Macroeconomic Dynamics, **22**, 2018, 620–639. Printed in the United States of America.
doi:10.1017/S1365100516000225

OIL PRICE SHOCKS, INVENTORIES, AND MACROECONOMIC DYNAMICS

Ana María Herrera
University of Kentucky

This paper investigates the time delay in the transmission of oil price shocks using disaggregated manufacturing data on inventories and sales. VAR estimates indicate that industry-level inventories and sales respond faster to an oil price shock than aggregate gross domestic product, especially in industries that are energy-intensive. In response to an unexpected oil price increase, sales drop and inventories are accumulated. This leads to future reductions in production. We estimate a modified linear–quadratic inventory model to inquire whether the patterns observed in the VAR impulse responses are consistent with rational behavior by the firms. Estimation results suggest that three mechanisms play a role in the industry-level dynamics. First, oil prices act as a negative demand shock. Second, the shock catches manufacturers by surprise, resulting in higher-than-anticipated inventories. Third, because of their desire to smooth production, manufacturers deviate from the target level of inventories and spread the decline in production over various quarters; hence the delay in the response of aggregate output.

Keywords: Oil Shocks, Macroeconomic Fluctuations, Inventories

1. INTRODUCTION

A puzzling aspect of the historical correlation between oil prices and the macroeconomy is the substantial time delay in the transmission of an oil price innovation [see Hamilton and Herrera (2004), Kilian and Lewis (2011), among others]. In contrast with the rather fast propagation of monetary policy or technology shocks, a slowdown in real gross domestic product (GDP) growth typically has not shown up until four quarters after an unexpected oil price increase.[1] This paper uses disaggregate manufacturing data to investigate this puzzle empirically.

I begin my analysis by estimating a vector autoregression on the real oil price change, sales growth, and the inventory–sales ratio for 21 manufacturing industries (19 two-digit and 2 three-digit SIC industries) plus three aggregates (total

This research was supported by the NSF under Grant SES-003840 and was partially completed while visiting Harvard's Kennedy School of Government under a Repsol-YPF research fellowship. I am thankful to Jim Hamilton, Bill Hogan, Lutz Kilian, Valerie Ramey, and three anonymous referees, as well as participants at numerous conferences and seminars, for helpful comments and suggestions. Address correspondence to: Ana María Herrera, Department of Economics, Gatton College of Business and Economics, Lexington, KY 40506-0034, USA; e-mail: amherrera@uky.edu.

 1365-1005/18

manufacturing, nondurables, and durables). I find that sales in energy-intensive industries (e.g., transportation equipment, petroleum products) respond to an unexpected oil price increase in less than a year, faster than the response of GDP to the same innovation. In addition, the initial effect on industry-level output is less pronounced because of an increase in finished goods inventories. These inventories are gradually worked down by a continuing period of curtailed production. This pattern is suggestive of the classic inventory-accelerator model of the business cycle.

To inquire whether this account of the oil price dynamics is consistent with a model of firm behavior, I estimate and test a linear–quadratic model of inventory accumulation. This model was originally developed by Holt et al. (1960) and has been extensively used in empirical analysis of inventory behavior.[2] Although this literature is impressively broad, it has not been very successful in producing economically plausible parameter values. In particular, parameter estimates are seldom statistically significant, sometimes have the wrong sign, and are often unsupportive of the underlying model [Fuhrer et al. (1995)].

In this paper I estimate a modified version of the linear–quadratic inventory model, in which I introduce two generalizations. First, we model the shock to the marginal cost of production as an $I(1)$ variable cointegrated with sales, as suggested by Hamilton (2002).[3] This departure from the common assumption that the cost shock is stationary has the benefit of accounting for stochastic trends in sales and inventories, while ensuring that both marginal costs of production and inventory carrying costs are stationary along the long-run equilibrium path. Second, I allow a more general specification of the cost and demand shocks faced by the firm than commonly assumed in the empirical literature.[4]

Estimates of this modified linear–quadratic inventory model are shown to produce industry-level impulse responses that resemble those implied by the VAR model. Moreover, the dynamics entailed by our estimates is consistent with two stylized facts about inventory behavior: procyclicality and persistence [Ramey and West (1999)]. In the wake of an oil price shock, economic activity contracts and inventories are drawn down. The rise in the inventory–sales ratio, resulting from a smaller decline in inventories than in sales is slowly worked down as adjustment to the steady state takes place.

This paper is organized as follows. Section 2 discusses the data as well as some measurement choices. Section 3 uses a VAR framework to study the dynamics of oil price innovations at the industry level. In Section 4 I inquire whether the uncovered dynamics is consistent with rational behavior of the firms by estimating and testing a modified version of the linear–quadratic inventory model. Section 5 concludes.

2. DATA AND MEASUREMENT

To investigate the nature of the time lag in the propagation of oil price innovations, I use data on manufacturing sales and finished goods inventories (hereafter inventories) from the Bureau of Economic Analysis (BEA). The series span the period between January 1959 and March 2000, are measured in chained dollars of 1996, and comprise three manufacturing aggregates (total manufacturing, durables, and nondurables), nineteen 2-digit SIC industries, and two 3-digit SIC sectors (motor vehicles and other transportation equipment).[5]

Although the data are available at a monthly frequency from the BEA, I choose to transform monthly data into quarterly series by aggregating monthly sales and using end-of-the-quarter inventories. Although this time aggregation constitutes a deviation from the inventory literature and a loss of higher-frequency information, it significantly diminishes the computational burden involved in the estimation of our inventory model and it facilitates comparison with the oil price shocks–macroeconomy literature.

The first data choice to be made here is how to characterize the data–generating process of inventories and sales. The leading approach in the inventory literature has been to model inventories and sales as stationary around a deterministic trend. However, results from a DF-GLS test reported in Table A.1 of the Online Appendix[6] indicate that we cannot reject the null hypothesis of a unit root at a 5% significance level for any sectors except tobacco inventories and sales. Furthermore, residual-based cointegration tests suggest that inventories and sales are cointegrated for more than half of the industries. Therefore, in my analysis I consider an industry where sales, inventories, and production have a stochastic trend[7] and the first two series are cointegrated.

A second choice is the measure of oil prices. We follow Mork (1989) and Lee and Ni (2002) in measuring oil prices by the refiners' acquisition cost (RAC) instead of the PPI, when possible, and make adjustments to account for the price controls of the 1970s. I deflate the RAC by the consumer price index (CPI) and then compute the rate of growth by taking the first difference in the logarithm of the real oil price.

3. DYNAMICS OF OIL PRICE INNOVATIONS AT THE INDUSTRY LEVEL

Assume the data-generating process for a particular industry to be given by a three-dimensional VAR(4), where $\mathbf{x}_t$ contains the log growth of the real oil price, the log growth of quarterly real sales, and the log difference between inventories and sales. The VAR is assumed to have a linear moving-average representation given by

$$\mathbf{x}_t = A(L)\,\mathbf{u}_t, \quad A(0) = A_0, \tag{1}$$

where $\mathbf{w}_t = [u_{o,t},\ u_{s,t},\ u_{h,t}]'$ is a vector of white noise structural innovations. The process in (1) is consistent with evidence of cointegration (see Table A.1 in the Online Appendix) and can be directly mapped into the usual error-correction model.[8] For identification purposes we assume A_0 is a lower triangular matrix. The ordering of the real oil price change before the manufacturing variables imposes the reasonable restriction that oil prices do not respond contemporaneously to changes in industry-level sales or inventories [see Kilian and Vega (2011)].

