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

Aubhik Khan[1]
Federal Reserve Bank of Philadelphia

Julia K. Thomas
University of Minnesota
University of Pennsylvania

January 2003

[1]We thank seminar participants at Carnegie-Mellon, The Federal Reserve Bank of Minneapolis, Indiana, Ohio State, UQAM and Yale, session participants at the 2002 Midwest Macroeconomics and Society for Economic Dynamics meetings and, in particular, Fernando Alvarez, V. V. Chari, Daniele Coen-Pirani, Larry Jones, Ed Prescott and Randy Wright, for their comments. The views expressed in this paper do not necessarily reflect those of the Federal Reserve Bank of Philadelphia or the Federal Reserve System.

Abstract

We develop an equilibrium business cycle model in which final goods producers 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 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. By 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.

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 beyond 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 study a dynamic stochastic general equilibrium model where, given nonconvex factor adjustment costs, producers follow generalized (S, s) inventory policies with regard to intermediate goods. 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 (i) maintain inventories of intermediate goods, and (ii) adjust these inventories only when their stock is sufficiently far from a target level. Our equilibrium analysis implies that 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 these 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. In particular,

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

[2] See Blinder and Maccini (1991).

[3] When 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.

we focus on the cyclicality 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 adjustment 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 explain half of the measured cyclical variability of inventory investment. Furthermore, inventory investment is procyclical, and production is more volatile than sales, as consistent with the data.

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. Thus our equilibrium analysis, which endogenizes final sales, alters our understanding of the role of inventory accumulation for cyclical movements in GDP. In particular, we find that the positive correlation between final sales and net inventory investment does not imply that inventories necessarily amplify aggregate fluctuations in production. The dynamics of final sales are altered by their presence. In the context of our equilibrium business cycle model, the introduction of inventories does not substantially raise the variability of production, because it lowers the variability of final sales.

2 Empirical regularities

In this section, we discuss the small 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 27 percent that of output.[5] Next, net inventory investment is procyclical; its correlation coefficient with GDP is 0.66. Moreover, as the correlation between inventory investment and final sales is itself positive, 0.42 for the 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 dynamic stochastic 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 why they exist. In our model, they arise as a result of assumed 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 that the associated costs are continuous functions of the change in production. This assumption underlies the traditional production smoothing model (and extensions that retain

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

[5] Net investment in private nonfarm inventories x_t, is detrended relative to GDP; the detrended series is $\frac{x_t - \overline{x}_t}{\overline{y}_t}$, where $\overline{x}_t$ is the HP-trend of the series and $\overline{y}_t$ is the trend for GDP.

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 firm-level production technology. 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. 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 examining the successes and limitations of 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). The aggregate implications of this alternative model have been left relatively unexplored; in fact, thusfar there has been no quantitative equilibrium

[6] One frequently mentioned 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.

[7] This 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.

examination of (S,s) inventories.[8] Moreover, we view the (S,s) model as applying to a wider 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, as is not unlikely, 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. Finally, there is empirical support for the (S,s) alternative. For example, 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.

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

[8] The only equilibrium analysis we know of is that of Fisher and Hornstein (2000). They focus on explaining the greater volatility in production relative to sales in an (S,s) model of retail inventories without capital.

finished goods. Given the primary cyclical role of manufacturing input inventories, we adopt a model that is consistent with these stocks. However, it is important to emphasize that we do not limit our analysis to the role of manufacturing inputs alone. In particular, we do not identify our inventories, or our firms, as belonging to a specific sector of the economy. Rather, we interpret our inventories as stocks of intermediate goods that broadly represent goods 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 types of optimizing agents in the economy, households, intermediate goods producers and final goods firms. Households supply labor to both types of producers and purchase consumption from final goods firms. They save using 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 provide an explicit role for inventories by assuming that final goods firms face fixed costs of ordering or accepting deliveries of intermediate goods. In particular, 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 point in time. 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 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. The distribution of firms over inventory levels evolves according to a mapping Γ_μ, $\mu' = \Gamma_\mu(z, A)$, and capital similarly evolves according to $K' = \Gamma_K(z, A)$.[9] For convenience, we assume that productivity follows a Markov Chain, $z \in \{z_1, \ldots, z_{N_z}\}$, where

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

and $\sum_{j=1}^{N_z} \pi_{ij} = 1$ for each $i = 1, \ldots, N_z$. 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.[10] 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 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).

