Discussion Paper 10

Institute for Empirical Macroeconomics
Federal Reserve Bank of Minneapolis
Research Department
250 Marquette Avenue
Minneapolis, Minnesota 55480

January 1989

THE MAGNITUDE OF THE SPECULATIVE MOTIVE FOR HOLDING INVENTORIES IN A REAL BUSINESS CYCLE MODEL

Lawrence J. Christiano

Federal Reserve Bank of Minneapolis
and National Bureau of Economic Research

Terry J. Fitzgerald

Federal Reserve Bank of Minneapolis
and University of Minnesota

ABSTRACT

The motive to hold inventories purely in the hope of profiting from a price increase is called the speculative motive. This motive has received considerable attention in the literature. However, existing studies do not have a clear implication for how large it is quantitatively. This paper incorporates the speculative motive for holding inventories into an otherwise standard real business cycle model and finds that empirically plausible parameterizations of the model result in an average inventory stock to output ratio that is virtually zero. For this reason we conclude that the quantitative magnitude of the speculative role for holding inventories in this model is quite small. This suggests the possibility that the study of aggregate economic phenomena can safely abstract from inventory speculation.

Prepared for presentation to the NBER Nonlinear Rational Expectations Modeling Study Group, Minneapolis, September 1988. We gratefully acknowledge the computing support of the Institute for Empirical Macroeconomics.

1. Introduction.

In the post war U.S., the stock of total inventories has on average equaled roughly one quarter's GNP. The inventory literature presents several motivations which could in principle account for these holdings of inventories. One of these is the so-called speculative motive. Inventories are said to be acquired for speculative reasons if the decision to do so is based exclusively on the prospects for profiting from a price increase. This might occur, for example, if a large positive output disturbance drives down the price of current output and creates an expectation of a capital gain from holding inventories. Another example of the speculative motive is the accumulation of inventories in order to avoid a stockout in the event that a high price state is realized in the future.

As the stockout-avoidance example suggests, the non-negativity constraint on inventory stocks plays a central role in the analysis of the speculative motive for holding inventories. Because of this, and out of a desire to maintain analytic tractability, existing studies of this motive have done so in a partial equilibrium context.[1] As a result, they do not have clear implications for the magnitude of the speculative motive in the aggregate economy. One class of models that has been useful for investigating quantitative issues in macroeconomics is that of the real business cycle model introduced by Kydland and Prescott (1982) and Long and Plosser (1982). We compute the magnitude of the speculative motive

[1]The literature on inventory speculation which recognizes explicitly the non-negativity constraint on inventories includes Abel (1985), Aiyagari, Eckstein and Eichenbaum (1988), Kahn (1987) and Reagan (1982). Several papers, including those of Blanchard (1983), Eichenbaum (1984) and Christiano and Eichenbaum (1987), recognize the speculative motive for holding inventories and attempt to mimick the effects of the non-negativity constraint by introducing a term in firms' objective function which penalizes square deviations of the stock of inventories from some function of sales. The high penalty associated with very low inventory stocks is intended to capture stockout costs. The high penalty resulting from high inventory stocks is designed to capture high inventory carrying costs.

in a particular real business cycle model. Our model has the property that there is no motive, apart from the speculative one, for holding inventories. The magnitude of the speculative motive in the model is measured by the average inventory stock to quarterly output ratio that occurs in the model's equilibrium. Since the solution to the model is analytically intractable, we find it using iterative numerical methods.

Our model is a version of the real business cycle model studied by Gary Hansen (1985), modified to give agents the option of accumulating inventories. In one competitive equilibrium market environment which decentralizes the Pareto optimum studied in the paper, there are three agents: households, firms and speculators. Households sell labor time and purchase output in an effort to maximize infinite horizon discounted utility. Firms operate the economy's production technology for converting physical capital and labor into the single output good. Speculators have access to a technology which they can use to store (i.e., hold in inventory) the output good. The only source of uncertainty in the model is a shock to firms' production technology. Labor markets meet and clear prior to the current date's realization of the shock, and commodity markets clear after the shock is observed. Supply in the commodity market comes from firms and also speculators if they believe the output price is temporarily high *and* they have positive inventory stocks on hand. Demand in the commodity market originates with speculators—if they think current price is temporarily low—and from households. Firms' gross capital investment is also a source of demand in the commodity market. However, in contrast with households and speculators, firms are required to make their investment decisions prior to observing the price that clears the current period's commodity market. The timing assumptions in the model are intended to capture the idea that there is momentum in employment and physical capital decisions which prevents them from reacting instantaneously to current information. By contrast, inventory investment and household consumption decisions do respond instantly to current period shocks in our model. Our timing assumptions, not a part of Gary Hansen's (1985) original model, have the effect—in the absence of speculators—of increasing the sensitivity of the current price to technology shocks. They therefore increase the likelihood that speculators will choose to hold inventories.

We find that when the technology shocks are of empirically reasonable magnitude, then inventories are never held. When we allow a technology shock equal to the biggest one observed in the past 30 years to hit the economy 11 percent of the time, the average inventory to quarterly gross output ratio is positive, but only .00064. This is 1304 times smaller than the corresponding empirical average value, .9. We conclude that the magnitude of the speculative motive for holding inventories in this model is quantitatively negligible.

We hope that the paper makes an independent contribution as a case study in the application of numerical solution methods. We have included a fairly careful description of our computational strategy, and present examples of some "tricks" which—at least in our example—are extremely effective in accelerating the calculations.

Because of the absence of externalities in our model, the competitive equilibrium allocations are the solution to a central planning problem. The next two sections describe this problem, while the competitive market economy is discussed in detail in Appendix A. Section 4 describes alternative strategies for solving the planning problem, and compares the amount of computer time used by each. In section 5 we argue that plausibly parameterized versions of our model produce no more than a negligible role for the speculative inventory holding motive. Section 6 contains concluding remarks.

2. The Model.

The economy produces the quantity Y_t of a single good in period t. This good is allocated among economy–wide consumption, C_t, gross investment, X_t, and inventory accumulation, ΔI_t as follows:

$$(1) \qquad C_t + X_t + \Delta I_t = Y_t,$$

Gross investment and inventory accumulation are related to capital and inventory stocks as follows:

$$X_t \equiv K_t - (1-\delta)K_{t-1},$$
$$\Delta I_t \equiv I_t - I_{t-1},$$

where K_t and I_t are the economy–wide stock of capital and inventories, respectively, at the end of period t. Also, δ is the rate of depreciation on a unit of capital, and inventories are assumed not to depreciate.

Gross output is related to total hours worked and capital by the following aggregate production function:

$$(2) \qquad Y_t = F(z_t,H_t,K_{t-1}) \equiv (z_tH_t)^{(1-\theta)}K_{t-1}^{\theta}.$$

Here θ is capital's share in output, H_t is economy–wide hours worked, and z_t is a technology shock which grows at the stochastic rate s_t, i.e.:

$$(3) \qquad z_t = z_{t-1}\exp(s_t).$$

The random shock s_t is distributed independently over time with mean μ and standard deviation σ_s. We assume s_t is a realization from the following discrete distribution over the M states s(1), s(2), ..., s(M):

$$\text{Prob}[s_t = s(\ell)] = p_\ell,\ \ell = 1,\ldots,M,$$

(4) $$\Sigma_{\ell=1}^{M} p_\ell = 1;\ p_\ell \geq 0,\ \ell = 1,\ldots,M,$$

$$\Sigma_{\ell=1}^{M} p_\ell s(\ell) = \mu;\ \ \Sigma_{\ell=1}^{M} p_\ell [s(\ell) - \mu]^2 = \sigma_s^2.$$

Let N_t denote the population at date t. We assume that this grows at the constant rate, n:

(5) $$N_t = nN_{t-1},\ \text{for all t.}$$

Let lower case letters denote per–capita quantities, i.e., $h_t \equiv H_t/N_t$, $k_t \equiv K_t/N_t$, $y_t \equiv Y_t/N_t$, $i_t \equiv I_t/N_t$. Then, the resource constraint and production technology can be represented in per capita terms as follows:

(6) $$c_t + k_t - \frac{1-\delta}{n} k_{t-1} + i_t - \frac{1}{n} i_{t-1} = n^{-\theta}(z_t h_t)^{(1-\theta)} k_{t-1}^{\theta}.$$

We assume that the structure of the economy is such that in equilibrium, per capita consumption, hours worked, capital investment and inventory investment solve the following planning problem. Maximize

$$(7) \qquad E\sum_{t=0}^{\infty} \beta^t u(c_t,h_t),\ u(c_t,h_t) = \log(c_t) + \gamma(T-h_t),$$
$$\gamma > 0, 0 < \beta < 1, T > 0,$$

over contingency plans for c_t, k_t and i_t, subject to (3)–(4), (6); to c_t, h_t, $T - h_t$, k_t, i_t being non–negative (T is the total endowment of time); and to the following information constraint. The plans for i_t and c_t can be contingent on s_t and the model's variables dated t–1 and earlier. However, we assume that the plans for h_t and k_t *cannot* be contingent on s_t, and must only feed back on variables date t–1 and earlier. Thus, the planner's date t decisions are made in two stages. In the first stage—which occurs prior to the realization of s_t—the date t hours and investment decisions are made. The consumption and inventory investment decisions are made in the second stage, after observing s_t. This two stage decision structure captures the idea that there is an element of precommitment in employment and capital investment decisions, but not in inventory investment and consumption decisions.

The Appendix provides a rigorous exposition of the role of speculation in a competitive market economy whose equilibrium coincides with the solution to the planning problem that was just described. A sketch of that economy was provided in the introduction to the paper. In the market economy of the appendix, gross population growth, n, is set to unity. Aiyagari (1986) discusses one way to decentralize a planning problem like the one we have just described when $n > 1$ using the constructive immortality idea in Barro (1974).

The Appendix shows that speculators follow a simple reservation price strategy. If the price of a sure claim on date t+1 consumption relative to date t consumption is less than the reservation ratio of unity, then $I_t = 0$. Inventories will only be held when this price ratio is equal to, or greater than, unity. The reciprocal of this ratio is the risk free rate of interest, R_t, which coincides with the planner's intertemporal marginal rate of substitution in a unit of consumption, i.e.,

(8) $$R_t = \frac{\partial u(c_t, h_t)}{\partial C_t} / [\beta E_t \frac{\partial u(c_{t+1}, h_{t+1})}{\partial C_{t+1}}]$$

$$= \frac{n}{\beta} \frac{c_t^{-1}}{E_t c_{t+1}^{-1}}.$$

Here, the expectation is conditioned on variables dated t and earlier, including s_t. Thus, we have

(9) $$R_t > 1 \text{ implies } I_t = 0$$
$$R_t \leq 1 \text{ implies } I_t \geq 0.$$

3. Stationary Representation of the Model.

A solution to the model is the set of contingency plans which solve the planning problem described in section 2. Because we assume $\mu > 0$ in (4), k_t, y_t, c_t grow at the rate μ on average and, in particular, do not remain within a bounded set. In order to apply solution procedures which require that the control variables belong to a bounded set, we must first transform the model into an alternative, equivalent, form which possesses the required boundedness property. In particular, let

$$(10) \qquad \bar{k}_t = k_t/z_{t-1},\ \bar{c}_t = c_t/z_t,\ \bar{y}_t = y_t/z_t.$$

The resource constraint, (6), can be expressed in terms of these variables simply by dividing (6) by z_t. This operation is guaranteed to be well defined since $z_t > 0$ according to (3). Thus,

$$(11) \qquad \bar{c}_t + \frac{1}{\exp(s_t)}\bar{k}_t - \frac{1-\delta}{n\exp(s_t+s_{t-1})}\bar{k}_{t-1} + \bar{i}_t - \frac{1}{n\exp(s_t)}\bar{i}_{t-1}$$

$$= [\exp(s_t+s_{t-1})n]^{-\theta}h_t^{(1-\theta)}\bar{k}_{t-1}^{\theta}.$$

Note also that (7) can be rewritten in terms of $\bar{c}_t$ as follows:

$$(12) \qquad \kappa + E\sum_{t=0}^{\infty}\beta^t u(\bar{c}_t,h_t).$$

Here, $\kappa = E\sum_{t=0}^{\infty}\beta^t\log(z_t)$ and can be dropped from the analysis since it is beyond the control of the planner.

