What do inventories tell us about news-driven business cycles?*

Nicolas Crouzet and Hyunseung Oh†

Kellogg School of Management, Northwestern University and Vanderbilt University

January 19th, 2015

Abstract

There is widespread disagreement on the quantitative contribution of news shocks to business-cycle fluctuations. This paper provides a simple identifying restriction, based on inventory dynamics, that tightly pins down the contribution of news shocks to business-cycle volatility. We show that finished-good inventories must fall when there is an increase in consumption and investment induced by news shocks. A structural VAR with these sign restrictions indicates that news shocks account for at most 20 percent of output volatility. Since inventories comove positively with consumption and investment in the data, shocks that generate negative comovement cannot account for the bulk of fluctuations.

*We thank Ricardo Reis for his support throughout this project, and Stephanie Schmitt-Grohé and Jon Steinssón for their advice. We also thank Dario Caldara, Yongsung Chang, James Kahn, Jennifer La'O, Neil Mehrotra, Virgiliu Midrigan, Tommaso Monacelli, Miguel Morin, Emi Nakamura, Serena Ng, John Roberts, Frank Schorfheide, Dmitriy Sergeyev, Jae W. Sim, Martin Uribe, Michael Woodford, anonymous referees, and numerous seminar participants at Columbia University, Federal Reserve Board, Expectations in Dynamic Macroeconomic Models, Midwest Macro Meetings and Society for Economic Dynamics for helpful comments. This paper supercedes the previous version titled "Can news shocks account for the business-cycle dynamics of inventories?". All errors are ours.

†Contact: n-crouzet@kellogg.northwestern.edu and hyunseung.oh@vanderbilt.edu (corresponding author)

1 Introduction

The sources of business cycles are an enduring subject of debate among macroeconomists. Recently, the literature has focused on news shocks — shocks that change agents' expectations about future economic fundamentals, without affecting current fundamentals — as a potential driving force of aggregate fluctuations. Starting with Beaudry and Portier (2006), this literature has argued that news shocks may provide a good account of expansions and recessions, stressing episodes such as the US and Asian investment booms and busts of the late 1990s as examples.

In the news view of business cycles, booms and busts arise as a result of changes in expectations (Beaudry and Portier, 2013). For example, when productivity is expected to increase in the future, investment increases today, in order to build up the capital stock and take advantage of lower future marginal costs. The boom in investment raises wages and hours worked, and the additional income leads to a consumption boom. Recent work has established the theoretical conditions under which business cycle models can successfully capture this mechanism. A prominent example is Jaimovich and Rebelo (2009), who show that, in a model with investment adjustment costs, variable capacity utilization, and weak wealth effects on hours worked, an expected rise in the marginal product of capital indeed leads to a boom in consumption, investment and output today.[1] With these elements, business cycle models can generate positive comovement of consumption, investment and output in response to shifts in expectations, making news shocks a plausible source of business-cycle fluctuations.

However, the literature has yet to come to a consensus regarding the quantitative importance of this type of shock. While some estimates suggest that they account for as much as 60 percent of output volatility, other equally plausible methods obtain numbers as low as 10 percent.[2] The goal of this paper is to use a new theoretical insight to pin down more precisely the contribution of news shocks to business-cycle volatility. Our innovation is to focus on a variable that is highly informative about news shocks, but so far has been neglected in the literature: investment in finished-good inventories.

[1] Adding variable capacity and weak wealth effects on hours worked allows output to rise on impact and satisfy current demand, while investment adjustment costs force firms to smooth the desired increase in the stock of capital over time and start investing today.

[2] For example, Beaudry and Portier (2006) and Schmitt-Grohé and Uribe (2012) estimate that the contribution of news shocks to the business-cycle volatility of output is above 50 percent, while Barsky and Sims (2011) and Khan and Tsoukalas (2012) find much smaller numbers. This will be discussed in detail in later sections.

We focus on investment in finished-good inventories for two reasons. First, there is empirical evidence that finished-good inventories are a forward-looking variable that reflect changes in expectations about future economic conditions. For instance, Kesavan, Gaur, and Raman (2010) find that finished-good inventory data help improve forecasts of future sales. Second, finished-good inventories should react differently to changes that happen today and to changes that are expected to happen in the future. Temporary changes in productivity provide a straightforward illustration of this intuition. When productivity increases today, current income increases, so that sales rise. Additionally, firms bunch production today, in order to make the most out of the productivity increase. As a result, finished-good inventories also rise. Thus, when productivity changes today, sales and inventories should comove *positively*.[3] When productivity is expected to increase tomorrow, permanent income increases, which also leads to a rise in sales. However, since firms expect future production to be cheaper than current production, they satisfy this increase in sales not by producing more, but by depleting inventories. Thus, when productivity is expected to change tomorrow, sales and inventories should comove *negatively*. The sign of the comovement between sales and inventories should therefore help us distinguish between current and future expected changes in productivity.

In section 2, we start our analysis by introducing inventories into a news-driven business-cycle model. The inventory block of our model follows Bils and Kahn (2000). In section 3, we establish analytically that, in this model, good news about the future indeed lead to a boom in the components of private sales (consumption and investment), but a fall in inventories. The key mechanism that accounts for this result is similar to the intuition described above: news shocks cause agents to want to substitute production over time, and they are able to do so easily through inventory investment. With good news about the future, marginal cost is expected to be lower in the future than today. Optimal inventory investment behavior then dictates that firms should delay production, and satisfy current demand by drawing down on existing inventories.

In section 4, we show that this result survives a number of extensions of the baseline model. First, we show that the fall in inventories after a positive news shock is deep and protracted. Second, we establish that this result holds for other types of news, especially news about demand. Third, we show that the result survives the introduction of various types of adjustment costs, in particular

[3] From now on, we will use the term "inventories" to indicate finished-good inventories when there is no confusion.

adjustment costs to inventories and price adjustment costs. In section 5, we show that the result also holds in models with other microfoundations for inventory holding than those of Bils and Kahn (2000), such as the stockout-avoidance model of Kahn (1992), Kryvtsov and Midrigan (2013) and Wen (2011) or the (S,s) inventory model of Khan and Thomas (2007b). The strong intertemporal substitution channel with regards to news shocks is thus a general feature of inventory models.

Since structural models indicate that inventories and sales should comove negatively in response to news shocks, we use this prediction to assess their quantitative contribution to business-cycle volatility. In section 6, we describe a simple empirical strategy, based on a structural VAR with sign restrictions, that identifies shocks which generate negative comovement between inventories and sales. We find that these shocks explain less than 20 percent of the forecast error variance of output at business-cycle frequencies. The reason we get a small and precise estimate is that, in the data, inventories and sales comove positively at those frequencies. Any shock that generates negative comovement between inventories and sales must have limited importance over the business cycle.

A cautious interpretation of this result is that it provides an upper bound to the contribution of news shocks to business-cycle volatility, since surprise innovations to other fundamentals could also be generating the same pattern of comovement.[4] However, section 7 shows that this small upper bound on the contribution of news shocks is consistent with results obtained using an estimated DSGE model. We estimate the model of Schmitt-Grohé and Uribe (2012) augmented with an inventory block, and using inventory data. This model includes a range of sources of news aside from TFP and the labor wedge, and stationary as well as non-stationary shocks. Our median posterior estimate of the contribution of news shocks to the volatility of output is 17 percent. Section 8 concludes.

Our work relates to a number of papers that examine the behavior of investment in response to news shocks. Jaimovich and Rebelo (2009), Christiano, Ilut, Motto, and Rostagno (2008), as well as Schmitt-Grohé and Uribe (2012) document the importance of investment adjustment costs for news shocks to generate an immediate boom in investment and output. Inventory investment

[4]In section 4, we show that surprise shocks to TFP and the labor wedge both generate positive comovement between inventories and sales in the structural models we analyze, in line with our discussion above. To the extent that the shocks we identify have a non-news component, it must therefore reflect surprise changes in other fundamentals than TFP and the labor wedge.

has been mostly neglected in this literature. One exception is Vukotic (2013), where inventories are introduced as a factor of production in the durable sector. Our approach is quite different from hers, since we examine inventories that are stored as finished goods. In her model, inventories play a role similar to investment. Our contribution to the news literature is to illustrate a new channel through which news shocks operate – the intertemporal substitution channel – by focusing on the investment behavior of finished-good inventories, a form of investment that is distinct from investment in fixed capital.

Our work also relates to the recent literature on inventories that matches their stylized business-cycle behavior using micro-founded models. These models differ in the reason for which firms hold inventories, aside from shifting production over time. One branch of the literature argues that inventories exist in order to facilitate sales, either by their use for displaying and advertising purposes (Bils and Kahn, 2000), or because they can be used as a buffer against stockouts (Wen, 2011; Kryvtsov and Midrigan, 2013). Another branch of the literature argues that inventories exist due to bunching behavior induced by fixed ordering costs (Fisher and Hornstein, 2000; Khan and Thomas, 2007b). Since our focus is on finished-good inventories, we follow more closely the former approach.[5] Nevertheless, the key theoretical result of the paper also applies to the latter approach, since a common feature of all these models is that inventories are producers' means of substituting production over time. Our contribution to this literature is to highlight this common mechanism, and draw its implications for business-cycle comovement in the context of news shocks.

Finally, our empirical approach is based on the literature on sign restrictions in a VAR framework. Sign restrictions have been used to identify monetary policy shocks (Faust, 1998; Uhlig, 2005), fiscal policy shocks (Mountford and Uhlig, 2008; Caldara and Kamps, 2012), and also news shocks (Beaudry, Nam, and Wang, 2011).

2 A finished-good inventory model

In this section, we describe a general equilibrium model of inventory dynamics based on the work of Pindyck (1994), Bils and Kahn (2000), and Jung and Yun (2006).

The key feature of the so-called "stock-elastic" demand model of this section is the assumption

[5] See Khan and Thomas (2007a) for a comparison of the two approaches.

that sales of a firm are elastic to the amount of goods available for sale, which we term "on-shelf goods." This assumption finds empirical support for many categories of goods, as documented by Pindyck (1994) or Copeland, Dunn, and Hall (2011). The positive elasticity of sales to on-shelf goods captures the idea that with more on-shelf goods, customers are more likely to find a good match and purchase the product.

The key results we obtain are common to structural models of inventories with alternative microfoundations for inventory use, such as stockout avoidance or (S,s) models. However, the tractability of this framework helps highlight the key mechanisms at play. We come back to alternative models in later sections.

2.1 Description of the stock-elastic demand model

The economy consists of a representative household and monopolistically competitive firms. The output of the firms are storable goods, of which they keep a positive inventory.[6] We start with the household problem.

Household problem A representative household maximizes the following expected sum of discounted utility,

$$E_0\left[\sum_{t=0}^{\infty}\beta^t U(c_t, n_t; \psi_t)\right], \tag{1}$$

where c_t is the consumption of the final good, n_t denotes the supply of labor services, and ψ_t is a exogenous variable that introduces a wedge between the marginal rate of substitution between consumption and leisure, and the real wage, and which we call a "labor wedge" shock. We assume that the household's period utility function takes the form proposed by Greenwood, Hercowitz, and Huffman (1988, henceforth GHH):

$$U(c,n;\psi) = \frac{1}{1-\sigma}\left(c - \psi\frac{n^{1+\xi^{-1}}}{1+\xi^{-1}}\right)^{1-\sigma},$$

[6]For clarity, we discuss our main results in a one-sector economy with storable final goods. However, a sector-specific news shock would trigger intertemporal substitution within that sector, as it does in our baseline model. The question of the transmission of such shocks to other sectors is beyond the scope of this paper; see, e.g., Beaudry et al. (2011) for a discussion of the inter-sectoral transmission of news.

where ξ is the Frisch elasticity of labor supply and σ denotes the inverse of the elasticity of the household's intertemporal substitution. This preference specification has been widely used in the literature on news shocks, and it implies zero wealth effects on labor supply.[7]

The household's maximization problem is subject to the following constraints:

$$\int_0^1 p_t(j)s_t(j)dj + \mathbb{E}_t\left[Q_{t,t+1}B_{t+1}\right] \leq W_t n_t + R_t k_t + \int_0^1 \pi_t(j)dj + B_t, \tag{2}$$

$$k_{t+1} = i_t\left[1 - \phi\left(\frac{i_t}{i_{t-1}}\right)\right] + (1-\delta_k)k_t, \tag{3}$$

$$c_t + i_t \leq x_t, \tag{4}$$

$$x_t = \left(\int_0^1 v_t(j)^{\frac{1}{\theta}} s_t(j)^{\frac{\theta-1}{\theta}} dj\right)^{\frac{\theta}{\theta-1}}. \tag{5}$$

Equation (2) is the household budget constraint. The household earns income each period by providing labor n_t at a given nominal wage W_t, lending capital k_t at a rate R_t, claiming the nominal profit $\pi_t(j)$ from each firm $j \in [0,1]$, and receiving nominal bond payments B_t. It spends its income in purchases of each variety in the amount $s_t(j)$ at a price $p_t(j)$, and in purchases of the state-contingent one-period bonds B_{t+1}. The probability-adjusted price of each of these nominal bonds is $Q_{t,t+1}$, for each state in period $t+1$.

Equation (3) is the law of motion of capital with adjustment costs to investment. The adjustment cost function $\phi(\cdot)$ is twice-differentiable, with $\phi(1) = \phi'(1) = 0$ and $\phi''(1) > 0$. When firms' desired future level of capital is high, this type of adjustment cost forces them to smooth out the desired increase over time and start investing today.

Equation (4) states that the household's consumption and investment cannot exceed its total absorption of final goods, x_t, which is constructed by aggregating their purchase of intermediate goods $\{s_t(j)\}_{j\in[0,1]}$. The aggregation of the intermediate goods $\{s_t(j)\}_{j\in[0,1]}$ into x_t is given by a Dixit-Stiglitz type aggregator (5) where $v_t(j)$ is the taste-shifter for each product j and θ is the elasticity of substitution across intermediate goods. It follows from expenditure minimization that

[7] As emphasized by Jaimovich and Rebelo (2009), news shocks are most likely to generate business-cycle comovement when wealth effects on labor supply are small. The GHH preference specification is thus the "best case" scenario in terms of the ability of the model to generate comovement in response to news.

the demand function for each good and the aggregate price level take the following forms:

$$s_t(j) = v_t(j) \left(\frac{p_t(j)}{P_t} \right)^{-\theta} x_t, \qquad P_t = \left(\int_0^1 v_t(j) p_t(j)^{1-\theta} dj \right)^{\frac{1}{1-\theta}}.$$

In the stock-elastic demand model, the taste shifter for variety j is assumed to depend on the amount of goods on shelf proposed by the firm producing variety j, $a_t(j)$, in the following fashion:

$$v_t(j) = \left(\frac{a_t(j)}{a_t} \right)^{\zeta}, \tag{6}$$

where the normalization by a_t, defined as the the economy-wide average of on-shelf goods, ensures that the mean of $\nu_t(j)$ across goods is equal to 1. The parameter $\zeta > 0$ controls the degree of the shift in taste due to the relative amount of goods on-shelf.

Finally, the household is given an initial level of capital k_0 and bonds B_0, and its optimization problem is subject to no-Ponzi conditions for both capital and stage-contingent bond holdings.

