CEP 19-09

What Drives Inventory Accumulation?
News on Rates of Return and Marginal Costs

Christoph Görtz
University of Birmingham

Christopher Gunn
Carleton University

Thomas A. Lubik
Federal Reserve Bank of Richmond

October 1th, 2019

CARLETON ECONOMIC PAPERS

Department of Economics

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

What Drives Inventory Accumulation? News on Rates of Return and Marginal Costs*

Christoph Görtz
University of Birmingham†

Christopher Gunn
Carleton University‡

Thomas A. Lubik
Federal Reserve Bank of Richmond§

October 2019

Abstract

We study the effects of news shocks on inventory accumulation in a structural VAR framework. We establish that inventories react strongly and positively to news about future increases in total factor productivity. Theory suggests that the transmission channel of news shocks to inventories works through movements in marginal costs, through movements in sales, or through interest rates. We provide evidence that changes in external and internal rates of return are central to the transmission for such news shocks. We do not find evidence of a strong substitution effect that shifts production from the present into the future.

Keywords: Structural VAR, News Shocks, Inventories
JEL Classification: C32, E22, E32, G31

*We are grateful to Felipe Schwartzman for very useful comments and suggestions. The views expressed in this paper are those of the authors and should not be interpreted as those of the Federal Reserve Bank of Richmond or the Federal Reserve System.

†Department of Economics. JG Smith Building, 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

Firms use inventories in a myriad of ways as part of the production and sales process. In a sense, inventories are the ultimate residual in that surprise movements in demand can be addressed by adding unsold products to the inventory stock or by running down this stock in the face of excess demand. Similarly, materials inventories serve as a buffer to fluctuating input demand and supply. At the same time, inventories can also have a strategic aspect for a firm in that they allow for demand and production smoothing by choice.

There is a long line of theoretical and empirical work that studies the determinants of inventories. In this paper we focus on an aspect that recently has garnered some attention in the macroeconomics literature, namely anticipation of future movements in technology. Such "news shocks" are arguably an important component of inventory management as firms have to forecast future sales and the costs of maintaining and adjusting the inventory stock. While the former can be addressed by drawing on inventory holdings the latter are a function of the costs of current and future production. News about future technological advancements can thus affect inventories through a variety of channels as a mechanism to shift economic activity over time.

We show empirically that news shocks have a significant impact on inventory movements from the moment the news arrives to the point when a technological improvement materializes. Furthermore, we provide evidence on the economic determinants of such movements. We do so in a structural VAR framework where we allow for news about future total factor productivity (TFP) movements to affect variables in the present. Such shocks are identified following the standard prescriptions in the news shock literature. We establish the baseline result that news shocks raise inventory holdings on impact before higher TFP is actually realized.

The documented expansion of the inventory stock in response to news about higher future TFP is not a priori self-evident. Conventional views about inventory behavior suggest that, on the one hand, such news provide incentives to run down the current inventory stock and increase stockholdings in the future when the high productivity is realized. This intertemporal substitution effect is closely related to movements in marginal costs, which are both costs of production and costs of restocking inventories and are thus expected to fall when TFP rises in the future. In addition to this substitution effect, the associated rise in sales of consumption and investment goods creates an incentive to increase inventories due to motives such as avoiding stockouts or enhancing demand. This second aspect of

inventory dynamics can be thought of as a demand effect. To the extent that both effects are present, our results suggest a strong positive demand effect which is not outweighed by a negative substitution effect. At least as far the response to news shocks is concerned, our findings thus provide support to the demand-enhancing motive for holding more inventories in light of rising sales as in Bils and Kahn (2000).

We arrive at this conclusion by investigating the transmission mechanism leading to the documented increase in inventories. In particular, we construct aggregate measures of debt and equity cost of capital as well as implied cost of capital measures from firm-level data. This is based on the findings of Jones and Tuzel (2013) who suggest that these rate of return measures move countercyclically with inventories. We find that all measures decline significantly in response to a TFP news shock and prior to the realization of higher TFP. This decline in the opportunity cost of holding inventories thus supports the documented expansion along this margin.

We further study the response of various measures of marginal cost to a TFP news shock. Declining marginal costs between the time the news about higher future TFP arrives and its actual realization suggests a negative substitution effect. However, once introduced in our structural VAR none of our marginal cost measures exhibits such a decline. This leads us to conclude that the behavior of inventories in response to news shock is likely due to a demand effect tied to generating sales and satisfying consumption decisions while the production side in terms of movements in marginal cost over time does not appear to play a large role. More specifically, our findings indicate that a negative substitution effect is not a key determinant.

Our results offer important implications for the existing literature. For one, our finding of a procyclical inventory response is further evidence in favor of the view that news about the future is an important determinant of aggregate fluctuations. Had our empirical estimates shown that the substitution effect is a dominant force, this would have gone against the grain of the insights in Beaudry and Portier (2004), for instance. The behavior of inventories thus serves as a litmus test for this branch of the literature. Second, we provide evidence on an important model component for introducing inventories. Specifically, our findings support the stock-elastic demand model of Bils and Kahn (2000) as an alternative to other inventory frameworks. Finally, our findings also address the relationship between inventories and interest rates which the previous literature, e.g. Maccini et al. (2004), found difficult to address. We show that the proper measure for the interest-rate component is the risk premium and not the level of real interest rates.

We proceed as follows. In the next section, we document the effects of identified news shocks on inventories in a structural VAR framework. Against this background, we disentangle the effects of news shocks on several determinants of inventory accumulation in section 3, specifically external and internal rates of return and marginal cost. The final section summarizes and concludes. An Appendix provides detail on the data construction and additional robustness checks.

2 TFP News Shocks and Their Effect on Inventories

Anticipation of future total factor productivity (TFP) movements is a potentially important source of aggregate fluctuations (e.g., Beaudry and Portier, 2004). A large empirical literature shows that such news shocks are a significant driver of macroeconomic variables, specifically output and investment, and move such variables contemporaneously.[1] A macroeconomic quantity that has not received much attention in this literature is inventories.[2] We consider inventory holdings as a key component for assessing the impact and transmission of news shocks in an economy. Inventories are essentially a forward-looking variable, but they also reflect the residual effects of news shocks on sales and production. If the effects of TFP news are such that they lead to differential responses of the latter variables, such a wedge shows up as a change in inventories. Moreover, future TFP movements would prompt anticipatory actions by firms to use inventories accumulation strategically. In this paper, we investigate these trade-offs and assess their importance.

We first provide empirical evidence that news shocks have a significant impact on aggregate inventory accumulation. We do so by estimating a Bayesian VAR that captures the joint evolution of aggregate quantities, including inventories, and a process for technology. The analysis is based on Görtz et al. (2019) to whom we refer for more details and evidence for robustness of our results. The VAR includes U.S. GDP, total hours worked, investment as the sum of fixed investment and durable consumption expenditure, consumption as the sum of expenditure on non-durable consumption and services, and the S&P500 stock market index as a proxy for an expectations process that captures forward-looking information. We include the utilization-adjusted TFP series provided by Fernald (2012) as a basis for identifying news shocks. Finally, we use non-farm private inventories as our inventory measure. They are defined as the physical volume of inventories owned by private non-farm

[1] For instance, see Barsky and Sims (2011) and Schmitt-Grohe and Uribe (2012). More recently, Görtz and Tsoukalas (2017) show that TFP news shocks are relevant drivers of the cycle.

[2] Exceptions are Crouzet and Oh (2016) and Vukotic (2019) who do not provide empirical evidence on the transmission mechanisms behind the inventory response.

businesses, valued at average prices of the period.

We identify a news shock by following the convention in the empirical literature. The news shock component is identified based on the Max Share method of Francis et al. (2014). We assume that, first, the news shock does not move TFP on impact, and second, that the news shock maximizes the variance of TFP at a specific long but finite horizon. We assume this horizon to be 40 quarters in line with the literature. All quantity variables enter in levels, are seasonally adjusted and in real per-capita terms, except for hours, which are not deflated. We estimate the VAR using quarterly data for the period 1985Q1 to 2015Q1.[3] The Appendix contains further details on the VAR specification and the identification strategy.