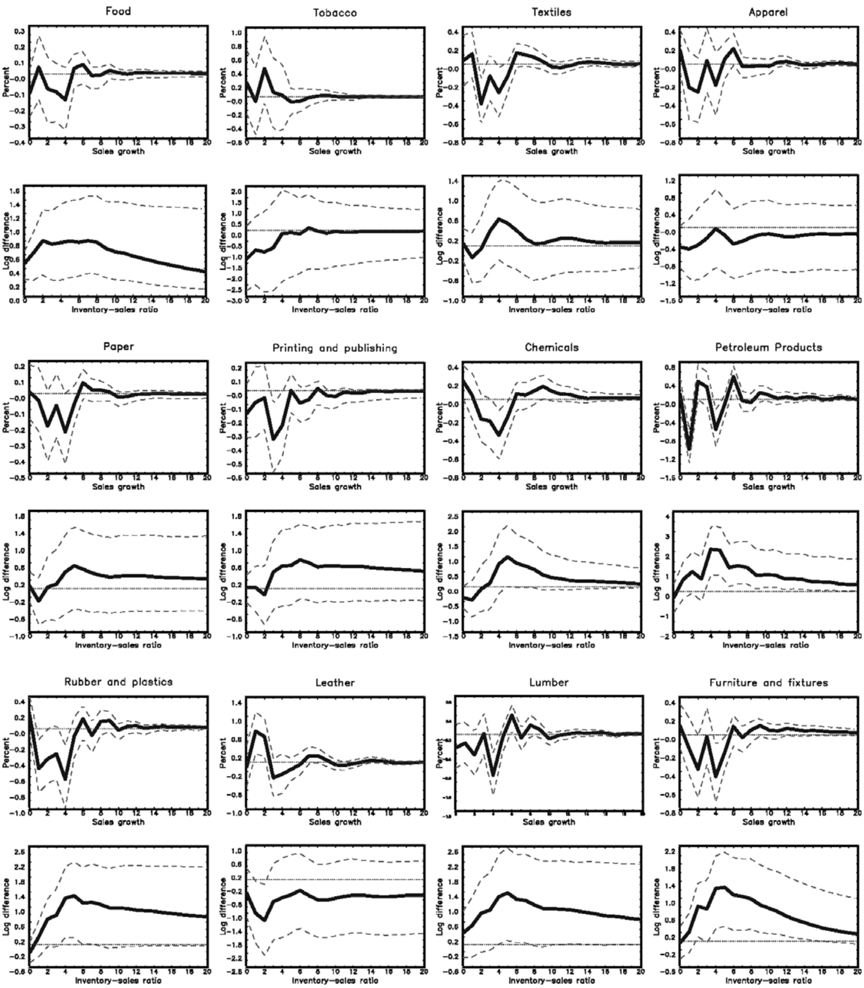

FIGURE 1. Responses to a 10% increase in the real oil price. Estimates based on the reduced-from VAR(4) system described in Section 3. The 90% confidence intervals were computed using Killan's (1998) bootstrap-after bootstrap method.

Note that, given the responses for the level of sales and inventories, we can infer the production response using the inventory identity

$$Q_t = S_t + H_t - H_{t-1}, \qquad (2)$$

where Q_t denotes output, S_t denotes sales, and H_t denotes inventories.[9] Figure 1 illustrates the impulse responses to an unexpected 10% increase in the real oil price. The 90% confidence intervals are computed using Kilian's (1998)

https://doi.org/10.1017/S1365100516000225 Published online by Cambridge University Press

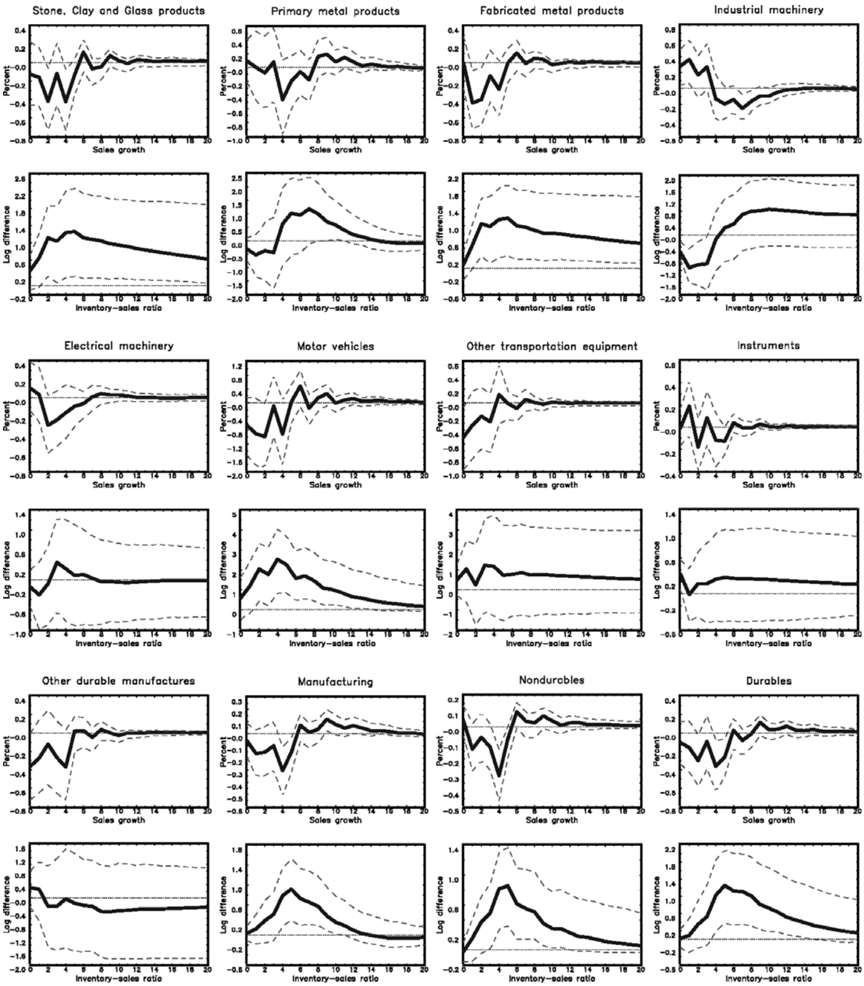

FIGURE 1. Continued.

bootstrap-after-bootstrap method. For the sake of brevity I relegate the cumulative impulse response to the Appendix (see Figures A.1a–A.1c).

Four important features of oil price innovations dynamics are apparent:

- *Industry-level sales decline in response to an oil price increase.* In particular, for industries that are energy-intensive in production (e.g., chemicals, rubber and plastics, petroleum products) or consumption (e.g., motor vehicles, other transportation equipment), a decline in sales occurs during the first year. Significant reductions follow in the remaining sectors and aggregate manufacturing.

https://doi.org/10.1017/S1365100516000225 Published online by Cambridge University Press

- *Industry-level output declines in response to an oil price increase*. Declines in production are observed within a year for chemicals, petroleum products, rubber and plastics, lumber, furniture and fixtures, stone, clay and glass, fabricated metal products, motor vehicles, and the three manufacturing aggregates. A decline in the remaining industries is not evident until a year later. The timing of the contraction for total manufacturing is consistent with the time delay in the response of aggregate GDP to oil price innovations.
- *Inventories usually decline at a slower pace than sales, leading to a hump-shaped response of the inventory–sales ratio*. Significant deviations from the benchmark inventory–sales ratio are observed for chemicals, petroleum products, rubber and plastics, lumber, furniture and fixtures, stone, clay, and glass, fabricated metal products, and motor vehicles. These are industries either that use petroleum intensively as an input or for which the automobile industry constitutes an important demand source.[10] A similar pattern is observed for the three manufacturing aggregates.
- *The contractionary effect is largest for motor vehicles but is also significant for industries that are energy-intensive or for which motor vehicles constitute an important demand factor*. The long-run elasticity of sales to oil prices is about twice as large for motor vehicles (−0.36) as for furniture and fixtures (−0.17), the sector with the second largest effect. Moreover, a year later when the economic slowdown spreads to aggregate manufacturing (and real GDP), the 10% increase in oil prices has resulted in a 4.2% decline in motor vehicles production and contractions of 0.7, 0.6, 1.0, 2.0, and 1.5% in apparel, chemicals, petroleum products, rubber and plastics, and stone, clay, and glass products, respectively.

