Working Paper Series

(S, s) Inventory Policies in General Equilibrium

Jonas D.M. Fisher and Andreas Hornstein

Working Papers Series
Research Department
(WP-96-24)

Federal Reserve Bank of Chicago

(S, s) Inventory Policies in General Equilibrium*

Jonas D.M. Fisher[a]   Andreas Hornstein[b]

November, 1994
Revised: February, 1998

Abstract

We study the aggregate implications of (S, s) inventory policies in a dynamic general equilibrium model with aggregate uncertainty. Firms in the model's retail sector face idiosyncratic demand risk, and (S, s) inventory policies are optimal because of fixed order costs. The distribution of inventory holdings affects the aggregate outcome in two ways: variation in the decision to order and variation in the rate of sale through the pricing decisions of retailers. We find that both mechanisms must operate to reconcile observations that orders are more volatile than, and inventory investment is positively correlated with, sales, while remaining consistent with other salient business cycle characteristics. The model exhibits strong amplification for some shocks, and persistence to a limited extent.

JEL Classification: E10, E22, E32

Keywords: Inventories, Business Cycles, Search, Matching

* Both authors thank SSHRC for financial support. We also thank Mark Bils, Hugo Hopenhayn, Ig Horstmann, Michael Parkin, Richard Rogerson, several anonymous referees and various seminar participants for helpful comments. Hornstein thanks the Institute for Empirical Macroeconomics for their hospitality during the time when this project was initiated. Any opinions expressed are those of the authors and do not necessarily reflect those of the Federal Reserve Banks of Chicago and Richmond or the Federal Reserve System. Our e-mail addresses are: jfisher@frbchi.org and hornstein@rich.frb.org.
[a] Federal Reserve Bank of Chicago.
[b] Federal Reserve Bank of Richmond.

1. Introduction

Variations in inventory investment form an integral part of the business cycle. In order to understand the role these variations play, considerable attention has been focused on (S, s) inventory policies. Starting with Blinder (1981) and Caplin (1985), a substantial body of literature has emerged which studies the aggregate implications of exogenously given (S, s) inventory policies in partial equilibrium settings. Conducting a complete study of the macroeconomic impact of these policies requires that they be analyzed in a general equilibrium framework. This, however, has proven to be difficult because of the presence of idiosyncratic risk and its implications for the aggregate distribution of inventory holdings. In this paper we describe a way to bridge this gap between partial and general equilibrium analysis.[1]

We study a dynamic general equilibrium model with aggregate uncertainty which includes a retail sector. Firms in the retail sector face idiosyncratic demand and use (S, s) inventory policies because of fixed order costs, as in Scarf (1960). The model builds on and extends the work of Caplin (1985) and Caballero and Engel (1991) who consider the aggregate implications of (S, s) inventory policies for a fixed number of sellers who follow potentially different, but fixed inventory policies, and who face exogenous sales.[2] Although we also restrict attention to the case of fixed inventory policies, these policies are optimal in our environment. In addition, we incorporate two general equilibrium feedback effects absent from the literature. First, the number of retailers ordering depends on the aggregate state of the economy via entry into the retail sector.[3] This contrasts with previous work in which the number of retailers ordering is a residual derived from sales and the distribution of inventories. Second, through the price setting behavior of retail firms there is feedback from the distribution of inventories to demand and sales.

We evaluate the general equilibrium model by how well it can account for two prominent

[1] Previous work on general equilibrium models with inventories uses reduced form representations in which inventories simply are another input to production. For examples, see Kydland and Prescott (1982) and Christiano (1988). Exceptions include Christiano and Fitzgerald (1991) and Bental and Eden (1993).

[2] Other studies along these lines include Blinder (1981) and Lovell (1993).

[3] This is similar to Benabou (1988) who studies a search economy without aggregate risk where sellers use (S, s) policies for setting prices and the number of sellers is endogenously determined through a free-entry condition.

stylized facts of inventory investment, namely that orders are more volatile than sales, and that inventory investment and sales are positively correlated.[4] We find that the model does account for these stylized facts; moreover the model also corresponds well to other salient characteristics of the business cycle. In order to account for the inventory and sales evidence both the entry and price feedback effects have to be operational. We also find that increasing fixed costs, thereby widening the (S, s) band, amplifies the responses of aggregate variables to some exogenous disturbances. To a limited extent propagation effects are present as well. Finally, we establish existence and uniqueness of a simplified competitive equilibrium version of our model in which entry is the only general equilibrium feedback. With the exception of inventory investment being positively correlated with sales, the business cycle implications of this model are similar to the more general model.

The rest of the paper proceeds as follows. In the next section we describe the inventory problem for a retailer with fixed ordering costs which we use to construct a simple partial equilibrium model of the retail sector. We discuss the aggregate implications of the optimal inventory policy for this example and compare them with previous work. In section 3 we embed the retail sector into a general equilibrium stochastic growth environment. In section 4 we employ the partial equilibrium model as a benchmark to illustrate and quantify the feedback effects in the general equilibrium model. In the fifth section we conclude.

2. (S, s) Inventory Policies in Partial Equilibrium

The foundation of our analysis is a simple example from Stokey and Lucas (1989, pp. 389-390) which captures the essential characteristics of (S, s) inventory models studied to date. In this section we extend the example to incorporate price setting by retailers and use it to formulate a benchmark partial equilibrium model which facilitates isolating the feedback effects in the general equilibrium model described in the next section. We exploit discreteness in the retailer

[4]Blinder and Maccini (1991) document the stylized facts of orders/production, sales and inventory investment for various sectors of the US economy. Fair (1989) has suggested that these stylized facts are an artifact of measurement error because production is not directly measured. Direct data on quantities are available for a limited number of goods producing industries and Fair (1989) and Krane and Braun (1991) find that in many of these industries production is less volatile than sales. Note, that these exceptions apply to a small number of manufacturing industries, not the trade sector on which we focus here.

environment as does Caplin (1985), and as in previous partial equilibrium work orders are a residual derived from sales and the inventory distribution. In the partial equilibrium model (S, s) inventory policies alone fail to account for the basic inventory facts and they have limited implications for business cycle analysis. These observations are consistent with previous work on one-sided (S, s) policies in partial equilibrium, for example Caplin and Spulber (1987).[5]

2.1. The Retailer Problem

Consider a risk-neutral retailer whose inventory stock can only take integer values, $x = 0, 1, \ldots, \bar{x}$. We assume that there is a capacity constraint on inventories, that is a retailer cannot hold more than $\bar{x}$ units of the good in inventory.[6] If a retailer places an order of size $x > 0$ she has to pay $\gamma_0 + \gamma_1 x$, with $\gamma_0, \gamma_1 > 0$. In any period a retailer meets at most one shopper, and the probability of this match is θ. Before a retailer is matched with a shopper she posts a price for the good, and a matched shopper may accept or reject the offer. The utility a shopper derives from the consumption of the good is random and not observed by the retailer. A shopper will accept the offer at price p and buy one unit of the good, if the utility from consumption is sufficiently high.[7] The acceptance probability δ is a decreasing function of the posted price. The probability that a retailer sells a good in any period is then $\theta\delta(p)$.

Let $R > 1$ be the gross real interest rate, then it is never optimal for the retailer to order before the inventory stock is depleted. This is so because postponing the order for one period does not affect revenue from sales, but reduces the present value cost of the order. The expected

[5]For a one-sided (S, s) policy the state variable monotonically declines without intervention and once the state variable reaches its lower bound s it is reset to the upper bound S. For a two-sided (S, s) policy the state variable may increase or decrease and when the state variable reaches its upper or lower bound it is reset to some interior point, $s < s^* < S$. Caplin and Leahy (1991) argue that one-sided and two-sided (S, s) policies can have very different implications for aggregate outcomes. For our inventory problem one-sided (S, s) policies are optimal.

[6]This is not a restriction since we can always choose $\bar{x}$ so that it does not constrain optimal decisions on inventory holdings. We introduce a capacity constraint because it simplifies the discussion of the general equilibrium environment in the next section.

[7]In a multistage model of bargaining where a seller makes sequential price offers to a buyer with unknown demand for the good, Sobel and Takahashi (1983) have shown that a take it or leave it offer with a posted price is optimal, if the seller can commit to a sequence of price offers, and the seller is at least as patient as the buyer.

capital value of a retailer with a positive inventory level x is

$$V_x = \max_p \left\{ \theta\delta(p) \left(p + R^{-1} V_{x-1} \right) + \left[1 - \theta\delta(p) \right] R^{-1} V_x \right\}$$

and the optimal order size solves

$$x^* = \arg\max_{x \leq \bar{x}} V_x - \gamma_0 - \gamma_1 x \text{ and } V_0 = V_{x^*} - \gamma_0 - \gamma_1 x^*.$$

Clearly an optimal inventory policy for the retailer is of the (S, s) variety. In particular the policy is to order nothing if x is positive and to order x^* units if $x = 0$.[8] There are no direct costs associated with inventory holdings, therefore V_x is positive and nondecreasing in x. In Proposition 1 (all propositions are stated and proved in the appendix) we show that under weak assumptions about demand $\delta(p)$, the incremental value of one unit of inventory $V_x - V_{x-1}$ is declining with higher inventory holdings.[9] This property is implied by the assumption of positive discounting and that an additional unit of inventory cannot be sold before all preceding inventory units have been sold.

The optimal price choice can be interpreted as a simple static profit maximization problem for a monopolist,

$$V_x = \max_p \left\{ \theta\delta(p) \left[p - R^{-1} \left(V_x - V_{x-1} \right) \right] + R^{-1} V_x \right\}, \tag{1}$$

where the marginal cost of a sale is the discounted present value of the capital loss of having one less unit of the good in inventory. Because the capital value of a retailer is a 'concave' function of the stock of inventories, the marginal cost of a sale is declining with higher inventories and retailers with higher inventories charge lower prices.

[8] In general, (S, s)-inventory policies allow for the possibility that the retailer places an order before all inventory is sold. In our environment the retailer will not order before she has sold all of her inventory, because, even though the event of a sale is random, the quantity the retailer can sell is fixed. If the quantity the retailer can sell is also random, the retailer may want to hold inventories to avoid a stock-out, Kahn (1987).

[9] In section 3 we describe an explicit model of demand and derive the acceptance probability from an environment where shoppers use reservation price strategies. The properties of demand needed for Proposition 1 then follow from assumptions on the preferences of shoppers.

2.2. Implications of (S, s) Inventory Policies in Partial Equilibrium

The study of (S, s) inventory models has been partly motivated by the observation that for these policies orders of a retailer are more volatile than sales. Aggregate orders and sales are, however, determined by the interaction of inventory policies and the distribution of inventory holdings. There is then no reason to expect that, in general, the behavior of microeconomic and aggregate variables is closely related.[10]

Consider our example from above, assume there is a continuum of identical retailers, and let the total measure of retailers be fixed at N. Suppose there is no aggregate risk and demand, that is the event of a sale, is constant and independent across retailers. In this case, the order and price policy derived above is optimal and it defines a Markov process on the inventory level for each retailer. Let $g_{x,t}$ denote the measure of retailers with inventory holdings x at the beginning of period t before any orders have been placed, let $\hat{g}_{x,t}$ denote the measure of retailers who have depleted their inventory stock in the previous period and order x units at the beginning of period t, and let $\bar{g}_{x,t} = g_{x,t} + \hat{g}_{x,t}$ denote the total measure of retailers with inventory holdings x in the market in period t. Note that $\sum_{x=1}^{\bar{x}} \bar{g}_{x,t} = N$.

In this environment retailers face idiosyncratic risk only, and, assuming a law of large numbers holds, the distribution of inventory stocks evolves as follows

$$\begin{aligned} g_{x,t+1} &= (1 - \theta\delta_x)\,\bar{g}_{x,t} + \theta\delta_{x+1}\bar{g}_{x+1,t} \\ \hat{g}_{x,t} &= \begin{cases} \theta\delta_1\bar{g}_{1,t-1} & \text{if } x = x^* \\ 0 & \text{if } x \neq x^* \end{cases} \end{aligned} \tag{2}$$

where $\delta_x = \delta(p_x)$ and p_x is the optimal price with inventory level x. The measure of retailers with inventory stock x declines because a fraction $\theta\delta_x$ of these retailers sell one unit and now have $x - 1$ units, and increases because a fraction $\theta\delta_{x+1}$ of retailers with inventory $x + 1$ sell one unit and now have x units. At the maximal inventory level x^* there is also the additional inflow of retailers who are placing orders. Based on the distribution of inventory stocks we can define aggregate orders as $Y_t = \sum_{x=1}^{\bar{x}} x\hat{g}_{x,t}$, aggregate sales as $C_t = \theta\sum_{x=1}^{\bar{x}} \delta_x\bar{g}_{x,t}$, and inventory

[10] This point has been made before, see for example Caballero (1992).

investment as $I_{t+1} - I_t = Y_t - C_t$.

