

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

Aubhik Khan¹

Federal Reserve Bank of Philadelphia

Julia K. Thomas

Federal Reserve Bank of Philadelphia and NBER

18 December 2006

¹We thank Patrick Kehoe, Robert King and Richard Rogerson for extensive discussions and suggestions that have substantially raised the quality of this work. We are also grateful to Fernando Alvarez, Ellen McGrattan, Edward Prescott, Valerie Ramey and two anonymous referees, as well as seminar participants at Alicante, Bank of Canada, Carnegie-Mellon, Cornell, Indiana, Iowa, Ohio State, Montreal, Penn, Pompeu Fabra, UQAM and Yale and The Federal Reserve Banks of Atlanta, Minneapolis, Philadelphia and St. Louis and session participants at the 2002 Midwest Macro and SED meetings, February 2003 NBER EF&G meetings, and June 2003 Madrid Conference on Lumpy Investment, for helpful comments. Thomas thanks the Alfred P. Sloan Foundation for a research fellowship supporting this work, and the National Science Foundation for research support under grant #0318163. The views expressed in this paper do not reflect those of the Federal Reserve Bank of Philadelphia or the Federal Reserve System.

Abstract

We develop an equilibrium business cycle model where nonconvex delivery costs lead producers of final goods to follow generalized (S,s) inventory policies with respect to intermediate goods. When calibrated to match the average inventory-to-sales ratio in postwar U.S. data, our model reproduces two-thirds 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. Our model contradicts this view. Despite the positive correlation between sales and inventory investment, we find that inventory accumulation has minimal consequence for the cyclical variability of GDP. In equilibrium, procyclical inventory investment diverts resources from the production of final goods; thus, it dampens cyclical changes in final sales, leaving GDP volatility essentially unaltered. Moreover, although business cycles arise solely from shocks to productivity and markets are perfectly competitive in our model, it nonetheless yields a countercyclical inventory-to-sales ratio.

1 Introduction

Inventory investment is a volatile component of GDP. Moreover, its comovement with final sales raises the variance of production above that of sales. Historically, such observations have led researchers to emphasize changes in inventories as central to an understanding of aggregate fluctuations.¹ Despite this, modern business cycle theory has been surprisingly silent on the topic.² Our goal in this paper is to develop a quantitative general equilibrium model of endogenous inventory investment and use it to formally evaluate several prominent claims regarding the cyclical role of inventories.

We begin by extending a basic equilibrium business cycle model to include fixed costs associated with the acquisition of intermediate goods used in final production, thereby inducing firms to maintain inventories of intermediate goods and manage them according to generalized (S,s) policies. When calibrated to match the average aggregate inventory-to-sales ratio in the postwar U.S. data, our baseline inventory model accounts for 64 percent of the measured cyclical variability of inventory investment, and it reproduces each of the following essential empirical regularities involving inventories: (i) inventory investment is procyclical, (ii) it co-moves with final sales, and thus (iii) the cyclical variability of total production exceeds that of sales. To our knowledge, no previous model with a micro-foundation for inventories has succeeded in reproducing these three regularities in quantitative general equilibrium; even absent this discipline, they have eluded most models.

We use our model as a laboratory to re-consider the following claims about the role of inventories over the business cycle. First, as noted above, it is widely believed that procyclical inventory accumulation exacerbates aggregate fluctuations. Second, as a corollary, it follows that reductions in aggregate inventory holdings accompanied by smaller changes in inventory investment will dampen the severity of business cycles. This view, strengthened by the predictions of reduced-form inventory models, has led some to argue that improvements in inventory management may have caused the substantial fall in cyclical GDP volatility observed in the U.S. since the mid-1980s.³

Our analysis challenges both of these claims. First, we compare simulated data from our calibrated inventory model to that from a control model where the fixed costs causing inventories are eliminated. The results of this exercise indicate that the aggregate business cycle would be essentially unchanged, and the percent standard deviation of GDP would fall by just 2.8 basis points, from 1.886 to 1.858, if inventories were to disappear entirely from the economy. Next, in an exercise intended to crudely

¹Blinder (1990, page viii) concludes "business cycles are, to a surprisingly large degree, inventory cycles." See also the survey by Ramey and West (1999) and the references cited therein.

²When 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 a factor of production, or Kahn, McConnell and Perez-Quiros (2001), where they are a source of household utility.

³This idea, originally espoused by Kahn, McConnell and Perez-Quiros (2000), has been widely studied. See, for example, McCarthy and Zakrajsek (2000), Ahmed, Levin and Wilson (2002), Irvine and Schuh (2002), Ramey and Vine (2004), Stock and Watson (2002) and Maccini and Pagan (2003).

mimic a substantial improvement in inventory management, we compare simulated data from two versions of our inventory model distinguished only by the size of their fixed delivery costs. Our results there indicate that the standard deviation of HP-filtered GDP would fall imperceptibly, from 1.896 to 1.886, if the frictions causing inventories were reduced sufficiently to yield a 15 percent decline in the size of these stocks.

These findings are specific to the outcomes of a particular inventory model, one where inventory accumulation arises from fixed order costs, markets are perfectly competitive, and technology shocks drive all aggregate fluctuations. Nonetheless, they raise the following broader observation about partial equilibrium analyses involving inventories. If movements in final sales are taken as exogenous, researchers need only note the positive correlation between sales and inventory investment to conclude that changes in inventory investment must raise the cyclical volatility of GDP. However, central to our general equilibrium analysis is the fact that both inventory investment and final sales are endogenous, and that their dynamics are inter-related. Because they effectively enter the same resource constraint, there is an important trade-off between inventory accumulation versus consumption and capital investment that is notably absent in partial equilibrium analyses. Thus, smaller fluctuations in inventory investment are accompanied by greater fluctuations in the sum of these other activities.

Because the above-mentioned equilibrium trade-off is central to understanding our inventory model's mechanics, it is instructive to consider its origin and implication in the model's response to a persistent rise in exogenous total factor productivity. During the resulting expansion, firms' efforts to avoid the delivery costs implied by excessively frequent orders lead them to increase their stocks of intermediate goods. This implies a larger increase in the demand for these goods than would otherwise occur, and thus a disproportionate rise in labor allocated to their production at some expense to the rise in labor allocated to final goods production. However, because the capital stock is both predetermined and slowly evolving, diminishing returns to labor discourages large changes in the production of intermediate goods. As a result, procyclical inventory investment diverts resources that might otherwise be used in the production of final goods, thereby dampening the rise in final sales. Conversely, in a recession, reduced inventory investment shifts more intermediate goods into production and moderates the fall in final sales. Thus, cyclical fluctuations in inventory investment do not substantially raise the variability of GDP because they lower the variability of final sales. Similarly, when a decline in fixed delivery costs makes inventory accumulation less important in the economy, the resulting fall in the variability of inventory investment is almost entirely offset by increased volatility in final sales.

Our work also has an important implication regarding the source of business cycle fluctuations. It responds to a challenge extended by the work of Bils and Kahn (2000), which suggests that a technology-shock driven business cycle is incompatible with the behavior of inventories. In particular, Bils and Kahn argue that, absent imperfect competition, business cycle models driven by technology

shocks can not reproduce the observed countercyclical inventory-to-sales ratio.⁴ Although its markets are perfectly competitive and its business cycles arise solely from productivity shocks, our model economy nonetheless exhibits both procyclical inventory investment and a countercyclical inventory-to-sales ratio. These two regularities of the data coexist quite naturally in our economy, again as a consequence of general equilibrium. Because capital adjusts gradually in equilibrium, changes in aggregate inventory holdings are themselves protracted. Thus, procyclical movements in the aggregate stock of inventories are gradual relative to those in the flow of final sales.

Finally, we use our model to explore the puzzlingly slow inventory adjustment speeds in many empirical studies.⁵ When viewed through the lens of a conventional estimation approach, our simulated model data exhibit persistence in the inventory-to-sales ratio consistent with existing empirical estimates, and thus an implied inventory adjustment rate neatly inside the range of values inferred from actual data. However, this estimate is far from the actual adjustment speed in our economy. We find that heterogeneity in firms' inventory and production levels breaks the linear mapping between the persistence of the inventory-sales ratio and the economywide adjustment rate implied by the standard stock-adjustment equation.

