THE UNIVERSITY OF SYDNEY

Economics Working Paper Series

2009 - 04

Inventory Mistakes and the Great Moderation

James Morley & Aarti Singh

February 2015

Inventory Shocks and the Great Moderation

James Morley[*]
University of New South Wales

Aarti Singh[†]
University of Sydney

February 5, 2015 [‡]

Abstract

Why did the volatility of U.S. real GDP decline by more than the volatility of final sales with the Great Moderation in the mid-1980s? One explanation is that firms shifted their inventory behavior towards a greater emphasis on production smoothing. We investigate the role of inventories in the Great Moderation by estimating an unobserved components model that identifies inventory and sales shocks and their propagation in the aggregate data. Our estimates provide little support for increased production smoothing. Instead, smaller transitory inventory shocks explain the excess volatility reduction in output relative to sales. These shocks behave like informational errors related to production that must be set in advance and their reduction also helps explain the changed forecasting role of inventories since the mid-1980s. Our findings provide an optimistic prognosis for a continuation of the Great Moderation despite the dramatic movements in output during the recent economic crisis.

Keywords: Great Moderation; inventories; production smoothing; unobserved components model. *JEL codes*: E32; E22; C32

---

[*]Corresponding author: School of Economics, UNSW Business School, University of New South Wales, Sydney 2052, Australia; Email: james.morley@unsw.edu.au.

[†]Email: aarti.singh@sydney.edu.au.

[‡]We thank the editor, an anonymous referee, Steve Fazzari, Ben Herzon, Jun Ma, David Papell, Krishna Pendakur, Herman Stekler, Rodney Strachen, Shaun Vahey, seminar and conference participants at George Washington University, Simon Fraser University, Australian National University, University of Adelaide, University of Cincinnati, University of Houston, University of Melbourne, University of New South Wales, University of Sydney, University of Wisconsin-Milwaukee, the SCE Meetings, the SNDE Meetings, the Australian Conference on Quantitative Macroeconomics, and the Australiasian Macroeconomics Workshop for helpful comments and suggestions. The paper was written in part while Morley was a Visiting Scholar in the Economics Department at the University of Sydney. The usual disclaimers apply.

1 Introduction

Lower volatility of the growth rate of the U.S. real GDP since the mid-1980s, first documented by Kim and Nelson (1999) and McConnell and Perez-Quiros (2000), has spurred extensive research into its causes. Better inventory management is often put forth as one of the leading explanations for this so-called "Great Moderation".[1] The emphasis on inventories is motivated by a striking but well-known feature of the aggregate data—output growth was more volatile than sales growth prior to the mid-1980s, but since then output and sales have shared a similar lower level of volatility. Given the accounting relationship between output, sales, and inventory investment, the excess volatility reduction in output relative to sales directly implies some role for inventories in the Great Moderation.

What is it about inventory behavior that has changed? One possible answer is that firms shifted towards a greater emphasis on production smoothing. Golob (2000) finds that the stylized facts in the aggregate data emphasized by Blinder and Maccini (1991) as being so challenging to the relevance of production smoothing have shifted in a more favourable direction in recent years. Kahn, McConnell, and Perez-Quiros (2002) focus on the durable goods sector and find evidence of an improved ability of inventories to forecast future sales, leading them to argue that better information has facilitated increased production smoothing. By contrast, Herrera and Pesavento (2005) consider industry-level manufacturing and trade data and find little evidence of a change in the relationship between inventories and sales.[2]

[1]Other explanations are better monetary policy and smaller macroeconomic shocks (a.k.a. "good luck"). See Clarida, Gali, and Gertler (2000), Stock and Watson (2003), and Ahmed, Levin, and Wilson (2004), among many others.

[2]Irvine and Schuh (2005) find that the unconditional covariances for sales, output, and inventories within and across key industrial sectors have declined in way that is consistent with increased production smoothing. However, Herrera, Murtazashvili, and Pesavento (2009) disentangle the sources of this decline by computing conditional covariances that control for changes in dynamics and find higher, not lower, correlations for sales and inventories across sectors. Meanwhile, McCarthy and Zakrajsek (2007) consider both aggregate and industry-level data in a structural vector autoregressive model and find some evidence

In this paper, we estimate an unobserved components model using aggregate data to help disentangle the role of inventories from that of sales in explaining the decline in the volatility of U.S. real GDP. We find that a change in the sales process explains about half of the overall decline. However, in terms of the excess decline in output volatility relative to sales, we find that it reflects smaller transitory inventory shocks rather than a shift towards greater production smoothing. These shocks behave like informational errors made by firms in response to noisy signals when setting production in advance of sales, and their reduction also helps explain the apparent changed forecasting role of inventories with the Great Moderation.

Our findings have important implications for the much-questioned continuation of the Great Moderation.[3] While inventory "mistakes" due to informational errors will likely continue to be made, their apparent reduction could reflect structural changes in the economy such as improved informational flows and the rise of "just-in-time" production. Thus, even if the Great Moderation were due to smaller shocks rather than changes in their propagation, as emphasized by Stock and Watson (2003), Ahmed, Levin, and Wilson (2004), and many others, the shocks may not just be those that fit under the ephemeral-sounding "good luck" hypothesis. In particular, despite large aggregate shocks during the recent economic crisis, the possible technological and structural reasons for smaller inventory shocks suggest that we should not expect a return to the persistent high levels of output volatility experienced during the 1970s and earlier.

The rest of this paper is organized as follows. Section 2 presents some stylized facts for the aggregate data that motivate our analysis and solves a standard cost minimization problem to provide a theoretical context for

---

of increased production smoothing, but conclude that other factors, such as better policy, explain the bulk of the decline in the volatility of aggregate economic activity.

[3]On this topic, Clark (2009) considers whether the increase in volatility during the Great Recession was as widespread across different sectors of the economy as with the decline in volatility with the Great Moderation. He finds that the increased volatility was largely driven by oil price and financial shocks. Thus, he argues the Great Moderation will continue as the effects of these oil price and financial shocks dissipate.

interpreting our empirical results. Section 3 develops the unobserved components model that we use to disentangle the roles of inventory and sales shocks and their propagation in explaining the Great Moderation. Section 4 reports empirical results. Section 5 discusses the implications of our findings for the continuation of the Great Moderation and concludes.

2 Background

2.1 Output volatility and its components

Output, sales, and inventories are related by the following identity:

$$y_t \equiv s_t + \Delta i_t, \tag{1}$$

where y_t is the natural logarithm of output, s_t is the natural logarithm of sales, and Δi_t is a residual measure of inventory investment.[4] For our empirical analysis, we use quarterly data for the sample period of 1960Q1-2014Q1 from the Bureau of Economic Analysis (BEA) on U.S. real GDP and final sales (NIPA Table 1.2.6) to measure the variables in equation (1). To investigate the role of inventories in the Great Moderation, we split the data into pre- and post-moderation sample periods of 1960Q1-1984Q1 and 1984Q2-2014Q1.[5]

Table 1 reports sample statistics related to the volatility of the first-differences of the variables in equation (1). The first stylized fact to emerge

[4]The true accounting identity is between the levels of output, sales, and inventory investment, while the residual measure is approximately equal to actual inventory investment as a percentage of sales (i.e., $\Delta i_t = ln(1 + \Delta I_t/S_t) \approx \Delta I_t/S_t$, where ΔI_t is the level of inventory investment and S_t is the level of sales). We consider the residual measure because it allows us to relate changes in the estimated sales and inventory processes from our unobserved components model directly to the change in output volatility, which is the primary aim of our analysis. This measure was also considered in Kahn, McConnell, and Perez-Quiros (2002). In Section 4.6, we consider the robustness of our results to using a direct measure of inventory investment based on the first differences of the log stock of inventories (i.e., inventory investment as a percentage of the lagged stock of inventories).

[5]Kim and Nelson (1999) and McConnell and Perez-Quiros (2000) both estimate the structural break in the variance of U.S. real GDP growth to have occurred in 1984Q1. In order to keep our analysis focused, we treat this break date as known for the purposes of estimation, although we note there is some degree of uncertainty about its exact timing (see, for example, Stock and Watson, 2003, and Eo and Morley, forthcoming).

TABLE 1. SAMPLE STATISTICS

	Pre-moderation (1960Q1-1984Q1)	Post-moderation (1984Q2-2014Q1)	Change across samples
s.d.(Δy_t)	1.06	0.61	-0.45
s.d.(Δs_t)	0.83	0.57	-0.26
s.d.($\Delta^2 i_t$)	0.67	0.38	-0.29
corr.($\Delta s_t, \Delta^2 i_t$)	-0.01	-0.23	-0.22

Table 1: Sample standard deviation (s.d.) and correlation (corr.) statistics are reported for the first differences of log output, log sales, and a residual measure of inventory investment based on the difference between log output and log sales. All series are multiplied by 100.

from these sample statistics is that real GDP growth stabilized dramatically in recent years, as has been widely reported in the literature. The second stylized fact is that output was more volatile than sales in the pre-moderation period, but both have a similar lower level of volatility in the post-moderation period, which has also been discussed previously (see, for example, Kahn, McConnell, and Perez-Quiros, 2002, and Golob, 2000).

One possible explanation for these changes in volatility is an increased emphasis on production smoothing by firms. Yet, the results in Table 1 provide mixed signals about the overall relevance of production smoothing. In the pre-moderation period, both the excess volatility of output relative to sales and the lack of a large negative contemporaneous correlation between sales and inventories directly undermine the idea that firms use inventories to buffer production from fluctuations in sales, as emphasized in the survey article by Blinder and Maccini (1991). By contrast, the shift to more similar levels of volatility and a negative contemporaneous correlation between sales and inventories in the post-moderation period is more consistent with production smoothing, as pointed out by Golob (2000). However, the finding that both sales and inventories also became less volatile in the post-moderation period clearly argues against production smoothing as the sole explanation for the Great Moderation. In addition, the fact that output is still no less volatile than sales in the post-moderation period continues to argue against production smoothing as the primary motive for holding inventories.[6]

These mixed signals from Table 1 motivate our development of an unobserved components model of the aggregate data in Section 3 to help disentangle the role of increased production smoothing from other factors in explaining the Great Moderation.

2.2 Inventories and forecasting

Beyond the well-known reduction in volatility, the Great Moderation also corresponded to a change in the forecasting role of inventories (see, for example, Kahn, McConnell, and Perez-Quiros, 2002). Figure 1 motivates why inventory investment is particularly useful for forecasting output and sales. The left panel plots log output and log sales using the BEA data discussed above. Both series are nonstationary, which is easily confirmed by standard unit root and stationarity tests. However, both series appear to share the same stochastic trend. The right panel plots the first-differences of the two series and the difference between the two series, which is the residual measure of inventory investment defined in equation (1). All of these series are stationary, which again is confirmed by standard tests.

More formally, the idea that the residual measure of inventory investment is stationary corresponds to cointegration between log output and log sales with a cointegrating vector of $[1, -1]'$.[7] Cointegration corresponds to the idea that output and sales share the same stochastic trend, which is important because it implies that the cointegrating error term (i.e., the residual

[6] Also, as emphasized by Blinder and Maccini (1991), changes in finished goods inventories, which can be most directly related to the production smoothing motive, are neither the largest nor most volatile component of inventory investment.

[7] Granger and Lee (1989) find evidence of "multicointegration" between output and sales (with vector $[1, -1]'$) and between inventories and sales (with an estimated vector). Their analysis is in terms of levels rather than logarithms and they consider sectoral data on sales and inventories. For the aggregate data considered here, we find stronger evidence of a cointegrating relationship (as measured by the strength of the adjustment to the long-run equilibrium) for log output and log sales than for the levels. Meanwhile, contrary to multicointegration, we do not find any evidence of cointegration between the accumulation of the residual measure of inventory investment and log sales using standard tests.