A simple application of results for the convergence of Markov chains shows that the distribution of inventory levels after orders have been placed converges to an invariant distribution[11]

$$\bar{g}_x = \frac{\delta_x^{-1}}{\sum_{i=1}^{x^*} \delta_i^{-1}} N \text{ for } 1 \leq x \leq x^* \text{ and } \hat{g}_x = \begin{cases} \theta \delta_1 \bar{g}_1 & \text{if } x = x^*, \\ 0 & \text{if } x = 0. \end{cases} \tag{3}$$

In general the inventory distribution is not uniform, but weighted towards lower inventory levels, because retailers with lower inventories charge higher prices and are therefore less likely to make a sale. In a similar environment in which the sales probabilities are independent of the retailer's inventory holdings, Caplin (1985) shows that the inventory distribution converges to a uniform invariant distribution, and that for a finite number of sellers aggregate orders are more volatile than aggregate sales. Nevertheless a law of large numbers still holds and as the number of sellers increases, the variability of average orders and sales converges to zero: aggregate orders and sales are constant.

We now introduce aggregate risk and develop a partial equilibrium example which will help us understand the mechanism of the general equilibrium model studied in the remainder of the paper. Let aggregate demand be random, in particular let the probability of a match θ_t be a stationary stochastic process with domain $\left[\underline{\theta}, \bar{\theta}\right]$, $0 < \underline{\theta} < \bar{\theta} < 1$. We also assume that the number of retailers in the economy is fixed, and that other economic variables, like the acceptance probability as a function of price or the interest rate, remain constant. In this environment it is still optimal for a retailer not to order before she has sold all of her inventory, but the order size and pricing policy will in general be state contingent. In the following sections we describe conditions for which a fixed order size and pricing policy remain optimal in a more general environment. We may assume that these conditions are also satisfied here and that the order size x^* and the sales probability δ_x do not depend on the state of the economy, in particular the probability of a match θ_t.[12] Our analysis is partial equilibrium in nature because we assume

[11] See Stokey and Lucas (1989), Theorem 11.4.

[12] This follows the literature on the aggregation of exogenously given (S, s) policies, see for example Caplin and Spulber (1987) or Caballero and Engel (1991).

that aggregate demand is exogenous to the retail market, and that there are no feedbacks from aggregate demand, sales or inventory investment to other aggregate economic variables.

We show in Proposition 2 that, for this example the distribution of inventory holdings after orders have been placed still converges to the invariant distribution defined in (3). This result corresponds to the work by Caballero and Engel (1991) on the convergence of the cross-sectional distribution of a fixed measure of agents who follow one-sided (S, s) policies. As shown in Proposition 3, after the distribution of inventories has converged, aggregate variables behave as follows: (i) the variance of orders equals the variance of sales, (ii) inventory investment and sales are negatively correlated; (iii) orders and sales are not more volatile than demand and the autocorrelations of orders and sales equal the autocorrelation of demand; and (iv) orders and sales are positively correlated only if demand is positively autocorrelated.

These findings suggest two conclusions. First, (i) and (ii) indicate the aggregate implications of (S, s) inventory policies are not necessarily consistent with the two key observations from the inventory literature. Second, (iii) and (iv) suggest that (S, s) inventory policies do not necessarily propagate or amplify shocks. We have noted above that these conclusions in large part reflect previous results on the aggregation of exogenously given one-sided (S, s) policies in partial equilibrium environments.

In the preceding analysis heterogeneity across retailers represented by the inventory distribution plays no role at all. In fact, despite the presence of aggregate uncertainty, the inventory distribution converges to a time invariant distribution. One reason why (S, s) policies have such a limited effect is that the inventory distribution affects aggregate variables only through the order decision of retailers who have stocked out in the previous period. We expect that a general equilibrium analysis of (S, s) inventory policies will yield different results than the partial equilibrium analysis, since there are more channels through which the inventory distribution, as part of the aggregate state of the economy, can affect outcomes in the retail sector.

At this point we can identify three channels through which general equilibrium feedbacks can operate. First, aggregate demand in the retail sector, as represented by the matching probability θ and the acceptance probability δ, will depend on the state of the economy. Second, the number of retailers N operating in the economy will depend on the state of the economy, since the profitability of retailers is state contingent. Third, the decisions of a retailer on price p_x and order size x^* will be, in general, state contingent.

3. (S, s) Inventory Policies in a General Equilibrium Model

Here we embed the retail sector described in the previous section into a variant of the standard stochastic growth model. There is an infinitely lived representative household whose utility depends on non-leisure activities and the consumption of goods which are bought from a large number of retailers. The household has income from the supply of labor to a manufacturing sector and dividends from ownership of manufacturers and retailers, and trades in state contingent claims. The manufacturing sector cannot sell its product directly to the household, but has to sell through the retail sector. Because of trading frictions in the retail sector the representative household has to search for consumption goods, and each retailer faces a stochastic demand for her good. A retailer posts a price for her product which will affect the probability of a sale, and the probability of a sale is consistent with the optimal search behavior of the household. We model the trading frictions using a matching type environment, where each retailer is randomly matched with the household.

Without aggregate uncertainty the economy converges to a steady state. Analogous to the previous section, in the steady state each retailer faces idiosyncratic uncertainty, but the aggregate distribution of retailers over inventory holdings is invariant. We introduce aggregate uncertainty through random shocks to the production of goods and to the time discount factor. In the remainder of this section we describe the environment in more detail, and then outline a recursive equilibrium. A tractable example is introduced and discussed to conclude the section.

3.1. The Environment

The preferences of the representative household are

$$E_0 \left[\sum_{t=0}^{\infty} \left(\Pi_{\tau=0}^{t} \beta_\tau \right) \ln \left(c_t - \frac{n_t^{1+\nu}}{1+\nu} \right) \right] \tag{4}$$

where c_t is an index of consumption goods, and $n_t \leq \bar{n}$ is time spent in non-leisure activities. The discount factor β_t is an exogenous, possibly random, variable with bounded support, $0 < \beta_t \leq \bar{\beta} < 1$. The wage elasticity of non-leisure activities is $1/\nu$ and $\nu > 0$.[13]

The consumption index is defined as follows. In each period there is a continuum of retailers of measure N_t. For the household each retailer provides a unique good, and a household can consume exactly one unit of each retailer's good in a period. The household cannot store the good of any retailer. Let $c_t(i) \in \{0,1\}$ be the consumption of a good obtained from retailer $i \in [0, N_t]$. The utility derived from the consumption of the i-th good, $u_t(i)$, is an iid random variable with support $\mathbf{U} = [\underline{u}, \overline{u}]$, $0 \leq \underline{u} < \overline{u} \leq \infty$, and twice differentiable cumulative density function F. The household is not informed about the quality of any retailer's good before it actually meets the retailer. The consumption index is the integral over all utilities

$$c_t \equiv \int_0^{N_t} u_t(i) c_t(i) di. \tag{5}$$

There are frictions in the market for retail goods due to a cost of contacting retailers. To contact retailers the household must search, which requires that it spend some time shopping, $n_{s,t}$. The quantity of search effort produced is

$$s_t = \zeta n_{s,t}, \tag{6}$$

where $\zeta > 0$ measures the efficiency of time spent shopping. Interpret this as the household splitting up into s_t sub-units, shoppers for short, who search independently. A shopper is randomly matched with at most one retailer and for any shopper the probability of a match with a retailer is ω_t. The total number of shoppers who contact retailers is then $\omega_t s_t$.

Shopping is one of the two non-leisure activities the household pursues. The other activity

---

[13] These preferences are used by Greenwood, Hercowitz and Huffman (1988) and others. They imply that wage changes have no wealth effect (a common implication of empirical labor supply studies). Except for procyclical average sales, preferences separable in consumption and leisure have the same implications as these preferences. The current formulation mimics the behavior of an environment with home production and logarithmically separable preferences, as in Benhabib, Rogerson and Wright (1991), which we studied in an earlier draft of this paper.

is working in the manufacturing sector, $n_{y,t}$. Total non-leisure activities are

$$n_t = n_{y,t} + n_{s,t}. \tag{7}$$

Labor is the only input to production in the manufacturing sector,

$$y_t = z_t n_{y,t}^{\alpha} \tag{8}$$

with $0 < \alpha \leq 1$ and productivity z_t is an exogenous, possibly random, variable with bounded support, $0 < z_t \leq \bar{z}$.[14] The goods produced in the manufacturing sector are distributed to the retail sector, which resembles the one described in section 2. Before orders are placed there is a measure $g_{x,t}$ of retailers with inventory holdings x. An order of size x uses up $\gamma_0 + \gamma_1 x$ units of the manufactured good. Unlike the previous section, we assume there is free entry to the retail sector. We also assume there are no explicit costs to entry besides the cost of making an order.[15] Let $\hat{g}_{x,t}$ denote the measure of retailers who order x units of the market good, then the resource constraint for the good produced in the manufacturing sector is

$$\sum_{x=1}^{\bar{x}} \hat{g}_{x,t} \left(\gamma_0 + \gamma_1 x\right) \leq y_t, \tag{9}$$

and $\bar{g}_{x,t} = g_{x,t} + \hat{g}_{x,t}$ is the measure of retailers with inventory holdings x after orders have been placed.[16] The total number of retailers in the market is $N_t = \sum_{x=1}^{\bar{x}} \bar{g}_{x,t}$.

The presence of search costs implies that shoppers cannot pick retailers at will. Instead they are randomly matched with retailers. In any period a retailer (shopper) is matched with at most one shopper (retailer). If there is a total measure of N_t retailers and S_t shoppers, then the

[14] Decreasing returns to labor may be due to a fixed factor of production, for example capital, which we do not model.

[15] Note that this means there is no difference between existing retailers who have sold all of their inventory and are now restocking, and new retailers coming into the market.

[16] Here we make use of the fact that in equilibrium retailers with positive inventory holdings do not place any orders.

number of matches is given by

$$M_t = \min\left\{\varphi S_t^{\mu} N_t^{1-\mu}, N_t, S_t\right\} \quad (10)$$

with $0 < \mu < 1$ and φ denotes the efficiency of matching. The probability that a shopper (retailer) meets a retailer (shopper) is $\omega_t = M_t/S_t$ ($\theta_t = M_t/N_t$). Retailers and shoppers take these probabilities as exogenous parameters.

The environment is subject to aggregate uncertainty since preferences and technologies are subject to exogenous stochastic shocks, $\mathbf{z}_t$. In what follows we focus on two cases, either the time discount factor is random, $\mathbf{z}_t = \beta_t$ and $z_t = z$, or the productivity level is random, $\mathbf{z}_t = z_t$ and $\beta_t = \beta$. We assume that the exogenous state follows a first order Markov process with transition probabilities $\Phi(\mathbf{z}_{t+1}|\mathbf{z}_t)$. A complete description of the environment in each period also includes the beginning of period distribution of inventory holdings before orders have been placed, $g_t = (g_{1,t}, g_{2,t}, \ldots, g_{\bar{x},t})$. This endogenous state variable can change because of differential changes in the probability of a sale for retailers with different inventories, or because of changes in the rate at which new retailers enter, equivalently the rate at which inventories are restocked. This concludes the description of the environment.

3.2. A Recursive Equilibrium

In this sub-section we describe an equilibrium in which consumers and producers of manufacturing goods behave competitively, and retailers behave as described in section 2.1. We are interested in a recursive equilibrium where prices and the decisions of agents are time invariant functions of the aggregate state of the economy. We assume that agents can write contracts which are contingent on the state of the economy, and that there is a complete set of markets in these state contingent claims.

The description of the household's problem proceeds in two steps. First we discuss the optimal search for consumption goods, and then describe the full dynamic optimization problem.[17] Recall from section 2.1 that a retailer with inventory x posts a price p_x for her product before any

[17] In the following we will drop time subscripts and let a prime denote next period's value of a variable.

shopper arrives. Since posted prices depend on inventory holdings, which differ across retailers, not all retailers post the same price. Thus, if a shopper is matched with the i-th retailer, he will face a price p_x with probability $\bar{g}_x/N$, the normalized distribution over inventory holdings after orders have been placed. The shopper then inspects the good, that is observes $u(i)$, and decides whether to purchase the good or not. Note that his decision cannot depend on the choices of other shoppers since shoppers do not communicate with each other. Also shoppers will not purchase more than one good since the household cannot store goods.

Under these circumstances the shopper will use a reservation utility policy, that is he will purchase the good if his utility is above some reservation utility r_x, which may depend on the posted price. For such a reservation utility policy the expected utility value from a match is

$$q = \sum_{x=1}^{\bar{x}} \int_{u \geq r_x} uF(du)(\bar{g}_x/N), \tag{11}$$

where $0 \leq q \leq E[u]$, and the expected expenditure for a match is

$$e = \sum_{x=1}^{\bar{x}} p_x \left[1 - F\left(r_x\right)\right] \left(\bar{g}_x/N\right). \tag{12}$$

The consumption good index implied by this policy is $c = q\omega s$.

The budget constraint of the representative agent depends on the wage rate w, non-labor income π, the price distribution for goods $\{p_x, \bar{g}_x\}_{x=1}^{\bar{x}}$, the match probability of a shopper ω, the prices Q for state contingent claims a, and wealth. In a recursive equilibrium all prices are time invariant functions of the aggregate state, and the distribution of inventory holdings evolves according to a time invariant law of motion. We denote the time invariant equilibrium function of an aggregate variable with the corresponding capital letter of the variable, for example $w = W(g, \mathbf{z})$, and the law of motion for the inventory holding distribution is $g' = G(g, \mathbf{z})$.

