Can news shocks account for the business-cycle dynamics of inventories?[*]

Nicolas Crouzet and Hyunseung Oh[†]

December 2012

Abstract

The procyclicality of inventory investment is a central feature of US business cycles. As such, it provides a test for the recent literature on news shocks, which argues that anticipated changes in fundamentals are important sources of aggregate fluctuations. We show that, in a range of inventory models, anticipated shocks to fundamentals generate booms of a peculiar kind: consumption and investment increase, but inventories fall persistently. During these booms, production and inventory investment are dominated by intertemporal substitution, as firms satisfy sales out of inventory stock and delay production until the realization of the anticipated shock. This mechanism is surprisingly difficult to overturn. We derive analytical parameter restrictions which guarantee procyclical inventory dynamics in response to news shocks, and show that standard calibrations considered in the literature do not come close to satisfying the restrictions. Furthermore, the introduction of the frictions studied by the news literature, such as variable capacity utilization and adjustment costs, is not sufficient to restore the procyclicality of inventories. We use the models' restrictions on the comovement of sales and inventories to identify news shocks in postwar US data. We find that the identified shock leads to a diffusion in TFP, but has a short implementation lag and accounts for a small fraction of long-run movements in TFP, inventories and sales.

Keywords: News shocks; Business cycles; Investment; Inventories; Intertemporal substitution; Adjustment costs.

JEL Classification Numbers: E13, E23, E37.

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 1990's as examples. Schmitt-Grohé and Uribe (2012) estimate that in the US economy, news shocks account for about half of all business-cycle fluctuations.

Of particular importance to theories of the business cycle based on news shocks is the behavior of investment. At the business cycle frequency, aggregate investment is procyclical and highly volatile. Moreover, the empirical literature on investment has documented its sensitivity to asset price changes that forecast future growth in cash flows (Fazzari, Hubbard, and Petersen, 1988). Jaimovich and Rebelo (2009) show that, in a neoclassical growth model with investment adjustment costs, variable capacity utilization, and weak wealth effects on hours worked, an expected rise in the marginal product of capital leads to a boom in investment today. Adding variable capacity and weak wealth effects on labor supply allows output to rise on impact and satisfy current demand, while adjustment costs lead firms to smooth the desired increase in the stock of capital over time and start investing today. Their suggestions carry through to other forms of investment, such as investment in material or work-in-progress inventories. So long as the corresponding stock enters the production function, as in the work of Christiano (1988), these will behave similarly as investment in productive capital.

In this paper, we study the behavior of an altogether different type of investment in response to news shocks: investment in finished-good inventories. Finished-good inventories do not affect the future marginal productivity of capital, and thus do not fall under the category described above. However, there is abundant evidence that finished-good inventories are a forward-looking variable that responds to changes in expectations about future economic conditions. For instance, Kesavan, Gaur, and Raman (2010) find that finished-good inventory data are valuable for forecasting sales.

Over the business cycle, two salient features of inventories have been extensively documented in the literature, and hold for all categories of inventories (see for example Ramey and West (1999)). First, both the stock of inventory and inventory investment are procyclical.[1] The left panel of figure 1 shows the movement of the stock of inventories during the 5 most recent NBER recessions. In all 5 recessions, inventories declined. The quarterly US postwar unconditional correlation of the cyclical component of inventory investment and

---

[1] By an accounting identity, output equals sales plus the change in the stock of inventory, net of inventory depreciation. We will refer the stock of inventory as "inventories", and the flow into this stock as "inventory investment".

output is 0.63[2]. Second, the ratio of the stock of inventory to sales, the IS ratio, is countercyclical. Although inventory investment is well-known to be volatile, it constitutes only a small fraction of the overall inventory stock. Thus, movements in the stock of inventories are typically slower than movements in final sales. The right panel of figure 1 shows the IS ratio during the most recent 5 NBER recessions. In all 5 except the 2001 recession, the IS ratio increased. The quarterly US postwar unconditional correlation between the cyclical component of the IS ratio and output is -0.35.[3]

In what follows, we ask whether, in response to news shocks, business cycle models can generate inventory fluctuations characterized by (i) procyclical inventory investment, and (ii) a countercyclical IS ratio. Our answer is a stark no.

In section 2, we start our analysis by embedding the stock-elastic demand model of Bils and Kahn (2000) into an otherwise standard Real Business Cycle model. In section 3, we use this model to show that good news about the future lead to a boom during which consumption and investment rise, but inventories fall. The intuition at the heart of our result is that news shocks lead to intertemporal substitution in production. As future marginal cost is expected to be lower than current marginal cost, firms face a downward sloping marginal cost profile. Optimal inventory investment behavior then dictates that they should delay production, and satisfy current demand by drawing down on existing inventories. Thus, news booms lead to inventory disinvestment. As we will discuss in detail, this is not the only force affecting inventories during a news boom. Absent movements in the growth rate of marginal cost, firms typically follow a target inventory to sales behavior, maintaining the ratio of inventories to sales at a constant value. Increases in sales thus lead to increases in inventories. However, we show that in response to a news shock, this force is quantitatively unlikely to be sufficient to overcome the effects of intertemporal substitution in production. We also establish that our results holds in an alternative model of inventories, the stockout-avoidance model of Kahn (1992) and Kryvtsov and Midrigan (2012).

In section 4, we show that this result extends strongly to the dynamic response of the two classes of models. In the period leading up to the realization of the anticipated increase in fundamentals (the "news" period), low expected future marginal cost generates substantial inventory disinvestment. The fall in inventories after the news shock is thus deep and protracted. In section 5, we introduce various forms of dynamic rigidities to the two classes of models, such as inventory adjustment costs and habit formation. We show that even the combination of all the rigidities we consider is insufficient to generate a procyclical response of inventories

---

[2]The cycle-trend decomposition is done using a Hodrick-Prescott filter with smoothing parameter set at $1,600$. Output is in logs. However, inventory investment cannot be expressed in logs, since it occasionally takes negative values. Instead, we divide inventory investment by the level of the trend output and filter the series. Details on the data used are consigned to appendix A.

[3]Other studies that report a stronger negative correlation focus on the manufacturing or retail sector, whereas we use the broader category of private sales of final goods to compute the ratio.

to news shocks.

Having established that the negative comovement of inventories and sales is specific to news shocks, we propose to use this structural prediction as a means to identify news shocks. In section 6, we describe an empirical strategy based on this idea, and we show that the shock identified in this fashion in postwar US data leads to a future increase in TFP and an immediate and persistent increase in sales. However, the decline in inventories in response to the shock is very short, at most 2 quarters, implying that the "news" period is short. Furthermore, forecast error variance decompositions suggest that the shock accounts for at most 26% of long-run movements in TFP and inventories, and up to 39% in sales. This contrasts with the empirical literature on news shocks, which finds that they account for a substantial fraction of the movements in TFP, especially in the long run. Section 7 concludes.

2 The stock-elastic demand model

In this section, we lay out 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 these models is the assumption 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. This may arise either because of greater availability of goods, or because more on-shelf goods may provide a wider variety within the same product. For example, a shoe store with more colors and size of all kinds are likely to attract more customers and sell more goods. Cachon and Olivares (2009) find empirical evidence supporting the view that greater product variety leads to higher sales at the industry level.

2.1 Description of the 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. 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):

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

where ξ is the Frisch elasticity of labor supply and σ denotes the inverse of the elasticity of intertemporal substitution in consumption. This preference specification has been widely used in the literature on news shocks, and it implies zero wealth effects on labor supply. None of our results depend crucially on the choice of preferences, and in fact we view the GHH specification as a best-case scenario for the procyclicality of inventory investment, since it limits any fall in output due to negative wealth effects of the news shock on the household labor supply curve.

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

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

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

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

Equation (2) is the household budget constraint. The household earns income each period by providing labor n_t at a given wage W_t, renting capital k_t at a rate R_t, claiming the 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 nominal bonds B_{t+1}. The probability-adjusted price of each of these bonds is $Q_{t,t+1}$, for each state in period $t+1$.

Equation (3) is the accumulation rule of capital in the presence adjustment costs to investment. The adjustment cost function $\phi(\cdot)$ is twice-differentiable with $\phi(1) = \phi'(1) = 0$, and $\phi''(1) > 0$. Following Jaimovich and Rebelo (2009), adjustment costs of this form guarantee that firms start accumulating capital ahead of the implementation period of the news shock, so that investment increases when good news about the future are announced.

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

of the form:

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

where $v_t(j)$ is the taste-shifter for each product j and θ is the elasticity of substitution across goods. It follows from expenditure minimization that 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 stock-elastic demand models, the taste-shifter for variety j is assumed to depend on the amounts 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 a no-Ponzi condition for both capital and stage-contingent nominal bond holdings.

Firm problem Each monopolistically competitive firm $j \in [0,1]$ maximizes the following 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 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 nominal bonds 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 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 current production $y_t(j)$. The parameter δ_i denotes the depreciation rate of inventories. Equation (10) states that on-shelf goods that is unsold is accounted as inventories. Equation (11) is the production function. Firms use a constant returns to scale production function, with capital and labor as inputs. The parameter α governs the degree of short-run decreasing returns to scale in labor. The variable z_t represents total factor productivity and is exogenous. Finally, monopolistically competitive firms face the demand function (12) stemming from solving the household problem.

Market clearing Labor and capital markets clear, and the net transaction of nominal bond is zero:

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

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

$$B_{t+1} = 0. \tag{15}$$

Sales of goods for each variety j also clears by the demand function described above. The average on-shelf goods in the economy a_t is defined by

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

2.2 Equilibrium

A market equilibrium of this economy is defined as follows.

Definition 1 (Market equilibrium of the stock-elastic model) *A market equilibrium of the stock-*

elastic model is a set of stochastic processes:

$$c_t, n_t, k_{t+1}, i_t, B_{t+1}, x_t, a_t, \{a_t(j)\}, \{n_t(j)\}, \{k_t(j)\}, \{v_t(j)\}, \{s_t(j)\}, \{y_t(j)\}, \{inv_t(j)\}, \{p_t(j)\}, W_t, R_t, P_t, Q_{t,t+1}$$

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 no-Ponzi conditions,*

- *each firm $j \in [0, 1]$ maximizes (7) subject to (8) - (12),*

- *markets clear according to (13) - (16).*

Note that the two exogenous processes in our economy are total factor productivity z_t and labor wedge ψ_t. We limit ourselves to these shocks following the results of Schmitt-Grohé and Uribe (2012). They find that among all anticipated shocks, it is anticipated shocks to these variables that are the primary contributors to aggregate fluctuations. It will nevertheless be clear that at least on impact, our results hold for anticipated innovations to all other exogenous variables studied in Schmitt-Grohé and Uribe (2012) (investment-specific shocks, government spending shocks and preference shocks).

2.3   The optimal choice of inventories

The full set of equilibrium conditions are provided in appendix B. As we show there, a market equilibrium of the stock-elastic model is symmetric, so that $a_t(j) = a_t$, $s_t(j) = s_t$, $inv_t(j) = inv_t$, $y_t(j) = y_t$, and $p_t(j) = p_t$ for all j. Here, we discuss the optimal stock choice of firms.

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

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

Using this, the optimal stock choice of firms is governed by the equation:[4]

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

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 benefits of adding this extra unit. First, there is a convenience yield: by producing and stocking an extra unit, the firm is able generate an additional fraction $\frac{\partial s_t}{\partial a_t}$ of sales. Second, the remainder of the extra unit,

---

[4]Here, $q_{t,t+1} = Q_{t,t+1}P_{t+1}/P_t$ denotes the real stochastic discount factor of the household.

$1 - \frac{\partial s_t}{\partial a_t}$, is not sold and instead kept as inventory until the next period, when it contributes to reducing future production costs.[5]

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 the convenience yield of inventories.

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 them.

3.1   A reduced-form framework

We analyze a first-order log-linear approximation of the model around its steady-state. The log-deviation of a variable x_t from its steady-state value is denoted by $\hat{x}_t$.

