

Discussion Paper 104

Institute for Empirical Macroeconomics
Federal Reserve Bank of Minneapolis
250 Marquette Avenue
Minneapolis, Minnesota 55480-0291

September 1995

(S,s) Inventory Policies in General Equilibrium

Jonas D. M. Fisher and Andreas Hornstein*

University of Western Ontario

ABSTRACT

We study the aggregate implications of (S,s) inventory policies in a dynamic general equilibrium model. Firms in the model's retail sector face idiosyncratic demand risk, and (S,s) inventory policies are optimal because of fixed order costs. The model economy replicates salient features of the business cycle and reconciles evidence that orders are more volatile than sales, and that inventory investment is positively correlated with sales. There are two main results. First, we find that general equilibrium effects and the optimal order size are important for the economy's response to exogenous shocks. Second, we find that key features of our results are independent of the presence of idiosyncratic risk.

*Both authors are affiliated with the University of Western Ontario and thank SSHRC for financial support. We also wish to thank Mark Bilal, Hugo Hopenhayn, Ig Horstmann, Michael Parkin, Richard Rogerson, 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. Our e-mail addresses are: fisher4@sscl.uwo.ca and hornstein@sscl.uwo.ca. The views expressed herein are those of the authors and not necessarily those of the Federal Reserve Bank of Minneapolis or the Federal Reserve System.

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 Scarf (1960)'s model of (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.

We study a dynamic general equilibrium model which includes a retail sector. Firms in the retail sector face idiosyncratic demand risk and use (S, s) inventory policies because of a fixed cost of ordering. The model economy's response to exogenous shocks is consistent with salient features of the business cycle. It is also consistent with evidence that orders are more volatile than sales and that inventory investment is positively correlated with sales for many firms, industries, and in the aggregate.¹

There are two main results. First, we find that general equilibrium effects and the size of the optimal order are important for the economy's response to exogenous shocks. Second, we find that key features of our results for the aggregate economy are independent of the presence of idiosyncratic risk. This latter finding suggests to us that working with general equilibrium (S, s) inventory models that eschew idiosyncratic risk may be a fruitful modelling strategy for future work on the relation between inventories and the business cycle.

Recently Blinder and Maccini (1991) have restated the view that the behavior of inventories and aggregate fluctuations are closely connected. They argue that most of the downturn

¹One source for this evidence is Blinder and Maccini (1991). We have redone their calculations using alternative detrending procedures and data sources with qualitatively similar results.

in output in a typical recession can be attributed to a decline in inventory investment. Furthermore, they find that components of aggregate inventory investment for which the key decision variable is deliveries dominate movements in aggregate inventory investment. Since (S, s) inventory policies are appropriate in these cases, it is reasonable to study their potential to amplify and propagate impulses in our model.

We consider two impulses, a productivity shock and a shock to the time discount factor. We interpret the latter as a reduced form representation of an interest rate shock. Three features of the stochastic properties of our model are of particular interest for business cycle analysis. First, hours worked, orders, sales, average sales per retailer, and inventory investment are all procyclical when the model economy is subjected to the two shocks. Second, (S, s) inventory policies have little impact on the response of aggregate variables to a productivity shock. Third, while (S, s) inventory policies have only a moderate effect on the propagation of an interest rate shock, they amplify considerably the response of the economy to such a shock.

Our paper builds on and extends the work of Caplin (1985) and Cabellero and Engel (1991), who consider the aggregate implications of (S, s) inventory policies in the context of partial equilibrium models with a fixed measure of sellers who follow potentially different, but fixed policies.² We extend their work along two dimensions. First, since the business cycle is characterized by a high degree of serial correlation, we need to allow for variation in the intervention points which define the inventory policy. As a step in this direction, we consider an indirect representation in which the number of sellers who order in a period varies. Second, since partial equilibrium analysis ignores potentially important general equilibrium interactions, we work within a general equilibrium framework.³

The foundation of our model is a simple example from Stokey and Lucas (1989), which appears to capture the main characteristics of (S, s) inventory models studied to date. In the example a retailer faces a stochastic, integer valued demand for her good, and a fixed cost of ordering. We modify the example to allow retailers to set prices for their goods, and we

²Other studies along these lines include Blinder (1981), Bertola and Cabellero (1990) and Lovell (1993).

³Previous 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).

provide a rationale for the presence of firm specific demand by introducing trading frictions. General equilibrium is achieved by embedding the retail sector into a standard stochastic growth model.

The rest of the paper proceeds as follows. In the next section we describe the optimal inventory problem for a retailer with fixed ordering costs. We also discuss the aggregate implications, in partial equilibrium, of the optimal inventory policy for this example and compare them with previous work. In section 3 we embed a retail sector where retailers use (S, s) policies into a general equilibrium model. Section 4 contains a description of our findings from a numerical analysis of the model. In the fifth section we conclude.

