Inventory Behavior with Permanent Sales Shocks

Louis J. Maccini
(Corresponding author)
Department of Economics
Johns Hopkins University
3400 N. Charles Street
Baltimore, MD 21218
(410) 516-7607
maccini@jhu.edu

Bartholomew Moore
Department of Economics
Fordham University
441 East Fordham Road
Bronx, NY 10458
(718) 817-4049
bmoore@fordham.edu

Huntley Schaller
Department of Economics
Carleton University
1125 Colonel By Drive
Ottawa, ON K1S 5B6
(613) 520-3751
huntley_schaller@carleton.ca

April, 2012

We thank Martin Eichenbaum, James Kahn and Adrian Pagan for comments on an earlier draft of the paper and seminar participants at the American Economic Association, Bank of Canada, Econometric Society, Fordham University, Institute for Advanced Studies (Vienna), International Society for Inventory Research, The Society for Computational Economics, the University of Kentucky, Mark Blanchette for excellent research assistance, and the SSHRC for financial assistance.

Inventory Behavior with Permanent Sales Shocks

Abstract

Empirically, sales are I(1). We derive a new model of inventories based on this fact. Our theory implies three startling results. First, the variance of production is equal to the variance of sales in the long run. Second, this result holds regardless of the strength of production smoothing, stockout avoidance, or cost shocks. Third, at business cycle horizons, the conditional variance of production is greater than sales. Our theory leads to a different way of estimating, testing, and calibrating inventory models. The calibrated model simultaneously accounts for four traditional inventory puzzles and three puzzles about inventories and monetary policy.

[JEL Classification Codes: E22, E23]
Keywords: Inventories, Production Smoothing, Stockout Avoidance,
Cointegration, Monetary Policy Effects

April, 2012

I. INTRODUCTION

Inventory movements are important. In the 2007-09 recession, inventories accounted for one-third of the fall in US GDP, a huge amount for such a small component of output.[1] This is typical: Inventory movements account for a wildly disproportionate share of macroeconomic fluctuations in most postwar US recessions – and in other countries, too.[2] Despite the importance of inventory fluctuations, there are large gaps in our understanding of the basic economics of inventories. Equally seriously, there are sharp contradictions between the predictions of standard theory and the response of inventories to the main macroeconomic policy tool, monetary policy.

It has long been thought that inventories act as a shock absorber for fluctuations in aggregate demand. Standard economic theories imply that inventories are used to smooth production. A long-standing puzzle is why production varies more than sales in the data. A variety of theoretical explanations have been proposed. Our theory implies that these explanations all turn out to be irrelevant, at least in the long run.

In the data, sales are I(1). We develop a new theory of inventories in which this fact plays a central role. Our model implies three startling new analytical results. First, in the empirically relevant case, the variance of production is <u>equal</u> to the variance of sales in the long run. Second, this result holds regardless of the strength of production smoothing, stockout avoidance, or cost shocks. Third, at business cycle horizons, the conditional variance of production is <u>greater</u> than sales. This implies that inventories amplify sales shocks during business cycles, rather than dampening shocks as production smoothing would imply.[3]

Our theory leads to a different way of estimating, testing, and calibrating inventory models. When sales are I(1), our theory implies that inventories will also be I(1), an implication that is consistent with the data. This means that we need to derive a cointegrating relationship.[4] This is more difficult than it might initially appear, because we need to linearize the firm's Euler equation around stationary variables, a task that has not been explicitly addressed in the existing literature. The workhorse linear-quadratic inventory model does not lend itself to the task. We

---

[1] According to NIPA data, over the six quarters 2008:1-2009:2, the cumulative change in inventory investment was 34.8% of the cumulative change in GDP.

[2] See Blinder and Maccini (1991) and Ramey and West (1999).

[3] On the theory side, a pioneering paper is West (1990), who, in the context of other issues, obtains a weak inequality on the relative variance of production and sales, allowing for both stationary and $I(1)$ sales. We build on West's work and obtain more specific results for $I(1)$ sales.

[4] Hamilton (2002), Kashyap and Wilcox (1993), Ramey and West (1999), and Rossana (1993, 1998) are early papers that consider the cointegrating relationship.

present a model that captures much of the flavor of the linear-quadratic model but that can be linearized around stationary variables. We test whether other variables that might affect inventories (e.g., input costs) are I(1), as well as testing whether the variables that are assumed to be stationary in our theory are, in fact, I(0).

Our theory leads to a specific cointegrating relationship that we estimate on aggregate US data. Past efforts to estimate the effects of the determinants of inventories, based on I(0) econometrics, have often suggested that key variables, such as input costs and the interest rate, have coefficients with the wrong sign or statistically insignificant coefficients. In contrast, the coefficients in the cointegrating relationship implied by our theory all have the theoretically predicted signs and are strongly significant.

Our theory shows how the underlying structural parameters can be calculated from the cointegrating relationship implied by the model. Using this novel approach to calibration, which flows directly from our theory, we simulate our model. Simulations of the calibrated model show that it provides a unified explanation for four traditional inventory puzzles. In addition, the model accounts for three puzzles about monetary policy and inventories that have been documented in the previous literature.

An important traditional puzzle that has plagued the literature for decades is the *Variance ratio puzzle*. If production costs are convex, then firms want to smooth production in response to demand shocks. This has long been thought to imply that production should vary less than sales. But, empirically, production tends to vary more than sales. As noted above, our theoretical model implies that the variance of production should be equal to the variance of sales in the long run. Why then do empirical studies typically find that production varies more than sales? Simulations of our model reveal that small sample bias is the culprit. The simulations indicate that, in samples of the size used in empirical studies in the inventory literature, conventional procedures will suggest that production moves more than sales. Asymptotically, however, the variance of production is equal to the variance of sales.

Other traditional puzzles include the following. *Slow adjustment puzzle*: As an influential survey of the inventory literature puts it, "One major difficulty with stock-adjustment models is that adjustment speeds generally turn out to be extremely low; the estimated adjustment speed is often less than 10 percent per month. This is implausible when even the widest swings in inventory stocks amount to no more than a few days of production. [Blinder and Maccini (1991, page 81)]". *Wen puzzle:* Wen (2005) distinguishes between the movements of output and sales at short horizons (less than three quarters) and medium horizons (about 8-40 quarters). At medium horizons, he finds that production is more volatile than sales. More surprisingly, he finds that production is less volatile than sales at short horizons. Wen argues that these stylized facts constitute a "litmus test" for inventory theory and concludes that none of the existing explanations for the variance ratio puzzle -- stockout avoidance, cost smoothing, or increasing returns to scale -- can account for the behavior of output and sales at both short and medium horizons. *Input cost puzzle*: When costs are low, firms have an incentive to produce more and build up their inventories. It has been surprisingly difficult, however, to find evidence of an empirical relationship between observable costs and inventories. Simulations of our calibrated model enable us to explain these traditional inventory puzzles as well.

Our theory yields a closed-form solution for the conditional variance ratio – the variance of output relative to that of sales over the short and medium run. The insights from the formula for the conditional variance ratio play a key role in the simultaneous solution of three of the traditional puzzles -- the variance ratio, slow adjustment, and Wen puzzles. One of the most important insights is that convexity of production costs, which provides the motive for production smoothing, is consistent with a higher medium-run variance for production than for sales. The formula for the conditional variance ratio shows that the relative importance of two key motives -- production smoothing and stockout avoidance -- depends on the steady-state real interest rate. Previous theoretical work has not elucidated this role of the interest rate. Another insight is that, despite the fact that the variance of production exceeds the variance of sales at business cycle horizons, the model is consistent with slow adjustment speeds in inventory investment equations due to the high convexity of production cost and thus the strong incentive for the firm to smooth production. Further, we are able to explain the Wen puzzle by calculating, using the calibrated model, that the production smoothing motive dominates high frequency inventory movements, but the stockout avoidance motive dominates business cycle movements. Finally, our cointegrating regression provides strong empirical evidence of the effect of input costs on inventories in the long-run, which explains the input cost puzzle. Despite the strong evidence of the effect of input costs on inventories, simulations of our calibrated model indicate that input costs have little effect on the conditional variance ratio.

Here are the three monetary policy puzzles. *Mechanism puzzle*: Monetary policy changes the interest rate and should affect inventories, since the interest rate represents the opportunity cost of holding inventories. In fact, VAR studies find that monetary policy affects inventories. But 40 years of empirical research on inventories based on $I(0)$ econometrics has generally failed to find any significant effect of the interest rate on inventories. *Sign puzzle*: Stimulative monetary policy lowers the interest rate. A decrease in the interest rate should **increase** inventories. VAR studies find that the short-term effect of stimulative monetary policy is just the opposite -- inventories **decline**. *Timing puzzle:* Expansionary monetary policy induces a transitory decline in the interest rate. The effect of monetary policy on the interest rate largely disappears within one year. But inventories begin to rise only after the transitory shock to the interest rate has largely dissipated.

In our model, the firm's response to an interest rate movement depends on the extent to which the firm believes the movement is persistent. This makes the transitional dynamics of the inventory response to a change in the interest rate complex and nonlinear – and therefore difficult to detect using $I(0)$ econometrics. In contrast, when we use $I(1)$ econometrics – specifically, the cointegrating regression implied by our model – the data provide strong evidence of the role of the interest rate. The combination of model and empirical evidence provides the solution to the mechanism puzzle.

The key to our model's success in explaining the sign puzzle is the role of inventories in buffering demand shocks. A stimulative monetary policy shock lowers the interest rate and increases sales, but the firm cannot immediately raise production, so inventories fall at the same time that monetary policy is pushing the interest rate down.

Two elements of our model explain the timing puzzle. First, the firm takes time to learn whether a movement in the interest rate is persistent (i.e., represents a regime switch). This delays the firm's response to the interest rate movement. Second, production smoothing plays a role. A stimulative monetary policy shock lowers the interest rate and increases the desired level of inventories. But, because of the convexity of the production cost function, the firm is reluctant to adjust production too sharply, so the change in inventories is gradual.

The paper is organized as follows. Section II introduces the model with $I(1)$ sales. Section III states the decision rule for inventories implied by the model. Section IV presents three startling results regarding the relative variance of output that emerge from the model. Section V describes the model's implications for estimation and testing. Section VI outlines our innovative approach to calibration, which flows from the model. Section VII explains how the model resolves four traditional inventory puzzles -- slow adjustment, variance ratio, Wen, and input cost. Section VIII examines how the model accounts for the three monetary policy puzzles. Section IX provides a summary and conclusion.

II. THE MODEL

The literature on inventory models has been dominated by the use of linear-quadratic approximations of an underlying cost function originally advanced in Holt, et al (1960)[5]. In this paper, we depart from the linear-quadratic literature by assuming a constant elasticity approximation to an underlying cost function. We utilize a constant elasticity approximation to ensure that the equilibrium conditions can be expressed in terms of stationary ratios.

The representative firm is assumed to minimize the present value of its expected costs over an infinite horizon.[6] Real costs per period consist of production costs and inventory holding costs. Production costs, PC_t, are defined as

$$PC_t = A_t Y_t^{\theta_1} W_t^{\theta_2} \tag{1}$$

with, $\theta_1 > 1$, $\theta_2 > 0$, where Y_t is real output and W_t is real input costs, which we will measure with real input prices of variable factors of production, and A_t is a shift variable that captures the state of technology and fixed factors of production.[7] Observe that average production costs, J_t, are

---

[5] Studies in the literature that have used the linear-quadratic model in work on inventories include, for example, Blanchard (1983), Blinder (1986-b), West (1986), Miron and Zeldes (1988), Eichenbaum (1989), Durlauf and Maccini (1996), Hamilton (2002), Humphreys, et al (2001), Kashyap and Wilcox (1993), and Wen (2005).

[6] We assume that the firm minimizes discounted expected costs and thereby abstract from market structure issues, because our key innovation is to recognize that sales are $I(1)$ and to analyze the implications of this empirical fact for the long-run behavior of inventories. See, e.g., Bils and Kahn (2000), Chang, Hornstein and Sarte (2009) and Jung and Yun (2011, 2012) for models that deal with market structure issues. Even though we abstract from market structure issues, our model is quite successful in capturing many aspects of the behavior of inventories. An interesting question for future research is whether our characterization of inventory behavior at business cycle horizons can be refined by incorporating market structure issues into the model.