Proposition 1 (Stock-elastic model) *On impact and with only news shocks, so that $\hat{z}_t = 0$ and $\hat{\psi}_t = 0$,*

---

[5]Because the shadow value of the stock of inventories, the Lagrange multiplier on equation (9), is equal to the marginal cost of production, inventories are never zero in this economy. This will not be the case with idiosyncratic production costs in the stockout avoidance model.

the following reduced-form equations represent the log-linearized market equilibrium of definition 1:

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

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

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

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

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

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

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

This framework will allow us to derive restrictions on reduced-form parameters that guarantee a positive response of inventories to news shocks. We first discuss briefly the reduced-form equations (20)-(24), in order to gain intuition about the meaning of the reduced-form parameters.

Equation (20) relates marginal cost to output. With $\omega > 0$, this 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. The literature often refers to this parameter as the degree of real flexibilities (or the inverse of the degree of real rigidities). 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 range of values for the parameter ω is thus:

$$\omega \in [0.3, 3].$$

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. In its log-linearized form, κ in (21) denotes the steady-state output to sales ratio. In NIPA, the time series average of inventory investment over total output is around 0.5 percent, so that a reasonable range of values for κ is:

$$\kappa \in [1, 1.01].$$

Equation (22) is the log-linearized form of the first-order optimality condition for inventories in (19). It first states that, holding other factors fixed, inventories $\hat{inv}_t$ and sales $\hat{s}_t$ move one to one. That is, firms follow a target inventory to sales ratio behavior, absent movements in prices and costs. Second, when markups $\hat{\mu}_t$ fall, sales increases, so that given a production level, inventories fall. The parameter $\tau > 0$ controls the intensity of this markup channel. Third, inventories respond to changes in the discounted growth rate of marginal cost $\hat{\gamma}_t$. When $\hat{\gamma}_t > 0$, future production is expensive relative to current production, firms bunch their production today, and inventories increase. The parameter $\eta > 0$ controls the intensity of this intertemporal substitution channel. Its value depends closely on the intertemporal costs of holding inventories: (i) the opportunity cost, governed by the discount rate β, and (ii) the storage cost, governed by the depreciation rate δ_i. Equation (27) indeed indicates that, since $IS > 0$:

$$\eta > \frac{1}{1 - \beta(1 - \delta_i)}.$$

With a quarterly discount factor $\beta = 0.99$ and depreciation rate $\delta_i = 0.02$, this suggests a lower bound of:

$$\eta > 33,$$

irrespective of the value of the IS ratio. That is, a 1 percent increase in the present value of future marginal cost will lead firms to accumulate more than 33 percent of inventories relative to sales, as expressed in (23).

Since we do not assume nominal rigidities, the firms set a constant markup of price over future marginal cost, as in (23). Consequently, the value of $\tau > 0$ is irrelevant to our analysis. Lastly, equation (24) is a consequence of 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 positive news shock has no impact on current fundamentals ($\hat{\psi}_t = 0$ and $\hat{z}_t = 0$), but it increases sales ($\hat{s}_t > 0$).

Proposition 2 (The impact response of inventories to a good news about the future) *After a positive news shock,*

1. *the IS ratio falls;*

2. *inventories increase, if and only if:*

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

The first part of this proposition is encouraging, since the inventory stock to sales ratio is countercyclical at business cycle frequencies. The second part of this proposition indicates that inventories respond positively to a news shock when ω is small and κ and η are large. Following our discussion on numerical values, a conservative upper bound on κ/ω is 3.36, which corresponds to a large degree of real rigidities ($\omega = 0.3$) along with a large output to sales ratio ($\kappa = 1.01$). On the other hand, our lower bound for η is 33, so that:

$$\eta > 33 \gg 3.36 \geq \frac{\kappa}{\omega}.$$

Therefore, the condition of the second part 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.

3.3 Discussion

Proposition 2 indicates that even under conservative calibrations of our reduced-form framework, inventories should fall in response to good news about future fundamentals. In order to understand this result, it is useful to separate the positive and negative effects that news shocks have on inventory investment.

On the one hand, a news shock leads to a boom in output and sales, which is partly channelled into positive inventory investment. Two separate mechanisms contribute to this. First, if sales were unchanged, an increase in output will increase inventory investment by the accounting identity linking output, sales and inventory investment. Second, since firms follow a target inventory-sales ratio behavior in the absence of price movements, an increase in sales will lead to a build-up of inventories. Because of these two mechanisms, the shock partially works to increase the stock of inventories.

On the other hand, the shock generates a downward sloping profile for the marginal cost. This is because either a future improvement in TFP, or a future fall in the labor wedge reduces future marginal cost relative

12

to current marginal cost. This induces intertemporal substitution in production: firms postpone production, lower their desired inventory-sales ratio, and satisfy current demand by drawing down on their inventories.

Figure 2 spells this argument in further detail. The left panel plots the inventory law of motion (LOM) and inventory optimality (IO) schedules, when $\hat{inv}_{t-1} = 0$:

$$\hat{inv}_t = \delta_i \frac{1}{\kappa - 1} \left(\kappa \hat{y}_t - \hat{s}_t \right), \tag{LOM}$$

$$\hat{inv}_t = -\eta \omega \hat{y}_t + \hat{s}_t, \tag{IO}$$

which are obtained from (20) and (21)-(24) respectively. The right panel plots similar schedules for the inventory stock to sales ratio. Shifts in those schedules are proportional to the increase in sales $\hat{s}_t > 0$ resulting from the news shock. Note that the slope of (IO) increase in absolute value with η, the elasticity of intertemporal substitution in production, and ω, the inverse of the degree of real rigidities. On the one hand, the increase in sales results in an upward shift of (IO), as, all other things equal, firms increase inventories proportionately to their sales. This shift captures the expansionary effects of the boom on inventory investment: absent any shifts in (LOM), this would lead to an increase in inventories. Moreover, the increase in inventories would be larger, the flatter the schedule. However, the rise in sales also shifts (LOM) downward - inventories cannot increase more than the net increase in output minus the net increase in sales. The shift is in general larger than the shift in the (IO) schedule, since $\frac{\delta_i}{\kappa - 1} = \frac{1}{IS} \geq 1$ when $IS \leq 1$. Thus, typically - and especially so if the schedule (IO) is steep, as suggested by our discussion - inventories fall after the shock.

When could inventories potentially respond positively to news shocks? The previous discussion suggests that the slope of the schedule (IO) would have to be small, so that one or both of the following conditions have to be met: (i) the degree of real rigidities is large, or equivalently, ω is small; (ii) the elasticity of intertemporal substitution in production is small, which occurs when intertemporal costs of holding inventories are large. The advantage of our proposition is that it quantifies how large real rigidities and intertemporal costs of holding inventories should be. We now turn to discussing these analytical bounds.

How large should real rigidities be? Substituting the structural relation given in proposition 2 for the reduced-form parameters η and κ, the following inequality must be satisfied for inventories to behave procyclically:

$$\omega \leq \overline{\omega} \equiv \frac{IS}{1 + IS}(1 - \beta(1 - \delta_i))(1 + \delta_i IS).$$

In other words, real rigidities must be higher than $1/\overline{\omega}$ for inventories to comove with sales. This has an intuitive interpretation: real rigidities tend to dampen movements in marginal cost; for a given elasticity of intertemporal of inventories to the growth rate of marginal cost, smaller movements in relative marginal cost imply a smaller degree of intertemporal substitution. Hence large real rigidities are needed to weaken the intertemporal substitution in production and avoid the fall in inventories after the news shock. How large? Assuming a standard value of $\beta = 0.99$ at quarterly frequency, $\overline{\omega}$ is a function of IS and δ_i. In figure 3, we plot $\overline{\omega}$ as a function of δ_i for three values of the inventory-sales ratio: $IS = 0.25$, $IS = 0.50$, and $IS = 0.75$. The message from figure 3 is that, regardless of the value of the IS ratio, even with a very large depreciation rate of inventories, the degree of real rigidities needed to achieve the positive response of inventories is very large. For example, for $IS = 0.75$, when we assume that 10 percent of inventories depreciates each quarter, the upper bound on ω is still as low as $\overline{\omega} = 0.05$, roughly a sixth of the value of 0.33 suggested by Dotsey and King (2006).

How large should intertemporal costs be? The reason why inventories are not procyclical with news shocks is that firms that produce to stock have a large intertemporal elasticity of substitution in production η. The elasticity of intertemporal substitution in production in turn depends importantly on the intertemporal costs of holding inventories, summarized by the opportunity cost $1/\beta - 1$ and storage cost δ_i. The steady-state growth rate of marginal cost, $\gamma = \beta(1 - \delta_i)$, is thus inversely related to this intertemporal cost. When both the opportunity cost and the storage cost are set to 0, γ attains its maximum value at 1. When $\gamma = 1$, the value η goes to infinity. The right panel of figure 2 shows the value of η as a function of γ. Even with large intertemporal costs of holding inventories ($\gamma = 0.9$, corresponding to $\delta_i = 0.1$ and $\beta = 0.99$), η is above 5. To reach the bound of $\frac{\kappa}{\omega} = 3.3$ suggested by our conservative calibration, the rate of depreciation of inventories would have to be at least $\delta = 0.15$ at the quarterly frequency, which in turn would give rise to counterfactually high steady-state IS ratios.

We thus conclude that parametrizations of this model based on a wide range of targets for the inventory-sales ratio lead to negative responses of inventories to a news shock, unless one is willing to commit to very large intertemporal costs of holding inventories along with a degree of real rigidities above the range commonly discussed in the literature.

3.4 News shocks in the stockout-avoidance model

A natural question is whether our results are specific to the particular inventory model we have chosen to analyze. Another branch of the literature on finished-good inventories motivates the existence of positive

output inventory stocks by the presence of uncertainty in demand, and the implied possibility of stocking out. In these models, firms are assumed to have imperfect information on the demand schedule for their variety at the time they make production and pricing decisions. When 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.[6]

In an appendix separate from the paper, we study the effects of news shocks in this class of models in detail. We show that a reduced-form framework similar to that of proposition 1 obtains, and moreover that our main result carries through: in response to good news about future fundamentals, under standard calibrations of the model, the IS ratio falls but inventories fall as well. We again obtain analytical restrictions on reduced-form parameters to precisely quantify the conditions under which this result holds. Additionally, we argue that, as in the stock-elastic demand model, the main mechanism dominating the response of inventories to anticipated shocks is intertemporal substitution in production. In both classes of models, a first order condition of the form of (19) holds, so that inventory decisions balance convenience yields, markups and intertemporal substitution motives. Our results thus indicate that the fact that the intertemporal substitution motive is quantitatively the stronger one for news shocks is a feature common to a range of inventory models, and does not depend on the precise micro-foundation for the convenience yields.

4   Dynamic responses

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 output, sales and inventories. We now turn to the dynamic response to the anticipated shock. We show that the negative response of inventories is persistent: it extends throughout the anticipation period, and even after the implementation of the shock.

4.1   Calibration

We calibrate the stock-elastic and stockout avoidance models at the quarterly frequency. The numerical values for the parameters are summarized in table 1. Standard model parameters are calibrated using estimates from the business cycle literature. Parameters specific to the inventory blocks of the models are calibrated to match sample averages of the IS ratio and the output to sales ratio. In particular, our

---

[6]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, 2012), and Wen (2011) for detailed analysis of the properties of this class of models.

calibration results in an annual rate of depreciation of inventories of 4.17%. Details on the calibration are given in appendix C.

For the stock-elastic demand model, our calibration implies that $\eta = 146$, $\omega = 0.562$ and $\kappa = 1.0053$, so that applying proposition 2, inventories respond negatively to news shocks on impact. For the stockout avoidance model, our calibration likewise implies a negative inventory response given by the proposition in the online appendix.

4.2   Impulse responses to news shocks

We first study the impulse responses of output, sales, inventories and the IS ratio to 4-period news shocks to TFP and labor wedge. The realized TFP process is AR(1) with the persistence $\rho_z = 0.99$. For the labor wedge process, the realized process is AR(1) with persistence $\rho_\psi = 0.95$.[7]

Figure 9 reports the impulse responses for the stock-elastic demand model. Note first that sales rise continuously, from the impact period to the materialization of the anticipated shock. This increase is due to a combined increase in consumption and investment, the former because of the wealth effect associated and the latter because of the presence of investment adjustment costs.

