

Is there news in inventories?

Gortz, Christoph; Gunn, Christopher; Lubik, Thomas

DOI:

[10.1016/j.jmoneco.2022.01.005](https://doi.org/10.1016/j.jmoneco.2022.01.005)

License:

Creative Commons: Attribution-NonCommercial-NoDerivs (CC BY-NC-ND)

Document Version

Peer reviewed version

Citation for published version (Harvard):

Gortz, C, Gunn, C & Lubik, T 2022, 'Is there news in inventories?', *Journal of Monetary Economics*, vol. 126, pp. 87-104. <https://doi.org/10.1016/j.jmoneco.2022.01.005>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Is There News in Inventories?*

Christoph Görtz
University of Birmingham[†]

Christopher Gunn
Carleton University[‡]

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

January 2022

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 mechanisms: sticky wages and the presence of knowledge capital accumulated through learning-by-doing. Only the latter moves the wedges to qualitatively match the empirical behaviour. 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 want to thank the editor Yuriy Gorodnichenko, the associate editor Eric Sims, and an anonymous referee for guidance and comments that improved the paper considerably. 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. Importantly, the response of inventories in our baseline model is consistent with and informative for the response of marginal cost.

The core of our full model is the framework of Jaimovich and Rebelo (2009), which is closely

¹We 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.

related to Schmitt-Grohe and Uribe (2012). It includes the trio of particular specifications of preferences, investment adjustment costs and variable capital utilization, which are features generally recognized in the news literature as needed for generating comovement of macroeconomic aggregates in response to a TFP news shock. We extend this model to include finished goods inventories based on the stock-elastic demand model of Bils and Kahn (2000). 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. (2002), 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

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

theoretical frameworks (see Ramey and West (1999), for a comprehensive survey and critical assessment). In our theoretical modeling of inventories, we are guided by Bills 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. Section 4 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

³This 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

⁴In 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.

⁵This 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.

⁶See, 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.

⁷We 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

⁸We 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)).

⁹The 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

¹⁰A 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).

¹¹Appendix 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.¹³ 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 we construct a time series for value weighted patents up to 2018Q2, which extends the sample used in Cascaldi-Garcia and Vukotic (2020). Due to data limitations at the time they conducted their study, they 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

¹²Kogan 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.

¹³The two identification schemes result in very similar shock series. When we identify a news shock from a VAR that corresponds to the one of Figure 3 either with our baseline max-share identification or with the one proposed by Cascaldi-Garcia and Vukotic (2020), the correlation between the two shock series is 0.985.

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 desired inventory-sales ratio, which counters this effect. It is along these margins that the success of a theoretical model to replicate the empirical findings rests.¹⁴

Jaimovich and Rebelo (2009) document the 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 4.¹⁵

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

¹⁴Görtz et al. (2019) construct aggregate measures of debt and equity cost of capital and implied cost-of-capital measures from firm-level data. In response to a TFP news shock, all measures decline significantly prior to the realization of higher TFP. We also study the response of various measures of marginal cost to a TFP news shock. However, none of these measures shows a decline in marginal costs that would point to a strong incentive to run down current inventories and build up stocks again once the higher productivity is realized. Overall, we find evidence against a strong negative substitution effect, but support for a strong positive demand effect. This finding serves further to motivate a demand-enhancing motive for holding more inventories in line with Bilal and Kahn (2000).

¹⁵The inventory-to-sales ratio is the ratio of private non-farm inventories and final sales of domestic business as in Lubik and Teo (2012). Utilization is provided by Fernald (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.

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, commensurate 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 Bilal 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 that we detail in the online

appendix.

3.1.1 Intermediate goods firm

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

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

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 \tilde{K}_t as in Jaimovich and Rebelo (2009) and Schmitt-Grohe and Uribe (2012).¹⁶ 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 \tilde{K}_t at rental rate r_t from the capital services market. It then sells its output Y_t at real price τ_t to the distributors. The firm's profit maximization problem results in standard demand functions for labor and capital services, respectively: $w_t = \alpha_n \tau_t \frac{Y_t}{N_t}$ and $r_t = \alpha_k \tau_t \frac{Y_t}{\tilde{K}_t}$. 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

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

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.¹⁷ 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.

¹⁶These 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.

¹⁷This structure follows Bilal and Kahn (2000) and is standard in modeling demand for goods drawn from inventories. It also supports a convenient decentralization of production.