While useful for understanding how the dynamics of inventory accumulation interact with the dynamics of sales and other broad macroeconomic aggregates, the simplicity of our baseline inventory model limits its quantitative fit in two respects. First, it understates the relative volatility of inventory investment. Second, it generates an excessively countercyclical relative price of inventories. These deficiencies are largely resolved in enhanced versions of the model, as we discuss in a final section of the paper. Nonetheless, we continue to find that the inventory-sales ratio is countercyclical, and that the cyclical volatility of GDP is essentially unaffected by the frictions causing inventories. Throughout these enhanced versions of our model, alongside a battery of other examples, our results repeatedly challenge the conventional wisdom about inventories.

2 Model selection

Given positive real interest rates, the first challenge in any formal analysis of inventories is to explain their existence. Within macroeconomics, by far the most common rationalization for these 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

⁴Bils and Kahn (2000) examine a partial equilibrium reduced-form inventory model and find that, without countercyclical markups (which are rejected by the industry data they examine), the model cannot simultaneously generate procyclical inventory investment and a countercyclical inventory-sales ratio when it is driven by technology shocks.

⁵See Ramey and West (1999). Several studies have argued that this puzzle in the data arises from the omission of important nonconvexities in the firm-level production technology. See, for example, Caplin (1985), Blinder and Maccini (1991) and McCarthy and Zakrajšek (2000).

its simplest form, the model assumes that final sales are an exogenous stochastic series, and that adjustments to the level of production involve 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 only 13 percent of total private nonfarm inventories (Ramey and West (1999), p. 869). Additionally, a number of researchers have found that this class of model fares 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 arise 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.

Given the extensive body of research already devoted to the production smoothing model, we instead base our analysis on the leading microeconomic model, the (S,s) inventory model originally solved by Scarf (1960). In our model, inventories arise as a result of nonconvex delivery costs. To economize on such costs, firms hold stocks, making active adjustments only when these stocks are sufficiently far from a target. We choose to explore this motive in part because it may explain a broad 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 approach. Mosser (1991) tests a simple fixed-band (S,s) model on aggregate retail trade data and finds it more successful in explaining the observed time series than the traditional linear quadratic model. More recently, McCarthy and Zakrajšek (2000) have isolated nonlinearities indicative of (S,s) 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, thus far 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), which focuses on explaining the

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. This has been addressed in several ways. For example, Ramey (1991) studies increasing returns to production, while Eichenbaum (1989) explores productivity shocks. The most common approach, motivated by the stockout avoidance model of Kahn (1987), has been to assume costs of deviating from a target inventory-to-sales ratio; see Ramey and West (1999).

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.⁷ This allows them to tractably study (S,s) inventory policies in general equilibrium without confronting substantial heterogeneity across firms.

In our model economy, changes in the aggregate stock of inventories occur through two channels. First, changes in the order sizes of firms engaged in inventory investment produce movements along the intensive margin. Next, changes in the fractions of firms placing orders from each given level of inventories (i.e., shifts in a nontrivial adjustment hazard) interact with a time-varying distribution of firms over inventory holdings to produce movements along the extensive margin. The assumptions made by Fisher and Hornstein (2000) preclude the first of these mechanisms, which we find is an important channel 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 above, the dynamics of capital accumulation play a central role in determining the aggregate effects of inventories in our 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 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). Manufacturing inventories as a whole are far more cyclical than retail and wholesale inventories (the other main components of private nonfarm inventories).⁸ However, within manufacturing, inventories of intermediate inputs are twice the size of finished goods inventories (see Ramey and West (1999)). Moreover, the results of a variance decomposition undertaken by Humphreys, Maccini and Schuh (2001) indicate that they are three times more volatile. Given these observations, we develop a model that includes inventories of manufacturing inputs. However, we do not limit our analysis to these stocks. 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 our model to deliver a relative magnitude of inventories matching that of total private nonfarm inventories.

⁷Specifically, 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.

⁸Over the postwar period, the contemporaneous correlation between detrended inventory investment and GDP is 0.65 for manufacturing, while it is 0.32 (0.35) for retail (wholesale) trade.

3 Model

3.1 Overview

There are three sets of agents in the economy: households, intermediate goods producers and final goods firms. Households supply labor to all producers, and they purchase consumption goods from final goods firms. Intermediate goods firms own capital and hire labor for production. They supply their output to final goods producers, from whom they purchase investment goods. Final goods firms use intermediate goods and labor to produce output used for consumption and capital accumulation. All firms are perfectly competitive.

We assume a continuum of final goods firms with measure one. Each produces output using intermediate goods, m , and labor, n , through a concave, decreasing returns to scale production function, $G(m, n)$. We provide an explicit motive for inventory accumulation by assuming that these 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 $s \in R_+$.

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 units of labor and drawn from a time-invariant distribution $H(\xi)$ common across firms. Before production, the firm can pay its fixed cost and adjust its stock of intermediate goods available for current production, $s_1 \geq 0$. Letting x_m denote the size of such an adjustment, the stock available for production is $s_1 = s + x_m$. Alternatively, the firm can avoid the cost, set $x_m = 0$, and enter production with its initial stock, $s_1 = s$.

Following its inventory adjustment decision, a final goods 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 $s' = s_1 - m$. (Throughout the paper, primes indicate one-period-ahead values.) Finally, inventories incur storage costs proportional to the level of inventories held. Given end of period inventories s' , the total cost of storage is $\sigma s'$, where $\sigma > 0$ is a parameter capturing the marginal cost of holding inventories in units of the final good. Thus, the firm's net production of final goods is $y = G(m, n) - \sigma s'$.

Intermediate goods are supplied by a large number of identical producers. The representative intermediate goods firm produces with capital, k , and labor, l , using a constant returns to scale technology, F . Its output is $x = zF(k, l)$, where z is exogenous stochastic total factor productivity. Capital depreciates at the rate $\delta \in (0, 1)$, and the firm augments its capital stock for the next period using final goods as investment; $k' = (1 - \delta)k + i$.

The aggregate total factor productivity shock follows a Markov Chain, $z \in \{z_1, \dots, z_{N_z}\}$, where $\Pr(z' = z_j | z = z_i) \equiv \pi_{ij} \geq 0$, and $\sum_{j=1}^{N_z} \pi_{ij} = 1$ for each $i = 1, \dots, N_z$. This is the sole source of aggregate uncertainty in the model. Its placement in the intermediate goods sector is dictated by the countercyclical relative price of inventories in the aggregate data, since the relative price of inventories in the equilibrium of our model will equal the relative price of intermediate goods (which

will fall only if a shock raises intermediate goods firms' productivity relative to that of final goods firms). Throughout the paper, we represent current productivity, z_i , by z except where necessary for clarity.

A unit measure of identical households value consumption and leisure in each period and discount future utility by $\beta \in (0, 1)$. Households are endowed with 1 unit of time in each period, and they supply labor to all firms in the economy. They own all intermediate and final goods firms, and they have access to a complete set of state-contingent claims.⁹ Denoting the representative household's total consumption and labor supply at date t by c_t and n_t^h , respectively, its expected discounted lifetime utility is $\mathbf{E}_0 \sum_{t=0}^{\infty} \beta^t u(c_t, 1 - n_t^h)$.

3.2 Competitive equilibrium

We now describe the behavior of producers and households, beginning with a summary of the aggregate state observed by all prior to their decisions.¹⁰ Recall that final goods firms carry stocks of intermediate goods to avoid frequent payment of fixed delivery costs. At any date, some firms adjust their stocks and others do not, given differences in delivery costs. Thus, the model yields an endogenous distribution of final goods firms over inventory levels, $\mu : \mathcal{B} \rightarrow [0, 1]$, where \mathcal{B} is the Borel algebra and $\mu(S)$ represents the measure of firms with start-of-period inventories in the set $S \in \mathcal{B}$.

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 their total factor productivity described in the section above.¹¹ 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 $K' = \Gamma_K(z, A)$. Below, we summarize the law of motion governing the endogenous aggregate state by $A' = \Gamma(z, A)$.

The final good is the numeraire, and equilibrium relative prices are functions of the aggregate state. Firms employ labor at real wage $\omega(z, A)$, and intermediate goods are traded at relative price $q(z, A)$. Finally, $Q_j(z, A)$ is the price of an Arrow security that will deliver one unit of the final good next period if $z' = z_j$; in equilibrium, all firms discount their future earnings using these state-contingent prices.

Problem of a final goods firm: Let $v^0(s, \xi; z, A)$ represent the expected discounted value of a final goods firm with current inventory stock, s , and fixed cost draw, ξ , given the aggregate state

⁹We introduce a complete set of Arrow securities only so as to derive the prices that firms use to discount their future profit flows.

¹⁰While we have chosen to examine the decentralized economy here, it should be noted that the competitive allocation corresponds to the solution of a planning problem, since markets are complete and perfectly competitive.