In line with our discussion of the previous sections, inventories fall. The fall is large and persistent, and reaches its through in the period preceding the materialization of the anticipated shock. At the same time, output remains mostly unchanged until period 5, when the anticipated shock materializes, so that the increase in sales associated to the shock is almost entirely met by inventory disinvestment. To build further intuition for the responses of inventories, note that the optimal labor supply and demand schedules in an economy with inventories is:

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

so that marginal cost is given by:

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

News about an increase in future productivity $(\hat{z}_{t+T})$ or a future decrease in the labor wedge $(\hat{\psi}_{t+T})$ lead to a decline in future marginal cost. This is reflected in the negative impulse response of the expected discounted marginal cost, which we report in figure 6. Through equation 22, this fall in expected discounted marginal cost leads to a fall in inventories which is sufficient to overcome the effect of the increase in sales. Thus, the

---

[7]These estimates of persistence are close to the empirical findings in the literature.

qualitative impulse responses to both types of shocks are similar.

Note that there still is a very small increase in output during the first four periods, which might seem somewhat puzzling, given that capital is fixed in the short run, and productivity is unchanged, so that the labor demand schedule of firms should not shift. While this logic holds in a flexible price model without inventories, it does not hold in a model with inventories. Indeed, in contrast to models without inventories, the optimal pricing policy of firms does not imply that marginal cost is fixed — instead, it is expected discounted marginal cost that is constant, and equal to the inverse markup. 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. The mechanism is somewhat similar to the effects of variable capacity utilization in a flexible price model without variable capacity utilization of the type studied by Jaimovich and Rebelo (2009). The effects, though, are smaller, because the current increase in marginal cost in response to the shock has more limited effects on labor demand than increases in capacity utilization brought about by a fall in the interest rate. We consider in the next section the effects of introducing variable capacity utilization into our models.

Figure 5 reports the impulse responses for the stockout-avoidance model. They are qualitatively similar to the impulse response of the stockout-avoidance model, with sales increasing gradually, inventories falling and output remaining almost constant until the materialization of the shock. There are three noticeable differences from the stock-elastic demand model. First, the fall in inventories is delayed and less persistent, as inventories return to their steady-state value ten periods after the shock. Second, output falls slightly in the periods between the announcement and the materialization of the shock. Third, the increase in sales is more gradual. Thus, the effects of intertemporal substitution in production on inventories are somewhat mitigated – though far from enough to generate a positive response of inventories to the shock. The differences between the models are due to the procyclical markup movements in the stockout avoidance model.[8] The markup increases, and marginal cost falls on impact, albeit by a small magnitude. This has the effect of shifting labor demand curve inwards, thus reducing hours worked and output. The increase in markup also limits the increase in consumption and investment, which results in a more gradual increase in sales.

4.3   Do suprise shocks generate comovement?

While anticipated shocks generate persistent negative comovement between inventories and sales, one may wonder whether this also occurs after surprise innovations to fundamentals. The impulse responses reported in figures 7 and 8 show that this is not the case. Inventories, sales and output all increase in

---

[8]In the stock-elastic demand model, under flexibles prices, markups are constant. In the stockout avoidance model, markups are not constant even under flexible prices; see our separate appendix on the stockout avoidance model for details.

response to suprise innovations to TFP and the labor wedge. The short-run response of the IS ratio is also consistent with its observed countercyclicality at business cycle frequencies, in line with the findings of Khan and Thomas (2007) and Wen (2011).[9] The models' predictions are thus broadly consistent with the observed behavior of inventories and sales over the business cycle. Thus, in the two models, the negative comovement of inventories and sales is an identifying feature of anticipated shocks to fundamentals.

5 Resolving the comovement problem

The dynamic responses in section 4 show that the additional sales generated from news shocks are satisfied by inventory depletion rather than more production. The key channel to this result is the strong intertemporal substitution motive induced by anticipations of low marginal costs in the future. In this section, we investigate whether this mechanism may be offset by allowing through variable capacity utilization, or forcing through adjustment costs, inventories, sales and output to increase in the short-run. We focus on news to TFP for clear exposition although similar results hold with news to the labor wedge.

5.1 Variable capacity utilization

Since capital is fixed in the short run, current production can only increase with a news shock by an increase in labor. The small response of output in figure 9 indicates that labor barely moves on impact. To overcome this problem, Jaimovich and Rebelo (2009) assume that the production function depends on capacity utilization, but that higher utilization bears higher depreciation of capital. Denoting u_t as the utilization of capital at period t, the production function and the law of motion for capital are modified respectively as follows:

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

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

where $\delta'(\cdot) > 0$ and $\delta''(\cdot) > 0$. With adjustment costs, investment increases on impact to smooth the future increase in investment due to the higher level of future productivity. In a model without inventories as in Jaimovich and Rebelo (2009), this leads to a fall in the marginal value of installed capital relative to the marginal value of income, since with a higher level of investment, it is less costly to replace existing capital.

---

[9]The countercyclicality of the IS 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 demand 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 IS ratio has motivated Kryvtsov and Midrigan (2012) to investigate the ability of countercyclical markup movements to mute inventory increases in response to demand-side shocks. However, in response to both TFP and demand shocks, the key business-cycle feature of inventories which news shocks fail to generate — their procyclicality — is robust: it holds regardless of the values of the persistence parameters ρ_χ and ρ_z.

Due to this fall in value, it becomes efficient to increase capital utilization today. By the same logic, one would expect capacity utilization to also increase in a model with inventories, thus allowing for a larger increase in output and a smaller fall in inventories.

In figures 9 and 10, we plot the impulse responses for the models with and without variable capacity utilization. Impulse responses are mostly unchanged: output stays the same and inventories still fall. Capacity utilization barely responds to the anticipated shock.

This puzzling neutrality of capacity utilization to anticipated shocks comes directly from the role of inventories in the economy. Note that both the marginal value of installed capital and the marginal value of income fall with the shock. In a model without inventories, the household cannot absorb more than what is currently produced. Thus, the fall in the marginal value of income is small, relative to the fall in the marginal value of installed capital, leading to a large increase of utilization. On the other hand, in a model with inventories, current sales increase more than current production, through the depletion of inventories. This implies a larger fall in the marginal value of income. In turn, the relative fall in the marginal value of installed capital is smaller, and utilization rises by less. Quantitatively, the effect on the marginal value of income is sufficient to essentially eliminate any significant rise in capacity utilization.

5.2 Adjustment costs to output and inventories

To increase current production and reduce the fall in inventories, we next introduce adjustment costs. We consider three possible types of adjustment costs: adjustment costs to inventories, output and on-shelf goods. Adjustment cost to inventories penalizes immediate inventory depletion and thus weakens the intertemporal substitution motive. Adjustment cost 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 cost to goods on shelf are 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 are 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 figures 11 and 12, we show the responses of each model with and without adjustment costs, where output adjustment cost is assumed for the stock-elastic demand model and on-shelf good adjustment cost is assumed for the stockout-avoidance model.[10] 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.

As discussed in detail by Jaimovich and Rebelo (2009), adjustment costs to investment are essential in order to generate a positive response of investment in capital goods in response to news shocks. With this form of adjustment cost, investment decisions depend solely on the discounted sum of future marginal values of capital, or future Tobin's Q. Anticipated shocks affect the marginal productivity of future capital, and thus raise future Tobin's Q, which directly translates into an increase in current investment.

This is not so for inventory investment decisions. The marginal value of inventory stock, the "inventory" Q depends on two factors: the sales yield of inventories, and the growth rate of marginal cost. Importantly, lower future marginal cost makes the inventory stock less, not more valuable in the future. Therefore, while adjustment costs to have some frontloading effects on the stock of inventories because of anticipations of high future sales, they are more limited than in the case of investment in capital goods. The intertemporal substitution channel remains strong, and inventories fall, although to a lesser extent than in the absence of adjustment costs.

5.3 Habits to consumption

Since the intertemporal substitution channel is hard to overcome from the production side, we next turn to reducing the initial response of current sales with respect to news shocks. Following Schmitt-Grohé and Uribe (2012), we assume that there is internal habit persistence in the period utility function, so that period utility is given by:

$$U(V_t) = \frac{V_t^{1-\sigma} - 1}{1 - \sigma},$$

$$V_t = c_t - bc_{t-1} - \psi_t \frac{n_t^{1+\xi^{-1}}}{1 + \xi^{-1}},$$

where b is the habit persistence parameter. By setting $b > 0$, households refer to their previous consumption level when setting their current consumption. This implies that with regards to a news shock, consumers will not abruptly change their consumption to a higher level. Figures 13 and 14 depict the obtained impulse

---

[10]Note that in the stockout-avoidance model, the specified form of adjustment cost for inventories cannot be imposed due to inventories occasionally becoming 0 at the firm-level.

responses for $b \in \{0, 0.4, 0.8\}$. We observe that the initial consumption and sales responses are smaller with higher degrees of habit persistence. However, output remains similar in all three cases. Therefore, in the presence of habit persistence, the fall in inventories is muted, but not sufficiently so to overcome intertemporal substitution in production.

5.4   Combining all the elements

We have shown that the above three resolutions do not work separately. However, there is still hope that combining the three elements will generate the positive comovement of inventories. By adding a production smoothing motive on the firm side, firms will start producing more today. At the same time, by reducing the initial consumption response through habit persistence, the fall in the marginal value of consumption will be smaller. In that case, the value of installed capital relative to the value of consumption will now decline, leading to a higher utilization of capital and hence even more production.

Figures 15 and 16 show the response with all three elements considered above included. There is no improvement in the stockout-avoidance model, while inventories are close to not moving in the stock-elastic demand model. Hence the three elements combined together still do not generate a significant positive comovement of inventories. Moreover, in the stock-elastic demand model, adding all these elements leads to inventories becoming a barely moving variable over the business cycle, counter to the high volatility of inventory investment documented in the literature.

6   What do inventories tell us about news shocks? An empirical investigation

Our analysis of inventory models suggests that the negative comovement of inventories and sales is a defining feature of anticipated shocks to fundamentals, be it TFP or the labor wedge. Indeed, as we have discussed at length, it holds for all plausible calibrations of the models, and moreover, under those calibrations, suprise shocks to fundamentals generate positive comovement between sales and inventories. In this section, we use this structural restriction to identify anticipated shocks to fundamentals in VAR framework.

6.1   Empirical strategy

Our empirical strategy identifies anticipated shocks from restrictions on impact matrices, obtained from the estimation of a reduced-form VAR in levels. Our VAR includes three observables: TFP, inventories and

sales. Formal tests for cointegration indicate that these series are integrated of order 1, at least. However, a specification of the VAR in levels produces estimates of impulse response functions that are robust to cointegration of an unknown form, and moreover since we focus on impact restrictions, a VECM specification is not necessary to identify structural impulse responses.[11] Formally, we estimate:

$$y_t = d + B(L)y_{t-1} + v_t, \quad v_t \sim N(0, \Omega),$$

where $B(L)$ is a matrix polynomial, $y_t = [\log(TFP_t), \log(Inv_t), \log(S_t)]'$, d is a 3×1 vector, Ω is a 3×3 positive definite matrix, and the errors v_t are independent and identically distributed. Given an estimate of Ω, denoted $\hat{\Omega}$, all the possible impact matrices identifying structural shocks are the invertible matrices A which satisfy:

$$\hat{\Omega} = AA'.$$

Our empirical strategy is based on the results of the previous sections. Namely, news shocks are *the only shock* to satisfy the following three restrictions:

(i) they are orthogonal to contemporaneous unforecastable innovations to TFP;

(ii) they generate an increase in sales on impact;

(iii) they generate a fall in inventories on impact.

Restriction (i) is a zero impact restriction, while restrictions (ii)-(iii) are sign restrictions. As described in Moon et al. (2011), these sign restrictions do not point-identify the impact matrix A, and therefore the impulse responses. Instead, they identify a set of impact matrices, each of which satisfies (i)-(iii). However, for our three-variable system, the set of impact matrices satisfying our combination of impact and sign restrictions can be characterized analytically, and is given in the lemma below, whichs follows from the work of Moon et al. (2011).

Lemma 2 (The set of identified impact matrices.) *Let C denote the unique Cholesky decomposition of $\hat{\Omega}$, and let e and $f > 0$ denote its $(3, 2)$ and $(3, 3)$ elements, respectively. For $\theta \in [0, 2\pi]$, define the proper rotation matrix $\overline{U}(\theta)$ by:*

$$\overline{U}(\theta) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos(\theta) & -\sin(\theta) \\ 0 & \sin(\theta) & \cos(\theta) \end{pmatrix}$$

---

[11]See chapter 9 of Lutkepohl (2007).