The planning problem posed above can now be restated in terms of the variables defined in (10). In particular, the objective is to find contingency plans for $\bar{c}_t$, $\bar{k}_t$, h_t and $\bar{i}_t$ to maximize (12) subject to (4), (11), the non-negativity constraints and the information constraints. Because of the recursive structure of the problem, it can be formulated as a dynamic programming problem. This will permit us to appeal to existing numerical methods for solving such problems. Before doing this, it is convenient to eliminate a control variable, $\bar{c}_t$, and the constraint, (11), by using (11) to substitute out for $\bar{c}_t$ in $u(\bar{c}_t,h_t)$. This yields the following instantaneous return function:

$$(12) \qquad r(\bar{k}, \bar{k}', \bar{i}, \bar{i}', h, s, s'),$$

where $\bar{k}$, $\bar{i}$, s are the variables that are known at the beginning of the period to the planner, $\bar{k}'$, $\bar{i}'$, h are to be chosen, and s' is the technology shock that is realized during the period. That is, in terms of the subscript t notation, at date t, $\bar{k}'$, $\bar{i}'$, h, s' correspond to $\bar{k}_t$, $\bar{i}_t$, h_t, and s_t, respectively. Also, $\bar{k}$, $\bar{i}$, s correspond to $\bar{k}_{t-1}$, $\bar{i}_{t-1}$, s_{t-1}. The assumed information structure results in a two stage decision. First, the variables $\bar{k}'$ and h are chosen after observing only $\bar{k}$, $\bar{i}$ and s. Then, $\bar{i}'$ is chosen based on also observing s'. Bellman's equation corresponding to this problem is:

$$(13) \qquad v(\bar{k},\bar{i},s) = \max_{(\bar{k}',h)\epsilon\Lambda(\bar{k},\bar{i},s)} \sum_{\ell=1}^{M} p_\ell \max_{\bar{i}'\epsilon B[\bar{k},\bar{i},s,s(\ell)]} \{r[\bar{k}, \bar{k}', \bar{i}, \bar{i}', h, s, s(\ell)] + \beta v[\bar{k}',\bar{i}',s(\ell)]\}.$$

Here, A and B are bounded feasibility sets constraining the relevant choice variables. The boundaries of these sets are determined by the non–negativity requirements on the variables of the problem.

The problem is solved once a function v is found that satisfies (13). When it is found, then the decision rules are simply the values of $\bar{k}'$, $\bar{i}'$ and h that solve the maximization problem in (13). We write these decision rules as follows:

$$\begin{aligned} \bar{k}' &= k(\bar{k}, \bar{i}, s) \\ h &= h(\bar{k}, \bar{i}, s) \qquad (14) \\ \bar{i} &= i(\bar{k}, \bar{i}, s, s'). \end{aligned}$$

It is straightforward to derive decision rules for the original problem of interest simply by expressing $\bar{k}_t$ and $\bar{i}_t$ in terms of k_t and i_t and z_t in (14).

4. Solving the Model

The first subsection describes how we concentrated h out of the maximization problem in (13) by exploiting the fact that h does not directly enter the value function. We then describe three value function iteration methods for solving the resulting concentrated dynamic programming problem. Our experience with these methods is described in the final subsection.

4.a Concentrating h Out of Bellman's Equation.

To save computation time, we decided to limit the value of M to 2. Also, let

$$(15) \qquad s(1) = \bar{s},\ s(2) = \underline{s},\ \bar{s} > \underline{s}, \text{ and } p_1 = \bar{p},\ p_2 = 1 - \bar{p}.$$

In addition, let $\bar{i}_{\bar{s}}$ denote the variable $\bar{i}'$ when the high shock ($\bar{s}$) occurs and let $\bar{i}_{\underline{s}}$ denote $\bar{i}'$ in the event of a low shock state. In this notation, (13) can be rewritten as follows:

$$(16) \qquad v(\bar{k},\bar{i},s) = \max_{(\bar{k}',\, h,\, \bar{i}_{\bar{s}}, \bar{i}_{\underline{s}}) \epsilon F(\bar{k},\bar{i},s)} \{ q(\bar{k}, \bar{k}', \bar{i}, \bar{i}_{\bar{s}}, \bar{i}_{\underline{s}}, h, s) + \beta[\bar{p} v(\bar{k}', \bar{i}_{\underline{s}}, \underline{s}) + (1-\bar{p}) v(\bar{k}', \bar{i}_{\bar{s}}, \bar{s})]\}$$

where

$$q(\bar{k}, \bar{k}', \bar{i}, \bar{i}_{\bar{s}}, \bar{i}_{\underline{s}}, h, s)$$

$$= \bar{p}\, r(\bar{k}, \bar{k}', \bar{i}, \bar{i}_{\bar{s}}, h, s, \bar{s}) + (1 - \bar{p})\, r(\bar{k}, \bar{k}', \bar{i}, \bar{i}_{\underline{s}}, h, s, \underline{s}).$$

In (16), F is defined by the requirement that $\bar{k}'$, $\bar{i}_{\bar{s}}$, $\bar{i}_{\underline{s}}$, and h be non–negative and also that consumption be non–negative in both the $\bar{s}$ and $\underline{s}$ states.

For a given set of values for all the other arguments of q, the optimizing choice of h,

$$\text{(17)} \qquad h = \psi(\bar{k}, \bar{k}', \bar{i}, \bar{i}_{\bar{s}}, \bar{i}_{\underline{s}}, s),$$

is implicitly defined by the following condition:

$$\text{(18)} \qquad q_6(\bar{k}, \bar{k}', \bar{i}, \bar{i}_{\bar{s}}, \bar{i}_{\underline{s}}, h, s) \geq 0.$$

Here, the strict inequality holds if, and only if $h = T$. In (18), q_6 denotes the partial derivative of q with respect to its 6th argument.

Using (17), we can eliminate h from (16) as follows:

$$\text{(19)} \qquad v(\bar{k}, \bar{i}, s) =$$

$$\max_{(\bar{k}', \bar{i}_{\bar{s}}, \bar{i}_{\underline{s}}) \epsilon G(\bar{k}, \bar{i}, s)} \{q(\bar{k}, \bar{k}', \bar{i}, \bar{i}_{\bar{s}}, \bar{i}_{\underline{s}}, s) + \beta[\bar{p} v(\bar{k}', \bar{i}_{\underline{s}}, \underline{s}) + (1-\bar{p}) v(\bar{k}', \bar{i}_{\bar{s}}, \bar{s})]\},$$

where

$$q(\bar{k}, \bar{k}', \bar{i}, \bar{i}_{\bar{s}}, \bar{i}_{\underline{s}}, s) = q[\bar{k}, \bar{k}', \bar{i}, \bar{i}_{\bar{s}}, \bar{i}_{\underline{s}}, \psi(\bar{k},\bar{k}',\bar{i},\bar{i}_{\bar{s}},\bar{i}_{\underline{s}},s), s].$$

In (19), G is defined by the condition that $\bar{k}'$, $\bar{i}_{\bar{s}}$, and $\bar{i}_{\underline{s}}$ are non–negative, and that consumption in each of the two states is non–negative. Equation (19) is (13) with h concentrated out.

4. b <u>Three Value Function Iteration Methods Described.</u>

The right side of (19) defines a functional, T, mapping from the space of value functions into itself. Thus, (19) can be written as follows:

(20) $$v = T(v).$$

The heart of the problem of finding a solution to our planning problem lies in finding the fixed point, v, in (20). We applied several value function iteration methods to do this. To apply these methods we had to discretize the set of possible values that $\bar{k}$, $\bar{k}'$, $\bar{i}$, $\bar{i}_{\bar{s}}$, $\bar{i}_{\underline{s}}$ can assume. Once this is done, we are free to think of v as a vector in $\mathbb{R}^m$, where m is the number of possible values that $(\bar{k}, \bar{i}, s)$ can take. In addition, we can think of T as a a vector valued function, mapping from $\mathbb{R}^m$ into $\mathbb{R}^m$. The objective, then, is to find the fixed point of this function.

Standard Value Function Iterations

All value function iteration methods share the characteristic that each is based on computing a sequence, $v_0, v_1,..$, with $\lim_{j\to\infty} v_j = v$. We call the simplest such algorithm *standard value function iteration.* It generates a sequence of value functions by iterating on T: $v_j = T(v_{j-1})$, $j = 1,2,\dots$, with $v_0 \equiv 0$. This method is also sometimes called the method of *successive approximation* (Bertsekas [1976,p.237].)

Newton Value Function Iterations

One alternative to iterating on T by standard value function iterations uses Newton's method. Given v_j, this method equates v_{j+1} with the fixed point of the linear Taylor series expansion of T about $v = v_j$. Specifically, the linearized T function is $T(v_j) + T'(v_j)(x - v_j)$, where $T'(v_j)$ is the derivative of $T(x)$ with respect to x, evaluated at $x = v_j$.[2] Then, v_{j+1} is by assumption the fixed point of this function, so that, $v_{j+1} = v_j + [I - T'(v_j)]^{-1}[T(v_j) - v_j] \equiv T_\infty(v_j)$.[3] The reason for using T_∞ to signify this operator is made clear below. A difficulty of the method is that it requires inverting the m×m matrix $[I - T'(v_j)]$. Unless the structure of the problem is such that some recursive algorithm for doing this rapidly

[2]The matrix T' is proportional to the state transition probability matrix implied by π_j, the decision rule implied by the maximization implicit in the definition of $T(v_j)$ (see the right side of equation [19].) The decision rule, π_j, associates a value for $\bar{k}'$, $\bar{i}_{\bar{s}}$, $\bar{i}_{\underline{s}}$ with each point, $(\bar{k},\bar{i},s)$, in the state space. To establish that T' is a state transition probability matrix, it is useful to express $T(v_j)$ in matrix form. First, associate each of the m possible points, $(\bar{k},\bar{i},s)$, in the state space with one of the integers $\ell = 1,2,\dots,m$, where $m = M \times n_k \times n_i$. Here, n_k and n_i are the number of points in the capital and inventory grid, respectively. Let the m×1 vector q_j be the vector of values of the function q defined after (19) associated with π_j. For example, the first element of q_j is the value of the function q evaluated at the first point in the state space, and at the value of $(\bar{k}',\bar{i}_{\bar{s}},\bar{i}_{\underline{s}})$ associated by π_j with this point. Also, let G_j be the m×m state transition matrix implied by π_j. In particular, the ℓ,n element of G_j is the probability of passing to the n^{th} point in the state space starting from the ℓ^{th} point. Each row of G_j has no more than M non–zero entries and all entries in each row sum to unity. In this notation, it is easy to confirm that $T(v_j) = q_j + \beta G_j v_j$. Since a small perturbation in the v_j's induces no change in π_j—and hence, q_j and G_j—it follows that $T'(v_j) = \beta G_j$.