[7] In the empirical work, we allow θ_1 to be freely estimated without imposing the assumption that $\theta_1 > 1$, though $\theta_1 > 1$ is required for positive and rising marginal production costs. A production cost function with rising marginal production costs, due to either the presence of fixed factors of production or diminishing returns to scale, has been

$$J_t = \frac{PC_t}{Y_t} = A_t Y_t^{\theta_1 - 1} W_t^{\theta_2} \tag{2}$$

and marginal production costs are $\theta_1 J_t$.

Inventory holding costs, HC_t, are

$$HC_t = \delta_1 \left(\frac{N_{t-1}}{X_t} \right)^{\delta_2} X_t + \delta_3 N_{t-1} \tag{3}$$

with $\delta_1 > 0$, $\delta_2 < 0$, and $\delta_3 > 0$, where N_t is the stock of finished goods inventories at the end of period t, and X_t is the level of real sales, which is given exogenously.[8] Inventory holding costs consist of two basic components. One, $\delta_1 \left(\frac{N_{t-1}}{X_t} \right)^{\delta_2} X_t$, which we refer to as stockout avoidance costs, captures the idea that, given sales, higher inventories reduce costs in the form of lost sales because they reduce stockouts.[9] The other, $\delta_3 N_{t-1}$, which we refer to as storage costs, captures the idea that higher inventories raise holding costs in the form of storage costs, insurance costs, etc.[10]

---

widely used in the inventory literature to capture the production smoothing motive. See, for example, the papers listed in footnote 4 as well as Kashyap and Wilcox (1993) and Hamilton (2002) who, as we do, use cointegration methods in their empirical work.

[8] The assumption that sales are exogenous is empirically consistent with the pioneering work on inventories and cointegration by Granger and Lee (1989), who conclude (page S151) that, "The sales series may be thought of as being largely exogenously determined." Theoretically, sales can be endogenized by specifying an inverse demand function. Industry equilibrium can be analyzed with such a demand curve, as in Eichenbaum (1989). Alternatively, Christiano and Eichenbaum (1989) and West (1990) derive such a linear inverse demand curve in general equilibrium. In linear-quadratic inventory models, this leads to a decision rule that is similar to the case with exogenous sales. See, e.g., Ramey and West (1999, Section 4). An alternative approach to endogenizing sales is to incorporate inventories into a general equilibrium model. See Jung and Yun (2005), Khan and Thomas (2007), Wen (2008), Wang and Wen (2009), Iacoviello, Schiantarelli and Schuh (2007), among others. A potentially interesting topic for future research is to take the model of firm behavior developed here and incorporate it into a general equilibrium model.

9 See Bils and Kahn (2000) for a model that utilizes a constant elasticity specification of the benefits of holding finished goods inventories, with the benefits embedded on the revenue side of the firm. As discussed in footnote 6, there are benefits to abstracting from market structure issues if the objective is to take account of the fact that sales are $I(1)$ and analyze the long-run behavior of inventories. See Maccini and Pagan (2009) for a recent paper that uses a specification of the benefits of holding finished goods inventories that is similar to equation (3).

[10] These two components underlie the rationale for the quadratic inventory holding costs in the standard linear-quadratic model. The formulation above separates the components and assumes constant elasticity functional forms which facilitates log-linearization around steady-state ratios. Observe that (3) implies a "target stock" of finished goods inventories that minimizes finished goods holding costs. The target stock is $N_t^{TS} = \left(-\delta_3 / \delta_1 \delta_2\right)^{\frac{1}{\delta_2 - 1}} X_t$ so that the implied stock is proportional to sales. This is analogous to the target stock assumed in the standard linear-quadratic

The firm's information set at time t, Ω_t, includes the current and past values of all relevant variables, but when the firm chooses N_t its information set is Ω_{t-1}. This assumption, which is standard in the inventory literature, captures the idea that inventories buffer demand shocks; for example, the firm may meet an unanticipated increase in sales out of its stock of inventories. Let β_t be a variable real discount factor, which is given by $\beta_t = 1/(1+r_t)$, where r_t denotes the real rate of interest. The firm's optimization problem is to minimize the present discounted value of expected total costs,

$$E_0 \sum_{t=1}^{\infty} \left[\prod_{j=1}^{t} \beta_j \right] C_t, \tag{4}$$

where $E_0 = E\{.\mid\Omega_0\}$, and

$$C_t = PC_t + HC_t = A_t Y_t^{\theta_1} W_t^{\theta_2} + \delta_1 \left(\frac{N_{t-1}}{X_t} \right)^{\delta_2} X_t + \delta_3 N_{t-1}, \tag{5}$$

subject to the inventory accumulation equation, which gives the change in inventories as the excess of production over sales,

$$N_t - N_{t-1} = Y_t - X_t. \tag{6}$$

The optimality conditions that result from this optimization problem are

$$E_{t-1}\left\{ \beta_t \left[\theta_1 A_t Y_t^{\theta_1 - 1} W_t^{\theta_2} - \xi_t \right] \right\} = 0 \tag{7}$$

and

$$E_{t-1}\left\{ \beta_t \left[\beta_{t+1} \left(\delta_2 \delta_1 \left(\frac{N_t}{X_{t+1}} \right)^{\delta_2 - 1} + \delta_3 - \xi_{t+1} \right) + \xi_t \right] \right\} = 0 \tag{8}$$

where ξ_t is the Lagrange multiplier associated with the constraint (6).

To interpret the optimality conditions, eliminate the Lagrange multiplier to reduce the optimality conditions to

$$E_{t-1}\beta_t \theta_1 A_t Y_t^{\theta_1 - 1} W_t^{\theta_2} + E_{t-1}\left\{ \beta_t \beta_{t+1} \left(\delta_2 \delta_1 \left(\frac{N_t}{X_{t+1}} \right)^{\delta_2 - 1} + \delta_3 \right) \right\} = E_{t-1}\beta_t \beta_{t+1} \theta_1 A_{t+1} Y_{t+1}^{\theta_1 - 1} W_{t+1}^{\theta_2}$$

---

model. Note that the target stock is not the steady-state stock of finished goods inventories. The steady-state stock minimizes total costs in steady state whereas the target stock merely minimizes inventory holding costs.

Now, $E_{t-1}\beta_t\theta_1 A_t Y_t^{\theta_1-1}W_t^{\theta_2}$ is the marginal cost of producing a unit of output today, $E_{t-1}\beta_t\beta_{t+1}\theta_1 A_t Y_{t+1}^{\theta_1-1}W_{t+1}^{\theta_2}$ is the discounted marginal cost of producing a unit of output tomorrow, and $E_{t-1}\left\{\beta_t\beta_{t+1}\left(\delta_2\delta_1\left(N_t / X_{t+1}\right)^{\delta_2-1}+\delta_3\right)\right\}$ is the discounted marginal holding cost. The Euler equation thus states that the firm should equate the marginal cost of producing a unit of output today and carrying it in inventories to the discounted marginal cost of producing the unit of output tomorrow.

In Appendix B, we show that linearizing the optimality conditions around steady-state values yields a linearized Euler equation:

$$
\begin{aligned}
&E_{t-1}\left\{(\theta_1-1)\theta_1\overline{J}\left[\ln Y_t-\overline{\beta}\ln Y_{t+1}\right]+\theta_2\theta_1\overline{J}\left[\ln W_t-\overline{\beta}\ln W_{t+1}\right]\right.\\
&\qquad\left.+\overline{\beta}\left(\delta_2-1\right)\delta_2\overline{\psi}\left[\ln N_t-\ln X_{t+1}\right]+\theta_1\overline{J}r_{t+1}+\theta_1\overline{J}\tilde{u}_{t+1}^{A}+c\right\}=0
\end{aligned}
\tag{9}
$$

where $\overline{J}$ is steady-state average production cost, $\theta_1\overline{J}$ is steady-state marginal production cost[11], $\overline{\psi}$ is steady-state value of average stockout avoidance costs[12], which are defined by $\psi_t=\delta_1\left[R_{Nt}\left(1-x_{t+1}\right)\right]^{\delta_2-1}$, $\delta_2\overline{\psi}$ is steady-state marginal stockout avoidance costs, $\overline{R}_N$ is the steady-state inventory/sales ratio, $\overline{\beta}=1/\left(1+\overline{r}\right)$, $\overline{r}$ is the unconditional mean real interest rate, $\overline{x}$ is the steady-state growth rate of sales, $\tilde{u}_{t+1}^{A}$ is a stationary shock, c is a constant, and a bar above a variable denotes a steady-state value.

In the data, sales are I(1), as shown in Table 1, which shows that N is also I(1). As we will show, the fact that sales are I(1) has startling implications.

[11] See Hamilton (2002) for a careful discussion of the stationarity properties of marginal productions costs that are implied by inventory models. In particular, Hamilton (2002) shows how stationarity of marginal production costs arises naturally when sales, costs, output, etc. are nonstationary.

[12] Note that $\overline{\psi}$ is average steady state stockout avoidance costs, not average total inventory holding costs. The latter is $\overline{\psi}+\delta_3$, which includes both stockout avoidance costs and storage costs.

Table 1
Unit Root Tests and Estimated Cointegrating Regression

Panel A: Unit Root Tests – Variables in Cointegrating Regression				
N	X	W	π_1	π_3
-2.808 [0.194]	-3.202 [0.084]	-2.236 [0.469]	-2.732 [0.223]	-3.101 [0.106]

Panel B: Unit Root Tests – Ratios Assumed to be Stationary			
N/X	Y/X	J	ψ
-3.878 [0.013]	-8.727 [0.000]	-6.491 [0.000]	-3.855 [0.014]

N is inventories, X is sales, Y is output, W is input costs, J is average production costs, ψ is related to marginal stockout avoidance costs, π_1 is the probability of being in the low-interest-rate state, and π_3 is the probability of being in the high-interest-rate state. All variables are log-linearly detrended. The cell entries are ADF tests for unit roots, p-values in brackets. (The number of lags in the ADF tests was chosen using a standard criterion; i.e., the lag length that minimizes the AIC plus 2. All of the unit root tests include a constant and a deterministic trend.)

Since sales and inventories are I(1), we need to linearize the optimality conditions around stationary variables. We assume that the ratios, $R_{Nt} = N_t / X_t$, $R_{Yt} = Y_t / X_t$, J_t, and ψ_t, are stationary. Table 1 presents unit root tests. Standard ADF tests reject the null hypothesis of a unit root for each of the four ratios.