Firm problem Each monopolistically competitive firm $j \in [0, 1]$ maximizes the expected discounted sum of profits

$$E_0 \left[\sum_{t=0}^{\infty} Q_{0,t} \pi_t(j) \right], \tag{7}$$

where

$$\pi_t(j) = p_t(j) s_t(j) - W_t n_t(j) - R_t k_t(j). \tag{8}$$

Note that the profit in each period is the revenue from sales net of the cost from hiring labor $n_t(j)$ and renting capital $k_t(j)$ at their respective prices W_t and R_t. The term $Q_{0,t}$ is the discount factor of between period 0 and t, so that $Q_{0,t} = \prod_{T=0}^{t-1} Q_{T,T+1}$. This discount factor is consistent with households being the final owners of firms. The firm faces the following constraints:

$$a_t(j) = (1-\delta_i)inv_{t-1}(j) + y_t(j), \tag{9}$$

$$inv_t(j) = a_t(j) - s_t(j), \tag{10}$$

$$y_t(j) = z_t k_t^{1-\alpha}(j) n_t^{\alpha}(j), \tag{11}$$

$$s_t(j) = \left(\frac{a_t(j)}{a_t}\right)^{\zeta} \left(\frac{p_t(j)}{P_t}\right)^{-\theta} x_t. \tag{12}$$

Equation (9) is the inventory stock accumulation equation. The stock (on-shelf goods) of the firm, $a_t(j)$, consists of the undepreciated stock of inventories from the previous period $(1-\delta_i)inv_{t-1}(j)$ and of current production $y_t(j)$. The parameter δ_i denotes the depreciation rate of inventories. Equation (10) states that on-shelf goods that are unsold are accounted as inventories.[8] Equation (11) is the production function. Firms use a constant returns to scale production function, with capital and labor as inputs. The variable z_t represents total factor productivity and is exogenous. Finally, monopolistically competitive firms face the demand function (12) stemming from the household problem.

Market clearing Labor and capital markets clear, and net bond holdings is zero:

$$n_t = \int_0^1 n_t(j)dj, \tag{13}$$

$$k_t = \int_0^1 k_t(j)dj, \tag{14}$$

$$B_t = 0. \tag{15}$$

Sales of goods for each variety j also clears, as is implicit in the expression of the demand function (12). The average level of on-shelf goods in the economy a_t is defined by:

$$a_t = \int_0^1 a_t(j)dj. \tag{16}$$

Since the price of the consumption good P_t is a numeraire, in what follows we will use the lower-case variables $w_t = W_t/P_t$ for real wage, $r_t = R_t/P_t$ for real rental rate of capital, $b_t = B_t/P_t$ for real

[8] In the data, this is recorded as the end-of-period inventory stock in each period.

bond holdings, $q_{t,t+1} = Q_{t,t+1}P_{t+1}/P_t$ for the real stochastic discount factor, and $\tilde{p}_t(j) = p_t(j)/P_t$ for the relative price of good j.

2.2 Equilibrium

A market equilibrium of this economy is defined as follows.[9]

Definition 1 (Market equilibrium of the stock-elastic demand model) *A market equilibrium of the stock-elastic demand model is a set of stochastic processes for aggregate variables*

$$c_t, n_t, k_{t+1}, i_t, b_{t+1}, x_t, a_t, w_t, r_t, q_{t,t+1},$$

and firm-level variables

$$\{a_t(j)\}, \{n_t(j)\}, \{k_t(j)\}, \{v_t(j)\}, \{s_t(j)\}, \{y_t(j)\}, \{inv_t(j)\}, \{\tilde{p}_t(j)\},$$

such that, given the exogenous stochastic processes z_t, ψ_t, as well as initial conditions k_0, b_0 and $\{inv_{-1}(j)\}$:

- *households maximize (1) subject to (2) - (6) and two no-Ponzi conditions,*
- *each firm $j \in [0,1]$ maximizes (7) subject to (8) - (12),*
- *markets clear according to (13) - (16).*

The two exogenous processes in our economy are total factor productivity z_t and the labor wedge ψ_t. We define news shocks to productivity in period t as any innovation to $\{z_t\}_{t\geq 0}$ such that $z_t = \mathbb{E}_{t-1}z_t$ but $\mathbb{E}_t z_{t+s} \neq \mathbb{E}_{t-1}z_{t+s}$ for some $s > 0$. News shocks to the labor wedge are defined similarly. News shocks to these two exogenous processes are the primary contributors to aggregate fluctuations in Schmitt-Grohé and Uribe (2012).[10] Surprise shocks to productivity, on the other hand, are defined as innovations to $\{z_t\}_{t\geq 0}$ such that: $\mathbb{E}_{t-1}z_t \neq z_t$, and they are similarly defined for the labor wedge.

[9]Since this is an economy with flexible prices, only relative prices and the real allocation are determined in equilibrium.

[10]Other types of shocks will be discussed in later sections.

2.3 The optimal choice of inventories

The full set of equilibrium conditions are provided in appendix A. We study a symmetric equilibrium: $a_t(j) = a_t$, $s_t(j) = s_t$, $inv_t(j) = inv_t$, $y_t(j) = y_t$, and $p_t(j) = p_t$ for all j. We first discuss the optimal stock choice of firms.

In the market equilibrium, marginal cost is the real wage divided by the marginal product of labor:

$$mc_t = \frac{w_t}{\alpha z_t (k_t/n_t)^{1-\alpha}}. \tag{17}$$

Using this, the optimal stock choice of firms is governed by the equation:

$$mc_t = \frac{\partial s_t}{\partial a_t} + \left(1 - \frac{\partial s_t}{\partial a_t}\right) \mathbb{E}_t[q_{t,t+1}(1-\delta_i) mc_{t+1}]. \tag{18}$$

The left hand side of this equation represents the cost of adding an extra unit of goods to the stock of goods on sale, a_t, which equals the current marginal cost of production. The right hand side represents the two benefits of adding this extra unit. First, by producing and stocking an extra unit, the firm is able generate an additional fraction $(\partial s_t/\partial a_t)$ of sales. Second, since some of the extra goods stocked will not be sold and will be stored as inventories for the next period, future production costs are reduced.

It is important to notice that at the nonstochastic steady state of the economy, the stock of inventories is positive. Since the real interest rate and the inventory depreciation rate are both positive at the steady state, holding inventories is costly. However, consistent with the first term on the right hand side of (18), there is a convenience yield associated with holding a positive amount of inventories in each period. In the model, the convenience yield is the additional sales created by holding a positive level of stock. Therefore, there will be a positive amount of inventories in steady state, despite the intertemporal costs that holding inventories implies.

Rearranging, (18) can be expressed as:

$$\frac{\partial s_t}{\partial a_t} = \frac{\gamma_t^{-1} - 1}{\mu_t - 1}, \tag{19}$$

where:

$$\mu_t \equiv \frac{1}{(1-\delta_i)\mathbb{E}_t[q_{t,t+1}mc_{t+1}]}, \quad \gamma_t \equiv (1-\delta_i)\mathbb{E}_t\left[\frac{q_{t,t+1}mc_{t+1}}{mc_t}\right].$$

The variable μ_t is the markup of price over expected discounted marginal cost. This is the relevant markup concept in an economy where firms produce to stock: indeed, the true cost of sales is not current but future marginal cost, since selling an extra unit reduces tomorrow's stock of goods. The variable γ_t is the expected discounted growth rate of marginal cost, which summarizes the firm's opportunity cost of producing today. The optimal stocking behavior of a firm balances these 3 margins: markup, discounted growth rate of marginal cost, and additional sales generated by extra inventory holdings.

In equilibrium, the optimal choice of inventories can be approximated up to first order as:[11]

$$\widehat{inv}_t = \hat{s}_t + \eta\hat{\gamma}_t,$$

where hatted variables represent log-deviations from its steady-state. This condition states that two factors determine the dynamics of inventories.

First, $\hat{s}_t$ represents the demand channel: firms in this economy build up their inventories when sales are high. For example, when there is an increase in aggregate demand, firms make the most out of it by stocking more goods on shelf to generate additional sales. However, since the additional unit on stock will not lead to a full amount of realized sales, (end-of-period) inventories also increase.

Second, $\eta\hat{\gamma}_t$ represents the intertemporal substitution channel, where $\eta > 0$ is a combination of structural parameters that will be specified in proposition 1. Intuitively, η captures the degree of intertemporal substitution of production in this economy. For example, when there is an increase in future expected discounted marginal cost relative to current marginal cost, then $\hat{\gamma}_t$ is positive and firms will increase their inventories. This happens because firms realize that it is cheaper to produce today than in the future and they now bunch their production today and store more inventories. When the value of η is large, then the degree of intertemporal substitution is such that even a small change in the perception of the marginal cost will result in a large change in inventories.

Thus, in this model, changes in inventories are governed by the relative strength of the demand channel and the intertemporal substitution channel.

[11] This equation is derived by combining (10), (19) and the optimal pricing condition $\hat{\mu}_t = 0$.

3 The impact effect of news shocks

We now turn to studying the effect of news shocks in this model economy. In this section, we focus on impact responses. We derive analytical conditions under which news shocks result in positive comovement on impact between sales and inventories, assess whether those conditions are likely to hold in reasonable calibrations of the model, and inspect the mechanisms underpinning the result.

3.1 A log-linearized framework

We analyze a first-order log-linear approximation of the model around its steady-state. The following framework summarizes the equilibrium conditions needed for the purpose of our subsequent analysis.

Proposition 1 (Stock-elastic demand model) *On impact and without surprise shocks, so that $\hat{z}_t = 0$ and $\hat{\psi}_t = 0$, the following equations represent the log-linearized market equilibrium of definition 1:*

$$\widehat{mc}_t = \omega \hat{y}_t, \tag{20}$$

$$\kappa \hat{y}_t = \hat{s}_t + \frac{\kappa - 1}{\delta_i}[\widehat{inv}_t - (1-\delta_i)\widehat{inv}_{t-1}], \tag{21}$$

$$\widehat{inv}_t = \hat{s}_t + \tau \hat{\mu}_t + \eta \hat{\gamma}_t, \tag{22}$$

$$\hat{\mu}_t = 0, \tag{23}$$

$$\hat{\mu}_t + \hat{\gamma}_t + \widehat{mc}_t = 0. \tag{24}$$

The mapping from the structural model parameters to the parameters of the reduced-form equations is given by:

$$\omega = \frac{1 + (1-\alpha)\xi}{\alpha \xi}, \tag{25}$$

$$\kappa = 1 + \delta_i IS, \tag{26}$$

$$\eta = \frac{1 + IS}{IS} \frac{1}{1 - \beta(1-\delta_i)}, \tag{27}$$

$$\tau = \theta \frac{1 + IS}{IS},$$

where IS is the steady-state inventory-sales ratio, given by

$$IS = \frac{(\theta - 1)(1 - \beta(1 - \delta_i))}{\zeta\beta(1 - \delta_i) - (\theta - 1)(1 - \beta(1 - \delta_i))}.$$

Equation (20) relates marginal cost to output, which is derived by combining the labor supply and demand conditions, and the production function. Importantly, this equation is not connected to the introduction of inventories in our model. With $\omega > 0$, the equation states that real marginal cost increases with output. The parameter ω is the elasticity of marginal cost with respect to output, keeping constant total factor productivity. In other words, ω represents the degree of decreasing returns in the economy due to predetermined capital in the short run (represented by α) and the disutility of labor supply (represented by ξ). The value of ω itself has been at the center of debate in the monetary economics literature, and there is disagreement about its magnitude. Woodford (2003) contrasts two values of ω: 1.25, from Chari, Kehoe, and McGrattan (2000), and 0.47, from Rotemberg and Woodford (1997). Moreover, Dotsey and King (2006) suggest a lower bound of 0.33 for ω. A conservative lower bound for ω is thus:

$$\omega \geq 0.3.$$

Equation (21) is the law of motion for the stock of inventories, obtained from combining equations (9) and (10). This law of motion states that output should equal sales plus inventory investment. The parameter κ in (21) denotes the steady-state output to sales ratio. In NIPA, the time series average of inventory investment over output is around 0.5 percent, so that:

$$\kappa = 1.005.$$

Equations (22) and (23) are the optimal stocking and pricing conditions, respectively. Combining these two equations, we see that inventories are determined by the demand channel ($\hat{s}_t$) and the intertemporal substitution channel ($\eta\hat{\gamma}_t$), as we discussed in section 2. Here we focus on the numerical value of η, the degree of intertemporal substitution in production. Equation (27) indicates that a lower bound for η is $(1 - \beta(1 - \delta_i))^{-1}$. The lower bound depends on two parameters, β and δ_i. First, the household discount factor β governs the opportunity cost of holding inventories.

In the limiting case where $\beta = 1$, there is no opportunity cost of holding inventories since the real interest rate $1/\beta - 1$ is 0. Second, the depreciation rate of inventories δ_i represent the physical cost of holding inventories. Therefore, the value $1 - \beta(1 - \delta_i)$ represents the overall intertemporal cost of adjusting inventories. In the limiting case where both the opportunity cost and the physical cost of inventories are zero, then the lower bound of η is infinity. At quarterly frequency, we set $\beta = 0.99$, which is standard. For δ_i, the logistics literature estimates the carrying cost to be around 12–15 percent in annual terms.[12] With a rather high value of $\delta_i = 0.04$, the lower bound is thus:

$$\eta > 20.$$

Lastly, equation (24) follows from the definition of μ_t and γ_t in section 2.

3.2 The impact response of inventories to good news about the future

Given sales $\hat{s}_t$, equations (20) - (24) relate the following four variables: output $\hat{y}_t$, inventories $\hat{inv}_t$, the discounted growth rate of marginal cost $\hat{\gamma}_t$, and markups $\hat{\mu}_t$. We adopt the following definition of a news shock in the context of this reduced-form framework: a news shock is a shock that has no impact on current fundamentals ($\hat{z}_t = 0$ and $\hat{\psi}_t = 0$), but which changes expectations about future fundamentals ($\mathbb{E}_t \hat{z}_{t+k} \neq 0$ or $\mathbb{E}_t \hat{\psi}_{t+k} \neq 0$ for some $k > 0$).

Proposition 2 (The impact response of inventories to a good news about the future) *When news arrive, inventories and sales positively comove on impact if and only if:*

$$\eta < \frac{\kappa}{\omega}.$$

This proposition indicates that the positive comovement between inventories and sales only depends on the three parameters discussed above, κ, ω and η. With $\kappa = 1.005$, the two parameters ω and η need to be sufficiently small for there to be positive comovement between inventories and sales on impact.. Following our previous discussion on numerical values, a conservative upper bound on κ/ω is 3.3. However, given that our lower bound of η with a large carrying cost of inventories is

[12] The overall carrying cost suggested in the literature is on average 25 percent in annual terms (Stock and Lambert, 2001). However, these include interest payments and clerical costs of managing inventories, which are also separately documented in Stock and Lambert (2001). Excluding those costs gives our numbers.

still 20, the condition of proposition 2 is not met and in fact fails by an order of magnitude. Thus, our framework indicates that following the arrival of good news about the future, the boom in sales associated to a news shock is accompanied by a fall in inventories. In other words, there is negative comovement between inventories and sales in response to news shocks.

3.3 Discussion

The numerical discussion of proposition 2 concludes that inventories must fall when good news about the future generates a current boom in sales. The two key parameters that drive this result are ω and η.

First, when ω is small, then a sales boom will also correspond to an inventory boom. This is because with a small ω, marginal cost barely responds to changes in the level of output produced by firms. Therefore, intertemporal substitution in production is less attractive for firms. In this case, inventories are mostly used to affect demand, and with a sufficient increase in demand, firms will optimally accumulate inventories.

Second, when η is small, the intertemporal substitution channel itself becomes weak. This is the case when the firm faces large costs in storing goods for the future. When the interest rate is high or the depreciation of inventories are high, then it is costly for firms to hold inventories. In this economy, even though marginal cost may respond sensitively to production, firms will be less willing to smooth this out by adjusting inventories. Therefore a sufficient increase in demand will also lead to an accumulation of inventories.

To be more precise on this connection between η and the cost of storing goods, recall that the lower bound of η is negatively related to the intertemporal cost of adjusting inventories, $1-\beta(1-\delta_i)$. In fact, we also find that the value of η itself is negatively related with the intertemporal cost. In figure 1, we fix the other structural parameters and change the value of $1-\beta(1-\delta_i)$ to show this relation.[13] In the extreme case with zero intertemporal cost of adjusting inventories, we see that the degree of intertemporal substitution, η, reaches infinity. With higher intertemporal cost imposed, the value of η becomes smaller, but far from satisfying the positive comovement condition of proposition 2 even for the upper bound of κ/ω, which is 3.3 (the horizontal line on the graph).