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

The producer takes as given that A evolves over time according to $\mu' = \Gamma_\mu(z, A)$ and $K' =$

[9] Throughout 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.

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

$\Gamma_K(z, A)$ where $A \equiv (K, \mu)$, and changes in productivity follow the law of motion described in (1).The following efficiency conditions describe its selection of employment and investment.

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

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

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 $(z, \omega(z, A), q(z, A))$, and $x(k; z, A) = zF(1, L(z, A))k$. This means that current profits are linear in k;

$$\pi(z, A) \equiv qzF(1, L(z, A)) + (1 - \delta) - \omega L(z, A).$$

It is straightforward to show that this property is inherited by 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.

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

When (5) is satisfied, the intermediate goods firm is indifferent to any level of k' and will purchase investment as the residual from final goods production after 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, $\xi \in [\underline{\xi}, \overline{\xi}]$. 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) as well as its current delivery cost ξ. Before production, it undertakes an inventory adjustment decision. In particular, the firm may absorb its fixed cost and adjust its stock of intermediate goods available for production, $s_1 \geq 0$. 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 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	$s_1 = s$	$s' = s - m$

Inventories incur storage costs; we assume that these are proportional to the level of inventories held. Given end of period inventories s', the 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 $(s; z, A)$.

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

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

3.2.1 Production decisions

Beginning with the second subperiod, 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

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

taking prices (p, ω and q), and the evolution of $A' \equiv (K', \mu')$ according to Γ_K and Γ_μ, as given. Given the production-time stock of intermediate goods, s_1, and the continuation value of inventories of these goods, $V(s'; z_j, A')$, equation (7) yields both the firm's employment (in production) decision, which satisfies $D_2 G(s_1 - s', n) = \omega$, and its use of intermediate goods. Let $N(s_1; z, A)$ describe its employment and $S(s_1; z, A)$ the stock of intermediate goods retained for future use. Current production of final goods is $Y(s_1; z, A) = G(s_1 - S(s_1; z, A), N(s_1; z, A))$. 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 that precedes production. At the start of the period, for a final goods firm with beginning of period 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 subperiod, 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.

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

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

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 this common target, which 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 $\widetilde{\xi}(s; z, A)$, the cost that equates the net value of inventory adjustment to the value of non-adjustment.

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

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

$$\xi^T(s;z,A) = \min\Big\{\max\Big(\underline{\xi}, \widetilde{\xi}(s;z,A)\Big), \overline{\xi}\Big\} \tag{11}$$

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

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

$$x_m(s,\xi;z,A) = s_1(s,\xi;z,A) - s \tag{13}$$

The common distribution of adjustment costs facing final goods firms, given their threshold adjustment costs, implies that $H\Big(\xi^T(s;z,A)\Big)$ 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

$$V(s;z,A) = pqs + H\left(\xi^T(s;z,A)\right)V^a(z,A) - p\omega\int_{\underline{\xi}}^{\xi^T(s;z,A)} \xi H(d\xi) \tag{14}$$
$$+\Big(1 - H\left(\xi^T(s;z,A)\right)\Big)\Big(V^1(s;z,A) - pqs\Big),$$

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, 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.

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

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

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

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

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

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.

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

We next define 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 $1 - H(\xi^T(s;z,A))$ do not adjust their inventories. Thus, $\mu(s)[1 - H\Big(\xi^T(s;z,A)\Big)]$ 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, $\mu(s^*(z,A)) + \int_{\mathbf{S}} H\Big(\xi^T(s;z,A)\Big)\mu\left(ds\right)$ 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 $S^{-1}(\widetilde{s};z,A)$ as the production-time inventory level that gives rise to next period inventories $\widetilde{s}$ in the solution to (7). For any stock $\widetilde{s}$ other than that arising from the target level of production-time inventories,

$S^{-1}(\widetilde{s}; z, A) \neq s^*(z, A)$,

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

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

$$\mu'(\widetilde{s}) = \mu\Big(s^*(z, A)\Big) + \int_{\mathbf{S}} H\Big(\xi^T(s; z, A)\Big)\mu(ds). \tag{19}$$

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 the numbers of new shares in final goods firms, $\lambda'_F(s)$, and intermediate goods firms, λ'_I, to purchase at prices $\rho_F(s; z, A)$ and $\rho_I(z, A)$ respectively.[11] Their expected lifetime utility maximization problem is described recursively below.