When $e > 0$, the set of impact matrices satisfying restrictions $(i) - (iii)$ is given by:

$$\mathcal{J}^+\left(\hat{\Omega}\right) = \left\{ A(\theta) = CU(\theta) \quad \Big| \quad \theta \in \left[0, \overline{\theta}\right], \quad \overline{\theta} = \cot^{-1}\left(\frac{e}{f}\right) \in \left[0, \frac{\pi}{2}\right] \right\}.$$

When $e < 0$, the set of impact matrices satisfying restrictions $(i) - (iii)$ is given by:

$$\mathcal{J}^-\left(\hat{\Omega}\right) = \left\{ A(\theta) = CU(\theta) \quad \Big| \quad \theta \in \left[\underline{\theta}, \frac{\pi}{2}\right], \quad \underline{\theta} = \cot^{-1}\left(-\frac{f}{e}\right) \in \left[0, \frac{\pi}{2}\right] \right\}.$$

In this result, we have adopted the convention that the column number 3 of the impact matrix corresponds to the impact effect of the identified news shock.

The proof of this lemma is in appendix D. Intuitively, this lemma indicates that we exclude the possibility that the remaining shock shares the same feature as our identified shock (zero impact on TFP and negative impact comovement of inventories and sales) which may happen for a range of estimates of C. This ensures that our scheme identifies a unique shock, which we call a news shock.

An important feature of our identification strategy is that it will identify a combination of anticipated changes in fundamentals not restricted to TFP. Indeed, in our model, an anticipated change in the labor wedge satisfies the restrictions (i)-(iii) and should therefore be recovered by the procedure. This is in contrast to other strategies proposed to identified anticipated shocks, such as Beaudry and Portier (2006), Beaudry and Lucke (2010) and Barsky and Sims (2011), that focus exclusively on anticipated shocks to TFP. The recent evidence in Schmitt-Grohé and Uribe (2012) suggests that a broader class of news shocks could play a role in the business cycle, and our method allows for this possibility.

Since our impact and dynamic model results suggest that anticipated shocks have qualitatively the same effects on sales and inventories, we view the results of our identification scheme as an upper bound on the joint contributions of all anticipated shocks on these variables. We now turn to providing Monte-Carlo evidence to support this claim.

6.2   Monte-Carlo evaluation

We test our identification strategy by Monte-Carlo simulation. We use data simulated the stock-elastic demand model of section 2, using the baseline calibration described in table 2. We simulate 10 blocks of 10000 time-series observations from the model, and keep the last 9000 observations from each block. In line with the empirical specification explored below, we estimate a VAR(3) in levels on each block of data, including a constant term. We construct the set of identified responses for the simulated data as the union of the (exactly identified) set obtained for each block. We construct the set of Forecast Error Variances (FEV)

in a similar fashion.

In figure 17, we report the results of our exercise when our model economy is hit by three shocks: surprise shocks to TFP, surprise shocks to the labor wedge, and anticipated shocks to TFP. We attribute a standard deviation of $\sigma_z = \sigma_\psi = \sigma_{z,N} = 1$ to all three shocks.[12] The model's true impulse responses mostly lie within the set of impulse responses identified by impact restrictions (the shaded blue area of figure 17). In table 2, we report the set of FEV identified using our sign restrictions. With the exception of short-run movements in inventories, the model's FEV are all within the identified sets.

We then conduct a second Monte-Carlo exercise, by adding anticipated shocks to the labor wedge. In this case, the set of identified impulse responses and FEV we identify should reflect the effects of both news shocks to TFP and to the labor wedge. We set the standard deviation of the additional news shock to $\sigma_{\psi,N} = 1$. The set of identified impulse responses are reported in figure 18. The set of identified TFP responses shifts down, reflecting the fact that some of the identified shocks have no effect on TFP. Likewise, the identified impulse response set for sales shifts up, reflecting the combined impact of the two types of news shocks. For this reason, the model impulse response with respect to labor wedge news is not included in the identified set. Table 3 reports the set of identified FEV. While the model FEV for each news shock separately do not always fall within the identified set of FEV, the sum of the FEV of both types of news shocks does. This Monte-Carlo experiment thus suggests that if there are several types of anticipated shocks generating the data we observe, then our procedure should provide an upper bound on the combined effect of the contribution of these shocks to the business-cycle movements of observables.

6.3 Empirical evidence

We now apply our identification scheme to quarterly postwar US data.

6.3.1 Data

We use three observables in our exercise: TFP, inventories and sales. For TFP, we use the Solow residual series constructed in Fernald (2012). Our baseline sample period is 1960Q1–2010Q4. It is well known that the Solow residual may overstate the volatility of TFP because it does not correct for cyclical variation in capacity utilization and labor hoarding. We come back to this issue in the robustness section below. The data for inventories and sales is described in appendix A.

---

[12]As in our baseline calibration, the persistence of the TFP shock is $\rho_z = 0.99$, and the persistence of the labor wedge shock is $\rho_\psi = 0.95$.

6.3.2 Specification and estimation

We estimate a specification of our system with four lags, obtained using the Schwartz information criterion. We estimate the model using Bayesian methods. We use a diffuse prior for both the coefficients of the autoregressive structure and the variance-covariance matrix of errors terms; we consider the robustness of our results to other prior specifications below. Each draw from the posterior identifies a set of possible impulse responses satisfying our impulse restrictions, and we use a uniform conditional prior on the identified set to draw from the posterior of the impulse responses, following Moon et al. (2011). We follow the same method for the posterior distribution of FEV.

6.3.3 Results

Figure 19 reports the estimated impulse responses for the baseline specification.

The first result is that, in this specification, the estimated impulse responses are qualitatively consistent with model predictions. In particular, the identified shock generates a diffusion in TFP, much as the identified shocks of Beaudry and Portier (2006) and Barsky and Sims (2011). Unlike these papers, we obtain these results without imposing any restrictions on the medium or long-run behavior of TFP. Quantitatively, the shock generates a persistent fall in inventories, and a near-permanent increase in sales. In response to a 0.2% increase in TFP in the long-run, sales initially jump by about 0.7%, roughly in line with the magnitudes of the model, where a 1% near-permanent increase in TFP generates a 2.5% increase in sales. However, the response of inventories, which falls by at most 0.5%, is much smaller than that predicted by the baseline calibration of the model, where the near-permanent increase in TFP generates a fall of roughly 5% of inventories in the quarter preceding the realization of the anticipated shock.

The second result is that the anticipation period of the news shock is small. In fact, TFP has already increased substantially, by 0.2% percent, two quarters after the shock. The shock identified by our restrictions is thus not a slow diffusion in TFP, but rather an almost-immediate, and near-permanent increase. Accordingly, the increase in sales is almost immediate. From the standpoint of the model, the fact that the anticipation horizon is short may also account for the small magnitude of the (negative) response of inventories. Note that inventory investment picks up after two quarters in the estimated response, in accordance with the model's prediction, whereby inventory investment becomes positive when the TFP increase is realized.

The FEV attributable to this shock are reported in table 4. First, the identified shock explains a relatively small fraction of TFP movements, even in the long-run: 40 quarters out, the median FEV attributable to the shock is only 11%. This is in contrast to results obtained by other identification schemes, such as Barsky

and Sims (2011), where, by construction, the forecastable increase in TFP accounts for a large fraction of long-run TFP movements. Second, for horizons of 5 quarters and more, this shock also accounts for less than 20% of the FEV of inventories and 18% of the FEV of sales. Thus, shocks causing a persistent negative comovement between sales and inventories seem to be of little relevance in accounting for the observed times series for TFP, inventories and sales in the long run. However, the shock does seem to account for a substantial fraction of short-run fluctuation in inventories (61% on impact in the baseline calibration), and sales (33%). These large contributions vanish after 2 quarters. This finding is consistent with the results of Wen (2005), who documents that inventories are countercyclical at very high frequencies, between 2 and 3 months per cycle, although unconditional correlations are dominated by the procyclicality of inventories in the medium to long run. Our identification scheme attributes a large portion of these short-run movements to our identified shock.

We summarize the key points of our baseline results as follows: (i) our identified shock generates a behavior qualitatively consistent with model impulse responses to anticipated shocks; (ii) the identified impulse response suggests that the horizon of implementation is possibly as short as 2 quarters; (iii) the contribution of the shock to movements in TFP, inventories and sales is substantial in the short run, but not in the long run, contrary to the model's predictions, where FEV are similar at all horizons. Thus, our conclusion is that while news shocks with short anticipation horizons may offer a rationale for the very high-frequency behavior of sales and inventories, the 4 to 8 quarter anticipation periods studied in the literature on news shocks is hard to justify. Moreover, the news shocks we identify do not account for the majority of long-run movements in TFP, inventories or sales. In this sense, the empirical behavior of inventories is a challenge for anticipated shocks to fundamentals.

6.4   Robustness

To check the robustness of our results, we conduct our empirical exercise using different measures of TFP, splitting samples, and assuming a different prior distribution.

First, in order to account for the fact that procyclical variation in capacity utilization may be driving TFP movements, we use an alternative TFP series, proposed by Fernald (2012) and which adjusts for capital utilization and labor hoarding. Figure 20 shows the impulse response of TFP, inventories and sales to identified news shocks using this measure of TFP. The initial response is muted but in the long-run, we observe that TFP increases. Moreover, sales also show a permanent increase. However, inventories decline only for the initial period, indicating again that the anticipation period is short. Additionally, since inventories are mostly held by investment and durable goods producing sectors, we use a third TFP series,

utilization adjusted and specific to the equipment and consumer durable sector. This series is also constructed by Fernald (2012). Figure 21 shows that our results also remain when using this series.

Second, as studied in detail by McCarthy and Zakrajšek (2007), inventory dynamics have substantially changed since the 1980's: while the procyclicality of inventories and the counteryclicality of IS ratios remains, the volatility of total inventory investment has fallen, possibly because of improvements in inventory management, contributing to the fall in output volatility. 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. The results are reported in figures 22 and 23. After 1984, we observe that our results mostly remain in terms of the dynamics of the posterior median, although the credible region is wider than before. However, before 1984, our identification strategy recovers a long-run decline in TFP, inventories and sales. In the pre-1984 sample, it is likely that the shocks recovered with our identification strategy may not be limited to anticipated shocks. In particular, oil shocks in the 70's may play an important role in shaping the dynamics of inventories and sales.

Third, we conducted our exercise using different prior specifications: the Minnesota prior, and a Normal-Wishart prior. The results obtained with both types of prior are similar, and we do not report them here.

In table 4, we also report forecast error variance decompositions obtained in these alternative specifications. The same message emerges as in the baseline specification: while identified news shocks may account for very short-run fluctuations our observables, they are minor contributors to their medium and long-run movements.

We thus conclude that the key features of the news shock identified in our baseline specification – the qualitative shape of impulse responses, the short "news period", and the small contribution to medium and long-run forecast error variances – are robust to alternative measures of TFP, alternative prior specifications, and that they survive in the post-1984 sample.

7   Conclusion

In this paper, we studied the response of inventories to anticipated improvements in fundamentals. We established restrictions on structural and reduced-form model parameters under which the response of inventories and sales to these shocks will be characterized by (i) inventory procyclicality and (ii) countercyclicality of the IS ratio. We showed that these restrictions are violated by standard calibrations of the two classes of models we study, resulting, in particular, in a fall in inventories after the shock. Our analysis highlighted the key mechanism behind this result: inventory disinvestment occurs because anticipated improvements in fundamentals generate a strong motive for intertemporal substitution 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, adjustment costs, and habit formation. Lastly, we used the negative comovement between inventories and sales to identify news shocks in postwar US data. We showed that the dynamic responses of inventories, sales, and TFP to the identified shocks are qualitatively consistent with structural impulse responses, but also suggest that news shocks play a small role in aggregate fluctuations, for two reasons: the identified "news period" is very short, at most 2 quarters; and the shock contributes little to medium and long-run movements of TFP and inventories. Given the procyclicality of inventories at business-cycle frequencies, this result may have been expected. However, it emphasizes that this procyclicality is hard to square with the effect of news shocks standard models of inventories, in which small movements in future marginal costs lead to large inventory adjustments.

