




Network

Rimini Centre for Economic Analysis Working Paper Series

wp 22-11

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

**Christoph Görtz
Christopher Gunn
Thomas A. Lubik**



RCEA aims to further independent, advanced research in Economics, Econometrics and related fields and to promote contact between economists, econometricians and scientists from other fields all over the world. Research at RCEA is conducted to enlighten scientific and public debate on economic issues, and not to advance any economic, political or social agenda. In this respect, RCEA mission is particularly concerned with raising awareness and stimulating discussion on the changes required to make capitalism sustainable along its economic, environmental, human, social and political dimensions

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

Christoph Görtz
University of Birmingham & RCEA[†]

Christopher Gunn
Carleton University[‡]

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

September 2022

Abstract

We study the determinants of inventory accumulation in a structural VAR framework with news shocks. Specifically, we investigate how news shocks affect two key determinants of inventory movements, namely rates of return and marginal costs. We establish that inventories react strongly and positively to news about future increases in total factor productivity. 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 for a dominant role of marginal costs.

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

^{*}We are grateful to the editor Sanjay Chugh, two anonymous referees, Pete Klenow and 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. 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

Inventories play a major role in the extent and volatility of business cycle fluctuations. 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) or Irvine and Schuh (2005)). Aggregate inventories, in their dual role as input and output inventories, are also central to business cycle transmission for instance via production networks (Iacoviello et al. (2011); Sarte et al. (2015)). Consequently, there is a long line of theoretical and empirical work that has studied the underlying driving forces and channels of inventory accumulation. Much of this work has focused on studying inventories in the context of standard disturbances such as surprise shifts in Total Factor Productivity (TFP) or various forms of shifts in demand. There has been relatively little work studying inventories in response to anticipated future movements in TFP, which have emerged as a prominent candidate driving force of business cycles.

Such “news shocks” are arguably an important component of inventory management and planning ahead since 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. In addition, 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. News about future technological advancements can thus affect inventories through a variety of channels as a mechanism to shift economic activity over time.

In this paper, we investigate two key determinants of inventory movements related to TFP news, namely rates of return and marginal costs, found to be important by the literature (see e.g. Bilal and Kahn (2000); Iacoviello et al. (2011); Lubik and Teo (2012); Kryvtsov and Midrigan (2013)). We do so in a structural VAR framework where we allow for news about future TFP movements to affect variables in the present. Such shocks are identified following standard approaches in the news shock literature. 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. We find that all measures decline significantly in response to a positive TFP news shock and prior to the realization of higher TFP. This decline in the opportunity cost of holding inventories is consistent with the rise in inventory and other macroeconomic aggregates that we document from the arrival of the news. We further study

the response of various measures of marginal costs to TFP news shocks. We find that all measures of marginal costs grow through the expansion phase toward the eventual arrival of higher TFP. This documented expansion of the inventory stock in response to positive TFP news is not a priori self-evident. The conventional view suggests that such news would provide incentives to run down the current inventory stock and increase stockholdings in the future when the high productivity is realized. This 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.

Our findings in this paper help reconcile the empirical evidence regarding inventories and TFP news with the conventional understanding of inventory behavior. Specifically, we trace out the standard mechanisms in theoretical inventory models and evaluate their empirical relevance. In that way, our work provides guidance for the modelling of inventories, particularly in light of our empirical result that inventories respond strongly to news shocks as they are forward looking, and news shocks themselves are found to be important for aggregate fluctuations. Our results on the response of marginal costs do not indicate support for the conventional view of a strong negative substitution effect that shifts production into the future and draws down current inventory in the face of anticipated future productivity growth. In contrast, news shocks reduce real rates of return and therefore the opportunity cost of holding inventories.

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 results shown that, conditional on TFP news shocks, inventories did not comove positively with the other macroeconomic aggregates, this countercyclical movement would not be consistent with the unconditional evidence. Hence, this would have gone against the grain of the insight in Beaudry and Portier (2004), who document a large role of news shocks as drivers of business cycles. The behavior of inventories thus serves as a litmus test for this branch of the literature. Although there is a large early empirical literature on the determinants of inventory investment (exemplified, for instance, by Humphreys et al. (2001)) that relies on reduced-form modeling, the more recent literature heavily relies on theoretical models. Our study is important as there is relatively limited empirical evidence of the kind we provide on the support of the different inventory models since the time when structural changes resulted in different inventory management approaches.

Finally, our findings also contribute to understanding the relationship between inventories and interest rates, which is present in many models on inventory behavior. Previous

empirical literature, e.g. Maccini et al. (2004), found this aspect difficult to resolve. We show that a predictive measure for the interest-rate component of inventory accumulation is the risk premium and not the level of real interest rates. In contrast, Copeland et al. (2019) find a relationship between real interest rates and inventories in a specific market, namely the automobile market for new light vehicles. However, our study documents that this relationship does not exist at the aggregate level and conditional on TFP news shocks. Our findings also relate to the results in Jones and Tuzel (2013), who show that convex adjustment costs and a countercyclical price of risk can explain the empirically observed positive relation between rates of return and inventory movements in a production-based asset pricing model. While they focus on the unconditional response of inventories, our work is concerned with explaining the response to TFP news shocks. Moreover, Jones and Tuzel (2013) do not consider marginal cost measures as alternative drivers of inventory accumulation.

There is now a substantial literature on TFP news shocks to which we cannot do full justice here. Several contributions stress the relevance of these anticipated shocks for aggregate fluctuations (e.g. Schmitt-Grohe and Uribe (2012), Kamber et al. (2017), Cascaldi-Garcia and Vukotić (2022)), while other contributions are more sceptical (e.g. Kurmann and Sims (2019)). Despite their importance for understanding business cycles, this literature abstracts almost entirely from inventory holdings. Two notable exceptions concerned with the relationship between news shocks and inventories are recent contributions by Crouzet and Oh (2016) and Vukotic (2019), which our work closely touches upon. The former introduce inventories into a variant of the standard news-shock model of Jaimovich and Rebelo (2009), utilizing a reduced-form stockout-avoidance specification. While they come to the conclusion that such TFP news shocks are of limited importance for aggregate fluctuations, Vukotic (2019) finds both theoretically and empirically that inventories are critical for reconciling the effects of news shocks in a two-sector model with the data. However, neither paper provides empirical evidence on the transmission mechanisms behind the inventory response.

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, marginal cost, and real interest rates. The final section summarizes and concludes. An online appendix provides detail on the data construction and additional robustness checks.

2 TFP News Shocks and Their Effect on Inventories

Anticipation of movements in TFP 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.¹ A macroeconomic quantity that has not received much attention in this literature is inventories. Firms use inventories as part of the production and sales process. In a sense, inventories serve a residual function 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 to buffer 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.