Without loss of generality we assume that the only intertemporal trades are in state contingent claims. Let $Q(g', \mathbf{z}'|g, \mathbf{z})$ be the price of a state contingent claim $a(g', \mathbf{z}'|g, \mathbf{z})$ which pays one unit of the manufacturing good if next period's aggregate state is $(g', \mathbf{z}')$, conditional on the current aggregate state $(g, \mathbf{z})$. The risk free rate of return is then $R(g, \mathbf{z}) =$

$1/\int Q\left(G\left(g,\mathbf{z}\right),\mathbf{z}'|g,\mathbf{z}\right)d\mathbf{z}'$. Let $\tilde{a}$ denote the number of state contingent claims which pay off in the current period. The state of the household is then given by the vector $(\tilde{a},g,\mathbf{z})$.

We are now in a position to describe the optimization problem of the representative household,

$$
\begin{aligned}
\mathcal{U}\left(\tilde{a},g,\mathbf{z}\right)= \quad & \max \quad \left\{U\left(c,n\right)+\beta E\left[\mathcal{U}\left(a\left[g',\mathbf{z}'|g,\mathbf{z}\right],g',\mathbf{z}'\right)|g,\mathbf{z}\right]\right\} \\
& \text{s.t.} \quad c=\zeta n_s\omega q, \\
& \qquad \zeta n_s\omega e+\textstyle\int Q(g',\mathbf{z}'|g,\mathbf{z})a(g',\mathbf{z}'|g,\mathbf{z})d\mathbf{z}'=\tilde{a}+wn_y+\pi, \\
& \qquad n=n_s+n_y,\text{ and} \\
& \qquad g'=G\left(g,\mathbf{z}\right),\omega=\Omega\left(g,\mathbf{z}\right),p_x=P_x\left(g,\mathbf{z}\right), \\
& \qquad w=W\left(g,\mathbf{z}\right),\pi=\Pi\left(g,\mathbf{z}\right),
\end{aligned}
\tag{13}
$$

where U is the period utility function. The household's static decision problem regarding the acquisition of consumption goods can be separated into the following two stage decision problem. First, for any desired average quality resulting from a match the agent chooses a reservation utility policy which minimizes the cost associated with that quality. The implied cost function, $e\left(q\right)$, represents a price index for average quality. Let λ be the marginal cost of quality for this policy. It is straightforward to show that the first order conditions for reservation utilities are sufficient to characterize the optimal reservation utility policy. Next, the agent chooses the optimal mix of shopping time and desired average quality. At this stage the objective function is not necessarily concave with respect to the choice variables n_s and q. In Proposition 4 we show that the problem can be transformed to a concave problem, and that the solution to this problem is unique.

Assuming an interior solution exists to the static decision problem regarding the acquisition of consumption goods and the choice of non-leisure time, an optimal decision has to satisfy the following first order conditions

$$
\begin{aligned}
U_c r_x &= \kappa p_x \\
U_c\zeta\omega q-U_n &= \kappa\zeta\omega e \\
U_n &= \kappa w
\end{aligned}
\tag{14}
$$

where $U_c=\partial U/\partial c$, $U_n=-\partial U/\partial n$, and the Lagrange multiplier κ is the marginal utility of income. The first and third equations of (14) combined state that the marginal rate of substitution between non-leisure time and the marginal accepted consumption good is equal to the price of the marginal accepted consumption good,

$$\frac{r(p)U_c}{U_n} = \frac{p}{w} \text{ for } r(p) \in (\underline{u}, \overline{u}).$$

This suggests that we rewrite the reservation utility policy as follows

$$r(p) = \frac{p}{\lambda} \text{ with } \lambda = w\frac{U_c}{U_n} \text{ and } p \in (\lambda\underline{u}, \lambda\overline{u}). \tag{15}$$

We can now modify the retailer's problem from section 2.1 for a recursive equilibrium with aggregate uncertainty. Since shoppers follow reservation price strategies the probability that a shopper will accept the offer at price p is $\delta(p/\lambda) = 1 - F(p/\lambda)$. This probability is also the demand function which a retailer faces, and it is a nonincreasing function of the posted price. Next period's returns are discounted using the state contingent prices Q. From the optimal intertemporal decision of the household we have that the normalized marginal rate of substitution between present and future non-leisure time is equalized to the contingent claim price

$$Q(g', \mathbf{z}'|g, \mathbf{z}) = \tilde{Q}(g', \mathbf{z}', g, \mathbf{z})\,\Phi(\mathbf{z}'|\mathbf{z}) \text{ and } \tilde{Q}(g', \mathbf{z}', g, \mathbf{z}) = \beta\frac{U_n(c', n')/W(g', \mathbf{z}')}{U_n(c, n)/W(g, \mathbf{z})}. \tag{16}$$

The capital value of a retailer which has x units of inventory and who does not reorder before the inventory is depleted is

$$\begin{aligned} V_x(g, \mathbf{z}) = \max\nolimits_p \quad & \left\{\theta\delta(p/\lambda)\left(p - E\left[\tilde{Q}(g', \mathbf{z}'|g, \mathbf{z})\, V_{x-1}(g', \mathbf{z}') \mid g, \mathbf{z}\right]\right)\right. \\ & \left. + [1 - \theta\delta(p/\lambda)]\, E\left[\tilde{Q}(g', \mathbf{z}'|g, \mathbf{z})\, V_x(g', \mathbf{z}') \mid g, \mathbf{z}\right]\right\}, \\ & \text{where } g' = G(g, \mathbf{z}),\ \theta = \Theta(g, \mathbf{z}) \text{ and } \lambda = \Lambda(g, \mathbf{z}). \end{aligned} \tag{17}$$

Given free entry, it is not optimal to reorder before all inventories are depleted if V_x defined in (17) is 'concave'. To see this, consider a retailer with current inventory x and an order of size Δx. If the retailer adds the order to the current inventory, the capital value is $V_{x+\Delta x}$.

Alternatively, because of our free entry assumption the retailer could simply open up another store, in which case her capital value would be $V_x + V_{\Delta x}$. Since V is concave the second option gives a higher capital value. We do not have a proof that $V_x(g, \mathbf{z})$ is a concave function in x, but in our computational experiments below we verify it is always concave.

The optimal order policy maximizes the gains from entry, and because of free entry the rate at which retailers restock their inventory will adjust such that these gains are zero in equilibrium,

$$V_x(g, \mathbf{z}) \leq \gamma_0 + \gamma_1 x \text{ and } V_x(g, \mathbf{z}) = \gamma_0 + \gamma_1 x \text{ if } \hat{g}_x(g, \mathbf{z}) > 0. \tag{18}$$

The law of motion for the inventory distribution G is then defined by

$$\begin{gathered} \bar{g}_x = g_x + \hat{g}_x \text{ and } g'_x = \left[1 - \theta\delta\left(p_x/\lambda\right)\right]\bar{g}_x + \theta\delta\left(p_{x+1}/\lambda\right)\bar{g}_{x+1} \\ \text{for } x = 1, \ldots, \bar{x} \text{ and } \bar{g}_{\bar{x}+1} = 0. \end{gathered} \tag{19}$$

Production of the manufacturing good is competitive so that manufacturers solve,

$$\max z n_y^{\alpha} - W(g, \mathbf{z})\, n_y. \tag{20}$$

Here we have normalized the price of the manufactured good at one. The solution to this problem determines the real wage.

The following definition of the recursive equilibrium summarizes this section.[18]

Definition 1. *A recursive equilibrium is a collection of functions* $\{\Lambda, \Theta, \Omega, \Pi, G, \hat{G}_x, A, S, N_s, N_y, P_x, Q, W, V_x, Y\}$ *of the state* $(g, \mathbf{z})$ *such that*

(i) *given* G, Ω, P_x, W, Π, *and* Q, *the decision rules* A, N_s, N_y, Λ *solve the household's problem (13);*

(ii) *given* W, *the decision rules* N_y *and* Y *solve the manufacturer's problem (20);*

[18] Note that we do not explicitly include trading of ownership claims to retailers or the manufacturing firm in our definition of an equilibrium. These claims are included implicitly in that (i) retailers maximize their asset value using a set of complete state contingent claims, and (ii) non-labor income includes profits from the manufacturing sector and existing retailers minus investment in new retailers.

(iii) *given* G, Θ, Λ, *and* Q, *the decision rule* P_x *and the value function* V_x *solve the retailer's problem (17);*

(iv) *the entry rule* $\hat{G}_x$ *and* V_x *satisfy the free entry condition (18), and* V_x is 'concave'*;*

(v) *given* $\hat{G}_x$, P_x, Λ, *and* Θ, *the law of motion* G *is consistent with (19);*

(vi) *given* S *and* $N = \sum_x \left(g_x + \hat{G}_x\right)$, *the matching probabilities* Ω *and* Θ *are consistent with (10);*

(vii) *given* Y *and* $\hat{G}_x$ *the market for manufacturing goods clears, (9) is satisfied with equality.*

(viii) *non-labor income is* $\Pi = \Theta \sum_{x=1}^{\overline{x}} P_x \delta \left(P_x/\Lambda\right) \left(g_x + \hat{G}_x\right) - \alpha z N_y^{\alpha}$, *and state contingent claims,* A, *are in zero net supply.*

3.3. A Tractable Example

The law of motion for the endogenous aggregate state of the economy, that is the inventory distribution, depends on the behavior of aggregate demand, entry of retailers, and the optimal order and pricing policies of retailers. These elements represent the three channels described in section 2.2 through which general equilibrium feedbacks in the retail sector operate. The last of these three elements can make it very difficult to compute an equilibrium. In order to facilitate the numerical analysis in the next section, we eliminate the third element and focus on aggregate demand and entry. In the following we introduce two additional conditions such that order and pricing policies are independent of the aggregate state.

The first condition simplifies the optimal pricing policy much as discrete inventories effect the order policy. It is designed to ensure the optimal price depends only to a limited extent on the aggregate state. We assume that the idiosyncratic utility shock for a retailer's good has finite support. In particular, we consider a two point support $\mathbf{U}$ with a symmetric uniform distribution,

$$u = \begin{cases} u_l = 1 - \eta, & \text{with probability } 1/2, \\ u_h = 1 + \eta, & \text{with probability } 1/2, \end{cases} \tag{21}$$

and $0 < \eta < 1$. In Proposition 5 we show that the reservation utility policy of shoppers

$$\text{`accept if } \lambda u_j \geq p_x \text{ with } j = l, h \text{ and } , \bar{x} = 1, \ldots, \bar{x}, \text{'} \tag{22}$$

and the pricing policy of retailers

$$\text{`}p_x = \lambda u_j \text{ with } j = h \text{ if } x \leq x' \text{ and } j = l \text{ if } x > x', \text{'} \tag{23}$$

are consistent in an equilibrium. For this equilibrium retailers post at most two different prices, and if two prices are posted, then retailers with inventories higher than the critical inventory level x' will post the lower price. Below we refer to this as an x' *pricing scheme.*[19]

The second condition contributes to keeping both the order size and the pricing scheme aggregate state independent. In the economy without aggregate uncertainty the order size and pricing scheme in a steady state are unique. With a small amount of variation in the aggregate shocks, there is no incentive to deviate from the steady state policies. Hence, the second condition is to assume that aggregate uncertainty is not 'too large' so that the steady state order size and pricing scheme remain optimal. We check that the steady state policies are optimal for our simulations.

The simplified model we have now introduced is a general equilibrium extension of the partial equilibrium model described in section 2.2. This model allows us to focus on the remaining two channels through which general equilibrium effects operate, aggregate demand and entry. Aggregate demand for current consumption goods, as reflected by the probability that a retailer is matched with a shopper, is in part determined by the search effort of the consumer. This search effort will depend, in part, on the average price a consumer has to pay for goods purchased. For the case that retailers charge different prices, this expected price depends on the relative number of retailers who hold low (high) inventories and charge a high (low) price. Thus when there are changes in the distribution of inventories incentives to shop also vary. We call this feedback the *price channel* due to the key role played by the price setting of retailers. Since our

[19] The case where only one price is offered in a given period corresponds to $x' = 0$ in this notation.

example explicitly rules out variations in the order size, entry is the only channel in the model for changing the aggregate supply of current and future consumption goods. In our tractable example, entry, equivalently the order volume, depends in a non-trivial way on the inventory distribution. We refer to this particular feedback from the inventory distribution to order volume as the *entry channel*.[20]

4. Implications of (S, s) Inventory Policies in General Equilibrium

In this section we investigate how general equilibrium feedback effects influence the relation between (S, s) inventory policies and aggregate variables. We do so by examining various parameterizations of the tractable example described in section 3.3. Equilibrium in the example is characterized by a system of functional equations. An analytical solution to this system is not available, but it can be solved using standard numerical methods as described in Fisher and Hornstein (1996).[21]

We arrive at the following main conclusions. (1) The general equilibrium model is consistent with evidence that orders are more volatile than sales, that inventory investment is positively correlated with sales, and, for some parameterizations, that average sales are procyclical. (2) The entry channel is sufficient to make orders more volatile than sales, but the price channel is necessary to make inventory investment positively correlated with sales and to make average sales procyclical. (3) The model predicts that the width of the (S, s) band is associated with amplified responses of aggregate variables to some exogenous impulses. (4) There exists a competitive version of the general equilibrium model in which all market participants take prices as given. For this version orders remain more volatile than sales, but inventory investment is no longer positively correlated with sales. The business cycle characteristics of this model are similar to the more general model.