[13]The value of η is a function of β and δ_i only in the form of $1-\beta(1-\delta_i)$. Hence there is no need to consider β and δ_i separately.

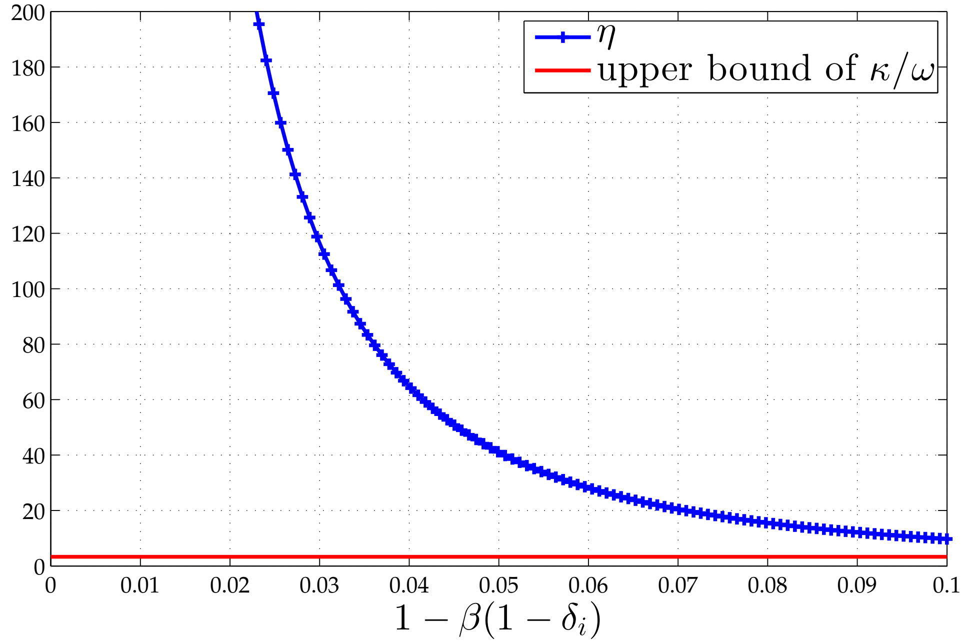

Figure 1: Value of η as a function of $1 - \beta(1 - \delta_i)$. Stock-elastic demand model; holding fixed all the other structural parameters.

4 Dynamic analysis

The analysis of the previous section focused on the impact responses to news shocks, in an effort to understand forces underlying the joint response of inventories and sales to news shocks. We found that news shocks generate negative comovement between inventories and sales. We now turn to several extensions of this result. We first show that the negative comovement between inventories and sales holds beyond impact, and we establish that introducing variable capacity utilization (which allows output to increase in the short run) does not overturn our result. Second, we study the response of inventories to surprise (non-news) shocks, and confirm that the negative comovement property is an identifying feature of news shocks. Third, we analyze whether news about other fundamentals than productivity or the labor wedge, in particular about news about demand, also generate negative comovement between inventories and sales. Fourth, we check the robustness of our result by introducing different types of adjustment costs to inventories.

Since the analysis will be numerical, we start with a brief discussion on the calibration of the model.

Parameter	Value	Description
β	0.99	Subjective discount factor
σ	1	Inverse elasticity of household intertemporal substitution
δ_k	0.025	Depreciation rate of capital
$\phi''(1)$	9.11	Investment adjustment costs
ξ	2.5	Frisch elasticity of labor supply
ψ	6.72	Steady-state hours worked = 0.2
α	0.67	Labor elasticity of production function
θ	5	Elasticity of substitution across intermediate goods
δ_i	0.025	Depreciation rate of inventories
ζ	0.25	Steady-state inventory-sales ratio = 0.75
ρ_z	0.99	Persistence of the productivity process
ρ_ψ	0.95	Persistence of the labor wedge process

Table 1: Calibration of the stock-elastic demand model.

4.1 Calibration

The calibration which we use in this section is summarized in table 1. Standard model parameters are calibrated using estimates from the business-cycle literature. For the exogenous variables we assume that the realization of these shocks follow AR(1) processes. For the persistence of the shocks, we assume that $\rho_z = 0.99$ and $\rho_\psi = 0.95$.

Our calibration implies that $\eta = 67.15$, $\omega = 1.09$ and $\kappa = 1.02$, so that applying proposition 2, inventories respond negatively to news shocks on impact.

4.2 Impulse response to news shocks and variable capacity utilization

We first study the impulse responses of output, sales and inventories to 4-period positive news shocks to productivity and the labor wedge. That is, at period 0, agents get signals that future productivity ($\mathbb{E}_0 z_4$) will increase or future labor wedge ($\mathbb{E}_0 \psi_4$) will decrease.[14]

Figure 2 reports the impulse responses. Note first that consumption and investment, which are the two components of private sales, increase immediately, and during all subsequent periods. Consumption increases because of the wealth effect associated with the good news, and investment increases because of the presence of investment adjustment costs.

In line with our previous discussion, inventories fall. The fall is large and persistent, and reaches its trough in the period preceding the realization of the change in fundamentals. At the same time,

[14] We define a positive news shock as a shock that generates an increase in sales. When the labor wedge is expected to decrease, then households expect to face less disutility of working in the future. This will also boost current sales.

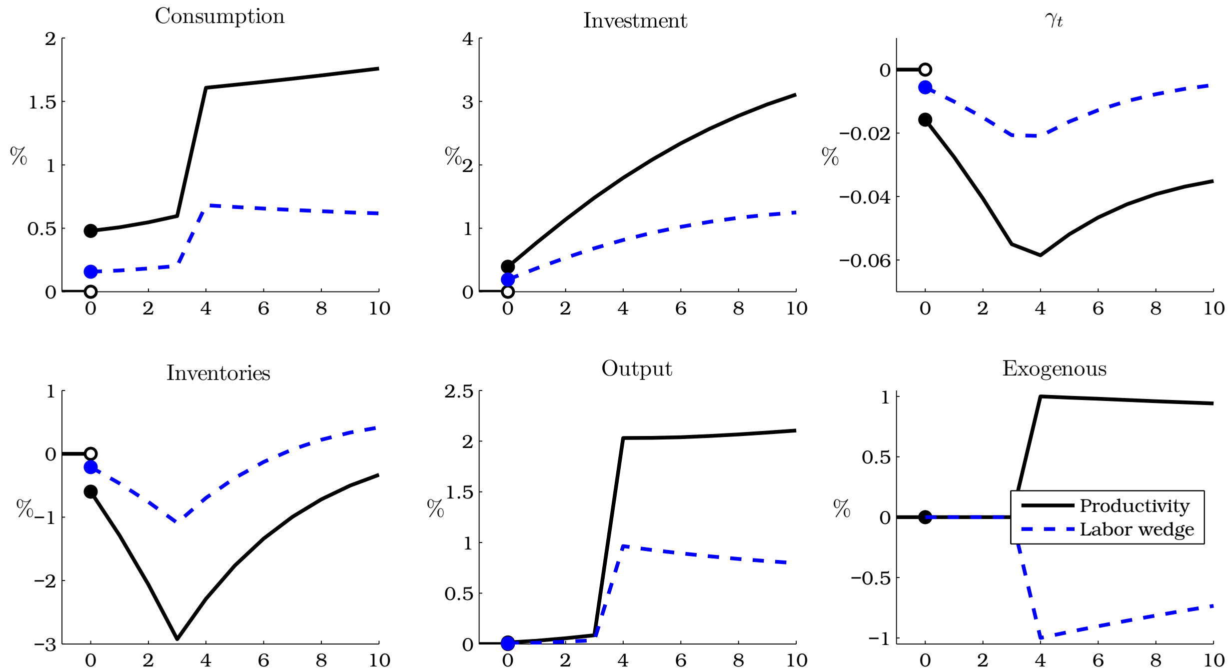

Figure 2: Impulse responses to news shocks in the stock-elastic demand model. Solid line: 4 period news about productivity; dashed line: 4 period news about the labor wedge. The time unit is a quarter. Impulse responses are reported in terms of percent deviation from steady-state values.

output remains mostly unchanged until period 4, when fundamentals change. That is, the increase in sales during the news period is almost entirely met by inventory disinvestment. To build further intuition for the responses of inventories, note that labor market clearing implies that:

$$\psi_t n_t^{\frac{1}{\xi}} = \alpha mc_t z_t k_t^{1-\alpha} n_t^{\alpha-1}, \tag{28}$$

so that marginal cost is given by:

$$\widehat{mc}_t = \omega \hat{y}_t - \hat{z}_t + \hat{\psi}_t - (\omega + 1)(1 - \alpha)\hat{k}_t. \tag{29}$$

This equation tells us that both news about an increase in future productivity and news about a decrease in future labor wedge would contribute to a decline in future marginal cost. In general equilibrium, this downward pressure in the marginal cost profile is reflected in the negative impulse response of the expected discounted marginal cost γ_t, which we report in the upper right panel of figure 2. Since inventories are used to smooth out the difference in marginal cost of production over time, this fall in the expected discounted marginal cost leads to a fall in inventories which is sufficient to overcome the effect of the increase in sales, as indicated by equation (22).

Note that we are not forcing output to be fixed during the news period and that there still is a small increase in output for the first four periods. Although capital is fixed in the short run, and both productivity and the labor wedge are unchanged during the news period, the labor demand schedule of firms may still shift with changes in marginal cost, as we see from the right hand side of equation (28). Indeed, in contrast to models without inventories, the optimal pricing policy of firms does not imply that marginal cost is fixed — instead, it is the expected discounted marginal cost that is constant. Through equation (28), the increase in demand is associated to a rise in marginal cost which shifts out the labor demand curve, resulting in a small increase in hours worked. However, since the marginal cost is effectively smoothed out by the strong inventory substitution channel in our economy, the actual movement in marginal cost is small and therefore labor only slightly increases in equilibrium. Therefore the small change in output is an optimal response of the economy with inventories.

To make this point more clear, we allow capacity utilization to vary. Denoting u_t as the utilization of capital at period t, the production function and the captial accumulation function are modified as follows:

$$\begin{aligned} y_t &= z_t(u_t k_t)^{1-\alpha} n_t^{\alpha}, \\ k_{t+1} &= (1-\delta(u_t))k_t + \left[1 - \phi\left(\frac{i_t}{i_{t-1}}\right)\right] i_t, \end{aligned}$$

where $\delta(1) = \delta_k$, $\delta'(\cdot) > 0$ and $\delta''(\cdot) > 0$. In words, higher utilization of capital increases output, but this comes at a cost of higher rate of depreciation of capital. In a model without inventories, such as Jaimovich and Rebelo (2009), capacity utilization increases with news about a future rise in productivity. This is because with a future rise in productivity, the presence of investment adjustment costs leads to an increase in capital investment today. The increase in capital investment generates a fall in the value of installed capital. At the same time, the positive income effect from the household generates a fall in the marginal value of income due to the concavity of the utility function. Overall, the fall in the value of installed capital is steeper than the fall in the marginal value of income, and therefore capacity is utilized more to satisfy the additional demand.

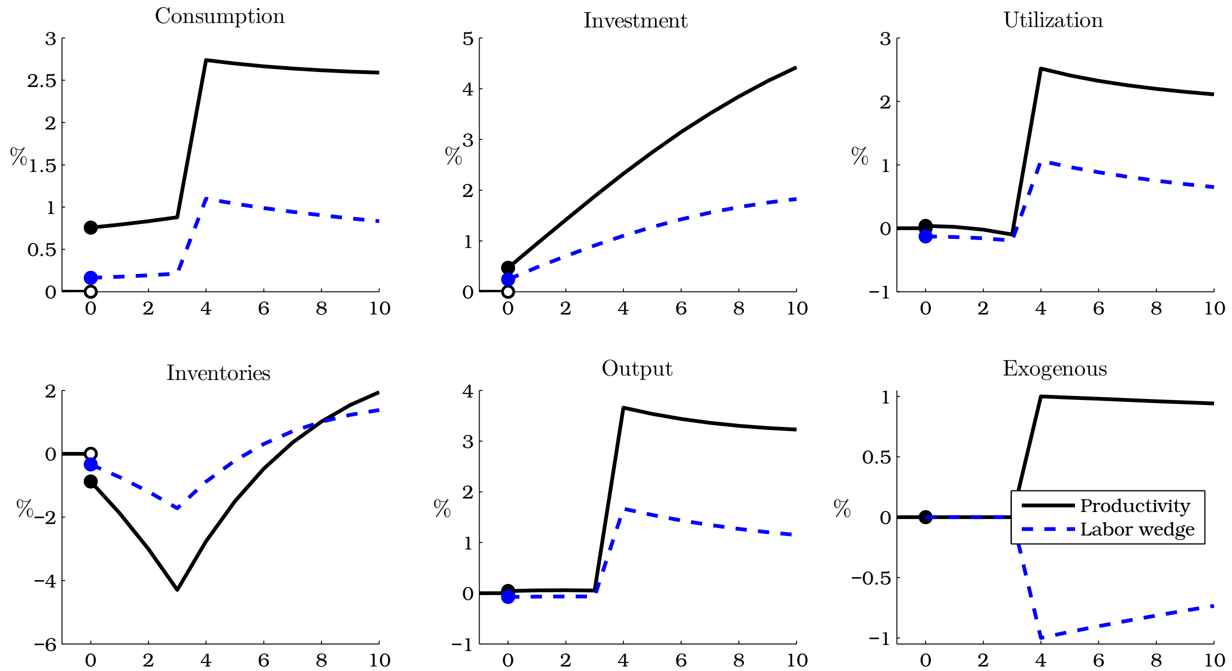

Figure 3: Impulse responses to news shocks in the stock-elastic demand model with variable capacity utilization. Utilization parameter: $\delta_k''(1) = 0.34$; solid line: 4 period news about productivity; dashed line: 4 period news about the labor wedge. The time unit is a quarter. Impulse responses are reported in terms of percent deviation from steady-state values.

In figure 3, we plot the impulse responses for the inventory model with variable capacity utilization. As we see, the quantitative response of capacity utilization during the news period is modest. Utilization significantly increases only after the shock realizes.

The small response of capacity utilization during the news period comes directly from the household preference and the role of inventories in the economy. The marginal value of income λ in our model with GHH preference is the following:

$$\lambda = \left(c - \psi \frac{n^{1+\xi^{-1}}}{1+\xi^{-1}} \right)^{-\sigma}.$$

In this economy, the increase in consumption and investment can be matched by depleting inventories, rather than by working more or utilizing capital more intensively. As in the baseline model, hours do not need to rise along with the increase in consumption. The response of hours generates a steeper fall in the marginal value of income than in a model without inventories. In turn, a steeper fall in the marginal value of income means that Tobin's Q falls by less, in relative terms.

This implies that increased capacity utilization (which uses up installed capital) is a less attractive way of meeting the increase in sales. Instead, firms prefer to sell out of stock in order to meet the increase in demand. Thus, the negative comovement between inventories and sales in response to news shocks is an equilibrium outcome, even when we allow for production to increase in the short-run, either through increased utilization of labor or capital.

4.3 Do surprise shocks generate positive comovement?

While news shocks generate a persistent negative comovement between inventories and sales, one may wonder whether this also occurs after surprise innovations to fundamentals. The impulse responses reported in figure 4 show that this is not the case. Inventories, consumption, investment and output all increase in response to surprise innovations to productivity and the labor wedge. The short-run response of the inventory-sales ratio is also consistent with its observed countercyclicality at business-cycle frequencies, in line with the findings of Khan and Thomas (2007a) and Wen (2011).[15] For this type of shock, the model's predictions are thus broadly consistent with the observed behavior of inventories and sales over the business cycle. Thus, the negative comovement of inventories and sales is an identifying feature of news shocks relative to non-news shocks.

