Production and Inventory Dynamics Under Ambiguity Aversion*

Yulei Luo†
The University of Hong Kong

Jun Nie‡
Wuhan University

Xiaowen Wang§
Shanghai University of Finance and Economics

Eric R. Young¶
University of Virginia and Federal Reserve Bank of Cleveland

April 3, 2024

Abstract

In this paper we propose a production-cost smoothing model with Knightian uncertainty due to ambiguity aversion to study the joint behavior of production, inventories, and sales. Our model can explain ten facts that previous studies find difficult to account for simultaneously including: the high volatility of production relative to sales, the low ratio of inventory-investment volatility to sales volatility, the positive correlation between sales and inventory investment, and the negative correlation between the inventory-to-sales ratio and sales. Our main results extend to a model of endogenous sales. Finally, we find that the stock-out avoidance motive emerges endogenously in our model, reconciling the long debate in the inventory literature over the production-cost smoothing and the stock-out avoidance models.

Keywords: Ambiguity Aversion, Knightian Uncertainty, Inventories, Production-Cost Smoothing.

JEL Classification Numbers: D83, E21, F41, G15.

---

*We are indebted to the co-editor (Yuriy Gorodnichenko), the associate editor (Cosmin Ilut) and the anonymous referee for detailed comments and suggestions which have substantially improved the paper. We are also grateful to Ken Kasa, Aubhik Khan, Tom Sargent, Chi-Wa Yuen, Hongsong Zhang, Lei Zhang, and conference participants for helpful discussions and comments. We also thank Chenchuan Shi for excellent research assistance throughout this project. Luo thanks the General Research Fund (GRF No. HKU17500619) in Hong Kong for financial support.

†Faculty of Business and Economics, The University of Hong Kong, Hong Kong. E-mail: `yulei.luo@gmail.com`.

‡School of Economics and Management, Wuhan University. E-mail: `niejun.econ@gmail.com`.

§School of Economics, Shanghai University of Finance and Economics. E-mail: `wang.xiaowen@sufe.edu.cn`.

¶Department of Economics, University of Virginia, Charlottesville, VA. E-mail: `ey2d@virginia.edu`. Research Department, Federal Reserve Bank of Cleveland, Cleveland, OH. Email: `eric.young@clev.frb.org`.

1 Introduction

In more than half a century, economists have made great strides toward a better understanding of firms' inventory decisions. This progress is unsurprising given the disproportionate role inventory investment plays in explaining fluctuations in output.[1] Many previous studies on inventory dynamics are based on the linear-quadratic framework developed by Holt et at. (1960), which assumes firms' production and inventory decisions are associated with convex costs and stochastic demand and costs. Three prominent theories have emerged from this framework: the production smoothing (PS) model, which assumes firms use inventories to minimize production fluctuations given stochastic demand; the production cost smoothing (PCS) model, which introduces shocks to the marginal cost of production and assumes that firms hold inventories to smooth the production costs rather than production levels; and the stock-out avoidance (SOA) model, which assumes firms have incentive to accumulate inventories because they cannot short sell their products without cost.

Researchers have questioned whether these theories can rationalize two important aspects of the data: the higher volatility of production relative to sales and the procyclicality of inventory investment. The consensus in the literature is that the PS model fails to generate a volatility of output higher than that of sales (which is known as the "production smoothing puzzle").[2] Although the PCS model can generate a higher output volatility with the help of large cost shocks to production, it fails to generate procyclical inventory investment. The SOA model appears to be more consistent with the data along these two dimensions. We provide an alternative mechanism by which firms can avoid negative inventories and makes the model to be consistent with data in more dimensions.[3]

One common feature of these models is that they assume firms know exactly the probability distributions of the demand and cost shock processes. For example, a typical assumption used in the literature is an AR(1) specification for both the demand (Blinder 1982) and cost shock processes (Eichenbaum 1989). Firms consider them as true data-generating processes for demand and cost

[1]Inventory investment constitutes less than 1 percent of real GDP in most countries, but accounts for a significant fraction of real GDP fluctuations. For example, in the United States, the decline in inventory investment in the two most recent severe recessions, $1990-1991$ and $2007-2009$, explains 49 percent and 42 percent of the total decline in output, respectively.

[2]The production smoothing puzzle was first raised by Blanchard (1983). If inventories are used to absorb fluctuations in demand, the model predicts that production will be smoother than sales. However, the data show the opposite. See Blinder (1986), Blinder and Maccini (1991), and Lai (1991) for detailed discussions on this puzzle.

[3]Section 6 provides more detailed discussions on the two different ways to model stockout avoidance motives.

shocks, and in turn use them to form expectations for next period's demand and cost values. In practice, however, the true demand and cost shock processes may be difficult to precisely measure. Thus, it is plausible that firms' reference models for these processes may be misspecified in some way, and that firms may be aware of this potential for error and alter their production and inventory decisions accordingly.

In this paper, we study the implications of firms' uncertainty about these demand- and cost-shock processes for the joint dynamics of inventories, production, and sales. We generalize an otherwise standard PCS model by assuming: (i) that firms face Knightian uncertainty (or model uncertainty) about the true dynamics of sales and production costs; and ii) that aversion to this uncertainty leads them to make production and inventory decisions under the worst-case scenario.[4] Our approach is also partially motivated by the increased concerns following the $2007-2009$ financial crisis about whether existing macroeconomic models are misspecified. Some recent papers argue that model uncertainty, the "unknown unknowns," played an important role in the recent economic and financial crises (see, for example, Caballero and Krishnamurthy, 2008).

We describe firms' aversion to ambiguity about the true dynamics of sales and production cost using the preference for robustness proposed by Hansen and Sargent (2007). In robustness models, agents have in mind a reference model that represents their best estimate of the model economy. However, because they are worried that this reference model may be incorrect in some hard-to-specify way, they make their optimal decisions under the worst-case scenario that they consider reasonable. We assume that firms are concerned about the possibility that their true model is misspecified in a manner that is difficult to detect statistically; consequently, they choose their decisions *as if* the subjective distribution over shocks was chosen by an evil agent in order to maximize their expected total costs.[5]

Our main result is that once we take ambiguity aversion into account, an otherwise standard PCS model can generate not only the two aspects of the data mentioned above but also other features of the data, including the relative volatility of inventory investment to sales and the contemporaneous

[4]Throughout this paper we use "Knightian uncertainty" and "model uncertainty" interchangeably.

[5]In this paper, we follow Hansen et al. (1999), Hansen and Sargent (2007), and Bidder and Smith (2012) in studying a model with ambiguity aversion described by "*multiplier*" preferences. For other specifications and applications of ambiguity aversion and robustness in macroeconomics, see Luo and Young (2010), Ilut and Schneider (2014), and Ilut et al. (2020).

correlation between the inventory-to-sales ratio and sales. In addition, introducing ambiguity aversion into our PCS model allows the SOA motive to emerge *endogenously*, reconciling the PCS model with the SOA model, due to a precautionary motive.[6] These findings suggest that embedding ambiguity aversion into this framework is a promising and perhaps necessary step for rationalizing the dynamics of inventories and production.

The key mechanism through which ambiguity aversion influences production and inventory dynamics are as follows. First, as firms with ambiguity aversion are concerned that their reference model (i.e., the best estimated model using available data and knowledge) is misspecified in some aspects and make production decisions under the worst-case scenario, production becomes *more sensitive* to shocks than without ambiguity aversion. For example, when firms with ambiguity aversion see a positive sales shock, they fear that future sales might be higher than what the reference model suggests, and thus produce more than without ambiguity aversion to better smooth production. Second, the uncertainty (about future sales and production costs) generates precautionary behavior for firms, implying that firms with ambiguity aversion on average have higher levels of production and inventories than without ambiguity aversion. This precautionary pattern can be labeled "*making hay while the sun shines*". It is worth noting that the "making hay while the sun shines" mechanism was also discussed in the literature on consumption and savings such as Hansen et al. (1999) and Luo and Young (2010).

This mechanism enables our model to explain ten dimensions of data including the two stylized ones mentioned above.[7] First, as firms adjust production more aggressively in response to sales, the relative volatility of production to sales will increase, helping to resolve the production-smoothing puzzle. Second, when the response of production to sales is larger than one-to-one (say, due to a high degree of ambiguity aversion), the response of inventories can move in the same direction as sales, producing procyclical inventory investment. Third, the response of inventory investment to sales changes from negative to positive as ambiguity aversion rises, the volatility of inventory investment first declines and then rises. Fourth, the inventory-to-sales ratio becomes countercyclical

[6] Our theory differs from the cost shock channel. As argued in Wen (2005), if inventories cannot be negative and demand and cost shocks are uncorrelated, allowing for cost shocks does not change the correlation between sales and inventory investment.

[7] It is not surprising that some of these ten dimensions are highly correlated. For example, if production andn sales are highly correlated and the contemporaneous correlation between inventory investment and sales is negative, then he contemporaneous correlation between inventory investment and production is also negative.

if ambiguity aversion is sufficiently strong, because a sufficiently strong degree of ambiguity aversion successfully prevents firm inventories from turning negative, an important point linked to the SOA motive.[8] In other words, this precautionary behavior due to ambiguity aversion can endogenously generate the stock-out avoidance motive even though we do not model the ad hoc SOA motive via adding a non-negativity constraint on inventories or an "accelerator" term in the inventory-holding cost function.

We use the GMM to estimate the four key model parameters: the persistence and volatility of the cost shock, the relative importance of inventory holding costs, and the degree of ambiguity aversion. To cover as more as possible the relevant data information, we consider ten moments in our estimation. (See the detailed definitions of these moments in Table 1 and Proposition 2.) At the estimated parameter values, our model matches all ten moments reasonably well. The uncovered value of the model uncertainty parameter implies a 22% probability that the best estimated processes (which are AR(1) processes) cannot be statistically distinguished from the worst-case models based on a likelihood ratio test. This value is in the reasonable range of the values reported in the robust control literature and implies that the model's success is not due to unreasonable fears of model misspecification.

As mentioned earlier, we show that the presence of Knightian uncertainty provides an endogenous mechanism for generating a stock-out avoidance motive. The stock-out avoidance motive was first proposed by Kahn (1987) and has been widely used in the inventory literature, as it successfully resolves the production-smoothing puzzle. Existing models incorporate the stock-out avoidance motive in two ways. The first is to directly impose a non-negative constraint on inventories (Kahn 1987; Wen 2005) so firms have to accumulate more inventories to avoid possible stock-outs. The second is to include an "accelerator" term that represents a quadratic cost associated with allowing inventories to deviate from some fixed proportion of sales (Eichenbaum, 1989; Ramey, 1991; and Ramey and West, 1999). Notice that this cost term may not completely rule out negative inventories.[9] We show that our model with model uncertainty endogenously avoids negative inventories with a reasonable degree of ambiguity aversion. The intuition is that firms with higher ambiguity aversion have a stronger

[8] When there is no ambiguity or the degree of ambiguity aversion is low, inventory is generally negative. For example, see lower panels of Figure 5. It is the negative inventory that makes the inventory-to-sales ratio countercyclical.

[9] In one simulation exercise, we show ruling out negative inventories requires the parameter on the targeted proportion of sales to be significantly higher than empirically plausible.

incentive to accumulate inventories as their decisions are based on the worst-case scenario, which generates a form of precautionary savings. Therefore, taking Knightian uncertainty into account can endogenously generate the stock-out avoidance motive even though the interpretation is completely different.

We choose the linear-quadratic Gaussian (LQG) framework to study the implications of ambiguity aversion on the dynamics of inventory and production mainly because this is the primary framework used in this literature. This structure makes it easier for us to contrast our results with the existing results from other models. In addition, as shown in Hansen and Sargent (2007, 2010), within the LQG setting using relative entropy to measure model misspecification leads to a simple generalization to the ordinary LQG dynamic program problem, which keeps our model tractable and allows us to fully inspect the mechanism through which ambiguity aversion affects the dynamics of production and inventories. There are also a few recent papers that study inventory dynamics and business cycles based on general equilibrium models (Khan and Thomas, 2007; Wen, 2011). These models include more elements so as to study a wide range of other important issues on business cycles, the complexities in these models obscure the intuition we want to highlight. The main point we make here is that once we allow for a moderate degree of ambiguity aversion, a standard LQG inventory model can explain the data reasonably well, even without explicitly assuming the stockout avoidance motive.