[20] Another interpretation of the free entry condition views variations in entry as variations in the ordering frequency. Suppose that there is an infinite supply of retailers, then at any point in time only a finite measure of these retailers has decided to order and is in the market selling goods. We can think of a retailer who has just sold the last good in inventory as deciding whether to reorder now or wait and order some time later.

[21] We can show existence and uniqueness of an equilibrium for a simplified version of the environment, see subsection 4.4 below.

4.1. Parameter Values and the Steady State

To implement our model we need to specify the following parameters

Preferences:	β, η, ν;
Retail Technology:	$\zeta, \varphi, \mu, \gamma_0, \gamma_1$;
Manufacturing Technology:	α, z;
Aggregate Uncertainty:	ρ, σ.

This list includes two parameters determining the nature of aggregate uncertainty to be explained below. We choose parameters to highlight the impact of the inventory and price policies and to achieve a degree of empirical plausibility. The examples analyzed below accordingly are based on five baseline choices for the parameters, summarized in table 1, which we now discuss.

The baseline parameter sets are chosen to ensure that in steady state, and in a region of the state space local to steady state, five particular combinations of order policy x^* and pricing scheme x', one for each parameter set, are optimal from the perspective of a representative retailer. These combinations are $(x^*, x') = (1, 0)$, $(2, 0)$, $(2, 1)$, $(4, 0)$, and $(4, 2)$. We calibrate η to fix x'.[22] The fixed cost parameter γ_0 determines x^*. We do check the optimality of the price and order strategies in our simulations.[23]

The intratemporal elasticity of labor supply with respect to the real wage is approximately equal to $1/\nu$. One feature we would like the model to have is that hours worked and shopping time are procyclical. In order for this to be the case, we work with $\nu = 0.2$, which implies a relatively large labor supply elasticity when compared to elasticities reported in empirical labor supply studies.[24] We calibrate ζ, which controls the efficiency of search effort, to fix the steady state cross sectional variance of sales for retailers.[25] Notice that the matching probability for θ determines this variance. Except for values of θ close to unity the variance of sales relative to its

[22] Fisher and Hornstein (1996) derive a formula for η such that a particular x' pricing scheme is optimal if there is no aggregate uncertainty. We use this formula to select values for η.

[23] The procedure for doing this is described in Fisher and Hornstein (1996).

[24] Similar results can be attained in a model with home production in which home and market goods are relatively easy to substitute, see for example Benhabib, Rogerson and Wright (1991).

[25] We are unable to calibrate ζ to match evidence on time devoted to shopping activities because once ν is selected the endogenous expected surplus from shopping fixes n_s.

mean is very large. For this reason we calibrate ζ so that in steady state $\theta = 0.85$. This choice also isolates the propagation effects of (S, s) inventory policies in our model, since it weakens a channel of persistence in the model that is unrelated to (S, s) inventory policies *per se*.[26]

In keeping with related work in the real business cycle literature we set labor's share of income α to 0.64. The steady state time discount factor β is set to $1.05^{-1/12}$, so that a period in the model is a month and the real interest rate in steady state is 5% at an annual rate. The constant term φ in the matching function is set to 0.85. This implies that in steady state the number of shoppers and retailers are identical, $N = S$, and the probability of a match from the perspective of a shopper is the same as for a retailer, $\omega = \theta$. We fix $z = 1$ and $\gamma_1 = 1.175$. These are normalizations which have no major consequence for the analysis. We do not have any direct evidence on the matching function elasticity, μ. Computational experiments show that a high value for μ can result in saw-tooth responses to exogenous shocks for some aggregate variables. Since we do not view these response patterns as empirically plausible, we select a relatively low value, $\mu = 0.25$, for our baseline parameterizations.[27]

With these parameters we can compute a non-stochastic steady state. At this point it is worth considering the implications of our parameter selections for the values of key variables in steady state. This helps to evaluate the extent to which our selections can be considered reasonable. Two variables are useful for this purpose: the steady state aggregate inventory-sales ratio and the steady state expected time between orders for a typical retailer. These variables are shown in table 2 for the five different baseline parameterizations.

The aggregate inventory-sales ratio (consistent with the analysis in section 2.2, we define inventories to be the total stock of inventories at the end of a period) varies from a low of 0.18 in the $(1, 0)$ model to a high of 2.80 in the $(4, 2)$ model and is positively related to both x^* and x'. Over the period 1967-1991 the average of the ratio of total retail inventories to total retail sales at the monthly frequency in U.S. data is 1.47.[28]

---

[26] Lower values of θ correspond to longer mean time between sales. An exogenous impulse that increases the number of retailers has a longer lasting impact on the economy the longer the mean time between sales because the stock of inventories takes longer to return to its steady state level.

[27] For more on saw-tooth response patterns see footnote 34.

[28] This number is based on the ratio of retail inventories to retail sales as given by the variables IVRR and RTRR, respectively, from the Citibase databank.

The expected time between ordering, not surprisingly, is also positively related to x^* and x'. It ranges from a low of 1.17 months in the $(1,0)$ model to a high of 7.04 months in the $(4,2)$ model. This variable is important for our analysis because it has a substantial impact on the degree of persistence implied by the model. Aguirregabiria (1994) finds that the mean time between ordering of all goods sold by a Spanish supermarket chain is about 1.3 months. The range of his estimates includes a low of less than a month, to a high of over 10 months. We are not aware of any other empirical studies that provide independent evidence on the mean time between ordering.

As described previously, we allow for the possibility of technology and discount rate shocks. We only include one shock in the model at a time, but use a common form for the forcing process,

$$\ln \mathbf{z}_t = (1-\rho)\ln \mathbf{z} + \rho \ln \mathbf{z}_{t-1} + \epsilon_t, \text{ with } \epsilon_t \sim N(0,\sigma^2), \,\forall\, t, \;|\rho| < 1,$$

for $(\mathbf{z}_t, \mathbf{z}) = (z_t, z)$ or (β_t, β). Since the period length in the model is a month we adopt the view that the disturbances are best modeled as persistent. In our baseline experiments, then, we fix $\rho = 0.9$. To ensure that time invariant price and inventory strategies are optimal for our parameterizations we assume the standard deviation of innovations, σ, is small enough. In all our experiments we use $\sigma = 0.01$. Consider the forcing process an approximation since the random variable it describes is not bounded. This approximation performs well since innovations are small.

4.2. Illustrating the General Equilibrium Effects

In this section we illustrate the model's general equilibrium feedback effects and clarify the differences between the general equilibrium model and partial equilibrium model studied in section 2.2. To do this we examine the dynamic responses of various endogenous variables to technology and discount rate shocks. The impulse response functions we are interested in are shown in figures 1 to 5 for the different (x^*, x') combinations and alternative specifications of aggregate uncertainty. After describing the layout of the figures and some details about measurement, we discuss the one-price examples followed by the two-price examples.

The first column of each figure displays dynamic responses for aggregate sales, aggregate orders, the number of shoppers, and the average sale price of consumption goods relative to manufactured goods. The second column shows dynamic responses for inventory investment, average sales per retailer, net entry in the retail sector, and the proportion of retailers in the inventory distribution that have inventories greater than x'. With two exceptions all the responses are per cent deviations from the "stochastic steady state" value of the variable in question computed by assuming the aggregate shock equals its mean forever and the inventory distribution has converged to its corresponding invariant distribution.[29] The exceptions are inventory investment and net entry, which are shown as absolute deviations from their stochastic steady state values. All the responses are displayed as a proportion of the standard deviation of the innovation to the aggregate shock. They are based on a one standard deviation positive innovation to the aggregate shock in question which occurs in period 3 in the figures.

We measure sales, orders and inventory investment according to the definitions given in section 2.2. Manufactured goods are the numeraire and we use nonstochastic steady state base period prices to value sales and the stock of inventories in any date. To be precise, the average retail price is used to value retail sales and we value the stock of inventories in terms of the average replacement cost. In the figures 'average retail price' is the marginal utility of the consumption index. Finally, net entry is defined in the figures as total entry in the period less the number of retailers that stock out at the end of the previous period.

One-price examples

As displayed in figure 1, a persistent, positive technology shock generates a persistent boom in retail sales, inventory investment, orders, the number of shoppers, and net-entry into the retail sector. These responses all follow from the fact that because of persistently higher wages households demand more current and future consumption for standard permanent income hypothesis reasons. This raises the demand for retail goods, and in particular the average retail price.[30] Notice that the magnitude of the retail price response is inversely related to the degree

[29] These stochastic steady state values are not quite the same as the corresponding non-stochastic steady state values. This is because they are based on the solution to the model incorporating aggregate uncertainty. In practice, the difference is very small.

[30] Higher real wages increase the household's willingness to pay for current consumption. See equation (15).

of net-entry. This illustrates the effect of the entry channel on the supply of current consumption goods. Entry tends to drive the retail price down and partially offsets the increase in the average retail price brought on by higher demand for retail goods.[31]

In contrast to the partial equilibrium case, the responses of sales, orders, and inventory investment in the general equilibrium model replicate the qualitative features of a boom. The only response inconsistent with a boom is average sales (in the one-price models, θ), which always drops.[32] The behavior of the distribution of inventories across retailers gives some indication that average sales in a two-price model will behave differently. Describing why this is so helps to illustrate the price channel. In the lower right hand corner of figure 1 the response of the proportion of retailers that hold inventories higher than $x' = 1$ for the $(2, 0)$ model, and higher than $x' = 2$ for the $(4, 0)$ model, are shown. Given the positive net-entry response, the proportion of firms toward the high end of the inventory distribution rises for both parameterizations. Below we show that this kind of response carries over to the two-price examples, so that in these cases a higher proportion of retail firms will charge the lower price in response to a positive z-shock. This increases the marginal return to shopping, encourages a stronger response of shoppers, and reverses the response of average sales. The behavior of the other main aggregates is not effected.

At this point it is interesting to study the amplification and propagation effects of (S, s) policies. Regarding amplification, notice that although the size of the response of inventory investment increases with the order size, the responses for sales and orders decline with the order size. This is due to a wealth effect. The fixed order cost is larger in higher x^* economies, so the marginal labor cost of additional consumption is greater as well. Moreover, part of the consumption benefits of additional labor effort accrue later in time. It follows that for the same size shock the consumption response should be smaller in economies with higher order sizes.

To consider propagation, it is helpful to examine, along with figure 1, the case of purely temporary z-shocks, $\rho = 0$, displayed in figure 2. Notice that the persistence of the response

[31] Despite higher retail prices the zero profit condition is maintained in these examples because interest rates (not shown) also respond positively to the technology shock.

[32] With positive net entry and a larger number of shoppers the response of the matching probability appears ambiguous. However, it is straightforward to show, using equations (14), that in the one-price models the proportion of retailers actually being matched with shoppers, θ, must fall, if an exogenous impulse leads to positive net entry. This fact underlies the fall in average sales per retail firm.

of sales is positively related to x^*, regardless of the value of ρ assumed. This is because, other things equal, a larger order size means a given positive change in orders increases the supply of consumption goods for a longer period of time. It seems to suggest that (S, s) inventory policies do in fact have the potential to contribute to the persistence of aggregate fluctuations. The humped responses of output (that is orders) and the behavior of inventory investment in figure 1 offer more indications of persistence. Yet, as figure 2 shows, much of the persistence in the model is due to the persistence of the exogenous shock.

The one-price response functions with discount rate shocks (β-shocks) are similar to the z-shock cases in terms of business cycle dynamics and persistence. The $\rho = 0.9$ functions are shown in figure 3 (for other results, here and below, see Hornstein and Fisher, 1996). These responses are qualitatively very close to figure 1, except that here the average retail price falls, the response of the number of shoppers is not always positive, and the response of sales is much more delayed.

The main difference with the z-shock examples evident from figure 3 is that the inventory policy has a large amplification effect. In particular, a wider (S, s) band is associated with amplified sales and orders responses, although the amplification effect of a wider (S, s) band is mitigated as ρ is reduced (not shown). Amplification is a direct result of the entry channel. A permanent increase in β increases the value of an order and encourages entry into the retail sector. In addition, the incentive to enter the retail sector is more interest rate sensitive the larger is the optimal order size.[33] With respect to figure 3 these considerations explain (i) the positive and amplified responses of net-entry, sales, orders and inventory investment, and (ii) the negative and amplified response of the average retail price. Implication (ii) follows since entry tends to lower retail prices due to a pure supply effect.

Two-price examples

We now consider briefly impulse responses implied by the two-price models. To conserve

[33] To see this, consider a retailer which faces a constant discount factor β and that can make a sale with probability one at a price equal to unity in every period. In this case the expected value of an order of size x^* is $V_{x^*} = (1 - \beta^{x^*})/(1 - \beta)$. The elasticity of V_{x^*} with respect to β is increasing in x^*. This implies the increased interest rate sensitivity of entry. Note that in this example temporary changes in β should have a smaller impact on the value of an order relative to permanent changes.

space we consider the case of $\rho = 0.9$ only. The relevant response functions are displayed in figures 4 and 5 for the z-shock and β-shock, respectively. We have included in these figures the responses for the $(1, 0)$ model to illustrate the overall impact of the price and entry channels as well as changes in the inventory policy on the responses of the endogenous variables.