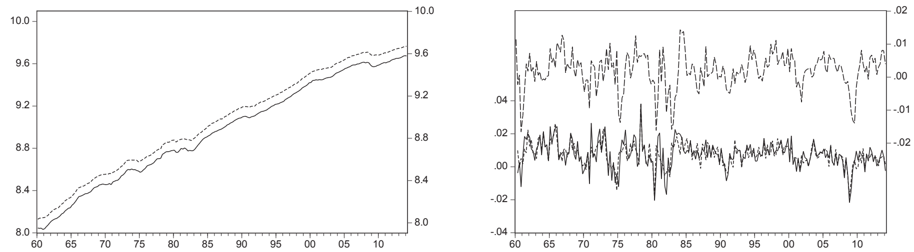

Figure 1: The left panel plots real GDP (solid line, left vertical axis) and final sales (dashed line, right vertical axis), both expressed in natural logarithms. The first differences of the two series (right vertical axis) along with the residual measure of inventory investment (thick dashed line, left vertical axis) are plotted in the right panel. The sample period is 1960Q1-2014Q1.

measure of inventory investment) must forecast future movements in output and/or sales in order for the long-run cointegrating relationship to be restored over time.

We demonstrate the change in the forecasting role of inventories with a simple vector error correction model (VECM) given as follows:

$$\Delta y_t = \gamma_{y,0} + \alpha_y \Delta i_{t-1} + \sum_{j=1}^{p} \gamma_{yy,j} \Delta y_{t-j} + \sum_{j=1}^{p} \gamma_{ys,j} \Delta s_{t-j} + \nu_{y,t}, \qquad (2)$$

$$\Delta s_t = \gamma_{s,0} + \alpha_s \Delta i_{t-1} + \sum_{j=1}^{p} \gamma_{ss,j} \Delta s_{t-j} + \sum_{j=1}^{p} \gamma_{sy,j} \Delta y_{t-j} + \nu_{s,t}, \qquad (3)$$

where the α parameters are error-correction coefficients given that $\Delta i_t = y_t - s_t$, the lagged differences of output and sales capture short-run dynamics, with the lag order p based on the Schwarz Information Criterion (SIC), and the ν shocks are assumed to be white noise.

Table 2 reports the estimates for the error-correction coefficients. In the pre-moderation period, the estimate $\hat{\alpha}_y = -0.68$ suggests that a positive change in inventories predicts a large decline in future output, all else equal. Meanwhile, inventory investment appears to have no significant predictive impact on future sales. The results for the post-moderation period are strikingly different. First, the estimates suggest that a positive

TABLE 2. ERROR CORRECTION COEFFICIENTS

	Pre-moderation (1960Q1-1984Q1)	Post-moderation (1984Q2-2014Q1)
α_y	-0.68 (0.18)	-0.21 (0.15)
α_s	-0.11 (0.16)	0.52 (0.14)

Table 2: OLS estimates are reported, with standard errors in parentheses. SIC selects a lag order of $p = 1$ for the pre-moderation sample and $p = 2$ for the post-moderation sample (and the full sample). The results are qualitatively robust for different numbers of lags and are reported here for $p = 2$, with estimates of the other parameters omitted for simplicity.

change in inventories still predicts a decline in future output, but there is a much smaller estimated effect that is not statistically significant at the 5% level. Second, the estimate $\hat{\alpha}_s = 0.52$ suggests that a positive change in inventories predicts an increase in future sales, all else equal. Therefore, inventories had a strong negative forecasting relationship with future output prior to the Great Moderation, but since then inventories have had a strong positive forecasting relationship with future sales.

At first glance, the finding that inventories forecast future sales in the post-moderation period might seem highly supportive of increased production smoothing. For example, Kahn, McConnell, and Perez-Quiros (2002) hypothesize that improvements in information technology have helped firms anticipate future sales, with inventories being more reflective of intentional production smoothing towards these future sales. However, the forecasting role of inventories might have simply changed due to a different composition of the underlying shocks driving inventory investment. Unfortunately, the role of production smoothing versus a change in the composition of shocks cannot be disentangled from the VECM estimates alone. Again, as with the stylized facts in Table 1, we are motivated by these competing explanations to develop an unobserved components model in Section 3.[8]

[8] Also, the VECM estimates suggest that both output and sales adjust to restore the long-run equilibrium, directly implying the presence of a common unobserved stochastic trend rather than one of the variables acting as the *de facto* trend. This provides an empirical motivation for the structure of our unobserved components model in Section 3.

2.3 Cost minimization

In order to be more formal about the motives for holding inventories and to provide some context for understanding our empirical results, we consider a standard linear-quadratic cost minimization problem to solve for optimal inventory management, similar to Blanchard (1983), West (1986, 1990), Ramey and West (1999), and Hamilton (2002), among many others.[9] Letting I_t denote the stock of inventories and C_t denote costs, the representative firm chooses a path for inventories to minimize its expected discounted costs over the infinite horizon:

$$\min_{\{I_{t+j}\}_{j=0}^{\infty}} E_t \sum_{j=0}^{\infty} \beta^j C_{t+j}, \tag{4}$$

where the discount factor $0 < \beta < 1$. The cost function is

$$\begin{aligned} C_t = 0.5a_1(\Delta Y_t)^2 + 0.5a_2(Y_t - S_t^*)^2 + 0.5a_3(\Delta I_t)^2 \\ + 0.5a_4(I_{t-1} - a_5 S_t^*)^2 + u_{c,t} Y_t, \end{aligned} \tag{5}$$

where Y_t is the level of output, S_t^* is the long-run level of sales (and the permanent component of the marginal cost of production, as in Hamilton, 2002), $u_{c,t}$ is a transitory marginal cost shock, and cost coefficients $a_i > 0$ for $i = 1, ..., 5$.

The costs motivating production smoothing are given by the first two terms in equation (5). Specifically, $(\Delta Y_t)^2$ captures the idea that it is costly to change production in the short run and $(Y_t - S_t^*)^2$ captures the idea that it is costly to have output far away from the long-run level of sales S_t^*.

[9] The cost minimization problem is a version of the Holt, Modigliani, Muth, and Herbert's (1960) partial equilibrium linear-quadratic framework characterizing inventory decisions at the firm level. Davis and Kahn (2008), Blinder and Maccini (1991), and others have pointed out that the linear-quadratic framework is more applicable for finished goods inventories than for inventories of materials and supplies held by manufactures, which are arguably better captured by an (S,s) model. However, Ramey and West (1999) argue against such a literal interpretation of the cost function for the representative firm. Also, as discussed in Blinder and Maccini (1991), the (S,s) model cannot be easily applied to study aggregate inventory dynamics. Meanwhile, see Wen (2005) for general equilibrium analysis of production smoothing and stockout avoidance motives for holding inventories.

The costs motivating stockout avoidance are given by the third and fourth terms, where $(\Delta I_t)^2$ captures the idea that it is costly to draw down from or add to the stock of inventories in the short run and $(I_{t-1} - a_5 S_t^*)^2$ captures the idea that it is costly to have inventories away from a long-run level that depends positively on the long-run level of sales.[10]

For our partial equilibrium analysis, we treat prices and sales as exogenous. Thus, given the level of sales, denoted by S_t, the representative firm's inventory choices determine output according to the inventory identity:

$$Y_t = S_t + \Delta I_t. \tag{6}$$

To complete the model, we need to specify the cost shock and the sales process, including its long-run level. First, we assume that the cost shock, $u_{c,t}$, is white noise and independent of sales.[11] Second, we assume the level of sales has permanent and transitory components, $S_t = S_t^* + e_{s,t}$, with $S_t^* = S_{t-1}^* + e_{p,t}$ corresponding to the stochastic trend in sales, where $e_{p,t}$ denotes a permanent sales shock, and transitory sales are driven by the transitory sales shock, $e_{s,t}$. As with the cost shock, the sales shocks are assumed to be white noise.

We can generate the key theoretical results by focusing on the long-run motives. Thus, for simplicity, we set $a_1 = a_3 = 0$ for the time being, although we will return to the short-run motives later when interpreting some of the parameters in our empirical model. Then, the first-order con-

[10] For simplicity, we consider a continuous and symmetric version of the stockout avoidance motive. Instead of just being concerned with a literal "stockout" (i.e., having insufficient inventories to satisfy a large positive sales shock), which would correspond to a discrete and asymmetric specification for the cost, we assume that the representative firm implicitly has a large enough stock of inventories to satisfy any given sales shock, but the cost of doing so increases exponentially with the size of the shock.

[11] Because the cost shocks is multiplied by output in the cost function, its impact will be related to the scale of the economy. Meanwhile, the assumption that the cost shock is independent of sales does not mean that all exogenous changes in marginal costs are independent of sales. Following Hamilton (2002), we assume that exogenous changes in costs that do impact sales are reflected in aggregate sales shocks. This abstracts from a likely asymmetry that negative sales shocks related to events such as spikes in oil prices may affect costs by more than equivalent sized positive sales shocks.

dition for the simplified cost minimization problem is

$$E_t \left[\begin{array}{c} a_2\{(Y_t - S_t^*) - \beta\left(Y_{t+1} - S_{t+1}^*\right)\} \\ +\beta a_4\{I_t - a_5 S_{t+1}^*\} + u_{c,t} - \beta u_{c,t+1} \end{array} \right] = 0. \tag{7}$$

Substituting for output using the inventory identity, setting future shocks to their expected value of zero, and dividing through by a_2 gives us

$$E_t \left[\Delta I_t - \beta \Delta I_{t+1} + \theta W_t + e_{s,t} + \frac{1}{a_2} u_{c,t} \right] = 0, \tag{8}$$

where $\theta = \beta a_4 / a_2$ and $W_t = I_t - a_5 S_t^*$ is the deviation of the stock of inventories from its long-run level. Rewriting the above equation in terms of W_t, we have

$$E_t[(1 + \beta + \theta) W_t - W_{t-1} - \beta W_{t+1}] = z_t, \tag{9}$$

where $z_t = -a_5 e_{p,t} - e_{s,t} - (1/a_2) u_{c,t}$. Finally, following Hansen and Sargent (1980), we solve the above polynomial to get

$$W_t = \varphi W_{t-1} + \varphi z_t, \tag{10}$$

where $\varphi = \frac{\delta - (\delta^2 - 4\beta^{-1})^{1/2}}{2}$ is the stable root of the polynomial, with $\delta = \beta^{-1}(1 + \beta + \theta)$. Because φ is the stable root, W_t will be stationary. Furthermore, it can be shown that $0 < \varphi < 1$ given that the assumptions on the cost coefficients and the discount factor.

Because W_t is stationary, the first key result is that inventories and sales will be cointegrated with vector $(1, -a_5)'$, similar to the theoretical result in Hamilton (2002).[12] It is straightforward, then, to show that log inventories

[12] In standard linear-quadratic models, such as Ramey and West (1999), the production cost using our notation is $0.5 a_2 Y_t^2 + u_{c,t} Y_t$ (assuming no cost of adjusting production, $a_1 = 0$). Then, when output is nonstationary, the marginal cost of production goes to infinity. In particular, the cointegrating relationship between inventories and sales is not $(1, -a_5)'$, but is instead $(1, (((1 - \beta) a_2)/(\beta a_4)) - a_5)'$. By contrast, abstracting from the short-run cost to inventory adjustment and the incorporation of transitory sales shocks, our model is identical to Hamilton's (2002) model and, therefore, the marginal cost of production remains finite and, again, the cointegrating relationship is $(1, -a_5)'$. Meanwhile, we incorporate the short-run cost to inventory adjustment in order to consider competing short-run motives for holding inventories in Section 3.3 and transitory sales shocks given their empirical relevance, as seen in Section 4.2. Of course, given their short-run nature, these modifications from Hamilton (2002) do not affect the cointegration result.

and log sales will be cointegrated with vector $[1, -1]'$, a theoretical result that informs the specification of our empirical model in the next section. Furthermore, the fact that W_t depends on permanent sales shocks in addition to transitory cost and sales shocks informs our allowance of permanent shocks to affect transitory deviations from trend in our empirical model.