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

subject to

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

$$A' = \Gamma(z, A) \tag{22}$$

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, $\Lambda_I(k, \lambda_I, \lambda_F; z, A)$ 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.

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

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,

$$\omega(z,A) = \frac{D_2U\Big(C(1,\mu;z,A), 1-N(1,\mu;z,A)\Big)}{D_1U\Big(C(1,\mu;z,A), 1-N(1,\mu;z,A)\Big)}, \tag{23}$$

and all firms must discount future profit flows with state-contingent discount factors that are consistent with households' marginal rate of intertemporal substitution, $\frac{D_1U\left(C(1,\mu;z,A),1-N(1,\mu;z,A)\right)}{\beta D_1U\left(C(1,\mu';z',A'),1-N(1,\mu';z',A')\right)}$. Following the approach outlined in Khan and Thomas (2002), 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.

$$p(z,A) = D_1U\Big(C(1,\mu;z,A), 1-N(1,\mu;z,A)\Big) \tag{24}$$

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 $\Gamma_K(z,A) \equiv (1-\delta)K + \overline{Y}(z,A) - C(1,\mu;z,A)$, where $\overline{Y}(z,A)$ 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;

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

For any particular output price p, the two requirements above directly imply a relative price for intermediate goods, $q(z, A; p)$, and a wage, $\omega(z, A; p)$, which in turn imply levels of output and consumption. Given these, the equilibrium output price $p(z, A)$ is that which satisfies condition (5), so that the intermediate goods firm is satisfied to invest what remains of final output after consumption.

$$\overline{Y}(z, A) = C(1, \mu; z, A) + \left[K' - (1 - \delta)K\right]$$

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 intermediate goods firm's supply and total usage by final goods firms, $q(z, A)\Big(x(K; z, A) - \overline{M}(z, A)\Big)$.

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 three values, $N_z = 3$, 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 $\bar{\xi} = 0$, the result is a model where no firm has an incentive to hold inventories.[12] 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 business capital-to-GDP ratio of 1.094, a value derived from averaging U.S. data between $1953 - 2000$. 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 set to equal the value implied by the NBER-CES Manufacturing database, 0.5; this lies in the range estimated by Jorgenson, Gollop and Fraumeni (1987) for U.S. manufacturing over the years 1947-79.[13] The remaining production parameter, θ_n, is taken to imply a labor's share of output averaging $\frac{2}{3}$, a value similar to that selected by Hansen (1985) and Prescott (1986). In terms of preferences, the subjective discount factor, β, is selected to yield a real interest rate of 4 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. Both values are taken from Prescott (1986).

[12] This is essentially the real business cycle model of Hansen (1985) generalized for an intermediate good.

[13] As the intermediate goods in our model are not intended to represent those used in manufacturing alone, and this value is probably substantially lower than that for intermediate goods in retail and wholesale, we intend to explore higher θ_m in future.

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 $\varepsilon_{t+1} \sim N\left(0, \sigma_\varepsilon^2\right)$, 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 $\left(\rho, \sigma_\varepsilon^2\right)$. 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 upper support for adjustment costs (uniformly distributed on $[0, \overline{\xi}]$) and the storage cost associated with inventories. We determine the upper support as follows. Using NIPA data, we compute that the quarterly real private nonfarm inventory-to-sales ratio has averaged 0.714 in the U.S. between $1947Q1$ and $1997Q4$, when the data series ends. This 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. Thus, given the storage cost parameter σ, we select the parameter $\overline{\xi}$ to reproduce this average inventory-to-sales ratio in our model. The storage cost parameter is difficult to identify in the data; for our baseline calibration, it is set to equal the rate of depreciation on capital.[14] Given $\sigma = \delta$, the upper support of the cost distribution is calibrated at $\overline{\xi} = 0.204$.