2. Optimal Inventory Policies in the Retail Sector

The theory of optimal inventory decisions in the presence of non-convex costs has been studied extensively in operations research. The major contribution is the work by Scarf (1960) on the optimality of an (S, s) inventory policy for a retailer who faces an exogenous demand process for her good, and a fixed cost of ordering. Stokey and Lucas (1989, pp. 389-390) provide a simple example which appears to capture the basic features of much of the research on (S, s) inventory policies along this line.

Since we want to introduce retailers into a general equilibrium environment we have to model the determination of prices. Building on the assumption that the quantity demanded of the good is specific to the retailer, we suppose that the price of the good is specific to that retailer. This means that a retailer has some monopoly power and sets prices, and sales depend on prices. In this section we modify the example of Stokey and Lucas (1989) accordingly and derive the optimal price and inventory policy for a retailer who can control the probability of a sale by the choice of price. We then discuss the implications of (S, s) inventory policies for the behavior of aggregate orders, sales, and inventory investment in a market with a large number of identical retailers.

2.1. The Retailer Problem

Consider a retailer whose inventory stock can only take integer values, $x = 0, 1, \dots, \bar{x}$. 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.⁴ If a retailer places an order of size $x > 0$ she has to pay $\gamma_0 + \gamma_1 x$, with $\gamma_0, \gamma_1 > 0$, that is there is a fixed cost for each order.

Demand for the good of the retailer has two components. 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. If the shopper accepts the offer he will buy one unit of the good. The utility a shopper derives from the consumption of the good is random and not observed by the retailer.⁵ The probability that a shopper will accept the offer at price p is $\delta(p)$, where 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 capital value of a retailer with a positive inventory level x is

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

or alternatively

$$V_x = \max_p \left\{ \theta\delta(p) [p - R^{-1}(V_x - V_{x-1})] + R^{-1}V_x \right\} \quad (2.1)$$

and the optimal order size solves

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

⁴This is not a restriction since we can always choose \bar{x} such that it does not constrain optimal decisions on inventory holdings. We introduce a capacity constraint because it will simplify the discussion of the general equilibrium environment in the next section.

⁵In 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.

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 $x = 0$.

From (2.1) it is apparent that the optimal price choice can be interpreted as a simple static profit maximization problem for a monopolist 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. There are no direct costs associated with inventory holdings, therefore V_x is positive and nondecreasing in x . In Proposition 1 we show that, under weak assumptions about demand, $\delta(p)$, $\{V_x\}$ is increasing at a decreasing rate, that is it is a ‘concave’ function of the stock of inventories.⁶ As Proposition 1 also shows, this implies that retailers with larger inventories charge lower prices.⁷

2.2. Aggregate 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. There is, however, no reason to expect that, in general, the behavior of inventories at the microeconomic level and at the aggregate level are closely connected.⁸

Consider our example and assume that there is a large number of identical retailers. In particular let there be a continuum of retailers, and let the total measure of retailers be N . Suppose also that demand, that is the event of a sale, is independent across retailers. In this case, the optimal (S, s) policy defines a Markov process on the inventory levels 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 $\sum_{x=1}^{\infty} \bar{g}_{x,t} = N$.

In this environment there is no aggregate risk, and retailers face idiosyncratic risk only.

⁶The propositions we refer to in the main text are stated formally and proved in appendix A.

⁷In 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.

⁸This point has been made before, see for example Caballero (1992).

Assume that a law of large numbers holds, then the optimal order and price policies determine the evolution of the aggregate distribution of inventory stocks as follows,

$$\begin{aligned} g_{x,t+1} &= (1 - \theta\delta_x) \bar{g}_{x,t} + \theta\delta_{x+1} \bar{g}_{x+1,t} \\ \hat{g}_{x,t} &= \begin{cases} \theta\delta_1 \bar{g}_{1,t-1} & \text{if } x = x^* \\ 0 & \text{if } x \neq x^* \end{cases} \end{aligned} \quad (2.2)$$

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.