We assume further that the real interest rate follows a three-state Markov switching process.[13] Specifically, we assume that the real interest rate follows

$$r_t = r_{S_t} + \sigma_{S_t} \cdot \varepsilon_t \tag{10}$$

where $\varepsilon_t \sim$ i.i.d. N(0,1) and where $S_t \in \{1,2,3\}$ follows a Markov switching process. Let $r_1 < r_2 < r_3$, so that, when $S_t = 1,2,3$, the real interest rate is in the low-interest-rate, moderate-interest-rate, and high-interest-rate regime, respectively. S_t and ε_t are assumed to be

[13] This is consistent with empirical patterns in real interest rates; see Garcia and Perron (1996) and Maccini, Moore and Schaller (2004). The latter paper describes how the firm uses its observations of the real interest rate to develop its probability assessments. For a comprehensive discussion of Markov switching processes, see Hamilton (1994, Chapter 22).

independent. Denote the transition probabilities governing the evolution of S_t by $p_{ij} = \text{Prob}(S_t = j \mid S_{t-1} = i)$. Collecting these probabilities into a matrix we have

$$P = \begin{bmatrix} p_{11} & p_{21} & p_{31} \\ p_{12} & p_{22} & p_{32} \\ p_{13} & p_{23} & p_{33} \end{bmatrix}. \tag{11}$$

The firm is assumed to know the structure and parameters of the Markov switching process but does not know the true real interest rate regime. The firm must therefore infer S_t from observed interest rates. We denote the firm's current probability assessment of the true state by π_t. That is,

$$\pi_t = \begin{bmatrix} \pi_{1t} \\ \pi_{2t} \\ \pi_{3t} \end{bmatrix} = \begin{bmatrix} \text{Prob}(S_t = 1 \mid \Omega_t) \\ \text{Prob}(S_t = 2 \mid \Omega_t) \\ \text{Prob}(S_t = 3 \mid \Omega_t) \end{bmatrix}. \tag{12}$$

Given π_{t-1}, the term $E_{t-1}r_{t+1}$ in equation (9) can be computed as

$$E_{t-1}r_{t+1} = \mathrm{r_v}' P^2 \pi_{t-1} = \gamma_1 \pi_{1t-1} + \gamma_2 \pi_{2t-1} + \gamma_3 \pi_{3t-1} \tag{13}$$

where $\mathrm{r_v}' = [r_1, r_2, r_3]$ and $\begin{bmatrix} \gamma_1 & \gamma_2 & \gamma_3 \end{bmatrix} \equiv \mathrm{r_v}' P^2$. Since $\pi_{1t-1} + \pi_{2t-1} + \pi_{3t-1} = 1$ by definition, we can eliminate π_{2t-1} from the right hand side of (13) to obtain

$$E_{t-1}r_{t+1} = \left(\gamma_1 - \gamma_2\right)\pi_{1t-1} + \left(\gamma_3 - \gamma_2\right)\pi_{3t-1} + \gamma_2 . \tag{14}$$

Then, substituting (14) into (9) yields

$$\begin{aligned} & E_{t-1}\Big\{\left(\theta_1 - 1\right)\theta_1 \overline{J}\left[\ln Y_t - \overline{\beta} \ln Y_{t+1}\right] + \theta_2 \theta_1 \overline{J}\left[\ln W_t - \overline{\beta} \ln W_{t+1}\right] \\ & + \overline{\beta}\left(\delta_2 - 1\right)\delta_2 \overline{\psi}\left[\ln N_t - \ln X_{t+1}\right] + \theta_1 \overline{J} \tilde{u}^A_{t+1}\Big\} + \theta_1 \overline{J}\left[\left(\gamma_1 - \gamma_2\right)\pi_{1t-1} + \left(\gamma_3 - \gamma_2\right)\pi_{3t-1} + \gamma_2\right] + c = 0 \end{aligned} \tag{15}$$

which is the log-linearized Euler equation incorporating the firm's learning process.

III. DECISION RULE

The log-linearized Euler equation implied by the model, equation (15), may be written as a second-order expectational difference equation. Solving this difference equation yields a decision rule, which is stated in the following proposition.

Proposition 1. *Decision Rule: The model implies that the firm's decision rule is*

$$\ln N_t \;=\; \Gamma_0 + \lambda_1 \ln N_{t-1} + \Gamma_X \ln X_{t-1} + \Gamma_W \ln W_{t-1} + \Gamma_{\pi 1}\pi_{1t-1} + \Gamma_{\pi 3}\pi_{3t-1} + u_t \tag{16}$$

where

$$0 \le \lambda_1 = 1 + \frac{\bar{r}+\zeta}{2} - \frac{1}{2}\left[\left(\bar{r}+\zeta\right)^2 + 4\zeta\right]^{\frac{1}{2}} \le 1, \tag{17-a}$$

$$\zeta = \frac{(\delta_2 - 1)\delta_2\bar{\psi}}{(\theta_1 - 1)\theta_1\bar{J}}\frac{\bar{R}_Y}{\bar{R}_N} > 0, \tag{17-b}$$

Γ_0 *is a constant,* u_t *is a stationary shock, and* $\bar{R}_Y$ *and* $\bar{R}_N$ *are the steady state values of* R_{Yt} *and* R_{Nt} *respectively.*

Proof: See Appendix C.

The coefficient on lagged inventories, λ_1, is the stable root of the relevant characteristic equation that emerges from solving the second-order expectational difference equation that is implied by the Euler equation. As we shall see below, $1-\lambda_1$ defines the speed of adjustment that governs the fraction of the gap between “desired” and actual inventories that is closed each period by inventory investment.

The decision rule shock, u_t, arises from unanticipated fluctuations in sales and output. In the short run, inventories act as a buffer, absorbing these unanticipated fluctuations. The resulting inventory movements involve only transitory deviations from the level of inventories dictated by the variables in the decision rule, so u_t is stationary.

The coefficients on sales, input costs and interest-rate-regime probabilities are defined in the following propositions.

Proposition 2. *Decision Rule Coefficient on Sales: The coefficient on sales in the decision rule,* Γ_X *, is*

$$\Gamma_X = \left[\frac{(\delta_2 - 1)\delta_2\bar{\psi}}{(\theta_1 - 1)\theta_1\bar{J}} - \bar{r}\right]\frac{\bar{R}_Y}{\bar{R}_N}\left(\frac{\lambda_1}{1+\bar{r}-\lambda_1}\right) \tag{18}$$

Further, $\Gamma_X \gtreqless 0$ *as* $\dfrac{(\delta_2 - 1)\delta_2\bar{\psi}}{\bar{r}} \gtreqless (\theta_1 - 1)\theta_1\bar{J}$.

Proof: See Appendix C.

The term $(\delta_2 - 1)\delta_2\overline{\psi}$ is the convexity of stockout avoidance costs and $(\theta_1 - 1)\theta_1\overline{J}$ is the convexity of production costs. Hence, an increase in sales will induce the firm to produce enough additional output to raise inventory holdings so long as the present value of the change in marginal stockout avoidance costs exceeds the change in marginal production costs, and vice-versa.

To understand the intuition for the sign of Γ_X, suppose the firm is making a marginal decision about output in response to an increase in sales. The cost of producing one more unit of output is a one-time cost. The benefit of an additional unit of inventory is the present value of the reduction in stockout avoidance costs. If the latter dominates, the firm will produce enough additional output to increase its inventory holdings. But, if the former dominates, the firm will increase output by less than the increase in sales, so its inventory holdings will fall.

Proposition 3. *Decision Rule Coefficient on Input Costs: The coefficient on input costs in the decision rule,* Γ_W *, is*

$$\Gamma_W = -\frac{\overline{r}\theta_2}{(\theta_1 - 1)}\frac{\overline{R}_Y}{\overline{R}_N}\left(\frac{\lambda_1}{1+\overline{r}-\lambda_1}\right)<0 \tag{19}$$

Proof: See Appendix C.

The model implies that the coefficient on input costs in the decision rule Γ_W will be negative. Intuitively, an increase in input costs raises production costs, which induces the firm to cut production and thereby reduce inventory holdings.

Proposition 4. *Decision Rule Coefficients on the Interest-Rate-Regime Probabilities: The model implies that the decision rule coefficients on the interest-rate-regime probabilities are*

$$\Gamma_{\pi_1} \equiv \frac{-\lambda_1}{(\theta_1 - 1)}\frac{\overline{R}_Y}{\overline{R}_N}\gamma'\left[I - \frac{\lambda_1}{1+\overline{r}}P\right]^{-1}\begin{bmatrix}1\\-1\\0\end{bmatrix} \tag{20-a}$$

$$\Gamma_{\pi_3} \equiv \frac{-\lambda_1}{(\theta_1 - 1)}\frac{\overline{R}_Y}{\overline{R}_N}\gamma'\left[I - \frac{\lambda_1}{1+\overline{r}}P\right]^{-1}\begin{bmatrix}0\\-1\\1\end{bmatrix} \tag{20-b}$$

where $\gamma' \equiv [\gamma_1 \quad \gamma_2 \quad \gamma_3]$. *Furthermore, if*

$$(p_{11} + p_{22})/2 > 0.5 \tag{21-a}$$

$$\left(p_{22} + p_{33}\right)/2 > 0.5 \tag{21-b}$$

$$p_{13} = p_{31} = 0, \tag{21-c}$$

then, $\Gamma_{\pi_1} > 0$ *and* $\Gamma_{\pi_3} < 0$.

Proof: See Appendix C.

Observe that the assumptions that $\left(p_{11} + p_{22}\right)/2 > 0.5$ and $\left(p_{22} + p_{33}\right)/2 > 0.5$ mean that the interest rate regimes are persistent, and the assumption that $p_{13} = p_{31} = 0$ means that the economy moves through the medium-interest-rate regime on its way from the high- to low-interest-rate regime, and vice versa. Intuitively, if the interest rate regimes are persistent (and they are in the data, as reported in (30) below), lower interest rates today imply that future interest rates will be lower, reducing the opportunity cost of holding inventories and thereby inducing the firm to hold more inventories.

IV. THE RELATIVE VARIANCE OF OUTPUT: THREE IMPLICATIONS

A key question about inventories is whether they amplify or dampen demand shocks. In a model with sales shocks, West (1990) obtains an inequality – $Var\left[Y\right]/Var\left[X\right] \leq 1$ – that applies both when sales follow a stationary stochastic process and when sales are $I\left(1\right)$.[14] We are able to establish a more specific result for the $I\left(1\right)$ case.

To start, we derive the decision rule for output, which takes the form of a first-order linear difference equation.[15] Using assumptions about the stochastic processes for sales and input costs that are consistent with our earlier assumption that they are $I\left(1\right)$ variables, we solve the difference equation backwards to obtain an equation for output as a function of output and sales at a fixed date in the past and subsequent sales and cost shocks. By taking the variances of $\ln Y_t$ and $\ln X_t$, conditional on output and sales at a fixed date in the past, we obtain Proposition 5.

Proposition 5*. Conditional Variance Ratio:*

---

[14] More precisely, West asserts that $E\left[X_t^2 - Y_t^2\right] \geq 0$. Under his assumption that all variables have zero unconditional mean, $E\left[X_t^2 - Y_t^2\right] = E\left[X_t^2\right] - E\left[Y_t^2\right] = Var\left[X_t\right] - Var\left[Y_t\right]$, so his result implies that $Var\left[Y\right]/Var\left[X\right] \leq 1$.

[15] To focus on production smoothing, stockout avoidance, and cost shocks, we here assume a constant interest rate.

$$\frac{\mathrm{Var}\left[\ln Y_t \middle| \ln Y_{t-n}\right]}{\mathrm{Var}\left[\ln X_t \middle| \ln X_{t-(n+1)}\right]} = \frac{1}{1+\frac{1}{n}} + \frac{\left(1-\lambda_1+\tilde{\Gamma}_X\right)}{\left(1+n\right)}\left[\left(1-\lambda_1+\tilde{\Gamma}_X\right)\left(\frac{1-\lambda_1^{2n}}{1-\lambda_1^2}\right)+2\left(\frac{1-\lambda_1^n}{1-\lambda_1}\right)\right]$$

$$+\left[\frac{\tilde{\Gamma}_W^2}{1+n}\right]\left(\frac{1-\lambda_1^{2n}}{1-\lambda_1^2}\right)\frac{\sigma_W^2}{\sigma_X^2} \tag{22}$$

where $\mathrm{Var}\left(\ln Y_t \mid \ln Y_{t-n}\right)$ *is the variance of* $\ln Y_t$ *conditional on* $\ln Y_{t-n}$, $\mathrm{Var}\left(\ln X_t \mid \ln X_{t-(n+1)}\right)$ *is the variance of* $\ln X_t$ *conditional on* $\ln X_{t-(n+1)}$, σ_X^2 *is the variance of the sales shock,* σ_W^2 *is the variance of the cost shock, and*

$$\tilde{\Gamma}_X = \frac{\bar{R}_N}{\bar{R}_Y}\Gamma_X = \left[\frac{\left(\delta_2-1\right)\delta_2\bar{\psi}}{\left(\theta_1-1\right)\theta_1\bar{J}} - \bar{r}\right]\frac{\lambda_1}{(1+\bar{r}-\lambda_1)} \gtrless 0 \tag{23-a}$$

$$\tilde{\Gamma}_W = \frac{\bar{R}_N}{\bar{R}_Y}\Gamma_W = -\frac{\bar{r}\theta_2}{\theta_1-1}\left(\frac{\lambda_1}{1+\bar{r}-\lambda_1}\right) < 0, \tag{23-b}$$

where $\tilde{\Gamma}_X$ *and* $\tilde{\Gamma}_W$ *are the elasticities of output with respect to sales and input costs, respectively.*

Proof: See Appendix C.

In the empirical literature, authors frequently compute what is referred to as the "variance ratio". The calculation takes the data for output and sales over a sample, computes the variance of each, and then takes the ratio of the variance of output to the variance of sales. The variance ratio is thus an unconditional statistic. Mathematically, we can obtain the unconditional variance ratio from the model by taking the limit of the conditional variance ratio, (22), as $n \to \infty$. Proposition 6 states the result.