5 Numerical method

(S, s) models of inventory accumulation are rarely examined in equilibrium. As these models are characterized by an aggregate state vector that includes the distribution of the stock of inventory holdings across firms, computation of equilibrium is 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

[14] We conduct sensitivity analysis with respect to σ, in each case adjusting $\overline{\xi}$ to maintain the inventory-sales ratio at 0.714. Higher storage costs raise the cost of holding inventories and thus require higher adjustment costs to match the measured inventory-sales ratio. As footnote 21 describes, reducing σ to a value 25 percent below that in the baseline calibration yields inconsequential changes in results. In current work, we are resolving the model using a value of σ consistent with conventional estimates of carrying costs.

argument of the value function.[15] We apply spline approximation to V and V^1, 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 method applied in Khan and Thomas (2002), which itself applied a variation on the method of Krussel and Smith (1998). In particular, we approximate the distribution in the aggregate state vector with a vector of moments, $m = (m_1, ..., m_I)$, 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 very 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, $K' = \widehat{\Gamma}_K\left(z, K, m; \chi_l^K\right)$, $m' = \widehat{\Gamma}_m\left(z, K, m; \chi_l^m\right)$, $p = \widehat{p}\left(z, K, m; \chi_l^p\right)$ and $pq = \widehat{pq}\left(z, K, m; \chi_l^{pq}\right)$, 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, $t = 1, \ldots, T$. To determine equilibrium in each period, we begin by calculating m_t using the actual distribution, μ_t, and then use $\widehat{\Gamma}_K$ and $\widehat{\Gamma}_m$ to specify expectations of K_{t+1} and m_{t+1}. This determines

[15] For additional details, see Khan and Thomas (2002).

$\beta \sum_{j=1}^{N_z} \pi_{ij} w\left(z_j, K_{t+1}, m_{t+1}\right)$, and $\beta \sum_{j=1}^{N_z} \pi_{ij} V\left(s'; z_j, K_{t+1}, m_{t+1}\right)$ for any s'. Given the second function, the conditional expected continuation value associated with any level of inventories, we are able to determine $s^*\left(z, K, m\right)$ and $\xi^T\left(s; K, m\right)$, hence recovering the decisions of final goods firms and thus next period's distribution, given any values for 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.[16] 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 capital tomorrow, $K_{t+1} = (1-\delta) K_t + Y_t - C_t$, that satisfies 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 $\left(\chi_l^K, \chi_l^m, \chi_l^p, \chi_l^{pq}\right)$ using OLS.

We repeat this two-step process, first solving for V given $\left(\chi_l^K, \chi_l^m, \chi_l^p, \chi_l^{pq}\right)$, 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 $\left(\chi_{l+1}^K, \chi_{l+1}^m, \chi_{l+1}^p, \chi_{l+1}^{pq}\right)$, until these parameters converge. The number of partition means used to proxy for the distribution μ, I, is increased until agents' forecasting rules are sufficiently accurate.

5.1 Forecasting functions

Table 5 displays the actual forecasting functions used for the baseline inventory model, that in which the model's inventory-sales ratio matches its measured counterpart when averaged over the simulation. We use a log-linear functional form for each forecasting rule that is conditional on the level of productivity, z_i, $i = 1, \ldots, N_z$.[17] In

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

[17] We 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. In future we plan to assess

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 2 shows that, for each of the three 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 0.997.

In that they provide a description of the behavior of equilibrium prices and the laws of motion for capital and inventories, the regressions in table 2 also offer some insight into the impact of inventories on the model. 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, in forecasting 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 $\overline{\xi} = 0.204$, there are 5 levels of inventories identifying firms.[18] This beginning of period distribution is in columns $1-5$, while the first column, labelled adjustors, lists those firms from each of these groups that undertake inventory adjustment prior to production.

$I = 2$ for robustness. We were unable to complete this experiment in the current draft because it implies 5−dimensional value functions, which given our nonlinear method, implies substantial additional computing cost.

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

The inventory level selected by all adjusting firms, referred to above as the target value s^*, is 1.221 in the steady state. Firms that adjusted their inventory holdings last period, those in column 1, begin the current period with 0.833 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 0.328, almost 40 percent, of their available stock of intermediate goods in current production.