¹¹As firms' delivery costs are iid draws from a time-invariant distribution, the joint distribution of firms over inventories and these costs may be constructed from the distribution over inventories alone. Hence, delivery cost draws are not part of the aggregate state.

(z, A) . Recall that any such firm chooses whether or not to undertake active inventory adjustment prior to production. Contingent on that decision, the firm selects its order for intermediate goods, $x_m \neq 0$, which determines its stock available for current production, $s_1 = s + x_m$. Given s_1 , the firm chooses its employment $n \geq 0$, and future inventories, $s' \geq 0$, thus determining its production net of storage costs, $y = G(s + x_m - s', n) - \sigma s'$.

We state the problem facing such a firm using equations (1) - (2). Although we suppress the arguments of q , ω and Q_j for ease of exposition, recall that all are functions of the aggregate state. The firm's problem is

$$v^0(s, \xi; z, A) = \max \left\{ -\omega \xi + \max_{x_m \geq -s} \left[-qx_m + v^1(s + x_m; z, A) \right], v^1(s; z, A) \right\}. \quad (1)$$

Here, $v^1(s_1; z, A)$ summarizes expected discounted profits gross of current order costs, conditional on the available stock of intermediate goods at production time:

$$v^1(s_1; z, A) = \max_{n \geq 0, s_1 \geq s' \geq 0} \left[G(s_1 - s', n) - \sigma s' - \omega n + \sum_{j=1}^{N_z} Q_j \int_{\underline{\xi}}^{\bar{\xi}} v^0(s', \xi; z_j, A') H(d\xi) \right], \quad (2)$$

given the aggregate law of motion $A' = \Gamma(z, A)$. Finally, in equation (2), $\int_{\underline{\xi}}^{\bar{\xi}} v^0(s', \xi; z_j, A') H(d\xi)$ represents the expected continuation value associated with future inventories s' if the aggregate state next period is (z_j, A') .

Intermediate goods firm's problem: Given its pre-determined capital stock, k , and the current aggregate state, (z, A) , the representative intermediate goods firm chooses current employment, l , and capital for the next period, k' . Its value, $w(k; z, A)$, solves the functional equation:

$$w(k; z, A) = \max_{l \geq 0, k' \geq 0} \left(qzF(k, l) - \omega l - (k' - (1 - \delta)k) + \sum_{j=1}^{N_z} Q_j w(k'; z_j, A') \right), \quad (3)$$

given $A' = \Gamma(z, A)$.

Household's problem: The representative household receives an aggregate dividend, $D(z, A)$, from the economy's firms in each period, and its net worth is held in the form of Arrow securities, a . In addition to its asset income, the household receives labor income $\omega(z, A)n^h$ given its choice of total hours worked, n^h .

In each period, the household allocates its initial wealth plus labor and dividend income across current consumption, c , and purchases of new securities, a'_j , $j = 1, \dots, N_z$, to maximize its expected discounted lifetime utility. Specifically, given its assets and the aggregate state, the household chooses

$\left(c, n^h, \left(a'_j \right)_{j=1}^{N_z} \right)$ to solve

$$h(a; z, A) = \max \left(u \left(c, 1 - n^h \right) + \beta \sum_{j=1}^{N_z} \pi_{ij} h \left(a'_j; z_j, A' \right) \right) \quad (4)$$

subject to

$$c + \sum_{j=1}^{N_z} Q_j(z, A) a'_j \leq a + D(z, A) + \omega(z, A) n^h. \quad (5)$$

To rule out Ponzi schemes, the following additional constraints are imposed on household purchases of securities: $a'_j \geq \underline{a}$, $j = 1, \dots, N_z$, where $\underline{a} < 0$. These constraints do not bind in equilibrium. Finally, the household also takes as given the evolution of the endogenous aggregate state, $A' = \Gamma(z, A)$.

Equilibrium: A *Recursive competitive equilibrium* is a set of functions, $(v, x_m, n, s', w, l, k', h, c, n^h, (a_j)_{j=1}^{N_z}, \omega, q, (Q_j)_{j=1}^{N_z}, D, \Gamma_K, \Gamma_\mu)$, satisfying the following conditions.¹²

1. Firm and household decisions are optimal;
2. Markets for final goods, intermediate goods, labor and securities clear;
3. Laws of motion for aggregate state variables are consistent with individual decisions:
 - (a) $\mu'(\tilde{S}) = \int_{\{(s, \xi) | s'(s, \xi; z, A) \in \tilde{S}\}} H(d\xi) \mu(ds)$ for all $\tilde{S} \in \mathcal{B}$ defines $\Gamma_\mu(z, A)$;
 - (b) $K' = k'(K; z, A)$ defines $\Gamma_K(z, A)$.

As there is no heterogeneity across households, there is zero net supply of Arrow securities in equilibrium, so $a = 0$. As a result, the representative household's consumption and total hours worked may be written simply as functions of the aggregate state, $C(z, A)$ and $N(z, A)$. We will use this result below.

3.3 Firm behavior and inventory adjustment

In this section, we develop several properties of firms' decision rules that will be used in discussing our model's results. We first simplify the description of their optimization problems using a result from the representative household's problem. In equilibrium, the household's choice of Arrow securities requires that $Q_j(z, A) = \pi_{ij} \frac{\beta U_1(C(z_j, A'), 1 - N(z_j, A'))}{U_1(C(z, A), 1 - N(z, A))}$. This allows us to reformulate firms' problems, eliminating the time-varying discount factor. To be more precise, we now require that each firm weight its current profits by the output price $p(z, A) = U_1(C(z, A), 1 - N(z, A))$ and discount its future expected earnings by β . The only implication of our reformulation is that value functions are

¹²To avoid additional notation, we use choice variables to denote decision rules. Thus (x_m, n, s') are functions of final goods firms' state vector (s, ξ, z, A) ; (l, k') are functions of the intermediate goods firm's state vector (k, z, A) ; and $(c, n^h, (a_j)_{j=1}^{N_z})$ are functions of the household state vector (a, z, A) .

now measured in units of marginal utility, rather than final output; the resulting Bellman equations yield the same decision rules as above.

We begin by presenting the intermediate goods firm's reformulated value function. Suppressing the arguments of p, q and ω for brevity, W solves

$$W(k; z, A) = \max_{k', l} \left(p \left[qzF(k, l) + (1 - \delta)k - k' - \omega l \right] + \beta \sum_{j=1}^{N_z} \pi_{ij} W(k'; z_j, A') \right). \quad (6)$$

Linear homogeneity of F immediately implies that the firm's decision rule for employment, and hence its production of the intermediate good, is proportional to its capital stock.

Next, we turn to final goods firms. Let $V^0(s, \xi; z, A)$ represent the reformulated value function 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 (7) - (9) below. We divide the period into two sub-periods, an adjustment sub-period and a production sub-period, and we break the exposition of the firm's problem into the distinct problems it faces as it enters into each of these sub-periods.

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, its inventories for next period, and hence its current production, to solve

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

which is the analogue to (2) above. Given the continuation value of inventories, $V(s'; z_j, A')$, equation (7) yields both the firm's employment (in production) decision, $N(s_1; z, A)$, and its stock of intermediate goods retained for future use, $S(s_1; z, A)$. Its net production of final goods is then $Y(s_1; z, A) = G(s_1 - S(s_1; z, A), N(s_1; z, A)) - \sigma S(s_1; z, A)$. Thus, we have decision rules for employment, production, and next-period inventories as functions of the production-time stock s_1 .

Given the middle-of-period valuation of the firm, V^1 , we now examine the inventory adjustment decision made by a final goods firm entering the period with inventories s and drawing adjustment cost ξ . Equations (8) - (9) describe the 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:

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

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

Note that the target inventory choice in (9) is independent of both the current inventory level, s , and the fixed cost, ξ . Thus, all firms that adjust their inventory holdings choose the same production-time stock 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). Equation (7) then implies common employment and intermediate goods use 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 inventories, it is immediate from equation (8) that a firm will do so if its fixed cost is at or below $\tilde{\xi}(s; z, A)$, the cost that equates the net value of inventory adjustment to the value of non-adjustment:

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

Note that the cost satisfying (10) depends upon the firm's initial stock, s , which we will refer to as its type. 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:

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

From the analysis above, we arrive at the following decision rules for production-time inventory holdings and stock adjustments:

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

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

Finally, as final goods firms face a common distribution of adjustment costs, H , the probability that a firm of type s will alter its inventory stock before production is given by $H(\xi^T(s; z, A))$.