Proposition 6. *Unconditional Variance Ratio:*

$$\lim_{n\to\infty}\frac{\mathrm{Var}\left[\ln Y_t \middle| \ln Y_{t-n}\right]}{\mathrm{Var}\left[\ln X_t \middle| \ln X_{t-(n+1)}\right]} = 1 \tag{24}$$

Proof: Proposition 6 follows from taking the limit of equation (22) as $n \to \infty$, noting that $0 \le \lambda_1 \le 1$. QED

Proposition 6 is our key result: When sales are $I(1)$, the variance ratio is 1. This means that the variance of production is equal to the variance of sales in the long run.

The intuition for our key result flows from the persistence of the shocks to sales. Suppose a firm has a convex cost function and faces uncertain demand. If the shocks to demand are transitory, it is optimal for the firm to produce at an intermediate level of output rather than to sometimes produce at a low level and sometimes at a high level. But suppose there is a permanent shock to sales. Then the firm increases output by the same amount as the permanent shock, because a permanent shock relocates the optimal level of output.

Proposition 7. *Effect of Production Smoothing, Stockout Avoidance, and Cost Shocks: The structural parameters θ_1, θ_2, δ_2 and the variance parameter in the stochastic process for input costs σ_W^2 have no effect on the unconditional variance ratio.*

Proof: Proposition 7 follows directly from Propositions 5 and 6. The structural parameters θ_1, θ_2, and δ_2 enter the second and third terms of equation (22) – through λ_1, $\tilde{\Gamma}_X$, $\tilde{\Gamma}_W$ -- and σ_W^2 directly enters the third term, but none of these parameters enters the first term. As $n \to \infty$, the second and third terms approach 0. QED

Proposition 7 means that the strength of production smoothing (reflected in θ_1), stockout avoidance (reflected in δ_2), and cost shocks (measured by σ_W^2 and reflected in θ_2) play no role in the long-run response of production to sales.

V. ESTIMATION AND TESTING

Since Panel A of Table 1 shows that the key variables are $I(1)$, we begin by deriving the cointegrating regression.[16]

Proposition 8. *Cointegrating Regression: The model in Section II implies that inventories, sales, input costs, and the interest-rate-regime probabilities are cointegrated, with cointegrating regression*

$$\ln N_t = b_0 + b_X \ln X_t + b_W \ln W_t + b_{\pi_1}\pi_{1,t-1} + b_{\pi_3}\pi_{3,t-1} + \nu_t , \tag{25}$$

where

---

[16] Kashyap and Wilcox (1993) and Ramey and West (1999) provide earlier derivations of a cointegrating vector for inventories (both under the assumption of a constant real interest rate).

$$b_X = 1 - \frac{\bar{r}(\theta_1 - 1)\theta_1 \bar{J}}{(\delta_2 - 1)\delta_2 \overline{\overline{\psi}}} \qquad b_W = -\frac{\bar{r}\theta_2\theta_1\bar{J}}{(\delta_2 - 1)\delta_2 \overline{\overline{\psi}}} \tag{26-a}$$

$$b_{\pi_1} = -(\gamma_1 - \gamma_2)\frac{(1+\bar{r})\theta_1 \bar{J}}{(\delta_2 - 1)\delta_2 \overline{\overline{\psi}}} \qquad b_{\pi_3} = -(\gamma_3 - \gamma_2)\frac{(1+\bar{r})\theta_1 \bar{J}}{(\delta_2 - 1)\delta_2 \overline{\overline{\psi}}}, \tag{26-b}$$

b_0 *is a constant, and* v_t *is a stationary error term.*

Proof: See Appendix C.

The equations in (26) are of crucial importance in understanding inventory behavior because they provide a mapping between the cointegrating regression coefficients and the underlying structural parameters.

Proposition 8 suggests an immediate test of the model, since it states that the variables in equation (25) will be cointegrated. The data are consistent with Proposition 8: The Johansen-Juselius test rejects the null hypothesis of no cointegrating vector, with a test statistic of 97.9 (p-value=0.001).[17]

We estimate equation (25) using the Stock and Watson (1993) Dynamic OLS (DOLS) estimator. Stock and Watson (1993) show how DOLS reduces the small sample bias in OLS estimates of cointegrating regressions.[18] Applied to estimating the cointegrating regression in (25), the DOLS empirical specification (including constant and time trend) is:

$$\begin{aligned} \ln N_t &= b_0 + b_T t + b_X \ln X_t + b_W \ln W_t + b_{\pi_1}\pi_{1,t-1} + b_{\pi_3}\pi_{3,t-1} \\ &+ \sum_{s=-p}^{p} B_{X,s}\Delta \ln X_{t-s} + \sum_{s=-p}^{p} B_{W,s}\Delta \ln W_{t-s} + \sum_{s=-p}^{p} B_{\pi 1,s}\Delta\pi_{1,t-1-s} + \sum_{s=-p}^{p} B_{\pi 3,s}\Delta\pi_{3,t-1-s} + \eta_t. \end{aligned} \tag{27}$$

The results are presented in Table 2.

---

[17] For reasons of data availability, the sample is 1959:01 to 2004:08. The number of lags used in the test is set to minimize the AIC.

[18] See also Caballero (1994), Caballero (1999), and Schaller (2006) on the small sample bias in OLS estimates of cointegrating regressions. Our own Monte Carlo simulations show that the OLS bias is severe enough in our context to yield estimates of b_X with the wrong sign. Based on our Monte Carlo simulations, we set $p = 48$ in equation (27).

Table 2
Estimated Cointegrating Regression

Constant	Time	b_X	b_W	b_{π_1}	b_{π_3}
11.589 (23.389)	1.5E-03 (10.911)	0.250 (3.098)	-0.753 (-5.244)	0.098 (10.974)	-0.028 (-4.216)

DOLS estimates of the cointegrating vector with t-statistics in parentheses.

The signs of the coefficients of the cointegrating regression that are implied by the model are presented in the next proposition.

Proposition 9. *Signs of the Coefficients in the Cointegrating Regression:*

A. $b_X \underset{<}{\overset{>}{-}} 0 \;\; as \;\; \frac{(\delta_2 - 1)\delta_2\overline{\psi}}{\overline{r}} \underset{<}{\overset{>}{-}} (\theta_1 - 1)\theta_1\overline{J}$

B. $b_W < 0$

C. If (21-a), (21-b), and (21-c) hold, then $b_{\pi_1} > 0$*, and* $b_{\pi_3} < 0$.

Proof: See Appendix C.

Proposition 9-A states that, in the long run, an increase in sales will increase (decrease) inventories if the present value of the convexity of stockout avoidance costs exceeds (is less than) the convexity of production costs. The intuition is the same as that for Proposition 2. As Table 2 shows, our results from estimating the cointegrating regression yield an estimate of b_X that is positive and statistically significant (The t-statistic is 3.1.), so our results indicate that empirically the stockout avoidance motive dominates the production smoothing motive.

Proposition 9-B states that an increase in input costs will reduce long-run inventories. Table 2 shows that the data are consistent with Proposition 9-B. The estimate of b_W is negative and strongly statistically significant. (The t-statistic is -5.2.)

Proposition 9-C states that a higher probability of the low-interest-rate regime increases inventories, and a higher probability of the high-interest-rate regime reduces inventories. Table 2 shows that the data are consistent with the predictions of Proposition 9-C. The estimate of b_{π_1} is positive and strongly statistically significant. (The t-statistic is about 11.0.) The estimate of b_{π_3} is negative and also statistically significant. (The t-statistic is -4.2.)

VI. CALIBRATION

Our approach to calibration is innovative and flows directly from the model. Since sales (and the other key variables) are $I(1)$, we use the cointegrating regression to calibrate the structural parameters of the model. Note from the definitions of b_X and b_{π_3} that

$$\frac{1-b_X}{b_{\pi_3}}=\frac{\bar{r}\left(\theta_1-1\right)}{\left(1+\bar{r}\right)\left(\gamma_2-\gamma_3\right)}. \tag{28}$$

Since $\bar{r}$, γ_2, and γ_3 are given from our estimates of the Markov switching model, we can use our estimates of b_X and b_{π_3} to obtain a unique value for θ_1 from (28). Similarly, note from the definitions of b_W and b_{π_3} that

$$\frac{b_W}{b_{\pi_3}}=\frac{\bar{r}\theta_2}{\left(1+\bar{r}\right)\left(\gamma_3-\gamma_2\right)}. \tag{29}$$

We can thus use our estimates of b_W and b_{π_3} to obtain a unique value for θ_2 from (29).

We obtain $\bar{r}$, γ_2, and γ_3 from our estimates of the parameters of the stochastic process for the real interest rate, that is, from our estimates of the elements of P and r_v. From our estimation of the three-state Markov-switching model, those estimates are

$$P=\begin{bmatrix} p_{11}=0.98 & p_{21}=0.02 & p_{31}=0.00 \\ p_{12}=0.02 & p_{22}=0.96 & p_{32}=0.05 \\ p_{13}=0.00 & p_{23}=0.02 & p_{33}=0.95 \end{bmatrix} \tag{30}$$

and $r_v' \equiv \begin{bmatrix} r_1 & r_2 & r_3 \end{bmatrix}=\begin{bmatrix} -1.37 & 1.77 & 5.04 \end{bmatrix}$. Together these estimates imply that the unconditional mean of the monthly real interest rate is $\bar{r}$ =0.001, which gives $\bar{\beta}=0.999$.

Finally, note from the definitions of b_{π_3} in (26-b) and recalling that $\bar{\psi}=\delta_1\left[\bar{R}_N\left(1-\bar{x}\right)\right]^{\delta_2-1}$

$$b_{\pi_3}=\frac{\left(1+\bar{r}\right)\left(\gamma_2-\gamma_3\right)\theta_1\bar{J}}{\left(\delta_2-1\right)\delta_2\delta_1\left[\bar{R}_N\left(1-\bar{x}\right)\right]^{(\delta_2-1)}} \tag{31}$$

Using the estimate of b_{π_3}, the normalization[19] $\delta_1 = 1$, and given values[20] for $\overline{R}_N, \overline{x}$, and $\overline{J}$, equation (31) gives a single restriction on the value of δ_2. We have assumed that $\delta_2 < 0$. We therefore search numerically over $\delta_2 \in (-\infty, 0]$ to find the value of δ_2 that satisfies (31).[21]

Thus, we obtain a unique value for each of the model's structural parameters – there are no free parameters. The values that we obtain are reported in Table 3, Panel A. The calibrated parameters are consistent with our theoretical predictions. In particular, $\theta_1 > 1$, $\theta_2 > 0$, and $\delta_2 < 0$.

Table 3
Calibrated Structural Parameters and Decision Rule Coefficients

Panel A: Cost Function Parameters			
θ_1	θ_2	δ_1	δ_2
65.097	64.354	1	-0.676

Panel B: Decision Rule Coefficients				
λ_1	Γ_X	Γ_W	Γ_{π_1}	Γ_{π_3}
0.949	0.0128	-0.0386	0.0011	-0.0008

As shown in equation (1), θ_1 and θ_2 are the elasticities of production cost with respect to output and with respect to input costs, respectively. As shown in equation (3), δ_2 is the elasticity of stockout avoidance costs with respect to the inventory/sales ratio. As is common in the inventory literature, δ_1 is normalized to 1. Consequently, θ_1 θ_2, and δ_2, are measured relative to δ_1. (The storage cost parameter δ_3 is not included in the table because it does not affect the decision rule coefficients.) The coefficients λ_1, Γ_X, Γ_W, $\Gamma_{\pi 1}$, and $\Gamma_{\pi 3}$ are the coefficients in the firm's decision rule on lagged inventories, sales, input costs, and the Bayesian probabilities of the low-interest-rate and high-interest-rate regimes, respectively.

Similarly, equations (18), (19) and (20) provide the mapping from the structural parameters to the decision rule coefficients. These can be used to derive calibrated values of the

---

[19] This normalization implies that we can only evaluate the relative magnitude of other structural parameters such as θ_1 and δ_2. A comparable situation exists with linear-quadratic inventory models, where the relative magnitude of key structural parameters determines the behavior of inventories. See, e.g., Ramey and West (1999), p. 894.

