Journal of Monetary Economics 79 (2016) 49–66

Contents lists available at ScienceDirect

Journal of Monetary Economics

journal homepage: www.elsevier.com/locate/jme

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

Nicolas Crouzet[a], Hyunseung Oh[b,*]

[a] Kellogg School of Management, Northwestern University, Evanston, IL 60208, USA
[b] Department of Economics, Vanderbilt University, Nashville, TN 37235, USA

ARTICLE INFO

Article history:
Received 30 January 2015
Received in revised form
21 March 2016
Accepted 23 March 2016
Available online 4 April 2016

Keywords:
News shocks
Business cycles
Investment
Inventories
Intertemporal substitution

ABSTRACT

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

1. Introduction

The sources of business cycles are an enduring subject of debate among macroeconomists. A recent literature argues that news shocks—shocks that change agents' expectations about *future* economic fundamentals without affecting *current* fundamentals—could play a leading role in aggregate fluctuations.[1] The theoretical conditions under which these shocks generate comovement among macroeconomic variables are, by now, well understood. However, no consensus has yet emerged regarding their empirical importance. While some estimates suggest that they account for as much as 60 percent of output volatility, other equally plausible methods obtain values as low as 10 percent.[2]

The goal of this paper is to use a new theoretical insight to pin down more precisely the quantitative contribution of news shocks to business-cycle volatility. This insight builds on the behavior of a business-cycle variable that is highly informative about news shocks but has so far been neglected in the literature: investment in finished-good inventories.

We focus on investment in finished-good inventories for two reasons. First, there is empirical evidence that finished-good inventories are a forward-looking variable that reflects changes in expectations about future economic conditions. For instance, Kesavan et al. (2010) find that finished-good-inventory data help improve forecasts of future sales. Second, finished-good inventories should react differently to changes in fundamentals that occur today than to changes that are expected to take place in the future.

* Corresponding author. Tel.: +1 615 322 8094.
E-mail address: hyunseung.oh@vanderbilt.edu (H. Oh).

[1] See Beaudry and Portier (2013) for a review of the contributions to this literature.

[2] For example, Beaudry and Portier (2006) and Schmitt-Grohé and Uribe (2012) estimate that the contribution of news shocks to the business-cycle volatility of output is above 40 percent, while Barsky and Sims (2011) and Khan and Tsoukalas (2012) find much smaller numbers.

http://dx.doi.org/10.1016/j.jmoneco.2016.03.005
0304-3932/

Temporary changes in productivity provide a straightforward illustration of this second claim. When productivity increases today, current income increases, so sales rise. Additionally, firms bunch production today in order to make the most out of the productivity increase. As a result, finished-good inventories also rise. Thus, when productivity changes today, sales and inventories should comove *positively*. On the other hand, when productivity is expected to increase tomorrow, the present value of income increases, which also leads to a rise in sales. However, since firms expect future production to be cheaper than current production, they satisfy this increase in sales not by producing more, but by depleting inventories. Thus, when productivity is expected to change tomorrow, sales and inventories should comove *negatively*. The sign of the comovement between sales and inventories should therefore help to distinguish between current and future expected changes in productivity.

Sections 2 and 3 formalize this intuition in the context of a standard news-driven business-cycle model. The model is identical to that of Jaimovich and Rebelo (2009) and Schmitt-Grohé and Uribe (2012), except that the economy also holds finished-good inventories. In this model, good news about the future indeed leads to a boom in the components of private sales (consumption and investment), but also generates a fall in inventories. The key mechanism behind this result is that good news causes agents to backload production and satisfy the boom in sales via inventory depletion—something they cannot do in an economy without inventories. This result holds across all plausible calibrations, and survives a number of extensions of the baseline model.

Section 4 uses this prediction to assess the quantitative contribution of news shocks to business-cycle volatility. We describe an empirical vector autoregression (VAR) strategy that identifies shocks using sign restrictions on inventories and sales, consistent with theoretical predictions. When applied to quarterly data, this strategy identifies shocks that explain less than 20 percent of the forecast error variance (FEV) of output at business-cycle frequencies. The intuition for this small estimate is that, in the data, inventories and sales tend to comove positively, so that any shock that generates negative comovement must have limited importance over the business cycle.

A cautious interpretation of this result is that it provides an upper bound to the contribution of news shocks to business-cycle volatility, since surprise innovations to other fundamentals could also be generating the comovement pattern.[3] Nevertheless, this result emphasizes that the business-cycle behavior of inventories is a potential challenge to the news view of business cycles.

Related literature: This paper emphasizes that news shocks can trigger intertemporal substitution in *production*, via inventory (dis)investment. This complements a strand of literature that shows how these shocks could cause intertemporal substitution in *consumption*, via capital (dis)investment (Jaimovich and Rebelo, 2009; Christiano et al., 2008). This literature has not analyzed inventory investment thus far.[4]

This paper also contributes to the recent business-cycle literature on inventory dynamics (Bils and Kahn, 2000; Fisher and Hornstein, 2000; Jung and Yun, 2006; Khan and Thomas, 2007; Wen, 2011; Kryvtsov and Midrigan, 2013). Although these models differ in their precise microfoundation for inventory holding, a feature shared by all is that inventories serve to shift production over time. The results of Section 3 help pinpoint the importance of this common mechanism. In particular, these results highlight a key parameter, the elasticity of intertemporal substitution in production (EISP), which summarizes the strength of this mechanism and is common to all these models.

Finally, the empirical approach of Section 4 is based on recent advances in the estimation of sign-restricted VARs (Arias et al., 2013; Moon et al., 2013). While sign restrictions have been used to identify monetary policy shocks (Faust, 1998; Uhlig, 2005), fiscal policy shocks (Mountford and Uhlig, 2008; Caldara and Kamps, 2012), and also news shocks (Beaudry et al., 2011), their application to inventory dynamics is, to our knowledge, new.

2. Inventories in a real business cycle framework

This section describes a simple but general model of inventory dynamics. The model has two main components. First, it borrows most of its structure from the Real Business Cycle (RBC) models of Jaimovich and Rebelo (2009) and Schmitt-Grohé and Uribe (2012). Second, it incorporates a simple inventory block: output produced today can be stored as inventory for use in subsequent periods. We use it as our baseline environment to study inventory dynamics in response to news shocks.

A central feature of inventory management is that it helps smooth out time variation in production costs. At the same time, inventories are costly to hold, both because of depreciation and foregone interest costs. In steady state, production costs are constant and there is no need to smooth them out. Thus, if cost smoothing were the sole purpose of inventory holdings, the steady-state level of inventories should be zero.

Since this is at odds with the large inventory stock held by firms in the data, the inventory literature has proposed various inventory-holding motives, corresponding to different microeconomic frictions. The leading microfoundations are demand for variety, stock-out risk, and inventory-ordering costs. A common feature of these microfoundations is that they

[3] Section 3 shows that surprise shocks to total factor productivity and the labor wedge both generate positive comovement between inventories and sales in the structural models analyzed in the paper, consistent with the intuition outlined above. To the extent that the identified shocks have a non-news component, that component must therefore reflect surprise changes in fundamentals other than total factor productivity and the labor wedge.

[4] To our knowledge, the only exception is Vukotic (2013), where inventories are introduced as a factor of production in the durable sector, and play a role similar to fixed capital investment.

imply a "target inventory-to-sales (IS) ratio" behavior: the IS ratio tends to revert back to its steady-state value, and transitory deviations occur only as shocks require.

The baseline model captures this feature by using a simple IS adjustment cost function. We view this as the most transparent way of introducing the fundamental mechanisms at play in leading inventory models into a standard RBC model, while retaining generality. Because of this generality, the main results carry through to the leading microfounded models of finished-good inventories, including the stock-elastic demand model of Bils and Kahn (2000) and the stockout-avoidance model of Kryvtsov and Midrigan (2013).[5]

2.1. Description of the baseline model

Economic environment: A representative agent maximizes expected 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, and ψ_t is an exogenous variable that shifts the marginal rate of substitution between consumption and leisure (a "labor wedge"). The period utility function takes the form proposed by Greenwood et al. (1988, henceforth GHH):

$$U(c, n; \psi) = \frac{1}{1-\sigma}\left(c - \psi\frac{n^{1+\xi^{-1}}}{1+\xi^{-1}}\right)^{1-\sigma}, \tag{2}$$

where ξ is the Frisch elasticity of labor supply and σ is the inverse of the elasticity of intertemporal substitution in consumption.[6] The production function is

$$y_t = z_t(u_t k_t)^{1-\alpha} n_t^{\alpha}, \tag{3}$$

where z_t is total factor productivity (TFP), $\alpha \in]0,1[$ is the degree of short-term decreasing returns to labor, k_t is physical capital available for production, and u_t is the utilization rate of capital. Physical capital evolves according to the law of motion:

$$k_{t+1} = i_t\left[1-\phi\left(\frac{i_t}{i_{t-1}}\right)\right] + (1-\delta_k(u_t))k_t. \tag{4}$$

The function $\phi(.)$ represents adjustment costs to investment. It is assumed to be twice differentiable and to satisfy $\phi(1) = \phi'(1) = 0$ and $\phi''(1) > 0$. Capital utilization costs are summarized by $\delta_k(.)$, which satisfies $\delta_k(1) = \overline{\delta}_k \in]0,1[$, $\delta_k'(.) > 0$ and $\delta_k''(.) > 0$.

Private sales, s_t, is the sum of consumption and investment:

$$s_t = c_t + i_t. \tag{5}$$

Goods to be used for consumption or investment can be obtained in two ways: either by producing more output, or by drawing down existing inventories. The relationship between inventories, output, and sales is captured by

$$s_t + inv_t + s_t\chi\left(\frac{inv_t}{s_t}\right) = y_t + (1-\delta_i)inv_{t-1}. \tag{6}$$

The stock of inventories, inv_t, depreciates at a rate of δ_i. Additionally, by adjusting the level of inventory relative to sales, the economy incurs a resource cost $\chi(\cdot)$ per unit of sales. The function $\chi(\cdot)$ captures the IS ratio adjustment costs mentioned above. It is assumed to be twice differentiable, decreasing, and convex. We discuss its economic interpretation below. This law of motion for inventories constitutes the only departure from standard news-driven RBC models; it replaces the resource constraint $s_t = y_t$.