The second key result is that the inventory investment is stationary and its persistence depends on the relative costs associated with production smoothing and stockout avoidance. In particular, we can use equation (10) and the sales process to solve for the following ARMA(1,1) process for inventory investment:

$$(1 - \varphi L)\Delta I_t = a_5(1 - \varphi)e_{p,t} - \varphi(1 - L)e_{s,t} - \frac{1}{a_2}\varphi(1 - L)u_{c,t}, \tag{11}$$

where L is the lag operator. Again, because φ is the stable root, inventory investment will be stationary, which directly implies from the inventory identity and the sales process that $Y_t - S_t^*$ will also be stationary. Furthermore, because $0 < \varphi < 1$, inventory investment will be persistent, with the persistence increasing in a_2, the long-run cost that motivates production smoothing, and decreasing a_4, the long-run cost that motivates stockout avoidance.[13]

The third key result is that both the relative volatility of output and sales and the correlation between sales and inventories depend on the importance of the underlying shocks. Based on the sales process, the inventory identity, and the optimal inventory investment process in equation (11), we can solve for the relative volatilities and correlations depending on which

[13]The persistence of inventory investment is strongly evident in the data. Specifically, the first sample autocorrelation for inventory investment constructed using the levels (rather than logs) of the quarterly output and sales data considered above is 0.46 for the pre-moderation sample and 0.68 for the post-moderation sample. Given the sample sizes, these estimates are strongly significant. In terms of the effects of the production smoothing and stockout avoidance motives on persistence, the comparative statics are based on the following partial derivatives: $\frac{\partial \varphi}{\partial a_2} = \frac{\partial \varphi}{\partial \delta}\frac{\partial \delta}{\partial \theta}\frac{\partial \theta}{\partial a_2} = -\frac{a_4}{2a_2^2}(1 - \delta(\delta^2 - 4\beta^{-1})^{-1/2})$ and $\frac{\partial \varphi}{\partial a_4} = \frac{\partial \varphi}{\partial \delta}\frac{\partial \delta}{\partial \theta}\frac{\partial \theta}{\partial a_4} = \frac{1}{2a_2}(1 - \delta(\delta^2 - 4\beta^{-1})^{-1/2})$. Given $(\delta^2 - 4\beta^{-1})^{1/2} > 0$ for a real root, $0 < \varphi < 1$, and the assumptions for the cost coefficients, it is straightforward to show that $\frac{\partial \varphi}{\partial a_2} > 0$ and $\frac{\partial \varphi}{\partial a_4} < 0$.

shocks prevail. First, consider only permanent sales shocks, setting the variances of the other shocks to zero. Then, we get the following variance ratio and covariance expressions:

$$\frac{var(\Delta Y_t)}{var(\Delta S_t)} = 1 + 2a_5(1-\varphi) + 2a_5^2\frac{(1-\varphi)^2}{1+\varphi} \tag{12}$$

$$\frac{cov(\Delta S_t, \Delta^2 I_t)}{var(\Delta S_t)} = a_5(1-\varphi) \tag{13}$$

It is straightforward to show that the variance ratio is strictly greater than one. Meanwhile, given that $a_5 > 0$ and $0 < \varphi < 1$, the covariance will be positive. Thus, a large role for permanent sales shocks could explain the excess volatility of output relative to sales for the pre-moderation period reported in Table 1, although it does not explain the correlation results in either of the pre- or post-moderation periods. Second, consider only transitory sales shocks:

$$\frac{var(\Delta Y_t)}{var(\Delta S_t)} = \frac{(1-\varphi)^2}{1+\varphi} \tag{14}$$

$$\frac{cov(\Delta S_t, \Delta^2 I_t)}{var(\Delta S_t)} = \frac{\varphi(\varphi-3)}{2} \tag{15}$$

Given $0 < \varphi < 1$, the variance ratio will be strictly less than one and the covariance will be negative. Thus, a large role for transitory sales shocks could explain the negative correlation in the post-moderation period, although it cannot explain the relative volatilities in either of the pre- or post-moderation periods. Finally, consider only cost shocks:

$$\frac{var(\Delta Y_t)}{var(u_{c,t})} = \frac{2\varphi^2(3-\varphi)}{a_2^2(1+\varphi)} \tag{16}$$

$$cov(\Delta S_t, \Delta^2 I_t) = 0 \tag{17}$$

In this case, we cannot standardize by the variance of sales because it is zero. But it is trivial to see that output will be more volatile than sales and the covariance of inventories and sales will be zero. Thus, a large role for cost shocks could explain the relatively high volatility of output and the lack of correlation between sales and inventories in the pre-moderation period, although it cannot explain the results in the post-moderation period.

The fourth key result is that the implied forecasting role of inventory investment depends on the importance of the underlying shocks. We calculate the partial effects of an unpredictable change in inventories on future output growth and future sales growth given a prevailing shock. To do this, we compute the ratio of the marginal effects of each shock on future output and sales growth and on inventories. First, consider a permanent sales shock:

$$\frac{\partial \Delta Y_{t+1}}{\partial e_{p,t}} \Big/ \frac{\partial \Delta I_t}{\partial e_{p,t}} = \varphi \tag{18}$$

$$\frac{\partial \Delta S_{t+1}}{\partial e_{p,t}} \Big/ \frac{\partial \Delta I_t}{\partial e_{p,t}} = 0 \tag{19}$$

Given $0 < \varphi < 1$, inventory investment will have a positive forecasting relationship with future output growth, while there is no forecasting relationship with future sales growth. Thus, a large role for permanent sales shocks is not consistent with the any of the VECM results in Table 2, except perhaps the small response of sales in the pre-moderation period. Second, consider a transitory sales shock:

$$\frac{\partial \Delta Y_{t+1}}{\partial e_{s,t}} \Big/ \frac{\partial \Delta I_t}{\partial e_{s,t}} = \frac{1 - \varphi + \varphi^2}{\varphi} \tag{20}$$

$$\frac{\partial \Delta S_{t+1}}{\partial e_{s,t}} \Big/ \frac{\partial \Delta I_t}{\partial e_{s,t}} = \frac{1}{\varphi} \tag{21}$$

Again, given $0 < \varphi < 1$, it is straightforward to show that the first expression is positive, while the second expression is positive and greater than one. Thus, a large role for transitory sales shocks is consistent with the post-moderation VECM results. Finally, consider a cost shock:

$$\frac{\partial \Delta Y_{t+1}}{\partial u_{c,t}} \Big/ \frac{\partial \Delta I_t}{\partial u_{c,t}} = \varphi - 1 \tag{22}$$

$$\frac{\partial \Delta S_{t+1}}{\partial u_{c,t}} \Big/ \frac{\partial \Delta I_t}{\partial u_{c,t}} = 0 \tag{23}$$

In this case, the first expression is negative. Thus, a large role for cost shocks is consistent with the pre-moderation VECM results for both output and sales growth.

Based on this cost minimization analysis, a change in inventory behavior could reflect a change in the relative costs motivating production smoothing versus stockout avoidance and/or a change in the sales process. For example, a simple explanation for the excess decline in output volatility presented in Table 1 would be a *relative* reduction in the costs motivating stockout avoidance (i.e., a reduction in costs of accessing inventory stocks compared to costs of changing production plans). A simple explanation for the change in the forecasting role of inventories with respect to future sales presented in Table 2 would be a change in composition of sales shocks, with permanent shocks becoming relatively more important.

As discussed in Blinder and Maccini (1991) and Kahn, McConnell, and Perez-Quiros (2002), the nature of informational flows in the production process is such that some changes in inventories will be unintentional and unrelated to actual sales rather than optimal responses to sales shocks–i.e., they will be informational errors that produce inventory "mistakes". Beyond the timing for the long-run stockout avoidance term in the cost function, it might seem that our simple cost minimization analysis abstracts from the fact that some production must be set in advance based on noisy signals about sales.[14] However, inventory mistakes provide a leading example of cost shocks that should not affect the sales process. Specifically, they are costly to make, but they should not alter the path of sales (if they are truly mistakes) and, under rational expectations, the path of sales (or anything else) should not imply any predictability in the informational er-

[14]The tradeoff between production smoothing and stockout avoidance can be seen as capturing the idea that it is less costly to set production in advance than at the moment sales are realized. Specifically, the costs associated with accumulating or depleting inventories (i.e., with the stockout avoidance motive) only need to be borne if a firm also finds it costly to change production when a sales shock is realized. Otherwise, the firm will simply adjust production in response to the shock, thus avoiding the costs associated with accessing inventories. Consequently, the key abstraction in the cost minimization analysis is in terms of the information flows about sales, rather than setting production in advance.

rors. Thus, inventory mistakes should have the same effects as the cost shocks in our cost minimization analysis, including possibly explaining the excess volatility of output relative to sales and the negative forecasting relationship between inventories and future output in the pre-moderation data. The key question addressed in this paper, then, is how important are the transitory inventory shocks that reflect inventory mistakes and any other cost shocks that do not affect sales in explaining the Great Moderation relative to changes in the sales process or increased production smoothing. Again, to answer this question and to help sort out the competing explanations for the basic sample statistics and the VECM results, we develop an unobserved components model in the next section to identify inventory shocks, changes in the sales process, and intentional responses of inventories to the sales process from the aggregate data.

3 Model

3.1 An unobserved components model

Our unobserved components (UC) model separates each of the observable series for log output, log sales, and a measure of log inventories (derived as an accumulation of the residual measure of inventory investment) into a permanent component and a transitory deviation from the permanent component:

$$y_t = \tau_t^* + (y_t - \tau_t^*), \tag{24}$$

$$s_t = \tau_t^* + (s_t - \tau_t^*), \tag{25}$$

$$i_t = i_t^* + (i_t - i_t^*). \tag{26}$$

The permanent components are specified as follows:

$$i_t^* = \tau_t^* + \kappa_t, \tag{27}$$

$$\tau_t^* = \mu_\tau + \tau_{t-1}^* + \eta_t, \qquad \eta_t \sim i.i.d.N(0, \sigma_\eta), \tag{28}$$

$$\kappa_t = \mu_\kappa + \kappa_{t-1} + \lambda_{\kappa\eta}\eta_t + \omega_t \qquad \omega_t \sim i.i.d.N(0, \sigma_\omega), \tag{29}$$

where i_t^* is the trend for inventories, τ_t^* is the common trend for output and sales, which also drives inventories, as in Section 2.3, and κ_t is an additional trend that allows for permanent changes in the inventory-sales ratio. The trends have deterministic drifts μ_τ and μ_κ, respectively, and they are driven by η_t, the permanent sales shock, and ω_t, the permanent shock to the inventory-sales ratio, respectively.

The specification of a common stochastic trend for output and sales is based on the empirical and theoretical results in Sections 2.2 and 2.3 that log output and log sales at the aggregate level are cointegrated with vector $[1-1]'$. The additional trend, κ_t, captures the empirical result that our measure of inventories is not cointegrated with output or sales and could reflect time variation in the accelerator parameter a_5 in our theoretical model. Because the residual-based measure of inventories (approximately) corresponds to an accumulation of the inventory-investment-to-sales ratio, we also allow for permanent sales shock to affect the additional trend via the impact coefficient $\lambda_{\kappa\eta}$ in equation (29) in order to capture the implied correlation between sales growth and the accumulated inventory-sales ratio.[15]

The transitory components follow stationary processes:

$$\psi_y(L)^{-1}(y_t - \tau_t^*) = \lambda_{y\eta}\eta_t + \lambda_{y\omega}\omega_t + \lambda_{y\epsilon}\epsilon_t + v_t, \tag{30}$$

$$\psi_s(L)^{-1}(s_t - \tau_t^*) = \lambda_{s\eta}\eta_t + \epsilon_t, \tag{31}$$

$$\psi_i(L)^{-1}(i_t - i_t^*) = \lambda_{i\eta}\eta_t + \lambda_{i\omega}\omega_t + \lambda_{i\epsilon}\epsilon_t + v_t, \tag{32}$$

where the $\psi(L)$ lag polynomials capture invertible Wold coefficients and $\lambda_{y\eta}$, $\lambda_{y\omega}$, $\lambda_{y\epsilon}$, $\lambda_{s\eta}$, $\lambda_{i\eta}$, $\lambda_{i\omega}$, and $\lambda_{i\epsilon}$ are the impact coefficients for transitory output, sales, and inventories in response to the shocks. The transitory shocks are $\epsilon_t \sim i.i.d.N(0, \sigma_\epsilon)$, and $v_t \sim i.i.d.N(0, \sigma_v)$, where ϵ_t is a transitory sales shock and v_t is a transitory inventory shock, which, as discussed in more detail in Section 3.2, could reflect informational errors.

