

Inventories and the business cycle:
An equilibrium analysis of (S,s) policies

Aubhik Khan¹

Federal Reserve Bank of Philadelphia

Julia K. Thomas

University of Minnesota

Federal Reserve Bank of Minneapolis

October 2003

¹We thank seminar participants at Carnegie-Mellon, The Federal Reserve Banks of Atlanta, Minneapolis and Philadelphia, Indiana, Iowa, Ohio State, Penn, UQAM and Yale, session participants at the 2002 Midwest Macro and SED meetings, the February 2003 NBER EF&G meetings, the June 2003 Madrid Conference on Lumpy Investment, and, in particular, Fernando Alvarez, V. V. Chari, Daniele Coen-Pirani, Ed Prescott and Valerie Ramey for their comments. The views expressed in this paper do not necessarily reflect those of the Federal Reserve Banks of Philadelphia or Minneapolis or the Federal Reserve System.

Abstract

We develop an equilibrium business cycle model in which the producers of final goods pursue generalized (S,s) inventory policies with respect to intermediate goods, a consequence of nonconvex factor adjustment costs. Calibrating our model to reproduce the average inventory-to-sales ratio in postwar U.S. data, we find that it explains over half of the cyclical variability of inventory investment. Moreover, inventory accumulation is strongly procyclical, and production is more volatile than sales, as in the data.

The comovement between inventory investment and final sales is often interpreted as evidence that inventories amplify aggregate fluctuations. In contrast, our model economy exhibits a business cycle similar to that of a comparable benchmark without inventories, though we do observe somewhat higher variability in employment, and lower variability in consumption and investment. Thus, equilibrium analysis, which necessarily endogenizes final sales, alters our understanding of the role of inventory accumulation for cyclical movements in GDP. The presence of inventories does not substantially raise the variability of production, because it dampens movements in final sales. Similarly, when reductions in adjustment costs lower, but do not eliminate, average inventory holdings, the variability of GDP is essentially unchanged, because the reduced costs cause an offsetting rise in the variability of final sales.

1 Introduction

Inventory investment is both procyclical and volatile. Changes in firms' inventory holdings appear to account for almost half of the decline in production during recessions.¹ Moreover, the comovement between inventory investment and final sales raises the variance of production above that of sales. Historically, such observations have often prompted researchers to emphasize inventory investment as central to an understanding of aggregate fluctuations.² Blinder (1990, page viii), for example, concludes that "business cycles are, to a surprisingly large degree, inventory cycles". By contrast, modern business cycle theory has been surprisingly silent on the topic of inventories.³

We derive inventory investment within a dynamic stochastic general equilibrium model. In particular, we extend the basic equilibrium business cycle model to include fixed costs associated with the acquisition of intermediate goods for use in final goods production. Given these costs, final goods firms optimally pursue generalized (S,s) policies; that is, they maintain inventories of intermediate goods, and they actively adjust these stocks only when they are sufficiently far from a target level. In our model, this target level varies endogenously with the aggregate state of the economy. Because adjustment costs differ across firms, in addition to productivity and capital, the aggregate state vector includes a distribution of producers over inventory levels.

Our objective is two-fold. First, we evaluate the ability of our equilibrium generalized (S, s) inventory model to reproduce salient empirical regularities. Specifically, we focus on the cyclical and variability of inventories, and the relative volatility of production and sales, as described below. Second, we examine the model's predictions for the role of inventories in aggregate fluctuations. This provides a formal analysis of the extent to which the existence of inventory investment amplifies or prolongs cyclical movements in production.

To assess the usefulness of our model in identifying the role of inventories in the business cycle, we evaluate its ability to reproduce (1) the volatility of inventory investment relative to production, (2) the procyclicality of inventory investment and (3) the greater volatility of production over that of sales. We view these three empirical regularities as essential characteristics of any formal analysis of the cyclical role of inventories. When we calibrate our equilibrium business cycle model of inventories to reproduce the average inventory-to-sales ratio in the postwar U.S. data, we find that it is able to

¹Ramey and West (1999) show that, on average, the decline in real inventory investment accounts for 49 percent of the decline in real gross domestic production during postwar U.S. recessions.

²See Blinder and Maccini (1991).

³When inventories are included in equilibrium models, their role is generally inconsistent with their definition. See, for example, Kydland and Prescott (1982) and Christiano (1988), where inventories are factors of production, or Kahn, McConnell and Perez-Quiros (2001), where they are a source of household utility.

explain roughly 54 percent of the measured cyclical variability of inventory investment. In addition, inventory investment is procyclical, and production is more volatile than sales, as consistent with the data. Moreover, our simulated model data exhibits persistence in the inventory-to-sales relationship consistent with empirical estimates. Beyond providing support for the model, this is of independent interest as it may help to explain the puzzlingly slow adjustment speeds found in empirical studies. We find that heterogeneity in the inventory levels held by nonadjusting firms breaks the linear mapping between the persistence of the inventory-sales relation and the economywide adjustment rate implied by the standard stock-adjustment equation.

Examining our model's predictions for the aggregate dynamics of output, consumption, investment and employment, we find that the business cycle with inventories is broadly similar to that generated by a comparable model without them. Nonetheless, the inventory model yields somewhat higher variability in employment, and lower variability in consumption and investment. Our central result is that the positive correlation between final sales and net inventory investment does not imply that inventories necessarily amplify aggregate fluctuations in production. In our equilibrium analysis, the dynamics of final sales are altered: the introduction of inventories does not substantially raise the variability of production because it lowers the variability of final sales. Similarly, when the fixed costs that cause inventories are raised to yield a substantial increase in the overall size of these stocks, the resulting rise in GDP variability is negligible. Again, this is because rises in fixed costs reduce the volatility of the endogenous final sales series enough to almost entirely offset the raised variability in inventory investment. Thus, beyond establishing the essentiality of equilibrium analysis, our findings also demonstrate the importance of focussing explicitly upon the economic fundamentals that cause inventories.

2 Empirical regularities and model selection

In this section, we discuss the set of empirical regularities concerning inventory investment that are most relevant to our analysis.⁴ Table 1 summarizes the business cycle behavior of GDP, final sales and changes in private nonfarm inventories in quarterly postwar U.S. data. Note first that the relative variability of inventory investment is large. In particular, though inventory investment's share of gross domestic production averages less than one-half of one percent, its standard deviation is 29.5 percent that of output.⁵ Next, net inventory investment is procyclical; its correlation coefficient with GDP is 0.67. Moreover, as the correlation between inventory investment and final sales is itself positive, 0.41 for the

⁴For more extensive surveys, see Fitzgerald (1997), Hornstein (1998) and Ramey and West (1999).

⁵Net investment in private nonfarm inventories is detrended as a share of GDP.

data summarized in table 1, the standard deviation of production substantially exceeds that of sales. It is this second positive correlation that is commonly interpreted as evidence that fluctuations in inventory investment increase the variability of GDP. For example, this, alongside supporting information from a bivariate VAR in inventories and final sales, leads Ramey and West (1999, page 874) to suggest that inventories “seem to amplify, rather than mute movements in production”. Our interest is in examining this thesis using quantitative general equilibrium analysis.

Inventories have received relatively little emphasis in general equilibrium models of aggregate fluctuations. Given positive real interest rates, the first challenge in any formal analysis of inventories is to explain their existence. In our model, they arise as a result of nonconvex order costs. To economize on such costs, firms choose to hold stocks and follow (S,s) policies in their management, adjusting only when they are sufficiently far from a target stock.

Within macroeconomics, by far the most common rationalization for inventory stocks has been the assumption that production is costly to adjust, and the associated costs are continuous functions of the change in production. This assumption underlies the traditional production smoothing model (and extensions that retain its linear-quadratic representative-firm structure). In its simplest form, the model assumes that final sales are an exogenous stochastic series, and that adjustments to the level of production incur convex costs. As a result, firms use inventories to smooth production in the face of fluctuations in sales.⁶ An apparent limitation of the model is that it applies to a narrow subset of inventories, finished manufacturing goods, which represents 13 percent of the total in table 2.⁷ Additionally, a number of researchers have suggested that this class of model has fared poorly in application to data. Blinder and Maccini (1991, page 85) summarize that it has been “distinctly disappointing, producing implausibly low adjustment speeds, little evidence that inventories buffer sales surprises, and a lack of sensitivity of inventory investment to changes in interest rates”. Blinder (1981) and Caplin (1985) conjecture that such weaknesses may have arisen from the model’s convex adjustment costs. In more recent work, Schuh (1996) estimates three modern variants of the model using firm-level data, and finds that each accounts for only a minor portion of the movements in firm-level inventories.

⁶A frequently noted difficulty with the original production smoothing model is its prediction that production is less variable than sales, and relatedly that sales and inventory investment are negatively correlated. These inconsistencies with the data have been addressed in several ways. For example, Ramey (1991) shows that they may be resolved if there are increasing returns to production, while Eichenbaum (1991) explores productivity shocks, and Coen-Pirani (2002) integrates the stockout avoidance motive of Kahn (1987) in a model of industry equilibrium.

⁷This interpretation of the model’s applicability is widespread, and is reinforced by the common empirical application to finished manufacturing goods alone. However, Ramey and West (1999) offer a counterargument suggesting that the model might be interpreted more broadly.

This he explains in part as the result of heterogeneity in the firm-level data that is necessarily omitted by the assumption of a representative firm.

Given the extensive body of research already devoted to the production smoothing model, we instead base our analysis on the leading microeconomic model of inventories, the (S,s) model originally solved by Scarf (1960). First, we view the (S,s) model as applying to a wide group of inventories. As Blinder and Maccini (1991) have argued, the decisions facing manufacturers purchasing inputs for production and wholesalers and retailers purchasing goods from manufacturers are similar in that they each involve decisions as to when and in what quantity orders should be undertaken from other firms. If there are fixed costs associated with moving items from firm to firm, then efforts to avoid such costs may explain why stocks of manufacturing inputs, as well as those of finished goods in retail and wholesale trade, are held. Next, there is empirical support for the (S,s) approach. Mosser (1991) tests a simple fixed-band (S,s) model on aggregate retail trade data and reports that it is more successful in explaining the observed time series than is the traditional linear quadratic model. More recently, McCarthy and Zakrajšek (2000) have isolated nonlinearities indicative of (S,s) inventory policies in firm-level inventory adjustment functions in manufacturing, and Hall and Rust (1999) have shown that a generalized (S,s) decision rule can explain the actual inventory investment behavior of a U.S. steel wholesaler.

The aggregate implications of the (S,s) inventory model have been largely unexplored; in fact, thusfar there has been no quantitative general equilibrium analysis of this environment. The only equilibrium study we know of is that by Fisher and Hornstein (2000), who focus on explaining the greater volatility of orders relative to sales in a model of retail inventories without capital. Building on the work of Caplin (1985) and Caballero and Engel (1991), who study the aggregate implications of exogenous (S,s) policies across firms, Fisher and Hornstein construct an environment that endogenously yields time-invariant one-sided (S,s) rules and a constant order size per adjusting firm.⁸ This allows them to tractably study (S,s) inventory policies in general equilibrium without confronting substantial heterogeneity across firms.

