(S,s) insights into the role of inventories in business cycles *and* high frequency fluctuations

Aubhik Khan
Ohio State University

Julia K. Thomas
Ohio State University

February 2016

ABSTRACT

We develop an equilibrium model to explain salient empirical regularities regarding the movements in aggregate production, sales and inventory investment over the business cycle while simultaneously explaining a distinct set of facts about these series at higher frequencies. We begin with a two-sector (S,s) model where fixed costs of ordering inputs lead firms to reduce their order frequencies by accumulating inventories. Adding idiosyncratic shocks to firms' total factor productivities, we discipline the distributions of order costs and firm-level productivity using average aggregate inventory data and firm-level data on sales and output growth volatility.

The model yields a distribution of firms over productivity and inventories, as only some firms adjust their stocks in any period. When a firm does not adjust, its existing inventory restricts its inputs available for production. This prompts firms to raise their average stocks in times of persistently high aggregate demand. The risks associated with potentially high future order costs and sharp changes in relative productivity also influence firms' sales and inventory decisions. These idiosyncratic risks can alter firms' decisions sufficiently to yield important changes in aggregate dynamics, particularly in the high frequencies.

Our model generates key patterns from U.S. data at business cycle frequencies. Inventory investment is procyclical and positively correlated with final sales, GDP varies more than sales, and the inventory-to-sales ratio is countercyclical. These successes are robust to changes in micro- parameters affecting the distributions of order costs and relative productivity. However, because the average distribution of firms over stocks is shaped by the persistence and variability of idiosyncratic order costs and productivities, these parameters influence the extent and speed of firms' responses to aggregate shocks, and thus the model's ability to reproduce high-frequency aspects of aggregate data. When firms are more certain about their order costs and shifts in their relative productivities are transitory, they adjust their inventories faster following an aggregate shock. In such cases, the model performs well not only in business cycle respects, but also in distinct high frequency respects documented here. At high frequencies, the relative volatility of inventory investment rises sharply, and sales and inventory investment are negatively correlated, while both series maintain positive correlations with GDP. Despite these distinctions, our model results suggest that aggregate fluctuations are surprisingly unaffected by inventories even at high frequencies.

KEYWORDS: Inventories, (S,s) policies, equilibrium business cycles, high frequency fluctuations

1 Introduction

Do changes in inventories affect movements in GDP and the components of final sales in important and predictable ways? Guided by reasoning from Metzler (1941), Lovell (1961), Blinder (1981) and beyond, policymakers and economic analysts informing the business press are convinced that they do. When the ratio of business inventories relative to sales rises to unusually high levels, as for instance at the end of 2015, many warn of an impending 'inventory correction' recession driven by declines in inventory investment. However, the arguments supporting such predictions are invariably hazy and incomplete.

The advance estimate of annualized real GDP growth for 2015Q4 was just 0.7 percent, with nonfarm inventory investment subtracting 0.38 percentage points from that growth. To predict whether the last two quarters' declines in inventory accumulation promise continued weak GDP growth over the coming months and quarters, we might consult a quantitative dynamic stochastic general equilibrium business cycle model that explains the salient empirical regularities. The problem with this strategy is that no such model exists. To date, no single endogenous inventory model has proved successful with respect to the cyclical facts regarding inventory investment while simultaneously generating empirically valid high-frequency behavior in the series.[1] Because inventory investment exhibits sizeable fluctuations in both frequency ranges, we see this as an important omission and aim to correct it here. The challenge in doing so lies in the fact that the relationships between inventory investment, GDP, and final sales at high frequencies differ markedly from those over the business cycle.

Table 1 presents a set of basic inventory facts for the business cycle. Here, we report relative volatilities and contemporaneous correlations of real quarterly domestic business GDP, final sales, net changes in private nonfarm inventories, and the inventory-to-sales ratio for the U.S. over 1954:1 - 2014:3. The first two series are detrended in the conventional way, first logged, then HP-filtered using weight 1600. As inventory investment passes through zero, we detrend it as a share of GDP; that is, we HP-filter the ratio NII/GDP. Since inventory-sales is itself a ratio, we directly apply the HP-filter to this series.[2]

[1] Sarte et al. (2015) develop a model with a multi-stage production technology to explore the extent to which altered inventory dynamics before versus after the mid-1980s may give clues to what caused sharp changes in the volatility and correlation patterns of GDP, labor productivity and hours worked beginning in the Great Moderation. Their estimated model is qualitatively consistent with the co-movement of U.S. inventory investment and sales in business cycle and high frequencies. However, as with the original Kydland and Prescott (1982) time-to-build technology it nests as a special case, the stocks therein are required factors of production.

[2] The overall level of volatility is reflected in the GDP column, where we report the percent standard deviation of filtered GDP in parentheses. Moments in Table 1 are similar when we replace the Hodrick-Prescott filter with a bandpass filter and isolate frequencies of 6 to 32 quarters per cycle.

TABLE 1. Business cycle volatilities and contemporaneous correlations

	GDP	Final Sales	NII	I-S Ratio
σ_x/σ_{GDP}: RELATIVE STD. DEV.	(2.122)	0.874	0.283	0.660
$corr(x, GDP)$: GDP CORRELATION	1.000	0.945	0.639	-0.442
$corr(x, NII)$: NII CORRELATION	0.639	0.397	1.000	-0.201

Table 1 shows that, over the business cycle: (a) inventory investment is procyclical, (b) inventory investment co-moves with final sales, (c) final sales is less volatile than GDP, and (d) the stock of inventories adjusts gradually, so the inventory-sales ratio is countercyclical despite the co-movement of sales and inventory investment. Note also that inventory investment is far more cyclically volatile than the 0.283 relative standard deviation might initially suggest. Recall that this series is detrended as a share of GDP, and consider that inventory investment averages only 0.56 percent of GDP over 1954:1 - 2014:3. By contrast, business fixed investment averages 17.6 percent of GDP over these dates; when detrended as a share of GDP, its relative volatility is 0.285 (versus 2.21 when log HP-filtered).

The basic observations (a) - (d) drawn from Table 1 are consistent with earlier findings in Khan and Thomas (2007) and the more recent findings of Kryvtsov and Midrigan (2013), which considers both business cycles overall and in the case where the data are conditioned to isolate patterns arising under monetary shocks. Khan and Thomas (2007) show that these relationships can be reproduced in an equilibrium business cycle model wherein inventories derive from an (S,s) motive, firms are perfectly competitive and fluctuations are driven by shocks to aggregate productivity. Kryvtsov and Midrigan (2013) instead consider monetary shocks as a source of business cycles in a New Keynesian environment. They show that a stockout avoidance motive for inventories succeeds with respect to the inventory facts in their model, but only so long as firms have decreasing returns to scale and their markups of price over marginal cost are sufficiently countercyclical. One deficiency of their stockout avoidance model is its prediction of a strong negative correlation between sales and inventory investment. Considering an (S,s) inventory motive instead in the same setting, they find this correlation moves nearer the data, while the basic findings elsewhere are unchanged. Despite very different environments, Kryvtsov and Midrigan (2013) and Khan and Thomas (2007) arrive at a common negative prediction regarding the quantitative importance of inventory adjustment for business cycle fluctuations in GDP. Both studies compare business cycle moments arising in the model with inventories to those from a counterpart model where the inventory motive is stripped away, and both find negligible differences.

The basic inventory facts change in two important ways when we shift attention to the high frequencies. We present the same series again in Table 2, but now band-pass filter the data to isolate frequencies of 2 to 6 quarters per cycle. Two key facts jump out in this table. First, the relative volatility of inventory investment is dramatically greater in the high frequencies than over the business cycle. Second, while inventory investment remains procyclical, it no longer co-moves with final sales; the correlation between sales and inventory investment is roughly -0.2. These two facts together paint a very different picture for the nature of fluctuations in the high frequency band relative to the business cycle, suggesting that inventory investment is the dominant driver of GDP movements at high frequencies. This in turn suggests that Kryvtsov and Midrigan (2013) and Khan and Thomas (2007) may have focused attention in the wrong frequency band when considering the importance of inventory investment for aggregate fluctuations.[3]

TABLE 2. High frequency volatilities and contemporaneous correlations

	GDP	Final Sales	NII	I-S Ratio
σ_x/σ_{GDP}: RELATIVE STD. DEV.	(0.547)	0.830	0.614	0.912
$corr(x, GDP)$: GDP CORRELATION	1.000	0.642	0.538	-0.331
$corr(x, NII)$: NII CORRELATION	0.538	-0.201	1.000	0.491

In this paper, we develop an equilibrium model capable of reproducing both sets of empirical regularities regarding the movements in aggregate production, sales and inventory investment. Despite its reliance on only a single (S,s) inventory motive, the model generates negative co-movement between sales and inventory investment at high frequencies often associated with a buffer stock motive for inventories, while it is simultaneously consistent with the positive GDP correlations from Table 2, and with the business cycle data wherein procyclical inventory investment rises and falls alongside sales. Beyond these aggregate successes, our model generates empirically valid volatility in sales and output growth at the level of the firm by design.

Production occurs across two sets of firms in our model. The first set combines intermediate goods with labor to produce a final good used for consumption and investment in capital. The second set supplies the first with intermediate goods, producing these using capital and labor which are, in turn, supplied by households. Inventories arise in our economy from the fact that final good firms face nonconvex costs associated with shipments of intermediate goods. These costs lead firms to place infrequent orders with their suppliers and maintain stocks on hand.

Firms follow generalized (S,s) policies with respect to their inventories; they place orders when their existing stock falls below a lower threshold, and their order level is determined by an upper

[3] Restricting attention to frequencies ranging 2 to 3 quarters per cycle, Wen (2005) also documents extreme volatility in inventory investment and a negative correlation with final sales. When we shift to this highest frequency band, no number in Table 2 changes sign; the same is true if we also truncate our sample to end at 2004:2 as did his. We speculate that his conflicting finding of countercyclical inventory investment and a relative volatility of final sales slightly exceeding 1 may stem from his inclusion of data over 1947:1 - 1953:4, a highly volatile episode omited from our sample. See figure 2 in Wen (2005).

threshold. So long as a firm's existing stock remains between these thresholds, it avoids paying order costs; hence its beginning-of-period inventory constrains the quantity of intermediate goods available for use in current production. As such, when a firm does order intermediate goods, it orders in sufficient quantity to carry it through several periods' production. If the firm places an order in a time when aggregate productivity is unusually high, expecting the boom to persist, it increases the quantity it orders beyond what is necessary to cover the rise in its current production (sales), because it recognizes that continued high demand for its output in coming periods will exhaust its inventory more rapidly than usual. It is in large part for this reason that our model succeeds in generating the salient empirical regularities characterizing the movements of GDP, final sales and inventory investment over the business cycle. In particular, the model explains the procyclicality of inventory investment, its positive correlation with final sales, and thus the higher variability of GDP relative to final sales. Moreover, depending upon the specification of micro-level shocks affecting firms' inventory and production decisions, the model reproduces between 50 and 60 percent of the observed cyclical variability of net inventory investment, and roughly 45 percent of the observed high frequency variation in the series.