Figure 1 reports the baseline result from inventories data. It shows impulse response functions to an identified TFP news shock from the eight-variable VAR with three lags as specified above. The graphs depicts the median responses and the 16-84% coverage regions from the posterior distribution of VAR parameters. All activity variables increase prior to the significant rise in TFP which occurs after 12 quarters. While comovement between output, consumption, investment and hours over this post-Great Moderation sample has been documented before, our new finding is the corresponding increase in the stock of private non-farm inventories in response to a news shock. Its hump-shaped adjustment pattern shows that inventory investment is positive until about three years out, shortly before the higher productivity level is actually realized.

This finding establishes the stylized fact that inventories rise on impact in response to news about higher future TFP. Görtz et al. (2019) show that this result is robust to alternative identification schemes, alternative data sets, and that it holds in various sectors and for various types of inventories. Moreover, the authors show that news shocks are central to explaining fluctuations in inventories and GDP. They capture between 43-59% (44-66%) of the forecast error variance in inventories (GDP) over a time horizon from 6-32 quarters.[4] In the remainder of this paper we disentangle these responses in terms of two mechanisms.

[3] Our choice of sample period is limited by several considerations. First, the end date of the sample is restricted by data availability for the cost of capital measures, in particular by data on new order to shipments of durable goods which is provided by Jones and Tuzel (2013). Moreover, we are limited by the availability of Lettau and Ludvigson's (2001) consumption-wealth ratio measure that figures prominently in the construction of the equity cost of capital. For comparability with the VAR based on these measures we therefore decided to restrict the aggregate VAR to the same sample period. Results using the most recent data do not show any notable difference and are available on request.

[4] In the Appendix we also report results on the effects of unanticipated TFP shocks as a consistency check and show that inventories increase alongside the other macro aggregates in response to a surprise shock.

3 News and the Forces behind Inventory Accumulation

The behavior of inventories in response to TFP shocks is not a priori self-evident. We can broadly distinguish two forces that affect inventory accumulation. The first channel operates directly on the cost of inventory accumulation over time. Other things being equal, increases in TFP reduce a firm's marginal cost of production and thereby the cost of re-stocking inventory holdings. This channel implies an intertemporal substitution effect. If news arrives that TFP is higher in the future relative to the present, it provides an incentive to reduce the current inventory stock and increase holdings in the future when the high productivity is realized. All else equal, it suggests that positive TFP news lowers inventories which is the opposite of the key finding in Figure 1.[5] We would therefore expect to find in the data that an identified news shock lowers marginal cost on impact to be consistent with the aggregate responses.

The second channel can be thought of as reflecting fluctuating demand. Given production, or assuming that production cannot respond quickly enough to satisfy changing demand conditions, inventories are simply residuals of excess supply over planned sales. For instance, inventories increase in light of rising consumption and investment to ensure a stable inventory-to-sales ratio and to avoid stockouts.[6] Positive news stimulates consumption and investment, thereby sales, to which firms respond by increasing inventory holdings. To the extent that both these effects are present, and that marginal costs fall over time, our results suggest that the intertemporal substitution effect is dominated by the positive demand effect.

This section sheds light on these transmission channels by providing evidence from detailed micro-data. We capture the first channel, the intertemporal substitution effect related to the supply of inventory-relevant output, by measuring the behavior of marginal cost directly using a production function approach as in Nekarda and Ramey (2013). The underlying assumption is that the anticipation of higher future TFP implies a decline in future marginal cost and hence an incentive to substitute production from today into the future by drawing down the inventory stock. Section 3.3 considers the response of marginal costs to a TFP news shock for a wide variety of specifications.

[5] The impulse responses in Figure 1 also show that TFP does not move in response to the news for 12 quarters which provides a somewhat clean window on the intertemporal substitution effect in that actual measured TFP matters for marginal cost.

[6] In addition, inventories by themselves can be productive for generating sales. This demand effect has been documented by Bils and Kahn (2000). For the purposes of this paper, we treat this aspect as indistinguishable from a pure demand effect as it operates in the same direction.

As far as the second channel is concerned we take guidance from Jones and Tuzel (2013) and utilize the relationship between internal and external rates of return and inventory accumulation. These authors show that there is a tight, negative relationship between inventory growth and the risk premium, as measured by the cost of capital. We extend their work by studying how news shocks affect the latter which reflects the risk of holding inventories, for instance, as a result of input inventories taking time to be transformed into final products, or finished goods inventories being subject to uncertainty about demand. We regard changes in this risk premium as indicative of the business cycle and thereby the demand for credit and, ultimately, sales.[7] The relationship between news about higher future TFP and a decline in the risk premium serves as an indicator of this effect.[8] We consider the debt and equity cost of capital as an external opportunity cost and the implied cost of capital as an internal measure in sections 3.1 and 3.2. The former is constructed from aggregate data, while the latter is constructed from firm level data.

3.1 News and the Debt and Equity Cost of Capital

Our analysis proceeds in two steps. First, we construct measures of risk premia, that is, the excess return on portfolios of either stocks or bonds, following the methodology of Jones and Tuzel (2013). They show that the debt and equity cost of capital is negatively related to inventory investment. As the former fall in an economic upturn, in line with an expansion of potential sales, inventory investment rises which also reflects lower holding costs. In order to assess the relevance of this channel, we add the equity and debt cost of capital measures separately in a seven-variable VAR system and identify news shocks in the same manner as before.

The risk premia are constructed from standard regressions of excess returns on a set of predictive variables. Specifically, we use as dependent variable either the return on the US stock market minus the one-month Treasury bill return (RMRF) or the return on corporate bonds minus the one-month Treasury bill return (RBRF). As regressors, we include seven independent variables based on their predictive power from previous work (Jones and Tuzel, 2013). These include: the term spread (TERM), the default spread (DEF), the dividend

[7] To put it differently, a rise in inventories in response to news is suggestive of a positive total effect. Since we do not find evidence of a negative substitution effect, what might be labeled as a demand effect is likely to be the dominate driver of a positive total effect.

[8] This finding also resolves a long-standing puzzle in the inventory literature discussed, for instance, by Maccini et al. (2004), namely the lack of an empirical relationship between real interest rates and inventory accumulation which virtually all theoretical models predict. That is, the relationship is between the risk premium and cost-of-capital measures and not the level of real interest rates.

yield (DP), the ratio of new orders to shipments of durable goods (NOS), the consumption-wealth ratio (CAY) of Lettau and Ludvigson (2001), as well as the real return on a nominally riskless asset (RF) and the four-quarter moving average of this variable (RF4).[9] We then use the fitted values from these regressions as measures of the equity cost of capital and debt cost of capital, respectively.[10]

Figures 2 and 3 show impulse response functions of selected variables from the two VAR specifications in response to a TFP news shock. We find that both cost-of-capital measures decline significantly for several years after the arrival of news. As in the baseline case, TFP rises significantly around the three-year mark after the news shocks. In both specifications, inventories increase on impact and remain strongly elevated over the full identification horizon. Excess returns thus move countercyclically to otherwise expansionary news shocks. This pattern can also be interpreted as a decline in the opportunity costs for holding inventories. At the same time, it is consistent with the interpretation of the rise in inventories prior to movements in TFP as driven by a strong demand effect.

This finding based on a structural VAR confirms the results of Jones and Tuzel (2013). At the same time, it adds an additional layer in that it shows that a driver of the negative relationship between inventory investment and the external cost of capital is news about future higher TFP. The latter stimulates sales and investment and leads to inventory accumulation to satisfy the additional current and future demand in line with the inventory framework of Bils and Kahn (2000). In addition, positive news reduces the risk premium and the cost of capital. We now turn to an alternative, internal measure of the cost of capital to investigate the robustness of this mechanism.