In our model, as in the generalized (S,s) investment model of Caballero and Engel (1999), there are three mechanisms that drive changes in the aggregate stock of inventories. First, there are movements in the intensive margin; that is, changes in the order sizes of firms engaged in inventory investment. Second, there are changes in the fractions of firms that actually place orders from each given level of inventories; in other words, shifts in a nontrivial adjustment hazard that produce extensive margin movements. Third, there is time-variation in the distribution of firms over inventory holdings; changes

⁸Specifically, they assume indivisible retail goods, one unit sold per successful retailer per period, and small aggregate shocks. Together, these assumptions imply that retailers place orders only when their stocks are fully exhausted, and that the common target inventory level to which they then adjust never varies.

in this distribution interact with the adjustment hazard to induce further fluctuations along the extensive margin. The assumptions made by Fisher and Hornstein (2000) permit only the third of these three mechanisms, which suggests that their analysis may have omitted important channels through which changes in firms' inventory decisions affect the aggregate economy. More broadly, our analysis is distinguished from theirs by our inclusion of capital. As we have noted, inventory models have had difficulty reproducing procyclical inventory investment. Fisher and Hornstein find that inventory investment is procyclical in their model, but only in general equilibrium. This suggests that the absence of capital accumulation may be important to their result, since inventory accumulation is the only mechanism for consumption smoothing in their model. Finally, our analysis is quantitative; our purpose is to examine the extent to which inventory investment alters aggregate fluctuations.

A further distinguishing feature of our model is that it does not focus exclusively on finished goods inventories. Both Blinder and Maccini (1991) and Ramey and West (1999) have emphasized that inventories of finished manufacturing goods have seen disproportionate attention in theoretical and empirical work relative to other, more cyclically important, components of private nonfarm inventories. Manufacturing inputs, the sum of materials and supplies and work-in-process, are a particularly notable omission, as first stressed by Ramey (1989). Table 2 shows that manufacturing inventories are far more cyclical than retail and wholesale inventories, the other main components of private nonfarm inventories. It also shows that, within manufacturing, inventories of intermediate inputs are twice the size of finished goods. Moreover, the results of a variance decomposition undertaken by Humphreys, Maccini and Schuh (2001) indicate that intermediate inputs in manufacturing are three times more volatile than finished goods. Given the primary cyclical role of manufacturing input inventories, we develop a model that includes these stocks. However, we do not limit our analysis to manufacturing inputs. In particular, we do not identify our intermediate goods, or our firms, as belonging to a specific sector. Rather, our inventories are stocks that broadly represent goods held in various stages of completion throughout the economy. Consequently, we calibrate the relative magnitude of inventories in our model to match that of total private nonfarm inventories.

3 Model

There are three sets of agents in the economy, households, intermediate goods producers and final goods firms. Households supply labor to both types of producers and purchase consumption goods from final goods firms. They save through asset markets where they trade shares that entitle them to the earnings of both intermediate and final goods producers. All firms in the economy are perfectly

competitive. First, identical intermediate goods producers own capital and hire labor for production. They sell their output to, and purchase investment goods from, final goods producers. Next, final goods firms use intermediate goods and labor to produce output that may be used for consumption or capital accumulation.

We derive inventories explicitly in our model by assuming that final goods firms face fixed costs of ordering or accepting deliveries of intermediate goods. As the costs are independent of order size, these firms choose to hold stocks of intermediate goods, s , where $s \in \mathbf{S} \subseteq R_+$. Further, the costs vary across final goods firms, so some will adjust their inventory holdings, while others will not, at any date. As a result, the model yields an endogenous distribution of final goods firms over inventory levels, $\mu : \mathcal{B}(\mathbf{S}) \rightarrow [0, 1]$, where $\mu(S)$ represents the measure of firms with start-of-period inventories in the set $S \in \mathcal{B}(\mathbf{S})$.

The economy's aggregate state is (z, A) , where $A \equiv (K, \mu)$ represents the endogenous state vector. K is the aggregate capital stock held by intermediate goods firms, and z is total factor productivity in the production of intermediate goods.⁹ The distribution of final goods firms over inventory levels evolves according to a mapping Γ_μ , $\mu' = \Gamma_\mu(z, A)$, and capital similarly evolves according to $K' = \Gamma_K(z, A)$.¹⁰ We assume that productivity follows a Markov Chain, $z \in \{z_1, \dots, z_{N_z}\}$, where

$$\Pr(z' = z_j \mid z = z_i) \equiv \pi_{ij} \geq 0, \quad (1)$$

and $\sum_{j=1}^{N_z} \pi_{ij} = 1$ for each $i = 1, \dots, N_z$. Except where necessary for clarity, we suppress the index for current productivity below.

All producers employ labor at the real wage, $\omega(z, A)$, and those involved in the production of final goods purchase intermediate goods at the relative price $q(z, A)$. Finally, all firms, whether producing intermediate or final goods, value current profits by the final output price $p(z, A)$ and discount future earnings by β .¹¹ For brevity, we suppress the arguments of ω , q and p where possible below.

3.1 Intermediate goods producers

The representative intermediate goods producer uses capital, k , and labor, l , in a constant returns to scale technology, $zF(k, l)$ to produce intermediate goods. These are sold to final goods firms at the

⁹This is the sole source of aggregate fluctuations in the model. Its placement in the production of intermediate goods allows consistency with the countercyclical relative price of inventories in the aggregate data, as described in section 6.2.

¹⁰Throughout the paper, primes indicate one-period ahead values. We define Γ_μ in section 3.2.3, following the description of firms' problems, and Γ_K in section 3.4. Below, we summarize the aggregate law of motion as $A' = \Gamma(z, A)$.

¹¹This is equivalent to requiring that firms discount by $1 + r_{t,t+k} = \frac{p_t}{\beta^k p_{t+k}}$ between the states in t and $t+k$, where p represents households' current valuation of output and β is their subjective discount factor. This discounting rule is an implication of equilibrium, as discussed in section 3.4.

relative price q . The producer may adjust next period's capital stock using final goods as investment. Capital depreciates at the rate $\delta \in (0, 1)$. Equation 2 below is the functional equation describing the intermediate goods producer's problem. The value function W is a function of the aggregate state (z, A) , which determines the prices p, q and ω .

$$W(k; z, A) = \max_{k', l} \left(p \left[qzF(k, l) + (1 - \delta)k - k' - \omega l \right] + \beta \sum_{j=1}^{N_z} \pi_{ij} W(k'; z_j, A') \right) \quad (2)$$

The producer takes as given that A evolves over time according to $A' = \Gamma(z, A)$, and changes in productivity follow the law of motion described in (1). The following efficiency conditions describe its selection of employment and investment.

$$zD_2F(k, l) = \frac{\omega}{q} \quad (3)$$

$$\beta \sum_{j=1}^{N_z} \pi_{ij} D_1 W(k'; z_j, A') = p \quad (4)$$

Because F is linearly homogenous, the producer's decision rules for employment and production are proportional to its capital stock; $l(k) \equiv L(z, A)k$, where $L(z, A)$ solves (3) as a function of z, ω and q , and $x(k; z, A) = zF(1, L(z, A))k$. This means that current profits, $\pi(z, A)k$, are linear in k , as is the value function; $W(k; z, A) = w(z; A)k$, where

$$w(z, A) \cdot k = \max_{k'} p(z, A) \left[\pi(z, A)k - k' \right] + \beta \sum_{j=1}^{N_z} \pi_{ij} w(z_j, A') k'.$$

Equation 4 then implies that an interior choice of investment places the following restriction on the equilibrium price of final output.

$$p(z, A) = \beta \sum_{j=1}^{N_z} \pi_{ij} w(z_j, A') \quad (5)$$

When (5) is satisfied, the intermediate goods firm is indifferent to any level of k' and will purchase investment equal to the final goods remaining after households' consumption.

3.2 Final goods producers

There are a large number of final goods firms, each facing time-varying costs of arranging deliveries or sales of intermediate goods. Given differences in delivery costs, some firms adjust their stocks, while others do not, at any date. Thus, firms are distinguished by their inventories of intermediate goods.

At the start of any date, a final goods firm is identified by its inventory holdings, s , and its current delivery cost, $\xi \in [\underline{\xi}, \bar{\xi}]$. This cost is denominated in hours of labor and drawn from a time-invariant

distribution $H(\xi)$ common across firms. Intermediate goods used in the current period, m , and labor, n , are the sole factors of final goods production, $y = G(m, n)$, where G exhibits decreasing returns to scale. Note that technology is common across these firms; the only source of heterogeneity in production arises from differences in inventories.

The timing of final goods firms' decisions is as follows. At the beginning of each period, any such firm observes the aggregate state (z, A) and its current delivery cost ξ . Before production, it undertakes an inventory adjustment decision. In particular, the firm can absorb its fixed cost and adjust its stock of intermediate goods available for production, $s_1 \geq 0$.¹² Letting x_m denote the chosen size of such an adjustment, the stock available for current production becomes $s_1 = s + x_m$. Alternatively, the firm can avoid the cost, set $x_m = 0$, and enter production with its initial stock; $s_1 = s$. Following the inventory adjustment decision, the firm determines current production, selecting $m \in [0, s_1]$ and $n \in R_+$. Intermediate goods fully depreciate in use, and the remaining stock with which the firm begins the next period is denoted s' . Measuring adjustment costs in units of final output using the wage rate, ω , the firm's order choice is summarized below.

Table 3

order size	total order costs	production-time stock	next-period stock
$x_m \neq 0$	$\omega\xi + qx_m$	$s_1 = s + x_m$	$s' = s_1 - m$
$x_m = 0$	0	$s_1 = s$	$s' = s_1 - m$

Finally, inventories incur storage costs that are proportional to the level of inventories held. Given end of period inventories s' , a firm's total cost of storage is $\sigma s'$ where $\sigma > 0$ is a parameter capturing the unit cost of holding inventories.

Let $V^0(s, \xi; z, A)$ represent the expected discounted value of a final goods firm with start-of-date inventory holdings s and fixed order cost ξ . We describe the problem facing such a firm using (6) - (9) below. First, for convenience, we define the beginning of period expected value of the firm, prior to the realization of its fixed cost, but given $(s; z, A)$.

$$V(s; z, A) \equiv \int_{\underline{\xi}}^{\bar{\xi}} V^0(s, \xi; z, A) H(d\xi) \quad (6)$$

¹²As the distinction between s and s_1 indicates, we avoid assuming that the stock of intermediate goods available for current production must be determined a period in advance. This is consistent with our quarterly calibration of the model, which is dictated by the frequency of aggregate data.

