

Inventory Behavior with Permanent Sales Shocks

Louis J. Maccini
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
(Corresponding author)
Department of Economics
Carleton University
1125 Colonel By Drive
Ottawa, ON K1S 5B6
(613) 520-3751
huntley_schaller@carleton.ca

May 16, 2013

We thank Martin Eichenbaum, James Kahn and Adrian Pagan for comments on an earlier draft of the paper; seminar participants at the American Economic Association, Bank of Canada, Econometric Society, Fordham University, Institute for Advanced Studies (Vienna), International Society for Inventory Research, Society for Computational Economics, and 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 $I(1)$. Starting from this fact, we derive 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 that of sales. We explain – analytically and intuitively – 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

May 16, 2013

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. We refer to these gaps as the *Traditional Inventory Puzzles*. 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. We refer to these contradictions as the *Monetary Policy Puzzles*.

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 traditional inventory puzzle is why production varies more than sales in the data – a puzzle which we refer to as the *Variance Ratio Puzzle*. A variety of theoretical explanations have been proposed. Our analysis implies that all of these explanations are irrelevant, at least in the long run.

In the data, sales are $I(1)$. Based on this fact, we obtain three startling new 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. This implies that inventories amplify sales shocks during business cycles, rather than dampening shocks as production smoothing would imply.³

The intuition for the first result – that the variance of production is equal to the variance of sales in the long run – flows directly from the inventory accumulation identity, once we take into account that sales are $I(1)$. If sales, output, and inventories are all $I(1)$, the identity implies

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.

that sales and output are cointegrated with cointegrating vector $(1, -1)$, which means that the variances of output and sales are the same in the long run. Because the proof flows from the inventory accumulation identity, it does not rely on production smoothing, stockout avoidance, or input costs. To the best of our knowledge, these two results are completely new in the inventory literature. We elaborate on these two points and their implications in the paper.

If the variances of sales and output are the same in the long run, why do empirical studies typically find that production varies more than sales? Here, the answer involves much more mathematics (but simple mathematics!). To analytically address this issue, we introduce a new concept – the conditional variance ratio – which we define as the variance of output divided by the variance of sales, where both variances are conditional on the level of the relevant variable at a fixed time in the past. Using our analytical result for the conditional variance ratio, we show that: 1) standard estimates of the variance ratio are biased; and 2) the direction of the bias depends on the relative strength of the production smoothing and stockout avoidance motives. Again, these appear to be new results.

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. We refer to these observations as the *Wen Puzzle*. Wen argues that these stylized facts constitute a "litmus test" for inventory theory and concludes that none of the existing accounts of the relative variability of output and sales – whether based on production smoothing, stockout avoidance, or increasing returns to scale – can account for the behavior of output and sales at both short and medium horizons.

Our analytical result for the conditional variance ratio provides the foundation for a complete solution to the Wen puzzle. The intuition flows from the short-run buffering role of inventories. A standard – and highly realistic – assumption in the inventory literature is that firms must set production before they know the state of current demand. As a result, a positive sales shock is initially met by running down inventories. Output is initially unchanged, which tends to push down the estimated variance of output over a short horizon. Output soon begins to adjust, but the conditional variance of output is initially smaller than the conditional variance of sales (where both are conditional on their levels before the shock). This is why production is measured as being less volatile than sales at short horizons.

Our analytical results show that, if stockout avoidance dominates production smoothing, the conditional variance ratio will be greater than 1 at medium horizons. Intuitively, the firm wants to raise output both to: 1) meet the new permanently higher level of demand; and 2) replenish the inventories that are run down right after the shock. This leads output to respond more than one-for-one to the sales shock, so the conditional variance ratio rises above 1 over medium horizons. Over the sample size typically available for empirical work, this effect dominates the low short-horizon conditional variance ratio and leads to the upward bias in the measured variance ratio.

Our analysis of the conditional variance ratio provides an important insight: Whether stockout avoidance dominates production smoothing depends on the mean real interest rate. Intuitively, 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. The higher the real interest rate, the smaller the weight the firm puts on stockout avoidance.

In view of the fact that sales are $I(1)$, we use the cointegrating regression that links inventories to sales, input costs, and the interest rate to calibrate the structural parameters. Our paper is the first to use a cointegrating regression to calibrate the structural parameters of any economic model (as far as we know). This works beautifully. When we plug the calibrated structural parameters into the equations derived from the model, we obtain the stylized facts that constitute the puzzles.

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)]". This is the *Slow Adjustment Puzzle*. Our analytical results show that an increase in the convexity of production costs, relative to the convexity of stockout avoidance costs, decreases the speed of adjustment. Intuitively, an increase in the convexity of production costs increases the incentive to smooth production. This makes for slower adjustment.

Our analysis implies that a key coefficient in the cointegrating regression will be positive if stockout avoidance dominates production smoothing (in the way that emerges from our equations; i.e., with the stockout avoidance motive weighted by the reciprocal of the mean real

interest rate). Our estimate of this parameter, which reflects the effect of a sales shock on desired long-run inventories, is positive. This is an additional piece of evidence that ties together with the results we obtain when we plug the calibrated structural parameters into the equations derived from the model.

We also provide solutions for three important monetary puzzles that involve the relationship between monetary policy and inventories. One such puzzle is the *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. 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 *Sign* and *Timing 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*: Lower interest rates are associated with lower inventories, instead of the reverse. The subsequent rise is the *Timing puzzle*: Inventories begin to rise after the fall in interest rates has largely disappeared.

We are successful in capturing the “double-hump” dynamic response of inventories to a monetary policy shock. 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 does not immediately raise production due to strong

production smoothing incentives, 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.

We make a conscious decision to use a partial equilibrium inventory model, partly to connect with a long theoretical and empirical inventory literature, dating back at least to Holt, Modigliani, Muth, and Simon (1960). More importantly, a partial equilibrium framework allows us to derive analytical results that bring out a series of new insights about the forces driving inventories. These analytical results lead to a rich new intuitive understanding of inventories. A DSGE model would add complexities that make it difficult to obtain analytical results and would tend to muddy the intuition. Furthermore, a DSGE model would require the specification and estimation of relationships for the household sector, the services sector, monetary policy rules, etc. Any failure to correctly specify all those relationships could produce biased estimates of the parameters of the model, including potentially those that capture the degree of production smoothing and stockout avoidance, which are critical to our ability to understand the empirical puzzles. Of course, there are disadvantages to a partial equilibrium approach. It does not account for feedback between different agents through equilibria in different markets. In general equilibrium, sales would be endogenous, a point that could be particularly important in a nominal model (e.g., a New Keynesian model) in which variations in the markup play a central role. We therefore view the current paper as a first step in better understanding inventory behavior, much like partial equilibrium models of non-convexities in adjustment costs (e.g., Bertola and Caballero (1994)) that have subsequently led to a rich vein of research in DSGE models (including, e.g., Khan and Thomas (2007)). If researchers involved in the DSGE literature are intrigued by the way in which $I(1)$ sales shocks lead to an understanding of inventories that accounts for so many inventory puzzles, perhaps DSGE models based on $I(1)$ sales shocks will become an active research area in the future.

The paper is organized as follows. Section II introduces a variation on the traditional linear-quadratic model that is better suited to the case of $I(1)$ sales. Section III focuses on the cointegrating regression. Section IV explains our solutions to four traditional inventory puzzles – variance ratio, Wen, slow adjustment, and input cost. In addition, Section IV also presents three startling results regarding the variance ratio (i.e., the variance of output divided by the variance of sales) that emerge from the model. Section V examines how the model accounts for the three monetary policy puzzles. Section VI 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)⁴. In this paper, we use 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.⁵ Real costs per period consist of production costs and inventory holding costs. Production costs, PC_t , are defined as

$$PC_t = A_t Y_t^{\theta_1} W_t^{\theta_2} \quad (1)$$

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

⁴ 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).