Görtz et al. (2022) provide empirical evidence that news shocks have a significant impact on aggregate inventory accumulation. We begin by confirming their result on the procyclical response of aggregate inventories to TFP news shocks in order to establish a baseline for the more disaggregated analysis performed in this paper. We estimate a Bayesian VAR that captures the joint evolution of aggregate quantities, including inventories, and a process for technology. 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. Non-farm private inventories serve as the inventory measure, defined as the physical volume of inventories owned by private non-farm businesses, valued at average prices of the period. News shocks are identified from the utilization-adjusted 2016 vintage of the TFP series provided by Fernald (2014).

We identify a news shock by following the convention in the empirical literature, specifically an extension of 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 the VAR in levels, are seasonally adjusted and in real per-capita terms, except for hours, which are in per-capita terms but not deflated. We estimate the VAR using quarterly data for the period 1985Q1

¹See Barsky and Sims (2011) and Schmitt-Grohe and Uribe (2012). More recently, Görtz and Tsoukalas (2017) and Görtz et al. (2021) argue that TFP news shocks are key drivers of the business cycle.

to 2015Q1.^{2,3} Online appendix A contains further details on the VAR specification and the identification strategy.

Figure 1 reports the baseline result from aggregate inventories data. It shows impulse response functions to an identified TFP news shock from the seven-variable VAR as specified above. The graphs depict 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 (e.g., Görtz and Tsoukalas (2018), Görtz et al. (2022)), a notable 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.⁴

The rise in all macroeconomic aggregates conditional on a TFP news shock is indicative of this shock being an important driver of aggregate fluctuations. The forecast error variance decomposition associated with the baseline VAR model in Figure 1, shows that at business cycle frequencies, between 6 and 32 quarters, the TFP news shock explains a substantial share in the variation in macroeconomic aggregates and stock prices. In particular, the news shock is an important driver of variations in GDP and inventories, explaining between 39-55% and 37-56% of their forecast error variance, respectively. Details are reported in online appendix B.1. Intuitively, inventory behavior is to a large extent driven by anticipated shocks since news are arguably an important determinant for forward looking inventory management and firms have to forecast future sales and the costly maintenance of stock

²Our 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 (2001) consumption-wealth ratio measure that figures prominently in the construction of the equity cost of capital. For comparability with the VAR systems that include these measures, we therefore decided to restrict the same sample period. Results using the most recent data, as far as data availability permits, do not show any notable difference and are available on request.

³Our data is based on a 2016Q2 vintage. Kurmann and Sims (2019), Bouakez and Kemoe (2022) and Clements and Galvão (2021) flag the potentially distortionary impact of data revisions in the measurement of the dynamic effects of news shocks. Most substantial data revisions occur in the three years after data release. Following a referee's suggestion, we have verified that our results are not materially affected when using a 2022 vintage instead of the 2016Q2 vintage (retaining the estimation horizon at 1985Q1-2010Q1).

⁴In online appendix B.3, we also show that these results are robust to employing an alternative news shock identification recently suggested Kurmann and Sims (2019). This appendix also shows that the results of the following sections on cost of capital measures and marginal cost are robust to using this alternative identification.

holdings. This makes it highly relevant to understand the underlying forces behind inventory accumulation conditional on news shocks and we turn to this next.




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.

While Görtz et al. (2022) focus on the response of aggregate inventories to TFP news shocks, it is important to go beyond this since finished goods inventories and input inventories are not necessarily as closely related as the similarity in name may suggest. In fact, the literature documents that unconditionally they are statistically rather different objects (see e.g. Humphreys et al. (2001)), which may also carry over to their response to TFP news shocks. We therefore consider the response of inventories to TFP news shocks in the retail, wholesale and manufacturing sectors. The former two sectors hold almost exclusively finished goods inventories, while in the latter sector their share reduces to approximately one third of all inventories. For the manufacturing sector, disaggregated data exists for different inventory types so that we can differentiate between finished goods inventories and input inventories. The non-farm private inventory measure considered in Figure 1 is not available at a disaggregated level. Instead, we consider business inventories for this exercise which differs from the non-farm private inventory measure in how stockholdings are valued.⁵

⁵Inventories can be valued in various ways, depending on the specific objective. While the business inventory measure allows for a more disaggregated investigation, it restricts our sample length as it is only available from 1992Q1. The difference between the business inventory and non-farm private inventory

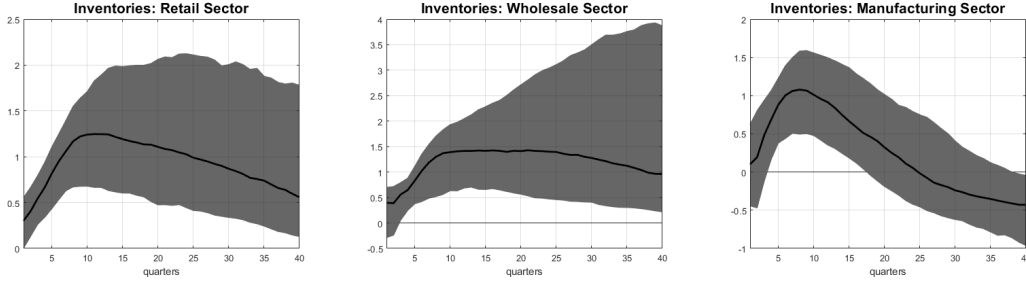


Figure 2: **IRF to TFP news shock.** Results based on a seven-variable VARs including business inventories in different sectors. Sample 1992Q1-2015Q1. The solid line is the median and the shaded gray areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

Figures 2 and 3 show responses of selected variables to a TFP news shock based on the seven-variable VAR in Figure 1, where the inventory variable is replaced (one-by-one) by sectoral or inventory-type measures.⁶ Figure 2 documents that inventories in the retail, wholesale and manufacturing sector rise well before TFP increases significantly. Figure 3 reports impulse responses of inventory types in the manufacturing sector. We show that finished goods inventories as well as input inventories rise. Our results suggest that the expansion in aggregate inventories is driven by finished goods and input inventories and is broad-based across the sectors that hold the vast majority of inventories. This allows us to proceed in the following by focusing on aggregate inventories to understand the transmission of news shocks, which also has the advantage that we are not restricted to the shorter sample imposed by the use of the business inventory measures.

3 The Forces Behind Inventory Accumulation

We motivate our empirical approach with a partial equilibrium model that highlights several key forces potentially affecting inventory accumulation. We follow Bilal and Kahn (2000) in modeling inventories as a mechanism to generate sales. At the same time it implies a target inventory-sales ratio that captures the idea of stockout avoidance. We also

measures is that the former is measured at cost of acquisition and the latter at average sales prices. As such, considering this alternative measure also serves as a robustness check for the response of inventories to a TFP news shock documented in Figure 1.