For this UC model, the transitory deviations from trend are driven not only by transitory shocks, but also by adjustments to permanent shocks,

[15] The impact of transitory sales on the additional trend was not significant in a more general version of the model and so was dropped for simplicity.

which is consistent, for example, with the optimal path for inventories solved in Section 2.3. By assuming this flexible structure, permanent and transitory movements are allowed to be correlated, even though the underlying shocks are specified to be mutually uncorrelated. As discussed in Morley, Nelson, and Zivot (2003), a UC model with correlated components is identified given sufficiently rich dynamics. For our application, we estimate the model for sales and inventories, assuming AR(2) dynamics for their transitory components (i.e., $\psi_s(L)^{-1} = 1 - \phi_{s,1}L - \phi_{s,2}L^2$ and $\psi_i(L)^{-1} = 1 - \phi_{i,1}L - \phi_{i,2}L^2$, with roots of the AR polynomials lying strictly outside the unit circle to ensure stationarity) and leaving the process for output implicit. The two-variable UC model has 15 independent parameters and corresponds to a reduced-form vector autoregressive moving-average (VARMA) process with 17 parameters.[16] As a result, the model is identified, although weak identification is still a potential problem, as discussed in more detail in Section 4.1. A state-space representation of the UC model is presented in Appendix A.

3.2 Interpretation of shocks

The economic interpretation of most shocks is straightforward. Permanent and transitory sales shocks, η_t and ϵ_t, should capture technology and demand factors in the aggregate economy. The permanent inventory shocks, ω_t, should capture any permanent changes in the inventory-sales ratio due to evolving inventory management practices, shifts in the nature of production (e.g., from goods to services), and changes in the costs of accessing and holding inventories. The transitory inventory shocks, v_t, should capture informational errors that arise due to noisy signals firms receive about sales in conjunction with the fact that some production must be set in advance of

[16]There are four AR parameters and two drift terms that are common to both specifications. In addition, the two-variable UC model has four variance parameters and five independent impact coefficients (see Section 3.3 below), while the VARMA model has three variance-covariance parameters and eight MA parameters associated with two-lags of vector MA terms. Note that, because sales and inventories are not restricted to be cointegrated, our multivariate UC model is more analogous to Sinclair (2009) than to Morley (2007).

sales. Kahn, McConnell, and Perez-Quiros (2002) also consider similar unintentional inventory shocks and note that their magnitude reflects both the flow of information about future sales and the extent to which production needs to be set in advance. For example, a firm may regard an order as a signal of future sales and begin production on this basis, but the order may be unrelated to actual sales and, therefore, subsequently cancelled. To the extent that the firm increased production based on this order, the cancellation was not predicted and the resulting inventory accumulation will be a mistake. Meanwhile, to the extent that production can be held off closer to the date of the actual sale, fewer inventory mistakes will be made.

The key distinction between the transitory sales shocks and transitory inventory shocks in the UC model is that inventory shocks are assumed to have no direct impact on future sales. Unexpected changes in inventories which do affect aggregate demand will be classified as sales shocks, as will temporary cost shocks that have aggregate effects, including temporary shocks to productivity (e.g., Miron and Zeldes, 1989, and Hamilton, 2002) or input cost shocks (e.g., Maccini, Moore, and Schaller, 2012). Consequently, any cost shocks that do not affect aggregate sales will also be categorized as transitory inventory shocks. However, our conjecture is that, at the aggregate level, most other cost shocks (e.g., oil price shocks) should have an impact on sales and would be captured by the sales shocks in our UC model.

3.3 The impact coefficients

Output, sales, and inventory investment are linked together by equation (1). As a result, only a subset of the impact coefficients are, in fact, independent. For the UC model, the following equations describe the relationships between the coefficients implied by equation (1):

$$\lambda_{y\eta} = 1 + \lambda_{\kappa\eta} + \lambda_{i\eta} + \lambda_{s\eta}, \tag{33}$$

$$\lambda_{y\epsilon} = 1 + \lambda_{i\epsilon}, \tag{34}$$

$$\lambda_{y\omega} = 1 + \lambda_{i\omega}. \tag{35}$$

Therefore, only five of the eight impact coefficients in the UC model are independently determined.

We impose additional restrictions on the values of the independent impact coefficients based on how output, sales, and inventories respond to exogenous shocks. For example, consider "scenario A" of a positive permanent sales shock to the common stochastic trend τ_t^*. Under this scenario, permanent sales will increase one for one. If actual sales do not change, sales will fall below trend and $\lambda_{s\eta} = -1$. By contrast, if sales increase by the same amount as permanent sales, either due to a ramping up of production and/or due to a running down of existing inventories, then $\lambda_{s\eta} = 0$. Based on these extreme cases, we can bound $\lambda_{s\eta} \in [-1, 0]$. This scenario implies that permanent inventories rise by $1 + \lambda_{\kappa\eta}$ with permanent sales. Meanwhile, if inventories adjust immediately, $\lambda_{i\eta} = 0$. Or, if inventories remain unchanged, then they will be below their long-run target and $\lambda_{i\eta} = -1 - \lambda_{\kappa\eta}$. It is even possible that inventories temporarily decrease if sales adjust but output does not, in which case $\lambda_{i\eta} = -2 - \lambda_{\kappa\eta}$. As a result, $\lambda_{i\eta} \in [-2 - \lambda_{\kappa\eta}, 0]$ if $\lambda_{\kappa\eta} >= -1$ and $\lambda_{i\eta} \in [0, -1 - \lambda_{\kappa\eta}]$ if $\lambda_{\kappa\eta} < -1$. From equation (33) and the bounds on $\lambda_{s\eta}$, this implies $\lambda_{y\eta} \in [-1, 1 + \lambda_{\kappa\eta}, 0]$ if $\lambda_{\kappa\eta} >= -1$ and $\lambda_{y\eta} \in [\lambda_{\kappa\eta,0}, 0]$ if $\lambda_{\kappa\eta} < -1$. For reference, the bounds are presented in Table 3.

The possible values of the impact coefficients for the ϵ_t and v_t shocks are more straightforward to analyze. A positive temporary sales shock, which we label as "scenario B", leads sales to rise temporarily above their long-run target. If $\lambda_{i\epsilon} = -1$, output remains unchanged and the increase in sales is entirely accommodated by a decline in inventories. However, if output rises and inventories remain unchanged, then $\lambda_{i\epsilon} = 0$. Thus, $\lambda_{i\epsilon} \in [-1, 0]$, which from equation (34) implies $\lambda_{y\epsilon} \in [0, 1]$. Meanwhile, a positive shock to the long-run level of inventories, which we label as "scenario C", raises i_t^* one for one. If output does not change then $\lambda_{i\omega} = -1$. However, if output does respond, $\lambda_{i\omega} = 0$. Thus, $\lambda_{i\omega} \in [-1, 0]$, which from equation

(35) implies $\lambda_{y\omega} \in [0, 1]$.

The cost function analysis in Section 2 allows us to relate the different reasons a firm holds inventories to the various impact coefficients. Table 3 reports the implied values of the impact coefficients that are consistent with the production smoothing and stockout avoidance motives under the different scenarios considered above. For simplicity, we focus our discussion on the long-run motives, although the table also reports the implied values of the impact coefficients for the short-run motives. As before, consider scenario A of a positive permanent shock to sales. Suppose actual sales increase such that $\lambda_{s\eta} = 0$ (see the left columns in panel (ii)). In this case, if a firm solely wants to smooth production in the long run, it will increase output and slowly adjust it to the new long-run target such that $\lambda_{y\eta} = 0$ and $\lambda_{i\eta} = -1 - \lambda_{\kappa\eta}$. But if a firm is solely guided by the stockout avoidance motive, it will increase output to accommodate the increase in sales and also restore inventories to their long-run target such that $\lambda_{y\eta} = 1 + \lambda_{\kappa\eta}$ and $\lambda_{i\eta} = 0$. Meanwhile, consider the case where actual sales remain unchanged after a positive permanent shock to sales and $\lambda_{s\eta} = -1$ (see the right columns in panel (ii)). To smooth production, a firm will increase output to minimize deviations from target with $\lambda_{y\eta} = 0$ and $\lambda_{i\eta} = -\lambda_{\kappa\eta}$, while to avoid stockouts, it will restore inventories to their long-run target, $\lambda_{i\eta} = 0$ and $\lambda_{y\eta} = \lambda_{\kappa\eta}$. The implications under scenario B of a temporary sales shock and scenario C of a shock to the long-run level of gross inventories are once again more straightforward. The impact coefficients will be $\lambda_{i\epsilon} = \lambda_{i\omega} = -1$ when a firm is guided solely by a desire to smooth production and $\lambda_{i_\epsilon} = \lambda_{i\omega} = 0$ when it is guided solely by fear of stockouts. The short-run motives reported in panel (i) are determined in a similar fashion.

3.4 Implied forecast errors and forecasting

To the extent that the transitory inventory shocks reflect informational errors, they should be related to forecast errors for inventories. However, there is an important distinction between inventory mistakes, as captured by transitory inventory shocks in the UC model, and the overall forecast

TABLE 3. INVENTORY SHOCKS AND IMPACT COEFFICIENTS

(i) Short-run motives

Scenario A: Permanent shock to sales

	$\lambda_{s\eta} = 0$		$\lambda_{s\eta} = -1$	
	PS	SA	PS	SA
$\lambda_{y\eta}$	-1	0	-1	-1
$\lambda_{i\eta}$	$-2 - \lambda_{\kappa\eta}$	$-1 - \lambda_{\kappa\eta}$	$-1 - \lambda_{\kappa\eta}$	$-1 - \lambda_{\kappa\eta}$

Scenario B: Temporary shock to sales

	PS	SA
$\lambda_{y\epsilon}$	0	1
$\lambda_{i\epsilon}$	-1	0

Scenario C: Permanent shock to inventories

	PS	SA
$\lambda_{y\omega}$	0	0
$\lambda_{i\omega}$	-1	-1

(ii) Long-run motives

Scenario A: Permanent shock to sales

	$\lambda_{s\eta} = 0$		$\lambda_{s\eta} = -1$	
	PS	SA	PS	SA
$\lambda_{y\eta}$	0	$1 + \lambda_{\kappa\eta}$	0	$\lambda_{\kappa\eta}$
$\lambda_{i\eta}$	$-1 - \lambda_{\kappa\eta}$	0	$-\lambda_{\kappa\eta}$	0

Scenario B: Temporary shock to sales

	PS	SA
$\lambda_{y\epsilon}$	0	1
$\lambda_{i\epsilon}$	-1	0

Scenario C: Permanent shock to inventories

	PS	SA
$\lambda_{y\omega}$	0	1
$\lambda_{i\omega}$	-1	0

Table 3: Implied impact coefficients for different shocks are presented for production smoothing (PS) versus stockout avoidance (SA) motives. In the short run, PS corresponds to $\Delta y_t = 0$ and SA corresponds to $\Delta i_t = 0$. In the long run, PS corresponds to $y_t - \tau_t^* = 0$ and SA corresponds to $i_t - i_t^* = 0$.

error for inventory investment. Indeed, this distinction explains why the UC model is particularly helpful in examining the role of inventories in the Great Moderation and the changed forecasting role of inventories.

We define an inventory forecast error, or period-to-period "unexpected" inventories as

$$\Delta i_t^u \equiv \Delta i_t - E_{t-1}[\Delta i_t], \tag{36}$$

where Δi_t is the actual change in inventories and $E_{t-1}[\Delta i_t]$ is the expected change in inventories. Assuming firms observe the underlying shocks hitting the economy and have rational expectations, the UC model implies the following structure for these forecast errors:

$$\Delta i_t^u = y_t - s_t - E_{t-1}[y_t - s_t] = (\lambda_{y\eta} - \lambda_{s\eta})\eta_t + (\lambda_{y\epsilon} - 1)\epsilon_t + \lambda_{y\omega}\omega_t + v_t. \tag{37}$$

The inventory forecast error depends on sales and inventory shocks at date t. However, only part of the forecast error would be due to informational errors based on noisy signals, as captured by the transitory inventory shock v_t. For the other shocks, firms implicitly choose how to respond via the impact coefficients, where these coefficients reflect a desire to smooth production versus a fear of stockouts, as discussed in the previous subsection. For instance, again consider scenario A of a positive permanent sales shock. Depending on how much sales immediately adjust to a permanent shock and firms' motives, there will be accumulation of inventories in the current period by a factor of $(\lambda_{y\eta} - \lambda_{s\eta})$ and this factor is what makes this accumulation intentional.

How does the UC model help in understanding the changed forecasting role of inventories captured by the VECM results in Table 2? One explanation for the results is that inventory changes are more predictable and they provide a better signal of future sales. We consider this possibility by calculating and comparing the variances of the inventory forecast errors and expected inventory investment (i.e., $\Delta i_t^e = \Delta i_t - \Delta i_t^u = E_{t-1}[\Delta i_t]$). Appendix B describes how we calculate these variances for our UC model.