4. CAN THE DYNAMICS OF OIL PRICE INNOVATIONS BE RIGOROUSLY RECONCILED WITH RATIONAL BEHAVIOR BY FIRMS?

These patterns are suggestive of the classic inventory-accelerator model of the business cycle. An increase in oil prices leads consumers to abstain from new purchases. Partly because the shock catches manufacturers by surprise, and partly out of a desire to smooth output fluctuations, manufacturers deviate from their target level of inventories and spread the decline in production over several quarters. By the fourth quarter, curtailed production in energy-intensive sectors has resulted in lower sales and income for other industries, thus leading the economy into a recession.

Although this account of the dynamics of an oil price innovation seems intuitively plausible, can it be rigorously reconciled with profit-maximizing behavior by firms and apparent production-cost schedules? To answer this question I estimate and test a linear–quadratic inventory model. My model relies on the traditional quadratic approximation to the costs faced by the firm, but I introduce two important changes.

First, I modify the setup to account for the presence of stochastic trends and comovement in inventories, sales, and the stochastic cost shock. The motivation for this modification is twofold: (a) statistical tests indicate that inventories and sales have a unit root and are cointegrated (see Table A.1 of the Online Appendix); (b) when sales have a unit root and the cost shock is stationary, the marginal production cost tends to infinity, so the firm minimizes costs by letting inventory management cost go to infinity [Hamilton (2002)]. This problem can be avoided by assuming that both the cost shock and sales have a unit root and are cointegrated. This assumption is motivated on the grounds that cost-saving technological progress generates an upward trend in sales.

Second, I use a less restrictive specification of the demand disturbances than is common in applications of the linear–quadratic inventory model. In particular, I assume that real oil prices have a direct effect on sales growth. The rationale for this modification is twofold. Although, in the linear–quadratic literature, energy prices are commonly modeled as an observable cost shifter, previous studies have rarely found energy prices to be statistically significant [Ramey and West (1999)]. In addition, VAR estimation results uncovered a statistically significant effect of oil price innovations on sales.

4.1. A Model of Inventory Behavior

Consider the following decision problem, similar to that Hamilton (2002):[11]

$$\max_{\{Q_t,H_t\}_{t=0}^{\infty}} E_0\left\{\sum_{t=0}^{\infty}\beta^t(P_tS_t - C_t)\right\} \tag{3}$$

subject to

$$C_t = (1/2)[a_0(\Delta Q_t)^2 + a_1(Q_t - U_{c,t})^2 + a_2(H_{t-1} - a_3S_t)^2], \tag{4}$$

$$Q_t = S_t + H_t - H_{t-1}, \tag{5}$$

where P_t is the price of the good in period t, S_t is real sales during period t, C_t is the cost of production, Q_t is the quantity produced during period t, H_t are inventories of finished goods at the end of period t, β is the discount rate, and $U_{c,t}$ is a stochastic exogenous shock to the marginal cost of production.

The first-order condition for cost minimization is derived by differentiating the objective function (3) with respect to H_t:

$$\begin{aligned} &E_t[a_0(\Delta Q_t - 2\beta\Delta Q_{t+1} + \beta^2\Delta Q_{t+2}) + a_1(Q_t - U_{c,t}) \\ &\quad -\beta a_1(Q_{t+1} - U_{c,t+1}) + \beta a_2(H_t - a_3S_{t+1})] = 0. \end{aligned} \tag{6}$$

Consider the case where inventories and sales have a unit root and are cointegrated with cointegrating vector $(1, -a_3)$. Further, assume that the unobserved shock to the marginal cost of production, $U_{c,t}$, has a unit root and is cointegrated with sales, so that

$$U_{c,t} - S_t - k_c = v_{c,t} \sim I(0), \tag{7}$$

where

$$v_{c,t} = \theta_{c1} v_{c,t-1} + \theta_{c2} v_{c,t-2} + \varepsilon_{c,t}, \tag{8}$$

k_c is a constant term, and the innovation $\varepsilon_{c,t}$ has a zero-mean normal distribution with variance σ_c^2. As I mentioned before, cointegration between $U_{c,t}$ and S_t can be motivated by the presence of an unobserved technology shock (an upward trend in $U_{c,t}$ or a downward trend in $U_{c,t}^*$) that generates an upward trend in sales and, given the inventory accumulation equation (5), also in production.[12]

I consider the data-generating process for sales of a particular industry to be given by

$$\Delta S_t = k_s + \lambda_{s1} \Delta S_{t-1} + \lambda_{s2} \Delta S_{t-2} + \lambda_{o1} o_{t-1} + \lambda_{o2} o_{t-2} + \lambda_{o3} o_{t-3} + \lambda_{o4} o_{t-4} + e_{s,t}, \tag{9}$$

where $\Delta S_t = v_{s,t} \sim I(0)$, and o_t is simply the change in real oil prices. In turn, the process for o_t is given by

$$o_t = k_o + \omega_{o1} o_{t-1} + \omega_{o2} o_{t-2} + e_{o,t} \tag{10}$$

and

$$\begin{bmatrix} e_{s,t} \\ e_{o,t} \end{bmatrix} = \begin{bmatrix} 1 & \lambda_{o0} \\ 0 & 1 \end{bmatrix} \begin{bmatrix} \varepsilon_{s,t} \\ \varepsilon_{o,t} \end{bmatrix},$$

where the innovations $\varepsilon_{s,t}$, and $\varepsilon_{o,t}$ are uncorrelated normally distributed processes. The constant terms k_s, k_c, and k_o are not separately identified, because they only affect the constant term in the expression for the firm's optimal level of inventories. Hence, without loss of generality, we can solve the firm's optimal inventory problem with all constants set to zero and then add the constants at the final step of the maximum-likelihood estimation.

The optimization problem can be stated as

$$\min_{\{u_t\}_{t=0}^{\infty}} E\left\{ \sum_{t=0}^{\infty} \beta^t \begin{bmatrix} u_t & \mathbf{x}_t' \end{bmatrix} \mathbf{G} \begin{bmatrix} u_t \\ \mathbf{x}_t \end{bmatrix} \mid \mathcal{F}_0 \right\} \tag{11}$$

subject to

$$\mathbf{x}_{t+1} = \mathbf{A}\mathbf{x}_t + \mathbf{B}u_t + \mathbf{C}\mathbf{w}_{t+1}, \tag{12}$$

where $\mathbf{x}_t = (H_{t-1}, H_{t-2}, S_{t-1}, v_{c,t}, v_{c,t-1}, v_{s,t}, v_{s,t-1}, o_t, o_{t-1}, o_{,t-2}, o_{,t-3})'$ denotes the state vector that summarizes the information relevant to the firm's decision, $u_t = H_t$ denotes the control variable, and $\mathcal{F}_0$ denotes the information set at $t = 0$.[13] The solution to this optimization problem takes the form

$$u_t = -\mathbf{F}\mathbf{x}_t, \tag{13}$$

where $\mathbf{F}$ can be computed following Anderson et al. (1996).

Equation (13), with the constant term, k_h, added back in,

$$H_t = k_h - f_1 H_{t-1} - f_2 H_{t-2} - f_3 S_{t-1} - f_4 v_{c,t} - f_5 v_{c,t-1} - f_6 v_{s,t} - f_7 v_{s,t-1} - f_8 o_t - f_9 o_{t-1} - f_{10} o_{t-2} - f_{11} o_{t-3}, \quad \textbf{(14)}$$

together with equations (8), (9), and (10), constitutes an observable state-space model in which (8) is the state equation and (9), (10), and (14) are the observation equations. After setting the discount factor $\beta = 0.98$ and normalizing the coefficient $a_1 = 1$,[14] we obtain estimates of the structural parameters via maximum likelihood. Then these estimates and the Kalman filter are used to trace the response of sales, inventories, and output to an innovation in the real oil price, $\varepsilon_{o,t}$.