⁵ 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.

state of technology, fixed factors of production, and the organizational structure of the firm.⁶

Observe that average production costs, J_t , are

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

and marginal production costs are $\theta_1 J_t$. Because A_t is a catchall for a variety of factors, some of which do not lend themselves readily to measurement, our specification does not impose any assumptions on (or even lead to any clear implications for) the cyclicity of marginal cost.

If rising marginal cost were the only reason for holding inventories, the standard analysis suggests that output would vary less than sales. Beginning with Holt, Muth, Modigliani and Simon (1960), many models have included a quadratic cost of deviating from a target level of inventories as a way of capturing the stockout avoidance motive and inducing output to vary more than sales. An important study by Kahn (1987) shows that explicitly modeling the non-negativity constraint on inventories is a way of introducing a stockout avoidance motive into the model and rationalizing a target level of inventories. An alternative way of incorporating a stockout avoidance motive is to make sales a function of inventories. Bils and Kahn (2000) illustrate the advantages of this approach when the objective is to analyze the cyclicity of marginal cost and the markup.⁷ Our primary objective, however, is to understand the relative variability of output and sales and how these can be reconciled with the slow adjustment of inventories. A particular focus is the Wen puzzle, because Wen (2005) argues that his findings cannot be explained by production smoothing, stockout avoidance, cost shocks, or increasing

⁶ 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 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.

⁷ Yet another way of introducing a stockout avoidance motive, which goes very far back in the inventory literature (e.g., Baumol (1952) and Tobin (1956)), is fixed costs. Khan and Thomas (2007) make an important contribution by incorporating fixed costs into a general equilibrium model of inventories. Unfortunately, if one wants to obtain analytical results that clarify how the interaction between production smoothing and stockout avoidance explains the major inventory puzzles, the assumption of fixed costs makes the analysis intractable.

returns to scale. It turns out that the economic mechanisms that explain inventory puzzles can be most clearly understood if we follow the more traditional literature and model stockout avoidance through the cost function. We therefore specify inventory holding costs as

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

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. 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.⁹

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 These two components underlie the rationale for the quadratic inventory holding costs in the standard linear-quadratic model. Observe that (3) implies a "target stock" of finished goods inventories that minimizes finished goods holding costs. The target stock is $N_t^{TS} = (-\delta_3 / \delta_1 \delta_2)^{\frac{1}{\delta_2-1}} X_t$ so that the implied stock is proportional to sales. As emphasized by Bils and Kahn (2000), the target inventory/sales ratio in this class of model is a constant, instead of proportional to a variable markup as in their model. A possibly interesting variation on our approach would be to incorporate this aspect of their framework. 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.

Hamilton (2002) points out that there is an awkward feature of earlier models that have motivated a cointegrating regression for inventories: Marginal production and inventory holding costs are non-stationary. He argues that this is economically implausible and suggests a solution that is consistent with the spirit of Kashyap and Wilcox (1993), West (1995), and Ramey and West (1999). In our case, log marginal production cost is $\ln \theta J_t = \ln \theta + \ln J_t = \ln \theta + \ln A_t + \ln Y_t^{\theta_1-1} + \ln W_t^{\theta_2}$ and log marginal inventory holding cost is $\ln (\partial HC_t / \partial N_{t-1}) = \ln \delta_1 + \ln \delta_2 + (\delta_2 - 1)(\ln N_{t-1} - \ln X_t)$ (ignoring δ_3 , which is non-stochastic). We assume that $\ln J_t = \ln A_t + \ln Y_t^{\theta_1-1} + \ln W_t^{\theta_2}$ and $\ln N_{t-1} - \ln X_t$ are each stationary, which is sufficient to address Hamilton's (2002) important point. Table 1 confirms this assumption by showing that the data strongly reject the null hypothesis that either $\ln N/X$ or $\ln J$ contains a unit root.

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 Y_t and 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,

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

where $E_0 = E\{\cdot | \Omega_0\}$, and

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

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

$$N_t - N_{t-1} = Y_t - X_t. \quad (6)$$

The optimality conditions that result from this optimization problem are

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

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

$$E_{t-1} \{ \beta_t [N_t - N_{t-1} - Y_t + X_t] \} = 0 \quad (9)$$

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

To interpret the optimality conditions, use (7) to eliminate the Lagrange multiplier in (8) to obtain

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

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,

$E_{t-1} \beta_t \beta_{t+1} \theta_1 A_{t+1} Y_{t+1}^{\theta_1-1} W_{t+1}^{\theta_2}$ is the discounted marginal cost of producing a unit of output tomorrow,

and $E_{t-1} \left\{ \beta_t \beta_{t+1} \left(\delta_2 \delta_1 \left(\frac{N_t}{X_{t+1}} \right)^{\delta_2-1} + \delta_3 \right) \right\}$ 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 log linearizing the optimality conditions around steady-state values yields a linearized Euler equation:

$$E_{t-1} \left\{ (\theta_1 - 1) \theta_1 \bar{J} [\ln Y_t - \bar{\beta} \ln Y_{t+1}] + \theta_2 \theta_1 \bar{J} [\ln W_t - \bar{\beta} \ln W_{t+1}] + \bar{\beta} (\delta_2 - 1) \delta_2 \bar{\psi} [\ln N_t - \ln X_{t+1}] + \theta_1 \bar{J} r_{t+1} + \theta_1 \bar{J} \tilde{u}_{t+1}^A + c \right\} = 0 \quad (10)$$

where $\bar{J}$ is steady-state average production cost, $\theta_1 \bar{J}$ is steady-state marginal production cost¹⁰, $\bar{\psi}$ is steady-state value of average stockout avoidance costs¹¹, which are defined by

$\psi_t = \delta_1 [R_{N_t} (1 - x_{t+1})]^{\delta_2-1}$, where x_{t+1} is the growth rate of X between t and $t+1$, $\delta_2 \bar{\psi}$ is steady-

¹⁰ 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.

¹¹ Note that $\bar{\psi}$ is average steady state stockout avoidance costs, not average total inventory holding costs. The latter is $\bar{\psi} + \delta_3$, which includes both stockout avoidance costs and storage costs.

state marginal stockout avoidance costs, $\bar{R}_N$ is the steady-state inventory/sales ratio, $\bar{\beta}=1/(1+\bar{r})$, $\bar{r}$ is the unconditional mean real interest rate, $\bar{x}$ 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.

Table 1
Unit Root Tests

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 (except π_1 and π_3) 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.¹² Specifically, we assume that the real interest rate follows

¹² 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).

$$r_t = r_{S_t} + \sigma_{S_t} \cdot \varepsilon_t \quad (11)$$

where $\varepsilon_t \sim \text{i.i.d. } N(0,1)$ 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 independent. Denote the transition probabilities governing the evolution of S_t by $p_{ij} = \text{Prob}(S_t = j | S_{t-1} = i)$. Collecting these probabilities into a matrix we have

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

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,

$$\pi_t = \begin{bmatrix} \pi_{1t} \\ \pi_{2t} \\ \pi_{3t} \end{bmatrix} = \begin{bmatrix} \text{Prob}(S_t = 1 | \Omega_t) \\ \text{Prob}(S_t = 2 | \Omega_t) \\ \text{Prob}(S_t = 3 | \Omega_t) \end{bmatrix}. \quad (13)$$

Given π_{t-1} , the term $E_{t-1}r_{t+1}$ in equation (10) can be computed as

$$E_{t-1}r_{t+1} = \mathbf{r}'_v P^2 \pi_{t-1} = \gamma_1 \pi_{1t-1} + \gamma_2 \pi_{2t-1} + \gamma_3 \pi_{3t-1} \quad (14)$$

where $\mathbf{r}'_v = [r_1, r_2, r_3]$ and $\begin{bmatrix} \gamma_1 & \gamma_2 & \gamma_3 \end{bmatrix} \equiv \gamma' \equiv \mathbf{r}'_v P^2$. 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 (14) to obtain

$$E_{t-1}r_{t+1} = (\gamma_1 - \gamma_2) \pi_{1t-1} + (\gamma_3 - \gamma_2) \pi_{3t-1} + \gamma_2. \quad (15)$$