⁶In online appendix B.2, we also show responses of all variables in a VAR with aggregate business inventories corresponding to Figure 1. Despite the shorter sample and the different definition used for the aggregate inventory measure, we find results are consistent with those shown in Figure 1. In particular, we document that business inventories rise on impact in response to the news shocks while TFP does so only with a delay of more than two years.

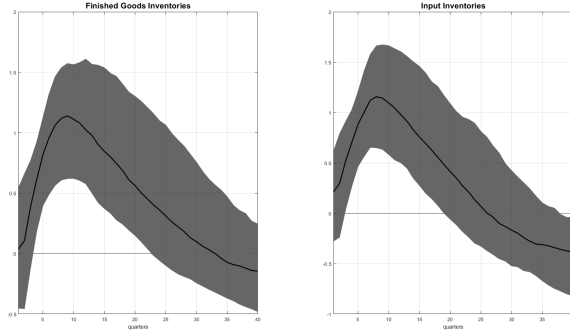


Figure 3: **IRF to TFP news shock.** Results based on a seven-variable VARs including different types of business inventories. Sample 1992Q1-2015Q1. The solid line is the median and the shaded gray areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

allow for nominal price rigidity since it leads to time-varying marginal costs of sales, which is the margin along which the interest rate and intertemporal substitution channels work. Finally, we allow for the presence of a generic financial friction to generate a distinct role for firm-specific opportunity cost of funds beyond the risk-free rate.

We consider the problem of “distributors” that acquire a homogeneous good Y_t from a perfectly competitive intermediate goods sector at real price τ_t . They differentiate Y_t into goods varieties Y_{it} , $i \in [0, 1]$, at zero cost, with a transformation rate of one-to-one. The distributors have market power over the sales of their differentiated varieties. Per the stock-elastic demand framework of Bils and Kahn (2000), the sales demand for a given variety i is increasing in the goods of that variety available for sale, A_{it} . Goods available for sale for the i th distributor are the sum of the differentiated output and the previous period’s inventories subject to depreciation $A_{it} = (1 - \delta_x) X_{it-1} + Y_{it}$, where the stock of inventories X_{it} are the goods remaining at the end of the period $X_{it} = A_{it} - S_{it}$, and $0 < \delta_x < 1$ is the rate of depreciation of the inventory stock. The i th distributor sets price P_{it} for sales S_{it} of its variety subject to its demand curve, and faces Rotemberg-style quadratic price adjustment costs $\Phi_{it} = \Phi(P_{it}, P_{it-1})$.

To introduce a role for the credit spread in inventory accumulation, we focus on a simple motivating example with minimal structure. We assume that distributors are forced to borrow from financial intermediaries at rate R_t^l to finance inventories, for reasons such as a mismatch between the timing of a distributor’s purchases of its homogeneous good and its revenues. In particular, assume that distributors must finance inventory each period in the

form of one-period loans L_{it} , such that

$$L_{it} \geq X_{it}. \quad (1)$$

We further assume the presence of financial frictions in a financial intermediation sector external to the distributor that results in a loan rate R^l as a premium to the risk-free rate, R^f , such that $R^l > R^f$. This assumption nests a number of popular frameworks in the literature such as Gertler and Karadi (2011).

Each period, the distributor then solves the problem of choosing P_{it} , S_{it} , Y_{it} , A_{it} and L_{it} to maximize discounted profits:

$$E_t \sum_{k=0}^{\infty} m_{t+k,t} \left[\frac{P_{it+k}}{P_{t+k}} S_{it+k} - \tau_t Y_{it+k} - \Phi_{it+k} + L_{it+k} - R_{t-1k}^l L_{it-1+k} \right],$$

subject to its demand curve, the law of motion for goods available for sale, the loan constraint and the definition of the inventory stock, and where $m_{t+k,t}$ is the shareholder/owner's stochastic discount factor.

Combining first-order conditions and imposing symmetric equilibrium over distributors yields an equilibrium optimal stocking condition of the form:

$$\frac{X_t}{S_t} = \chi(\tau_t, \mu_t^x), \quad (2)$$

where $\frac{\partial \chi(\tau_t, \mu_t^x)}{\partial \tau_t} < 0$, $\frac{\partial \chi(\tau_t, \mu_t^x)}{\partial \mu_t^x} > 0$, and where μ_t^x is the shadow value of inventory, given by:

$$\mu_t^x = 1 + E_t m_{t+1,t} \left\{ (1 - \delta_x) \tau_{t+1} - R_t^l \right\}. \quad (3)$$

Since any increase in sales results in a reduction of stock holdings, μ_t^x has the equivalent interpretation as the *marginal cost of sales*. Similarly, the real price of output τ_t can be interpreted as the *marginal cost of output*.

Equation (3) allows us to interpret the marginal cost of sales μ_t^x as the expected discounted value of future marginal costs, net of the expected discounted loan rate. Increasing sales by drawing down inventories in order to forgo production today means that the distributor will lower loan costs R_t^l due in the future, but they also need to eventually increase production in the future. In addition to the two equations above, the distributor's problem results in a pricing condition such that the distributor optimally sets its relative price as a markup over the marginal cost of sales μ_t^x (see Görtz et al. (2022) for additional details). Under sticky prices, this markup is dynamic and the marginal cost of sales μ_t^x is time-varying.

The optimal stocking condition (2) is the key equation governing inventory dynamics. It implies that distributors target an inventory-sales ratio $\frac{X_t}{S_t}$ for given marginal cost of output

τ_t and marginal cost of sales μ_t^x . All else equal, the distributor increases inventory holdings with a rise in sales in order to maintain the optimal inventory-sales ratio. We label this the *demand channel* of inventory accumulation. Thus, as also argued in Crouzet and Oh (2016), a TFP news-driven increase in sales implies a motive to increase inventories in turn. Similarly, inventory holdings decline with a rise in current marginal costs as the high cost of current output makes it more attractive to run down inventories to satisfy sales. We label this the *current cost channel*.

For a pure news shock, where TFP increases in the future but not in the present, the impact on inventories then depends on the direction of the endogenous response of current marginal costs to the news shock. Additionally, inventory holdings are also positively related to expected marginal costs through μ_t^x . If marginal costs are expected to be lower in the future relative to the present, inventories decline in what may be labelled the *intertemporal substitution effect*. This arises since it is optimal to defer inventory accumulation to the future when the cost of new output required to accumulate inventory is lower relative to the present.⁷ A scenario where an expected increase in TFP leads to a fall in expected future marginal costs thus implies a motive to decrease inventories in the present.⁸

In addition to these channels, the earlier literature has also identified the interest rate as a key driver of stock holdings and inventory accumulation, for instance Blinder and Maccini (1991) and especially Maccini et al. (2004). In this context, the interest rate broadly serves as the opportunity cost of holding inventory. Although there is no explicit "interest rate effect" in equation (2), there is an indirect link to both the risk-free rate and the loan rate through μ_t^x . Defining the risk-free rate as $R_t^f = \frac{1}{E_t\{m_{t+1,t}\}}$, we can see this clearly by assuming certainty equivalence, where equation (3) then implies that, all else equal, the marginal cost of sales μ_t^x is decreasing in *both* the risk-free rate R_t^f and the credit spread $\frac{R_t^l}{R_t^f}$. Inventory X_t is decreasing in both variables also through the dependence of X_t on μ_t^x in the optimal stocking condition (2), in what may be labelled as *opportunity cost effects*. Similar to the marginal cost channel above, the impact on inventories of the opportunity cost channel depends on the direction of the endogenous response of the interest rate and spread to the news boom.

