+h} = \sum_{\tau=0}^{h} A_\tau \widetilde{B}_0 D \varepsilon_{t+h-\tau}. \tag{30}$$

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

$$V_{i,j}(h) = \frac{e_i'\left(\sum_{\tau=0}^{h} A_\tau \widetilde{B}_0 D e_j e_j' D' \widetilde{B}_0' A_\tau'\right) e_i}{e_i'\left(\sum_{\tau=0}^{h} A_\tau \Sigma A_\tau'\right) e_i} = \frac{\sum_{\tau=0}^{h} A_{i,\tau} \widetilde{B}_0 \gamma\gamma' \widetilde{B}_0' A_{i,\tau}'}{\sum_{\tau=0}^{h} A_{i,\tau} \Sigma A_{i,\tau}'}, \tag{31}$$

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

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

$$V_{1,1}(h) + V_{1,2}(h) = 1, \tag{32}$$

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

The Max Share approach chooses the elements of $\widetilde{B}_0$ to make this restriction on forecast error variance share hold as closely as possible. This is equivalent to choosing the impact matrix so that contributions to $V_{1,2}(h)$ are maximized. Consequently, we choose the second column of the impact matrix to solve the following optimization problem:[25]

$$\begin{aligned} \arg\max_{\gamma} V_{1,2}(h) &= \frac{\sum_{\tau=0}^{h} A_{i,\tau} \widetilde{B}_0 \gamma\gamma' \widetilde{B}_0' A_{i,\tau}'}{\sum_{\tau=0}^{h} A_{i,\tau} \Sigma A_{i,\tau}'}, \\ \text{s.t. } \gamma\gamma' &= 1,\ \gamma(1,1) = 0,\ \widetilde{B}_0(1,j) = 0,\ \forall j > 1. \end{aligned} \tag{33}$$

We restrict γ to have unit length to be a column vector of an orthonormal matrix. The second and third constraints impose that a TFP news shock cannot affect TFP contemporaneously. We therefore identify a TFP news shock from the estimated VAR as the shock that: (i) does not move TFP on impact; and (ii) maximizes the share of variance explained in TFP at a long but finite

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

horizon h.

A.2 Specification for the Minnesota Prior in the VAR

We estimate the VAR using a Bayesian approach. The prior for the VAR coefficients A a standard Minnesota prior as commonly used in the literature. It is of the form

$$vec(A) \sim N\left(\underline{\beta}, \underline{V}\right),$$

where $\underline{\beta}$ is one for variables which are in log-levels, and zero for hours, the E5Y and inflation. The prior variance $\underline{V}$ is diagonal with elements,

$$\underline{V}_{i,jj} = \begin{cases} \frac{\underline{a}_1}{p^2} \text{ for coefficients on own lags} \\ \frac{\underline{a}_2 \sigma_{ii}}{p^2 \sigma_{jj}} \text{ for coefficients on lags of variable } j \neq i \\ \underline{a}_3 \sigma_{ii} \text{ for intercepts} \end{cases}$$

where p denotes the number of lags. Here σ_{ii} is the residual variance from the unrestricted p-lag univariate autoregression for variable i. The degree of shrinkage depends on the hyperparameters $\underline{a}_1, \underline{a}_2, \underline{a}_3$. We set $\underline{a}_3 = 1$ and we choose $\underline{a}_1, \underline{a}_2$ by searching on a grid and selecting the prior that maximizes the in-sample fit of the VAR, as measured by the Bayesian Information Criterion.[26]

B Additional Empirical Evidence

B.1 Forecast Error Variance Decomposition

Figure 10 displays the variance shares explained by the TFP news shock.

B.2 Longer Sample Periods

Changes in the behavior of inventories that coincide with the onset of the Great Moderation have been widely documented in the literature (e.g. McCarthy and Zakrajsek (2007)). This aspect motivates our focus on the Great Moderation sample in addition to data availability issues highlighted in the main body. However, it is interesting to evaluate whether the rise of inventories in anticipation of higher future TFP is present also when considering longer samples. Figure 11 shows this is indeed the case for the 1960Q1-2018Q2 sample. The figure reports strong comovement of

[26]The grid of values we use is: $\underline{a}_1$ = (1e-4:1e-4:9e-4, 0.001:0.001:0.009, 0.01:0.01:0.1, 0.1:0.1:1), $\underline{a}_2$ = (0.01,0.05,0.1,0.5,1,5). We consider all possible pairs of $\underline{a}_1$ and $\underline{a}_2$ in the above grids.

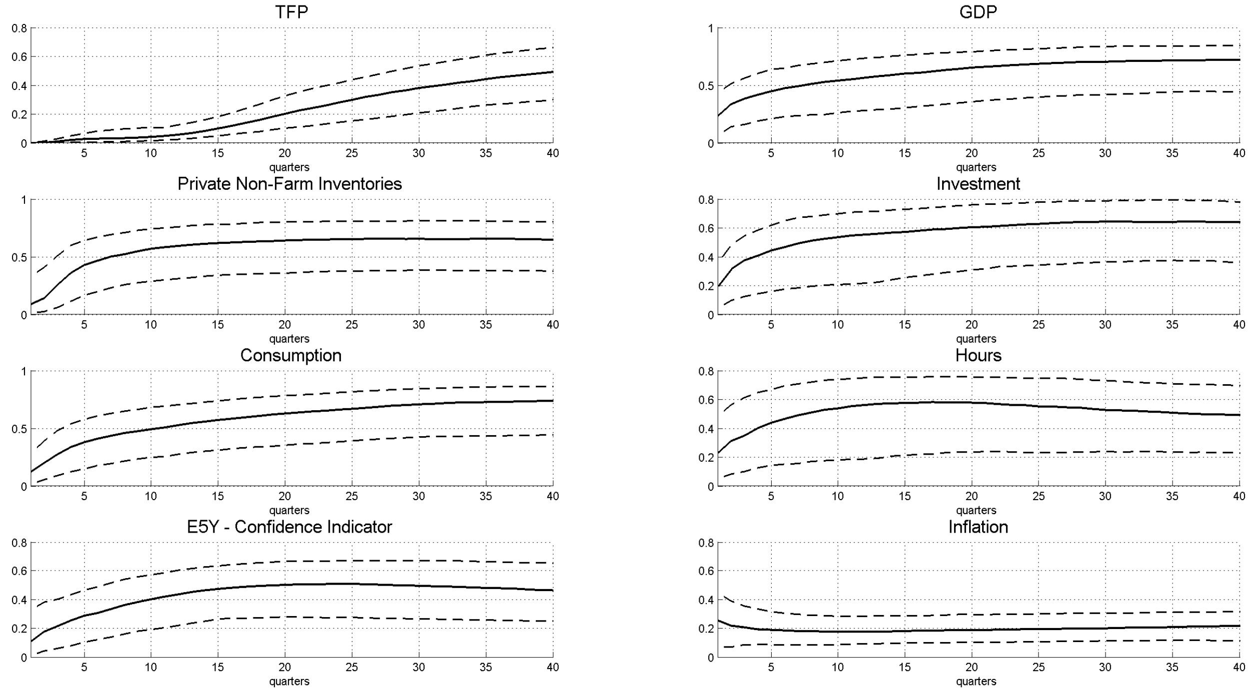

Figure 10: Forecast error variance decomposition (FEVD) of variables to the TFP news shock. Sample 1983Q1-2018Q2. Baseline identification. The solid line is the median and the dashed lines are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters.

all macroeconomic aggregates, including inventories, several quarters before TFP increases significantly. The sample is restricted by the availability of the E5Y. If we use the S&P500 instead we can consider a 1948Q1-2018Q2 sample. Figure 12 shows that IRFs based on this sample are qualitatively and largely also quantitatively very similar to the results based on our 1983Q1-2018Q2 baseline sample and the 1960Q1-2018Q2 sample. Overall, the fact that inventories rise in response to a TFP news shock is very robust if our baseline sample is extended.

B.3 The Response of Inventories across Sectors and Types of Inventories

This section provides additional evidence on the robustness of the procyclical response of inventories established in section 2.2 of the main text.

Figure 13 reports the impulse response functions of the specification with business inventories. By necessity, this sample is shorter as the inventory series and its subcomponents are only available since 1992Q1. We consider this alternative specification important as it is not a priori obvious at which prices inventories should be measured. The figure shows that the rise in inventories prior to TFP is robust when we use the business inventory series. All variables exhibit qualitative responses that are very similar to the baseline, although the shorter sample results in somewhat

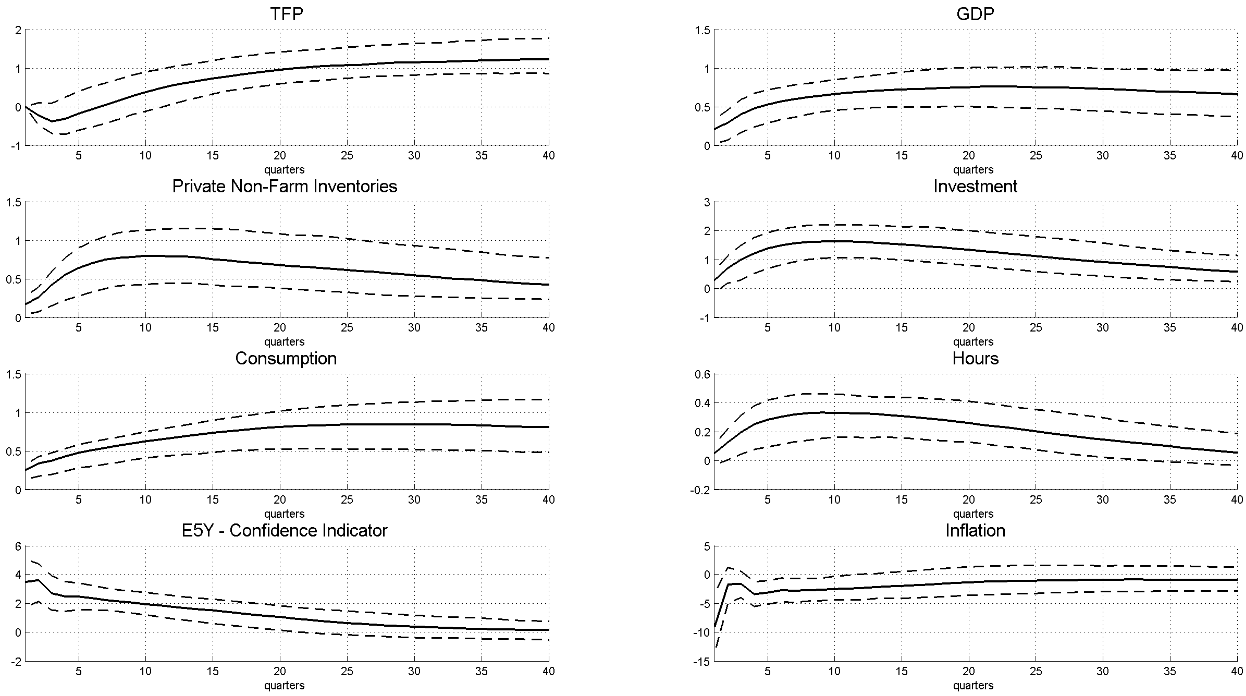

Figure 11: **IRF to TFP news shock. Sample 1960Q1-2018Q2.** The solid line is the median and the dashed lines are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

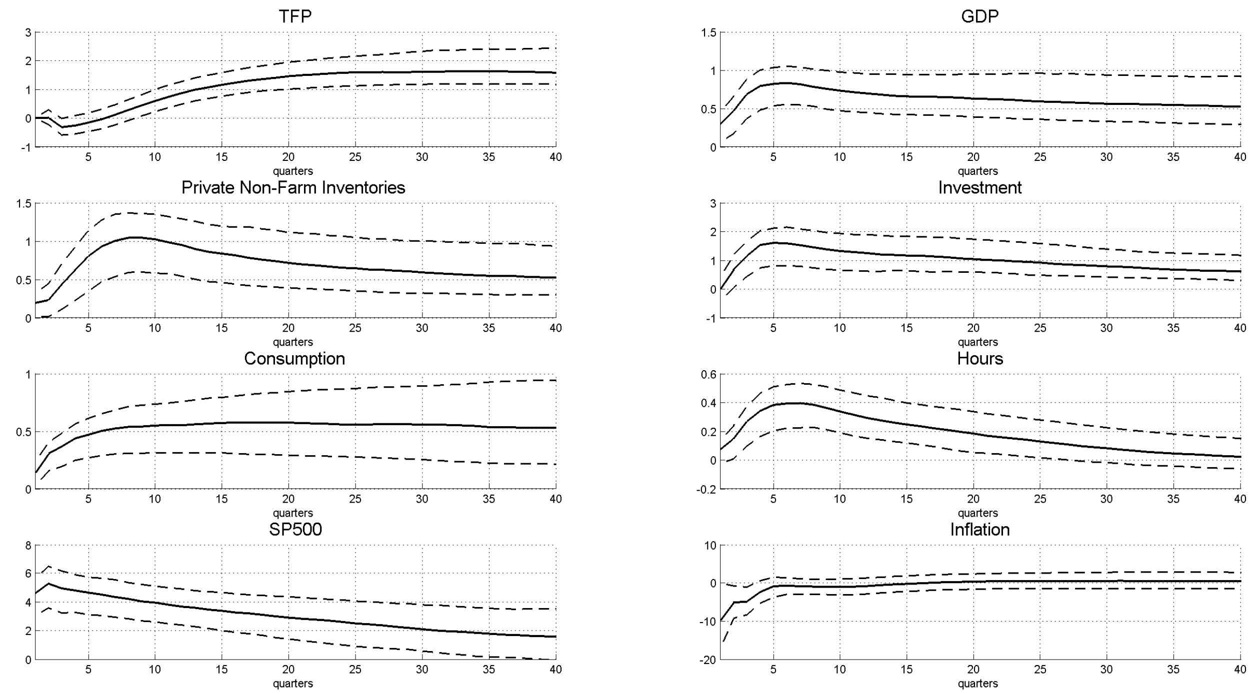

Figure 12: **IRF to TFP news shock. Sample 1948Q1-2018Q2.** The solid line is the median and the dashed lines are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

wider confidence bands. Overall, this specification confirms the comovement of macroeconomic aggregates, including inventories, in response to an anticipated TFP shock and prior to the rise in TFP itself.