Then, substituting (15) into (10) yields

$$\begin{aligned} E_{t-1} \left\{ (\theta_1 - 1) \theta_1 \bar{J} [\ln Y_t - \bar{\beta} \ln Y_{t+1}] + \theta_2 \theta_1 \bar{J} [\ln W_t - \bar{\beta} \ln W_{t+1}] \right. \\ \left. + \bar{\beta} (\delta_2 - 1) \delta_2 \bar{\psi} [\ln N_t - \ln X_{t+1}] + \theta_1 \bar{J} \tilde{u}_{t+1}^A \right\} + \theta_1 \bar{J} [(\gamma_1 - \gamma_2) \pi_{1t-1} + (\gamma_3 - \gamma_2) \pi_{3t-1} + \gamma_2] + c = 0 \end{aligned} \quad (16)$$

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

The log-linearized Euler equation implied by the model, equation (16), may be written as a second-order expectational difference equation. Solving this difference equation for $E_{t-1} \ln N_t$, assuming I(1) processes for sales and real input costs, taking into account the Markov-switching process for the real interest rate, and using the inventory identity, (6), the decision rule for inventories is

$$\ln N_t = \Gamma_0 + \lambda_1 \ln N_{t-1} + \Gamma_X \ln X_{t-1} + \Gamma_W \ln W_{t-1} + \Gamma_{\pi 1} \pi_{1t-1} + \Gamma_{\pi 3} \pi_{3t-1} + u_t \quad (17)$$

where

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

$$\Gamma_W = - \frac{\bar{r} \theta_2}{(\theta_1 - 1)} \frac{\bar{R}_Y}{\bar{R}_N} \left(\frac{\lambda_1}{1 + \bar{r} - \lambda_1} \right) \quad (19)$$

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

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

$$\lambda_1 = 1 + \frac{\bar{r} + \zeta}{2} - \frac{1}{2} \left[(\bar{r} + \zeta)^2 + 4\zeta \right]^{\frac{1}{2}} \quad (21-a)$$

$$\zeta = \frac{(\delta_2 - 1) \delta_2 \bar{\psi}}{(\theta_1 - 1) \theta_1 \bar{J}} \frac{\bar{R}_Y}{\bar{R}_N} \quad (21-b)$$

$$u_t = (\bar{R}_Y / \bar{R}_N) (u_{t-1}^A + u_t^Y - u_t^X) \quad (21-c)$$

where λ_1 is the stable root of the relevant characteristic equation, u_t is a stationary shock that is a combination of the sales shock, u_t^X , a production shock, u_t^Y , and a technology shock, u_{t-1}^A , Γ_0 is a constant, and $\bar{R}_Y$ and $\bar{R}_N$ are the steady state values of R_Y and R_N respectively. The decision rule shock, u_t , arises from unanticipated fluctuations in sales, output and the state of technology.

In the short run, inventories act as a buffer, absorbing these unanticipated fluctuations. See Appendix B for a detailed derivation of the decision rule.

III. THE COINTEGRATING REGRESSION

Table 1 shows that $\ln N_t$, $\ln X_t$, $\ln W_t$, π_1 , and π_3 are $I(1)$.¹³ We therefore follow a number of papers in the literature by focusing on the cointegrating regression for inventories.¹⁴ In Appendix B, we show that the model in Section II implies that inventories, sales, input costs, and the interest-rate-regime probabilities are cointegrated, with cointegrating regression

$$\ln N_t = b_0 + b_X \ln X_t + b_W \ln W_t + b_{\pi_1} \pi_{1,t-1} + b_{\pi_3} \pi_{3,t-1} + v_t, \quad (22)$$

where

$$b_X = 1 - \frac{\bar{r}(\theta_1 - 1)\theta_1 \bar{J}}{(\delta_2 - 1)\delta_2 \bar{\psi}}, \quad (23-a) \quad b_W = -\frac{\bar{r}\theta_2\theta_1 \bar{J}}{(\delta_2 - 1)\delta_2 \bar{\psi}} \quad (23-b)$$

$$b_{\pi_1} = -(\gamma_1 - \gamma_2) \frac{(1 + \bar{r})\theta_1 \bar{J}}{(\delta_2 - 1)\delta_2 \bar{\psi}}, \quad (23-c) \quad b_{\pi_3} = -(\gamma_3 - \gamma_2) \frac{(1 + \bar{r})\theta_1 \bar{J}}{(\delta_2 - 1)\delta_2 \bar{\psi}} \quad (23-d)$$

b_0 is a constant, and v_t is a stationary error term. As discussed in Appendix B, this equation can be derived from the Euler equation for inventories. An important motivation for our focus on the cointegrating regression is the Maccini, Moore, and Schaller (2004) finding that the decision rule does poorly if one wants to measure the effect of the interest rate on inventories, often yielding point estimates that imply that a higher interest rate increases inventories and never yielding a significantly negative effect of the interest rate on inventories.

Equation (22) suggests an immediate test of the model, since it states that the variables in the equation will be cointegrated. The data are consistent with equation (22): The Johansen-

¹³ Since π_{1t} and π_{3t} have a restricted range, one might wonder whether it is better to model them as $I(0)$ or $I(1)$. We note two points. First, in careful applied econometric research, variables with restricted ranges, such as the nominal interest rate, are modeled as $I(1)$ variables when they are highly persistent. (See, e.g., Stock and Watson (1993) and Caballero (1994).) Second, unit root tests indicate that π_{1t} and π_{3t} are $I(1)$.

¹⁴ See, e.g., Kashyap and Wilcox (1993), Ramey and West (1999), Hamilton (2002), and Maccini, Moore, and Schaller (2004).

Juselius test rejects the null hypothesis of no cointegrating vector, with a test statistic of 97.9 (p-value=0.001).¹⁵

Stock and Watson (1993) explain the econometric problems associated with estimating a cointegrating regression using OLS and explain how Dynamic OLS (DOLS) addresses these problems. The DOLS specification of equation (22) is

$$\ln N_t = b_0 + b_T t + b_X \ln X_t + b_W \ln W_t + b_{\pi_1} \pi_{1,t-1} + b_{\pi_3} \pi_{3,t-1} + \sum_{s=-p}^p B_{X,s} \Delta \ln X_{t-s} + \sum_{s=-p}^p B_{W,s} \Delta \ln W_{t-s} + \sum_{s=-p}^p B_{\pi_1,s} \Delta \pi_{1,t-1-s} + \sum_{s=-p}^p B_{\pi_3,s} \Delta \pi_{3,t-1-s} + \eta_t. \quad (24)$$

Caballero (1994, 1999) provides econometric theory and Monte Carlo simulations showing that:

1) the bias associated with OLS estimation is particularly severe when adjustment frictions are important; and 2) a large value of p may be needed in equation (24) to correct the bias.¹⁶

Although the model in Section II does not involve any explicit adjustment frictions, the curvature of the production cost function, relative to the curvature of stockout avoidance costs, leads to slow adjustment of inventories (as we discuss in more detail later). With these econometric issues in mind, we conduct our own Monte Carlo simulations to choose p . Our Monte Carlo simulations of the econometric properties of Static OLS (SOLS; i.e., standard OLS) and DOLS estimators for the cointegrating regression for inventories are an innovation relative to previous work. Our simulations show that the econometric issues discussed by Stock and Watson (1993) and Caballero (1994, 1999) are relevant for the estimation of the cointegrating regression for inventories. Based on our Monte Carlo simulations, we set $p = 48$. The coefficient estimates are presented in Table 2.

We use the mappings in (23) – between the cointegrating regression coefficients and the structural parameters – in calibrating the structural parameters. In addition, we use estimates of parameters like γ_1 , which we obtain by estimating a Markov switching specification for the real interest rate, in calibrating the structural parameters. The details are in Appendix B. We assume

¹⁵ For reasons of data availability, the sample is 1959:01 to 2004:08. The variables enter as shown in equation (24) without detrending. The number of lags used in the test is set to minimize the AIC.