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 relative marginal cost. The approach we have taken in this paper is to compute the elasticity implied by existing models of output 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, or using an altogether different type of inventories.[13] Theoretically, anticipated shocks to fundamentals 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.

---

[13]See Chang, Hornstein, and Sarte (2009) for example.

References

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

Beaudry, P. and B. Lucke (2010). Letting Different Views about Business Cycles Compete. *NBER Macroeconomics Annual 24*(1), 413–456.

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

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

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

Broda, C. and D. E. Weinstein (2006). Globalization and the Gains from Variety. *Quarterly Journal of Economics 121*(2), 541–585.

Cachon, G. P. and M. Olivares (2009). Drivers of Finished-Goods Inventory in the U.S. Automobile Industry. *Management Science 56*(1), 202–216.

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*, 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.

Chetty, R., A. Guren, D. Manoli, and A. Weber (2011). Are Micro and Macro Labor Supply Elasticities Consistent? A Review of Evidence on the Intensive and Extensive Margins. *American Economic Review 101*(3), 471–475.

Cho, J. and T. Cooley (1995). The Business Cycle with Nominal Contracts. *Economic Theory 6*(1), 13–33.

Christiano, L. J. (1988). Why does Inventory Investment Fluctuate so Much? *Journal of Monetary Economics 21*(2–3), 247–280.

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.

Fazzari, S. M., R. G. Hubbard, and B. C. Petersen (1988). Financing Constraints and Corporate Investment. *Brookings Papers on Economic Activity 19*(1), 141–206.

Fernald, J. (2012). A Quarterly, Utilization-Adjusted Series on Total Factor Productivity. *Federal Reserve Bank of San Francisco working paper 19*.

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

Gruen, T., D. Corsten, and S. Bharadwaj (2002). Retail Out-of-Stocks: A Worldwide Examination of Extent, Causes and Consumer Responses.

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*.

Justiniano, A., G. E. Primiceri, and A. Tambalotti (2010). Investment shocks and business cycles. *Journal of Monetary Economics 57*(2), 132–145.

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

King, R., C. Plosser, and S. Rebelo (1988). Production, Growth and Business Cycles: I. The Basic Neoclassical Model. *Journal of Monetary Economics 21*(2–3), 195–232.

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 (2012). Inventories, Markups, and Real Rigidities in Menu Cost Models. *forthcoming, Review of Economic Studies*.

Lutkepohl, H. (2007). *New Introduction to Multiple Time Series Analysis*. Springer Verlag.

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, E. Granziera, and M. Lee (2011). Inference for VARs Identified with Sign Restrictions. *NBER Working Paper* (17140).

Nekarda, C. and V. Ramey (2010). The Cyclical Behavior of the Price-Cost Markup. *UC San Diego mimeo*.

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

Ramey, V. A. and K. D. West (1999). Inventories. In J. B. Taylor and M. Woodford (Eds.), *Handbook of Macroeconomics*, Volume 1 of *Handbook of Macroeconomics*, Chapter 13, pp. 863–923. Elsevier.

Richardson, H. (1995). Control Your Costs - Then Cut Them. *Transportation & Distribution 36*(12), 94–94.

Rotemberg, J. and M. Woodford (1997). An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy. In *NBER Macroeconomics Annual 1997, Volume 12*, NBER Chapters, pp. 297–361. National Bureau of Economic Research, Inc.

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

Smets, F. and R. Wouters (2007). Shocks and Frictions in U.S. Business Cycles: A Bayesian DSGE Approach. *American Economic Review 97*(3), 586–606.

Wen, Y. (2005). Understanding the Inventory Cycle. *Journal of Monetary Economics 52*(8), 1533–1555.

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*. Priceton: Princeton University Press.

A Appendix to section 1

A.1 Data Sources

The postwar quarterly data (1947Q1 – 2012Q1) used in section 1 to produce the unconditional moments as well as figure 1 come from the following sources:

- **Output:** Gross Domestic Product, NIPA Table 1.1.6 (Real Gross Domestic Product, Chained 2005 Dollars, Seasonally Adjusted).

- **Inventory investment:** Change in private inventories, NIPA Table 1.1.6 (Real Gross Domestic Product, Chained 2005 Dollars, Seasonally Adjusted).

- **Inventory-Sales ratio:** Nonfarm inventories to final sales of goods and structures, NIPA Table 5.7.6A (Real Private Inventories and Real Domestic Final Sales of Business by Industry, Chained 2005 Dollars, Seasonally Adjusted) and Table 5.7.6B (Real Private Inventories and Real Domestic Final Sales by Industry, Chained 2005 Dollars, Seasonally Adjusted). The series from the two tables have been ratio spliced at the value in 1996Q4.

All moments are computed by detrending the log series by HP-filtering.

A.2 Cyclicality of output inventories

The procyclical behavior of inventories is also true when focusing on a narrower type of inventories that are in the form of finished goods. Using the postwar quarterly data (1947Q1 – 2012Q1) for retail and wholesale trade inventories from from NIPA Table 5.7.6A and 5.7.6B and ratio spliced as discussed above, the correlation between the log-detrended retail trade inventories and output is 0.68, while that of wholesale trade inventories is 0.49. Since these type of inventories are mainly in finished goods, the positive correlation suggests that output inventories are also procyclical. For reference, the correlation of the overall private inventories and output is 0.59, where the overall private inventories are also taken from Table 5.7.6A and 5.7.6B.

B Appendix to section 2

B.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.2 Equilibrium symmetry

In the formulation of the first order conditions above, we dropped the subscript j, since a market equilibrium of the stock-elastic demand model is symmetric. One can see this as follows. First, combining the

labor and capital demand schedules of firms, one sees that the capital-labor ratio is identical across firms. This in turn implies (using either the capital demand or labor demand schedule) that marginal cost is constant across firms. Using equation (43), the first order condition governing the optimal choice of inventories, the inventory to sales ratio, or equivalently the stock-to-sales ratio, is constant across firms. Furthermore, using the optimal pricing condition (44), price is constant across firms. Thus, substituting out price and the stock-to-sales ratio in the demand schedule for variety j, we have:

$$s_t(j) = \left(\frac{a_t(j)}{a_t}\right)^\zeta s_t = \left(\frac{\Xi_t s_t(j)}{a_t}\right)^\zeta s_t,$$

which in turn implies that sales are constant across firms, $s_t(j) = s_t$.

C   Appendix to section 4

Table 1 summarizes our baseline calibrations of the stock-elastic and stockout avoidance models. In this appendix, we explain in more detail our calibrations and in particular the targets we match for inventory-related parameters. As discussed in detail in section 3, our results on the negative comovement between inventories and sales hold for a large range of parameters; the precise calibration pursued in this appendix of course falls within this range.

In each model, we choose the value of the elasticity of output to labor in order to match a steady-state labor share of income of 68.5 percent, the sample average of the labor income share in the data we use in our empirical estimates of section 6.[14] In both models, we choose the elasticity of substitution across varieties θ so that the steady-state markup of price over marginal cost is 25 percent. This is in the range of 24 to 30 percent computed by Nekarda and Ramey (2010) using a panel of industries.[15] We choose the steady-state labor disutility shifter ψ to match a fraction of hours worked in steady-state of 0.2. Under our preference specification, the parameter ξ is the Frisch elasticity of labor supply. We follow the bulk of the real business cycle literature and choose a relatively high Frisch elasticity of $\xi = 2.5$. This is the value used by Jaimovich and Rebelo (2009), and it is also in the range of the values considered by Cho and Cooley (1995) and King, Plosser, and Rebelo (1988).[16] Finally, we use the median estimate of 9.11 reported by Schmitt-Grohé and

---

[14]The steady-state labor share in our model is defined as $s_n = wn/py$, with w and p the relative prices of labor and output, respectively. In our setup, because of monopolistic competition, the steady-state labor share does not equal the elasticity of output to labor.

[15]It is also close to the median estimates of 21 percent in Smets and Wouters (2007) and 23 percent in Justiniano, Primiceri, and Tambalotti (2010), and consistent with the value used in Kryvtsov and Midrigan (2010, 2012). Broda and Weinstein (2006) use import data and estimate an elasticity of substitution across varieties of 3.0 to 3.7 for the lowest level of good disaggregation, which would imply a higher steady-state markup of 50 percent. Using their higher estimates of the elasticity of substitution across varieties would not affect our results.

[16]As emphasized by Chetty, Guren, Manoli, and Weber (2011), the large Frisch elasticity which the real business cycle literature has found to be necessary to fit the business-cycle dynamics of aggregate hours is at odds with micro estimates of the

Uribe (2012) for the calibration of the curvature of the adjustment cost function.[17]

We choose the remaining structural parameters of our calibrations to match data moments directly related to inventories. In the two models we consider, we choose to match the average quarterly IS and output-to-sales ratios in NIPA data, which are $IS = 0.5$ and $\kappa = 1.0053$, respectively.

To our knowledge, there are no existing empirical estimates of the quarterly rate of depreciation of inventories, δ_i. The macroeconomic literature has used annualized values ranging from 0 (Kahn, 1987) to 10 percent (Kryvtsov and Midrigan, 2012). A commonly cited source for these values is Richardson (1995), who suggests that inventory carrying costs are smaller than 1 percent per month. However, carrying costs do not directly map into the geometric depreciation rate of our model, since they only reflect storage and financial costs, and not the costs associated to product decay. Given the uncertainty on this parameter, we proceed by noting that in the steady-state of our model, the rate of depreciation of inventories is related to our two targets, the IS and output to sales ratios, by $\delta_i = (\kappa - 1)/IS$. This implies an annualized depreciation rate of 4.17 percent, in the range used in the macroeconomic literature.

Finally, given other structural parameters (and, in particular, δ_i), the IS ratio depends on the elasticity of sales to stock ζ (in the stock-elastic demand model) and on the variance of the demand shock σ_d (in the stockout avoidance model). In the stock-elastic demand model, our baseline calibration targeting an IS ratio of 0.5 implies an elasticity of sales to stock of $\zeta = 0.126$. Bils and Kahn (2000) report estimates of this elasticity for six production to stock industries. Their estimates range from 0.011 (for the tobacco industry) to 0.494 (for the petroleum industry). Our calibration thus implies an elasticity in the range of their estimates. To compute impulse responses for the stockout avoidance model, we use a log-normal demand shock distribution, with a mean normalized to 1. Our calibration results in a standard deviation of $\sigma_d = 0.32$. In turn, this standard deviation yields a steady-state stockout probability of 8.4 percent. Bils (2004) studies CPI data comprising 63 categories of consumer durables and finds a weighted stockout probability of 5.2 percent after eliminating non-temporary stockouts. On the other hand, Gruen, Corsten, and Bharadwaj (2002) survey existing empirical evidence on the frequency of stockouts for retailers. They combine results covering 661 outlets and 71000 consumers, and estimate the stockout probability to be 7.9 percent in the US. Our calibration thus delivers a stockout probability that is slightly above direct empirical estimates.[18]

---

intensive margin Frisch elasticity. Their survey of existing micro evidence on the intensive margin Frisch elasticity points to a lower value of 0.5. However, our choice of a high Frisch elasticity is in fact conservative, from the standpoint of the behavior of inventories. Indeed, a lower Frisch elasticity would result in a smaller degree of real rigidities in our model, amplifying the response of wages to the shock. In turn, the negative response of inventories would also be magnified.

[17]Smets and Wouters (2007) report a mean estimate of 5.74 for this parameter, while the estimate of Justiniano, Primiceri, and Tambalotti (2010) is 3.14. A lower value of 4 for the curvature of the adjustment cost function would not change our results for sales, output and investment, but it would somewhat weaken the initial response of investment.

[18]Calibrations of the stockout-avoidance model resulting in lower stockout probabilities would deliver almost similar impulse responses, while not allowing us to match the NIPA IS and output to sales ratios simultaneously.

D Appendix to section 6

Proof of lemma 2.

Let C denote the unique Cholesky decomposition of $\hat{\Omega}$. Then, for any impact matrix A, it is easy to show that

$$U = C^{-1}A$$

is an orthogonal real matrix, that is, $U^T U = U U^T = I$. Thus, given the Cholesky decomposition of the variance-covariance matrix, one can obtain all impact matrices through a rotation of the Cholesky decomposition; that is, the set of impact matrices $\mathcal{I}(\hat{\Omega})$ is given by:

$$\mathcal{I}(\hat{\Omega}) = \{A \in M_3(\mathbb{R}) | \exists U \in O_3(\mathbb{R}) \text{ s.t. } A = CU\},$$

where C denotes the Cholesky decomposition of $\hat{\Omega}$. As detailed in Moon et al. (2011), sign restriction identification schemes, whether through simualtion or analytically, typically exploit this construction of the set of impact matrices.