In the next step, we study the effects of news shocks on inventories in the manufacturing, wholesale, and retail sectors, which comprise the overwhelming majority of inventory stocks. Figure 14 shows the responses of business inventories in each of these sectors to the aggregate TFP news shock. The VAR is estimated by including the sectoral inventories one by one instead of the aggregate inventory measure. The sectoral impulse responses exhibit almost identical hump-shaped patterns: a rise on impact towards a peak response around 10 quarters before declining gradually over the forecast horizon. These results support the finding from the aggregate baseline specification in that the expansion of the inventory stock and other variables is broad-based across sectors.

We also dig deeper into the composition of inventory holdings. The two trade sectors, wholesale and retail, hold almost entirely finished goods inventories, while the inventory stock in the manufacturing sector is split across finished goods inventories (36%), work in process (30%) and input inventories in the form of materials and supplies (34%) over the restricted 1992:Q2 - 2018:Q2 sample period for business inventories and their components. Figure 15 shows the responses of inventory types in the manufacturing sector when included one by one in the VAR.[27] Finished goods and input inventories in the manufacturing sector rise strongly before the realization of anticipated higher productivity as in the baseline specification and all other variations considered above.

Overall, we find the results documented in section 2.2 on the procyclicality of the inventory response, conditional on a TFP news shock, are very robust across the considered dimensions.

B.4 Evidence from Alternative Identification Schemes

The results in the main body of the paper are generated using the Max-share method proposed by Francis et al. (2014). This method is widely used in the literature and identifies a news shock as the shock that (i) does not move TFP on impact and (ii) maximizes the variance of TFP at the 40 quarter horizon. In addition, Section 2.3 provides robustness for our results using the method proposed by Kurmann and Sims (2019) that abstracts from the zero-impact restriction, and a patent

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

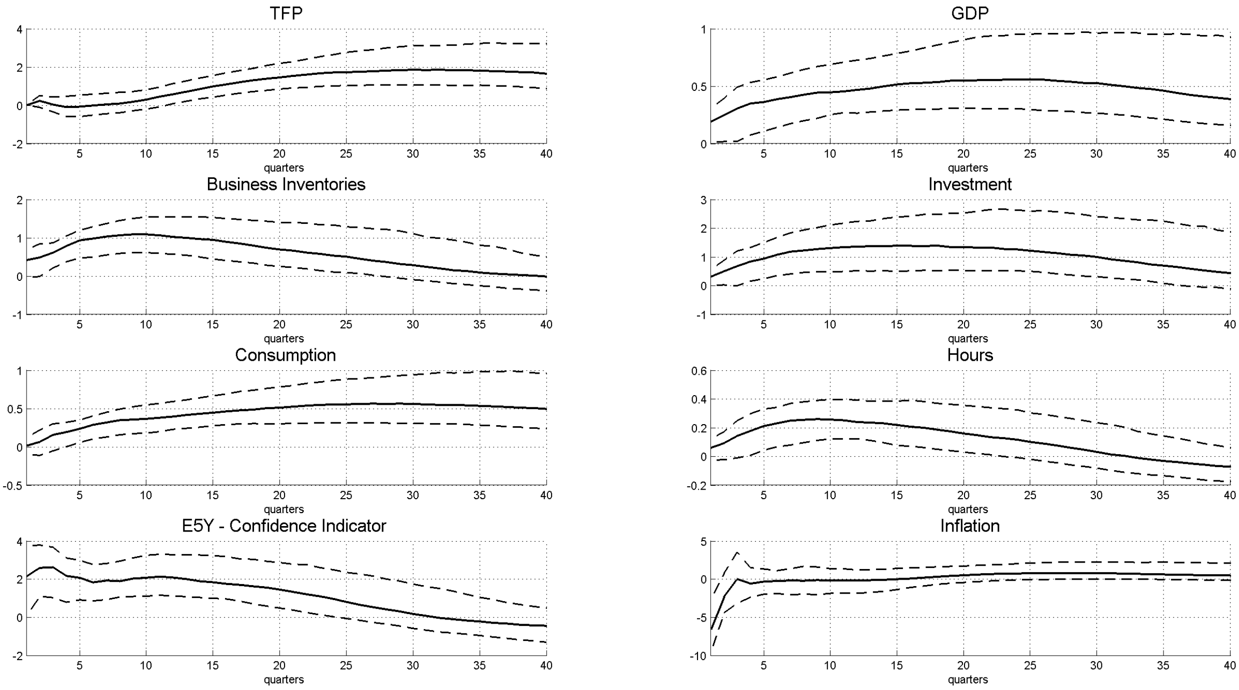

Figure 13: **IRF to TFP news shock – including Business Inventories. Max Share identification.** Sample 1992:Q1-2018:Q2. The solid line is the median and the dashed lines are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

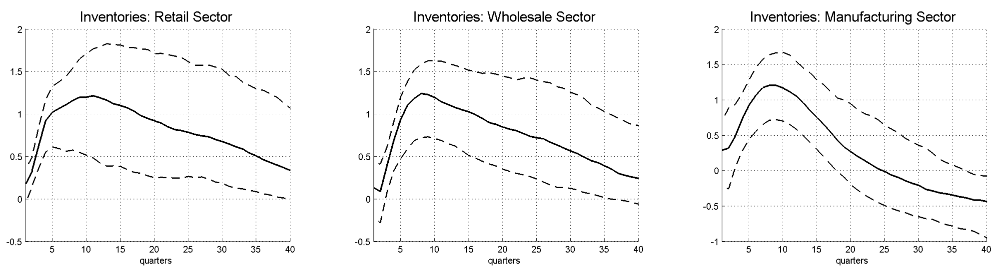

Figure 14: **IRF of business inventories by sector to TFP news shock. Max Share identification.** Sample 1992:Q1-2018:Q2. Subplots result from eight variable VARs comprising TFP, GDP, consumption, investment, hours, inventory measure, inflation, E5Y. The inventory measures were included one-by-one in the VAR system. The solid line is the median and the dashed lines are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

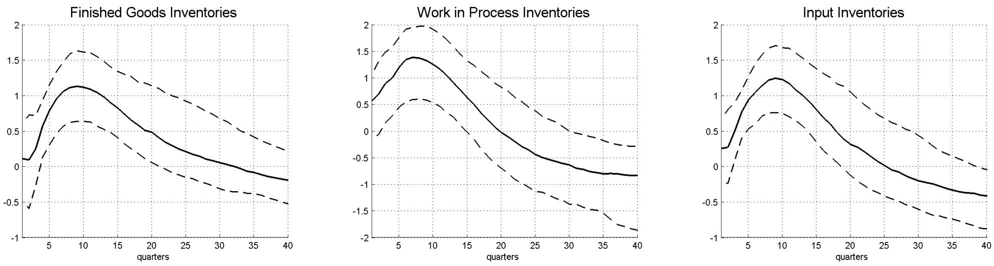

Figure 15: **IRF of business inventories in the manufacturing sector by inventory type to TFP news shock. Max Share identification.** Sample 1992:Q1-2018:Q2. Subplots result from eight variable VARs comprising TFP, GDP, consumption, investment, hours, inventory measure, inflation, E5Y. The inventory measures were included one-by-one in the VAR system. The solid line is the median and the dashed lines are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

based identification proposed by Cascaldi-Garcia and Vukotic (2020).

This section shows robustness of findings using two alternative approaches. First, the identification scheme suggested in Barsky and Sims (2011) that recovers the news shock by maximizing the variance of TFP over horizons from zero to 40 quarters, and the restriction that the news shock does not move TFP on impact; second, the Forni et al. (2014) long-run identification scheme, which is similar in spirit to the Max Share method. This method identifies the news shock by imposing the zero impact restriction on TFP and seeks to maximise the impact of the news shock on TFP in the long run. Both are closely related to the baseline and Kurmann and Sims (2019) identification in the sense that they also rely on the TFP measure to identify the news shock. Figure 16 provides a comparison between the median responses based on the Max share method and the methods proposed by Barsky and Sims (2011) and Forni et al. (2014). The median responses of the Max Share methodology and the Forni et al. (2014) methodology are virtually indistinguishable and also the median based on the Barsky and Sims (2011) methodology is very similar. Importantly, all macroeconomic aggregates, including inventories, rise in response to a TFP news shock.

B.5 Further Robustness Results from the Baseline Identification

Figure 17 shows the response of inventories from an eight-variable VAR that corresponds to Figure 1. When we vary the news identification horizon h, it is evident that the positive response of inventories obtained using $h = 40$ in the main body is robust for $h = 20$, $h = 30$, $h = 50$, $h = 60$ and $h = 80$. For different specifications of h, responses of all other variables are also very similar

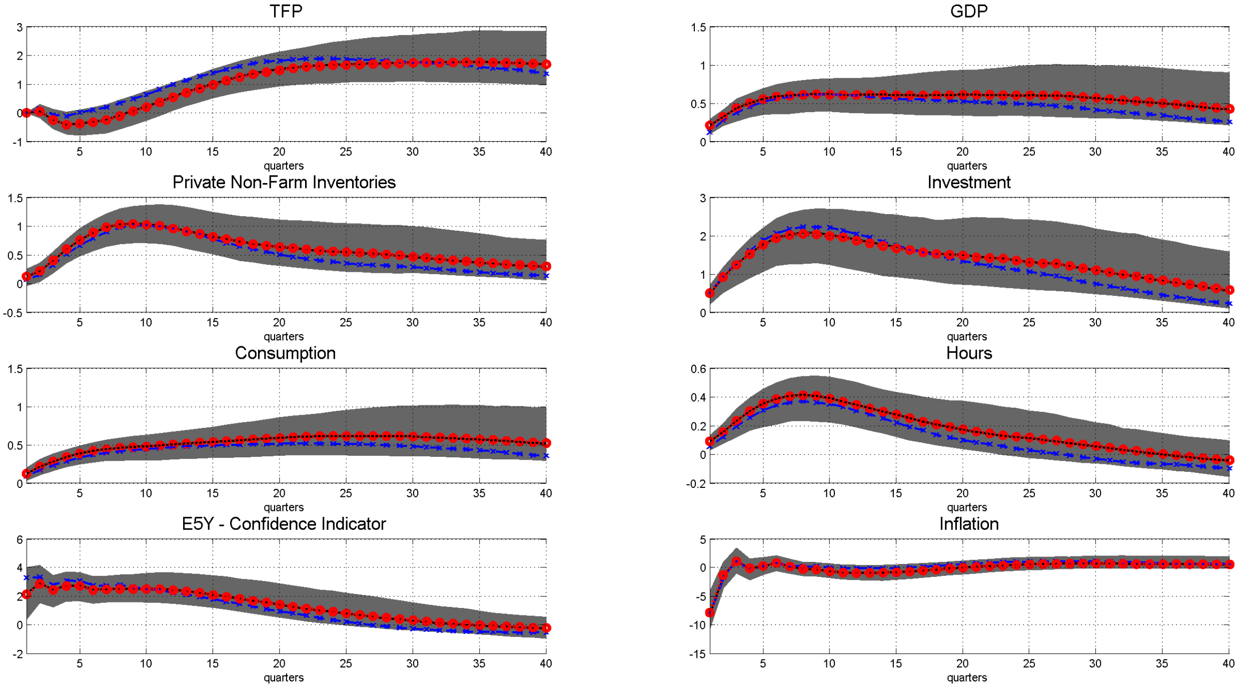

Figure 16: **IRF to TFP news shock.** Sample 1983Q1-2018Q2. The black solid line is the median response identified using the Max-share method. The shaded gray areas are the corresponding 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The blue line with crosses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

to the ones reported in Figure 1 and are available upon request.