¹⁶ Some researchers might be inclined to cite the superconvergence property of cointegrating regressions as an advantage, relative to stationary econometric techniques. In light of the work of Stock and Watson (1993) and Caballero (1994, 1999) on the bias in coefficient estimates when OLS is used as the estimator of a cointegrating regression, we are hesitant to overemphasize this potential advantage of cointegrating regressions.

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 simplest form of $I(1)$ stochastic process for sales and input cost – a geometric random walk – with the variance calibrated to the data. (The variance of u_t in (17) is also calibrated to the data.) Using the decision rule, the calibrated values of the structural parameters, and the stochastic processes for sales, input costs, and the interest rate, we simulate the model.

IV. THE TRADITIONAL INVENTORY PUZZLES

A. Variance Ratio Puzzle

Table 1 shows that sales, output, and inventories are $I(1)$. The first difference of an $I(1)$ variable is $I(0)$, so $N_t - N_{t-1}$ is $I(0)$. From equation (6), this implies that $Y_t - X_t$ is equal to an $I(0)$ variable. This is equivalent to saying that Y_t and X_t are cointegrated with cointegrating vector $(1, -1)$. This, in turn, implies that $Var[Y_t] = Var[X_t]$ asymptotically; i.e., that the variance ratio $Var[Y_t]/Var[X_t] = 1$. The same holds if the variance ratio is defined in terms of the logs of output and sales; i.e., $Var[\ln Y_t]/Var[\ln X_t] = 1$.

The foregoing proof that the variance ratio is equal to 1 is very straightforward and flows directly from the starting point of our paper, which is that sales are $I(1)$. The proof relies on minimal assumptions – just the inventory accumulation identity. Despite the simplicity of the proof, we have not been able to find a statement of this result in the literature. The closest we

have been able to find is an inequality obtained by West (1990) – $Var[Y]/Var[X] \leq 1$ – that applies both when sales follow a stationary stochastic process and when sales are $I(1)$.¹⁷

It may be helpful to provide some intuition. We begin by explaining the standard production smoothing intuition and show how it is fundamentally changed when sales shocks are permanent. Suppose sales shocks are transitory and that sales take one of two possible values – Low (Y_0^L) or High (Y_0^H) – with equal probability, as shown in Figure 1. It is then optimal for the firm to always produce at a Middle (Y_0^M) level of output, since this is cheaper than producing at output levels Y_0^L half of the time and Y_0^H half of the time (which corresponds to cost at the point on the straight line directly above Y_0^M).

Figure 1
Transitory Sales Shocks: Intuition for Production Smoothing





Figure 2
Intuition for the Effect of a Permanent Sales Shock



Suppose instead that the firm faces a permanent sales shock that shifts the whole distribution of sales so that the middle point is now Y_1^M , rather than Y_0^M , as illustrated in Figure 2. Now the firm always wants to produce at the middle point Y_1^M , since this is the point at which


¹⁷ More precisely, West asserts that $E[X_t^2 - Y_t^2] \geq 0$. Under his assumption that all variables have zero unconditional mean, $E[X_t^2 - Y_t^2] = E[X_t^2] - E[Y_t^2] = Var[X_t] - Var[Y_t]$, so his result implies that $Var[Y]/Var[X] \leq 1$.

it minimizes cost. In the long run, the firm will therefore change output by the same amount as the permanent shock to sales.

If the variance ratio is 1, why do empirical researchers typically obtain estimates of the variance ratio that are greater than 1? We can use a simple figure to explain the intuition. In Figure 3, the solid line shows the path of sales in response to a one-standard-deviation shock. Because the shock is permanent, sales move to a higher level in period 1. Because we are illustrating a one-time shock, there is no further movement in sales.

The simulated path of output is shown by the line with open circles in Figure 3. Because output and sales are cointegrated with cointegrating vector $(1, -1)$, they move by the same amount in the long run. With a long enough sample, an empirical researcher would therefore obtain a variance ratio of 1.

Figure 3
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 the ratio θ_1/δ_2 were 2.5 times its calibrated value.

In period 1, production barely moves. This reflects the buffering role of inventories, which is captured in our model by the assumption that the firm does not know the state of demand at the time it sets production for period 1. In the initial period, the increase in sales is met by depleting inventories. If the stockout avoidance motive dominates the production smoothing motive, in period 2, output rises by more than the amount of the sales shock. This is

illustrated in the line with open circles. Because output rises more than the amount of the sales shock, an empirical researcher with several years of data will obtain a variance ratio greater than 1.

Table 3 reports the variance ratio for simulations with different sample sizes. For a fairly small sample size (50 observations), the variance ratio is 1.09. As the sample size increases, the variance ratio falls. For a sample size of 548 monthly observations (the length of our sample), the simulations yield a variance ratio of 1.02, which is in line with the typical empirical finding of a variance ratio greater than 1. For a large enough sample (5000 monthly observations), the small sample bias disappears and the variance ratio is 1.00.

Table 3
Simulation Evidence on Small Sample Bias in the Variance Ratio

Sample Size (Number of Monthly Observations)	50	125	250	548	1000	2500	5000
Median Variance Ratio	1.09	1.05	1.03	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 10,000 repetitions of the simulation.

Intuitively, the amount by which production changes in period 2 will depend on the relative strength of the production smoothing and stockout avoidance motives. The production smoothing motive tends to lead the firm to change output by less than the change in sales in the short run. On the other hand, the firm does not want to lose sales by running out of inventory, so the stockout avoidance motive tends to lead the firm to change output by more than the change in sales – to satisfy the new, permanently higher level of demand and to replenish the inventories that were depleted by the initial sales shock.

The simulation results confirm this intuition. The line with solid triangles in Figure 3 shows what the response of production would be if the production smoothing motive dominated the stockout avoidance motive. In this case, output would rise by less in period 2 than the amount of the original sales shock. Production gradually creeps upward, asymptotically approaching the new, higher level of sales. If the production smoothing motive dominates the stockout avoidance motive, small sample bias causes the variance ratio to approach 1 from below. For example, if we increase the ratio θ_1/δ_2 by a factor of 2.5, the variance ratio is 0.99 for sample sizes of 50 or 125. The fact that empirical researchers typically find that the variance ratio is greater than 1

provides indirect evidence that our approach to calibrating the inventory model is on the right track. Later in the paper, we derive analytical results that provide further insight into the way in which the balance between the production smoothing and stockout avoidance motives influences the behavior of inventories.

B. Wen Puzzle

i. Basic Results

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."¹⁸

We begin by explaining the intuition for our solution to the Wen puzzle. Suppose the firm is hit with a permanent sales shock, as illustrated in Figure 3. At first, because the firm set production before it knew what demand would be in period 1, output is essentially unchanged and the shock is absorbed by running down inventories. Thus, in the first period, the volatility of output is low relative to the volatility of sales. In the second period, output responds to the sales

¹⁸ 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. It is not clear that recent papers such as Khan-Thomas (2007) or Wen (2011) are able to explain the Wen Puzzle any better than earlier papers. These two papers certainly do not offer an explicit solution for the Wen Puzzle.

shock. This increases the measured variance of output, but the conditional variance of output is still smaller than the conditional variance of sales (where both are conditional on their levels before the shock). Over the medium term, output continues to move in a delayed response to the original shock, but sales are unchanged, since the initial shock was permanent. Eventually, the ongoing movements in output lead to a measured variance of output that is greater than the variance of sales, so the conditional variance ratio exceeds 1. (Just below, we explain further what we mean by the conditional variance ratio.) In the long run, output and sales move by the same amount, so the conditional variance ratio approaches the unconditional variance ratio, which, as we showed earlier, is equal to 1.