We now turn to the characterization of the subset of $\mathcal{I}(\hat{\Omega})$ that satisfies restrictions (i)-(iii). First, letting:

$$C = \begin{pmatrix} a & 0 & 0 \\ b & d & 0 \\ c & e & f \end{pmatrix}$$

where by definition of the Cholesky decomposition, $a > 0$, $d > 0$ and $f > 0$, and letting $U = (U_{i,j})$, the first line of A is given by $[aU_{11}, aU_{12}, aU_{13}]$, so that restriction (i) implies that $U_{12} = U_{13} = 0$. Furthermore, if we impose that the $(1,1)$ element of A is strictly postive, then $U_{11} > 0$. In turn, the equalities $U^T U = I$ then impose that $U_{21} = U_{31} = 0$ and $U_{11}^2 = 1$, so that $U_{11} = 1$. The equalities also imply that $U_{32}U_{22} + U_{23}U_{33} = 0$ and $U_{22}^2 + U_{23}^2 = U_{32}^2 + U_{33}^2 = 1$. This in turn means that the orthogonal matrix U must take the form:

$$\overline{U}(\theta) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos(\theta) & -\sin(\theta) \\ 0 & \sin(\theta) & \cos(\theta) \end{pmatrix} \quad \text{or} \quad \underline{U}(\theta) = \begin{pmatrix} 1 & 0 & 0 \\ 0 & -\cos(\theta) & \sin(\theta) \\ 0 & \sin(\theta) & \cos(\theta) \end{pmatrix}, \quad \theta \in [0, 2\pi[,$$

where the former matrix is a proper rotation $(\det(\underline{U})(\theta) = 1)$, and the latter matrix is an improper rotation $(\det(\underline{U})(\theta) = -1)$. The orthogonal matrices associated to the impact matrices satisfying our first restriction are thus rotation matrices, either proper or improper.

We now turn to applying our second restriction to the subset of $\mathcal{I}(\theta)$ composed of proper rotation matrices $\overline{U}(\theta)$. The impact matrix $\overline{A}(\theta)$ corresponding to a particular $\overline{U}(\theta)$ is given by:

$$\overline{A}(\theta) = \begin{pmatrix} a & 0 & 0 \\ b & d\cos(\theta) & -d\sin(\theta) \\ c & e\cos(\theta) + f\sin(\theta) & -e\sin(\theta) + f\cos(\theta) \end{pmatrix}.$$

We label the third shock the news shock. In order to ensure that this shock generates negative comovement between inventories and sales, we impose that inventories respond positively and sales respond negatively, that is:

$$-d\sin(\theta) < 0 \text{ and } -e\sin(\theta) + f\cos(\theta) > 0.$$

Moreover, in order to ensure that it is the *only* shock orthogonal to own TFP innovations with those properties, we impose that θ is such that:

$$d\cos(\theta) > 0 \quad \text{and} \quad e\cos(\theta) + f\sin(\theta) > 0,$$

so that the shock generates an increase in sales and positive comovement between inventories and sales.

Depending on the sign of $\frac{e}{f}$ (which is the sign of e, since f is strictly positive), these restrictions define different sets of matrices. When $e > 0$, restrictions on elements of the third column are verified when $\theta \in [0, \pi]$ and $\cot(\theta) > \frac{e}{f}$, that is, for:

$$\theta \in \left[0, \overline{\theta}\right], \quad \overline{\theta} = \cot^{-1}\left(\frac{e}{f}\right) \in \left[0, \frac{\pi}{2}\right]. \tag{46}$$

The restrictions on the elements of the second column are always satisfied (as $\theta \in \left[0, \frac{\pi}{2}\right]$).

Things are slightly more complicated when $e < 0$. The set of values of θ satisfying our restrictions on the third column is still $\theta \in \left[0, \overline{\theta}\right]$, $\overline{\theta} = \cot^{-1}\left(\frac{e}{f}\right)$, but now $\overline{\theta} \in \left[\frac{\pi}{2}, \pi\right]$. But in this range, our restrictions on the second column may not always hold. For example, in the range $\left[\frac{\pi}{2}, \overline{\theta}\right]$, we have that:

$$d\cos(\theta) < 0, \quad e\cos(\theta) + f\sin(\theta) > 0,$$

so that our restrictions on the second column are violated; in effect, our restrictions on the third column alone are not sufficient to identify news shocks in this range. In is straightforward to check that the restrictions

on the second column are only satisfied when:

$$\theta \in \left[\underline{\theta}, \frac{\pi}{2}\right], \quad \underline{\theta} = \cot^{-1}\left(-\frac{f}{e}\right).$$

Thus, since:

$$0 < \underline{\theta} < \frac{\pi}{2} < \overline{\theta} < \pi,$$

the set of θ satisfying our restrictions on the second and third column is given by:

$$\theta \in \left[\underline{\theta}, \frac{\pi}{2}\right], \quad \underline{\theta} = \cot^{-1}\left(-\frac{f}{e}\right) \in \left[0, \frac{\pi}{2}\right]. \tag{47}$$

Equation (46) characterizes the subset of impact matrices statisfying our sign restrictions obtained using proper rotations when $e > 0$, and equation (47) characterizes this subset when $e < 0$.

What about improper rotation matrices? It turns out that none can satisfy our impact restrictions. To see this, first note that the impact matrix associated to an improper rotation $\underline{U}(\theta)$ is given by:

$$\underline{A}(\theta) = \begin{pmatrix} a & 0 & 0 \\ b & -d\cos(\theta) & d\sin(\theta) \\ c & -e\cos(\theta) + f\sin(\theta) & e\sin(\theta) + f\cos(\theta) \end{pmatrix}.$$

First, it must be the case that $\theta \in \left[\pi, \frac{3\pi}{2}\right]$ in order to guarantee that $d\sin(\theta) < 0$ and $-d\cos(\theta) > 0$. When $e < 0$, this immediately implies that $-e\cos(\theta) + f\sin(\theta) < 0$, so that no value of θ satisfies our restrictions. When $e > 0$, our restriction on the response of sales to the news shock is satisfied if and only if:

$$\theta > \underline{\theta} = \cot^{-1}\left(-\frac{e}{f}\right) \in \left[\frac{3\pi}{2}, 2\pi\right].$$

But given that we must have $\theta \in \left[\pi, \frac{3\pi}{2}\right]$, again no value of θ can satisfy our restrictions.

Thus, no improper rotation matrix can generate an impact matrix satisfying our restrictions. Therefore, the set of impact matrices satisfying our restrictions is fully characterized by equations (46) and (47).

∎

Tables and figures

Tables and figures for section 1
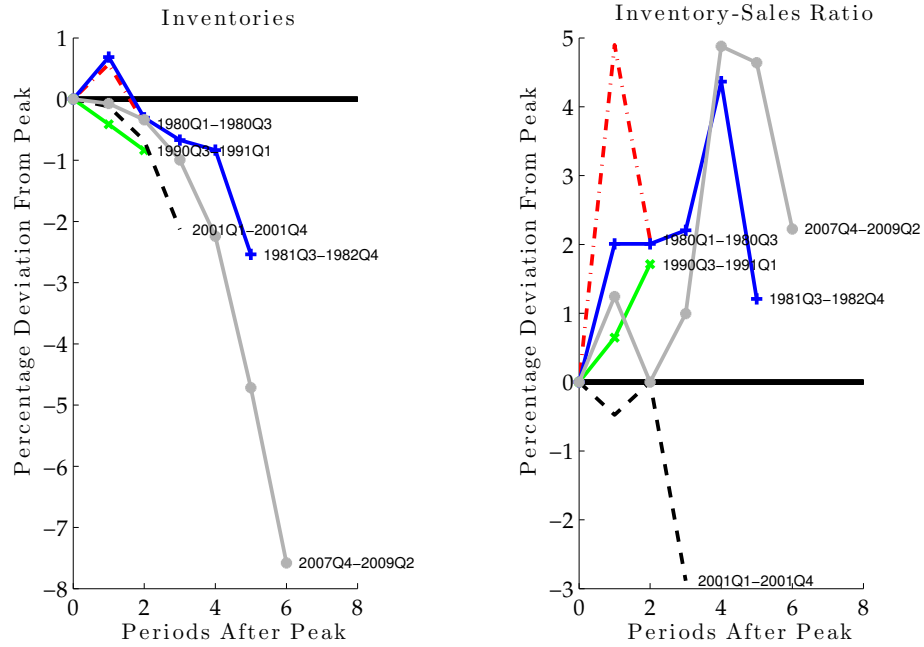
Figure 1: Inventories after NBER peaks (quarterly). See data apendix for data sources.

# Tables and figures for section 3

Figure 2: The impact response of inventories (left panel) and the IS ratio (right panel) to a news shock.


Figure 3: Stock-elastic demand model. The left panel provides the upper bound on ω for procyclical inventories, derived from targeting the steady-state IS ratio. The right panel provides the value of η as a function of $\gamma (= \beta(1-\delta_i))$, holding fixed all the other structural parameters.

Tables and figures for section 4

Parameter	Value	Description
Parameters common to both models		
β	0.99	Subjective discount factor
σ	1	Elasticity of intertemporal substitution in consumption
δ_k	0.025	Rate of depreciation of capital
ξ	2.5	Frisch elasticity of labor supply
$\phi''(1)$	9.1	Curvature of investment adjustment costs
δ_i	0.011	Rate of depreciation of inventories
ρ_z	0.99	Persistence of TFP process
ρ_ψ	0.95	Persistence of labor wedge process
Stock-elastic demand model		
α	0.896	Labor elasticity of production function
θ	5	Elasticity of substitution across varieties
ψ	7.457	Steady-state labor shifter
ζ	0.126	Elasticity of sales to on-shelf goods
Stockout avoidance model		
α	0.893	
θ	5.718	
ψ	7.419	
σ_d	0.315	Standard deviation of idiosyncratic demand shock

Table 1: Calibration of the stock-elastic demand and stockout avoidance models.

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

Figure 5: **Impulse responses to 4-period news shocks in the stockout avoidance model.** Solid line: TFP news; dashed line: labor wedge news. The time unit is a quarter. Impulse responses are reported in terms of percent deviation from steady-state values.