Related Literature Our paper is related to two branches of literature. First, it is related to the long branch of literature examining the joint dynamics of production, inventories, and sales based on the linear-quadratic framework dating back to Holt et al. (1960). This literature in general studies different theories' implications for inventory and production dynamics focusing on two aspects. Blanchard (1983) is the first to show the production-smoothing model cannot explain why output volatility is higher than sales volatility in the data. To increase the relative volatility of output to sales, Blinder (1986), West (1987), and Eichenbaum (1989) convert the production smoothing model into a production-cost smoothing model by introducing shocks to technology and production cost. Eichenbaum (1989) provides empirical evidence for the PCS model over the PS model. Kahn (1987) and Maccini and Zabel (1996) show that incorporating the SOA motive can also help explain the relative high output volatility. Wen (2005) shows that the PCS model cannot explain the procyclical inventory investment, while the SOA model can. Bils and Kahn (2000) use the cyclicality of the inventory-to-sales ratio to derive the implications for the cyclicality of markups. Kryvstov and Midrigan (2013), Sarte et al. (2015), and Görtz et al. (2022) use inventory models to make inferences about drivers of business cycles. Our paper contributes to this literature by showing that incorporating ambiguity aversion not only helps the LQG inventory framework account for the data, but it also endogenously generates the SOA motive.

Second, our paper contributes to the branch of literature on Knightian uncertainty, ambiguity aversion, and robustness. Many recent papers have shown the usefulness of viewing agents as having (potentially) misspecified models and being aware of this fact. See, for example, Hansen et al. (1999), Hansen and Sargent (2010), Bidder and Smith (2012), Luo et al. (2012), Djeutem and Kasa (2013), Ilut and Schneider (2014), and Ilut et al. (2020). In addition, the implied distortion due to ambiguity aversion in our model is similar to the AR(1) ambiguity shocks proposed in Bhandari *et al.* (2022). They identify AR(1) ambiguity shocks using U.S. survey data, and find that empirically the ambiguity shocks are an important source of variation in labor market variables. To the best of our knowledge, we are the first to introduce Knightian uncertainty and ambiguity aversion into the inventory literature.

The rest of paper is organized as follows. Section 2 presents the data and facts. Section 3 first examines the theoretical implications of the RE-PCS model, and then introduces our benchmark model of robust production and inventory decisions. Section 4 delivers the model's quantitative implications and show how the model fits the data well along the key aspects. Section 5 extends our model by considering endogenous sales. Section 6 compares our model with the traditional SOA model. Section 7 provides further discussions on the dynamics of inventories and production. Section 8 concludes.

2 Data and Facts

In this section, we document some facts regarding the joint behavior of production, inventories, and sales. We focus primarily on the manufacturing sector, where inventories play an important role in the production process. As a robustness check, we also report results using data from the wholesale and retail sectors in Online Appendix A.

Our monthly data comes from the Bureau of Economic Analysis (BEA) and covers the 1967m1 – 2018m12 period. To exclude the volatile period around the Global Financial Crisis (GFC), we also report statistics based on a sample prior to the GFC, i.e., 1967m1–2007m12.[10] Following Blinder and Maccini (1991), we define the inventory stock as the sum of the final goods inventory plus work-in-progress (WIP) inventory. We include WIP following the arguments in Ramey and West (1999) that changes in final goods inventory only accounts for a small and relatively-smooth fraction of inventory investment, and therefore does not present a full picture.

Following Blinder (1986) and Kryvstov and Midrigan (2013), we define our production measure as the sum of sales and inventory investment. We take logarithms of the raw data and detrend using the HP filter.[11] Sales, x_t, are from different sector, and we take the logarithm and detrend it to get HP filtered $\log(x_t)$. Monthly output are constructed from $y_t = x_t + \Delta i_t$, where i_t is the sectorial inventory stock. For wholesale and retail sectors, y_t is associated with deliveries. We take logarithm of y_t and detrend using HP filter. Since inventory investment can be negative, the measurement of Δi_t is different. We define $\Delta i_t = (i_t - i_{t-1})/x_t$. The other exception is the inventory-to-sales ratio, i_t/x_t, which is defined as level of inventory stock divided by level of sales. The data moments corresponding to: (i) the relative volatility of production to sales, (ii) the relative volatility of inventory investment to sales, (iii) the contemporaneous correlation between production with sales, (iv) the contemporaneous correlation between inventory investment and sales, (v) the contemporaneous correlation between the inventory-to-sales ratio and sales, (vi) the contemporaneous correlation between the change in inventories and production, (vii) the contemporaneous correlation between the inventory-to-sales ratio and production, (viii) the first-order autocorrelation of production, (ix) the first-order autocorrelation of the change in inventories, and (x) the first-order autocorrelation of the inventory-to-sales ratio are then calculated based on the treated data accordingly. The persistence of sales are estimated from detrended $\log(x_t)$.

Table 1 summarize the key moments we focus on with GMM Newey-West standard errors in parentheses. It reports the statistics for the manufacturing sector. We will mainly focus on Table 1 when explaining the results because they are related to production side, but the basic pattern is similar across all sectors. As the first row of the table shows, the ratio of production volatility to

[10]The $1967-1996$ raw data are measured as chained 1996 dollars to the SIC level, and data in 1997 and onwards are measured as chained 2009 dollar to the NAICS level. After adjusting for deflation and possible small jumps because of different estimation standards, we connect the two sub-sample periods.

[11]Blinder and Maccini (1991) detrend the data by regressing on a constant, a time trend, and an OPEC variable. They find that the OPEC variable has a very small effect on the results. In our data, the persistence of linearly-detrended sales is close to 1, suggesting a random walk process.

sales volatility is larger than 1. The statistics are all significant at the 1% level. The second row of Table 1 shows the relative volatility of inventory investment to sales. Although the ratio looks small, once one notes that inventory investment is less than 1 percent of sales, it looks much bigger.

The third row shows the correlation of output and sales. They move with each other almost one-to-one and the standard error is quite small. The fourth row displays the correlation between inventory investment and sales. Inventory investment is procyclical – its correlation with sales is 0.3 in the manufacturing sector.[12]

The fifth row reports the correlation between the inventory-to-sales ratio and sales, which is negative. The counter-cyclical inventory-to-sales ratio is also documented in Khan and Thomas (2007, 2016), Wen (2011), and Kryvtsov and Midrigan (2013).

The sixth and seventh rows display the correlation between inventory investment and output, as well as the correlation between inventory-to-sales and output. They follow the same pattern with the fourth and fifth rows about the cyclicity of inventory investment and inventory-to-sales ratio.

The eighth row reports the autocorrelation of output. Its autoregressive coefficient is close to that of sales. The ninth row displays the autocorrelation of inventory investment, which is quite small compared to other macroeconomic aggregates.

The tenth row gives the autocorrelation of inventory-to-sales ratio. It demonstrates that inventories not only adjust partially to the change in sales, but also the adjustment of inventory-to-sales ratio to its steady state quite slow.

3 A Production-Cost Smoothing Model With Ambiguity Aversion

In this section, we will introduce our production-cost smoothing model with ambiguity aversion. We first describe the basic model environment and then introduce Knightian uncertainty and ambiguity aversion.

[12]The excessive variance ratio and positive correlation between inventory investment and sales over the busisness cycle are widely documented in the literature using different filtering or detrending methods, including those mentioned in the main text.

3.1 The Basic Model Environment

The basic model follows the literature on the production-cost smoothing model (Blinder 1986 and Eichenbaum 1989) closely.[13] Our ultimate goal is to examine how uncertainty about the sales process and the production cost interact with the preference for robustness and affects the joint dynamics of production, inventories, and sales. Specifically, following Blinder (1986) and Eichenbaum (1989), we assume that the production cost function of the representative firm is

$$C\left(y_t\right) = \alpha_y \Gamma_t y_t + \frac{1}{2}\alpha_y y_t^2, \tag{1}$$

where y_t denotes the production level for the firm in period t. The condition, $\alpha_y > 0$, indicates that we assume the firm operates in a region of rising marginal costs. This function thus embodies the production *level* smoothing motive. Γ_t is a stochastic shock to the marginal cost of output, so firms have the incentive to raise production when this cost is low, provided the shock persists into the future. That is, (1) also embodies the production *cost* smoothing motive. The firm observes the cost shock before choosing the production plan.

The cost shock is generated by the AR(1) process

$$\Gamma_{t+1} = \rho_\Gamma \Gamma_t + w_{t+1}, \tag{2}$$

where $\rho_\Gamma \in [0, 1)$ is the persistence coefficient and w_t is an iid Gaussian innovation to the cost with mean 0 and variance Ψ. Both the persistence of the shock ρ_Γ and the variance Ψ are unknown to the econometrician and will be calibrated from the data. We model this cost shock as an unobservable disturbance as the cost shock explanation works best this way in the the literature, while observable cost shifters such as real wages and interest rates seem to play little role in inventory dynamics.[14]

The inventory holding cost function is

$$H\left(i_t\right) = \frac{1}{2}\alpha_i i_t^2, \alpha_i > 0, \tag{3}$$

where i_t denotes the inventory level for the firm at the end of period t. The coefficients α_i and α_y govern the relative importance of the backlog cleanup motive and the production-smoothing

[13] Eichenbaum (1989) finds supportive empirical evidence for the production-cost smoothing model, in which inventories serve to smooth production cost rather than as a buffer stock to production levels. Blinder (1986) and West (1987) retain the assumption of convex adjustment cost functions and considered production-cost smoothing models rather than the production-level smoothing models.

[14] See Ramey and West (1999) for a discussion.

motive in the firm's cost structure, respectively. We define $\alpha_{iy} \equiv \alpha_i/\alpha_y$ to represent this relative importance, as only this ratio plays a role in the model. In some studies, such as Kahn (1987), Eichenbaum (1989), Ramey (1991), and Ramey and West (1999), the inventory holding cost takes the form $H(i_t) = \alpha_i \left(i_{t-1} - \alpha_x x_t\right)^2/2$, where x_t is the amount of sales for the firm in t and $\alpha_x > 0$ governs the stockout avoidance motive, which induces the firm to hold inventories even without production cost considerations. To keep our baseline model as simple as possible, we focus only on the case $\alpha_x = 0$. In Section 6, we show that model uncertainty can provide a microfoundation for the stockout avoidance motive.

The accounting identity relating output, sales, and inventories is

$$y_t = x_t + i_t - i_{t-1}, \tag{4}$$

where x_t is the amount of sales for the firm in t. Given sales and initial inventories, the firm makes optimal production and inventories to minimize the following discounted expected cost:

$$v\left(i_{t-1}, x_t, \Gamma_t\right) = \min_{\{y_t\}} \mathbb{E}_t \left\{ \sum_{j=t}^{\infty} \beta^{j-t} \left[C\left(y_t\right) + H\left(i_t\right)\right] \right\}, \tag{5}$$

subject to (1), (3), and (4).

To close the model, we need to specify the exogenous sales processes. The sales process is governed by the AR(1) process

$$x_{t+1} = \rho x_t + \varepsilon_{t+1}, \tag{6}$$

where $\rho \in [0, 1)$ is the persistence coefficient and ε_{t+1} is an iid Gaussian innovation to sales with mean $(1-\rho)\overline{x}$ and variance Ψ. Here, ε_t can be interpreted as the demand shock. In addition, we assume $\operatorname{corr}\left(\varepsilon_t, w_t\right) = 0$, i.e., there is no correlation between the innovations to the cost and sales shocks.

Finally, the timing of our model economy is as follows. In period t, a firm currently holding an inventory stock i_{t-1} observes the demand and cost shocks, produces an amount of y_t of a single good, and sells an amount of x_t at the prevailing price. Production is instantaneous, and there is no delay between production and sales. In this sense, inventory is not a speculative tool to meet unexpected demand, but rather a buffer tool to smooth costly production. Finally, we assume that there are also no intermediate goods in our model and all goods are finished goods or works in progress and are free of physical depreciation.

3.2 The RE Solution and Its Implications

Before we examine the effects of ambiguity aversion on the joint behavior of inventory and production, we first present the analytical expressions for the optimal inventory and production rules under RE. Under the RE assumption, firms have complete confidence and do not worry about model misspecification. Using these expressions, we can clearly see how inventory and production respond to the sales and the cost shocks, and thus explain why the RE version of the PCS model fails to reproduce two stylized facts widely documented in the literature: (i) the relative volatility of production to sales, and (ii) the positive contemporaneous correlation between sales and inventories.