A simple application of results for the convergence of Markov chains shows that the cross-sectional distribution of inventory levels after orders have been placed converges to an invariant distribution⁹

$$\bar{g}_x = \frac{\delta_x^{-1}}{\sum_{i=1}^{x^*} \delta_i^{-1}} N \text{ for } 1 \leq x \leq x^* \text{ and } \hat{g}_x = \begin{cases} \theta\delta_1 \bar{g}_1 & \text{if } x = x^* \\ 0 & \text{if } x = 0 \end{cases} . \quad (2.3)$$

This means that aggregate orders and sales are constant. In a similar environment Caplin (1985) proves 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.¹⁰

This example assumes that there is no aggregate risk. With random aggregate demand, aggregate orders and sales should be random, and it may be the case that orders are more volatile than sales. Consider the following modification of the law of motion for the cross-sectional distribution of inventory holdings. Let the probability of a match θ_t be a stationary stochastic process with domain $[\underline{\theta}, \bar{\theta}]$, $0 < \underline{\theta} < \bar{\theta} < 1$. Assume that the order policy x^* and the sales probability δ_x implied by the optimal pricing rule do not depend on the state of

⁹See Stokey and Lucas (1989), Theorem 11.4.

¹⁰Caplin (1985) also shows that in his environment the invariant distribution over inventories is uniform. This arises in our example only if all retailers post the same price.

the economy, in particular the probability of a match θ . In Proposition 2 we show that in this case the cross-sectional distribution of inventory holdings after orders have been placed still converges to the invariant distribution defined in (2.3). 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.

In Proposition 3 we describe the long run behavior of aggregate sales, orders, and inventory investment with random aggregate demand. We show that if the distribution of inventories has converged then: (1) the variance of orders equals the variance of sales, (2) inventory investment and sales are negatively correlated; (3) orders and sales are not more volatile than demand and the autocorrelations of orders and sales equal the autocorrelation of demand; and (4) orders and sales are positively correlated only if demand is positively correlated.

These findings suggest two conclusions. First, as (1) and (2) indicate, the aggregate implications of (S, s) inventory policies are not necessarily consistent with two key observations from the inventory literature – orders are more volatile than sales, and inventory investment is positively correlated with sales. Second, (3) and (4) suggest that (S, s) inventory policies do not necessarily propagate or amplify shocks, at least when studied in partial equilibrium.

The convergence result for aggregate uncertainty will not apply if the optimal pricing rule is not fixed, but is a function of the aggregate state, in particular of θ . This is immediate since the invariant inventory distribution is a function of the sales probabilities δ_x implied by the optimal pricing policy. Furthermore sales probabilities themselves may depend on the aggregate state of the economy. The inventory distribution may also fail to converge to an invariant distribution if the number of sellers is not constant, as is assumed in the existing literature on the aggregation of (S, s) policies.

In the absence of convergence, sales and orders can have different volatilities. For the next step in a study of the aggregate implications of (S, s) inventory policies we have two choices. Either we stay in the partial equilibrium framework and assume empirically reasonable laws of motion for variables like θ , δ , \hat{g} , and R , which are exogenous in this framework, or we include the retail sector in a general equilibrium model and derive the laws of motion from the underlying economic environment. We pursue the second alternative.

3. The Retail Sector in a General Equilibrium Model

In this section we embed the retail sector described in the previous section into a variant of a standard stochastic growth model. There is an infinitely lived representative household whose utility depends on the consumption of goods which are bought from a large number of retailers, and non-leisure activities. The household earns income from the supply of labor to a manufacturing sector. 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 market 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. Productivity shocks are included because they play a major role in real business cycle models. Shocks to the time discount rate are included in order to study how endogenous variables respond to interest rate shocks. In the remainder of this section we describe the environment in more detail, and then outline a recursive equilibrium. A discussion of a simplified version of the model concludes this section.

3.1. The Environment

The preferences of the representative household are

$$E_0 \left[\sum_{t=0}^{\infty} [\Pi_{\tau=0}^t \beta_{\tau}] U(c_t, n_t) \right] \quad (3.1)$$

where c_t is an index of consumption goods, and n_t is time spent in non-leisure activities. The discount factor β_t is an exogenous, possibly random, variable with bounded support,

$0 < \beta_t < 1$. We assume the following parametric form for U ,

$$U(c, n) = \ln \left[c - \frac{n^{1+\nu}}{1+\nu} \right] \quad (3.2)$$

with $\nu > 0$ and $1/\nu$ is the wage elasticity of non-leisure activities.¹¹

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. 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 $\mathbf{U} = [\underline{u}, \bar{u}]$, $0 \leq \underline{u} < \bar{u} \leq \infty$, and twice differentiable cumulative density function F . The consumer is not informed about the quality of any retailer's good before he actually meets the retailer. The consumption index is the integral over all utilities

$$c_t \equiv \int_0^{N_t} u_t(i) c_t(i) di. \quad (3.3)$$

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

$$s_t = \zeta n_{s,t}, \quad (3.4)$$

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

¹¹This particular parametric form has the property that wage changes have no wealth effect, see for example Greenwood, Hercowitz and Huffman (1988). We have chosen this utility function because it allows for a relatively strong procyclical response of shopping time by the household, which results in procyclical average sales for retailers. An alternative route, which we have explored in a previous version of this paper, is to introduce a home-production structure with non-market goods, see for example Benhabib, Rogerson and Wright (1991). For our purposes the properties of the two frameworks are essentially the same.