[3]This method is the one used in Rust (1987,1988a,1988b). The global convergence of this method is discussed by Bertsekas (1976,pp.245–47), who calls it the *policy iteration algorithm.* See footnote 4 for further details.

is available, then direct application of the procedure is computationally prohibitive. We therefore seek computationally efficient ways to approximate this inverse. The following value function iteration method accomplishes this.

Hybrid Value Function Iterations

Let π_j denote the policy function relating $\bar{k}'$, $\bar{i}_{\bar{s}}$, $\bar{i}_{\underline{s}}$ to the state variables, $\bar{k}$, $\bar{i}$, s computed during the standard value function iteration step, $T(v_j)$. We denote the value of using policy π_j for one period given that next period's state variables are valued according to the value function x by $T_{\pi_j}(x)$. By construction, $T_{\pi_j}(v_j) \equiv T(v_j)$. Let $T^2_{\pi_j}(v_j) \equiv T_{\pi_j}[T_{\pi_j}(v_j)]$ be the value of following policy π_j for two periods given that the subsequent period's state is valued according to v_j, and define $T^p_{\pi_j}(v_j)$ similarly, for p = 2,3,... . Our third value function iteration method computes v_{j+1} from v_j as follows:

(21) $$v_{j+1} = T_p(v_j) \equiv T^p_{\pi_j}(v_j).$$

We call this method *hybrid value function iteration.* When p = 1, then this method reduces to the standard value function iteration method, i.e., $T_1 \equiv T$. In addition, it is easy to show that

(22) $$T_\infty = \lim_{p \to \infty} T_p.$$ [4]

---

[4] To see this, apply the framework in footnote 2. Thus, $T_{\pi_j}(v_j) \equiv T(v_j) = q_j + \beta G_j v_j$ and $T^2_{\pi_j}(v_j) \equiv T_{\pi_j}[q_j + \beta G_j v_j] = q_j + \beta G_j q_j + \beta^2 G_j^2 v_j$. Proceeding recursively, $T_p(v_j) \equiv T^p_{\pi_j}(v_j) = [I - \beta G_j]^{-1}[I - (\beta G_j)^p]q_j + (\beta G_j)^p v_j$. This converges to the limit $[I - \beta G_j]^{-1} q_j = z_j$, say. The existence of the indicated inverse is assured by the fact that z_j is the fixed point of T_{π_j},

That is, as the number of hybrid value function iteration steps increases, the method converges to the Newton method described above. Roughly, T_p approximates the matrix $[I - T'(v_j)]^{-1}$ by its p^{th} order series expansion: $I + T'(v_j) + T'(v_j)^2 + ... + T'(v_j)^p$.

In applying the hybrid value function iteration method, we have found it useful to vary the value of p, with p starting out small ($\cong$ 30) during the initial iterations and then jumping to a very large number (i.e., 10,000)—at which point the method is basically the Newton method–upon reaching a trigger point.[5] The jump in p is triggered when the number of states in which the current iteration decision coincides with the previous iteration's optimal decision falls below some value. We document these findings below.

4. c Computational Results.

In order to solve the model, we require values for the parameters, and grids for (detrended) capital and inventories, $\bar{k}_t$ and $\bar{i}_t$. We think of the time

---

which is known to exist and be unique as a consequence of the fact that T_{π_j} is a contraction mapping. To show that $T_\infty(v_j) = z_j$, some simple substitutions are required. Recall from footnote 2 that $T(v_j) = q_j + \beta G_j v_j$, so that $q_j = T(v_j) + [I - \beta G_j]v_j - v_j$. Upon substituting, we find $z_j = v_j + [I - \beta G_j]^{-1}[T(v_j) - v_j]$. But, according to footnote 2, βG_j is $T'(v_j)$, so that $z_j = T_\infty(v_j)$, as claimed. The value function iteration algorithm, defined by equating v_{j+1} to z_j, is called the *policy iteration algorithm* by Bertsekas (1976,p.246). That z_j and $T_\infty(v_j)$ are mathematically identical therefore establishes the equivalence of policy iterations and newton value function iterations.

[5]In practice, T_p may involve fewer than p applications of the T_{π_j} operator, as in (21). In particular, let $w_0 = v_j$ and define $w_i = T_{\pi_j}(w_{i-1})$, $i = 1,2,...,p$. Then, according to (21), $T_p(v_j) \equiv w_p$. When, however, we found w_r to be close to w_{r-1} for some $r < p$, we set $T_p(v_j) = w_r$. Thus, for example, when p = 10,000, we rarely found it necessary to literally iterate on T_π 10,000 times.

interval in our model as quarterly, and set the discount rate at a 3 percent annual rate. That is $\beta = 1.03^{-.25}$. The parameters n, δ, θ, γ, μ where set at 1.00325, .0183, .3606, .00275, and .004, respectively. These values are taken from Christiano (1988,Table 2, Model 1). These parameter values cause selected steady state properties of the nonstochastic version of our model to roughly match the corresponding sample averages from post–war U.S. data. (See Christiano [1988] for a discussion of the data.) The sample averages used include those of c_t/y_t, x_t/y_t, k_t/y_t, h_t and $\log(y_t)-\log(y_{t-1})$, and are reported in column 1 of Table 3. Here, x_t denotes $k_t - [(1-\delta)/n]k_{t-1}$ in the model and per capita gross fixed capital formation in the data. Column 2 of Table 3 reports the steady state properties of the model. The rationale for this method of selecting parameters lies in the assumption that the steady state properties of the nonstochastic version model correspond closely to the corresponding first moment properties of the stochastic version of the model. Under this assumption, the method of assigning parameter values which equates steady state properties with sample averages is an exactly identified, first moment estimator.

We set $\sigma_s = .019$, which is the standard deviation of $\log(z_t)-\log(z_{t-1})$ when z_t is computed using (2), the relevant post–war U.S. data, and the value of θ used in the model. (See Christiano [1988,Table 3].) Given the mean and variance restriction on s_t, we still have one degree of freedom in parameterizing its distribution given that M = 2. In doing so we where guided by the histogram of 112 empirical values of s_t covering the period 1956.2 – 1984.1 which are graphed in Figure 1. Based on that graph, we chose three models of s_t. In the first, s_t is skewed to the right, with $\bar{s} = .057$, the largest empirical value of s_t. In the second model, s_t is symmetrically distributed about μ, with $\underline{s} = -.015$ (15%) and $\bar{s} = .023$ (14%), where the numbers in parentheses are the percent of empirical values of s_t lying in the tail. These values of s_t are indicated in Figure 1. The third model specifies the s_t to be skewed to the left, with $\underline{s} = -.049$, the smallest empirical value in the sample.

We used these three distributions to define four versions of our model. Model 1 incorporates the right skewed distribution of s_t's, Model 2 the symmetric distribution, and Model 3 the left skewed distribution. For reasons to be made clear in the next section, each of these models was solved subject to the restriction that inventory holdings are always identically zero. We do, however, impose the two stage decision structure on those models. Model 4 incorporates the right skewed distribution of s_t's and permits inventories to be held.

The only thing left to specify are the grids for capital and inventories. Obviously, no grid for the latter is required in the case of Models 1 – 3. In Models 1 – 3, the capital grid contains 10,000 grid points, and covers the range indicated in Table 1. Table 1 also reports the distribution of s_t associated with each model, as just discussed. In the case of Model 4, the number of grid points on capital was reduced to 800, and the number on inventories was set at 11. Computing times for each model are reported at the bottom of the Table. The decision rules for each of the four models were stored for use in the simulations discussed in the next section.

Results of experiments with versions of the solution procedures discussed in the previous subsection are reported in Table 2. In each case, our convergence criterion was that the maximum percent difference between elements in v_j and v_{j-1} be less than $.1\times10^{-9}$. In addition, we always started the iterations with $v_0 = 0$. The calculations involved solving a version of Model 4 in which the capital and inventory grid contain 100 and 11 points, respectively. The first row shows what happened when we applied standard value function iterations. We canceled the job after the program had used up 44 CPU minutes of Cray–2 supercomputer time. Based on our examination of the output from that run, it seemed that the program still needed at least another 33 minutes to converge. The output revealed that the decision rule had stopped changing on each iteration. Consequently, although at each step the relatively expensive maximization problem implicit in the T_1 operator was being carried out, in practice the program was simply applying the T_{π_j} operator repeatedly. This suggested that convergence could be speeded up by starting with hybrid value function iterations with p = 1 (i.e., standard value iterations) and then raising p to 10,000 when the decision rule stopped changing very much. The results of this are reported in the second row of Table 2, which shows that the routine converged reasonably quickly, in about 4.6 minutes. Thus, this one change improved convergence speed by a factor exceeding 10. Row three of Table 2 reports the effects of starting the iterations with p = 30 rather than 1. That resulted in a further decline in computing time by a factor of 3.5. We then considered the effect of triggering Newton value iterations (i.e., p = 10,000) sooner, when the decision rule stopped changing at 75 percent of the points in the state space. That resulted in very little change in computer time (see row 4.) Row 5 reports the results of starting with Newton iterations immediately. Evidently, that dominates by a factor of over 10 the method of applying standard value iterations from start to finish (eg., row 1). However, results with setting p = 10,000 throughout are clearly inferior to setting p to a lower number at the beginning. Finally, row 6 reports the result of starting with p = 30 and not triggering Newton iterations at all. Comparing the results of this row with those in rows 3 and 4 shows that it is better, when p = 30 initially, to trigger Newton iterations at the end.

In sum, the results in Table 2 suggest that hybrid value function iteration, with $p > 1$ is far more efficient than standard value function iterations. Second, the results suggest that it is best to start out with a low value of p and then trigger into Newton value iterations before the elements of the decision rules stop changing. Clearly, we cannot expect the performance results associated with $p > 1$

reported in Table 2 to apply in all problems. We expect that setting $p > 1$ will be most advantageous when the maximization problem to the right of the equality in (19) is particularly costly in terms of computer time.

5. The Magnitude of the Speculative Motive for Holding Inventories

The results of this section are based on simulating 100 realizations of length 10,000 each, of data from each of our four models. Each simulation was carried out by drawing randomly from the relevant model's distribution for s_t. We first discuss the simulations of Model 2, the one with a symmetric distribution for the s_t's (see Table 1 for Model definitions.) Figure 2 plots the histogram of R_t generated by Model 2. The mean value of the gross risk free rate is 1.0146, or close to 6% at an annual rate. Interestingly, 1.0146 is also the risk free rate that obtains in the steady state of the nonstochastic version of the model. The distribution in Figure 2 has two humps, reflecting the two humps in the distribution of s_t. The higher risk free rates are associated with $s_t = \underline{s}$, and the lower ones with $s_t = \bar{s}$. Significantly, from our perspective, R_t is never less than 1 in Model 2. As a consequence, the fact that I_t is constrained to be zero during the simulations is never binding (recall [9].) Therefore, in Model 2 the speculative motive for holding inventories is quantitatively negligible.