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

[10] All seven independent variables enter with one lag, whereby we select those predictors that minimize the Akaike Information Criterion (AIC). For the regression on excess stock market returns, RMRF, this criterion selects DP, which has a coefficient of 1.76*, and the intercept is -0.02 (significance at the 10% (1%) level is indicated by * (***)). For the excess corporate bond return RBRF the regression includes TERM (3.5931***), RRF4 (1.1270***), DP (0.6617***), CAY (0.2527***) and the intercept (0.0433***) where the coefficients are given in parentheses.

3.2 News and the Implied Cost of Capital

The implied cost of capital (ICC) is a firm's internal rate of return that equates the present value of expected future cash flows with the current stock price. We now construct measures of the ICC from firm level data as a proxy of the opportunity costs of holding inventories. We follow the literature and consider a variety of specifications based on different identification assumptions.[11] We use quarterly firm-level data of listed non-financial corporations from Compustat and CRSP to estimate expected earnings and use these to construct the firm-level ICC measures.[12] The actual procedure follows the methodologies summarized in Hou et al. (2012) closely.[13] We aggregate quarterly firm level observations of a particular ICC measure to a quarterly time series by taking the average per quarter. The resulting time series for the four ICCs are then used one-by-one in the seven-variable VAR, as in the previous subsection.

Figure 4 shows that all measures decline significantly in response to a TFP news shock, in a manner similar to the behavior of the external rate of return as measured by the debt and equity cost of capital. Moreover, there are no notable qualitative differences between the responses of the four measures which suggests that the results are robust to changes in the data construction procedure. The behavior of the other variables in the VAR to the news shocks remains unchanged from the baseline.

Overall, we find that external and internal rates of return decline in response to a positive news shock. This finding is broad-based across aggregate and micro-level data and robust across various specifications. It indicates a decline in the opportunity costs of inventories. Specifically, a news shock increases demand which firms respond to by increasing inventory holdings. This is reflected in a decline in the cost of holding inventories, as established in this analysis. We now turn to studying the other plausible channel, namely an intertemporal substitution effect as captured by marginal cost.

[11] These ICC measures can be broadly classified in three categories: (i) Easton (2004) and Ohlson and Juettner-Nauroth (2005) are based on so-called abnormal earnings growth models; (ii) Gebhardt et al. (2001) is based on the individual income valuation model; and (iii) Gordon (1997) is based on a generic growth model. The models differ in terms of assumptions about short- and long-term growth rates, their use of forecasted earnings, and the explicit forecast horizon.

[12] Our dataset contains all firms at the intersection of the CRSP return files and the Compustat fundamentals files. We explain how the dataset is constructed and cleaned in detail in Appendix B.2.

[13] Details of the ICC construction can be found in Appendix B.

3.3 News and Marginal Cost

A firm's marginal cost is a measure of the resources required to produce an additional unit of output. Movements in TFP are thereby a key driver of marginal cost and as such can be expected to be sensitive to news about future TFP increases. Standard models on the effect of news shocks (e.g., Jaimovich and Rebelo, 2009; Crouzet and Oh, 2016; Görtz et al., 2019) identify an intertemporal production smoothing channel. A future increase in TFP implies *ceteris paribus* lower marginal cost relative to their level today so that it becomes relatively cheaper to produce at the time the higher productivity is realized. Firms may therefore shift production intertemporally into the future. Similarly, the marginal cost of production is related to the marginal cost of inventory investment.[14] Therefore, a news shock gives an incentive to lower inventory holdings in the present as re-stocking in the future becomes less costly.

Our results so far show that current inventories rise in response to news. This finding suggests that the intertemporal substitution effect via production smoothing is not present in the data or is not strong enough to overcome the demand effect identified in the previous exercise. To investigate this question we follow the template in Nekarda and Ramey (2013) of constructing several measures for marginal costs and estimate their response to identified news shocks in our baseline VAR.

In a competitive market, real marginal cost MC is given by:

$$MC_t = \frac{W_t/P_t}{F_h\left(K_t, H_t\right)}, \tag{1}$$

where W/P is the real wage and $F_h\left(K, H\right)$ is the marginal product of labor. The specific functional form of marginal cost depends on assumptions about the production function. Under Cobb-Douglas technology the natural logarithm of real marginal cost is proportional to the labor share:

$$\log\left(MC_t\right) \approx \log(s_t), \tag{2}$$

where the labor share $s = \frac{(W_t/P_t)H_t}{F_h(K_t,H_t)}$. Alternatively, we consider a CES production function, where real marginal cost can be written as:

$$\log\left(MC_t\right) \approx \log(s_t) - \left(\frac{1}{\sigma} - 1\right)\left[\log(Y_t) - \log\left(Z_t H_t\right)\right]. \tag{3}$$

Technology is denoted by Z_t, σ is the elasticity of substitution between capital and labor and Y_t is output in value added terms.[15]

[14] The two marginal cost concepts differ when inventories serve the purpose of generating sales as in the Bils and Kahn (2000) framework. See Lubik and Teo (2012) for further discussion.

[15] If output is measured as gross output, the same expression obtains as long as the production function is

Marginal cost measures based on the two technology specifications are constructed with alternative definitions of the labor share. We consider the labor share in the private business sector and, alternatively, the nonfarm business version, both provided by the BLS. As a measure for technology, we use John Fernald's utilization-adjusted TFP series, where we set σ at a baseline value of 0.5 in line with Nekarda and Ramey (2013).[16] We use non-financial corporate business gross value added as measure for output which we divide by population. Hours H is defined as hours worked of all persons in the non-farm business sector. Any nominal values are deflated by the GDP deflator. We also consider two additional measures that correct the labor share for overhead labor based on the approach in Nekarda and Ramey (2013). We multiply BLS data on employees, average weekly hours and average hourly wages (all of production and nonsupervisory employees in the private sector) and then divide by current dollar output in private business.

Figure 5 shows the responses of the marginal cost measures when they are included one by one in our baseline VAR. The first panel depicts the response of marginal cost based on a Cobb-Douglas production function and the private business sector labor share. The measure does not move in anticipation of news about higher future TFP - it declines, in fact - but increases around the time when TFP rises significantly.[17] Under the assumption of a CES production function, this marginal cost measure in the second panel rises once the increase in TFP is realized. For the first few quarters after the arrival of the news there is a decline in this marginal cost measure, similar to the previous case. While this does provide evidence for intertemporal substitution it is in the opposite direction of what we postulated.

The third and fourth panels of Figure 5 show responses of the marginal cost measures when accounting for overhead labor. For either production function these measures do not move upon the arrival of news about higher future TFP. Moreover, they increase after several quarters. While the CES-based measure starts to decline from its peak only very slowly, the Cobb-Douglas specification declines somewhat earlier and even falls below zero after about 8 years. When using the alternative labor share measure based on the nonfarm business sector following Galí et al. (2007), responses are qualitatively and quantitatively very similar. These are shown in Appendix C where we provide further evidence on robustness

either (i) generalized CES where the elasticities of substitutions are equal across all inputs; or (ii) nested CES where the elasticity of substitution between labor input and a composite of the other inputs (see Nekarda and Ramey, 2013). For this reason we use the value added measure for Y.

[16] Appendix C.2 contains an extensive robustness analysis with respect to this parameter.

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

of the exercises related to marginal cost measures.

Overall, we note that none of the marginal cost measures exhibits a significant decline upon the arrival of the TFP news shock or even in the first quarters when the increase in TFP is realized around 3 years later.[18] Only one measure that accounts for overhead labor falls below the zero line after higher TFP has been realized. However, it is significant only after about eight years which is arguably a long time after the realization of higher TFP. We conclude that none of the marginal cost measures indicate support for a strong negative substitution effect that shifts production into the future and draws down the inventory stock upon arrival of news about higher future TFP. Taken together with the evidence in the preceding section on the presence of a strong demand effect, this behavior of marginal cost is thus consistent with the increase in inventories in response to higher future TFP.