[20] $\overline{R}_N, \overline{x}$, and $\overline{J}$ are steady-state ratios. For $\overline{R}_N$ and $\overline{x}$, we therefore use the sample mean values of N_t/X_t and $\Delta X_{t+1}/X_t$, respectively, which gives $\overline{R}_N = 0.468$ and $\overline{x} = 0.00108$. Note from $J_t = A_t Y_t^{\theta_1} W_t^{\theta_2} / Y_t$ that $\overline{J}$ denotes the steady-state value of average production costs. Based on data from the 1992 Census of Manufacturing, we estimate production costs to be 73.4% of total output and set $\overline{J} = 0.734$.

[21] Our numerical search shows that only one value of δ_2 satisfies (31).

decision rule coefficients. The calibrated values of the decision rule coefficients are reported in Table 3, Panel B. From Proposition 2, Γ_X will be positive if the present value of the change in marginal stockout avoidance costs exceeds the change in marginal production costs. The calibrated structural parameters imply that Γ_X is indeed positive. Further, from Propositions 3 and 4 and the sufficient conditions stated in (21), the model implies that $\Gamma_W < 0$, $\Gamma_{\pi 1} > 0$ and $\Gamma_{\pi 3} < 0$. That is, our model implies that an increase in costs or an increase in the probability that the economy is in the high-interest-rate regime lowers inventories. The calibrated structural parameters imply calibrated values of the decision rule coefficients that are consistent with these predictions.

VII. TRADITIONAL PUZZLES

A. Slow Adjustment Puzzle

In early empirical work on inventories, a common specification was the stock-adjustment equation. Lovell (1961), for example, developed a model which yielded an inventory investment relationship of the form:

$$\mathrm{N}_t - \mathrm{N}_{t-1} = \vartheta\left(\mathrm{N}_t^* - \mathrm{N}_{t-1}\right) + \mathrm{u}_t^N \qquad (32)$$

where u_t^N is a shock. In the Lovell framework, inventory investment is proportional to the gap between the actual and desired stock of inventories. The proportionality factor, ϑ, measures the speed of adjustment, as it captures the fraction ϑ of the deviation between desired and actual inventories that is closed each period. The *slow adjustment puzzle* is that estimated values of ϑ appear to be implausibly low. Blinder and Maccini (1991, page 82) summarize the puzzle as follows, "Theory strains to explain low adjustment speeds unless the incentive to smooth production is extremely strong, which is hard to reconcile with the fact that production is more variable than sales. So the puzzle remains." [22]

---

[22] A number of possible explanations have been put forward to explain the slow adjustment puzzle. One explanation emphasized econometric problems -- either omitted variables or problems with the econometric procedure – see Maccini and Rossana (1984) and Blinder (1986a). Another explored the effect of aggregation bias – see Christiano and Eichenbaum (1989), Seitz (1988), Blinder (1986a), and Coen-Pirani (2004). See Blinder and Maccini (1991) and Ramey and West (1999) for surveys.

To derive an inventory investment relationship in our model, subtract $\ln N_{t-1}$ from both sides of (26) to get

$$\ln N_t - \ln N_{t-1} = (1-\lambda_1)\left[\ln N_t^* - \ln N_{t-1}\right] + u_t \tag{33}$$

where

$$\ln N_t^* = \frac{1}{1-\lambda_1}\left[\Gamma_X \ln X_{t-1} + \Gamma_W \ln W_{t-1} + \Gamma_{\pi 1}\pi_{1t-1} + \Gamma_{\pi 3}\pi_{3t-1} + \Gamma_o\right] \tag{34}$$

is the "desired" stock of (log) inventories. Comparing (33) to (32) we see that $1-\lambda_1$ measures the speed of adjustment. Using the definition of λ_1 in (17-a), this speed of adjustment term can be written as

$$1-\lambda_1 = -\frac{1}{2}\left\{\bar{r} + \zeta - \left[\left(\bar{r}+\zeta\right)^2 + 4\zeta\right]^{\frac{1}{2}}\right\} \tag{35}$$

where ζ is defined in (17-b).

Straightforward calculations reveal the relationship between the speed of adjustment, the convexity of production costs, and the convexity of stockout avoidance costs. To see this, differentiate (35) and use (17-b) to get $\partial(1-\lambda_1)/\partial\theta_1 < 0$, which states that greater convexity of production costs reduces the speed of adjustment. Similarly, $\partial(1-\lambda_1)/\partial\delta_2 > 0$, which states that greater convexity of stockout avoidance costs increases the speed of adjustment.

Intuitively, greater convexity of production costs increases the incentive to smooth production, which makes the firm slow to change the level of production. With respect to stockout avoidance, the intuition is as follows. When the firm pays a cost as a result of not having enough inventory, the firm wants to increase inventories when sales go up, so long as sales are positively serially correlated. The desired inventory level rises immediately when sales increase. The stronger the stockout avoidance motive, the more rapidly the firm wants to adjust output. Table 4 illustrates the effect of θ_1 and δ_2 on the speed of adjustment.

In general, it is not possible to recover the transition dynamics of a variable from a cointegrating regression. Intuitively, this is because the cointegrating regression captures the long-run behavior of the variable, abstracting from transition dynamics. Our model is an exception, because the "stickiness" of inventories arises from the structure of the model, rather than from an ad hoc adjustment cost function. By deriving the cointegrating regression from the model, we are able to recover the structural parameters from the cointegrating regression coefficients. The structural parameters (which are reported in Table 3) imply that the speed of adjustment $\vartheta = (1-\lambda_1) = 0.051$. This is consistent with estimates from the large empirical literature on the speed of adjustment of inventories. (See Blinder and Maccini (1991) and Ramey and West (1999).) The slow speed of adjustment that flows from the model implies that the incentive to smooth production is quite strong.

Table 4
Production Smoothing, Stockout Avoidance, and the Speed of Adjustment

	δ_2					
θ_1		−0.15	−0.35	−0.676	−1.4	−3
	2	0.5309	0.7255	0.8712	0.9690	0.9973
	10	0.1087	0.1859	0.3020	0.5501	0.9001
	65.097	0.0164	0.0294	**0.0513**	0.1130	0.3395
	100	0.0105	0.0190	0.0335	0.0748	0.2368
	200	0.0050	0.0093	0.0166	0.0378	0.1261

The entry in each cell is the speed of adjustment, $(1-\lambda_1)$. The calibrated values are $\theta_1 = 65.097$ and $\delta_2 = -0.676$.

Equations (33) and (34) and Table 4 help to explain why $I(0)$ econometrics has largely been a failure in the study of inventories. When the speed of adjustment of inventories is slow, it takes so long for movements in the variables that determine long-run desired inventories , N_t^* , to have an effect on current inventories, N_t, that traditional $I(0)$ approaches, such as adding 6 or 12 lags, typically fail to pick up the effects of variables like the interest rate and input costs.[23]

B. Variance Ratio Puzzle

Proposition 6 shows that, when sales are $I(1)$ (as they are in the data), the variance ratio is 1. Why then do empirical researchers typically obtain estimates of the variance ratio that are greater than 1? Our simulations of the calibrated model reveal that this occurs because of small sample bias in the variance ratio, as shown in Table 5. Our empirical results are based on a sample of 548 monthly observations, a relatively large number of observations (and long time

[23] A similar point has been made in the literature on fixed capital by Caballero (1994, 1999) and Schaller (2006).

span) for empirical work in macroeconomics. The simulations show that the variance ratio, calculated over this number of monthly observations, is about 1.02. For a sample of 2500 monthly observations, the variance ratio is still slightly greater than 1, about 1.01. It is only with a sample size of 5000 observations, roughly an order of magnitude larger than the usual sample size, that the variance ratio converges to 1.00. In our discussion of the Wen puzzle below, we explain why this small sample bias arises.

Table 5
Simulation Evidence on Small Sample Bias in the Variance Ratio

Sample Size (Number of Monthly Observations)	548	1000	2500	5000
Median Variance Ratio	1.02	1.01	1.01	1.00

This table is based on simulations of the calibrated model for different sample sizes. The second row reports the median variance ratio over 1000 repetitions of the simulation. Structural parameters are calibrated to the data as shown in Table 3.

C. Wen Puzzle and the Conditional Variance Ratio

i. Accounting for the Wen Puzzle

Wen (2005) distinguishes between the movements of production and sales at medium horizons (about 8-40 quarters) and short horizons (less than three quarters). At medium horizons, he finds that production is more volatile than sales. More surprisingly, he finds that production is less volatile than sales at short horizons. His empirical work shows that these stylized facts hold for the US, a number of other industrialized countries (Australia, Austria, Canada, Denmark, France, Finland, Great Britain, Japan, the Netherlands, and Switzerland), Europe as a whole, and the OECD as a whole. Wen (2005, p. 1533) argues that, "The stylized fact that production and inventories exhibit drastically different behaviors at the high- and low-cyclical frequencies offers a litmus test for [inventory] theories."[24]

Define the empirical conditional variance ratio as

[24] Wen's argument runs as follows. The short-horizon behavior of output and sales is consistent with production smoothing but not with stockout avoidance. The medium-horizon behavior of output and sales is consistent with stockout avoidance but not with production smoothing. The medium-horizon behavior of output and sales is consistent with increasing returns to scale (i.e., concavity of the production cost function), but the short-horizon behavior is not. Finally, if cost shocks are incorporated into a model with a production-smoothing motive, cost shocks can make output more variable than sales, but: 1) cost shocks make output more variable than sales at both short and medium horizons; or 2) when non-negativity constraints on inventories dominate, cost shocks have no effect on the correlation between inventory investment and sales. Wen (2005) thus concludes that none of the existing explanations for the variance ratio puzzle -- stockout avoidance, cost shocks, or increasing returns to scale -- can simultaneously account for the behavior of output and sales at both short and medium horizons.

$$CVR = \frac{\mathrm{Var}\left(\ln Y_t - \ln Y_{t-n}\right)}{\mathrm{Var}\left(\ln X_t - \ln X_{t-(n+1)}\right)}. \quad (36)$$

To motivate the empirical CVR, note that $\ln X_t$, conditional on $\ln X_{t-(n+1)}$, is the sum of subsequent sales shocks:

$$\ln X_t = \ln X_{t-(n+1)} + \sum_{j=1}^{n+1} u^X_{t-j}. \quad (37)$$

Thus, the variance of log sales, conditional on $\ln X_{t-(n+1)}$, is

$$\mathrm{Var}\left(\ln X_t \,\middle|\, \ln X_{t-(n+1)}\right) \equiv \mathrm{Var}\left(\ln X_t - \ln X_{t-(n+1)}\right) = \left(n+1\right)\sigma^2_X. \quad (38)$$

In our data, the empirical CVR shows the pattern documented by Wen (2005). As shown in the first row of Table 6, at short horizons, CVR <1; for example, for $n=1$ month, CVR = 0.62; for $n=2$ months, CVR = 0.71. At business cycle horizons, CVR >1; for example, for $n=50$ months, CVR = 1.02; for $n=70$ months, CVR = 1.01.

Using Proposition 5, we can calculate the CVR in the model as a function of the horizon n. The solid line in Figure 1 shows the CVR from our calibrated model which is less than 1 at short horizons and greater than 1 at business cycle horizons. The second row of Table 6 reports the numerical magnitudes, based on the calibrated model, for selected n. The model is successful in accounting for the Wen puzzle.

Table 6
Conditional Variance Ratio for Selected n

	Short Horizons		Business Cycle Horizons	
	$n=1$	$n=2$	$n=50$	$n=70$
Data	0.62	0.71	1.02	1.01
Model	0.56	0.74	1.02	1.02

The row labeled "Data" reports the empirical conditional variance ratio, which is defined in equation (36). The row labeled "Model" reports the conditional variance ratio from the model, which is calculated from equation (22) with the structural parameters calibrated to the data as shown in Table 3.