Following the same procedure used in Eichenbaum (1989) and Ramey and West (1999), we can derive the analytical inventory and production rules under RE,

$$\begin{aligned} i_t &= \lambda i_{t-1} \\ &+ \left\{ \underbrace{-\lambda}_{<0} x_t + \underbrace{(1-\lambda)}_{>0} \underbrace{\mathbb{E}_t \left[\sum_{j=1}^{\infty} \left(\frac{\lambda}{R} \right)^j x_{t+j} \right]}_{\text{expected future sales}} \right\} + \left\{ \underbrace{-\lambda}_{<0} \Gamma_t + \underbrace{(1-\lambda)}_{>0} \underbrace{\mathbb{E}_t \left[\sum_{j=1}^{\infty} \left(\frac{\lambda}{R} \right)^j \Gamma_{t+j} \right]}_{\text{expected future cost shocks}} \right\}, \end{aligned} \tag{7}$$

$$\begin{aligned} y_t &= (\lambda - 1)\, i_{t-1} \\ &+ \left\{ \underbrace{(1-\lambda)}_{>0} x_t + \underbrace{(1-\lambda)}_{>0} \underbrace{\mathbb{E}_t \left[\sum_{j=1}^{\infty} \left(\frac{\lambda}{R} \right)^j x_{t+j} \right]}_{\text{expected future sales}} \right\} + \left\{ \underbrace{-\lambda}_{<0} \Gamma_t + \underbrace{(1-\lambda)}_{>0} \underbrace{\mathbb{E}_t \left[\sum_{j=1}^{\infty} \left(\frac{\lambda}{R} \right)^j \Gamma_{t+j} \right]}_{\text{expected future cost shocks}} \right\}, \end{aligned} \tag{8}$$

where λ and $\lambda^{-1}R$ are the roots of $X + R/X = 1 + R + R\alpha_i/\alpha_y$, with $\lambda \in (0,1)$. Inventory and production decision rules can be decomposed into a linear combination of the lagged level of inventories, current and expected future sales, and current and expected future costs.

It is clear from (7) and (8) that under RE, both inventory and current production increase with expected future sales. The intuition behind this result is that when firms expect higher sales in the future, they increase their current production and replenish inventories to smooth production levels. From (7), we can also see that inventory depends negatively on the current level of sales. This dependence captures the notion that, in the presence of increasing marginal costs, firms would prefer to meet an additional increase in current sales by both reducing existing inventories by λ and increasing current production by $1 - \lambda$.

For the effects of the cost shock, we can also see from (7) and (8) that both inventory and current production increase with expected future cost shocks. This motive for inventory accumulation reflects the fact that firms will build up inventories (via production) in periods when costs are relatively low, and meet future sales out of these stocks of inventories. In this sense, inventories can serve to smooth production costs. In contrast, both desired inventories and production depend negatively on the current cost shock. Other things being equal, firms will meet sales out of inventories rather than producing new output in periods when marginal production costs are high.[15]

Given the sales and the cost shocks are generated by the AR(1) processes (2) and (6), the coefficients on i_{t-1}, x_t, and Γ_t are

$$y_t = \mu_i i_{t-1} + \mu_x x_t + \mu_\Gamma \Gamma_t, \tag{9}$$

$$i_t = \lambda_i i_{t-1} + \lambda_x x_t + \lambda_\Gamma \Gamma_t, \tag{10}$$

where

$$\lambda_i = \lambda \in (0,1)\,,\ \lambda_x = -\lambda \frac{R-\rho}{R-\rho\lambda} \in (-1,0)\,,\ \lambda_\Gamma = -\lambda \frac{R-\rho_\Gamma}{R-\rho_\Gamma \lambda} < 0;$$

$$\mu_i = \lambda - 1 \in (-1,0)\,,\ \mu_x = 1 + \lambda_x \in (0,1)\,,\ \mu_\Gamma = \lambda_\Gamma < 0.$$

Clearly, from (10), the net effect of the sales on inventories (λ_x) is negative. The negative effect of an increase in current sales will be partly offset by the positive effect of this increase on expected future sales. Similarly, the net effect of the cost shock on inventories λ_Γ is also negative and the negative effect of an increase in current costs will be partly offset by the positive effect of this increase on expected future costs. Using (9) the net effect of sales on production is $\mu_x = 1 + \lambda_x \in (0,1)$. In the presence of increasing marginal costs, firms would prefer to meet any additional increase in current sales by both reducing existing inventories by λ_x and increasing current production by μ_x. Similarly, the net effect of the cost shock on current production μ_Γ is also negative because the negative effect of an increase in current costs will be partly offset by the positive effect of this increase on expected future costs.

Equation (10) predicts that inventories are countercyclical since $\lambda_x < 0$, which is inconsistent with empirical evidence. Inventory investment and sales are generally positively correlated (see Blinder 1991 and Ramey and West 1999). Furthermore, using the identity equation linking production, sales

[15] Given (1), it is straightforward to show that the marginal cost of production is $\alpha_y (\Gamma_t + y_t)$.

and inventory investment, we can infer that the negative correlation between inventory investment and sales predicted by the RE model will result in a relative volatility of production to sales smaller than one. In our quantitative analysis in Section 4, we find that production is less volatile than sales even with cost shocks, which is also inconsistent with the data. In this US economy production is more volatile than sales in all major sectors and in most industries.

3.3 Introducing Model Uncertainty Due to Robustness

In this section, we model firms' aversion to ambiguity about the true dynamics of sales and production costs using the preference for robustness (RB) proposed by Hansen and Sargent (2007).[16] Specifically, we assume that firms do not know the true probability model generating the data, and are thus concerned about the possibility that their estimated model is misspecified in a manner that is difficult to detect statistically. In other words, the firms accept the approximating (reference) model governed by (2) and (6) as the best estimated model using available information, but are still concerned that it is misspecified. They therefore consider a range of models (i.e., the distorted model) surrounding the reference model when computing the continuation payoff, and then make optimal decisions given the worst-case scenario.

A tractable version of robust control considers the question of how to make decisions when the agent does not know the true data-generating model and considers a range of models surrounding the given approximating model, (2) and (6):

$$x_{t+1} = \rho x_t + \varepsilon_{t+1} + \omega_{t+1}^1, \tag{11}$$

$$\Gamma_{t+1} = \rho_\Gamma \Gamma_t + w_{t+1} + \omega_{t+1}^2, \tag{12}$$

where $\omega_{t+1} \equiv \left\{\omega_{t+1}^1, \omega_{t+1}^2\right\}$ distorts the mean of the innovations to the sales and cost shocks, and makes decisions that maximize expected total costs given this worst possible model (known as the distorted model).[17] To ensure that the true model, (2) and (6), is a good approximation when the distorted model, (11) and (12), generates the data, we constrain the approximation errors by an

[16] Both a preference for "wanting robustness" proposed by Hansen and Sargent and "ambiguity aversion" proposed by Epstein and his coauthors can be used to capture the same idea of the multiple priors model of Gilboa and Schmeidler (1989).

[17] Formally, this setup is a game between the decision-maker and a malevolent nature that chooses the distortion process ω_{t+1}.

upper bound η_0:

$$\mathbb{E}_0\left[\sum_{t=0}^{\infty}\beta^{t+1}\omega_{t+1}^T\omega_{t+1}\right]\leqslant\eta_0. \tag{13}$$

where $\mathbb{E}_0[\cdot]$ denotes conditional expectations evaluated with the approximating model. The left side of this inequality is a statistical measure of the discrepancy between the distorted and approximating models. Note that the standard RE case corresponds to $\eta_0 = 0$. In the general case in which $\eta_0 > 0$, an evil agent is given an intertemporal entropy budget $\eta_0 > 0$ which restricts the set of models that the agent views as plausible alternatives to the approximating model. Following Hansen and Sargent (2007), we compute robust decision rules by solving the following two-player zero-sum game: a minimizing decision maker chooses the optimal inventory and production decisions and a maximizing evil agent chooses the model distortion process $\{\omega_{t+1}\}$:

$$v\left(i_{i-1},x_t,\Gamma_t\right)=\min_{\{\omega_{j+1}\}}\max_{\{y_j\}}\left\{-\left(C(y_t)+H(i_t)\right)+\beta\mathbb{E}_t\left[\frac{1}{2}\theta\omega_{t+1}^T\omega_{t+1}+v(i_i,x_{t+1},\Gamma_{t+1})\right]\right\} \tag{14}$$

subject to the distorted state transition equation (i.e., the worst-case model) (11) and (12), where $\vartheta > 0$ is the Lagrangian multiplier on the entropy constraint specified in (13) and controls how bad the error can be. (13) and (14) are called "*constraint preference*" and "*multiplier preference*," respectively. As shown in Hansen and Sargent (2010), there is a one-to-one correspondence between η_0 in (13) and ϑ in (14).

Following Hansen and Sargent (2007, 2010), we use the risk-sensitivity operator, a special case of recursive utility, to characterize the preference for robustness. The robust control version of the inventory model proposed in Section 3.1 can be expressed as

$$v(i_{i-1},x_t,\Gamma_t)=\min_{\{y_t\}}\left\{\frac{1}{2}\alpha_y y_t^2+\alpha_y\Gamma_t y_t+\frac{1}{2}\alpha_i i_t^2+\beta\mathcal{R}_t v(i_t,x_{t+1},\Gamma_{t+1})\right\}, \tag{15}$$

subject to (2), (4), and (6), and the distorted expectation operator

$$\mathcal{R}_t\left[v(i_t,x_{t+1},\Gamma_{t+1})\right]=-\frac{2}{\alpha}\log\mathbb{E}_t\left[\exp\left(-\frac{\alpha}{2}v\left(i_t,x_{t+1},\Gamma_{t+1}\right)\right)\right], \tag{16}$$

where $\alpha > 0$ measures the degree of robustness in the model. The operator $\mathcal{R}$ distorts the usual conditional expectation with a single parameter $\alpha > 0$. The reason for the equivalence between risk-sensitive and robust control is that imposing a constraint on conditional relative entropy is equivalent to an agent having a preference that applies an exponential transform to the continuation value.[18] Replacing with the form of $\mathcal{R}$ delivers the risk-sensitive evaluation used in control theory, and the same results can be obtained as when solving the above max-min problem with a penalty term defined in robust control theory, by setting $\alpha = -\theta^{-1}$.[19] If $\alpha = 0$, the (15) specification reduces to the standard RE expected utility specification, which is equivalent to solving a max-min problem with θ diverging to $+\infty$.

Denote $s_t = \begin{bmatrix} i_{t-1} & x_t & \Gamma_t \end{bmatrix}^T$ as the state vector. Using (4), (15) can be rewritten as

$$v\left(i_{t-1}, x_t, \Gamma_t\right) = \min_{\{y_t\}} \left\{ R y_t^2 + 2 s_t^T W y_t + s_t^T Q s_t + \beta \mathcal{R}_t \left[v\left(i_t, x_{t+1}, \Gamma_{t+1}\right) \right] \right\}, \tag{17}$$

subject to the state transition equations

$$s_{t+1} = A s_t + B y_t + C \overrightarrow{\varepsilon}_{t+1}, \tag{18}$$

where $R = 1 + \alpha_y/\alpha_i$, $W = \begin{bmatrix} 1 \\ -1 \\ \alpha_y/\alpha_i \end{bmatrix}$, $Q = \begin{bmatrix} 1 & -1 & 0 \\ -1 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix}$, $A = \begin{bmatrix} 1 & -1 & 0 \\ 0 & \rho & 0 \\ 0 & 0 & \rho_\Gamma \end{bmatrix}$,
$B = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}$, $C = \begin{bmatrix} 0 & 0 \\ \omega & 0 \\ 0 & \omega_\Gamma \end{bmatrix}$, and $\overrightarrow{\varepsilon}_{t+1} = \begin{bmatrix} \varepsilon_{t+1} \\ w_{t+1} \end{bmatrix}$.

To solve the linear quadratic robust control problem, we use the adapted ordinary optimal linear regulator approach proposed in Hansen and Sargent (2007). The first-order conditions of the problem requires us to solve an algebraic Ricatti equation for P:

$$P = Q + \beta A' \mathcal{D}(P) A - (\beta A' \mathcal{D}(P) B + W)(R + \beta B' \mathcal{D}(P) PB)^{-1} (\beta B' \mathcal{D}(P) PA + W'), \tag{19}$$

where $\mathcal{D}(P) = P + \alpha PC \left(I - \alpha C' PC\right)^{-1} PC'$. Under some regularity conditions on the cost functions, the Ricatti equation has a unique positive semidefinite solution.[20]

[18] Although risk-sensitive control and robust control lead to the same decision rules, they are based on different preferences and give the decision different interpretations. Risk-sensitive control suggests an (enhanced) risk averse agent, whereas robust control suggests an agent who is uncertain about the model that generates the data. These two preferences can be quite different with respect to welfare calculations.

