

EXPLAINING INVENTORIES: A BUSINESS CYCLE ASSESSMENT OF THE STOCKOUT AVOIDANCE AND (S,s) MOTIVES

AUBHIK KHAN

Federal Reserve Bank of Philadelphia

JULIA K. THOMAS

*Federal Reserve Bank of Philadelphia
and
NBER*

We evaluate two leading explanations for inventories, the (S,s) and stockout avoidance motives, examining each within dynamic stochastic general equilibrium environments. We find that the (S,s) model is far more consistent with the cyclical behavior of aggregate inventories in the postwar United States when fluctuations arise from technology shocks, rather than preference shocks, whereas the converse is true for the stockout avoidance model. The (S,s) model succeeds in explaining the average magnitude of inventories in the U.S. economy and in reproducing the cyclical regularities involving inventories and other aggregate series. The stockout avoidance model does not. Even with idiosyncratic risk added to strengthen it, the stockout avoidance motive is insufficient to generate stocks near the data without destroying model performance along other important margins. Moreover, it appears incapable of sustaining inventories alongside capital. These findings suggest a fundamental flaw in reduced-form inventory models where stocks are loosely rationalized by this motive.

Keywords: Inventory Investment, (S,s) Policies, Buffer-Stock Motive, Technology Shocks, Preference Shocks

1. INTRODUCTION

Macroeconomics has seen a reawakening of interest in inventories in recent years, largely in connection with research exploring possible explanations for changes in the severity of the overall business cycle.¹ Despite this, inventories play no

We thank Satyajit Chatterjee, George Hall, Robert King and session participants at the 2004 ASSA and Midwest Macro meetings for comments and suggestions. Thomas thanks the Alfred P. Sloan Foundation for a research fellowship supporting this work, the National Science Foundation for research support under Grant 0318163, and the Minnesota Supercomputing Institute. The views expressed in this paper do not represent those of the Federal Reserve Bank of Philadelphia or the Federal Reserve System. Address correspondence to: Julia K. Thomas, Research Dept, FRB, Philadelphia, Ten Independence Mall, Philadelphia, PA 19106 USA; e-mail: mail@juliathomas.net.

role in modern quantitative business cycle theory. On those rare occasions when inventories appear in quantitative dynamic stochastic general equilibrium models, their existence is assumed.² Consequently, these existing models cannot be used to understand the cyclical role of inventories, nor how changes in inventory investment may affect other aggregate series, because the essential mechanism inducing firms to hold stocks of inputs or finished goods (rather than using or selling them) is absent.

In this paper, we seek useful models of aggregate fluctuations with inventory investment. There are three basic motives used to explain the holding of these zero return assets within the inventory literature: (i) in the production smoothing model, the motive is the avoidance of costs associated with changing production levels; (ii) in the stockout avoidance model, it is the avoidance of risk implied by a delay between the commitment of factor inputs and the realization of shocks affecting firms' marginal costs or revenues; (iii) in the (S,s) model, it is the avoidance of nonconvex costs associated with moving goods from firm to firm. For reasons explained below, we examine the latter two motives, locating conditions under which each model is able to generate the measured ratio of inventories to sales in the actual economy and exploring the ability of each to reproduce some basic inventory facts summarized in Section 2. As we assess our models' predictions for the cyclical role of inventories, we study two separate sources of aggregate fluctuations, persistent shocks either to total factor productivity or to preferences, in each case explaining the dynamic response of the model.

Our focus on alternative sources of the business cycle reflects a historical emphasis within the inventory literature. For example, as discussed by Blinder and Maccini (1991) and Ramey and West (1999), much research has been devoted to modifications of the traditional production smoothing model to allow for procyclical inventory investment in response to demand shocks. More recently, Bils and Kahn (2000) write, "Researchers have studied inventory behavior because it provides clues to the nature of the business cycle" (p. 458, column 1). Using a model where inventories are a direct input in generating sales, they find that, absent imperfect competition, the countercyclical inventory-to-sales ratio observed in aggregate and industry-level data requires a procyclical marginal cost of production, which suggests that technology shocks are relatively unimportant in the business cycle.³ Theirs is an important and influential example of research that uses differences in the dynamic response of an inventory model to different shocks to gain insight into the business cycle. Motivated by this approach, we study the extent to which each of our models is able to reproduce the salient inventory facts under technology versus preference shocks.

The sharp contrast between the predictions of the equilibrium (S,s) inventory model of Khan and Thomas (2005) and reduced-form models wherein stocks are motivated by stockout avoidance suggests that different mechanisms for generating inventory investment offer quite different predictions about the cyclical role of inventories.⁴ Thus, it may matter a great deal whether the (S,s) or stockout avoidance motive better describes the majority of inventory holdings in the economy.

Moreover, the extent to which each model is able to reproduce the inventory facts may depend crucially on whether we assume that aggregate fluctuations originate through exogenous fluctuations in total factor productivity or shocks to preferences.

Our first basic inventory model is the generalized (S,s) model of Khan and Thomas (2005). Inventories exist in the (S,s) model because firms face fixed costs of orders. To economize on these costs, firms choose to order infrequently and carry stocks of the good in question. Here we review this model's ability to reproduce the inventory facts when aggregate fluctuations arise from exogenous changes in firms' total factor productivity, and we introduce new results on the behavior of the model when the business cycle is instead driven by preference shocks. The cost-avoidance motive underlying inventories in the (S,s) model is sufficiently strong so that the model is able to reproduce the measured average inventory-to-sales ratio in the U.S. economy. When so calibrated, it goes far in matching the inventory facts if the business cycle is driven by technology shocks. However, we find that the model is substantially less successful when confronted with preference, or demand, shocks. In that case, the (S,s) model shares the same basic failings that have historically plagued the production smoothing literature: countercyclical inventory investment, a negative correlation between sales and inventory investment, and sales volatility exceeding that of production.

Although stockout avoidance has been prominent in rationalizing stocks within the recent inventory literature, the model itself has not previously seen a quantitative dynamic stochastic general equilibrium inspection. In this model, firms hold inventories because they must commit to their production plans in advance and cannot vary factor inputs in response to current shocks. Thus, in contrast to the (S,s) motive, inventories exist because they allow firms insurance against shocks to technology or marginal utility that cause changes in the equilibrium relative price of their output. Given the equity-premium puzzle [Mehra and Prescott (1985)], it is reasonable to question whether the risk motive is sufficient to yield nontrivial inventories in a quantitative DSGE version of the model. Our examination of a basic representative firm formulation indicates that it is not; specifically, the model generates zero inventory holdings in the presence of capital, and, absent capital, requires extreme variation in aggregate shocks. Thus, we are led to consider a multisector generalization of the model that adds idiosyncratic shocks to strengthen the stockout avoidance motive. We find that the inclusion of idiosyncratic risk in the generalized model mitigates the problem somewhat, but does not solve it; the gap between average stocks in the United States and those explained by the model remains large. Even in the absence of capital accumulation, either the inventory-to-sales ratio in the model is far too low, or the correlation between inventory investment and final sales is negative. Thus, the stockout avoidance model either cannot reproduce the measured level of inventories present in the actual economy or fails with regard to other important inventory facts. In each of these respects, the model is more consistent with the inventory facts when the business cycle is driven by preference, rather than technology, shocks. Nonetheless, our results suggest

that the stockout avoidance motive is too weak to provide a foundation upon which to build models designed to study the aggregate implications of inventory investment.

Our analysis excludes the canonical model of inventory investment, the production smoothing model, which has been the basis of much research in the inventory literature. We do this for two reasons. First, the production smoothing model has been extensively studied; see Ramey and West (1999) for an authoritative review of both the model and its well-established shortcomings.⁵ Second, extensions of the model that preserve its original motivation for inventory holdings have corrected its counterfactual prediction of countercyclical inventory investment by assuming convex costs of deviating from a target inventory-to-sales ratio [Ramey and West (1999)].⁶ The informal motivation for these costs is that they capture the costs associated with an increased probability of stockout for firms whose inventories deviate too far from some constant proportion of sales. Thus we study the stockout avoidance model.

The essential inability of the basic production smoothing model to explain procyclical inventory investment under demand-driven fluctuations led researchers to consider the leading microeconomic model of inventories, the (S,s) model originally solved by Scarf (1960), and it motivated Kahn's (1987) influential paper developing the stockout avoidance model. Thus, in some sense, we are studying the successors to the production smoothing model. With respect to the (S,s) model, Blinder (1981) and Caplin (1985) showed how partial equilibrium versions of the model with exogenous (S,s) bands could be consistent with the comovement of aggregate inventory investment and sales.⁷ In our calibrated (S,s) models, where the adjustment triggers vary endogenously as functions of the aggregate state of the economy, this is not a foregone conclusion; it occurs only when fluctuations are driven by productivity shocks.