4 Conclusion

Our paper contains three key findings. First, we show that news of higher future TFP increases inventory holdings on impact before reaching a peak close to when the productivity advance is realized. Second, we show that a news shock leads to an extended decline in internal and external rates of return. These are key variables in a firm's decision to hold and accumulate inventories and they move countercyclically. We interpret this relationship as consistent with a demand effect that expands the inventory stock. Third, we provide evidence of an increase in several marginal cost measures in response to a TFP news shock, but only at the time when the higher productivity is realized. At best, marginal cost somewhat declines on impact, but not significantly across various specifications.

These findings are mutually consistent in that the latter two offer an explanation for the first. The increase of inventories in response to news is consistent with a demand effect where increased sales drive inventory accumulation. At the same time, the marginal cost channel whereby lower production and re-stocking cost drive inventory accumulation is at best inoperative or moves in the opposite direction of what standard inventory models might predict. Overall, the findings in this paper strongly suggest that news about future TFP are a key driver of inventories and that the main transmission channel is through their role in satisfying demand and generating sales.

[18] In the stock-elastic demand model of Bils and Kahn (2000), there is a one-to-one mapping between marginal costs and the inventory-sales ratio. If marginal costs were countercyclical, this would imply a procyclical inventory-sales ratio, which is at odds with the data. That fact that we find marginal costs are not countercyclical is consistent with the inventory-sales ratio not being procyclical, consistent with the data.

References

[1] Barsky, Robert B., and Eric R. Sims (2011): "News shocks and business cycles". *Journal of Monetary Economics*, 58(3), pp. 273-289.

[2] Beaudry, Paul, and Franck Portier (2004): "An exploration into Pigou's theory of cycles". *Journal of Monetary Economics*, 51(6), pp. 1183-1216.

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

[4] Crouzet, Nicolas, and Hyunseung Oh (2016): "What do inventories tell us about news-driven business cycles?" *Journal of Monetary Economics*, 79, pp. 49-66.

[5] Easton, Peter D. (2004): "PE Ratios, PEG ratios, and estimating the implied expected rate of return on equity capital". *The Accounting Review*, 79(1), pp. 73-95.

[6] Fernald, John (2012): "A quarterly, utilization adjusted series on total factor productivity". Federal Reserve Bank of San Francisco Working Paper Series 2012-19.

[7] Francis, Neville, Michael Owyang, Jennifer Roush, and Riccardo DiCecio (2014): "A flexible finite-horizon alternative to long-run restrictions with an application to technology shocks". *Review of Economics and Statistics*, 96, pp. 638-647.

[8] Galí, Jordi, Mark Gertler, and J. David López-Salido (2007): "Markups, gaps, and the welfare costs of business fluctuations". *Review of Economics and Statistics*, 89, pp. 44-59.

[9] Gebhardt, William R., Charles M. C. Lee, and Bhaskaran Swaminathan (2001): "Toward an implied cost of capital". *Journal of Accounting Research*, 39(1), pp. 135-176.

[10] Gordon, Joseph R. (1997): "The finite horizon expected return model". *Financial Analysts Journal*, 53(3), pp. 52-61.

[11] Görtz, Christoph, Christopher Gunn, and Thomas A. Lubik (2019): "Is there news in inventories?" Manuscript.

[12] Görtz, Christoph, and John Tsoukalas (2017): "News and financial intermediation in aggregate fluctuations". *Review of Economics and Statistics*, 99(3), pp. 514 530.

[13] Hou, Kewei, Mathijs A. van Dijk, and Yinglei Zhang (2012): "The implied cost of capital: A new approach". *Journal of Accounting and Economics*, 53(3), pp. 504-526.

[14] Jaimovich, Nir, and Sergio Rebelo (2009): "Can news about the future drive the business cycle?" *American Economic Review*, 99(4), pp. 1097-1118.

[15] Jones, Christopher S., and Selale Tuzel (2013): "Inventory investment and the cost of capital". *Journal of Financial Economics*, 107, pp. 557-579.

[16] Lettau, Martin and Sydney Ludvigson (2001): "Consumption, aggregate wealth, and expected stock returns". *Journal of Finance*, 56(3), pp. 815-849.

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

[18] Maccini, Louis, Benjamin Moore, and Huntley Schaller (2004): "The interest rate, learning, and inventory investment". *American Economic Review*, 94, pp. 1303-1327.

[19] Nekarda, Christopher J., and Valerie A. Ramey (2013): "The Cyclical Behavior of the Price-Cost Markup". *NBER Working Paper* #19099, National Bureau of Economic Research.

[20] Ohlson, James A., and Beate E. Juettner-Nauroth (2005): "Expected EPS and EPS growth as determinants of value". *Review of Accounting Studies*, 10(2-3), pp. 349-365.

[21] Schmitt-Grohe, Stephanie and Martin Uribe (2012): "What's news in business cycles?" *Econometrica*, 80(6), pp. 2733-2764.

[22] Vukotic, Marija (2019): "Sectoral effects of news shocks". *Oxford Bulletin of Economics and Statistics*, 81, pp. 215-249.

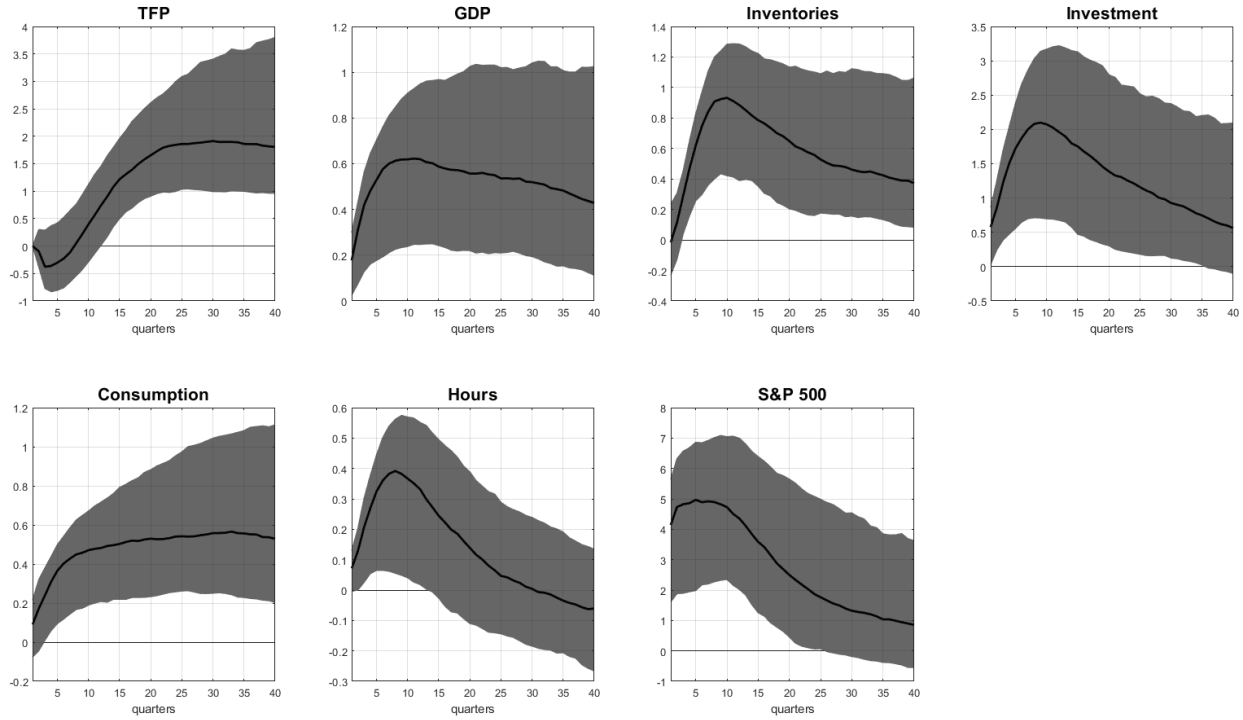

Figure 1: **IRF to TFP news shock.** Results based on a seven-variable VAR. Sample 1985Q1-2015Q1. The solid line is the median and the shaded gray areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