Our environment builds on the (S,s) model in Khan and Thomas (2007).[4] We introduce idiosyncratic shocks to firm-level productivity, we consider alternative specifications for the distribution of order costs, and we include a second aggregate shock affecting productivity in the final goods sector, alongside that in the intermediate goods sector, to more closely reproduce observed cyclical dynamics of the relative price of inventories. These generalizations allow us to address the high frequency fluctuations in aggregate production, sales and inventory investment, as well as firm-level observations on the volatility of output and sales, while maintaining the model's aggregate successes in the business cycle frequencies.

All else equal, high variability in order costs implies considerable variance in the length of an inventory episode, the number of periods between one order of inputs and the next. Over such episodes, firms' inventories decline, while their sales also tend to fall. As such, high variance in the length of inventory episodes tends to yield high variability in sales and output across firms. Absent idiosyncratic productivity shock, if order costs are distributed uniformly between 0 and an upper bound dictated by the average aggregate inventory-to-sales ratio in postwar U.S. data, the model predicts excessive variation in sales and output growth relative to the firm-level

[4] See the discussion there for a survey of related literature. More recently, Wen (2011) studies business cycle fluctuations in an equilibrium model with both input and output inventories arising from a stockout avoidance motive. As noted above, Kryvtsov and Midrigan (2013) embed both (S,s) and stockout avoidance motives in a New Keynesian model driven by monetary policy shocks, also confining attention to business cycle frequencies. Wen (2005) instead considers production smoothing and stockout avoidance models targeted at high frequency moments isolated therein.

data. We find we can correct this by reducing the uncertainty in order costs; in particular, the model delivers empirically plausible micro-level variability when firms' order costs come from a right-skewed distribution with substantial probability mass associated with relatively high costs.

Our introduction of shocks to firm-level total factor productivity also influences the way firms manage their inventories, and thus the distribution of firms over stocks. This, in turn, has non-trivial implications for the high-frequency dynamics of our model. Absent idiosyncratic shocks, firms' inventory policies are effectively one-sided (S,s) rules. By this we mean that, under ordinary aggregate conditions, firms' stocks never exceed the threshold level they select in times of active inventory adjustment. However, when we introduce idiosyncratic productivity shocks, each firm's shock history since the time of its last order determines its history of sales over the episode, and thus its current inventory holding. Some firms find themselves with excess inventories but, given nonconvex shipment costs, choose not to adjust their stocks. In this setting, a positive aggregate shock implies an increase in the expected real interest rate, and a fall in the relative price of intermediate goods. Both reduce the value of inventories and increase the returns to current production. Firms holding excess inventories find it worthwhile to increase production, and are able to do so without ordering. In the aggregate, the result is a more rapid response in the production of final goods, and a somewhat muted response in net inventory investment. By contrast, when we reduce uncertainty regarding future order costs, those firms that respond to an aggregate shock by ordering intermediate goods order far more than they otherwise would and retain a large fraction of their raised stock, because they place high probability weight on a lengthy delay before their next order. This tends to produce very sharp adjustments in the aggregate stock of inventories that weaken the co-movement between sales and inventory investment, particularly in the high frequencies. When a second aggregate shock is included to eliminate excess countercyclicality in the relative price of inventories, the co-movement weakens somewhat further, delivering the weak negative correlation in the high frequency band.

Considering business cycle frequencies, our findings on the role of procyclical inventory investment for the fluctuations in GDP are consistent with earlier findings in Khan and Thomas (2007) and Kryvtsov and Midrigan (2012). We show that an aggregate shock to total factor productivity in the intermediate good sector increases total production and final sales, as well as both the number of firms placing orders and the average size of their orders. Because the gradual accumulation of capital slows the rise in intermediate goods supply, however, there is a trade-off between firms' increased inventory accumulation and their increased production of final goods. Each activity comes at some expense to the other, so the rise in final sales is dampened by, and dampens, the rise in inventory investment. In this sense, the frictions that induce inventories also offset much of the volatility in GDP that might otherwise arise from procyclical fluctuations in inventory investment.

More surprisingly, despite much higher relative volatility of inventory investment in the high frequency band, our model predicts that the dynamics of inventory investment do not drive increased GDP volatility even in the high frequencies. In fact, the standard deviation of GDP falls marginally when the frictions causing inventories are introduced. Here again, the volatility arising directly from inventory investment is partly offset by accompanying reductions in the volatility of final sales, given a fundamental trade-off between storing the goods held in inventory and using them in production. Moreover, whereas the positive correlation between sales and inventory investment in the business cycle frequencies drives slightly higher GDP volatility in the presence of inventories, the mildly negative correlation between these series in the high frequencies delivers the reverse implication.

The remainder of the paper is organized as follows. Section 2 presents the model, and section 3 highlights aspects of its decision rules we find useful in arriving at a numerical algorithm to solve for competitive equilibrium. Sections 4 and 5 discuss our solution method and model calibration. Section 6 presents results, first exploring steady state, then considering business cycle fluctuations and aggregate fluctuations at high frequencies. Section 7 concludes.

2 Model

Production in the economy occurs across two sets of firms. Final goods, which may be consumed or used for investment in capital, are produced by a unit measure of firms that operate a decreasing returns to scale technology using labor and intermediate goods. Intermediate goods are produced by a representative firm that operates a constant returns to scale technology using capital and labor. Final good firms face fixed costs of receiving a delivery of the intermediate input, and, in response, choose to hold inventories of the good. All firms are perfectly competitive, in both factor and output markets. Our baseline model considers a single aggregate shock to productivity in the intermediate goods sector. However, to achieve consistency with the high frequency behavior of inventories below, we will extend the model to include a second independent aggregate shock affecting productivity in the final goods sector. We allow for both shocks as we lay out the notation here.

The representative intermediate good firm is subject to a shock to its total factor productivity, z_I, which follows a Markov Chain. Each period, $z_I \in \{z_1, \ldots, z_{N_I}\}$ with $\Pr\{z_I' = z_c \mid z_I = z_r\} = \pi^I_{rc} \geq 0$, $c = 1, \ldots, N_I$, and $\sum_{c=1}^{N_I} \pi^I_{rc} = 1$ for $r = 1, \ldots, N_I$. Letting $k \in \mathbf{K}$ denote its capital stock at the start of a period, the firm's state is (k, z_I). Final good firms are identified at the start of a period in part by what stock of intermediate goods they have on hand, $s \in \mathbf{S} \subset \mathbf{R}_+$. They are also subject to idiosyncratic shocks to their total factor productivity, $\varepsilon \in \mathbf{E}$, which follow a Markov Chain. Each period, $\varepsilon \in \{\varepsilon_1, \ldots, \varepsilon_{N_\varepsilon}\}$ with $\Pr\{\varepsilon' = \varepsilon_j \mid \varepsilon = \varepsilon_i\} = \pi^\varepsilon_{ij} \geq 0$,

$j = 1, \ldots, N_\varepsilon$, and $\sum_{j=1}^{N_\varepsilon} \pi_{ij}^\varepsilon = 1$ for $i = 1, \ldots, N_\varepsilon$. Beyond their individual states, (s, ε), firms in the final goods sector are also affected by a common productivity shock $z_F \in \{z_1, \ldots, z_{N_F}\}$ with $\Pr\{z_F' = z_c \mid z_F = z_r\} = \pi_{rc}^F \geq 0,\ c = 1, \ldots, N_F$, and $\sum_{c=1}^{N_F} \pi_{rc}^F = 1$ for $r = 1, \ldots, N_F$.

For convenience below, we summarize the economy's exogenous aggregate state by $z = (z_I, z_F)$ and its transition probabilities by $\pi_{lm}^z = \Pr(z' = (z_I, z_F)_m \mid z = (z_I, z_F)_l)$. Each π_{lm}^z is derived from the transition probabilities π_{rc}^I and π_{rc}^F, and we denote the two sectors' aggregate productivity shocks in realized state z_l as z_{Il} and z_{Fl}, for each $l = 1, ..., N_z$ (where $N_z = N_I N_F$). The economy's endogenous aggregate state includes the aggregate capital stock, $\overline{k}$, alongside the distribution of final good firms over start-of-period inventory holdings and individual productivities. We summarize the distribution of firms over (s, ε) using the probability measure μ defined on the Borel algebra, $\mathcal{S}$, generated by the open subsets of the product space $\mathbf{S} \times \mathbf{E}$. We study recursive competitive equilibrium, wherein all households and firms take the evolution of the aggregate state, $(\overline{k}, \mu, z)$, as given, understanding that the endogenous component evolves according to the following laws of motion.

$$\overline{k}' = \Gamma_k\left(\overline{k}, \mu, z\right) \tag{1}$$

$$\mu' = \Gamma_\mu\left(\overline{k}, \mu, z\right) \tag{2}$$

These laws of motion are, in equilibrium, determined by individuals' actions.

2.1 Final good firms

We begin with the firms that produce final goods. As any such firm enters the period, it has expected value $v\left(s, \varepsilon_i; \overline{k}, \mu, z_l\right)$ and makes the following decisions. First, the firm observes a fixed delivery cost associated with stock adjustment, and it chooses whether to pay that cost to order or sell intermediate goods. Following the order decision, we define s_0 as the stock of intermediate goods that the firm holds at the time of production. From this stock, the firm chooses what goods to use in current production, m, where $m \leq s_0$. Given m, and given the employment choice n, the firm's production is $z_{Fl}\varepsilon_i F(m, n)$.

Let e denote the (mid-period) production-time value of a firm identified by (s_0, ε_i);

$$\begin{aligned} e\left(s_0, \varepsilon_i; \overline{k}, \mu, z_l\right) &= \max_{m,n,s'} \Big(z_{Fl}\varepsilon_i F(m, n) - \sigma s' - \omega\left(\overline{k}, \mu, z_l\right) n \\ &\quad + \sum_{m=1}^{N_z} d_m\left(\overline{k}, \mu, z_l\right) \sum_{j=1}^{N_\varepsilon} \pi_{ij}^\varepsilon v\left(s', \varepsilon_j; \overline{k}', \mu', z_m\right)\Big) \\ &\quad \text{subject to} \\ 0 &\leq s' \leq s_0 - m, \text{ and (1) - (2).} \end{aligned} \tag{3}$$

Inventories held for the start of the next period are s', and $\sigma > 0$ is a real storage cost per unit inventory. The real wage, ω, is a function of the aggregate state, as are d_m, $m = 1, \ldots, N_z$. These stochastic discount factors are contingent on the future value of z; in equilibrium, they match the prices of Arrow securities for the representative household. To shorten equations, we will suppress the explicit aggregate state notation for prices often below. We also suppress indices for aggregate and idiosyncratic productivity shocks where possible while avoiding confusion.