4.4 Other types of news shocks

Although the two types of shocks we have considered up to now have been found to be significant sources of news in the literature (Schmitt-Grohé and Uribe, 2012), we do not need to limit our result to these shocks. In fact, proposition 2 implies that the negative comovement holds for any type of news shocks, since on impact, all news shocks share the feature that no fundamentals change.

In this section, we consider two other types of news shocks: discount factor shocks and government spending shocks. First, consider a news shock to the discount factor. When the discount factor is expected to increase in the future, then households expect that in the future they will consume more and save less. They will then consume less today since they now discount the future

[15]The countercyclicality of the inventory-sales ratio is not completely robust to the calibration of the shock, as it depends partly on the magnitude of the initial increase in sales. For a smaller persistence of productivity shocks of $\rho_z = 0.8$, for example, the response of sales is more muted, and the IS ratio becomes procyclical. This behavior of the inventory-sales ratio has motivated Kryvtsov and Midrigan (2013) to investigate the ability of countercyclical markup movements to mute inventory increases in response to demand-side shocks, since in the data, the inventory-sales ratio is countercyclical. However, in response to both productivity and demand shocks, the procyclicality of inventories holds regardless of the values of the persistence parameters ρ_z and ρ_ψ.

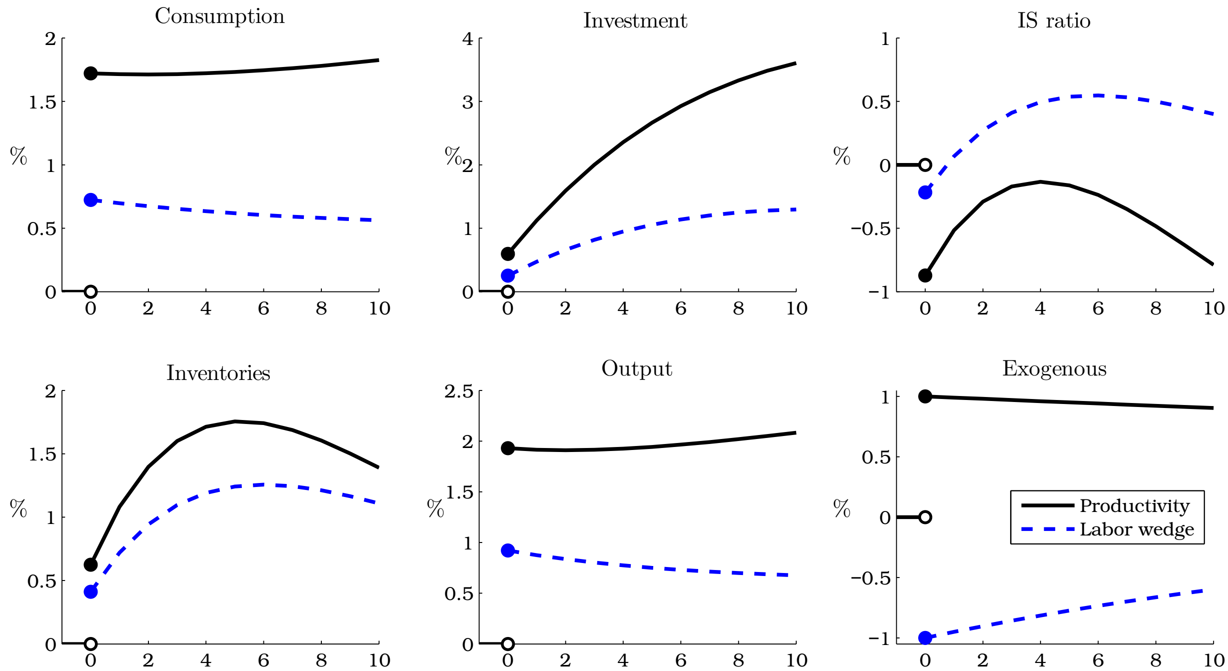

Figure 4: Impulse responses to surprise shocks in the stock-elastic demand model. Solid line: productivity; dashed line: labor wedge. The time unit is a quarter. Impulse responses are reported in terms of percent deviation from steady-state values.

less. Moreover since savings and hence investment will decrease in the future, with investment adjustment costs, investment will also start decreasing today. Therefore, news about an increase in future discount factor generates a fall in sales. At the same time, the fall in investment leads to a decrease in future capital, which generates an increase in future marginal cost. Therefore, inventories will increase, confirming that the negative comovement property holds.

Second, when there is a future increase in government spending, then inventories will increase to build up for the demand from government spending, since marginal cost is expected to rise in the future with the additional demand from the government. At the same time, since the households in the end take the burden of this spending, consumption and investment fall. Again, there is negative comovement between inventories and sales with this type of news shock.

Figure 5 shows the impulse responses to these two shocks.[16] As discussed, the negative comovement property also holds on impact and afterwards.

[16]The persistence of each process is 0.17 for the discount factor, and 0.95 for government spending. These values are those estimated by Schmitt-Grohé and Uribe (2012).

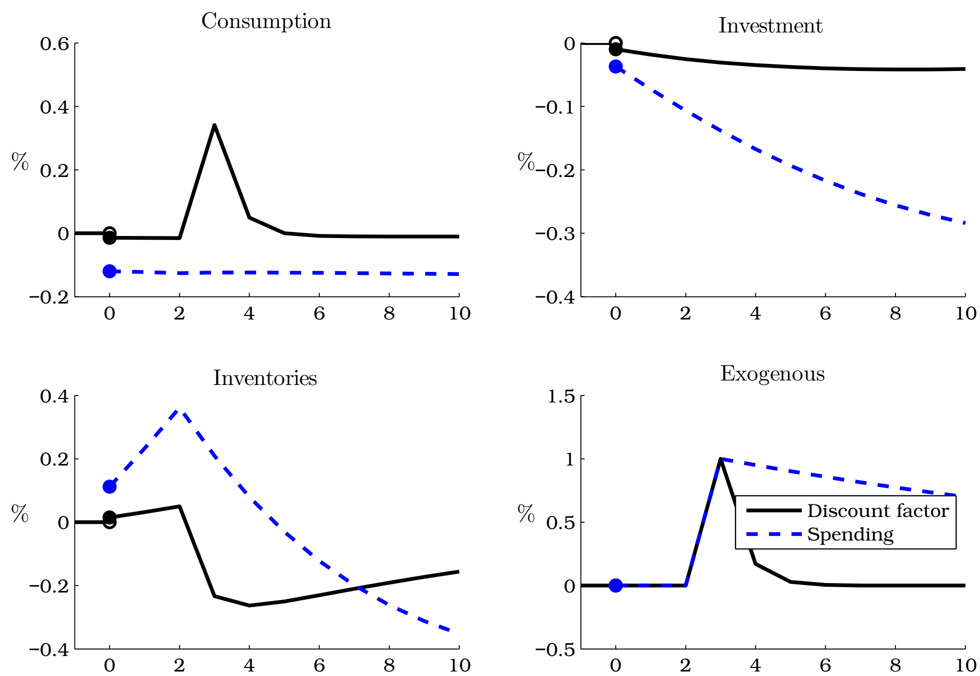

Figure 5: Impulse responses to other news shocks in the stock-elastic demand model. Solid line: 4 period news about the discount factor; dashed line: 4 period news about government spending. The time unit is a quarter. Impulse responses are reported in terms of percent deviation from steady-state values.

4.5 Adding adjustment costs

Adding adjustment costs to capital investment is a crucial ingredient to generate an investment boom in response to news shocks (Jaimovich and Rebelo, 2009). Capital is slow to adjust, and with this form of adjustment cost, investment decisions depend on the whole expected future path of the marginal value of capital, or future Tobin's Q. News shocks affect the marginal productivity of future capital, and thus raise future Tobin's Q, which directly translates into an increase in current investment.

The rationale for imposing adjustment on inventory investment is less clear, in particular in the case of finished-good inventories. First, whereas building a factory or machinery takes time and hence requires adjustment periods, stocking or depleting an already existing product should be the most flexible adjustment that firms can take. Second, as we discussed in the previous sections, it is not the level, but the growth rate of marginal cost that is important for finished-good inventory investment decisions. Therefore, adding adjustment costs to finished-good inventory investment is a less appealing approach.

However, adjustment costs to the stock of inventories may have a better justification: the total stock of inventories do seem large and slowly moving. Moreover, our analysis so far suggests that with a positive news shock, we need additional channels for production to increase and adjustment costs may be helpful in that respect. We consider three possible types of adjustment costs: adjustment costs to inventories, output and on-shelf goods. Adjustment costs to inventories penalize immediate inventory depletion and thus weaken the intertemporal substitution motive. Adjustment costs to output force firms to smooth out the response of output to the shock, and in turn reduce the incentive to deplete inventories to satisfy sales. Finally, adjustment costs to goods on shelf affect the sum of output and past inventories. Making adjustment costs bear on this variable might have effects that combine both types of adjustment costs described above.

These adjustment costs are introduced by assuming that the law of motion for inventories is modified as follows:

$$inv_t = (1-\delta_i)inv_{t-1} + y_t - s_t - ADJ_t,$$

where ADJ_t is the adjustment cost of each type. We assume the following form:

$$ADJ_t = \phi_x\left(\frac{x_t}{x_{t-1}}\right)x_t, \qquad x \in \{inv, y, a\},$$

where $\phi_x(1) = \phi_x'(1) = 0$ and $\phi_x''(1) > 0$. In figure 6, we show the responses of the model with and without adjustment costs. We experiment with different levels of adjustment costs, and for all values, we observe that the initial fall in inventories are smaller in both models with adjustment costs, but not close to being positive. We conclude that adjustment costs to inventories and output are not sufficient to generate a procyclical response of inventories.

The logic behind this result is that with adjustment costs to inventories or production, firms are now more willing to smoothly adjust their stock of inventories, and hence produce more today when there is good news. However, to make this happen, wages must increase to induce households to work more. With an increase in wages, households have more income, and consumers will increase their current consumption level. This additional increase in consumption, arising from an increase

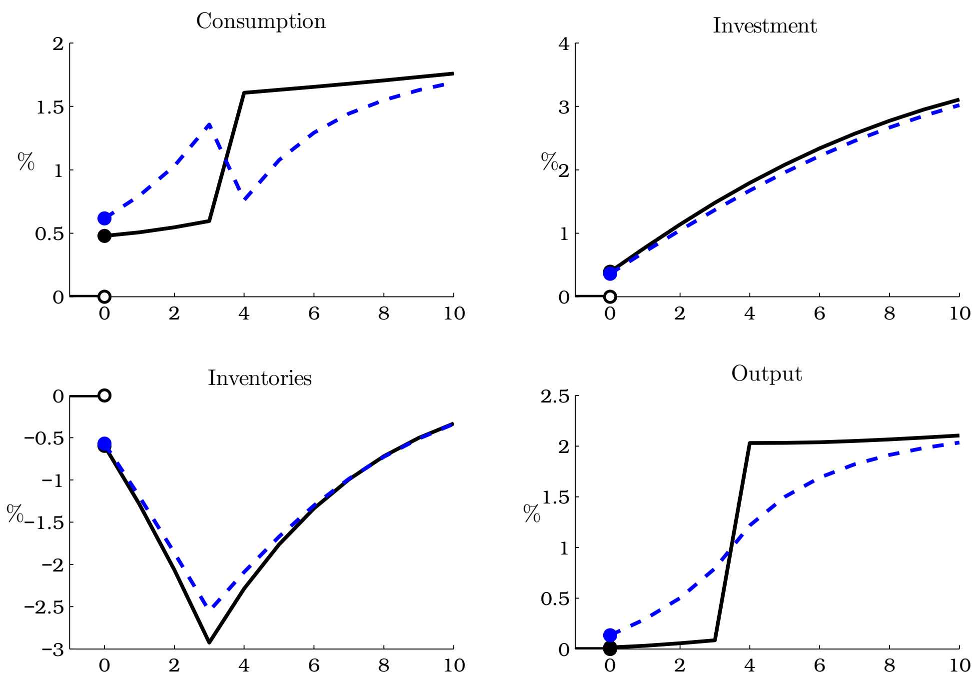

Figure 6: Robustness of impulse responses to output adjustment cost. Impulse responses to 4 period productivity news shock. Solid line: without output adjustment cost; dashed line: with output adjustment cost. The time unit is a quarter. Impulse responses are reported in terms of percent deviation from steady-state values.

in wages, partially offsets the effect of the increase in output on total inventory accumulation.

4.6 Introducing nominal rigidity

In line with the bulk of the literature on news shocks, we have so far studied economies in which firms are allowed to costlessly change their price every period. However, in an economy with incomplete price adjustment, end-of-period inventories may accumulate (or decumulate) not only because of intertemporal substitution motives, but also because firms charge excessively high (or low) prices, resulting in sales below (or above) target. This may have implications for the results on the comovement between inventories and sales obtained so far.

We study this issue by introducing quadratic price adjustment costs in the firm's profit function into our baseline setup, following Rotemberg (1982). Details are reported in the online appendix.[17] The first result is that even in the presence of nominal rigidities, both impact and dynamic co-

[17] The online appendix reports results for price adjustment costs which imply a slope of the New Keynesian Philips Curve identical to Calvo models with price change frequencies ranging from 10% to 50% at the quarterly frequency, which nests the estimates of this parameter obtained in the literature.

movements of inventories and sales in response to news shocks remain negative. Nominal rigidities in fact amplify the negative response of inventories to good news about the future. In response to good news, demand and marginal cost increase. With incomplete price flexibility, firms are unable to fully offset the increase in marginal cost by increasing their price, and their markup falls. This implies an even larger increase in demand. Output responses remain muted even with nominal rigidities, hence inventories decline more than under flexible prices. For high degrees of nominal rigidity (equivalent to Calvo price change probabilities of 10% per quarter), the initial drop in inventories is roughly twice as large as in the model with flexible prices.

Our second result is that, even with incomplete price adjustment, inventories and sales comove positively on impact in response to surprise changes in TFP and the labor wedge. Again, output response is dominated by the intertemporal substitution channel and with a temporary reduction in marginal cost induced by an increase in TFP, firms substitute production towards the present and accumulate inventories today.

We conclude that in our baseline model, introducing nominal rigidities does not reverse our two claims: news shocks generate negative comovement between sales and inventories, while surprise shocks generate positive comovement.[18]

5 Robustness: Other inventory models

A natural question is whether our result is specific to the inventory model we have chosen to analyze. In this section, we discuss models with alternative micro-foundations for inventory holdings. Leading business-cycle models introduce inventories either as buffers against uncertainties in demand at the firm level (stockout-avoidance models), or as stocks associated with nonconvex delivery costs at the firm level (Ss inventory models). We will focus more on the first approach since it provides a better account of finished-good inventory behavior (Khan and Thomas, 2007a). Nevertheless, we also discuss the second approach for completeness.

A summary of the discussion is that our result holds for other standard inventory models as well. This is because one important role for inventories in all of these models is the intertemporal

[18] We limit ourselves to this simple model in order to illustrate the fact that nominal rigidities need not overturn our results. The question of whether different forms of nominal rigidities, such as for example wage rigidities, affect the transmission of news has not, so far, been definitively addressed in the literature, and is beyond the scope of this paper.

substitution channel. With inventories, producers are allowed to flexibly change their production schedule based on their perception of the intertemporal cost profile. Since news shocks directly affect this perception, the other margins which differ across models matter less.