Next, consider the R_t's generated by Models 1 and 3. Recall that the s_t's in the former are skewed to the right, while those in the latter are skewed to the left. The histogram of risk free rates from these models is plotted in Figure 3. In each case, the mean value of R_t is 1.0146, as in Model 2 and in the nonstochastic version of the model. Both interest rate distributions in Figure 3 are bi-modal, with one peak being considerably higher than the other. The histogram marked by the solid line corresponds to the left-skewed technology shocks (Model 3.) Its smaller peak is located in the right corner of the graph, and is associated with $s_t = \underline{s}$, which occurs with low probability (see Table 1.) Large negative shocks create a rise in the relative price of goods in the current period relative to the future, and therefore create an incentive to decumulate inventories, if indeed any are being held.

In fact, in Model 3 it is optimal to never hold inventories, because negative values of R_t-1 never occur. High values of the technology shock—which have the effect of reducing the current price of goods—are never big enough to induce speculators to hold inventories. The low values of the technology shock do not have a high enough probability to induce speculators to accumulate inventories in advance. Thus, the speculative motive for holding inventories is also quantitatively negligible in Model 3.

Next, consider Model 1, which has a right skewed technology shock distribution. In this model, R_t-1 is always driven into the negative region when (and, only when) a big positive technology shock is realized, so that the restriction $I_t \equiv 0$ is binding. The average value of R_t, conditional on $s_t = \underline{s}$, is .99278 and the lowest value of R_t ever computed is .99138. These correspond to annualized rates of interest of −2.9 and −3.4 percent, respectively. Thus, in this model there is a quantitatively identifiable speculative motive for holding inventories.

Figure 4 graphs the histogram of R_t generated by Models 1 and 4. The Model 1 histogram is reproduced from Figure 3 for convenience. Recall that Model 4 is the same as Model 1, with the exception that $I_t \equiv 0$ is not imposed. In both Models 1 and 4, the average value of R_t is 1.0146. At first glance this was surprising to us since, as expected, the lower left tail of the R_t distribution was shifted to the right in Model 4 relative to Model 1. For example, the lower bound on R_t was shifted up from −3.4 percent (annualized rate) for Model 1 to −.6 percent (annualized rate) for Model 4.[6] The reason the mean value of R_t is nevertheless identical in the two models is that the mean of R_t conditional on $s_t = \underline{s}$ is shifted

[6]Of the 1,000,000 simulated values of R_t-1 for Model 4, 5.8 percent were negative. In each of these cases, $\bar{i}_t$ was positive. Presumably, had our $\bar{i}_t$ grid been finer, $\bar{i}_t$ would have been chosen slightly larger in these cases, and a negative value of R_t-1 would never have been observed.

down in Model 4 relative to that in Model 1.[7] That, in turn, reflects that the marginal utility of consumption associated with $s_t = \bar{s}$ is higher in Model 4 than in Model 1 due to holdings of inventories in the high shock state. This marginal utility term appears in the denominator of R_t (see [8]), reducing its magnitude in Model 4 when $s_t = \underline{s}$.

Table 3 shows how much inventories are actually being held in Model 4. The first column of numbers in that table reports the empirical values of, among other variables, the ratios i_t/y_t and di_t/y_t, where $di_t \equiv i_t - \frac{1}{n}i_{t-1}$. In the data these are .9 and .0065, respectively. In the model these are .00064 (.00013) and .000014 (.000023), respectively. Numbers in parentheses are the standard deviation of the associated statistic across the 100 simulations. Thus, the speculative motive for holding inventories, as captured by Model 4, can only account for .07% (= 100×.00064/.94) of aggregate inventories. Moreover, in that model people only hold positive inventories in periods in which the high technology shock is realized. Recall that shock is the largest empirical value of s_t observed in the last 30 years, and in this model it occurs 11 percent of the time. Thus, even with very large shocks (although keeping the variance of s_t empirically reasonable), the speculative motive for holding inventories is far too small to account for an appreciable fraction of inventories held using the real business cycle model of this paper.

Table 3 also permits one to compare the mean values of c_t/y_t, x_t/y_t, k_t/y_t h_t and $\log(y_t) - \log(y_{t-1})$ implied by the models with the associated steady states (see column 1). Note in particular how close mean and steady state values are. Therefore, in the framework of this model it is valid to interpret the parameter

[7]The frequency of $s_t = \underline{s}$ and $s_t = \bar{s}$ is the same in the two models, being 11.3537 and 88.6453 percent, respectively (identical shocks were used in the simulations.) The mean value of R_t conditional on $s_t = \underline{s}$ and $s_t = \bar{s}$ is 1.01743 and .99278 for Model 1, respectively, and 1.01653 and 1.00000 for Model 4, respectively.

estimation procedure described earlier—under which one equates model steady state and sample properties—as a first moment estimator.

Another useful indicator of the role of inventory investment in this model is the role it plays in buffering consumption from technology shocks. Table 4 reports the volatility of consumption, capital investment, and inventory investment, relative to the volatility of income. These are defined as follows. The volatility of a variable, say w_t, is defined as v_w, where $v_w = (1/T)\Sigma_{t=0}^{T}|\Delta w_t|/y_t$, where Δ is the first difference operator and T is the number of observations in the sample (T = 112 in the data and 10,000 in the simulations.) Then, the volatility of w_t relative to the volatility of y_t is v_w/v_y. Note that in the case of Models 1 – 3, the relative volatility of consumption is quite high compared with what it is in the actual data. In Model 4 the relative volatility of consumption is lower than it is in Model 1, showing that the role of speculation is to stabilize consumption (and, hence, the Arrow–Debreu price of consumption), but that it does not make it as smooth as it is in the data. Put simply, the return on holding inventories is so low (on average, it is negative) that using them to smooth out consumption, though feasible, is simply not worth it. Note also that the relative volatility of capital investment is higher in Model 4 than in Model 1. We suspect this reflects that inventories accumulated in high shock states in Model 4 are invested in physical capital in the following period. Such investments can be expected to generate a high return because of the random walk nature of the shocks.

6. Conclusion.

This paper has taken a standard model used to study aggregate fluctuations, and modified it to allow speculators to hold goods in storage (inventory) in the hope of profiting from a price increase. When the model is parameterized in an empirically reasonable way, inventories of the model's one good are practically never held.

One way to explain our finding views our competitive model economy from the perspective of a central planner who, because of the assumed absence of externalities, chooses the same levels of consumption, inventory holdings, hours worked, etc., as occur in competitive equilibrium. A measure of the planner's willingness to trade goods intertemporaly is given by the risk free rate of interest: the number of consumption goods the planner is willing to give up next period in exchange for one extra good in the present. As long as the risk free rate exceeds 1, the planner prefers zero to positive stocks of inventories. This is because by reducing stocks to zero the planner can increase welfare by raising current consumption at the cost of only an equal amount of consumption in the next period. Only when the risk free rate is less than, or equal to 1 would the planner be willing to hold a positive amount of inventories. The reason inventory stocks are almost never held is that when the planner is constrained to hold zero inventories, the rate of interest almost always exceeds 1, so that the constraint practically never binds. In this constrained version of the model the net risk free rate of interest is on average quite high—around 6% at an annual rate—and has a fairly small variance. The fact that the average value of the risk free rate is quite high in a model like ours has received a great deal of attention in the literature.[8] One paper that has drawn attention to this fact particularly forcefully is Mehra and Prescott (1985).

We suspect that a heterogeneous good version of our model, in which each industry's production function has a large idiosyncratic component in its productivity shock and/or tastes shift randomly between goods, may produce larger speculative inventory stocks in equilibrium. This possibility seems interesting to us, however exploring it using the strategy of this paper involves substantial computational challenges. This is because heterogeneity increases the number of endogenous state variables, greatly increasing the number of calculations required to solve the model. In the mean time we think it fair to conclude from our results that one can safely abstract from inventory speculation in models of aggregate fluctuations.

In particular, in models specified at the level of aggregation in this paper, more plausible explanations of the observed holdings of aggregate inventories are motives that focus on the service flow yielded by holding inventories, in addition to possible price changes. One channel by which inventories produce a service flow is by facilitating increased efficiency in the allocation of labor and physical capital. For example, Blinder (1981,p.453) points out that economies of scale in transportation and fixed costs associated with placing and processing orders would be expected to result in manufacturers bunching their shipments to retailers—thereby resulting in positive inventory holdings—in order to reduce costs. Similarly, costs—in the form of machine down time and labor time—of switching

[8]When utility is time separable and separable across consumption and leisure—as we assume—then the risk free rate which obtains in non-stochastic steady state is $\beta^{-1} + \gamma\Delta\log(c) + n-1$, where β^{-1} is the gross rate of time preference, $\Delta\log(c)$ is the net rate of growth in per capita consumption, n is the gross population growth rate, and γ is the coefficient of relative risk aversion on instantaneous utility. In our model the annual rate of time preference is 3%, $\gamma = 1$, and $\Delta\log(c)$ and $n-1$ are around 1.5% at an annual rate, implying a steady state risk free rate of about 6% annually. This is also the risk free rate in the stochastic version of our economy.

production among heterogeneous goods could induce manufacturers to make production runs that exceed their planned sales for the current period (Lovell [1981,p.508].) As noted by Blinder (1981), these considerations are reminiscent of features of the simple (S,s) model of inventory accumulation associated with Arrow, Harris and Marschak (1951). Other existing models of aggregate inventories, including those of Kydland and Prescott (1982) and Christiano (1988) attempt to capture the service flow motive for holding inventories by including them as a factor of production in the production function. One interesting, as yet unanswered, question is how accurate this approximation is.

Table 1

Decision Rule Information[1]

	Model 1	Model 2	Model 3	Model 4
$\mathcal{E}_K$	{9744,15530}	{11012,17414}	{12266,19120}	{9772,15457}
K Grid	{7000,17000}	{9500,19500}	{11000,21000}	{9000,16000}
#K Grid pts	10,000	10,000	10,000	800
$\underline{s}$	–0.0028	–0.015	–0.049	–0.0028
$\bar{s}$	0.057	0.023	0.011	0.057
$p_{\underline{s}}$	0.89	0.50	0.11	0.89
$p_{\bar{s}}$	0.11	0.50	0.89	0.11
$\mathcal{E}_I$	0[2]	0	0	{0,10}
I Grid	0	0	0	{0,20}
#I Grid pts	0	0	0	11
CPU[3] Time	2472.	1542.	1770.	1191.4

[1] $\mathcal{E}_K$ and $\mathcal{E}_I$ denote the ergodic sets of $\bar{k}_t$ and $\bar{i}_t$, respectively. The points in these sets were found by simulating 1,000,000 observations on the variables of the model by drawing an equal number of s_t's from the indicated distribution. The lower point in the ergodic set of a variable is its minimum value in the 1,000,000 observations and the upper point is the maximum value. K and I grid are the boundaries of the grid on $\bar{k}_t$ and $\bar{i}_t$, respectively, used in the computations. Also reported is the number of grid points.

[2] Zeros reflect that $\bar{i}_t \equiv 0$ in the computations.

[3] CPU time refers to central processing unit time, in seconds, used in solving the indicated model. For models 1, 2, 3 the computations were carried out on the Federal Reserve Bank of Minneapolis' Amdahl mainframe computer. The computations for Model 4 were carried out on the Cray–2.

Table 2

Computational Costs of Alternative Versions of Hybrid Policy Iterations[1]

p[2]	Trigger[3]	Steps[4]	CPU[5] Time
1	no	1740	2640.0[6]
1	3	121	287.7
30	3	26	81.3
30	1,650	22	84.9
∞	no	46	227.8
30	no	112	279.8

[1]Calculations based on a model identical to Model 4 in Table 1, except that the capital grid contained only 100 points.