Figure 18 shows that our result on the procyclicality of inventories to a TFP news shock is also robust when considering a very small-scale VAR.

Figure 19 shows IRFs from a VAR that corresponds to Figure 1, but where we replace GDP with sales. Overall, the results are very similar to those in Figure 1. Sales rises in response to the news shock and increase upon impact more than inventories.

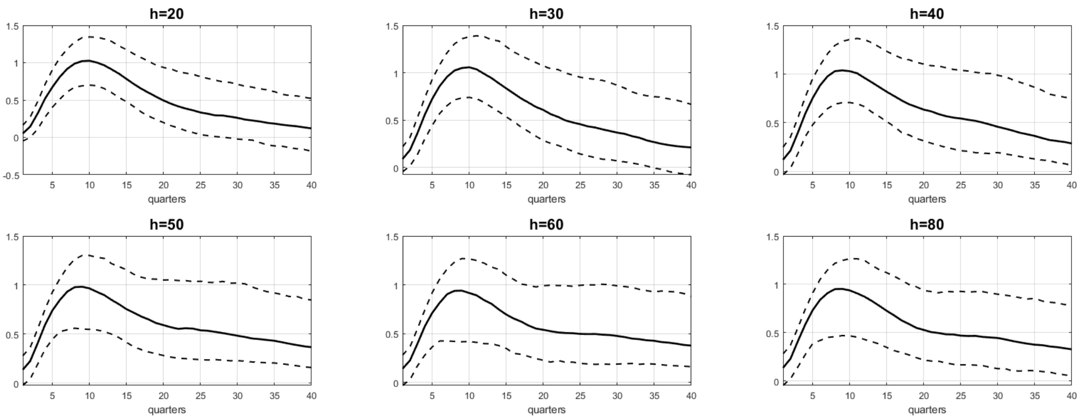

Figure 17: **Response of inventories to TFP news shock. Baseline identification.** The figure shows the response of private non-farm inventories in the eight-variable VAR in (main body) Figure 1 for different maximisation horizons *h* using the baseline Max Share identification. The solid line is the median and the dashed lines are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations

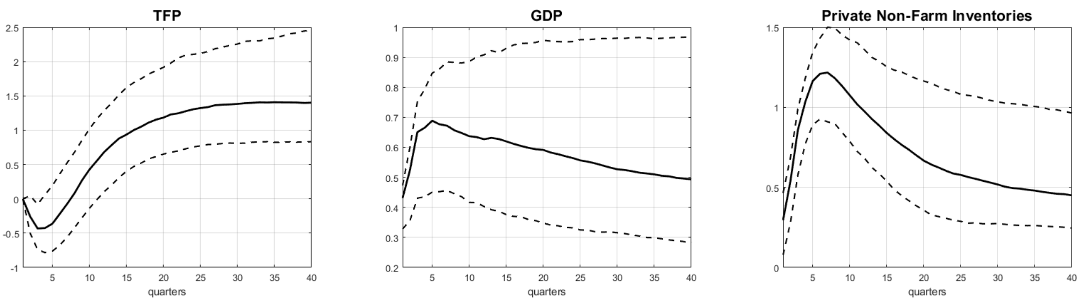

Figure 18: **IRF to TFP news shock. Baseline identification.** The shock is identified using the Max Share approach in a three-variable VAR. The solid line is the median and the dashed lines are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations

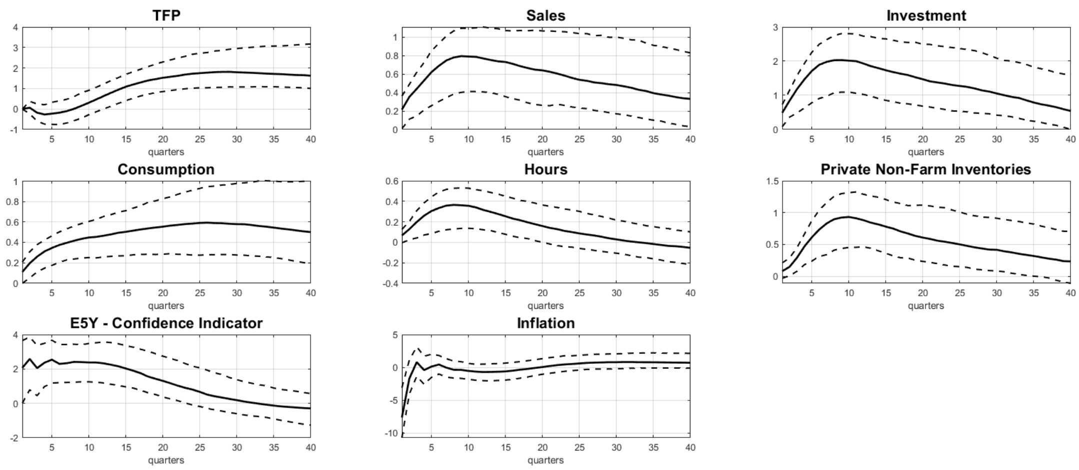

Figure 19: **IRF to TFP news shock.** The shock is identified using the Max Share approach. The solid line is the median and the dashed lines are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations

C Additional Model Details: Baseline Model

This appendix section details elements of the *Baseline Model* not shown in the main text.

C.1 Model Description: Baseline Model

C.1.1 Households and Government

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

$$E_0 \sum_{t=0}^{\infty} \beta^t \Gamma_t \frac{\left(V_t^{1-\sigma}-1\right)}{1-\sigma}, \tag{34}$$

where $0<\beta<1$, $\sigma>0$, and where Γ_t is a stationary stochastic preference shock process. The argument V_t is given by

$$V_t = C_t - bC_{t-1} - \psi N_t^{\xi} J_t, \tag{35}$$

where

$$J_t = \left(C_t - bC_{t-1}\right)^{\gamma_j} J_{t-1}^{1-\gamma_j} \tag{36}$$

is a preference component that makes consumption and labor non-time-separable and is consistent with the balanced-growth path in a growing economy. This preference structure, which follows Schmitt-Grohe and Uribe (2012) and is based on Jaimovich and Rebelo (2009), nests the no-income effect structure of Greenwood et al. (1988) in the limit as the parameter $0<\gamma_f\leq 1$ tends toward zero. The parameter $0\leq b<1$ allows for habits in consumption; and $\xi>1$ is related to the Frisch elasticity of labour supply (and is equal to it when $\gamma_j=b=0$).

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

$$K_{t+1} = \left[1-\delta\left(u_t\right)\right] K_t + m_t I_t \left[1 - S\left(I_t/I_{t-1}\right)\right], \tag{37}$$

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

The household's budget constraint is given by:

$$C_t + \Upsilon_t I_t + T_t = w_t N_t + r_t u_t K_t + \Pi_t, \tag{38}$$

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

The household chooses sequences of C_t, I_t, N_t, u_t and K_{t+1} to maximize intertemporal utility subject to the constraints above, resulting in standard first-order conditions.

C.1.2 Stochastic Exogenous Processes

The model includes six exogenous stochastic processes: a shock to the level of stationary TFP (z_t), a shock to the growth rate of non-stationary TFP (g_t^{Ω}), a shock to the growth rate of non-stationy IST (g_t^{Υ}), a marginal efficiency of investment (MEI) shock (m_t), a preference shock (Γ_t) and a government spending shock (ε_t). We assume that these stochastic processes follow individually stationary first-order processes and are mutually uncorrelated, given as

$$\ln\left(\frac{\vartheta_t}{\vartheta}\right) = \rho_{\vartheta} \ln\left(\frac{\vartheta_{t-1}}{\vartheta}\right) + e_{\vartheta,t}, \tag{39}$$

for $\vartheta = \{z, g^{\Omega}, g^{\Upsilon} m, \Gamma, \varepsilon\}$.

We allow for news shocks to both the stationary and non-stationary TFP shocks and assume that the innovations in these two stochastic proceses contain both anticipated and unanticipated components. Moreover, news signals arrive with horizons of 4, 8 and 12 quarters as is standard in the literature (see e.g. Görtz et al. (2021). The innovations are thus given by:

$$e_{\vartheta,t} = \begin{cases} e^0_{\vartheta,t} + e^4_{\vartheta,t-4} + e^8_{\vartheta,t-8} + e^{12}_{\vartheta,t-12}, & \vartheta = \{z, g_{\Omega}\} \\ e^0_{\vartheta,t}, & \vartheta = \left\{m, g^{\Upsilon}, \varepsilon_t, \Gamma_t\right\} \end{cases},$$

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

C.2 Model equilbrium, stationary and solution method: Baseline Model

In a symmetric equilibrium, $Y_{it} = Y_t^*$, $A_{it} = A_t^*$, $X_{it} = X_t^*$, $P_{it} = P_t^*$ and $S_{it} = S_t^*$ $\forall i$. It then follows that $Y_t = \int_0^1 Y_t^* di = Y_t^*$, $A_t = \int_0^1 A_t^* di = A_t^*$, $X_t = \int_0^1 X_t^* di = X_t^*$, 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

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

and implying that $\frac{P_{it}}{P_t} = 1 \quad \forall i$.

The resulting equilibrium model system consists of a symmetric competitive equilibrium as a set of stochastic processes $\{C_t, V_t, I_t, G_t, S_t, Y_t, N_t, u_t, J_t, K_t, X_t, A_t, w_t, r_t, \tau_t, \mu_t^j, \mu_t^k, \lambda_t\}_t^{\infty}$, given initial conditions and exogenous stochastic processes, and where μ_t^j, μ_t^k, and λ_t respectively denote the multipliers on the definition of J_t , physical capital accumulation, and the household budget constraint.

In the following, we list these equations and detail how to transform the non-stationary system, which is driven by stochastic trends, into a stationary counterpart amenable to solution and estimation.

C.2.1 Equilibrium system

The equilibrium system is as follows:

$$
\begin{aligned}
V_t &= C_t - bC_{t-1} - \psi N_t^{\xi} J_t, && (40)\\
J_t &= (C_t - bC_{t-1})^{\gamma_j} J_{t-1}^{1-\gamma_j}, && (41)\\
\Gamma_t V_t^{\sigma} + \mu_t^f \gamma_f \frac{J_t}{C_t - bC_{t-1}} &= \lambda_t + b\beta E_t \left\{ \Gamma_{t+1} V_{t+1}^{-\sigma} + \mu_{t+1}^j \gamma_j \frac{J_{t+1}}{C_{t+1} - bC_t} \right\}, && (42)\\
\xi \psi \Gamma_t V_t^{-\sigma} N_t^{\xi-1} J_t &= \lambda_t w_t, && (43)\\
r_t &= \frac{\mu_t^k}{\lambda_t} \delta'(u_t), && (44)\\
\Upsilon_t \lambda_t &= \mu_t^k m_t \left\{ 1 - S\left(\frac{I_t}{I_{t-1}}\right) - S\prime\left(\frac{I_t}{I_{t-1}}\right) \frac{I_t}{I_{t-1}} \right\} + && (45)\\
&\quad + \beta E_t \mu_{t+1}^k m_{t+1} S\prime\left(\frac{I_{t+1}}{I_t}\right) \left(\frac{I_{t+1}}{I_t}\right)^2, \\
\mu_t^j &= -\psi \Gamma_t V_t^{-\sigma} N_t^{\xi} + \beta(1-\gamma_j) E_t \mu_{t+1}^j \frac{J_{t+1}}{J_t}, && (46)\\
\mu_t^k &= \beta E_t \left\{ \lambda_{t+1} r_{t+1} u_{t+1} + \mu_{t+1}^k [1 - \delta(u_{t+1})] \right\}, && (47)\\
K_{t+1} &= [1 - \delta(u_t)] K_t + m_t I_t \left[1 - S\left(\frac{I_t}{I_{t-1}}\right) \right], && (48)\\
Y_t &= z_t (\Omega_t N_t)^{\alpha_n} (u_t K_t)^{\alpha_k} (\Omega_t H)^{1-\alpha_n-\alpha_k}, && (49)\\
w_t &= \alpha_n \tau_t \frac{Y_t}{N_t}, && (50)\\
r_t &= (1-\alpha_k) \tau_t \frac{Y_t}{u_t K_t}, && (51)\\
A_t &= (1-\delta_x) X_{t-1} + Y_t, && (52)\\
X_t &= A_t - S_t, && (53)\\
\frac{\theta - 1}{\theta} &= \beta(1-\delta_x) E_t \frac{\lambda_{t+1}}{\lambda_t} \tau_{t+1}, && (54)\\
\tau_t &= \frac{\zeta}{\theta} \frac{S_t}{A_t} + \frac{\theta - 1}{\theta}, && (55)\\
G_t &= \left(1 - \frac{1}{\varepsilon_t}\right) Y_t, && (56)\\
C_t + \Gamma_t I_t + G_t &= S_t. && (57)
\end{aligned}
$$

In addition, we have laws of motion for the exogenous processes z_t, Γ_t, m_t ε_t, $g_t^{\Upsilon} = \Upsilon_t / \Upsilon_{t-1}$ and $g_t^{\Omega} = \Omega_t / \Omega_{t-1}$ as described above.