Another explanation for the changed forecasting role is that the com-

TABLE 4. MARGINAL EFFECTS OF SHOCKS ON FORECASTS

	Permanent shocks		Transitory shocks	
	η_t	ω_t	ϵ_t	v_t
$\frac{\partial \Delta y_{t+1}}{\partial \Delta i_t^u}$	$\frac{\lambda_{s\eta}(\phi_{s,1}-1)+\lambda_{i\eta}(\phi_{i,1}-2)-1-\lambda_{\kappa\eta}}{1+\lambda_{i\eta}+\lambda_{\kappa\eta}}$	$\frac{\lambda_{i\omega}(\phi_{i,1}-2)-1}{1+\lambda_{i\omega}}$	$\frac{(\phi_{s,1}-1)+\lambda_{i\epsilon}(\phi_{i,1}-2)}{\lambda_{i\epsilon}}$	$\phi_{i,1}-1$
$\frac{\partial \Delta s_{t+1}}{\partial \Delta i_t^u}$	$\frac{\lambda_{s\eta}(\phi_{s,1}-1)}{1+\lambda_{i\eta}+\lambda_{\kappa\eta}}$	0	$\frac{(\phi_{s,1}-1)}{\lambda_{i\epsilon}}$	0

Table 4: Marginal effects of the underlying shocks on forecast errors and forecasts of future output and sales growth are presented.

position of underlying shocks in an inventory forecast error has changed, with inventory mistakes, as captured by transitory inventory shocks, playing a smaller role. In order to investigate the effects of a change in the composition of shocks and, therefore, relate the UC model to the VECM results, we solve for the partial effects of an inventory forecast error on future output growth and future sales growth: $\frac{\partial \Delta y_{t+1}}{\partial \Delta i_t^u}$ and $\frac{\partial \Delta s_{t+1}}{\partial \Delta i_t^u}$. To do this, we first analytically compute the following marginal effects: (i) impact of each shock on future output and sales growth and (ii) the impact of each shock on an inventory forecast error. Taking the ratio of these marginal effects, we calculate the impact of an inventory forecast error on output growth and sales growth due to a particular shock, holding all else equal. This is similar to the analysis in Section 2.3 on the forecasting implications of different shocks in the cost minimization problem. Table 4 presents the implied partial effects of a forecast error, which are clearly different for the various underlying shocks. Thus, a change in the relative importance of these shocks directly implies a change in the reduced-form forecasting implications of inventories.

4 Empirical results

4.1 Data and methods

As considered in Section 2, the raw data are quarterly U.S. real GDP and final sales from the BEA for the sample periods of 1960Q1-1984Q1 and 1984Q2-2014Q1. We estimate the UC model for sales and inventories, leaving the estimated process for output implicit. Our measure for sales is 100 times the natural logarithms of real sales and our measure for inventories is calculated by i) constructing a residual measure of the change in inventories based on the identity given in equation (1) for 100 times log output and 100 times log sales and ii) accumulating changes given an arbitrary initial level of log inventories.

We estimate the UC model using Bayesian posterior simulation based on Markov-chain Monte Carlo (MCMC) methods. Specifically, we consider a multi-block random-walk chain version of the Metropolis-Hastings (MH) algorithm with 500,000 draws after a burn-in of 20,000 draws. We check the robustness of our posterior moments to different runs of the chain and for different starting values. The multi-block setup allows us to obtain relatively low correlations between parameter draws, suggesting the sampler is working well. See Chib and Greenberg (1995) for more details on the MH algorithm.

There are two reasons why we consider Bayesian estimation. First, UC models can suffer from weak identification. In particular, UC models are closely related to VARMA models, which are notoriously difficult to estimate due to the problem of near cancellation of AR and MA terms and multiple modes for the likelihood surface. A particularly troublesome estimation difficulty is the so-called "pile-up problem" whereby maximum likelihood estimates tend to hit boundaries even when true parameters are not equal to the boundary values. Preliminary analysis via maximum likelihood estimation (MLE) confirmed multiple modes and some pile-up problems. By contrast, Bayesian estimation with relatively uninformative priors reveals a clear interior mode for the posterior function. Our main inferences about the Great Moderation turn out to be robust to whether we consider the MLE results or the interior mode. However, Bayesian estimation provides a sense of parameter uncertainty that we cannot easily obtain for the MLE results given that some parameters hit boundaries. The second reason why we consider Bayesian estimation is that it provides posterior moments not only for the model parameters, but also for some complicated functions of the model parameters that are of particular interest such as counterfactual standard deviations for output growth and implied error-correction parameters.

Our priors are specified as follows: 1) the AR coefficients have standard Normal distributions (i.e., $N(0,1)$), truncated to ensure stationarity (i.e., the roots of the characteristic equations for the AR lag polynomials lie outside the unit circle); 2) the drift for the additional trend in gross inventories has a diffuse $N(0,100)$ distribution, while the drift for long-run sales (and output) is concentrated out of the likelihood by recentering the growth rate data; 3) the precisions (inverse variances) have $\Gamma(0.01, 0.01)$ distributions, which correspond to highly diffuse priors for the variances; 4) the impact coefficients have standard Normal distributions with means recentered to be the midpoints of the bounds described in Section 3.3 (assuming $\lambda_{\kappa\eta} = 0$) and truncation to ensure the coefficients lie within or on the bounds (for any value of $\lambda_{\kappa\eta}$); and 5) the initial values for the permanent levels of sales and inventories in the pre-moderation period have diffuse Normal distributions that are centered at initial observations (minus one-period drifts) and have variances of $10,000$. All of these priors are relatively uninformative in the sense that the posteriors are dominated by the likelihood and our main qualitative inferences are robust to a range of different priors, including the flat/improper priors implicit in the consideration of MLE.

4.2 Estimates

Table 5 reports means and standard deviations of the posterior distributions of the parameters for the UC model and their changes from the pre-moderation period to the post-moderation period. The estimated volatilities become smaller for all of the shocks, although the declines are only significant for transitory shocks. The persistence of transitory sales appears to be stable across the two periods, while the persistence of transitory inventories has declined. The estimated drift for the additional trend in inventories is negative, reflecting the fact that the unconditional expectation of inventory investment implied by the UC model is $E[\Delta i_t] = \mu_\tau + \mu_\kappa$ based on equations (26), (27), and (32) and the fact that inventory investment is close to zero on average (see Figure 1). In particular, if $E[\Delta i_t] \approx 0$, the positive drift for output and sales captured by μ_τ must be offset by a negative value of μ_κ of similar magnitude (i.e., $\mu_\tau \approx -\mu_\kappa$).[17] Correspondingly, the change in magnitude in the drift for the additional trend in inventories from the pre-moderation to post-moderation period is very similar to (but opposite sign of) the change in the avrerage growth rate of sales, although it is not significant. Meanwhile, the estimated impact coefficients are generally similar in the two periods and their changes are not significant even when the estimates appear to differ.

Because it can be difficult to interpret some of the individual parameters in Table 5, especially the impact coefficients, we calculate implied volatilities, measured by standard deviations, and correlations of the underlying variables and key components. Table 6 reports means and standard deviations of the posterior distributions for these implied volatilities and correlations and their changes. The estimated volatilities become smaller in all cases and the corresponding results are qualitatively similar to the sample statistics reported in Table 1.[18] The changes are significant in all cases, except for the volatility of expected inventory investment and the correlation coefficients. Consistent with a countercyclical relationship between inven-

[17]There are small differences in the magnitudes of the drifts that could reflect secular changes in the inventory-sales ratio, perhaps due to the shift in aggregate production from manufacturing to services, as well as improvements in inventory and supply chain management, as discussed in Davis and Kahn (2008).

[18]The exact numbers in Tables 1 and 6 are different because Table 6 reports posterior means, which tend to different than posterior modes for variances. Indeed, the posterior modes are almost identical to the statistics in Table 1, confirming that our priors are largely uninformative.

TABLE 5. PARAMETERS FOR THE UC MODEL

	Pre-moderation (1960Q1-1984Q1)	Post-moderation (1984Q2-2014Q1)	Change across samples
		Sales process	
σ_η	2.26 (1.27)	1.24 (0.37)	-1.02 (1.24)
σ_ϵ	0.55 (0.09)	0.31 (0.06)	-0.24 (0.11)
ϕ_s^*	0.78 (0.12)	0.80 (0.08)	-0.02 (0.12)
$\lambda_{s\eta}$	-0.76 (0.18)	-0.75 (0.11)	0.01 (0.21)
		Inventory process	
σ_ω	1.11 (0.55)	0.65 (0.19)	-0.46 (0.57)
σ_υ	0.37 (0.07)	0.16 (0.03)	-0.21 (0.08)
ϕ_i^*	0.86 (0.06)	0.68 (0.08)	-0.18 (0.09)
μ_κ	-0.70 (0.12)	-0.44 (0.06)	0.26 (0.13)
$\lambda_{\kappa\eta}$	-0.08 (0.42)	-0.43 (0.17)	-0.35 (0.45)
$\lambda_{y\eta}$	-0.86 (0.12)	-0.74 (0.10)	0.12 (0.15)
$\lambda_{y\epsilon}$	0.76 (0.15)	0.72 (0.14)	-0.04 (0.20)
$\lambda_{i\omega}$	-0.83 (0.16)	-0.87 (0.08)	-0.04 (0.17)

Table 5: Posterior means of the parameters of the UC model are reported, with posterior standard deviations in parentheses. The ϕ^* parameters refer to sums of autoregressive coefficients for the AR(2) specifications.

TABLE 6. IMPLIED VOLATILITIES AND CORRELATIONS

	Pre-moderation (1960Q1-1984Q1)	Post-moderation (1984Q2-2014Q1)	Change across samples
s.d. (Δy_t)	1.14 (0.12)	0.64 (0.05)	-0.50 (0.13)
s.d. (Δs_t)	0.94 (0.12)	0.57 (0.05)	-0.37 (0.12)
s.d. ($\Delta^2 i_t$)	0.73 (0.06)	0.40 (0.03)	-0.34 (0.07)
corr. (Δs_t, $\Delta^2 i_t$)	-0.08 (0.07)	-0.15 (0.07)	-0.07 (0.11)
s.d. (Δi_t)	0.78 (0.12)	0.47 0.05)	-0.32 (0.13)
s.d. (Δi_t^u)	0.46 (0.06)	0.22 (0.03)	-0.24 (0.06)
s.d. (Δi_t^e)	0.62 (0.14)	0.41 (0.06)	-0.21 (0.15)
corr. (Δs_t, Δi_t^u)	-0.27 (0.11)	-0.25 (0.13)	0.02 (0.16)

Table 6: Posterior means of implied volatilities, measured in terms of standard deviations of variables, and correlations are reported, with posterior standard deviations in parentheses.

tories and sales, the correlation between unexpected inventories and sales is negative.[19] Meanwhile, the volatility estimates suggest an increase in the relative importance of expected inventories in overall inventory investment. At first glance, this change appears consistent with increased production smoothing and potentially explains the changed forecasting role of inventories in the recent sample. We investigate these possibilities in the next few subsections.

4.3 Increased production smoothing?

Given the decline in output volatility, it is natural to ask whether firms have increased their use of inventories to smooth production in the post-moderation period. Comparing the impact coefficient estimates in Table 5 with theoretical values in Table 3 in Section 3.3, the only relevant cases that we can consider are the following: the short-run scenario B, and the long-run scenarios A, B and C. In the pre-moderation period, the impact coefficients are not particularly informative for scenario A, because $\hat{\lambda}_{s\eta}$ is

[19]To the extent that informational errors might contribute to this countercyclical relationship, they will be classified as sales shocks. Thus, again, our transitory inventory shocks specifically capture informational errors and other cost shocks that do not affect sales.

reasonably close to -1 and $\hat{\lambda}_{\kappa\eta}$ is close to zero, at which point the other relevant coefficients are the same for both motives. For this sample period, based on scenario B for both the long-run and the short-run, the estimated impact coefficient is $\hat{\lambda}_{y\epsilon} = 0.76$, closer to the predicted value of 1 if firms were only concerned about avoiding stockouts. However, the long-run scenario C is more consistent with a focus on production smoothing, given the estimated parameter $\hat{\lambda}_{i\omega} = -0.83$. Based on these coefficients, the results for the pre-moderation period are ambiguous. In the post-moderation period, both $\hat{\lambda}_{y\epsilon}$ and $\hat{\lambda}_{i\omega}$ have decreased to 0.72 and -0.87, respectively. The decline in $\hat{\lambda}_{y\epsilon}$ suggests that the stockout avoidance has become less important, while a decrease in $\hat{\lambda}_{iv}$ suggests that production smoothing has become more important in the post-moderation period. However, neither change is significant. Moreover, based on scenario A, $\hat{\lambda}_{y\eta}$ equals -0.74 which suggests that in the post-moderation period firms were guided by the stockout avoidance motive. Taken together, the results for the impact coefficients provide weak support at best for the notion that production smoothing has become more relevant in the recent sample.