The intuition for the model's explanation of the Wen puzzle involves the balance between the production smoothing and stockout avoidance motives. The solid line in Figure 1 shows the CVR from the model when the structural parameters are calibrated to the data. If the production-smoothing motive were stronger (relative to the stockout avoidance motive), the CVR would be lower at all horizons. The dashed line in Figure 1 illustrates the case where θ_1 is substantially higher than the value in the data.[25] This makes the production smoothing motive stronger; in fact, so much stronger that the production smoothing motive dominates at all horizons and output varies less than sales at all finite horizons. The explanation for the Wen puzzle therefore runs as follows. The firm engages in both production smoothing and stockout avoidance. At short horizons, production smoothing dominates, so the variance ratio is less than 1. At business cycle horizons, stockout avoidance dominates, so the variance ratio is greater than 1.

Figure 1
Conditional Variance Ratio

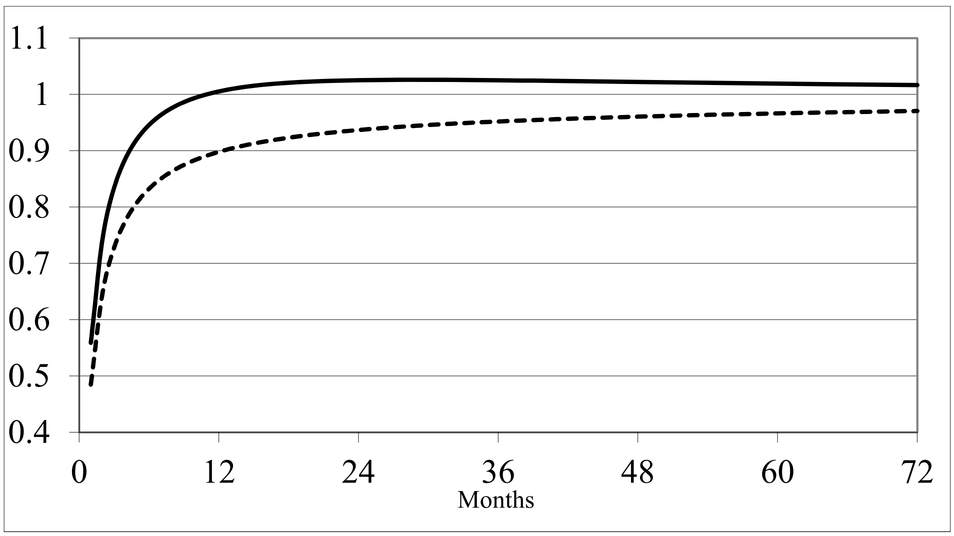

The solid line shows the conditional variance ratio (CVR) calculated from equation (22), where the structural parameters are calibrated using the cointegrating regression. The dashed line shows the conditional variance ratio for θ_1 equal to 2.5 times the calibrated value. A larger value of θ_1 implies greater convexity of the production cost function and thus a stronger production-smoothing motive. The horizontal axis shows the horizon (n) in months.

Our analysis of the Wen puzzle provides insight into why empirical researchers typically find that the (unconditional) variance ratio is greater than 1. In the model, starting at about the 9-month horizon, the conditional variance ratio is greater than 1. As $n \to \infty$ (i.e., as the horizon gets very long), the variance ratio approaches 1. But, in a finite sample, empirical researchers are effectively taking the average of the conditional variance ratio, which is above 1 for most values of n. The result is the finite sample bias that we document above.

---

[25] The relevant structural parameters, θ_1 and δ_2 , enter λ_1 and $\tilde{\Gamma}_X$ in equation (22) as a ratio. See equation (17) for ζ and how it enters λ_1, and equation (23-a) for $\tilde{\Gamma}_X$. Thus, setting θ_1 higher is equivalent to setting δ_2 proportionately lower.

ii. Reconciling the Conditional Variance Ratio with Slow Adjustment

The solid line in Figure 2 shows the path of sales in response to a one-time permanent sales shock. The line with open circles shows the response of production, based on the calibrated model. In the long run, production moves by the same amount as sales. In the short run, production moves more than sales. This short-run movement of production relative to sales is the issue highlighted by Blinder and Maccini (1991): How can we reconcile the fact that production varies more than sales with the slow adjustment of inventories, which, according to the standard analysis, is the result of a strong incentive to smooth production? If the production smoothing motive were the only force, production would change by less, in the short run, than the amount of a sales shock. Why, then, does production vary more than sales? The answer is the balance between production smoothing and stockout avoidance. The line with triangles in Figure 2 illustrates what would happen if the production smoothing motive were substantially stronger. In the long run, production would still rise by the same amount as sales, but, in the short run, production would rise by less than sales.

Figure 2
Response of Production to a Sales Shock

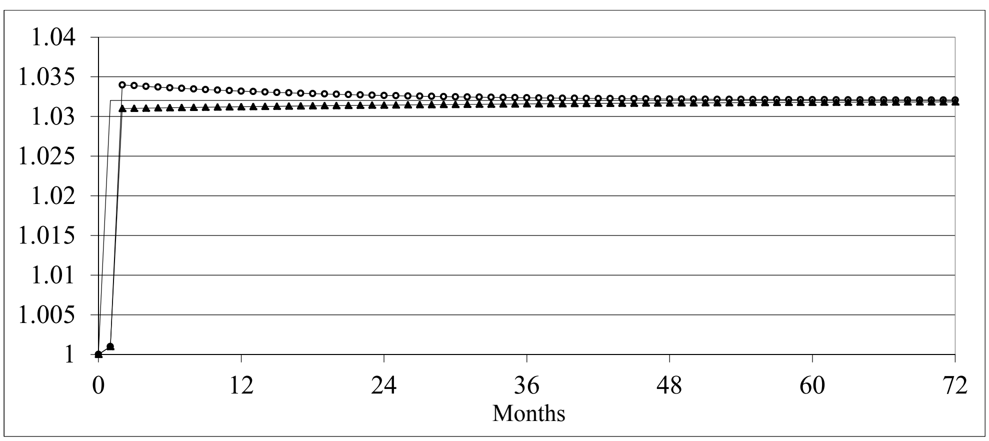

The solid line shows the path of sales in the wake of a one-time, one-standard-deviation permanent sales shock. The line with open circles shows the response of output in the model when the structural parameters are calibrated using the cointegrating regression for inventories. The line with solid triangles shows what the response of output would be if θ_1 were 2.5 times its calibrated value.

iii. Role of the Interest Rate

In a model with a stockout avoidance motive, the firm has a long-run desired stock of inventories, N_t^*, which is shown in equation (34). Proposition 2 gives the mathematical expression for Γ_X, which determines whether or not a positive sales shock increases N_t^*. Proposition 2 reveals that the interest rate plays a key role: The criterion for N_t^* to increase in response to a positive sales shock depends on the present value of the convexity of stockout avoidance costs, and the present value depends on $\bar{r}$. To the best of our knowledge, our work is the first to highlight the role of the real interest rate in understanding the relative variances of output and sales.

The role of the real interest rate in determining the CVR is illustrated in Figure 3. In our data, the mean annual real interest rate is 1.2%, which corresponds to a monthly real interest rate of 0.1% (or $\bar{r} = 0.001$). This implies that the production-smoothing motive dominates the CVR at short horizons, as shown by the solid line. If the real interest rate were lower, the stockout avoidance motive would be even more important and the CVR would be greater than 1 at still shorter horizons, as shown by the line with open circles. If the real interest rate were higher, the stockout avoidance motive would be less important. For example, for a mean annual real interest rate of 5.0%, production smoothing would dominate at all horizons and the CVR would never be greater than 1, as illustrated by the line with triangles.

Figure 3
The Effect of the Real Interest Rate on the Conditional Variance Ratio

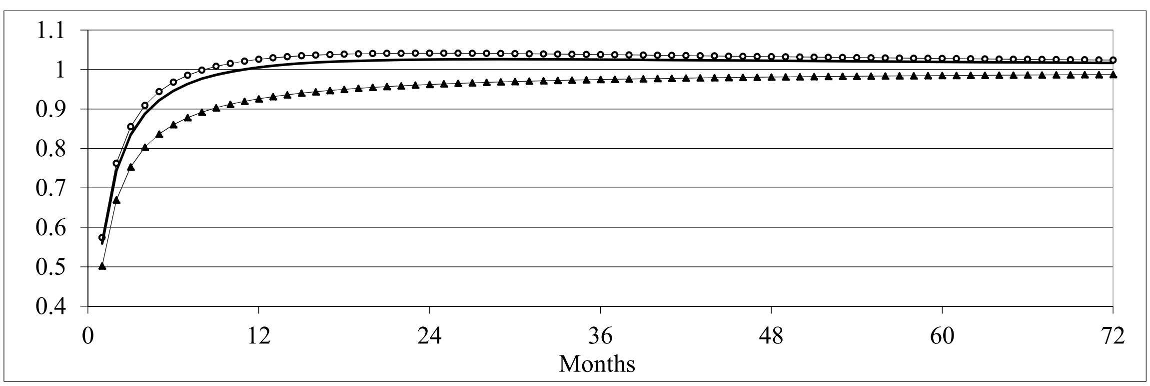

The solid line shows the conditional variance ratio calculated from equation (22) where the annual mean real interest rate is equal to our estimate of 1.2%. The line with open circles shows the conditional variance ratio for a mean real interest rate of 0.3% and the line with triangles shows the conditional variance ratio for a mean real interest rate of 5%.

D. Input Cost Puzzle

An increase in input costs should cause firms to reduce their inventories. However, in the past it has been difficult to find evidence of a significant relationship between inventories and observable costs. Proposition 8 shows that real input costs appear in the cointegrating regression for inventories with coefficient b_W. Table 2 reports that the estimated value of b_W is -0.753, which is negative, as predicted by Proposition 9. The t-statistic on b_W is -5.244. Thus, the cointegrating regression provides significant evidence that observable costs affect inventories.

Why does the cointegrating regression provide stronger evidence than previous studies? Each period, the firm is hit by sales shocks (an unanticipated change in demand) and production shocks (e.g., due to a supply chain disruption), represented in the model by u_t^X and u_t^Y, respectively. These can immediately move inventories by a large amount, with no change in input costs. In contrast, changes in input costs affect N^*, and the firm adjusts very slowly towards N^*. $I(0)$ econometric techniques tend to emphasize short-run fluctuations in the variables, where the relationship between inventories and input costs is very noisy. Without the introduction of lagged variables, $I(0)$ techniques only capture the contemporaneous relationship. When adjustment is slow, very long lags may be needed, and there is no guarantee that the lag structure will be stable over time. In contrast, $I(1)$ econometric techniques (specifically, the cointegrating regression) emphasize the relationship between inventories and the variables that determine inventories in the long run.

In the literature, cost shocks have been a leading potential explanation for the variance ratio puzzle.[26] In our model, as in other models, production-cost shocks tend to increase the conditional variance ratio. However, numerical results based on equation (22) show that the contribution of production cost shocks to the conditional variance ratio is negligible – less than 1%.

Our analysis explains how cost shocks can have a statistically significant effect on inventories in the long run but little effect on the conditional variance ratio. In a model where the variables are $I(0)$, a shock that reduces marginal production costs today tends to induce intertemporal substitution of production from the future to the present, leading production to change without a change in sales. As Table 1 shows, W_t is $I(1)$, so input cost shocks are permanent. But, if the shock is permanent, there is no reason for intertemporal substitution. The

[26] In their survey, for example, Ramey and West (1999) discuss highly persistent shocks to production cost as an explanation for the slow adjustment and variance ratio puzzles.

shock still has an effect on production because it changes N_t^*, but this effect is slow and subtle (hard to even detect with $I(0)$ econometric techniques, as discussed above). This is why input cost shocks have little effect on the conditional variance ratio.

VIII. MONETARY POLICY PUZZLES

A. Monetary Policy Shocks

To identify monetary policy shocks we follow Bernanke and Mihov (1998) and estimate a vector autoregression whose variables are divided into a policy block and a non-policy block. In our version of the Bernanke-Mihov VAR, the non-policy block consists of the natural logarithms of real sales ($\ln X_t$), the GDP deflator, real input prices ($\ln W_t$), and real inventories ($\ln N_t$).[27] Our policy block, which is the same as Bernanke and Mihov's, consists of total reserves, non-borrowed reserves, and the Fed funds rate and is restricted using plausible assumptions about the market for bank reserves. Details of this Bernanke-Mihov VAR are provided in Appendix D.