5.1 Stockout-avoidance model

One branch of the literature on finished-good inventories motivates inventories by introducing a lag between production and the realization of sales. These "stockout-avoidance models" assume that firms have imperfect information on their idiosyncratic demand schedule at the time they make production decisions. When realized demand for their product is unusually high, firms may run out of available product — a "stockout" — and lose potential sales. This motivates firms to put, on average, more on-shelf goods than they expect to sell, and carry over excess goods as inventory into the next period.[19]

In an online appendix, we study the effects of news shocks in this class of models in detail. We show that a log-linearized framework similar to that of proposition 1 obtains, and moreover that our main result carries through: in response to good news about the future, under standard calibrations of the model, sales increase while inventories fall.[20] Additionally, we argue that, as in the stock-elastic demand model, the main mechanism dominating the response of inventories to news shocks is intertemporal substitution in production. In figure 7, we plot the value of η, the degree of intertemporal substitution, as a function of the intertemporal cost. Again, we see that even with large intertemporal substitution costs, the degree of intertemporal substitution is strong.

The similarity of the two classes of models comes from the fact that the optimal stocking condition (18) also holds in the stockout-avoidance model. The cost of stocking is the marginal cost of production. The benefits of stocking are twofold: (i) in case sales turn out to be higher than expected, the firm can increase its sales by stocking more output; (ii) in case sales turn out to be lower than expected, the firm can lower future production costs by stocking output as inventories. It turns out that even in this class of models, the intertemporal substitution motive is quantitatively

[19]This mechanism is consistent with existing evidence that stockouts occur relatively frequently at the firm level. Bils (2004) uses data from the BLS survey underlying the CPI and estimates that stockout probabilities in this dataset are roughly 5 percent. More recently, using supermarket-level data for a large retailer, Matsa (2011) suggests that stockout probabilities are in the range of $5-10$ percent. See Kahn (1987, 1992), Kryvtsov and Midrigan (2010, 2013), and Wen (2011) for detailed analysis of the properties of this class of models.

[20]This follows from obtaining analytical restrictions on reduced-form parameters to precisely quantify the conditions under which this result holds.

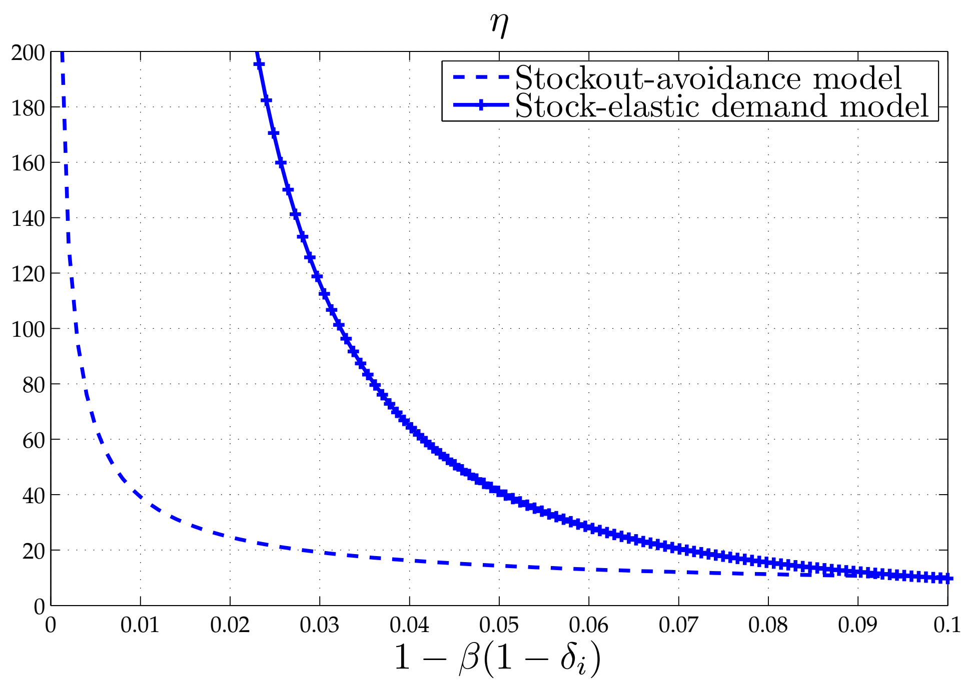

Figure 7: Value of η as a function of $1-\beta(1-\delta_i)$, in the stockout-avoidance model, holding fixed all the other structural parameters. For comparison, η for the stock-elastic demand model, same as figure 1, is also plotted.

stronger for news shocks.

5.2 (S,s) inventory model

The existence of nonconvex delivery costs at the firm level has also been claimed to be a potentially important motive for holding inventories, albeit mainly input inventories, as opposed to finished-good inventories on which our analysis has so far focused. In the model of Khan and Thomas (2007b), the firm pays a fixed cost when placing an order for intermediate inputs. The fixed delivery cost is stochastic, which generates a steady-state distribution of firms with different levels of input inventories. In this model, the optimal stocking condition for stock-adjusting firms also balances the cost and benefit of ordering inputs, along the lines of the discussion of equation (18). To be precise, the cost of stocking is the total cost of inputs and a fixed delivery cost. The benefits of stocking are twofold. (i) In case the delivery cost turns out to be higher than expected, firms will not order in that period. The total production capacity of the firm is then constrained by the amount of input inventories it previously stocked. Hence, more input inventories allows the firm to produce more goods when demand is high but the delivery cost of new intermediate input becomes too high. (ii) In case the delivery cost turns out to be lower than expected, firms can order intermediate inputs in that period. In this case, the firm will bunch their orders for intermediate inputs and stock them as inventories if it expects that the unit cost of inputs will be high in the future.

In response to news about an increase in future productivity, firms understand that future demand will increase. At the same time, they understand that future unit costs of intermediate inputs are also cheaper. We solved for the perfect foresight transition dynamics with a news shock to productivity in Khan and Thomas (2007b).[21] Along this path, as in the case of the other models, inventories fall in the short-run, while sales increase.

Overall, these alternative models also have the feature that, in response to good news about the future, inventories fall while sales increase, especially right before the realization of the shock. Therefore, we conclude that the strong intertemporal substitution channel with news shocks is a common feature of other standard inventories models as well.

6 Estimating the importance of news shocks I: SVAR approach

Our analysis of inventory models suggests that the negative comovement of inventories and sales is a defining feature of news shocks. In this section, we use this structural restriction to estimate the importance of news shocks.

The approach we take in this section is to estimate a structural VAR with sign restrictions. Since the robust prediction of our theoretical analysis is that news shocks generate negative comovement between inventories and sales, we will use this prediction directly to estimate the explanatory power of news shocks. The appealing aspect of our sign restriction VAR approach is that we can remain agnostic as to other identifying features, and therefore robustly identify shocks without misspecification concerns. On the other hand, a drawback of this approach is that identification of shocks is only partial, in the sense that we may also be capturing non-news shocks that drive negative comovement between inventories and sales. As such, our results can be interpreted as providing an upper bound on the contribution of news shocks to business-cycle volatility.

[21] Refer to Khan and Thomas (2007b) for the solution algorithm.

6.1 Data

We use four observables in our exercise: inventories, consumption, investment and output. Consumption includes nondurables and services, investment includes fixed investment and durables, and output is GDP. For inventories, we used nonfarm private inventories as a whole, or only retail trade inventories to focus on finished-good inventories.[22] However, our results are not sensitive to the type of inventories used for estimation. Therefore, in this section, we present results for nonfarm private inventories. All data are seasonally adjusted, and expressed in real per capita terms. Our sample period is 1955Q1–2006Q4.[23]

6.2 Baseline specification and estimation

Our baseline identification strategy imposes that on impact, there is disinvestment in inventories, whereas consumption and fixed investment increase.[24] The VAR model we estimate is the following:

$$X_t = A + B(L)X_{t-1} + U_t.$$

For X_t, we use log levels of each variable, in order for the estimation to be robust to the presence of cointegrating relations. We estimate the model with a constant term and four lags.[25] We use Bayesian methods, with a diffuse prior for both the coefficients of the autoregressive structure and the variance-covariance matrix of the error terms. Each draw from the posterior identifies a set of possible impulse responses satisfying our impact restriction, and we use a uniform conditional prior on the identified set to draw from the posterior of the impulse responses, following Moon,

[22] Nonfarm private inventories include input inventories as well as finished-good inventories. While a portion of these input inventories are work-in-progress inputs since production takes time for some industries, the remaining material input inventories are goods that are *not* used for current production and hence are different from capital. As we discussed in the (S,s) inventory model, positive news shocks to productivity would lead firms not to overstock materials that will not be used immediately since with higher productivity, the cost of these goods become cheap in the future. Since material inventories and finished-good inventories together is on average 80 percent of nonfarm private inventories, we use nonfarm private inventories as our benchmark variable. However, we are also aware that the complexity of the production structure in some industries may obscure the definition of intermediate inventories from that of intermediate inputs and hence we also use retail trade inventories separately. We thank an anonymous referee for pointing this out.

[23] The sources of the data are NIPA tables 1.1.5 and 5.7.6.

[24] On impact, a fall in inventories is equivalent to a fall in inventory investment, since the impulse response is from the steady state. The joint restriction on consumption and investment is not restrictive since in the data, the two series are highly positively correlated.

[25] The Schwartz information criterion suggests two lags. The results to follow are not sensitive to the number of lags.

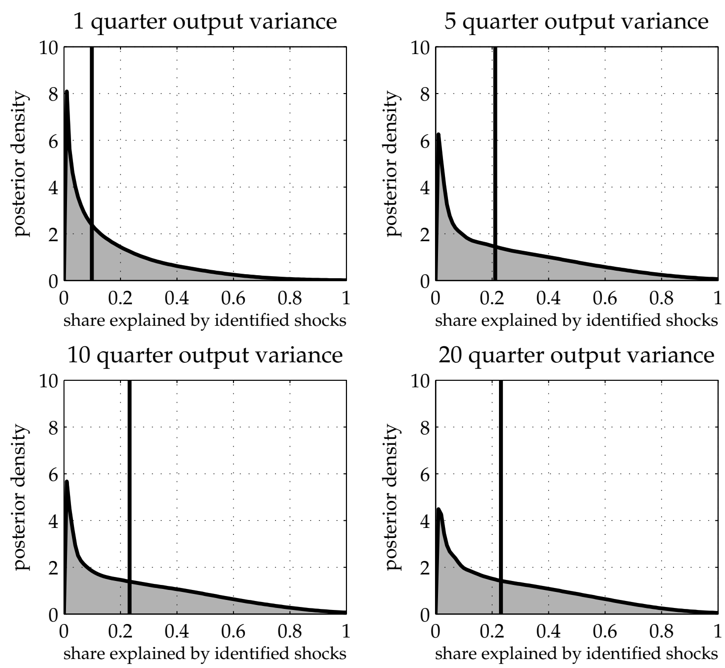

Figure 8: Output variation accounted for by shock identified using impact restrictions. Posterior probability density and the median (vertical line) for the share of forecast error variance at each horizon.

Schorfheide, and Granziera (2013). Using 20000 draws, the posterior distribution of the forecast error variance (FEV) of output accounted for by these identified shocks is computed.[26]

6.3 Baseline result

Figure 8 reports the posterior distribution of the FEV of our identified shocks on output, for different horizons.[27] The posterior has a sharp mode close to zero, and the median is close to 20 percent at most horizons. In figure 9, we plot the set of identified impulse responses. The median response to our identified shock generates a persistent boom in consumption and investment, and a moderate boom in output. The fall in inventories is short lived; on average, positive inventory investment occurs immediately after the initial disinvestment, and the stock of inventories becomes positive after 3 quarters. Note that in our model, this is the case when good news are expected to materialize in the near future. Therefore, our identified shock resembles short-horizon news, with

[26]Our result to follow is not sensitive to adding more draws.

[27]As noted above, we plot the case for nonfarm private inventories but the plot is similar with retail trade inventories as well.

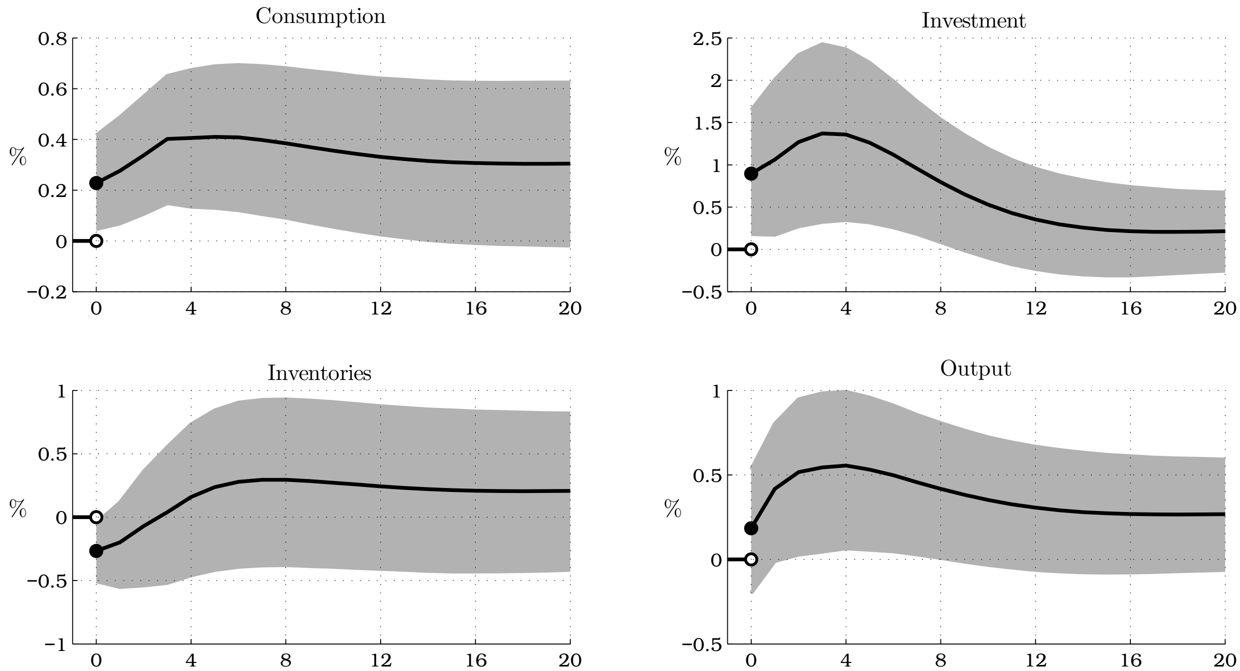

Figure 9: Impulse responses to shock identified using impact restrictions. Median (solid line) and 80% credible set.

news lasting for only 1 period.

Our identification strategy only imposes impact restrictions, and therefore we are not able to distinguish between short and long-horizon news shocks. Since the focus of the news literature is not on one or two quarter news shocks, but rather longer horizons, our next step is to impose restrictions beyond impact.

6.4 Dynamic restrictions

Figure 10 reports the posterior distribution of the FEV of output attributable to our identified shocks, when the following restriction is used to identify them: inventory disinvestment occurs for 2 periods, and at the same time both consumption and investment increase for 2 periods.[28] We see that the posterior has a sharp mode close to zero, and the median is now close to 10 percent at all horizons, about half smaller than the result with impact restrictions only. To get a sense of the information that inventories deliver, figure 10 also plots the posterior distribution of the FEV when only consumption and investment are above the steady state for 2 periods. As we see, without the

[28] Simulating the model with a range of different parameters, we find that this 2 period sign restriction indeed holds for the news shocks we consider.