At the start of the period, a firm must decide whether to adjust its inventory by undertaking a shipment of intermediate goods. Let $q\left(\overline{k}, \mu, z\right)$ denote the relative price of intermediate goods. We assume that inventory adjustment requires the payment of a fixed delivery cost, ξ, which is denominated in units of labor. While delivery costs are independent of the size of an order, and in this sense fixed, they may vary across firms and over time. Let $\xi \in [\xi_l, \xi_u]$ be drawn from a distribution, $H(\xi)$. If H is degenerate, there is no uncertainty in inventory adjustment costs.

Given its productivity ε, its beginning of period inventory s, and the aggregate state, a final good firm chooses whether to pay its delivery cost and adjust its stock of intermediate goods. The firm may avoid its ξ by placing no order for intermediate goods and achieving expected value $e\left(s, \varepsilon; \overline{k}, \mu, z\right)$. Alternatively, it may pay the delivery cost, choose an order of size $s_0 - s$, and achieve value $e_A\left(\varepsilon; \overline{k}, \mu, z\right) + qs - \omega\xi$. The production-time stock for any adjusting firm with current productivity ε is that which solves:

$$e_A\left(\varepsilon; \overline{k}, \mu, z\right) = \max_{s_0}[e\left(s_0, \varepsilon; \overline{k}, \mu, z\right) - qs_0]. \tag{4}$$

Notice that the gross value of stock adjustment, e_A, is independent of initial inventories, because the firm faces a constant unit price to buy or sell intermediate goods. Finally, the firm's value at the start of the period after the realization of its current delivery cost, $v_1(s, \varepsilon, \xi; \cdot)$, is determined by its choice of whether or not to undertake a stock adjustment:

$$v_1\left(s, \varepsilon, \xi; \overline{k}, \mu, z\right) = \max\left\{e_A\left(\varepsilon; \overline{k}, \mu, z\right) + qs - \omega\xi,\ e\left(s, \varepsilon; \overline{k}, \mu, z\right)\right\}. \tag{5}$$

Having described the binary adjustment decision, we can now define the value function in the right-hand side of (3);

$$v\left(s, \varepsilon; \overline{k}, \mu, z\right) = \int_{\xi_l}^{\xi_u} v_1\left(s, \varepsilon, \xi; \overline{k}, \mu, z\right) H(d\xi). \tag{6}$$

Let $s_0\left(s, \varepsilon, \xi; \overline{k}, \mu, z_l\right)$ describe the decision rule for a firm with beginning of period inventories s, idiosyncratic shock ε, and current stock adjustment cost ξ,. The production time stock of intermediate goods will equal the beginning of period level, s, if the firm does not pay ξ to adjust its stock. Let $\chi\left(s, \varepsilon, \xi; \overline{k}, \mu, z_l\right) = 1$ if the firm pays ξ, 0 otherwise. Further, let $n_0\left(s, \varepsilon, \xi; \overline{k}, \mu, z_l\right)$ and $m_0\left(s, \varepsilon, \xi; \overline{k}, \mu, z_l\right)$ denote the firm's decisions for employment and the quantity of intermediate goods used in production. Its sales may then be expressed as:

$$y_0\left(s, \varepsilon, \xi; \overline{k}, \mu, z_l\right) = z_{Fl}\varepsilon F\left(m_0\left(s, \varepsilon, \xi; \overline{k}, \mu, z_l\right), n_0\left(\xi, \varepsilon, s; \overline{k}, \mu, z_l\right)\right) - \sigma s', \tag{7}$$

where s' is the end of period inventory, $s' = s_0\left(s, \varepsilon, \xi; \overline{k}, \mu, z_l\right) - m_0\left(s, \varepsilon, \xi; \overline{k}, \mu, z_l\right)$. Finally, the firm's net inventory investment in units of the final good is $(s' - s)q$.

2.2 Intermediate goods production

As mentioned above, a representative firm produces intermediate goods using capital and labor. Given our focus on inventory investment, we assume this firm is able to frictionlessly adjust both factors of production subject to one period time-to-build capital accumulation. Given k, the firm chooses n and k' to solve

$$w_I\left(k; \overline{k}, \mu, z_l\right) = \max_{n,k'}\left(q z_{Il} k^{\alpha} n^{1-\alpha} - \omega n - \left[k' - (1-\delta)k\right] + \sum_{m=1}^{N_z} d_m w_I\left(k'; \overline{k}', \mu', z_m\right)\right) \quad (8)$$
$$\text{subject to (1) and (2).}$$

Let $n_I\left(k; \overline{k}, \mu, z_l\right)$ describe the employment chosen by the intermediate producer, and let $k_I\left(k; \overline{k}, \mu, z_l\right)$ be its choice of capital for next period. The firm's production is

$$x_I\left(k; \overline{k}, \mu, z_l\right) = z_{Il} k^{\alpha} n_I\left(k; \overline{k}, \mu, z_l\right)^{1-\alpha}. \quad (9)$$

As noted above, in the baseline version of our model, $N_z = N_I$, $N_F = 1$, and $z_I = z$. There, aggregate shocks to total factor productivity directly affect only the intermediate good producer and only indirectly affect final good firms through movements in relative prices. Having a single shock in the intermediate goods sector ensures that the relative price of inventories is countercyclical, as it is in the data, but it exaggerates the strength of that cyclicality (see Khan and Thomas (2007)). By including a separate aggregate shock in the final goods sector, we bring the model-generated q series closer to the observed relative price of inventories.

2.3 Households

Households own all firms in the economy. However, as they are identical, we do not model a market where shares in firms are traded. Instead we assume that a representative household owns all firms through a mutual fund, and this asset yields aggregate dividends which are the sum of all profits generated by final and intermediate good firms. Given perfect competition and their constant returns to scale production function, intermediate good firms earn no profits. However, because we assume decreasing returns to scale in F, there are profits among final good firms. These profits allow them to pay their fixed costs, and all residual profits are redistributed to the household.

Let a represent the household's beginning of period stock of Arrow securities. The household solves:

$$\begin{aligned} w_h\left(a;\overline{k},\mu,z_l\right) &= \max_{c,n,(a'_m)_{m=1}^{N_z}} \left(u\left(c,1-n\right)+\beta\sum_{m=1}^{N_z}\pi_{lm}^{z}w_h\left(a'_m;\overline{k}',\mu',z_m\right)\right) \\ & \quad \text{subject to} \\ c+\sum_{m=1}^{N_z} d_j\left(\overline{k},\mu,z_l\right)a'_j &\leq a+\omega n+\phi\left(\overline{k},\mu,z_l\right) \\ c &\geq 0,\ a'_m\geq \underline{a} \\ & \quad \text{(1) and (2).} \end{aligned} \tag{10}$$

Where c is current consumption and n is hours of work, $\phi\left(\overline{k},\mu,z_l\right)$ represents the aggregate dividend and $\underline{a}$ is a borrowing limit. Let $c\left(a;\overline{k},\mu,z_l\right)$ define the decision rule for current consumption, $a_m\left(a;,\mu,z_l\right)$ the contingent claims purchased for $z'=z_m$, $m=1,\ldots,N_z$, and $n\left(a;\overline{k},\mu,z_l\right)$ total hours worked.

2.4 Equilibrium

In equilibrium, households maximize utility and firms in both sectors maximize profits. All profits net of inventory adjustment and storage costs are paid to the household as dividends, and the markets for final goods, intermediate goods and labor clear. As there is a representative household, the equilibrium supply of contingent claims, a, is 0. Hence consumption and hours worked are functions only of the aggregate state. This leads to the following equilibrium condition in the market for final goods.

$$c\left(0;k,\mu,z\right)+\left(k_I\left(k;k,\mu,z\right)-\left(1-\delta\right)k\right)=\int_{\xi_l}^{\xi_u}\int_{\mathbf{S}\times\mathbf{E}} y_0\left(s,\varepsilon,\xi;k,\mu,z\right)\mu\left(d\left[s\times\varepsilon\right]\right) \tag{11}$$

In the above, we have also imposed the equilibrium condition that the capital stock of the intermediate good producer is equal to the aggregate capital stock, $k=\overline{k}$. Next, the market-clearing condition for intermediate goods is:

$$x_I\left(k;k,\mu,z\right)=\int_{\xi_l}^{\xi_u}\int_{\mathbf{S}\times\mathbf{E}}\left[s_0\left(s,\varepsilon,\xi;k,\mu,z\right)-s\right]\mu\left(d\left[s\times\varepsilon\right]\right). \tag{12}$$

Recall that $s_0\left(s,\varepsilon,\xi;k,\mu,z\right)\neq s$ only for those firms that pay their labor-denominated delivery costs to adjust their inventories before production, and recall that the indicator function $\chi\left(s,\varepsilon,\xi;k,\mu,z\right)$ identifies such firms. Equilibrium in the labor market then requires:

$$\begin{aligned} n\left(0;k,\mu,z\right) &= n_I\left(k;k,\mu,z\right) \\ &\quad +\int_{\xi_l}^{\xi_u}\int_{\mathbf{S}\times\mathbf{E}}\left[n_0\left(s,\varepsilon,\xi;k,\mu,z\right)+\chi\left(s,\varepsilon,\xi;k,\mu,z\right)\xi\right]\mu\left(d\left[s\times\varepsilon\right]\right). \end{aligned} \tag{13}$$

Given the inventory and production decisions of final good firms, and recalling their sales, y_0, are recorded net of inventory storage costs, total profits are

$$\begin{aligned}\phi(k,\mu,z) \quad = \quad & \int_{\xi_l}^{\xi_u}\int_{\mathbf{S}\times\mathbf{E}}\Big(y_0(s,\varepsilon,\xi;k,\mu,z) \\ & -\omega(k,\mu,z)\left[n_0(s,\varepsilon,\xi;k,\mu,z)+\chi(s,\varepsilon,\xi;k,\mu,z)\,\xi\right] \\ & -q(k,\mu,z)\left[s_0(s,\varepsilon,\xi;k,\mu,z)-s\right]\Big)\mu\left(d\left[s\times\varepsilon\right]\right).\end{aligned} \tag{14}$$

The evolution of the distribution of final good firms, as described by Γ_μ in (2), is given by

$$\mu'\left(S'\times E\right)=\sum_{\varepsilon_j\in E}\pi_{ij}^{\varepsilon}\int_{\{(s,\xi)|s_0(s,\varepsilon,\xi;\cdot)-m_0(s,\varepsilon,\xi;\cdot)\in S'\}}\mu\left(ds,\varepsilon_i\right)H\left(d\xi\right), \tag{15}$$

for all $S'\times E$ measurable. The evolution of the aggregate capital stock, Γ_k in (1) is given by

$$k'=k_I\left(k;k,\mu,z\right). \tag{16}$$

A *recursive competitive equilibrium* is defined by

(i) a value function $v_1(s,\varepsilon,\xi;k,\mu,z)$ that solves the final good firm problem described in (3) - (6) with associated decision rules $s_0(s,\varepsilon,\xi;k,\mu,z)$, $n_0(s,\varepsilon,\xi;k,\mu,z)$, and $m_0(s,\varepsilon,\xi;k,\mu,z)$,

(ii) a value function $w_I(k;k,\mu,z)$ that solves the intermediate good firm problem (8) with associated decision rules $n_I(k;k,\mu,z)$ and $k_I(k;k,\mu,z)$,

(iii) a value function $w_h(a;k,\mu,z)$ that solves the household problem (10) with associated decision rules $c(a;k,\mu,z)$, $n(a;k,\mu,z)$, $(a_m(a;k,\mu,z))_{m=1}^{N_z})$,

(iv) relative prices $\omega(k,\mu,z)$, $q(k,\mu,z)$, and $(d_m(k,\mu,z))_{m=1}^{N_z})$ such that (11), (12) and (13) hold, and $a_m(0,k,\mu,z)=0$ for $m=1,\ldots,N_z$,

where (a) $\phi(k,\mu,z)$ satisfies (14), (b) $y_0(s,\varepsilon,\xi;k,\mu,z)$ is given by (7), (c) $x_I(k;k,\mu,z)$ is given by (9) and (d) Γ_μ and Γ_k are defined by (15) and (16).