As inventory holdings decline with the time since their last order, firms are willing to accept larger adjustment costs as they move from group 1 across the distribution to group 5. 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, 0.31, 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, 1.22. 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 4, the entire stock is exhausted. Nonetheless, firms' ability to replenish their stocks prior to production in the next period implies that, even here, the adjustment probability is less than one. However, for the 0.061 firms that begin the period with zero inventories, all adjust prior to production, adopting the common target. Hence, while the columns labelled $1-5$ 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.

6.1.1 Comparison to average empirical adjustment patterns

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 will be closed in each period. As discussed in Ramey and West (1999), when estimates of this gap are based on aggregate data, they typically uncover a first-order autocorrelation coefficient between 0.8 and 0.9, which implies that between 0.1 and 0.2 of the gap between target and actual inventories is closed in any given period. A number of researchers have objected that these rates of inventory adjustment are implausibly low.

Schuh (1996) provides evidence suggesting that aggregate estimates may indeed 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. Moreover, a persistent problem with empirical estimates of inventory adjustment is that they require knowledge of target inventory levels. These unobservable endogenous variables are necessarily model-specific in practice.

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.[19]

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 they are extensive, rather than intensive; since firms optimally follow (S, s) policies, they 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 noted above, these adjustment rates rise with the inventory gap; the model implies the rising adjustment hazard characteristic of (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 prediction is substantially higher than the typical aggregate estimate, but lower than the firm-level estimates of Schuh (1996). 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 27.4 percent of firms in the column labelled 1 as having 0 duration, the 26.4 percent of firms in the column labelled 2 as having a duration of 1 period, and so on. The mean duration of an inventory gap, measured in this way, is 1.46 quarters in our model, slightly less than 4.5 months. Finally, in comparison with the empirical adjustment hazards of McCarthy and Zakrajšek (2000), we find that only 22 percent of our firms have adjustment rates exceeding 0.6.

[19] It must be noted, however, that these results rely upon a measure of the inventory target that is almost certainly biased downward, in that it fails to account for forward-looking precautionary motives such as those highlighted in our description of table 6 above.

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 cyclicality 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.[20] 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 52 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.[21]

[20] The exception is net inventory investment, which is detrended relative to GDP, as described in footnote 5.

[21] To examine the sensitivity of our results to our choice of storage cost, we have also explored a value 25 percent lower than that used with our baseline calibration, $\sigma = 0.0144$, while raising $\overline{\xi}$ to maintain the average inventory-sales ratio at its calibration target. We found no qualitative change in our results, and quantitative changes were small. Specifically, the contemporaneous correlation between inventory investment and GDP falls to 0.754, the relative standard deviation of final sales rises to 0.910, and the relative standard deviation of inventory investment falls slightly to 0.124. Moreover, given strong tradeoffs between σ and $\overline{\xi}$ for the value of inventories, the comparative business cycle results discussed in the remainder of this section were also insentive to this parameter

Certainly, there are differences between the model and data. The most pronounced departures in the model are its understated variability of inventory investment and its exaggerated counter-cyclicality of the ratio of inventories to final sales. However, the degree of 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, 0.781 in our model.

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 current implicit price deflator for final sales.[22]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 (0.667) and our benchmark model without inventories (0.723), as seen in table 9. Both models predict a strongly countercyclical relative price, −0.981 in the inventory model and −0.993 in the model without inventories, an immediate consequence of our assumption of procyclical shocks to the productivity of firms supplying intermediate goods. While the measured relative price is also countercyclical, a finding that partly motivated our choice of the location of the technology shock, its correlation with GDP is substantially weaker, −0.23.[23]

change.

[22] The 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, (which in our model equals those at the start of the subsequent quarter), while our relative price corresponds to the beginning of the current quarter.

[23] Our 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.

6.2.2 Aggregate implications of inventory investment

In tables 8A and 8B, we begin to assess the role of inventories in the business cycle using our model. The first row of each table presents results for the benchmark model without inventories, the second row reports the equivalent moment from the inventory model based on the same simulated shock series. 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 cyclicality 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. The familiar features of household consumption smoothing continue 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 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.