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 demands for intermediate goods and for labor, their use of intermediate goods, and their production of the final good. The aggregate order for intermediate goods is the sum of the stock adjustments from each start-of-period inventory level s , weighted by the numbers of firms undertaking these adjustments:

$$\mathbf{X}(z, A) = \int H(\xi^T(s; z, A)) (s^*(z, A) - s) \mu(ds). \quad (14)$$

Total use of intermediate goods, $\mathbf{M}(z, A)$, is the total production-time stock across adjusting and non-adjusting firms, less that retained by each firm as inventories for the subsequent period:

$$\begin{aligned} \mathbf{M}(z, A) = & \left[s^*(z, A) - S(s^*(z, A); z, A) \right] \int H(\xi^T(s; z, A)) \mu(ds) \\ & + \int \left[s - S(s; z, A) \right] \left[1 - H(\xi^T(s; z, A)) \right] \mu(ds). \end{aligned} \quad (15)$$

Aggregate production of the final good is a weighted sum of the output of adjusting and nonadjusting firms, and, similarly, total employment among final goods firms is a weighted sum of labor in their

production, together with the total time costs of adjustment:

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

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

In concluding this section, it is useful to explain how the aggregates above determine our counterparts to production, sales and inventory investment in the NIPA. In the model, final sales is the aggregate production of final goods, $\mathbf{Y}(z, A)$. Next, aggregate net inventory investment is defined as the change in the total value of inventories. In our model, this is the difference between total orders and use of intermediate goods weighted by their relative price, $q(z, A)(\mathbf{X}(z, A) - \mathbf{M}(z, A))$. Finally, in the model, as in the data, GDP is the sum of sales and inventory investment.

4 Calibration and solution

We examine the implications of inventory accumulation for our otherwise standard equilibrium business cycle model using numerical methods. In calibrating the 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 constant-returns-to-scale Cobb-Douglas production function with capital share α . Their productivity follows a Markov Chain with nine values, $N_z = 9$, that is itself the result of discretizing an estimated log-normal process for technology with persistence ρ and variance of innovations, σ_ε^2 . Each final goods firm's production function is $G(m, n) = m^{\theta_m} \eta^{\theta_n}$, with intermediate goods' share θ_m and labor's share θ_n . The adjustment costs that provide the basis for their inventory holdings are uniformly distributed with lower support 0 and upper support $\bar{\xi}$. Finally, we assume that the representative household has period utility $u(c, 1 - n^h) = \log c + \eta \cdot (1 - n^h)$.¹³

¹³This specification may be derived from a model with indivisible labor and employment lotteries, as shown by Rogerson (1988).

4.1 Control model

If we set $\bar{\xi} = 0$, the result is a model where no firm has an incentive to hold inventories.¹⁴ 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 as a *control* model against which to evaluate the effect of introducing inventory accumulation. The parameterizations of the control and inventory models are identical, with the already noted exception of the cost distribution associated with adjustments to intermediate goods holdings.

The first set of parameters common to both the control and inventory models ($\alpha, \theta_m, \theta_n, \delta, \beta, \eta$) are derived as follows. 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 and 2002. The depreciation rate δ is equal to the average ratio of business investment to business capital over the same time period. The distinguishing feature of the control model, relative to the Indivisible Labor Economy of Hansen (1985), is the presence of intermediate goods. The new parameter associated with this 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, of 0.499. 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 one-third of available time.

We determine the stochastic process for productivity using the Crucini Residual approach described in King and Rebelo (1999). In contrast to the Solow method, where total factor productivity shocks are inferred using data on aggregate output, capital and labor together with the assumption of an aggregate production function, the Crucini approach infers these shocks from a linear approximation to the full solution of a model alongside data on aggregate output and each series present in the model's endogenous aggregate state vector. When applied to our control model, this approach allows us to estimate a shock series using data on only aggregate output and capital. A continuous shock version of the control 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

¹⁴When $\bar{\xi} = 0$, any final goods firm can order exactly the quantity of intermediate goods it will use in current production without suffering delivery costs. In this case, the expected return to holding inventories is simply the appreciation of the relative price of intermediate goods adjusted for storage costs, $\frac{\mathbf{E}(q'|z, \mu)}{(1+\sigma)q} - 1$. Our calibration implies persistent aggregate shocks that are never sufficiently large that this return exceeds the expected real interest rate. Thus inventories are never held in simulations of this special case of our model.

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 yield new estimates of $(\rho, \sigma_\varepsilon^2)$. The process is repeated until these estimates converge. Resulting values for the shock's persistence and the implied variability of model GDP are similar to those found in comparable business cycle studies (for example, Prescott (1986)). Table 1 lists the complete baseline parameter set.

4.2 Inventory model

For all parameters that are also present in the control 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 particular, the capital-output ratio, hours worked, and the shares to factors of production) are close.

The two parameters that distinguish the inventory model from the control are the proportional storage cost associated with inventories and the upper support for adjustment costs. 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 calculate that storage costs should represent 12 percent of the annual value of inventories in our model.¹⁵ Next, using NIPA data, we compute that the quarterly real private nonfarm inventory-to-sales ratio has averaged 0.7155 in the United States between 1954:1 and 2002:4. (As noted by Ramey and West (1999), the real series, in contrast to its nominal counterpart, exhibits no trend.) Given the parameters specified above, these calibration targets jointly determine our remaining two parameters, ($\sigma = 0.012$ and $\bar{\xi} = 0.220$), and imply that the relative price valuing inventories is $q = 0.417$ in the model's steady state.

4.3 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 (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

¹⁵Excluded components are cost of money, taxes, physical handling and clerical and inventory control; see Richardson (1995). The last component is already reflected in our model by labor-denominated adjustment costs.

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 defer adjustment and exhaust their entire current stock in 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 (1989).

Second, equilibrium prices are functions of a large state vector, given 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 in proxy for the distribution as they forecast the future state to make decisions given current prices. In choosing this proxy, we extend the method applied in Khan and Thomas (2003), which itself applied a variation on the method of Krusell and Smith (1997, 1998). In particular, we assume that agents 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 choices by producers, 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. In the results reported here, $I = 1$. This means that, alongside z and K , agents use only the mean of the current distribution of firms over inventory levels, the start-of-period aggregate stock, to forecast the relevant features of the future endogenous state.

The actual distribution of firms over inventories in our model is a large but finite object. In simulations, we use this actual distribution in each period, alongside firms' value functions (derived using the forecasting rules described above), to determine equilibrium prices and quantities and thus the subsequent period's distribution. Next, the simulation data are used in a regression step to revise agents' forecasting rules. Finally, based on the revised forecast rules, value functions are resolved, and the model is simulated again, with this iterative process repeating until the forecasting rules converge. The approximation implied by this numerical approach would be unacceptable if it generated large errors in forecasts; however, the resulting forecast rules prove to be highly accurate. Standard errors across all regressions are small, and R^2 's all exceed 0.999. A complete description of our solution algorithm, along with a table of agents' forecasting rules, is provided in a technical appendix available on request.

5 Steady state results

Table 2 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. In our baseline calibration, where $\bar{\xi} = 0.22$, firms are distributed over 6 levels of inventories at the

start of the period. This distribution is in columns labeled 1 – 6, while the first column, labeled adjusters, represents those firms from each of these groups that undertake inventory adjustment prior to production. Of course, 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 1.694 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. Because 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 6. Thus, their probability of undertaking an order rises as their inventory holdings fall further from the target, and the model exhibits a rising adjustment hazard in the sense of Caballero and Engel (1999).

The steady state table exhibits evidence of precautionary behavior among final goods firms, given their uncertainty about the length of time until they will next undertake adjustment. While the representative firm in the control 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, although final goods firms substitute labor for the scarcer factor of production as their inventory holdings decline, the fraction of inventories used in production 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 no inventory, only 84 percent adjust prior to production. The remainder, a group representing 0.28 percent of all plants, forego current production and await lower adjustment costs.

6 Business cycle results

We begin this section with a brief review of the empirical regularities concerning inventory investment that are most relevant to our analysis.¹⁶ Table 3 summarizes the business cycle behavior of GDP, final sales and net inventory investment in quarterly postwar U.S. data.¹⁷ All series are detrended using a Hodrick-Prescott filter with a weight of 1600. We HP-filter the logarithms of GDP and final sales. However, as net inventory investment is often negative, the same approach cannot be used for this series. Instead, it is detrended as a share of GDP; that is, we apply the HP-filter to

¹⁶For more extensive surveys, see Fitzgerald (1997), Hornstein (1998) and Ramey and West (1999).