The firm acquires each variety i from the distributors at relative price $p_{it} = P_{it}/P_t$, where $P_t = \left[\int_0^1 v_{it} P_{it}^{1-\theta} di \right]^{\frac{1}{1-\theta}}$ 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 i th variety:

$$S_{it} = v_{it} p_{it}^{-\theta} S_t. \quad (2)$$

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:

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

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

$$X_{it} = A_{it} - S_{it}, \quad (4)$$

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 i th 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:

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

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:¹⁸

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

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.¹⁹

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

¹⁸We discuss details of the other shock processes in the online appendix, where we estimate the model.

¹⁹In Appendix F we estimate a full version of the model including a suite of shocks and all structural mechanisms that we examine in the main body of the paper.

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.²⁰

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

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

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

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

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.²¹ 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 in $\tau(\chi_t)$ governs the trade-off between these two margins.

We now define $\chi(\tau_t) = \tau^{-1}(\tau_t)$, 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:

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

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

²⁰Our 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.

²¹The 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:

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

where $Q_t = z_t \Omega_t^{1-\alpha_k} (\tilde{K}_t)^{\alpha_k}$, 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.²² We combine (10) and (11) to eliminate sales S_t :

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

such that the output demand equation reads:

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

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):

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

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 comovement condition (see appendix C.4 for the detailed derivations):

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

where $\hat{y}_t > 0$, $\varepsilon_x = \left| \frac{\chi'(\tau)}{\chi(\tau)} \tau \right|$ 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}_t^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}_t^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

²²The following discussion is closely related to the theoretical results in Crouzet and Oh (2016). An important difference is that we focus on non-stationary technology news shocks rather than on their stationary counterparts. The former has received considerably more empirical support than the latter (see e.g. Schmitt-Grohe and Uribe (2012) and Görtz and Tsoukalas (2018)). We further consider the effect of variations in capital utilization in our analytical analysis as it is a potentially important factor to facilitate expansions in stock holdings.

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 (23). Consider an alternative model, where the production technology is now given by

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

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:

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

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

$$MRS_t = \Phi_t \tau_t F_{Nt}, \quad (19)$$

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:

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

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:

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

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:

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

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.²³

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.²⁴ 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. (2002) 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

²³We 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).

²⁴We detail the values of the additional parameters unique to the sticky wage and price model in the Appendix D.3.

technological knowledge through their experiences in engaging labor in the production process.²⁵

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:

$$Y_t = z_t (\Omega_t N_t)^{\alpha_n} \tilde{K}_t^{\alpha_k} (\Omega_t H_t)^{1-\alpha_n-\alpha_k}, \quad (23)$$

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

$$H_{t+1} = (1 - \delta_h)H_t + H_t^{\gamma_h} N_t^{1-\gamma_h}, \quad \text{where} \quad 0 \leq \delta_h \leq 1, \quad 0 \leq \gamma_h < 1, \quad \nu_h > 0. \quad (24)$$

The knowledge capital accumulation (24) nests a log-linear specification for $\delta_h = 1$ common in the literature such as in Chang et al. (2002), Cooper and Johri (2002) and d'Alessandro et al. (2019), but also allows for a more general linear formulation for $0 < \delta_h < 1$.²⁶

The intermediate goods firm's optimization problem now involves choosing N_t , \tilde{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 $\Pi_t^Y = \tau_t Y_t - w_t N_t - r_t \tilde{K}_t$. 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 (24), these are given by, respectively:

$$w_t = \tau_t \alpha \frac{Y_t}{N_t} + q_t^h (1 - \gamma_h) \frac{H_t^{\gamma_h} N_t^{1-\gamma_h}}{N_t}, \quad (25)$$

$$q_t^h = \beta E_t \frac{\lambda_t}{\lambda_{t+1}} \left\{ (1 - \alpha_n - \alpha_h) \tau_{t+1} \frac{Y_{t+1}}{H_{t+1}} + q_{t+1}^h \left(1 - \delta_h + \gamma_h \frac{H_{t+1}^{\gamma_h} N_{t+1}^{1-\gamma_h}}{H_t} \right) \right\}. \quad (26)$$

The presence of internalized knowledge capital in the firm's technology adds an additional term into the firm's hours-worked first order condition (25) 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

²⁵The 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. (2002) and Cooper and Johri (2002) study the propagation properties of learning-by-doing in the context of business cycle models. Since then various researchers have exploited these properties to help business cycle models better fit various features of the data. This includes Gunn and Johri (2011), who show how learning-by-doing can yield comovement of consumption, investment, hours worked, and stock prices in response to TFP news. More recently, d'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.