Next, we divide the period into two sub-periods, an adjustment sub-period and a production sub-period, and we break the description of the firm's problem into the distinct problems it faces as it enters into each of these sub-periods.¹³

3.2.1 Production decisions

Beginning with the second sub-period, let $V^1(s_1; z, A)$ represent the value of entering production with inventories s_1 . Given this stock available for production, the firm selects its current employment, n , and inventories for next period, s' , (hence the amount of its stock to use in current production, $m = s_1 - s'$) to solve

$$V^1(s_1; z, A) = \max_{s' \geq 0, n \geq 0} \left(p \left[G(s_1 - s', n) - \omega n - \sigma s' \right] + \beta \sum_{j=1}^{N_z} \pi_{ij} V(s'; z_j, A') \right), \quad (7)$$

taking prices (p , ω and q) and the evolution of A' as given. Given the production-time stock of intermediate goods, s_1 , and the continuation value of inventories, $V(s'; z_j, A')$, equation (7) yields both the firm's employment (in production) decision and its use of intermediate goods. Let $N(s_1; z, A)$ describe its employment and $S(s_1; z, A)$ its stock of intermediate goods retained for future use. Its current production of final goods is then $Y(s_1; z, A) = G(s_1 - S(s_1; z, A), N(s_1; z, A))$. Thus, we have decision rules for employment, production, and next-period inventories as functions of the production-time stock s_1 .

3.2.2 Inventory adjustment decisions

Given the middle-of-period valuation of the firm, V^1 , we now examine the inventory adjustment decision. At the beginning of the period, consider the problem of a final goods firm with inventories s and adjustment cost ξ . Equations (8) - (9) describes the (s, ξ) firm's determination of (i) whether to place an order and (ii) the target inventory level with which to begin the production sub-period, conditional on an order. The first term in the braces of (8) represents the net value of stock adjustment, (the gross adjustment value less the value of the payments associated with the fixed delivery cost,) while the second term represents the value of entering production with the beginning of period stock.

$$V^0(s, \xi; z, A) = pqs + \max \left\{ -p\omega\xi + V^a(z, A), -pqs + V^1(s; z, A) \right\} \quad (8)$$

$$V^a(z, A) \equiv \max_{s_1 \geq 0} \left(-pqs_1 + V^1(s_1; z, A) \right) \quad (9)$$

¹³This division of the period is for expositional convenience only; no uncertainty is resolved between the two sub-periods.

Note that the target inventory choice in (9) is independent of both the current inventory level, s , and fixed cost, ξ . Thus, all firms that adjust their inventory holdings choose the same production-time level and achieve the same gross value of adjustment, $V^a(z, A)$. Let $s^* \equiv s^*(z, A)$ denote the common target that solves (9) as a function of the aggregate state of the economy. Equation (7) then implies common employment and intermediate goods use choices across all adjusting firms, as well as identical inventory holdings among these firms at the beginning of the next period.

Turning to the decision of whether to adjust to the target level of inventories, it is immediate from equation (8) that a firm will place an order if its fixed cost falls at or below $\tilde{\xi}(s; z, A)$, the cost that equates the net value of inventory adjustment to the value of non-adjustment.

$$-p\omega\tilde{\xi}(s; z, A) + V^a(z, A) = -pqs + V^1(s; z, A) \quad (10)$$

Given the support of the cost distribution, and using (10) above, we define $\xi^T(s; z, A)$ as the type-specific threshold cost separating those firms that place orders from those that do not.

$$\xi^T(s; z, A) = \min\left\{\max\left(\underline{\xi}, \tilde{\xi}(s; z, A)\right), \bar{\xi}\right\} \quad (11)$$

Thus, we arrive at the following decision rules for production-time inventory holdings and stock adjustments.

$$s_1(s, \xi; z, A) = \begin{cases} s^*(z, A) & \text{if } \xi \leq \xi^T(s; z, A) \\ s & \text{if } \xi > \xi^T(s; z, A) \end{cases} \quad (12)$$

$$x_m(s, \xi; z, A) = s_1(s, \xi; z, A) - s \quad (13)$$

The common distribution of adjustment costs facing final goods firms, given their threshold adjustment costs, implies that $H(\xi^T(s; z, A))$ is the probability that a firm of type s will alter its inventory stock before production. Using this result, the start-of-period value of the firm prior to the realization of its fixed delivery cost, (6), may be simplified as

$$V(s; z, A) = pqs + H(\xi^T(s; z, A)) V^a(z, A) - p\omega \int_{\underline{\xi}}^{\xi^T(s; z, A)} \xi H(d\xi) \quad (14)$$

$$+ \left(1 - H(\xi^T(s; z, A))\right) \left(V^1(s; z, A) - pqs\right),$$

where $\int_{\underline{\xi}}^{\xi^T(s; z, A)} \xi H(d\xi)$ is the conditional expectation of the fixed cost ξ .

3.2.3 Aggregation

Having described the inventory adjustment and production decisions of final goods firms as functions of their type, s , and cost draw, ξ , we can now aggregate their demand for the production of intermediate goods firms, their demand for labor, their use of intermediate goods, and their production of the final good. First, the aggregate demand for intermediate goods is the sum of the stock adjustments from each start-of-period inventory level s , weighted by the measures of firms undertaking these adjustments.

$$\bar{X}(z, A) = \int_{\mathbf{S}} H(\xi^T(s; z, A)) (s^*(z, A) - s) \mu(ds) \quad (15)$$

Second, the total usage of these intermediate goods, $\bar{M}(z, A)$, is the total production-time stock less that which remains at the end of the period, held as inventories for the subsequent date.¹⁴

$$\bar{M}(z, A) \equiv \int_{\mathbf{S}} \left[\int_{\underline{\xi}}^{\bar{\xi}} \left(s_1(s, \xi; z, A) - S(s_1(s, \xi; z, A); z, A) \right) H(d\xi) \right] \mu(ds)$$

Next, the production of final goods is the population-weighted sum of production across adjusting and non-adjusting firms.

$$\begin{aligned} \bar{Y}(z, A) &= Y(s^*(z, A); z, A) \int_{\mathbf{S}} H(\xi^T(s; z, A)) \mu(ds) + \\ &\quad \int_{\mathbf{S}} Y(s; z, A) [1 - H(\xi^T(s; z, A))] \mu(ds) \end{aligned} \quad (16)$$

Finally, employment demand by final goods firms is the weighted sum of labor employed in production by adjusting and non-adjusting firms together with the total time costs of adjustment.

$$\begin{aligned} \bar{N}(z, A) &= N(s^*(z, A); z, A) \int_{\mathbf{S}} H(\xi^T(s; z, A)) \mu(ds) \\ &+ \int_{\mathbf{S}} [1 - H(\xi^T(s; z, A))] N(s; z, A) \mu(ds) + \int_{\mathbf{S}} \left[\int_{\underline{\xi}}^{\xi^T(s; z, A)} \xi H(d\xi) \right] \mu(ds) \end{aligned} \quad (17)$$

We next examine Γ_μ , the evolution of the distribution of final goods firms using (10) - (11). Of each group of firms sharing a common stock $s \neq s^*$ at the start of the current period, fraction $1 - H(\xi^T(s; z, A))$

¹⁴This may be equivalently expressed as the population-weighted sum of the usage of intermediate goods across adjusting and non-adjusting firms:

$$\begin{aligned} \bar{M}(z, A) &= [s^*(z, A) - S(s^*(z, A); z, A)] \int_{\mathbf{S}} H(\xi^T(s; z, A)) \mu(ds) \\ &\quad + \int_{\mathbf{S}} [s - S(s; z, A)] [1 - H(\xi^T(s; z, A))] \mu(ds). \end{aligned}$$

do not adjust their inventories. Thus, with some abuse of notation, $\mu(s)[1 - H(\xi^T(s; z, A))]$ firms will begin the next period with $S(s; z, A)$ as defined in section 3.2.1. Those firms that either enter the period with the current target or actively adjust to it for production, $\mu(s^*(z, A)) + \int_{\mathbf{S}} H(\xi^T(s; z, A))\mu(ds)$ in all, will move to the next period with $S(s^*(z, A); z, A)$.

Given the preceding discussion, the evolution of the distribution of final goods firms may be described as follows. Define $S^{-1}(\tilde{s}; z, A)$ as the production-time inventory level that gives rise to next period inventories \tilde{s} in the solution to (7). For any stock \tilde{s} other than that arising from the target level of production-time inventories, $S^{-1}(\tilde{s}; z, A) \neq s^*(z, A)$,

$$\mu'(\tilde{s}) = \left[1 - H\left(\xi^T\left(S^{-1}(\tilde{s}; z, A)\right)\right)\right]\mu\left(S^{-1}(\tilde{s}; z, A)\right). \quad (18)$$

For the stock arising from the target inventory level, $S^{-1}(\tilde{s}; z, A) = s^*(z, A)$,

$$\mu'(\tilde{s}) = \mu\left(s^*(z, A)\right) + \int_{\mathbf{S}} H\left(\xi^T(s; z, A)\right)\mu(ds). \quad (19)$$

3.3 Households

The economy is populated by a unit measure of identical households who value consumption and leisure and discount future utility by $\beta \in (0, 1)$. Households have fixed time endowments in each period, normalized to 1, and they receive real wage $\omega(z, A)$ for their labor. Their wealth is held as one-period shares in final goods firms, denoted by the measure λ_F , and as shares in the unit measure of identical intermediate goods firms, λ_I .

At each date, households must determine their current consumption, C , hours worked, N , as well as what new shares in final goods firms, λ'_F , and intermediate goods firms, λ'_I , to purchase at prices $\rho_F(s; z, A)$ and $\rho_I(z, A)$ respectively.¹⁵ Their expected lifetime utility maximization problem is described recursively below.

$$R(\lambda_I, \lambda_F; z, A) = \max_{C, N, \lambda'_I, \lambda'_F} \left(U(C, 1 - N) + \beta \sum_{j=1}^{N_z} \pi_{ij} R(\lambda'_I, \lambda'_F; z_j, A') \right) \quad (20)$$

subject to

$$\begin{aligned} C + \rho_I(z, A)\lambda'_I + \int_{\mathbf{S}} \rho_F(s; z, A)\lambda'_F(ds) \\ \leq \omega(z, A)N + \rho_I(z, A)\lambda_I + \int_{\mathbf{S}} \rho_F(s; z, A)\lambda_F(ds) \end{aligned} \quad (21)$$

$$A' = \Gamma(z, A) \quad (22)$$

¹⁵In equilibrium, these prices are $\frac{V(s; z, A)}{p(z, A)}$ and $\frac{W(K; z, A)}{p(z, A)}$.

Let $C(\lambda_I, \lambda_F; z, A)$ summarize their choice of current consumption, $N(\lambda_I, \lambda_F; z, A)$ their allocation of time to work, $\Lambda_I(\lambda_I, \lambda_F; z, A)$ their purchases of shares in the representative intermediate goods firm, and $\Lambda_F(s, \lambda_I, \lambda_F; z, A)$ the quantity of shares they purchase in final goods firms that will begin next period with inventories s .