¹⁷Reported series are GDP of domestic business less housing, final sales of domestic business, and changes in the total value of private nonfarm inventories, 1954:1 - 2002:1. Data are seasonally adjusted and chained in 1996 dollars.

the ratio NII/GDP .

Note first that the relative variability of net inventory investment is large. In particular, though its share of gross domestic production averages roughly one-half of one percent, its standard deviation is 29.5 percent that of output. Next, 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 data summarized here, the standard deviation of production substantially exceeds that of sales. Again, it is for this reason that the comovement of sales and inventory investment is commonly interpreted as evidence that fluctuations in inventory investment increase the variability of GDP.

6.1 Accounting for the inventory facts

Our model's predictions for the volatility and cyclicity of GDP, final sales and inventory investment, as well as the inventory-sales ratio, are reported in table 4. There, we compare the results of a 10,000 period simulation to the corresponding values taken from the data, with each model-generated series detrended exactly as is its empirical counterpart. Together, the panels of this table 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 accounts for 64 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.

Certainly, there are differences between the model and data, most notably the model's understated variability of inventory investment, an issue that we address below in section 8. 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, 0.83 in the simulated economy. 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.

The model's success in reproducing the basic inventory facts arises because aggregate inventories rise and fall with total production and final sales. Consider the economic dynamics following a persistent rise in productivity. This causes a rise in both current and planned future consumption, and hence in final goods production. As a result, final goods firms deplete their stocks of intermediate goods more rapidly. If average inventory holdings were left unchanged, this higher rate of use would necessarily require more frequent orders and a persistent rise in adjustment costs. To avoid this, firms increase their average inventory holdings.

At the shock's impact, the fall in the marginal cost of intermediate goods leads adjusting firms to place larger orders and accumulate more inventories. This, reinforced by a transitory rise in

adjustment rates that increases the number of ordering firms, more than offsets the more rapid decumulation among nonadjusting firms.¹⁸ In the aggregate, inventories begin to rise. However, given any level of intermediate goods production, there is a trade-off between accumulating these goods as inventories and using them in current production. Unless intermediate goods production increases sufficiently, increases in final sales are necessarily dampened by the rise in inventory investment. This will be important to our discussion in the next section.

Before proceeding further, it is useful to note the relation of the relative price of inventories in our model to its empirical counterpart. In the data, we measure the relative price of inventories using the one-period lagged implicit price deflator for end-of-period private nonfarm inventories divided by the implicit price deflator for final sales. Detrending the series, we find that its percentage standard deviation is 0.653 that of output, a value somewhat larger than that in our inventory model, 0.535. Our model predicts a strongly countercyclical relative price (its contemporaneous correlation with GDP is -0.973), an immediate consequence of our assumption of a single shock 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, -0.257 . This correlation would move further from the data if we had instead assumed a single shock evenly affecting intermediate and final goods production. In that case, it would be strongly positive, at 0.96, and even more so if we had assumed the shock affected only the production of final goods. This discrepancy between model and data may reflect additional shocks present in the data but absent our model economy, as explained in section 8.

6.2 Aggregate implications of inventory investment

In table 5, 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 control model without inventories; the second row reports the equivalent moment from the baseline inventory model driven by the same sequence of shocks. (We will defer discussion of the third rows until section 6.4.) 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.

We introduced our paper by discussing the view that inventories exacerbate fluctuations in production. Table 5 provides little support for this view. Though the baseline inventory economy has a higher standard deviation of GDP than the control economy, the difference is only 2.8 basis points. This is because the rise in GDP volatility implied by procyclical inventory investment is almost en-

¹⁸Given decreasing returns, efficiency requires that any rise in production be spread across the distribution of firms. However, the rise in production among nonadjusting firms is constrained by relatively low initial stocks. A rise in adjustment rates, by reducing the number of such firms, alleviates this problem.

tirely offset by a reduction in the volatility of final sales. (Recall that final sales in the control model is equivalent to production, given the absence of inventory investment.) Given that the level of inventories in our model is calibrated to reproduce their intensity of use in the U.S. economy, our result suggests that inventories do not amplify fluctuations in production.

In comparing the control and inventory economies, it is useful to note that agents in the former choose to invest only in one asset, capital, while agents in the latter choose to invest both in capital and in inventories of intermediate goods. Because production of final goods draws upon intermediate goods, procyclical inventory investment crowds out some capital investment, given that labor allocated to the production of intermediate goods does not rise sufficiently to fully accommodate the accumulation of these goods while consumption remains strongly procyclical. Thus, we see that capital is less procyclical in the inventory economy. Also consistent with reduced responses in the production of final goods, the relative variability of consumption is somewhat lower when inventories are accumulated. The relative variability of total hours worked, by contrast, is higher in the economy with inventories.

In both economies, a positive productivity shock lowers the relative price of intermediate goods, q , and predicts a persistent increase in production. This leads final goods firms to anticipate high use of intermediate goods. While firms in the control economy increase their orders by exactly the rise in use at each date, fixed adjustment costs make this policy suboptimal. As we discussed in section 6.1, firms in the inventory economy increase their average stocks in order to avoid more frequent orders and a persistent rise in adjustment costs. Consequently, the increase in the production of intermediate goods must supply not only their higher use, but also additional inventory accumulation. Relative to the control economy, this implies a disproportionate share of the rise in employment is allocated to intermediate goods production.

Our discussion suggests that employment among intermediate goods firms should be more responsive to aggregate shocks in the model with inventories, while the converse should hold for employment among final goods firms. Indeed, in the inventory model, the relative volatilities of labor in intermediate goods production, L , and labor in final goods production, N , are 0.890 and 0.537, respectively, while they are equal at 0.674 in the control model. Overall, we see that the inventory economy's higher variance in total hours arises from increased volatility in hours worked in the production of intermediate goods, and despite reduced hours volatility in final goods production. These patterns of relative variability in employment across sectors carry over into the production and use of intermediate goods. Thus, in the final columns of table 5, production of intermediate goods, X , is more volatile in the presence of inventories, while the use of these goods in final production, M , is less volatile.

Although the increase in intermediate goods production is larger in the inventory economy, it does not fully accommodate inventory accumulation because diminishing marginal product of labor hinders large increases in production, given capital. Thus, inventory accumulation diverts some intermediate

goods from being used in current production, thereby reducing the rise in employment and production among final goods firms. As a result, final sales is dampened relative to the control model. This reduces the rise in not only consumption, but also investment, which slows capital accumulation in the inventory economy and prolongs the dampening of final sales relative to the control. In other words, because the stock of capital constrains increases in the supply of intermediate goods, there is a tradeoff between increased production of final goods versus inventory accumulation that perpetuates itself by diverting some resources away from the production of investment goods that are used to increase the future stock of capital.

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 implies that inventories must increase the variability of production. However, as clear from table 5 and the discussion above, final sales are not exogenous; they are affected by the introduction of inventories. Our general equilibrium analysis suggests that procyclical inventory investment reduces cyclical fluctuations in final output. The percentage standard deviation of final sales falls from 1.858 to 1.583 when inventories are introduced into our model economy. This reduction in final sales variability largely offsets the effects of inventory investment for the variance of total production.

6.3 Role of capital

In the previous section, we found that procyclical inventory investment dampens the changes in final production following productivity shocks, and consequently inventories cause little increase in GDP volatility. In this section, we establish that the presence of capital in our model is essential to this prediction. Our reasoning is as follows. A substantial capital share in production, given diminishing marginal product of labor, slows the fall in the marginal cost of producing intermediate goods following a positive productivity shock. This decline is complete only after sufficient capital has been accumulated. Moreover, the rise in capital is itself gradual due to households' preference for smooth consumption profiles. This slows the rise in the supply of intermediate goods, so increases in inventory investment necessarily reduce intermediate goods used in final production.

Here, we verify the essentiality of capital in our results by considering a model with a substantially smaller role for capital in production. Specifically, we examine the effect of inventory accumulation when capital's share in intermediate goods production is reduced to imply an average capital-output ratio one-quarter that observed in the data. This involves a capital share in intermediate goods production of roughly 0.09.¹⁹

When capital is made largely irrelevant in production, cyclical variations in inventory investment

¹⁹To isolate the effects of the reduced capital-output ratio, we hold returns to scale constant, which implies a rise in the economywide share to labor from 0.64 to 0.78. Elsewhere, parameters are selected to preserve the control model's fit to the remaining calibration targets, and the shock process is recalibrated accordingly. As a result, there is no change in average hours worked or the mean inventory-to-sales ratio.

yield substantial increases in GDP volatility; the standard deviation of GDP, at 2.349, is roughly 51 basis points higher in the inventory model than it is in the corresponding control model without inventories, 1.837. (Recall that this difference was only 2.8 basis points in our comparison of the calibrated models above.) Here, inventory investment is far more volatile than in the calibrated inventory economy; in fact, its relative standard deviation, 0.375, exceeds that in the data. Moreover, much of this high volatility in inventory accumulation translates into high GDP volatility, because the dampening of final sales is comparatively minor; the percent standard deviation of sales, 1.84 in the control model, falls only to 1.64 in the presence of inventories.