As noted in Section 2.3, the autoregressive coefficient, φ, for inventory investment in the cost function analysis depends on the cost coefficients a_2 and a_4. Therefore, we can look at the autoregressive coefficients for transitory inventories in our UC model to infer the relative costs associated with (long-run) production smoothing versus stockout avoidance.[20] The estimate $\hat{\phi}_i^*$ is 0.86 in the pre-moderation period, suggesting that the costs motivating production smoothing were relatively high. However, the relative costs have decreased, as the estimate $\hat{\phi}_i^*$ is 0.68 in the post-moderation period, suggesting somewhat less of a need to emphasize production smoothing in recent years.[21] This change is significant. Thus, the decline in the per-

[20]The autoregressive coefficients for transitory inventories in our UC model directly correspond to the autoregressive coefficients for inventory investment measured as the first difference in log inventories. Specifically, the UC model implies an ARMA(2,2) model for inventory investment, with the same AR coefficients as transitory inventories. Thus, we can relate the autoregressive coefficients for the UC model to the autoregressive coefficient for inventory investment (based on the level of inventories) in Section 2.3.

[21]The coefficient ϕ_i^* is the sum of the two autoregressive coefficients for an AR(2) specifi-

sistence of transitory inventories runs contrary to an increased relevance of production smoothing. Overall, then, the UC model estimates provide little support for an increase in production smoothing.

4.4 Counterfactuals

Next, we conduct counterfactual experiments to help disentangle the role of inventories from that of sales in explaining the decline in overall output volatility.[22] Our main objective is to determine whether changes in the inventory process—(i) less volatile shocks and/or (ii) changes in the propagation mechanism (autoregressive and impact coefficients)—could have accounted for the Great Moderation. To do this, we hold the parameters of the sales process fixed at their pre-moderation values and let the parameters associated with inventories (σ_ω, σ_v, ϕ_i^*, $\lambda_{\kappa\eta}$, $\lambda_{y\eta}$, $\lambda_{y\epsilon}$, and $\lambda_{i\omega}$) change to their post-moderation values. We also try to isolate the role of different inventory shocks (σ_ω and σ_v) or the propagation mechanism (ϕ_i^*, $\lambda_{\kappa\eta}$, $\lambda_{y\eta}$, $\lambda_{y\epsilon}$, and $\lambda_{i\omega}$) by changing only subsets of parameters at a time. For completeness, we also consider an experiment in which the inventory process is fixed and the sales process is allowed to change. Table 7 reports means and standard deviations of the posterior distributions for the actual and

cation of transitory inventories. Thus, we are implicitly using the sum of the AR coefficients as our measure of persistence. However, the estimated reduction in persistence is also evident if we consider the largest inverse root of the characteristic equation for the AR lag polynomial or the half-life based on an impulse response function.

[22] See Stock and Watson (2003), Ahmed, Levin, and Wilson (2004), Sims and Zha (2006), and Kim, Morley, and Piger (2008), among many others, for counterfactual experiments with VAR models. Of particular relevance to the analysis here, Kim, Morley, and Piger (2008) discuss the benefits of Bayesian inference for counterfactual quantities. Specifically, Bayesian analysis produces posterior moments for the counterfactual quantities, thus providing a sense of estimation uncertainty that is not available in the classical context. Meanwhile, Benati and Surico (2009) are critical of counterfactual analysis with reduced-form VAR models given an underlying dynamic stochastic general equilibrium (DSGE) structure generating the data. However, unlike a finite-order reduced-form VAR model, our UC model includes contemporaneous transmission within the propagation mechanism and captures VARMA dynamics, as would be implied by a DSGE structure. So, our counterfactual analysis is robust to Benati and Surico's critique of counterfactual analysis based on VAR models, although, of course, it is an open question whether our UC model parameters are "structural" in the Lucas-critique sense that a subset of parameters could have changed without all of the other parameters changing too.

TABLE 7. COUNTERFACTUAL EXPERIMENTS

	Change in s.d.(Δy_t)
Actual	-0.50 (0.13)
Sales process	-0.29 (0.15)
Inventory process	-0.08 (0.21)
Shocks	-0.12 (0.07)
Transitory shocks only	-0.11 (0.05)
Propagation	0.07 (0.21)

Table 7: Posterior means of implied changes in volatility, measured in terms of the standard deviations of output growth, are reported, with posterior standard deviations reported in parentheses. The counterfactual experiments involve changing a subset of parameters to obtain implied counterfactual changes in volatilities in the post-moderation period.

counterfactual changes in output growth volatility.

According to the results in Table 7, a change in the sales process on its own could have generated about half of the overall actual decline in the standard deviation of output growth, with the change being significant. Given that the autoregressive dynamics for sales are quite similar in the pre- and post-moderation periods, this result conforms to the "good luck" hypothesis in the sense that smaller sales shocks rather than a change in their propagation appears to be a key element of the Great Moderation. Also, the finding in Table 5 that the autoregressive dynamics did not change much suggests that it is possible to think about changing the values of some parameters of the UC model without other parameters necessarily changing too, thus perhaps mitigating concerns that the Lucas critique is empirically relevant in this setting.

In terms of inventories, the results in Table 7 suggest that their primary role in the Great Moderation was in generating the excess reduction in output volatility relative to sales. Furthermore, the results clearly support the idea that this excess reduction in volatility was driven by smaller inventory shocks rather than a change in their propagation. Consistent with the weak findings on the role of production smoothing discussed in the previous subsection, a change in inventory propagation alone would not have generated a reduction in volatility. Instead, almost the entire excess reduction in volatility that can be related to inventories appears to be due to a reduction in transitory inventory shocks, with this change being significant. Meanwhile, the sum of the counterfactual reductions in volatility is less than the overall reduction, suggesting there was an important interaction between the changes in the sales and inventory processes in explaining the Great Moderation.

4.5 Implied forecasting role of inventories

Even if increased production smoothing is not responsible for the reduction in output volatility, a question remains as to whether it is necessary to explain the changed forecasting role of inventories with the Great Moderation. Based on Table 6, a larger proportion of overall inventory investment is predictable from one period to the next, which is certainly consistent with increased production smoothing in advance of future sales. However, the analysis in Section 3.4 suggests that the forecasting role of inventories can also change with the composition of inventory forecast errors. Therefore, we consider whether the reduction in the transitory inventory shocks that explains so much of the excess reduction in output volatility relative to sales can also explain the changed forecasting role of inventories.

We determine the implied forecasting role of inventories given a change in the composition of shocks by calculating the marginal effects presented in Table 4 based on our parameter estimates. Then, we weight these marginal effects by the contribution of each underlying shock to the overall forecast error, where the weights are calculated as the ratio of the standard deviation of a shock relative to the standard deviation of the overall inventory forecast error. This calculation provides us with implied error correction coefficients (in the absence of predictable inventory changes). Table 8 reports posterior means and standard deviations for the implied error correction coefficients and their changes.

The results in Table 8 are qualitatively in line with the VECM estimates in Table 2. Specifically, there is a diminished negative forecasting relation-

TABLE 8. IMPLIED ERROR CORRECTION COEFFICIENTS

	Pre-moderation (1960Q1-1984Q1)	Post-moderation (1984Q2-2014Q1)	Change across samples
$\frac{\partial \Delta y_{t+1}}{\partial \Delta i_t^u}$	-1.11 (0.22)	-0.77 (0.40)	0.33 (0.44)
$\frac{\partial \Delta s_{t+1}}{\partial \Delta i_t^u}$	-0.09 (0.18)	0.22 (0.33)	0.31 (0.37)

Table 8: Posterior means of error correction coefficients implied by the UC model are reported, with posterior standard deviations in parentheses. The marginal impacts of the underlying shocks are weighted by their relative standard deviations.

ship between inventories and future output growth and an increased positive forecasting relationship between inventories and future sales growth in the post-moderation period. The changes are not significant and the estimated quantitative effects are somewhat different than the VECM results in Table 2, but this likely reflects the fact that the predictability of inventory investment has also changed along with the composition of shocks. The main point is that the results in Table 8 are consistent with the changing composition of shocks (specifically smaller transitory inventory shocks) explaining the changed forecasting role of inventories with the Great Moderation, without needing to rely on increased production smoothing.

4.6 Robustness

When analyzing inventory behavior, there is always a question of which data to consider. In this paper we focus on whether changes in inventory behaviour can shed light on the Great Moderation. As a result, we use our residual measure of inventory and focus on the aggregate data. Using inventory stock data directly or looking more closely at the durable good sector for which inventories are most relevant would not allow us investigate the role of inventories in moderating aggregate output growth so directly.

A reasonable question, though, is whether the findings reported above are robust to consideration of inventory stock data instead of our residual measure and durable goods data instead of the aggregate data. The short answer is yes. Using the inventory stock data (NIPA Table 5.8.6A for the sample period of 1969Q1-2014Q1 and extrapolating back to 1960Q1 using actual inventory investment), we found the following results: First, the sample statistics for inventory investment are similar to our residual measure in Table 1. Second, sales and inventories are not cointegrated, supporting the inclusion of an additional trend in our UC model, for which the estimates of the variance are again significant. Third, the estimated $\lambda_{\kappa\eta}$ is not significant when using the inventory stock measure, suggesting that the permanent inventory-sales shock, ω_t, is indeed independent of the permanent sales shock, as is assumed in our model. Fourth, the forecasting relationships as captured by error-correction coefficients using the first differences of the log stock of inventories, which is approximately the change in inventories as a percentage of the lagged stock of inventories, are not as strong as for our residual measure. Therefore, focusing on the residual measure of inventories is not only important because it allows us to directly relate changes in the sales and inventory processes in our UC model to the change in output volatility, but it is also more relevant for understanding the changed forecasting role of inventories for future output and sales growth. Overall, then, the specification of the UC model and the qualitative results were robust to using the inventory stock data.

When we estimated our UC model using quarterly output and sales data for durable goods from the BEA (NIPA Tables 1.2.3 and 1.2.5), the results were also robust and some our key findings were even more pronounced than for the aggregate data. Again, we find that sales and inventories are not cointegrated.[23] Meanwhile, the residual measure of inventory investment is responsible for a larger portion of the overall decline in output volatility than for the aggregate data. Consistent with this finding, the counterfactual analysis for the durable goods data suggests that inventories played a larger role than sales in the overall decline in volatility of durable

[23]Notably, this suggests that the lack of cointegration for the aggregate data is not simply due to a compositional effect, noted by Ramey and Vine (2004), where the inventory-sales ratio changes as services become a larger component of sales.

goods output. As with the aggregate data, inventory shocks played the primary role in the excess volatility reduction of output, with smaller transitory shocks accounting for most of this excess reduction. Meanwhile, the VECM results and forecasting implications from the UC model were quite similar to those for the aggregate data.

5 Conclusion

In this paper, we have investigated the role of inventories in the Great Moderation. Based on an unobserved components model that identifies inventory and sales shocks and their propagation in the aggregate data, we find little evidence for increased production smoothing in recent years. Instead, we find that smaller transitory inventory shocks explain the bulk of the excess volatility reduction of output relative to sales. These smaller transitory inventory shocks also potentially explain the changed forecasting role of inventories with the Great Moderation.

We note that, despite a very different approach and data, our results are closely in line with Herrera and Pesavento (2005). Specifically, they consider sectoral data and find that the reduction in the volatility of inventories with the Great Moderation was larger and more prevalent among input inventories than for finished-goods inventories. Given that production smoothing primarily relates to finished-goods inventories, their finding also argues against increased production smoothing explaining the Great Moderation, while it is entirely consistent with smaller inventory shocks.