3 Decisions

In this section, we re-examine the problems above to isolate several properties that simplify the model's solution. Beginning with final goods firms, we first establish that all firms that place an order and share a common idiosyncratic shock value will make common production and inventory choices. In other words, for all firms that pay ξ to buy or sell intermediate goods, the beginning of period stock, s, is irrelevant to the stock at production-time, s_0, and at the end of the period, s'.

For expositional convenience, we define the conditional expectation of a firm's discounted future value as a function of current productivity ε_i and the current exogenous aggregate state z_l:

$$v^e\left(s'|\varepsilon_i, z_l\right) \equiv \sum_{m=1}^{N_z} d_m\left(k, \mu, z_l\right) \sum_{j=1}^{N_\varepsilon} \pi_{ij}^{\varepsilon} v\left(s, \varepsilon_j; k', \mu', z_m\right) \tag{17}$$

where each $d_m\left(k, \mu, z_l\right)$ again denotes the state-contingent discount factor $d\Big(\left(k', \mu', z_m\right)|\left(k, \mu, z_l\right)\Big)$. Next, using (4) and (3), we have the gross adjustment value of a firm with productivity ε_i:

$$\begin{aligned} e_A\left(\varepsilon_i; k, \mu, z_l\right) &= \max_{s_0, n, s'} \Big(z_{Fl}\varepsilon_i F\left(s_0 - s', n\right) - \omega n - \sigma s' - q s_0 + v^e\left(s'|\varepsilon_i, z_l\right)\Big) \\ & \quad \text{subject to} \\ 0 &\leq s' \leq s_0, \end{aligned}$$

where $\omega = \omega\left(k, \mu, z_l\right)$ and $q = q\left(k, \mu, z_l\right)$. The first-order conditions for s_0 and n are as follow.

$$\begin{aligned} z_{Fl}\varepsilon_i D_1 F\left(s_0 - s', n\right) &= q \\ z_{Fl}\varepsilon_i D_2 F\left(s_0 - s', n\right) &= \omega \end{aligned} \tag{18}$$

Together, these equations imply that, for any firm adjusting its inventory at the start of the period, employment and the use of intermediate goods in production are determined only by its productivity and current relative prices; $n = n\left(\varepsilon_i, k, \mu, z_l\right)$, $m = m\left(\varepsilon_i, k, \mu, z_l\right)$

Because the firm's delivery cost is independent of the size of its order, the past stock s also has no influence on its end-of-period stock in periods when it undertakes an order. In any period that a firm places an order for intermediate goods, its end-of-period stock solves:

$$\max_{s'>0}[-(q+\sigma)s' + v^e\left(s'|\varepsilon_i, z_l\right)].$$

If the firm adopts an interior s', it will satisfies the first-order condition:

$$Dv^e\left(s'|\varepsilon_i, z_l\right) = q + \sigma. \tag{19}$$

In any case, s' is clearly independent of s for a stock adjustor, and may be denoted by $g\left(\varepsilon_i, k, \mu, z_l\right)$. This implies that the production-time stock s_0 is also independent of s, since we know from above that m has this same property.

Once s' is known, (18) immediately yields $s_0 = g_0\left(\varepsilon_i, k, \mu, z_l\right)$, the production-time level of inventories for a firm currently adjusting its inventories. Given these observations, final good firms may be viewed as experiencing inventory episodes, with each such episode beginning at the date of a stock adjustment and ending at the date of the next one. This is useful in that a firm's decisions inside each inventory episode are independent of its past episodes.

The next observation is that the binary decision regarding stock adjustment is determined by reference to a threshold order cost that varies as a function of the aggregate state and the firm's current inventory and productivity. Within the set of firms sharing common s and ε, each will pay its order cost $\xi \in [\xi_l, \xi_u]$ if the following inequality is satisfied.

$$-\xi\omega + e_A\left(\varepsilon;\overline{k},\mu,z\right) \geq e\left(s,\varepsilon;\overline{k},\mu,z\right) \tag{20}$$

Define the threshold order cost, $\xi^T\left(s,\varepsilon;\overline{k},\mu,z\right)$, as that ξ at which (20) holds with equality. The decision rule for the production-time stock is

$$s_0 = \begin{cases} g_0\left(\varepsilon_i, k, \mu, z_l\right) & \text{if } \xi \leq \xi^T\left(s,\varepsilon;\overline{k},\mu,z\right) \\ s & \text{otherwise} \end{cases}.$$

The fraction of firms of type (s,ε) that order intermediate goods is given by $H\left(\xi^T\left(s,\varepsilon;\overline{k},\mu,z\right)\right)$. The remaining type (s,ε) firms adopt common employment, production and end-of-period stocks. Their decisions depend on the full firm-level state $(s,\varepsilon;\overline{k},\mu,z)$, and the constraint $0 \leq s' \leq s - m$ may bind.

Turning to the intermediate good sector, recall that the representative firm there has constant returns production and thus earns no profits. The linearity of its problem means it must be indifferent to any level of future capital k' in equilibrium, a familiar results we briefly establish in the context of our model. Given its capital and the aggregate state, the firm's optimal choice of employment solves $\max\limits_n \left(qz_{Il}k^\alpha n^{1-\alpha} - \omega n\right)$, which leads to the following linear decision rules for employment and the supply of intermediate goods.

$$\begin{aligned} n_I\left(k;k,\mu,z_l\right) &= \left((1-\alpha)\frac{q}{\omega}z_{Il}\right)^{\frac{1}{\alpha}} k \\ x_I\left(k;k,\mu,z_l\right) &= z_{Il}\left((1-\alpha)\frac{q}{\omega}z_{Il}\right)^{\frac{1-\alpha}{\alpha}} k \end{aligned}$$

These decision rules directly imply a linear profit function, which in turn implies the firm's value function is linear in k. Once we use the firm's employment choice to simplify (8), the first-order condition for k' is:

$$1 = \sum_{m=1}^{N_z} d_m\left(k,\mu,z_l\right)\left(\alpha\left(\frac{(1-\alpha)}{\omega_m}\right)^{\frac{1-\alpha}{\alpha}}\left(q_m z_m\right)^{\frac{1}{\alpha}} + 1 - \delta\right), \tag{21}$$

where $\omega_m = \omega(\Gamma_k\left(k,\mu,z_l\right), \Gamma_\mu\left(k,\mu,z_l\right), z_m)$ and $q_m = q(\Gamma_k\left(k,\mu,z_l\right), \Gamma_\mu\left(k,\mu,z_l\right), z_m)$. As long as this restriction on prices is satisfied, the intermediate good firm is indifferent to k', and its investment is determined as the residual of final goods after consumption demand is satisfied.

Finally, we turn to households for restrictions on the equilibrium wage and firms' state-contingent discount factors. As noted above, household consumption and hours worked are functions of only the aggregate state in equilibrium. The household efficiency condition with respect to hours worked implies that the equilibrium real wage is given as:

$$\omega\left(k,\mu,z_l\right)=\frac{D_2u\left(c\left(0,k,\mu,z_l\right),1-n\left(0,k,\mu,z_l\right)\right)}{D_1u\left(c\left(0,k,\mu,z_l\right),1-n\left(0,k,\mu,z_l\right)\right)}.$$

Next, to ensure equilibrium in the markets for firm shares, each firm's state-contingent discount factor must agree with the equilibrium price of an Arrow security for state z_m;

$$d_m\left(k,\mu,z_l\right)=\frac{\beta\pi^z_{lm}D_1u\left(c\left(0,\Gamma_k\left(k,\mu,z_l\right),\Gamma_\mu\left(k,\mu,z_l\right),z_m\right),1-n\left(0,\Gamma_k\left(k,\mu,z_l\right),\Gamma_\mu\left(k,\mu,z_l\right),z_m\right)\right)}{D_1u\left(c\left(0,k,\mu,z_l\right),1-n\left(0,k,\mu,z_l\right)\right)}.$$

4 Numerical solution

The aggregate state of the model includes a distribution of firms over idiosyncratic productivity levels and inventories, as well as capital and total factor productivity. Furthermore, firm-level decisions involve discrete choices. Given occasionally binding non-negativity constraints on their stocks, a standard linear solution method is infeasible. Turning to a nonlinear method, we must immediately confront the problem of a large aggregate state vector. Following a similar approach to that in Khan and Thomas (2008), we solve the model using the state space approximation method of Krusell and Smith (1997). In particular, we impose an assumption that all agents in our economy forecast equilibrium prices as functions of m, a vector of moments drawn from μ. We use $(k,m;z_l)$ as a proxy for the aggregate state, replacing $\Gamma_\mu\left(k,\mu;z_l\right)$ with a law of motion for Γ_m, $\Gamma_m\left(k,m;z_l\right)$ and replacing $\Gamma_k\left(k,\mu;z_l\right)$ with $\Gamma_k\left(k,\mu;z_l\right)$.

Our solution iterates to arrive at accurate forecasting rules. Let $\left(\Gamma^i_k,\Gamma^i_m\right)$ represent the $i-th$ iteration. Given these rules, we reformulate the problem of final good firms, having them weight their current profits by $p\left(k,m;z_l\right)$, which in equilibrium is the household marginal utility of consumption, and discount their future values by β. This eliminates Arrow securities from the analysis, but does not affect decision rules (see Khan and Thomas (2008)).