4.2. Inventories, Oil Price Shocks, and Industry Dynamics

The model of optimal inventory behavior just described is most appropriate for the six industries identified as "production-to-stock" (food, tobacco, apparel, chemicals, petroleum products, and rubber and plastics). Nevertheless, to the extent that the so-called "production-to-order" industries hold substantial inventories of finished goods, the desire to smooth production might explain movements in inventories. In this section we focus our discussion on the six production-to-stock industries, the motor vehicles sector, and the three manufacturing aggregates.

Inventories and production costs: magnitude and interpretation of the cost parameters. The usual linear–quadratic inventory model embodies two different motives for holding inventories. The cost of adjusting production, $a_0 \Delta Q_t$, and the cost of producing, $a_1 Q_t$, represent a production–smoothing motive. That is, a firm may hold inventories because they facilitate the intertemporal allocation of production. A second motive for holding inventories is reflected in the term $a_2 (H_{t-1} - a_3 S_t)$, which is the accelerator term. This term reflects the trade-off between the physical cost of holding inventories and the cost of avoiding stock-outs. Yet an important implication of assumption (7) is that now the quadratic cost is directly associated with inventory investment. Hence, with the exception that here $a_0 \neq 0$, the model is closer to the flexible accelerator model than the usual linear–quadratic setup. As a result, higher values of a_1 imply greater output flexibility.

Table 1 reports maximum-likelihood estimates and associated asymptotic standard errors under the heading "Structural model." The magnitudes of a_0 and a_2 relative to a_1—which we normalize to 1—suggest that output should track sales closely in response to a demand shock. Note how a_0 and a_2 are estimated to be positive but less than 1 for all sectors. Interestingly, the degree of precision of these cost estimates seems to be higher for sectors where the oil price shock enters significantly into the sales process (e.g., chemicals, petroleum products, motor vehicles, manufacturing, and durable manufactures). In addition, whereas the

TABLE 1. Parameter estimates for inventory models

	Food				Tobacco				Apparel				Chemicals				Petroleum products			
	Structural model		Behavioral model		Structural model		Behavioral model		Structural model		Behavioral model		Structural model		Behavioral model		Structural model		Behavioral model	
Parameter	Coefficient	S.E.	Coefficient	S.E.	Coefficient	S.E.	Coefficient	S.E.	Coefficient	S.E.	Coefficient	S.E.	Coefficient	S.E.	Coefficient	S.E.	Coefficient	S.E.	Coefficient	S.E.
a_0	0.000	0.036	0.481	0.831	0.037	0.024	1.102	1.227	0.225*	0.132	0.000	0.140	0.189***	0.067	0.815	0.767	0.032**	0.015	0.049**	0.022
a_1				26.852	16.986	0.907	1.010													
a_2	0.017	0.015	0.015	0.013	0.555	0.602	0.025	0.021	0.220	0.152	0.200	0.155	0.003	0.004	0.000	0.000	0.034	0.024	0.025	0.018
a_3	0.167***	0.016	0.162***	0.017	−0.090	0.068	−0.046	0.036	0.437***	0.018	0.429***	0.021	0.366***	0.057	0.321	13.298	0.084***	0.019	0.077***	0.020
λ_{s1}	−0.074	0.087	−0.085	0.077	−0.280***	0.079	−0.281***	0.079	0.046	0.078	0.009	0.087	0.314***	0.075	0.352***	0.080	−0.360***	0.076	−0.356***	0.077
λ_{s2}	−0.221***	0.079	−0.230***	0.077	0.096	0.078	0.092	0.078	−0.205***	0.068	−0.166**	0.080	−0.022	0.068	−0.090	0.083	−0.227***	0.079	−0.238***	0.077
λ_{o1}	0.655	0.875	0.515	0.870	−0.216	0.218	−0.214	0.219	−0.110	0.219	−0.028	0.316	−1.348	0.730	−1.234	0.951	−4.966***	0.662	−5.029***	0.667
λ_{o2}	−1.461	1.215	−1.462	1.345	0.367	0.320	0.367	0.321	−0.193	0.331	−0.141	0.429	0.694	1.084	0.113*	1.838	4.876***	1.074	4.986***	1.074
λ_{o3}	0.387	1.220	0.310	1.352	−0.135	0.324	−0.147	0.330	−0.150	0.337	−0.071	0.406	−0.632	1.075	−0.611**	1.990	−0.236	1.215	−0.551	1.202
λ_{o4}	0.382	0.868	0.615	0.880	−0.096	0.223	−0.079	0.229	0.346	0.224	0.244	0.283	0.921	0.747	1.375	1.066	−0.087	0.905	0.167	0.832
ω_{o1}	1.063***	0.079	1.065***	0.079	1.065***	0.079	1.065***	0.079	1.064***	0.079	1.065***	0.079	1.065***	0.079	1.065***	0.079	1.064***	0.079	1.065***	0.079
ω_{o2}	−0.099	0.079	−0.100	0.079	−0.099**	0.078	−0.100	0.079	−0.100	0.079	−0.100	0.079	−0.100	0.079	−0.100	0.079	−0.100	0.079	−0.100	0.079
θ_{c1}	0.105	0.092	−0.194	0.255	−0.205**	0.095	−0.545***	0.157	0.187*	0.104	0.396**	0.157	0.027	0.081	−0.236	0.174	−0.121	0.095	−0.133	0.092
θ_{c2}	−0.081	0.088	−0.118	0.119	−0.131	0.095	−0.305**	0.121	0.029	0.097	0.033	0.118	−0.044	0.086	−0.122	0.105	−0.052	0.093	−0.092	0.090
k_h	0.453*	0.246	0.328**	0.165	0.242*	0.134	0.142**	0.065	−0.772***	0.199	−0.891***	0.283	0.133	0.130	0.117***	0.045	0.476***	0.162	0.415**	0.151
k_s	0.608**	0.273	0.597**	0.282	0.089	0.074	0.082	0.073	0.205***	0.073	0.071	0.094	0.759***	0.223	0.762***	0.283	0.756***	0.219	0.776***	0.227
k_o	0.045*	0.026	0.045*	0.026	0.044*	0.026	0.045*	0.026	0.045*	0.026	0.045*	0.026	0.045	0.026	0.045*	0.027	0.045*	0.026	0.045*	0.026
σ_c^2	0.234***	0.071	0.712	1.053	0.020***	0.004	0.135*	0.184	0.115*	0.065	0.062	0.044	0.393***	0.095	1.045	1.013	0.129***	0.024	0.131***	0.024
σ_s^2	2.265***	0.253	2.262***	0.252	0.156***	0.017	0.156***	0.017	0.264***	0.030	0.259***	0.029	2.123***	0.239	2.098***	0.235	1.461***	0.163	1.462***	0.163
σ_o^2	0.021***	0.002	0.021***	0.002	0.021***	0.002	0.021***	0.002	0.021***	0.002	0.021***	0.002	0.021***	0.002	0.021***	0.002	0.021***	0.002	0.021***	0.002
σ_{so}	−0.020	0.018	−0.020	0.018	0.001	0.005	0.002	0.005	0.005	0.006	0.005	0.006	0.015	0.017	0.015	0.022	0.015	0.014	0.016	0.015
λ_{s1}^*			3.886**	1.684			0.927***	0.131			−0.943	3.327			0.663***	0.209			0.001	0.143
λ_{s2}^*			−0.252	0.183			0.070	0.065			−1.271	4.728			0.303**	0.115			−0.450	0.534
LR test (p-value)																				
Unobserved cost	0.028				0.148				0.161				0.840				0.481			
Effect of oil price	0.750				0.484				0.065				0.030				0.000			
Behavioral	0.000				0.546				1.000				0.022				1.000			