In contemplating whether or not the Great Moderation is now over, it is important to consider what might have caused a reduction in transitory inventory shocks in the first place. To the extent that these shocks reflect informational errors about future sales and arise due to the fact that some production must be set in advance, their reduction could correspond to improved information flows about future sales or to greater flexibility in terms of setting production closer to sales. Distinguishing between these two hypotheses is difficult. However, we might expect improved informational flows to reflect a change in the predictability of the sales process. Thus, our finding that the dynamics of transitory sales remain unchanged with the Great Moderation does not lend itself to an "improved forecast" hypothesis.[24] Also, somewhat contrary to improved forecasts, which presumably occur gradually due to learning, is the fact that the volatility reduction appears to have been sudden (see Kim and Nelson, 1999, and McConnell and Perez-Quiros, 2000). Therefore, the rise of "just-in-time" production (see McConnell, Mosser, and Perez-Quiros, 1999) is the more compelling explanation for smaller transitory inventory shocks due to fewer inventory mistakes, as it is more plausible that new production processes were implemented quickly, especially after the deep recessions of the early 1980s. Also, our finding that the implied costs motivating production smoothing have declined relative to the costs motivating stockout avoidance is consistent with the idea that less production needs to be set in advance.

Although inventory mistakes might have become smaller for structural and technological reasons, it is unlikely that they will disappear altogether. In particular, the extra volatility in U.S. output relative to sales during the 2007-2009 recession strongly supports the idea that some production must be set in advance and inventory mistakes will continue to be made.[25] At the same time, given their links to technology and despite some large changes in inventories during the recent recession, a smaller variance for transitory inventory shocks provides a much more optimistic prognosis for the continuation of the Great Moderation than the "good luck" hypothesis (or, for that matter, the "good policy" hypothesis).

On a related note, it has long been understood that the role of invento-

[24] Ramey and Vine (2006) find some evidence of a change in sales dynamics for the U.S. automobile industry, which is the archetypal industry involving production that must be set in advance.

[25] The dramatic depletion of inventories in late 2008 and early 2009 is also consistent with inventory adjustments in the face of severe cash flow problems for firms in the middle of a deep recession. Carpenter, Fazzari, and Petersen (1994, 1998) highlighted the role of financing constraints in the inventory cycle. In terms of our analysis, it suggests that some of what we ascribe to inventory "mistakes" may, in fact, be deliberate temporary run-downs of inventory stocks during recessions. However, the volatility and forecasting implications of such inventory run-downs should be the same as for inventory mistakes.

ries in output fluctuations is asymmetric in terms of business cycle phases, with a much larger role being played in recessions than in expansions (see, for example, Blinder and Maccini, 1991, and Golob, 2000).[26] However, the analysis in this paper is based on a linear model and, therefore, does not capture this asymmetry. Thus, given the predominance of expansions in the sample periods covered in this paper, our results likely reflect the past and possibly future behavior of output, sales, and inventories in expansions more than in recessions (over 80% of the observations in our sample are from NBER-dated expansions). This could, in part, explain some of the differences between our conclusions and those in a recent paper by Maccini and Pagan (2009). They explicitly measure movements in output related to business cycle phases and find little role for inventories in the changed behavior of output with the Great Moderation.[27] It also means that we cannot draw strong conclusions about possible changes in recession and recovery dynamics due to inventories (see Camacho, Perez-Quiros, and Rodriguez-Mendizabal, 2009). Modeling business cycle asymmetries associated with inventories presents its own challenges and opportunities, which we leave for future research.

References

[1] Ahmed, S., A. Levin and B. Wilson. 2004. "Recent U.S. Macroeconomic Stability: Good Luck, Good Policies, or Good Practices?" *Review of Economics and Statistics* 86(3): 824-32.

[26]Granger and Lee (1989) find evidence of asymmetric error correction effects depending on the sign of inventory investment relative to its mean and the sign of the cointegrating error for inventories and sales.

[27]Somewhat more consistent with our findings, Maccini and Pagan (2009) find that increased production smoothing does not play a role in the Great Moderation. Instead, they find that an estimated structural model based on pre-moderation data could only have generated the observed reduction in output volatility if the volatilities of the sales process and technology shocks declined by about half. In this sense, their results are strongly supportive of the "good luck" hypothesis. However, their model does not allow for inventory mistakes. As a robustness check, they do consider a modified version of their model in which only past values of sales are observed by firms when setting production. However, this is different from inventory mistakes that arise from noisy signals about sales.

[2] Benati, L. and P. Surico. 2009. “VAR Analysis and the Great Moderation” *American Economic Review* 99(4): 1636-1652.

[3] Blanchard, O.J. 1983. “The Production and Inventory Behavior of the American Automobile Industry” *Journal of Political Economy* 91: 585-614.

[4] Blinder, A.S. and L.J. Maccini. 1991. “Taking Stock: A Critical Assessment of the Recent Research on Inventories” *Journal of Economic Perspectives* 5(1): 73-96.

[5] Camacho, M., G. Perez-Quiros, and H. Rodriguez-Mendizabal. 2009. “Are the High-Growth Recovery Periods Over?.” Mimeo.

[6] Carpenter, R.E., S.M. Fazzari, and B.C. Petersen. 1994. “Inventory Investment, Internal-Finance Fluctuations, and the Business Cycle.” *Brookings Papers on Economic Activity* 2: 75-138.

[7] Carpenter, R.E., S.M. Fazzari, and B.C. Petersen. 1998. “Financing Constraints and Inventory Investment: A Comparative Study with High-Frequency Panel Data.” *Review of Economics and Statistics* 80: 513-518.

[8] Chib, S. and E. Greenberg. 1995. “Understanding the Metropolis-Hastings Algorithm.” *American Statistician* 49: 327-335.

[9] Clarida, R., J. Gali, and M. Gertler. 2000. “Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory.” *Quarterly Journal of Economics* 115(1): 147-180.

[10] Clark, T. 2009. “Is the Great Moderation Over? An Empirical Analysis.” *Federal Reserve Bank of Kansas City Economic Review* 4: 5-42.

[11] Davis, S. and J.Kahn. 2008. “Interpreting the Great Moderation: Changes in the Volatility of Economic Activity at the Macro and Micro Levels.” *Journal of Economic Perspectives* 22(4): 155-180.

[12] Eo, Y. and J. Morley. Forthcoming. "Likelihood-Ratio-Based Confidence Sets for the Timing of Structural Breaks." *Quantitative Economics*.

[13] Golob, J. 2000. "Post-1984 Inventories Revitalize the Production Smoothing Model" Mimeo.

[14] 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 Econometrics* 4: S145-S159.

[15] Hamilton, J. D. 2002. "On the Interpretation of Cointegration in the Linear-Quadratic Inventory Model" *Journal of Economic Dynamics and Control* 26: 2037-2049.

[16] Hansen, L. P. and T. J. Sargent. 1980. "Formulating and Estimating Dynamic Linear Rational Expectations Model." *Journal of Economic Dynamics and Control* 2(1): 7-46.

[17] Herrera, A. M., I. Murtazashvili, and E. Pesavento. 2008. "The Comovement in Inventories and in Sales: Higher and Higher" *Economics Letters* 99: 155-158.

[18] Herrera, A.M. and E. Pesavento. 2005. "The Decline in U.S. Output Volatility: Structural Changes and Inventory Investment" *Journal of Business & Economic Statistics* 23(4): 462-472.

[19] Holt, C.C., F. Modigliani, J.F. Muth, and S.A. Herbert. 1960. *Planning, Production, Inventories, and Work Force*. Englewood Cliffs, Prentice-Hall, N.J.

[20] Irvine, F. O. and S. Schuh. 2005. "The Roles of Comovement and Inventory Investment in the Reduction of Output Volatility" Federal Reserve Bank of Boston, Research Paper Series, Working paper No. 05-9.

[21] Kahn, J., M. McConnell, and G. Perez-Quiros. 2002. "On the Causes of the Increased Stability of the U.S. Economy."*FRBNY Economic Policy Review* 8: 183-202.

[22] Kim, C.-J., J. Morley, and J. Piger. 2008. "Bayesian Counterfactual Analysis of the Sources of the Great Moderation." *Journal of Applied Econometrics* 23: 173-191.

[23] Kim, C.-J., and C. Nelson. 1999. "Has the U.S. Economy become more Stable? A Bayesian Approach Based on a Markov-Switching Model of the Business Cycle." *Review of Economics and Statistics* 81(4): 608-616.

[24] Maccini, L., B. Moore and H. Schaller. 2012. "Inventory Behaviour with Permanent Sales Shocks" Mimeo.

[25] Maccini, L. and A. Pagan. 2009. "Inventories, Fluctuations, and Business Cycles" Mimeo.

[26] McCarthy, J. and E. Zakrajsek. 2007. "Inventory Dynamics and Business Cycles: What has Changed?"*Journal of Money Credit and Banking* 39(2-3): 615-638.

[27] McConnell, M., P. Mosser, and G. Perez-Quiros. 1999. "A Decomposition of the Increased Stability of GDP Growth."*FRBNY Current Issues in Economics and Finance* 5, no. 13.

[28] McConnell, M., and G. Perez-Quiros. 2000. "Output Fluctuations in the United States: What Has Changed Since the Early 1980s?" *American Economic Review* 90(5): 1464-1476.

[29] Miron, J., and S. Zeldes. 1988. "Seasonality, Cost Shocks, and the Production Smoothing Model of Inventories." *Econometrica* 56(4): 877-908.

[30] Morley, J.C.. 2007."The Slow Adjustment of Aggregate Consumption to Permanent Income" *Journal of Money, Credit, and Banking* 39(2-3): 615-638.

[31] Morley, J.C., C.R. Nelson and E. Zivot. 2003."Why are the Beveridge-Nelson and Unobserved-Component Decompositions of GDP so Different?" *Review of Economics and Statistics* 85(4): 235-243.

[32] Ramey, V.A. and D.J. Vine. 2004. "Why Do Real and Nominal Inventory-Sales Ratios Have Different Trends?Declining Volatility in the U.S. Automobile Industry?" *Journal of Money, Credit and Banking* 36(5): 959-963.

[33] Ramey, V.A. and D.J. Vine. 2006. "Declining Volatility in the U.S. Automobile Industry." *American Economic Review* 96(5): 1876-1889.

[34] Ramey, V.A. and K.D. West. 1999. "Inventories." In *Handbook of Macroeconomics*, edited by John B. Taylor and Michael Woodford, pp. 863-923. New York: North Holland, Elsevier.

[35] Sims C., and T. Zha. 2006. "Were There Regime Switches in US Monetary Policy?" *American Economic Review* 96: 54-81.

[36] Sinclair, T. 2009. "The Relationship between Permanent and Transitory Movements in U.S. Output and the Unemployment Rate" *Journal of Money Credit and Banking* 41(2-3): 529-542.

[37] Stock, J., and M. Watson. 2003. "Has the Business Cycle Changed and Why?" *NBER Macroeconomics Annual 2002* 17: 159-218.

[38] Wen, Y. 2005. "Understanding the Inventory Cycle" *Journal of Monetary Economics* 52: 1533-1555.

[39] West, K. 1986. "A Variance Bounds Test of the Linear Quadratic Inventory Model" *Journal of Political Economy* 94: 374-401.

[40] West, K. 1990. "The Sources of Fluctuations in Aggregate Inventories and GNP" *Quarterly Journal of Economics* CV: 939-972.

A Appendix

In this appendix, we present the state-space representation of the UC model.