One noteworthy difference between the benchmark business cycle economy and the baseline inventory economy is the latter's higher standard deviation of GDP. We introduced our paper by discussing the view that inventories exacerbate fluctuations in production. Table 8 appears to provide some equilibrium support for this view. However, the increase in GDP volatility is rather small, only 0.092 percentage 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, table 8A shows that the variability of final sales actually falls in the presence of inventory investment.[24] This is further evident in the relative variability of consumption and

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

investment, both of which are reduced in the inventory model. The variability of total hours worked, by contrast, is raised relative to the economy without inventories.

Tables 9A and 9B provide additional observations that may help in explaining the differences across models, particularly with regard to the hours series. Note that the inventory economy's higher 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, (reflected by the countercyclical K/L series), are stronger in the inventory model, 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 one-third of firms actively adjust their inventories in any given period in the steady state.[25] Staggered factor 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, as reflected in the reduced variability of consumption, investment and final 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 in-

[25] Nonetheless, the rate of adjustment is procyclical in the inventory model, and relatively variable. Its percentage standard deviation relative to output is 0.941, and the contemporaneous correlation between the number of firms undertaking adjustment and GDP is 0.961.

ventories. This appears to explain the increased variability of hours worked, both in total and in the production of intermediate goods, and the reduced variability of final sales. 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, 1.35 for the benchmark model, falls to 1.28 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 the inventory-to-sales ratio

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, increasing the level of inventories relative to final output in the economy, should further mitigate the business cycle. We explore this claim by increasing the upper support of the cost distribution, $\overline{\xi}$, from the baseline value of 0.204 to 0.3. This pushes the average inventory-to-sales ratio up by approximately 15 percent to 0.83. We interpret this change as a rise in the average level of inventory holdings in the economy. Maintaining all other parameters, and using the same simulated shock series as above, we contrast the behavior of this *high inventory* economy to the baseline inventory economy where the inventory-to-sales ratio is 0.714, the average quarterly value observed between $1947Q1$ and $1997Q4$ in the data.

Table 10A reveals that higher inventory levels are associated with a fall in the variability of consumption, investment and final sales, and also a reduction in the percentage standard deviation of GDP.[26] Moreover the volatility of hours worked in intermediate goods production rises, though, with lesser responses in intermediate goods usage, the decline in the variability of labor employed by final goods firms more than offsets any 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 firms to the exogenous changes in productivity that drive the business cycle, both because of the staggered nature of their factor adjustments and because of their reluctance to deplete or over-accumulate their stocks in response to shocks. As a result, larger average adjustment costs associated with a higher average inventory-to-sales ratio reduce, rather than amplify, the severity of business cycles.

The increased prevalence of inventories in the model economy certainly raises the variability of net inventory investment. Its standard deviation relative to GDP is now much closer, at 0.222, to the measured value in the data, 0.271. However, the volatility of final sales declines, its relative standard deviation moving closer to its empirical counterpart, 0.824. As a result the positive correlation between final sales and net inventory investment, 0.702, fails to raise the variance of production. In fact, GDP volatility actually falls relative to the economy with the lower inventory-to-sales ratio. Thus, when viewed in reverse as a switch from the high inventory to baseline inventory economy, our results do not support arguments that improvements in inventory management, in reducing average inventory-sales ratios, are responsible for the dampening of recent U.S. business cycles.[27] Instead, they highlight a potentially *stabilizing role of inventories,* one that is necessarily overlooked when the endogeneity

[26] Although the relative volatility of consumption rises in the high inventory economy, the percent standard deviation in consumption falls slightly.

[27] See Kahn, McConnell and Perez-Quiros (2001).

of final sales is ignored, and likely to be obscured when the existence of inventories is assumed rather than derived.

7 Concluding remarks

In the pages above, 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 higher variance of production relative to sales, the countercyclicality of the inventory-to-sales ratio (qualitatively), and approximately one-half of the relative variability of net inventory investment. Using this model to assess the role of inventory investment 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. However, the adjustment costs that induce inventory holdings also dampen fluctuations in final output, which substantially limits the effects of inventory accumulation for the variability of total production, despite the positive correlation between final sales and inventory investment. Reexamining the model's predictions in the presence of higher adjustment costs, we have seen that an increased presence of inventories in the economy actually reduces aggregate fluctuations.