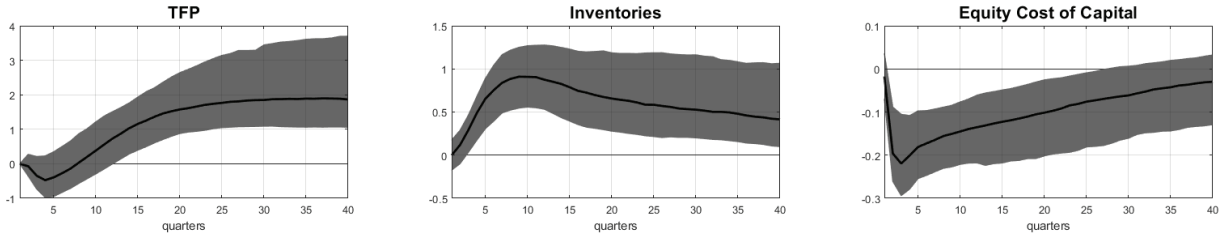

Figure 2: **IRF of Equity Cost of Capital measure to TFP news shock.** Selected variables based on a seven-variable VAR including TFP, GDP, consumption, hours, inventories, equity cost of capital, S&P 500. Sample 1985Q1-2015Q1. The solid line is the median and the shaded gray areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

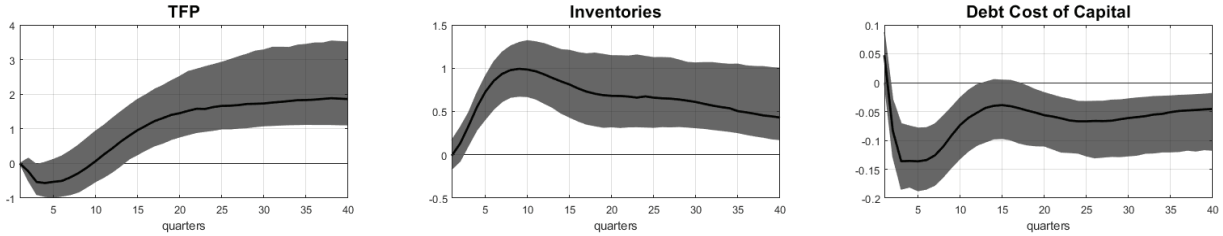

Figure 3: **IRF of Debt Cost of Capital measures to TFP news shock.** Selected variables based on a seven-variable VAR including TFP, GDP, consumption, hours, inventories, debt cost of capital, S&P 500. Sample 1985Q1-2015Q1. The solid line is the median and the shaded gray areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

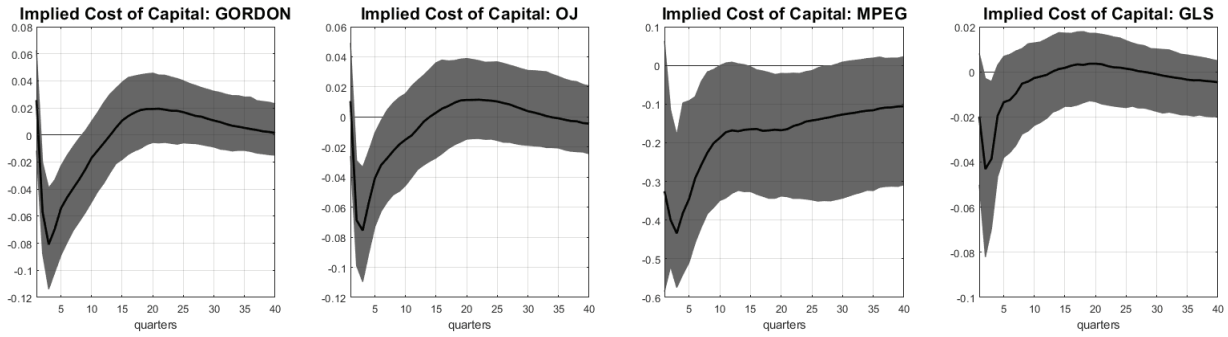

Figure 4: **IRF of Implied Cost of Capital measures to TFP news shock.** Each subplot results from a seven-variable VAR including TFP, GDP, consumption, hours, inventories, one particular measure for implied cost of capital (ICC), S&P 500. The ICC measures are constructed according to Gordon (1997) (GORDON), Ohlson and Juettner-Nauroth (2005) (OJ), Easton (2004) (MPEG), Gebhardt et al. (2001) (GLS). Sample 1985Q1-2015Q1. The solid line is the median and the shaded gray areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

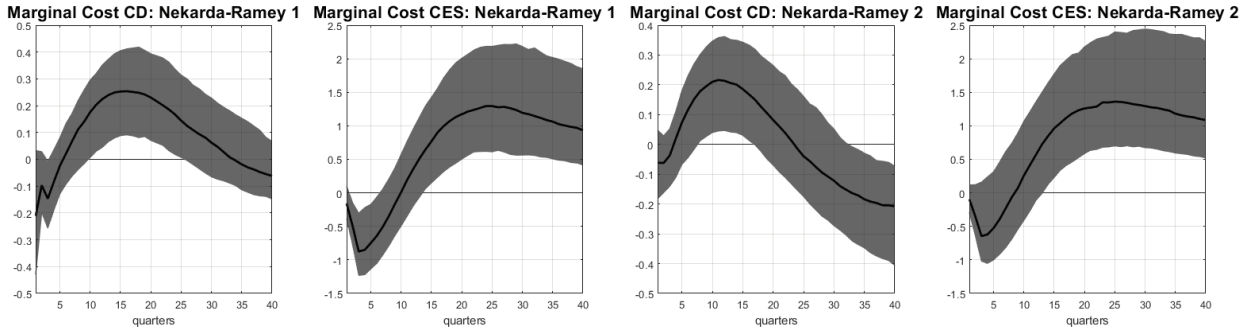

Figure 5: **IRF of marginal cost measures to TFP news shock.** Sample 1985Q1-2015Q2. Each subplot results from a seven-variable VAR including TFP, GDP, consumption, hours, inventories, one particular measure for marginal cost, S&P 500. Marginal cost construction is based on Nekarda and Ramey (2013): CD/CES indicates the use of a Cobb-Douglas/CES production function and 1/2 refers to the use of the private business sector labor share/measure for overhead labor. The solid line is the median and the shaded gray areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

Appendix

A Details on the VAR

We specify the following reduced-form VAR of lag length p:

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

where y_t is an $n \times 1$ vector and $A(L)$ is a lag polynomial of order p over comformable coefficient matrices $\{A_p\}_{i=1}^{p}$. u_t is an error term with covariance matrix Σ. We define the structural errors ε_t from the mapping:

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

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

$$y_t = C(L)u_t, \tag{A.3}$$

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

We can define the h-step ahead forecast error as:

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

The share of the forecast error variance of variable i that can be attributed to shock j at horizon h is then:

$$v_{i,j}(h) = \frac{e_i'\left(\sum_{\tau=0}^{h} A_\tau \widetilde{B}_0 D e_j e_j' D' \widetilde{B}_0' A_\tau'\right) e_i}{e_i'\left(\sum_{\tau=0}^{h} A_\tau \Sigma A_\tau'\right) e_i} = \frac{\sum_{\tau=0}^{h} A_{i,\tau}\widetilde{B}_0 D\gamma\gamma'\widetilde{B}_0' A_{i,\tau}'}{\sum_{\tau=0}^{h} A_\tau \Sigma A_\tau'}, \tag{A.5}$$

where e_i denotes a selection vector with one in the i-th position and zeros elsewhere, while the e_j vector picks out the j-th column of D, denoted by γ. $\widetilde{B}_0\gamma$ is an $n \times 1$ vector that corresponds to the j-th column of a possible orthogonalization of the estimation error covariance matrix. It therefore can be interpreted as an impulse response vector.

In the following, we discuss the methodology that identifies the TFP news shock from the VAR model. This so called Max Share methodology is based on Francis et al. (2014) who isolate unanticipated productivity shocks by maximizing the forecast error variance

share of TFP at a long but finite horizon. At a long enough horizon h all variations in TFP are either accounted for by anticipated or unanticipated shocks to this variable. We can then write:

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

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

Hence, we choose the second column of the impact matrix to solve the following optimization problem:[1]

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

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

B Constructing Implied Cost of Capital Measures

We use firm-level data from Compustat and the Center for Research in Security Prices (CRSP) to estimate implied cost of capital measures. Section B.1 provides details on the construction of different implied cost of capital measures. Section B.2 documents the underlying dataset construction.