C.2.2 Stationarity and Solution Method

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

The stochastic trend components of output and capital are given by $X_t^y = \Upsilon_t^{\frac{\alpha^*-1}{\alpha^*}} \Omega_t$ and $X_t^k = \Upsilon^{\frac{-1}{\alpha^*}} \Omega_t$, respectively, where $\alpha^* = 1 - \alpha_k$. The stochastic trend components of all another non-stationary variables can be expressed as some function of X_t^y and X_t^k. In particular, define the following stationary variables as transformations of the above 18 endogenous variables: $c_t = \frac{C_t}{X_t^y}$,$v_t = \frac{V_t}{X_t^y}$, $i_t = \frac{I_t}{X_t^y}$, $g_t = \frac{G_t}{X_t^y}$, $s_t = \frac{S_t}{X_t^y}$, $y_t = \frac{Y_t}{X_t^y}$, $n_t = N_t$, $u_t = u_t$, $j_t = \frac{J_t}{X_t^y}$, $k_t = \frac{K_t}{X_{t-1}^k}$, $\mathrm{x}_t = \frac{X_t}{X_t^y}$, $a_t = \frac{A_t}{X_t^y}$, $\bar{w}_t = \frac{w_t}{X_t^y}$, $\bar{r}_t = \frac{X_t^k}{X_t^y} r_t$, $\tau_t = \tau_t$, $\bar{\mu}_t^f = \left(X_t^y\right)^\sigma \mu_t^f$, $\bar{q}_t^k = \frac{X_t^k(\mu_t^k/\lambda_t)}{X_t^y}$, and $\bar{\lambda}_t = \left(X_t^y\right)^\sigma \lambda_t$. In addition, define the two additional stationary variables, $g_t^y = \frac{X_t^y}{X_{t-1}^y}$ and $g_t^k = \frac{X_t^k}{X_{t-1}^k}$ as the growth-rates of the stochastic trends in output and capital.

The stationary system is then given by:

$$v_t = c_t - b\frac{c_{t-1}}{g_t^y} - \psi N_t^{\xi} j_t, \tag{58}$$

$$j_t = \left(c_t - b\frac{c_{t-1}}{g_t^y}\right)^{\gamma_j}\left(\frac{j_{t-1}}{g_t^y}\right)^{1-\gamma_j}, \tag{59}$$

$$\Gamma_t v_t^{\sigma} + \bar{\mu}_t^j \gamma_j \frac{j_t}{c_t - b\frac{c_{t-1}}{g_t^y}} = \bar{\lambda}_t + b\beta E_t\left(g_{t+1}^y\right)^{-\sigma}\left\{\Gamma_{t+1}v_{t+1}^{-\sigma} + \bar{\mu}_{t+1}^j\gamma_j\frac{j_{t+1}}{c_{t+1} - b\frac{c_t}{g_t^y}}\right\}, \tag{60}$$

$$k_{t+1} = \left[1-\delta(u_t)\right]\frac{k_t}{g_t^k} + m_t i_t\left[1 - S\left(\frac{i_t g_t^k}{i_{t-1}}\right)\right], \tag{61}$$

$$\xi\psi\Gamma_t v_t^{-\sigma} n_t^{\xi-1}\frac{f_t}{\bar{\lambda}_t} = \bar{w}_t, \tag{62}$$

$$\bar{r}_t = q_t^k \delta\prime(u_t), \tag{63}$$

$$\begin{aligned} 1 = {} & q_t^k m_t\left\{1 - S\left(\frac{i_t g_t^k}{i_{t-1}}\right) - S\prime\left(\frac{i_t g_t^k}{i_{t-1}}\right)\frac{i_t g_t^k}{i_{i-1}}\right\} + \\ & + \beta E_t g_{t+1}^{\Upsilon}\left(g_{t+1}^y\right)^{-\sigma}\frac{\bar{\lambda}_{t+1}}{\bar{\lambda}_t} q_{t+1}^k m_{t+1} S\prime\left(\frac{i_{t+1}g_{t+1}^k}{i_t}\right)\left(\frac{i_{t+1}g_{t+1}^k}{i_t}\right)^2, \end{aligned} \tag{64}$$

$$\bar{\mu}_t^j = -\psi\Gamma_t v_t^{-\sigma} n_t^{\xi} + \beta(1-\gamma_f)E_t\left(g_{t+1}^y\right)^{1-\sigma}\bar{\mu}_{t+1}^j\frac{j_{t+1}}{j_t}, \tag{65}$$

$$q_t^k = \beta E_t g_{t+1}^{\Upsilon}\left(g_{t+1}^y\right)^{-\sigma}\frac{\bar{\lambda}_{t+1}}{\bar{\lambda}_t}\left\{\bar{r}_{t+1}u_{t+1} + q_{t+1}^k\left[1-\delta(u_{t+1})\right]\right\}, \tag{66}$$

$$y_t = (n_t)^{\alpha}\left(u_t\frac{k_t}{g_k^k}\right)^{1-\alpha} H^{1-\alpha_n-\alpha_k}, \tag{67}$$

$$\bar{w}_t = \alpha\tau_t\frac{y_t}{n_t}, \tag{68}$$

$$\bar{r}_t = (1-\alpha)\tau_t\frac{y_t}{u_t\frac{k_t}{g_t^k}}, \tag{69}$$

$$a_t = (1-\delta_x)\frac{x_{t-1}}{g_t^y} + y_t, \tag{70}$$

$$x_t = a_t - s_t, \tag{71}$$

$$\frac{\theta-1}{\theta} = \beta(1-\delta_x)E_t\left(g_{t+1}^y\right)^{-\sigma}\frac{\bar{\lambda}_{t+1}}{\bar{\lambda}_t}\tau_{t+1}, \tag{72}$$

$$\tau_t = \frac{\zeta}{\theta}\frac{s_t}{z_t} + \frac{\theta-1}{\theta}, \tag{73}$$

$$g_t = \left(1 - \frac{1}{\varepsilon_t}\right)y_t, \tag{74}$$

$$c_t + i_t + g_t = s_t, \tag{75}$$

$$g_t^y = g_t^{\Omega}\left(g_t^{\Upsilon}\right)^{(\alpha-1)/\alpha}, \tag{76}$$

$$g_t^k = g_t^y / g_t^{\Omega}, \tag{77}$$

in addition to the exogenous processes z_t, Γ_t, m_t ε_t, g_t^{Υ} and g_t^{Ω}.

C.3 Illustrative Calibration: Baseline Model

Our choice of parameter values is guided by the existing literature, where we maintain comparability with Jaimovich and Rebelo (2009) and Schmitt-Grohe and Uribe (2011) for the aspects of the news shock mechanism and Lubik and Teo (2012) for the inventory component. In some instances, we choose values of parameters to give the Baseline model the best chance of delivering procyclical inventory. The calibration is intended for illustrative purposes only. Later we estimate the parameters using Bayesian methods, and specify prior values located well within central ranges establish in the literature.

We report the illustrative calibration in Table 2. We set the household's discount factor β to 0.9957, which is implied by the real interest rate computed from average inflation and the federal funds rate over our sample period. The elasticity of intertemporal substitution is as in Jaimovich and Rebelo (2009), $\sigma = 1$. The disutility of working parameter ξ is set to 1.1, which implies a relatively elastic Frisch elasticity of labor supply of 10 in order to give the a good chance of delivering procyclical inventory. Finally, we set γ_f, the preference parameter that determines the strength of the income effect, to 0.01 based on Schmitt-Grohe and Uribe (2012).

On the firm side, we set the elasticity parameter in the production function to $\alpha = 0.64$ as in Jaimovich and Rebelo (2009), and the degree ofof decreasing-returns-to-scale (DRS) to labor and capital in production, $1 - \alpha_n - \alpha_k$, to 0.1, following Jaimovich and Rebelo (2009) and Schmitt-Grohe and Uribe (2011). For the parameters related to physical capital, we fix steady-state physical capital depreciation at $\delta = 0.025$ and the elasticity of marginal utilization $\delta_k''(1)/\delta_k'(1) = 0.15$. There is a wide range of values for this elasticity to be found in the literature. For example, Christiano et al. (2005) find estimates of 0.01, while Schmitt-Grohe and Uribe (2012) have 0.34, and Smets and Wouters (2007) report 0.54. We choose a value of 0.15 within this range, close to the value of 0.25 used in Jaimovich and Rebelo (2009). As with the Frisch elasticity, choose the value of this to give the model a good chance of delivering procyclical inventory. Similarly, the literature also finds a wide range of values for the investment adjustment cost parameter s''. Smets and Wouters (2007) estimate it to be 5.7, Christiano et al. (2005) find 2.48, and Schmitt-Grohe and Uribe (2012) 9.1. We choose a relatively low value of $s'' = 1.3$, but as well, show robustness of the results to variation in this parameter as part of our inventory comovement analysis.

The parameters related to inventories are based on the empirical estimates in Lubik and Teo

(2012). The inventory depreciation rate δ_x is set to 0.05. The taste shifter curvature ζ is chosen as 0.67 to yield a steady-state sales-to-stock ratio of 0.55, as in Lubik and Teo (2012). The goods aggregator curvature parameter θ is set to 6.8, which results in a steady-state goods markup of 10%.

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

Table 2: Illustrative Calibration: Baseline model

Description	Parameter	Value
Subjective discount factor	β	0.9957
Household elasticity of intertemporal substitution	σ	1
Determinant of Frisch elasticity of labor supply	ξ	1.1
Habit persistence in consumption	b	0.7
Wealth elasticity parameter (GHH/KPR pref)	γ_f	0.001
Labor elasticity in production	α_n	0.64
DRS to N and K in production	$1-\alpha_n-\alpha_k$	0.1
Elasticity of capacity utilization	$\delta_k''(1)/\delta_k'(1)$	0.15
Capital depreciation	δ_k	0.025
Investment adjustment cost	s''	1.3
Inventory depreciation	δ_x	0.05
Goods aggregator curvature	θ	6.8
Taste shifter curvature	ζ	0.67
TFP growth process persistence	ρ_Ω	0.4
Steady state government spending over output	g/y	0.18
Steady state hours	n	0.2
Steady state capacity utilization	u	1
Steady state GDP growth rate (in %)	g^y	0.42545
Steady state RPI growth rate (in %)	g^{rpi}	-.58203

C.4 Conditions governing inventory comovement: Baseline Model

We examine the key equations of the supply and demand for output and develop analytical expressions to characterize the conditions governing inventory comovement. First, to gain insight into the connect between τ_t, inventory and production inputs it is helpful on the demand side to combine (10) and (11) and eliminate sales S_t, yielding

$$\left(1+\frac{1}{\chi(\tau_t)}\right)X_t = (1-\delta_x)X_{t-1} + Y_t. \tag{78}$$

We focus on our analysis on what we refer to as the "news period", which is the range of time periods defined from $t = 1$ when the news shocks is received, to the period $t + p - 1$, one period before TFP actually changes in period $t + p$. As such, during this period, there are no changes in stationary or non-stationary TFP (and of course, no changes in IST or any other shock other than the news shock). On the supply side, we focus our analysis on the "near-GHH" case with no habits in consumption, where $\frac{\partial MRS_t}{\partial C_t} = \frac{\partial MRS_t}{\partial C_{t-1}} = 0$, such that $MRS_t = MRS(N_t)$ is a function of N_t only. Imposing these restrictions on labor market equilibrium then results in

$$MRS(N_t) \approx \tau_t F_n(N_t, u_t K_t), \tag{79}$$

where the notation $F(N_t, u_t K_t)$ represents the production function over the news-boom period with no shifts in technology, $F(N_t, \widetilde{K}_t) = F(N_t, \widetilde{K}_t; H, z, \Omega)$, and where we have explicitly notated capital services $\widetilde{K}_t$ as its component $u_t K_t$. Utilization u_t is in turn defined by the capital services equilibrium condition

$$F_{\widetilde{K}}(N_t, u_t K_t) = \frac{q_t^k}{\tau_t}\delta'(u_t). \tag{80}$$

Given predetermined capital K_t, (79) and (80) imply a specific value of hours N_t and utilization u_t for a given value of the ratio $\frac{q_t^k}{\tau_t}$, which we can interpret as the relative price of new capital K_{t+1} to homogeneous output Y_t.