Equilibrium: An equilibrium of this economy is a set of stochastic processes $\{c_t, n_t, i_t, k_{t+1}, u_t, y_t, s_t, inv_t\}_{t\geq 0}$ that maximize (1) subject to (3)–(6), given a law of motion for $\{z_t, \psi_t\}_{t\geq 0}$ and initial conditions k_0, inv_{-1}, z_{-1}, and ψ_{-1}.[7]

News shocks and surprise innovations: The two exogenous processes in our economy are total factor productivity z_t and the labor wedge ψ_t.[8] News shocks in period t are innovations to $\{x_t\}_{t\geq 0}$ such that $x_t = \mathbb{E}_{t-1}x_t$ but $\mathbb{E}_t x_{t+s} \neq \mathbb{E}_{t-1}x_{t+s}$ for some $s > 0$, for $x \in \{z, \psi\}$. Surprise innovations in period t are $\{x_t\}_{t\geq 0}$ such that $\mathbb{E}_{t-1}x_t \neq x_t$, for $x \in \{z, \psi\}$.

[5] This is discussed in more detail in Section 3.4.3. See also the online appendix for a detailed analysis of these models.

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

[7] The full set of first-order conditions characterizing an equilibrium are reported in the online appendix.

[8] The baseline model focuses on these shocks for two reasons. First, TFP data are the focus of many empirical papers that attempt to identify news shocks. Second, Schmitt-Grohé and Uribe (2012) find that news about the labor wedge is an important driver of business cycles.

2.2. Key features of the baseline model

Features unrelated to inventories: Readers familiar with the literature on news shocks will recognize three important features of this model: variable capacity utilization, investment adjustment costs, and the choice of preferences. As established by Jaimovich and Rebelo (2009), these features are in general necessary for news to generate comovement among business-cycle variables. With variable capacity utilization, output can increase in the short run even without a shift in current productivity. Investment adjustment costs force any future increase in capital to be smoothed out over time and thus guarantee a short-run increase in investment. Finally, GHH preferences allow current hours worked to increase in the short run, by muting the wealth effects on labor supply induced by positive news about the future. The baseline model thus generates positive short-run comovement of consumption, investment, output, and hours to news shocks.

Introducing inventories: The key departure of the baseline model from a standard environment is the law of motion for inventories (6). It changes the standard model in two main ways. First, in this model, sales can be met not only by producing, but also by drawing down inventories. Hence, demand does not necessarily drive output. Second, the IS ratio adjustment cost generates output losses, as captured by the term $s_t\chi(inv_t/s_t)$. The adjustment cost function $\chi(x)$ is assumed to be decreasing in x. This captures the idea that there are benefits to holding a sufficient amount of inventories relative to sales, which creates a yield to holding positive amounts of inventories.[9]

Optimal choice of inventories: The first-order condition for inventory choice is

$$mc_t = -mc_t\chi'\left(\frac{inv_t}{s_t}\right) + \mathbb{E}_t\left[\beta(1-\delta_i)mc_{t+1}\right], \tag{7}$$

where mc_t denotes the shadow cost of output, that is, the marginal cost of production expressed in utility units.[10] The left-hand side of this equation represents the cost of stocking an additional unit of inventory. This is equal to the marginal cost of production, mc_t. The right-hand side represents the two benefits of adding this extra unit into the stock of inventories. The first term represents the convenience yield: with $-\chi'(\cdot)>0$, the additional unit of inventory reduces current resource costs. The second term represents the reduction of future production costs since more goods become available next period.

In steady state, the IS ratio satisfies:

$$-\chi'(IS) = 1-\beta(1-\delta_i). \tag{8}$$

The right-hand side of this equation represents the steady-state (or average) cost of holding inventories. It consist of two parts: (i) the opportunity cost of holding goods, captured by the steady-state gross interest rate $1/\beta$, and (ii) depreciation costs. The right-hand side is the steady-state marginal benefit of holding inventory: a reduction in IS adjustment costs. The steady-state value of the IS ratio is such that this marginal benefit is equated to inventory holding costs.

Away from the steady-state, the IS ratio will only deviate from this value in response to variation in marginal cost over time. To see this, note that Eq. (7) can be approximated up to a first order as

$$\widehat{inv}_t = \hat{s}_t + \eta\hat{\gamma}_t, \tag{9}$$

where hatted variables represent log deviations from its respective steady state. The parameter $\eta>0$ will be discussed below. The variable γ_t is the expected discounted growth rate of marginal cost,

$$\gamma_t \equiv \mathbb{E}_t\left[\beta(1-\delta_i)\frac{mc_{t+1}}{mc_t}\right]. \tag{10}$$

The inventory choice condition (9) captures the main forces driving inventories in this model. First, absent movements in γ_t, inventories move one for one with sales, and the IS ratio remains constant and equal to the steady-state value pinned down by Eq. (8). This can be thought of as a "demand channel" introduced by the IS ratio adjustment cost $\chi(\cdot)$, which induces inventories to move in lockstep with sales.

The second force governing inventory dynamics is the term $\eta\hat{\gamma}_t$. This term reflects changes in inventories due to changes in the relative cost of production over time—an "intertemporal substitution channel." When there is an increase in the future expected marginal cost relative to the current marginal cost, $\hat{\gamma}_t$ is positive and end-of-period inventories increase relative to sales: it is indeed cheaper to produce tomorrow than today.

The strength of this channel is governed by the parameter η, which plays a crucial role in what follows. This parameter determines the degree of intertemporal substitution in production. Given the difference between the marginal cost today and tomorrow, a large value of η corresponds to a strong intertemporal substitution channel and hence large a inventory adjustment. For brevity, we refer to it as the elasticity of intertemporal substitution in production (EISP) in the rest of the paper. The relationship between the EISP and the structural parameters is given by

$$\eta = \frac{1}{|\epsilon_{\chi'}|}\frac{\beta(1-\delta_i)}{1-\beta(1-\delta_i)} > 0, \tag{11}$$

[9] Additionally, the adjustment cost function is assumed to be convex, since otherwise it may be optimal to produce in infinite amounts.

[10] Formally, mc_t is the Lagrange multiplier on constraint (6) in the representative agent's problem; see the online appendix for details.

Table 1
Calibrated parameters and target moments in the baseline model.

Parameter	Value	Description
β	0.99	Subjective discount factor
σ	1	Household elasticity of intertemporal substitution
ξ	2.5	Frisch elasticity of labor supply
α	0.67	Labor share of income
$\overline{\delta}_k$	0.025	Capital depreciation rate
$\phi_I''(1)$	9.11	Investment adjustment cost
u	1	Capacity utilization
$\delta_k''(1)/\delta_k'(1)$	0.34	Depreciation elasticity of capacity utilization
n	0.2	Steady state hours worked
ρ_z	0.99	Persistence of the TFP process
ρ_ψ	0.95	Persistence of the labor wedge process
δ_i	0.025	Inventory depreciation rate
IS	0.75	Inventory-to-sales ratio
η	10	Elasticity of intertemporal substitution in production (EISP)
χ	0.01	Steady-state resource cost

where $|\epsilon_{\chi'}|$ is the absolute value of the steady-state elasticity of $\chi'(\cdot)$. Two things are worth noting. First, the EISP is large whenever inventory holding costs are small, that is, when $\beta(1-\delta_i)$ is close to 1. In the extreme case where the holding cost of inventories is zero, the EISP is infinitely large. Second, the EISP depends on the elasticity of $\chi'(\cdot)$, the convenience yield of inventories. With a high elasticity, the convenience yield quickly falls, and inventory adjustment is limited.

In sum, inventory choices in this model are governed by a demand channel and an intertemporal substitution channel. Moreover, the strength of the intertemporal substitution channel is captured by the EISP, η.

3. How should inventories respond to news? A structural sign restriction

This section shows that in plausible calibrations of the model of the previous section, news about future economic fundamentals generate negative comovement between inventories and sales. We first illustrate this result with a baseline calibration and then establish it across plausible calibrations.

3.1. A baseline calibration

The baseline calibration of the model of Section 2 is reported in Table 1. For parameters unrelated to inventories, it follows the values calibrated and estimated in Schmitt-Grohé and Uribe (2012). Exogenous variables follow an AR(1) process with surprise and news shocks. Inventory-related parameters are set to match three targets: (i) the physical rate of depreciation of the inventory stock, (ii) the steady-state IS ratio, and (iii) the EISP, η.[11]

The EISP is calibrated as follows. The inventory block of the model is simple and abstracts from exact microfoundations for inventory-holding. Nevertheless, in the microfounded inventory models of Bils and Kahn (2000) and Kryvtsov and Midrigan (2013), the condition governing optimal inventory choice is similar to Eq. (9), and there is an equivalent reduced-form parameter $\tilde{\eta}$ capturing the EISP. The value of the EISP is a function of the structural parameters, either of the elasticity of sales to stock (Bils and Kahn, 2000) or of the stockout probability (Kryvtsov and Midrigan, 2013).[12]

This equivalence result is used to calibrate the EISP of the baseline model; specifically, the EISP is set to the lower bound implied by these microfounded models. The appendix discusses the derivation of this lower bound; it shows that in these models, η is at least 10 and typically much larger. The choice of the lower bound value in the baseline calibration is motivated by the fact that, as will become clear in a few paragraphs, any value above this lower bound generates similar predictions.

3.2. The effect of news in the baseline calibration

Fig. 1 reports the impulse responses of several variables to four-period positive news shocks to TFP and the labor wedge. To be precise, at period 0, the agent receives a signal that future TFP ($\mathbb{E}_0 z_4$) will increase or that the future labor wedge ($\mathbb{E}_0 \psi_4$) will decrease.[13]

[11] While the steady-state resource cost χ does not not have a clear empirical counterpart, its precise value has no bearing on our results.
[12] This result is formally established in the online appendix.
[13] A positive news shock is defined as a shock that generates an increase in sales. When the labor wedge is expected to decrease, the agent expects to face less disutility of working in the future. This will boost current sales.