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

$$n_t = n_{y,t} + n_{s,t}. \quad (3.5)$$

Goods sold by retailers are produced in the manufacturing sector, where labor is the only input to a decreasing returns production function

$$y_t = z_t n_{y,t}^\alpha \quad (3.6)$$

with $0 < \alpha \leq 1$ and $z_t > 0$ is random exogenous productivity. There may be a fixed factor of production, for example capital, which we do not model. 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 that there is free entry to the retail sector. We also assume that there are no explicit costs to entry besides the cost of making an order. 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. 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

$$\sum_{x=1}^{\bar{x}} \hat{g}_{x,t} (\gamma_0 + \gamma_1 x) \leq y_t, \quad (3.7)$$

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.¹² 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,

¹²Here we make use of the fact that in equilibrium retailers with positive inventory holdings do not place any orders.

then the number of matches is given by

$$M_t = \min \{ \varphi S_t^\mu N_t^{1-\mu}, N_t, S_t \} \quad (3.8)$$

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, $g_t = (g_{1,t}, g_{2,t}, \dots, g_{\bar{x},t})$.¹³ 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. 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.¹⁴ Recall from section 2 that a retailer with inventory x posts a price p_x for her

¹³Note that because there is a capacity constraint \bar{x} on inventories the aggregate state is finite dimensional.

¹⁴In the following we will drop time subscripts and let a prime denote next period's value of a variable.

product before any 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, the shopper 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 $u_s(i)$, and decides whether to purchase the good or not. Note that the shopper's decision cannot depend on the choices of other shoppers since shoppers do not communicate with each other. Also a shopper will not purchase more than one good since the household cannot store market 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

$$q = \sum_{x=1}^{\bar{x}} \int_{u \geq r_x} u F(du) (\bar{g}_x/N), \quad (3.9)$$

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

$$e = \sum_{x=1}^{\bar{x}} p_x (1 - F[r_x]) (\bar{g}_x/N). \quad (3.10)$$

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 distribution of prices for market 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 $g' = \Gamma(g, \mathbf{z})$.

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 $R(g, \mathbf{z}) = [\int Q(\Gamma(g, \mathbf{z}), \mathbf{z}'|g, \mathbf{z}) d\mathbf{z}']^{-1}$. 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,

$$\begin{aligned}
\mathcal{U}(\tilde{a}, g, \mathbf{z}) = & \max \{U(c, n) + \beta E[\mathcal{U}(a[g', \mathbf{z}'|g, \mathbf{z}], g', \mathbf{z}') | g, \mathbf{z}]\} \\
\text{s.t. } & c = \zeta n_s \omega q, \\
& \zeta n_s \omega e + \int Q(g', \mathbf{z}'|g, \mathbf{z}) a(g', \mathbf{z}'|g, \mathbf{z}) d\mathbf{z}' = \tilde{a} + w n_y + \pi, \\
& n = n_s + n_y, \text{ and} \\
& g' = \Gamma(g, \mathbf{z}), \omega = \Omega(g, \mathbf{z}), p_x = P_x(g, \mathbf{z}), \\
& w = W(g, \mathbf{z}), \pi = \Pi(g, \mathbf{z}).
\end{aligned} \tag{3.11}$$

It is not obvious that the household's static decision problem regarding the acquisition of consumption goods is a concave maximization problem. Consider 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. Let λ be the marginal cost of quality for this policy, $\lambda = \partial e / \partial q$. The cost function represents a price index for average quality. It is straightforward to show that the first order conditions for reservation utilities are sufficient to characterize the optimal reservation utility policy. In a second stage 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 , but the problem can be transformed into a concave problem.