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.

References

[1] Blinder, A. S. (1990) *Inventory theory and consumer behavior*, University of Michigan Press.

[2] 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.

[3] 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.

[4] Christiano, L. J. (1988) 'Why does inventory investment fluctuate so much?' *Journal of Monetary Economics* 21: 247-80.

[5] Coen-Pirani, D. (2002) 'Microeconomic inventory behavior and aggregate inventory dynamics' GSIA working paper, Carnegie-Mellon University.

[6] Fisher, J. D. M. and A. Hornstein (2000) '(S, s) inventory policies in general equilibrium' *Review of Economic Studies* 67: 117-145.

[7] Fitzgerald, T. J. (1997) 'Inventories and the business cycle: an overview' *Economic Review* **33 (3)**: 11-22, Federal Reserve Bank of Cleveland.

[8] Hall, G. and J. Rust (1999) 'An empirical model of inventory investment by durable commodity intermediaries' Yale University working paper.

[9] Hansen, G. D. (1985) 'Indivisible labor and the business cycle' *Journal of Monetary Economics* 16, 309-327.

[10] Hornstein, A. (1998) 'Inventory investment and the business cycle' *Economics Quarterly* 84, 49-71, Federal Reserve Bank of Richmond.

[11] Humphreys, B. R., L. J. Maccini, and S. Schuh (2001) 'Input and output inventories' *Journal of Monetary Economics* 47: 347-375.

[12] Johnson, S. A. (1989) *Spline approximation in discrete dynamic programming with application to stochastic multi-reservoir systems* Unpublished dissertation (Cornell, Ithaca, NY).

[13] Kahn, J. A. (1987) 'Inventories and the volatility of production' *American Economic Review* 77, 667-679.

[14] 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 (2001).

[15] Khan A.and J. K. Thomas (2002) 'Nonconvex factor adjustments in equilibrium business cycle models: do nonlinearities matter?' *Journal of Monetary Economics* forthcoming.

[16] 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.

[17] Krusell, P. and A. A. Smith Jr. (1997) 'Income and wealth heterogeneity, portfolio choice, and equilibrium asset returns' *Macroeconomic Dynamics* 1, 387-422.

[18] Krusell, P. and A. A. Smith Jr. (1998) 'Income and wealth heterogeneity in the macroeconomy' *Journal of Political Economy* 106, 867-896.

[19] Kydland, F. and E. C. Prescott (1982) 'Time to build and aggregate fluctuations' *Econometrica* 50, 1345 - 1370.

[20] 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.

[21] Mosser, P. C. (1991) 'Trade inventories and (S,s)' *Quarterly Journal of Economics* 106, 1267-86.

[22] Prescott, E. C. (1986), 'Theory ahead of business cycle measurement' *Quarterly Review* 10, Federal Reserve Bank of Minneapolis: 9-22.

[23] Ramey, V. A. (1989) 'Inventories as factors of production and economic fluctuations' *American Economic Review* 79, 338-54.

[24] Ramey, V. A. (1991) 'Nonconvex costs and the behavior of inventories' *Journal of Political Economy* 99, 306-334.

[25] 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.

[26] Rogerson, R. (1988) 'Indivisible labor, lotteries and equilibrium' *Journal of Monetary Economics* 21, 3-16.

[27] 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.

[28] 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	1.675	0.824	0.271
correlation with GDP	1.000	0.951	0.658
correlation with NII	0.658	0.417	1.000

Data are quarterly U.S., 1953.1 – 2001.2, 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 investment in private nonfarm inventories x_t, is detrended relative to GDP; the detrended series is $(x_t - \bar{x}_t) / \bar{y}_t$, where $\bar{x}_t$ is the HP-trend of the series and $\bar{y}_t$ is the trend for 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 3, 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	σ	δ	$\underline{\xi}$	$\overline{\xi}$	ρ	σ_ε
0.990	2.185	0.252	0.500	0.293	0.019	0.019	0.000	0.204	0.981	0.014

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, σ: per-unit inventory storage cost, δ: capital depreciation rate, $\underline{\xi}$: adjustment cost lower bound, $\overline{\xi}$: adjustment cost upper bound, ρ: technology shock persistence, σ_ε: standard deviation of technology innovations.