As this intuition illustrates, we need a counterpart to the variance ratio that takes into account the horizon. This is where the conditional variance ratio fits in. The numerator of the conditional variance ratio is the variance of output, conditional on output n periods in the past. The denominator is the variance of sales, conditional on sales $n+1$ periods in the past. (Both output and sales are measured in logs for analytical tractability.) For specificity, assume the I(1) process for sales is $\ln X_t = \ln X_{t-1} + u_t^X$, where u_t^X is i.i.d. with mean 0 and variance σ_X^2 . Based on the model in Section II, we obtain the following result.¹⁹

¹⁹ To obtain the conditional variance ratio, first solve for the decision rule for output. To do so, use the solution for $E_{t-1} \ln Y_t$ and the inventory identity and assume that sales and input costs follow I(1) processes (specifically, random walks). To simplify the analysis and to focus on the traditional explanations for the variance ratio puzzle, assume further that real interest rates are constant, which implies that the probabilities are fixed over time. Then, the decision rule for output is:

$$\ln Y_t = \lambda_1 \ln Y_{t-1} + (1 - \lambda_1) \ln X_{t-1} + (1 + \tilde{\Gamma}_X) u_{t-1}^X + \tilde{\Gamma}_W u_{t-1}^W + c_Y$$

where u_{t-k}^X is the sales shock and u_{t-i}^W is the input cost shock. Next, using backward substitution, the above equation can be written as:

$$\ln Y_t = \lambda_1^n \ln Y_{t-n} + (1 - \lambda_1) \left(\sum_{i=0}^{n-1} \lambda_1^i \right) \ln X_{t-(n+1)} + \sum_{k=0}^{n-1} \left[(1 - \lambda_1) \sum_{i=0}^k \lambda_1^i + (1 + \tilde{\Gamma}_X) \lambda_1^k \right] u_{t-k-1}^X + \tilde{\Gamma}_W \sum_{i=1}^n \lambda_1^{i-1} u_{t-i}^W + \tilde{c}_Y.$$

Taking the variance of this equation appropriately yields the Analytical Conditional Variance Ratio. See Appendix B for the details.

$$\begin{aligned}
\text{Analytical Conditional Variance Ratio} &= \frac{\text{Var}[\ln Y_t | \ln Y_{t-n}]}{\text{Var}[\ln X_t | \ln X_{t-(n+1)}]} \\
&= \frac{1}{1+1/n} + \frac{(1-\lambda_1 + \tilde{\Gamma}_X)}{(1+n)} \left[(1-\lambda_1 + \tilde{\Gamma}_X) \left(\frac{1-\lambda_1^{2n}}{1-\lambda_1^2} \right) + 2 \left(\frac{1-\lambda_1^n}{1-\lambda_1} \right) \right] + \left[\frac{\tilde{\Gamma}_W^2}{1+n} \right] \left(\frac{1-\lambda_1^{2n}}{1-\lambda_1^2} \right) \frac{\sigma_W^2}{\sigma_X^2},
\end{aligned} \tag{25}$$

where $\text{Var}(\ln Y_t | \ln Y_{t-n})$ is the variance of $\ln Y_t$ conditional on $\ln Y_{t-n}$, $\text{Var}(\ln X_t | \ln X_{t-(n+1)})$ 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

$$\tilde{\Gamma}_X = \frac{\bar{R}_N}{\bar{R}_Y} \Gamma_X = \left[\frac{(\delta_2 - 1)\delta_2 \bar{\psi}}{(\theta_1 - 1)\theta_1 \bar{J}} - \bar{r} \right] \frac{\lambda_1}{(1 + \bar{r} - \lambda_1)} \tag{26-a}$$

$$\tilde{\Gamma}_W = \frac{\bar{R}_N}{\bar{R}_Y} \Gamma_W = -\frac{\bar{r}\theta_2}{\theta_1 - 1} \left(\frac{\lambda_1}{1 + \bar{r} - \lambda_1} \right), \tag{26-b}$$

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

Equations (25) and (26) confirm a key insight in the existing literature. The variability of output, relative to sales, depends on the relative strength of the production smoothing and stockout avoidance motives. More precisely, it depends on the convexity of production costs divided by the convexity of stockout adjustment costs (i.e., $(\delta_2 - 1)\delta_2 \bar{\psi} / ((\theta_1 - 1)\theta_1 \bar{J})$). A key new insight of our paper is that this is true in the “short run”, but not in the long run. As we prove in our explanation of the variance ratio puzzle, in the long run, the variance ratio is 1.

Using equation (25) and the calibrated parameters, we can calculate the analytical conditional variance ratio as a function of the horizon n . The solid line in Figure 4 shows that the analytical conditional variance ratio is less than 1 at short horizons and greater than 1 at business cycle horizons.


We define the empirical conditional variance ratio as

$$\text{Empirical Conditional Variance Ratio} = \frac{\text{Var}(\ln Y_t - \ln Y_{t-n})}{\text{Var}(\ln X_t - \ln X_{t-(n+1)})}. \tag{27}$$

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

$$\ln X_t = \ln X_{t-(n+1)} + \sum_{j=1}^{n+1} u_{t-j}^X. \quad (28)$$

Figure 4
Analytical Conditional Variance Ratio



The solid line shows the analytical conditional variance ratio calculated from equation (25), where the structural parameters are calibrated using the cointegrating regression. The dashed line shows the analytical conditional variance ratio if the relative strength of the production smoothing and stockout avoidance motives (i.e., the ratio of the convexity of production costs to the convexity of the stockout avoidance costs, $(\delta_2 - 1)\delta_2\bar{\psi}/(\theta_1 - 1)\theta_1\bar{J}$), were equal to 2.5 times the calibrated value. The horizontal axis shows the horizon (n) in months.

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

$$\text{Var}\left(\ln X_t \mid \ln X_{t-(n+1)}\right) \equiv \text{Var}\left(\ln X_t - \ln X_{t-(n+1)}\right) = (n+1)\sigma_X^2. \quad (29)$$

In the data, the empirical conditional variance ratio shows the pattern documented by Wen (2005). As shown in the first row of Table 4, at short horizons, the empirical conditional variance ratio is less than 1; for example, for $n = 1$ month, it is 0.62; for $n = 2$ months, it is 0.71. At business cycle horizons, the empirical conditional variance ratio is greater than 1; for example, for $n = 50$ months, it is 1.02; for $n = 70$ months, it is 1.01.

The analytical conditional variance ratio has the same pattern as the empirical conditional variance ratio. At short horizons, both are substantially below 1, as shown in Table 4. At business cycle horizons, both are larger than 1.

ii. The Role of the Interest Rate

Equations (25) and (26) provide an important new insight: The variability of output, relative to the variability of sales, depends on the mean real interest rate. We begin by explaining the intuition. Next, we use the equations of the model to precisely describe the details. Finally, we use the model to illustrate the effect of the mean real interest rate.

Table 4
Conditional Variance Ratio for Selected n

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

The row labeled “Analytical” reports the analytical conditional variance ratio from the model, which is calculated from equation (25) with the structural parameters calibrated to the data. The row labeled “Empirical” reports the empirical conditional variance ratio, calculated from the data using the definition in equation (27).

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 production smoothing dominates, the firm’s desired long-run inventory stock decreases. But, if stockout avoidance dominates, the desired long-run inventory stock increases. Of course, the present value of the reduction in stockout avoidance costs depends on the mean real interest rate, so the reciprocal of the mean real interest rate (i.e., $1/\bar{r}$) weights the stockout avoidance motive. The higher the mean real interest rate, the less weight the firm places on stockout avoidance.

From the decision rule for inventories, we can derive the long-run desired stock of inventories²⁰, N_t^* .

$$\ln N_t^* = \frac{1}{1 - \lambda_1} \left[\Gamma_X \ln X_{t-1} + \Gamma_W \ln W_{t-1} + \Gamma_{\pi 1} \pi_{1t-1} + \Gamma_{\pi 3} \pi_{3t-1} + \Gamma_o \right] \quad (30)$$

²⁰ Specifically, the desired stock of inventories is derived by setting $\ln N_t = \ln N_{t-1} = \ln N_t^*$ and $u_t = 0$ in equation (17).