The above discussion highlights several key forces of inventory accumulation, which individually may not necessarily work in the same direction in response to TFP news. In

⁷In the special case of flexible prices, μ_t^x is constant since the markup between relative price and the marginal cost of sales μ_t^x is constant. The time-varying intertemporal substitution effect is therefore not operative in (2). The same applies to the interest rate effects described below.

⁸Further details on the channels discussed above can be found in Crouzet and Oh (2016) and Görtz et al. (2022).

fact, the positive response of inventories to TFP news in Figure 1 suggests that any individual mechanism implying inventory de-accumulation in response to positive TFP news is not a dominant channel.⁹

4 Empirical Evidence on Forces Behind Inventory Accumulation

We now shed light on the key forces that affect inventory accumulation by providing evidence from aggregate and detailed firm-level data. We provide evidence on the role of current and expected future marginal costs and the opportunity costs in the context of expected future changes in TFP. Furthermore, we discuss how this evidence relates to the channels we isolated in the context of the theoretical model.

Opportunity costs. To study the opportunity costs of holding and carrying inventory we take guidance from Jones and Tuzel (2013) and utilize the relationship between internal and external rates of return and inventory accumulation. They 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 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 4.1 and 4.2, respectively. The former is constructed from aggregate data, while the latter is constructed from firm-level data.

Marginal costs. We study the marginal cost channel by measuring the behavior of marginal cost directly using a production function approach as in Nekarda and Ramey (2013). Section 4.3 considers the response of marginal costs to a TFP news shock for a wide variety of specifications.

4.1 News and the Debt and Equity Cost of Capital

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 unconditionally the debt and equity cost of capital is negatively related to inventory in-

⁹The impulse responses in Figure 1 also show that TFP does not move in response to the news for 12 quarters. This provides suggestive evidence for the relative unimportance of the intertemporal substitution effect in that actual measured TFP matters for marginal cost.

vestment, which can reflect lower holding costs. In order to assess the relevance of this mechanism for inventory accumulation in response to news, 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 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).¹⁰ We then use the fitted values from these regressions as measures of the equity cost of capital and debt cost of capital, respectively.¹¹

The top and bottom panels of Figure 4 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 thus be interpreted as a decline in the opportunity costs for holding inventories.

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

¹⁰The 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.

¹¹All 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.

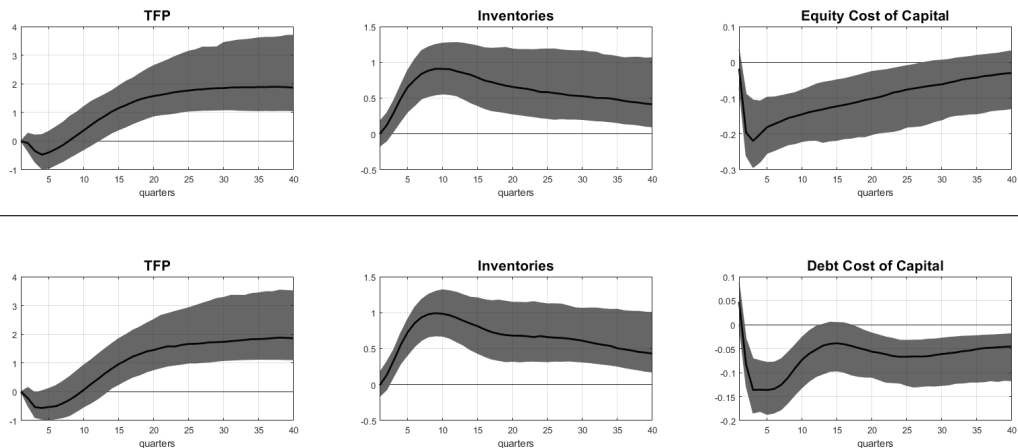


Figure 4: **IRF of Equity (Debt) Cost of Capital measure to TFP news shock** — **top (bottom) row**. Selected variables based on two seven-variable VAR systems including TFP, GDP, consumption, hours, inventories, equity (debt) cost of capital, S&P 500. Variables from the respective VAR are shown in the top (bottom) row. 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.

relationship between inventory investment and the external cost of capital is news about future higher TFP. Changes in this risk premium are indicative of the business cycle, thereby the demand for credit and thus sales. The decline in rates of return is consistent with the expansion in macroeconomic aggregates we find empirically in the VAR and as economic intuition would suggest.¹² We now turn to an alternative, internal measure of the cost of capital, to investigate the robustness of this mechanism.

4.2 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 construct measures of the ICC from firm-level data as a proxy for the opportunity costs of holding inventories. Following the literature, we consider four specifications based on different identification assumptions.¹³ We use quarterly firm-level data of listed non-financial corporations

¹²Although our results do not formally demonstrate this link, they are suggestive of a demand channel from news to inventories. Specifically, increased credit stimulates sales and investment and leads to inventory accumulation to satisfy the additional current and future demand in line with the inventory framework of Bilal and Kahn (2000).

¹³These ICC measures are widely used and 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) Joseph R. Gordon (1997)

from Compustat and CRSP to estimate expected earnings and use these to construct the firm-level ICC measures.¹⁴ The actual procedure follows the methodologies summarized in Hou et al. (2012) closely.¹⁵ 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 5 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 in response to the news shocks remains unchanged from the baseline.

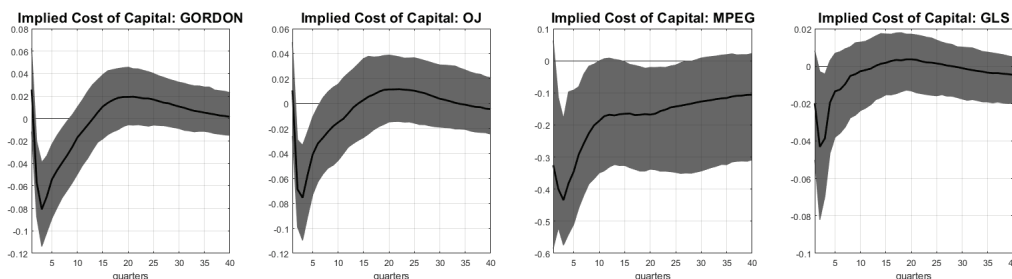


Figure 5: **IRF of Implied Cost of Capital measures to TFP news shock.** Each subplot results from a seven-variable VAR including TFP, GDP, consumption, hours, inventories, one particular measure for implied cost of capital (ICC), S&P 500. The ICC measures are constructed according to Gordon (1997) (GORDON), Ohlson and Juettner-Nauroth (2005) (OJ), Easton (2004) (MPEGE), 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.