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

B.1 General Approach

The estimation of firm-level implied cost of capital (ICC) measures requires a measure for earnings forecasts. Based on Hou et al. (2012) and closely related to Fama and French (2000, 2006), we generate such forecasts by estimating the following pooled cross-sectional regression for each quarter from 1985Q1, using the previous ten years of data. Specifically, we estimate the regression:

$$E_{i,t+\tau} = \beta_0 + \beta_1 A_{i,t} + \beta_2 D_{i,t} + \beta_3 DD_{i,t} + \beta_4 E_{i,t} + \beta_5 NegE_{i,t} + \beta_6 AC_{i,t} + \varepsilon_{i,t+\tau}. \quad \text{(B.1)}$$

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

We construct four different, but widely used ICC measures based on Easton (2004), Ohlson and Juettner-Nauroth (2005), Gebhardt et al. (2001) and Gordon (1997).[2] For this purpose, we merge the Compustat data with information from CRSP on market equity (MVAL) defined as the product of Number of Shares Outstanding (CSHO) and the Stock Price at the end of the quarter (PRCC). We further use the 1-Year Treasury Constant Maturity Rate as risk free rate. Prior to computing earnings forecasts and ICC measures we apply the cleaning procedures outlined in Section B.2 below to the Compustat-CRSP dataset.

We use this dataset to compute the different ICC measures at time t for firm i. In particular, the measure according to Gordon (1997) is computed using:

$$MVAL_{i,t} = \frac{E_t\left[EA_{i,t+1}\right]}{ICC_{i,t}}, \quad \text{(B.2)}$$

where the implied cost of capital is denoted by $ICC_{i,t}$, $MVAL_{i,t}$ is market equity and $EA_{i,t+1}$ is the earnings forecast for $t+1$ based on information available at time t. E_t is the expectations operator associated with the earnings forecast.

[2] See e.g. Ashbaugh-Skaife et al. (2009), Hail and Leuz (2009) and Chava and Purnanandam (2010).

The ICC measure according to Easton (2004) is computed using:

$$MVAL_{i,t} = \frac{E_t\left[EA_{i,t+2}\right] + ICC_{i,t} \times E_t\left[D_{i,t+1}\right] - E_t\left[EA_{i,t+1}\right]}{ICC_{i,t}^2}, \tag{B.3}$$

where $D_{i,t+1}$ denotes the dividend in $t+1$, which is computed using the using the current dividend payout ratio (for firms with positive earnings), or the current dividends divided by 6% of the total assets as an estimate of the payout ratio (for firms with negative earnings).

The ICC measure according to Ohlson and Juettner-Nauroth (2005) is computed using:

$$ICC_{i,t} = 0.5\left(\frac{E_t\left[D_{i,t+1}\right]}{MVAL_{i,t}} + (\gamma_t - 1)\right) + \left[0.25\left(\frac{E_t\left[D_{i,t+1}\right]}{MVAL_{i,t}} + (\gamma_t - 1)\right)^2 + \frac{E_t\left[EA_{i,t+1}\right]}{ICC_{i,t}}\left(g_t - (\gamma_t - 1)\right)\right]^{1/2}, \tag{B.4}$$

with the short-term growth rate given by:

$$g_t = 0.5\left(\frac{E_t\left[EA_{i,t+3}\right] - E_t\left[EA_{i,t+2}\right]}{E_t\left[EA_{i,t+2}\right]} + \frac{E_t\left[EA_{i,t+5}\right] - E_t\left[EA_{i,t+4}\right]}{E_t\left[EA_{i,t+4}\right]}\right), \tag{B.5}$$

as in Gode and Mohanram (2003). γ_t t is the perpetual growth rate in abnormal earnings beyond the forecast horizon which is set to the current risk-free rate minus 3%.

The ICC measure according to Gebhardt et al. (2001) is computed using:

$$MVAL_{i,t} = B_{i,t} + \sum_{\tau=1}^{11} \frac{E_t\left[(ROE_{i,t+\tau} - ICC_{i,t}) \times B_{i,t+\tau-1}\right]}{(1 + ICC_{i,t})^{\tau}} + \frac{E_t\left[(ROE_{i,t+12} - ICC_{i,t}) \times B_{i,t+\tau+11}\right]}{ICC_{i,t} \times (1 + ICC_{i,t})^{11}}, \tag{B.6}$$

where $B_{i,t}$ is book equity and $ROE_{i,t}$ is the return on book equity. The expected return on book equity is determined based on clean surplus accounting as $B_{i,t+\tau} = B_{i,t+\tau-1} + EA_{i,t+\tau} - D_{i,t+1}$.

Each of the four different firm-level ICC estimates is aggregated to a time series. We thereby follow the convention in the literature and replace any firm-time ICC estimates below zero by a missing value. We further set the top one percentile of all firm-time observations for a particular ICC measure to missing prior to aggregating the firm observations by taking averages over each quarter.

B.2 Cleaning the Compustat-CRSP Dataset

Our dataset contains all firms at the intersection of the CRSP return files and the Compustat fundamentals files. We select the sample by making the following adjustments to the data retrieved from Compustat-CRSP:

- We delete all regulated, quasi-public or financial firms (primary SIC classification is between 4900-4999 and 6000-6999).

- We delete firms that reported earnings in a currency other than USD.
- We account for the effects of mergers and acquisitions by deleting all observations that include firms with (i) acquisitions (ACQ) exceeding 15% of total assets (ATQ), or (ii) sales growth exceeding 50% in any year due to a merger.
- We drop companies with all values for total assets (AT) or investment in plant, property and equipment (CAPX) that are missing or zero. We drop missing observations for CAPX if they are at the beginning or end of a company's reported data. If CAPX is missing in the middle of a company's reported data we drop the entire company.
- We drop firms with less than three quarters of data.
- We apply the following filters to key variables:
  - We replace missing values of DPQ with zero.
  - We set negative values of CHEQ, DLCQ, DPQ and DVPQ to missing.
  - We set values smaller or equal to zero of ACTQ, LCTQ, ATQ and MVAL to missing.
  - We winsorize, that is, we limit outliers or extreme values, of IBQ at the top and bottom percentile.
  - We winsorize ATQ, ACTQ, LCTQ, CHEQ, DLCQ, DPQ, DVPQ and MVAL at the top percentile.
- ATQ, ACTQ, LCTQ, CHEQ, MVAL, DLCQ, IBQ and DPQ are deflated applying the Gross Domestic Product: Implicit Price Deflator. DPQ is deflated applying the Gross Private Domestic Fixed Investment: Nonresidential Implicit Price Deflator.

The cleaned dataset consists of 19,599 firms and 781,478 observations for the time horizon 1985Q1-2015Q2.

C Additional VAR Evidence

We report two sets of additional evidence from the structural VAR. First, we present results on the effects of surprise TFP shocks which gives additional insight into the role of the transmission channels of TFP movements, whether anticipated or unanticipated. Second, we offer additional evidence on the response of marginal cost measures.

C.1 Surprise TFP Shocks

We report a selected set of impulse response function from our baseline VAR with aggregate inventories in Figures A.1 - A.4. In addition, we also consider the effects of a surprise innovation to TFP for the specification with various marginal cost measures. The identification of these unanticipated TFP shocks is as discussed in section A of the appendix. We summarize insights from these additional exercises in the following.

Figure A.1 shows strong comovement of the key macroeconomic aggregates with the exception of inventories who do not rise on impact, but do so only gradually. This suggests that, in fact, news shocks are a key driver of inventory accumulation. The hours response is initially negative and then largely insignificant. Figure A.2 reports the responses of equity and debt cost of capital measures. They decline in response to a surprise TFP shock, whereas the responses of the internal cost of capital measures in Figure A.3 are more mixed.