We see from the figures that the responses of the main aggregates to both kinds of shock are qualitatively very similar to the one-price responses, with the exception of average sales and shopping time. Through the price channel, higher net-entry makes the price distribution more attractive for the household, which increases shopping time and thereby average sales. This in turn means that total sales respond faster, especially for time discount shocks. For the z-shock, notice that the optimal price and ordering policies of individual retailers have only a minor impact on the behavior of aggregate orders. For β-shocks, notice the amplification effect of a wider (S, s) band indicated in figures 3 is evident here also.

Although the qualitative features of the one-price and two-price examples are very similar, this does not mean that the price channel is unimportant and that entry is the main channel through which aggregate variables are influenced by the inventory distribution. As we now show, the faster response of sales to shocks in the two-price example is necessary to account for the positive correlation of inventory investment with sales and the procyclicality of average sales.[34]

[34] The price channel also has a significant impact on the behavior of aggregate variables for other parameterizations of the model. For example, recall our discussion of parameter selection where mention was made of how saw-tooth impulse responses can emerge from the model when μ is 'large'. Three factors explain this. First, the elasticity of θ with respect to N and S is increasing in μ. Second, the price channel gives rise to a shopping time decision rule that depends positively on the number of retailers with $x > x'$ and negatively on the number of retailers with $x \leq x'$. Third, the price channel amplifies the response of the number of shoppers to an exogenous disturbance. To see how the saw-tooth patterns emerge, consider the case of a positive innovation to z in an otherwise baseline $(2, 1)$ version of the model with $\rho = 0.9$ and, say, $\mu = 0.5$. The innovation tends to increase the number of shoppers more than the number of retailers, which, because of the relatively high value of μ, leads to a large positive response of the matching probability θ. The high value of θ shifts the price distribution down enough so that in the following period shopping time drops to induce a significantly lower θ. With still positive net entry, the price distribution shifts back in favor of encouraging the substitution of time toward shopping and a cyclical pattern emerges. This pattern is reminiscent of the oscillations which emerge from traditional inventory-accelerator models (see Metzler 1941). However the stable high frequency dynamics implied by some parameterizations are derived from an equilibrium model and not from mechanical equations, as is the case in the traditional models.

4.3. Quantifying the General Equilibrium Effects

To quantify the general equilibrium effects we analyze summary statistics based on simulating the model under the baseline parameterizations. Table 3 reports statistics related to sales (C), orders (Y), inventory investment (ΔI), and average sales (C/N) based on monthly data.[35] The notation σ_x denotes the standard deviation of the variable $x = C, Y, \Delta I$ relative to the standard deviation of the exogenous shock. The notation $\rho(C, x)$ denotes the correlation coefficient for variables C and $x = Y, \Delta I, C/N$. Finally, except for ΔI, we log all variables before computing statistics.

The relationship between statistics and model parameterizations for each of the two sources of aggregate uncertainty are in many respects quite similar. First, in contrast to the partial equilibrium model, orders are more volatile than sales and inventory investment is positively correlated with sales in the two-price examples. The excess volatility of orders over sales is positively related to both x^* and x'. This phenomenon seems to be due to a combination of sales volatility being dampened while fluctuations in inventory investment are amplified, as we change parameters so that x^* and x' are increased. The price channel is important for generating a positive unconditional correlation of inventory investment with sales. In the absence of the price channel, inventories and sales are essentially uncorrelated for technology shocks and negatively correlated for discount rate shocks. For the two-price model with an active price channel, inventories and sales are positively correlated because as discussed above sales respond faster due to the stronger decline in prices.

Second, orders are uniformly positively correlated with sales and this correlation is generally inversely related in equilibrium to x^* and x'. The correlation of average sales with sales is never positive. However, this correlation becomes less negative when the price channel is operational and the (S, s) band is widened and it is positive if, in addition, ρ is small enough (for example, $\rho = 0$). These observations suggest both feedback effects contribute to procyclical average sales.[36]

[35] In results not reported here we find the qualitative conclusions drawn from these monthly statistics translate to quarterly data.

[36] Some empirical evidence supporting the view that average sales are procyclical is that the per dealer average unit sales of new US made cars is procyclical for the sample 1962-1988. The dealership data is from Ward's

Third, a wider (S, s) band has a mixed effect on generating persistence. It tends to increase the persistence of sales and inventory investment fluctuations, but it has little impact on the persistence of orders fluctuations.

One feature of the discount rate shock cases does stand out, namely the strong association of larger values for x^* and x' with increased volatility in sales, orders and inventory investment. In the $(x^*, x') = (1, 0)$ model $\sigma_C = 0.19$, $\sigma_Y = 0.19$, and $\sigma_{\Delta I} = 0.01$. In the $(4, 0)$ version of the model the volatility of sales and orders increases by a factor of seven, and inventory investment increases by a factor of thirty. In the $(4, 2)$ version of the model the volatility statistics increase again so that $\sigma_C = 1.66$, $\sigma_Y = 1.87$, and $\sigma_{\Delta I} = 0.35$.

4.4. Isolating Entry: The Competitive Equilibrium as a Special Case

We now examine a special case in which we abstract from trading frictions and focus on the entry channel. This example demonstrates that inventory holdings due to fixed ordering costs can be incorporated into a competitive equilibrium model. This is useful because the competitive equilibrium is Pareto-optimal and we can use the social planning problem to show existence of an equilibrium and characterize its properties. A quantitative analysis shows that for the competitive equilibrium example, orders remain more volatile than sales, but that inventory investment is no longer positively correlated with sales.

In the absence of trading frictions retailers behave competitively, that is they take demand and prices as given. Consider the limiting case of our model in which the representative household is a very efficient shopper, that is $1/\zeta = 0$. In an equilibrium for this economy, the household sends out shoppers to visit all retailers, that is $\theta = 1$, and buys only from retailers with the lowest price. A typical retailer takes the price of the consumption good as given: she has no incentive to raise her price since the probability of a sale would be zero, and she has no incentive to lower the price since she already sells the good for sure. In terms of our discussion of x' pricing schemes, the competitive equilibrium is a one-price equilibrium ($x' = 0$) in which the matching probability for retailers is one.

Automotive Annual (various issues) and total unit sales of domestic new cars is Citibase variable RCARD. We calculate the contemporaneous correlation between Hodrick-Prescott filtered real GDP (Citibase variable GDPQF) and average sales per dealer to be 0.59.

In Proposition 6 we show the competitive equilibrium exists and is unique. The version of this model without aggregate uncertainty illustrates an important general equilibrium characteristic. In the partial equilibrium of section 2.2 the inventory distribution does not converge to the invariant distribution if there is no risk, $\theta = \delta = 1$. This absence of convergence is associated with fluctuations in orders, production and wages. In the comparable general equilibrium setting, convergence to the invariant distribution takes place because agents want to smooth non-leisure activities over time, and wage changes induce convergence.

We follow the same strategy as before to analyze the stochastic equilibrium, but do not display the results. With the exception of γ_0, ζ and η, the parameterizations of the competitive economy are exactly the same as in table 1.[37] The competitive economy has predictions for all the same variables as the costly search model, except that average sales are constant and the number of shoppers is indeterminate. With z-shocks, the impulse response functions for the comparable variables are essentially the same as in the similarly parameterized one-price costly search model. With one qualification the same can be said for the responses to a β-shock. The responses to this shock in the competitive model are somewhat less persistent than in the costly search model.[38] Simulations of the competitive economy also imply qualitatively similar values for the comparable statistics reported for the one-price costly search model in table 3. We find that with $x^* > 1$ orders are more volatile than sales, but that for all x^* values considered and for all $\rho \geq 0$ the dynamics of inventory investment and sales are such that these variables are essentially uncorrelated or negatively correlated unconditionally.

5. Conclusion

We have constructed a dynamic general equilibrium model to analyze the aggregate implications of (S, s) inventory policies and have identified two general equilibrium effects: the endogenous variation in the order volume and the interaction of the inventory distribution with different

[37] The values of the fixed cost parameter for the $x^* = 1$, 2 and 4 versions of the competitive model are $\gamma_0 = 0$, 0.008 and 0.035, respectively. The competitive case is defined by $1/\zeta = 0$. Finally, we choose the normalization $\eta = 0$.

[38] This is mostly due to the fact that the expected time between ordering for a typical retailer is lower in the competitive cases relative to their comparable costly search cases.

pricing policies across retailers. We have shown that both general equilibrium effects are necessary to account for the observations that orders are more volatile than sales and that inventory investment is positively correlated with sales. Moreover, the model is consistent with other salient features of the business cycle.

A competitive equilibrium version of the model, in which the entry channel is the only general equilibrium feedback effect, mimics the costly search economy described in the main part of the paper, except for the inventory investment-sales correlation and the procyclicality of average sales. The competitive equilibrium has the advantage of being much more amenable to numerical analysis since it eschews the treatment of idiosyncratic risk for retailers and implications of this risk for optimal price setting. We interpret this result as providing some justification for abstracting from microeconomic detail in future work on the relation between lumpy inventory adjustment and the business cycle.[39] Another advantage of the competitive model is that it is relatively easy to extend. For example it can be used to study at least two important extensions. First, explicit variation in the intervention points of the inventory policy over the business cycle, and second incorporating (S, s) inventory policies into the standard neoclassical growth model with capital accumulation.

Finally, the inventory problem we have studied is just one example of the general problem of optimal control with fixed adjustment costs. Other examples that have been studied in the literature include the adjustment of houses, consumer and producer durables, prices, employment, and cash balances. The basic strategy for analyzing the general equilibrium effects of (S, s) policies we pursue in this paper is to combine discreteness in the choices of retailers with a free-entry condition to ensure that aggregate adjustment takes place in the mass of retailers and not in the decisions of individual retailers for small aggregate shocks. This technique can be applied in other settings with (S, s)-type polices. For example, in the price-setting model of Benabou (1988, 1992) adding discreteness to the preferences of consumers would ensure that prices take on discrete values. With small aggregate shocks we expect that the free-entry condition in that model would imply that aggregate adjustment would take place in the mass of

[39] The work of Caplin and Leahy (1991) suggests the similarity between the competitive and the search model may depend on the one-sided nature of the (S, s) policies we study and that with two-sided policies idiosyncratic risk may play a more important role.

firms and not in the price-setting policies of individual firms.

Appendix

In this appendix we state and prove the propositions referred to in the main text.

Assumption 1. *Let $R > 1$, and demand is elastic, $\Psi(p) = -\delta'(p)p/\delta(p) > 1$, nonincreasing, $\delta' \leq 0$, bounded, $0 \leq \delta(p) \leq 1$, and it attains the upper bound for some positive price p_l, $\delta(p_l) = 1$. Either (a) the elasticity of demand is constant, or (b) the demand function attains its lower bound for some finite price p_h.*

Remark 1. *The representative household's optimal reservation utility strategy described in section 3 implies that (a) will be satisfied if good specific utility shocks are distributed according to a Pareto distribution, and (b) will be satisfied if the shocks follow a uniform distribution.*

Proposition 1. *Assume that δ satisfies Assumption 1. Then the change in the capital value of a retailer, that is $\Delta V_x = V_x - V_{x-1}$, and the optimal price p_x, are both decreasing in x.*

Proof. The first order condition for an optimal price implies that marginal cost is less than or equal to marginal revenue

$$\Delta V_x/R \leq f_1(p_x) = \frac{\Psi(p_x) - 1}{\Psi(p_x)} p_x \text{ and} \tag{24}$$

Given that $\Psi(p) > 1$, this is also sufficient for a maximum. From Assumption 1 it follows that f_1 is increasing in p_x. Thus the optimal price is increasing with the discounted capital loss associated with the sale of one unit of the good. From the definition of V_x it follows that

$$\Delta V_x/R = f_2(p_x, V_{x-1}) = \frac{\theta\delta\left(p_x\right)p_x - \left(1 - R^{-1}\right)V_{x-1}}{\theta\delta\left(p_x\right) + R - 1}. \tag{25}$$

We want to show that for any V_{x-1} there exist unique p_x and V_x which satisfy equations (24) and (25).