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. In addition, lower rates of return are also consistent with an increase in demand, which firms respond to by increasing inventory holdings. We now turn to studying the other plausible

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.

¹⁴Our 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 online appendix C.2.

¹⁵Details of the ICC construction can be found in online appendix C.1.

channel, namely an intertemporal substitution effect as captured by marginal cost.

4.3 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. (2022)) 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 into the future. Similarly, the marginal cost of production is related to the marginal cost of inventory investment.¹⁶ Therefore, as standard theory would suggest, a news shock gives an incentive to lower inventory holdings in the present since 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 if an intertemporal substitution effect via production smoothing is present in the data it is not strong enough to overcome the effect of a declining risk premium, which we identified in the previous section. Another possibility is that the intertemporal effect is present, but dominated by other forces triggered by the TFP news shock. That is, if marginal cost actually rises through a news-driven boom, firms have an additional incentive to increase production now to build up inventories since it will be relatively more costly to do so in a few periods. 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. We use multiple marginal cost measures for robustness reasons since each depends on structural assumptions about technology specifications and the labor share. Details on the construction of all marginal cost measures are provided in Online Appendix D.

Figure 6 shows the responses of four marginal cost measures when they are included one by one in our baseline VAR. The first panel depicts the response of a marginal cost measure 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, but increases significantly only after ten quarters, just before the time when the rise in TFP is observed.¹⁷

¹⁶The two marginal cost concepts differ when inventories serve the purpose of generating sales as in the Bilal and Kahn (2000) framework. See Lubik and Teo (2012) for further discussion.

¹⁷The behavior of the variables in the VAR that are not shown is very similar to the ones in the baseline

The second panel shows the marginal cost response when accounting for overhead labor. Again, marginal cost do not move initially, but rise after about eight quarters. A similar pattern can be observed in the third and fourth subplot, where we consider marginal cost measures based on a CES production function, either with the private business sector labor share or using overhead labor. Marginal costs are insignificant in the short run, or decline slightly in the case of subplot three, but increase in the medium run. Both measures decline from the peak slowly so that marginal cost remain at an elevated level up to the 40 quarter horizon. The measures based on the Cobb-Douglas production function shown in the first two subplots decline somewhat faster, but only the second one falls below zero after about 8 years. When using the alternative labor share measure based on the non-farm business sector following Galí et al. (2007), responses are qualitatively and quantitatively very similar. These are shown in online appendix B.4 where we provide further evidence on robustness of the exercises related to marginal cost measures.

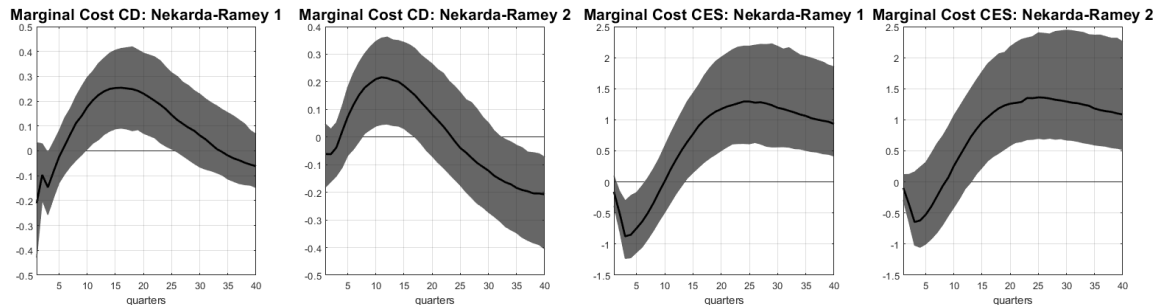


Figure 6: **IRF of marginal cost measures to TFP news shock.** Sample 1985Q1-2015Q2. Each subplot results from a seven-variable VAR including TFP, GDP, consumption, hours, inventories, one particular measure for marginal cost, S&P 500. Marginal cost construction is based on Nekarda and Ramey (2013): 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.

Overall, we note that none of the marginal cost measures except one move upon the arrival of the TFP news shock for at least two years.^{18,19} All measures eventually rise

in Figure 1, where TFP increases significantly after about 12 quarters.

¹⁸In 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. The fact that we find marginal costs are not countercyclical is consistent with the inventory-sales ratio not being procyclical, consistent with the data.

¹⁹This finding runs counter to some news-shock models in the literature. It is, however, consistent with

during the boom-phase leading into the eventual rise in TFP. Only one measure 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 therefore 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, this behavior of marginal cost is thus consistent with the increase in inventories in response to higher future TFP. Specifically, a marginal cost channel, via a negative substitution effect, appears less likely as a driver of aggregate inventory movements in response to TFP news shocks than many inventory models would suggest.²⁰

4.4 The Real Interest Rate Response

Our finding that measures of the risk premium move countercyclical to otherwise expansionary TFP news shocks, 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. In related work, Copeland et al. (2019) show there is, in fact, a relationship for a specific market, namely the automobile market for new light vehicles.

We construct the real interest rate measure in a manner similar to Copeland et al. (2019) as the difference between the nominal interest rate and inflation expectations. The former is defined as the BAA-bond yield which is the interest rate earned on investment-grade bonds.²¹ Figure 7 shows the responses to a TFP news shock when the real interest rate measure is included in the seven-variable VAR. The finding of Copeland et al. (2019) for the automotive market does not carry over when we apply their measure to the aggregate economy. The real interest rate based on the BAA bond yield remains insignificant throughout the entire horizon, while inventories increase strongly much before TFP rises significantly after about the marginal cost path in Görtz et al. (2022). This framework suggests that dampening the rise in marginal costs in standard news-shock models is a means to achieving a procyclical response of inventories to TFP news.

²⁰At best, it is possible that the upward path of marginal costs creates a *positive* substitution effect that provides the incentive to pull production forward and increase inventories.

²¹The BAA bond yield is obtained from the H.14 release of the Board of Governors. Inflation expectations are from the Federal Reserve Bank of Richmond, computed as the forecast from a univariate ARMA-model of inflation.

12 quarters. Consistent with Figure 1 the other activity variables also rise on impact.




Figure 7: **IRF to TFP news shock.** Results based on a seven-variable VAR including inventories and a real interest rate measure based on the BAA bond yield. 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.