2. INVENTORY FACTS

Here, we briefly review the behavior of output, final sales, and inventory investment, as well as the inventory-to-sales ratio, over the postwar business cycle. Table 1 summarizes the cyclical behavior of GDP, final sales, changes in private nonfarm inventories, and the inventory-to-sales ratio in quarterly postwar U.S. data. From this table, we arrive at a series of stylized facts that will be used to evaluate the performance of the candidate inventory models studied below.

Note first that the relative variability of inventory investment is large. In particular, though inventory investment's share of gross domestic production averages only about one-half of one percent, its standard deviation is 29.5% of that of output.⁸ Next, the series is procyclical; its correlation coefficient with GDP is 0.67. These two features of the data, the procyclicality and high variability of inventory investment, partly explain the emphasis that many researchers and policymakers have placed on examining inventories toward better understanding aggregate

TABLE 1. Postwar U.S. inventory facts

<i>Avg. inventory sales: 0.716</i>	GDP	Final sales	Net inventory investment	Inventory/sales
Percent standard deviation relative to GDP	2.237	0.710	0.295	0.545
Correlation with GDP		0.943	0.669	-0.381
Correlation with NII	0.669	0.411		

Note: Data series are domestic business GDP less housing, final sales of domestic business, changes in private nonfarm inventories and private nonfarm inventory-to-sales ratio. Data are quarterly U.S., 1953:1–2002:1, seasonally adjusted and chained in 1996 dollars. GDP, final sales, and the inventory-to-sales ratio are reported as percentage standard deviations, detrended using a Hodrick–Prescott filter with a weight of 1600. Net inventory investment is detrended as a share of GDP.

fluctuations. The positive correlation between inventory investment and final sales (0.41 for the data summarized in Table 1) is a third empirical regularity reinforcing this attention. Given the accounting identity, this is sufficient to imply that the standard deviation of production substantially exceeds that of sales (our fourth fact). As such, it has been interpreted by some as evidence that changes in inventory investment amplify the business cycle.⁹ Finally, the inventory-to-sales ratio is countercyclical; its contemporaneous correlation with GDP is -0.381 . As noted above, this has been taken as evidence against technology shocks as a primary source of the business cycle.

The above represent a core set of empirical regularities that any useful model of inventories should seek to address. Beyond these, an additional test of each candidate model we consider will be its ability to reproduce the average size of stocks in the U.S. economy, as reflected by the measured level of inventories relative to sales. Between 1951.1 and 2002.4, the real quarterly inventory-to-sales ratio averages 0.716, whereas the nominal ratio averages 0.809. As noted by Ramey and West (1999), the real ratio shows no clear trend, although the nominal ratio, now at around 0.613, has been declining since the early 1980s, when it stood at a high of 1.0.¹⁰

3. (S,s) MODEL

In this section, we generalize the Khan and Thomas (2005) (S,s) model of inventory investment to allow for both productivity and preference shocks. The first, a shift in the total factor productivity of intermediate good producers, may be interpreted as either a demand or a supply shock. Although it increases the productivity of intermediate goods firms, and hence may be seen by them as a conventional supply shock, it also increases the relative price of the final good in our general equilibrium model. In the latter respect, the shock resembles a conventional demand shock to final goods producers. By contrast, the second shock is less ambiguous. As it shifts households' marginal utility of current consumption with no direct effect

on any firm's marginal cost schedule, it is a conventional demand shock. As seen below, this shock has very different implications for the dynamics of the model.

Our description of this model follows as a planning problem, and thus appears quite different from our exposition in Khan and Thomas (2005). However, the model is identical in every respect other than its allowance for preference shocks. There are three sets of agents in the economy, a representative household and two types of perfectly competitive firms. The household purchases consumption goods from a unit measure of final goods producers, and it supplies labor both to these firms and to a representative intermediate goods producer. The intermediate goods producer combines labor and capital to supply an intermediate good used in production by final goods firms and purchases investment goods from them.

Inventories arise in the economy because final good firms face time-varying costs of undertaking orders with their intermediate goods suppliers. Because the costs are independent of order size, final goods firms choose to hold stocks of the intermediate good to reduce the frequency of their orders and hence the payments of these fixed costs.¹¹ In any given period, only a fraction of firms choose to undertake active inventory adjustment by placing an order. Thus, there is a nontrivial distribution of final goods producers, identified by their beginning-of-period inventory holding, $s \in \mathcal{S} \subseteq \mathbf{R}_+$, and fixed cost draw, $\xi \in [\underline{\xi}, \bar{\xi}]$. However, as each firm's order cost is an independent draw from a time-invariant distribution, $H(\xi)$, at each date, the aggregate state involves only the distribution of final goods firms over s , which we denote by $\mu : \mathcal{B}(\mathcal{S}) \rightarrow [0, 1]$.

The current exogenous state, z , takes on one of N_z values and follows a Markov chain with $\Pr\{z' = z_j \mid z = z_i\} \equiv \pi_{ij}^z \geq 0$, where $\sum_{j=1}^{N_z} \pi_{ij}^z = 1$ for each $i \in \{1, \dots, N_z\}$. This exogenous stochastic process may affect the marginal utility of current consumption, representing a preference shock. Alternatively, it may be a technology shock shifting the total factor productivity of the intermediate good producer.¹²

At the start of any period, each final goods firm draws a current adjustment cost, which is a time cost representing the number of labor hours that the firm must hire to undertake an order for intermediate goods, irrespective of the size of the order. Now identified by its beginning-of-period inventory holdings and cost draw, a firm of type (s, ξ) determines whether to pay its fixed cost and order the intermediate good, prior to current production. Letting $x(s, \xi)$ represent its chosen order size, the firm's available stock of the intermediate good at the time of production is $s + x(s, \xi)$. Next, at production time, the firm determines what portion of its available stock to use in current production, $m(s, \xi)$, and its labor for production, $n(s, \xi)$, and hence its output, $G(m(s, \xi), n(s, \xi))$.¹³ Intermediate goods retained for future use, $s'(s, \xi) = s + x(s, \xi) - m(s, \xi)$, incur linear storage costs, σ per unit.

Given the preceding overview of our model economy, we now describe in more detail the elements of the associated planning problem listed below in equations (1)–(7). Equation (1) limits the available quantity of final goods, Y , to the total

produced across all final goods firms, less the output lost to technological storage costs:

$$\begin{aligned}
 Y \leq & \int_{\mathcal{S} \times [\underline{\xi}, \bar{\xi}]} G(m(s, \xi), n(s, \xi)) H(d\xi) \mu(ds) \\
 & - \sigma \int_{\mathcal{S} \times [\underline{\xi}, \bar{\xi}]} [s + x(s, \xi) - m(s, \xi)] H(d\xi) \mu(ds). \tag{1}
 \end{aligned}$$

Equation (2) is the aggregate resource constraint on final goods. These goods are used for both household consumption, C , and investment by the intermediate good producer, whose capital depreciates at rate $\delta \in (0, 1)$:

$$C + [K' - (1 - \delta)K] \leq Y. \tag{2}$$

Next, from the set of equations in (3), the quantity of the intermediate good used in production by any firm of type (s, ξ) cannot exceed its stock available after its order decision, $s + x(s, \xi)$; equivalently, each firm’s stock of inventories for the next period must be nonnegative:

$$m(s, \xi) \leq s + x(s, \xi), \quad \forall s \in \mathcal{S}, \xi \in [\underline{\xi}, \bar{\xi}]. \tag{3}$$

The representative intermediate goods producer uses capital and labor, K and L , in a linearly homogeneous technology F to produce the intermediate good ordered by final good firms. Total production of the intermediate good, the left-hand side of (4), must satisfy the total quantity ordered across all final goods firms. Technology shock versions of the model allow shifts to total factor productivity in the production of intermediate goods, $D\psi(z_i) \neq 0$; under preference shocks, $\psi(z_i) = 1$.

$$\psi(z_i) F(K, L) \geq \int_{\mathcal{S} \times [\underline{\xi}, \bar{\xi}]} x(s, \xi) H(d\xi) \mu(ds). \tag{4}$$

Equation (5) constrains the household’s total hours of work, N , to be no less than its time spent in the production of the intermediate good, L , and its time spent working for final goods firms. The latter includes total time allocated to production in each firm, $n(s, \xi)$, along with that allocated to costly stock adjustments. (Recall that the nonconvex adjustment cost, ξ , is denominated in units of time.) $I : \mathbf{R}_+ \rightarrow \{0, 1\}$ is an indicator function taking on a value of 1 for firms deferring orders; $I(x) = 1$ if $x = 0$, whereas $I(x) = 0$ if $x \neq 0$.