We would expect marginal costs to fall on account of the higher realized productivity, all other things being equal. The evidence on this somewhat mixed, however. Measures based on Cobb-Douglas production tend to decline on impact. The Nekarda-Ramey 2 measure declines strongly on impact and is the only one that is in line with prior expectations, whereas Nekarda-Ramey 1 declines with a long delay. In contrast, CES-based MC measures increase strongly on impact, tracking the TFP response closely (as for the news shock). Naturally, this observation runs counter to what theory would suggest. This is arguably due to the construction of the CES-based measures. These are a function of TFP and other variables (in levels). However, when constructing the series there is no guidance on the level of TFP (in relation to the other variables). We note that Fernald provides the TFP series in growth rates, whereby we reconstruct the level by arbitrarily choosing an initial level. Alternatively, theory does not offer much guidance either as the simple intuition may be incorrect.

C.2 Additional Evidence on Marginal Costs

Figure A.5 shows the response of two marginal cost measures to a TFP news shock when they are included one-by-one in a seven-variable VAR. The two measures in the figure are constructed using the preferred measure for the labor share by Galí et al. (2007), namely the BLS labor share in the non-farm business sector. They are either based on the CES (CES: Gali et al.) or Cobb-Douglas (CD: Gali et al.) production function. Qualitatively and quantitatively the responses of these two marginal cost measures to a TFP news shock are

very similar to the responses shown in Figure 5 in the main text when using the labor share measure preferred by Nekarda and Ramey (2013) (CES: Nekarda-Ramey 1, CD: Nekarda-Ramey 1). In line with the discussion in the main text, neither of the two marginal cost measures in Figure A.5. provides evidence for a strong negative substitution effect through a fall in marginal costs. This is consistent with the rise in inventories we report in response to a TFP news shock.

Table A.1 shows the unconditional correlations of HP-filtered GDP with all our considered measures for marginal costs. Marginal costs are acyclical or mildly countercyclical which is in line with the evidence in Nekarda and Ramey (2013). They report that markups are acyclical or mildly procyclical. In addition to the abbreviations explained in the paragraph above, we note that CD: Nekarda-Ramey 2 and CES: Nekarda-Ramey 2 refer to the marginal cost measures which are constructed by considering a measure for overhead labor under the assumption of either a Cobb-Douglas or a CES production function. The results shown in Figures 5 and A.5 are robust to variations of the elasticity of substitution σ between capital and labor in the construction of the marginal cost measures. Based on the empirical literature, Chirinko (2008) concludes that plausible values for σ lie in a range between 0.4 and 0.6. Our baseline calibration is 0.5. Robustness checks using these two values yield very similar responses of all marginal cost measures to a TFP news shock. Qualitatively they are virtually unchanged. More detailed results are available upon request.

Table A.1: GDP-MC Correlations

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

Notes: Time series are HP-filtered with smoothing parameter 1,600.Sample period is 1985Q1-2015Q2.

References

[1] Ashbaugh-Skaife, H., Collins, D. and Kinney Jr, W. and Lafond, R. (2009): "The effect of SOX internal control deficiencies on firm risk and cost of equity". *Journal of Accounting Research*, 47, pp. 1-43.

[2] Barsky, Robert B. and Eric R. Sims (2011): "News shocks and business cycles". *Journal of Monetary Economics*, 58(3), pp. 273-289.

[3] Chava, S. and Purnanandam, A. (2010): "Is default risk negatively related to stock returns?" *Review of Financial Studies*, 23, pp. 2523-2559.

[4] Chirinko, Robert S. (2008): "σ: The long and short of it". *Journal of Macroeconomics*, 30(2), pp. 671-686.

[5] Easton, Peter D. (2004): "PE ratios, PEG ratios, and estimating the implied expected rate of return on equity capital". *The Accounting Review*, 79, pp. 73-95.

[6] Fama, Eric, and Kenneth French (2000): "Forecasting profitability and earnings". *Journal of Business*, 73, pp. 161-175.

[7] Fama, Eric, and Kenneth French (2006): "Profitability, investment and average returns". *Journal of Financial Economics*, 82, pp. 491-518.

[8] Francis, Neville, Michael Owyang, Jennifer Roush, and Riccardo DiCecio (2014): "A flexible finite-horizon alternative to long-run restrictions with an application to technology shocks". *Review of Economics and Statistics*, 96, pp. 638-647.

[9] Galí, Jordi, Mark Gertler, and J. David López-Salido (2007): "Markups, gaps, and the welfare costs of business fluctuations". *Review of Economics and Statistics*, 89, pp. 44-59.

[10] Gebhardt, William R., Charles M. C. Lee, and Bhaskaran Swaminathan (2001): "Toward an implied cost of capital". *Journal of Accounting Research*, 39, pp. 135-176.

[11] Gode, D. and Mohanram, P. (2003): "Inferring the cost of capital using the Ohlson-Juettner model". *Review of Accounting Studies*, 8(4), pp. 399-431.

[12] Gordon, Myron J. (1997): "The finite horizon expected return model". *Financial Analysts Journal*, 53(3), pp. 52-61.

[13] Görtz, Christoph, Christopher Gunn, and Thomas A. Lubik (2019): "Is there news in inventories?" Manuscript.

[14] Hail, L. and Leuz, C. (2009): "Cost of capital effects and changes in growth expectations around U.S. cross-listings". *Journal of Financial Economics*, 93, pp. 428-454.

[15] Hou, Kewei, Mathijs A. van Dijk, and Yinglei Zhang (2012): "The implied cost of capital: A new approach". *Journal of Accounting and Economics*, 53(3), pp. 504-526.

[16] Hribar, Paul and Daniel W. Collins (2002): "Errors in estimating accruals: Implications for empirical research". *Journal of Accounting Research*, 40(1), pp. 105-134.

[17] Nekarda, Christopher J., and Valerie A. Ramey (2013): "The Cyclical Behavior of the Price-Cost Markup". *NBER Working Paper* #19099, National Bureau of Economic Research.

[18] Ohlson, James A. and Beate E. Juettner-Nauroth (2005): "Expected EPS and EPS growth as determinants of value". *Review of Accounting Studies*, 10, pp. 349-365.