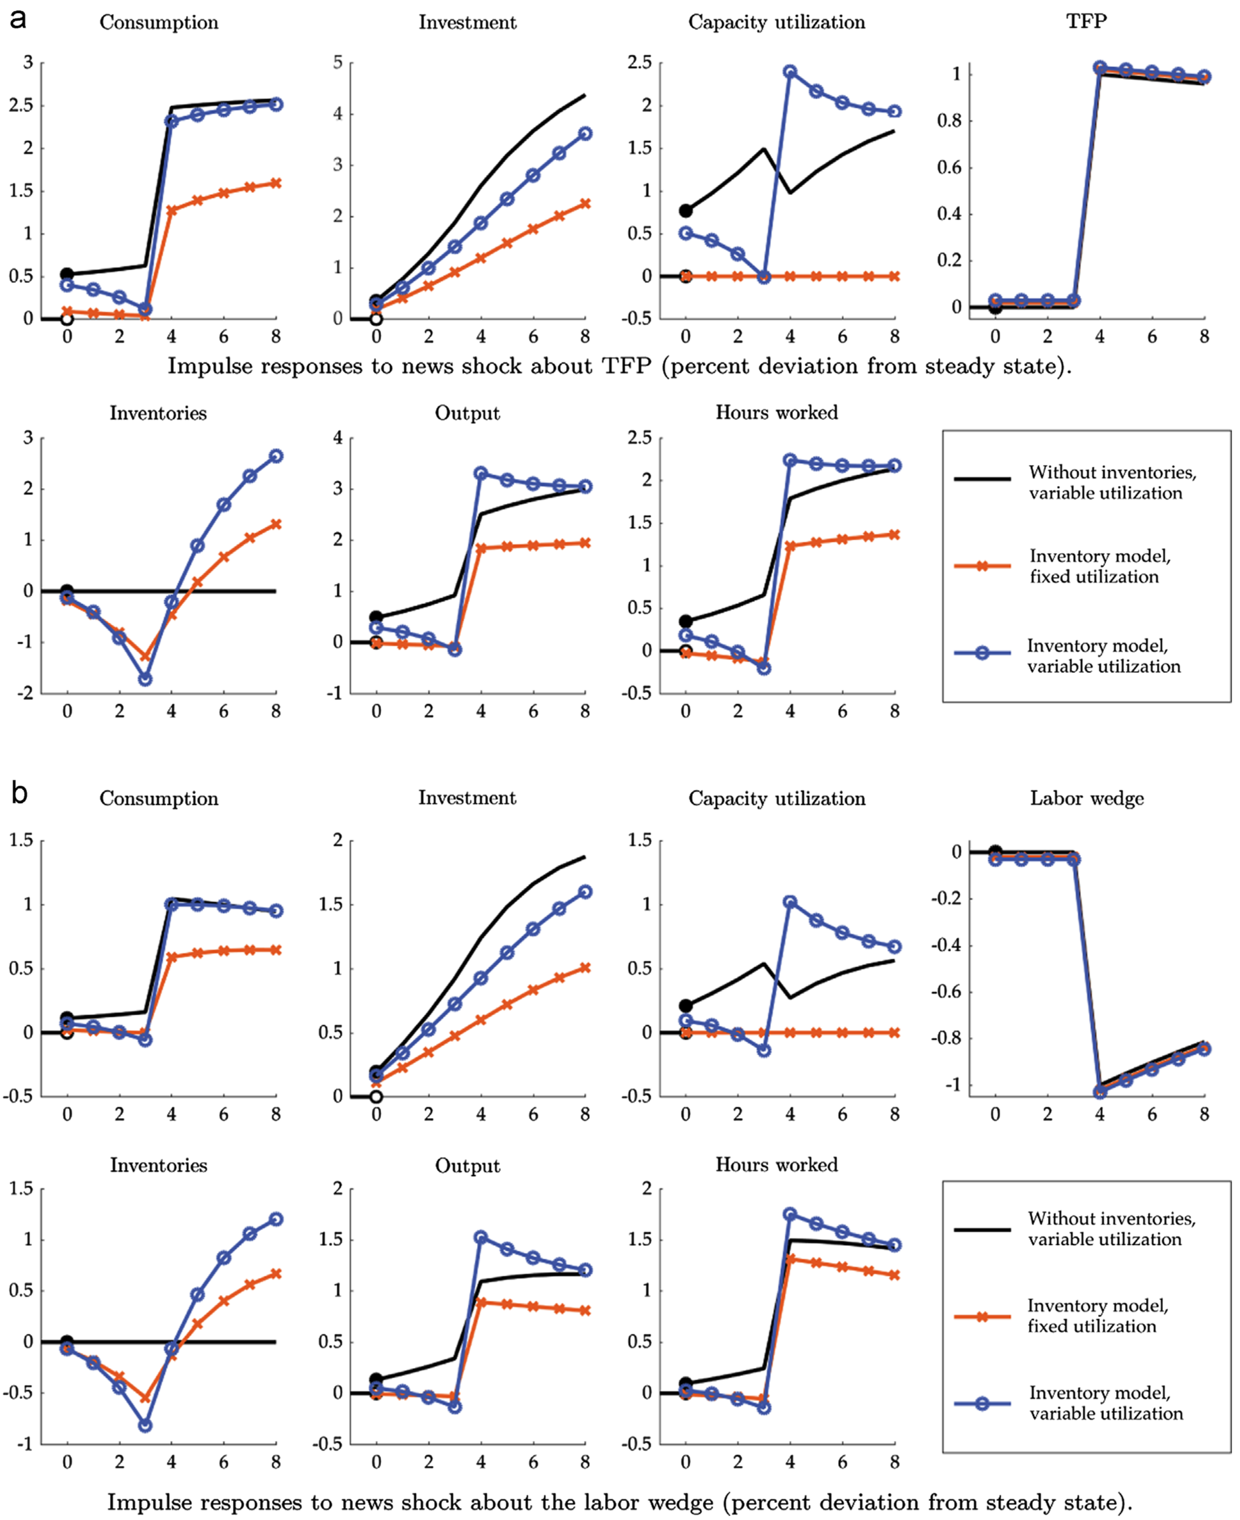

Fig. 1. Impulse responses to news shocks in three variants of the model. *Note*: The blue circled line is the baseline model. The orange crossed line is the baseline model without variable capacity utilization. The black line is the baseline model without inventories. (a) Impulse responses to news shock about TFP (percent deviation from steady state). (b) Impulse responses to news shock about the labor wedge (percent deviation from steady state) (For interpretation of the references to color in this figurecaption, the reader is referred to the web version of this paper.).

Note first that consumption and investment increase immediately and during all the subsequent periods. Consumption increases because of the wealth effect associated with the good news, and investment increases because of the presence of investment adjustment costs.

On the other hand, inventories fall. The fall is large, persistent, and reaches its trough in the period preceding the realization of the change in fundamentals. At the same time, output remains mostly unchanged until period 4, when the fundamentals change. That is, the increase in sales during the news period is almost entirely met by inventory disinvestment.

Thus, in this calibration, in response to news, inventories fall, while private sales (the sum of consumption and investment) increase.

3.3. Understanding the result

In order to understand the negative comovement between inventories and sales in response to news, we start by analyzing the impact response when capacity utilization is fixed: $u_t = 1$.

3.3.1. No variable capacity utilization: an analytical restriction

On impact, news shocks do not change fundamentals ($\hat{z}_t = \hat{\psi}_t = 0$). Since capital is fixed in the short run, the following relationship among output, marginal cost, and marginal utility holds:

$$\widehat{mc}_t - \hat{\lambda}_t = \omega \hat{y}_t, \tag{12}$$

where λ_t is the marginal utility of consumption. This relationship is derived by combining the optimal consumption and labor decisions with the production function. In a model without inventories, sales and output are the same. As a result, the marginal cost of output would be equal to the marginal utility of consumption ($\widehat{mc}_t = \hat{\lambda}_t$), and output would be fixed. In an economy with inventories, instead, marginal cost can respond to output movements in the short run. The short-run elasticity of marginal cost to output, ω, is given by

$$\omega = \frac{1 + (1-\alpha)\xi}{\alpha \xi}. \tag{13}$$

This parameter captures marginal-cost movements due to short-run decreasing returns to scale (represented by α) and to the disutility of labor supply (represented by ξ). In combination with the EISP, it governs the short-run response of the economy to news, as established in the following proposition.

Proposition 1 (*Impact response to news*). *Assuming fixed capacity utilization, when news arrive, inventories and sales positively comove on impact if and only if*

$$\eta \omega < \frac{\kappa}{\kappa + IS(1-\beta)(1-\delta_i)}, \tag{14}$$

where $\kappa > 0$ is the steady-state output-to-sales ratio in the economy. In particular, a necessary condition for positive comovement is

$$\eta \omega < 1. \tag{15}$$

Proof. See the appendix.□

To understand this result, it is useful to think through the effects of good news about future productivity. Because of a combination of the wealth effect on consumption and investment adjustment costs, sales increase on impact. Additionally, the news indicate that marginal cost will be lower in the future. When the EISP is large, it is optimal for production to be delayed and to satisfy the sales boom out of inventories. As a result, with a high EISP, inventories are likely to fall in response to the shock.

The proposition indicates that the only potential countervailing force to this intertemporal substitution channel would be a low value of ω. When ω is very small, marginal cost barely responds to changes in the level of output produced. In this case, there is no need for inventory investment to smooth out variation in production costs, since these fluctuations are relatively small.

As discussed above, and explained in detail in the appendix, the EISP is at least $\eta \geq 10$ and typically much larger in microfounded inventory models. The value of ω, on the other hand, has been at the center of debate in the monetary economics literature, and the lower bound suggested in the literature is $\omega \geq 0.33$, following Dotsey and King (2006). Thus, a lower bound for the product $\eta\omega$ is 3.3.

This violates the condition of Proposition 1. In other words, even at the lower bound of ω and the EISP, the model cannot generate positive comovement of inventories and sales in response to news. Thus, the negative-impact comovement of inventories and sales is a structural restriction implied by the model, in the sense that it occurs for any plausible calibration of the model.

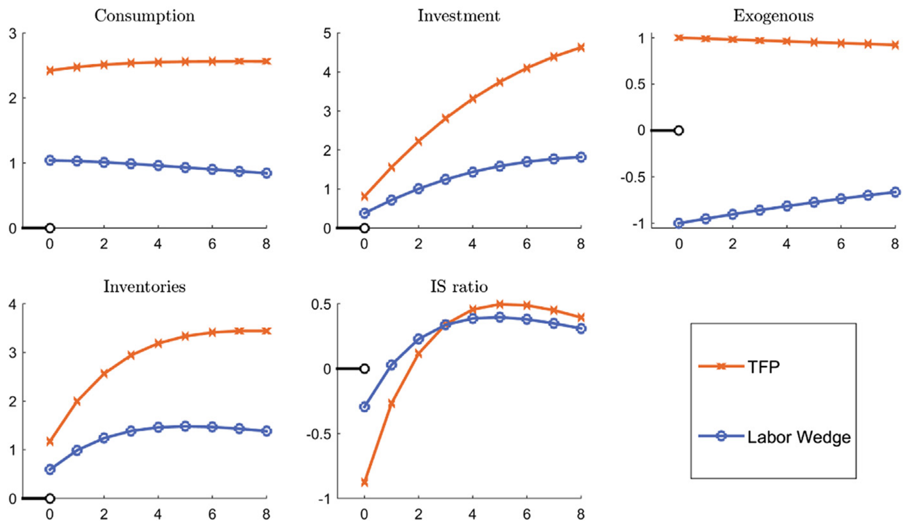

Fig. 2. Impulse responses to non-news shocks in the baseline model (percent deviation from steady state).

3.3.2. The role of variable capacity utilization

As indicated by Fig. 1, the negative comovement of inventories and sales in response to news also occurs when capacity utilization is allowed to vary.[14] Instead, relative to a model without inventories, the response of utilization is muted.