$$N \geq L + \int_{\mathcal{S} \times [\underline{\xi}, \bar{\xi}]} (n(s, \xi) + [1 - I(x(s, \xi))]\xi) H(d\xi) \mu(ds). \tag{5}$$

Next, equation (6) describes the evolution of the distribution of final goods firms over inventory levels. For $B \in \mathcal{B}(\mathcal{S})$,

$$\mu'(B) = \int_{\{(s, \xi) \mid s+x(s, \xi)-m(s, \xi) \in B\}} H(d\xi) \mu(ds). \tag{6}$$

Equation (7) defines the indirect utility function of the representative household using a recursive representation. Here, $\beta \in (0, 1)$ represents the household subjective discount factor. Preference shock versions of the model shift the marginal utility of current consumption, $D_{13}U(C, 1 - N; z) \neq 0$.

$$V(z_i, \mu, K) = \max_{\Omega} \left[u(C, 1 - N; z_i) + \beta \sum_{j=1}^{N_z} \pi_{ij}^z V(z_j, \mu', K') \right]. \tag{7}$$

Given the current aggregate state, (z_i, K, μ) , the planner’s problem is then to choose $\Omega \equiv \{C, N, L, K', [m(s, \xi), n(s, \xi), x(s, \xi)]_{s \in \mathcal{S}, \xi \in [\underline{\xi}, \bar{\xi}]}\}$ to solve (7) subject to (1)–(6) above and the time endowment constraint, $0 \leq N \leq 1$, as well as nonnegativity constraints on $s, m(s, \xi), n(s, \xi), C, L$, and K' . Note that these constraints do not prevent final goods firms from adjusting their stocks downward; that is, $x(s, \xi)$ may be negative.

The solution to this model is somewhat involved for two reasons. First, final goods firms face occasionally binding nonnegativity constraints in selecting their future inventories, which necessitates a nonlinear solution. Next, the aggregate state vector is large, as it includes the distribution of these firms over inventory levels. These problems are simplified by solving for the competitive equilibrium directly, as described in Khan and Thomas (2005) for the case of technology shocks.

4. STOCKOUT AVOIDANCE MODELS

In contrast to the (S,s) avoidance of fixed costs seen above, the stockout avoidance model generates inventories as a means to shield firms from risk arising from the timing of their production decisions. Here, we require that firms determine their factor inputs for current production before the state is fully known. As a result, they may find that they have either too little or too much output relative to that which would be selected, were it possible, after the observation of the current state.¹⁴ In such environments, inventories allow firms some flexibility in choosing current sales after the state is revealed. We explore two general equilibrium variants of the stockout avoidance model below, beginning with a basic one-sector formulation of the model.

4.1. Basic

Our basic stockout avoidance model is essentially that developed by Kahn (1987), modified only in its allowance for capital accumulation and general equilibrium. Here, we assume a representative firm that must allocate some level of labor for current production, N , before the current aggregate state is fully known. Its production is $y = \psi(z)F(N, K)$, where, as in Section 3, $\psi(z)$ allows for shocks to total factor productivity, and K is the firm’s current capital stock, predetermined by its decisions at the previous date.

Because employment must be chosen before the current shock, z_j , is known, the firm accumulates finished goods inventories as a buffer stock for periods when either its productivity is low or households' marginal utility of consumption, and hence the demand for its output, is high. Given S and K , the aggregate stocks of inventories and capital at the start of the period, and given the realized exogenous state, z_j , total available output is $\psi(z_j)F(N, K) + S + (1 - \delta)K$. This output is used for household consumption, C_j , and for investment in future inventories and capital, S'_j and K'_j .

Define z_i as the previous date's realization of the exogenous state (useful in predicting that of the current period). Then the aggregate state at the start of the period is given by (z_i, K, S) , and the planning problem for this basic stockout avoidance model is described by the functional equation in equations (8)–(10). The choice variables, $\Omega \equiv \{N, (K'_j, S'_j, C_j)_{j=1}^{N_z}\}$, reflect the fact that investments in capital and inventories, as well as consumption, are selected after the current shock is known, while employment must be committed beforehand. As in Section 3, the inclusion of z_j in the household momentary utility function allows z to take the form of a preference shock affecting the current marginal utility of consumption.

$$V(z_i, K, S) = \max_{\Omega} \sum_{j=1}^{N_z} \pi_{i_j}^z [u(C_j, 1 - N; z_j) + \beta V(z_j, K'_j, S'_j)] \tag{8}$$

subject to

$$C_j + [K'_j - (1 - \delta)K] + S'_j \leq \psi(z_j)F(N, K) + S \tag{9}$$

$$S'_j \geq 0, \quad K'_j \geq 0. \tag{10}$$

Equation (9) shows that current output, and hence its use in consumption and investment, is constrained by the level of employment selected prior to the observation of z_j . This provides an explicit motive for inventory accumulation. Note that, in each of the parameterized examples examined below, inventories would disappear if z_j were known before N was allocated, as the real interest rate is almost never zero.

4.2. Generalized

The inventory motive in the basic stockout avoidance model derives from the variability of an aggregate shock. As will be seen below, this has stark implications for the relation between the volatility of aggregate production and the average level of inventories in the economy. To alleviate this problem, we now generalize the basic model, introducing idiosyncratic shocks across firms, in order to strengthen the stockout avoidance motive.

We assume that there are now three types of firms, each identified with a distinct good. First, there are two sets of intermediate goods producers. Each period, the relative price of the good produced by the first set is affected by an exogenous

shock, γ , whereas the exogenous shock affecting the second set takes on a value of $1 - \gamma$.¹⁵ As γ affects the relative prices of the intermediate goods, it is easily interpretable as an idiosyncratic demand shock faced by the producers of these inputs. Both the aggregate shock, denoted z as above, and the idiosyncratic shock, γ , follow Markov chains. We define $\Pr\{\gamma' = \gamma_l \mid \gamma = \gamma_k\} \equiv \pi_{kl}^\gamma \geq 0$, where $\sum_{l=1}^{N_\gamma} \pi_{kl} = 1$ for each $k \in \{1, \dots, N_\gamma\}$.¹⁶

As before, all firms are perfectly competitive. Production of the final good uses the output of both sets of intermediate goods firms in a constant-returns to scale production function, $G(m_1, m_2; \gamma)$. Its allocation to current consumption and capital investment is constrained in the aggregate by

$$C_{j,l} + [K'_{j,l} - (1 - \delta)K] \leq G(m_{1,j,l}, m_{2,j,l}; \gamma_l). \tag{11}$$

Intermediate goods firms of each type a ($a = 1, 2$) produce using capital, K_a , and labor, N_a , in a constant-returns technology F . Their output, along with any stocks they currently hold (s_a), may be used immediately in the production of final goods or stored in inventory for future use, subject to the constraints

$$m_{a,j,l} + s'_{a,j,l} \leq \psi(z_j)F(N_a, K_a) + s_a \quad \text{for } a = 1, 2, \tag{12}$$

where $m_{a,j,l}$ denotes goods of type a used in the current period (given realized state j, l).

Each intermediate goods firm must determine its current employment $N_a \in (0, 1)$, as well as its capital rental $K_a \in R_+$, before the aggregate and idiosyncratic shocks are known. However, the level of each intermediate good used in production of the final good, $m_{a,j,l}$, and thus current consumption, $C_{j,l}$, and aggregate capital investment, $K'_{j,l} - (1 - \delta)K$, as well as inventory investment in each type of intermediate good, $s'_{a,j,l} - s_a$, are chosen after the realization of the shocks. Thus, as in the basic stockout avoidance model above, inventories allow the economy to buffer against shocks to productivity or preferences that are only known after labor and capital have been allocated for production. In addition to the type-specific restrictions noted above, these precommitted factor allocations must satisfy the following aggregate constraints:

$$K_1 + K_2 \leq K, \tag{13}$$

$$N_1 + N_2 \leq 1. \tag{14}$$

The planning problem for this generalized formulation of the stockout avoidance model is listed below. Here, the aggregate state vector includes aggregate capital and the stocks of intermediate goods of each type held as inventories, s_1 and s_2 , as well as the previous date's realizations of z and γ , given their usefulness in predicting current values of these shocks. As before, z may affect either total factor

productivity in (12) or households' marginal utility of consumption in (15).

$$\begin{aligned}
 V(z_i, \gamma_k, K, s_1, s_2) = & \max_{\Omega} \sum_{j,l} \pi_{ij}^z \pi_{kl}^\gamma [u(C_{j,l}, 1 - N_1 - N_2; z_j) \\
 & + \beta V(z_j, \gamma_i, K'_{j,l}, s'_{1,j,l}, s'_{2,j,l})], \tag{15}
 \end{aligned}$$

subject to (11)–(14), $K'_{j,l} \geq 0$, and $s'_{a,j,l} \geq 0$ for $a = 1, 2$, where the choice set Ω is

$$\Omega \equiv \{N_1, N_2, K_1, K_2, (m_{1,j,l}, m_{2,j,l}, s'_{1,j,l}, s'_{2,j,l}, K'_{j,l}, C_{j,l})_{j,l=1}^{N_z, N_y}\}.$$

5. PARAMETER VALUES

We must begin this section by noting that, although the (S,s) model is calibrated, the stockout avoidance models we study here are instead parameterized examples. At present, these are too stylized to allow useful calibration. (The reason for this will be clear below.) Thus, we set most common parameters to the values selected for the (S,s) model, and then, in the case of the generalized stockout avoidance model, we select the parameters governing idiosyncratic shocks and elasticity of substitution in final goods production to maximize the model's fit to the inventory facts described in Section 2.