3.4 Equilibrium

In equilibrium, households will hold a portfolio of all firms, ($\Lambda_I(1, \mu; z, A) = 1$ and $\Lambda_F(s, 1, \mu; z, A) = \mu'(s)$), and will supply a level of labor consistent with employment across these firms, at each date. Consequently, the real wage must equal households' marginal rate of substitution between leisure and consumption,

$$\omega(z, A) = \frac{D_2 U\left(C(1, \mu; z, A), 1 - N(1, \mu; z, A)\right)}{D_1 U\left(C(1, \mu; z, A), 1 - N(1, \mu; z, A)\right)}, \quad (23)$$

and all firms must discount future profit flows with state-contingent discount factors that are consistent with households' marginal rate of intertemporal substitution, $\frac{\beta D_1 U\left(C(1, \mu'; z', A'), 1 - N(1, \mu'; z', A')\right)}{D_1 U\left(C(1, \mu; z, A), 1 - N(1, \mu; z, A)\right)}$. Following the approach outlined in Khan and Thomas (2003), we have already imposed the latter restriction in describing firms' problems above. Specifically, we have assumed that all firms value current profit flows at the final output price $p(z, A)$, which represents the household marginal utility of equilibrium consumption, and that firms discount their future values by the subjective discount factor β .

$$p(z, A) = D_1 U\left(C(1, \mu; z, A), 1 - N(1, \mu; z, A)\right) \quad (24)$$

When p and ω are evaluated at the equilibrium values of consumption and total work hours, we are able to recover all equilibrium decision rules by solving firms' problems alone.

Because there is no heterogeneity in intermediate goods production, in equilibrium, $K = k$ at each date. Thus, the evolution of the aggregate capital stock, summarized above by $K' = \Gamma_K(z, A)$, is defined as $\Gamma_K(z, A) \equiv (1 - \delta)K + \bar{Y}(z, A) - C(1, \mu; z, A)$, where $\bar{Y}(z, A)$ is given by (16). Next, the aggregate demand for intermediate goods by final goods firms adjusting their holdings of inventories must equal the production of these inputs, and household labor supplied must fulfill total employment demand across intermediate and final goods firms;

$$\bar{X}(z, A) = x(K; z, A) \quad \text{and} \quad N(1, \mu; z, A) = L(z, A)K + \bar{N}(z, A).$$

Finally, it is convenient to describe equilibrium inventory investment in terms of total use and production of intermediate goods. Aggregate inventory investment is defined as the change in total inventories,

weighted by the relative price of the intermediate good. In equilibrium, this is the q -weighted difference between the supply and total use of intermediate goods, $q(z, A)\left(x(K; z, A) - \overline{M}(z, A)\right)$.

4 Parameter choices

We examine the implications of inventory accumulation for an otherwise standard equilibrium business cycle model using numerical methods. In calibrating our model, we choose the length of a period as one quarter and select functional forms for production and utility as follows. We assume that intermediate goods producers have a Cobb-Douglas production function with capital share α , and that their productivity follows a Markov Chain with two values, $N_z = 2$, that is itself the result of discretizing an estimated log-normal process for technology with persistence ρ and variance of innovations, σ_ε^2 . Final goods firms also have Cobb-Douglas technology, with intermediate goods' share θ_m , $G(m, n) = m^{\theta_m}n^{\theta_n}$. The adjustment costs that provide the basis for inventory holdings in our model are assumed to be distributed uniformly with lower support 0 and upper support $\overline{\xi}$. Finally, we assume that households' period utility is the result of indivisible labor decisions implemented with lotteries (Rogerson (1988), Hansen (1985)), $u(C, 1 - N) = \log C + \eta \cdot (1 - N)$.

4.1 Benchmark model

If we set $\overline{\xi} = 0$, the result is a model where no firm has an incentive to hold inventories. With no adjustment costs, final goods firms buy intermediate goods in every period; hence there are two representative firms, an intermediate goods firm and a final goods firm. We take this model as a benchmark against which to evaluate the effect of introducing inventory accumulation. The parameterization of the benchmark and inventory models is identical, with the already noted exception of the cost distribution associated with adjustments to intermediate goods holdings.

The parameters that are common to both the benchmark and inventory models, $(\alpha, \theta_m, \theta_n, \delta, \beta, \eta)$, are derived, wherever possible, from standard values. The parameter associated with capital's share, α , is chosen to reproduce a long-run annual nonfarm business capital-to-output ratio of 1.415, a value derived from U.S. data between 1953 – 2002. The depreciation rate δ is equal to the average ratio of investment to business capital over the same time period. The distinguishing feature of the benchmark model, relative to the Indivisible Labor Economy of Hansen (1985), is the presence of intermediate goods. The single new parameter implied by the additional factor of production, the share term for intermediate goods, is selected to match the value implied by the updated Jorgenson, Gollop and Fraumeni (1999) input-output data from manufacturing and trade. From this data set, we obtain an annual weighted

average of materials' share across 21 2-digit manufacturing sectors and the trade sector, averaged over 1958-1996, at 0.499.¹⁶ The remaining production parameter, θ_n , is taken to imply a labor's share of output averaging 0.64, as in Hansen (1985) and Prescott (1986). Turning to preferences, the subjective discount factor, β , is selected to yield a real interest rate of 6.5 percent per year in the steady state of the model, and η is chosen so that average hours worked are $\frac{1}{3}$ of available time.

We determine the stochastic process for productivity using the Crucini Residual approach described in King and Rebelo (1999). A continuous shock version of the benchmark model, where $\log z_{t+1} = \rho \log z_t + \varepsilon_{t+1}$ with $\varepsilon_{t+1} \sim N(0, \sigma_\varepsilon^2)$, is solved using an approximating system of stochastic linear difference equations, given an arbitrary initial value of ρ . This linear method yields a decision rule for output of the form $Y_t = \pi_z(\rho) z_t + \pi_k(\rho) k_t$, where the coefficients associated with z and k are functions of ρ . Rearranging this solution, data on GDP and capital are then used to infer an implied set of values for the technology shock series z_t . Maintaining the assumption that these realizations are generated by a first-order autoregressive process, the persistence and variance of this implied technology shock series yields new estimates of $(\rho, \sigma_\varepsilon^2)$. The process is repeated until these estimates converge. The resulting values for the persistence and variance of the technology shock process are not uncommon.

4.2 Inventory model

Table 4 lists the baseline calibration of our inventory model. For all parameters that are also present in the benchmark model, we maintain the same values as there. This approach to calibrating the inventory model is feasible, as the steady states of the two model economies, in terms of the capital-output ratio, hours worked, and the shares of the three factors of production, are close.

The two parameters that distinguish the inventory model from the benchmark are the storage cost associated with inventories and the upper support for adjustment costs (uniformly distributed on $[0, \bar{\xi}]$). Conventional estimates of inventory storage costs (or *carrying costs*) average 25 percent of the annual value of inventories held (Stock and Lambert (1987)). Excluding those components accounted for elsewhere in our model (for instance, the cost of money reflected by discounting) and those associated with government (taxes), we calibrate σ to yield storage costs at 12 percent of the

¹⁶For each year, we obtain sector-specific values of materials' share by computing the ratio of the value of materials relative to the (producer price) value of output for each sector. Next, each sector's θ_m is weighted by the value of its output relative to the total, and the results summed to yield the year's average θ_m across sectors. The resulting average over 1958-1996 is remarkably close to the annual average annual value of materials' costs, excluding energy, in the NBER-CES Manufacturing Database of 4-digit SIC manufacturing industries compiled by Bartlesman, Becker and Gray (Bartlesman and Gray 1996) for the years 1958-1997, which is 0.50.

annual value of inventories.¹⁷ In our calibrated model, where the steady-state value of q is 0.417, this implies a proportional cost of $\sigma = 0.012$. Next, using NIPA data, we compute that the quarterly real private nonfarm inventory-to-sales ratio has averaged 0.7155 in the U.S. between 1947:1 and 2002:1.¹⁸ Given the storage cost parameter σ , we select the upper support on adjustment costs, $\bar{\xi}$, at 0.220 to reproduce this average inventory-to-sales ratio in our model.

5 Numerical method

The (S, s) inventory model developed above is characterized by an aggregate state vector that includes the distribution of the stock of inventory holdings across firms, which makes computation of equilibrium nontrivial. Our solution algorithm involves repeated application of the contraction mapping implied by (6), (7), (8) and (9) to solve for final goods firms' start-of-period value functions V , given the price functions $p(z, A)$, $\omega(z, A)$ and $q(z, A)$ and the laws of motion implied by Γ and (π_{ij}) . This recursive approach is complicated in two ways, as discussed below.

First, the nonconvex factor adjustment here requires that we solve for firms' decision rules using nonlinear methods. This is because firms at times find themselves with a very low stock of intermediate goods relative to their production-time target, but draw a sufficiently high adjustment cost that they are unwilling to replenish their stock in the current period. At such times, they will exhaust their entire stock in production, deferring adjustment until the beginning of the next period, before further production. Thus, a non-negativity constraint on inventory holdings occasionally binds, and firms' decision rules are nonlinear and must be solved as such. This we accomplish using multivariate piecewise polynomial splines, adapting an algorithm outlined in Johnson (1987). In particular, our splines are generated as the tensor product of univariate cubic splines, with one of these corresponding to each argument of the value function.¹⁹ We apply spline approximation to V , using a multi-dimensional grid on the state vector for these functions.

Second, equilibrium prices are functions of a large state vector, given the presence of the distribution of final goods firms in the endogenous aggregate state vector, $A = (K, \mu)$. For computational feasibility, we assume that agents use a smaller object to proxy for the distribution in forecasting the future state and thereby determining their decisions rules given current prices. In choosing this proxy, we extend the

¹⁷Excluded components are: cost of money, taxes, physical handling and clerical and inventory control. The latter components are already reflected in our model by the presence of labor-denominated adjustment costs.

¹⁸This value lies just above the Ramey and West (1999) average for G7 countries of 0.66. Moreover, as noted by these authors, the real series, in contrast to its nominal counterpart, exhibits no trend.

¹⁹For additional details, see Khan and Thomas (2003).

method applied in Khan and Thomas (2003), which itself applied a variation on the method of Krusell and Smith (1998). In particular, we approximate the distribution in the aggregate state vector with a vector of moments, $m = (m_1, \dots, m_I)$, drawn from the distribution. In our work involving discrete heterogeneity in production, we find that sectioning the distribution into I equal-sized partitions and using the conditional mean of each partition is efficient in that it implies small forecasting errors.

The solution algorithm is iterative, applying one set of forecasting rules to generate decision rules that are used in obtaining data upon which to base the next set of forecasting rules. In particular, given I , we assume functional forms that predict next period's endogenous state (K', m') , and the prices p and pq , as functions of the current state, $K' = \widehat{\Gamma}_K(z, K, m; \chi_l^K)$, $m' = \widehat{\Gamma}_m(z, K, m; \chi_l^m)$, $p = \widehat{p}(z, K, m; \chi_l^p)$ and $pq = \widehat{pq}(z, K, m; \chi_l^{pq})$, where χ_l^K , χ_l^m , χ_l^p , and χ_l^{pq} are parameter vectors that are determined iteratively, with l indexing these iterations. For the class of utility functions we use, the wage is immediate once p is specified; hence there is no need to assume a wage forecasting function.