Why is this so? In a model without inventories, such as that in Jaimovich and Rebelo (2009), capacity utilization increases with news about a future rise in productivity. This is because with a future increase in productivity, the presence of investment adjustment costs leads to an increase in capital investment today. The increase in capital investment generates a fall in the value of installed capital. At the same time, the shock generates a fall in the marginal value of income. Overall, the fall in the value of installed capital is steeper than the fall in the marginal value of income, and therefore capacity is utilized more to satisfy the additional demand.

In an economy with inventories, on the other hand, capacity utilization increases far less in the short run. This is because of the interaction between preferences and the role of inventories. With GHH preferences, the marginal value of income is

$$\lambda = \left(c - \psi \frac{n^{1+\xi^{-1}}}{1+\xi^{-1}} \right)^{-\sigma}. \tag{16}$$

With inventories, the increase in consumption and investment can be matched by depleting inventories, rather than by working more or utilizing capital more intensively. Therefore, hours do not necessarily rise along with the increase in consumption. The limited response of hours generates a steeper fall in the marginal value of income. In turn, a steeper fall in the marginal value of income means that Tobin's Q falls less, in relative terms. This implies that increased capacity utilization (which uses up installed capital) is a less attractive way of matching the increase in demand.

3.4. Generalizing the results

We next discuss the generality and scope of the finding that news shocks generate negative comovement between inventories and sales.

3.4.1. Do surprise innovations generate positive comovement?

The impulse responses reported in Fig. 2 indicate that, in response to surprise innovations in fundamentals, inventories, consumption, investment, and output all comove positively. The short-run response of the inventory-to-sales ratio is also consistent with its observed countercyclicality at business-cycle frequencies. For surprise innovations, the model's predictions are broadly consistent with the observed behavior of inventories and sales over the business cycle. Thus, the negative comovement of inventories and sales does not hold in response to surprise innovations, so that it is an identifying feature of news shocks.

3.4.2. Can adjustment costs induce inventory accumulation in the short run?

The previous analysis suggests that additional channels are necessary in order to generate a short-run increase in production in response to news. Adding adjustment costs to output is one way achieve this goal: with an additional production-smoothing channel, the future increase in production may be front-loaded, resulting in an early accumulation of

[14] In that case, however, no analytical result analogous to Proposition 1 exists.

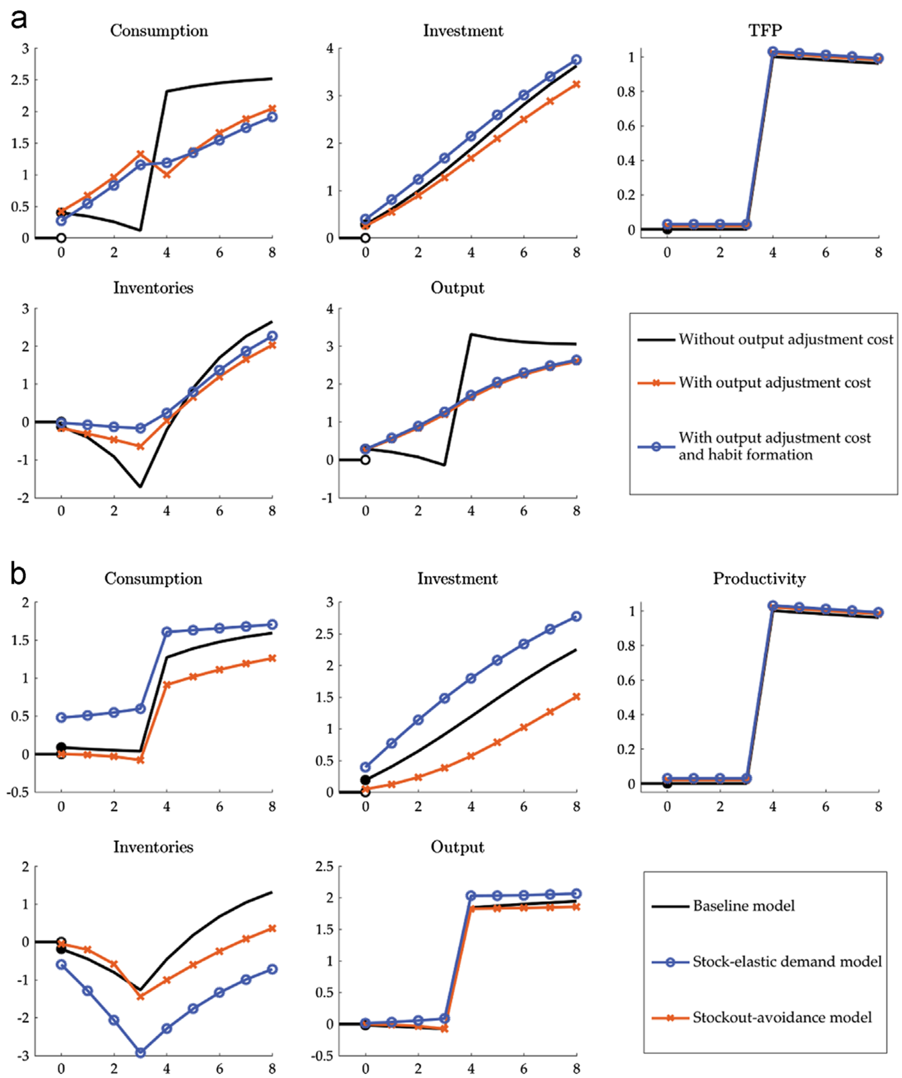

Fig. 3. Robustness of impulse responses: adjustment costs (top), microfoundations for inventory holding motive (bottom). (a) Impulse responses to news shock about TFP (percent deviation from steady state), in the baseline model (thick black line), with output adjustment cost ($\phi_y''(1)=6$, crossed orange line), and with both output adjustment costs and consumption habit formation ($\phi_y''(1)=6$, b=0.6, circled blue line). (b) Impulse responses to news shock about TFP (percent deviation from steady state), in the baseline model (thick black line), and in two models with micro-founded inventory holding motives: the stock-elastic demand model (blue circled line) and the stockout-avoidance model (orange crossed line) (For interpretation of the references to color in this figurecaption, the reader is referred to the web version of this paper.).

inventories. This production-smoothing channel can be introduced in the model of Section 2 using the following output adjustment cost function:

$$\phi_y\left(\frac{y_t}{y_{t-1}}\right)y_t, \tag{17}$$

where $\phi_y(1)=\phi_y'(1)=0$ and $\phi_y(1)''>0$. Panel (a) of Fig. 3 reports the responses of the model with and without adjustment costs. Production increases with adjustment costs, but inventories still fall.

The logic behind this result is that with adjustment costs to production, output increases steadily even in the short run. This is achieved by increasing capacity utilization and hours worked. With higher utilization, capital depreciates quickly and investment demand increases. Moreover, since the marginal product of labor is high, labor income is higher, which increases consumption.[15] The increase in production thus generates a higher demand for investment and consumption, leaving little room for inventory accumulation.

[15] In a competitive market equilibrium, real wages increase, which leads to more labor income for the household.

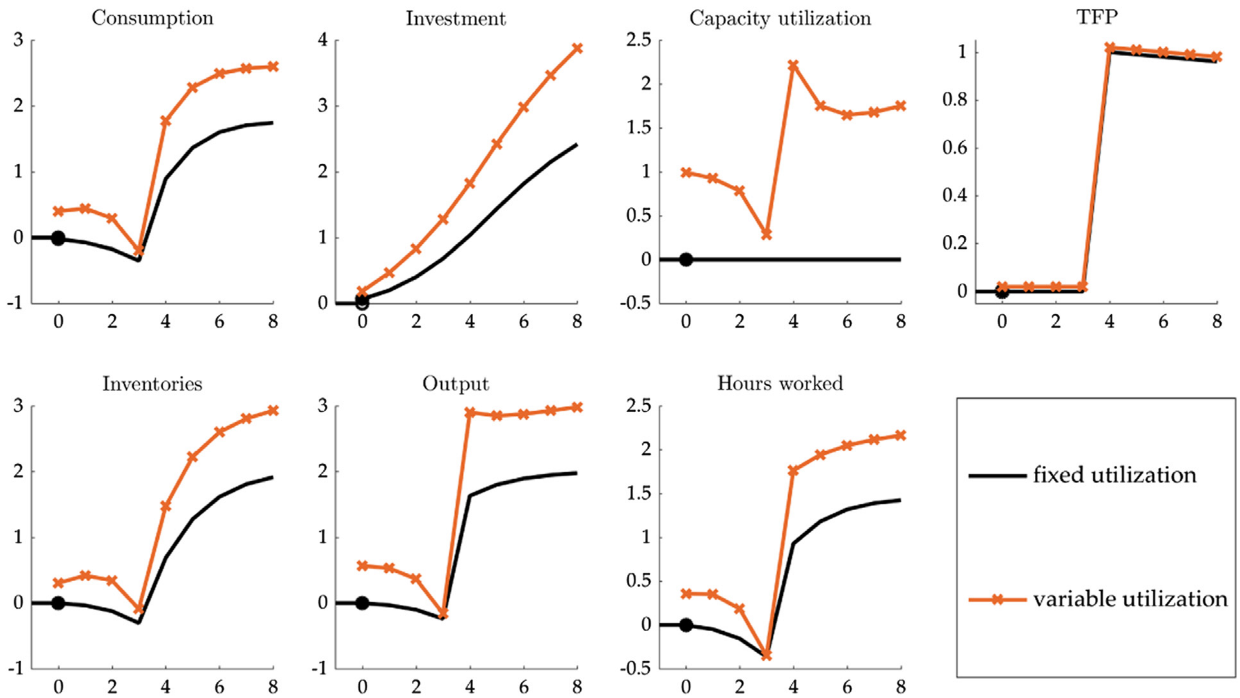

Fig. 4. Impulse responses (percent deviation from steady state) in the baseline model, when the EISP is small ($\eta = 0.5$).

In order to moderate this initial increase in consumption, habit formation can be introduced in the period utility function, as follows:

$$U_t = \frac{1}{1-\sigma}\left(c_t - bc_{t-1} - \psi_t \frac{n_t^{1+\xi^{-1}}}{1+\xi^{-1}}\right)^{1-\sigma}, \tag{18}$$

where b is the consumption-habit-formation parameter. Panel (a) of Fig. 3 indicates that the negative comovement between inventories and sales survives the introduction of habit formation.

Habit formation changes the dynamics of consumption with regards to news shocks in two ways. On the one hand, with consumption habits, the future increase in consumption tends to be front-loaded. On the other hand, habits imply that jumps in consumption should be smaller. The two channels offset each other, and compared with the model without habit formation, consumption changes only by a small amount. At the same time, the change in investment offsets the change in consumption, and sales remain at a similar level.