As is well known in the news literature based on the work of Jaimovich and Rebelo (2009), the "flow-form" of investment adjustment costs leads to a fall in the relative price of capital q_t^k in response to TFP news, thereby lowering the cost of utilization in (80), resulting in a rise in utilization. [28]. This in turn results in a rise in utilization From (79) and (80), This rise in u_t in turn leads to a rise in N_t, which we can interpret as a utilization-induced increase in labour demand in response to TFP news. Adding inventories however introduces a wedge into this equation through time variation in τ_t. When the value of output is high - such as when there is when there is an increase in demand for sales S_t upon receipt of news - the rise in τ_t both lowers the marginal cost of utilization $\frac{q_t^k}{\tau_t}$ in (80) on top of any drop in q_{kt}, and as well, increases the value of the marginal

[28] See Jaimovich and Rebelo (2009) and Christiano et al. (2007) for in-depth discussions related to this mechanism for models without inventories, and GHH for discussion of a similar margin of adjustment due to exogenous movements in q_t^k.

product in putting putting upward pressure on u_t and N_t.

We can then use (79), (80), (14) and the production function $Y_t = F(N_t, u_t K_t)$ to characterize the dynamics of Y_t and X_t for given values of q_t^k and K_t, without necessarily determining the values of q_t^k and K_t consistent with general equilibrium. To do this, we focus on the linear approximation of the de-trended equivalents of these equations about steady state.

Beginning with the demand side of output, we have the output demand curve (with sales substituted out) given by

$$(1+\frac{s}{x})\hat{x}_t = \frac{1-\delta_x}{g^y}\hat{x}_{t-1} + \frac{y}{x}\hat{y}_t - \frac{s}{x}\varepsilon_x\hat{\tau}_t, \tag{81}$$

where $\varepsilon_x = |\frac{\chi'(\tau)}{\chi(\tau)}\tau|$, and where "hats" denote percent deviations from the detrended stationary steady state. We are interested in the conditions under which inventory co-moves with output. As such, we wish to isolate the conditions under which $\hat{x}_t > 0$ for $\hat{y}_t > 0$.

Using (81), this $\hat{x}_t > 0$, $\hat{y}_t > 0$ comovement condition is then expressed as:

$$\hat{\tau}_t < \frac{1}{\varepsilon_x}\left(\frac{x}{s}\frac{(1-\delta_x)}{g^y}\hat{x}_{t-1} + \frac{y}{s}\hat{y}_t\right), \tag{82}$$

where $\hat{y}_t > 0$. Intuitively, all else equal we require a small change in the price of output τ_t relative to the change in Y_t for inventory to comove positively, consistent with our intuitive discussion earlier from the market for output.

To understand how τ_t responds to a change in production, we combine the linearized versions of (79), and the production function $F(N_t, u_t K_t)$ to get:

$$\tau_t = \left(\frac{\xi}{\alpha_n} - 1\right)\hat{y}_t - \frac{\xi}{\alpha_n}\alpha_k\left(\hat{u}_t + \hat{k}_t\right), \tag{83}$$

and then use the linearized version of (80) to replace u_t, resulting in the output supply curve,

$$\hat{\tau}_t = \frac{\left(\frac{\xi}{\alpha_n} - 1\right) - \theta_u}{1+\theta_u}\hat{y}_t - \frac{\theta_u}{1+\theta_u}\varepsilon_u\hat{k}_t + \frac{\theta_u}{1+\theta_u}\hat{q}_t^k, \tag{84}$$

where $\theta_u = \frac{\xi}{\alpha_n}\frac{\alpha_k}{1+\varepsilon_u}$.

The first term on the right-hand side describes the slope of the output supply curve, which is flatter for a higher labor supply elasticity (lower ξ), a higher elasticity of labor in production (higher α_n), or a higher value of θ_u stemming from a lower cost of utilization ε_u. The second and third terms capture the shifts in the supply of output curve due to changes in the capital stock k_t and the price of capital q_t^k respectively. The shifts from both of these factors are ultimately due to shifts in labour demand: an increase in K_t shifts the marginal product of labour directly, and a fall in q_t^k shifts it indirectly through increasing utilization by lowering its cost.

Combining (84) with the inventory $\hat{x}_t > 0$ inequality condition (82) above yields

$$\left(\frac{\left(\frac{\xi}{\alpha_n} - 1\right) - \theta_u}{1 + \theta_u} - \frac{y}{s}\frac{1}{\varepsilon_x} \right) \hat{y}_t - \frac{\theta_u}{1+\theta_u} \varepsilon_u \hat{k}_t + \theta_u \hat{q}_t^k - \frac{x}{s}\frac{1}{\varepsilon_x}\frac{(1-\delta_x)}{g^y}\hat{x}_{t-1} < 0. \tag{85}$$

where $\hat{y}_t > 0$. This inequality describes the equilibrium response consistent with $\hat{x}_t > 0$ for $\hat{y}_t > 0$ through the lens of the market for output, conditional on the general equilibrium response of $\hat{q}_k^k$, $\hat{K}_t$ and $\hat{x}_{t-1}$ (recall $\hat{y}_t > 0$). As such, the sign of the expression on the left-hand is a function of both the sign of the coefficients, as well as the sign and magnitude of the particular general equilibrium response of $\hat{y}_t$, $\hat{k}_t$, $\hat{q}_k^k$, and $\hat{x}_{t-1}$. In principle, one could drill down further into other equations of the model outside of the market for output to characterize the general equilibrium response of $\hat{k}_t$, $\hat{q}_k^k$, and $\hat{x}_{t-1}$ and then frame this expression in terms of a potentially large set of parameters across the model. Instead, we think it is more informative to focus only on the block of equations within the market for output, exploiting the dynamic structure of the model to characterize parameter conditions where possible, and reducing the analysis to separate important special cases.

C.4.1 Impact period $t = 1$

We begin our analysis by focusing on the impact period $t = 1$ when the news shock arrives. By virtue of $\hat{x}_{t-1}$ and $\hat{k}_t$ being pre-determined, $\hat{x}_{t-1}, \hat{k}_t = 0$ in period 1, and thus these two terms drop out of the condition (85). To understand the respective role played by the various elements in this condition, we proceed in three steps, each case examining a special case of this condition, beginning with the most restrictive, keeping our focus on $t = 1$ through all the steps.

1. No capacity utilization. The first step involves involves assuming very high costs of capacity utilization, approximating a model without variable capacity utilization. We can represent this case with $\varepsilon_u \to \infty$, such that $\theta_u \to 0$, reducing the condition (17) down to a pure parameter restriction of the form:

$$\frac{\xi}{\alpha_n} - 1 < \frac{y}{s}\frac{1}{\varepsilon_x}. \tag{86}$$

This condition says that for inventory to co-move with output on impact in the absence of utilization, the slope of the output supply curve, represented on the left-hand side, must be less than the absolute value of the slope of the output demand curve, represented on the right-hand side. In other words, given an outward shift in the output demand curve (due to an increase in sales), the price of output τ_t must rise less than proportionately than output y_t . How restrictive is this condition?

We can show that in steady state, $\frac{1}{\varepsilon_x} = \frac{1-\beta^*(1-\delta_x)}{1-\gamma}$, where γ can be pinned down to the data through $\gamma = (1+\frac{s}{x})$. For anything other than an unrealistically high inventory depreciation rate, $\frac{1}{\varepsilon_x}\frac{y}{s}$ is a very small number, primarily on account of the term $1-\beta^*(1-\delta_x)$. Indeed, for the calibrated case, with $\gamma = 0.55$, an inventory depreciation rate of 5%, and $\frac{y}{s} = 1.04$, $\frac{1}{\varepsilon_x}\frac{y}{s} = 0.12$. In contrast, even for highly elastic labor supply, the slope of the output supply curve will be much larger. Indeed, for $\xi = 1.2$ and $\alpha_n = 0.64$, $\frac{\xi}{\alpha_n} - 1 = 0.88$, which is not close to satisfying the positive inventory condition on impact.

2. Variable capacity utilization, zero adjustment costs to investment. In the second step we now examine to what extent variable capacity utilization on its own can loosen this condition. We now assume a smaller cost of utilization, such that capacity utilization is variable, but also assume near-zero investment adjustment costs, $s'' \approx 0$. This implies $q_t^k \approx 0$, such that the cost of utilization is not impacted by variation in the price of capital[29]. In this case, (17) reduces to

$$\frac{\left(\frac{\xi}{\alpha_n} - 1\right) - \theta_u}{1+\theta_u} < \frac{y}{s}\frac{1}{\varepsilon_x}. \tag{87}$$

As with (86), this equation again compares the slope of the output supply and demand curves. Incorporating utilization now however flattens the output supply curve by the amount through $\frac{1}{1+\theta_u}$ in the denominator and $-\theta_u$ in the numerator, increasing the range over which the other parameters can satisfy the inequality. Recalling that $\theta_u = \frac{\xi}{\alpha_n}\frac{\alpha_k}{1+\varepsilon_u}$, we note that even with a very small cost of utilization represented through $\varepsilon_u = 0.01$, using the same numbers for the parameters common to the previous step yields $\frac{\theta_u}{1+\theta_u} = 0.32$, resulting in the slope of the output slope curve being $\frac{\left(\frac{\xi}{\alpha_n}-1\right)-\theta_u}{1+\theta_u} = 0.56$, still a sufficient distance from satisfying (88).

We conclude from our analysis in the previous two steps that the respective slopes of the output supply and demand curves are unlikely on their own to allow satisfy the inventory comovement condition. Indeed, our analysis above suggests that the coefficient $\left(\frac{\left(\frac{\xi}{\alpha_n}-1\right)-\theta_u}{1+\theta_u} - \frac{y}{s}\frac{1}{\varepsilon_x}\right) > 0$ on $\hat{y}_t$ in (85) is positive for realistic parameter values.

3. Variable capacity utilization, positive adjustment costs to investment. In the third step

[29]Note that in the absence of mechanism (such as investment adjustment costs) which make the capacity utilization cost time-varying, variable capacity utilization works to effectively amplify the effect of labor in production. Indeed, as shown in Wen (1998), one can use the utilization optimality condition to substitute out utilization in production, resulting in a reduced-form production function with increased elasticity to labor, which in our framework here, shows up as a reduction in the slope of the output supply curve. Finally, note that unlike the the corresponding model without inventories where hours-worked cannot respond to news without positive investment adjustment costs which increase utilization, utilization and thus hours can vary in response to news in this model on account of time-variation in τ.

we now assume that adjustment costs are non-zero, $s'' > 0$, giving (17) in the impact period as

$$\left(\frac{\left(\frac{\xi}{\alpha_n} - 1\right) - \theta_u}{1 + \theta_u} - \frac{y}{s} \frac{1}{\varepsilon_x} \right) \hat{y}_t + \theta_u \hat{q}_t^k < 0 \quad \text{for} \quad t = 1. \tag{88}$$

Relative to (87) where the condition related to the impact of the parameters on the slopes of the supply and demand for output curves, in (88) time-variation in q_t^k shifts the output supply curve. In particular, a fall in q_t^k due to news shifts the output supply curve outwards, lowering the rise in τ_t for a given shift in the demand curve due to the increase in sales. The positive coefficient on $\hat{y}_t$ in (88) combined with $\hat{y}_t > 0$ means that only a large enough fall in q_t^k on impact could potentially satisfy the condition. We investigate this general equilibrium effect through simulation by recording the response of inventory for a range of values of s''. Figure 20 shows the results of this exercise. As is clear from the figure, changes in s'' result in different responses of capacity utilization on impact, stemming from the different effect of s'' on q_t^k, however, there is very little effect on the response of inventory on impact. Clearly, variation in q_t^k on its own is not enough to satisfy the comovement condition on its own.