Now that capital has such a small share in production, its important role in slowing aggregate responses to shocks is largely eliminated; adjustments in both models, control and inventory, are more rapid. Relative to our baseline model, the more rapid rise in the supply of intermediate goods following a positive technology shock allows faster inventory investment with less crowding out of final goods production. As a result, sharp quantitative differences emerge between the control and inventory models. In the control model, the percent standard deviations of orders and use of intermediate goods are equal at 3.16. Moving to the inventory economy, a sharp rise in the volatility of orders to 4.38 allows rapid inventory adjustments at little expense to the volatility of intermediate goods use in production, 2.91.

Given a very low capital share in intermediate goods production, marginal product of labor schedules are much flatter than those in the calibrated inventory economy. In this case, a positive shock to productivity generates larger increases in employment, and thus in the supply of intermediate goods, before diminishing marginal productivity discourages further increase. Thus, intermediate goods production responds much more sharply when capital's share is low, and the relative price of these inputs in final production falls further and more rapidly. This alleviates the trade-off between inventory accumulation and intermediate goods use, allowing final goods firms to raise their inventories faster, with little crowding out of final production. The result is an episode of rapid inventory accumulation during which GDP rises substantially more than its counterpart in the model without inventories. Thus, inventory investment can cause sharp increases in GDP volatility when capital is sufficiently unimportant in production. Comparing these results with those discussed in section 6.2, we conclude that capital plays a central role in determining the aggregate effects of inventories in our model.

We have shown here that our model's central trade-off between inventory investment and final sales arises because gradual capital accumulation slows changes in intermediate goods production. We argue that this result would remain if we assumed that capital was used to produce not only intermediate goods but also final goods. Consider such an alternative model. Assuming an unchanged aggregate capital to output ratio, the share to capital in intermediate goods production would be reduced, and thus capital would now be less important in constraining changes in the supply of intermediate goods. However, the level of capital in the production of final goods would now be an important determinant of the marginal product of intermediate goods. As such, capital used by final

goods firms (and hence aggregate capital) would directly limit the demand for intermediate goods in current production and thus, given the slow rise in capital, the demand for these goods as inventories toward production in nearby dates. As a result, gradual capital accumulation would again slow the increase in the production of intermediate goods, just as in our model.²⁰

6.4 Changes in average inventory holdings

In this section, we briefly consider what our analysis might contribute to recent discussions regarding the large drop in U.S. GDP volatility in the mid-1980s, and, in particular, the Kahn, McConnell and Perez-Quiros (2001) argument that improvements in inventory management were responsible for this change. Ramey and Vine (2004) identify a structural break at 1984:1 where they find the variance of GDP growth decreases by 50 percent. Examining the aggregate data before and after this date, we find that the standard deviation of (log HP-filtered) 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, 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 was not the leading force behind the dampened fluctuations in GDP.

To explore this question further, we increase the upper support of the cost distribution in our model, $\bar{\xi}$, from the baseline value of 0.220 to 0.333, which pushes the average inventory-to-sales ratio up by 15 percent to 0.8315. Maintaining all other parameters, and using the same simulated shock series as above, we contrast the behavior of this *high inventory* economy to our *baseline inventory* economy where the inventory-to-sales ratio is 0.7155, the average quarterly value observed in the data. The results of this exercise may be seen by comparing the third rows to the second rows in table 5.

Starting from the high inventory economy, and moving to the baseline, the reduced prevalence of inventories is associated with less cyclically volatile inventory investment; its standard deviation relative to GDP falls slightly from 68 to 64 percent of that measured in the data. However, for reasons described in section 6.2, the underlying reduction in adjustment frictions responsible for this lower inventory volatility also raises the volatility of final sales. As a result, GDP volatility falls by only 1 basis point.

Based on these results, we find little support for the suggestion that technological improvements in inventory management, by reducing average inventory-sales ratios, are responsible for dampened U.S. business cycles. In the data, the average real (nominal) inventory-sales ratio was 0.719 (0.858) during

²⁰This argument is supported by the results of a reduced-form model where inventories are assumed as a factor of production. There, as the share to capital in final goods production is raised (and its share in intermediate goods production lowered to maintain the aggregate capital-output ratio), we find no change in the model's aggregate dynamics. In particular, the minor increase in GDP volatility that occurs with the introduction of inventories is unaffected. Details are available upon request.

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 that we have just examined. Our theory predicts that the cyclical volatility of GDP would be reduced by far less than even 1 percent if adjustment frictions were reduced to yield a 15 percent decline in the average inventory-sales ratio. Moreover, absent other changes in fundamentals, this decline would be accompanied by a rise in the volatility of final sales and a fall in the relative volatility of inventory investment. We conclude that, irrespective of changes in inventory-sales ratios, the direct explanation for dampened business cycles lies elsewhere.²¹

7 Two puzzles about inventory adjustment

We noted in section 6.1 that our inventory model is consistent with the data in its prediction of a countercyclical inventory-to-sales ratio. This happens because the trade-off between inventory accumulation and intermediate goods use slows the rise in the aggregate stock following a rise in productivity. While the largest increase in sales occurs immediately, the rise in inventories, like the familiar response in capital, is far more gradual. 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 keep inventory stocks in pace with sales. This difficulty arises quite immediately in their environment because the particular function through which inventories are assumed to generate sales leads these two series to move closely together over time.²² 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. Yet the inventory-sales ratio is strongly countercyclical. For models

²¹ Focusing on the automobile industry, *Ramey and Vine (2004)* suggest an alternative explanation based upon reduced sales volatility (and persistence) and nonconvexities in firms' cost functions arising from institutional arrangements. Estimating a model of inventory holding behavior, *Maccini and Pagan (2003)* also reject the inventory-driven explanation. Finally, *Stock and Watson (2003)* examine several possible causes of reduced cyclical volatility in the United States and other G7 countries, among them reduced inventory holdings. Their results indicate that the fall in volatility may be a largely transitory result of smaller shocks experienced over the past two decades.

²² Sales are assumed to be a state-invariant power function of the stock of available goods. Specifically, $s = d(p) \cdot a^\phi$, where s is sales, a is the stock of goods available for sale, (the sum of inventories and current output), and p is price.

designed to examine inventories, this illustrates that the method used to introduce these stocks can substantially influence the results obtained. Further, reconsidering why inventories move more gradually than sales in our model economy, it is important to understand that aggregate stocks, whether of capital or inventories, tend to adjust slowly in general equilibrium. This is an immediate consequence of households' preference for smooth consumption paths, which prevents sharp movements in investment, and it suggests that general equilibrium analysis may be essential in understanding the inventory-to-sales relationship.²³

Our model may also offer some insight into a second puzzle, one involving inventory adjustment speeds. Much of the empirical inventory literature has estimated 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.

As indicated above, the common method of inferring inventory adjustment rates from the data rests on an assumption of partial adjustment toward a target inventory-to-sales ratio. Here, we illustrate the difficulties that can arise with this approach if the data is in fact generated by heterogeneous firms pursuing endogenous (S,s) inventory policies. Equation (18) 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:

$$S_t = \rho S_t^* + (1 - \rho)S_{t-1} + \varepsilon_t. \quad (18)$$

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

$$S_t^* = \theta X_t. \quad (19)$$

In 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 0.825. Ramey and West show that, given (18) and (19), this

²³Maintaining the assumptions of technology shocks and perfect competition, a reduced-form model where inventories are assumed as a factor of production exhibits the same countercyclical inventory-to-sales ratio when solved in general equilibrium. Details are provided in an appendix available at <http://www.juliathomas.net/inventoryappendices>.

autocorrelation is equal to $(1 - \rho)$, which would imply an adjustment rate of $\hat{\rho} = 0.175$ for our model economy. Note that this lies well inside the range of previous empirical estimates from aggregate data. However, interpreting the number of firms adjusting in our model as the aggregate adjustment rate, this is only about one-half of the true value, 0.27. If we instead weight firms by the difference between their target and actual end-of-period inventories, $[s_t^* - m_t(s_t^*)] - [s_t - m_t(s_t)]$, and compute the fraction of this difference closed in the aggregate, the implied adjustment rate rises to 0.35.