3.4.3. *Is the negative comovement prediction robust to microfoundations?*

One may wonder whether the precise microfoundations for inventory holding matter for our result. Panel (b) of Fig. 3 reports the impulse responses to good news about TFP in the two leading microfounded models of finished-good inventories, the stock-elastic demand model (Bils and Kahn, 2000) and the stockout-avoidance model (Kryvtsov and Midrigan, 2013). Both models predict similar responses to the model of Section 2: the drop in inventories is large and persistent. In fact, it is larger than in that model. This is because, in both the stock-elastic demand model and the stockout-avoidance model, the EISP is larger than in the model of Section 2. The online appendix derives analytical results similar to Proposition 1 for both models and establishes formally that no plausible calibration of either model can generate positive comovement between inventories and sales. Thus, the main finding of this section is robust to the precise microfoundations for inventory holdings.

3.4.4. *What happens when EISP is small?*

Finally, we study the predictions of the model when the EISP is outside of the range predicted by microfounded models. Fig. 4 reports the impulse responses to news about TFP with $\eta = 0.5$. By virtue of Proposition 1, on impact, inventories and sales should comove positively in the model without variable capacity utilization. They also comove positively in the model with variable capacity utilization, as reported in Fig. 4. However, there is a strong tendency of inventory disinvestment after impact. In fact, inventory disinvestment persists throughout the anticipation period. Thus, even when the EISP is outside its plausible range, inventory investment and sales (consumption plus investment) comove negatively after impact. The intertemporal substitution mechanism highlighted in this section is thus remarkably robust.

4. An empirical evaluation of news using inventory data

The analysis of Section 3 suggests that inventories must fall when there is an increase in consumption and investment induced by news shocks. This section uses this structural restriction to empirically estimate the importance of news shocks.

We take the approach of estimating a structural VAR with sign restrictions. This approach is appealing because it allows one to remain agnostic about other identifying features of news shocks and therefore to identify them without misspecification concerns. The main drawback is that the identification of shocks is only partial: it may also be capturing non-news shocks that cause inventories to comove negatively with consumption and investment. As discussed in Section 3, none of the "surprise" innovations analyzed so far have this feature, which may alleviate concerns about partial identification. Nevertheless, a strict interpretation of the results to follow is that they provide an upper bound on the contribution of news shocks to business-cycle volatility.

4.1. Data

The main empirical exercise uses four observables: inventories, consumption, investment, and output. Consumption includes nondurables and services, investment includes fixed investment and durables, and output is GDP. Inventories are measured using nonfarm private inventories as a whole.[16] However, the results are not sensitive to the type of inventories used for estimation: in particular, they carry through using only retail trade inventories, as discussed below. All data are seasonally adjusted and expressed in real per-capita terms. The sample period is 1955Q1–2006Q4.[17]

4.2. Baseline specification and estimation

The baseline identification strategy imposes that on impact, when there is disinvestment in inventories, consumption and fixed investment increase.[18] The following VAR is estimated:

$$X_t = A + B(L)X_{t-1} + U_t. \tag{19}$$

For X_t, log levels of each variable are used, in order for the estimation to be robust to the presence of cointegrating relations. The main specification includes a constant term and four lags.[19] We use the Bayesian sign-restriction approach developed by Arias et al. (2013) and Moon et al. (2013).[20] The prior for the reduced-form coefficients of the VAR is a multivariate normal inverse Wishart. Draws from the unrestricted posterior of structural parameters are obtained, and only those draws that satisfy the sign restriction are kept.[21]

4.3. Impact sign restrictions

Panel (a) of Fig. 5 reports the posterior distribution of the FEV of output with respect to the identified shock, at different horizons. The posterior has a sharp mode close to zero, and the median is close to 20 percent at most horizons. Panel (a) of Fig. 6 plots the set of identified impulse responses. The median response to the identified shock generates a persistent boom in consumption and investment and a moderate boom in output. The fall in inventories is short lived; on average, positive inventory investment occurs immediately after the initial disinvestment, and the stock of inventories becomes positive after three quarters. In the model, this occurs when when good news are expected to materialize in the near future. Therefore, the shock identified by the sign restriction resembles short-horizon news, with news lasting for only one period.

The identification scheme used so far only imposes impact restrictions, and is therefore not able to distinguish between short- and long-horizon news shocks. The results of the previous section however suggest that with longer-horizon news, the negative comovement between inventories and sales persists after impact. The next step is therefore to impose restrictions beyond impact.

[16] Nonfarm private inventories include input inventories as well as finished-good inventories. While a portion of these input inventories are work-in-progress inventories, the remainder are goods that are *not* used for current production or consumption, and thus potentially fall under the purview of the model. Since material inventories and finished-good inventories together make up 80 percent of nonfarm private inventories on average, nonfarm private inventories can be used as a benchmark. However, the complexity of the production structure may make it difficult to separate unused input inventories from work-in-progress inventories. The robustness section therefore looks at retail trade inventories separately. We thank an anonymous referee for pointing this out.

[17] The online appendix contains the exact sources of the data used to construct consumption, output, investment, and inventories.

[18] On impact, a fall in inventories is equivalent to a fall in inventory investment, since the impulse responses are computed relative to the steady state.

[19] The results to follow are not sensitive to the number of lags.

[20] In particular, this approach does not use penalty functions because, as pointed out by Arias et al. (2013), they may implicitly impose additional restrictions on variables, beyond those that constitute the identifying scheme.

[21] The results reported here use 10,000 posterior draws.

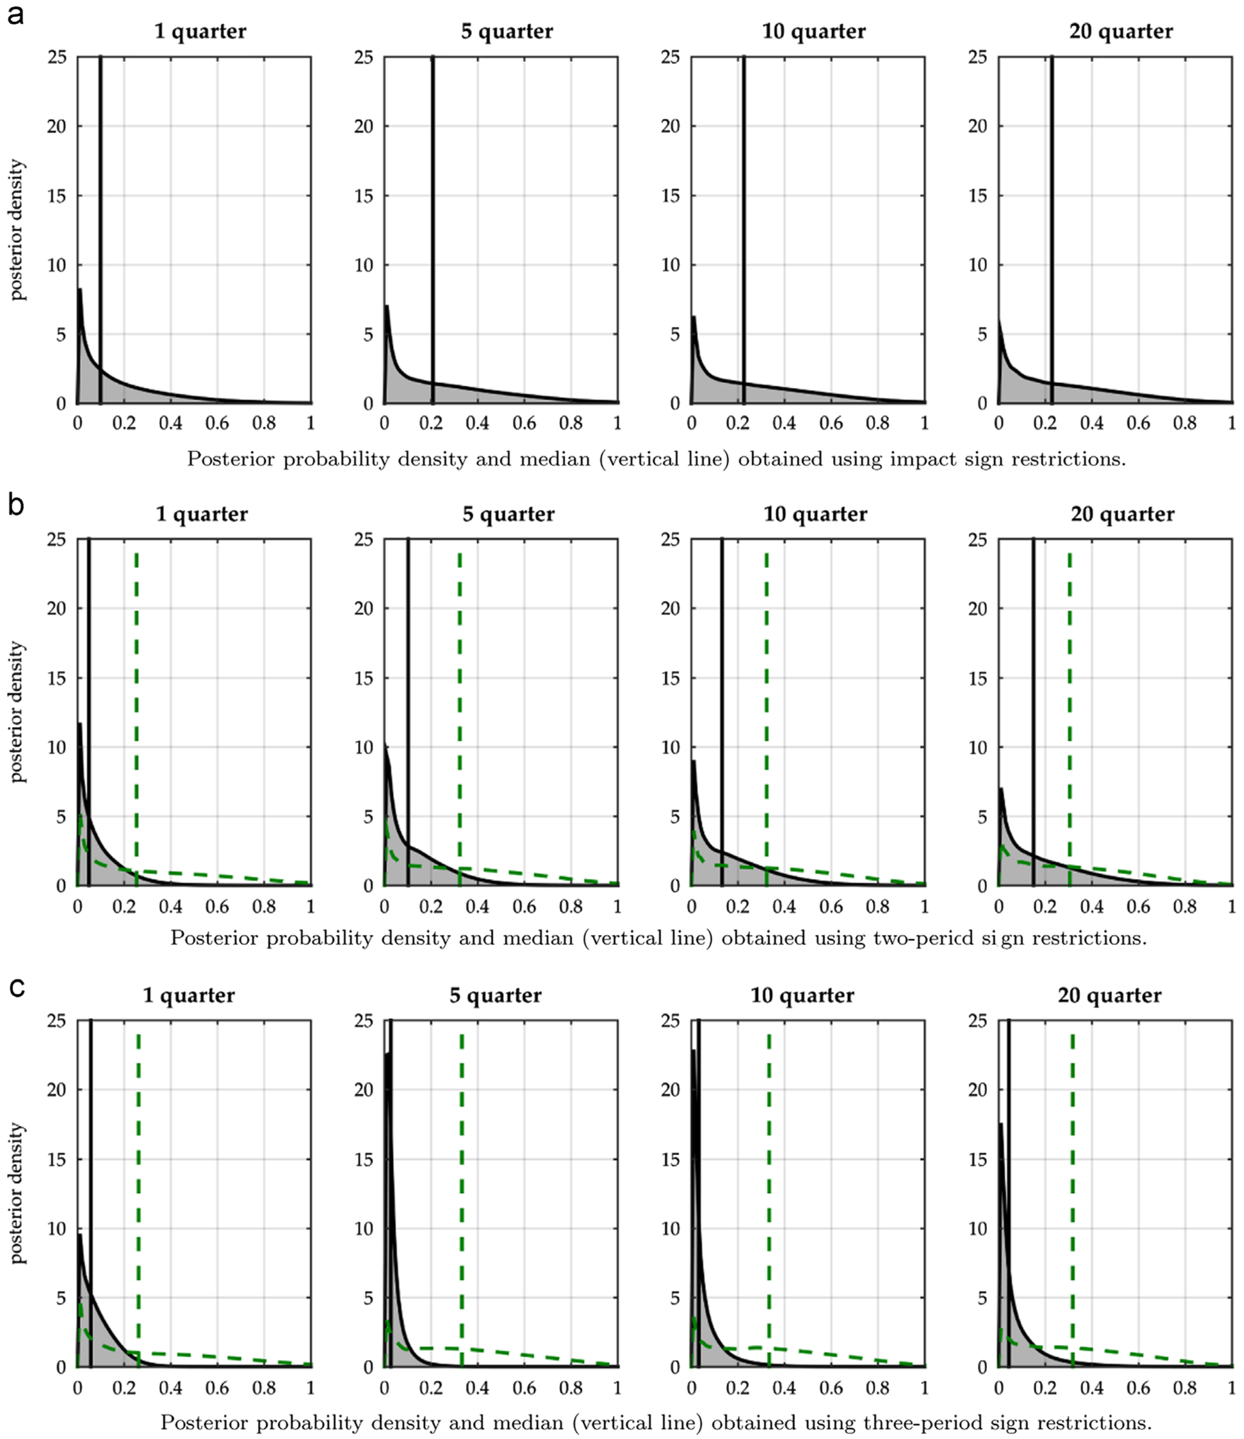

Fig. 5. Share of FEV of output at horizons 1, 5, 10, 20, attributable to news shocks. Note: News shocks are identified using sign restrictions on inventories, consumption, and investment. *Top panel*: impact restrictions; *middle panel*: two-period dynamic restrictions; *bottom panel*: three-period dynamic restrictions. The solid black lines correspond to the full identification scheme. The dashed green lines correspond to a shock identified using only restrictions on consumption and investment. See main text for details on the identification schemes. (a) Posterior probability density and median (vertical line) obtained using impact sign restrictions. (b) Posterior probability density and median (vertical line) obtained using two-period sign restrictions. (c) Posterior probability density and median (vertical line) obtained using three-period sign restrictions.