Figure 4
Empirical Response of the Probabilities to a Stimulative Monetary Policy Shock

A. Probability of Low-Interest-Rate Regime **B. Probability of High-Interest-Rate Regime**

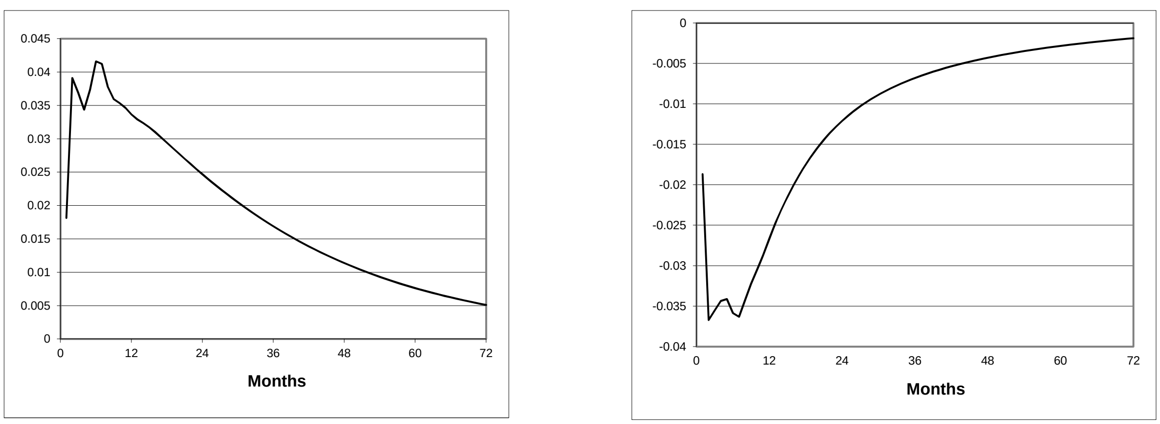

The solid lines in Figures 4-A and 4-B present the impulse response function of π_1 (the probability of the low-real-interest-rate regime, as perceived by the firm) and π_3 (the probability of the high-real-interest-rate regime, as perceived by the firm), respectively, to a one-standard-deviation stimulative monetary policy shock. The horizontal axis shows time in months.

---

[27] We found that the inclusion of input prices was sufficient to address the price puzzle and so do not add a commodity price index.

Having obtained the monetary policy shocks from the Bernanke-Mihov VAR, we then estimate a three-variable VAR with the monetary policy shocks, π_1, and π_3.[28] Figure 4 shows the impulse response functions of π_1 and π_3 to a one-standard-deviation easing of monetary policy. As Figure 4-A shows, easing monetary policy increases the probability of the low-interest-rate state, with the peak response occurring about six months after the shock. Results are similar for π_3. The ergodic probability of the high-interest-rate state is about 0.19. The peak decline in π_3 is 0.036, which represents a decrease of about 19% in the likelihood of the high interest rate regime. The effect of monetary policy on π_1 and π_3 is quite persistent, with more than half the peak effect on π_1, for example, still present two years after the shock.

B. The Mechanism Puzzle

Previous empirical studies have found little evidence that the interest rate affects inventories.[29] If the interest rate doesn't affect inventories, how does monetary policy influence inventories?[30] If the interest rate does affect inventories, why have more than 40 years of empirical studies failed to find the relationship?

In our theoretical model, the real interest rate is subject to both transitory and persistent shocks. Purely transitory shocks have little effect on inventories, but firms do react to shocks that may be persistent. In the past, empirical inventory research has primarily used $I(0)$ techniques.[31] These techniques tend to emphasize high-frequency movements in the data, where there is much transitory variation in the interest rate without corresponding variation in inventories – and much transitory variation in inventories (due to their role in buffering sales shocks) without corresponding variation in the interest rate.

Table 2 reports our estimates of the cointegrating regression. The key coefficients for the mechanism puzzle are those on π_1 (the probability of the high-interest-rate regime) and π_3 (the

[28] We use six lags of each variable. We do not include the probabilities in the Bernanke-Mihov VAR because there is too much collinearity between the probabilities and the interest rate.

[29] See Blinder and Maccini (1991, page 82). One exception is Maccini, Moore, and Schaller (2004), who also use $I(1)$ econometrics. In contrast to the current paper, they do not address the sign and timing puzzles.

[30] VAR-based studies that find that monetary policy shocks affect inventories include Bernanke and Gertler (1995), Christiano, Eichenbaum, and Evans (1996), and Jung and Yun (2011).

[31] There are a few exceptions, including Granger and Lee (1989), Kashyap and Wilcox (1993), and Rossana (1993, 1998), but none of these papers estimate the effect of the real interest rate on inventories. Rossana (1993) comes closest by providing separate point estimates of the effects of the nominal interest rate and inflation.

probability of the low-interest-rate regime). Theory predicts that the coefficient on π_1 should be positive and the coefficient on π_3 should be negative. The data confirm both of these theoretical predictions. The coefficients on both π_1 and π_3 are significantly different from zero.

The decision rule for the firm's choice of inventories, equation (16), shows that monetary policy shocks can affect inventories through their effects on sales, input costs, π_1 and π_3. Having calibrated our model to the cointegrating regression, we can use the calibrated decision rule to measure the economic importance of the effect that interest-rate movements have on inventories. In general, the previous literature has treated the interest rate as constant and so has been unable to measure the effect of interest-rate movements.

We define the opportunity-cost effect as the change in inventories that results from a monetary policy shock, holding sales and input costs constant. To measure this effect we generate the theoretical response of inventories to a monetary policy shock. We first use our Bernanke-Mihov VAR to find the response of sales and input costs to a one-standard-deviation stimulative monetary policy shock. We then use the response of sales and input costs together with the response of π_1 and π_3 (as shown in Figure 4) in our calibrated decision rule to calculate the theoretical response of log inventories to the monetary policy shock. Using this theoretical response we can measure the peak effect of a monetary policy shock on log inventories. Repeating this exercise, but holding sales and input costs constant, we find that the opportunity-cost effect is equal to 78% of the peak effect. Thus, although the opportunity-cost effect has been extremely difficult to detect using $I(0)$ econometric techniques, our calibrated model suggests that it is economically important.

C. The Sign Puzzle

Stimulative monetary policy reduces the interest rate and should, therefore, **increase** inventories. However, VAR studies find that the short-term effect of stimulative monetary policy is to **decrease** inventories. This is the sign puzzle. To verify that the sign puzzle exists in our data, we use our Bernanke-Mihov VAR to calculate the empirical response of inventories to a monetary policy shock. The responses of the Fed funds rate and inventories to a one-standard-deviation stimulative monetary policy shock are shown in Figure 5. As found in other studies,

Figure 5
Empirical Responses to a Stimulative Monetary Policy Shock

A. Federal Funds Rate **B. Inventories**

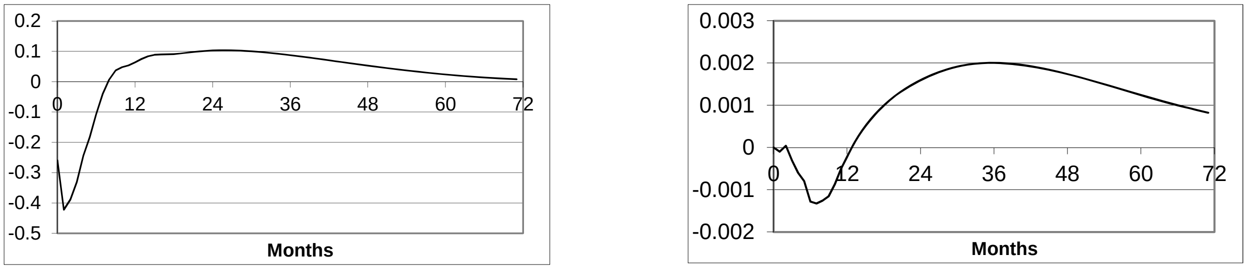

Figures 5-A and 5-B present the empirical impulse response function of the Fed funds rate and inventories, respectively, to a one-standard-deviation stimulative monetary policy shock. The horizontal axis shows time in months.

the Bernanke-Mihov VAR estimated with our data shows that the initial response to a stimulative monetary policy shock is a decline in both the Fed Funds rate and inventories.

Does our model generate this negative short-term decline in inventories for a stimulative monetary policy shock? As explained in our discussion of the mechanism puzzle above, we use the empirical response of sales, input costs, π_1, and π_3 in our calibrated decision rule, equation

Figure 6
Theoretical Response of Inventories to a Stimulative Monetary Policy Shock and the Effect of the Convexity of the Production Cost Function

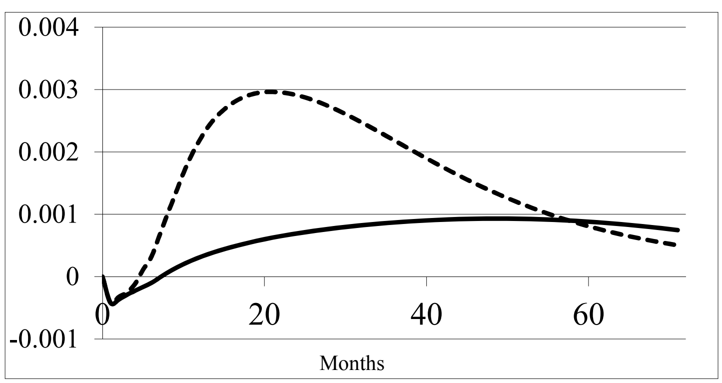

The solid line displays the theoretical response of inventories to a one-standard-deviation stimulative monetary policy shock, based on the model presented in Section II, calibrating the structural parameters using the cointegrating regression, as shown in Table 3. The dashed line shows the theoretical response of inventories based on setting θ_1 (the parameter that controls the convexity of the production cost function) equal to 0.5 times the value obtained when the parameters are calibrated using the cointegrating regression. The horizontal axis shows time in months.

(16), to find the theoretical response of log inventories to a monetary policy shock.[32] The solid line in Figure 6 presents this theoretical response of inventories to a stimulative monetary policy shock. As the figure shows, the initial response is for inventories to decline. The key to understanding our model's success in matching the empirical sign puzzle is the role of inventories in buffering demand (sales) shocks. Sales rise in the wake of a stimulative monetary policy shock. Production does not respond immediately, so inventories fall as they buffer the positive sales shock.

D. The Timing Puzzle

The transitory effect of a monetary policy shock on the Fed funds rate is shown by the empirical impulse response function in Figure 5-A. Within eight months, the Fed funds rate returns to its pre-shock level. It is only many months later that inventories rise above their pre-shock level, as shown in Figure 5-B. The peak effect of the monetary policy shock on inventories occurs years after the shock.[33] This is the timing puzzle.

Regime switching and learning provide part of the explanation for the timing puzzle. Because of learning, the Bayesian probabilities of being in a given interest rate régime respond slowly to a change in the interest rate. This can be seen in Figure 4-A, where more than one-third of the effect of the monetary policy shock on π_1 is still present three years after the shock. Simulations of the calibrated model show that learning delays the response of inventories by about one quarter (three months).

Production smoothing also plays a role. An interest rate shock changes the desired long-run inventory level. However, changing output away from the usual level is expensive because of the convexity of the cost function. If firms recognize that the interest rate shock is transitory, they will adjust output and the stock of inventories little, if at all. Because firms are reluctant to adjust output, the change in the stock of inventories is delayed.

The convexity of the cost function is measured by the parameter θ_1. In Figure 6, we illustrate the effect of changes in θ_1 on the theoretical impulse response function for inventories. If we set θ_1 equal to half the value implied by the cointegrating regression estimates, the peak effect on inventories occurs twenty-eight months earlier (the dashed line in Figure 6). Figure 6

[32] Since the monetary policy shock is by definition unanticipated we assume that the initial increase in sales is also unanticipated.