Figure 6: **Impulse response of the growth rate of marginal cost to 4-period news shocks in the stock-elastic demand model.** Solid line: TFP news; dashed line: labor wedge news. The time unit is a quarter. Impulse responses are reported in terms of percent deviation from steady-state values.
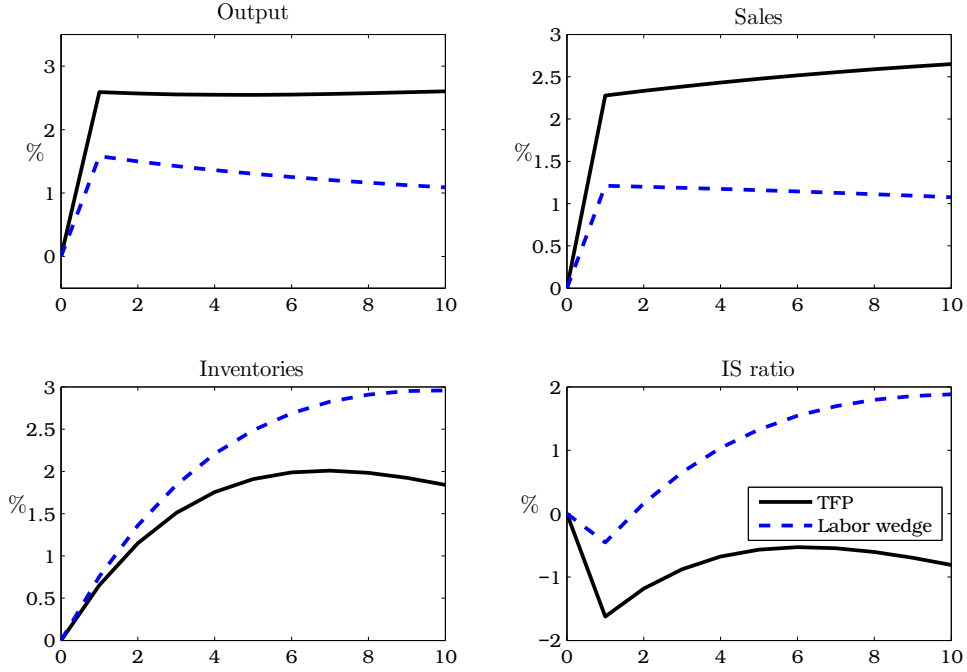
Figure 7: **Impulse responses to suprise shocks in the stock-elastic demand model.** Solid line: 1% suprise increase in TFP; dashed line: 1% surprise fall in the labor wedge. The time unit is a quarter. Impulse responses are reported in terms of percent deviation from steady-state values.
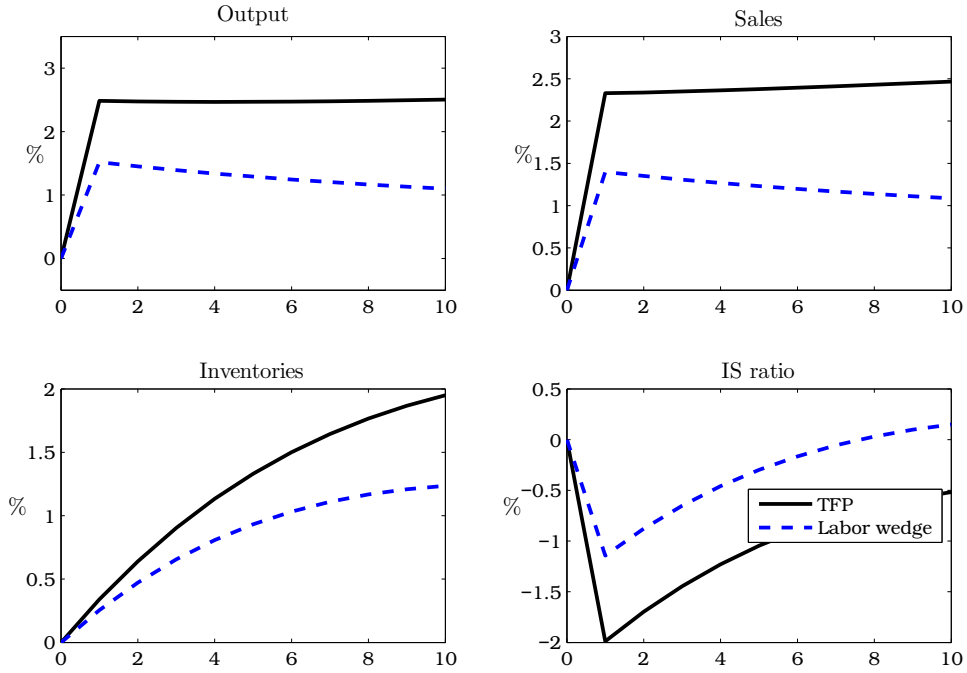
Figure 8: **Impulse responses to suprise shocks in the stockout avoidance model.** Solid line: 1% suprise increase in TFP; dashed line: 1% surprise fall in the labor wedge. The time unit is a quarter. Impulse responses are reported in terms of percent deviation from steady-state values.

Tables and figures for section 5
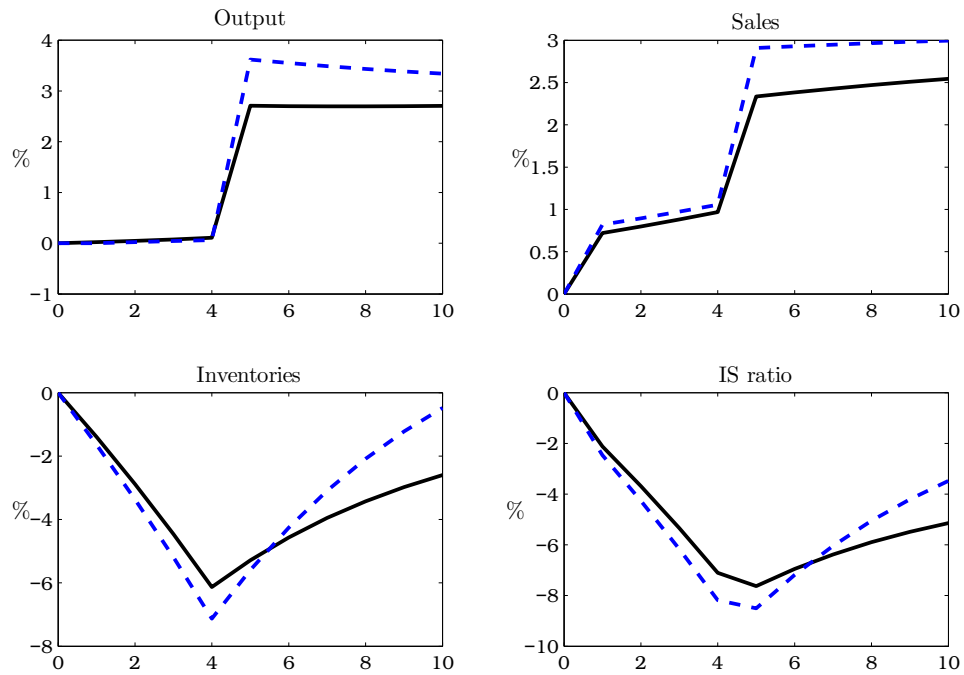
Figure 9: **Impulse responses to 4-period TFP news shock with capacity utilization in the stock-elastic demand model.** Solid line: $\delta''(1) = 0$; dashed line: $\delta''(1) = 0.34$.
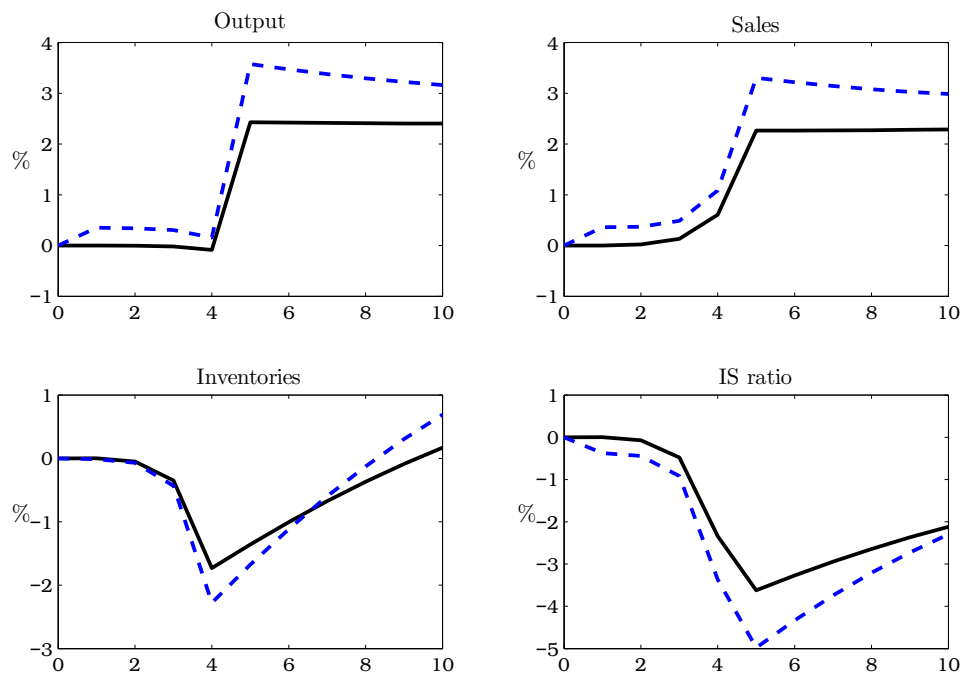
Figure 10: **Impulse responses to 4-period TFP news shock with capacity utilization in the stockout avoidance model.** Solid line: $\delta''(1) = 0$; dashed line: $\delta''(1) = 0.34$.
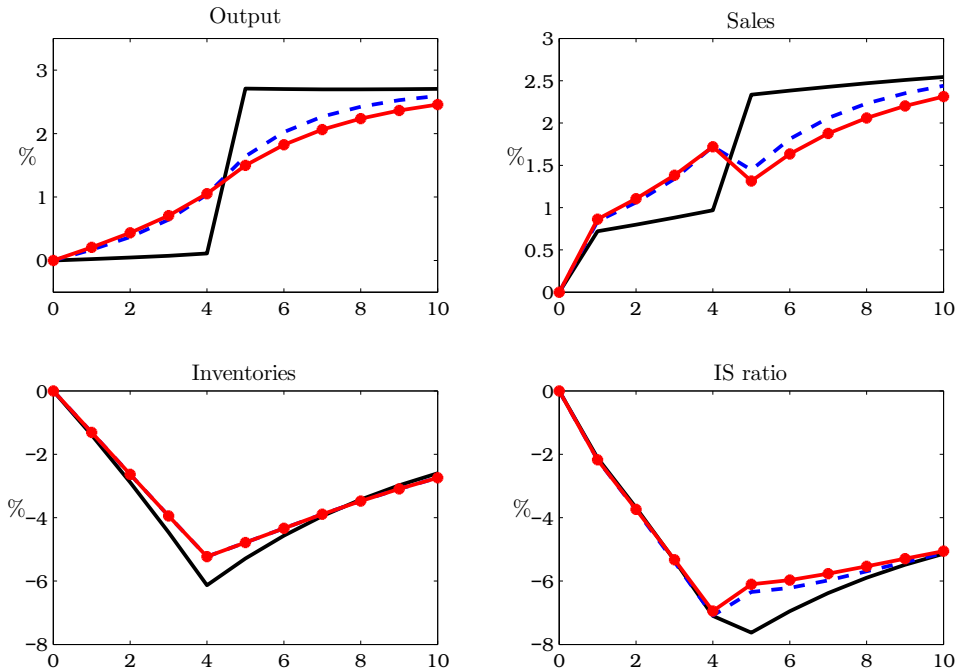
Figure 11: **Impulse responses to 4-period TFP news shock with output adjustment cost in the stock-elastic demand model.** Solid line: $\phi_y = 0$; dashed line: $\phi_y = 3$; circled line: $\phi_y = 6$.

Figure 12: **Impulse responses to 4-period TFP news shock with stock adjustment cost in the stockout avoidance model.** Solid line: $\phi_a = 0$; dashed line: $\phi_a = 3$; circled line: $\phi_a = 6$.