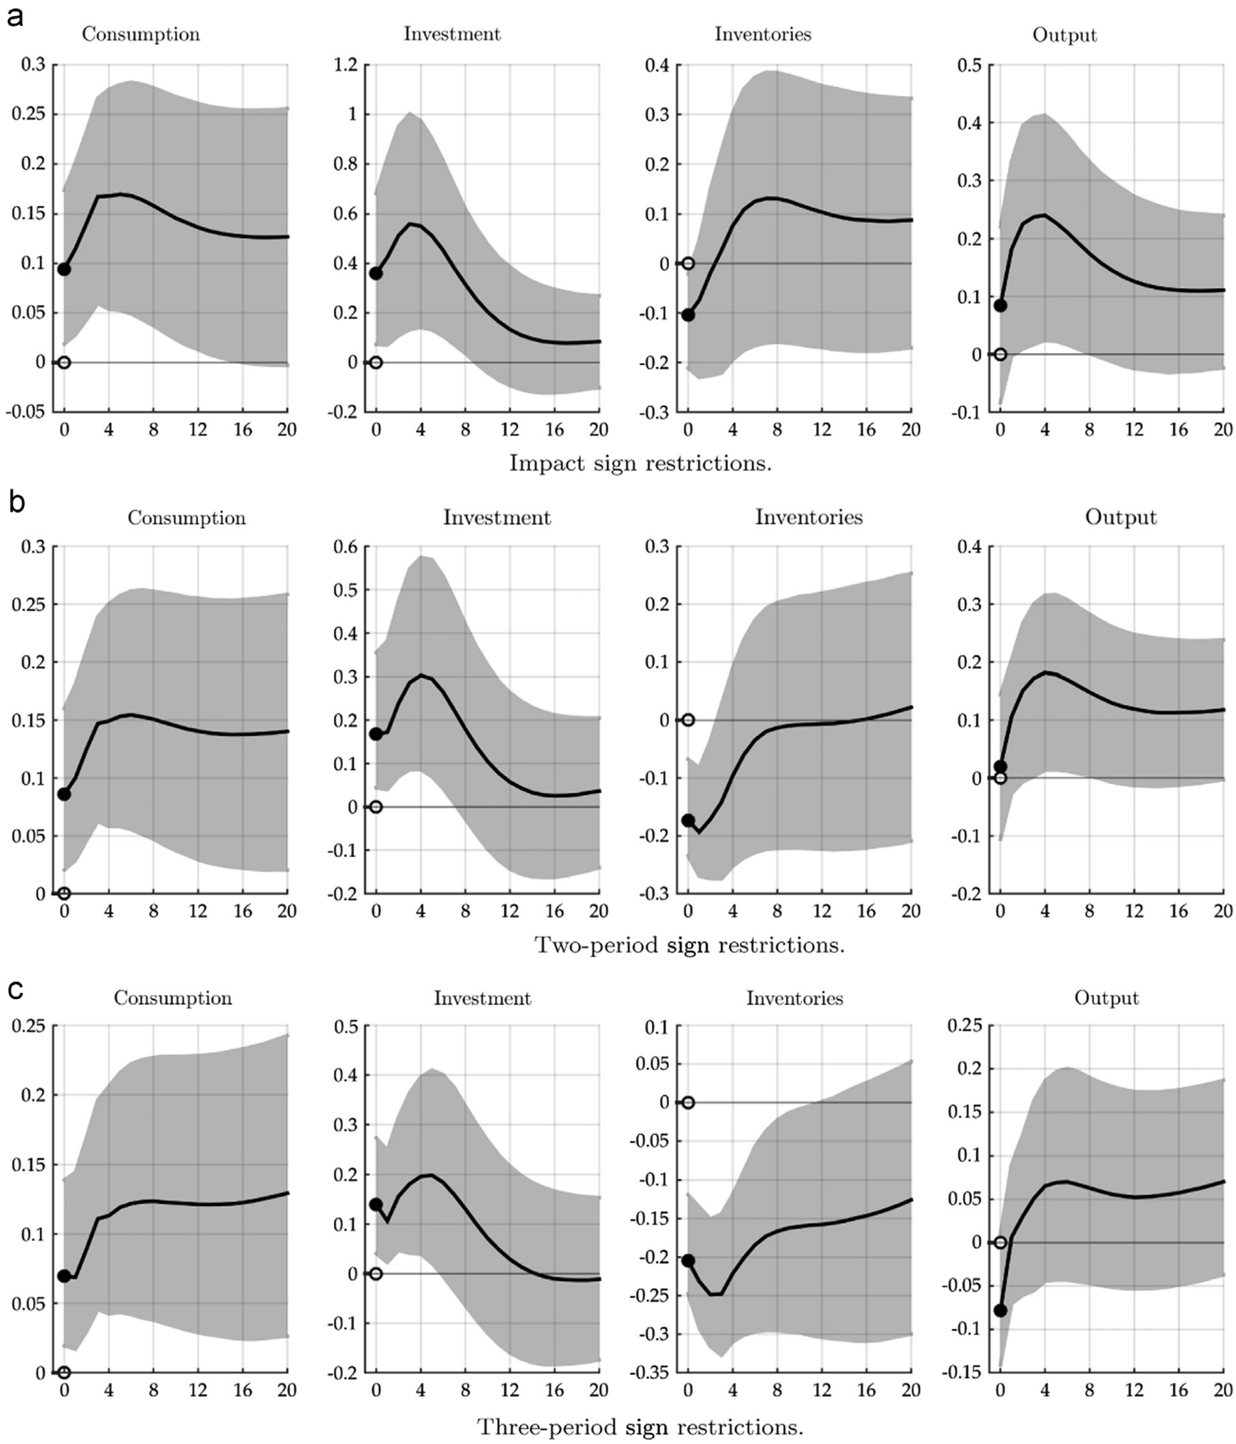

Fig. 6. Median (solid line) and 80% credible set (shaded area) for the impulse responses to news shocks. Note: News shocks are identified using sign restrictions on inventories, consumption, and investment. *Top panel*: impact restrictions; *middle panel*: two-period dynamic restrictions; *bottom panel*: three-period dynamic restrictions. See main text for details on the identification scheme. (a) Impact sign restrictions. (b) Two-period sign restrictions. (c) Three-period sign restrictions.

4.4. Dynamic sign restrictions

Panel (b) of Fig. 5 reports the posterior distribution of the FEV of output attributable to our identified shocks using dynamic sign restrictions. The sign restrictions are the following: inventories fall but consumption and investment increase for two periods after the shock.[22]

The resulting posterior distribution of the FEV has a sharp mode close to zero, and the median is now close to 10 percent at all horizons, about half smaller than the result with impact restrictions only. To get a sense of the contribution of inventories to this result, panel (b) of Fig. 5 also plots the posterior distribution of the FEV, using alternative dynamic restrictions: only consumption and investment are required to increase for two periods. Without the inventory restrictions, the posterior distribution is relatively flat and the median share of FEV attributable to the identified shock is 30 percent overall. Including the inventory sign restrictions causes the median share of the shock to fall by about two-thirds.

Panel (b) of Fig. 6 reports the impulse responses to the shock identified using the dynamic restriction. Inventory disinvestment occurs for two periods, but after that, there is again investment in inventories. Output increases more modestly than consumption and investment increase in the short run, consistent with model predictions.

The dynamic sign restrictions can also be extended to three periods, that is, three periods of inventory disinvestment accompanied by increases in consumption and fixed investment. As reported in panel (c) of Fig. 5, the median share of FEV explained by the identified shock is now below 5 percent at most horizons, and tightly estimated, with basically no probability assigned above 20 percent. Therefore, the news shocks identified with three-period restrictions *at most* account for 20 percent of output variation. The role of the dynamic restrictions on inventories is again central: without it, the two- and three-period schemes deliver almost identical FEV contributions to output. Panel (c) of Fig. 6 reports the impulse responses of the identified shock with three-period restrictions. Although the movement in output is modest, it actually declines on impact.

4.5. Understanding the results

We summarize the key points of these empirical results as follows: (i) the identified impulse response with impact restrictions suggests that most news shocks are short-lived, with a recovery of inventories after the impact disinvestment; (ii) the identified news shocks based on impact restrictions explain on average 20 percent of output variation at all horizons; (iii) dynamic restrictions generate a tighter posterior distribution of the fraction of output variation explained by the shock; (iv) based on dynamic restrictions, long-horizon news shocks explain on average less than 5 percent, and at most 20 percent, of output variation at all horizons.

The reason why the contribution of news, in particular long-horizon news, turns out to be small is that inventories are a procyclical variable. In the data, the unconditional contemporaneous correlation between inventories and sales (consumption plus investment) at business-cycle frequencies is 0.50. Since the identification is based on negative comovement of these generally positively comoving variables, it would come as a surprise if the identified shocks were able to generate the bulk of business cycles.

4.6. Robustness

The main results survive a number of robustness checks, which are summarized in Table 2. First, as mentioned above, they are not sensitive to using different types of inventory data. Panel (b) reports the posterior median value of FEV using retail trade inventories instead of total private inventories. It is very similar to the baseline results of Fig. 5 reported in panel (a). While retail inventories might in general be a less procyclical variable, this evidence suggests that its persistent negative comovement with sales is rarely found in the data.