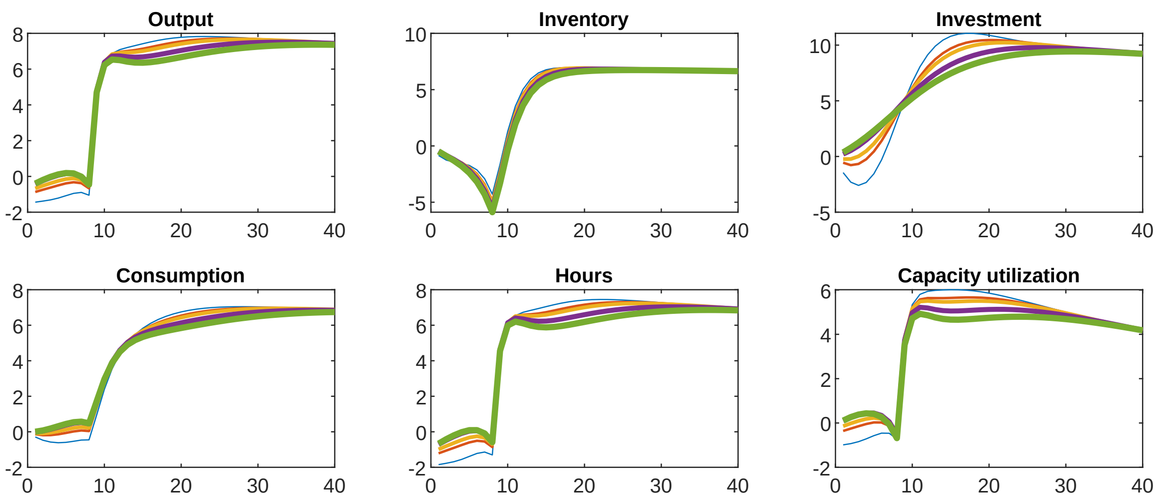

Figure 20: **IRF to 8-period out non-stationary TFP news shock: baseline model** - Sensitivity to s''. $s'' = \{0.5, 1, 1.3, 2, 5, 10\}$ (ordered from thin to thick line).

C.4.2 Periods $t = 2$ to $t + p - 1$

From period $t = 2$ to $t + p - 1$, according to (17), a rise in $\hat{k}_t$ and $\hat{x}_{t-1}$ or a fall in q_t^k can potentially shift the output supply curve to enough to loosen the condition. We make several remarks regarding these periods.

First, for x_{t-1} to help satisfy the condition requires of course that $x_{t-1} > 0$. In period 2, this requires that $x_1 > 0$, which we ruled above as unlikely, so for period $t = 2$ at least, the burden lies with k_t and q_t^k. Second, assuming a business-cycle like boom whereby output growth is positive for several periods such that $\hat{y}_{t+1} > \hat{y}_t$, the positive coefficient on $\hat{y}_t$ in (85) means that any factors that shift the output supply curve will have to increasingly shift it over time to overcome the increasing rise in $\hat{y}_t$ over time.

We again investigate this general equilibrium effect through simulation. Periods $t = 2$ to $t = 11$ in Figure 20 show the response of the model for the periods in question. As the simulation shows, the rise and k_t and fall in q_t^k are not enough to satisfy the comovement condition. Moreover, since inventories fall more and more over time, the rise in $\hat{y}_t$ is outpacing the response of these other factors.

In summary, our analysis for the baseline model concludes that the respective slopes of the output supply and demand curves are unlikely on their own to allow satisfy the inventory comovement condition in any of the periods in the news-period. Instead, the analysis points to the endogenous response of factors that will shift either of these curves on impact and in subsequent periods. In the context of this baseline model, in the impact period, only one factor offers this possibility: investment adjustment costs, yet our simulations suggest that variation in q_{t^k} on its own is unable to satisfy the comovement condition. In subsequent periods, $\hat{k}_t$, $\hat{x}_{t-1}$ and q_t^k offer the potential to shift the output supply curve, however, our simulations suggest that these factors are not enough, and that their combined effect is outpaced by the increasing rise in $\hat{y}_t$ over time.

D Additional Model Details: Sticky Wage and Price Model

This appendix section details elements of the *Sticky Wage and Price Model* not shown in the main text.

D.1 Model Description: Sticky Wage and Price Model

We introduce sticky prices by following Lubik and Teo (2012), whereby we assume that distributors face convex adjustments costs in setting prices. The sticky-wage component follows the decentralization of Schmitt-Grohe and Uribe (2012) and Smets and Wouters (2007). We add a continuum of monopolistically competitive labor unions, indexed by $j \in [0, 1]$, and a competitive employment agency to the baseline setting. Monopolistic unions buy homogeneous labor from households, transform it into differentiated labor inputs, and sell them to the employment agency, which aggregates the differentiated labor into a composite and sells it to the intermediate goods producer. The unions face Calvo-type frictions in setting wages for each labor type and re-set their wage according to an indexation rule when unable to reoptimize. Since this particular decentralization of wage stickiness implies that consumption and hours are identical across households, we can continue to refer to a stand-in representative household as with the baseline model. Finally, we close the model with a standard monetary policy nominal interest rate rule.

D.1.1 Employment unions and employment agency

Our sticky-wage framework follows the decentralization of Schmitt-Grohe and Uribe (2012) and Smets and Wouters (2007). To our baseline model, we add a continuum of monopolistically competitive labor unions indexed by $j \in [0,1]$, and a competitive employment agency. 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. The unions face Calvo frictions in setting their wages for each labour type, and re-set their wage according to an indexation rule when unable to reoptimize. Since this particular decentralization of wage stickiness implies that consumption and hours are identical across households, we can continue to refer to a stand-in representative household as with the baseline model.

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 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 $W_{jt} = W_{jt-1}\pi_{t-1}^{\iota_w}\pi^{1-\iota_w}, \quad 0 \leq \iota_w \leq 1$, where $\pi_t = P_t/P_{t-1}$ and π is its steady state, and where $0 \leq \iota_w \leq 1$. A union that can re-optimize its wage in period t chooses its wage W_{jt}^* to maximize

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

subject to the demand curve for N_{jt}.

The employment agency acquires each jth 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 jth 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

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

The sticky wage framework results in a time-varying markup μ_t^w between the wage W_t paid by the intermediate goods firm and the wage W_t^h paid to the household, such that

$$\mu_t^w = \frac{w_t}{w_t^h}, \tag{89}$$

where $w_t = \frac{W_t^h}{P_t}$ and $w_t = \frac{W_t^h}{P_t}$. The dynamics of μ_t^w is captured by a resulting equilibrium wage Phillips curve derived from imposing equilibirum on the combination of the employment agency and union's problem.

D.1.2 Distributors

Distributors now face frictions in setting their prices, and as in Lubik and Teo (2012) , we assume that the ith 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 ith distributor then faces the problem of choosing P_{it}, S_{it}, Y_{it} and A_{it} to maximize

$$E_t \sum_{k=0}^{\infty} \beta^k \frac{\lambda_{t+k}}{\lambda_t} \left\{ \frac{P_{it+k}}{P_{t+k}} S_{it+k} - \tau_t Y_{t+k}(j) - \frac{\kappa}{2}\left[\frac{P_{it+k}}{\pi_{t-1}^{\iota_p}\pi^{1-\iota_p}P_{it+k-1}} - 1\right]^2 S_t \right\}, \tag{90}$$

subject to the same constraints as in the baseline model. The distributor's Y_{it}, X_{it} and A_{it} first-order conditions are the same as in the baseline model, but now the P_{it} condition is given by

$$(1-\theta)\frac{S_{it}}{P_t} - \kappa\left[\frac{P_{it}}{\pi_{t-1}^{\iota_p}\pi^{1-\iota_p}P_{it-1}} - 1\right]\frac{S_t}{\pi_{t-1}^{\iota_p}\pi^{1-\iota_p}P_{it-1}} + \beta E_t \frac{\lambda_{t+1}}{\lambda_t}\kappa\left[\frac{P_{it+1}}{\pi_t^{\iota_p}\pi^{1-\iota_p}P_{it}} - 1\right]\frac{P_{it+1}S_{t+1}}{\pi_t^{\iota_p}\pi^{1-\iota_p}P_{it}^2} + \mu_t^x\theta\frac{S_{it}}{P_{it}} = 0. \tag{91}$$

This equation 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 describes the adjustment of the first markup through inventories; the optimal pricing condition describes the adjustment of the second markup through price-setting.

Unlike in the flexible price baseline model where the markup between the marginal cost of sales and price is constant, under sticky prices, the Distributor's pricing condition implies that this markup is time-varying. This in turn means that the value of forgone inventory, μ_t^x, which we previously interpreted as the marginal cost of sales, is no longer constant. As such, this introduces μ_t^x as a time-varying wedge into the firm's optimal stocking equation,

$$\tau_t = \zeta p_{it} \frac{S_{it}}{A_{it}} + \mu_t^x \left(1 - \zeta \frac{S_{it}}{A_{it}}\right). \tag{92}$$

Imposing equilibrium, and solving for $\chi_t = \frac{X_t}{S_t}$ yields

$$\chi_t = \zeta \frac{1-\mu_t^x}{\tau_t - \mu_t^x} - 1 = \chi(\tau_t, \mu_t^x) \tag{93}$$

where $\chi_\tau(t) = \frac{\partial \chi(\tau_t, \mu_t^x)}{\partial \tau_t} < 0$ and $\chi_{\mu^x}(t) = \frac{\partial \chi(\tau_t, \mu_t^x)}{\partial \mu_t^x} < 0$, and where as in the baseline model, μ_t^x is equal to the expected discounted value of future marginal costs, $\mu_t^x = (1-\delta_x)\beta E_t \frac{\lambda_{t+1}}{\lambda_t} \tau_{t+1}$. The derivative $\chi_{\mu^x}(t)$ represents an intertemporal substitution effect on the inventory decision: all else equal, if marginal costs are expected to be lower in the future relative to the present, it is optimal to defer inventory accumulation to the future and run down inventory levels today. Thus compared to the baseline model where we identified a demand channel and cost channel to the inventory decision, we can now think about their being both a current and expected future cost channel in addition to the demand channel.

Adding sticky prices as a result adds an additional term to our comovement condition, now given by

$$\left(\frac{\left(\frac{\xi}{\alpha_n} - 1\right) - \theta_u}{1+\theta_u} - \frac{y}{s}\frac{1}{\varepsilon_x} \right) \hat{y}_t - \frac{\theta_u}{1+\theta_u} \varepsilon_u \hat{k}_t + \theta_u \hat{q}_t^k - \frac{x}{s}\frac{1}{\varepsilon_x}\frac{(1-\delta_x)}{g^y} x_{t-1} - \frac{1+\frac{\xi}{\alpha_n}}{1+\theta_u}\hat{\phi}_t^e + \frac{\theta_u}{1+\theta_u}\hat{\phi}_t^l - \mu^x \hat{\mu}_t^x < 0, \tag{94}$$

such that if all else equal discounted expected future marginal costs are expected are low relative to today (such as due to the effect of a future expected increase in TFP), Distributors have an incentive to run down inventories in the present, making the comovement condition more difficult to satisfy.[30]

D.1.3 Monetary Policy Rule

We close the model with a standard monetary policy rule where the interest rate, R_t^n, is set by the monetary authority according to a feedback rule,

$$\frac{R_t^n}{R^n} = \left(\frac{R_{t-1}^n}{R^n}\right)^{\rho_r} \left(\left(\frac{\pi_t}{\pi}\right)^{\phi_\pi} \left(\frac{Y_t}{Y_t^*}\right)^{\phi_y}\right)^{(1-\rho_r)} e^{\eta_t}, \tag{95}$$

where η_t is a monetary policy shock, and Y_t^* is level of output that would preside under flexible prices and without wage or price markup shocks.