For any I , $\widehat{\Gamma}_K$, $\widehat{\Gamma}_m$, \widehat{p} , and \widehat{pq} , we solve for V on a grid of values for $(s; z, K, m)$. Next, we simulate the economy for T periods, recording the actual distribution of final goods firms, μ_t , at the start of each period, $t = 1, \dots, T$. To determine equilibrium in each date, we begin by calculating m_t using the actual distribution, μ_t , and then we use $\widehat{\Gamma}_K$ and $\widehat{\Gamma}_m$ to specify expectations of K_{t+1} and m_{t+1} . This determines $\beta \sum_{j=1}^{N_z} \pi_{ij} w(z_j, K_{t+1}, m_{t+1})$, and $\beta \sum_{j=1}^{N_z} \pi_{ij} V(s'; z_j, K_{t+1}, m_{t+1})$ for any s' . Given the second function, the conditional expected continuation value associated with any level of inventories, we can determine $s^*(z, K, m)$ and $\xi^T(s; K, m)$, hence recovering the decisions of final goods firms and thus next period's distribution, for any values of p and q . Given any p , the equilibrium q is solved to equate the intermediate goods producer's supply, $x(K; z, A)$, to the demand generated by final goods firms.²⁰ The equilibrium output price, $p(z, A; \chi_l^K, \chi_l^m, \chi_l^p, \chi_l^{pq})$, is that which generates production of the final good such that, given $C = \frac{1}{p}$, the residual level of investment, $Y_t - C_t$, implies a level of future capital, $K_{t+1} = (1 - \delta)K_t + Y_t - C_t$, satisfying the restriction in (5). Finally, (18) and (19) determine the distribution of final goods firms over inventory levels for next period, μ_{t+1} . With the equilibrium K_{t+1} and μ_{t+1} , we move into the next date in the simulation, again solving for equilibrium, and so forth. Once the simulation is completed, the resulting data, $(p_t, p_t q_t, K_t, m_t)_{t=1}^T$, are used to re-estimate $(\chi_l^K, \chi_l^m, \chi_l^p, \chi_l^{pq})$ using OLS.

We repeat this two-step process, first solving for V given $(\chi_l^K, \chi_l^m, \chi_l^p, \chi_l^{pq})$, next using our solution for firms' value functions to determine equilibrium decision rules over a simulation, storing the equilibrium results for $(p_t, p_t q_t, K_t, m_t)_{t=1}^T$, and then updating $(\chi_{l+1}^K, \chi_{l+1}^m, \chi_{l+1}^p, \chi_{l+1}^{pq})$, until these parameters

²⁰This demand depends on the target inventory level $s^*(z, K, m)$, the start-of-period distribution of firms $\mu(s)$, and the adjustment thresholds of each firm type $\xi^T(s; K, m)$.

converge. The number of partition means used to proxy for the distribution μ, I , is chosen such that agents' forecasting rules are sufficiently accurate.

5.1 Forecasting functions

Table 5 displays the actual forecasting functions used for the baseline inventory model, based on a 4000 period simulation. We use a log-linear functional form for each forecasting rule that is conditional on the level of productivity, $z_i, i = 1, \dots, N_z$.²¹ In the results reported here, $I = 1$. This means that, alongside z and K , only the mean of the current distribution of firms over inventory levels, start-of-period aggregate inventory holdings, is used by agents to forecast the relevant features of the future endogenous state. This degree of approximation would be unacceptable if it yielded large errors in forecasts. However, table 5 shows that, for each of the two values of productivity, the forecast rules for prices and both elements of the approximate state vector are extremely accurate. The standard errors across all regressions are small, and the R^2 's are high, all above 0.999.

The regressions in table 5 also offer some insight into the impact of inventories on the model, as they provide a description of the behavior of equilibrium prices and the laws of motion for capital and inventories. In particular, note that there is relatively little impact of inventories, m_1 , on the valuation of current output, p , and capital, K . Inventories have somewhat larger influence in determining the price of intermediate goods and, of course, their own future value.

6 Results

6.1 Steady state

Table 6 presents the steady state behavior of final goods firms when we suppress stochastic changes in the productivity of intermediate goods producers, the sole source of aggregate uncertainty in our model. This table illustrates the mechanics of our generalized (S,s) inventory adjustment and its consequence for the distribution of production across firms. In our baseline calibration, where $\bar{\xi} = 0.22$, there are 6 levels of inventories identifying firms.²² This beginning of period distribution is in columns labelled 1–6, while the first column, labelled adjustors, represents those firms from each of these groups that undertake inventory adjustment prior to production.

The inventory level selected by all adjusting firms, referred to above as the target value s^* , is 1.694

²¹We have tried a variety of alternatives including adding higher-order terms and a covariance term. None of these significantly altered the forecasts used in the model.

²²The number of final goods firm types varies endogenously outside of the model's steady state.

in the steady state. Firms that adjusted their inventory holdings last period, those in column 1, begin the current period with 1.155 units of the intermediate good. Given the proximity of their stock to the target value, they are unwilling to suffer substantial costs of adjustment and, as a result, their probability of adjustment is low, 0.036. Thus the majority of such firms do not undertake inventory adjustment; these firms use 0.450, almost 40 percent, of their available stock of intermediate goods in current production.

Inventory holdings decline with the time since their last order, so firms are willing to accept larger adjustment costs as they move from group 1 across the distribution to group 6. Thus, their probability of undertaking an order rises as their inventory holdings decline, and the model exhibits a rising adjustment hazard in the sense of Caballero and Engel (1999). Firms optimally pursue generalized (S, s) inventory policies, undertaking factor adjustment stochastically, and the probability of an inventory adjustment rises in the distance between the current stock and the target level associated with adjustment.

The steady state table exhibits evidence of some precautionary behavior among final goods firms, as they face uncertainty about the length of time until they will next undertake adjustment. First, while the representative firm in the benchmark model orders exactly the intermediate goods it will use in current production, 0.42, ordering firms in the baseline inventory economy prepare for the possibility of lengthy delays before the next order, selecting a much higher production-time stock, 1.69. Next, as these firms' inventory holdings decline, the amount of intermediate goods used in production falls, as does employment and production. The intermediate goods-to-labor ratio, $\frac{m}{n}$, also falls, as firms substitute labor for the scarcer factor of production. However, the fraction of inventories used in production actually rises until, for firms with very little remaining stock, those in column 5, the entire stock will be exhausted in production unless adjustment is undertaken. Nonetheless, firms' ability to replenish their stocks prior to production in the next period implies that the adjustment probability is less than one. In fact, even among the 0.017 firms that begin the period with zero inventories, not all adjust immediately. Roughly 84 percent of them adjust prior to production, adopting the common target. The remainder, a group representing 0.28 percent of all plants, forego current production and await lower adjustment costs.²³ Hence, while the columns labelled 1 – 6 reflect the beginning of period distribution of firms over inventory levels, the final column is not relevant in the production-time distribution. The first column, reflecting the behavior of adjusting firms, replaces it in production.

²³Each member of this group re-enters production upon realizing a fixed cost at or below 0.184, roughly 85 percent of the maximum cost.

6.1.1 Comparison to estimated adjustment rates

Much of the empirical inventory literature has estimated linear inventory adjustment equations derived from linear-quadratic (LQ) models of firm behavior. Typically, these models predict that target inventory holdings are a function of expected sales and other variables, and that some constant fraction of the gap between actual and target inventory holdings is closed in each period. As discussed in Ramey and West (1999), estimates of this gap based on aggregate data typically uncover a first-order autocorrelation coefficient between 0.8 and 0.9, which implies that between 0.1 and 0.2 of the distance between target and actual inventories is closed in any given quarter. A number of researchers have objected that these rates of inventory adjustment are implausibly low.

Schuh (1996) provides evidence suggesting that aggregate estimates may be biased downwards. Estimating three versions of the linear stock adjustment model using monthly M3LRD data, he reports a mean duration of firm-level inventory gaps of 2.5 months. Next, he shows that this mean duration rises to between 4 and 6.5 months when he re-estimates using aggregated data. However, it is somewhat difficult to determine the usefulness of these estimates, since each of the empirical models examined explains very little of overall variation in firms' inventory levels.

Using quarterly COMPUSTAT data, McCarthy and Zakrajšek (2000) estimate a general adjustment hazard describing the average adjustment rate as a function of the inventory gap, the empirical counterpart to our $\alpha(s)$ in table 6. In contrast to the LQ model, which predicts linear stock-adjustment equivalent to a constant hazard, their estimation reveals a rising hazard in the firm-level data. Given their model-specific estimate of target inventory levels, McCarthy and Zakrajšek find that 99 percent of the firms in their sample have estimated adjustment rates between 0.6 and 0.8.²⁴

We evaluate the inventory adjustment predicted by our model against some of the aggregate and micro-evidence discussed above. The inventory adjustments here differ from those in the LQ model in that firms adjust completely (eliminating the entire gap between actual and target inventories) if they adjust at all. Thus, in table 6, the fractions of firms undertaking adjustment from each group, $\alpha(s)$, represent average adjustment rates as a function of the gap between actual and target inventories, $s - s^*$. As was evident from the table, these adjustment rates rise with the inventory gap; the model implies the rising adjustment hazard characteristic of generalized (S,s) adjustment. On average, approximately 27 percent of our firms undertake inventory adjustment in each period. Interpreting this as our counterpart to the percentage of the inventory gap that is closed each period, we find that our model's actual adjustment rate is substantially higher than the typical aggregate estimate, but lower than the firm-

²⁴These results rely upon an estimated target inventory level that is biased downward by its failure to allow for forward-looking precautionary motives such as those highlighted in our discussion of table 6 above.

level estimates of Schuh (1996). Nonetheless, the estimated persistence of the inventory-to-sales relation in our model, at 0.85, is consistent with its estimated counterpart from the aggregate data. This, when viewed through the lens of the standard stock adjustment equation, would imply an estimated adjustment rate substantially lower than the true one, as we discuss further in section 7. To compute the average duration of an inventory gap in our model, we use the population distribution in table 6 to obtain the duration probabilities for any given firm. Since adjustments occur within the period, we take the 26.8 percent of firms in the column labelled 1 as having 0 duration, the 25.8 percent of firms in the column labelled 2 as having a duration of 1 quarter, and so on. The mean duration of an inventory gap, measured in this way, is 1.57 quarters in our model, roughly 4.7 months. Finally, in comparison with the empirical adjustment hazards of McCarthy and Zakrajšek (2000), we find that only 25 percent of our firms have adjustment rates exceeding 0.6.

6.2 Business cycles

6.2.1 Inventory investment and final sales

Our first goal was to generalize an equilibrium business cycle model to reproduce the empirical regularities involving inventory investment. We saw this as a necessary first step in developing a model useful for analyzing the role of inventories in the business cycle. Table 7 presents our inventory model's predictions for the volatility and cyclicity of GDP, final sales, inventory investment and the inventory-to-sales ratio. These predictions, derived from model simulations, are contrasted with the corresponding values taken from postwar U.S. data. All series are Hodrick-Prescott filtered.