Second, as studied in detail by McCarthy and Zakrajšek (2007), inventory dynamics have changed since the 1980s: the procyclicality of inventories has weakened and the volatility of total inventory investment has fallen, possibly because of improvements in inventory management. In order to take into account the possibility of different "inventory regimes" in the data, panel (c) reports results for the post-1984 sample. The numbers are similar to the baseline result. This suggests that the nature of the comovement between aggregate inventories and private sales did not change substantially around this period.

Third, we loosen the two- and three- period dynamic sign restrictions, by allowing inventories to move freely on impact (two-period restriction) or to move freely both on impact and in the second period (three-period restriction). Since the identification is weaker, one should expect a higher FEV. However, as reported in panel (d), the implied FEV is not sensitive to this change. This is because inventory movements tend to be persistent: an inventory decline is typically the result of several periods of inventory disinvestment. Regardless of its short-term nature, a longer-term negative comovement between inventories and sales is hard to detect in the data. Finally, the results are also not sensitive to HP-filtering the data

[22] Unlike the results of Proposition 1, it cannot be established theoretically that, in the model of Section 2, news shocks satisfy this dynamic restriction. However, this can be tested by Monte-Carlo simulation. When applied to simulated data, the two period identification scheme correctly recovers the contribution of news shocks to output volatility, as well as the impulse-response functions, under different assumptions about the true contribution of news shocks to aggregate volatility. This is also true of the three-period sign restrictions. These results are reported in the online appendix.

Table 2
Posterior median share of FEV of output attributable to news (percentage), for alternative identification schemes.

Forecast horizon:	1	5	10	20
(a) Baseline				
Impact sign restrictions	9.82	20.84	22.61	22.87
Two-period sign restrictions	4.89	10.18	13.20	14.95
Three-period sign restrictions	5.72	2.74	3.22	4.39
(b) Retail inventories				
Impact sign restrictions	14.09	25.93	26.75	25.36
Two-period sign restrictions	12.17	5.00	5.66	8.31
Three-period sign restrictions	16.85	4.10	3.07	3.61
(c) Data after 1984				
Impact sign restrictions	12.11	19.37	21.24	23.29
Two-period sign restrictions	5.93	5.76	8.66	12.84
Three-period sign restrictions	8.71	3.11	4.29	8.97
(d) Loose identification				
Sign restrictions in second period only	5.64	11.43	14.42	16.05
Sign restrictions in third period only	5.03	2.51	2.93	4.09

or to using an alternative measure of output consistent with our model, which subtracts government spending and net exports from GDP.[23]

4.7. Relation with earlier VAR approaches

A number of empirical papers on news shocks have used structural VARs to quantify their contribution to business-cycle fluctuations.[24] This literature has emphasized the importance of using identifiying assumptions that are robust to model mis-specification. The approach proposed in this paper contributes to this effort by using a sign restriction which is shown to be robust across structural inventory models.

The approach of this paper nevertheless differs from earlier work in two fundamental ways. First, it does not rely on direct measures of TFP, or in general of the fundamentals affected by the news. Second, it can be used to provide an upper bound to the contribution of news to business-cycle fluctuations.[25] For these reasons, the identification scheme proposed in this paper provides a different perspective on the contribution of news from the existing literature, and can be used as a further robustness check on existing identification schemes.

To illustrate this, we constructed the inventory dynamics implied by news shocks identified using the short-run scheme described in Beaudry and Portier (2006), by estimating a three-variable VAR with utilization-adjusted TFP, the S &P 500 index and inventories. As documented by Beaudry and Portier (2006), the identified news shock generates a long-run increase in TFP in this three-variable system. However, the sign of the impact response of inventories to the shock, which is reported in Fig. 7, is inconclusive: it is weakly negative, but not significantly different from zero. Thus, while the identified shock does generate long-run increases in TFP, it also fails to generate a significant short-run decline in inventories, as predicted by structural models. This evidence indicates that, in conjunction with its news component, the shock identified using the short-run Beaudry and Portier (2006) scheme likely also has a non-news component, which drives inventories up in the short run. Under this interpretation, using the information contained in inventory movements could help further pin down the exact contribution of news to business-cycle fluctuations.[26]

5. Conclusion

This paper studied the response of inventories to news shocks. It established that structural models predict that inventories should fall in response to good news about the future. The key mechanism behind this theoretical result is that news shocks generate strong intertemporal substitution in production. Furthermore, this result can be used to identify news shocks in postwar US data, using sign restrictions in a VAR framework. The central empirical result is that news shocks play

[23] These last two robustness checks are reported in the online appendix.

[24] Some of the most prominent examples include Beaudry and Portier (2006), Beaudry and Lucke (2010), and Barsky and Sims (2011).

[25] Section 3.4.1 shows that surprise innovations to the two fundamental processes in the baseline economy do not generate negative comovement between inventories and sales. However, the possibility that surprise innovations to fundamentals not captured by the model may also generate this comovement pattern cannot be ruled out a priori. In that case, the identification scheme would capture the total contribution of these shocks, along with news shocks.

[26] Related points are made in Arias et al. (2013) with regards to the penalty-function approach used by Beaudry et al. (2011), and in Forni et al. (2014) on the non-fundamentalness of small-scale VAR models in identifying news shocks.

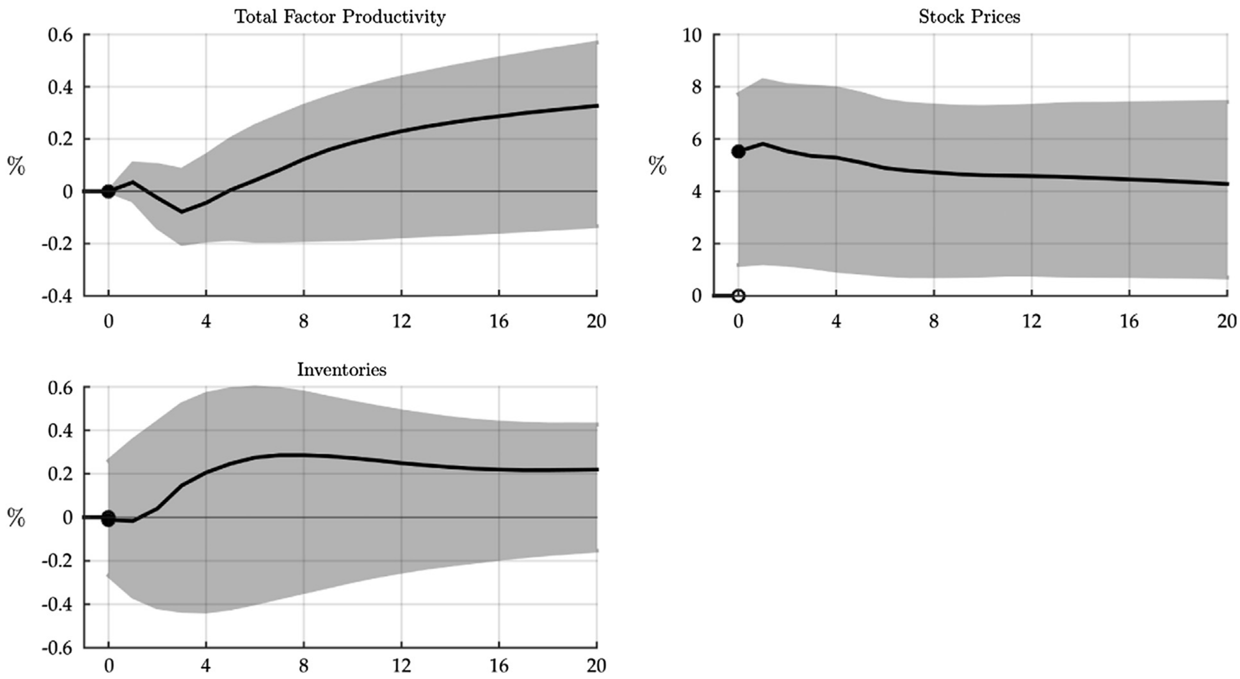

Fig. 7. Impulse responses to a news shock identified using short-run restrictions on TFP and stock prices. Note: The identified shock is required to have zero impact effect on TFP and a positive impact effect on stock prices, as in Beaudry and Portier (2006).

a small role in aggregate fluctuations, contributing to less than 20 percent of output variation. The intuition for this empirical result is that inventories tend to comove positively with sales at business-cycle frequencies.

The results suggest two future directions for progress. First, the analysis of Section 3 highlighted the role of the EISP for inventory dynamics. While the approach of this paper consisted in deriving this elasticity from existing microfounded models, a complementary approach would be to find methods to estimate it directly. Second, one could ask whether the results of this paper also hold at an industry level.[27] A news shock in one particular sector may lead to negative comovement of inventories and sales in that sector, but positive comovement in the aggregate. We leave this to future research.

Acknowledgments

We thank Ricardo Reis for his support throughout this project, and Stephanie Schmitt-Grohé and Jon Steinssón for their advice. We thank Jonas Arias for sharing his code, and Dario Caldara, Yongsung Chang, James Kahn, Jennifer La'O, Neil Mehrotra, Virgiliu Midrigan, Tommaso Monacelli, Miguel Morin, Emi Nakamura, Serena Ng, Dimitris Papanikolaou, John Roberts, Frank Schorfheide, Dmitriy Sergeyev, Jae W. Sim, Martin Uribe, Michael Woodford, and numerous seminar participants at Columbia University and the Federal Reserve Board for helpful comments at the early stage of this project. This paper supersedes the previous version titled, "Can News Shocks Account for the Business-Cycle Dynamics of Inventories?"

Appendix A. Calibration of the EISP in Section 2.1

The EISP in the stock-elastic demand model: In the stock-elastic demand (SE) model of Bils and Kahn (2000), the log-linearized first-order condition for inventory choice is

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

where γ_t is the growth rate of marginal cost. The EISP in this model and its lower bound are given by:

$$\eta^{SE} = \frac{1+IS}{IS}\frac{1}{1-\beta(1-\delta_i)} \geq \frac{1}{1-\beta(1-\delta_i)} \equiv \underline{\eta}^{SE},$$

since $IS \geq 0$. In the baseline calibration, $\delta_i = 0.025$ and $\beta = 0.99$, so that $\underline{\eta}^{SE} = 28.8$.[28]

[27] For example, Chang et al. (2009) study the employment effect of TFP shocks at an industry level by using inventory data.

[28] The online appendix provides a detailed exposition of an RBC version of the model and a derivation of the EISP.