[19] See Online Appendix B for a proof on the equivalence between the risk-sensitivity specification and a more primitive multiplier preference specification of the robust control problem.

[20] In addition to the usual detectability and stabilizability conditions, robust control requires that the evil agent's budget for misspecification is not too large (the breakdown condition). See Hansen and Sargent (2007) for details.

Since the model with ambiguity aversion implies that firms behave *as if* they hold worst-case beliefs, optimal outcomes are identical to that obtained in the model with expected cost minimizers who are *endowed* with those worst-case beliefs. However, this observational equivalence does not imply that the two models are identical – worst-case beliefs would not generally be invariant to changes in policies or structural features of the economy.[21]

3.4 Robust Policy Rules and the Stochastic Properties of the Joint Dynamics of Production, Inventories, and Sales

The following proposition summarizes the solutions to the optimization problem described by (31)-(32).

Proposition 1 *Given α, the production and inventory policy functions can be written as:*

$$y_t^* = \mu_i i_{t-1} + \mu_x x_t + \mu_\Gamma \Gamma_t, \tag{20}$$

$$i_t^* = \lambda_i i_{t-1} + \lambda_x x_t + \lambda_\Gamma \Gamma_t, \tag{21}$$

where $\lambda_i = 1 + \mu_i$, $\lambda_x = -1 + \mu_x$, $\mu_\Gamma = \lambda_\Gamma$.

Proof. In the presence of RB, the PCS model cannot be solved analytically, so we solve the above robust linear-quadratic problem numerically. Solving the RB-PCS model numerically yields the optimal control (20). Combining the accounting equation, $i_t - i_{t-1} = y_t - x_t$, with (20) yields (21). ■

We make several comments based on the solutions above to help the reader understand how ambiguity aversion influences the joint dynamics of production, sales, and inventories. First, as a general comment, note that as $\alpha \to 0$, $\mathcal{R}_t$ becomes the ordinary conditional expectation operator $\mathbb{E}[\cdot|\mathcal{I}_t]$, and the robust policies, (20) and (21), reduces to the corresponding RE policies, (9) and (10), in Section 3.2, respectively.

Second, given that the RB-PCS model can only be solved numerically, we cannot exactly inspect the mechanism about how RB affects the joint dynamics of production, inventories and sales. We instead use Figure 1 to illustrate how the response coefficients in the production and inventory

[21] Under worst-case beliefs policy could change the evolution of states, and therefore would change state-dependent worst-case distortions.

policies vary with the degree of RB (α) for different values of ρ when $R = 1.01$, $\rho_\Gamma = 0.118$, and $\alpha_i/\alpha_y = 0.552$.[22] Specifically, In the top of Figure 1, it shows how the three response coefficients, μ_i, μ_x, and μ_Γ, in the production rule vary with α; the bottom of the figure shows how the three response coefficients, λ_i, λ_x, and λ_Γ, in the inventory rule vary with α. It is clear from the top-second panel of the figure that μ_x increases with α for different values of ρ. Under RB, we no longer have explicit solutions; instead we have a numerical solution which is a linear combination of the lagged level of inventory, the current sales, and the current cost shock. The coefficients, μ_x and μ_Γ, in the production rule measure the net effects of current and expected future shocks. The intuition for the result that μ_x increases with α is as follows. Firms would prefer to meet an additional increase in current sales by increasing current production by $\mu_x = 1 - \lambda_x$ and changing (reducing or increasing) existing inventories by λ_x. Note that when the degree of RB is sufficiently strong, the value of μ_x is greater than one and the value of λ_x is positive, which means that current production responds to changes in current sales by more than one for one, and inventories and sales are positively correlated.

We can see from the bottom-second panel of Figure 1 that the response of inventory investment to sales (λ_x) changes from negative to positive as α becomes large. Firms' decisions are based on the worst-case scenario (i.e., they fear that future sales might be higher than what the approximating model generates, forcing them to incur high marginal cost), which generates precautionary behavior. Consequently, under RB, both production and inventory are more sensitive to changes in sales. This response pattern is often referred to as "*making hay while the sun shines.*" This mechanism allows the model to generate a procyclical inventory behavior and more volatile production relative to sales. Furthermore, this precautionary behavior due to ambiguity aversion endogenously generates the stock-out avoidance motive even though we do not model a stockout avoidance motive explicitly (such as imposing a non-negativity constraint on inventories or an "accelerator" term in the inventory holding cost function).

Third, for the effects of the cost shock, we can see the right panels of Figures 1 that both μ_Γ and λ_Γ are monotonically decreasing with α. If there is a reduction in current cost, firms build up inventories (via production) more aggressively in periods when costs are relatively low, and meet future sales out of these accumulated stocks of inventories. Again, the reason behind is that their

[22] In Section 4.1, we will discuss the choice of these preference and technology parameters and provide more details about how to estimate the sales and cost processes using the U.S. data. The main results in this section are robust to reasonable variation in these parameter values.

decisions are based on the worst-case scenario (i.e., they fear that future production cost might be higher than what the reference model generates), which also generates precautionary behavior.[23]

Finally, the fact that the absolute value of (λ_x) first declines and then increases suggests that the variance of inventory investment is likely to follow the same pattern. The volatility of inventory investment therefore has a non-monotonic relationship with the degree of ambiguity aversion – the volatility first declines and then increases as the degree of ambiguity aversion rises. We will explore this issue more in the quantitative section.

Using the numerical robust decision rules, we can readily compute the model's predictions on the key targeted moments of the joint dynamics of production, inventories, and sales. The following proposition summarizes the main results.

Proposition 2 *In the RB-PCS model, using (20) and (21), we can compute the following ten targeted moments:*

(i) the relative volatility of production to sales ($\mu_{yx} \equiv \operatorname{var}(y_t)/\operatorname{var}(x_t)$)*,*

(ii) the relative volatility of inventory investment to sales ($\mu_{ix} \equiv \operatorname{var}(\Delta i_t)/\operatorname{var}(x_t)$)*,*

(iii) the contemporaneous correlation between production with sales ($\rho_{yx} \equiv \operatorname{cov}(y_t, x_t)/\left(\sqrt{\operatorname{var}(y_t)}\sqrt{\operatorname{var}(x_t)}\right)$)*,*

(iv) the contemporaneous correlation between the change in inventories and sales ($\rho_{ix} \equiv \operatorname{cov}(\Delta i_t, x_t)/\left(\sqrt{\operatorname{var}(\Delta i_t)}\sqrt{\operatorname{var}(x_t)}\right)$)*,*

(v) the contemporaneous correlation between the inventory-to-sales ratio and sales ($\rho_{i/x,x} \equiv \operatorname{cov}(i_t/x_t, x_t)/\left(\sqrt{\operatorname{var}(i_t/x_t)}\sqrt{\operatorname{var}(x_t)}\right)$)*,*

(vi) the contemporaneous correlation between the change in inventories and production ($\rho_{iy} \equiv \operatorname{cov}(\Delta i_t, y_t)/\left(\sqrt{\operatorname{var}(\Delta i_t)}\sqrt{\operatorname{var}(y_t)}\right)$)*,*

(vii) the contemporaneous correlation between the inventory-to-sales ratio and production ($\rho_{i/x,y} \equiv \operatorname{cov}(i_t/x_t, y_t)/\left(\sqrt{\operatorname{var}(i_t/x_t)}\sqrt{\operatorname{var}(y_t)}\right)$)*,*

(viii) the first-order autocorrelation of production ($\rho_y \equiv \operatorname{cov}(\Delta y_t, \Delta y_{t-1})/\operatorname{var}(\Delta y_t)$)*,*

(ix) the first-order autocorrelation of the change in inventories ($\rho_i \equiv \operatorname{cov}(\Delta i_t, \Delta i_{t-1})/\operatorname{var}(\Delta i_t)$)*,*

(x) the first-order autocorrelation of the inventory-to-sales ratio ($\rho_{i/x} \equiv \operatorname{cov}(i_t/x_t, i_{t-1}/x_{t-1})/\operatorname{var}(i_t/x_t)$)*.*

Proof. See Online Appendix C for the detailed expressions for the first three targeted moments. Note that there are no explicit expressions for some moments. ■

[23]Note that Figure 1 shows that the coefficients of the decision rules are not sensitive to the persistence of sales given plausible degrees of ambiguity aversion.

In the next section we will quantitatively evaluate the model's ability to match these ten features of the data. In Section 7.1, we also derive some analytical solutions to the model without ambiguity aversion and robustness to help explain why the basic model fails.

4 Model's Quantitative Implications

In this section, we estimate and calibrate the model parameters and quantitatively examine the model's ability to fit the data.

4.1 Estimating and Calibrating the Parameters

There are four key parameters we need to pin down in this model: the degree of ambiguity aversion (α), the persistence of cost shock (ρ_c), the volatility of the cost shock innovation(Ψ), and the relative importance of the inventory holding cost (α_{iy}). We jointly estimate these parameters to match the ten moments (i) to (x) defined in the above proposition. The estimation is based on a Generalized Method of Moments (GMM) approach. (See Online Appendix D for details.)

The estimated values of the parameters are reported in the top panel of Table 2. As the table shows, the estimated values of the parameters are similar under these two approaches, except for the persistence of the cost shock. The optimal weighting matrix estimate of the persistence coefficient of the cost shock is a bit lower than the values estimated in Eichenbaum (1989), while the estimate with the identity matrix is close to that in Eichenbaum (1989). The GMM estimate of the relative importance of inventory holdings is within the range summarized in Ramey and West (1999). The bottom panel of Table 2 compares the RB model's performance on the ten moments with two RE models with different calibration strategies. We first obtain the RE model-implied moments by setting the RB parameter to be zero. (All the other parameter values are set to be the same for both the RB and RE models.) We then do a validation test denoted by "the RE model*" that is estimated to target the same moments in the data.

To provide a meaningful interpretation for the ambiguity aversion and robustness parameter (α), we follow Anderson et al. (2003) and Hansen and Sargent (2007) by calculating the detection error probability (DEP). (See Online Appendix E for the details.) The resulting DEP is 0.22, which means that there is 22 percent chance that a likelihood ratio test will improperly select between the approximating and distorted models. This value is within the reasonable range of the DEP in the literature. Hansen and Sargent (2007) argue that the DEP values between 0.1 and 0.3 are plausible. In the recent studies, Djeutem and Kasa (2013) show that to match the observed volatility of six U.S. dollar exchange rates (the Australian dollar, the Canadian dollar, the Danish dollar, the Japanese yen, the Swiss franc, and the British pound), the detection error probability should be set between 7.5 percent to 13.1 percent. In contrast, Kasa and Lei (2017) uses values above 40 percent in their models. As a higher DEP means a lower degree of model uncertainty, it suggests that our model does not require unreasonable fears of model misspecification to fit the data.

To further help understand how α influences the key moments, Figure 2 illustrates how the selected six moments vary with α. The figure shows that the degree of ambiguity aversion has a non-monotonic effect on the relative volatilities of production and inventory investment (the top-left two charts), and also it can alter the sign of the correlations for inventory investment and inventory-to-sales ratio with sales. We have briefly discussed these using the decision rules in the previous section, but now we will dig in deeper.

First, as discussed in the previous section, the U-shaped relationship between μ_{ix} and α (the top-second panel) is mainly driven by the effect of α on λ_x. As α increases, the absolute value of λ_x first declines and then rises (see Figure 1), which makes inventory investment less sensitive to sales fluctuations for moderate values of α. As this effect dominates other effects (such as those coming from changes in λ_i and λ_Γ), μ_{ix} first declines and then increases with α.

The U-shaped relationship between the volatility of inventory investment and the degree of ambiguity aversion also helps explain a similar U-shape between μ_{yx} and α (the top-left panel of Figure 2). We can use the identity equation (4) to decompose μ_{yx} into

$$\mu_{yx} = \frac{\operatorname{var}(\Delta i_t)}{\operatorname{var}(x_t)} + 2\frac{\operatorname{cov}(\Delta i_t, x_t)}{\operatorname{var}(x_t)} + 1. \tag{22}$$

The upper panel of Figure 3 shows that the initial decline in the variance ratio $\operatorname{var}(y_t)/\operatorname{var}(x_t)$ (when α is small) is mainly driven by the component $\operatorname{var}(\Delta i_t)/\operatorname{var}(x_t)$, while the latter increase (when α is relatively large) is a joint effect of both $\operatorname{cov}(\Delta i_t, x_t)$ and $\operatorname{var}(\Delta i_t)/\operatorname{var}(x_t)$, which both increase with α on the right side of the figure.