Panel A of the table reports percentage standard deviations for each series relative to that of GDP.²⁵ Contemporaneous correlations with GDP are listed in panel B. Together, the two panels of table 7 establish that our baseline inventory model is successful in reproducing both the procyclicality of net inventory investment and the higher variance of production when compared to final sales. Further, this simple model with nonconvex factor adjustment costs as the single source of inventory accumulation is able to explain 54 percent of the measured relative variability of net inventory investment. Finally, note that the inventory-to-sales ratio is countercyclical in our model, as in the data. We take these results to imply that the predictions of the model are sufficiently accurate to validate its use in exploring the impact of inventory investment on aggregate fluctuations.

Certainly, there are differences between the model and data. The most pronounced departures in the model are its understated variability of inventory investment and exaggerated countercyclicality of

²⁵The exception is net inventory investment, which is again detrended as a share of GDP.

the ratio of inventories to final sales. However, the strong procyclicality in inventory investment, as well as the excess variability of production over sales, are well reproduced by the model. The latter arises from the positive correlation between inventory investment and final sales, 0.87, in the simulated economy.

Before proceeding further, it is useful to note the relation of the relative price of goods held as inventories in our model, q , to its empirical counterpart. In the data, we measure the relative price of inventories using the one-period lagged implicit price deflator for private nonfarm inventories divided by the implicit price deflator for final sales.²⁶ Detrending the series, we find that its percentage standard deviation is 0.87 that of output, a value slightly larger than that in our inventory model (0.563) and our benchmark model (0.606), as seen in table 9. Both models predict a strongly countercyclical relative price (the contemporaneous correlation with GDP is -0.976 in the inventory model and -0.984 in the model without inventories), an immediate consequence of our assumption of shocks to the productivity of firms supplying intermediate goods. While the measured relative price is also countercyclical, a finding that motivated our choice of the location of the technology shock, its correlation with GDP is substantially weaker, -0.23 .²⁷

6.2.2 Aggregate implications of inventory investment

In table 8, we begin to assess the role of inventories in the business cycle using our model. The first row of each panel presents results for the benchmark model without inventories; the second row reports the equivalent moment from the inventory model driven by the same sequence of shocks. The most striking aspect of this comparison is the broad similarity in the dynamics of the two model economies. At first look, the introduction of inventories into an equilibrium business cycle model does not appear to alter the model's predictions for the variability or cyclicity of production, consumption, investment or total hours in any substantial way. The differences that do exist are quantitatively minor, and the qualitative features of the equilibrium business cycle model are unaltered. Household consumption smoothing continues to imply an investment series that is substantially more variable than output, allowing a consumption series that is less variable than output. Furthermore, the variability

²⁶The one-period lag in the inventory deflator is necessary in computing an empirical relative price series comparable to our model. This is because the inventory deflator in the data corresponds to inventories held at the end of a quarter, while our relative price corresponds to the beginning of the current quarter.

²⁷Our results are essentially unchanged if we replace the deflator for final sales in the data series' denominator with that for GDP or a weighted average of that corresponding to consumer nondurables and services. The percentage standard deviation of the ratio of the implicit price deflator for private nonfarm inventories to that of GDP, final sales or consumption is 1.46, 1.46 or 1.25, respectively, while the contemporaneous correlation with real GDP is -0.24 , -0.23 , or -0.25 .

of total hours remains lower than that of production. Likewise, panel B shows little difference in the contemporaneous correlations with output across the two models. The most apparent divergence appears with respect to capital, which is less procyclical in the inventory economy due to its reduced responsiveness of final sales.

We introduced our paper by discussing the view that inventories exacerbate fluctuations in production. Table 8 appears to provide some support for this view, as the baseline inventory economy has a higher standard deviation of GDP than the benchmark economy. However, the increase in GDP volatility is small, only 2.6 basis points. Given that the level of inventories in our model is calibrated to reproduce their intensity of use in the US economy, we may conclude from this that inventories are of minimal consequence in amplifying fluctuations in production. Furthermore, panel A shows that the variability of final sales actually falls in the presence of inventory investment.²⁸ This is further evident in the reduced relative variability of consumption and investment in the inventory model. The relative variability of total hours worked, by contrast, is raised relative to the economy without inventories.

Table 9 provides additional observations that may help in explaining the differences across models. Note that the inventory economy's higher relative variance in total hours arises entirely from increased variability in hours worked in the production of intermediate goods, L . Moreover, shifts toward more labor-intensive production of intermediate goods in times of high productivity are stronger in the inventory model, as reflected by its more countercyclical K/L series. This is partly because procyclical inventory investment diverts some resources away from the production of final goods, and hence from investment in capital. Total hours worked in final goods firms, N , are actually less variable in the presence of inventories. In both model economies, the use of intermediate goods per worker is procyclical, as technology shocks to intermediate goods production make the relative price of intermediate goods, q , countercyclical. However, this effect is weaker in the inventory economy; consequently M/N is less variable and less procyclical there.

Inventories exist in our model because of fixed adjustment costs. These costs imply state-dependent (S, s) adjustment policies for final goods firms maintaining stocks of intermediate goods. In table 6, we saw that only about 27 percent of firms actively adjust their inventories in any given period in the steady state.²⁹ Staggered adjustment reduces the average response of final goods firms to changes in relative prices associated with the business cycle. As a result, the response in final goods is dampened relative to the benchmark economy, resulting in the reduced variability of consumption, investment and final

²⁸Recall that final sales in the benchmark model is equivalent to production, given the absence of inventory investment.

²⁹Nonetheless, the rate of adjustment is strongly procyclical in the inventory model; its contemporaneous correlation with GDP is 0.95.

sales, the sum of these two series. One consequence of this dampened response is that efforts to increase production of intermediate goods following a positive productivity shock must rely relatively more on employment, and less on capital. This makes hours worked in intermediate goods production rise by more in such times than in the benchmark economy without inventories. Moreover, as productivity shocks are persistent, part of the raised level of intermediate goods delivered to adjusting final goods firms is retained by these firms as inventory investment, which increases in times of high productivity. Because this retained portion does not immediately translate into higher production of final output, fluctuations in final sales are dampened. Thus, inventory accumulation implies a second restraint on the volatility of final sales beyond that directly implied by the scarcity of inputs among those firms deferring orders.

In concluding this section, we emphasize what we see as a central result of our study. *All else equal*, a positive covariance between final sales and inventory investment must increase the variability of production. However, as was clear in table 8 and in the discussion above, final sales are not exogenous; they are affected by the introduction of inventories. Our general equilibrium analysis suggests that nonconvex costs, the impetus for the accumulation of inventories, tend to dampen changes in final output. The percentage standard deviation of final sales, 1.57 for the benchmark model, falls to 1.37 when inventories are present in the economy. This reduction in final sales variability largely offsets the effects of introducing inventory investment for the variance of total production.

6.2.3 Changes in average inventory holdings

The results of the previous section indicate that, when nonconvex costs induce firms to hold inventories, cyclical fluctuations in final goods production are reduced relative to those that would occur if the costs could be eliminated. It follows that higher levels of these costs should further mitigate the business cycle. We explore this claim by increasing the upper support of the cost distribution, $\bar{\xi}$, from the baseline value of 0.220 to 0.336. This pushes the average inventory-to-sales ratio up by 15 percent to 0.8315.³⁰ Maintaining all other parameters, and using the same simulated shock series as above, we contrast the behavior of this *high inventory* economy to the calibrated baseline inventory economy where the inventory-to-sales ratio is 0.7155, the average quarterly value observed between 1947:1 and 2002:1 in the data.

Table 10A reveals that higher inventory levels are associated with a fall in the variability of consumption, investment and final sales, while the volatility of hours worked in intermediate goods production is raised. However, with less responsiveness in the use of intermediate goods, the decline in the vari-

³⁰It may be useful to note that this is the average nominal inventory-to-sales ratio in the data over our sample period.

ability of labor employed by final goods firms largely offsets the impact of this increase on the standard deviation of total hours worked. As we have argued, nonconvex adjustment costs tend to dampen the response of final goods firms to the exogenous changes in productivity that drive the business cycle, both because of the staggered nature of their adjustments and because of their reluctance to deplete or over-accumulate their stocks in response to shocks. Thus, although we have increased adjustment costs to imply a fairly substantial rise in the average inventory-to-sales ratio, we find almost no change in the cyclical variability of GDP.

The increased prevalence of inventories in the high inventory economy is associated with more cyclically volatile inventory investment; its standard deviation relative to GDP rises to 62 percent that measured in the data. However, for the reasons described above, the underlying rise in adjustment frictions also causes the volatility of final sales to decline. As a result, although inventory investment continues to have positive correlation with final sales (0.867), GDP volatility rises by only 0.3 basis points relative to the baseline inventory economy. Based on these findings (viewed in reverse), we find little support for recent suggestions that technological improvements in inventory management, by reducing average inventory-sales ratios, are responsible for dampened U.S. business cycles.³¹ Instead, our results highlight a potentially stabilizing role of inventories that is easily overlooked when the endogeneity of final sales is ignored, or when the existence of inventories is assumed rather than derived.

7 Two puzzles about inventory adjustment

Our calibrated inventory model matches the data qualitatively in its prediction of a countercyclical inventory-to-sales ratio, but, as we noted in section 6.2.1, it overstates this countercyclicity. This happens because the relative price of intermediate goods in our model is too countercyclical given the single technology shock. We begin the section by relating this result to a puzzle raised in recent work by *Bils and Kahn (2000)*.

Based on a model in which inventories are assumed to be directly productive in generating sales, *Bils and Kahn* conclude that a business cycle model driven by technology shocks is incapable of delivering a countercyclical inventory-sales ratio in the absence of imperfect competition. The puzzle, they emphasize, is not that inventory investment is procyclical, but rather that it is not sufficiently so as to

³¹*Kahn, McConnell and Perez-Quiros (2001)* argue that reduced inventories are important in explaining the halving of GDP volatility since the mid '80s. This is disputed by *Ramey and Vine (2001)* in their study of the automobile industry. *Maccini and Pagan (2003)* also reject this thesis based on their experiments with an estimated model of inventory holding behavior.

keep inventory stocks in pace with sales.³² This difficulty arises quite immediately in their environment because the imposition of inventories as an input into sales leads these two series to move closely together over time. To break this tendency, and hence obtain the desired regularity, the authors find that they must introduce either procyclical marginal costs or countercyclical markups.

Here, by contrast, we have developed a business cycle model in which perfectly competitive final goods firms choose to hold inventories in order to reduce the fixed costs they incur in obtaining deliveries from their perfectly competitive suppliers. Moreover, business cycles in our model are driven by technology shocks alone. Nonetheless, our model has no difficulty in delivering a countercyclical inventory-sales ratio. In fact, it is excessive in this respect precisely because real marginal costs for final goods firms are too countercyclical. That said, for models designed to examine inventories, we view the current finding as an illustration of the central importance of providing a microfoundation for the presence of these stocks and studying them in general equilibrium.