The solution algorithm iterates between an inner loop step, where we solve the analogue to (3) - (6) and (8) given equilibrium price functions $p^i\left(k,m;z_l\right)$ and $q^i\left(k,m;z_l\right)$, and an outer loop step where, given these value functions, and the forecasting rules $\left(\Gamma^i_k,\Gamma^i_m\right)$, we simulate the economy.[5] The simulation step determines equilibrium prices p and q as a function of the actual aggregate state (k,μ,z_l). Given the intertemporal relative price, p, the relative price q is determined to clear the market for intermediate goods. The equilibrium level of p is that which implies a level of consumption that leads to a (k',m') such that the restriction implied by the intermediate good

[5] A separate price function for the real wage follows immediately from $p^i\left(k,m;z\right)$ as we assume a representative household whose period utility function has constant marginal utility of leisure. This is the familar result of an environment where households face indivisible labor supply decisions and have access to employment lotteries (Rogerson (1988), Hansen (1985)).

firm's first-order condition for capital in (21) is satisfied. The moments m' are derived from the actual distribution, μ', which is stored between one date and the next over the simulation using a two-dimensional grid. We use a weighted grid approximation method which, alongside multivariate spline approximation of firms' value function, does not constrain individual decisions to grid values. Using simulation data, new laws of motion $\left(\Gamma_k^{i+1}, \Gamma_m^{i+1}\right)$ are derived.

5 Calibration

We evaluate the plant-level and aggregate implications of inventory investment using several numerical examples across which we vary the stochastic process for idiosyncratic shocks to plants' total factor productivity and the distribution of order costs. All other production parameters, as well as preferences, are held constant throughout. Below, we discuss functional forms and parameter values for technology and preferences that are common across models. In section 5.2, we explain our choices for idiosyncratic shocks and the distribution of stock adjustment costs.

5.1 Common parameters

Across our model economies, we assume that the representative household's period utility is the result of indivisible labor (Hansen (1985), Rogerson (1988)): $u(c, 1-n) = \log c + \eta\,(1-n)$, and the final good firm production function takes a Cobb-Douglas form, $F(m,n) = m^{\theta_m} n^{\theta_n}$.

Model parameters, other than those involving idiosyncratic productivity shocks and order costs, are selected to ensure agreement with observed long-run values for key postwar U.S. aggregates in a nested *control* version of our model without order costs. We assume the length of a period is one-quarter, and set the discount factor, β, to imply an average annual real interest rate of 4 percent. The depreciation rate, δ, is selected to match an average investment-to-capital ratio of 10 percent, corresponding to the average value for the private capital stock between 1954 and 2013 in the U.S. Fixed Asset Tables. This implies an annual depreciation rate of 0.069.

Intermediate good share, $\theta_m = 0.499$ (see Khan and Thomas (2007)), and α and θ_n are jointly chosen to imply a labor's share of 0.6 and an average aggregate capital-to-output ratio of 2.33 as in the data. Next, the parameter governing the preference for leisure, η, is taken to imply an average of one-third of available time spent in market work.

In the baseline version of our model, we assume a single aggregate productivity shock affecting the intermediate good sector; $N_z = N_I$, $N_F = 1$, and $z_I = z$. In specifying the exogenous stochastic process there, we begin by assuming a continuous shock following a mean zero AR(1) process in logs: $\log z_I' = \rho_I \log z + \eta_I'$ with $\eta_z' \sim N\left(0, \sigma_{\eta_I}^2\right)$. Next, we estimate the values of ρ_I and σ_{η_I} from Solow residuals measured using NIPA data on US real GDP and private capital, together with the total employment hours series constructed by Prescott, Ueberfeldt, and Cociuba (2005) from CPS household survey data, over the years 1959-2002. Finally, we discretize the resulting productivity process using a grid with 9 shock realizations; $N_I = 9$. When we extend the model to include an independent aggregate shock to final good firms' productivity, we assume the second shock is also log-AR(1) and discretize it. There, we set $N_F = 2$ and target ρ_F and σ_{η_F} to best match the observed cyclical behavior of the relative price of inventories in our inventory model. Table 3 summarizes the parameters obtained to this point.

TABLE 3. Common parameter set

β	δ	θ_m	α	θ_n	η	ρ_z	σ_{η_I}	ρ_F	$\sigma_{\eta F}$
0.99	0.017	0.499	0.514	0.358	2.19	0.956	0.015	0.99	0.006

5.2 Firm-level productivity, order and storage costs

The remaining parameters involve the distribution of firm-specific productivity and the fixed costs incurred with shipments of intermediate goods, alongside the per-unit inventory storage cost, σ. To illustrate the effect of idiosyncratic shocks and uncertainty in order costs, we consider several examples. Across these, our goal is to match the average aggregate inventory-to-sales ratio and total storage costs and then examine each model's ability to reproduce firm-level measures of the standard deviation of output and sales growth. The average real quarterly nonfarm inventory-to-sales ratio in the U.S. over the period 1954Q1 - 2014Q3 is 0.81. In each of our examples, we select an upper bound on the distribution of order costs, ξ_u, to reproduce this level of inventories. We jointly select the storage cost parameter, σ, so that storage costs average 12 percent of the annual value of inventories as in the data (see Khan and Thomas (2007)).

We implement firm-level productivity shocks using a Markov Chain $(\varepsilon_i)_{i=1}^{N_\varepsilon}$ and $\left(\pi_{ij}^{\varepsilon}\right)_{i,j=1}^{N_\varepsilon}$ with $N_\varepsilon = 15$ to discretize the log-normal process, $\log \varepsilon' = \rho_\varepsilon \log \varepsilon + \eta'$. There is little agreement about the persistence of the idiosyncratic shock process, ρ_ε.[6] In the lead case we will pursue below, we fix persistence at 0.5; in every case, we set the volatility of innovations to imply $\sigma_\varepsilon = 0.02$. As explained in Khan and Thomas (2008), equilibrium movements in relative prices imply that the elasticity of response of firms to idiosyncratic shocks is roughly 12 times greater than that to aggregate shocks. Thus, while the standard deviation of idiosyncratic shocks is only 40 percent that of aggregate shocks in our examples, idiosyncratic shocks have a much larger impact on the production and inventory decisions of individual firms.

Comin and Phillipon (2005) use COMPUSTAT data to calculate the median standard deviation of 10-year centered moving averages of output growth, and find that this measure has been rising over time. At the beginning of their sample, in 1955 the standard deviation was approximately 0.1; by 2000, it was above 0.2. Bloom uses the COMPUSTAT over the period 1981 -

[6] Compare, for example, the values in Comin and Phillipon (2005) to those of Cooper and Haltiwanger (2006).

2000 and reports an average standard deviation of annual sales growth of 0.165. We evaluate our model's consistency with these firm-level growth volatility statistics below.

We consider two types of order cost distribution. The first is uniform between 0 and ξ_u. The second is generalized Beta with shape parameters (a, b) and upper support ξ_u. The special case with $a = b = 1$ is a uniform distribution; as we move further from that starting point, uncertainty regarding order costs is reduced.

Table 4 overviews a variety of cases. The first, model 0, is a special case without firm-level productivity shocks where order costs are drawn from the uniform distribution. This special case, based on Khan and Thomas (2007) serves as a baseline against which to consider the performance of our more general variants. In comparison with the empirical measures from Bloom (2009) and Comin and Phillipon (2005), this model generates far too much firm sales and output growth volatility. Continuing to focus on uniform order costs, models 1 - 3 allow firm productivity shocks ranging in persistence, with similarly poor firm-level results.

TABLE 4. Candidate models

	σ_ε	ρ_ε	a	b	ξ_u	σ	Bloom statistic	C & P statistic
DATA	−−	−−	−−	−−	−−	−−	**0.165**	**0.100 − 0.200**
model 0	**0.00**	**1.0**	**1**	**1**	**0.313**	**0.0088**	**0.971**	**0.408**
model 1	0.02	0.0	0	0	0.301	0.0089	0.967	0.442
model 2	**0.02**	**0.5**	**1**	**1**	**0.301**	**0.0089**	**1.515**	**0.429**
model 3	0.02	0.9	1	1	0.301	0.0089	0.920	0.438
model 4	0.02	0.0	5	2	0.174	0.0089	0.173	0.160
model 5	**0.02**	**0.5**	**5**	**2**	**0.164**	**0.0089**	**0.194**	**0.183**
model 6	0.02	0.9	5	2	0.163	0.0089	0.185	0.173
model 7	0.02	0.5	2	5	0.507	0.0089	1.116	0.370

In models 4 - 6, we eliminate much of the uncertainty over order costs, adopting a right skewed Beta distribution. Each of these models reproduces the observed volatility of firm-level sales and output growth fairly closely. Given their similar results and the lack of consensus regarding ρ_ε, we will adopt the intermediate value ('model 5') as our lead candidate going forward. Results for model 7 appear in the table only to emphasize that it is right-skewness in the distribution of order costs, not skewness alone, that allows the model to reproduce the firm-level moments; models we consider with left-skewed distributions consistency fail.

6 Results

6.1 Steady state

We begin with a series of steady state results designed to highlight aspects of firm-level behavior unique to our leading model (model 5). These help explain why our model so outperforms the other bolded models in Table 4, and also shed light on some of the cross-model differences we will see as we examine aggregate fluctuations below. In Figure 1, we examine how uncertainty in order cost draws influences a firm's inventory management and sales. Figure 1 compares 50 periods in the life of a firm facing uniform cost draws (model 2) to that for a firm drawing order costs from the right-skewed beta distribution (model 5).