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 $(q_j, e_j)_{j=1}^J$, 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 now 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.¹⁵

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

$$\begin{aligned} U_c r_x &= \kappa p_x \\ U_c \zeta \omega q - U_n &= \kappa \zeta \omega e \\ U_n &= \kappa w \end{aligned} \tag{3.12}$$

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 (3.12) 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

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

This suggests that we rewrite the reservation utility policy as follows

$$r(p) = p/\lambda \text{ with } \lambda = w \frac{U_c}{U_n} \text{ and } p \in (\lambda \underline{u}, \lambda \bar{u}). \tag{3.13}$$

We can now modify the retailer's problem from section 2 for a recursive equilibrium in an economy 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

¹⁵This approach is analogous to the use of lotteries in the context of non-convex problems, see for example Rogerson (1987) or Hansen (1985).

marginal rate of substitution between present and future non-leisure time is equalized to the contingent claim price

$$Q(g', \mathbf{z}' | g, \mathbf{z}) = \tilde{Q}(g', \mathbf{z}', g, \mathbf{z}) \Phi(\mathbf{z}' | \mathbf{z}) \text{ and } \tilde{Q}(g', \mathbf{z}', g, \mathbf{z}) = \beta \frac{U_n(c', n') / W(g', \mathbf{z}')}{U_n(c, n) / W(g, \mathbf{z})}. \quad (3.14)$$

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

$$V_x(g, \mathbf{z}) = \max_p \left\{ \theta \delta (p/\lambda) \left(p - E \left[\tilde{Q}(g', \mathbf{z}' | g, \mathbf{z}) V_{x-1}(g', \mathbf{z}') \mid g, \mathbf{z} \right] \right) \right. \\ \left. + [1 - \theta \delta (p/\lambda)] E \left[\tilde{Q}(g', \mathbf{z}' | g, \mathbf{z}) V_x(g', \mathbf{z}') \mid g, \mathbf{z} \right] \right\}, \quad (3.15)$$

where $g' = \Gamma(g, \mathbf{z})$, $\theta = \Theta(g, \mathbf{z})$ and $\lambda = \Lambda(g, \mathbf{z})$.

The optimal order size $x^* = X^*(g, \mathbf{z})$ maximizes the gains from entry, and because of free entry the rate at which retailers restock their inventory $\hat{g}_x(g, \mathbf{z})$ will adjust such that these gains continue to be zero in equilibrium,

$$x^* \in \arg \max_x V_x(g, \mathbf{z}) \text{ and } V_{x^*}(g, \mathbf{z}) = \gamma_0 + \gamma_1 x^*. \quad (3.16)$$

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

$$\bar{g}_x = g_x + \hat{g}_x \text{ and } g'_x = [1 - \theta \delta (p_x/\lambda)] \bar{g}_x + \theta \delta (p_{x+1}/\lambda) \bar{g}_{x+1} \\ \text{for } x = 1, \dots, \bar{x} \text{ and } \bar{g}_{x+1} = 0. \quad (3.17)$$

Given free-entry it is not optimal to reorder before all inventories are depleted if V_x defined in (3.15) is 'concave.' Consider a retailer with existing inventory x and an order of size Δx . If the retailer adds the order to the existing inventory, the capital value is $V_{x+\Delta x}$. Alternatively, because of our free entry assumption the retailer could simply open up another store, and 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 we verify it is always concave.

At the current stage of our research we would like to interpret the free entry assumption as a simple way to capture variations in the upper intervention point of the (S, s) policy.

In the partial equilibrium framework of the previous section the optimal inventory policy and the number of retailers are fixed. However, following a change in demand for the retail good, we expect that either the order size or the number of retailers who place an order will change.¹⁶ In the computational experiments that follow we focus on equilibria in which the order size is constant so that only variation along the latter margin occurs.¹⁷

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

$$\max z n_y^\alpha - W(g, \mathbf{z}) n_y. \quad (3.18)$$

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.

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

- (i) *given $\Gamma, \Omega, P_x, W, \Pi,$ and $Q,$ the decision rules S, N_s, N_y, Λ solve the household's problem (3.11);*
- (ii) *given $W,$ the decision rules N_y and Y solve the manufacturer's problem (3.18);*
- (iii) *given $\Gamma, \Theta, \Lambda,$ and $Q,$ the decision rule P_x and the value function V_x solve the retailer's problem (3.15);*
- (iv) *the entry rule \hat{G}_x satisfies (3.16);*
- (v) *given $\hat{G}_x, P_x, \Lambda,$ and $\Theta,$ the law of motion Γ is consistent with (3.17);*
- (vi) *given S and $N = \sum_x (g_x + \hat{G}_x),$ the matching probabilities Ω and Θ are consistent with (3.8);*

¹⁶This is independent of variations in the lower intervention point.

¹⁷Another 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.

(vii) given Y and \hat{G}_x the market for manufacturing goods clears, (3.7) is satisfied with equality.

3.3. A Tractable Example

The law of motion for the equilibrium distribution of retailers depends on the optimal order policy and on the pricing strategy of retailers. This can make it very difficult to compute an equilibrium. To facilitate the characterization of an equilibrium in our model we therefore impose the following two conditions.