Our model may also offer some insight into a puzzle raised in section 6.1.1, the surprisingly sluggish inventory adjustment speeds found in the data. Here we illustrate difficulties that can arise in inferring adjustment rates using an approach common in the empirical inventory literature that relies on partial adjustment towards a target inventory to sales ratio. We find that the estimated target relationship between inventory holdings and sales may fail to uncover state-dependence in the true target. Moreover, the law of motion assumed to govern aggregate adjustment towards this target may omit important terms that arise because of heterogeneity across firms.

Equation (25) is a version of the familiar stock-adjustment model, which assumes that actual economywide inventory holdings, S_t , adjust gradually toward a desired level of inventories, S_t^* , with ρ representing the rate at which the gap between the actual and target levels is closed in each quarter.

$$S_t = \rho S_t^* + (1 - \rho)S_{t-1} + \varepsilon_t \tag{25}$$

The stock-adjustment equation is operationalized by assuming that the unobservable desired stock is linearly related to sales,

$$S_t^* = \theta X_t, \tag{26}$$

where X_t is final sales.³³ As we have already discussed, typical estimates for the convergence rate, ρ , are between 0.1 and 0.2, and they are deemed implausibly low.

³²Recall that the procyclicality of inventory investment has been a central focus throughout the production-smoothing literature, given that the microfoundation for inventories there tends to generate the reverse prediction.

³³In some applications, cost variables are appended to the model. For example, Schuh (1996) includes a real interest rate. However, such terms are generally found to be insignificant.

We obtain an implied estimate of the adjustment rate ρ in our model as follows. First, we estimate θ using the cointegration approach described in Ramey and West (1999), which yields $\hat{\theta} = 0.7177$ for our simulated data. With this in hand, we then estimate the first-order autocorrelation of the inventory to sales relation, $S_t - \hat{\theta}X_t$, at 0.85. Ramey and West show that, given (25) and (26), this autocorrelation is equal to $(1 - \rho)$, which would imply an adjustment rate of $\hat{\rho} = 0.15$ for our model economy. Note that this lies in the center of the range of previous empirical estimates from aggregate data. However, it is only about one-half of the true value, 0.27.

There are several reasons why the persistence of the inventory-sales relation does not reveal the true average adjustment rate in our model economy. One reason is that equation (25) does not hold in our model. To see this, define $S_{t+1}^* \equiv s_t^* - m_t(s_t^*)$ as the common target inventory level held at the end of the period by each firm adjusting its stock in date t . Recall that the economy's true date t adjustment rate is the fraction of firms that are adjusters, $\rho_t \equiv \int H(\xi^T(s; z_t, A_t)) \mu_t(ds)$. Writing the aggregate inventory stock at the end of date t , S_{t+1} , as the sum of end-of-period inventories held by adjusters together with those held across all firms not adjusting, we arrive at the following relationship between true and target inventories.

$$S_{t+1} = \rho_t S_{t+1}^* + (1 - \rho_t) S_t + \int \left[1 - H(\xi^T(s; z_t, A_t)) \right] (s - m_t(s) - S_t) \mu_t(ds) \quad (27)$$

Equation 27 includes a weighted sum, across all firms not actively adjusting their stocks, of the differences between current end-of-period inventories and the average stock held at the end of the previous period. This time-varying term is missing in equation (25). A second reason that equation 25 fails to identify the true adjustment rate is that the relationship between target inventories and sales in our model is a nonlinear function of the aggregate state that is not captured in the first step of our estimation. Finally, in our model economy, the adjustment rate ρ_t is not only state-dependent, but co-moves positively with the target S_{t+1}^* .

8 Concluding remarks

In the preceding pages, we generalized an equilibrium business cycle model to allow for endogenous (S, s) inventories of an intermediate good in final goods production. We showed that our calibrated baseline model of inventories accounts for the procyclicality of inventory investment, the comovement of final sales and inventory investment (and hence the higher variance of production relative to sales), and slightly more than one-half of the relative variability of inventory investment. Using this model to assess the role of inventories in the aggregate business cycle, we found that the inventory economy

exhibits a business cycle that is broadly similar to that of its benchmark counterpart without inventory investment. The adjustment costs that induce inventory holdings also dampen fluctuations in final sales, which substantially limits the effects of inventory accumulation for the variability of total production, despite the positive correlation between final sales and inventory investment. Similar results appeared when we reexamined the model's predictions in the presence of higher adjustment costs; the increased variability of inventory investment was almost completely offset by reduced fluctuations in final sales.

To conclude, we briefly consider what our analysis might contribute to recent discussions regarding the large drop in U.S. GDP volatility in the mid-1980's. Evaluating the Kahn, McConnell and Perez-Quiros (2000) argument that improvements in inventory management were responsible for this change, Ramey and Vine (2001) identify a structural break at 1984:1 where the variance of GDP growth halves, and they provide a summary of the pre- and post-break dynamics of production, final sales and inventory investment in the durable goods sector, where they find the variance of production growth fell most sharply (by 80 percent).³⁴ We produce similar statistics for the aggregate series in table 11.

Panel A of our table shows that the cyclical volatility in U.S. domestic business production less housing dropped by 72 percent between 1954:1 - 1983:4 and 1984:1 - 2002:4. Variability in final sales and inventory investment showed lesser reductions, 64 and 27 percent, respectively. Thus, in panel B, the relative volatility of final sales rose, and, most importantly, the relative volatility of inventory investment rose substantially. This in itself suggests that a decline in inventories did not cause the dampened fluctuations in GDP. Finally, consistent with the rise in the two relative volatilities, the covariance between sales and inventory investment fell sharply, and their correlation coefficient dropped from roughly 0.49 to 0.08.

Based on our model, we view improvements in inventory management as an unlikely explanation for the drop in GDP volatility. First, in the aggregate data, the average real (nominal) inventory-sales ratio was 0.719 (0.858) during 1954:1 - 1983:4, and fell to 0.709 (0.731) during 1984:1 - 2002:4. Thus, the real ratio changed very little, roughly 1.4 percent, while the fall in the nominal ratio, at 16 percent, was quite comparable to the change examined in Table 10. From there, we see that the cyclical volatility in GDP is reduced by far less than even 1 percent when adjustment frictions are reduced to yield a 15 percent decline in the average inventory-sales ratio. Moreover, absent other changes in fundamentals, our theory predicts that this decline will be accompanied by a rise in the volatility of final sales, a fall in the relative volatility of inventory investment, and no change in the correlation between sales and

³⁴Their primary focus is more specifically on the automobile industry, which they use to consider an alternative explanation based upon reduced sales volatility (and persistence) coupled with nonconvexities in firms' cost functions implied by institutional constraints.

inventory investment.³⁵ We conclude that, irrespective of changes in inventory-sales ratios, the direct explanation for dampened business cycles must lie elsewhere in the economy.³⁶

In future work, we will consider additional sources of fluctuations. This is particularly important, as we know that the source of shocks has proved critical for the implications of the traditional inventory model. The technology shock studied here is ordinarily interpreted as a supply shock, since it raises productivity among intermediate goods producers. However, it may also be viewed by final goods firms as a demand shock, as it is essentially a rise in the relative price of their output. Thus, as in any general equilibrium model, the demand or supply origin of the current disturbance appears ambiguous. Nonetheless, when fluctuations arise from demand shocks that do not directly alter the relative price of intermediate goods, the cyclical role of inventories may differ from that seen here.

³⁵In moving from the high inventory economy to the baseline inventory economy, the percent standard deviation of final sales rises from 1.34 to 1.37, the relative volatility of inventory investment falls from 0.18 to 0.16, and the correlation between sales and inventory investment remains at 0.87.

³⁶Stock and Watson (2003) overview several proposed explanations and attempt to quantify the extent to which each has independently contributed to reduced cyclical volatility in the U.S. and other G7 countries. Their results suggest that the phenomenon may be a largely transitory result of smaller shocks experienced over the past two decades.

References

- [1] Bartelsman, E. J. and W. Gray (1996) 'The NBER Manufacturing Productivity Database' NBER Technical Working Paper No. 205.
- [2] Blinder, A. S. (1990) *Inventory theory and consumer behavior*, University of Michigan Press.
- [3] Blinder A. S. and L. J. Maccini (1991) 'Taking stock: a critical assessment of recent research on inventories' *Journal of Economic Perspectives* 5, 73-96.
- [4] Caballero, R. J. and E. M. R. A. Engel (1991) 'Dynamic (S,s) economies' *Econometrica* 59, 1659-86.
- [5] Caballero, R. J. and E. M. R. A. Engel (1999) 'Explaining investment dynamics in U.S. manufacturing: a generalized (S, s) approach' *Econometrica* 67, 783-826.
- [6] Caplin, A.S. (1985) 'The Variability of Aggregate Demand with (S,s) Inventory Policies' *Econometrica* 53, 1395-1410.
- [7] Christiano, L. J. (1988) 'Why does inventory investment fluctuate so much?' *Journal of Monetary Economics* 21: 247-80.
- [8] Coen-Pirani, D. (2002) 'Microeconomic inventory behavior and aggregate inventory dynamics' GSIA working paper, Carnegie-Mellon University.
- [9] Fisher, J. D. M. and A. Hornstein (2000) '(S, s) inventory policies in general equilibrium' *Review of Economic Studies* 67: 117-145.
- [10] Fitzgerald, T. J. (1997) 'Inventories and the business cycle: an overview' *Economic Review* **33** (3): 11-22, Federal Reserve Bank of Cleveland.
- [11] Hall, G. and J. Rust (1999) 'An empirical model of inventory investment by durable commodity intermediaries' Yale University working paper.
- [12] Hansen, G. D. (1985) 'Indivisible labor and the business cycle' *Journal of Monetary Economics* 16, 309-327.
- [13] Hornstein, A. (1998) 'Inventory investment and the business cycle' *Economics Quarterly* 84, 49-71, Federal Reserve Bank of Richmond.
- [14] Humphreys, B. R., L. J. Maccini, and S. Schuh (2001) 'Input and output inventories' *Journal of Monetary Economics* 47: 347-375.