Turning to the top-right panel of Figure 2, two factors help explain why an increase in α could switch the sign of the correlation between inventory investment and sales. The lower panel of Figure 3 plots how different parts of this correlation changes with α based on the definition $\rho_{ix} = \operatorname{var}(\Delta i, x)/(\sigma_{\Delta i}\sigma_x)$. Notice that σ_x is given so we only plot the other two components. The change in the sign of the correlation is due to the same change in the covariance term. As we explained in the previous section, this sign switch is due to the change in λ_x, which flips from negative to positive as α rises. Furthermore, the jump in the magnitude reflects the fact that both $|\operatorname{cov}(\Delta i_t, x_t)|$ and $\sigma_{\Delta i}$ drop to zero but the latter declines faster than the former, making the ratio jump up when the sign switches. The effect of α on the correlation between inventory investment and output also follows the same logic.

For the correlation between the inventory-to-sales ratio and sales, the sign of inventories generated by the model matters. In the RE case, when there is a positive cyclical shock to sales, the model predicts that inventories are used as a buffer stock and will be depleted, which implies countercyclical inventory investment. As we will show in Section 6, without the stock-out avoidance assumption the RE model generates negative inventories in most periods. As the reduction in inventories is less proportional than the increase of the sales level, we find that the inventory-to-sales ratio actually increases after the shock, suggesting a procyclical inventory-sales ratio. As shown in the left panel of Figure 4, after a positive sales shock in the middle of periods, the level of sales increases, while that of inventories decreases under a RE-PCS model, and the overall inventory-sales ratio increases. The right panel of Figure 4 clearly shows that a positive sales shock increases the firm's inventories under RB, though the increase is proportionally less than the increase of the sales level, and the inventory-to-sales ratio displays countercyclicality. Here the bottom line is, the presence of RB reverses the cyclicality of inventory investment without making the firm accumulate too much inventory, and preserves the right signs of inventories simultaneously. These results guarantee a right sign on the correlation between the inventory-to-sales ratio and sales, as well as that on the correlation between the inventory-to-sales ratio and output.

Finally, ambiguity also decreases the persistence of the inventory-sales ratio, as now the adjustment of inventories is faster.

4.2 Evaluation of Models' Key Performance

We now evaluate our RB model's quantitative ability to fit the data, and compare it explicitly to the standard RE model. We report the results in the lower panel of Table 2. The first column reports the empirical moments over the $1967 - 2018$ period, the second column reports the RB model's predictions, the third column reports the RE model's predictions, and the fourth column reports the

RE validation test results.

It is clear from the last three columns of Table 2 that the RB model did a better job than the two calibrated RE models along most of the key moments. Specifically, the first row of the lower panel compares the different models' predictions for the relative volatility of production to sales. It shows that the RB model generates a relative volatility of production to sales significantly above 1 as in the data, which solves the production-smoothing puzzle. In contrast, both RE models predict a value less than 1. As reported in the previous section, the required degree of model uncertainty, as measured by the DEP, is 0.22, which means that firms in our RB model economy face a probability of 22% of being unable to distinguish the distorted model from the approximating model. We have argued above that this value is reasonable.

The second row of the lower panel compares the models' predictions for the relative volatility of inventory investment to sales. It is clear from this row that the results of the RE and RB models are close. As we explained in the previous section, the RB model has the potential to reduce the volatility of inventory investment because the response of inventory investment to sales (λ_x) increases towards zero when the degree of RB (α) increases.

The third and fourth rows compare the models' predictions for the correlations between inventory investment and the ratio of sales to production. As explained before, both the PS and PCS models fail to match this fact; the RE model predicts the wrong sign of the correlation between inventory investment and sales. In contrast, the RB model generates the correct sign and a size very close to the empirical counterpart. Again, the change from a negative sign to positive relies on the fact that an increase in ambiguity aversion leads to production changing more than sales following a demand shock and thus causes inventory investment to move in the same direction as sales.

In the fifth and sixth rows, our RB model produces a countercyclical inventory-to-sales ratio, while the RE model does not. In the RB model inventories move less than one-for-one with sales. In contrast, in the RE model, production can not respond to sales adequately. Without incorporating the stockout-avoidance constraint, it generates negative inventories and a procyclical inventory-to-sales ratio.

The seventh row compares the correlation between sales and output. The correlation is slightly higher under RB since production responds more aggressively to sales. Both models produce a reasonably high correlation between the two aggregates, considering the small portion of inventories in the economy.

Eighth and ninth rows are about the autocorrelations of output and inventory investment. As output and inventories become more volatile and are more responsive to the sale shock, their autocorrelations naturally decrease. The RB model hence generates a smaller ρ_i, which is closer to the empirical counterpart; and a smaller ρ_y, which is somewhat lower than in the data.

Finally, as noted in Coen-Pirani (2004), the inventories-to-sales ratio for finished goods is persistent over the business cycle. It turns out that RB increases inventories proportionally more than the reduction of inventories under RE, so the adjustment speed of the inventory-to-sales ratio under RB is comparatively faster. The persistence of inventory-to-sales ratio is therefore lower under RB. Nevertheless, both models deliver a highly persistent inventory-sales ratio.

5 Extension to Endogenous Sales

In this section, we extend our benchmark model with exogenous sales to a model with endogenous sales and prices. We then show the main results obtained in our benchmark model still hold in this extension, and this model also outperforms the corresponding RE model both qualitatively and quantitatively.

5.1 The Model with Ambiguity Aversion and Endogenous Sales

In this section, we assume that sales is now *endogenously* determined. Specifically, following Bils and Khan (2000), we assume that for a given price, the amount of sales is an increasing function of its available stock of products (with an elasticity of $\phi \in (0, 1]$):

$$x_t = d(p_t)a_t^{\phi}, \tag{23}$$

where $d(\cdot)$ is the demand function and $a_t = i_{t-1} + y_t$ is the stock of products available for sale at time t.[24] Firms optimally choose the price of product, p_t, to maximize its expected profits. When $\phi = 1$, the specification is consistent with a competitive market that allows for the possibility of stockout, as firms cannot sell more than its inventories remained from the previous period plus current output. The demand for the single good, x_t, is determined by

$$d(p_t) = A_0 + A_1 p_t + v_t, \tag{24}$$

[24]Note that i_{t-1} is end-of-period inventory held at time $t - 1$.

where v_t is a demand shock. v_t follows an AR(1) process:

$$v_{t+1} = \rho_v v_t + C_\varepsilon \varepsilon_{t+1}. \tag{25}$$

where ε_{t+1} is an iid Gaussian innovation with mean 0 and variance 1, and C_ε is a positive constant. As in the benchmark model, the representative firm faces both production costs and inventory holding costs:

$$C(y_t, i_t, \Gamma_t) = \alpha_y \Gamma_t y_t + \frac{1}{2}\alpha_y y_t^2 + \frac{1}{2}\alpha_i i_t^2,\ \alpha_i, \alpha_y > 0, \tag{26}$$

with

$$\Gamma_{t+1} = \rho_\Gamma \Gamma_t + C_w w_{t+1}, \tag{27}$$

where w_{t+1} is an iid Gaussian innovation with mean 0 and variance 1, and C_w is a positive constant.

Following the same argument as in the benchmark model, to incorporate Knightian uncertainty due to ambiguity aversion we assume that the firm has concerns about model misspecification about both demand and cost shocks, v_t and Γ_t, and seeks robust decision rules. When the firm optimally chooses its price p_t and production y_t subject to Knightian uncertainty, the problem can be formulated as the following two-player zero-sum game:

$$\min_{\{\omega_{t+1}\}} \max_{\{y_t,\, p_t\}} \mathbb{E}_0 \left\{ \sum_{t=0}^{\infty} \beta^t \left[p_t x_t - C(y_t, i_t, \Gamma_t) + \beta \eta \omega_{t+1}^2 \right] \right\}, \tag{28}$$

subject to (4), (23), and

$$\Gamma_{t+1} = \rho_\Gamma \Gamma_t + C_w \left(w_{t+1} + \omega_{t+1}^1 \right), \tag{29}$$

$$v_{t+1} = \rho_v v_t + C_\varepsilon \left(\varepsilon_{t+1} + \omega_{t+1}^2 \right), \tag{30}$$

with i_{t-1} is given and $\eta > 0$ is the penalty coefficient. By choosing $\omega_{t+1} = \left(\omega_{t+1}^1, \omega_{t+1}^2\right)$ the hypothetical 'evil agent' selects the worst-case model. By choosing the optimal price and output, the ambiguity-averse firm pins down current sales and revenue.

We linearize the above model around the steady state and solve for linear decision rules. The key distinction between our benchmark model with exogenous sales and this model with endogenous price and sales is that in this model we need to linearize the revenue function $d(p_t)\, a_t^\phi$ around the steady states of price, inventories and production, and then we can transform it into a linear quadratic regulator problem. Denote $s_t = \begin{bmatrix} 1 & i_{t-1} & \Gamma_t & v_t \end{bmatrix}^T$ as the state vector and $u_t = \begin{bmatrix} y_t & p_t & \omega_{t+1} \end{bmatrix}^T$ as the control vector. We then get the following lemma:

Lemma 3 *Given (4), (23), (29), and (30), (28) can be rewritten as:*

$$V(s_t) = \min_{\{\omega_{t+1}\}} \max_{\{y_t,\, p_t\}} \left\{ Ru_t^2 + 2s_t^T W u_t + s_t^T Q s_t + \beta \mathbb{R}_t V(s_{t+1}) \right\}, \tag{31}$$

subject to the state transition equations:

$$s_{t+1} = As_t + Bu_t + C\overrightarrow{\varepsilon}_{t+1}, \tag{32}$$

where the elements in Q, W, R, A, B, C are specified by the steady state values of the model variables.

Proof. See Online Appendix F for the details. ■

The following proposition summarizes the solutions to the above optimization problem.

Proposition 4 *Given α, the production and inventory policy functions with endogenously sales can be written as:*

$$y_t^* = constant_0 + \mu_i i_{t-1} + \mu_\Gamma \Gamma_t + \mu_v v_t, \tag{33}$$

$$p_t^* = constant_1 + \varphi_i i_{t-1} + \varphi_\Gamma \Gamma_t + \varphi_v v_t, \tag{34}$$

$$i_t^* = constant_2 + \lambda_i i_{t-1} + \lambda_\Gamma \Gamma_t + \lambda_v v_t, \tag{35}$$

where $constant_0$, $constant_1$, and $constant_2$ are some constants, and μ, φ and λ are the corresponding response coefficients.

Proof. As mentioned in Section 3.3, solving the robust dynamic programming numerically yields the optimal controls (33) and (34). Combining the accounting equation, $i_t - i_{t-1} = y_t - x_t$, with (33) yields (35). ■

As in the benchmark model, given the complexity of the decisions, we cannot learn much from them. In the next subsection we will quantitatively evaluate the model's implications on explaining the same key dimensions of the joint dynamics of production, inventories and sales.

5.2 Quantitative Evaluation

In Bils and Khan (2007), the empirically estimated elasticity ϕ ranges from -0.049 to 0.486 in the six manufacturing industries: tobacco, apparel, lumber, chemicals, petroleum, and rubber. The average of ϕ is approximately 0.1, so we choose that value. We construct monthly output using the inventories and sales data as before, detrend it using the HP filter, and take the exponential of the cyclical component. We also take logs of Consumer Price Index (CPI), sales, inventory, and then take exponential of the HP-filtered cycles. This provides us with stationary positive aggregate numbers that facilitates our calibration. In the next step, we regress detrended sales to scaled detrended stocks (s_t/a_t^{ϕ}) and price (p_t) to obtain the demand function $d(p_t)$ and recover the values of A_0 and A_1. The volatility of the demand shock ε_t can also be obtained at the same time. Note that the first-order conditions with respect to y_t and p_t can be used to derive the following two equations:

$$\mathbb{E}_t\left[\frac{p_t s_t}{(i_t+y_t)\,MC_t}+\frac{i_t+y_t-\beta MC_{t+1}}{MC_t}\left(\frac{\phi s_t}{i_t+y_t}-1\right)-1\right]=0, \tag{36}$$

$$\left(1+\frac{A_0+A_1 p_t+\varepsilon_t}{A_1}\right)p_t=\beta\mathbb{E}_t\left[MC_{t+1}\right], \tag{37}$$

where MC_t is the time-t marginal cost, which can be utilized to obtain the steady state values of the states and controls in our model economy. Using the the steady states values of variables, the coefficient matrices in Lemma 1 can be readily computed. Finally, solving the robust LQ problem quantitatively generates the robust decision rules: (33), (34), and (35).