The coefficient Γ_x determines whether or not a positive sales shock increases N_t^* .²¹

$$\Gamma_x \begin{matrix} > \\ < \end{matrix} 0 \quad \text{as} \quad \frac{(\delta_2 - 1)\delta_2\bar{\psi}}{\bar{r}} \begin{matrix} > \\ < \end{matrix} (\theta_1 - 1)\theta_1\bar{J} \quad (31)$$

The left-hand side of the inequality is the present value of the change in marginal stockout avoidance costs. The right-hand side is the change in marginal production costs. Thus, condition (31) formalizes the intuition presented above. 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.

Once we take into account the role of the real interest rate, we find that the stockout avoidance motive dominates the production smoothing motive. This is directly reflected in the estimate of the cointegrating regression coefficient on sales b_x , which will be positive, as we can see from equation (23-a), if the stockout avoidance motive dominates the production smoothing motive (i.e., if condition (31) is satisfied). In fact, the point estimate of b_x from the cointegrating regression is positive. This provides direct evidence that stockout avoidance dominates production smoothing, even before we calculate the structural parameters and substitute them into condition (31) or use them in equation (25) to calculate the analytical conditional variance ratio.


Condition (31) plays a role in equations (25) and (26), since Γ_x also appears there. Equations (25) and (26) are the equations that describe how the volatility of output, relative to the volatility of sales, evolves in the “short run” (i.e., over any finite horizon).

The role of the real interest rate in determining the analytical conditional variance ratio is illustrated in Figure 5. In our data, the mean annual real interest rate is 1.2%, which corresponds to a monthly real interest rate of 0.1% (i.e., $\bar{r} = 0.001$). This implies that the production-smoothing motive dominates the analytical conditional variance ratio at short horizons, as shown

²¹ Γ_x is defined in equation (18). $\lambda_1/(1 + \bar{r} - \lambda_1)$ is positive because: 1) equation (21-b) shows that λ_1 is strictly positive and less than or equal to 1; and 2) $\bar{r}$ is positive. Recall that $\bar{\psi} = \delta_1[(1 - \bar{x})\bar{R}_N]^{\delta_2 - 1}$, $\delta_1 > 0$, and $\delta_2 < 0$, so, under the very mild assumption that the steady-state growth rate of sales $\bar{x}$ is less than 100%, $\bar{\psi} > 0$. Thus $(\delta_2 - 1)\delta_2\bar{\psi} > 0$. Recall that $\theta_1 > 1$ and $\bar{J} > 0$, so $(\theta_1 - 1)\theta_1\bar{J} > 0$. $\bar{R}_Y > 0$ and $\bar{R}_N > 0$. Thus, the sign of Γ_x depends on the inequality in equation (31).

by the solid line. If the real interest rate were lower, the stockout avoidance motive would be even more important and the analytical conditional variance ratio 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 analytical conditional variance ratio would never be greater than 1, as illustrated by the line with triangles.

Figure 5
The Real Interest Rate and the Analytical Conditional Variance Ratio



The solid line shows the analytical conditional variance ratio, calculated from equation (25), where the annual mean real interest rate is set equal to its value in the data, which is 1.2%. The line with open circles shows the analytical conditional variance ratio for a mean real interest rate of 0.3% and the line with triangles shows the analytical conditional variance ratio for a mean real interest rate of 5%.

iii. The Cyclicity of Inventory Investment

An important question in the literature has been the cyclicity of inventory investment; see, e.g., Ramey and West (1999) and Wen (2005). The inventory identity implies that

$$Var[Y] = Var[X] + Var[\Delta N] + 2Cov[X, \Delta N]. \quad (32)$$

Equation (32) illustrates the close relationship between the variance ratio puzzle and the cyclicity of inventory investment. Suppose inventory investment were acyclical, so $Cov[X, \Delta N] = 0$. Then $Var[Y] \geq Var[X]$, since $Var[\Delta N] \geq 0$. Thus, the only way that the variance of output could be less than the variance of sales would be if inventory investment were countercyclical (i.e., $Cov[X, \Delta N] < 0$). In the data, researchers tend to find that inventory

investment is procyclical. If production smoothing were the dominant motive, the empirical result that inventories are procyclical would be another puzzle.

As we emphasize in this paper, sales shocks are $I(1)$. As discussed above, this implies that the stock of inventories is $I(1)$, so inventory investment is $I(0)$. In the long run, the variance of an $I(0)$ variable becomes arbitrarily small relative to the variance of a $I(1)$ variable, so, in the long run, $Var[\Delta N]$ becomes arbitrarily small relative to $Var[X]$. Similarly, in the long run, $Cov[X, \Delta N]$ becomes arbitrarily small relative to $Var[X]$. Equation (32) therefore provides another way of seeing the result that the variance ratio is 1 in the long run, since $Var[Y] = Var[X]$ asymptotically.

At business cycle horizons, the conditional variance ratio is greater than 1, both theoretically (in our calibrated model) and empirically. This is consistent with the empirical evidence that inventory investment is procyclical, since $Cov[X, \Delta N] > 0$ implies that $Var[Y] > Var[X]$. The simplest way to see the intuition is to look at Figure 3. Focus on the case where the stock avoidance motive dominates the production smoothing motive (the case consistent with our estimates of the cointegrating regression). The path of output is represented by the line with open circles. Output hardly moves at the time of the shock, because the firm has set production before it knows the current state of demand. The only way it can meet the increase in demand is by depleting its inventories. If the stockout avoidance motive dominates, in the next period, output increases by more than the increase in sales because the firm needs to increase output to meet the new, permanently higher level of demand and because it needs to replace the inventories that were depleted in the initial period. Thus, inventory investment is high at the same time that sales are high; i.e., inventory investment is procyclical (at business cycle horizons).

C. Slow Adjustment Puzzle

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

$$N_t - N_{t-1} = \vartheta(N_t^* - N_{t-1}) + u_t^N \quad (33)$$

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, since 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

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. In the data, sales shocks are permanent, so the desired long-run inventory level rises when sales go up. The stronger the stockout avoidance motive, the more quickly the firm wants to adjust output to reach its new desired long-run inventory level.

To formally analyze the role of production smoothing and stockout avoidance, we begin by deriving the inventory investment relationship by subtracting $\ln N_{t-1}$ from both sides of (17) to get

$$\ln N_t - \ln N_{t-1} = (1 - \lambda_1) [\ln N_t^* - \ln N_{t-1}] + u_t \quad (34)$$

where $\ln N_t^*$ is the desired long-run stock of (log) inventories as defined in equation (30).

Comparing (34) to (33) we see that $1 - \lambda_1$ measures the speed of adjustment. Using the definition of λ_1 in (21-a), this speed of adjustment term can be written as

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

where ζ is defined in equation (21-b).

²² A number of possible explanations have been put forward for slow adjustment. One emphasizes econometric problems -- either omitted variables or problems with the econometric procedure -- see Maccini and Rossana (1984) and Blinder (1986a). Another explores 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.

The term $(\delta_2 - 1)\delta_2\bar{\psi}$ is the convexity of stockout avoidance costs and $(\theta_1 - 1)\theta_1\bar{J}$ is the convexity of production costs. Note from (21-b) that they enter ζ as a ratio. This implies that it is the relative strength of production smoothing and stockout avoidance that matters. As we have seen, this is a recurring theme that emerges from our analysis.²³ Straightforward mathematics shows that an increase in the convexity of production costs relative to the convexity of stockout avoidance costs decreases the speed of adjustment.²⁴

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 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 fairly strong.

D. Input Cost Puzzle

Economic theory suggests that higher input costs should reduce inventory holdings. Intuitively, this is because higher input costs make it more expensive for the firm to produce a marginal unit of output and therefore make it more expensive to build up inventories. Our model is consistent with previous economic theory in this respect. The effect of input costs on desired

²³ In the model section, we explain that we choose to represent the stockout avoidance motive in the cost function because this makes the analysis far clearer. Equation (21-b) is an example of how our modeling choice is helpful in clarifying that it is the relative importance of production smoothing and stockout avoidance that matters.