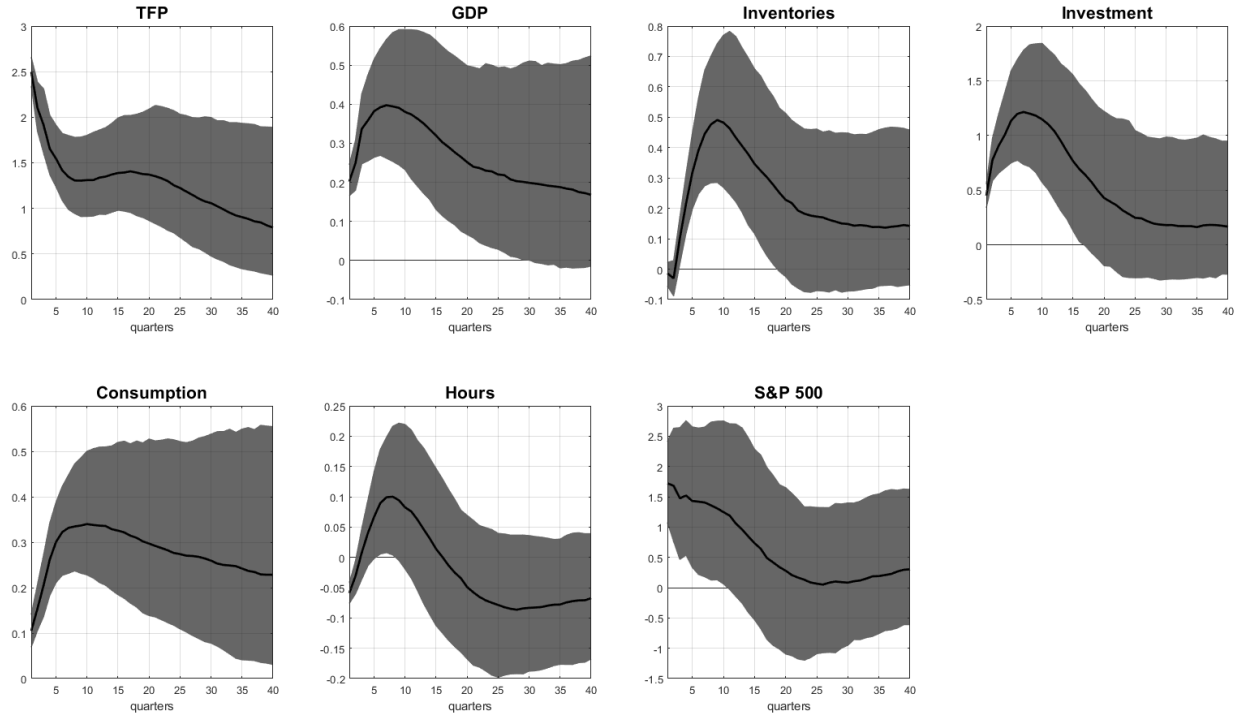

Figure A.1. **IRF to TFP surprise shock.** Results based on a seven-variable VAR. Sample 1985Q1-2015Q1. The solid line is the median and the shaded gray areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

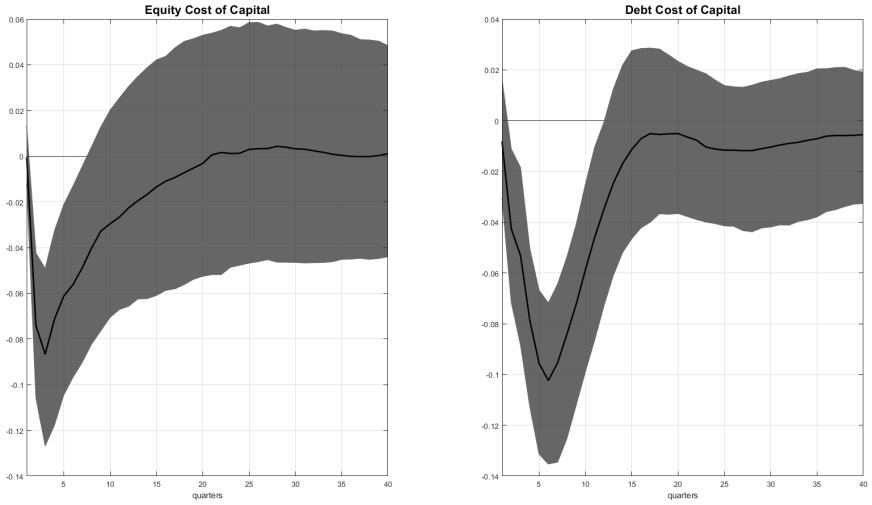

Figure A.2. **IRF of Equity and Debt Cost of Capital measure to TFP surprise shock.** Selected variables based on a seven-variable VAR including TFP, GDP, consumption, hours, inventories, equity or debt cost of capital measure, S&P 500. Sample 1985Q1-2015Q1. The solid line is the median and the shaded gray areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

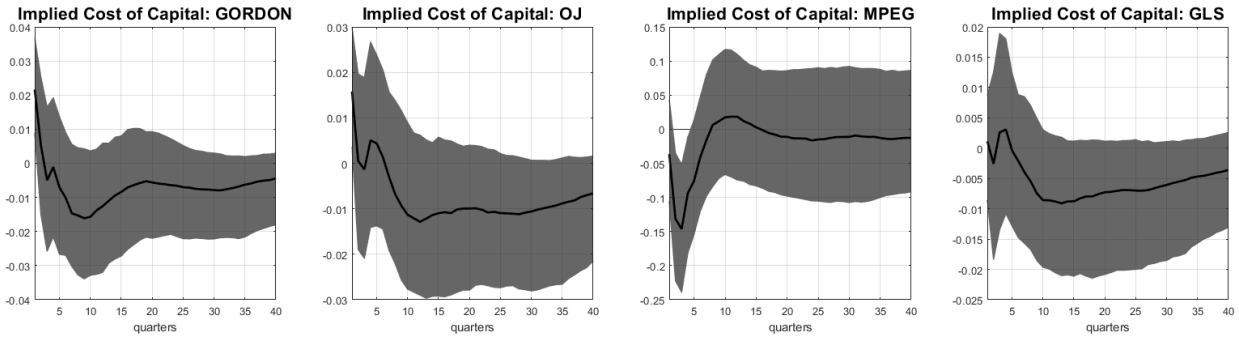

Figure A.3. **IRF of Implied Cost of Capital measures to TFP surprise shock.** Each subplot results from a seven-variable VAR including TFP, GDP, consumption, hours, inventories, one particular measure for implied cost of capital (ICC), S&P 500. The ICC measures are constructed according to Gordon (1997) (GORDON), Ohlson and Juettner-Nauroth (2005) (OJ), Easton (2004) (MPEG), Gebhardt et al. (2001) (GLS). Sample 1985Q1-2015Q1. The solid line is the median and the shaded gray areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

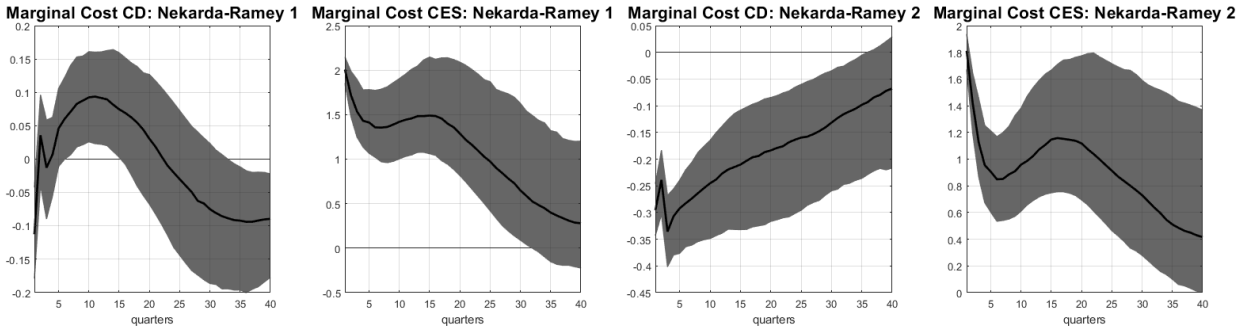

Figure A.4. **IRF of marginal cost measures to TFP surprise shock.** Sample 1985Q1-2015Q2. Each subplot results from a seven-variable VAR including TFP, GDP, consumption, hours, inventories, one particular measure for marginal cost, S&P 500. Marginal cost construction is based on Nekarda and Ramey (2013): CD/CES indicates the use of a Cobb-Douglas/CES production function and 1/2 refers to the use of the private business sector labor share/measure for overhead labor. The solid line is the median and the shaded gray areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

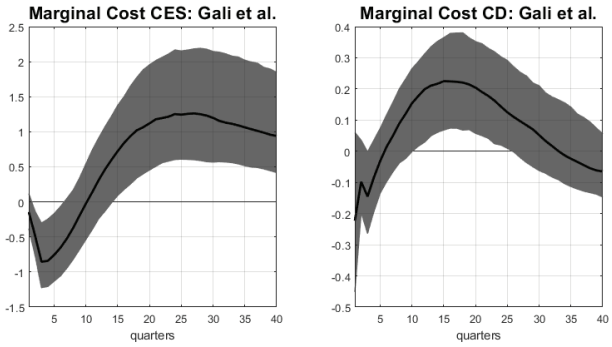

Figure A.5. **IRF of marginal cost measures to TFP news shock.** Sample 1985Q1-2015Q2. Each subplot results from a seven-variable VAR including TFP, GDP, consumption, hours, inventories, one particular measure for marginal cost, S&P 500. Marginal cost construction is based on Galí et al. (2007): CD/CES indicates the use of a Cobb-Douglas/CES production function. The solid line is the median and the shaded gray areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.