[2]Value of p in hybrid value function iteration for j = 1, 2, 3, ..., until the decision rule changes at less than *trigger* points in the state space, which is itself composed of 2,200 points. After this, $p = \infty$. When $p = 1$ the method corresponds to standard value function iterations. When $p = \infty$ the method represents Newton value function iterations.

[3]A no in this column means no trigger was used.

[4]Number of value function iteration steps.

[5]Central Processing Unit seconds on the Cray–2 supercomputer.

[6]This job was canceled before convergence because progress was so slow. Our guess is that convergence would have required at least another 2,000 seconds of CPU time.

Table 3

First Moment Properties[1]

	U.S. Data (1)	Steady State (2)	Model 1 (3)	Model 2 (4)	Model 3 (5)	Model 4 (6)
c_t/y_t	.725	.721	.722	.722	.722	.722
			(.010)	(.010)	(.010)	(.010)
x_t/y_t	.269	.279	.278	.278	.278	.278
			(.010)	(.010)	(.010)	(.010)
di_t/y_t	.0065	0	0	0	0	.000014
			(0)	(0)	(0)	(.000023)
k_t/y_t	10.59	10.98	11.05	11.04	11.04	11.05
			(.312)	(.314)	(.320)	(.310)
i_t/y_t	.90	0	0	0	0	.00064
			(0)	(0)	(0)	(.00013)
h_t	320.5	322.14	321.85	321.93	321.95	321.85
			(4.68)	(4.66)	(4.66)	(4.66)
$\Delta \log y_t$	.0041	.004	.0040	.0040	.0039	.0039
			(.0017)	(.0017)	(.0018)	(.0017)

[1]Numbers in column (1) taken from Table 1 in Christiano (1988). Numbers in column (2) are the values of the indicated variables in the steady state of the nonstochastic version of the model. Since steady state R_t is 1.0146, $i_t = 0$ in the steady state. Results in the columns (3) – (6) are based on simulating 100 artificial data sets, each with 10,000 observations, from the indicated model. For model definitions, see Table 1. The reported numbers all but column (2) the table are averages of the variable indicated in the first column, across the 100 data sets. Numbers in parentheses are the corresponding standard deviation.

Table 4

Relative Volatility Measures[1]

	U.S. Data[2]	Model 1	Model 2	Model 3	Model 4
Consumption (c)	0.47	0.86 (0.077)	1.00 (0.079)	0.81 (0.052)	0.75 (0.047)
Investment (x)	0.47	0.81 (0.047)	0.94 (0.045)	0.66 (0.063)	0.97 (0.062)
Inventory (di) Investment	0.51	0 (0)	0 (0)	0 (0)	0.29 (0.032)

[1]The volatility of a variable, say w_t, is defined as $v_w = (1/T)\ \Sigma_{t=1}^{T} |\Delta w_t|/y_t$, where $|\Delta w_t|$ denotes the absolute value of the first difference in w_t. The relative volatilities reported in the table are v_w/v_y for w = c, x and di, respectively.

[2]Taken from Table 4 in Christiano (1988). Other columns were computed using the artificial data underlying the results in Table 3.

References

Abel, Andrew, 1985, "Inventories, Stock–Outs and Production Smoothing," Review of Economic Studies LII, 283–293.

Aiyagari, S. Rao, 1986, "Overlapping Generations and Infinitely Lived Agents," Federal Reserve Bank of Minneapolis Working Paper 328, December.

Aiyagari, S. Rao, Zvi Eckstein and Martin Eichenbaum, 1988, "Inventories and Price Fluctuations Under Perfect Competition and Monopoly," in T. Kollintzas, ed., The Rational Expectations Equilibrium Inventory Model: Theory and Applications, Springer–Verlag, forthcoming.

Arrow, Kenneth J., Theodore Harris and Jacob Marschak, 1951, "Optimal Inventory Policy," Econometrica, vol. 19, July, pp. 250–72.

Barro, Robert J., 1974, "Are Government Bonds Net Wealth?," Journal of Political Economy, 82, pp. 1095–1117.

Bertsekas, Dimitri P., 1976, Dynamic Programming and Stochastic Control, Academic Press.

Blanchard, Olivier J., 1983, "The Production and Inventory Behavior of the American Automobile Industry," Journal of Political Economy, June, 91, 365–400.

Blinder, Alan S., 1981, "Retail Inventory Behavior and Business Fluctuations," Brookings Papers on Economic Activity, 2, pp.443–505.

Christiano, Lawrence J., 1988a, "Why Does Inventory Investment Fluctuate So Much?", Journal of Monetary Economics 21, pp. 247–280.

Christiano, Lawrence J., and Martin S. Eichenbaum, 1987, "Temporal Aggregation and Structural Inference in Macroeconomics," Carnegie–Rochester Conference Series on Public Policy, vol. 26, Spring.

Eichenbaum, Martin S., 1984, "Rational Expectations and the Smoothing Properties of Inventories of Finished Goods," Journal of Monetary Economics, November, 8, 277–304.

Kahn, James A., 1987, "Inventories and the Volatility of Production," American Economic Review, vol. 77, no.4.

Kydland, Finn, and Edward Prescott, 1982, "Time to Build and Aggregate Fluctuations," Econometrica, 50, pp. 1345–1371.

Long, John, and Charles Plosser, 1982, "Real Business Cycles," Journal of Political Economy.

Lovell, Michael C., 1981, "Comment on Blinder," Brookings Papers on Economic Activity, 2, pp.506–513.

Lucas, Robert E., Jr., 1984, "Money in a Theory of Finance," Carnegie-Rochester Series on Public Policy, 21, 9–46.

Mehra, Rajnish, and Edward C. Prescott, 1985, "The Equity Premium: A Puzzle," Journal of Monetary Economics 15 (2):145–162.

Reagan, Patricia B., 1982, "Inventory and Price Behaviour," Review of Economic Studies XLIX, 137–142.

Rust, John, 1987, "Optimal Replacement of GMC Bus Engines: An Empirical Model of Harold Zurcher," Econometrica.

Rust, John, 1988a, "Maximum Likelihood Estimation of Discrete Control Processes," SIAM Journal of Control and Optimization, vol. 26, no.5, November.

Rust, John, 1988b, "A Dynamic Programming Model of Retirement Behavior," in David Wise, ed., The Economics of Aging, National Bureau of Economic Research, University of Chicago Press.

Appendix A: A Decentralized Market Interpretation of Our Planning Problem

A competitive economy whose equilibrium allocations coincide with those of the planning problem in the text is presented here under the restriction that there is zero population growth (n=1). We adopt the normalization $N_t = 1$ for all t. In order to highlight the speculative motive for holding inventories, access to the economy's storage technology is restricted to a single agent called the speculator. Firms are the type of agent who have access to the economy's production technology, F. Households own firms and speculators and use the profits from this ownership (which are zero in equilibrium) together with income from labor services to purchase consumption goods. In the real world the distinction between households, speculators and firms are not as sharp as they are in our model. We maintain the distinction for expositional clarity.

In the first of the following two subsections we describe the market economy and define a competitive equilibrium. We use the time 0 Arrow–Debreu framework spelled out in Lucas (1984) for an economy very similar to ours. In the second subsection we establish that the equilibrium allocations in our market economy coincide with those of the planning problem studied in the text.

A.1 The Market Economy and Its Competitive Equilibrium.

Denote the entire history of realizations $s_0, s_1, s_2, \ldots, s_t$ by s^t. The set of all possible s^t is denoted S^t and consists of M^{t+1} elements. Finally, the joint density of s^t is $f^t(s^t)$. To simplify notation, we assume there is a single household, a single firm and a single speculator. We begin by considering each of these in turn. After that we define a competitive equilibrium.

Households

Household preferences are given by the following utility function:

$$\text{(A1)} \quad \sum_{t=0}^{\infty} \sum_{s^t \epsilon S^t} \beta^t f^t(s^t)\{u(c_t(s^t)) + \ell(h_t(s^{t-1}))\},$$

where $c_t(s^t)$ is date t consumption, contingent on history s^t, and $h_t(s^{t-1})$ is date t hours worked which is contingent on s^{t-1}. Here, $u(\cdot) \equiv \log(\cdot)$ and $\ell(h_t(s^{t-1})) \equiv -\gamma h_t(s^{t-1})$. Consumption is, and hours worked is not, a function of s_t. This reflects our assumption that date t consumption is set after, and hours worked before, s_t is realized.

Following is the household budget constraint:

$$\text{(A2)} \quad \sum_{t=0}^{\infty} \sum_{s^t \epsilon S^t} \{p_t(s^t)c_t(s^t) - w_t(s^{t-1})h_t(s^{t-1})\} = \pi_{\text{firm}} + \pi_{\text{speculator}}.$$

Here, π_{firm} and $\pi_{\text{speculator}}$ denote the profits of firms and speculators, respectively, and are defined below. In addition, $p_t(s^t)$ is the date t price of the consumption good, given history s^t and $w_t(s^{t-1})$ is the wage.

Households maximize (A1) subject to (A2) and consumption and hours worked being non-negative. Let λ denote the Lagrange multiplier on (A2) in this constrained maximization. Then, the first order condition for $h_t(s^{t-1})$ is:

$$\text{(A3)} \quad \beta^t \ell'(h_t(s^{t-1}))f^{t-1}(s^{t-1}) = -\lambda\, Mw_t(s^{t-1}).$$

The corresponding condition for $c_t(s^t)$ is:

(A4) $\beta^t f^t(s^t) u'(c_t(s^t)) = \lambda p_t(s^t)$.

Firms

Firms maximize their cash flow:

$$\text{(A5)}\quad \sum_{t=0}^{\infty} \sum_{s^t \epsilon S^t} \{p_t(s^t) y_t(s^t) - w_t(s^{t-1}) \mathbf{h}_t(s^{t-1})$$

$$- p_t(s^t)[k_t(s^{t-1}) - (1-\delta)k_{t-1}(s^{t-2})] \},$$

subject to

(A5a) $k_t(s^{t-1}) \geq (1-\delta)k_{t-1}(s^{t-2})$,

(A5b) k_{-1} given.

Here, $k_t(s^{t-1}) - (1-\delta)k_{t-1}(s^{t-2})$ is the date t, history s^{t-1} gross capital investment undertaken by the firm. Equation (A5a) is the assumption that investment is "irreversible". The fact that date t investment is not a function of s_t reflects our assumption that that decision has to be made before the realization of s_t. Also, $y_t(s^t)$ is the output produced by the firm, using production technology F:

$$\text{(A6)}\quad y_t(s^t) = F[s^t, h_t(s^{t-1}), k_{t-1}(s^{t-2})]$$

$$= [\exp(\Sigma_{j=0}^{t} s_j) z_{-1} h_t(s^{t-1})]^{(1-\theta)} [k_{t-1}(s^{t-2})]^{\theta}$$

The first order conditions for the firm's maximization problem are:

$$\text{(A7)}\quad \sum_{s_t} p_t(s^{t-1}, s_t) = \sum_{s_t, s_{t+1}} p_{t+1}(s^{t-1}, s_t, s_{t+1})[y^k_{t+1}(s^{t-1}, s_t, s_{t+1}) + 1-\delta],$$

for $k_t(s^{t-1})$ and

$$\text{(A8)}\quad Mw_t(s^{t-1}) = \sum_{s_t} p_t(s^{t-1}, s_t) y^n_t(s^{t-1}, s_t),$$

for $h_t(s^{t-1})$. In (A8), $y^h_t(s^t)$ and $y^k_t(s^t)$ are the derivative of F with respect to n and k, respectively. We denote the maximized value of (A5) by π_{firm}.