²⁴ To see this, differentiate (35) with respect to $(\delta_2 - 1)\delta_2\bar{\psi}/(\theta_1 - 1)\theta_1\bar{J}$, recalling: 1) the parameter restrictions $\theta_1 > 1$ and $\delta_2 < 0$; 2) the definitions of $\bar{J}$ (steady-state average production cost) and $\bar{\psi}$ (steady-state average stockout avoidance costs), which imply that $\bar{J} > 0$ and $\bar{\psi} > 0$; and 3) the definitions of $\bar{R}_Y$ (steady-state output/sales ratio) and $\bar{R}_N$ (steady-state inventory/sales ratio), which imply that $\bar{R}_Y > 0$ and $\bar{R}_N > 0$.


long-run inventories is reflected in Γ_w in equation (30). From the definition of Γ_w in equation (19), it is straightforward to show that the effect of higher input costs is to reduce desired long-run inventories under the assumptions of the model.²⁵

Although economic theory is clear on this point, the empirical inventory literature provides limited evidence that input costs significantly reduce inventory holdings. (See, e.g., the survey of empirical studies in Ramey and West (1999).) In contrast, when we estimate the cointegrating regression, b_w is -0.753 with a t-statistic of -5.244. Thus, the cointegrating regression provides significant evidence that an increase in input costs reduces inventories.

Why does the cointegrating regression provide stronger evidence than previous studies? The culprit for the weak evidence in earlier studies is the slow adjustment of inventories. The adjustment speed is $1 - \lambda_1 = 0.051$, which implies that only about 5% of the gap between desired long-run inventories $\ln N_t^*$ and actual inventories $\ln N_{t-1}$ is closed each month, as shown in equation (34). This means that the response of inventories to a decrease in input costs is spread over many months. This is illustrated in Figure 6, which shows the dynamic response of inventories to a permanent input cost shock, based on the calibrated model. It takes about two years (25 months, to be precise) for inventories to get $\frac{3}{4}$ of the way to the new steady state after a shock. This means that a very large number of lags would be needed to capture the effect of the shock with traditional techniques. But the lag coefficients are not structural parameters. In fact, they will be affected by shocks to other variables, so they will not be stable over time. This means it is nearly hopeless to try to estimate the effect of input cost shocks using traditional techniques (i.e., the econometrics of stationary variables). This is illustrated in Table 5. The first column of the table presents estimates of the coefficients on lagged input costs from a regression of $\ln N_t$ on $\ln N_{t-1}$ and three lags each of $\ln X_t$ and $\ln W_t$. The second column presents similar information for a specification that uses 12 lags each of $\ln X_t$ and $\ln W_t$. The first two columns report the medians of the point estimates and standard errors over 10,000 repetitions of the simulation. None of the coefficients on input costs are statistically significant. In fact, in every

²⁵ Under the assumptions of the model, $\theta_1 > 1$ and $\theta_2 > 0$. We assume that $\bar{r}$, the unconditional mean real interest rate, is positive. Thus $\bar{r}\theta_2/(\theta_1 - 1) > 0$. We have that $0 < \lambda_1 < 1$, so $\lambda_1/(1 + \bar{r} - \lambda_1) > 0$. Since $\bar{R}_y > 0$ and $\bar{R}_N > 0$, it follows from (19) that $\Gamma_w < 0$.

Figure 6
Response of Inventories to a Permanent Input Cost Shock



case, the coefficient estimate is much smaller in absolute value than the standard error. The third and fourth columns present the same information as the second column, this time for two randomly selected repetitions of the simulation. The last two columns illustrate how different the estimates of the lag coefficients and the cumulative effect can be. For example, the coefficients on the first two lags have different signs in the two sample repetitions (not to mention very different magnitudes). The sum of the first three coefficients is about 0.04 in the first sample repetition, compared to about -0.24 in the second sample repetition.

Intuitively, the cointegrating regression provides strong evidence on the effect of input cost shocks because it is based on permanent shocks and therefore reflects the long-run relationship between input costs and inventories. In the long run, input costs do have a significant effect, but, because the response of inventories is so slow, traditional (stationary) econometric techniques are not well suited to capturing this effect.

While on the subject of cost shocks, we briefly address another issue. In the literature, cost shocks have been a leading potential explanation for the variance ratio puzzle.²⁶ In our model, as in other models, cost shocks tend to increase the conditional variance ratio. However, numerical results based on equation (25) show that the contribution of production cost shocks to the conditional variance ratio is negligible – less than 1%. The production smoothing and

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

stockout avoidance motives, rather than cost shocks, are the main determinants of the variance ratio in the “short run” (i.e., at any finite horizon).

Table 5
The Effect of Input Costs on Inventories:
An Illustration of the Results from Stationary Econometrics
(Based on Simulations of the Calibrated Model)

Input Costs	Median Coefficient Estimate (Standard Error)		Two Sample Repetitions Coefficient Estimate (Standard Error)	
	3 Lags	12 Lags	12 Lags	12 Lags
$\ln W_{t-1}$	-0.0383 (0.1781)	-0.0371 (0.1812)	-0.2045 (0.1829)	0.0035 (0.1942)
$\ln W_{t-2}$	0.0020 (0.2499)	0.0031 (0.2542)	0.2620 (0.2521)	-0.0646 (0.2741)
$\ln W_{t-3}$	-0.0036 (0.1784)	0.0034 (0.2541)	-0.0180 (0.2520)	-0.1817 (0.2732)
$\ln W_{t-4}$		0.0018 (0.2541)	-0.3948 (0.2512)	0.2675 (0.2743)
$\ln W_{t-5}$		0.0013 (0.2542)	0.4724 (0.2516)	0.3162 (0.2748)
$\ln W_{t-6}$		-0.0013 (0.2542)	0.1930 (0.2524)	-0.3526 (0.2767)
$\ln W_{t-7}$		-0.0020 (0.2542)	-0.7395 (0.2526)	-0.0484 (0.2770)
$\ln W_{t-8}$		-0.0025 (0.2542)	0.5551 (0.2547)	0.2666 (0.2754)
$\ln W_{t-9}$		0.0023 (0.2542)	-0.0597 (0.2561)	0.1921 (0.2756)
$\ln W_{t-10}$		-0.0021 (0.2543)	-0.2228 (0.2563)	-0.5220 (0.2749)
$\ln W_{t-11}$		0.0012 (0.2543)	0.1397 (0.2563)	0.3477 (0.2767)
$\ln W_{t-12}$		0.0010 (0.1815)	-0.0281 (0.1851)	-0.2628 (0.1956)

The coefficients and standard errors come from a regression of $\ln N_t$ on $\ln N_{t-1}$ and lags of $\ln X_t$ and $\ln W_t$, based on simulated data from the calibrated model. The entries in the first two columns represent the median over 10,000 repetitions of the simulation.

V. 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 C.


Figure 7

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 lines in Figures 7-A and 7-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 what is known in the VAR literature as 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 7 shows the impulse response functions of π_1 and π_3 to a one-standard-deviation easing of monetary policy. As Figure 7-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.

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.

In this paper, we do something new: We estimate the relationship between inventories and the interest rate with both $I(0)$ and $I(1)$ econometric techniques.³² Table 6 presents estimates based on applying traditional $I(0)$ techniques to simulated data. The first column of the table presents estimates of the coefficients on lags of the interest rate in a regression of $\ln N_t$ on $\ln N_{t-1}$ and three lags each of $\ln X_t$, $\ln W_t$, and r_t . The second column presents similar information for a specification with 12 lags. None of the coefficients on the interest rate are statistically significant. In fact, in every case, the coefficient estimate is much smaller in absolute value than the standard error.

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 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 (17), 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 7) 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.

³² Maccini, Moore, and Schaller (2004) estimate the cointegrating regression for inventories on actual data, but do not compare $I(0)$ and $I(1)$ econometric techniques on simulated data.

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.