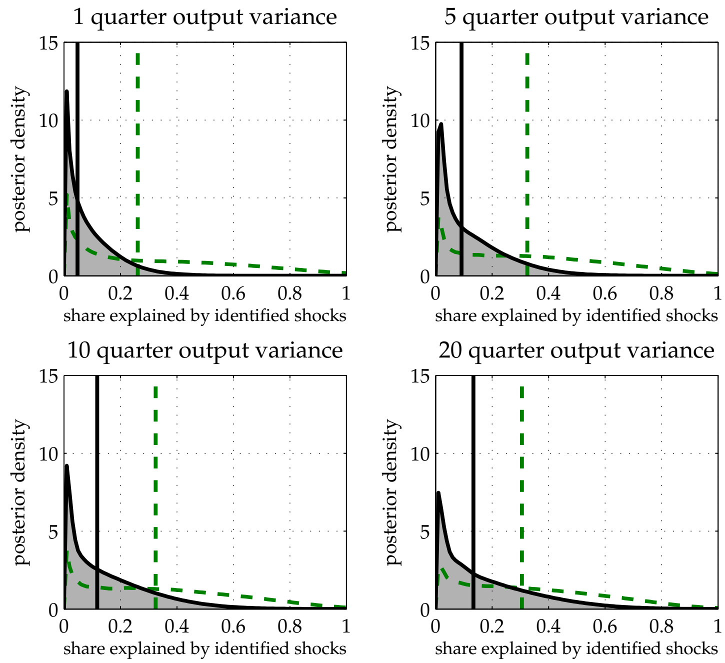

Figure 10: Output variation accounted for by shocks identified using 2 period restrictions. Posterior probability density and the median (vertical line) for the share of forecast error variance at each horizon. Solid line: 2 period negative comovement between Δinv_t and (c_t, i_t). Dashed line: 2 period positive comovement between c_t and i_t.

inventory restriction, the distribution is disperse and the median share of FEV for the set of shocks that drive positive comovement of consumption and investment is 30 percent overall. Hence with inventories, the posterior density becomes much tighter, and the median share of the shock falls by about 67 percent.

Figure 11 reports the impulse responses to the identified shock with 2 period restrictions. Inventory disinvestment occurs for 2 periods, but after that, there is again investment in inventories. Consumption and investment increase, but the increase in output is now modest.

We also extend our dynamic restriction to 3 periods, that is 3 periods of inventory disinvestment accompanied by increases in consumption and fixed investment. As in figure 12, the median share of FEV explained by the identified shock is now below 5 percent in most horizons, and tight with basically no probability assigned above 20 percent. Therefore, our news shocks identified with 3 period restrictions *at most* account for 20 percent of output variation. Figure 13 reports the impulse responses of the identified shock with 3 period restrictions. Although the movement in output is modest, it actually declines on impact.

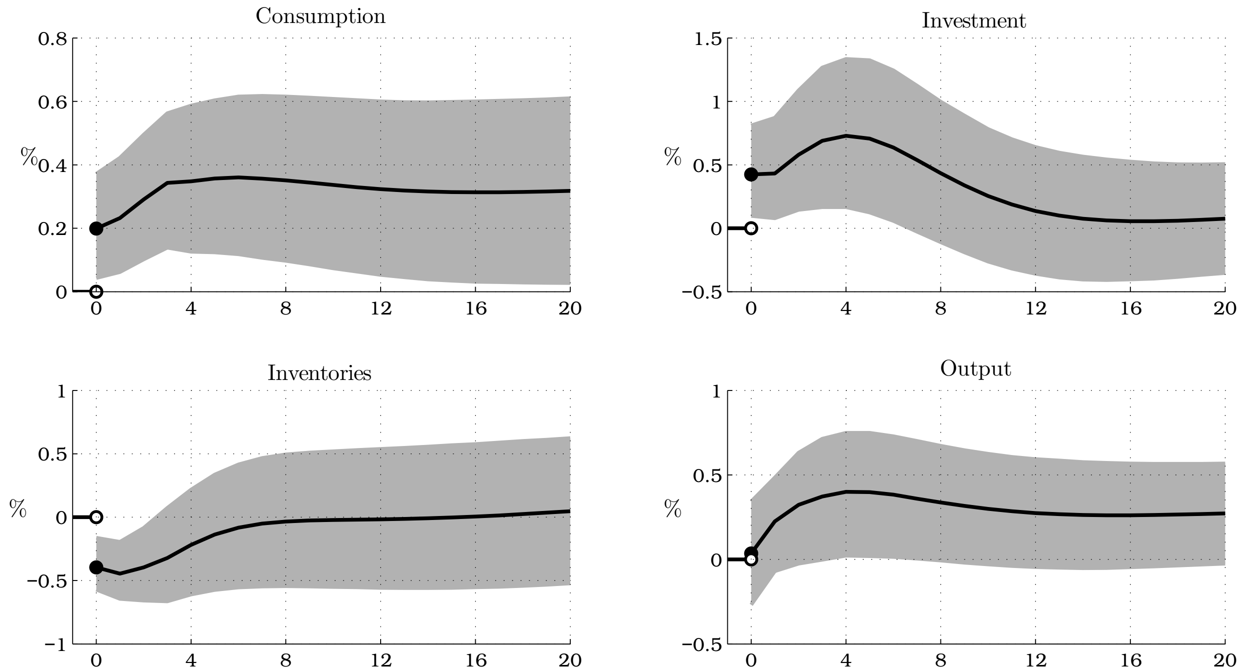

Figure 11: Impulse responses to shocks identified using 2 period restrictions. Median (solid line) and 80% credible set.

We summarize the key points of our empirical results as follows: (i) the identified impulse response with impact restrictions suggest that most news shocks are short-lived, with an immediate investment in inventories after the impact disinvestment; (ii) the identified news shock based on impact restrictions explain on average 20 percent of output variations at all horizons; (iii) restrictions beyond impact generate a tighter posterior distribution of output variations; (iv) long-horizon news shocks explain on average 5 percent, and at most 20 percent of output variations at all horizons.

The reason why the FEV turns out to be small is that inventories are a procyclical variable. In the data, the unconditional contemporaneous correlation between inventories and sales (consumption plus investment) is 0.50.[29] Since our identification is based on negative comovement of these generally positively comoving variables, it would come as a surprise if the identified shocks were able to generate the bulk of business cycles.

[29]This is based on HP filtered data but the result is not sensitive to filtering methods.

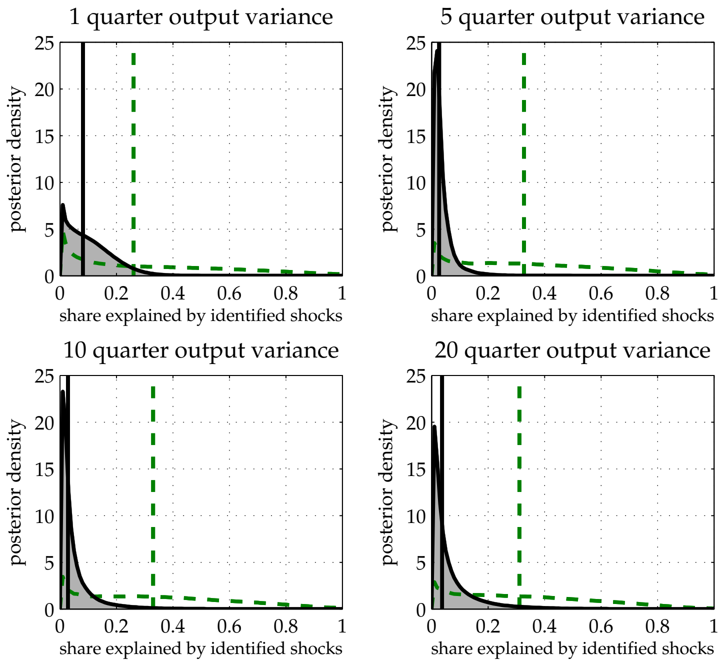

Figure 12: Output variation accounted for shocks identified using 3 period restrictions. Posterior probability density and the median (vertical line) for the share of forecast error variance at each horizon. Solid line: 3 period negative comovement between Δinv_t and (c_t, i_t). Dashed line: 3 period positive comovement between c_t and i_t.

6.5 Robustness

Since our identifying assumption only relates to the sign of the responses of inventories and components of sales, it is robust to changes in specification. We have nevertheless performed several robustness checks. First, we used different priors for the coefficients, such as the Minnesota prior or the Normal-Wishart prior. None of these priors alter significantly the results.[30] Second, when imposing our dynamic restriction, we also tried to be less restrictive by not imposing the negative comovement on impact or second period, in order to control for any demand effects that may remain in the short run with long-horizon news shocks. The result is not sensitive to this change since the stock of inventories move in a persistent manner. For example, by imposing that inventories are below average only in the third period, it mostly follows that inventories are below average for the first and second period as well. Third, as we mentioned above, our result is not sensitive to

[30]Since our focus is mainly on FEV decomposition, it might be more desirable to set a uniform prior on this moment. However, forecast error variance is a highly nonlinear transformation of the VAR coefficients, and existing methodologies do not allow us to easily adress this "inverse" problem. As a way to overcome this issue, we are showing our result with and without the negative comovement assumption in order to control for the role of the prior.

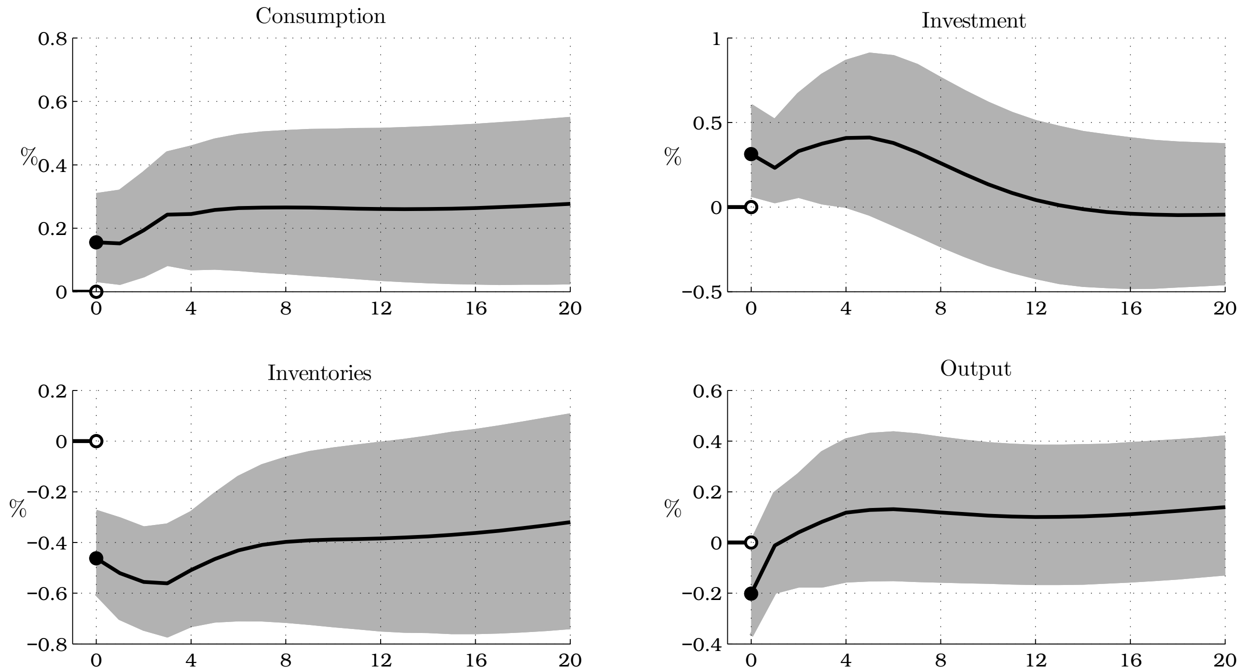

Figure 13: Impulse responses to shocks identified using 3 period restrictions. Median (solid line) and 80% credible set.

using different types of inventory data. Fourth, as studied in detail by McCarthy and Zakrajšek (2007), inventory dynamics have changed since the 1980s: while the procyclicality of inventories remains, the volatility of total inventory investment has fallen, possibly because of improvements in inventory management, contributing to the fall in output volatility. To address this issue, we take into account the possibility of different "inventory regimes" in the data by creating two separate samples, before and after 1984, and conduct our empirical exercise on each of the sub-samples. Our results are not sensitive to these various sample splits. This suggests that the nature of the comovement between aggregate inventories and private sales did not change substantially around this period.

6.6 Relation with earlier VAR approaches

A large part of the empirical literature on news shocks has used VAR identification strategies in order to pin down the contribution of news shocks to business-cycles fluctuations (some of the most prominent examples include Beaudry and Portier (2006), Beaudry and Lucke (2010), and Barsky and Sims (2011)). These papers share a common concern about misspecification bias, and aim to provide identification schemes that are as robust as possible to it. The approach presented in this section contributes to this effort since, as we have argued, the identification restriction we use is satisfied by news shocks in a variety of inventory models.

Our scheme chiefly differs from earlier approaches in that it does not rely on information about movements of TFP. We view this as a strength, since it allows us to provide an upper bound to the contribution of news about all fundamentals, even absent a direct measure of them, such as the TFP series.

It is nevertheless useful to compare our results to this earlier literature. In order to do this, we construct the inventory dynamics implied by news shocks identified using the short-run scheme described in Beaudry and Portier (2006). We run a 3 variable VAR with utilization-adjusted productivity, the S&P 500 index as stock prices, and inventories. We impose a zero impact restriction on productivity, and trace out the dynamics of inventories following an increase in stock prices. We find that the sign of the impact and short-run responses of inventories are inconclusive.[31] We interpret this result as indicating that the information contained in inventory movements about news shocks is not fully incorporated in short-run movement in stock prices and TFP, so that including inventory as part of the identication scheme is helpful to further pin down these shocks.[32] This additional information matters for the implications drawn for the contribution of news to output variation: our approach suggests that both the short and long-run contributions of news to output volatility do not exceed 20%, in contrast to the $50-60\%$ number obtained by Beaudry and Portier (2006). Our estimates are instead closer to the results of Barsky and Sims (2011), who find that in the short run, TFP news accounts to only 10% of output volatility.

7 Estimating the importance of news shocks II: DSGE approach

In this section, we estimate a structural DSGE model incorporating inventories by Bayesian methods, and use the estimates to assess whether news shocks are important contributors to business-cycle fluctuations. The purpose of this section is twofold. First, while the agnostic VAR method is robust to misspecification, it is still a partial identification strategy. Using additional

[31] This plot is reported in the appendix.

[32] Similar points are made in Arias, Rubio-Ramirez, and Waggoner (2013) with regards to the penalty function approach in Beaudry, Nam, and Wang (2011), or in Forni, Gambetti, and Sala (2014) on the non-fundamentalness of small-scale VAR models in identifying news shocks. Our information could add to this debate as well.

information based on the structure of our economy is in principle helpful in identifying news shocks more precisely. Second, our discussion is so far limited to shocks that are stationary. However, an important component of news shocks may be nonstationary and the importance of these nonstationary components are better understood when we directly model them.

It is however important to keep in mind that estimating a structural DSGE model has its own limitations. Our theoretical analysis did not require us to take a stand on a specific view of the structure of the economy. However, to estimate a DSGE model, we need to select a specific model to estimate, and the results we obtain are potentially subject to misspecification issues.

7.1 Model specification

The model we estimate in this section is an extended version of Schmitt-Grohé and Uribe (2012) with inventories introduced as in Bils and Kahn (2000). The model we estimate is similar to that of section 2, and its details are described in an online appendix. However, there are several differences that are worth mentioning here.

First, we allow for two types of nonstationary shocks in the model: nonstationary productivity and nonstationary investment-specific productivity shocks. By allowing for these shocks, we will be able to separately estimate the importance of stationary versus nonstationary news shocks.

Second, we allow for the price markup to change over time. That is, the demand function in (12) is now written as

$$s_t(j) = \left(\frac{a_t(j)}{a_t}\right)^{\zeta} \left(\frac{p_t(j)}{P_t}\right)^{-\theta_t} x_t,$$

where the price markup is assumed to be an AR(1) process.[33]

Third, on top of the seven observables used in Schmitt-Grohé and Uribe (2012), we also use the inventory series described in the previous section as an additional observable in our estimation procedure.