Speculators

Speculators' objective is to maximize:

$$\text{(A9)}\quad \sum_{t=0}^{\infty} \sum_{s^t \epsilon S^t} p_t[i_{t-1}(s^{t-1}) - i_t(s^t)],$$

subject to

(A9a) $i_t(s^t) \geq 0$ for all s^t, t,

(A9b) i_{-1} given.

Speculators' first order condition is:

$$\text{(A10)} \quad \sum_{s_{t+1}} p_{t+1}(s^t,s_{t+1}) - p_t(s^t) \qquad \begin{array}{l} \geq 0 \text{ then } I_t(s^t) \geq 0 \\ < 0 \text{ then } I_t(s^t) = 0. \end{array}$$

Equation (A10) formalizes the remarks about speculation made in the introduction. For example, inventories may be held if there is a state, s^*_{t+1}, in period t+1 in which the price, $p(s^t,s^*_{t+1})$, of the commodity is sufficiently high. They may also be held if $p_t(s^t)$ is sufficiently low. Of course, if the object to the right of (A10) is positive, then the demand for $i_t(s^t)$ by speculators is infinite so that in equilibrium we do not expect to see it exceed zero.

For purposes of comparison with the text, it is convenient to rewrite the object in (A10) as follows:

$$\text{(A11)} \quad R_t(s^t) = \frac{p_t(s^t)}{\sum_{s_{t+1}} p_{t+1}(s^t,s_{t+1})}.$$

Here, $R_t(s^t)$ is the rate at which a unit of $c_t(s^t)$ can be exchanged risklessly for a unit of time t+1 consumption. The formula for the risk free rate of interest given in (A11) can be seen to coincide with (8) after substituting (A4) into (A11).

Competitive Equilibrium

We are now in a position to define a competitive equilibrium:

Definition 1: A **competitive equilibrium** is a specification of $\{c_t(s^t)\}$, $\{h_t(s^{t-1})\}$, $\{\mathbf{h}_t(s^{t-1})\}$, $\{k_t(s^{t-1})\}$, $\{i_t(s^t)\}$, $\{y_t(s^t)\}$, $\{p_t(s^t)\}$, $\{w_t(s^{t-1})\}$ satisfying:

Market Clearing: for each t and s^t

$$c_t(s^t) + k_t(s^{t-1}) - (1-\delta)k_{t-1}(s^{t-2}) + i_t(s^t) - i_{t-1}(s^{t-1}) = y_t(s^t),$$

$$h_t(s^{t-1}) = \mathbf{h}_t(s^{t-1}).$$

Consumer Maximization: The quantities $\{c_t(s^t)\}$ and $\{h_t(s^{t-1})\}$ maximize (A1) subject to (A2).

Firm Profit Maximization: The quantities $\{y_t(s^t)\}$, $\{k_t(s^{t-1})\}$, and $\{\mathbf{h}_t(s^{t-1})\}$ maximize (A5) subject to (A5a)–(A5b).

Speculator Profit Maximization: The quantities $\{i_t(s^t)\}$ maximize (A9) subject to (A9a) – (A9b).

A.2 Equivalence of the Competitive Equilibrium and the Solution to the Planning Problem.

To demonstrate the equivalence between the competitive equilibrium and the solution to our planning problem, (13), we need only show that both satisfy the same first order conditions. In showing this we assume the value function, (13), is differentiable in its first two arguments, implicitly assuming that the "grid" on $\bar{k}_t$ and $\bar{i}_t$ is continuous. The result then follows from the fact that the two problems satisfy the same initial conditions and resource constraint.

Consider hours worked first. Substituting (A8) into (A3), we get

$$\beta^t \ell'(h_t(s^{t-1}))f^{t-1}(s^{t-1}) = -\sum_{s_t} p_t(s^{t-1},s_t)y_t^h(s^{t-1},s_t)$$

$$= \beta^t \sum_{s_t} f^t(s_t,s^{t-1})u'(c_t(s_t,s^{t-1}))y_t^h(s^{t-1},s_t),$$

after substituting from (A4). Rearranging:

$$\text{(A12)}\quad \gamma = \sum_{s_t} f^t(s_t|s^{t-1})u'(c_t(s_t,s^{t-1}))y_t^h(s^{t-1},s_t)$$

$$= \sum_{s_t} f^t(s_t|s^{t-1})u'(\bar{c}_t(s_t,s^{t-1}))(1-\theta)\bar{y}_t(s^{t-1},s_t)/h_t(s^{t-1}),$$

where $f^t(s_t|s^{t-1}) = f^t(s_t,s^{t-1})/f^{t-1}(s^{t-1})$, the conditional probability of s_t given s^{t-1}. The identity $y_t^h = (1-\theta)y_t/h_t$ and the labor market clearing condition, $h_t(s^{t-1}) = h_t(s^{t-1})$ were used in passing from the first to the second equality in (A12). In the notation of the text,

$$f^t(s(\ell)|s^{t-1}) = p_\ell,\ \ell = 1,\ldots,M, \text{ for all } t,\ s^{t-1}.$$

Also, the variables with over bars in (A12) are defined in (10). It is easily confirmed that the first order condition for h, (18), coincides with (A12).

Next, consider the first order condition for physical capital, $k_t(s^{t-1})$. Substituting (A4) into (A7):

$$\sum_{s_t} f^t(s_t|s^{t-1})u'(c_t(s_t,s^{t-1}))$$

(A13)

$$= \sum_{s_t,s_{t+1}} \beta f^{t+1}(s_t,s_{t+1}|s^{t-1})u'(c_{t+1}(s_t,s_{t+1},s^{t-1}))[y^k_{t+1}(s_t,s_{t+1},s^{t-1})+1-\delta],$$

after dividing by $\beta^t f^{t-1}(s^{t-1})$. Multiply both sides of this expression by z_{t-1} and use the facts $u'(c_t) = u'(\bar{c}_t)z_t^{-1}$, $y^k_{t+1} = \theta y_{t+1}/k_t = \theta(z_{t+1}/z_{t-1})\bar{y}_{t+1}/\bar{k}_t$, $z_{t-1}/z_{t+1} = \exp[-(s_{t+1}+s_t)]$ to rewrite (A13) as follows:

$$\sum_{s_t} f^t(s_t|s^{t-1})u'(\bar{c}_t(s_t,s^{t-1}))\exp(-s_t)$$

(A14)

$$= \sum_{s_t,s_{t+1}} \beta f^{t+1}(s_t,s_{t+1}|s^{t-1})u'(\bar{c}_{t+1}(s_t,s_{t+1},s^{t-1}))$$

$$\times\{\bar{y}^k_{t+1}(s_t,s_{t+1},s^{t-1}) + (1-\delta)\exp[-(s_{t+1}+s_t)]\}.$$

Here, $\bar{y}^k_{t+1} \equiv \theta\bar{y}_{t+1}/\bar{k}_t$. After some algebra, it can be confirmed that the first order condition for $\bar{k}'$ in (13) coincides with (A14). Doing this requires computing the derivative of v in (13) with respect to its first argument.

Finally, consider $i_t(s_t)$. Substituting (A4) into (A10) and rearranging:

$$\{[\sum_{s_{t+1}} \beta f^{t+1}(s_{t+1}|s^t)u'(\overline{c}_{t+1}(s^t,s_{t+1}))\exp(-s_{t+1})] - u'(\overline{c}_t(s^t))\}$$

(A15)

$$\times \lambda\beta^t f^t(s^t)/z_t.$$

Apart from the last expression, this is the derivative of the expression to the right of the equality in (13) with respect to $\bar{I}'$. The expression in square brackets is the derivative of the planner's value function with respect to its second argument, while $-u'(\overline{c}_t(s^t))$ corresponds to the derivative of r with respect to its fourth argument.