Using the robust decision rules, we can do the same GMM estimation as in our benchmark model to obtain the remaining parameters, $\{\rho_c, \alpha_i, \alpha_y, \alpha\}$.[25] The model comparisons are presented in Table 3. We can see from the table that the RB-PCS model with endogenous sales matches the data reasonably well. The calibrated detection error probability (DEP) is 0.08, which means that there is a 8 percent chance that a likelihood ratio test will improperly select between the approximating and distorted models. The third column of the table compares the robust control model to the standard rational expectations version when we set $\alpha = 0$; the RE-SOA model still underperforms the corresponding RB model along several dimensions. The fourth column is the validation test calibrated without ambiguity.

With endogenous sales, the mechanism of how RB affects the model's dynamics resembles but is not the same as that in the exogenous-sales model. We first examine the robust production rule, (33). A positive demand shock v_t in the sales function will give rise to higher sales; as a consequence, production will be increasing in v_t. Concerns about the specification of v_t reinforces this effect and its response coefficient is therefore larger than in the RE model. Likewise, fears of a higher current cost shock reduces the response coefficient on Γ_t (with a greater absolute value). Lower stocks available for sales further induce fears of lower sales so production decreases further compared with the RE

[25]Note that the steady state values of the model change with parameters. See Online Appendix G for discussion.

case. These two effects jointly contribute to a higher production volatility.

Second, the price responds to both the demand and cost shocks positively in the price rule (34). Profit-maximizing firms increase price in anticipation of shortages due to higher demand or lower production, and this price effect partially offsets the rise and reinforces the fall in sales, respectively. Under RB, firms increase the price more in response to an increase in demand and increase less in response to a cost shock. The latter is because that fears of lower production also induces fears of lower sales. Robust price rules therefore contributes to a less volatile sales during expansions as well as in recessions. In the end, output as well as sales are still more volatile under the RB model, and they generate a output-sales variance ratio greater than 1. The cyclical behavior of inventory investment and inventory-sales ratios are also consistent with the empirical counterparts.

As documented in Khan and Thomas (2007), cyclical fluctuations in inventory investment do not substantially raise the variability of GDP because they lower the variability of final sales as a general equilibrium effect. Here final sales can also be endogenously moderated due to the pricing effect in expansions and recessions. However, the price effect does not reduce the RB model's ability to match the empirical evidence. As reasoned above, output as well as inventory investment are more volatile in our RB model. The main disadvantage of the RE model is that it can only generate a procyclical or weakly countercyclical inventory-to-sales ratio, and the persistence of the inventory-to-sales ratio also falls short of the empirical counterpart, as shown in the sixth, seventh and tenth rows of Table 3. Under RB, the volatility of sales also increases due to the more responsive production, though the volatility of production is still greater than that of sales. The firm relies more on current production to meet sales, instead of holding more inventories. Inventory investment then increases more slowly than the increase of sales in expansions, and dropped less quickly than the decrease of sales in recessions. Inventories can thus move less proportionately than sales, which leads to a lower adjustment rate of inventories relative to sales and a less persistent inventory-to-sales ratio. Also note that in the validation test, the firm has some inventory holding motive, yet the fraction of inventory over sales is almost independent of the sales process, and displays little serial correlation.

6 Comparison With the Stockout Avoidance Motive

6.1 How to Endogenize the Stockout Avoidance Motive from the Perspective of Model Misspecification?

In this section, we compare our RB-PCS theory with the stockout avoidance motive widely examined in the literature on inventories. While our model shows that the uncertainty due to concerns about model misspecification can significantly improve the PCS model's predictions on these important aspects of the data, Wen (2005) shows that incorporating the stockout avoidance motive also helps explain the data along several dimensions.[26] The way Wen (2005) models the stockout avoidance motivation is to assume that firms cannot allow inventory to fall below zero (i.e., $i_t \geq 0$), which means firms want to avoid the situation where they must produce more to meet current demand. Our model with model uncertainty can endogenously generate this outcome: as firms have stronger concerns about model misspecifications, their production decisions become more responsive to both demand and cost shocks and therefore accumulate more inventories; consequently, it is more likely to avoid the negative inventory levels (although the model does not explicitly impose any restrictions on negative inventories).[27] Sarte et al. (2015) show that variation in the discount rates in consumers's utility function plays a key role in explaining the shifts in U.S. business cycles observed after the mid-1980s. The estimated high degree of substitutability between stages of production delivers an SOA motive. Fluctuations in their model alter intertemporal valuations and therefore are more apt to generate substantial investment. In our paper, in contrast to the SOA literature, firms are overly pessimistic about the demand and cost shock processes because they take the worst-possible case into account, and therefore reinforces intertemporal substitution.

6.2 Quantitative Analysis for the Endogenous SOA Mechanism

In this section, we first provide details about how to simulate the fraction of inventory stocks given different degree of model uncertainty. Given different values of the RB parameter α and other set of parameters calibrated in the benchmark model, we simulate the model 1000 times and calculate

[26] The discussion on the stockout avoidance motive goes back to Kahn (1987) and Maccini and Zabel (1996).

[27] Note that allowing for negative inventory stocks need not require that stocks are actually negative (which is of course impossible); rather, one can interpret a negative inventory stock as indicating future orders in excess of current stocks.

the average fraction of negative inventories in all periods. Like in the previous sections, in each simulation path we again simulate according to (2) and (6), and use the numerical decision rules, (20) and (21) to obtain the level of inventory stocks. Particularly, we set the initial value of inventory stocks here at the constant ones. Also, we do not add back the trend of log sales when calculating the level of sales, because otherwise when taking exponential it will dominate the zero mean cost shocks and make the signs of sales positive. The upper-right panel of Figure 5 plots the average share of negative inventory based on simulated series for different values of the RB parameter, α.[28] It is clear from the figure that our calibrated plausible values of α can help generate positive inventories.

As the upper-right panel of Figure 5 shows, the share of negative inventory is very high if $\alpha = 0$ (the RE model). As α increases, the fraction of negative inventory holdings starts to decrease, though initially it declines very gradually. As α approaches the threshold that leads the response of inventory to sales (λ_x) to change from negative to positive (see the bottom-second panel in Figure 1), the share of negative inventory declines rapidly to a level close to zero. If firms have stronger concerns about model misspecification, they adjust their production more aggressively in response to sales changes and therefore endogenously avoid negative inventory stocks. It is also clear from the lower-right panel of Figure 5 that the (mean) stock of inventories also increases with α. When α is sufficiently large, the level of inventories is always positive. This result has the flavor of precautionary behavior; with higher degrees of ambiguity aversion, firms naturally hold buffer stocks of inventory and avoid stockouts, even though having a negative inventory is not explicitly punished.

In the literature such as Kahn (1987), Eichenbaum (1989), Ramey (1991), and Ramey and West (1999), there is another way to explicitly model the stockout avoidance motive in the PS and PCS models. Specifically, they assume that the inventory holding cost takes the following form:

$$H(i_t) = \frac{1}{2}\alpha_i \left(i_{t-1} - \alpha_x x_t\right)^2, \tag{38}$$

where i_{t-1} is the inventory level at $t-1$, x_t is the amount of sales in t, and $\alpha_x > 0$ governs the stockout avoidance motive. This motive induces the firm to hold inventories even without production cost considerations, and (38) embodies both inventory holding and backlog costs. If $\alpha_x = 0$, it can be interpreted as an inventory holding cost function. In contrast, if $\alpha_x \neq 0$, the so-called accelerator motive takes effect and this function reflects stockout (backlog) as well as inventory holding costs,

[28] The average is based on 1000 simulations. We have checked that increasing the number of simulations does not change the results.

and thus captures a revenue-related motive for holding inventories (as α_x increases, stockout costs rise relative to backlog costs). Stockout costs arise when sales exceed the available stock. In other words, the higher the stock of inventories, the less likely is a stockout and the lower are stockout costs.

The upper-left panel of Figure 5 plots the average share of negative inventory based on simulated series for different values of the degree of the SOA motive, α_x. Other parameters are set at the values calibrated to match the data without ambiguity. We can see from this figure that if α_x is sufficiently large ($\alpha_x > 1.60$), the average share of negative inventory in the model with the SOA motive drops to zero. In addition, it is also clear from the lower-left panel of Figure 5 that the (mean) stock of inventories also increases with α_x. When the accelerator motive is operative, inventories will initially rise along with sales when there is a positive demand shock, suggesting that production rises more than sales and is more variable. One can obtain procyclical movements and an output-sales variance ratio above one through this specification. More specifically, if we let production cost be zero, mute the cost shocks and leave the inventory holding cost as the entire cost, the optimal inventory stock in period t will be $i_t = \alpha_x \mathbb{E}_t [x_{t+1}]$. Inventories will covary positively with expected sales, and thus with serially correlated sales themselves. However, in Ramey and West (1999), the median estimate of α_x varies from 0.4 (Ramey 1991) to 1.15 (Eichenbaum 1999) using two-digit manufacturing data from the US, which are obviously below the required critical value to generate nonzero inventory stocks in the model.

The calibrated results using the SOA model are presented in Table 4. The GMM estimation that best fits the data gives a threshold of $\alpha_x = 2.09$. It requires firms to hold a large amount of inventories relative to sales at the beginning of the period, and the SOA model-generated data still falls short of the empirical counterpart. For example, with such a large value of α_x, both output and inventory investment fluctuate too much, with an autocorrelation of the latter much larger than the empirical counterpart.

7 Further Discussions

In this section, we will further explain why the RB model outperforms the RE model. First, we will show that the RE model cannot generate a positive correlation between inventory investment and sales. In addition, we will examine the behavior of production and inventories before and during the

Great Moderation (GM) in the US economy.

7.1 Further Inspection of the RE Model's Prediction for the Inventory-Sale Correlation

Under RE, the inventory policy, (10), can be written as a combination of two AR(2) processes:

$$i_t = \lambda_x \frac{\varepsilon_t}{(1-\rho\cdot\mathbb{L})(1-\lambda_i\cdot\mathbb{L})} + \lambda_\Gamma \frac{w_t}{(1-\rho_\Gamma\cdot\mathbb{L})(1-\lambda_i\cdot\mathbb{L})}, \tag{39}$$

which shows that the stochastic properties of inventories are governed by the following two channels: i) ρ and ρ_Γ, which govern the external propagation mechanisms, and ii) $\lambda_i = \lambda_1$, which is an internal and endogenous propagation mechanism determined by the relative importance of the production smoothing motive to the backlog cleanup motive, α_i/α_y. Also note that λ is determined by R and α_i/α_y, but is independent of the values of ρ and ρ_Γ. In addition, as can be seen in (10) and (39), adding a cost shock does not change the response of the current inventory stock to sales and lagged inventory stocks, so it does not change the covariance of inventory investment and sales. Accordingly, compared with the production-level smoothing model, adopting the production-cost smoothing specification changes the value of var(Δi_t), but has no effect on the inventories-sales correlation; as a result, moderate cost shocks cannot generate the observed higher variance ratio of production to sales. Finally, since $\lambda_x < 0$ in this case $(R - \rho_\Gamma > 0)$, the contemporaneous covariance of inventory and sales will always be negative, independent of the cost shock process. In addition, the production policy, (9), implies that production is negatively correlated with the lagged level of inventories and is positively correlated to the current level of sales. If firms have plenty of inventory stock from the last period, they do not need to increase current period production. Similarly, if demand is high in the current period, firms will produce more.

The fourth column in the lower panel of Table 2 shows a validation test based on the RE model. Specifically, we estimate the three unknown parameters, $\Theta = \{\rho_c, \Psi, \alpha_{iy}\}$, using the two-step GMM procedure, so that the RE model-generated moments matches the ten moments (from (i) to (x)) defined in Proposition 2 of Section 3.4 as closely as possible. We choose these moments as our targets as they capture important features of the joint dynamics of production, inventory, and sales. One may note that if Facts (iii) and (v) are true given Fact (ii) (i.e., inventory investment is procyclical, while the inventory-to-sales ratio is countercyclical, inventory investment is much less volatile than sales), it implies that $\log(i)$ is less volatile than $\log(x)$. This is not equivalent to saying that Fact

(ii) is true, for $\log(i)$ can move faster than $\log(x_t)$ with inventory investment is still much less volatile. The bottom line is, the RB model matches the very low volatility of inventory investment relative to sales. On the other hand, the RE model fails to match the data well on the overall performance. We report the results in Table 2. The upper-third column gives estimated parameters. In the lower panel, the second column reports the empirical evidence during the $1967-2018$ period, and the third column reports the RE model's predictions. It is clear from the table that the RE model can merely generate an output-sales variance ratio around 0.98, while inventory investment is countercyclical and the inventory-to-sales ratio is procyclical. All these results are far apart with the empirical evidence. We then examine if the RB model can outperform the original RE model. A similar graph like Figure 2 shows that the RE model's performance improves when the degree of RB strengthens (i.e., α rises). Like in the previous figure, as the degree of RB rises, the firm's optimal production and inventory decisions generate a higher output-sales variance ratio, the comovement of inventory investment and the inventory-to-sales ratio switches signs and gets closer to the empirical counterpart, and the inventory-to-sales ratio is also persistent. The values of remaining moments are also reasonably close.