Table 5: Forecasting rules with one partition

		β_0	β_1	β_2	S.E.	adj. R^2
z_1:	pq	0.646	-0.291	-0.101	0.55e-003	0.9984
	p	1.353	-0.270	-0.033	0.03e-003	0.9999
	K'	0.024	0.886	0.015	1.45e-003	0.9999
	m_1'	-0.312	0.161	0.691	1.12e-003	0.9978
z_2:	pq	0.591	-0.325	-0.068	0.64e-003	0.9982
	p	1.341	-0.285	-0.010	0.04e-003	0.9999
	K'	-0.016	0.926	-0.034	1.66e-003	0.9998
	m_1'	-0.151	0.037	0.830	1.15e-003	0.9979
z_3:	pq	0.420	-0.226	-0.144	0.88e-003	0.9979
	p	1.269	-0.232	-0.041	0.06e-003	0.9999
	K'	0.056	0.846	0.019	1.60e-003	0.9997
	m_1'	-0.290	0.246	0.699	1.13e-003	0.9988

Forecasting rules conditional on current productivity: $\log(y) = \beta_0 + \beta_1 [\log(K)] + \beta_2 [\log(m_1)]$. Number of observations for z_1, z_2 and z_3 are 693, 1691 and 616, respectively.

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

	adjustors	1	2	3	4	5
$\mu(s)$: start-of-period distribution		0.279	0.269	0.232	0.158	0.061
s: start-of-period inventories		0.833	0.504	0.238	0.055	0.000
$\alpha(s)$: fraction adjusting		0.036	0.140	0.318	0.611	1.000
s_1: production-time inventories	1.221	0.833	0.504	0.238	0.055	0.000
production-time distribution	0.279	0.269	0.232	0.158	0.061	0.000
m: intermediate goods usage	0.389	0.328	0.266	0.183	0.055	0.000
n: labor	0.186	0.165	0.142	0.109	0.047	0.000
y: production	0.365	0.328	0.287	0.223	0.096	0.000
m/n	2.091	1.990	1.871	1.677	1.176	n/a

Table 7: Inventory dynamics for the baseline model

	GDP	Final Sales	Net Inventory Investment	Inventory/Sales
A: percent standard deviations relative to GDP				
data	1.675	0.824	0.271	0.721
baseline inventory	1.441	0.885	0.141	0.921
B: contemporaneous correlations with GDP				
data		0.951	0.658	-0.396
baseline inventory		0.996	0.834	-0.964

Table 8: Baseline inventory model

	GDP	Final Sales	Consumption	Investment	Total Hours	Capital
A: percent standard deviations relative to GDP						
benchmark	1.349	1.000	0.538	6.658	0.501	0.418
baseline inventory	1.441	0.885	0.471	6.323	0.575	0.407
B: contemporaneous correlations with GDP						
benchmark		1.000	0.965	0.961	0.959	0.158
baseline inventory		0.996	0.939	0.968	0.964	0.127

Table 9: Baseline inventory model continued

	L	N	X	M	q	K / L	M / N
A: percent standard deviations relative to GDP							
benchmark	0.501	0.501	1.721	1.721	0.723	0.693	1.258
baseline inventory	0.696	0.441	1.765	1.527	0.667	0.853	1.128
B: contemporaneous correlations with GDP							
benchmark	0.959	0.959	0.999	0.999	-0.993	-0.599	0.984
baseline inventory	0.955	0.962	0.998	0.991	-0.981	-0.719	0.966

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 model

	GDP	Final Sales	C	I	TH	L	N	M	NII	K/L	M/N
A: percent standard deviations relative to GDP											
high inventory	1.382	0.831	0.485	5.560	0.560	0.739	0.366	1.478	0.222	0.906	1.141
baseline inventory	1.441	0.885	0.471	6.323	0.575	0.696	0.441	1.527	0.141	0.853	1.128
B: contemporaneous correlations with GDP											
high inventory		0.988	0.939	0.963	0.959	0.936	0.968	0.981	0.804	-0.742	0.960
baseline inventory		0.996	0.939	0.968	0.964	0.955	0.962	0.991	0.834	-0.719	0.966

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.