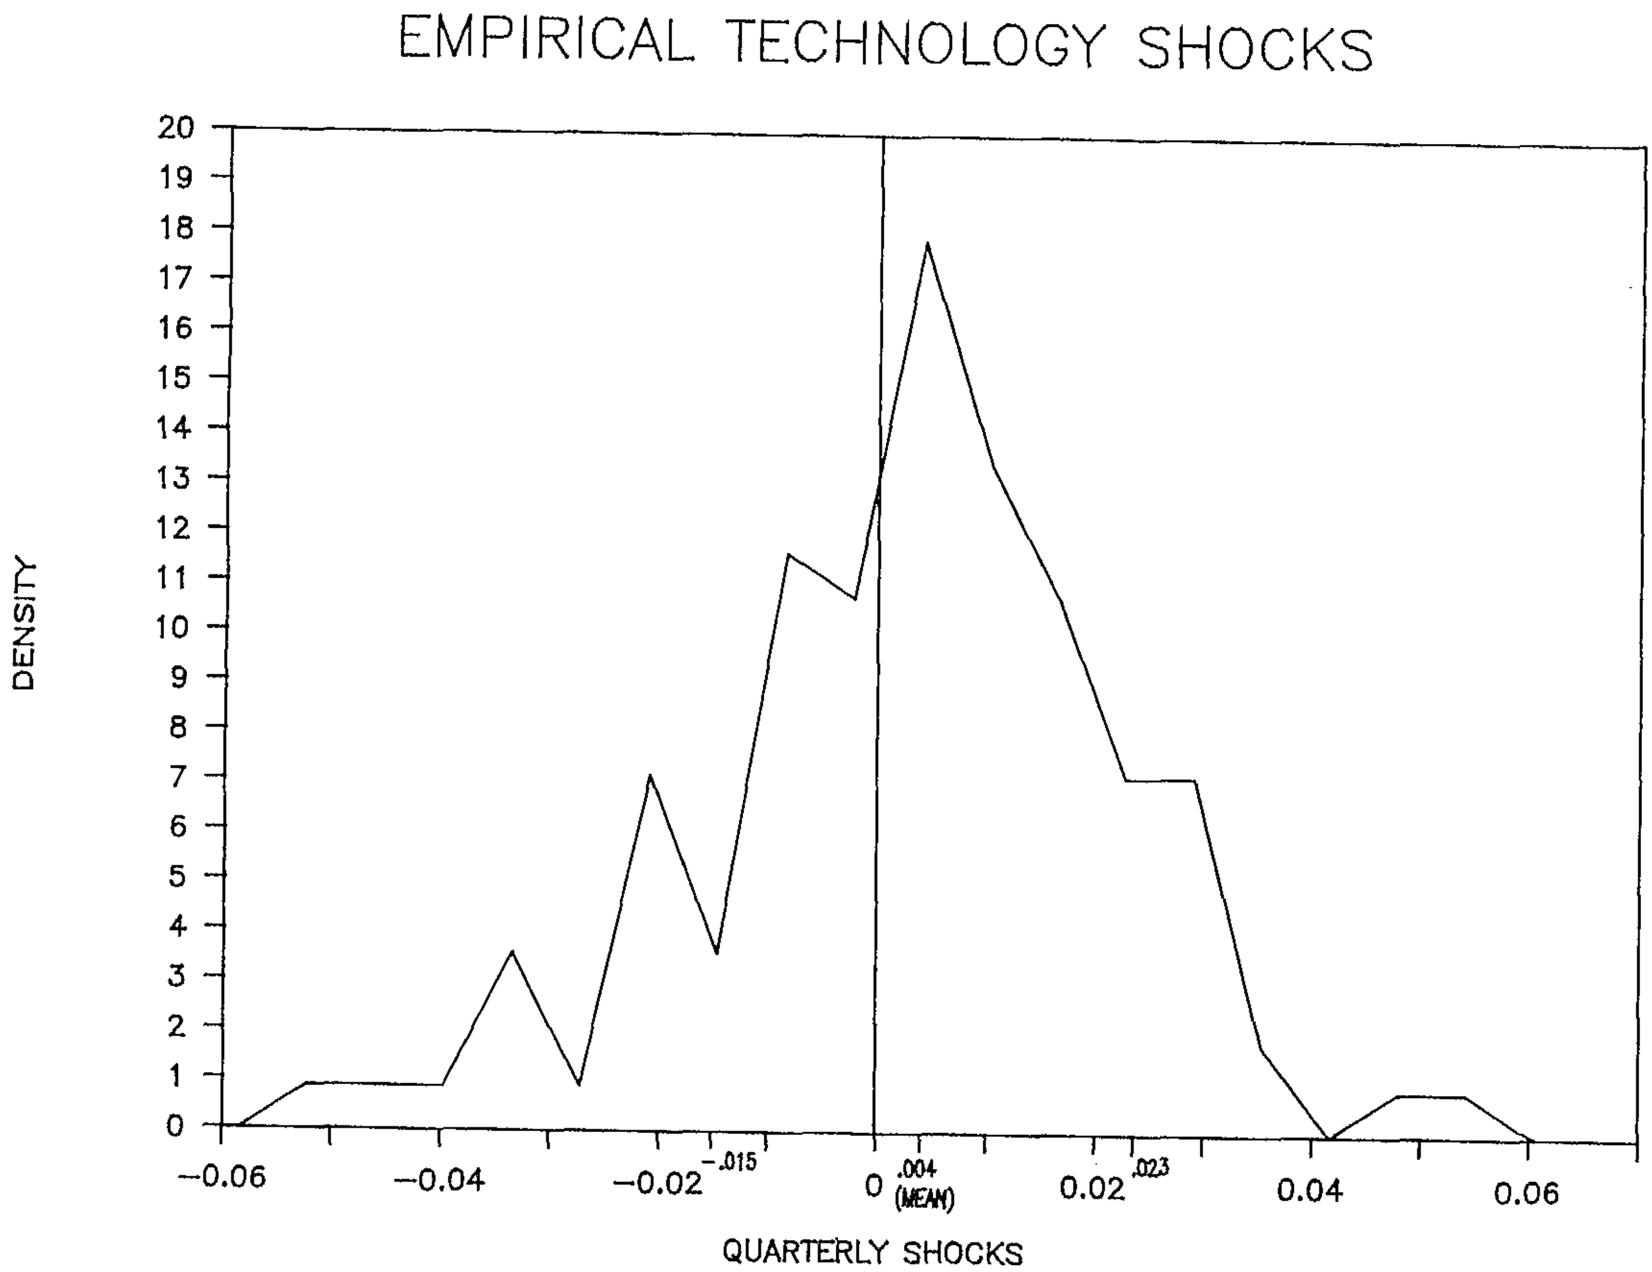

EMPIRICAL TECHNOLOGY SHOCKS
DENSITY
20
19
18
17
16
15
14
13
12
11
10
9
8
7
6
5
4
3
2
1
0
−0.06
−0.04
−0.02
−.015
0
.004
(MEAN)
0.02
.023
0.04
0.06
QUARTERLY SHOCKS

Figure 1

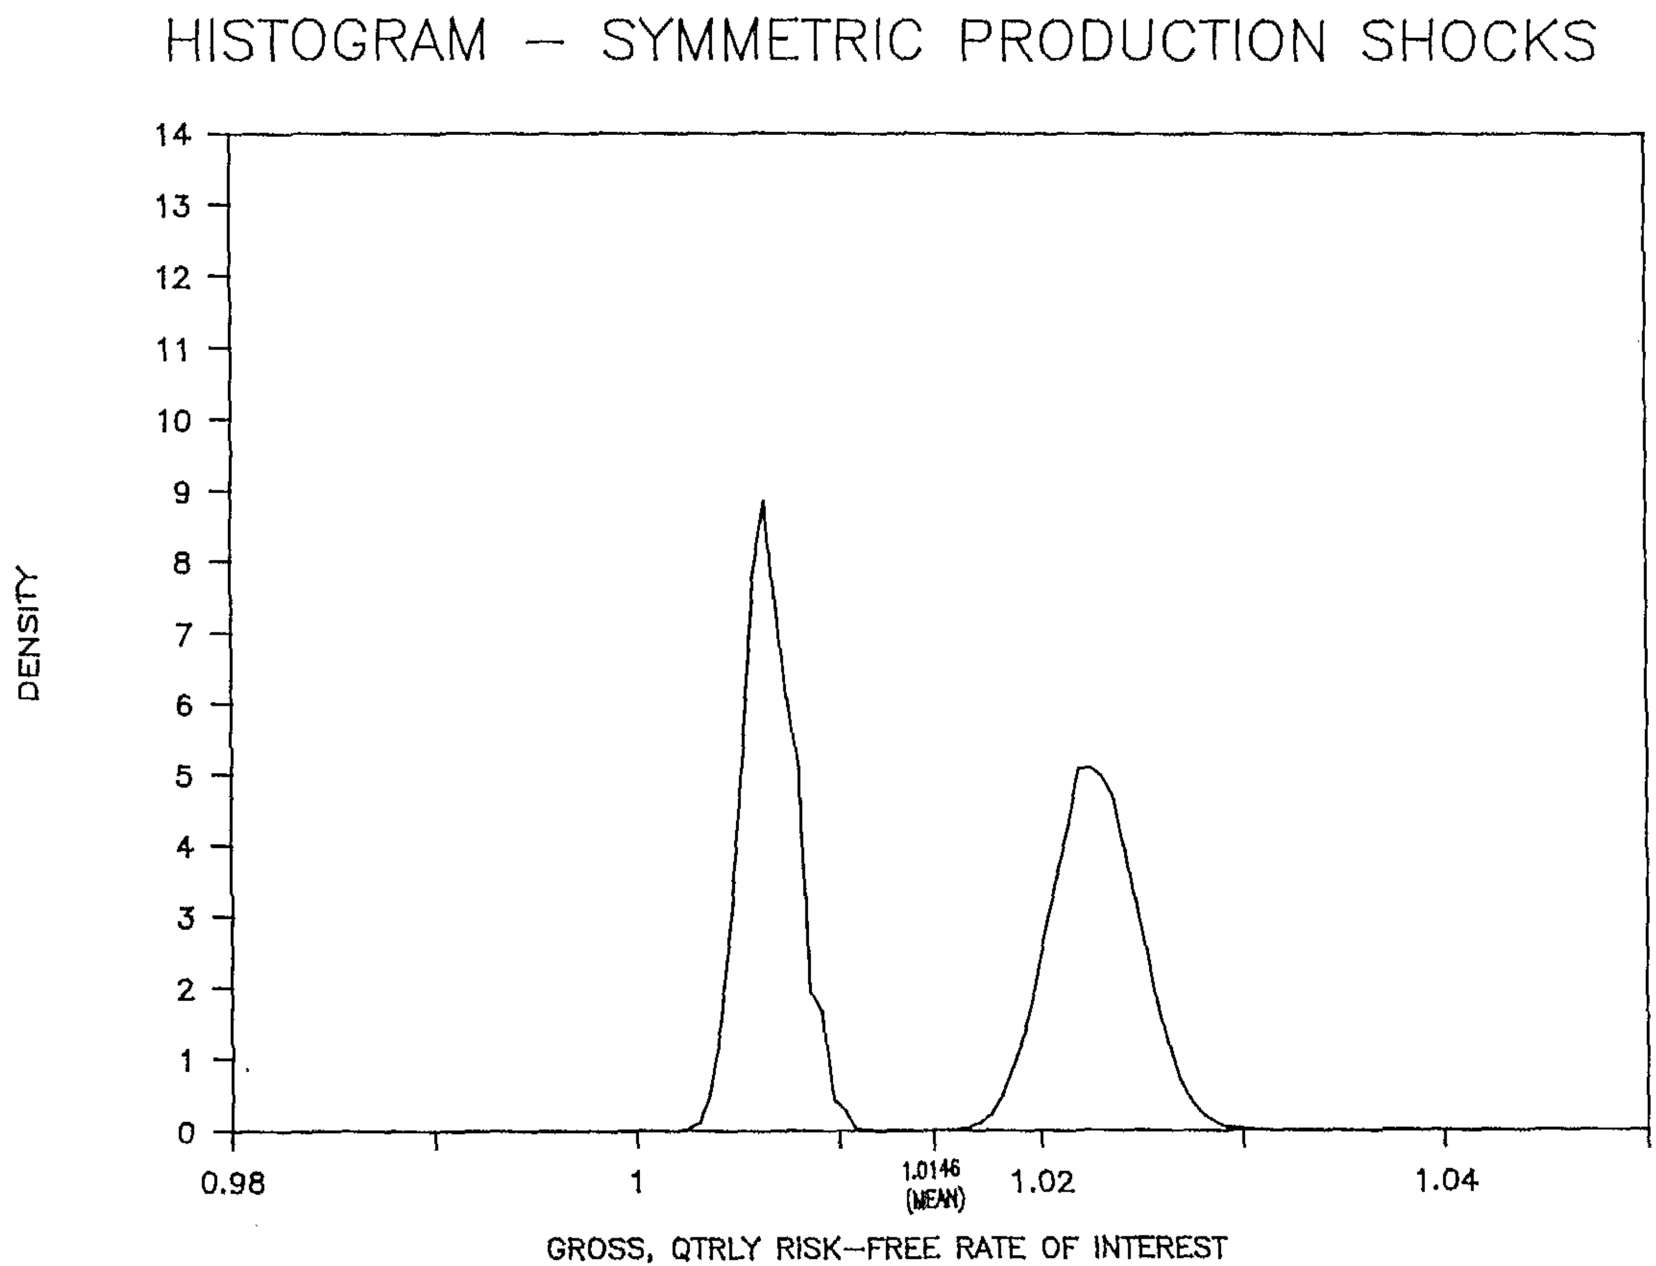

HISTOGRAM – SYMMETRIC PRODUCTION SHOCKS
DENSITY
14
13
12
11
10
9
8
7
6
5
4
3
2
1
0
0.98
1
1.0146
(MEAN)
1.02
1.04
GROSS, QTRLY RISK-FREE RATE OF INTEREST