Suppose that $f_2(p, V) \leq f_1(p)$ for all p, then $p_x = p_l$ and V_x is determined from (25). Alternatively suppose that there exists a p such that $f_2(p, V) > f_1(p)$. The function f_2 is bounded above by

$$f_2(p, V) \leq \frac{\delta(p)p}{R - 1} - \frac{V}{R}.$$

If condition (a) from Assumption 1 applies then for p sufficiently large $f_2(p, V) \leq f_1(p)$. Both f_1 and f_2 are continuous, therefore there exists a p_x such that $f_2(p_x, V) = f_1(p_x)$. Some algebra shows that $\partial f_2(p_x, V)/\partial p = 0$, and since f_1 is strictly increasing in p, the price p_x is unique. Now suppose that condition (b) of Assumption 1 applies. Thus for $p \geq p_l$, $f_2(p, V_{x-1}) \leq 0$, and $f_1(p) \geq 0$. By the same argument as above there exists a unique $p_x > p_l$ which solves (24) and (25). Since f_2 is decreasing in V, the increment ΔV_x is declining with x.□

Assumption 2. *All firms in the retail industry follow $(x^*, 0)$ inventory policies, $x^* \geq 1$, price strategies such that the probability of a sale is $\{\delta_x : x = 1, \ldots, x^*\}$, and $\theta_t \in [\underline{\theta}, \overline{\theta}], 0 < \underline{\theta} < \overline{\theta} < 1$, for all $t \geq 0$.*

Proposition 2. *Given Assumption 2, from any initial distribution $\bar{g}_0$, $\bar{g}_t$ converges to the invariant distribution (3),*

$$\lim_{t\to\infty} \bar{g}_{x,t} = \frac{\delta_x^{-1}}{\sum_{\tilde{x}=1}^{x^*} \delta_{\tilde{x}}^{-1}} N.$$

Proof. Consider a given realization of $\{\theta_t\}_{t=0}^{\infty}$. Denote the matrix of transition probabilities at time t,

$$\Pi_t = \begin{bmatrix} (1-\theta_t\delta_1) & \theta_t\delta_2 & 0 & 0 & \cdots & 0 & 0 \\ 0 & (1-\theta_t\delta_2) & \theta_t\delta_3 & 0 & \cdots & 0 & 0 \\ 0 & 0 & (1-\theta_t\delta_3) & \theta_t\delta_4 & \cdots & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ \theta_t\delta_1 & 0 & 0 & 0 & \cdots & 0 & (1-\theta_t\delta_{x^*}) \end{bmatrix},$$

and $\widetilde{\Pi}_\tau = \Pi_{\tau+(x^*-1)} \cdots \Pi_{\tau+1}\Pi_\tau$, $\tau \geq 0$. It is easy to verify that for any $\tau \geq 0$, $\widetilde{\Pi}_\tau$ is a positive stochastic matrix. The restriction of θ_t to a compact strict subset of the unit interval guarantees that the sum over the minimum elements in each column of $\widetilde{\Pi}_\tau$ is always greater than some $\varepsilon > 0$. Also, for any $t \geq 0$, the invariant distribution (3) is a unique fixed point of the mapping $\bar{g}_{t+1} = \Pi_t \bar{g}_t$. The result can be established now using the proof of Theorem 11.4 in Stokey and Lucas (1989). □

Proposition 3. *Normalize the number of firms in the retail industry to unity. Given Assumption 2, if the inventory distribution has converged to its invariant distribution, then*

(i) The variance of orders equals the variance of sales.

(ii) Inventory investment and sales are negatively correlated.

(iii) Orders and sales are not more volatile than demand and the autocorrelations of orders and sales equal the autocorrelation of demand.

(iv) Orders and sales are positively correlated only if demand is positively autocorrelated.

Proof. At time t, aggregate sales, C_t, equals the proportion of retailers that make a sale, that is $C_t = \sum_x \theta_t \delta_x \bar{g}_x = \theta_t \left(x^* / \sum_x \delta_x^{-1}\right)$. Aggregate orders, Y_t equal the proportion of retailers with $x = 1$ at time $t-1$ who make a sale times the size of the order. Thus $Y_t = \theta_{t-1}\delta_1\bar{g}_1 = \theta_{t-1}\left(x^* / \sum_x \delta_x^{-1}\right)$. Aggregate inventories at the beginning of time $t+1$, I_{t+1}, satisfies the identity $I_{t+1} = I_t + Y_t - C_t$. Thus aggregate inventory investment during period t, $\Delta I_t = I_{t+1} - I_t$, equals $(\theta_{t-1} - \theta_t)\left(x^* / \sum_x \delta_x^{-1}\right)$. Thus (i), (iii) and (iv) follow trivially. From the inventory identity

$$\sigma_Y^2 = \sigma_C^2 + \sigma_{\Delta I}^2 + 2\,\mathrm{cov}(C, \Delta I)\ ,$$

where σ_z^2, denotes the unconditional variance of $z = Y, C, \Delta I$. Since $\sigma_Y^2 = \sigma_C^2$, it follows that $\mathrm{cov}(C, \Delta I) = -\sigma_{\Delta I}^2/2$, which establishes (ii).□

Proposition 4. *There exists a unique solution to the optimization problem*

$$\begin{array}{ll} \max_{n_s,q} & U(c, n_s) \\ s.t. & c = \zeta n_s \omega q, \\ & \zeta n_s \omega e(q) = m, \end{array} \tag{26}$$

Proof. The household has the option of sending out shoppers with different reservation utility policies. This means that the household does not choose any particular quality q, but a signed measure on the set of feasible qualities $[0, E[u]]$. For simplicity assume that the household can choose among a finite set of policies $\{(q_j, e_j) : j = 1, \ldots, J\}$, in which e_j is the minimal expenditure when the expected match quality is q_j. The household can send out a measure ζn_{sj} of shoppers all of whom follow the j-th reservation utility policy. The utility index for market goods from this policy is $\zeta\omega \sum_j q_j n_{sj}$ and the cost from this policy is $\zeta\omega \sum_j e_j n_{sj}$. Clearly the household is solving a concave optimization problem. Suppose that the expenditure function is strictly concave in quality q, that is marginal cost of quality λ is increasing in quality, as is the case here. Applying the same arguments as in Hornstein and Prescott (1993), one can show that for the limiting case of bounded signed measures on $[0, E[u]]$ the optimal policy will put positive mass on one policy only and this is the solution to (26). This approach is analogous to the use of lotteries in the context of non-convex problems, see for example Rogerson (1987) or Hansen (1985).□

Proposition 5. *Let F be as defined in (21). An equilibrium satisfies (i) if households use the reservation price policy (22) then the offer price strategy (23) is optimal for a retailer; and (ii) if retailers use the offer price strategy (23) the reservation price policy (22) is optimal for households.*

Proof. (i) The optimal price policy of the retailer is based on the following argument. In equilibrium no price lower than λu_l will ever be offered by a retailer because the probability of acceptance by a shopper is a constant at these prices, and expected revenues are increasing in the offer price for all prices $p < \lambda u_l$. For a similar reason, no price in the region $\lambda u_l < p < \lambda u_h$ will ever be offered. Finally, no price higher than λu_h will ever be offered, unless the retailer expects revenues in the following period to be exceptionally high. The preceding argument suggests that by judicious choice of u_l and u_h at most two prices, λu_l and λu_h, will ever be offered by retailers in any given period. Notice, however, that if the difference between u_l and u_h is sufficiently small, then only one price, λu_l, will be offered in a given period. This is because for u_h sufficiently close to u_l, the gain in revenues from a sale is more than offset by the lower probability that a sale will actually take place. Households only go shopping if the activity has a positive net return. This means there will never be a one-price equilibrium with the price equal to λu_h because in this case shopping would have a zero net return.

(ii) We show that if the household consumes the good, then the household is indifferent among all mixed strategies of the form 'accept if $u_j > p_i/\lambda$ and accept with probability $a_i \in [0, 1]$ if $u_i = p_i/\lambda$ for $i, j = l, h$'. This implies that the reservation price policy (22) is optimal.

Let the retailer's price strategy be characterized by $\left(\bar{\lambda}, x'\right)$. Note that the household takes $\bar{\lambda} = \lambda$ as given. For the specified mixed strategies the household's minimal expected expenditure per expected match quality is

$$e = \bar{\lambda}\left(q - \underline{g}\eta\right). \tag{27}$$

with $\underline{g} = \sum_{x>x'} (\bar{g}_x/N)$. Analogous to the proof of Proposition 4 we can imagine the household sending out shoppers with different mixed strategies. For simplicity again assume that the household has a choice among a finite set of mixed strategies resulting in expenditure-quality pairs $(e_j, q_j)_{j=1}^{J}$. Let ζn_{sj} be the measure of shoppers with the j-th mixed strategy, then the optimal choice of n_{sj} has to satisfy the following first order condition

$$\zeta \omega q_j U_c - \left(\zeta \frac{\omega}{w} e_j + 1 \right) U_n \geq 0 \text{ for } n_{sj} \geq 0.$$

Using the definition of the expenditure function (27) and the equilibrium condition that $\bar{\lambda} = \lambda = w U_c / U_n$, this condition simplifies to

$$\left(\underline{g} \zeta \omega \lambda \eta / w - 1 \right) U_n \geq 0 \text{ for } n_{sj} \geq 0.$$

If the household consumes goods this condition has to be satisfied for some mixed strategy. Since this condition is independent of the choice of mixed strategy, it is satisfied for all mixed strategies. Therefore the household is indifferent between choosing any one of the mixed strategies.□

Remark 2. *We restrict attention to pure strategies where a shopper accepts an offer even if he is exactly indifferent between accepting and rejecting. Otherwise, if the price must be strictly less than the reservation utility to be accepted, there is no equilibrium because retailers would have an incentive to increase offer prices.*

Proposition 6. *When the representative household is very efficient, $1/\zeta = 0$, there exists a competitive equilibrium and it is unique.*

Proof. We show that our environment corresponds to that of a standard optimal growth model when the representative household is a very efficient shopper. This environment satisfies all the usual conditions of the welfare theorems and a solution to the social planning problem can be decentralized as a competitive equilibrium, as described in section 3. In the text we only consider equilibria for which retailers with positive beginning-of-period inventory holdings do not place orders. We do not impose this restriction in the following.

In any period the aggregate state of the economy is given by the measure of retailers over inventory holdings $g_t = [g_{x,t} : x = 1, \ldots, \bar{x} - 1] \in \Re_+^{\bar{x}-1}$, and the exogenous state $\mathbf{z}_t$. As noted before $\bar{x}$ is the capacity constraint, that is retailers cannot hold more than $\bar{x}$ units of goods in inventory within a period. The measure of inventory holdings changes because of order decisions and sales. Let $\hat{g}_{xx',t}$ be the measure of retailers who start out with x units of inventory at the beginning of the period and place an order of x' units, where $x = 0, \ldots, \bar{x} - 1$, and $x' = 1, \ldots, \bar{x}$. Let $\hat{g}_t = [\hat{g}_{xx',t} : x = 0, \ldots, \bar{x} - 1,\ x' = 1, \ldots, \bar{x}, \text{ and } x + x' \leq \bar{x}] \in \Re_+^{\bar{x}(\bar{x}+1)/2}$ denote the vector of order measures which satisfy the capacity constraint. Entry occurs if $\hat{g}_{0x',t} > 0$. The number of retailers who order before they sell all of their inventory cannot be larger than the number of retailers with the given initial inventory holding, that is the following consistency constraint applies for the measure of order decisions

$$\sum_{x'} \hat{g}_{xx',t} \leq g_{x,t} \text{ for } x \geq 1. \tag{28}$$

Given the measure of order decisions the resource constraint for the manufactured good is

$$\sum_{xx'} \left(\gamma_0 + \gamma_1 x' \right) \hat{g}_{xx',t} \leq z_t n_{y,t}^{\alpha}, \tag{29}$$

and the distribution of inventory holdings after orders have been filled is

$$\bar{g}_{x,t} = g_{x,t} - \sum_{x'} \hat{g}_{xx',t} + \sum_{x'+x''=x} \hat{g}_{x'x'',t}. \tag{30}$$

Since each retailer is matched with a shopper, consumption of the representative agent is

$$c_t = \sum_x \bar{g}_{x,t} \tag{31}$$

and next period's initial inventory holdings are

$$g_{x,t+1} = \bar{g}_{x+1,t}. \tag{32}$$

Because the household is very efficient at shopping there is no leisure loss associated with contacting retailers and $n_t = n_{y,t}$ and preferences U are as defined in (4). The constraints are summarized in the set $\Gamma(g_t, \mathbf{z}_t) = \{(c_t, n_t, g_{t+1})$: such that there exist $\hat{g}_t, \bar{g}_t \geq 0$ and c_t, n_t, $\hat{g}_t$, $\bar{g}_t$, g_{t+1} satisfy (28), (29), (30), (31), and (32)$\}$. The set Γ is closed and convex, and since n_t and z_t are bounded, the set is also compact. Applying standard theorems from Stokey and Lucas (1989), we can show that the dynamic programming problem for this dynamic optimization problem

$$\mathcal{U}(g, \mathbf{z}) = \max_{g'} U(c, n) + E[\beta \mathcal{U}(g', z')] \text{ s.t. } (c, n, g') \in \Gamma(g, \mathbf{z})$$

is well defined, and since U is strictly concave, the optimal policy is unique.□

Remark 3. *An extension of the competitive model which includes capital is immediate.*

References

Aguirregabiria, Victor. 1994. A Dynamic Structural Model of Price and Inventory Decisions with Fixed Adjustment Costs. University of Western Ontario mimeo.

Benabou, Roland. 1988. Search, Price Setting and Inflation. *Review of Economic Studies* 55, 353-373.

Benabou, Roland. 1992. Inflation and Efficiency in Search Markets. *Review of Economic Studies* 59, 299-329.

Bental, Benjamin and Eden, Benjamin. 1993. Inventories in a Competitive Environment. *Journal of Political Economy* 101, 863-86.

Benhabib, Jess, Rogerson, Richard and Wright, Randall. 1991. Homework in Macroeconomics: Household Production and Aggregate Fluctuations. *Journal of Political Economy* 99, 1166-87.