- [15] Johnson, S. A. (1989) *Spline approximation in discrete dynamic programming with application to stochastic multi-reservoir systems* Unpublished dissertation (Cornell, Ithaca, NY).
- [16] Jorgenson, D. W., F. M. Gollop, and B. M. Fraumeni (1999) *Productivity and U.S. Economic Growth* (toExcel, New York).
- [17] Kahn, J. A. (1987) 'Inventories and the volatility of production' *American Economic Review* 77, 667-679.
- [18] Kahn, J., M. McConnell and G. Perez-Quiros (2001) 'Inventories and the information revolution: implications for output volatility' Federal Reserve Bank of New York working paper.
- [19] Khan A. and J. K. Thomas (2003) 'Nonconvex factor adjustments in equilibrium business cycle models: do nonlinearities matter?' *Journal of Monetary Economics* 50, 331-360.
- [20] King, R. G. and S. T. Rebelo (1999) 'Resuscitating real business cycles' in: M. Woodford and J. Taylor, eds., *Handbook of Macroeconomics* IB (Elsevier Science, Amsterdam): 927 - 1007.
- [21] Krusell, P. and A. A. Smith Jr. (1997) 'Income and wealth heterogeneity, portfolio choice, and equilibrium asset returns' *Macroeconomic Dynamics* 1, 387-422.
- [22] Krusell, P. and A. A. Smith Jr. (1998) 'Income and wealth heterogeneity in the macroeconomy' *Journal of Political Economy* 106, 867-896.
- [23] Kydland, F. and E. C. Prescott (1982) 'Time to build and aggregate fluctuations' *Econometrica* 50, 1345 - 1370.
- [24] Maccini, L. J. and A. Pagan (2003) 'Exploring the role of inventories in the business cycle' Johns Hopkins University working paper.
- [25] McCarthy, J. and E. Zakrajšek (2000) 'Microeconomic Inventory Adjustment: Evidence from U.S. Firm-Level Data' Finance and Economics Discussion Paper 2000-24, Board of Governors of the Federal Reserve System.
- [26] Mosser, P. C. (1991) 'Trade inventories and (S,s)' *Quarterly Journal of Economics* 106, 1267-86.
- [27] Prescott, E. C. (1986), 'Theory ahead of business cycle measurement' *Quarterly Review* 10, Federal Reserve Bank of Minneapolis: 9-22.
- [28] Ramey, V. A. (1989) 'Inventories as factors of production and economic fluctuations' *American Economic Review* 79, 338-54.

- [29] Ramey, V. A. (1991) 'Nonconvex costs and the behavior of inventories' *Journal of Political Economy* 99, 306-334.
- [30] Ramey, V.A. and D. J. Vine (2001) 'Tracking the Source of the Decline in GDP Volatility: An Analysis of the Automobile Industry' UCSD working paper.
- [31] Ramey, V. A. and K. D. West (1999) 'Inventories' in: M. Woodford and J. Taylor, eds., *Handbook of Macroeconomics* IB. (Elsevier Science, Amsterdam), 863 - 922.
- [32] Richardson, H. (1995), 'Control Your Costs then Cut Them' *Transportation and Distribution*, 94-96.
- [33] Rogerson, R. (1988) 'Indivisible labor, lotteries and equilibrium' *Journal of Monetary Economics* 21, 3-16.
- [34] Scarf, H. (1960) 'The optimality of (S,s) policies in the dynamic inventory problem' in: K. J. Arrow, S. Karlin and H. Scarf, eds., *Mathematical Methods in the Social Sciences* (Stanford University Press, Stanford), 196-202.
- [35] Stock, J. R. and D. M. Lambert (1987) *Strategic Logistics Management* (Irwin, Homewood, Illinois).
- [36] Stock, J. H. and M. W. Watson (2003) 'Has the business cycle changed? Evidence and explanations' Harvard University working paper.
- [37] Schuh, Scott (1996) 'Evidence on the Link Between Firm-Level and Aggregate Inventory Behavior' Finance and Economics Discussion Series 96-46, Board of Governors of the Federal Reserve System.

Table 1: GDP, Final sales and inventories

	GDP	Final Sales	Net Inventory Investment
percent standard deviation relative to GDP*	2.237	0.710	0.295
correlation with GDP	1.000	0.943	0.669
correlation with NII	0.669	0.411	1.000

* Column 1 of row 1 reports percentage standard deviation of domestic business GDP less housing; columns 2 – 3 report volatilities relative to GDP for final sales of domestic business and changes in private nonfarm inventories. Data are quarterly U.S., 1953:1 – 2002:1, seasonally adjusted and chained in 1996 dollars. GDP and final sales are reported as percentage standard deviations, detrended using a Hodrick-Prescott filter with a weight of 1600. Net inventory investment is detrended as a share of GDP.

Table 2: Sectoral distribution of private non-farm inventories

	percentage of total stock of inventories*	STD(inventory investment)	correlation(inventory investment, GDP)
Manufacturing	37	0.14	0.65
finished goods	13		
work in process	12		
materials & supplies	12		
Trade			
retail	26	0.12	0.32
wholesale	26	0.09	0.35
Other	11		

* Column 1, the percentages of the total stock of inventories, is taken from Ramey and West (1999), page 869, table 4.

Table 4: Baseline calibration

β	η	α	θ_m	θ_n	δ	ρ	σ_ε	σ	ξ	$\bar{\xi}$
0.984	2.128	0.374	0.499	0.328	0.017	0.956	0.015	0.012	0.000	0.220

β : household subjective discount factor, η : preference parameter for leisure, α : capital's share in intermediate goods production, θ_m : intermediate goods' share in final goods production, θ_n : labor's share in final goods production, δ : capital depreciation rate, ρ : technology shock persistence, σ_ε : standard deviation of technology innovations, σ : per-unit inventory storage cost, ξ : adjustment cost lower bound, $\bar{\xi}$: adjustment cost upper bound.

Table 5: Forecasting rules in the baseline inventory economy

	β_0	β_1	β_2	S.E.	adj. R ²
Z ₁ :					
pq	0.563	- 0.471	- 0.102	0.65e-003	0.9998
p	1.357	- 0.328	- 0.050	0.01e-003	1.0000
K'	0.079	0.882	0.025	0.20e-003	1.0000
m₁'	- 0.227	0.154	0.766	1.13e-003	0.9997
Z ₂ :					
pq	0.538	- 0.522	- 0.065	0.93e-003	0.9997
p	1.341	- 0.337	- 0.036	0.02e-003	1.0000
K'	0.069	0.896	0.006	0.20e-003	1.0000
m₁'	- 0.240	0.190	0.744	0.86e-003	0.9998

Forecasting rules conditional on current productivity: $\log(X) = \beta_0 + \beta_1 [\log(K)] + \beta_2 [\log(m_1)]$. Number of observations for z_1 and z_2 are 2409 and 2591 respectively.

Table 6: Distribution of final goods firms in steady-state

	adjustors	1	2	3	4	5	6
$\mu(s)$: start-of-period distribution		0.268	0.258	0.224	0.159	0.074	0.017
s : start-of-period inventories		1.155	0.705	0.343	0.094	0.003	0.000
$\alpha(s)$: fraction adjusting		0.036	0.132	0.292	0.534	0.806	0.838
s_1 : production-time inventories	1.694	1.155	0.705	0.343	0.094	0.003	0.000
production-time distribution	0.268	0.258	0.224	0.159	0.074	0.014	0.003
m : intermediate goods usage	0.539	0.450	0.361	0.249	0.092	0.003	0.000
n : labor	0.225	0.196	0.167	0.127	0.060	0.004	0.000
y : production	0.437	0.385	0.331	0.253	0.121	0.008	0.000
m/n	2.401	2.291	2.165	1.966	1.520	0.595	na

Table 7: Inventory results for the baseline model

	GDP	Final Sales	Net Inventory Investment	Inventory/Sales
A: percent standard deviations relative to GDP*				
data	2.237	0.710	0.295	0.545
baseline inventory	1.598	0.859	0.158	0.807
B: contemporaneous correlations with GDP				
data		0.943	0.669	- 0.381
baseline inventory		0.997	0.906	- 0.933

* Column 1 of panel A reports percent standard deviation of GDP in the data and baseline inventory model. Data series are domestic business GDP less housing, final sales of domestic business, changes in private nonfarm inventories and private nonfarm inventory-to-sales ratio. Data are quarterly U.S., 1953:1 – 2002:1 seasonally adjusted and chained in 1996 dollars. GDP, final sales and the inventory-sales ratio are reported as percentage standard deviations, detrended using a Hodrick-Prescott filter with a weight of 1600. Net inventory investment is detrended as a share of GDP. Simulated model data are detrended in the same manner.

Table 8: Baseline inventory economy

	GDP	Final Sales	Consumption	Investment	Total Hours	Capital
A: percent standard deviations relative to GDP*						
benchmark	1.572	1.000	0.413	7.175	0.642	0.419
baseline inventory	1.598	0.859	0.374	6.240	0.666	0.377
B: contemporaneous correlations with GDP						
benchmark		1.000	0.919	0.967	0.967	0.132
baseline inventory		0.997	0.888	0.978	0.972	0.092

* Column 1 of panel A reports percent standard deviation of GDP.

Table 9: Baseline inventory economy continued

	L	N	X	M	q	K / L	M / N
A: percent standard deviations relative to GDP							
benchmark	0.642	0.642	1.600	1.600	0.606	0.809	1.012
baseline inventory	0.826	0.526	1.680	1.391	0.563	0.965	0.913
B: contemporaneous correlations with GDP							
benchmark	0.967	0.967	0.998	0.998	- 0.984	- 0.700	0.964
baseline inventory	0.964	0.978	0.999	0.989	- 0.976	- 0.789	0.943

L: labor in intermediate goods production, N: labor employed by final goods firms, X: total production of intermediate goods, M: total usage of intermediate goods, q: relative price of intermediate goods, K/L: capital-to-labor ratio in intermediate goods production, M/N: intermediate goods-to-labor ratio in final goods production.

Table 10: High inventory economy

	GDP	Final Sales	C	I	TH	L	N	M	NII	K/L	M/N
A: percent standard deviations relative to GDP*											
baseline inventory	1.598	0.859	0.374	6.240	0.666	0.826	0.526	1.391	0.158	0.965	0.913
high inventory	1.601	0.836	0.364	6.177	0.674	0.863	0.511	1.360	0.184	0.999	0.896
B: contemporaneous correlations with GDP											
baseline inventory		0.997	0.888	0.978	0.972	0.964	0.978	0.989	0.906	-0.789	0.943
high inventory		0.996	0.879	0.978	0.974	0.965	0.981	0.987	0.909	-0.803	0.939

* Column 1 of panel A reports percent standard deviation of GDP. Abbreviated series are C: consumption, I: capital investment, TH: total hours worked, L: labor in intermediate goods production, N: labor employed by final goods firms, X: total production of intermediate goods, M: total usage of intermediate goods, NII: net inventory investment, K/L: capital-to-labor ratio in intermediate goods production, M/N: intermediate goods-to-labor ratio in final goods production.

Table 11: Inventories and the structural break in GDP volatility

	GDP	Final Sales	Net Inventory Investment
A: percent standard deviations*			
U.S. 1954:1 - 1983:4	2.655	1.869	0.718
U.S. 1984:1 - 2002:4	1.249	0.960	0.548
B: percent standard deviations relative to GDP			
U.S. 1954:1 - 1983:4		0.704	0.270
U.S. 1984:1 - 2002:4		0.768	0.439
C: contemporaneous correlations with net inventory investment			
U.S. 1954:1 - 1983:4	0.705	0.486	
U.S. 1984:1 - 2002:4	0.540	0.080	

* Data series are domestic business production less housing, final sales of domestic business and changes in private nonfarm inventories. Data are seasonally adjusted and chained in 1996 dollars. GDP and final sales are reported as percentage standard deviations, detrended using a Hodrick-Prescott filter with a weight of 1600. Net inventory investment is detrended as a share of GDP.