FIGURE 1. Order cost uncertainty, firm inventory management and sales

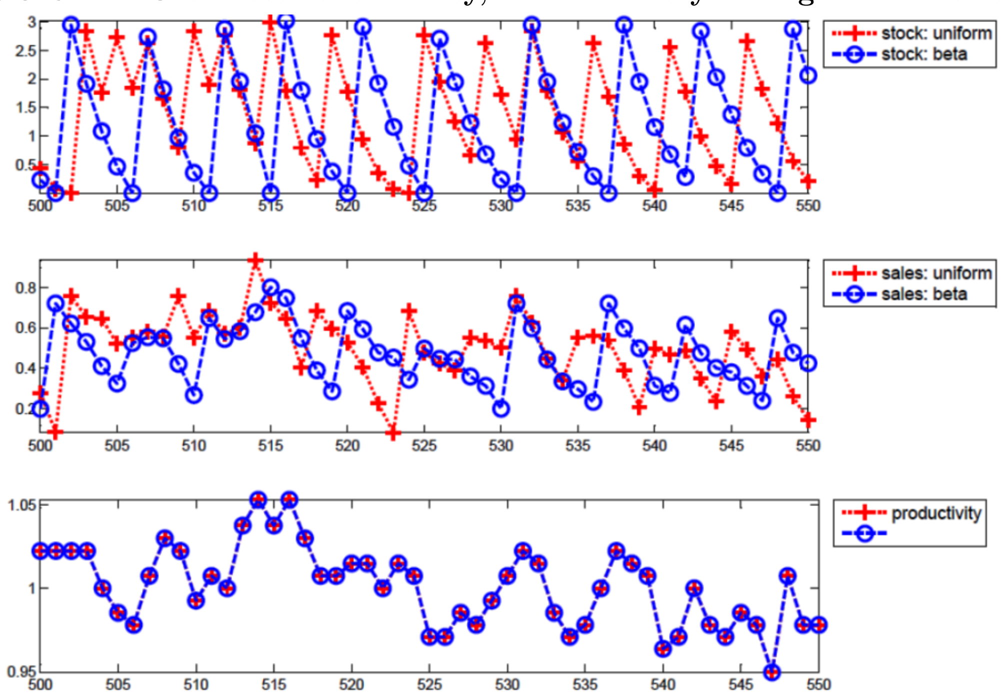

Notice that, as we move away from the uniform order cost distribution, the firm's pattern of inventory adjustment has a more regular Baumol-Tobin appearance. Under the skewed order cost distribution, the mean cost realization is near the maximum one, so the firm's timing of an order is fairly predictable. That observation is confirmed in the first two columns of Table 5 below, which presents firm-averaged moments from a panel of firms residing in model 2 versus model 5. Moving to the right-skewed distribution leads to a longer average duration between inventory adjustments, and much less variability in that duration, as reflected in the substantially lower coefficient of variation.

Beyond its more predictable inventory cycle, the firm facing beta cost draws in Figure 1 has a smoother sales series, with turning points more closely linked the relative productivity series in the bottom panel. These points are reinforced by the firm-averaged moments in Table 5. When firms face less uncertainty about order costs, variation in firm-level sales falls, and changes in firms' sales are less affected by their inventories, and more responsive to changes in their productivities.

TABLE 5. Averages from a common panel across models

	avg(duration)	cv(duration)	cv(sales)	corr(ε, y)	corr(s, y)
model 2 (uniform)	3.3 qtrs	0.47	0.38	0.38	0.36
model 5 (skewed)	4.0 qtrs	0.16	0.27	0.53	0.12

Aside from its more predictable order costs, our model is distinguished from the inventory model of Khan and Thomas (2007) by its firm-specific productivity shocks. Their presence implies a very different steady state distribution of firms over inventory levels, as may be seen below in Figure 2. The left panel is 'model 0' from Table 4, where all final good firms have the same productivity.[7] That model has a single rising adjustment hazard. As inventories are depleted, the gap between a firm's current stock and the steady state target stock rises. As that gap widens, firms grow willing to incur larger order costs to correct it, so an increasing fraction of firms pay order costs and adjust their stock back to the common target around $s = 2.8$. Thus, we see ever fewer firms as we look from the target stock leftward.

FIGURE 2. Steady state distribution without/with firm-level productivity shocks

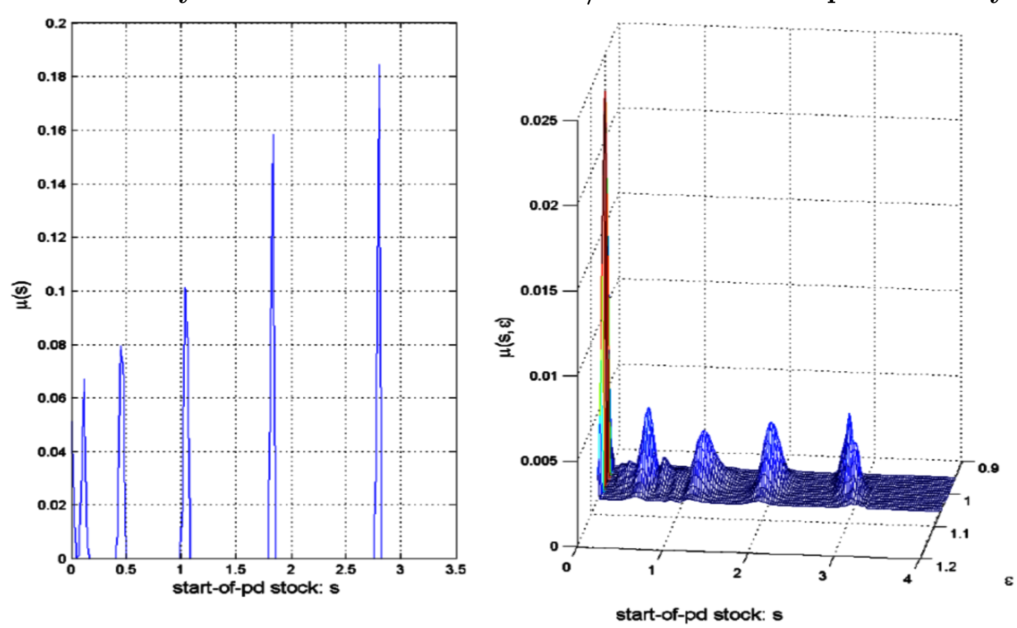

[7] The left panel figure has the same basic features when we replace the uniform order costs with our right-skewed beta distribution.

By contrast, 'model 5' in the right panel has a richer distribution, with firms spread more continuously over inventory levels. Firms' stock adjustments are still governed by a rising adjustment hazard; but now we have a distinct hazard for each firm-specific productivity level. If a firm has recently adjusted its stock, and its relative productivity falls, it can find itself with inventory exceeding the target level consistent with its new productivity. Thus, we see firms appearing on the right ramps of hazards, some of which will pay fixed costs to sell a part of their stock. In other words, (S,s) inventory adjustments are 2-sided in even ordinary times in our model.

6.2 Business cycle fluctuations

Table 6 recalls key aspects of the dynamics of HP-filtered inventory investment, final sales and GDP in the postwar U.S., and compares these to the corresponding moments from three cases of our inventory model. The columns labeled 'model 0, 'model 2' and 'model 5' are the cases bolded above in Table 4. The new case, 'model 5B' is the first time the independent aggregate shock to the final goods sector is active; otherwise, it shares a common parameter set with model 5, and thus the same steady state distribution of firms.

Table 6 shows that each of our inventory models is broadly successful with respect to the basic business cycle facts involving inventories. Net inventory investment is strongly procyclical in every instance, and our models explain between 51 and 58 percent of the empirical relative volatility of this series. Next, given success in generating a positive correlation between final sales and inventory investment, each model also reproduces the fact that sales is less cyclically volatile than GDP. Finally, given that the aggregate inventory stock responds to shocks quite gradually relative to the flow of final sales, the model also succeeds in generating a countercyclical inventory-to-sales ratio, though this is over-stated relative to the data.

TABLE 6. Confronting the business cycle facts

	DATA	model 0	model 2	model 5	model 5B
sd(GDP)	2.122	1.883	1.859	1.928	2.305
sd(NII)/sd(GDP)	0.283	0.143	0.144	0.163	0.155
corr(NII,GDP)	0.639	0.790	0.885	0.721	0.432
corr(NII,FS)	0.397	0.726	0.847	0.628	0.294
sd(FS)/sd(GDP)	0.874	0.891	0.875	0.890	0.943
corr(FS,GDP)	0.945	0.995	0.997	0.992	0.989
sd(I-S)/sd(GDP)	0.660	1.041	0.996	1.027	0.896
corr(I-S,GDP)	−0.442	−0.956	−0.956	−0.945	−0.905
corr(q,GDP)	−0.097	−0.978	−0.978	−0.928	−0.142

On the basis of the moments in Table 6, it appears that our inclusion of firm-level productivity shocks and greater certainty about order costs have minimal implication for the movements in inventories, total production and final sales over the business cycle. The relative volatility of inventory investment rises somewhat, and the overly strong co-movement of final sales and inventory weakens, but otherwise the moments are close.

Our purpose in including an aggregate productivity shock influencing final goods firms directly was to weaken excessively strong positive and negative co-movement arising in our single shock models. Without resetting the properties of the original aggregate shock, we have chosen the second shock in model 5B to eliminate the almost perfectly negative relationship between GDP and the relative price of the goods held in inventory. Thus, by design, this model achieves a much better result in the final row of the table. The new shock weakens correlations elsewhere, in most cases taking model closer to data.

6.2.1 baseline inventory model versus no-inventory control

We noted above that earlier work by Kryvtsov and Midrigan (2013) and Khan and Thomas (2007) studied the implications of inventories for cyclical fluctuations in GDP and found minimal effects. We consider here whether that invariance results still stands in our model. Table 7 compares HP-filtered business cycle moments from three versions of our model (rows bolded in Table 4) to those from an otherwise identical control model with no inventory motive.

TABLE 7. Cyclical role of inventories in models with one aggregate shock

	GDP	NII	FS	C	I	N	N(i)	N(fg)
A. percent std. deviations								
control	1.94	−−	1.94	0.57	9.56	1.45	1.45	1.45
inventory model 0	1.88	0.27	1.68	0.51	8.93	1.46	1.70	1.33
inventory model 2	1.86	0.27	1.63	0.51	8.40	1.42	1.69	1.26
inventory model 5	1.93	0.31	1.72	0.52	8.90	1.47	1.76	1.32
B. correlations with GDP								
control		−−	1.000	0.889	0.985	0.984	0.984	0.984
inventory model 0		0.790	0.995	0.866	0.978	0.986	0.979	0.984
inventory model 2		0.885	0.997	0.870	0.988	0.986	0.978	0.988
inventory model 5		0.721	0.992	0.867	0.978	0.975	0.952	0.980

The control model is a special case of our model(s) where there are no fixed order costs, and hence there are no inventories. Because this setting has no nonconvexity, its aggregate dynamics are impervious to the presence of idiosyncratic productivity shocks (see Khan and Thomas (2008)). Thus, a single set of control model results suffices for comparison with the inventory models without (model 0) and with (models 2 and 5) firm-specific productivities.

The control model bears several familiar features of a real business cycle. First, given household preferences for smooth consumption profiles, the volatility of consumption is exceeded by that of final sales (a series identical to GDP in this model), which is, in turn, exceeded by that of capital investment. Investment is by far the most volatile series, both because it facilitates household consumption smoothing and because the persistence of aggregate shocks implies that an unanticipated rise (fall) in the marginal products of labor and capital arising from a shock leads households to anticipate a protracted rise (fall) in the returns to investment, causing strong substitution effects on savings. Moreover, the second mechanism encouraging procyclical investment compounds the substitution effect of the procyclical wage on labor supply. Thus, at the impact of a positive aggregate productivity shock, consumption, investment and hours worked rise alongside GDP, explaining the strong contemporaneous correlations with GDP in panel B of the table. Finally, because there is no friction in the flow of intermediate goods to final goods producers in this model, there is effectively no separate intermediate goods sector. As such, the dynamics of the labor input in final versus intermediate good production are the same.