Blinder, Alan S. 1981. Retail Inventory Behavior and Business Fluctuations. In *Brookings Papers on Economic Activity*, ed. William C. Brainard and George L. Perry, 2: 443-505.

Blinder, Alan S.,and Maccini, Louis J. 1991. Taking Stock: A Critical Assessment of Recent Research on Inventories. *Journal of Economic Perspectives* 5, 73-96.

Caballero, Ricardo J. 1992. A Fallacy of Composition. *American Economic Review* 82, 1279-92.

Caballero, Ricardo J. and Engel, Eduardo M.R.A. 1991. Dynamic (S,s) Economies. *Econometrica* 59, 1659-86.

Caplin, Andrew S. 1985. The Variability of Aggregate Demand with (S,s) Inventory Policies. *Econometrica* 53, 1395-1410.

Caplin, Andrew S. and Spulber, Daniel F. 1987. Menu Costs and the Neutrality of Money. *Quarterly Journal of Economics* 102, 703-725.

Caplin, Andrew and Leahy, John. 1991. State-dependent pricing and the dynamics of output and prices. *Quarterly Journal of Economics* 106, 683-708.

Christiano, Lawrence J. 1988. Why Do Inventories Fluctuate So Much? *Journal of Monetary Economics* 21, 247-80.

Christiano, Lawrence J. and Fisher, Jonas D.M. 1994. Algorithms for Solving Dynamic Models with Occasionally Binding Constraints. Federal Reserve Bank of Minneapolis. Research Department. Staff Report 171.

Christiano, Lawrence J. and Fitzgerald, Terry. 1991. The Magnitude of the Speculative Motive for Holding Inventories in a Real Business Cycle Model. Federal Reserve Bank of Minneapolis. Institute of Empirical Macroeconomics. Discussion Paper 10.

Fair, Ray C. The production-smoothing model is alive and well. *Journal of Monetary Economics* 24, 353-64.

Fisher, Jonas D.M. and Hornstein, Andreas. 1996. (S, s) Inventories in General Equilibrium. Federal Reserve Bank of Chicago. Working Paper 96-24.

Greenwood, Jeremy, Hercowitz, Zvi, and Huffman, Gregory W. 1988. Investment, Capacity Utilization, and the Real Business Cycle. *American Economic Review* 78: 402-17.

Hansen, Gary. 1985. Indivisible Labor and the Business Cycle. *Journal of Monetary Economics* 16, 309-27.

Hornstein, Andreas, and Prescott, Edward C. 1993. The Firm and the Plant in General Equilibrium Theory. In *General Equilibrium, Growth,and Trade Vol. 2*, ed. Robert Becker et al., 393-410. San Diego: Academic Press.

Judd, Kenneth L. 1992. Projection Methods for Solving Aggregate Growth Models. *Journal of Economic Theory* 58: 410-52.

Kahn, James A. 1987. Inventories and the volatility of production. *American Economic Review* 77, 667-79.

Krane, Spencer D. and Braun, Steven N. 1991. Production smoothing evidence from physical-product data. *Journal of Political Economy* 99, 558-81.

Kydland, Finn, and Prescott, Edward C. 1982. Time to Build and Aggregate Fluctuations. *Econometrica* 50, 1345-70.

Lovell, Michael C. 1993. Simulating the Inventory Cycle. *Journal of Economic Behavior and Organization* 21, 147-79.

Metzler, Lloyd. 1941. The Nature and Stability of Inventory Cycles. *Review of Economics and Statistics* 23, 113-129.

Rogerson, Richard. 1987. An Equilibrium Model of Sectoral Reallocation. *Journal of Political Economy* 95, 824-34.

Scarf, Herbert. 1960. The Optimality of (S,s) Policies in the Dynamic Inventory Problem. In Kenneth J. Arrow, Samuel Karlin, and Patrick Suppes. Mathematical Methods in the Social Sciences, 1959. Proceedings of the First Stanford Symposium. pp. 196-202. Stanford University Press: Stanford CA.

Sobel, Joel, and Takahashi, Ichiro. 1983. A multistage model of bargaining. *Review of Economic Studies* 50, 411-26.

Stokey, Nancy and Lucas, Robert E. 1989. *Recursive Methods in Economic Dynamics*. Cambridge MA: Harvard University Press.

Table 1. Parameter Values for the Baseline Experiments.

Panel A: Parameters Fixed Across Experiments

$\upsilon = 0.2$	$\alpha = 0.64$	$\beta = 1.05^{-1/12}$
$\varphi = 0.85$	$z = 1$	$\gamma_1 = 1.175$
$\mu = 0.25$	$\rho = 0.9$	$\sigma = 0.01$

Panel B: Parameters Which Vary Across Experiments

(x^*, x')	γ_0	η	ζ
(1,0)	0	0.135	701.3
(2,0)	0.010	0.135	699.6
(2,1)	0.015	0.303	933.9
(4,0)	0.046	0.135	697.1
(4,2)	0.055	0.538	524.0

Note: The table entries for η are in per cent and all other entries are in levels.

Table 2. Steady State Values of Selected Variables for the Baseline Experiments.

	$(x^*, x') =$				
Statistic	$(1,0)$	$(2,0)$	$(2,1)$	$(4,0)$	$(4,2)$
Inventories / Sales	0.18	0.76	1.35	1.93	2.79
Expected Time Between Orders	1.17	2.32	3.47	4.71	7.04

Note: Expected time between orders is the mean time between ordering for the typical firm in months.

Table 3. Statistics Implied by the Model.

	Technology Shocks, $(x^*, x') =$					Discount Rate Shocks, $(x^*, x') =$				
Statistic	$(1,0)$	$(2,0)$	$(2,1)$	$(4,0)$	$(4,2)$	$(1,0)$	$(2,0)$	$(2,1)$	$(4,0)$	$(4,2)$
σ_C	2.11	2.00	1.93	1.83	1.74	0.19	0.72	1.07	1.41	1.66
σ_Y/σ_C	1.02	1.03	1.09	1.05	1.14	1.03	1.03	1.08	1.05	1.13
$\sigma_{\Delta I}$	0.10	0.23	0.24	0.34	0.34	0.01	0.09	0.15	0.28	0.35
$\rho(C,Y)$	0.99	0.97	0.97	0.93	0.94	0.99	0.97	0.97	0.92	0.93
$\rho(C,\Delta I)$	0.18	0.02	0.24	-0.04	0.15	0.15	-0.01	0.20	-0.08	0.11
$\rho(C,C/N)$	-0.99	-0.99	-0.13	-0.99	-0.06	-0.99	-0.97	-0.46	-0.92	-0.34
$\rho(C_t,C_{t-1})$	0.93	0.97	0.97	0.99	0.99	0.94	0.98	0.97	0.99	0.99
$\rho(Y_t,Y_{t-1})$	0.90	0.94	0.92	0.95	0.94	0.90	0.95	0.92	0.95	0.95
$\rho(\Delta I_t,\Delta I_{t-1})$	0.14	0.43	0.68	0.69	0.73	0.14	0.44	0.67	0.69	0.73

Note: Volatility statistics are presented as a proportion of the standard deviation of the exogenous shock.