Innovation	Y	C	I	N	G	INV
Prior News Total	37	47	48	39	49	40
Posterior News Total	17	49	10	38	52	14
Stationary Productivity Shock						
News	1	0	0	1	0	1
Current	16	3	8	10	0	14
Nonstationary Productivity Shock						
News	1	1	0	0	0	0
Current	15	10	6	3	4	7
Stationary Investment-Specific Shock						
News	1	1	5	1	0	2
Current	22	4	63	9	0	8
Nonstationary Investment-Specific Shock						
News	0	0	0	0	0	0
Current	0	0	1	0	0	0
Government Spending Shock						
News	1	0	0	1	51	0
Current	1	0	0	1	44	0
Household Preference Shock						
News	4	41	0	6	0	1
Current	3	28	0	5	0	1
Labor Wedge Shock						
News	8	6	3	27	0	4
Current	7	5	3	28	0	4
Price Markup Shock						
News	1	0	1	0	0	6
Current	18	1	9	7	0	47

Table 2: Variance decomposition from estimated model. All values are rounded and are in percentage terms. Y, C, I, N, G, INV refer to the growth rates of output, consumption, fixed investment, hours worked, government spending and inventories, respectively.

7.2 Estimation results

Table 2 summarizes the variance decomposition of the estimated model. While the prior median parameter values imply that news shocks account for 37 percent of output variation, we find that the median posterior estimate drops to 17 percent. This contrasts the result in a model without inventories where 41 percent of output variation could be accounted for by news shocks (Schmitt-Grohé and Uribe, 2012). Therefore, when firms are allowed to adjust inventories in the model, news shocks now play a smaller role. This small contribution of news shocks also holds for fixed investment and inventory investment. For all these variables, news shocks now account for around 10

[33] For θ_t, we transform it into the markup $\mu_t = \theta_t/(\theta_t - 1)$, and assume that it follows an AR(1) process.

percent of total variation. However, for other variables such as consumption, government spending, and hours, we still find a large role for news shocks, consistent with Schmitt-Grohé and Uribe (2012).

To summarize, structural estimation of a DSGE model including inventories indicates that news shocks account for less than 20 percent of output variation over the business cycle, consistently with the results obtained using the VAR approach.

8 Conclusion

In this paper, we studied the response of inventories to news shocks. We derived analytical restrictions on structural parameters under which inventories and sales will positively comove in response to news shocks. We showed that these conditions are violated by standard calibrations of the classes of models we study, resulting in negative comovement between inventories and sales in response to news shocks. Our analysis highlighted the key mechanism behind this result: news shocks generate a strong intertemporal substitution motive in production. Moreover, we showed that this mechanism persists during the "news period", even after introducing various frictions analyzed by the news literature, such as variable capacity utilization and adjustment costs. Lastly, we used the negative comovement between inventories and sales to identify news shocks in postwar US data. We found that news shocks play a small role in aggregate fluctuations, for two reasons: the identified "news period" is short, on average 1 quarter; and the long-horizon shock contributes less than 20 percent of output variation. The insight behind this result is that inventories are procyclical at business-cycle frequencies.

Our work suggests two future directions for progress. First, one contribution of our analysis was to highlight that a key parameter governing the response of inventories to news shocks is the elasticity of inventories to the discounted growth rate of marginal cost. The approach we have taken in this paper is to compute the elasticity implied by existing models of finished-good inventories. An alternative approach is to obtain empirical estimates of this elasticity, and explore modifications of existing models that may match those estimates. Second, we proposed a new way of identifying news shocks, using aggregate data on inventories and sales. An interesting question is whether our theoretical and empirical results could be modified if we were to take a more disaggregated view of inventories, with different sectors having different inventory intensities (Chang, Hornstein, and Sarte, 2009). Theoretically, news shocks in one particular sector may lead to negative comovement of inventories and sales in that sector, but this need not be so in the aggregate. Empirically, differences in the comovement of sales and inventories across sectors, using industry-level data, could be used to identify these sectoral news shocks. We leave this to future research.

References

Arias, J. E., J. F. Rubio-Ramirez, and D. F. Waggoner (2013). Inference Based on SVARs Identified with Sign and Zero Restrictions: Theory and Applications. *Mimeo, Duke University*.

Barsky, R. B. and E. R. Sims (2011). News Shocks and Business Cycles. *Journal of Monetary Economics 58*(3), 273–289.

Beaudry, P., M. Dupaigne, and F. Portier (2011). Modeling news-driven international business cycles. *Review of Economic Dynamics 14*(1), 72–91.

Beaudry, P. and B. Lucke (2010). Letting Different Views about Business Cycles Compete. *NBER Macroeconomics Annual 24*.

Beaudry, P., D. Nam, and J. Wang (2011). Do Mood Swings Drive Business Cycles and is it Rational? *NBER Working Paper 17651*.

Beaudry, P. and F. Portier (2006). Stock Prices, News, and Economic Fluctuations. *American Economic Review 96*(4), 1293–1307.

Beaudry, P. and F. Portier (2013). News Driven Business Cycles: Insights and Challenges. *NBER Working Paper 19411*.

Bils, M. (2004). Studying Price Markups from Stockout Behavior. *Mimeo, University of Rochester*.

Bils, M. and J. A. Kahn (2000). What Inventory Behavior Tells Us About Business Cycles. *American Economic Review 90*(3), 458–481.

Caldara, D. and C. Kamps (2012). The Analytics of SVAR: A Unified Framework to Measure Fiscal Multipliers. *Discussion Paper, Federal Reserve Board*.

Chang, Y., A. Hornstein, and P.-D. Sarte (2009). On the Employment Effects of Productivity Shocks: The Role of Inventories, Demand Elasticity, and Sticky Prices. *Journal of Monetary Economics 56*(3), 328–343.

Chari, V. V., P. J. Kehoe, and E. R. McGrattan (2000). Sticky Price Models of the Business Cycle: Can the Contract Multiplier Solve the Persistence Problem? *Econometrica 68*(5), 1151–1180.

Christiano, L. J., C. Ilut, R. Motto, and M. Rostagno (2008). Monetary Policy and Stock Market Boom-Bust Cycles. *European Central Bank Working Paper 955.*

Copeland, A., W. Dunn, and G. Hall (2011). Inventories and the Automobile Market. *RAND Journal of Economics 42*(1), 121–149.

Dotsey, M. and R. G. King (2006). Pricing, Production, and Persistence. *Journal of the European Economic Association 4*(5), 893–928.

Faust, J. (1998). The Robustness of Identified VAR Conclusions About Money. *Carnegie–Rochester Conference Series on Public Policy 49*, 207–244.

Fisher, J. D. M. and A. Hornstein (2000). (S,s) Inventory Policies in General Equilibrium. *Review of Economic Studies 67*, 117–145.

Forni, M., L. Gambetti, and L. Sala (2014). No News in Business Cycles. *The Economic Journal 124*(581), 1168–1191.

Greenwood, J., Z. Hercowitz, and G. W. Huffman (1988). Investment, Capacity Utilization, and the Real Business Cycle. *American Economic Review 78*(3), 402–417.

Jaimovich, N. and S. Rebelo (2009). Can News About the Future Drive the Business Cycle? *American Economic Review 99*(4), 1097–1118.

Jung, Y. and T. Yun (2006). Monetary Policy Shocks, Inventory Dynamics, and Price-Setting Behavior. *Federal Reserve Bank of San Francisco Working Paper*.

Kahn, J. A. (1987). Inventories and the Volatility of Production. *American Economic Review 77*(4), 667–679.

Kahn, J. A. (1992). Why is Production More Volatile Than Sales? Theory and Evidence on the Stockout-Avoidance Motive for Inventory-Holding. *Quarterly Journal of Economics 107*(2), 481–510.

Kesavan, S., V. Gaur, and A. Raman (2010). Do Inventory and Gross Margin Data Improve Sales Forecasts for U.S. Public Retailers? *Management Science 56*(9), 1519–1533.

Khan, A. and J. K. Thomas (2007a). Explaining Inventories: A Business Cycle Assessment of the Stockout Avoidance and (S,s) Motives. *Macroeconomic Dynamics 11*(5), 638–664.

Khan, A. and J. K. Thomas (2007b). Inventories and the Business Cycle: An Equilibrium Analysis of (S,s) Policies. *American Economic Review 97*(4), 1165–1188.

Khan, H. and J. Tsoukalas (2012). The Quantitative Importance of News Shocks in Estimated DSGE Models. *Journal of Money, Credit and Banking 44*(8), 1535–1561.

Kryvtsov, O. and V. Midrigan (2010). Inventories and Real Rigidities in New Keynesian Business Cycle Models. *Journal of the Japanese and International Economies 24*(2), 259–281.

Kryvtsov, O. and V. Midrigan (2013). Inventories, Markups, and Real Rigidities in Menu Cost Models. *Review of Economic Studies 80*(1), 249–276.

Matsa, D. A. (2011). Competition and Product Quality in the Supermarket Industry. *Quarterly Journal of Economics 126*(3), 1539–1591.

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

Moon, H., F. Schorfheide, and E. Granziera (2013). Inference for VARs Identified with Sign Restrictions. *Mimeo, University of Pennsylvania.*

Mountford, A. and H. Uhlig (2008). What are the Effects of Fiscal Policy Shocks? *Journal of Applied Econometrics 24*(6), 960–992.

Pindyck, R. S. (1994). Inventories and the Short-Run Dynamics of Commodity Prices. *RAND Journal of Economics 25*(1), 141–159.

Rotemberg, J. and M. Woodford (1997). An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy. *NBER Macroeconomics Annual 27.*

Rotemberg, J. J. (1982). Sticky prices in the united states. *The Journal of Political Economy*, 1187–1211.

Schmitt-Grohé, S. and M. Uribe (2012). Whats News in Business Cycles? *Econometrica 80*(6), 2733–2764.

Stock, J. R. and D. M. Lambert (2001). *Strategic Logistics Management.* McGraw-Hill/Irwin.

Uhlig, H. (2005). What are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure. *Journal of Monetary Economics 52*(2), 381–419.

Vukotic, M. (2013). A Sectoral Approach to News Shocks. *Mimeo, University of Warwick*.

Wen, Y. (2011). Input and Output Inventory Dynamics. *American Economic Journal: Macroeconomics 3*(4), 181–212.

Woodford, M. (2003). *Interest and Prices: Foundations of a Theory of Monetary Policy.* Princeton University Press.

A Detailed equilibrium conditions of the stock-elastic demand model

A.1 List of equilibrium conditions

A market equilibrium of the stock-elastic demand model is characterized by the following set of equations:

$$\left(c_t - \psi_t \frac{n_t^{1+\xi^{-1}}}{1+\xi^{-1}}\right)^{-\sigma} = \lambda_t \tag{30}$$

$$w_t = \psi_t n_t^{\xi^{-1}} \tag{31}$$

$$\xi_t \left(1 - \phi\left(\frac{i_t}{i_{t-1}}\right) - \left(\frac{i_t}{i_{t-1}}\right)\phi'\left(\frac{i_t}{i_{t-1}}\right)\right) + \beta \mathbb{E}_t \left[\xi_{t+1}\left(\frac{i_{t+1}}{i_t}\right)^2 \phi'\left(\frac{i_{t+1}}{i_t}\right)\right] = \lambda_t \tag{32}$$

$$i_t \left(1 - \phi\left(\frac{i_t}{i_{t-1}}\right)\right) + (1-\delta_k)k_t = k_{t+1} \tag{33}$$

$$\beta \mathbb{E}_t \left[(1-\delta_k)\xi_{t+1} + \lambda_{t+1} r_{t+1}\right] = \xi_t \tag{34}$$

$$c_t + i_t = x_t \tag{35}$$

$$z_t k_t^{1-\alpha} n_t^{\alpha} = y_t \tag{36}$$

$$mc_t \alpha \frac{y_t}{n_t} = w_t \tag{37}$$

$$mc_t (1-\alpha)\frac{y_t}{k_t} = r_t \tag{38}$$

$$(1-\delta_i) inv_{t-1} + y_t = s_t + inv_t \tag{39}$$

$$s_t + inv_t = a_t \tag{40}$$

$$\mathbb{E}_t \left[\beta \frac{\lambda_{t+1}}{\lambda_t}(1-\delta_i)\frac{mc_{t+1}}{mc_t}\right] = \gamma_t \tag{41}$$

$$\frac{1}{\mathbb{E}_t\left[(1-\delta_i) q_{t,t+1} mc_{t+1}\right]} = \mu_t \tag{42}$$

$$\zeta \frac{1}{1 + \frac{inv_t}{s_t}} = \frac{\frac{1}{\gamma_t} - 1}{\mu_t - 1} \tag{43}$$

$$\frac{\theta}{\theta - 1} = \mu_t \tag{44}$$

$$s_t = x_t \tag{45}$$

Conditions (30)-(35) characterize the optimum of the household's problem, conditions (36)-(43) characterize that of the firm, and condition (45) reflects market clearing for goods. Condition (43) characterizes its optimal choice of inventory holdings, while conditions (42) and (44) characterize optimal pricing by monopolistic firms in this environment. Conditions (39) and (40) are the law of motion for inventories, and the definition of goods on shelf, respectively.

B Robustness of the results of section 6

We first check whether our result are sensitive to the long-run properties of the data. As an initial verification, we remove non business-cycle frequencies from the data by applying an HP filter to each series. In figure 14, we observe that the impulse responses are quickly mean reverting. However, impulse response patterns are similar to the main text. Moreover, our result on forecast error variances are also similar to our benchmark since in the very short run, the shock accounts for 10 percent of output variation on average, and 30 percent of that in the long run. Comparing the result with no restrictions on inventories, we see that the short run (1 quarter) output variation becomes significantly more precise with a downward shift in the mean.

Second, in our benchmark estimation, we used real GDP as a measure of output. To be consistent with our model definition of output $y = c + i + \delta inv$, we also construct an alternative output series which subtracts government spending and net exports from the GDP series. That is, the alternative output measure is nominal GDP net of government spending and net exports, deflated by the GDP deflator, expressed in per capita terms. Figures 15 and 16 again confirm that our result is not sensitive to this extension. In figure 17, we see that by imposing 2 period restrictions, the mean output variation explained by the identified shock shifts significantly downwards at all horizons.

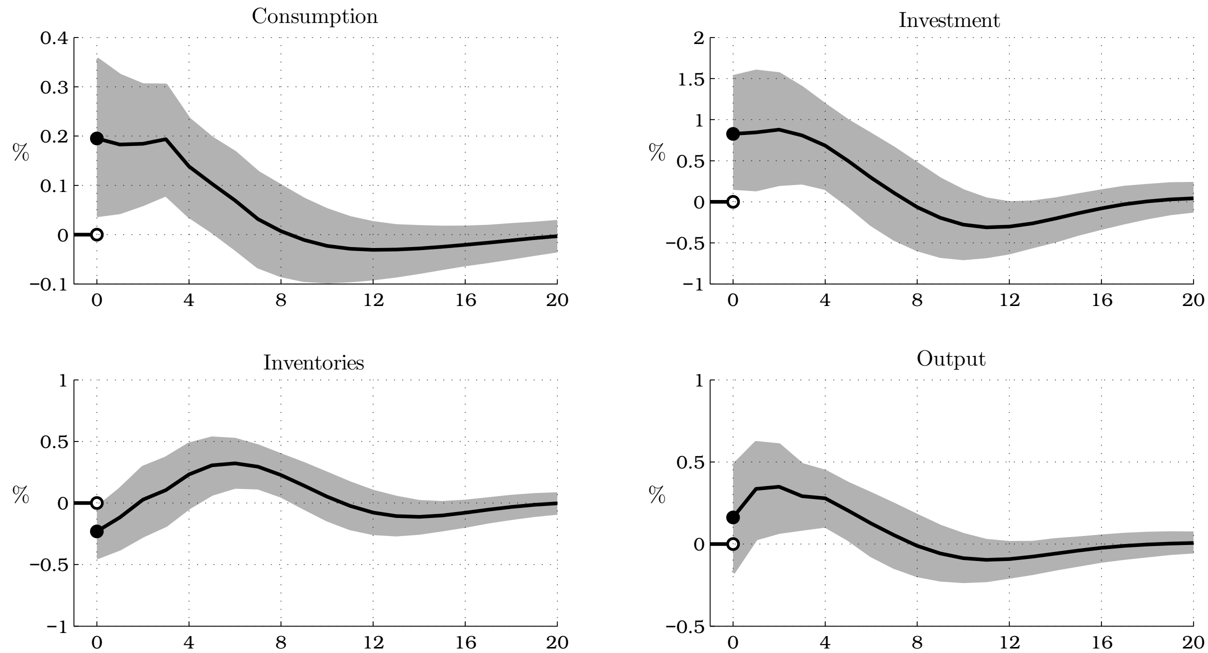

Figure 14: Robustness of impulse responses 1. Median and 80% credible set impulse responses to the shock identified using impact (1 period) restrictions for the HP filtered series.