5.1. (S,s) Model

In calibrating the (S,s) inventory model, we choose the length of a period as one quarter and select functional forms for production and utility as follows. We assume that intermediate goods producers have a Cobb–Douglas production function with capital share α , and, where applicable, that their total factor productivity $\psi(z)$ follows a Markov chain with two values, $N_z = 2$, that is itself the result of discretizing an estimated lognormal process for technology with persistence ρ and variance of innovations, σ_ε^2 . Final goods firms also have Cobb–Douglas technology, $G(m, n) = m^{\theta_m} n^{\theta_n}$, with intermediate goods share θ_m . The adjustment costs that provide the basis for inventory holdings in our model are assumed to be distributed uniformly with lower support 0 and upper support $\bar{\xi}$. Finally, we assume that households' period utility is the result of indivisible labor decisions implemented with lotteries [Rogerson (1988); Hansen (1985)]. For versions of the model driven by technology shocks, we assume that utility is independent of z and set $U(C, 1 - N, z) = \log C + \eta \cdot (1 - N)$. For versions with preference shocks, we assume instead that $U(C, 1 - N, z) = z \log C + \eta \cdot (1 - N)$.

Aside from parameters associated with preference shocks, the calibration described here is identical to that described in Khan and Thomas (2005). If we set $\bar{\xi} = 0$, there are no fixed costs of adjustment, and the result is a benchmark model where no firm has an incentive to hold inventories. The parameters of this benchmark model, $(\alpha, \theta_m, \theta_n, \delta, \beta, \eta, \rho, \sigma_\varepsilon^2)$, are derived according to standard calibration methods, as in Prescott (1986). The resulting parameter values are

then used for the inventory model with positive adjustment costs. This approach is necessitated by the occasionally binding nonnegativity constraints on inventory holdings, which preclude the possibility of a linear solution method for the inventory model. Given the distribution of final goods firms over inventory levels, the nonlinear solution required is expensive, prohibiting calibration by simulation. Thus, we instead solve the benchmark model linearly and calibrate it. The same parameter values in the inventory model imply very similar or identical average values for the capital-to-output ratio, the share of intermediate goods in production, labor’s share of output, the investment-to-capital ratio, the real interest rate, and hours worked.

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 taken to yield the average ratio of investment to business capital over the same period. The distinguishing feature of this benchmark model, relative to the indivisible labor economy of Hansen (1985), is the presence of intermediate goods. The single new parameter implied by the additional factor of production, the share term for intermediate goods, is selected to match the value implied by the updated Jorgenson, Gollop, and Fraumeni (1999) input–output data from manufacturing and trade. From this data set, we obtain an annual weighted average of materials’ share across 21 two-digit manufacturing sectors and the trade sector, averaged over 1958–1996, at 0.499. The remaining production parameter, θ_n , is taken to imply a total labor’s share 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% per year in the steady state of the model, and η is chosen so that average hours worked are one-third of available time. Resulting parameter values are listed in Table 2.

For versions of the model with a technology shock, we determine the stochastic process for total factor productivity $\psi(z)$ using the Crucini residual approach described by King and Rebelo (1999). A continuous shock version of the benchmark model, where $\log \psi(z') = \rho \log \psi(z) + \varepsilon'$ with $\varepsilon' \sim N(0, \sigma_\varepsilon^2)$, 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

TABLE 2. (S,s) inventory calibration

β	η	α	θ_m	θ_n	δ	ρ	$\sigma_\varepsilon^{\text{tech}}$	$\sigma_\varepsilon^{\text{pref}}$	σ	$\bar{\xi}$
0.984	2.128	0.374	0.499	0.328	0.017	0.956	0.015	0.020	0.012	0.220

Note: 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 : aggregate shock persistence, $\sigma_\varepsilon^{\text{tech}}$: standard deviation of innovations for model driven by technology shocks, $\sigma_\varepsilon^{\text{pref}}$: standard deviation of innovations for model driven by preference shocks, σ : per-unit inventory storage cost, $\bar{\xi}$: adjustment cost upper bound.

$Y = \pi_z(\rho)\psi(z) + \pi_k(\rho)k$, where the coefficients associated with z and k are functions of ρ . Rearranging this solution, data on GDP and capital are then used to infer an implied set of values for the technology shock series. 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. The resulting values for the persistence and variance of the technology shock process are not uncommon: $\rho = 0.956$ and $\sigma_\varepsilon = 0.015$.

For versions of the (S,s) model driven by preference shocks, we repeat the procedure described above, under the assumptions that $\psi(z) = 1$ and the shock to marginal utility now follows $\log z' = \rho \log z + \varepsilon'$ with $\varepsilon' \sim N(0, \sigma_\varepsilon^2)$. This yields the same estimated persistence, $\rho = 0.956$, but higher variability in the innovations, $\sigma_\varepsilon = 0.020$.

The two parameters that distinguish the (S,s) inventory model from the benchmark are the storage cost associated with inventories and the upper support for adjustment costs (uniformly distributed on $[0, \bar{\xi}]$). Conventional estimates of inventory storage costs (or *carrying costs*) average 25% of the annual value of inventories held [Stock and Lambert (1987)]. Excluding those components accounted for elsewhere in our model (for instance, the cost of money reflected by discounting) and those associated with government (taxes), we calibrate σ to yield storage costs at 12% of the annual value of inventories. In our calibrated model, where the steady-state value of the relative price of intermediate goods is 0.417, this implies a proportional cost of $\sigma = 0.012$. Next, using NIPA data, we compute that the quarterly real private nonfarm inventory-to-sales ratio has averaged 0.7155 in the United States between 1947:1 and 2002:1. Given the storage cost parameter σ , we select the upper support on adjustment costs, $\bar{\xi}$, at 0.220 to reproduce this average inventory-to-sales ratio in the steady state of our model.

5.2. Stockout Avoidance Models

For the basic stockout avoidance model, we use the same utility function and stochastic processes for the technology and preference shocks as calibrated above. Although this model allows for capital accumulation, we find that, in its presence, rate of return dominance drives inventories out of the economy, given even high levels of aggregate uncertainty. Thus, for the results presented here, we eliminate capital by assuming that the production of final goods is $F(N) = N^{1-\alpha}$, and we set $\alpha = 0.36$ to imply a labor's share of 0.64, as in the (S,s) model.

Table 3 summarizes the baseline parameters for the generalized stockout avoidance model with two sets of intermediate goods firms. We continue to assume the same utility function and stochastic process for the aggregate technology and preference shocks as before. Further, we assume that the final goods production function is a CES aggregate of the two intermediate goods with stochastic shares,

TABLE 3. Generalized stockout avoidance baseline parameterization

β	η	$1 - \alpha$	$1/(1 - \phi)$	δ	ρ	$\sigma_{\varepsilon}^{\text{tech}}$	$\sigma_{\varepsilon}^{\text{pref}}$	ρ_{γ}	σ_{γ}/σ_z
0.984	2.128	0.640	1.000	0.017	0.956	0.015	0.020	0.750	1.000

Note: β : household subjective discount factor, η : preference parameter for leisure, $1 - \alpha$: labor’s share in production, $1/(1 - \phi)$: elasticity of substitution between intermediate goods in final production, δ : capital depreciation rate, ρ : aggregate shock persistence, $\sigma_{\varepsilon}^{\text{tech}}$: standard deviation of innovations for model driven by technology shocks, $\sigma_{\varepsilon}^{\text{pref}}$: standard deviation of innovations for model driven by preference shocks, ρ_{γ} : idiosyncratic shock persistence, σ_{γ}/σ_z : standard deviation of idiosyncratic, relative to aggregate, shocks.

where γ takes on one of two values, $N_{\gamma} = 2$:

$$G(m_{1,j,l}, m_{2,j,l}; \gamma_l) = [\gamma_l m_{1,j,l}^{\phi} + (1 - \gamma_l) m_{2,j,l}^{\phi}]^{\frac{1}{\phi}}.$$

As already noted, it is difficult to generate inventory–sales ratios close to the data in the basic stockout avoidance model, where aggregate risk motivates inventory holdings. Although mitigated somewhat by the addition of idiosyncratic risk, this problem persists in our generalized model and is compounded by the presence of a competing asset with positive rate of return. Thus, for the results here, we again exclude capital and assume that intermediate goods firms produce according to $F(N) = N^{1-\alpha}$, where $\alpha = 0.36$.¹⁷ Despite this, we find that the model is only successful in generating a sufficiently large level of inventory holdings in cases with (i) high substitutability between the two types of intermediate goods (ϕ near one) and (ii) a highly variable idiosyncratic shock, γ . Unfortunately, in such cases, the resulting inventory investment series is far more variable than that seen in the data, and it is negatively correlated with final sales. (We discuss these results further in Section 6.3.)