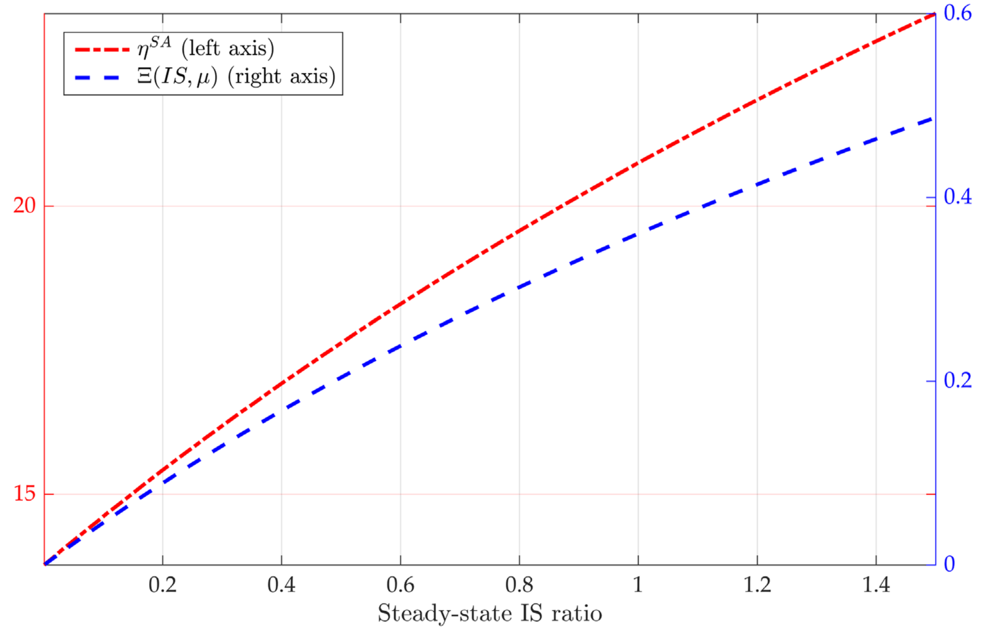

Fig. 8. Mapping between the EISP η^{SA} and the steady-state IS ratio in the stockout-avoidance model. Note: The EISP is plotted for $\mu = 1{:}25$. See the online appendix for details on the derivation of η^{SA} and $\Xi(\mu, IS)$.

The EISP in the stockout-avoidance model: The first-order condition for optimal inventory choice in the stockout-avoidance (SA) model of Kryvtsov and Midrigan (2013) is

$$\widehat{inv}_t = \hat{s}_t + \eta^{SA}\hat{\gamma}_t + \tau^{SA}\hat{\mu}_t.$$

In the SA model, the combination of firm monopoly power and stockout-avoidance behavior generates an endogenously time-varying markup $\hat{\mu}_t$, even in the absence of nominal rigidities. This time variation in markups can affect inventory dynamics, as captured by the term $\hat{\mu}_t$ in the equation above. The EISP in this model is given by

$$\eta^{SA} = \Xi(IS,\mu)\frac{1+IS}{IS}\frac{1}{1-\beta(1-\delta_i)},$$

where $0 \le \Xi(IS,\mu) < 1$ is a reduced-form parameter that depends on the steady-state IS ratio, IS, and the steady-state markup, μ.[29]

Since the lower bound on η^{SA} has no closed-form solution, Fig. 8 plots the relationship between η^{SA} and μ as a function of the IS ratio. The value of μ is set to 1.25, but the magnitude of the resulting lower bound does not depend on the choice of μ.[30] The lower bound on the EISP in this model is $\eta^{SA} \ge \underline{\eta}^{SA} = 13.8$.

Choice of EISP in the baseline model: Given the two lower bounds $\underline{\eta}^{SE} = 28.8$ and $\underline{\eta}^{SA} = 13.8$, we proceed conservatively and pick the steady-state elasticity of $\chi'(\cdot)$ so that $\eta = 10$, below the range implied by the two models.

Appendix B. Proof of Proposition 1

On impact and without surprise shocks, the following log-linearized equations describe the equilibrium of the model:

$$\kappa\hat{y}_t = (\kappa-\delta_i IS)\hat{s}_t + IS\widehat{inv}_t - IS(1-\beta(1-\delta_i))(\widehat{inv}_t - \hat{s}_t), \tag{B.1}$$

$$\omega\hat{y}_t = (\omega+1)(1-\alpha)\hat{u}_t + (\widehat{mc}_t - \hat{\lambda}_t), \tag{B.2}$$

$$\widehat{mc}_t - \hat{\lambda}_t = \frac{f(\kappa, IS, \beta, \delta_i)}{\eta}\left(\widehat{inv}_t - \hat{s}_t\right). \tag{B.3}$$

[29] The online appendix provides a precise derivation of the mapping $\Xi(IS,\mu)$. Its form depends on the choice of distribution for the idiosyncratic demand shock. In general, for any two-parameter distribution, the mapping between the first and second moments of that distribution and the steady-state markup and IS ratio is unique. Derivations are provided in two cases: log-normal demand and Pareto demand shocks.

[30] The online appendix discusses comparative statics of the EISP when the steady-state value of the markup changes, and shows that, ceteris paribus, a lower steady-state markup implies a higher EISP. For very high steady-state markups, e.g., $\mu = 2$ (a 100% markup in steady-state), the lower bound for the EISP is $\eta^{SA} = 7.0$.

The reduced-form parameters in this system are given by

$$\omega=\frac{1+(1-\alpha)\xi}{\alpha\xi}, \quad \kappa=1+\chi(IS)+\delta_i IS, \tag{B.4}$$

$$f(\kappa, IS, \beta, \delta_i)=\frac{IS\beta(1-\delta_i)}{\kappa-\delta_i IS+IS(1-\beta(1-\delta_i))}. \tag{B.5}$$

Combining (B.1)–(B.3), and setting $\hat{u}_t=0$, one can derive the following relationship between $\widehat{inv}_t$ and $\hat{s}_t$:

$$\left[\frac{f(\kappa, IS, \beta, \delta_i)}{\omega\eta}-\frac{IS\beta(1-\delta_i)}{\kappa}\right]\widehat{inv}_t=\left[\frac{f(\kappa, IS, \beta, \delta_i)}{\omega\eta}+\frac{\kappa-\delta_i IS+IS(1-\beta(1-\delta_i))}{\kappa}\right]\hat{s}_t.$$

Since the coefficient on the right-hand side is always positive, positive comovement between the two variables $\widehat{inv}_t$ and $\hat{s}_t$ occurs if and only if the coefficient on the left-hand side is positive. This in turn immediately gives the result of proposition 1.

Appendix C. Supplementary data

Supplementary data associated with this article can be found in the online version at http://dx.doi.org/10.1016/j.jmoneco.2016.03.005.

References

Arias, J.E., Rubio-Ramirez, J.F., Waggoner, D.F., 2013. Inference Based on SVARs Identified With Sign and Zero Restrictions: Theory and Applications. Duke University, Mimeo.
Barsky, R.B., Sims, E.R., 2011. News shocks and business cycles. J. Monet. Econ. 58 (3), 273–289.
Beaudry, P., Lucke, B., 2010. Letting different views about business cycles compete. NBER Macroecon. Ann., 24.
Beaudry, P., Nam, D., Wang, J., 2011. Do Mood Swings Drive Business Cycles and is it Rational? NBER Working Paper 17651.
Beaudry, P., Portier, F., 2006. Stock prices, news, and economic fluctuations. Am. Econ. Rev. 96 (4), 1293–1307.
Beaudry, P., Portier, F., 2013. News driven business cycles: insights and challenges. J. Econ. Lit. 52 (4), 993–1074.
Bils, M., Kahn, J.A., 2000. What inventory behavior tells us about business cycles. Am. Econ. Rev. 90 (3), 458–481.
Caldara, D., Kamps, C., 2012. The Analytics of SVAR: A Unified Framework to Measure Fiscal Multipliers. Federal Reserve Board Working Paper Series, No. 2012-20.
Chang, Y., Hornstein, A., Sarte, P.-D., 2009. On the employment effects of productivity shocks: the role of inventories, demand elasticity, and sticky prices. J. Monet. Econ. 56 (3), 328–343.
Christiano, L.J., Ilut, C., Motto, R., Rostagno, M., 2008. Monetary Policy and Stock Market Boom-Bust Cycles. European Central Bank Working Paper Series 955.
Dotsey, M., King, R.G., 2006. Pricing, production, and persistence. J. Eur. Econ. Assoc. 4 (5), 893–928.
Faust, J., 1998. The Robustness of Identified VAR Conclusions About Money. Carnegie–Rochester Conference Series on Public Policy, 49, pp. 207–244.
Fisher, J.D.M., Hornstein, A., 2000. (S,s) Inventory Policies in General Equilibrium. Rev. Econ. Stud. 67, 117–145.
Forni, M., Gambetti, L., Sala, L., 2014. No news in business cycles. The Economic Journal 124 (581), 1168–1191.
Greenwood, J., Hercowitz, Z., Huffman, G.W., 1988. Investment, capacity utilization, and the real business cycle. Am. Econ. Rev. 78 (3), 402–417.
Jaimovich, N., Rebelo, S., 2009. Can news about the future drive the business cycle? Am. Econ. Rev. 99 (4), 1097–1118.
Jung, Y., Yun, T., 2006. Monetary Policy Shocks, Inventory Dynamics, and Price-Setting Behavior. Federal Reserve Bank of San Francisco Working Paper Series, No. 2006-02.
Kesavan, S., Gaur, V., Raman, A., 2010. Do inventory and gross margin data improve sales forecasts for U.S. public retailers? Manag. Sci. 56 (9), 1519–1533.
Khan, A., Thomas, J.K., 2007. Inventories and the business cycle: an equilibrium analysis of (S,s) policies. Am. Econ. Rev. 97 (4), 1165–1188.
Khan, H., Tsoukalas, J., 2012. The quantitative importance of news shocks in estimated DSGE models. J. Money Credit Bank. 44 (8), 1535–1561.
Kryvtsov, O., Midrigan, V., 2013. Inventories, markups, and real rigidities in menu cost models. Rev. Econ. Stud. 80 (1), 249–276.
McCarthy, J., Zakrajšek, E., 2007. Inventory dynamics and business cycles: what has changed?. J. Money Credit Bank. 39 (2–3), 591–613.
Moon, H., Schorfheide, F., Granziera, E., 2013. Inference for VARs Identified with Sign Restrictions. University of Pennsylvania, Mimeo.
Mountford, A., Uhlig, H., 2008. What are the effects of fiscal policy shocks? J. Appl. Econ. 24 (6), 960–992.
Schmitt-Grohé, S., Uribe, M., 2012. Whats news in business cycles? Econometrica 80 (6), 2733–2764.
Uhlig, H., 2005. What are the effects of monetary policy on output? results from an agnostic identification procedure. J. Monet. Econ. 52 (2), 381–419.
Vukotic, M., 2013. A Sectoral Approach to News Shocks. University of Warwick, Mimeo.
Wen, Y., 2011. Input and output inventory dynamics. Am. Econ. J.: Macroecon. 3 (4), 181–212.