First, we restrict attention to an equilibrium in which the optimal order size is state independent. In the economy without aggregate uncertainty the optimal order size in a steady state is unique. Our characterization of the equilibrium involves an approximation of the economy around the steady state. This means that if the aggregate uncertainty is not ‘too large’ the optimal order size should remain at its steady state value. We do check that this is indeed true for our simulations.

Second, we assume that the idiosyncratic utility shock for a retailer’s good has finite support, which considerably simplifies the characterization of the optimal pricing strategy. In particular, we consider a two point support \mathbf{U} with a symmetric uniform distribution

$$u = \begin{cases} u_l = 1 - \eta, & \text{with probability } 1/2 \\ u_h = 1 + \eta, & \text{with probability } 1/2 \end{cases}, \quad (3.19)$$

and $0 < \eta < 1$. The reservation utility policy, accept a good if $u \geq p/\lambda$ where $\lambda = wU_s/U_n$, continues to be optimal, and prices offered by retailers remain non-increasing in the inventory level. We now show that the optimal price strategy is such that retailers will offer at most two different prices, and given this price strategy the reservation utility policy is optimal.

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.¹⁸

A retailer's optimal price offer continues to be nonincreasing in the inventory level. Together with the preceding argument this implies that there is some critical inventory level x' such that,

$$p_x = \lambda u_j \text{ with } j = h \text{ if } x \leq x' \text{ and } j = l \text{ if } x > x'. \quad (3.20)$$

In the following we refer to this as an x' pricing scheme. Notice that the case where only one price is offered in a given period corresponds to $x' = 0$ in this notation.

The reservation utility policy of the household is optimal given the price offers of retailers. For the policy described above the household uses pure strategies concerning the acceptance of price offers. In Proposition 4 we show that the household is indifferent between all mixed strategies of the form 'accept if $u_j > p_i/\lambda$ and accept with probability a_j if $u_j = p_i/\lambda$ for $i, j = l, h$ '. We assume that the household picks the pure strategy such that shoppers accept an offer even if they are 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.

In appendix B we show how the equilibrium of this version of the model can be characterized in terms of the solution to a particular system of functional equations. We also show how the system of functional equations can be solved using numerical methods. With a solution of this system in hand we can compute endogenous variables for a given sequence of exogenous shocks.

4. Properties of the Model

In this section we describe our findings from a numerical analysis of the tractable example. We focus on four issues. First, we investigate the ability of the model to mimic some basic

¹⁸This 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.

features of the business cycle. Second, we examine whether the model can reconcile evidence that orders are more volatile than sales and sales and inventory investment are positively correlated. Third, we consider the extent to which (S, s) inventory policies propagate and amplify exogenous disturbances to the economy. Fourth, we introduce a competitive economy as a special case of our environment and study the role of idiosyncratic risk. The section begins with a description of how we selected parameter values and the experiments that we conducted to shed light on these issues. We then describe the outcome of our experiments by examining impulse response functions and unconditional statistics for key variables.

4.1. Design of the Experiments

To implement our model we must specify $\eta, \nu, \zeta, \varphi, \mu, \alpha, \gamma_0, \gamma_1$, the unconditional means of the discount rate, β , and the technology shock, z , and the nature of aggregate uncertainty. The experiments we conducted are based on five baseline parameterizations of the model. These are summarized in table 1. By the end of this sub-section, the entries in this table will have been explained.

The distribution of idiosyncratic utility shocks is calibrated to obtain particular x' pricing schemes. In appendix C we derive a formula for η such that a particular x' pricing scheme is optimal if there is no aggregate uncertainty. With aggregate uncertainty we then take the pricing scheme as given in our computations because the selected η 's ensure the optimality of particular x' pricing schemes in a region of the state space local to steady state. We do check the optimality of the pricing scheme in our simulations.

The specification of η and the implied possibilities for x' pricing schemes form an integral part of the design of our experiments. 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 x^* and x' , one for each parameter set, are optimal from the perspective of a representative retailer.¹⁹ These combinations are $(x^*, x') = (1, 0), (2, 0), (2, 1), (4, 0)$, and $(4, 2)$.²⁰ From before, we calibrate η to fix x' . To fix x^* we use the fixed cost parameter

¹⁹In appendix B we describe the procedure we use to verify that along a given simulation the assumed ordering and pricing strategies are in fact optimal.

²⁰For example, a $(4, 2)$ economy has an optimal order size equal to 4 and a pricing scheme in which $p_x = \lambda u_h, x = 1, 2$ and $p_x = \lambda u_l, x = 3, 4$.

γ_0 . The γ_0 and η pairs which underlie the (x^*, x') combinations that form the basis of our experiments are shown in table 1.