²⁶In specification (24), 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.

(26) 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 (25) 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:

$$\phi_t^{ld} = 1 - q_t^h(1 - \gamma_h) \left(\frac{H_t^{\gamma_h} N_t^{1-\gamma_h}}{w_t N_t} \right). \quad (27)$$

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:

$$\phi_t^e = \frac{Y_t}{z_t (\Omega_t N_t)^{\alpha_n} \tilde{K}_t^{\alpha_k} (\Omega_t H)^{1-\alpha_n-\alpha_k}} = \left(\frac{H_t}{H} \right)^{1-\alpha_n-\alpha_k}. \quad (28)$$

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.²⁷ Notably, inventories now rise on impact and then increase in the ensuing periods along with the other major macroeconomic variables.²⁸ We can again understand this response through the perspective of our wedges

²⁷We detail the values of the additional parameters unique to the knowledge capital model in the Appendix E.4. We estimate the full version of the model featuring both knowledge capital and sticky wages and prices in Appendix F, where we also compare the sticky wage and price model with knowledge capital to a version without knowledge capital. The knowledge capital version scores considerably higher on account of the (log) marginal data density.

²⁸Figure 8 shows a relative scale between output and exogenous TFP compared to the VAR-based responses in section 2.2. Note however that the TFP shown in Figure 8 is not the model counterpart to that in the VAR-based response which is based on Fernald's growth accounting methodology which does not account for intangible capital. Rather, applying Fernald's growth accounting methodology to the model corresponds to equation (23) $z_t \Omega_t^{1-\alpha_k} \left(\frac{H_t}{K_t} \right)^{1-\alpha_n-\alpha_k}$, which we call measured TFP. The scale of the model-based response of measured TFP is in line with the empirical responses in section 2.2.

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, (26)). 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 capital 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 capital by using additional labor (see the firm's first-order condition (25)). In effect, the rise in the value of knowledge capital 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.²⁹

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.³⁰

²⁹The expansion in knowledge capital, which is a key component for the described model dynamics, is consistent with the empirical evidence on the response of proxies for knowledge capital discussed in section 2.4.

³⁰It is well known that theoretical models struggle to replicate the empirically observed short-lived decline in inflation documented in section 2.2 (see e.g. Kurmann and Otrok (2017)). While many standard frameworks almost necessitate inflation to rise to generate an expansion in response to a positive news shock, the presence of knowledge capital and its dampening effect on the rise in marginal costs allows for an expansion in our model that comes with an extremely mild increase in inflation. This flat path for inflation is consistent with the VAR-based inflation response, with the exception that the empirical inflation response shows a short lived decline at the time the news about higher future technology arrives.

4 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.

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. (2002). 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. and Tsoukalas, J. (2018). Sectoral TFP news shocks. *Economics Letters*, 168(C):31–36.
- Görtz, C., Tsoukalas, J., and Zanetti, F. (2021). News Shocks under Financial Frictions. *American Economic Journal: Macroeconomics*, forthcoming.
- 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 Otrok, C. (2017). News Shocks and Inflation: Lessons for New Keynesians. Technical report.
- 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. (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