TABLE 1. Continued

	Rubber and plastics				Motor vehicles				Manufacturing				Nondurables				Durables			
	Structural model		Behavioral model		Structural model		Behavioral model		Structural model		Behavioral model		Structural model		Behavioral model		Structural model		Behavioral model	
Parameter	Coefficient	S.E.	Coefficient	S.E.	Coefficient	S.E.	Coefficient	S.E.	Coefficient	S.E.	Coefficient	S.E.	Coefficient	S.E.	Coefficient	S.E.	Coefficient	S.E.	Coefficient	S.E.
a_0	0.003	0.021	0.126	0.457	0.008	0.005	0.011***	0.004	0.083***	0.032	0.046***	0.014	0.017	0.025	3.044	2.414	0.035*	0.019	0.012	0.009
a_1																				
a_2	0.015	0.009	0.021	0.014	0.044**	0.022	0.031*	0.017	0.020**	0.009	0.016**	0.007	0.004	0.004	0.006	0.007	0.013**	0.006	0.008**	0.003
a_3	0.203***	0.011	0.203***	0.010	0.043***	0.003	0.043***	0.004	0.209***	0.006	0.209***	0.007	0.237***	0.022	0.235***	0.020	0.190***	0.009	0.195***	0.011
λ_{s1}	0.033	0.077	0.034	0.077	−0.136*	0.079	−0.131*	0.079	0.284***	0.077	0.292***	0.080	0.243***	0.076	0.250***	0.077	0.270***	0.084	0.234***	0.081
λ_{s2}	−0.085	0.078	−0.081	0.078	−0.132*	0.079	−0.170**	0.080	0.025	0.067	−0.037	0.080	−0.157**	0.078	−0.163**	0.077	0.076	0.071	0.045	0.084
λ_{o1}	−0.863***	0.311	−0.858	0.312	−5.209***	2.026	−5.581**	2.379	−10.673**	4.863	−15.297**	6.066	−8.193***	2.262	−8.463**	2.250	−7.064*	4.069	−7.663	5.432
λ_{o2}	0.521	0.462	0.517	0.461	0.866	2.987	1.000	3.585	6.435	7.678	9.858	10.859	8.810***	3.368	8.937***	3.359	1.704	5.608	1.649	9.859
λ_{o3}	−0.385	0.479	−0.454	0.474	2.211	2.957	2.113	3.186	1.200	10.061	−1.293	14.326	−6.622*	3.541	−7.156**	3.455	3.723	5.181	4.662	12.176
λ_{o4}	0.653**	0.337	0.744**	0.331	1.880	2.027	1.989	2.165	2.997	6.645	4.145	8.728	4.879**	2.467	5.753**	2.366	2.414	3.738	−0.441	7.772
ω_{o1}	1.066***	0.079	1.065***	0.079	1.059***	0.078	1.065***	0.079	1.056***	0.079	1.065***	0.079	1.063***	0.079	1.065***	0.079	1.063***	0.079	1.065***	0.081
ω_{o2}	−0.099	0.078	−0.100	0.079	−0.097	0.078	−0.100	0.079	−0.098	0.079	−0.100	0.079	−0.095	0.079	−0.100	0.079	−0.106	0.079	−0.100	0.081
θ_{c1}	0.030	0.084	−0.048	0.304	0.016	0.087	−0.023	0.089	0.174**	0.087	0.195**	0.090	0.120	0.086	−0.423***	0.111	0.325***	0.087	0.358***	0.084
θ_{c2}	0.108	0.084	0.134	0.098	−0.200**	0.084	−0.243***	0.085	0.036	0.098	0.041	0.091	−0.111	0.085	−0.289**	0.108	−0.113	0.085	−0.128	0.082
k_h	0.235***	0.072	0.242	0.117	0.130***	0.050	0.111**	0.046	0.161	0.516	0.296	0.478	0.140	0.405	−0.077	0.197	0.158	0.310	0.356	0.271
k_s	0.332***	0.100	0.302***	0.109	0.951*	0.561	1.242*	0.692	3.495**	1.547	6.807***	2.154	2.962***	0.731	2.723***	0.807	1.167	1.211	4.460***	1.525
k_o	0.043	0.026	0.045*	0.026	0.048*	0.025	0.045*	0.026	0.052**	0.026	0.045*	0.026	0.041	0.026	0.045	0.026	0.055**	0.026	0.045*	0.026
σ_c^2	0.020***	0.003	0.031	0.041	0.040***	0.006	0.037***	0.006	2.810***	0.500	2.412***	0.337	1.060***	0.181	27.384	32.914	0.791***	0.115	0.718***	0.090
σ_s^2	0.322***	0.036	0.322***	0.036	13.833***	1.547	13.812***	1.543	115.505***	13.037	112.905***	12.598	16.791***	1.874	16.770***	1.870	60.362***	6.855	58.243***	6.535
σ_o^2	0.021***	0.002	0.021***	0.002	0.021***	0.002	0.021***	0.002	0.021***	0.002	0.021***	0.002	0.021***	0.002	0.021***	0.002	0.021***	0.002	0.021***	0.002
σ_{so}	0.006	0.007	0.007	0.007	−0.074	0.045	−0.076*	0.045	−0.106	0.137	−0.101	0.132	0.035	0.049	0.038	0.048	−0.120	0.093	−0.118	0.106
λ_{s1}^*			0.629	0.545			−0.003	0.224			−0.003	0.074			0.899***	0.046			−0.002	0.160
λ_{s2}^*			−0.100	0.298			0.668**	0.323			0.396	0.274			0.060	0.040			0.488	0.426
LR test (p-value)																				
Unobserved cost	0.419				0.063				0.099				0.191				0.001			
Oil price	0.013				0.010				0.136				0.001				0.078			
Behavioral	0.613				0.163				0.161				0.021				1.000			

Note: This table reports maximum likelihood estimates and associated standard errors (S.E.) for of the each sectors and the two models of interest. ***, **, and * denote significance at the 1%, 5% and 10% level, respectively. Three likelihood ratio tests are also reported: (1) Unobserved cost is the LR rest for the null that θ_{c1} and θ_{c2} are jointly insignificant; (2) Oil price is the LR test for the null that λ_{o1}, λ_{o2}, λ_{o3}, and λ_{o4} are jointly insignificant; (3) Behavioral is the LR test for the null that λ_{s1}^* and λ_{s2}^* are jointly insignificant.

assumption of cointegration between inventories and sales allows us to estimate the cointegration parameter, a_3, precisely for all industries but tobacco, the remaining cost parameters are statistically significant only for some sectors. All in all, these results point toward a strong accelerator motive in all the aggregates and most of the industries. The only possible exception is tobacco, where the cost of holding inventories, a_2, exceeds that of adjusting production, a_0; however, the estimates are not statistically significant.

As for the role of the cost shock, the data seem to fit our specification where $\upsilon_{c,t}$ affects the cost of production in four industries and two manufacturing aggregates. We reject the null that θ_{c1} and θ_{c2} are jointly insignificant for food, motor vehicles, manufacturing and durables (see p-value for LR test in the row "Unobserved cost" of Table 1). In addition, our finding that $\theta_{c1} + \theta_{c2} < 1$, $\theta_{c2} - \theta_{c1} < 1$ and $|\theta_{c2}| < 1$ for all sectors supports our assumption that $\upsilon_{c,t}$ in (7) is stationary.

To conclude this section, let us compare our parameter estimates with those found in the literature for the linear–quadratic inventory model. To do so, we divide the parameter estimates reported in Table 1 by the second derivative of the objective function (4) with respect to H_t [i.e., $c = \left(1 + 4\beta + \beta^2\right) a_0 + (1 + \beta)\, a_1 + \beta a_2$] evaluated at the estimated values $\widehat{a}_0$, $\widehat{a}_2$ and the value of $a_1 = 1$ corresponding to our normalization. We then compute the median across the sectors and compare it with the estimates reported by Ramey and West (1999). First, note that the estimated slope of the marginal production cost is found to be positive (see the third column of Table 2). This is consistent with all studies but Ramey (1991). Second, as found by previous studies, the cost of adjustment a_0 contributes only slightly to this upward slope. Finally, estimates of a_3 are consistent with observable patterns of average inventory–sales ratios across industries and are comparable to estimates obtained by other authors. For instance, a_3 is smaller for industries with lower average inventory–sales ratios, such as motor vehicles (average inventory–sales ratio = 0.057) and petroleum products (0.162), but larger for industries with higher ratios, such as apparel (0.256), chemicals (0.293), and rubber and plastics (0.319). Not surprisingly, our results are more in line with studies that allow for serially correlated cost variables.