The intratemporal elasticity of labor supply with respect to the real wage is approximately equal to $1/\nu$.²¹ One feature we would like the model to have is that the two market related household activities, 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.²²

We calibrate ζ to fix the steady state cross sectional variance of sales for retailers.²³ Notice that the matching probability for θ determines this variance. For all values of θ except those close to unity the cross sectional variance of sales relative to its mean is very large. For this reason we work with a steady state value of θ fairly close to unity. In particular, we calibrate ζ so that in steady state $\theta = 0.85$. This criterion has the additional desirable property that it weakens a channel of persistence in the model that has little to do with (S, s) inventory policies, relative to a lower value of θ .²⁴

In keeping with related work in the real business cycle literature we set $\alpha = 0.64$. We let $\beta = 1.05^{-1/12}$, so that a period in the model is a month and the real interest rate in steady state is 5% at an annual rate. We select parameters so that $N = S$ in steady state. In conjunction with the calibration criterion for ζ this requires that we set $\varphi = 0.85$. By choosing this value of φ the steady state probability of a match from the perspective of a shopper is the same as for a retailer, that is $\omega = \theta$ in steady state. We fix $z = 1$ and $\gamma_1 = 1.175$. These are normalizations which have no major consequence for the analysis. Particular choices for the matching exponent, μ , have a significant impact on the analysis so our approach to selecting this parameter deserves some discussion.

The matching exponent governs the responsiveness of ω and θ to changes in S and N .

²¹This is true for our parameterization where shopping time is small relative to work-time.

²²As mentioned in footnote 10, we could have achieved the same result by formulating a home production version of the household's problem in which home and market produced goods are relatively easy to substitute.

²³We are unable to calibrate ζ to match evidence in Juster and Stafford (1991) on time devoted to shopping activities because n_s is implicitly determined by n_y in steady state conditional on the expected surplus from shopping – once ν is selected the expected surplus from shopping fixes n_s .

²⁴Lower 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.

We find that in our two-price experiments ($x' > 0$), several key impulse response functions display a saw-tooth pattern when μ is relatively large. The reason for this will be elaborated upon later, but for now it suffices to say that it is due to the implied high elasticity of θ with respect to N and S . While these saw-tooth patterns are indicative of the rich array of implications (S, s) inventory policies have in our model, we view them as being empirically implausible. For this reason, we set μ to a relatively small value. In particular $\mu = 0.25$ in all our experiments.

With the parameters described above in hand it is possible to 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 helpful for this purpose: the steady state aggregate inventory-sales ratio and the steady state expected time between orders for a typical retailer. These values are shown in table 2 for the five different baseline parameterizations.

The aggregate inventory-sales ratio (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 1.47.²⁵

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.

It remains to describe how we specify aggregate uncertainty in our experiments. As

²⁵This number is based on the ratio of retail inventories to retail sales as given by the variables IVRR and RTRR, respectively, from the Citibase data bank.

described in the last section, we allow for the possibility of technology shocks and discount rate shocks. We study the implications of these shocks in isolation and use a common form for their forcing processes,

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

for $(\mathbf{z}_t, \mathbf{z}) = (z_t, \beta), (\beta_t, z)$.

Since the period length in the model is a month we adopt the view that the exogenous disturbances are best modeled as persistent. In our baseline experiments, then, we fix $\rho = 0.9$. However, we are interested in learning about the internal propagation mechanisms in the model so we also consider experiments with $\rho = 0$. To keep our computations as simple as possible, we need to ensure that fixed pricing schemes and ordering strategies are optimal for our parameterizations. We accomplish this by assuming that the standard deviation of innovations to the exogenous random variables, σ , is small. In all our experiments we assume $\sigma = 0.01\%$.

4.2. Impulse Response Functions in the Various Experiments

In figures 1 to 6 we display impulse response functions for the different (x^*, x') combinations and the alternative specifications of aggregate uncertainty. All these figures are laid out in the same way. The first column displays impulse 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 impulse 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 impulse 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.²⁶ The two exceptions are inventory investment and net entry, which are shown as absolute deviations from their stochastic steady state values. All the impulse responses are displayed

²⁶These 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.

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.

Several measurement issues must be addressed. First, sales are measured using a base period average price for retail sales. We take the average retail sale price from assuming the aggregate shock equals its mean forever as the base period price (measured in terms of manufactured goods). Second, inventories are measured in a way consistent with the national income and product accounts. In particular, we measure the stock of inventories at the end of the period in terms of replacement cost. The replacement cost is defined to be the average cost of replacement measured in terms of manufactured goods.²⁷ Third, we measure the average sale price of consumption goods on a period by period basis. Finally, net entry is defined as total entry in the period less the number of retailers that stock out at the end of the previous period.