[33] This is also documented in Christiano, Eichenbaum, and Evans (1996) and Jung and Yun (2005).

illustrates another interesting point. In the inventory literature, it has been very hard to pin down the convexity of the cost function. In their survey paper, Ramey and West (1999) report a wide range of estimates. Using the cointegrating regression to calibrate θ_1, we obtain a value of θ_1 that leads to a theoretical impulse response function that is similar to the empirical impulse response function. As Figure 6 illustrates, using a value of θ_1 that is 50% smaller leads to a theoretical impulse response function that no longer matches the empirical response: the peak response of inventories is too early and too large. Using a value of θ_1 that is much larger than the calibrated value (e.g., 50% larger) leads to a theoretical impulse response function that no longer even qualitatively resembles the empirical response. Methodologically, this suggests that good estimates of a well-specified cointegrating regression may provide a better technique for calibrating model parameters. Economically, it narrows the range of plausible estimates of the convexity of the cost function.[34]

IX. SUMMARY AND CONCLUSIONS

Less priority has been given to research on inventories in recent years than in the preceding decades. An important reason is probably the belief that inventories tend to cushion shocks (particularly, demand shocks). Since macroeconomists have been searching for mechanisms that amplify shocks, inventories have not seemed like a particularly promising research avenue.

We begin with an empirical fact: Sales are $I(1)$. We build a new model of inventories in which this fact plays a central role. Although our model retains some of the familiar elements of the long-established linear-quadratic model (such as production smoothing and stockout avoidance), we obtain startling results.

Our most startling result is that inventories do not cushion demand shocks. In the long run, Proposition 6 states that production moves as much as sales. Based on Proposition 5, we find that, at business cycle horizons, production moves more than sales.

We derive three propositions from the model that involve empirical predictions. Proposition 8 states that inventories are cointegrated with sales and the other variables that determine long-run inventories in our model. The data support this prediction. The statistical

---

[34] For example, it rules out the possibility of a concave cost function, an explanation suggested by Ramey (1991) for the variance ratio puzzle.

evidence is strong. (The p-value is 0.001.) Proposition 9-B states that an increase in input costs decreases inventories in the long run. The relevant empirical coefficient (b_W) is negative and strongly statistically significant. Proposition 9-C states that a higher probability of the low-interest-rate regime increases inventories and a higher probability of the high-interest-rate regime reduces inventories in the long run. The data support both predictions, and the relevant coefficients are strongly statistically significant.

In addition to passing these direct empirical tests, our model explains the four traditional inventory puzzles – the variance ratio, slow adjustment, Wen, and input cost puzzles. Moreover, our model accounts for three important puzzles about the effect of monetary policy on inventories.

The monetary policy puzzles involve the dynamic response of inventories to a monetary policy shock. Over the past two decades, an important challenge for macroeconomic models has been to account for the hump-shaped response of many aggregate variables to a monetary policy shock. A series of papers have shown that inventories display a more complex "double-hump" response. In reaction to a stimulative monetary policy shock, inventories decline in the first few months, rise until they reach a peak about three years after the shock, and then decline again. The initial decline is the sign puzzle: Low interest rates are associated with low inventories, instead of the reverse. The subsequent rise is the timing puzzle: Inventories begin to rise after the fall in the interest rate has largely disappeared. Our calibrated model is successful in capturing the "double-hump" dynamic response of inventories to a monetary policy shock.

Two further points should be emphasized. First, we do not allow ourselves any free parameters. The key structural parameters are calibrated using the cointegrating regression derived from the model. There are no free parameters that we can use to match empirical moments.

Second, in the previous papers that attempt to explain inventory puzzles, the objective has been to explain "static moments" such as the relative variance of production and sales or the correlation of inventory investment with output. In this paper, we set the bar higher: We explain both static moments and the dynamic response of inventories.

REFERENCES

Bernanke, Ben S. and Mark Gertler (1995), "Inside the Black Box: The Credit Channel of Monetary Policy Transmission", *Journal of Economic Perspectives*, Winter, 9(1), pp. 27-48.

Bernanke, Ben S. and Ilian Mihov (1998), "Measuring Monetary Policy", *Quarterly Journal of Economics*, 113(3), pp. 869-902.

Bils, Mark and James A. Kahn (2000), "What Inventory Behavior Tells Us about Business Cycles", American Economic Review, 90(3), pp. 458-481.

Blanchard, Olivier (1983), "The Production and Inventory Behavior of the American Automobile Industry", *Journal of Political Economy*, 91(3), pp. 365-400.

Blinder, Alan S. (1986-a), "More on the Speed of Adjustment in Inventory Models", *Journal of Money, Credit, and Banking*, 18(3), pp. 355-365.

Blinder, Alan S. (1986-b), "Can the Production Smoothing Model of Inventory Behavior be Saved?", *Quarterly Journal of Economics*, 101(3), pp. 431-453.

Blinder, Alan S. and Louis J. Maccini (1991), "Taking Stock: A Critical Assessment of Recent Research on Inventories", *Journal of Economic Perspectives*, Winter, 5(1), pp. 73-96.

Caballero, Ricardo J. (1994), "Small Sample Bias and Adjustment Costs," *Review of Economics and Statistics*, 76(1), pp. 52-58.

Caballero, Ricardo J. (1999), "Aggregate Investment," in John B. Taylor and Michael Woodford (eds.), *Handbook Of Macroeconomics*, vol. 1B, Amsterdam: North Holland, pp. 813-862.

Chang, Yongsung, Andreas Hornstein and Pierre-Daniel Sarte, "On the Employment Effects of Productivity Shocks: The Role of Inventories, Demand Elasticity and Sticky Prices", *Journal of Monetary Economics,* v 56, pp. 328-343.

Christiano, Lawrence J. and Martin Eichenbaum (1989), "Temporal Aggregation and the Stock Adjustment Model of Inventories", in T. Kollintzas (ed), *The Rational Expectations Inventory Model*, Springer-Verlag, New York, pp. 70-109.

Christiano, Lawrence J., Martin Eichenbaum and Charles Evans (1996), The Effects of Monetary Policy Shocks: Evidence form the Flow of Funds", *The Review of Economics and Statistics,* 78 (1), pp. 16-34.

Coen-Pirani, Daniele (2004), "Markups, Aggregation, and Inventory Investment", *American Economic Review*, 94(5), pp. 1328-1353.

Eichenbaum, Martin (1989), Some Empirical Evidence on the Production Level and Production Cost Smoothing Models of Inventory investment", *American Economic Review*, 79(4), pp. 853-864.

Garcia, Rene and Pierre Perron (1996), "An Analysis of the Real Interest Rate under Regime Shifts", *Review of Economics and Statistics*, 78(1), pp. 111-125

Granger, C.W.J., and T.H. Lee (1989), "Investigation of Production, Sales and Inventory Relationships Using Multicointegration and Non-Symmetric Error Correction Models," *Journal of Applied Economics* 4, pp. S145-S159.

Hamilton, James D. (1989), "A New Approach to the Economic Analysis of Nonstationary Time Series and the Business Cycle", *Econometrica*, 57(2), pp. 357-84.

Hamilton, James D. (1994), *Time Series Analysis*, Princeton: Princeton University Press.

Hamilton, James D. (2002), "On the Interpretation of Cointegration in the Linear- Quadratic Inventory Model", *Journal of Economic Dynamics and Control*, 26(12), pp. 2037-49.

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

Humphreys, Brad, Louis Maccini, and Scott Schuh (2001), "Input and Output Inventories," *Journal of Monetary Economics*, v 47, pp. 347-375.

Iacoviello, M., F. Schiantarelli, and Scott Schuh (2007), "Input and Output Inventories in General Equilibrium", Working Paper, Boston College.

Jung, Yongseung and Tack Yun (2011), "Inventory and Dynamic Effects of Monetary Policy Shocks", Working Paper.

Jung, Yongseung and Tack Yun (2012), Inventory Investment and the Empirical Phillips Curve", Working Paper.

Kahn, James (1987), "Inventories and the Volatility of Production", *American Economic Review*, 77(4), pp. 667-679.

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

Kashyap, Anil K. and David W. Wilcox (1993), "Production and Inventory Control at the General Motors Corporation in the 1920s and 1930s", *American Economic Review*, 83(3), pp. 383-401.

Khan, Aubhik and Julia Thomas (2007), "Inventories and the Business Cycle: An Equilibrium

Analysis of (S,s) Policies", *American Economic Review*, 97, 1165- 1188.

Lovell, Michael C. (1961), "Manufacturer's Inventories, Sales Expectations, and the Acceleration Principle", *Econometrica*, 29(3), pp. 293-314.

Maccini, Louis J., Bartholomew J. Moore, and Huntley Schaller (2004), "The Interest Rate, Learning, and Inventory Investment", *American Economic Review*, 94(5), pp.1303-1327.

Maccini, Louis J. and Adrian Pagan (2009), "Inventories, Fluctuations and Business Cycles", forthcoming in *Macroeconomic Dynamics*.

Maccini, Louis J. and Robert J. Rossana (1984), "Joint Production, Quasi-Fixed Factors of Production, and Investment in Finished Goods Inventories", *Journal of Money, Credit, and Banking*, 16, pp. 218-236.

Ramey, Valerie A. (1989, "Inventories as Factors of Production and Economic Fluctuations", *American Economic Review*, 79(3), pp. 338-54.

Ramey, Valerie A. (1991), "Nonconvex Costs and the Behavior of Inventories", *Journal of Political Economy*, 99(2), pp. 306-334.

Ramey, Valerie A. and Daniel Vine (2006), "Tracking the Source of the Decline in GDP Volatility: An Analysis of the Automobile Industry", *American Economic Review,* 96, 1876-1889.

Ramey, Valerie A. and Kenneth D. West (1999), "Inventories", in John B. Taylor and Michael Woodford (eds.), *Handbook of Macroeconomics*, Volume 1B, Amsterdam: North-Holland, pp. 863-923.

Rossana, Robert J. (1993), "The Long-Run Implications of the Production Smoothing Model of Inventories: An Empirical Test", *Journal of Applied Econometrics,* 8, pp. 295-306.

Rossana, Robert J. (1998), "Structural Instability and the Production Smoothing Model of Inventories", *Journal of Business and Economic Statistics,* 16(2), pp. 206-215.

Schaller, Huntley, (2006), "Estimating the long-run user cost elasticity," *Journal of Monetary Economics*, 53(4), 725-736.

Seitz, Helmut (1988), "Still More on the Speed of Adjustment in Inventory Models: A Lesson in Aggregation", University of Manheim Discussion Paper.

Stock, James H., and Mark W. Watson (1993), "A Simple Estimator of Cointegrating Vectors in Higher Order Integrated Systems", *Econometrica,* 61(4), pp. 783-820.

Wang, Peng-fei and Yi Wen (2009), "Inventory Accelerator in General Equilibrium", Working Paper, Federal Reserve Bank of St. Louis.

Wen, Yi (2005), "Understanding the Inventory Cycle", *Journal of Monetary Economics*, 52, 1533-1555.

Wen, Yi (2008), "Input and Output Inventories", Working Paper, Federal Reserve Bank of St. Louis.

West, Kenneth D. (1986), "A Variance Bound Test of the Linear-Quadratic Inventory Model", *Journal of Political Economy*, 94(2), pp. 374-401.

West, Kenneth D. (1990), "The Sources of Fluctuations in Aggregate inventories and GNP", *Quarterly Journal of Economics*, 115(6), pp. 939-972

APPENDIX A. Data and Sources

The real inventory and shipments data are produced by the Bureau of Economic Analysis and are derived from the Census Bureau's Manufacturers' Shipments, Inventories, and Orders survey. They are seasonally adjusted, expressed in millions of 1996 chained dollars, and cover the period 1959:01-2004:08 (due to issues of data availability). An implicit price index for shipments is obtained from the ratio of nominal shipments to real shipments.

The observable cost shocks include average hourly earnings of production and nonsupervisory workers for the nominal wage rate; materials price indexes constructed from two-digit producer price indexes and input-output relationships (See Humphreys, Maccini and Schuh (2001) for details); and crude oil prices as a measure of energy prices. Nominal input prices were converted to real values using the shipments deflator. The nominal interest rate is the 3-month Treasury bill rate. Real rates were computed by deducting the three-month inflation rate calculated by the Consumer Price Index.

The Fed funds rate and reserves data are from the Federal Reserve Bank of St. Louis's FRED II database. The monthly interpolation of the GDP Deflator uses the seasonally adjusted quarterly deflator from FRED II (GDPCTPI) and the seasonally adjusted monthly producer price indices for crude materials, capital equipment, finished goods, and intermediate materials and supplies (PWCMSA, PWFPSA, PWFSA, and PWIMSA, respectively) from DRI Basic Economics.