5 Conclusion

Our paper contains two key findings. First, we show that a TFP news shock leads to an extended decline in internal and external rates of return, which are key variables in a firm’s decision to hold and accumulate inventories. Second, 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. However, marginal cost measures tend to be insignificant on impact across various specifications.

These findings provide empirical support for the idea that rates of return are an important channel for the positive effects of news shocks on inventory accumulation. We find no evidence of a dominant negative intertemporal substitution effect in terms of marginal costs of production as suggested by theory. That is, 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. In that sense, the two forces are mutually consistent in that they offer an explanation for aggregate inventory behavior in response to news. Furthermore, this is consistent with a demand effect where increased sales drive inventory accumulation. 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 generating sales.

References

- Barsky, R. B. and Sims, E. R. (2011). News shocks and business cycles. *Journal of Monetary Economics*, 58(3):273–289.
- Beaudry, P. and Portier, F. (2004). An exploration into Pigou’s theory of cycles. *Journal of Monetary Economics*, 51(6):1183–1216.
- 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.
- Bouakez, H. and Kemoe, L. (2022). News Shocks, Business Cycles, and the Disinflation Puzzle. *Journal of Money, Credit and Banking*, forthcoming.
- Cascaldi-Garcia, D. and Vukotić, M. (2022). Patent-Based News Shocks. *The Review of Economics and Statistics*, 104(1):51–66.
- Clements, M. P. and Galvão, A. B. (2021). Measuring the effects of expectations shocks. *Journal of Economic Dynamics and Control*, 124(C).
- Copeland, A., Hall, G., and Maccini, L. (2019). Interest rates and the market for new light vehicles. *Journal of Money, Credit and Banking*, 51(5):1137–1168.
- Crouzet, N. and Oh, H. (2016). What do inventories tell us about news-driven business cycles? *Journal of Monetary Economics*, 79:49–66.
- Easton, P. D. (2004). Pe ratios, peg ratios, and estimating the implied expected rate of return on equity capital. *The Accounting Review*, 79(1):73–95.
- Fernald, J. (2014). A quarterly, utilization–adjusted series on total factor productivity. *Working Paper*, (2012-19).
- 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., Gertler, M., and López-Salido, J. D. (2007). Markups, gaps, and the welfare costs of business fluctuations. *The Review of Economics and Statistics*, 89(1):44–59.

- Gebhardt, W. R., Lee, C. M. C., and Swaminathan, B. (2001). Toward an implied cost of capital. *Journal of Accounting Research*, 39(1):135–176.
- Gertler, M. and Karadi, P. (2011). A model of unconventional monetary policy. *Journal of Monetary Economics*, 58(1):17 – 34.
- Görtz, C. and Tsoukalas, J. (2017). News and financial intermediation in aggregate fluctuations. *Review of Economics and Statistics*, 99(3):514–530.
- Görtz, C., Tsoukalas, J., and Zanetti, F. (2021). News Shocks under Financial Frictions. *American Economic Journal: Macroeconomics*, forthcoming.
- Görtz, C., Gunn, C., and Lubik, T. A. (2022). Is there news in inventories? *Journal of Monetary Economics*, 126:87–104.
- Görtz, C. and Tsoukalas, J. D. (2018). Sectoral tfp news shocks. *Economics Letters*, 168:31–36.
- Hou, K., van Dijk, M. A., and Zhang, Y. (2012). The implied cost of capital: A new approach. *Journal of Accounting and Economics*, 53(3):504 – 526.
- Humphreys, B. R., Maccini, L. J., and Schuh, S. (2001). Input and output inventories. *Journal of Monetary Economics*, 47(2):347–375.
- Iacoviello, M., Schiantarelli, F., and Schuh, S. (2011). Input And Output Inventories In General Equilibrium. *International Economic Review*, 52(4):1179–1213.
- Irvine, F. O. and Schuh, S. (2005). Inventory investment and output volatility. *International Journal of Production Economics*, 93-94:75 – 86. Proceedings of the Twelfth International Symposium on Inventories.
- Jaimovich, N. and Rebelo, S. (2009). Can news about the future drive the business cycle? *American Economic Review*, 99(4):1097–1118.
- Jones, C. S. and Tuzel, S. (2013). New orders and asset prices. *Review of Financial Studies*, 26(1):115–157.
- Joseph R. Gordon, M. J. G. (1997). The finite horizon expected return model. *Financial Analysts Journal*, 53(3):52–61.
- Kamber, G., Theodoridis, K., and Thoenissen, C. (2017). News-driven business cycles in small open economies. *Journal of International Economics*, 105(C):77–89.

- Kryvtsov, O. and Midrigan, V. (2013). Inventories, markups, and real rigidities in menu cost models. *The Review of Economic Studies*, 80(1):249–276.
- Kurmann, A. and Sims, E. (2019). Revisions in utilization-adjusted tfp and robust identification of news shocks. *Review of Economics and Statistics*, 103(2):216–235.
- Lettau, M. and Ludvigson, S. (2001). Consumption, aggregate wealth, and expected stock returns. *The Journal of Finance*, 56(3):815–849.
- Lubik, T. A. and Teo, W. L. (2012). Inventories, inflation dynamics and the new keynesian phillips curve. *European Economic Review*, 56(3):327–346.
- Maccini, L. J., Moore, B. J., and Schaller, H. (2004). The interest rate, learning, and inventory investment. *American Economic Review*, 94(5):1303–1327.
- Nekarda, C. J. and Ramey, V. A. (2013). The Cyclical Behavior of the Price-Cost Markup. NBER Working Papers 19099, National Bureau of Economic Research, Inc.
- Ohlson, J. A. and Juettner-Nauroth, B. E. (2005). Expected eps and eps growth as determinants of value. *Review of Accounting Studies*, 10(2-3):349–365.
- 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.
- Vukotic, M. (2019). Sectoral effects of news shocks. *Oxford Bulletin of Economics and Statistics*, 81(2):215–249.

Online Appendix

A Details on the VAR

We identify a news shock using an extension of the Max Share method of Francis et al. (2014). In particular, we specify the following reduced-form VAR of lag length p

$$y_t = A(L)u_t,$$

where y_t is an $n \times 1$ vector and $A(L)$ is a lag polynomial of order p over conformable 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

$$u_t = B_0\varepsilon_t,$$

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

$$y_t = C(L)u_t,$$

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

We can define the h -step ahead forecast error as

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

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

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

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 . $\tilde{B}_0\gamma$ 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. The methodology is an extension of the so-called Max Share method of Francis

et al. (2014), who isolate productivity shocks by maximizing the forecast error variance share of TFP at a long but finite horizon. Our approach assumes that 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

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

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 \tilde{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 problem²²

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

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 . We assume this horizon to be 40 quarters in line with the literature.