D.1.4 Stochastic Exogenous Processes: Sticky Wage and Price Model

Relative to the baseline model, there are three additional stochastic processes in the sticky wage and price model: a wage markup shock (ν_t^w), a price markup shock (ν_t^p) and a monetary policy shock η_t. The stochastic processes are thus given by

$$\ln\left(\frac{\vartheta_t}{\vartheta}\right) = \rho_\vartheta \ln\left(\frac{\vartheta_{t-1}}{\vartheta}\right) + e_{\vartheta,t}, \tag{96}$$

for $\vartheta = \{z, g^\Omega, g^\Upsilon m, \Gamma, \varepsilon\, \nu_t^w, \nu_t^p, \eta$. The innovations are defined as

$$e_{\vartheta,t} = \begin{cases} e^0_{\vartheta,t} + e^4_{\vartheta,t-4} + e^8_{\vartheta,t-8} + e^{12}_{\vartheta,t-12}, & \vartheta = \{z, g_\Omega\} \\ e^0_{\vartheta,t}, & \vartheta = \left\{m, g^{\Upsilon,}, \varepsilon_t, \Gamma_t, \nu_t^w, \nu_t^p, \eta\right\}. \end{cases}$$

D.2 Model equilibrium, stationary and solution method: Sticky Wage and Price Model

In addition to the symmetric equilibrium defined in the baseline model, $W_{jt}^* = W_t^*$, $N_{jt} = N_t^*$ $\forall j$. It then follows that $N_t^h = \int_0^1 n_t^* dj = N_t^*$.

In additional to the equilibrium definition for the baseline model, the sticky wage and price model results in an additional set of stochastic processes $\{\mu_t^w, \mu_t^x, w_t^h, R_t^n, \pi_t\}_t^\infty$.

[30] We emphasize that this additional $\hat{\mu}_t^x$ term in (22) is due to sticky prices, not sticky wages. In a version of the model with sticky wages but flexible prices, the distributor's pricing condition implies that the markup between marginal cost of sales and price is constant, as in the baseline model and thus the additional $\hat{\mu}_t^x$ term would drop out of (22).

The equilibrium system for the sticky wage and price model is the same as that of the baseline model, with the addition of the Distributor's pricing condition (91), the monetary policy rule (95), the wage markup definition (89), and the standard wage-setting and aggregate wage equation resulting from the sticky wage framework. Additionary, w_t^h replaces w_t in the household labor first-order condition (43). Stationarity proceeds as with the baseline model, where the nominal interest rate, inflation rate wage markup are stationary.

D.3 Illustrative Calibration: Stick Wage and Price Model

The illustrative calibration for the sticky wage and price model uses the same calibration as that of the Baseline model, with the addition of the parameters related to the nominal side of the economy, where we choose values consistent with the literature, including those from Lubik and Teo (2012) related to sticky pricing under inventory. Table 3 details these parameter choices.

Table 3: Illustrative Calibration: Sticky Wage and Price model - Additional parameters

Description	Parameter	Value
Taylor rule smoothing	ρ_{rn}	0.5
Taylor rule inflation	ϕ_π	1.5
Taylor rule output	$\phi_p i$	0.05
Price adjustment costs	κ	250
Calvo wage parameter	ζ_w	0.8
Price indexation	ι_p	0.5
Wage indexation	ι_w	0.5
Steady state wage markup	λ_w	1.1

E Additional Model Details: Learning-by-doing Model

This appendix section details elements of the *Learning-by-doing Model* not shown in the main text.

E.1 The labor demand wedge and stock prices

We can gain more insight into the labor demand wedge by manipulating (26) to give:

$$\phi_t^{ld} = 1 - q_t^h(1-\gamma_h)\left(\frac{H_{t+1}-(1-\delta_h)H_t}{w_t N_t}\right) = 1-(1-\gamma_h)\frac{q_t^h H_{t+1}}{w_t N_t}\left(1-(1-\delta_h)\frac{H_t}{H_{t+1}}\right). \quad (97)$$

Additionally, using the fact that the stock-price value the firm SP_t is given by:

$$SP_t = q_t^h H_{t+1}, \tag{98}$$

we can write (97) as:

$$\phi_t^{ld} = 1 - \frac{SP_t}{w_t N_t} \psi_t^h, \tag{99}$$

where $\psi_t^h = (1-\gamma_h)\left(1-(1-\delta_h)\frac{H_t}{H_{t+1}}\right)$. The labor demand wedge is a function of the ratio of stock prices over the wage bill. Indeed, under the log-linear case of Chang et al. (2002a) for $\delta_h = 1$, $\psi_t^h = 1-\gamma_h$, and

$$\phi_t^{ld} = 1 - \frac{SP_t}{w_t N_t}(1-\gamma_h). \tag{100}$$

The term $\frac{SP_t}{w_t N_t}$ acts like a type of "labor Tobin's Q". When the value of the firm is high relative to the cost of labour, the firm lowers its markup in order to increase labor and acquire more knowledge. Under the more general case for $0 < \delta_h < 1$, the same is true, except that the term $1-(1-\delta_h)\frac{H_t}{H_{t+1}}$ scales this effect, reinforcing the above when knowledge growth is expected to be high.

E.2 The importance of internalization

The above learning-by-doing model results in a labor demand wedge ϕ_t^{ld} that impacts the markup on impact, and a slower-moving efficiency wedge ϕ_t^e that doesn't move on impact, but gradually impacts the marginal cost of production. Importantly, the labor demand wedge ϕ_t^{ld} stems from our assumption that the firm internalizes the impact of its use of hours on knowledge in production. To see this, we can consider an alternative set-up that involves external-effects learning-by-doing only, whereby the firm acquires knowledge by the joint-action of other firms through the impact of the average level of labor $\bar{N}_t$ presiding in the economy. The production function and knowledge-accumulation equation under such an alternative scenario would then be given by:

$$Y_t = z_t \left(\Omega_t N_t\right)^{\alpha_n} \widetilde{K}_t^{\alpha_k} \left(\Omega_t \bar{H}_t\right)^{1-\alpha_n-\alpha_k}, \tag{101}$$

and

$$\bar{H}_{t+1} = (1-\delta_h)\bar{H}_t + \bar{H}_t^{\gamma_h} \bar{N}_t^{1-\gamma_h}, \tag{102}$$

where $\bar{N}_t$ and $\bar{H}_t$ are the economy-wide average levels of labor and knowledge respectively. Since the effect of learning-by-doing is now external to the firm however, the firm's problem is now essentially the same as in the baseline model, such that the firm chooses N_t and $\widetilde{K}_t$ to maximize

$\Pi_t^Y = \tau_t Y_t - w_t N_t - r_t \widetilde{K}_t$ subject to the production function, resulting in the standard demand functions for labor, $w_t = \alpha_n \tau_t \frac{Y_t}{N_t}$. Only the production technology changes. As such, in the context of our wedges framework in the labor market, the external effects model corresponds to $\phi_t^{ls} = 1$, $\phi_t^{ld} = 1$ and $\phi_t^e = \bar{H}_t^{1-\alpha_n-\alpha_k}$. In contrast to the learning-by-doing model, the external effects learning-by-doing model results only a time-varying efficiency wedge ϕ_t^e. The labor demand wedge ϕ_t^{ld} and its associated markup are constant.

E.2.1 Stochastic Exogenous Processes

The stochastic process in the learning by doing model are the same as in the baseline model.

E.3 Model equilibrium, stationary and solution method: Learning-by-Doing Model

In additional to the equilibrium definition for the baseline model, the sticky wage and price model results in an additional set of stochastic processes $\{h_t, q_t^h\}_t^\infty$.

The equilibrium system for the learning-by-doing model is the same as that of the baseline model, with the following additions:

$$H_{t+1} = (1-\delta_h)H_t + H_t^{\gamma_h} N_t^{1-\gamma_h}, \qquad \text{where} \qquad 0 \leq \delta_h \leq 1, \quad 0 \leq \gamma_h < 1, \quad \nu_h > 0. \tag{103}$$

and

$$q_t^h = \beta E_t \frac{\lambda_t}{\lambda_{t+1}} \left\{ (1-\alpha_n-\alpha_h)\tau_{t+1} \frac{Y_{t+1}}{H_{t+1}} + q_{t+1}^h \left(1 - \delta_h + \gamma_h \frac{H_{t+1}^{\gamma_h} N_{t+1}^{1-\gamma_h}}{H_t} \right) \right\}. \tag{104}$$

As well,

$$Y_t = z_t \left(\Omega_t N_t\right)^{\alpha_n} \widetilde{K}_t^{\alpha_k} \left(\Omega_t H_t\right)^{1-\alpha_n-\alpha_k},$$

replaces the baseline model production function (49), and

$$w_t = \tau_t \alpha \frac{Y_t}{N_t} + q_t^h (1-\gamma_h) \frac{H_t^{\gamma_h} N_t^{1-\gamma_h}}{N_t}, \tag{105}$$

replaces the intermediate goods firm's labour first order condition (50) in the baseline model.

Stationarity proceeds as with the baseline model, where now we define $\hat{q}_t^h = \frac{q_t^h}{X_t^y}$. As described in the main text, H_t is already stationary.

E.4 Illustrative Calibration: Learning-by-doing Model

The illustrative calibration for the Learning-by-doing Model uses the same calibration as that of the Baseline model, with the addition of the parametes related to learning-by-doing. There are two parameters related to learning-by-doing in the the model: the exponent on labor in knowledge capital accumulation, ν, and, the depreciation of knowledge capital, δ_h. We choose a prior of 0.3 for ν, consistent with values in the literature such as Gunn and Johri (2011), Cooper and Johri (2002) and Chang et al. (2002b). There is little guidance in the literature for the depreciation parameter δ_h, other than the implicit assumption of 100% depreciation with log-linear specifications of the knowledge capital accumulation equation in the specification of Chang et al. (2002b) and others. We choose a value of 0.2 for δ_h, reflecting the assumption of a higher depreciation rate of knowledge relative to physical capital, as discussed in the literature on learning-by-doing.

F Additional Details: Bayesian Estimation

This appendix section details elements of the Bayesian Estimation not shown in the main text. We estimate a "full" version of the model which is the combination of the learning-by-doing model and the sticky wage and price model (in other words, the baseline model with sticky wage and prices and learning-by-doing). In addition to the shocks 9 key shock processes, we also include an iid measurement error on the resource constraint. We calibrate a subsection of the parameters and estimate the remaining parameters as described in section 4. The calibrated parameters are summarized in Table 4. These choice and values of the calibrated parameters are standard, consistent with our illustrative calibration, and in general, not key parameters for the inventory comovement capabilities of the model.

We report prior distributions and posterior estimates in Table 5. Prior distributions conform to assumptions in Schmitt-Grohe and Uribe (2012) and Smets and Wouters (2007). However, we draw attention to a few key parameters. First, unlike the illustrative calibrations where we pushed some key parameters values to a range that would give the baseline model the best chance of delivering procyclical inventory, in our prior choice we remain agnostic to this and specify prior values located well within central ranges establish in the DSGE literature not concerned with inventory. In particular: (i) we specify a prior mean of 3 for the disutility of working parameter ξ, implying a Frisch elasticity of labor supply of 0.5 (compared to 10 in the illustrative calibration); (ii) we specify a prior mean of 0.5 for $\delta_k''(1)/\delta_k'(1)$, the elasticity of capital depreciation (compared to 0.15 in the illustrative calibration); (iii) we specify a prior mean of 4 for the s'', the investment adjustment cost parameter (compared to 1.3 in the illustrative calibration). Second, following Schmitt-Grohe and Uribe (2011), we assign a uniform prior over the GHH/KPR preference parameter γ_j over the interval $(0,1)$ to keep it largely uninformative as to the importance of TFP news in the posterior estimations, given the importance of this parameter to the comovement capabilities of consumption, invest and hours-worked in reponse to TFP news. Third, there are two parameters related to learning-by-doing in the the model: the exponent on labor in knowledge capital accumulation, ν, and the depreciation of knowledge capital, δ_h. We choose a prior of 0.3 for ν, consistent with values in the literature such as Gunn and Johri (2011), Cooper and Johri (2002) and Chang et al. (2002b). There is little guidance in the literature for the depreciation parameter δ_h, other than the implicit assumption of 100% depreciation with log-linear specifications of the knowledge capital accumulation equation in the specification of Chang et al. (2002b) and others. We choose a prior of 0.5 for δ_h, approximately midpoint between the 100% depreciation rate case implied by the log-linear specification and a rate more in line with physical capital depreciation (0.025).[31].

Posterior estimates of parameters common with these studies are broadly in line with the literature on medium-scale New Keynesian models.Consequently, we do not discuss them in detail beyond the discussion in section 4. Again, we draw attention to several key parameters. First, despite a choice of prior values that remains relative agnostic to inventory considerations, the resulting posterior means of the parameter values most critical to inventory comovement, end up being favorable for inventory comovement and relatively close to the illustrative calibration. These parameters include the disutility of working parameter ξ, the elasticity of capital depreciation parameter $\delta_k''(1)/\delta_k'(1)$, and the GHH/KPR preference parameter γ_j. Inventory procyclically aside, a very low posterior mean of this later parameter γ_j is suggestive of the important of TFP news overall.

F.1 Additional Results

In this section, we provide some additional evidence on impulse response functions from the estimated model.