Figure 13: **Impulse responses to 4-period TFP news shock with habit persistence in the stock-elastic demand model.** Solid line: $b = 0$; dashed line: $b = 0.4$; circled line: $b = 0.8$.
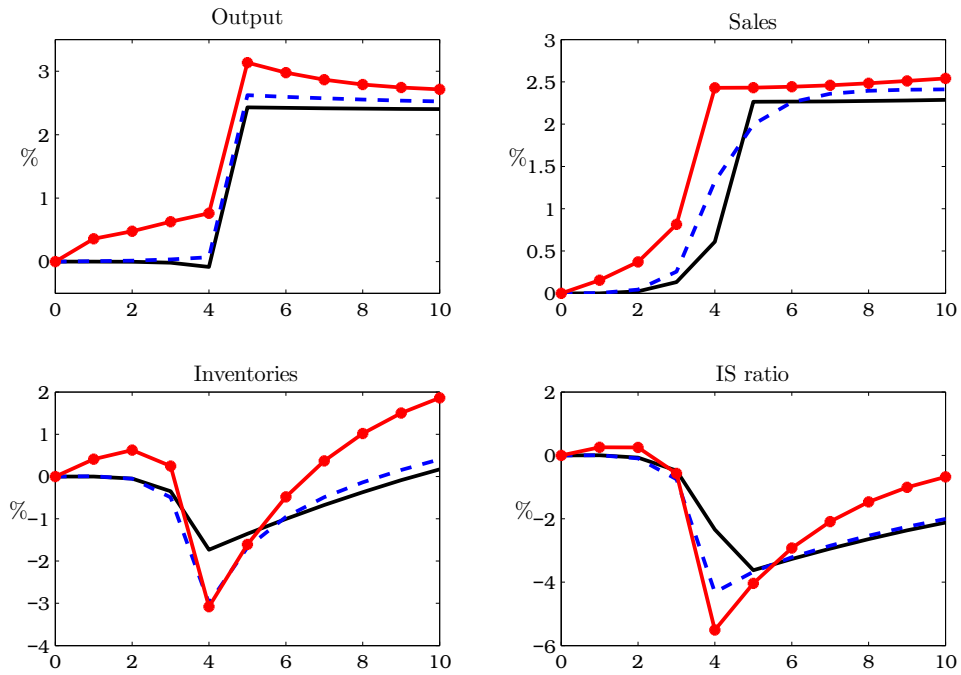

Figure 14: **Impulse responses to 4-period TFP news shock with habit persistence in the stockout avoidance model.** Solid line: $b = 0$; dashed line: $b = 0.4$; circled line: $b = 0.8$.
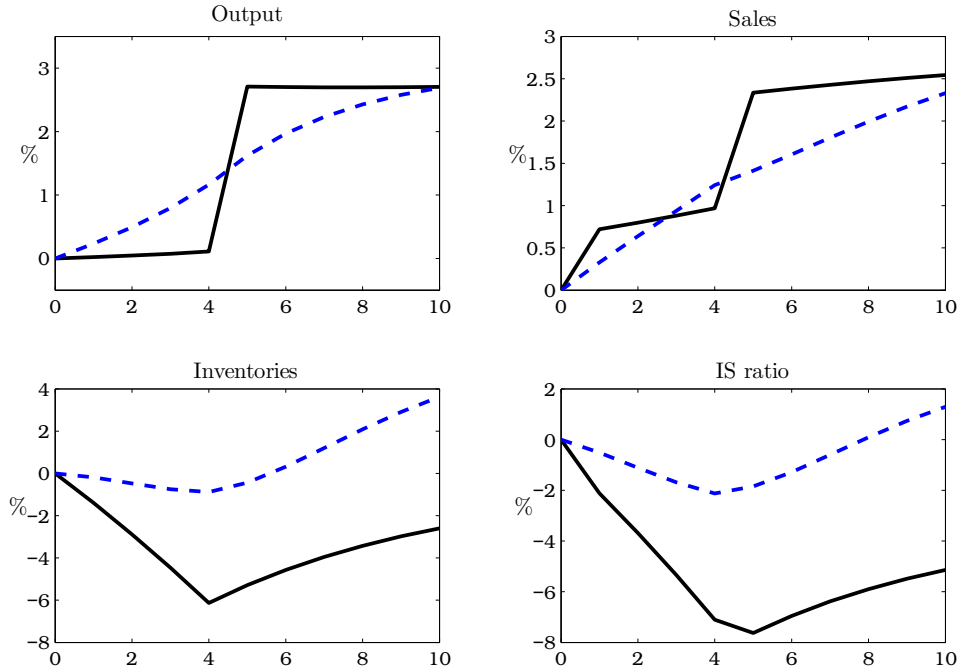
Figure 15: **Impulse responses to 4-period TFP news shock with all frictions in the stock-elastic demand model.** Solid line: baseline, no frictions; dashed line: $\delta''(1) = 0.34$, $\phi_y = 6$, $b = 0.8$.


Figure 16: **Impulse responses to 4-period TFP news shock with all frictions in the stockout avoidance model.** Solid line: baseline, no frictions; dashed line: $\delta''(1) = 0.34$, $\phi_a = 6$, $b = 0.8$.

Tables and figures for section 6

Figure 17: **Monte-Carlo experiment with news shock to TFP only**. TFP, inventories and sales, from left to right. The shaded blue area is the set of identified impulse responses using impact zero and sign restrictions. The solid line is the model impulse response to the TFP news shock.


Figure 18: **Monte-Carlo experiment with news shock to TFP and the labor wedge**. TFP, inventories and sales, from left to right. The shaded blue area is the set of identified impulse responses using impact and sign restrictions. The solid line is the model impulse response to the TFP news shock, and the dashed line is the model impulse response to the labor wedge news shock.
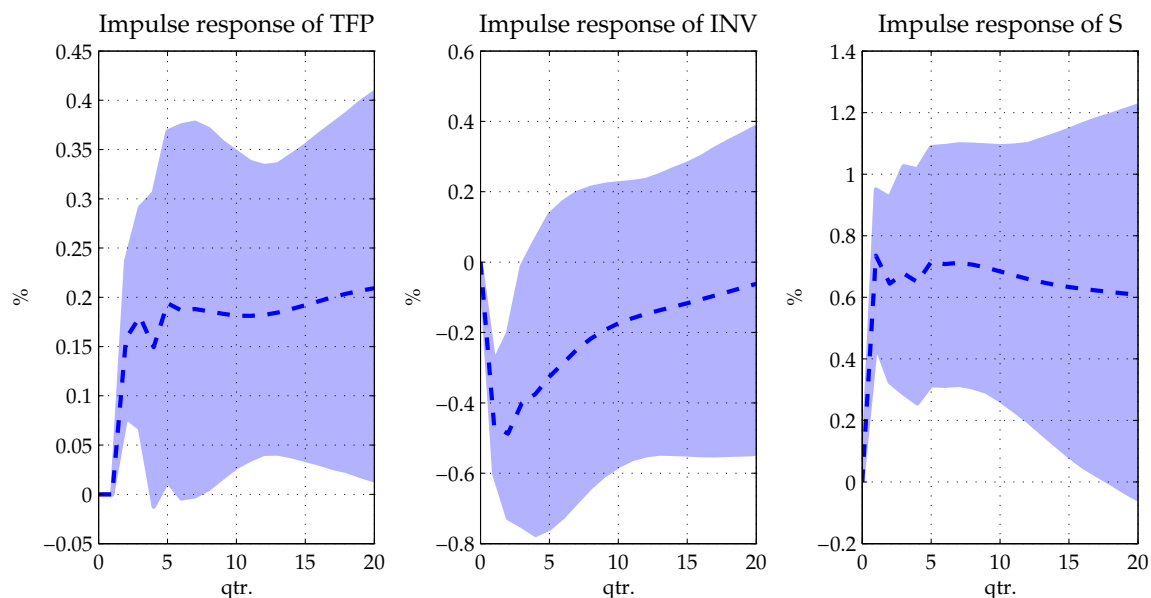
Figure 19: **Estimated impulse responses in the baseline specification.** The shaded blue area is the 10-90% coverage interval of the posterior distribution of identified impulse responses. The dashed line is the median of the posterior distribution of impulse reponses. Posterior intervals are constructed using 10000 draws from the posterior distribution.


Figure 20: **Estimated impulse responses with TFP adjusted for capacity utilization.**

Figure 21: **Estimated impulse responses with equipment and consumer durable TFP adjusted for capacity utilization.**

Figure 22: **Estimated impulse responses for the post-1984 sample.**

Figure 23: **Estimated impulse responses for the pre-1984 sample.**

	Horizon	1	5	10	20	40
TFP	Simulated	[0;0]	[0.04;0.25]	[0.13;0.57]	[0.17;0.68]	[0.17;0.73]
	Model	0	0.17	0.38	0.45	0.48
Inventories	Simulated	[0.02;0.59]	[0.03;0.71]	[0.02;0.72]	[0.11;0.55]	[0.28;0.50]
	Model	0.68	0.77	0.61	0.42	0.33
Sales	Simulated	[0.01;0.30]	[0.15;0.50]	[0.30;0.66]	[0.38;0.73]	[0.37;0.77]
	Model	0.07	0.18	0.34	0.40	0.45

Table 2: **Monte-Carlo experiment with news shock to TFP only.** The numbers in brackets in the "Simulated" lines are the set of identified FEV at different horizons; horizon 1 is the date of impact of the shock. The numbers in the "Model" lines are the model's FEV. Numbers are fractions of total forecast error variance.

	Horizon	1	5	10	20	40
TFP	Simulated	[0;0]	[0.01;0.22]	[0.05;0.49]	[0.08;0.59]	[0.08;0.64]
	Model (TFP)	0	0.17	0.38	0.45	0.48
	Model (LW)	0	0	0	0	0
Inventories	Simulated	[0.00;0.71]	[0.01;0.81]	[0.02;0.81]	[0.17;0.60]	[0.28;0.61]
	Model (TFP)	0.60	0.68	0.57	0.37	0.26
	Model (LW)	0.11	0.12	0.07	0.12	0.19
Sales	Simulated	[0.00;0.38]	[0.13;0.55]	[0.27;0.70]	[0.34;0.75]	[0.38;0.76]
	Model (TFP)	0.07	0.17	0.31	0.38	0.42
	Model (LW)	0.01	0.05	0.08	0.08	0.06

Table 3: **Monte-Carlo experiment with news shock to TFP and labor wedge.** The numbers in brackets in the "Simulated" lines are the set of identified FEV at different horizons; horizon 1 is the date of impact of the shock. The numbers in the "Model (TFP)" lines are the model's FEV with respect to TFP news shocks. The numbers in the "Model (LW)" lines are the model's FEV with respect to labor wedge shocks. Numbers are fractions of total forecast error variance.

	Horizon	1	5	10	20	40
TFP	Baseline	0	0.04	0.07	0.09	0.11
		[0;0]	[0.01;0.09]	[0.01;0.15]	[0.02;0.18]	[0.02;0.24]
	(1960Q1–1983Q4)	0	0.07	0.22	0.38	0.43
		[0;0]	[0.02;0.13]	[0.07;0.42]	[0.13;0.63]	[0.14;0.71]
	(1984Q1–2010Q4)	0	0.04	0.08	0.09	0.11
		[0;0]	[0.00;0.11]	[0.01;0.19]	[0.01;0.23]	[0.01;0.28]
	TFP adjusted	0	0.02	0.03	0.18	0.19
		[0;0]	[0.00;0.04]	[0.01;0.06]	[0.01;0.23]	[0.02;0.42]
	Equip-TFP adjusted	0	0.02	0.05	0.17	0.26
		[0;0]	[0.00;0.04]	[0.01;0.12]	[0.02;0.36]	[0.04;0.53]
Inventories	Baseline	0.61	0.20	0.11	0.09	0.09
		[0.22;0.94]	[0.03;0.43]	[0.02;0.26]	[0.01;0.20]	[0.01;0.21]
	(1960Q1–1983Q4)	0.66	0.26	0.14	0.26	0.38
		[0.31;0.94]	[0.08;0.47]	[0.05;0.26]	[0.08;0.51]	[0.10;0.70]
	(1984Q1–2010Q4)	0.56	0.13	0.12	0.14	0.15
		[0.12;0.95]	[0.04;0.28]	[0.03;0.26]	[0.02;0.34]	[0.02;0.36]
	TFP adjusted	0.58	0.20	0.15	0.14	0.15
		[0.16;0.94]	[0.06;0.43]	[0.04;0.32]	[0.03;0.30]	[0.04;0.32]
	Equip-TFP adjusted	0.57	0.20	0.18	0.18	0.20
		[0.13;0.96]	[0.06;0.42]	[0.05;0.37]	[0.04;0.41]	[0.03;0.47]
Sales	Baseline	0.33	0.18	0.16	0.16	0.16
		[0.11;0.51]	[0.04;0.32]	[0.04;0.31]	[0.03;0.32]	[0.03;0.34]
	(1960Q1–1983Q4)	0.36	0.31	0.28	0.33	0.37
		[0.16;0.54]	[0.16;0.48]	[0.14;0.45]	[0.16;0.52]	[0.17;0.59]
	(1984Q1–2010Q4)	0.33	0.12	0.12	0.13	0.14
		[0.07;0.56]	[0.02;0.28]	[0.01;0.29]	[0.01;0.33]	[0.01;0.35]
	TFP adjusted	0.45	0.38	0.38	0.36	0.36
		[0.12;0.74]	[0.04;0.75]	[0.04;0.73]	[0.05;0.70]	[0.06;0.69]
	Equip-TFP adjusted	0.46	0.39	0.39	0.39	0.37
		[0.10;0.77]	[0.03;0.78]	[0.03;0.79]	[0.03;0.78]	[0.04;0.74]

Table 4: **Estimated forecast error variance attributable to the identified news shock.** For each specification of the estimated model, the first line presents the median of the posterior distribution of the forecast error variance. The numbers in brackets are the 10-90% coverage interval of the posterior distribution of the forecast error variance.