Figure 2

HISTOGRAM — SKEWED PRODUCTION SHOCKS

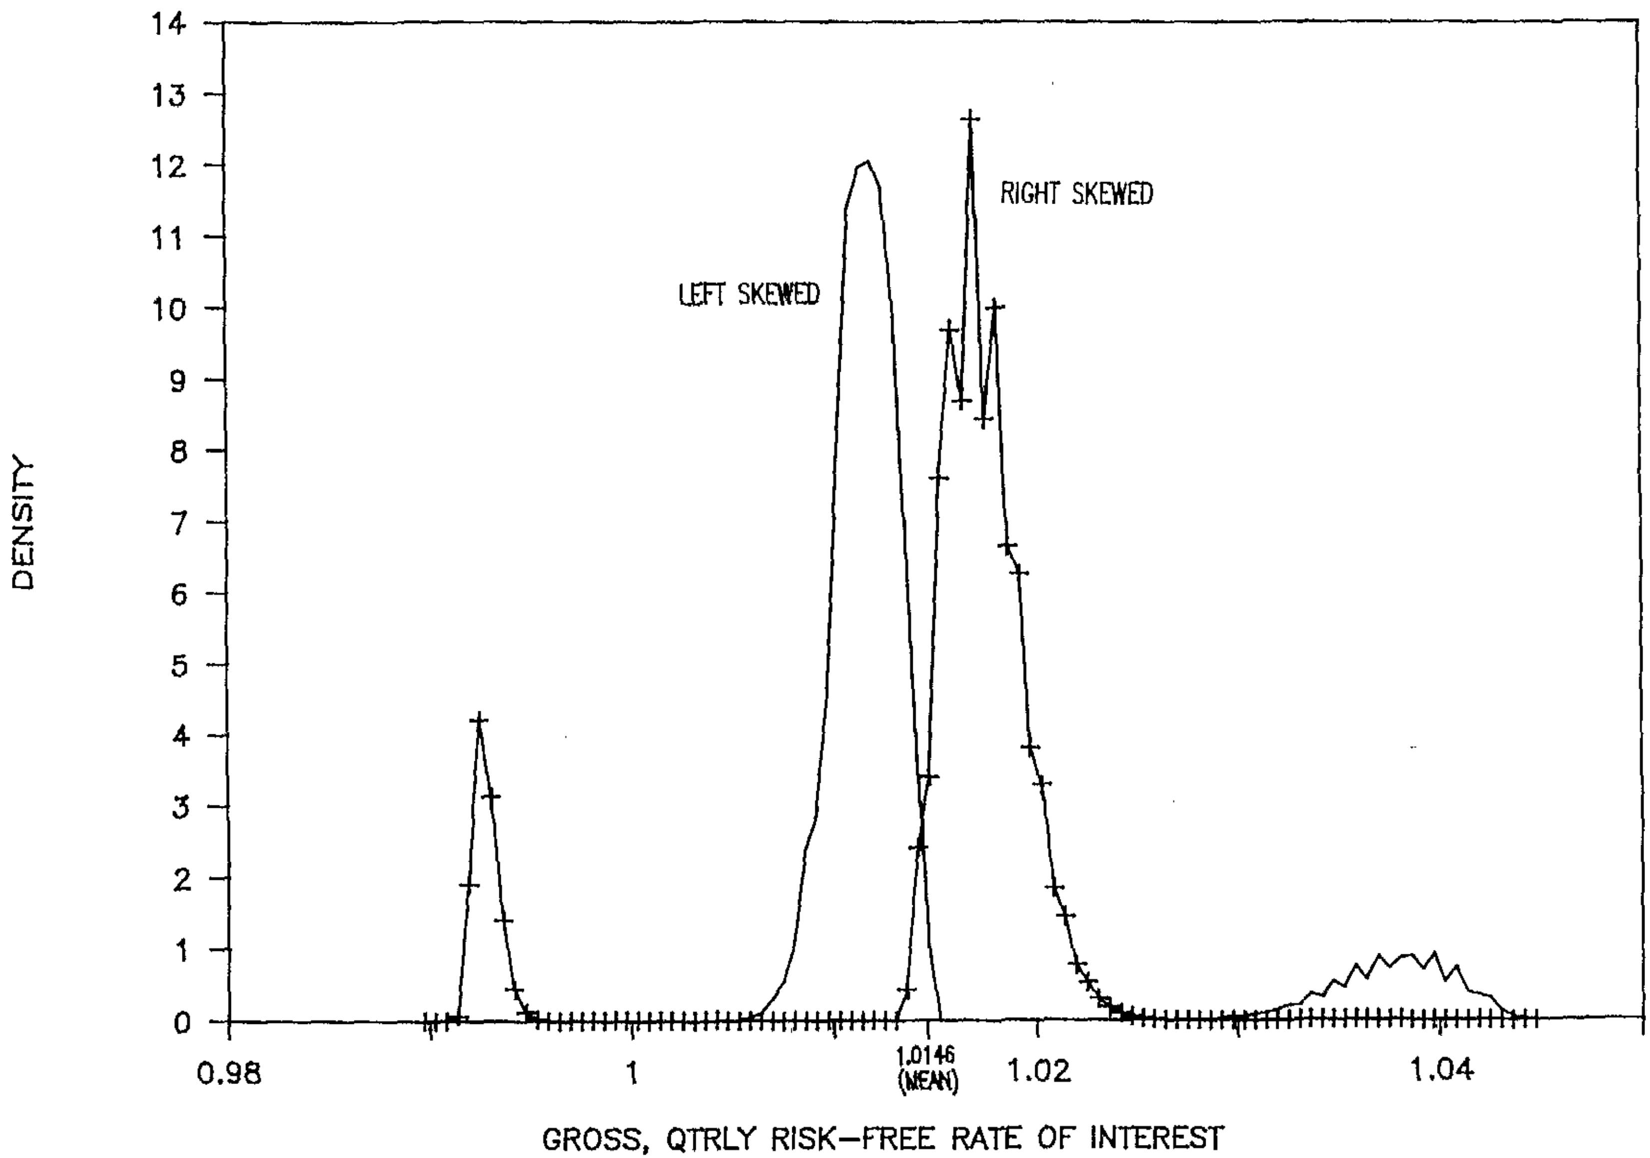

Figure 3

HISTOGRAM
RIGHT—SKEWED PRODUCTION SHOCKS

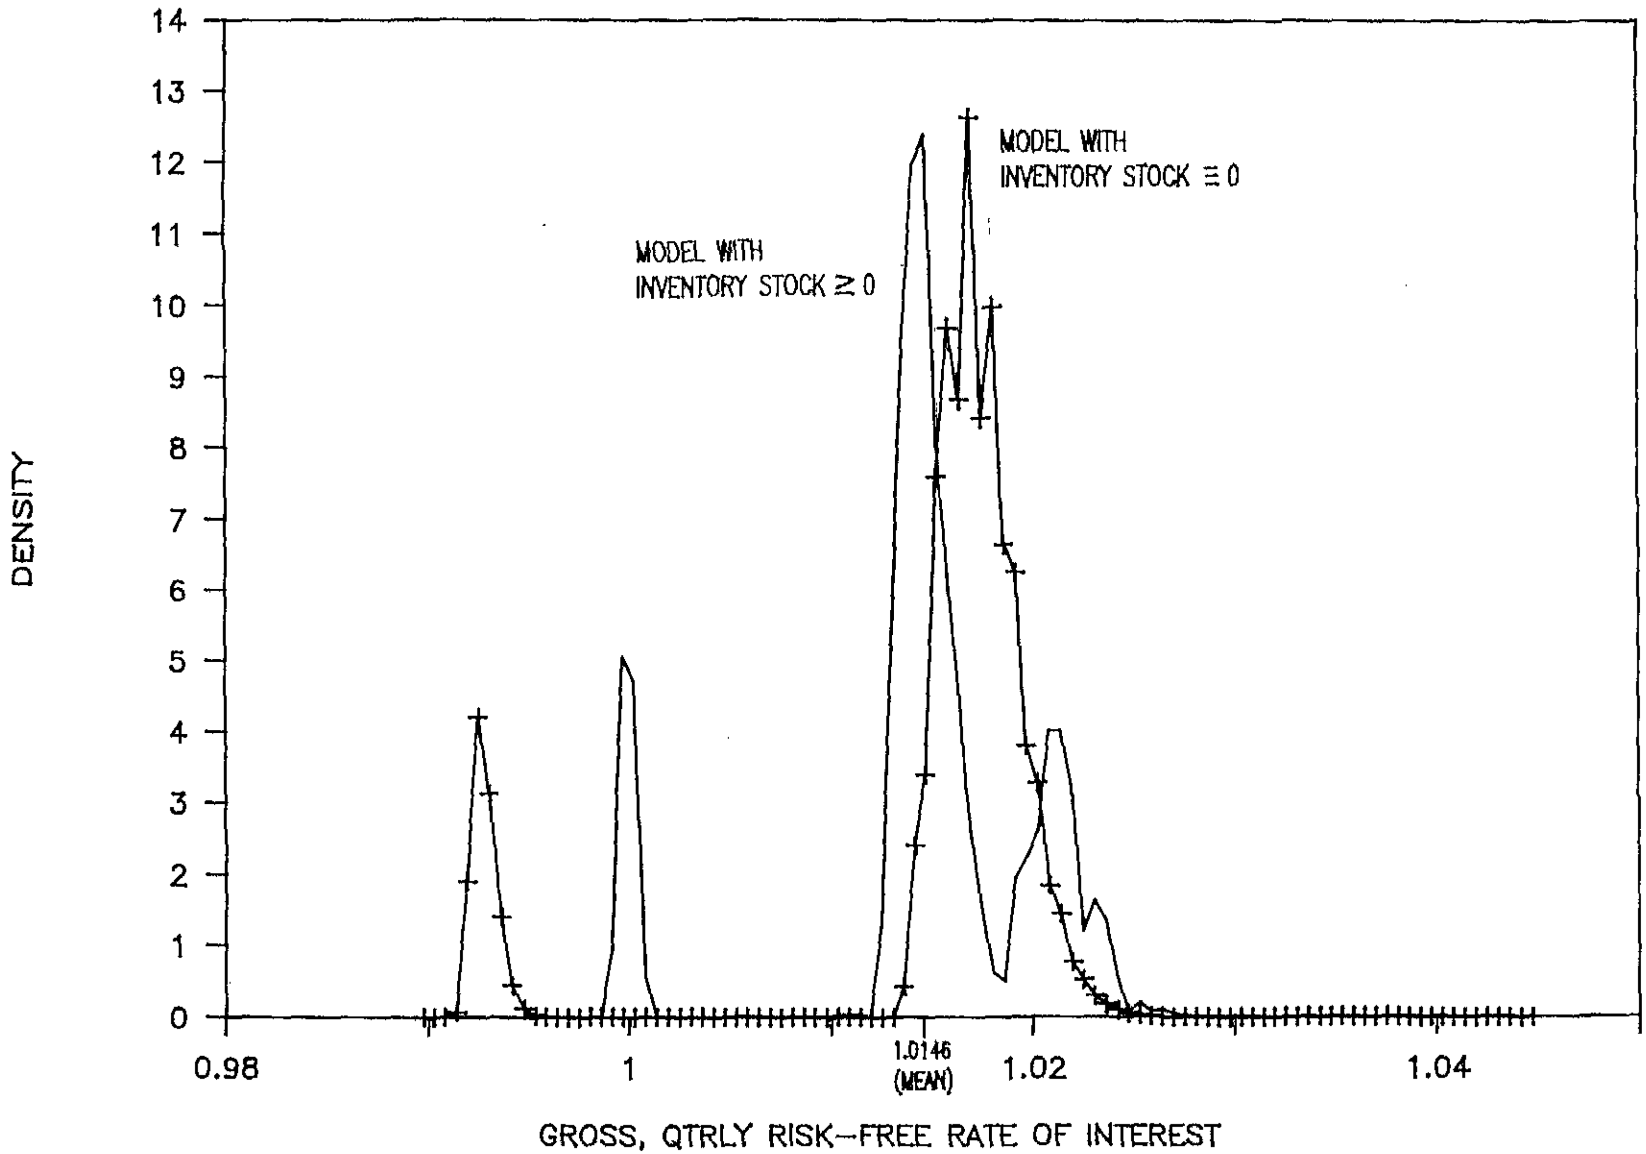

Figure 4