In summary, existing parameter estimates of the linear–quadratic inventory model cover a wide range (see Table 2), are seldom statistically significant, change with the normalization, sometimes have the wrong sign, and are often unsupportive of the underlying model [Fuhrer et al. (1995); Ramey and West (1999)]. Our estimates of the cost function are invariant to normalization and almost always have the correct sign, yet they are statistically insignificant in a few cases.

Oil price increases as negative demand shocks. Estimates of the structural model suggest that an unexpected oil price increase has a contractionary effect on sales. Note the negative sign and the statistical significance of the λ_{oi} ($i = 1, \ldots, 4$) for petroleum products, rubber and plastics, motor vehicles, manufacturing, nondurables, and durables. For these industries—as well as apparel and

Table 2. Comparison of industry and median point estimates of cost parameters

	a_0/c	a_1/c	$[(1+\beta)a_0+ a_1]/c$	a_2/c	a_3[a]	Number of industries
			Own estimates[b]			
Food	0.00	0.50	0.50	0.01	0.17	
Tobacco	0.01	0.36	0.39	0.20	−0.09	
Apparel	0.06	0.28	0.41	0.06	0.44	
Chemicals	0.06	0.32	0.44	0.00	0.37	
Petroleum products	0.01	0.45	0.48	0.02	0.08	
Rubber and plastics	0.00	0.50	0.50	0.01	0.20	
Motor vehicles	0.00	0.48	0.49	0.02	0.04	
Manufacturing	0.03	0.40	0.47	0.01	0.21	
Nondurables	0.01	0.48	0.50	0.00	0.24	
Durables	0.02	0.45	0.49	0.01	0.19	
			Median estimates[c]			
Models with serially correlated cost variables						
Herrera (2018)	**0.01**	**0.45**	**0.48**	**0.02**	**0.17**	**7**
Durlauf and Maccini (1995)	0	0.43	0.43	0.15	0.55	5
Eichenbaum (1989)	0	0.21	0.21	0.58	1.15	7
Kollintzas (1995)	−0.16	0.83	0.64	−0.09	1.14	6
Ramey (1991)	0.15	−0.63	−0.43	1.69	0.4	6
Models without serially correlated cost variables						
Fuhrer et al. (1995)	0.13	0.12	0.38	0	0.67	1
West (1986)	0.05	0.34	0.44	0.01	1.12	10

[a] Estimates in all studies but this one use monthly inventories and sales instead of quarterly sales and end-of-quarter inventories.
[b] Calculations are based on the estimates reported in Table 1. In the column definitions $c = (1 + 4\beta + \beta 2)a_0 + (1 + \beta)a_1 + \beta a_2$.
[c] Herrera (2018) denotes the median point estimates for all 2- and 3-digit industries reported in Table 1. The median estimates for other studies are taken from Table 10 in Ramey and West (1999).

chemicals—we reject the null that the coefficients on the oil price lags, λ_{oi}, are jointly insignificant (see p-value for the LR test in the row "Oil price" of Table 1).

To develop intuition for how oil price shocks are transmitted to inventories and output, we first relate the industries' estimated cost patterns to their responses to a negative demand shock. Table 1 shows that we can classify the industries in two groups according to the magnitude of the inventory holding cost, a_2, relative to the cost of adjusting production, a_0. For instance, the motor vehicles industry is more resistant to deviations from its target level of inventories, as suggested by the high value of a_2 relative to a_0. (That is, adjusting production, a_0, is less costly than adjusting inventories, a_2.) The estimated value of a_3 suggests that motor vehicles reduce inventories by $43 for every $100 drop in quarterly sales. In contrast, the larger value of a_0 relative to a_2, which is statistically equal to zero, for petroleum products suggests that the costs of adjusting production exceed those of adjusting inventories. As for the aggregates, durables and total manufacturing exhibit costs of adjusting production that exceed those of holding inventories.

Even though the production-smoothing motive appears to be operative in some industries ($a_0 > 0$), recall that there is evidence of a strong accelerator influence. Hence, with low positive values of a_3 and a_2 relative to a_1, we would expect an increase in oil prices to result in a decline in sales, production cutbacks, and procyclical movements in inventories. With convex production costs, the last would be the upshot of the accelerator motive dominating the incentive to smooth production.

Industry-level dynamics. We now turn to the question posed earlier: Can our recounting of the dynamics of an oil price innovation be rigorously reconciled with rational behavior by firms and apparent production-cost schedules? To address this issue, we use the Kalman filter to trace the impact of a one-time 10% increase in oil prices on sales, S_{t+j}, inventories, H_{t+j}, and output, Q_{t+j}. Figure 2 illustrates the impulse responses computed using this structural model (dashed line) as well as the cumulative responses generated by the VAR (solid line). The structural responses of sales and output roughly resemble those implied by the VAR estimates. Nevertheless, some differences are evident at long horizons. More specifically,

- *Industry-level sales decline in response to an oil price increase*. This negative correlation is evidenced in the sales contraction implied by the structural model. The negative sign and statistical significance of the oil price coefficients λ_{oi} in the sales equation, and the likelihood ratio test (see "Oil price" row in Table 1), provide additional evidence of this relationship. For tobacco, apparel, chemicals, petroleum products, and durables, the structural model generates a greater medium-run response of sales than the VAR.
- *Industry-level output declines in response to an oil price increase*. A slowdown in production is apparent for all industries and the manufacturing aggregates. Because output traces sales closely, the structural model generates larger output medium-run responses than the VAR for tobacco, apparel, chemicals, petroleum products, and durables.
- *Inventories usually decline at a slower pace than sales, leading to a hump-shaped response of the inventory–sales ratio*. This buildup is slowly worked down as inventories and sales adjust to their new steady state level. For some industries, inventories appear to exacerbate the negative effect of oil price innovations on output relative to that on sales. As we mentioned earlier, this pattern is consistent with a strong accelerator motive.
- *The contractionary effect is largest for motor vehicles but is also significant for industries that are energy-intensive or for which motor vehicles constitute an important demand factor*. According to the structural model, sales (production) of new motor vehicles decline about 2.5% by the fourth