Given the difficulty of reproducing the measured inventory-to-sales ratio without doing violence to the model’s second moments, we instead select the substitutability of intermediate goods and the stochastic process for γ to best match these second moments irrespective of the ratio. In particular, we assume a Cobb–Douglas aggregator ($\phi = 0$), and we select the persistence and variability of γ to best reproduce the variability of net inventory investment relative to GDP, the procyclicality of inventory investment, and its correlation with final sales. This leads us to set the persistence of γ at 0.75 and a standard deviation relative to that of the aggregate shock equal to 1. Note that this is a very different approach to determining the parameters governing inventory investment from that pursued for the (S, s) model. There, we selected the range of fixed costs to reproduce the measured inventory-to-sales ratio, an approach allowing us to formally evaluate the extent to which the inventory model is able to reproduce the observations of Table 1. Here, by contrast, we must instead interpret our baseline set of results as providing an upper bound on the extent to which the stockout avoidance model can explain the level of inventories held in the economy, subject to the remaining inventory facts.

TABLE 4. Calibrated (S,s) inventory model

	GDP	Final sales	Net inventory investment	Inventory/sales	Total hours
(A) Technology shocks (<i>avg. inventory/sales: 0.716</i>)					
Percent standard deviation relative to GDP	1.598	0.859	0.158	0.807	0.666
Correlation with GDP		0.997	0.906	-0.933	0.972
Correlation with NII		0.869			
(B) Preference shocks (<i>avg. inventory/sales: 0.716</i>)					
Percent standard deviation relative to GDP	1.692	1.026	0.051	0.818	1.502
Correlation with GDP		0.999	-0.495	-0.999	1.000
Correlation with NII		-0.532			

6. RESULTS

6.1. (S,s) Model

Table 4 summarizes the results of a 5,000-period simulation for both versions of the equilibrium (S,s) inventory model. In panel A, the business cycle is driven by technology shocks, whereas in panel B, it arises from shocks to the marginal utility of consumption (preference shocks). Again, the results for this model are distinct from those of the stockout avoidance models below in that they involve a calibrated average inventory-to-sales ratio. Thus we are in a position to ask to what extent the model is able to reproduce the cyclical regularities involving inventories, given its calibration to match the average presence of these stocks in the U.S. economy.

We begin by examining the results for this model under the assumption that cyclical fluctuations are driven by changes in total factor productivity, in panel A. The first row of the table reports percentage standard deviations for each series relative to that of GDP.¹⁸ Contemporaneous correlations with GDP are listed in the second row. Together, these two rows establish that the (S,s) inventory model, under technology shocks, is successful in reproducing both the procyclicality of net inventory investment and the higher variance of production when compared to final sales. The latter arises from the positive correlation between inventory investment and final sales, 0.87, in the simulated economy. Further, this simple model with nonconvex factor adjustment costs as the single source of inventory accumulation is able to explain 54% of the measured relative variability of net inventory investment. Finally, contrary to the findings of Bils and Kahn (2000), the inventory-to-sales ratio in this model driven by technology shocks is countercyclical, as in the data. In fact, perhaps the model's most pronounced quantitative departure from the data is the exaggerated countercyclicality of this ratio. This arises from the

overly countercyclical relative price of intermediate goods that is used to value inventories, which in turn results from the single productivity shock that directly affects only the producers of intermediate goods.

A persistent positive shock in this economy results in a persistent fall in the price faced by final goods firms for the intermediate goods used in their production, which increases both aggregate orders and use of these goods, as well as demand for the complementary labor input, thus raising production of final goods, which is final sales.¹⁹ The procyclicality of net inventory investment, as well as its comovement with final sales, arises from the fact that final goods firms raise their orders by more than their use of intermediate goods in response to such a shock. Given the shock's persistence, final goods firms anticipate increased demand for their output to persist, given raised demand for investment by intermediate goods producers and consumption by households smoothing the effects of their increased permanent income. To accommodate this, while avoiding payment of fixed order costs again in nearby dates, currently ordering firms retain an increased portion of their increased stock for use in future production. This precautionary accumulation by an increased number of ordering firms is large relative to the rise in intermediate goods use among nonordering firms; thus aggregate inventory investment increases alongside rises in the production of intermediate and final goods. Finally, given that procyclical inventory investment diverts intermediate goods from current final goods production, the initial rise in final sales is smaller than that in total production.

Turning to panel B, we next discuss the effects of preference shocks in the (S,s) model. As reviewed by Blinder and Maccini (1991), the leading macroeconomic model of inventory investment, the production smoothing model, predicts that production is less variable than sales when there are shocks to demand, which we interpret as preference shocks. They and other researchers have suggested that the (S,s) model of inventory investment might resolve this inconsistency with the data by yielding a positive covariance between sales and inventory investment.²⁰ Ironically, our (S,s) model of inventories fails along this and other margins when the business cycle is driven by preference shocks. Inventory investment is both countercyclical and negatively correlated with final sales. Consequently, sales are more variable than production. Thus, in contrast to its success with technology shocks, the (S,s) model driven by shocks to preferences is unable to resolve this long-standing problem in the inventory literature. Moreover, the relative variability of net inventory investment is only 5%; thus the model is able to explain only 17% of the observed variation. In fact, under preference shocks, the model's only success is its ability to generate a countercyclical inventory-to-sales ratio. However, this countercyclicality is even more overstated than it was above when the business cycle was driven by technology shocks.

A persistent positive shock to the marginal utility of current consumption increases households' willingness to work and their demand for current consumption. Consumption rises, as does current investment, due to the persistence of the shock. Given exogenous shifts in the marginal valuation of consumption relative

to leisure, the response in total hours is much sharper under preference shocks. Further, in contrast to the technology shock model, the rise in hours worked in final goods production is greater than the rise in hours worked in intermediate goods production. The sharp increase in final goods production, given households' urgency for current consumption, initially drives down firms' stocks of the intermediate good. As a result, net inventory investment has a negative contemporaneous correlation with both final sales and GDP. As the capital stock increases, investment in inventories recovers. At a four-quarter lag, its correlation with GDP is 0.54.

Although we have seen that the (S,s) model under technology shocks is able to address the inventory facts, its performance is far less successful when the business cycle is driven by preference shocks. One might question the choice of these shocks, in that they shift the marginal rate of substitution between leisure and consumption. However, we note that the natural alternative, shocks to the discount factor that do not distort this margin, yield counterfactually high variation in aggregate investment. In fact, when we drove our benchmark model with an estimated stochastic process for these shocks, the nonnegativity constraint on aggregate investment was binding several times in a 5,000-period simulation. This motivated our decision to focus instead on shocks to the marginal utility of consumption.

6.2. Basic Stockout Avoidance Model

Table 5 summarizes the results for both versions of the basic stockout avoidance model with a representative firm. In panel A, the business cycle is driven by technology shocks. In panel B, aggregate fluctuations arise from shocks to the marginal utility of consumption. In both cases, the stochastic processes for shocks

TABLE 5. Stockout avoidance: One good model ($5 \times$ variability in innovations)

	GDP	Final sales	Net inventory investment	Inventory/sales	Total hours
(A) Technology shocks (<i>avg. inventory/sales: 0.036</i>)					
Percent standard deviation relative to GDP	9.5	0.955	0.144	0.233	0.130
Correlation with GDP		0.989	0.439	0.334	0.356
Correlation with NII		0.329			
(B) Preference shocks (<i>avg. inventory/sales: 0.042</i>)					
Percent standard deviation relative to GDP	8.2	0.893	0.541	0.695	1.582
Correlation with GDP		0.831	0.487	-0.341	0.997
Correlation with NII		-0.069			

have the same persistence as those in the calibrated (S,s) model, but five times the standard deviation.