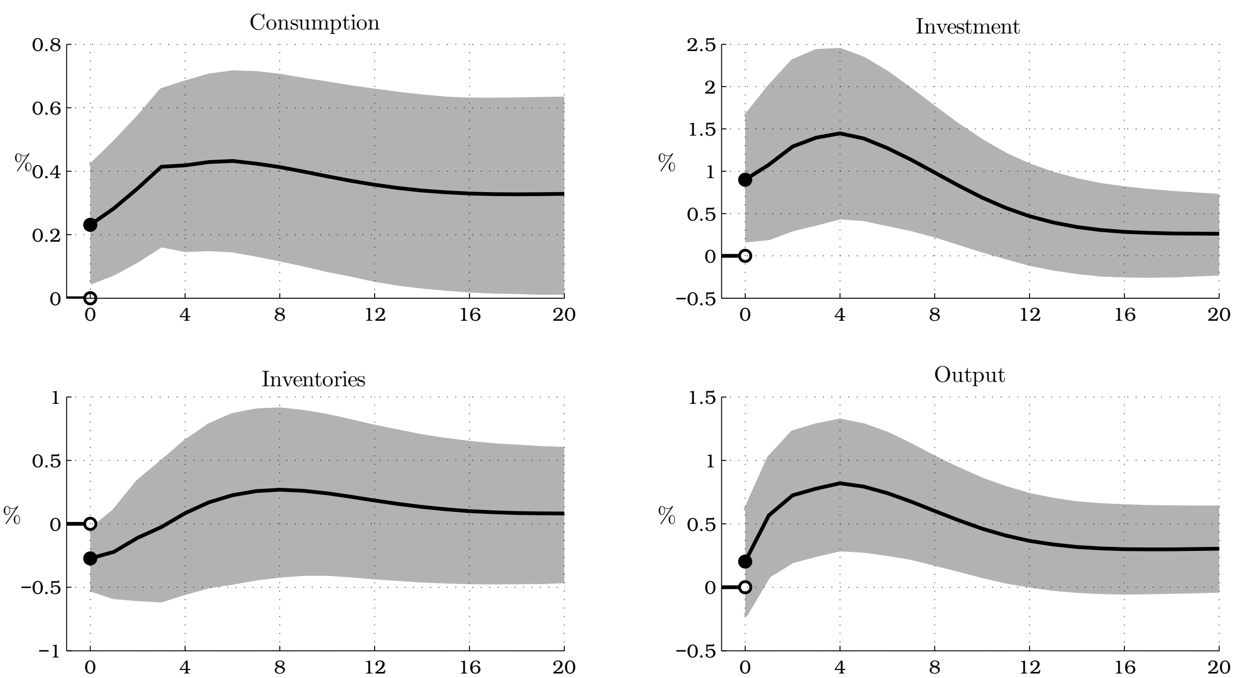

Figure 15: Robustness of impulse responses 2. Median and 80% credible set impulse responses to the shock identified using 1 period restrictions for the alternative output series (without government spending and net exports), with 1 period restrictions applied on inventories, consumption and investment.

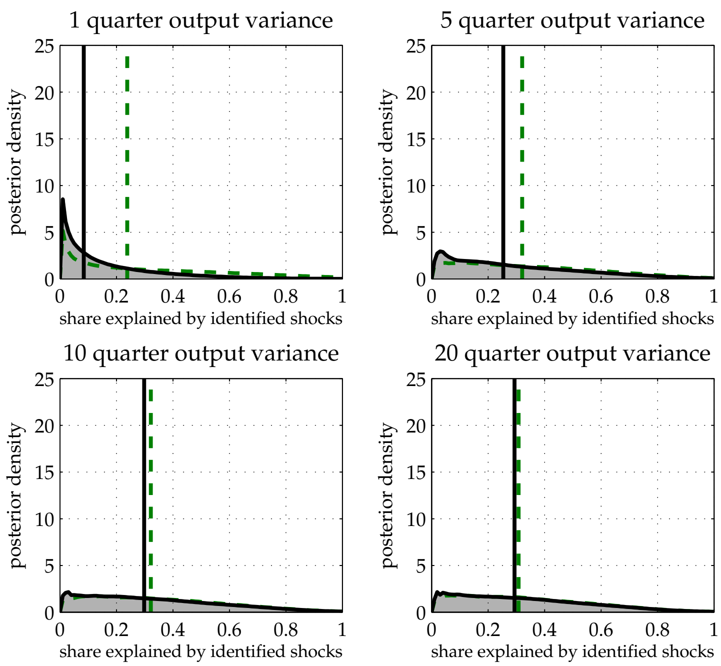

Figure 16: Robustness of forecast error variance 2. Posterior probability density and median (vertical line) for the share of forecast error variance of output at each horizon explained by shocks identified using the alternative output series (without government spending and net exports), with 1 period restrictions. Solid line: 1 period negative comovement between Δinv_t and (c_t, i_t). Dashed line: 1 period positive comovement between c_t and i_t.

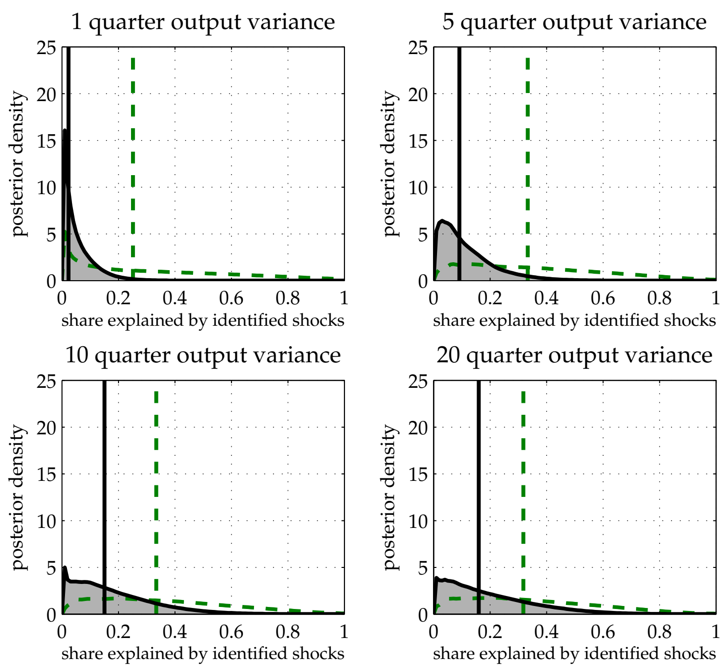

Figure 17: Robustness of forecast error variance 3. Posterior probability density and median (vertical line) for the share of forecast error variance of output at each horizon explained by shocks identified using the alternative output series (without government spending and net exports), with 2 period restrictions. Solid line: 2 period negative comovement between Δinv_t and (c_t, i_t). Dashed line: 2 period positive comovement between c_t and i_t.

Parameter	Value	Description
β	0.99	Subjective discount factor
σ	1	Household intertemporal elasticity of substitution
α_K	0.225	Capital share
α_N	0.675	Labor share
δ_k	0.025	Capital depreciation rate
u	1	Capacity utilization rate
μ^Y	1.0045	Gross per capita GDP growth rate
μ^A	0.9957	Gross investment price growth rate
G/Y	0.2	Government consumption to GDP
n	0.2	Hours
μ	1.15	Price markup
δ_i	0.025	Inventory depreciation rate
IS	0.75	Inventory-sales ratio

Table 3: Calibrated parameters.

C Details on the structural estimation of section 7

The estimation strategy is Bayesian, and mostly follows section 5 of Schmitt-Grohé and Uribe (2012). Readers should refer to that section for a detailed discussion. In table 3, we present calibrated parameters. Standard parameters follow the values used in Schmitt-Grohé and Uribe (2012), and inventory-specific parameters are discussed in the main paper.

The period of data we use is 1955Q2-2006Q4. For the measurement equations, we use the same 7 observables (output growth, consumption growth, investment growth, hours growth, government consumption growth, productivity growth, investment price growth) as in Schmitt-Grohé and Uribe (2012), where measurement errors are only allowed on output growth. On top of that, we also use the per capita real growth rate of inventories as an additional observable, with measurement errors also allowed on this series. The source of measurement error on inventories is due to different valuations in GDP computation and inventory measurement. That is, real stock of inventories in NIPA are computed by taking the average price during the period, using various valuation methods (FIFO, market value). On the other hand, inventory investment used to produce GDP is computed by the end-of-period price of inventories.[34] We allow for persistence in the measurement error for inventories.

It is important to note that adding data on inventories as an observable is not crucial to

[34]We thank Michael Cortez at the Bureau of Economic Analysis for clarifying this.

	Bayesian Estimation						
	Prior				Posterior		
Parameter	Distribution	Median	5%	95%	Median	5%	95%
$1/\xi$	Gamma	3.92	2.51	5.77	1.70	1.13	2.25
γ_h	Uniform	0.50	0.05	0.95	0.00	0.00	0.01
ϕ''_I	Gamma	3.92	2.51	5.77	9.23	7.34	10.35
δ''_k/δ'_k	Igamma	0.75	0.32	0.96	0.31	0.24	0.38
b	Beta	0.50	0.17	0.83	0.92	0.91	0.94
ρ_{xg}	Beta	0.73	0.32	0.96	0.79	0.64	0.91
ϕ''_y	Gamma	3.67	1.37	7.75	0.52	0.20	0.87
ϕ''_{inv}	Gamma	3.67	1.37	7.75	2.68	2.02	3.28
ϕ''_a	Gamma	3.67	1.37	7.75	0.65	0.21	1.13
ρ_z	Beta	0.73	0.32	0.96	0.95	0.92	0.97
σ^0_z	Gamma	1.15	0.27	3.06	0.64	0.54	0.74
σ^4_z	Gamma	0.66	0.16	1.76	0.12	0.00	0.26
σ^8_z	Gamma	0.66	0.16	1.76	0.09	0.00	0.21
ρ_{μ^A}	Beta	0.50	0.17	0.83	0.46	0.37	0.55
$\sigma^0_{\mu^A}$	Gamma	0.24	0.06	0.64	0.18	0.01	0.33
$\sigma^4_{\mu^A}$	Gamma	0.14	0.03	0.37	0.15	0.04	0.24
$\sigma^8_{\mu^A}$	Gamma	0.14	0.03	0.37	0.24	0.10	0.32
ρ_g	Beta	0.73	0.32	0.96	0.96	0.94	0.99
σ^0_g	Gamma	0.81	0.19	2.15	0.71	0.24	1.01
σ^4_g	Gamma	0.46	0.11	1.02	0.73	0.48	0.92
σ^8_g	Gamma	0.46	0.11	1.02	0.25	0.00	0.70
$\rho_{\mu^X}{}^*$	Beta	0.23	−0.18	0.46	0.30	0.17	0.47
$\sigma^0_{\mu^X}$	Gamma	0.35	0.08	0.94	0.47	0.30	0.63
$\sigma^4_{\mu^X}$	Gamma	0.20	0.05	0.53	0.08	0.00	0.18
$\sigma^8_{\mu^X}$	Gamma	0.20	0.05	0.53	0.09	0.00	0.18
ρ_ψ	Beta	0.73	0.32	0.96	0.99	0.98	1.00
σ^0_ψ	Gamma	0.92	0.22	2.45	1.42	0.90	1.80
σ^4_ψ	Gamma	0.53	0.13	1.42	1.57	1.24	1.91
σ^8_ψ	Gamma	0.53	0.13	1.42	0.31	0.00	0.84
ρ_{ζ_h}	Beta	0.50	0.17	0.83	0.19	0.11	0.30
$\sigma^0_{\zeta_h}$	Gamma	4.82	1.15	12.87	6.03	1.15	9.13
$\sigma^4_{\zeta_h}$	Gamma	2.78	0.66	7.43	6.10	1.08	8.51
$\sigma^8_{\zeta_h}$	Gamma	2.78	0.66	7.43	3.77	1.17	6.10
ρ_{z^k}	Beta	0.50	0.17	0.83	0.85	0.77	0.93
$\sigma^0_{z^k}$	Gamma	13.14	3.14	35.07	6.41	4.25	7.89
$\sigma^4_{z^k}$	Gamma	7.59	1.81	20.26	0.68	0.00	1.69
$\sigma^8_{z^k}$	Gamma	7.59	1.81	20.26	1.62	0.01	3.30
ρ_μ	Beta	0.50	0.17	0.83	0.77	0.71	0.83
σ^0_μ	Gamma	0.86	0.20	2.29	2.85	2.34	3.31
σ^4_μ	Gamma	0.50	0.12	1.33	0.63	0.26	0.93
σ^8_μ	Gamma	0.50	0.12	1.33	0.22	0.00	0.47
σ^{me}_{gy}	Uniform	0.15	0.02	0.29	0.27	0.25	0.30
ρ^{me}_{ginv}	Beta	0.50	0.17	0.83	0.21	0.09	0.33
σ^{me}_{ginv}	Uniform	0.15	0.02	0.29	0.30	0.29	0.30

Table 4: Parameter Estimation on US Data. Posterior is the result of estimation with using inventories as an additional observable. Hence 8 observable series (output, consumption, fixed investment, government spending, hours worked, TFP, investment price, inventories) are used. All numbers are rounded. A transformed parameter $\rho_{\mu^X} + 0.5$ is estimated for ρ_{μ^X}.

our estimation. Inventory investment is implicitly included in the existing observables used for estimation (output, consumption, investment and government spending) by the resource constraint (output net of consumption, investment, and government spending is inventory investment in a

Statistic	Y	C	I	N	G	TFP	A	INV
				Standard Deviations				
Data	0.91	0.51	2.28	0.84	1.14	0.75	0.41	0.88
Model	0.89	0.63	3.56	0.82	1.08	0.77	0.38	1.39
				Correlations With Output Growth				
Data	1.00	0.50	0.69	0.72	0.25	0.40	–0.12	0.44
Model	1.00	0.45	0.59	0.53	0.20	0.47	0.01	0.20
				Autocorrelations				
Data	0.28	0.20	0.53	0.60	0.05	–0.01	0.49	0.55
Model	0.39	0.39	0.75	0.21	0.02	0.06	0.46	0.80

Table 5: Model estimation result is based on posterior median estimates. The columns are output (Y), consumption (C), fixed investment (I), hours (N), government spending (G), total factor productivity (TFP), relative price of investment (A), and inventories (INV) all in growth rates.

closed economy). However, in the actual output data, net exports are also included and may potentially mask the dynamics of inventories. By directly including the stock of inventories as an observable, the inventory adjustment mechanism is likely to be more precisely estimated.

Table 4 summarizes the priors and posteriors in the model. Notice that for the priors on the standard deviations, we let the contemporaneous shock account for 75 percent of the total variance of the shocks. That is, priors are set such that news shocks account for 25 percent of the total variance.

Table 5 summarizes the predictions of the model. For standard deviations, most values are close to the data, but for fixed investment and inventories, the standard deviations are about 50 percent higher. Second, the model also predicts that inventories are positively correlated with output growth, with a correlation of 0.21. Lastly, we observe that the model autocorrelation is quite similar to the data, with hours (N) showing the most trouble, which is also discussed in Schmitt-Grohé and Uribe (2012).