Next, we consider the dynamics of a baseline inventory model like that examined in Khan and Thomas (2007). Recall that, in model 0, the fixed delivery costs causing inventories are distributed uniformly, and there are no idiosyncratic productivity shocks. Aside from the obvious distinction arising between final sales and GDP, the basic features of the business cycle discussed above are largely unchanged. One difference deserves comment. Despite its positive correlation between inventory investment and final sales (0.73), the baseline inventory model actually exhibits lower cyclical volatility in GDP than does the model without inventories. As was pointed out in Khan and Thomas (2007), the frictions that cause inventories need not necessarily amplify business cycles, because procyclical inventory accumulation causes a dampening of fluctuations in final goods production. This happens due to a trade-off between the use of intermediate goods in production (raising final sales) versus their storage in inventory, as we explain below. Figures 3 - 5 complement this discussion, presenting impulse responses from the model 0 inventory economy following a persistent 1-standard deviation rise in aggregate productivity.[8]

When a positive aggregate productivity shock hits the inventory economy, the marginal cost of intermediate goods production is reduced relative to normal, which in turn reduces the relative

[8] This tradeoff is more prominent in our current baseline model relative to the analogous model in Khan and Thomas (2007), because the capital share in intermediate goods production is larger here. For this reason, the dampening of final sales volatility outweighs the volatility arising from the inventory investment series, despite a quite similar contemporaneous correlation between final sales and inventory investment.

price of intermediate goods. In response, an increased fraction of final good firms place orders, as seen in the top panel of Figure 3. The resulting rise in total demand for intermediates is compounded by the fact that those final good firms that do place orders also order more than they would ordinarily, which raises the stock they have available at production time (middle panel). However, only part of their extra stock goes into current production. Because these firms anticipate that the raised demand for their goods used for consumption and investment will persist, they expect that any given stock of intermediate goods will be exhausted more quickly than usual. To avoid more frequent payments of fixed delivery costs, these firms raise their end-of-period inventories above normal levels (bottom panel). Indeed, they do so in sufficient quantity as to more than offset the raised use of intermediate goods among non-ordering firms that are constrained by their existing stocks. Thus, in the aggregate, total orders of intermediate goods rises by more than total use, as seen in the middle panel of Figure 4. This explains why inventory investment is procyclical despite the countercyclical relative price (q) valuing stocks.

FIGURE 3. Firm-level inventory responses in baseline model 0

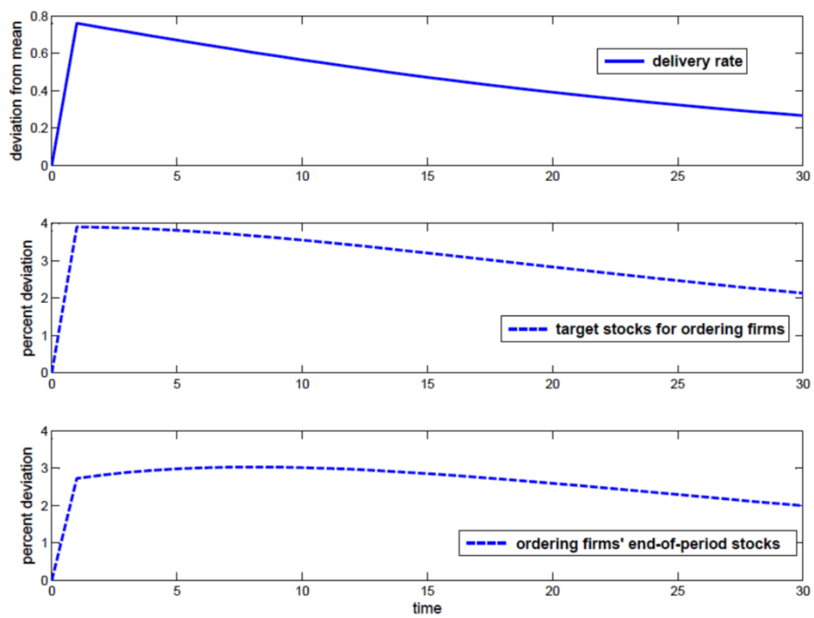

Given any fixed stock of intermediate goods, there must be a trade-off between their use and storage. The more responsive is the production of these goods to a positive productivity shock, the weaker will be this trade-off. However, whenever the predetermined capital input has an empirically plausible share in production (as it does here), the fall in the marginal cost of intermediate goods production following a positive productivity shock does not happen at once. Instead, this fall happens gradually as the capital stock is increased following the shock, so that the initial rise in total supply of intermediate goods is gradualized. This has several related implications. First, despite a strong rise in the demand for intermediate goods, any rise in inventory investment comes at the expense of rises in the production of final goods, as is reflected in the final columns of Table 7 and in the top and bottom panels of Figure 4.

FIGURE 4. Aggregate flow variable responses in baseline model 0

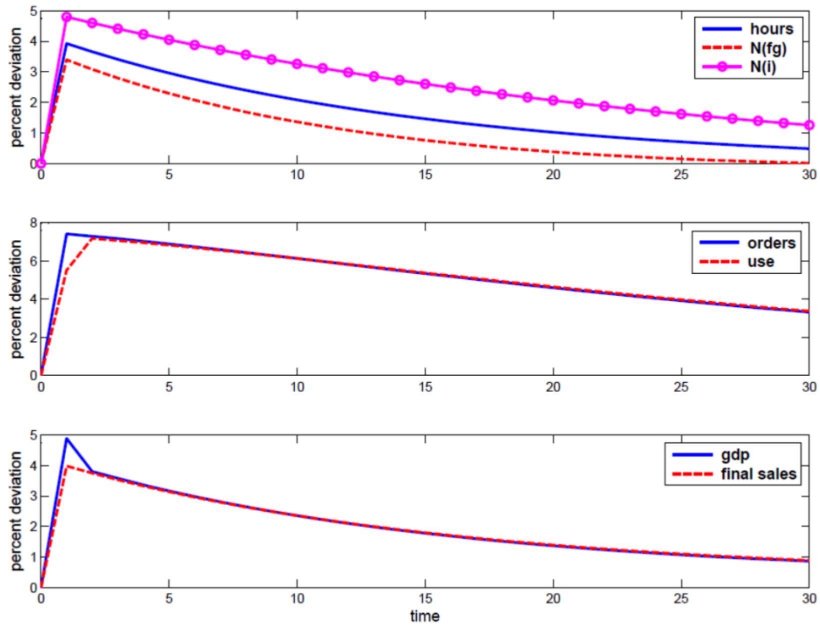

While a considerably greater rise in the demand for intermediate goods leads to a much sharper rise in employment and production in that sector relative to the control model, employment in final goods production rises by less. In other words, the use of intermediate goods rises by less in the inventory model. This implies a dampened rise in final sales, which reduces both the rise in consumption and in capital investment, explaining the reduced volatility in both series relative to the control model. Turning to Figure 5, that dampened response in capital investment protracts the episode over which the rise in intermediate goods production is insufficient to prevent a trade-off between final goods production and inventory accumulation, because it gradualizes the rise in the capital stock (top panel). Throughout this entire episode, the same trade-off that reduces volatility in final sales also limits the volatility of inventory investment. Thus, the rise in aggregate inventory holdings (middle panel) does not happen all at once, but instead takes roughly 10 periods. That, in turn implies that the ratio of the real aggregate stock of inventories relative to the flow of final sales, which naturally drops sharply at the impact of the shock, recovers very slowly thereafter. On balance, the fluctuations in final sales are sufficiently reduced by the frictions causing inventories in our model as to more than offset the procyclical movements in inventory investment, despite the co-movement of sales and inventory investment.

FIGURE 5. Aggregate stock responses in baseline model 0

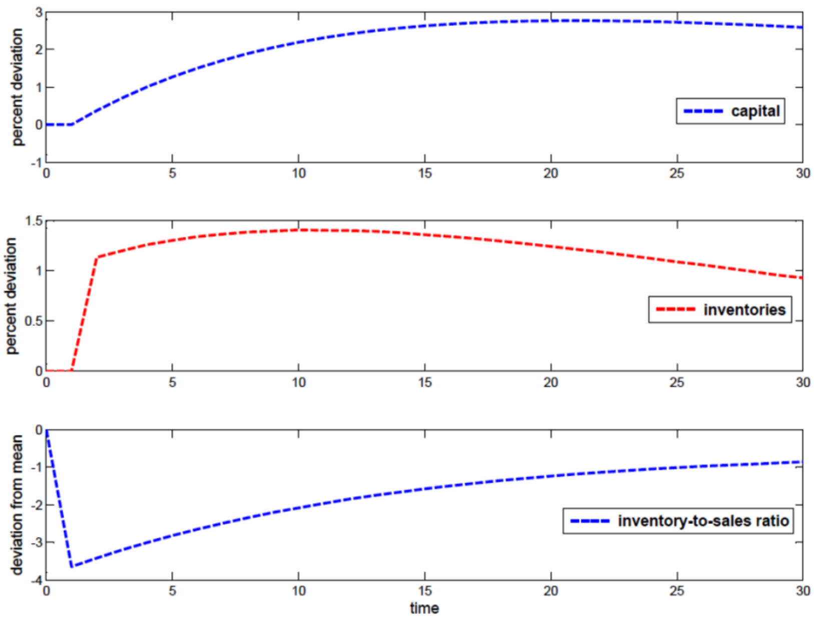

6.2.2 adding firm-level productivity differences

The third rows of Table 7 are from the model where final goods firms again face uniformly distributed fixed costs to order inputs, but now experience idiosyncratic productivity shocks. In model 2, the basic trade-off between the use and accumulation of intermediate continues to apply, and we see similar features in second moments as before. Returning to the impulse responses following an aggregate productivity shock discussed above, the one notable difference here is a lesser rise in employment and production in the final goods sector, which in turn slows capital accumulation. The trade-off between use and storage of goods is stronger when the location of production matters. With moderately persistent idiosyncratic productivity shocks, the very firms that should be responding most in their production and sales at the impact of a positive aggregate shock also know their relative productivity is likely to remain high in coming periods of high aggregate demand. Given large uncertainty about the order costs they may face in future to replenish their intermediate goods, they are particularly cautious in managing their inventories.

6.2.3 reducing uncertainty regarding order costs

In the final rows of Table 7, we have the leading case of our model (recall its much closer fit with respect to the firm level data in Table 4). Here, we maintain our firm-level productivity shocks, but replace the distribution from which fixed order costs are drawn with the right-skewed beta distribution. As we noted above, this distribution implies more certainty about the timing of firms' orders, because it concentrates more probability mass at the right-tail in the range of relatively high order costs. In this case, firms can predict the timing of their stock adjustments reasonably well. This means firms generally leave each of their inventory cycles with very little remaining stock, as was evident above in the top panel of Figure 1. In other words, they do not carry large precautionary stocks.