B Additional VAR Evidence

In this section, we report four sets of additional evidence from the structural VAR. First, we show forecast error variance decomposition results. Second, we present a robustness check on the result that inventories rise in response to news about higher future TFP. Third,

²²The optimization problem is formulated in terms of choosing γ conditional on any arbitrary orthogonalization \tilde{B}_0 to ensure the resulting identification belongs to the space of possible orthogonalizations of the reduced form.

we show robustness evidence on the results reported in the main body by considering an alternative news shock identification. Fourth, we offer additional evidence on the response of marginal cost measures.

B.1 Forecast Error Variance Decomposition

Figure 1.B reports the forecast error variance decomposition associated with the baseline VAR model shown in Figure 1. It shows that at business cycle frequencies, between 6 and 32 quarters, the TFP news shock explains a substantial share in the variation in macroeconomic aggregates and stock prices.




Figure 1.B: **Forecast Error Variance Decomposition 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.

B.2 Robustness on the Rise in Inventories

This section provides a robustness exercise on the rise in aggregate inventories in response to a TFP news shock. Figure 2.B is the equivalent to Figure 1 in the main body, but uses aggregate business inventories as a measure for stockholdings. It is not clear a priori how inventories should be measured, however any measure based on business inventories is only

available from 1992Q1 which restricts our sample used in this section. Figure 2.B documents, consistent with the evidence provided in the main body, that aggregate inventories rise in response to news about higher future TFP when an alternative definition is used to measure the inventory stock.

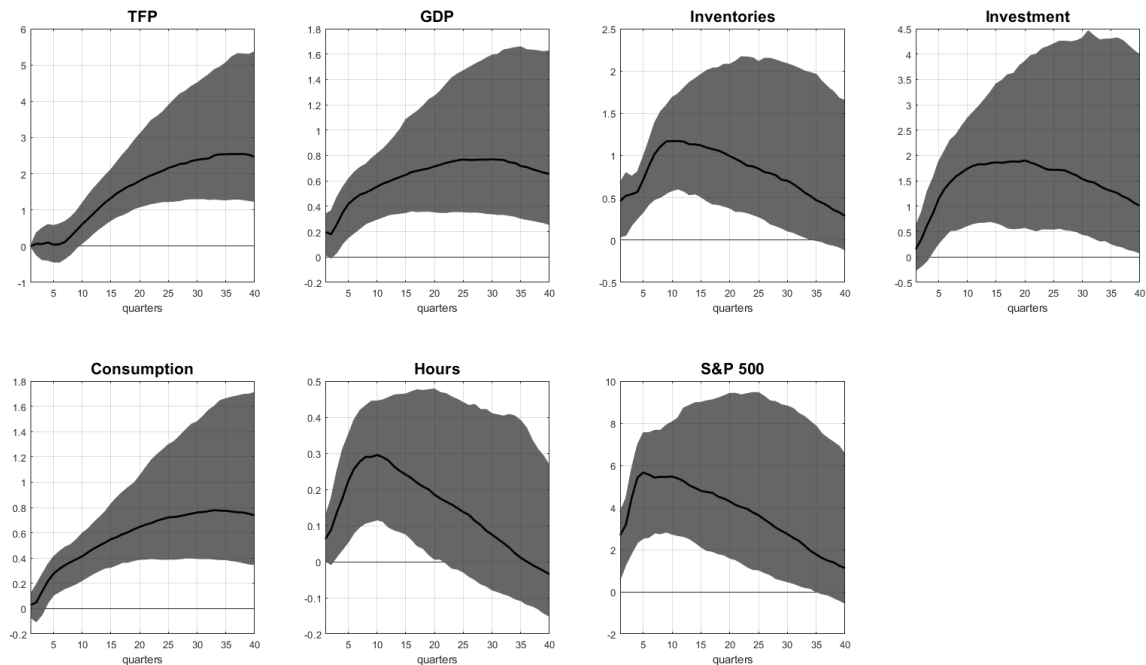


Figure 2.B: **IRF to TFP news shock.** Results based on a seven-variable VAR including business inventories. Sample 1992Q1-2015Q1. The solid line is the median and the shaded gray areas are the 16% and 84% posterior bands generated from the posterior distribution of VAR parameters. The units of the vertical axes are percentage deviations.

B.3 Robustness on the News Shock Identification

Our baseline max-share identification for news shocks is widely used in the literature, yet it may suffer from a certain degree of measurement error. For this reason, we subject our empirical findings above to the alternative identifications recently suggested by Kurmann and Sims (2019). They argue that the TFP measure may 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. Consistent with the choice in their study we set $h = 80$. 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 3.B shows that the results shown in Figure 1 in the main body are robust to using this alternative identification as responses are extremely similar. Also under the identification suggested by Kurmann and Sims (2019), even without the impact restriction, a news shock triggers a broad based expansion in macroeconomic aggregates, including inventories, and a delayed response of TFP. Figures 4.B, 5.B and 6.B show that when using the alternative news shock identification, results are extremely similar to those reported in the main body (they correspond to Figures 4, 5 and 6 in the main body). Notably, we still find that the ECC, DCC and ICC measures decline in response to news about higher future productivity. Also the conclusions considering the responses of the various marginal cost measures are robust to using the alternative shock identification.




Figure 3.B: **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.

B.4 Additional Evidence on Marginal Costs

Figure 7.B 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

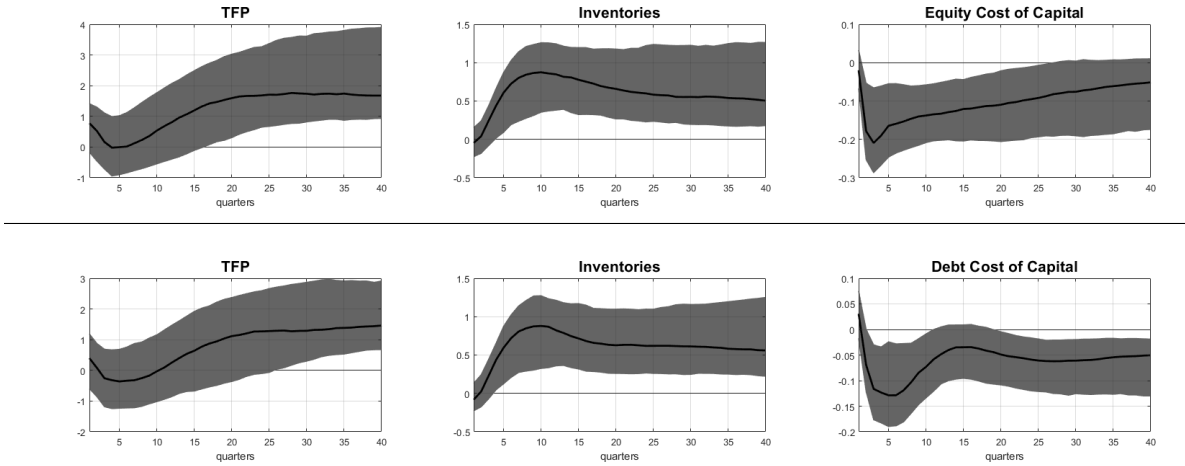


Figure 4.B: IRF of Equity (Debt) Cost of Capital measure to TFP news shock — top (bottom) row. Kurmann-Sims Identification. Selected variables based on two seven-variable VAR systems including TFP, GDP, consumption, hours, inventories, equity (debt) cost of capital, S&P 500. Variables from the respective VAR are shown in the top (bottom) row. 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.B: IRF of Implied Cost of Capital measures to TFP news shock. Kurmann-Sims Identification. 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.