We first investigate the correspondence of empirical and model-based responses to a TFP surprise shock. Our baseline Max Share identification identifies a TFP news and a TFP surprise shock.

[31] Overall, the results are relatively robust to alternatively specifying much lower or or higher priors on δ_h. Compared to a log-linear specification of knowledge capital accumulation, our linear specification (which nests the log-linear specification) proved to be much more stable under estimation.

Table 4: Calibrated Parameters

Parameter	**Description**	**Value**
β	Household subjective discount factor	0.996
σ_L	Intertemporal elasticity of substitution	1
N_{ss}	Steady state hours-worked	0.2
δ	Capital depreciation rate	0.025
α	Elasticity of labor in production	0.64
$\frac{G}{Y}$	Steady state government spending-GDP ratio	0.18
λ_w	Steady state wage markup	1.1
u	Steady state capital utilization rate	1
θ	Goods aggregator curvature	6.8
g^y	Steady state output growth rate	1.00425
g^{Υ}	Steady state growth rate of relative price of investment	0.9942

Figure 21 shows responses to a TFP surprise shock. It is evident from this figure that the TFP surprise shock triggers a broad based expansion in, amongst others, output, investment, consumption and inventories. Figure 22 shows the corresponding impulse responses to an unanticipated stationary TFP shock from the estimated DSGE model. The model generates, broadly consistent with the VAR-based responses, an expansion in the macroeconomic aggregates and some typical, hump-shaped patterns. One caveat to this specific comparison is that the VAR model does not separately identify a stationary and non-stationary TFP shock, whereas this distinction is enforced in the DSGE model. The VAR response to the surprise shock therefore likely conflates the two types of shocks. Since the responses appear to be stationary, we therefore contrast them to the stationary TPF shock responses in the DSGE model.

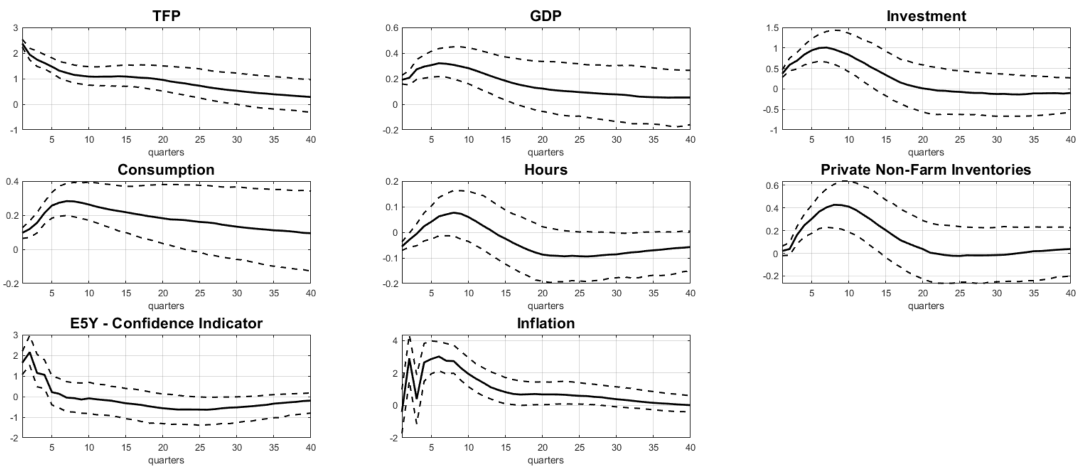

Figure 21: **IRF to TFP surprise shock.** The solid line is the median and the dashed lines are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations

Table 5: Estimated Parameters

Parameter	Description	Prior Mean	Posterior Distribution Mean	10%	90%	Prior Distrib.	Prior Std.dev.
γ_j	GHH/KPR pref	0.5	0.002	0.002	0.0021	unif	0.2878
b	Consumption habits	0.7	0.8682	0.8123	0.9268	beta	0.1
ξ	Determinant of Frisch elasticity	3	1.0732	1.0593	1.0865	gamm	1
ν	Labor in knowledge capital	0.3	0.0547	0.0372	0.072	beta	0.1
δ_h	Knowledge capital depreciation	0.5	0.5053	0.3559	0.6492	beta	0.1
ϕ_k''	Investment adjustment cost	4	12.8448	7.967	17.65	gamm	1.5
$\delta_k''(1)/\delta_k'(1)$	Elasticity of capacity utilization	0.5	0.189	0.0729	0.3111	gamm	0.25
δ_x	Inventory depreciation	0.05	0.0515	0.0488	0.0543	beta	0.0025
ζ	Taste shifter curvature	0.67	0.6719	0.6552	0.6879	gamm	0.01
ρ_{rn}	Taylor rule smoothing	0.5	0.5905	0.5538	0.6283	beta	0.025
ϕ_π	Taylor rule inflation	1.5	1.2112	1.1017	1.3144	gamm	0.25
ϕ_y	Taylor rule output	0.05	0.0252	0.0172	0.0332	gamm	0.01
κ	Price-adjustment costs	250	237.1511	195.3579	277.4332	norm	25
ζ_w	Calvo wage parameter	0.75	0.7912	0.7413	0.8429	beta	0.05
ι_p	Price indexation	0.5	0.5283	0.3859	0.6689	beta	0.1
ι_w	Wage indexation	0.5	0.2957	0.1674	0.4214	beta	0.1
Parameters relating to stochastic processes:							
ρ_z	Stationary TFP shock persistence	0.5	0.8355	0.6398	0.9846	beta	0.2
ρ_Γ	Preference shock persistence	0.5	0.4649	0.2742	0.6503	beta	0.2
ρ_m	MEI shock persistence	0.5	0.8914	0.8654	0.9184	beta	0.2
ρ_{ε_g}	Gov't spending shock persistence	0.5	0.9929	0.9889	0.9971	beta	0.2
ρ_{g^Ω}	Non-stationary TFP shock persistence	0.2	0.3574	0.2934	0.4248	beta	0.05
ρ_{g^Υ}	Non-stationary IST shock persistence	0.2	0.25	0.1541	0.3479	beta	0.05
ρ_η	Monetary policy shock persistence	0.5	0.791	0.7468	0.8352	beta	0.2
ρ_{ε_p}	Price markup shock persistence	0.5	0.1141	0.013	0.2153	beta	0.2
ρ_{ε_w}	Wage markup shock persistence	0.5	0.465	0.3245	0.6066	beta	0.2
σ_{e_z}	Stationary TFP shock SD	0.5	0.1755	0.1119	0.2373	invg	1
$\sigma_{e_z^4}$	Stationary TFP shock (4p news) SD	0.289	0.1195	0.0692	0.1689	invg	1
$\sigma_{e_z^8}$	Stationary TFP shock (8p news) SD	0.289	0.1184	0.0687	0.1665	invg	1
$\sigma_{e_z^{12}}$	Stationary TFP shock (12p news) SD	0.289	0.1168	0.0696	0.1634	invg	1
σ_{e_Γ}	Preference shock SD	0.5	13.4991	7.4055	20.0224	invg	1
σ_{e_m}	MEI shock SD	0.5	5.9187	3.9918	7.7519	invg	1
$\sigma_{e_{\varepsilon_g}}$	Gov't spending shock SD	0.5	2.007	1.5797	2.4213	invg	1
$\sigma_{e_{g^\Omega}}$	Non-stationary TFP growth shock SD	0.5	0.3982	0.2722	0.5199	invg	1
$\sigma_{e^4_{g^\Omega}}$	Non-stationary TFP shock (4p news) SD	0.289	0.1711	0.0822	0.2585	invg	1
$\sigma_{e^8_{g^\Omega}}$	Non-stationary TFP shock (8p news) SD	0.289	0.1957	0.0921	0.294	invg	1
$\sigma_{e^{12}_{g^\Omega}}$	Non-stationary TFP shock (12p news) SD	0.289	0.3331	0.2329	0.4366	invg	1
$\sigma_{e_{g^\Upsilon}}$	Non-stationary ISTC growth shock SD	0.5	0.9072	0.4468	1.3653	invg	1
σ_{e_η}	Monetary policy shock SD	0.5	0.1472	0.1292	0.1652	invg	1
$\sigma_{e_{\varepsilon_p}}$	Price markup shock SD	0.5	0.1627	0.1378	0.1878	invg	1
$\sigma_{e_{\varepsilon_w}}$	Wage markup shock SD	0.5	0.3683	0.2996	0.4352	invg	1
$\sigma_{e_{\mathrm{msrt}}}$	Measure error SD	0.251	0.473	0.4396	0.5	unif	0.144

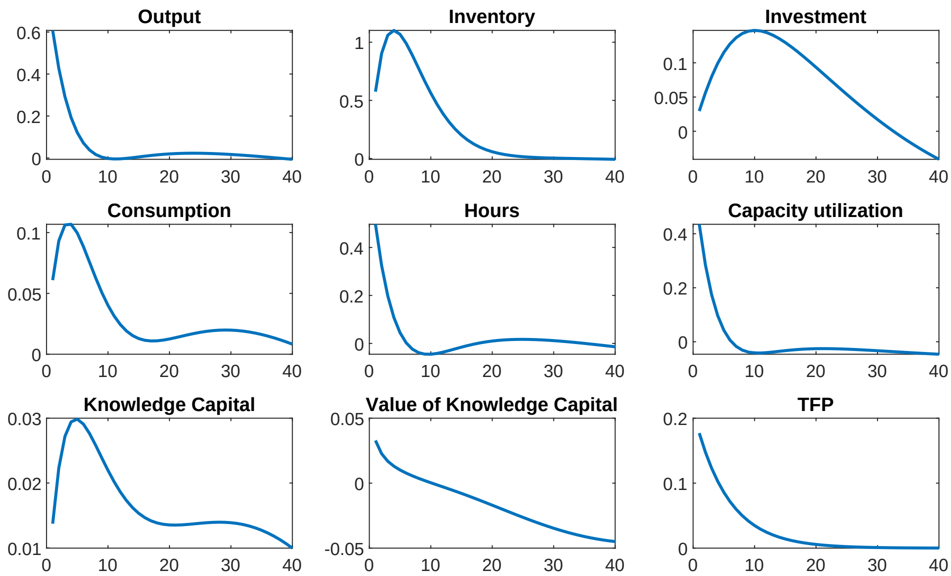

Figure 22: **IRF to an unanticipated stationary TFP shock: Estimated model (*Learning-by-doing + sticky wages and prices*).**

References in the Appendix

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

Cascaldi-Garcia, D. and Vukotic, M. (2020). Patent-based news shocks. *Review of Economics and Statistics*, forthcoming.

Chang, Y., Gomes, J. F., and Schorfheide, F. (2002). Learning-by-Doing as a Propagation Mechanism. *American Economic Review*, 92(5):1498–1520.

Christiano, L., Ilut, C., Motto, R., and Rostagno, M. (2007). Monetary policy and stock market boom-bust cycles. *manuscript, Northwestern University*.

Christiano, L. J., Eichenbaum, M., and Evans, C. L. (2005). Nominal rigidities and the dynamic effects of a shock to monetary policy. *Journal of Political Economy*, 113(1):1–45.

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

Francis, N., Owyang, M., Roush, J., and DiCecio, R. (2014). A flexible finite-horizon alternative to long-run restrictions with an application to technology shocks. *Review of Economics and Statistics*, 96:638–647.

Görtz, C., Tsoukalas, J., and Zanetti, F. (2021). News Shocks under Financial Frictions. *American Economic Journal: Macroeconomics*, forthcoming.

Greenwood, J., Hercowitz, Z., and Huffman, G. (1988). Investment, capacity utilization, and the real business cycle. *The American Economic Review*, 78:402–217.

Jaimovich, N. and Rebelo, S. (2009). Can news about the future drive the business cycle? *American Economic Review*, 99(4):1097–1118.

Kurmann, A. and Sims, E. (2019). Revisions in utilization-adjusted tfp and robust identification of news shocks. *Review of Economics and Statistics*, forthcoming.

Lubik, T. A. and Teo, W. L. (2012). Inventories, inflation dynamics and the new keynesian phillips curve. *European Economic Review*, 56(3):327–346.

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

Schmitt-Grohe, S. and Uribe, M. (2011). Business cycles with a common trend in neutral and investment-specific productivity. *Review of Economic Dynamics*, 14:122–135.

Schmitt-Grohe, S. and Uribe, M. (2012). What's news in business cycles? *Econometrica*, 80(6):2733–2764.

Smets, F. and Wouters, R. (2007). Shocks and frictions in US business cycles: A Bayesian DSGE approach. *American Economic Review*, 97(3):586–606.