As was true in model 2 above, final goods firms with high relative productivities are the most responsive in moving up the timing of their stock adjustments in response to the aggregate shock in model 5. However, they are now more certain about how long their next inventory cycle will last. Thus, they place larger orders than they would otherwise, and they retain less excess precautionary inventory. In the aggregate, firms placing orders raise their production-time stocks much more than did firms in model 2. Moreover, given far greater certainty that they will wait a full 5 quarters before their next stock adjustment, these firms exit the period having raised their end of period inventory by more. At the same time, final sales (and thus consumption and investment) also rises slightly more. Referring back to Table 5, this is because final good firms are much more responsive to their relative productivities in their production and sales when their inventory cycle is more predictable. In other words, the allocation of production is more efficient. Thus, overall GDP volatility in Table 7 rises to almost exactly that in the control model.

6.2.4 effects of a second aggregate shock

Does the inclusion of a second aggregate shock overturn the basic invariance result discussed above? No. Table 8 presents the same moments as in Table 7, this time comparing inventory model 5B to a control model with the same simulation of z_I and z_{FG}. Here, we see the same basic trade-off between use and storage of intermediate goods as in Table 7, with the presence of the inventory motive again dampening fluctuations in final sales sufficiently to all but completely offset procyclical inventory investment.

TABLE 8. Cyclical role of inventories with two aggregate shocks

	GDP	NII	FS	C	I	N	N(i)	N(fg)
A. percent std. deviations								
control	2.29	−−	2.29	0.83	10.85	1.59	1.59	1.59
inventory model 5B	2.31	0.357	2.18	0.81	10.76	1.60	1.85	1.49
B. correlations with GDP								
control		−−	1.000	0.894	0.973	0.972	0.972	0.972
inventory model 5B		0.432	0.989	0.875	0.962	0.960	0.919	0.968

6.3 High frequency results

In this section, we report how the specifications of firm-level shocks affect the aggregate dynamics of our (S,s) inventory model in higher frequencies corresponding to 2-6 quarters per cycle. Here again, we filter our model-generated data just as we do the actual data.

TABLE 9. Confronting the facts in high frequencies

	DATA	model 0	model 2	model 5	model 5B
sd(GDP)	0.547	0.783	0.778	0.769	0.921
sd(NII)/sd(GDP)	0.614	0.162	0.153	0.290	0.278
corr(NII,GDP)	0.538	0.831	0.936	0.468	0.206
corr(NII,FS)	−0.201	0.768	0.912	0.197	−0.074
sd(FS)/sd(GDP)	0.830	0.870	0.859	0.901	0.981
corr(FS,GDP)	0.642	0.995	0.998	0.959	0.961
corr(q,GDP)	0.200	−0.996	−0.997	−0.771	−0.165

Focusing first on the empirical moments, we note again the ways in which the basic inventory facts change as we move into the higher frequencies. First, the relative volatility of inventory investment is now much greater (0.61 versus 0.28 in the business cycle frequencies). Next, while inventory investment and final sales still have strong positive correlations with GDP, they no longer co-move with each other. Their significantly positive contemporaneous correlation over the business cycle (0.40) is replaced by a weakly negative one (−0.20). Table 9 examines whether our single-motive (S,s) inventory model successful with respect to the business cycle inventory facts can simultaneous deliver on these changes in the high frequencies.

Table 9 rules out model 0 and model 2 as good high frequency inventory models. In contrast to the data, these models have even stronger co-movement between GDP, final sales and inventory investment in the high frequencies than over the business cycle. Furthermore, the relative volatility of inventory investment hardly changes as we more from the business cycle band into higher frequencies. Thus, they explain very little of the high frequency variation observed in the data. Interestingly, these two models also performed poorly with respect to firm-level growth volatility in Table 4. What precisely the connection is remains to be explored.

As we move to the leading case of our model, model 5, where moderately persistent relative productivities are accompanied by fairly predicable firm-level inventory cycles, the fit to the data improves considerably. The high frequency volatility of inventory investment roughly doubles, while its correlation with GDP weakens considerably and lies near the data. Most notably, the strong correlation between sales and inventory investment is almost completed eliminated. Finally, when we add an aggregate shock to the final goods sector in the final column, the correlation turns slightly negative and nears its empirical counterpart.

We take the results in tables 4, 6 and 9 as grounds for hope that a carefully disciplined two-sector (S,s) model will ultimately succeed in simultaneously matching the inventory facts over the business cycle *and* at high frequencies. While our best candidate so far performs reasonably well in the business cycle band, it suffers three notable quantitative deficiencies in the high frequencies: its correlation between inventory investment and GDP is too weak, sales and GDP co-move too strongly, and the relative volatility of inventory investment is inadequate, at 45 percent that in the data. Aside from a thorough exploration of alternative combinations of idiosyncratic productivity shocks and order costs guided by additional moments from the firm-level data, a clear next step for this investigation is to replace our current independence assumption on the model's two aggregate shocks with a careful joint calibration targeting movements in the relative price of inventories at both business cycle frequencies and high frequencies.

Notwithstanding the quantitative flaws noted above, our leading model succeeds in reversing the positive correlation of sales and inventory investment (model: -0.07; data: -0.20), while maintaining positive co-movement of GDP with both series, and in yielding much higher relative volatility of inventory investment in the high frequencies in proportion to that in the business cycle frequencies (model: 1.79; data: 2.18). Thus, we permit ourselves a preliminary look at the extent to which high frequency movements in GDP and other key series are driven by inventory cycles. Table 10 compares our inventory model's results in the high frequency band to those from an otherwise identical control model lacking order costs and thus any inventory motive.

TABLE 10. High frequency role of inventories with two aggregate shocks

	GDP	NII	FS	C	I	N
% std. dev.						
control	**0.933**	−−	0.934	0.293	4.709	0.665
inventory model 5B	**0.921**	0.256	0.904	0.285	5.146	0.665
GDP corrn.						
control		−−	1.000	0.941	0.977	0.989
inventory model 5B		0.206	0.961	0.922	0.893	0.915

Interestingly, inventory investment does not drive increased GDP volatility even in the high frequencies. In fact, the standard deviation of GDP falls marginally when the frictions causing inventories are introduced. As was the case in the business cycle band, the volatility contributed directly by inventory investment is partly offset by accompanying reductions in the volatility of final sales, given the fundamental trade-off between storing the goods held in inventory and using them in production. Whereas the positive correlation between sales and inventory investment in the business cycle frequencies helped drive slightly higher GDP volatility in the presence of inventories, the mildly negative correlation between these series in the high frequencies delivers the reverse implication. The most interesting aspect of Table 10 is how the introduction of inventories alters the high-frequency behavior of capital investment, raising its volatility and weakening its co-movement with GDP. This will merit an under-the-hood exploration if it persists in future versions of the model with more carefully calibrated aggregate shocks.

7 Concluding remarks

One can never know from a given quarter's movements in inventories whether one is observing variation ultimately ascribable to business cycle frequencies or high frequencies. For this reason, and because inventory investment has substantial volatility within both frequency ranges, we have little hope of predicting what such movements imply for coming changes in employment and output without a coherent framework capable of reproducing a key set of empirical regularities involving this series both over the business cycle and at higher frequencies. In the pages above, we have sought to develop such a framework. We have seen that, when combined with plant-specific variation in productivity and relative certainty regarding the fixed costs that induce inventories, our (S,s) model is capable of generating the quite opposite co-movements observed between inventory investment, sales and GDP across these two frequency ranges, while maintaining consistency with what we know about firm-level variations in sales and output and the aggregate business cycle facts. As of now, our model's quantitative fit to these target regularities is imperfect. However, given the possible configurations of idiosyncratic productivity shocks and order costs yet to be explored, and pending careful joint calibration of the model's two aggregate shocks, we view the current results with optimism.

References

[1] Blinder, Alan S. (1981) 'Retail Inventory Behavior and Business Fluctuations,' *Brookings Papers on Economic Activity* 2: 443-505.

[2] Bloom, Nicholas (2009) 'The Impact of Uncertainty Shocks,' *Econometrica* **77**(3): 623–685

[3] Comin, D. and T. Philippon (2006) 'The Rise in Firm-Level Volatility: Causes and Consequences,' Chapter 3, pages 167 - 228, in Mark Gertler and Kenneth Roggoff (eds.) *NBER Macroeconomics Annual 2005*, Volume 20, MIT Press

[4] Cooper, Russell W. and John C. Haltiwanger (2006) 'On the Nature of Capital Adjustment Costs,' *Review of Economic Studies* **73**: 611-633.

[5] Hansen, Gary D. (1985) "Indivisible Labor and the Business Cycle,' *Journal of Monetary Economics* 16: 309-327.

[6] Khan, Aubhik and Julia K. Thomas (2007) 'Inventories and the business cycle: An equilibrium analysis of (S, s) policies,' *American Economic Review* **97**(4): 1165 – 88.

[7] Khan, Aubhik and Julia K. Thomas (2008) 'Idiosyncratic shocks and the role of nonconvexities in plant and aggregate investment dynamics,' *Econometrica.* **76**(2): 395–436.

[8] Krusell, Per and Anthony A. Smith Jr. (1997) 'Income and Wealth Heterogeneity, Portfolio Choice, and Equilibrium Asset Returns,' *Macroeconomic Dynamics* **1**: 387-422.

[9] Kryvtsov, Oleksiy and Virgiliu Midrigan (2013) 'Inventories, Markups, and Real Rigidities in Menu Cost Models,' *Review of Economic Studies* **80**(1): 249-76.

[10] Kydland, Finn E. and Edward C. Prescott (1982) 'Time to build and aggregate fluctuations,' *Econometrica* **50**(6): 1345-70.

[11] Lovell, Michael (1961) 'Manufacturers' Inventories, Sales Expectations, and the Acceleration Principle,' *Econometrica* **29**(3): 293–314.

[12] Metzler, Lloyd A. (1941) 'The Nature and Stability of Inventory Cycles,' *Review of Economics and Statistics* **23**(4): 113-29.

[13] Prescott, Edward C., Alexander Ueberfeldt and Simona E. Cociuba (2005): 'U.S. Hours and Productivity Behavior Using CPS Hours Worked Data: 1959-I to 2007-I,' Working Paper, University of Minnesota.

[14] Rogerson, Richard (1988) ‘Indivisible Labor, Lotteries and Equilibrium,’ *Journal of Monetary Economics* **21:** 3-16.

[15] 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**: 264-83.

[16] Wen, Yi (2005) ‘Understanding the Inventory Cycle,’ *Journal of Monetary Economics* **52***: 1533-55.*

[17] Wen, Yi (2011) ‘Input and Output Inventory Dynamics,’ *American Economic Journal: Macroeconomics* **3**(4): 1-33.