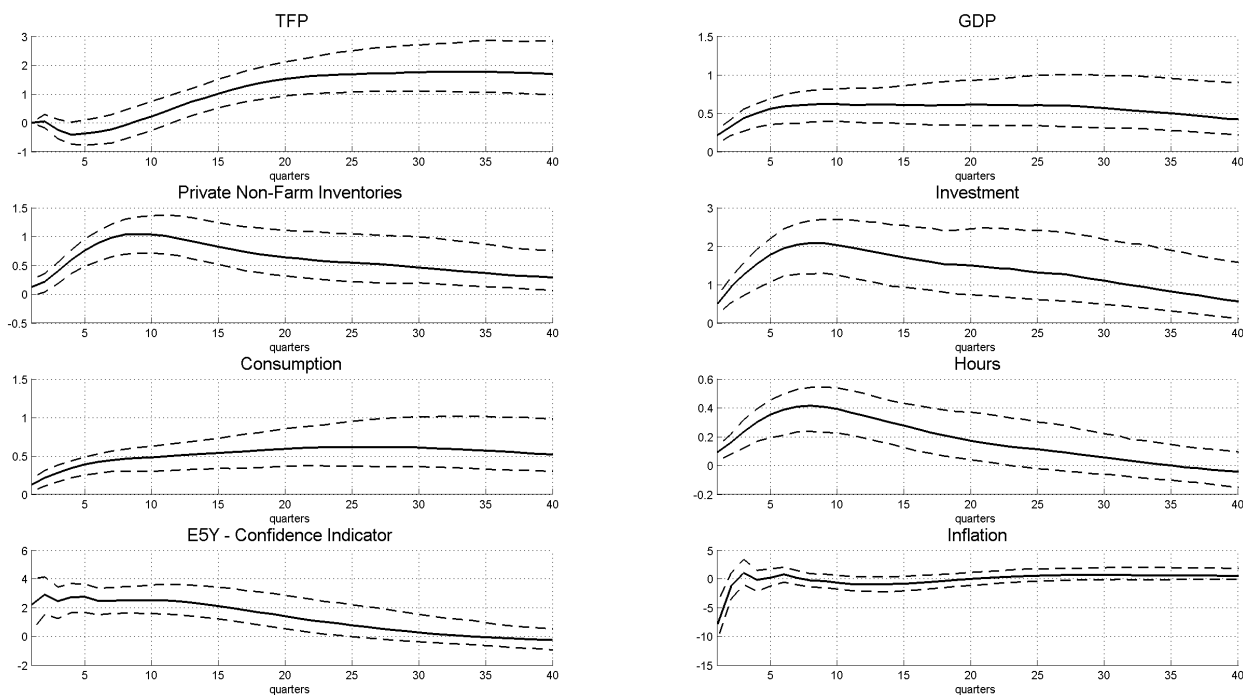


Figure 1: **IRF to TFP news shock – including Private Non-Farm Inventories.** Sample 1983Q1-2018Q2. The solid line is the median and the dashed lines are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.




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.

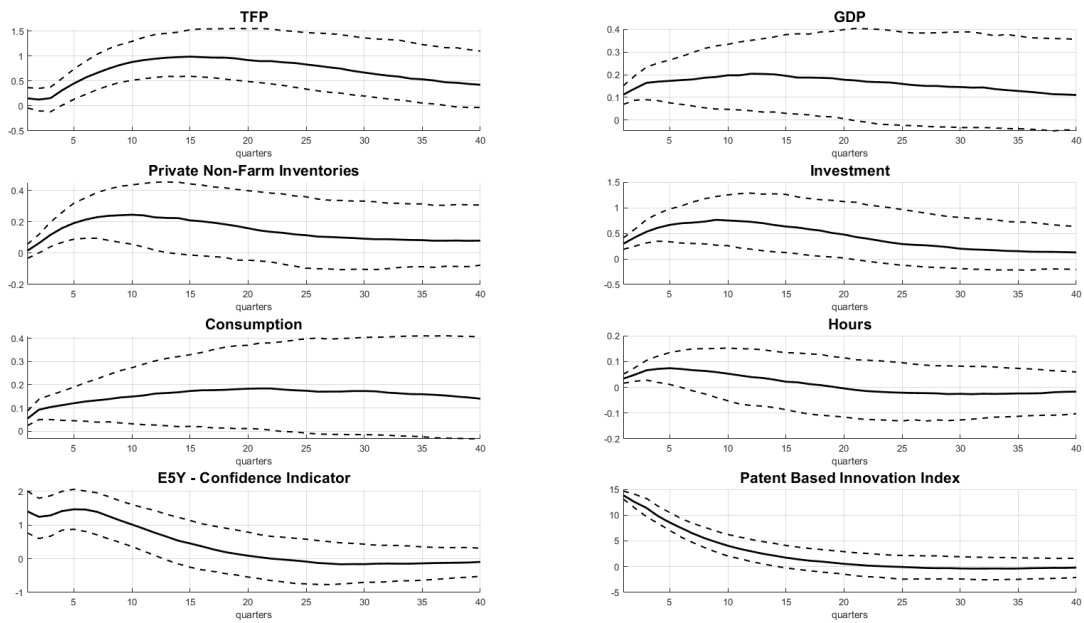


Figure 3: **IRF to patent based innovation shock.** The solid line is the median and the dashed lines are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