There are several reasons why the persistence of the inventory-sales relation does not reveal the true average adjustment rate in our model economy. First, equation (18) 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:

$$S_{t+1} = \rho_t S_{t+1}^* + (1 - \rho_t) S_t + \int \left[1 - H(\xi^T(s; z_t, A_t)) \right] (s - m_t(s) - S_t) \mu_t(ds). \quad (20)$$

Equation (20) 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 (18). A second reason that equation (18) 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 the 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 Robustness

In the preceding sections, we have developed the first quantitative general equilibrium model with endogenous inventory accumulation successful in reproducing the cyclical regularities involving inventories. Our equilibrium business cycle model simultaneously delivers procyclical inventory investment, the comovement of inventory investment and sales and the greater volatility of total production relative to sales, without forfeiting capital as a competing accumulable stock. Thus, it provides a unique laboratory within which to formally assess several prominent claims regarding the cyclical implications of inventories. Using it as such, we have found that endogenous inventory accumulation does not amplify aggregate fluctuations, nor do smaller average stocks imply reduced GDP volatility, and nor is it necessary to assume imperfect competition in order to reconcile the observation of a countercyclical inventory-to-sales ratio with technology-shock-driven business cycles. These findings represent a direct challenge to previous claims in the literature, and hence may be

viewed with some initial scepticism. In this section, we consider whether they are likely to survive in more elaborate versions of the model that sharpen its empirical performance.

Our baseline model’s relative simplicity makes it useful for understanding the inter-related dynamics of final sales and inventory investment in an actual economy. However, it also limits the model’s quantitative fit in two notable respects: (i) the relative price of inventories is too countercyclical, and (ii) the relative variability in net inventory investment is somewhat weak. In assessing the relevance of our findings, it is useful to know whether they would still arise if these two model-generated series were made to more closely resemble the data.²⁴

We explore each of these issues in table 6. First, we consider how the model’s fit to the cyclical behavior of the relative price of inventories might be improved, beginning with an explanation of the problem. To facilitate analysis in the sections above, we have assumed aggregate fluctuations in our baseline inventory economy arise from a single source, technology shocks that directly affect only the production of intermediate goods. Under this single-shock assumption, the largest rises in GDP occur in times of the largest declines in the marginal cost of producing intermediate goods. Thus, our baseline model generates a contemporaneous correlation between q and GDP near -1 . By contrast, the empirical correlation is quite weak, at about -0.257 . As we suggested in section 6.1, multiple shocks may largely offset each other’s influence on this correlation in the actual economy.

To see what information has been lost with our abstraction, we add an independent shock to the productivity of final goods firms, and we target the parameters of the Markov process governing this shock to match the cyclical dynamics of q (while maintaining all existing parameter values). Row 3 of table 6 reveals that adding the second shock is successful in this respect. Moreover, it improves the fit to the inventory facts by reducing the correlations of inventory investment with GDP and final sales. As to our model’s implications about the cyclical role of inventories, the final columns show that the inventory-sales ratio is again strongly countercyclical, and the addition to GDP volatility with the presence of inventories remains minimal, at 5 basis points.

Turning to the second issue, we believe that our model reproduces only 64 percent of the empirical relative volatility of inventory investment because it abstracts from stage-of-completion distinctions across stocks held in the actual economy. As our inventories represent stocks of an intermediate good, we have calibrated the share parameter governing their usefulness in production, θ_m , at 0.5 (the average share to intermediate goods in manufacturing and trade). At the same time, to study the effect of inventory accumulation on cyclical fluctuations in GDP and other broad aggregates, we have calibrated the fixed costs causing inventories in our model to imply an overall stock matching *total* private nonfarm inventories in the data, which include stocks of finished goods. This creates a tension for our calibration of θ_m .

Consider a retail firm producing with the technology $y = m^{\theta_m} n^{\theta_n}$, where m is a finished good

²⁴We thank two anonymous referees for raising each of these questions to our attention. Additional robustness exercises from a previous draft of this paper are available at <http://www.juliathomas.net/inventoryappendices>.

drawn from its inventories. For such a firm, it is likely that θ_m is closer to 1. As such, our baseline value represents a conservative lower bound for the overall production share of goods held as inventories. It would seem that further substantive progress toward understanding inventories may require the construction of a model with both intermediate and finished goods inventories, and thus far more heterogeneity. Here, to explore whether such an extension is likely to overturn the conclusions obtained from our relatively simple model, we consider a compromise designed to reflect some use of finished goods stocks in our final goods sector by raising θ_m to an intermediate value between 0.5 and 1 (and recalibrating our remaining parameters accordingly). Selecting the value of this share parameter at 0.65 to maximize the fit to relative inventory investment volatility in the data, we re-examine the model's predictions in row 4 of table 6.

Our proxy for the presence of distinct finished goods stocks has the desired effect, yielding a relative volatility of inventory investment roughly 92 percent that in the aggregate data, and it moves the inventory investment correlations with GDP and final sales closer to the data. Nonetheless, our conclusions about the role of inventories are again unaffected. Of course, the negative correlation between the relative price of inventories and GDP is too strong here as in our baseline results, as we have allowed no second shock process in this case. Row 5 of our table reports results for this same exercise when a second shock is included. There, the countercyclicality of q is corrected, while the correlations of inventory investment with GDP and final sales are moved even nearer the data. These improvements come at some expense to inventory investment volatility. Nonetheless, at 78 percent of the data, the model continues to reproduce a large fraction of observed aggregate inventory fluctuations. Viewing table 6 as a whole, we conclude that extensions to our model that substantially improve its quantitative match to the observed cyclical behavior of inventories and their relative price in no way alter our central findings and the challenge they present to the conventional wisdom about inventories.

9 Concluding remarks

In the preceding pages, we generalized an equilibrium business cycle model to allow for endogenous (S, s) inventories of an intermediate good. Assuming that aggregate fluctuations result from technology shocks in the intermediate goods sector, 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 almost two-thirds 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 in a counterpart model without inventory investment.

In our model, fixed delivery costs induce final goods firms to maintain stocks of intermediate goods and to increase these holdings during expansions. As a result, inventory investment co-moves

with final sales and GDP. However, in equilibrium, there is a trade-off between inventory investment and the production of final goods. Procyclical accumulation of inventories diverts both labor and intermediate goods that would otherwise be used in final production, and thereby mutes changes in final sales. This substantially limits the effect of inventory investment on the overall variability of GDP.

Our results represent a challenge to previous conclusions about the role of inventories in the business cycle. Contrary to the common belief among researchers and policymakers, we find that inventories do not amplify aggregate fluctuations. This, in turn, implies that improvements in inventory management are unlikely to explain substantial reductions in GDP volatility. Moreover, our results show that the observation of a countercyclical inventory-to-sales ratio, alongside procyclical inventory investment, is not, in itself, inconsistent with a business cycle driven by shocks to technology. While the presence of fixed adjustment costs is essential to the (S, s) policies that endogenize inventory investment in our model, general equilibrium is central to the trade-off between inventory investment and final sales. This suggests that the development of equilibrium models with different motives for inventories may lead to similar predictions, provided that such models reproduce the procyclicality of inventory investment.

References

- [1] Ahmed, S., A. Levin and B. Wilson (2002), 'Recent U.S. Macroeconomic Stability: Good Policies, Good Practices, or Good Luck' Federal Reserve Board, International Finance Division Working Paper.
- [2] Bills, M. and J. A. Kahn (2000) 'What Inventory Behavior Tells Us About Business Cycles' *American Economic Review* 90: 458-481.
- [3] Blinder, A. S. (1981) 'Retail Inventory Behavior and Business Fluctuations' *Brookings Papers on Economic Activity* 2: 443-505.
- [4] Blinder, A. S. (1990) *Inventory theory and consumer behavior*, University of Michigan Press.
- [5] 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.
- [6] Caballero, R. J. and E. M. R. A. Engel (1991) 'Dynamic (S,s) economies' *Econometrica* 59, 1659-86.
- [7] 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.
- [8] Caplin, A.S. (1985) 'The Variability of Aggregate Demand with (S,s) Inventory Policies' *Econometrica* 53: 1395-1410.
- [9] Christiano, L. J. (1988) 'Why does inventory investment fluctuate so much?' *Journal of Monetary Economics* 21: 247-80.
- [10] Eichenbaum, M. S. (1989) 'Some Empirical Evidence on the Production Level and Production Cost Smoothing Models of Inventory Investment' *American Economic Review* 79: 853-864.
- [11] Fisher, J. D. M. and A. Hornstein (2000) '(S, s) inventory policies in general equilibrium' *Review of Economic Studies* 67: 117-145.
- [12] Fitzgerald, T. J. (1997) 'Inventories and the business cycle: an overview' *Economic Review* 33 (3): 11-22, Federal Reserve Bank of Cleveland.
- [13] Hall, G. and J. Rust (1999) 'An empirical model of inventory investment by durable commodity intermediaries' Yale University working paper.
- [14] Hansen, G. D. (1985) 'Indivisible labor and the business cycle' *Journal of Monetary Economics* 16, 309-327.