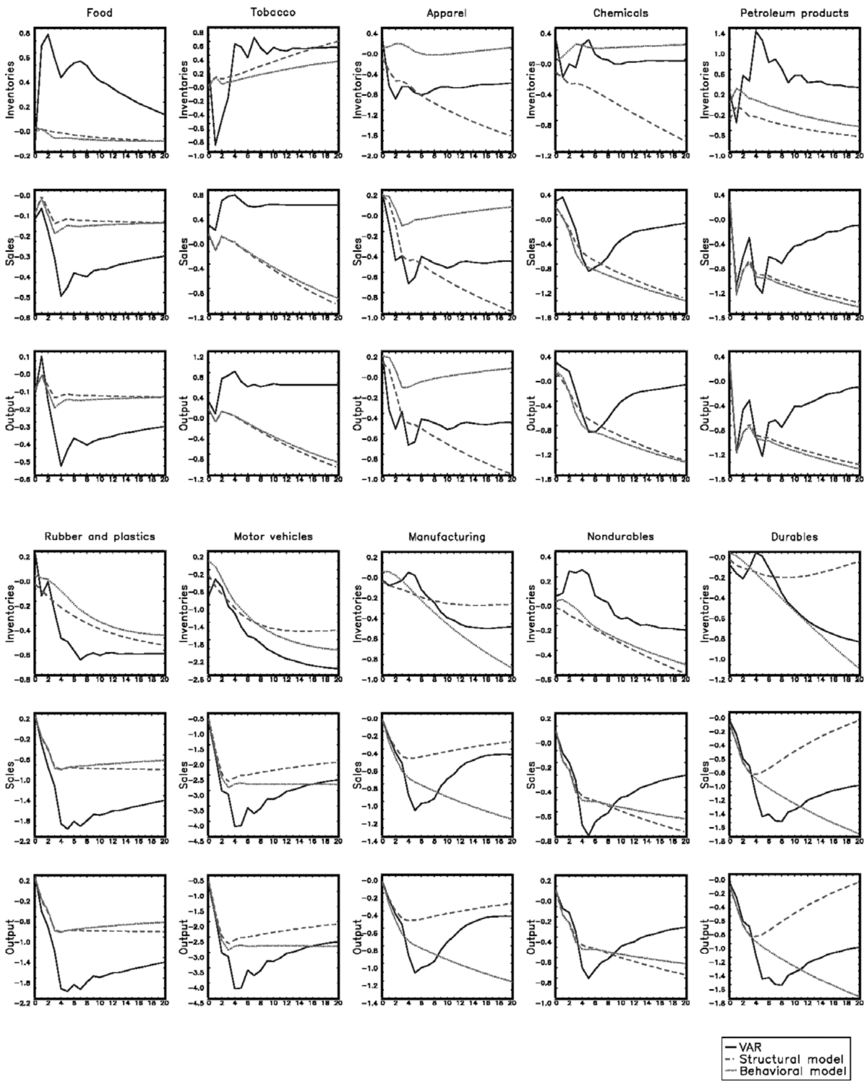

FIGURE 2. Effect of a 10% increase in the price of oil: Comparison across models.

quarter. This contraction is more than twice as large as the drop experienced by rubber and plastics (0.86%), the "production-to-stock" industry with the second largest contraction. The corresponding contractions in the production of chemicals, petroleum products, rubber and plastics, and apparel are considerably smaller (0.6, 0.9, 0.86, and 0.48%, respectively).

On the whole, the structural results are consistent with the VAR responses and suggestive of the old inventory-accelerator model of the business cycle. Consumer anxiety about oil prices leads households to cut back purchases. The firms' optimal policy response is to deviate from their target level of inventories and spread the decline in production over several quarters. In turn, the magnitude of this deviation is a function of the cost of holding inventories, a_2, the strength of the accelerator motive, a_3, and the cost of adjusting production, a_0, relative to the marginal production cost, a_1. Further, notice that this framework implies a permanent effect on the output level, although the growth rate of output returns to normal about two years after the innovation.

It is worth noting here that for all industries where oil price increases lead to a decline in sales, the response of inventories is consistent with two stylized facts documented in Ramey and West (1999): procyclicality and persistence of inventories. First, in the wake of an oil price innovation, sales fall and inventories are depleted. Second, the buildup in the inventory–sales ratio is worked down over a period of roughly two years. What leads to this procyclical movement of inventories and the persistence in the inventory–sales ratio? Given a convex cost function, a positive cost of adjusting production, $a_0 > 0$, and a positive production cost, $a_1 > 0$, the accelerator motive dominates the incentive to smooth production and thus leads to procycical inventories. Similarly, by allowing for a strong accelerator motive, the response of the inventory–sales ratio to a negative demand shock is persistent.

4.3. Can Unanticipated Changes in Inventories Exacerbate the Slowdown in Economic Activity?

Even though there are great similarities between the structural model and the VAR responses, there are also some differences. First, the structural estimates imply that firms respond immediately by reducing inventories, as both inventories and production smoothly decline to the new steady-state values. Thus, the structural responses exhibit a larger initial decline in inventories than estimated by the VARs for chemicals, rubber and plastics, and nondurables. Second, in the medium run, the structural responses of output appear to be slightly smaller than implied by the VAR, especially for food, apparel, rubber and plastics, and motor vehicles.

These differences suggest the possibility of an unanticipated and undesired accumulation of inventories, accompanied by a larger output drop in the following quarters. Yet this scenario is ruled out by construction in the structural model. One possibility worth considering is that firms do not correctly anticipate the effect that oil price innovations would have on sales. For instance, firms may rely on a simple rule of thumb when forecasting sales and making production decisions. Hence, they ignore factors believed to have only a small effect on profits [Akerlof (2002)]. We estimate a model where the process for sales, $\Delta S_t = v_{s,t}$, is given by

$$\Delta S_t = \lambda^*_{s1}\Delta S_{t-1} + \lambda^*_{s2}\Delta S_{t-2} + \varepsilon_{s,t}. \qquad \textbf{(15)}$$

In contrast with (9), here oil prices do not enter directly in the equation for sales. Note that the firm will eventually respond to the effects of an oil price innovation simply by adapting to the observed values of sales; they use (15) rather than (9) to form future sales forecasts. Instead, the econometrician uses (15) and (9) to construct the matrix **A** in (12) and from this finds the implied value of $\mathbf{F}(\lambda_{s1}^*, \lambda_{s2}^*)$. This corresponds to an econometric perspective in which oil prices really do matter for sales, but firms do not use this fact in making their production and inventory plans.

Maximum-likelihood estimates and associated standard errors for this modified framework are reported in Table 1 under the heading “Behavioral model.” The two new parameters lead to a significant increase in the log likelihood for food, chemicals, and nondurables (see p-values for the LR test in Table 1 in the row labeled “Behavioral”). Additional evidence that this behavioral story is consistent with the observed data can be gathered by comparing the impulse response functions in Figure 2. Note that for all sectors, except food and tobacco, the behavioral model (dotted red line) implies a more sluggish initial response of inventories to an oil price innovation. As these inventories are liquidated, they amplify the effect of the oil price innovation on production.

I conclude this section with a caveat. Clearly, the behavioral model is not the only alternative to the linear–quadratic inventory model. For some industries, other specifications might fit the data better. For instance, the assumption that sales are exogenous might be too strong for some firms, as could be the assumption of quadratic adjustment costs. I leave the study of alternative modifications for future research.

5. CONCLUSIONS

A puzzling aspect of the historical correlation between oil prices and aggregate economic activity is the substantial time lag between the increase in crude oil prices and the slowdown in real GDP growth. Typically, a decline in economic activity does not show up until four quarters after an unexpected oil price increase. This paper uses disaggregated manufacturing data to inquire into the causes of this time delay.

Using a VAR framework, I uncovered four features of the dynamics of oil price innovations at the industry level: (1) oil price innovations lead to a faster slowdown in industry-level output than in aggregate GDP; (2) industry-level sales decline in response to an oil price increase; (3) the response of the inventory–sales ratio is “hump-shaped,” with inventories exhibiting procyclical behavior; (4) the negative effect of an oil price increase is largest for motor vehicles output, yet significant contractions also occur in industries that are energy-intensive and for which motor vehicles constitute an important demand factor.

I then inquired whether these patterns were consistent with a model of firm behavior. Estimates of a modified linear–quadratic inventory model revealed a potential role for oil price innovations as a negative demand shifter. With convex costs and a strong accelerator motive, firms respond to this negative demand

https://doi.org/10.1017/S1365100516000225 Published online by Cambridge University Press

shock by depleting inventories and curtailing production. Partly because the shock catches manufacturers by surprise and partly because of their desire to balance the accelerator and production-smoothing incentives, manufacturers deviate from the target level of inventories and spread the decline in output over various quarters. By the end of the first year, further declines in production are evident across various industries, thus leading the economy into a recession.

NOTES

1. Estimates based on multivariate VARs indicate a two-quarter lag in the response of output to monetary policy shocks [Christiano et al. (2000)] and an immediate response to technology shocks [Christiano et al. (2003)].

2. See, for example, Blanchard (1983), West (1986), Eichenbaum (1989), Krane and Braun (1991), Ramey (1991), Kashyap and Wilcox (1993), West and Wilcox (1994, 1996), Durlauf and Maccini (1995), and Fuhrer et al. (1995), and—for excellent surveys—West (1995) and Ramey and West (1999).