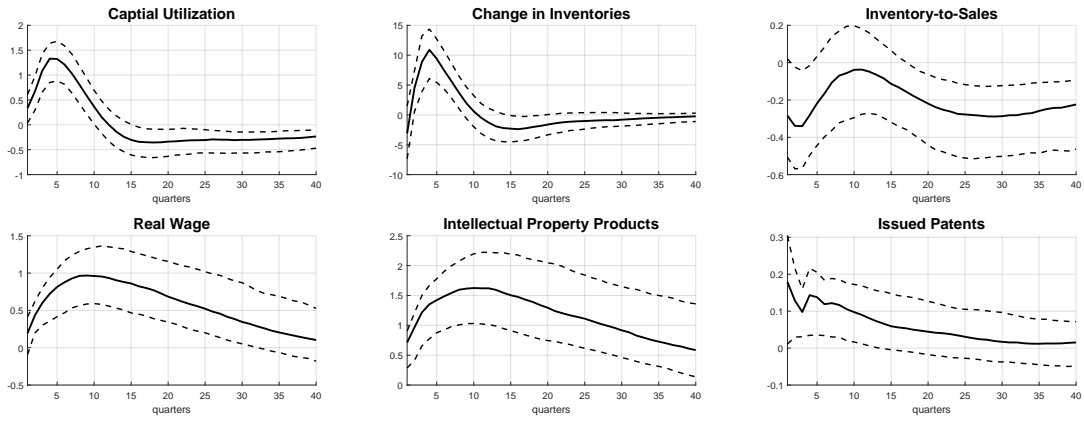


Figure 4: **IRF to TFP news shock. Max Share identification.** Subplots result from VARs comprising TFP, GDP, investment, hours, inflation and one of the plotted variables above at a time. The solid line is the median and the dashed lines are the 16% and 84% posterior bands generated from the posterior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