7.2 The Great Moderation

Beside the above stochastic properties, the literature on inventory dynamics also finds that the relationship between production, inventories and sales has changed somehow during the Great Moderation (GM). For example, Sarte et al. (2015) show that variations in the discount rates in consumers's utility function plays a key role in explaining the shifts in U.S. business cycles observed after the mid-1980s. The estimated high degree of substitutability between stages of production, similar to Bils and Kahn (2000), delivers an SOA motive. Fluctuations in the model alter intertemporal valuations and therefore are more apt to generate substantial investment. In our paper, in contrast to the SOA literature, firms are overly pessimistic about the cost shock process and therefore reinforces intertemporal substitution too. Furthermore, the shift of comovement properties and the relative volatilities of the data prior to and after 1984 can be explained by changes in the degree of RB and the amount of model uncertainty faced by firms. With more and more information, agents may be better able to draw inference about the underlying model-generating-process from a narrower set of surrounding models, hence reducing the degree of model uncertainty.

To investigate into this channel, we estimate our endogenous-sales model over the two separate subsamples split in 1984, targeting the data moments in each subsamples respectively. The whole data sample covers 1967-2007. The four chosen moments are those mentioned in Sarte et.al (2015): (i) the relative standard deviation of output to sales ($\sigma_{yx} \equiv \text{std}\,(y_t)\,/\,\text{std}\,(x_t)$), (ii) the relative standard deviation of inventories to output ($\sigma_{iy} \equiv \text{std}(i_t)/\,\text{std}(y_t)$), (iii) the relative standard deviation of inventory-to-sales ratio to output ($\sigma_{i/x,y} \equiv \text{std}(i_t/x_t)/\,\text{std}(y_t)$), and (iv) the correlation between inventory-to-sales ratio and output ($\rho_{i/x,y} \equiv \text{cov}(i_t/x_t, y_t)/\left(\sqrt{\text{var}(i_t/x_t)}\sqrt{\text{var}(y_t)}\right)$). Correlations and standard deviations are calculated using monthly data and HP filtering as in Kryvstov and Midrigan (2013). Table 5 reports the results. The inventory-sales ratio became less countercyclical in the second sub-period, the Great Moderation period, consistent with the results of Kryvtsov and Midrigan (2013). In addition, we also see from the table that inventory-output variance ratio is higher during the GM period. The calibrated DEPs are 0.02 for the period before 1984 (the first sub-period) and 0.22 for the period after 1984 (the second sub-period). These results indicate a decreased amount of model uncertainty during the GM period.

Despite the fact that output volatility decreased during the GM, sales volatility decreased even further, contributing to a slightly increased output-sales ratio after the shift. This result also suggests that fundamental uncertainty (risk) is decreased. Note that the fundamental uncertainty may lead to Knigtian uncertainty when agents are ambiguity averse, as shown in (29)-(30). In contrast, the deep parameter α has increased from 208.6 before the GM period to 379.0 during this period, indicating that the degree of robustness has increased. To sum up, the overall amount of model uncertainty decreased, reducing the sensitivity of production to the demand and the cost shocks, and resulting in more inventory holdings as a buffer stock. This reduction leads to the increased relative volatility of inventories to output and the less countercyclical inventory-to-sales ratio during the GM period.

8 Conclusions

In this paper, we construct a version of the PCS model with ambiguity aversion to study the joint dynamics of production, inventories, and sales. Our model can explain ten facts that previous studies find difficult to account for simultaneously. We also show that the stock-out avoidance motive (Kahn 1987) emerges endogenously in our model due to the ambiguity-aversion-induced precautionary behavior. Our analysis shows that allowing for a moderate degree of ambiguity aversion enables a standard linear-quadratic framework based on Blinder (1986), Eichenbaum (1989), and Ramey and West (1999) to account for the production and inventory dynamics reasonably well.

References

[1] Bhandari, A., Borovička, J. and Ho, P. (2022), "Survey Data and Subjective Beliefs in Business Cycle Models," Working Paper.

[2] Bidder, Rhys and Matthew E. Smith (2012), "Robust Animal Spirits," *Journal of Monetary Economics* 59(8), 738-750.

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

[4] Blanchard, Olivier. J. (1983), "The Production and Inventory Behavior of the American Automobile Industry," *Journal of Political Economy* 91: 585-614.

[5] Blinder, Alan S.(1986), "Can the Production Smoothing Model of Inventory Behavior be Saved," *Quarterly Journal of Economics* 101, 431-453.

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

[7] Caballero, R. and Krishnamurthy, A. (2008), "Collective Risk Management in a Flight to Quality Episode," *Journal of Finance* 63(5), 2195-2230.

[8] Coen-Pirani, Daniele (2004), "Markups, Aggregation, and Inventory Adjustment," *American Economic Review* 94(5), 1328-1353.

[9] Djeutem, Edouard and Kasa, Kenneth (2013), "Robustness and Exchange Rate Volatility," *Journal of International Economics* 91(1), 27-39.

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

[11] Gilboa, Itzhak and Schmeidler, David (1989), “Maxmin Expected Utility with Non-unique Prior,” *Journal of Mathematical Economics* 18(2), 141-153.

[12] Gï¿œrtz, Christoph, Christopher Gunnb, and Thomas A. Lubik (2022), “Is There News in Inventories?” *Journal of Monetary Economics* 126, 87-104.

[13] Hansen, Lars Peter and Thomas J. Sargent (2007), *Robustness*, Princeton University Press.

[14] Hansen, Lars Peter and Thomas J. Sargent (2010), “Wanting Robustness in Macroeconomics,” Chapter 20 in *Handbook of Monetary Economics* vol. 3, 1097-1157.

[15] Hansen, Lars Peter, Thomas J. Sargent, and Thomas D. Tallarini, Jr. (1999), “Robust Permanent Income and Pricing,” *Review of Economic Studies* 66(4), 873-907.

[16] Holt, Charles C., Modigliani, Franco, John F. Muth, and Herbert A. Simon (1960), *Planning, Production, Inventories and Work Force*, Prentice-Hall.

[17] Ilut, Cosmin and Martin Schneider (2014), “Ambiguous Business Cycles,” *American Economic Review* 104, 2368-2399.

[18] Ilut, Cosmin, Rosen Valchev, and Nicolas Vincent (2020), “Paralyzed by Fear: Rigid and Discrete Pricing under Demand Uncertainty,” *Econometrica* 88(5), 1899-1938.

[19] Kahn, James A. (1987), “Inventories and the Volatility of Production,” *American Economic Review* 77(4), 667-679.

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

[21] Kasa, Kenneth, and Xiaowen Lei (2018), “Risk, Uncertainty, and the Dynamics of Inequality,” *Journal of Monetary Economics* 94, 60-78.

[22] Khan, Aubhik and Thomas, Julia K. (2007a), “Inventories and the Business Cycle: An Equilibrium Analysis of (S, s) Policies,” *American Economic Review* 97(4), 1165-1188.

[23] Khan, Aubhik and Thomas, Julia K. (2007b), "Explaining Inventories: A Business Cycle Assessment of the Stockout Avoidance and (S, s) Motives," *Macroeconomic Dynamics* 11(5), 638 - 664.

[24] Knight, Frank H. (1921), *Risk, Uncertainty, and Profit*, Pantianos Classics.

[25] Kryvtsov, O., and Midrigan, V. (2013), "Inventories, Markups, and Real Rigidities in Menu Cost Models," *Review of Economic Studies* 80(1), 249-276.

[26] Lai, Kon S. (1991), "Aggregation and Testing of the Production Smoothing Hypothesis," *International Economic Review* 32(2), 391-403.

[27] Luo, Yulei and Eric R. Young (2010), "Risk-sensitive Consumption and Savings under Rational Inattention," *American Economic Journal: Macroeconomics* 2(4), 281-325.

[28] Luo, Yulei, Jun Nie, and Eric R. Young (2012), "Robustness, Information-processing Constraints, and the Current Accounts in Small Open Economies," *Journal of International Economics* 88(1), 104-120.

[29] Maccini, Louis J. and Zabel, Edward (1996), "Serial Correlation in Demand, Backlogging and Production Volatility," *International Economic Review* 37(2), 423-452.

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

[31] Ramey, Valerie and Kenneth West (1999), "Inventories." In *Handbook of Macroeconomics*, edited by John B. Taylor and Michael Woodford, pp. 863-923. New York: North Holland, Elsevier.

[32] Sarte, Pierre-Daniel, Felipe Schwartzman, and Thomas A. Lubik (2015), "What Inventory Behavior Tells Us about How Business Cycles Have Changed," *Journal of Monetary Economics* 76, 264-283.

[33] Wen, Yi (2005), "Understanding the Inventory Cycle," *Journal of Monetary Economics* 52(8), 1533-1555.

[34] Wen, Yi (2011), "Input and Output Inventory Dynamics," *American Economic Journal: Macroeconomics* 3(4), 181-212.