3. Note that whereas the term shock usually refers to an independent and identically distributed (i.i.d.) innovation in the VAR literature, the inventory literature defines a cost shock, $U_{c,t}$, as a stochastic exogenous variation in the cost of production. This stochastic process may have a unit root and could have both observed and unobserved components [see, e.g., Ramey and West (1999) and Hamilton (2002)].

4. Although the methods for estimating models with more general cost structures are well known [Anderson et al. (1996)], they are usually not implemented in the inventory literature because of their higher computational burden.

5. To convert the inventory data from cost to market prices we follow West (1983). Because of the change in industry classification from SIC to NAICS in the late 1990s, there is no concordance between the older 2-digit SIC data and the newer NAICS data. Hence we are not able to expand our sample beyond 2000:Q1.

6. The Online Appendix is available at http://gatton.uky.edu/faculty/herrera/documents/OilAppendix.pdf.

7. Because output, Y_t, is defined as the sum of sales, S_t, and inventory investment, ΔH_t, if $S_t \sim I(1)$ and $H_t \sim I(1)$, then $Y_t \sim I(1)$.

8. Let $\begin{bmatrix} \Delta o_t \\ \Delta s_t \\ h_t - s_t \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & -1 & 1 \end{bmatrix} \begin{bmatrix} \Delta o_t \\ \Delta s_t \\ \Delta h_t \end{bmatrix} + \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} (h_{t-1} - s_{t-1})$;
then the system in (1) can be rewritten as

$$\begin{bmatrix} \Delta o_t \\ \Delta s_t \\ \Delta h_t \end{bmatrix} = \begin{bmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 1 & 1 \end{bmatrix} A(L) \begin{bmatrix} u_{o,t} \\ u_{s,t} \\ u_{h,t} \end{bmatrix} - \begin{bmatrix} 0 \\ 0 \\ 1 \end{bmatrix} (h_{t-1} - s_{t-1}).$$

9. Note that when chain-aggregated data are used, the arithmetic sum of real sales and real investment in finished goods inventories constitutes only an approximate measure of output, given that the price deflators of the two series might differ [Whelan (2000)].

10. See Tables A.2 and A.3 of the Online Appendix.

11. The specification here is similar to that in Ramey and West (1999). However, the notation differs from theirs in that here production costs are given by $(1/2)a_1 Q_t^2 - a_1 Q_t U_{c,t} + U_{c,t}^2$, whereas Ramey and West specify production costs as $(1/2)a_1 Q_t^2 + Q_t U_{c,t}^*$. From the point of view of the firm, the term $U_{c,t}^2$ is a constant that does not affect the first-order conditions. The normalization $-a_1 Q_t U_{c,t} = Q_t U_{c,t}^*$ only simplifies the algebra.

12. See Hamilton (2002) for a detailed discussion of the interpretation of cointegration in the linear–quadratic model.

13. A detailed description of the optimization problem in the matrix form can be found in the Online Appendix.

14. The parameters a_0, a_1, and a_2 in the cost function (4) are only identified to a scale; thus the need for the normalization.

REFERENCES

Akerlof, George A. (2002) Behavioral macroeconomics and macroeconomic behavior. *American Economic Review* 92(3), 411–433.

Anderson, Evan W., Lars P. Hansen, Ellen R. McGrattan, and Thomas J. Sargent (1996) Mechanics of forming and estimating dynamic linear economies. In H.M. Amman, D.A. Kendrick, and J. Rust (eds.), *Handbook of Computational Economics*, Vol. 1, pp. 171–252. Amsterdam: North-Holland.

Blanchard, Olivier J. (1983) The production and inventory behavior of the American automobile industry. *Journal of Political Economy* 91, 365–400.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans (2000) Monetary policy shocks: What have we learned and to what end? In John B. Taylor and Michael Woodford (eds.), *Handbook of Macroeconomics* Vol. 1, pp. 65–148. Amsterdam: North-Holland.

Christiano, Lawrence J., Martin Eichenbaum, and Robert J. Vigfusson (2003) What Happens after a Technology Shock? International finance discussion paper 768, Washington, DC: Board of Governors of the Federal Reserve System.

Durlauf, Steven N. and Louis J. Maccini (1995) Measuring noise in inventory models. *Journal of Monetary Economics* 36, 65–89.

Eichenbaum, Martin S. (1989) Some empirical evidence on the production level and production cost smoothing models of inventory investment. *American Economic Review* 79, 853–864.

Fuhrer, Jeffrey C., George R. Moore, and Scott D. Schuh (1995) Estimating the linear–quadratic inventory model: Maximum likelihood versus generalized method of moments. *Journal of Monetary Economics* 35, 115–157.

Hamilton, James D. (2002) On the interpretation of cointegration in the linear–quadratic inventory model. *Journal of Economic Dynamics and Control* 26, 2037–2049.

Hamilton, James D. and Ana María Herrera (2004) Oil shocks and aggregate economic behavior: The role of monetary policy. *Journal of Money, Credit and Banking* 36(2), 265–286.

Holt, Charles C., Franco Modigliani, John F. Muth, and Herbert A. Simon (1960) *Planning Production, Inventories, and Work Force*. Englewood Cliffs, NJ: Prentice-Hall.

Kashyap, Anil K. and David W. Wilcox (1993) Production and inventory control at the General Motors Corporation during the 1920's and 1930's. *American Economic Review* 83, 383–401.

Kilian, Lutz (1998) Small-sample confidence intervals for impulse response functions. *Review of Economics and Statistics* 80(2), 218–230.

Kilian, Lutz and Logan T. Lewis (2011) Does the fed respond to oil price shocks? *Economic Journal* 121, 1047–1072.

Kilian, Lutz and Clara Vega (2011) Do energy prices respond to U.S. macroeconomic news? A test of the hypothesis of predetermined energy prices. *Review of Economics and Statistics* 93(2), 660–671.

Kollintzas, Tryphon (1995) A generalized variance bounds test with an application to the Holt et al. inventory model. *Journal of Economic Dynamics and Control* 19, 59–89.

Krane, Spencer D. and Steven N. Braun (1991) Production-smoothing evidence from physical-product data. *Journal of Political Economy* 99, 558–577.

Lee, Kiseok and Shawn Ni (2002) On the dynamic effects of oil price shocks: A study using industry level data. *Journal of Monetary Economics* 49, 823–852.

Mork, Knut A. (1989) Oil and the macroeconomy when prices go up and down: An extension of Hamilton's results. *Journal of Political Economy* 91, 740–744.

Ramey, Valerie A. (1991) Nonconvex costs and the behavior of inventories. *Journal of Political Economy* 99, 306–334.

Ramey, Valerie A. and Kenneth D. West (1999) Inventories. In John B. Taylor and Michael Woodford (eds.), *Handbook of Macroeconomics*, Vol. IB, pp. 863–923. Amsterdam: North-Holland.

West, Kenneth D. (1983) A note on the econometric use of constant dollar inventory series. *Economic Letters* 13, 337–341.

West, Kenneth D. (1986) A variance bounds test of the linear quadratic inventory model. *Journal of Political Economy* 94, 374–401.

West. Kenneth D. (1995) Inventory models. In M. Pesaran and M. Wickens (eds.), *Handbook of Applied Econometrics*, Vol. I (Macroeconomics), pp. 188–220. Oxford, UK: Basil Blackwell.

West, Kenneth D. and David W. Wilcox (1994) Estimation and inference in the linear–quadratic inventory model. *Journal of Economic Dynamics and Control* 18, 897–908.

West, Kenneth D. and David W. Wilcox (1996) A comparison of alternative instrumental variables estimators of a dynamic linear model. *Journal of Business and Economic Statistics* 14, 281–293.

Whelan, Karl (2000) A Guide to the Use of Chain Aggregate NIPA Data. Federal Reserve Board of Governors finance and economics discussion series 2000–35.