Table 6
The Effect of the Interest Rate on Inventories:
An Illustration of the Results from Stationary Econometrics
(Based on Simulations of the Calibrated Model)


Real Interest Rate	Median Coefficient Estimate (Standard Error)	
	3 Lags	12 Lags
$\ln r_{t-1}$	-0.1222 (0.9435)	-0.1201 (1.0056)
$\ln r_{t-2}$	-0.0957 (0.9693)	-0.0874 (1.0481)
$\ln r_{t-3}$	-0.0660 (0.9436)	-0.0440 (1.0660)
$\ln r_{t-4}$		-0.0198 (1.0734)
$\ln r_{t-5}$		-0.0210 (1.0768)
$\ln r_{t-6}$		-0.0114 (1.0778)
$\ln r_{t-7}$		-0.0525 (1.0778)
$\ln r_{t-8}$		0.0035 (1.0766)
$\ln r_{t-9}$		-0.0130 (1.0736)
$\ln r_{t-10}$		0.0084 (1.0662)
$\ln r_{t-11}$		-0.0090 (1.0486)
$\ln r_{t-12}$		0.0016 (1.0060)

The coefficients and standard errors come from a regression of $\ln N_t$ on $\ln N_{t-1}$ and lags of $\ln X_t$, $\ln W_t$, and r_t , based on simulated data from the calibrated model. Each entry represents the median over 10,000 repetitions of the simulation.

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 8. As found in other studies, 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.

Figure 8
Empirical Responses to a Stimulative Monetary Policy Shock




Figures 8-A and 8-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.

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 (17), to find the theoretical response of log inventories to a monetary policy shock.³³ The solid line in Figure 9 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

³³ Since the monetary policy shock is by definition unanticipated we assume that the initial increase in sales is also unanticipated.

policy shock. Production does not respond immediately, so inventories fall as they buffer the positive sales shock.

Figure 9
Theoretical Response of Inventories to a Stimulative Monetary Policy Shock



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. The dashed line shows the theoretical response of inventories based on setting θ_1/δ_2 equal to 0.5 times the value obtained when the parameters are calibrated using the cointegrating regression. The horizontal axis shows time in months.

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 8-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 8-B. The peak effect of the monetary policy shock on inventories occurs years after the shock.³⁴ 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 7-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).

³⁴ This is also documented in Christiano, Eichenbaum, and Evans (1996) and Jung and Yun (2005).

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.

In Figure 9, we illustrate the effect of changes in the convexity of production cost, relative to the convexity of stockout avoidance cost, on the theoretical impulse response function for inventories. If we set θ_1/δ_2 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 9).

VI. 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 find that output responds more than one-for-one to a demand shock over the time horizon that is relevant for business cycles. This response is driven by the firm's inventory policy. In view of the fact that inventory movements account for a large share of the drop in output during recessions (including the Great Recession!), it may be time for us to pay more attention to inventories.

Our model is not radically innovative relative to the workhorse linear-quadratic model. We do extend the standard model by assuming constant elasticity cost functions, which permit log-linear approximations around steady state values, but the economic mechanisms are essentially the same.³⁵ The key innovative aspect of our approach is to take seriously the

³⁵ Bilts and Kahn (2000) argue that this class of models (including the linear-quadratic model and our variant) is ill-suited to studying the cyclical properties of marginal cost or the markup. Our objective, however, is to explain the four traditional puzzles and three monetary policy puzzles. For this objective, sticking close to the linear-quadratic model – and putting the stockout avoidance motive in the firm's cost function – is tremendously helpful in clarifying the economic mechanisms that are at work. Just think about how many of the puzzles we explain in terms of the convexity of production cost relative to the convexity of stockout avoidance cost.

empirical fact that sales shocks are $I(1)$. This leads to a series of new analytical insights about inventories. In the long run, the variance ratio (variance of output/variance of sales) is 1. The long-run variance ratio is not affected by the production smoothing or stockout avoidance motives. However, at business-cycle horizons, our analytical results show that the conditional variance ratio does depend on the relative strength of these two motives. The steady-state real interest rate plays a key role in determining whether production smoothing or stockout avoidance dominates the behavior of inventories at business-cycle horizons. All of these results, which flow from our analytical approach, are new.

Once we take into account the fact that sales shocks are $I(1)$, we are able to explain a wide range of inventory puzzles. Although our model is not radically new, the fact that sales shocks are $I(1)$ does suggest a different way of calibrating an inventory model – using the cointegrating regression that links inventories to sales, input costs, and the probability of being in a given real-interest-rate regime.

The model and the empirical work together shed light on four traditional inventory puzzles that have plagued the literature for decades. Our resolution of two of the puzzles hinges on the interaction between the production smoothing and stockout avoidance motives. One important traditional puzzle is the Slow Adjustment Puzzle (inventories adjust very slowly to their desired levels). Our analytical results show that the adjustment speed is slow because the convexity of production costs is high, relative to the convexity of stockout avoidance costs. Part of the Wen Puzzle is that the variance of production exceeds the variance of sales at business cycle horizons.³⁶ Our analytical results show that this occurs because the present value of the change in marginal stockout avoidance costs exceeds the one-time change in marginal production costs. This is where the real interest rate plays a key role.

Another traditional puzzle that is crucial both to economic policy and to our understanding of the aggregate economy is the effect of input costs on inventories. Economic theory predicts that higher input costs should decrease the level of inventories. But it has been devilishly difficult to find empirical evidence of the effect of observable cost shocks on inventories. The key empirical problem is the slow adjustment of inventories to shocks. This

³⁶ The other part of the Wen Puzzle is the fact that the variance of production is less than the variance of sales at short horizons.

means that the effect of an input cost shock is spread over several years. It is very difficult to pick up this effect with stationary econometric techniques. When we use a conventional partial adjustment specification (a long-established technique, based on stationary econometrics, that has been used both in the inventories literature and other areas of economics) to estimate the effect of input costs on inventories using simulated data from the model, none of the coefficients are significant. This is not surprising, since the effect of the cost shock is spread over several years and the lag coefficients are not structural parameters. In fact, the lag coefficients depend on the sequence of shocks, both to the interest rate and other variables. In contrast, the cointegrating regression captures the long-run effect of input costs. Using actual US data, the cointegrating regression shows that input costs have a highly significant effect on inventories.

The cointegrating regression itself helps us solve several important monetary policy puzzles. A puzzle that is of crucial importance for economic policy is the Mechanism Puzzle. Empirical work based on stationary econometrics has found little evidence that the interest rate affects inventories. If the interest rate doesn't affect inventories, how does monetary policy influence inventories? In fact, we find that persistent changes in the interest rate have a strong effect on inventories. This comes through clearly in the cointegrating regression, where π_1 and π_3 , the probabilities of being in the low and high interest-rate regime, respectively, are highly significant. We make two further contributions. First, we find that more than 75% of the peak change in inventories in response to a monetary policy shock in the simulations is accounted for by the pure "opportunity cost" effect of the change in the interest rate (i.e., the direct effect of the interest rate on inventories, as opposed to indirect effects through changes in sales or input costs). Second, we show that, despite the fact that the interest rate has a strong effect on inventories, the coefficients on lagged values of the interest rate are statistically insignificant if we use $I(0)$ econometrics to estimate the relationship between inventories and the interest rate in the simulated data.

Our model and empirical work are able to rationalize the puzzling "double-hump" response of inventories to a stimulative monetary policy shock, which reflects the Sign Puzzle and the Timing Puzzle. In response to a stimulative monetary policy shock, inventories initially decline, then rise for several years, and then decline again in the data. Our results indicate that the buffering role of inventories accounts for the initial decline and that the slow adjustment speed of the subsequent increase in inventories is due to a combination of production smoothing

and the rational reluctance of firms to interpret a transitory decline in the interest rate as permanent. Inventories decline again as the effect of the transitory monetary policy shock ultimately fades away.

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

Second, in 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.

Finally, the model and empirical work are purposely set in a partial equilibrium environment. We show that a basic model of inventory holding behavior by a representative firm together with $I(1)$ sales shocks is capable of reconciling the various traditional and monetary policy puzzles that have occupied the inventory literature for decades. An important task for future work is to incorporate permanent demand shocks in a general equilibrium framework with endogenous sales and input prices.

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 from 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.
- Wen, Yi (2011), "Input and Output Inventory Dynamics", *American Economic Journal: Macroeconomics*, 3, 181-212.
- 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.