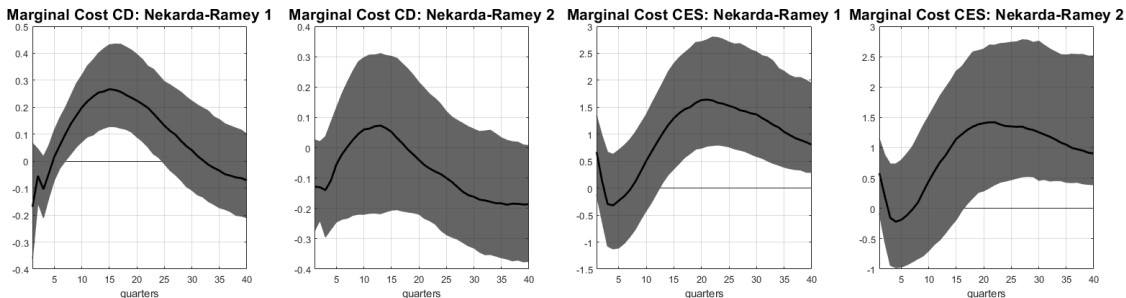


Figure 6.B: **IRF of marginal cost measures to TFP news shock. Kurmann-Sims Identification.** Sample 1985Q1-2015Q2. Each subplot results from a seven-variable VAR including TFP, GDP, consumption, hours, inventories, one particular measure for marginal cost, S&P 500. Marginal cost construction is based 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.

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 7.B 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 7.B 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 B.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 6 and 7.B 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




Figure 7.B: **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.

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 B.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.

C 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 C.1 provides details on the

construction of different implied cost of capital measures. Section C.2 documents the underlying dataset construction.

C.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) and Fama and French (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 Joseph R. Gordon (1997).²³ 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 C.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:

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

²³See e.g. Ashbraugh-Skaife et al. (2009), Hail and Leuz (2009) and Chava and Purnanandam (2010).

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.

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

$$MVAL_{i,t} = \frac{E_t [EA_{i,t+2}] + ICC_{i,t} \times E_t [D_{i,t+1}] - E_t [EA_{i,t+1}]}{ICC_{i,t}^2}, \quad (\text{B.3})$$

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:

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

with the short-term growth rate given by:

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

as in Gode and Mohanram (2003). γ_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:

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

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.

C.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.

D Constructing Marginal Cost Measures

We follow the template in Nekarda and Ramey (2013) of constructing several measures for marginal costs. This section provides the detail on their construction.

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

$$MC_t = \frac{W_t/P_t}{F_h(K_t, H_t)}, \quad (4)$$

where W/P is the real wage and $F_h(K, H)$ 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:

$$\log(MC_t) \approx \log(s_t), \quad (5)$$

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:

$$\log(MC_t) \approx \log(s_t) - \left(\frac{1}{\sigma} - 1\right) [\log(Y_t) - \log(Z_t H_t)]. \quad (6)$$

Technology is denoted by Z_t , σ is the elasticity of substitution between capital, K_t , and labor, H_t , and Y_t is output in value added terms.²⁴

We construct marginal cost measures based on the two technology specifications 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 Bureau of Labor Statistics (BLS). As a measure for technology, we use John Fernald’s utilization-adjusted TFP series, and we set σ at a baseline value of 0.5 in line with Nekarda and Ramey (2013).²⁵ 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.

References in the Appendix

Ashbraugh-Skaife, H., Collins, D. W., Kiney Jr, W. R., and Lafond, R. (2009). The effect of sox internal control deficiencies on firm risk and cost of equity. *Journal of Accounting Research*, 47(1):1–43.

²⁴If output is measured as gross output, the same expression obtains as long as the production function is 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 .

²⁵Online appendix B.4 contains an extensive robustness analysis with respect to this parameter.

- Barsky, R. B. and Sims, E. R. (2011). News shocks and business cycles. *Journal of Monetary Economics*, 58(3):273–289.
- Chava, S. and Purnanandam, A. (2010). Is default risk negatively related to stock returns? *The Review of Financial Studies*, 23(6):2523–2559.
- Chirinko, R. S. (2008). σ : The long and short of it. *Journal of Macroeconomics*, 30(2):671 – 686.
- Easton, P. D. (2004). Pe ratios, peg ratios, and estimating the implied expected rate of return on equity capital. *The Accounting Review*, 79(1):73–95.
- Fama, E. F. and French, K. R. (2000). Forecasting profitability and earnings. *The Journal of Business*, 73(2):161–175.
- Fama, E. F. and French, K. R. (2006). Profitability, investment and average returns. *Journal of Financial Economics*, 82(3):491–518.
- 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., Gertler, M., and López-Salido, J. D. (2007). Markups, gaps, and the welfare costs of business fluctuations. *The Review of Economics and Statistics*, 89(1):44–59.
- Gebhardt, W. R., Lee, C. M. C., and Swaminathan, B. (2001). Toward an implied cost of capital. *Journal of Accounting Research*, 39(1):135–176.
- Gode, D. and Mohanram, P. (2003). Inferring the cost of capital using the ohlson–juettner model. *Review of Accounting Studies*, 8:399—431.
- 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(3):428–454.
- Hou, K., van Dijk, M. A., and Zhang, Y. (2012). The implied cost of capital: A new approach. *Journal of Accounting and Economics*, 53(3):504 – 526.
- Hribar, P. and Collins, D. W. (2002). Errors in estimating accruals: Implications for empirical research. *Journal of Accounting Research*, 40(1):105–134.
- Joseph R. Gordon, M. J. G. (1997). The finite horizon expected return model. *Financial Analysts Journal*, 53(3):52–61.

- Nekarda, C. J. and Ramey, V. A. (2013). The Cyclical Behavior of the Price-Cost Markup. NBER Working Papers 19099, National Bureau of Economic Research, Inc.
- Ohlson, J. A. and Juettner-Nauroth, B. E. (2005). Expected eps and eps growth as determinantsof value. *Review of Accounting Studies*, 10(2-3):349–365.