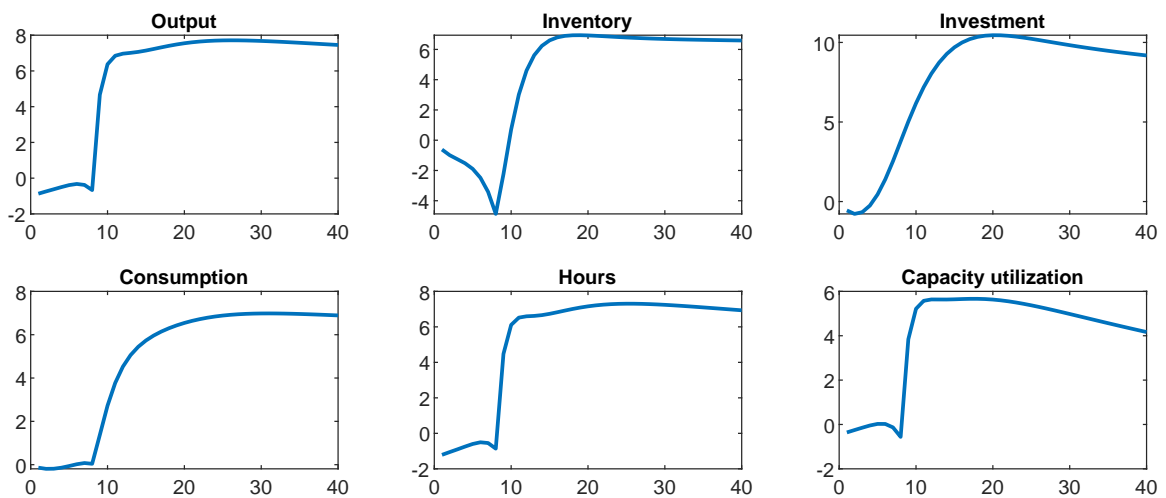


Figure 5: **IRF to 8-period out non-stationary TFP news shock: baseline model**




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

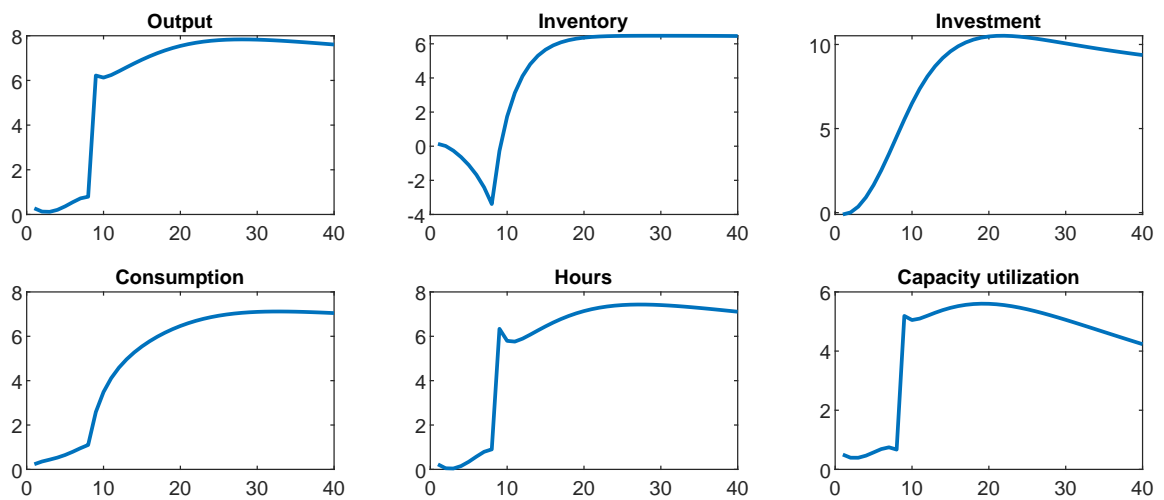


Figure 7: IRF to 8-period out non-stationary TFP news shock: *Sticky wage and price model*

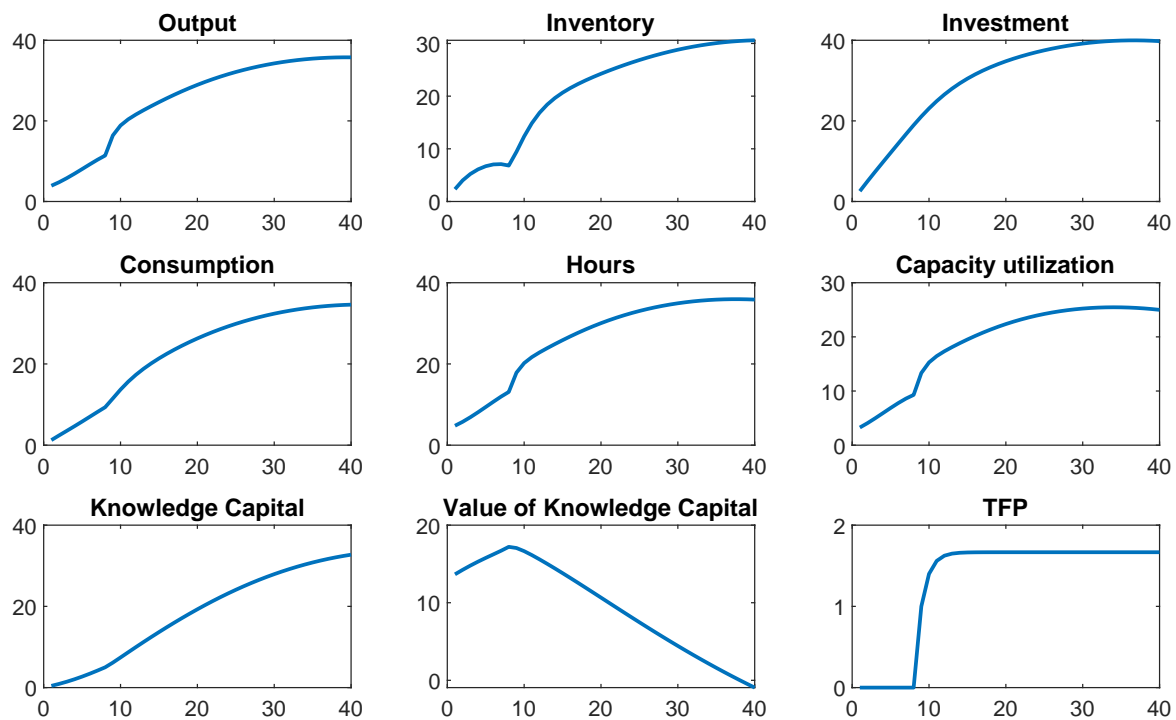


Figure 8: **IRF to 8-period out non-stationary TFP news shock: *Learning-by-doing model***