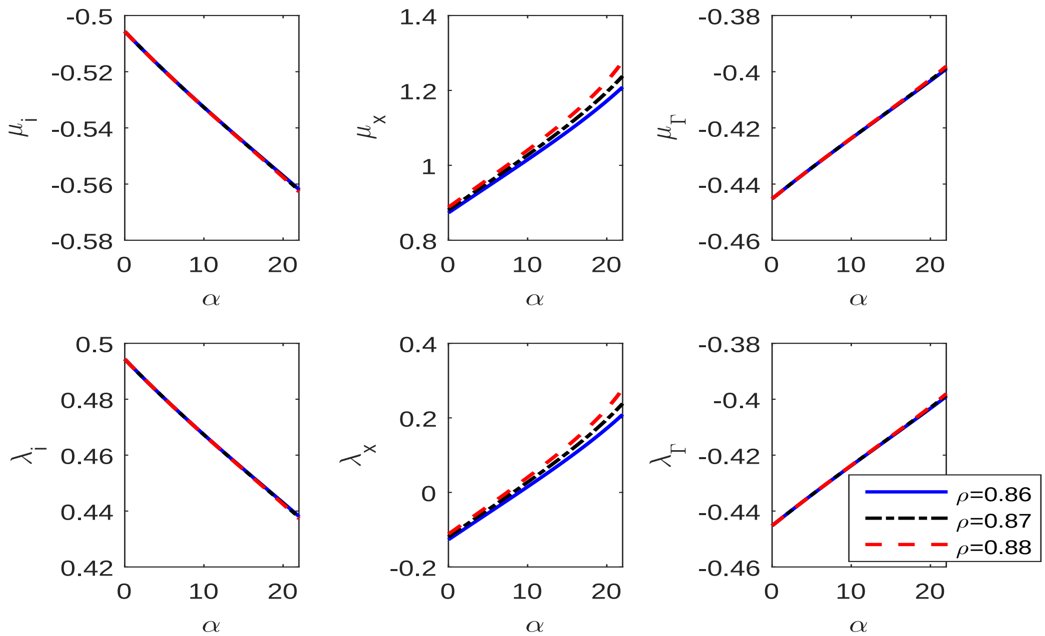

Figure 1: Effects of RB on Production and Inventory Decisions

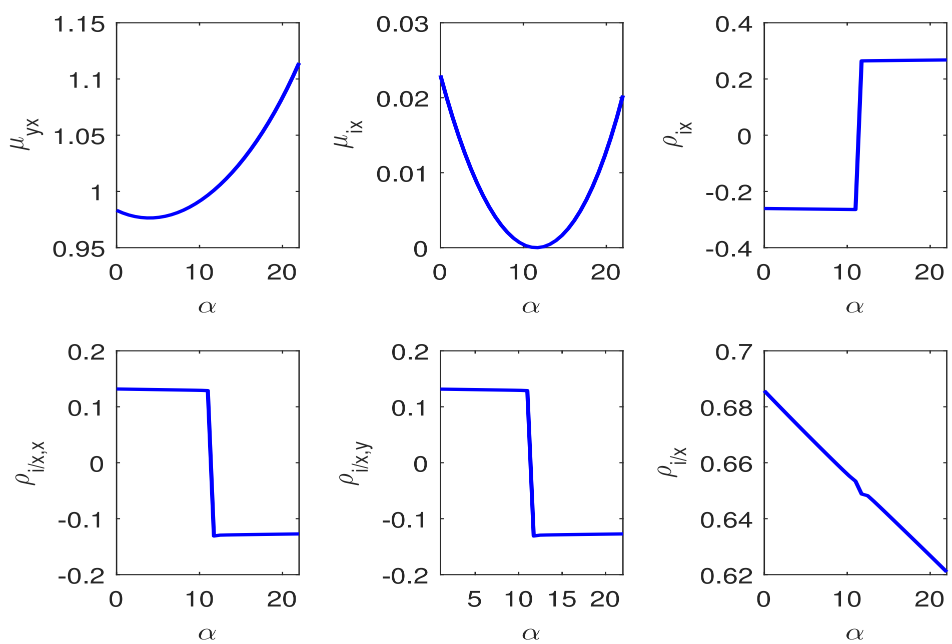

Figure 2: Changes of Moments When Varying the Degree of Model Uncertainty

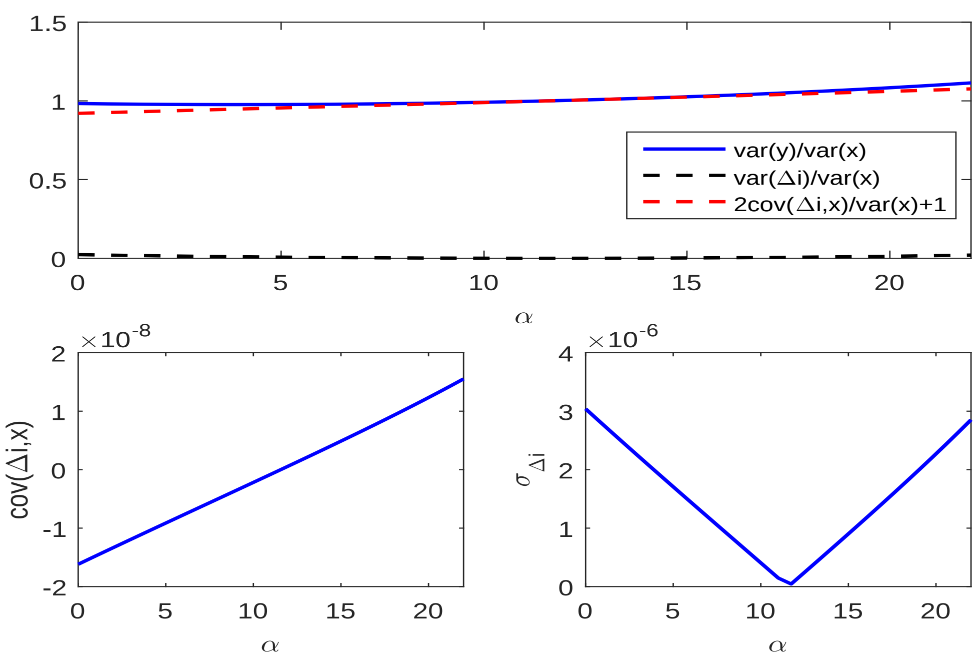

Figure 3: Decomposition of μ_{yx} and μ_{ix}

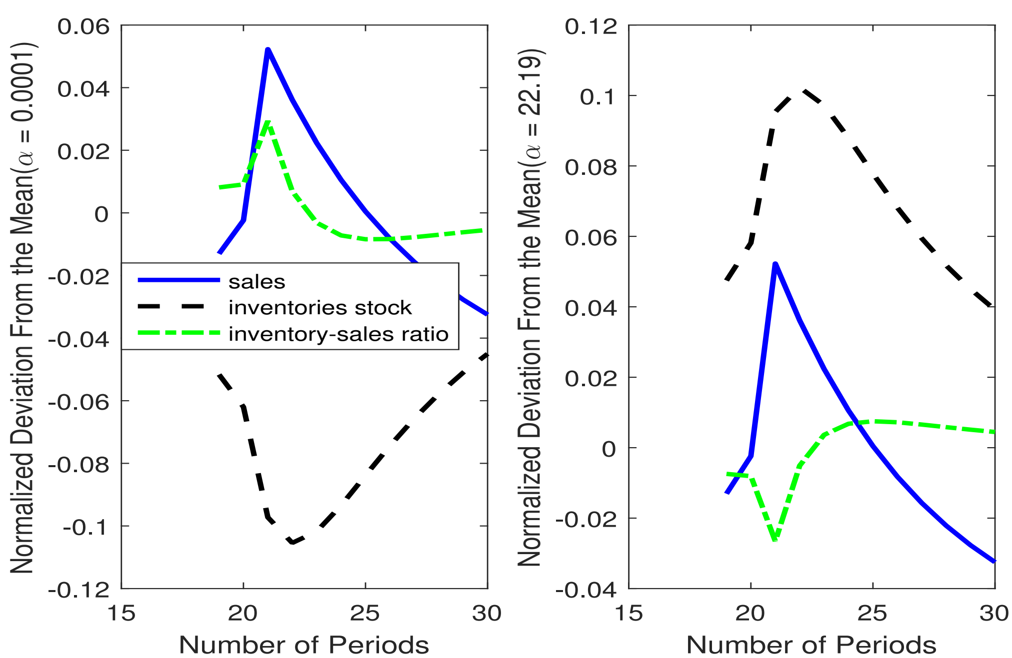

Figure 4: Normalized Deviation from the Mean of Sales, Inventories, and Inventory-sales Ratio After A Sales Shock

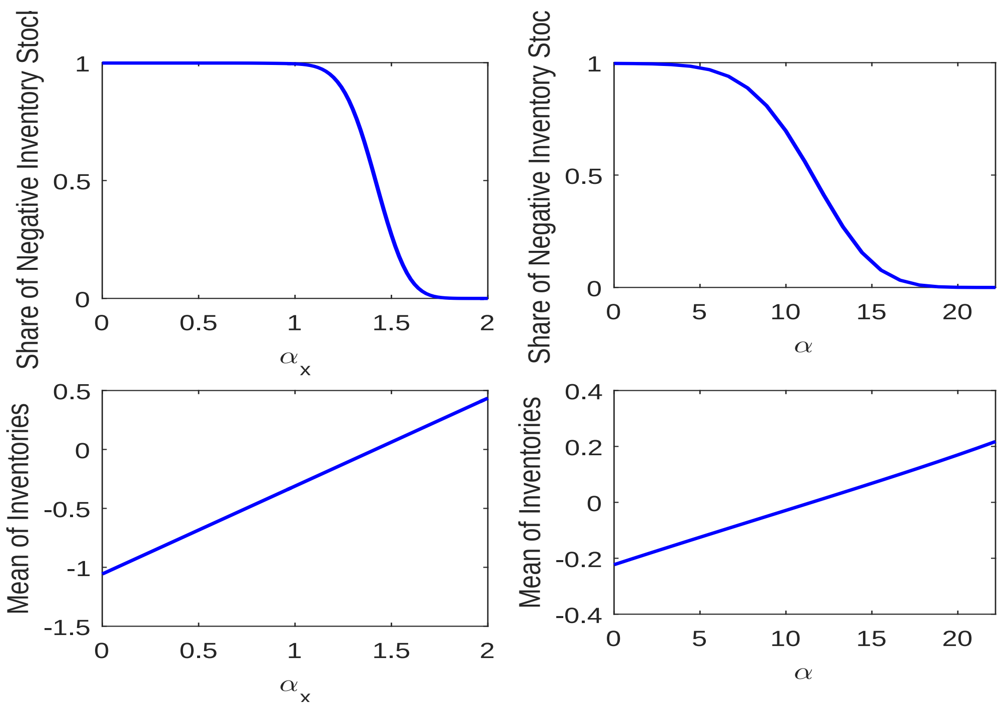

Figure 5: The Simulated Average Fraction of Negative Inventory Stocks and Simulated Mean of Inventories a)When α_x Rises b)When α Rises

Table 1: The Joint dynamics of Inventories, Production and Sales

	Whole Sample	Before Financial Crisis
	1967 – 2018	1967 – 2007
Manufacturing Sector		
Relative Volatility of Production to Sales (μ_{yx})	1.10 (.0151)	1.10 (.0192)
Relative Volatility of Inventories to Sales (μ_{ix})	0.029 (.0047)	0.033 (.0059)
Correlation between Ouput and Sales (ρ_{yx})	0.99 (.0018)	0.99 (.0021)
Correlation between Inventory Investment and Sales (ρ_{ix})	0.30 (.0543)	0.29 (.0551)
Correlation between Inventory-to-Sales Ratio and Sales ($\rho_{i/x,x}$)	-0.47 (.0621)	-0.49 (.0774)
Correlation between Inventory Investment and Output (ρ_{iy})	0.42 (.0478)	0.42 (.0503)
Correlation between Inventory-to-Sales Ratio and Output ($\rho_{i/x,y}$)	-0.46 (.0634)	-0.48 (.0798)
Autocorrelation of Output (ρ_y)	0.88 (.0190)	0.86 (.0226)
Autocorrelation of Inventory Investment (ρ_i)	0.26 (.0580)	0.25 (.0668)
Autocorrelation of Inventory-Sales Ratio ($\rho_{i/x}$)	0.98 (.0027)	0.98 (.0033)

Table 2: Benchmark Model Comparison($p = 0.22$)

Parameter		RB Model	RE Model	RE Model*
ρ_c		0.1183	0.1183	0.0124
Ψ		0.0118	0.0118	0.0162
α_{iy}		0.5520	0.5520	0.6863
α		22.1925	0	0
Moment	Data	RB Model	RE Model	RE Model*
μ_{yx}	1.10	1.11	0.98	0.98
μ_{ix}	0.029	0.021	0.023	0.015
ρ_{yx}	0.99	0.98	0.97	0.98
ρ_{ix}	0.30	0.27	-0.26	-0.27
$\rho_{i/x,x}$	-0.47	-0.13	0.13	0.12
ρ_{iy}	0.42	0.34	-0.18	-0.20
$\rho_{i/x,y}$	-0.46	-0.30	0.07	-0.05
ρ_y	0.88	0.75	0.83	0.82
ρ_i	0.26	0.62	0.69	0.63
$\rho_{i/x}$	0.98	0.62	0.69	0.63

Table 3: Endogenous-Sales Model Comparision($p = 0.08$)

Parameters		RB	RE	RE*
ϕ		0.1	0.1	0.1
ρ_c		0.9686	0.9686	0.9871
Ψ		0.0128	0.0128	0.0060
α_i		0.1704	0.1704	0.2397
α_y		0.2466	0.2466	0.1251
α		111.6	0	0
Moment	**Data**	**RB**	**RE**	**RE***
μ_{yx}	1.10	1.09	1.04	1.10
μ_{ix}	0.025	0.027	0.010	0.038
ρ_{yx}	0.99	0.99	1.00	0.98
ρ_{ix}	0.23	0.20	0.16	0.15
ρ_{iy}	0.36	0.35	0.25	0.33
$\rho_{i/x,x}$	-0.87	-0.83	0.70	-0.03
$\rho_{i/x,y}$	-0.85	-0.85	0.69	-0.04
ρ_y	0.88	0.88	0.92	0.91
ρ_i	0.08	0.28	0.29	0.25
$\rho_{i/x}$	0.85	0.83	0.97	0.08

Table 4: Performance of the SOA model

Parameter	Estimated Value
ρ_c	0.8732
Ψ	0.0161
α_{iy}	0.2781
α_s	2.0940

Moment	Data	SOA Model
μ_{yx}	1.10	1.32
μ_{ix}	0.029	0.116
ρ_{ix}	0.30	0.25
ρ_{iy}	0.42	0.38
$\rho_{i/x,x}$	-0.47	-0.14
$\rho_{i/x,y}$	-0.46	-0.42
ρ_{yx}	0.99	0.94
ρ_y	0.88	0.73
ρ_i	0.26	0.77
$\rho_{i/x}$	0.98	0.77

Table 5: Performance of the Model Prior to and During Great Moderation: Endogenous-Sales Model

Moment	Data (1967-1983)	Model $p = 0.02$ ($\alpha = 208.6$)	Data (1984-2007)	Model $p = 0.23$ ($\alpha = 379.0$)
σ_{yx}	1.04	1.02	1.06	1.03
σ_{iy}	0.31	0.37	0.75	0.66
$\sigma_{i/x,y}$	1.07	1.06	1.17	1.23
$\rho_{i/x,y}$	-0.94	-0.99	-0.75	-0.55