When the business cycle in this basic model arises from shocks (whether to technology or to the marginal utility of consumption) with the variability measured from the data, firms almost never hold any inventories. Simply put, without extreme variability in aggregate productivity, there is not sufficient risk in the economy to compensate for the zero net return on inventory investment. Stocks are accumulated only in those few periods when there is a large change in total factor productivity or the marginal utility of consumption. The mean inventory-to-sales ratio is less than 0.001 in both cases; essentially, inventories do not exist. When we increase the variability of the shock fivefold, average inventory-to-sales ratios rise above 0.035 in each model. Although this is still roughly 20 times lower than in the data, inventories are at least not so rare. Hence, we report results for these cases.

For the technology shock model, in panel A, the percentage standard deviation of GDP, 9.5, is more than four times its empirical counterpart.²¹ However, given the unit intertemporal elasticity of substitution assumed in household preferences, alongside the necessary absence of capital accumulation, there is almost no movement in employment. The variability of hours worked, relative to that of GDP, is 0.13. As a result, almost all the variation in output is directly due to changes in technology. The lack of sufficient hours variability in response to changes in the marginal product of labor underlies the mild procyclicality of inventory investment; its contemporaneous correlation with GDP is 0.439, and its relative variability is slightly below half that in the data. Moreover, as inventory investment is essentially uncorrelated with final sales, the variability of production only slightly exceeds that of sales. Finally, the inventory-to-sales ratio is procyclical.

For the preference shock model, by contrast, there is considerable movement in hours worked in response to shocks shifting the marginal rate of substitution between consumption and leisure, as seen in panel B. There, relative variability of total hours worked rises to 1.58. (The corresponding value in postwar U.S. data is roughly 0.95.) This in turn raises the relative variability of inventory investment by a factor of almost 4. At 0.541, it substantially exceeds its empirical counterpart. When a positive shock occurs in this model economy, recalling that hours, and hence output, cannot respond immediately, unexpectedly high demand for consumption (final sales) initially reduces inventory investment, thereby yielding a temporary decline in the ratio of inventories relative to sales. However, given the persistence of the rise in the marginal rate of substitution of leisure for consumption, labor inputs respond sharply in the next period, raising both output and inventory investment. As a result, production in this economy is substantially more variable than sales, and the inventory-to-sales ratio has about the same correlation with GDP as in the data. Unfortunately, these results arise in a model that, by contrast to the (S,s) model, can only explain a small fraction of the stock of inventories held in the economy. Moreover, achieving even this low average inventory-to-sales ratio requires a variability of GDP more than three and a half times larger than that in the data.

TABLE 6. Generalized stockout avoidance model and inventory dynamics

	GDP	Final sales	Net inventory investment	Inventory/sales	Total hours
(A) Technology shocks (<i>avg. inventory/sales: 0.043</i>)					
Percent standard deviation relative to GDP	1.8	0.993	0.197	0.266	0.121
Correlation with GDP		0.980	0.148	-0.011	0.240
Correlation with NII		-0.049			
(B) Preference shocks (<i>avg. inventory/sales: 0.074</i>)					
Percent standard deviation relative to GDP	1.6	0.824	0.473	0.470	1.547
Correlation with GDP		0.884	0.582	-0.730	0.979
Correlation with NII		0.135			

We conclude that the representative firm stockout avoidance model, even in the absence of capital, is incapable of producing inventory holdings, whether the business cycle arises from technology or preference shocks. This motivates our examination of the generalized stockout model below.

6.3. Generalized Stockout Avoidance Model

Baseline parameterization. Panels A and B of Table 6 report results for the baseline parameterization of the generalized stockout model, as described in Section 5.2. Here we use the measured stochastic processes for technology shocks, in panel A, and for shocks to the marginal utility of consumption, in panel B. The principal result is that now, with a reasonable variance for the aggregate shock and thus GDP, the generalized stockout model exhibits a nontrivial level of inventories. The mean inventory-to-sales ratios (0.043 in the case of technology shocks and 0.074 under preference shocks), although still quite low relative to the data, exceed those generated by the basic model above with five times the volatility in aggregate shocks.

Aside from its greater ability to sustain inventories alongside plausible GDP volatility, the technology shock model in panel A exhibits only minor improvements relative to its earlier representative firm counterpart (in Table 5A). There is still almost no movement in total hours worked, and, as a result, the relative variability in inventory investment is too low. Sales are only slightly less variable than production, the inventory-to-sales ratio is essentially acyclical, and the positive correlation between final sales and inventory investment is too weak to survive the HP-filter.

When a persistent positive shock to the productivity of intermediate goods producers unexpectedly raises available intermediate goods in this economy, the increase is almost entirely used to raise current consumption (final sales), and

there is essentially no propagation.²² In the case of the relatively more productive intermediate good (given γ), all extra output is used immediately and no inventories are held. In the case of the other intermediate good, a small portion of the increased production is retained as increased stock to be used (immediately) when its idiosyncratic state switches. This alone generates a minor rise in inventory investment at the date of the aggregate productivity shock. Thereafter, given only a very small rise in labor hours in subsequent dates, these initial features are essentially unchanged. GDP rises marginally above its impact date value for a short time, and consumption (final sales) absorbs most of the increase, given only small and gradual accumulation of the relatively less productive of the two intermediate goods.

In the absence of capital accumulation, to correct the lack of labor responsiveness that appears largely responsible for this model's poor performance, we find that we must move away from our current specification of preferences (logarithmic in consumption). When we reduce the elasticity of intertemporal substitution below 1, the results improve in several respects. Inventory investment becomes more procyclical and more strongly correlated with final sales, reducing the relative variability of final sales. In addition, the average level of inventories, relative to sales, rises somewhat. Unfortunately, these improvements come at the expense of an important business cycle regularity. In the absence of capital investment, the income effect on leisure dominates, so that hours move countercyclically. By contrast, although raising the elasticity of substitution does yield a strongly procyclical and volatile hours series, it causes the relative volatility in inventory investment to drop nearly to zero.

In panel B, we examine the generalized stockout avoidance model under shocks to marginal utility. Here, we find that hours worked become much more responsive, as was the case in the model's representative firm counterpart above (in Table 5B). In this case, the generalized stockout model can explain about half of the excess variability of production relative to sales, which is slightly more than that explained by the (S,s) model driven by technology shocks. Moreover, the relative variability of net inventory investment continues to actually exceed that observed in the data. Finally, the inventory-to-sales ratio becomes strongly countercyclical, and there is a weak positive correlation between final sales and net inventory investment.

Consider the effects of a persistent negative shock to the marginal utility of consumption in this economy, and suppose that the idiosyncratic shock has been at its high value for some periods, so that intermediate goods of type 1 have been, and continue to be, more useful in final goods production than goods of type 2. (In this case, there is no existing stock of good 1 at the date of the aggregate shock.) In midperiod, when the shock is observed, labor inputs, and hence production of both types of intermediate goods, have already been determined. Thus, total hours are initially unaffected, as is total production, though not its allocation. Given unanticipated low demand for consumption, use of intermediate goods in the final goods sector falls short of their production, causing surprise accumulation of both s_1 and s_2 , implying a rise in net inventory investment.²³ Thus, final sales and

inventory investment initially move in opposite directions, and the inventory-to-sales ratio rises sharply. In the next date, however, given anticipated persistence in the low aggregate state, the labor hired to produce each type of intermediate goods is reduced substantially (although n_1 continues to exceed n_2), and hence so is GDP. In the case of intermediate good 1, reduced production completely offsets the surprise inventory accumulation of the previous date; these intermediate goods suppliers sell their entire stock and maintain no inventories again until the idiosyncratic state switches. Suppliers of intermediate good 2 undertake similar decumulation; however, their reversals are more gradual, given good 2's lesser current usefulness together with the anticipation of a rise in its relative price at some future date when the idiosyncratic state switches. Nonetheless, the aggregate effect of these stock reductions is a sufficiently large decline in inventory investment accompanying the drop in total production to explain the strong positive correlation between these two series. Further, despite the initial rise in inventory investment, its sharp decline at the subsequent date with decreased production, alongside continued low consumption, is sufficient to produce the weakly positive correlation between final sales and inventory investment.

High inventory parameterizations. As we noted above, even when generalized to allow for idiosyncratic risk (and absent capital accumulation), the stockout avoidance model continues to have difficulty in generating empirically viable inventory holdings. In our baseline results, we chose to focus on a parameter set that could best match the cyclical behavior of inventories subject only to some minimal average level of these stocks. In Table 7, we present a series of results showing that it is possible to generate average inventory holdings similar to the data from this model by combining high substitutability of intermediate goods in final goods production, together with a highly variable idiosyncratic shock, but that this improvement comes at the cost of remaining inventory regularities, most notably the comovement of final sales and inventory investment.