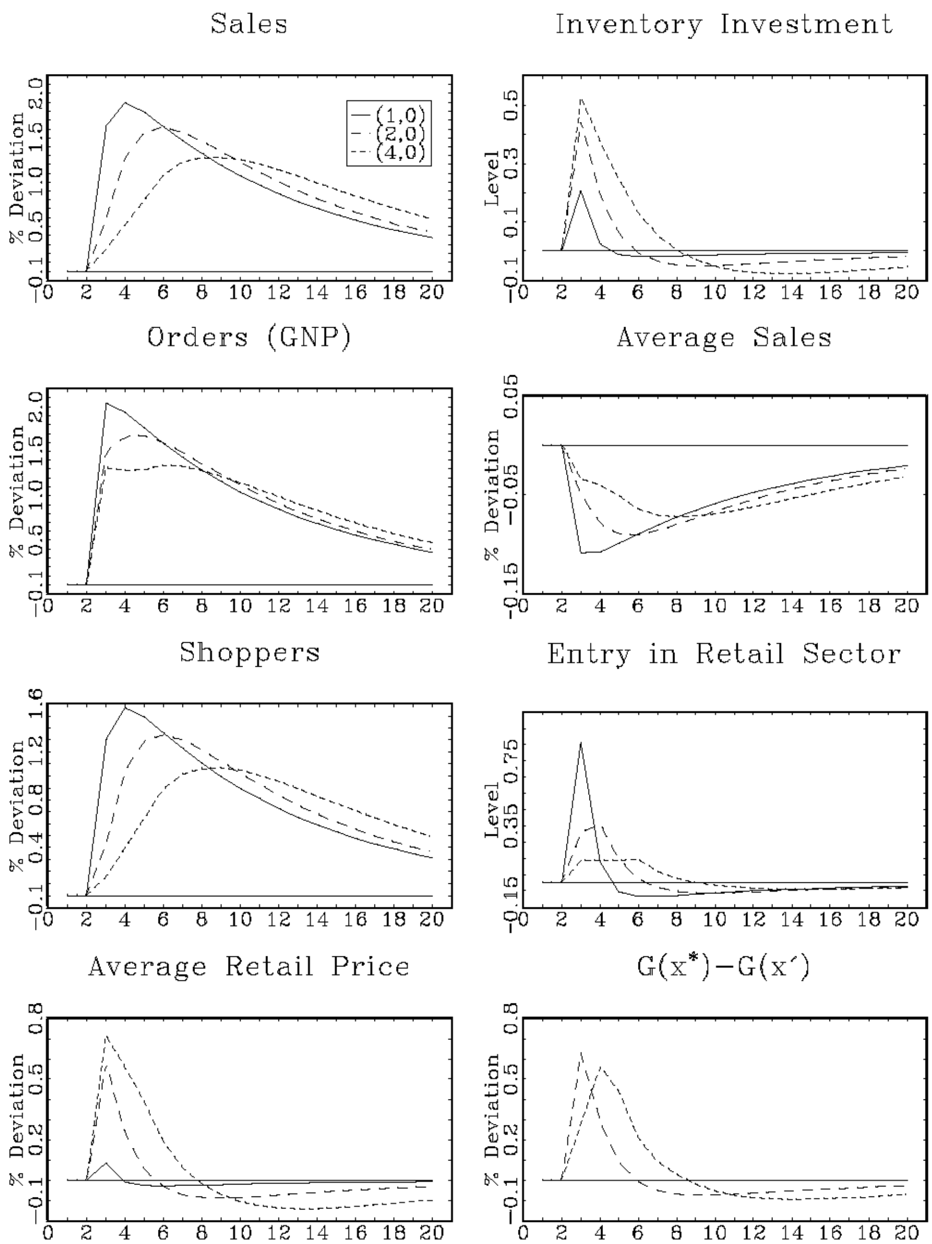

Figure 1: One-Price Impulse Responses to Technology Shocks, $\rho = 0.9$

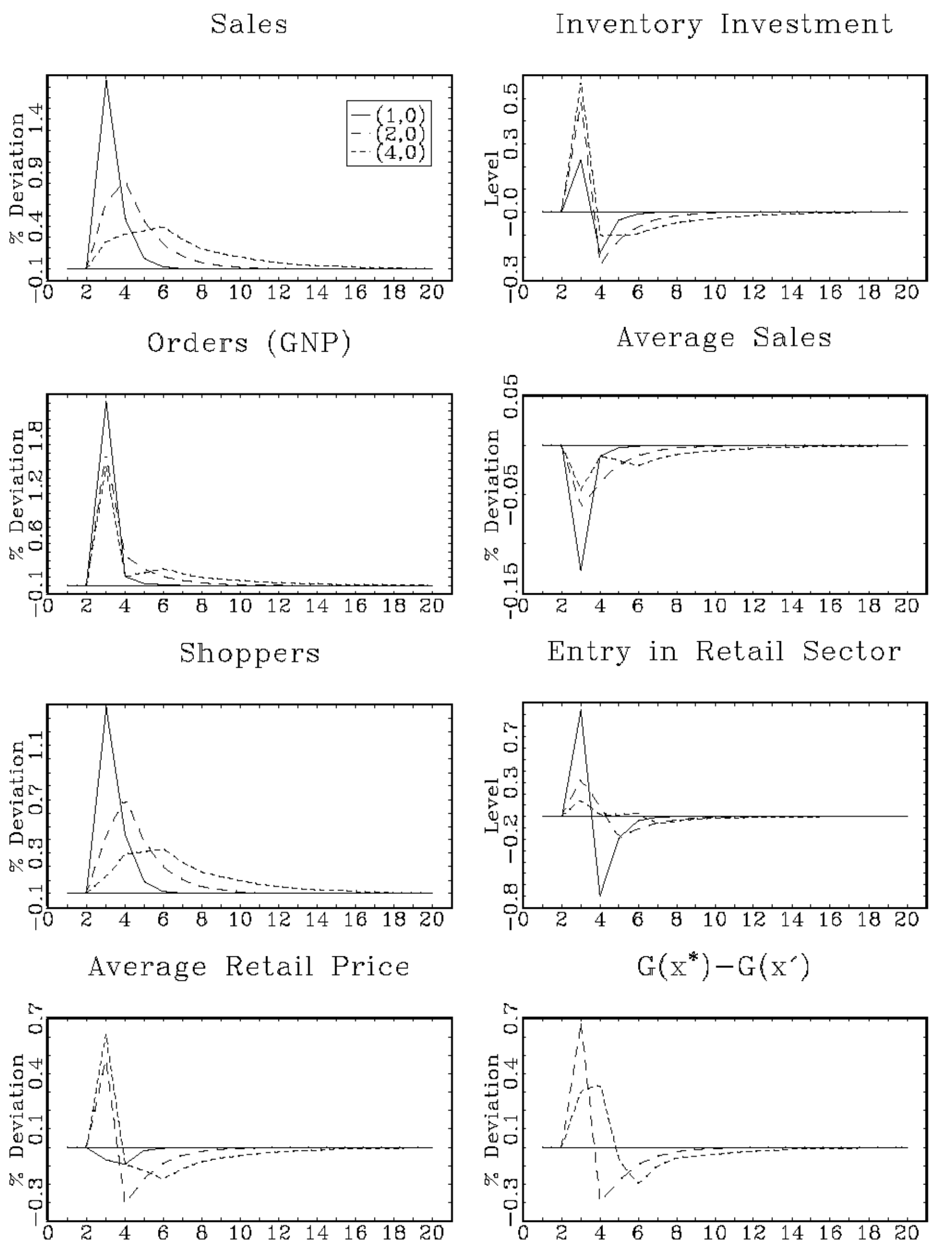

Figure 2: One-Price Impulse Responses to Technology Shocks, $\rho = 0$

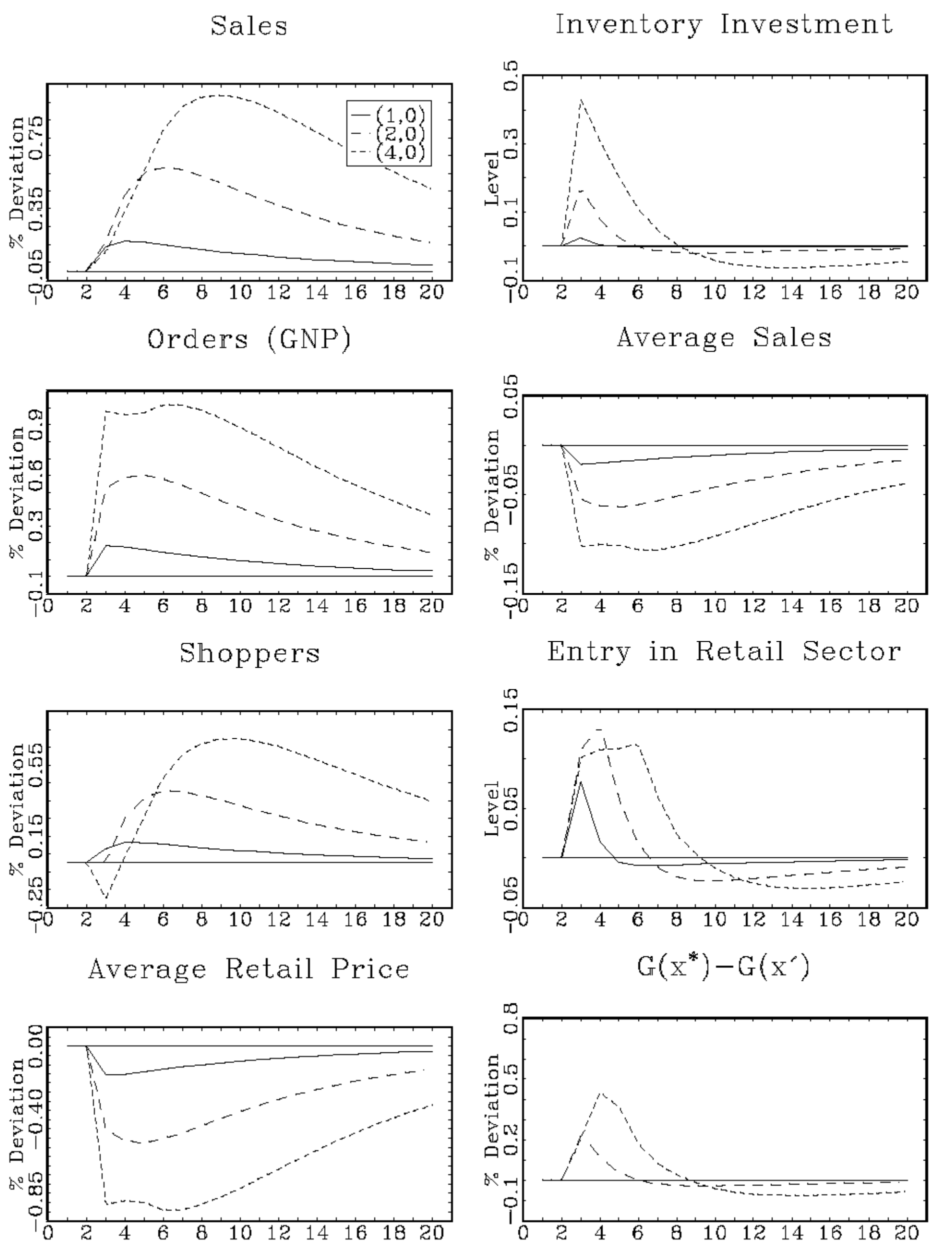

Figure 3: One-Price Impulse Responses to Discount Rate Shocks, $\rho = 0.9$

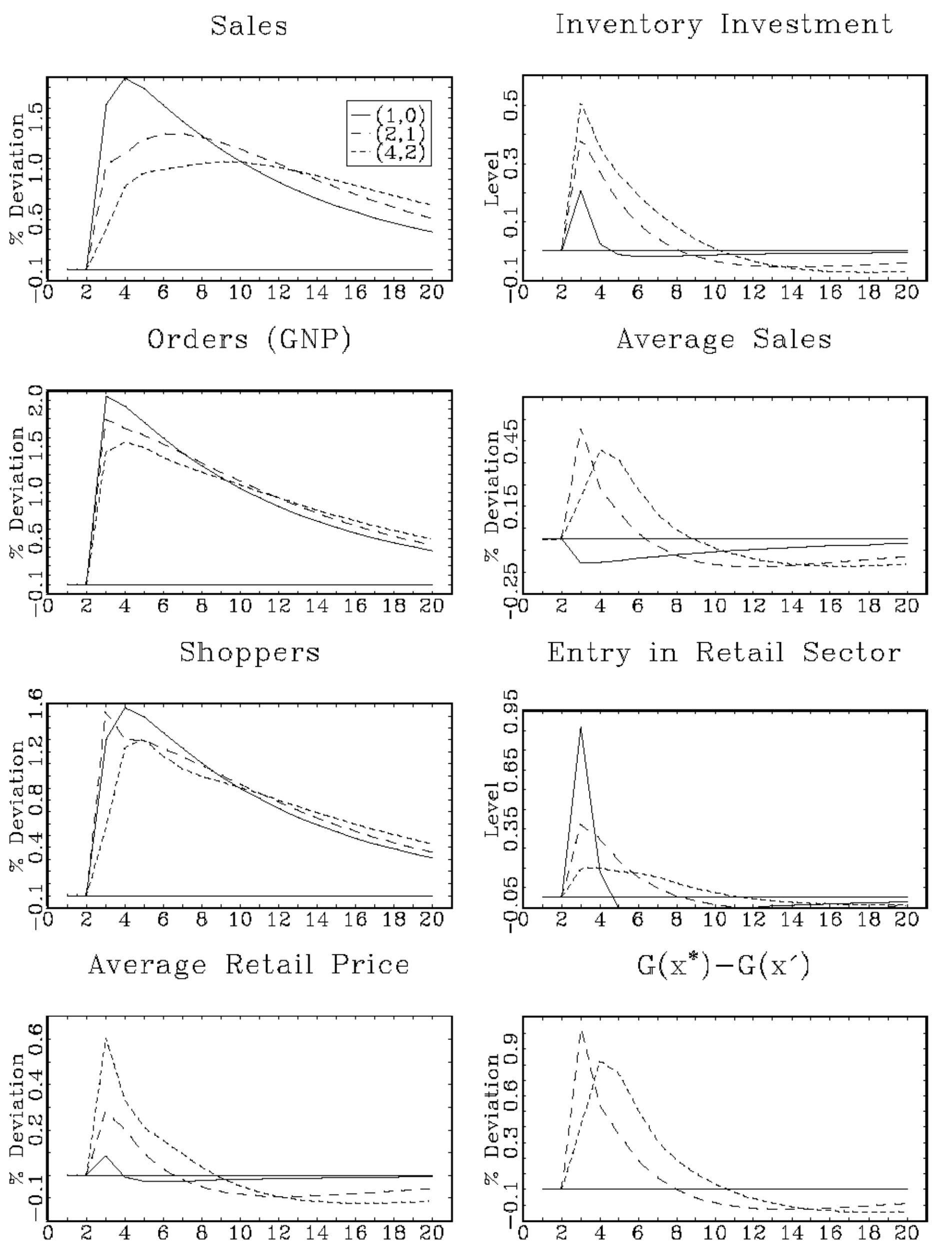

Figure 4: Two-Price Impulse Responses to Technology Shocks, $\rho = 0.9$

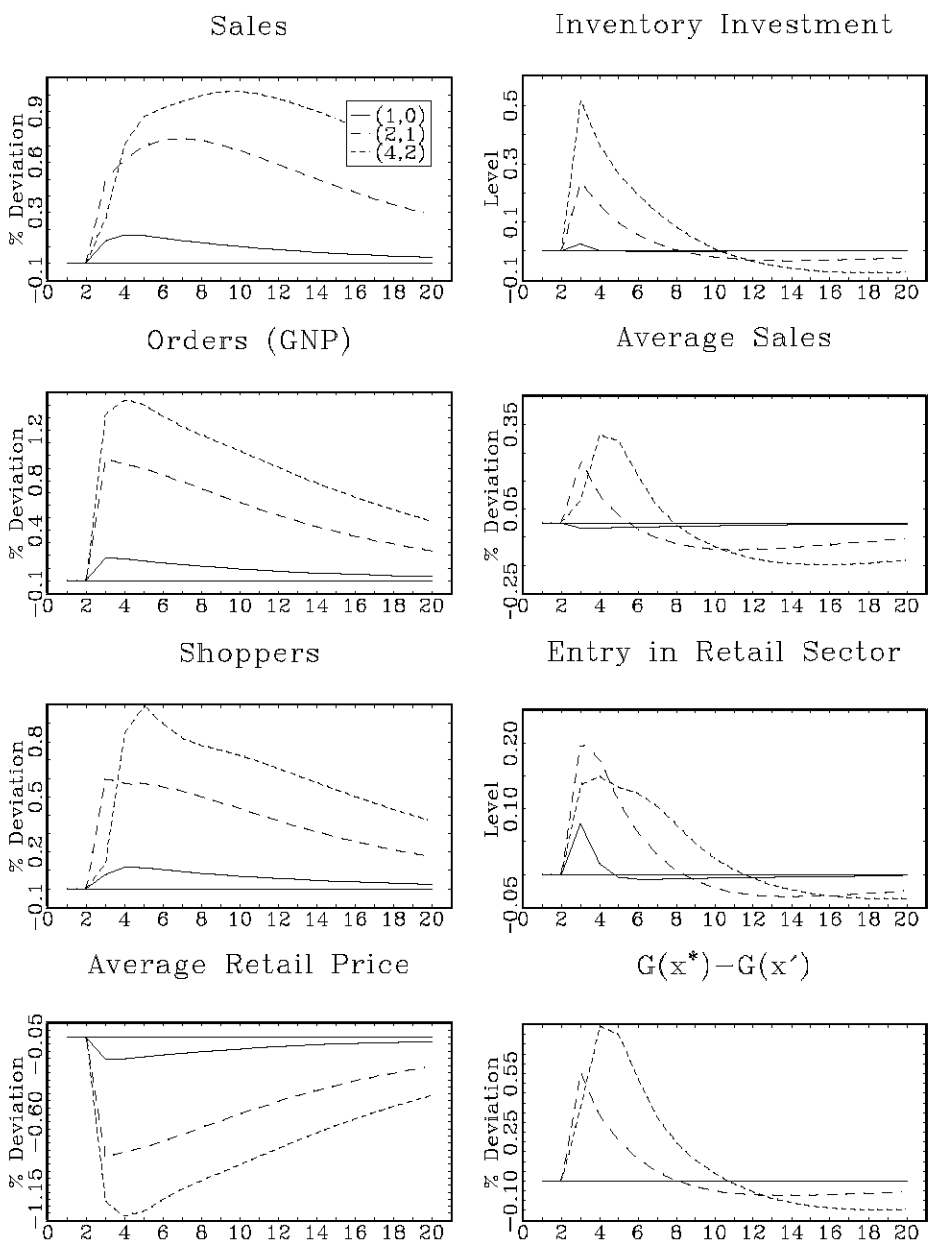

Figure 5: Two-Price Impulse Responses to Discount Rate Shocks, $\rho = 0.9$