The observation equation is

$$\tilde{\mathbf{y}}_t = \mathbf{H}\,\boldsymbol{\beta}_t$$

where

$$\tilde{\mathbf{y}}_t = \begin{bmatrix} s_t \\ i_t \end{bmatrix}, \; \mathbf{H} = \begin{bmatrix} 1 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 & 1 \end{bmatrix} \text{ and } \boldsymbol{\beta}_t = \begin{bmatrix} s_t - \tau_t^* \\ s_{t-1} - \tau_{t-1}^* \\ i_t - i_t^* \\ i_{t-1} - i_{t-1}^* \\ \tau_t^* \\ \kappa_t \end{bmatrix}$$

The state equation is

$$\boldsymbol{\beta}_t = \tilde{\boldsymbol{\mu}} + \boldsymbol{F}\boldsymbol{\beta}_{t-1} + \tilde{\boldsymbol{v}}_t$$

where

$$\tilde{\boldsymbol{\mu}} = \begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \\ \mu_\tau \\ \mu_\kappa \end{bmatrix}, \; \mathbf{F} = \begin{bmatrix} \phi_{s,1} & \phi_{s,2} & 0 & 0 & 0 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \phi_{i,1} & \phi_{i,2} & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}, \; \tilde{\boldsymbol{v}}_t = \begin{bmatrix} \lambda_{s\eta}\eta_t + \epsilon_t \\ 0 \\ \lambda_{i\eta}\eta_t + \lambda_{i\omega}\omega_t + \lambda_{i\epsilon}\epsilon_t + v_t \\ 0 \\ \eta_t \\ \omega_t + \lambda_{\kappa\eta}\eta_t \end{bmatrix}$$

and the covariance matrix of $\tilde{\boldsymbol{v}}_t$, $\mathbf{Q}$, is given by

$$\mathbf{Q} = \begin{pmatrix} \lambda_{s\eta}^2\sigma_\eta^2 + \sigma_\epsilon^2 & 0 & \lambda_{s\eta}\lambda_{i\eta}\sigma_\eta^2 + \lambda_{i\epsilon}\sigma_\epsilon^2 & 0 & \lambda_{s\eta}\sigma_\eta^2 & \lambda_{s\eta}\lambda_{\kappa\eta}\sigma_\eta^2 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ \lambda_{s\eta}\lambda_{i\eta}\sigma_\eta^2 + \lambda_{i\epsilon}\sigma_\epsilon^2 & 0 & \lambda_{i\eta}^2\sigma_\eta^2 + \lambda_{i\omega}^2\sigma_\omega^2 + \lambda_{i\epsilon}^2\sigma_\epsilon^2 + \sigma_v^2 & 0 & \lambda_{i\eta}\sigma_\eta^2 & \lambda_{i\eta}\lambda_{\kappa\eta}\sigma_\eta^2 + \lambda_{i\omega}\sigma_\omega^2 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ \lambda_{s\eta}\sigma_\eta^2 & 0 & \lambda_{i\eta}\sigma_\eta^2 & 0 & \sigma_\eta^2 & \lambda_{\kappa\eta}\sigma_\eta^2 \\ \lambda_{s\eta}\lambda_{\kappa\eta}\sigma_\eta^2 & 0 & \lambda_{i\eta}\lambda_{\kappa\eta}\sigma_\eta^2 + \lambda_{i\omega}\sigma_\omega^2 & 0 & \lambda_{\kappa\eta}\sigma_\eta^2 & \sigma_\omega^2 \end{pmatrix}$$

B Appendix

In this appendix, we solve the UC model for inventory investment, sales growth, and output growth. We then show how to calculate the implied variances of inventory investment, unexpected inventory investment, expected inventory investment, sales growth, and output growth for the UC model. In addition we also calculate $cov(\Delta s_t, \Delta^2 i_t)$ and $cov(\Delta s_t, \Delta i_t^u)$.

The change in inventories is given by

$$\Delta i_t = \Delta i_t^* + (1-L)(i_t - i_t^*) = \mu_\tau + \mu_\kappa + (1+\lambda_{\kappa\eta})\eta_t + \omega_t + z_t^i$$

where $(1-\phi_{i,1}L - \phi_{i,2}L^2)z_t^i = (1-L)x_t^i$ and $x_t^i = \lambda_{i\eta}\eta_t + \lambda_{i\omega}\omega_t + \lambda_{i\epsilon}\epsilon_t + \upsilon_t$. The process of sales growth is given by

$$\Delta s_t = \eta_t + z_t^s$$

where $(1-\phi_{s,1}L - \phi_{s,2}L^2)z_t^s = (1-L)x_t^s$ and $x_t^s = \lambda_{s\eta}\eta_t + \epsilon_t$. Then, using the inventory identity, the change in output can be re-written as

$$\Delta y_t = \Delta s_t + (1-L)\Delta i_t = (\eta_t + z_t^s) + (1+\lambda_{\kappa\eta})\eta_t + \omega_t + z_t^i - (1+\lambda_{\kappa\eta})\eta_{t-1} - \omega_{t-1} - z_{t-1}^i$$

Note that the state equation for z_t^s and z_t^i is

$$\mathbf{z}_t = \mathbf{K}\mathbf{z}_{t-1} + \mathbf{w}_t$$

where

$$\mathbf{z}_t = \begin{bmatrix} z_t^s \\ z_{t-1}^s \\ z_t^i \\ z_{t-1}^i \\ x_t^s \\ x_t^i \end{bmatrix}, \mathbf{K} = \begin{bmatrix} \phi_{s,1} & \phi_{s,2} & 0 & 0 & -1 & 0 \\ 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & \phi_{i,1} & \phi_{i,2} & 0 & -1 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \end{bmatrix}, \mathbf{w}_t = \begin{bmatrix} x_t^s \\ 0 \\ x_t^i \\ 0 \\ x_t^s \\ x_t^i \end{bmatrix}$$

Letting $\mathbf{W}$ be the covariance matrix with the following non-zero entries $\mathbf{W}[1,1] = \mathbf{W}[1,5] = \mathbf{W}[5,1] = \mathbf{W}[5,5] = \lambda_{s\eta}^2\sigma_\eta^2 + \sigma_\epsilon^2$, $\mathbf{W}[1,3] = \mathbf{W}[3,1] = \mathbf{W}[1,6] = \mathbf{W}[6,1] = \mathbf{W}[3,5] = \mathbf{W}[5,3] = \mathbf{W}[5,6] = \mathbf{W}[6,5] = \lambda_{s\eta}\lambda_{i\eta}\sigma_\eta^2 + \lambda_{i\epsilon}\sigma_\epsilon^2$, and $\mathbf{W}[3,3] = \mathbf{W}[3,6] = \mathbf{W}[6,3] = \mathbf{W}[6,6] = \lambda_{i\eta}^2\sigma_\eta^2 + \lambda_{i\omega}^2\sigma_\omega^2 + \lambda_{i\epsilon}^2\sigma_\epsilon^2 + \sigma_\upsilon^2$

and $vec(var(\mathbf{z}_t)) = ((I - \mathbf{K} \otimes \mathbf{K})^{-1} vec(\mathbf{W}))$. Then, the variance of inventory investment is given by

$$
\begin{aligned}
& var(\Delta i_t) = var((1+\lambda_{\kappa\eta})\eta_t + \omega_t + z_t^i) \\
& = (1+\lambda_{\kappa\eta})^2\sigma_\eta^2 + \sigma_\omega^2 + var(z_t^i) + 2cov((1+\lambda_{\kappa\eta})\eta_t, z_t^i) + 2cov(\lambda_{\kappa\epsilon}\epsilon_t, z_t^i + 2cov(\omega_t, z_t^i) \\
& = (1+\lambda_{\kappa\eta})^2\sigma_\eta^2 + \lambda_{\kappa\epsilon}^2\sigma_\epsilon^2 + \sigma_\omega^2 + var(z_t^i) + 2(1+\lambda_{\kappa\eta})\lambda_{i\eta}\sigma_\eta^2 + 2\lambda_{\kappa\epsilon}\lambda_{i\epsilon}\sigma_\epsilon^2 + 2\lambda_{i\omega}\sigma_\omega^2
\end{aligned}
$$

where $var(z_t^i)$ is the $[3,3]$ element of $var(\mathbf{z}_t)$. The variances of the two expectational components of inventory investment are given by

$$var(\Delta i_t^u) = (\lambda_{y\eta} - \lambda_{s\eta})^2\sigma_\eta^2 + (\lambda_{y\epsilon} - 1)^2\sigma_\epsilon^2 + \lambda_{y\omega}^2\sigma_\omega^2 + \sigma_v^2.$$

and

$$var(\Delta i_t^e) = var(\Delta i_t) - var(\Delta i_t^v)$$

The variance of sales growth is given by

$$var(\Delta s_t) = var(\eta_t + z_t^s) = \sigma_\eta^2 + var(z_t^s) + 2\lambda_{s\eta}\sigma_\eta^2$$

and $var(z_t^s)$ is the $[1,1]$ element of $var(\mathbf{z}_t)$. Finally, the variance of output growth is given by

$$
\begin{aligned}
& var(\Delta y_t) = var(\Delta s_t + \Delta i_t - \Delta i_{t-1}) \\
& = var(\Delta s_t) + 2var(\Delta i_t) + 2cov(\Delta s_t, \Delta i_t) - 2cov(\Delta s_t, \Delta i_{t-1}) - 2cov(\Delta i_t, \Delta i_{t-1})
\end{aligned}
$$

where

$$
\begin{aligned}
cov(\Delta s_t, \Delta i_t) &= cov(\eta_t + z_t^s, (1+\lambda_{\kappa\eta})\eta_t + \omega_t + z_t^i) \\
&= (1+\lambda_{\kappa\eta})\sigma_\eta^2 + (1+\lambda_{\kappa\eta})\lambda_{s\eta}\sigma_\eta^2 + cov(z_t^s, z_t^i) + \lambda_{i\eta}\sigma_\eta^2,
\end{aligned}
$$

$$
\begin{aligned}
cov(\Delta s_t, \Delta i_{t-1}) &= cov(\eta_t + z_t^s, (1+\lambda_{\kappa\eta})\eta_{t-1} + \omega_{t-1} + z_{t-1}^i) \\
&= cov(z_t^s, z_{t-1}^i) + cov(z_t^s, (1+\lambda_{\kappa\eta})\eta_{t-1} + \omega_{t-1}) \\
&= cov(z_t^s, z_{t-1}^i) + (\phi_{s,1} - 1)\lambda_{s\eta}(1+\lambda_{\kappa\eta})\sigma_\eta^2
\end{aligned}
$$

and

$$
\begin{aligned}
cov(\Delta i_t, \Delta i_{t-1}) &= cov((1+\lambda_{\kappa\eta})\eta_t + \omega_t + z_t^i, (1+\lambda_{\kappa\eta})\eta_{t-1} + \omega_{t-1} + z_{t-1}^i) \\
&= cov(z_t^i, (1+\lambda_{\kappa\eta})\eta_{t-1} + \omega_{t-1} + z_{t-1}^i) \\
&= (\phi_{i,1} - 1)(\lambda_{i\eta}(1+\lambda_{\kappa\eta})\sigma_\eta^2 + \lambda_{i\omega}\sigma_\omega^2) + cov(z_t^i, z_{t-1}^i)
\end{aligned}
$$

where $cov(z_t^s, z_t^i)$, $cov(z_t^s, z_{t-1}^i)$ and $cov(z_t^i, z_{t-1}^i)$ are the $[1,3]$, $[1,4]$ and $[3,4]$ element of $var(\mathbf{z}_t)$ respectively.

The two additional covariance terms are given below

$$
\begin{aligned}
cov(\Delta s_t, \Delta i_t^2) &= cov(\eta_t + z_t^s, (1+\lambda_{\kappa\eta})\eta_t + \omega_t + z_t^i - (1+\lambda_{\kappa\eta})\eta_{t-1} - \omega_{t-1} - z_{t-1}^i)) \\
&= cov(z_t^s, z_{t-1}^i) + cov(z_t^s, (1+\lambda_{\kappa\eta})\eta_{t-1} + \omega_{t-1}) \\
&= (1+\lambda_{\kappa\eta})\sigma_\eta^2 + \lambda_{i\eta}\sigma_\eta^2 + \lambda_{s\eta}(1+\lambda_{\kappa\eta})\sigma_\eta^2 - (\phi_{s,1}-1)\lambda_{s\eta}(1+\lambda_{\kappa\eta})\sigma_\eta^2 + cov(z_t^s, z_t^i) - cov(z_t^s, z_{t-1}^i)
\end{aligned}
$$

and

$$
\begin{aligned}
cov(\Delta s_t, \Delta i_t^u) &= cov(\eta_t + z_t^s, (\lambda_{y\eta} - \lambda_{s\eta})\eta_t + (\lambda_{y\epsilon} - 1)\epsilon_t + \lambda_{y\omega}\omega_t + v_t) \\
&= (\lambda_{y\eta} - \lambda_{s\eta})\sigma_\eta^2 + (\lambda_{y\epsilon} - 1)\sigma_\epsilon^2 + \lambda_{s\eta}(\lambda_{y\eta} - \lambda_{s\eta})\sigma_\eta^2.
\end{aligned}
$$