For the sake of comparison, each panel in Table 7 presents results under only one change to the baseline parameterization above. In panels A and B, we reexamine the technology shock model of Table 6A with a large rise in the variability of γ and a rise in the elasticity of substitution between intermediate goods, respectively. Panels C and D repeat this exercise for the preference shock model of Table 6B.

From panels A and C, we see that, if the idiosyncratic shock's volatility were raised high enough above that of the aggregate shock, the increased γ risk would eventually be sufficient to yield average inventory-to-sales ratios near 0.716, the average value in the data. Moreover, comparison of these two panels indicates that the preference shock model would require a lesser such rise than the model driven by technology shocks. However, in either case, we can infer from these panels that the required rise would yield extreme GDP volatility relative to the data and a strong negative correlation between final sales and inventory investment. (In the case of preference shocks, with a rise in γ volatility yielding an average

TABLE 7. Raising average inventory/sales in the generalized stockout avoidance model

(A) Technology shocks with high idiosyncratic volatility: $sd(\gamma)/sd(z) = 8/3$					
<i>Avg. inventory/sales: 0.270</i>	GDP	Final sales	NII	Inventory/sales	Total hours
Relative percent std. dev.	2.5	0.947	1.098	1.514	0.790
Correlation with GDP		0.402	0.584	-0.290	0.586
Correlation with NII		-0.509			
(B) Technology shocks with high substitutability: $1/(1 - \phi) = 3$					
<i>Avg. inventory/sales: 0.105</i>	GDP	Final sales	NII	Inventory/sales	Total hours
Relative percent std. dev.	1.8	0.931	0.388	0.375	0.105
Correlation with GDP		0.924	0.368	-0.084	0.455
Correlation with NII		-0.015			
(C) Preference shocks with high idiosyncratic volatility: $sd(\gamma)/sd(z) = 8/3$					
<i>Avg. inventory/sales: 0.466</i>	GDP	Final sales	NII	Inventory/sales	Total hours
Relative percent std. dev.	3.0	1.010	1.650	2.549	1.602
Correlation with GDP		-0.226	0.783	-0.404	0.968
Correlation with NII		-0.782			
(D) Preference shocks with high substitutability: $1/(1 - \phi) = 3$					
<i>Avg. inventory/sales: 0.209</i>	GDP	Final sales	NII	Inventory/sales	Total hours
Relative percent std. dev.	1.7	0.994	1.119	1.265	1.511
Correlation with GDP		0.394	0.556	-0.669	0.958
Correlation with NII		-0.545			

inventory-to-sales ratio at about 65% of that in the data, this correlation is already sufficiently negative to make GDP less volatile than final sales.) In both of these cases, the effects of aggregate shocks are similar to those in the baseline results discussed above, and continue to be distinct for technology versus preference shocks. However, unlike the baseline cases, the combination of high relative variability in the idiosyncratic shock, alongside the increased average stocks it implies, makes changes in γ now relevant to aggregate dynamics.²⁴

To understand the negative correlation between sales and inventory investment with high idiosyncratic variability, consider the effect of an unanticipated rise in γ , which makes good 1 now the more productive of the two intermediate goods. With labor inputs initially committed ($n_2 > n_1$), one might expect final sales to fall at the date of this switch in relative productivities. However, if it has been some time since the previous idiosyncratic shock, then there is a large stock of good 1 available for use in final goods production. At such times, some, but not all, of this stock is used to partly augment current production, actually yielding a rise in final sales. At the same time, excess good 2 is accumulated, but more gradually than the decumulation of good 1 stocks. Thus, alongside the rise in final sales, aggregate

inventory investment is negative. With a raised labor input, and hence production, in good 1 during the subsequent date (while good 2 production falls), final sales continues to rise. In this period, the large stock of good 1 is further reduced, and that of good 2 rises slowly. Thus, final sales and inventory investment continue to move oppositely. After several dates, as the stock of good 1 is depleted, good 2 inventories continue to rise, yielding raised inventory investment at the same time final sales begins to fall.

In panels B and D, we see that a raised elasticity of substitution between goods 1 and 2 has an effect on aggregate dynamics similar to the raised variability in γ discussed above. Here, too, there is a stronger incentive for the producers of the currently less productive intermediate good to accumulate stocks, as a switch in γ will imply larger effects on the demands for intermediate goods, and hence their relative prices, than in the baseline cases of the preceding section. Again, when the relative price switches in favor of good 1, the producers of this good draw upon their large existing stock to deliver a large rise in its use in final goods production sufficient to raise final sales, although they do not fully deplete this stock. At the same time, producers of good 2 begin to increase their stocks, but too slowly to prevent a negative inventory investment in the aggregate. Thereafter, the humped-shaped final sales response, accompanied by a u-shaped inventory investment response, is essentially that described above. In sum, although a raised elasticity of substitution across intermediate goods does raise average stocks in the model economy, it effectively makes the idiosyncratic shock more prominent in aggregate fluctuations, and thus shares the same negative aspects as the direct increase in the volatility of the idiosyncratic shock, most particularly, the negative correlation between final sales and inventory investment.

7. CONCLUDING REMARKS

In the preceding sections, we have evaluated the two leading models of inventory investment using a dynamic stochastic general equilibrium analysis. We find that the (S,s) model with capital, when calibrated to exhibit the observed average level of inventories relative to sales in the U.S. economy, is able to reproduce about half the measured variability of inventory investment when the business cycle is driven by technology shocks. Moreover, it is successful in predicting strongly procyclical inventory investment, a higher cyclical volatility in production relative to sales, a positive correlation between final sales and inventory investment, and a countercyclical inventory-to-sales ratio. By contrast, when the (S,s) model's business cycle is instead driven by shocks to preferences, the model fails in nearly all of these respects. Most notably, inventory investment is no longer procyclical, and its correlation with final sales becomes negative.

Abstracting from capital accumulation, we find that a generalized stockout avoidance model where the primary mechanism inducing firms to hold inventories is the risk associated with an idiosyncratic shock, succeeds in explaining several of the inventory facts when aggregate fluctuations result from shocks to the marginal

utility of consumption. This model generates a variability of inventory investment actually exceeding that in the data, total production more variable than sales, and a countercyclical inventory-to-sales ratio. Unfortunately, these successes are achieved only when the average inventory-to-sales ratio is about one-tenth the measured value. As seen above, the model is capable of much higher inventory-to-sales ratios, but to achieve them requires high variability in the idiosyncratic shock and a high elasticity of substitution between intermediate goods. The combination of these elements generates a strong negative correlation between inventory investment and final sales, and thus the lesser volatility of final sales relative to production is also forfeited.

The omission of capital in our results for the generalized stockout model is potentially important. First, its inclusion would allow us to pursue formal calibration of the model. Moreover, we know that, for the technology shock version of the model, introducing capital would increase the variability of hours worked under the current specification of preferences, and also maintain the procyclicality of this series in specifications involving lower elasticity of intertemporal substitution. Each of these might improve the performance of the technology shock model.

From the results of Section 6.2, we know that the inclusion of idiosyncratic risk is essential to generating inventories in the equilibrium stockout avoidance model. Recall that, in the basic representation with only aggregate risk, there were zero inventories in the presence of capital, and, in its absence, average stocks relative to sales were less than one-tenth of one percent given plausible volatility of the aggregate shock. In Section 6.3, we saw that, absent capital, the inclusion of idiosyncratic shocks mitigates this problem somewhat. Considering the inclusion of capital in this generalized model, then, the obvious question is whether idiosyncratic risk can continue to yield positive inventories in its presence. Absent additional frictions in the model, this is not likely. Returning to the formulation with capital in Section 4.2, note that capital investment allows a direct, positive return means of smoothing the effects of both aggregate and idiosyncratic shocks. Thus, it essentially eliminates any role for the zero return stocks. Moreover, increased variability in idiosyncratic shocks only translates into a raised volatility of capital investment, rather than a motive for inventories, as we have verified in simulations of the model without variable leisure.

In concluding, given the poor performance of the stockout avoidance models seen here, we must note that we began with what we saw as the most natural formulation including the stockout avoidance risk motive. Given the failure of that basic model, we then generalized it to strengthen the motive by including an additional element of risk. Obviously, we have not exhausted all possible formulations. Nonetheless, until some variant has been devised that can be demonstrated consistent with the data in a dynamic, stochastic general equilibrium setting, reduced-form inventory models appealing to this motive appear unfounded. The far more appropriate model, based on our analysis, is an (S,s) model with aggregate fluctuations arising from technology shocks.