- [15] Hornstein, A. (1998) 'Inventory investment and the business cycle' *Economics Quarterly* 84, 49-71, Federal Reserve Bank of Richmond.
- [16] Humphreys, B. R., L. J. Maccini, and S. Schuh (2001) 'Input and output inventories' *Journal of Monetary Economics* 47: 347-375.
- [17] Irvine, O. and S. Schuh (2002), 'Inventory Investment and Output Volatility' FRB Boston Series, paper no. 02-6.
- [18] Johnson, S. A. (1989) *Spline approximation in discrete dynamic programming with application to stochastic multi-reservoir systems* Unpublished dissertation (Cornell, Ithaca, NY).
- [19] Jorgenson, D. W., F. M. Gollop, and B. M. Fraumeni (1999) *Productivity and U.S. Economic Growth* (toExcel, New York).
- [20] Kahn, J. A. (1987) 'Inventories and the volatility of production' *American Economic Review* 77, 667-679.
- [21] 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.
- [22] 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.
- [23] Khan A. and J. K. Thomas (2004) 'Modeling Inventories Over The Business Cycle' NBER Working Paper No. 10652.
- [24] 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.
- [25] Krusell, P. and A. A. Smith Jr. (1997) 'Income and wealth heterogeneity, portfolio choice, and equilibrium asset returns' *Macroeconomic Dynamics* 1, 387-422.
- [26] Krusell, P. and A. A. Smith Jr. (1998) 'Income and wealth heterogeneity in the macroeconomy' *Journal of Political Economy* 106, 867-896.
- [27] Kydland, F. and E. C. Prescott (1982) 'Time to build and aggregate fluctuations' *Econometrica* 50, 1345 - 1370.
- [28] Maccini, L. J. and A. Pagan (2003) 'Exploring the role of inventories in the business cycle' Johns Hopkins University working paper.
- [29] 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.

- [30] Mosser, P. C. (1991) 'Trade inventories and (S,s)' *Quarterly Journal of Economics* 106, 1267-86.
- [31] Prescott, E. C. (1986), 'Theory ahead of business cycle measurement' *Quarterly Review* 10, Federal Reserve Bank of Minneapolis: 9-22.
- [32] Ramey, V. A. (1989) 'Inventories as factors of production and economic fluctuations' *American Economic Review* 79, 338-54.
- [33] Ramey, V. A. (1991) 'Nonconvex costs and the behavior of inventories' *Journal of Political Economy* 99, 306-334.
- [34] Ramey, V.A. and D. J. Vine (2004) 'Tracking the Source of the Decline in GDP Volatility: An Analysis of the Automobile Industry' NBER Working Paper No. 10384.
- [35] 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.
- [36] Richardson, H. (1995), 'Control Your Costs then Cut Them' *Transportation and Distribution*, 94-96.
- [37] Rogerson, R. (1988) 'Indivisible labor, lotteries and equilibrium' *Journal of Monetary Economics* 21, 3-16.
- [38] 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.
- [39] Schuh, S. (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.
- [40] Stock, J. R. and D. M. Lambert (1987) *Strategic Logistics Management* (Irwin, Homewood, Illinois).
- [41] Stock, J. H. and M. W. Watson (2003) 'Has the business cycle changed? Evidence and explanations' Harvard University working paper.

Table 1: Baseline parameter values*

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

*Columns 1 – 8 list parameters common across control and inventory models; 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. Columns 9 – 10 list parameters specific to inventory model; σ : per-unit inventory storage cost, $\bar{\xi}$: upper bound for fixed delivery costs.

Table 2: Distribution of final goods firms in baseline inventory steady-state

	adjustors	1	2	3	4	5	6
start-of-period distribution: $\mu(s)$		0.268	0.258	0.224	0.159	0.074	0.017
start-of-period inventories: s		1.155	0.705	0.343	0.094	0.003	0.000
fraction adjusting: $H(\xi^T(s))$		0.036	0.132	0.292	0.534	0.806	0.838
production-time inventories: s_1	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

Table 3: GDP, final sales and inventories in the postwar U.S.

	GDP	Final Sales	Net Inventory Investment
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 GDP in parentheses; columns 2 – 3 report standard deviations relative to GDP for final sales and net inventory investment.

Table 4: Inventory results for the baseline model*

	GDP	Final Sales	Net Inventory Investment	Inventory/Sales
		A: standard deviations relative to GDP*		
data	(2.237)	0.710	0.295	0.545
baseline inventory	(1.886)	0.839	0.188	0.742
		B: contemporaneous correlations with GDP		
data		0.943	0.669	- 0.381
baseline inventory		0.994	0.880	- 0.911

* Column 1 of panel A reports percent standard deviation of GDP in parentheses; columns 2 – 3 report standard deviations relative to GDP for final sales, net inventory investment and the inventory-sales ratio. Contemporaneous correlation between final sales and net inventory investment in baseline inventory model: 0.825.

Table 5: Business cycles with no inventories, baseline inventories and high inventories*

	GDP	FS	C	I	TH	K	X	M
		A: standard deviations relative to GDP						
control	(1.858)	1.000	0.384	7.470	0.674	0.423	1.579	1.579
baseline inventory	(1.886)	0.839	0.345	6.318	0.722	0.376	1.677	1.347
high inventory	(1.896)	0.823	0.338	6.296	0.733	0.373	1.688	1.322
		B: contemporaneous correlations with GDP						
control		1.000	0.902	0.970	0.969	0.118	0.998	0.998
baseline inventory		0.994	0.864	0.982	0.973	0.065	0.999	0.985
high inventory		0.994	0.856	0.981	0.975	0.066	0.999	0.985

* Column 1 reports percent standard deviation of GDP in parentheses. FS: final sales, C: consumption, I: capital investment, TH: total hours worked, K: capital stock, X: total orders of intermediate goods, M: total use of intermediate goods. In high inventory model: $\text{corr}(\text{FS}, \text{NII}) = 0.846$, $\text{corr}(\text{NII}, \text{GDP}) = 0.899$, and relative standard deviation of NII = 0.202.

Table 6: Robustness Results*

	$\sigma(\text{GDP})$	Inventory Facts				Relative Price		Implications for Claims	
		$\rho(\text{NII}, \text{GDP})$	$\rho(\text{NII}, \text{FS})$	rel. sd NII	rel. sd. FS	$\rho(q, \text{GDP})$	rel. sd. q	$\rho(\text{I-S}, \text{GDP})$	$\sigma(\text{GDP})$ Rise
1: U.S. data	2.237	0.669	0.411	0.295	0.710	-0.257	0.653	-0.381	
2: baseline model	1.886	0.880	0.825	0.188	0.839	-0.973	0.535	-0.911	0.028
3: second shock	2.223	0.671	0.567	0.160	0.901	-0.295	0.595	-0.894	0.050
4: raised intermed. goods share	1.902	0.744	0.579	0.270	0.819	-0.978	0.334	-0.861	0.044
5: second shock and raised θ_m	2.151	0.611	0.435	0.229	0.879	-0.234	0.400	-0.866	0.031

* Column 1: percent standard deviation of GDP in data and models, (rows 2 – 5 report standard deviation in each inventory model; subtract column 9 to obtain standard deviation in each corresponding control model). Columns 2-5 (Inventory Facts): contemporaneous correlations of inventory investment with GDP and final sales, respectively, then relative volatilities of inventory investment and final sales. Columns 6-7 (Relative Price): contemporaneous correlation with GDP, then relative standard deviation, of relative price of inventories. Columns 8-9 (Implications for Claims): $\text{corr}(\text{GDP}, \text{inventory-sales ratio})$, then the rise in percent standard deviation of GDP in moving from control model to inventory model. All series based on 10000 period simulations and filtered as in Table 1. Second shock in rows 3 and 5 is lognormal; $\rho_2 = 0.992$ and $\sigma_{\epsilon_2} = 0.006$ in row 3, and $\rho_2 = 0.956$ and $\sigma_{\epsilon_2} = 0.005$ in row 5.