NOTES

1. Kahn et al. (2002) suggest that improvements in inventory management methods are an important source of reduced GDP volatility, a finding that agrees well with conventional wisdom regarding the destabilizing role of inventories in the economy. Empirical analyses undertaken by Ramey and Vine (2004a) and Stock and Watson (2003) suggest otherwise, as does the theoretical analysis of Khan and Thomas (2005).

2. See, for example, Kydland and Prescott (1982) and Christiano (1988), where inventories enter as a factor of production.

3. Because inventories are required to produce sales, the two series have a strong tendency to move together in their partial equilibrium model. Thus, to generate a countercyclical inventory-to-sales ratio, Bils and Kahn find that there must be either countercyclical markups or procyclical marginal costs of production.

4. Recent examples of such models include Kahn et al. (2002), where inventories are a source of household utility, and Bils and Kahn (2000) and Coen-Pirani (2004), where they are required for sales.

5. In its original formulation, the production smoothing model assumed that firms hold inventories as a buffer against exogenous fluctuations in sales, given convex adjustment costs of varying production. This leads to the immediate prediction that sales and inventory investment are negatively correlated, which is inconsistent with both aggregate and firm-level evidence [see Schuh (1996)].

6. Alternative resolutions of the problem of countercyclical inventory investment include the introduction of relatively large technology shocks [Eichenbaum (1989)] or increasing returns in production [Ramey (1991)]. Each of these extensions reduces the relative importance of the convex costs and promotes production-bunching to generate an inventory investment series that comoves with sales. However, in doing so, they reduce the role of the very friction that causes inventories in the model.

7. Fisher and Hornstein (2000) undertake a general equilibrium analysis of an (S,s) model with time-invariant bands, arriving at similar results.

8. Note that net investment in private nonfarm inventories is detrended as a share of GDP.

9. See Khan and Thomas (2005) for further discussion on this topic.

10. Ramey and Vine (2004b) explain why the real inventory-to-sales ratio is the more appropriate series for intertemporal comparisons.

11. Because inventories here are stocks of intermediate goods, the model accommodates not only finished manufacturing goods, which are inputs in retail and wholesale trade, but also inventories of materials and supplies and work in process, which are inputs in manufacturing. Our decision to shift away from an exclusive focus on inventories of finished manufacturing goods is motivated by the larger size and variability of intermediate input inventories within manufacturing relative to finished goods. See Khan and Thomas (2005) for empirical evidence.

12. Our choice to model the technology shock as influencing the total factor productivity of only intermediate goods producers is motivated by a countercyclical relative price of private nonfarm inventory stocks in the data, as documented in Khan and Thomas (2005).

13. Note that our final goods firms may be interpreted as producers in any industry using intermediate inputs, or they may be interpreted as retail and wholesale firms that purchase essentially a finished good and combine it with labor to sell it.

14. In our equilibrium model, stockout is associated with inventories being reduced to zero.

15. As will be clear below, increases in γ raise the productivity of the intermediate good produced by the first set of firms relative to that produced by the second set.

16. As was the case with the (S,s) model, the stockout avoidance models exhibit occasionally binding nonnegativity constraints on inventories, and hence must be solved nonlinearly. Consequently, both shocks are discretized for numerical tractability.

17. For similar reasons, we exclude storage costs. Our results below will clarify the difficulty associated with capital's inclusion in this model. Once they have been presented, we will return, in Section 7, to discuss the implications of capital accumulation for the model's aggregate inventory holdings.

18. All model-generated data are treated just as the U.S. data presented in Table 1; as there, net inventory investment is detrended as a share of GDP.

19. As was noted above, our model solutions rely on discretized shocks for computational tractability. Given this discretization, our explanation of the economic dynamics in each model is based on careful reading of simulated time series rather than impulse response figures.

20. The essential assumption here is that, with an increase in sales, there will be a sufficient rise in the number of firms hitting their (S,s) adjustment triggers, and hence placing orders, to offset the declines in stocks among those that do not.

21. This arises immediately from our decision to examine cases with extreme aggregate shocks, as discussed above.

22. As in Section 6.1, we report the observed mechanics of each model following a change in the exogenous aggregate state, rather than impulse response figures, given the discretized shocks. To simplify the analysis, we focus here on simulation dates over which the idiosyncratic shock remains constant.

23. Note that, as this is an equilibrium model, the accidental stock accumulation is not imposed, and hence is not entirely an accident. Rather, it results from the valuations that intermediate goods producers place on these stocks toward reducing production in future dates relative to an unexpectedly low current sale price.

24. Under the baseline parameterization, such changes left consumption essentially unaltered (excepting a small decline at the date of the idiosyncratic shock's switch), and, although its composition shifted, aggregate inventory investment was similarly unaffected.

REFERENCES

- Bils, Mark and James A. Kahn (2000) What inventory behavior tells us about the business cycle. *American Economic Review* 90, 458–481.
- Blinder, Alan S. (1981) Inventories and the structure of macro models. *American Economic Review* 71(2), 11–16.
- Blinder, Alan S. and Louis J. Maccini (1991) Taking stock: a critical assessment of recent research on inventories. *Journal of Economic Perspectives* 5, 73–96.
- Caplin, Andrew S. (1985) The variability of aggregate demand with (S,s) inventory policies. *Econometrica* 53, 1395–1410.
- Christiano, Lawrence J. (1988) Why does inventory investment fluctuate so much? *Journal of Monetary Economics* 21, 247–280.
- Coen-Pirani, Daniele (2004) Markups, aggregation, and inventory adjustment. *American Economic Review* 94(5), 1328–1353.
- Eichenbaum, Martin S. (1984) Rational expectations and the smoothing properties of inventories of finished goods. *Journal of Monetary Economics* 14(1), 71–96.
- Fisher, Jonas D.M. and Andreas Hornstein (2000) (S,s) inventory policies in general equilibrium. *Review of Economic Studies* 67, 117–145.
- Hansen, Gary D. (1985) Indivisible labor and the business cycle. *Journal of Monetary Economics* 16, 309–327.
- Jorgenson, Dale W., F.M. Gollop and B.M. Fraumeni (1999) *Productivity and U.S. Economic Growth*. New York: toExcel.
- Kahn, James A. (1987) Inventories and the volatility of production. *American Economic Review* 77, 667–679.
- Kahn, James A., McConnell, Margaret and G. Perez-Quiros (2002) On the causes of the increased stability of the U.S. economy. *FRBNY Economic Policy Review* 8, 183–202.
- Khan, Aubhik and Julia K. Thomas (2005) Inventories and the Business Cycle: An Equilibrium Analysis of (S,s) Policies. Working paper, University of Minnesota.
- King, Robert G. and Sergio T. Rebelo (1999) Resuscitating real business cycles. In M. Woodford and J. Taylor (eds.), *Handbook of Macroeconomics* IB, pp. 927–1007. Amsterdam: Elsevier Science.

- Kydland, Finn E. and Edward C. Prescott (1982) Time to build and aggregate fluctuations. *Econometrica* 50, 1345–1370.
- Mehra, Rajnish and Edward C. Prescott (1985) The equity premium: a puzzle. *Journal of Monetary Economics* 15, 145–161.
- Prescott, Edward C. (1986) Theory ahead of business cycle measurement. *Quarterly Review, Federal Reserve Bank of Minneapolis* 10, 9–22.
- Ramey, Valerie A. (1991) Nonconvex costs and the behavior of inventories. *Journal of Political Economy* 99, 306–334.
- Ramey, Valerie A. and Daniel J. Vine (2004a) Tracking the Source of the Decline in GDP Volatility: An Analysis of the Automobile Industry. Working paper, University of California, San Diego.
- Ramey, Valerie A. and Daniel J. Vine (2004b) Why Do Real and Nominal Inventory–Sales Ratios Have Different Trends? Working paper, University of California, San Diego.
- Ramey, Valerie A. and Kenneth D. West (1999) Inventories. In M. Woodford and J. Taylor (eds.). *Handbook of Macroeconomics* IB, pp. 863–922. Amsterdam: Elsevier Science.
- Rogerson, Richard (1988) Indivisible labor, lotteries and equilibrium. *Journal of Monetary Economics* 21, 3–16.
- Scarf, Herbert E. (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*, pp. 196–202. Stanford, CA: Stanford University Press.
- Schuh, Scott (1996) *Evidence on the Link Between Firm-Level and Aggregate Inventory Behavior*. Finance and Economics Discussion Series 96-46, Board of Governors of the Federal Reserve System.
- Stock, James R. and Douglas M. Lambert (1987) *Strategic Logistics Management*. Homewood, IL: IRWIN.
- Stock, James H. and Mark W. Watson (2003) Has the Business Cycle Changed? Evidence and Explanations. Working paper, Harvard University.