We begin our discussion by considering impulse response functions associated with the one-price versions of the model. Consider first the implications of technology (z) shocks with $\rho = 0.9$, as displayed in figure 1. Retail sales, inventory investment, orders, the number of shoppers, and net entry into the retail sector, all exhibit positive and persistent responses. With the exception of the response of average sales per retailer (in the one-price models θ), which uniformly exhibits a persistent drop, the responses suggest the model can replicate the qualitative features of a boom. It is unclear from this figure whether orders are more volatile than sales, but it does appear that sales and inventory investment are positively correlated. Finally, notice that although the response of inventory investment increases with the order size, the responses of sales and orders decline with the order size.

Contrary to the conventional view of technology shocks, a demand interpretation of z shocks may be valid for this model since the relative price of retail goods responds positively in the periods following an innovation.²⁸ One problem with the z model, and its demand shock interpretation, however, is that average sales drop in a boom.²⁹ The behavior of

²⁷For example the replacement cost for a retailer with inventory holding x is $(\gamma_0 + \gamma_1 x) / x$.

²⁸On impact the real wage responds positively so the household's willingness to pay for consumption goods, given by equation (3.13) and indicated in figure 1, also increases.

²⁹With positive net entry and a larger number of shoppers the response of the matching probabilities

the distribution of inventories across retailers gives some indication that average sales in a two-price model will behave differently.

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 in both these experiments. If this kind of response carries over to the two-price versions of the model it would mean that a higher proportion of retail firms charge the lower price in response to a positive z shock. This *price distribution mechanism* would help to increase the marginal return to shopping, encourage a stronger response of shoppers, and, potentially, reverse the response of average sales.

We now discuss what the impulse response functions suggest about the potential for (S, s) inventory policies to generate persistence. To help our discussion it is useful to examine figure 2 in which impulse responses to z shocks with $\rho = 0$ are displayed, along with figure 1. To simplify the exposition, we focus on the responses of sales, orders and inventory investment. Notice that the persistence of the response of sales is positively related to x^* , regardless of the value of ρ assumed.³⁰ This seems to suggest that (S, s) inventory policies do in fact have the potential to contribute to the persistence of aggregate fluctuations. However, these figures also show that while higher order sizes generate some persistence, most of the persistence in orders and inventory investment appears to be due to the assumed process of the exogenous shock. This suggests that the potential for (S, s) inventory policies to generate persistence in output and inventory investment may be limited.

In figures 3 and 4 the one-price response functions with discount rate (β) shocks for $\rho = 0.9$ and $\rho = 0$, respectively, are displayed. In qualitative terms the responses are similar to those in figures 1 and 2, except that here the average sale price of consumption goods falls and the response of the number of shoppers is not always positive. Thus, in terms of

appears ambiguous. However, it is straightforward to show, using equations (3.12), 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.

³⁰This is an implication of the permanent income hypothesis, which is easiest to see in the context of figure 2. In response to a temporary positive innovation to z the household, as in standard business cycle models, wants to smooth consumption. This encourages the household to save most of the extra output produced due to the innovation. In our model inventory investment is the only avenue for saving. Inventories are a capital good and, analogous to standard business cycle models, the household uses them to smooth consumption.

generating the qualitative features of a boom this version of the model performs roughly the same as the z shock case. Notice however, that for β shocks higher order sizes magnify the response of sales and orders to shocks.

Since the average retail price falls with a β shock, a supply-side explanation of the latter finding seems reasonable. Consider a retailer currently facing a constant discount factor β 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 $v_{x^*} = (1 - \beta^{x^*})/(1 - \beta)$. Now consider a permanent change in β . It is straightforward to show that the elasticity of v_{x^*} with respect to β is increasing in x^* . This says that the value of an optimal order is more sensitive to perturbations in the inventory holding cost, the larger is the optimal order size. In addition temporary changes in β should have a smaller impact on the value of the retailer relative to permanent changes. Since excess profits in the retail sector encourage entry, the positive relationship between the responsiveness of the value of a retailer to perturbations in β and the order size suggests that the amplification results may not be very surprising.

Of course this is just a partial equilibrium argument. However, the impulse responses in figures 3 and 4 suggest that this effect is dominant in the dynamic general equilibrium of our model. In figure 3, where innovations have a relatively long lived effect on the discount rate, $\rho = 0.9$, the amplification effect is pronounced. In fact it is so strong that in order to keep the value of a retail firm at the time of ordering equal to zero, net entry in the periods following a β innovation is amplified as x^* is made larger, contrary to our findings for the z case. Higher net entry means the size of the response of aggregate orders is positively related to x^* , again contrary to the z case. In figure 4 where $\rho = 0$, the amplification effect of a wider (S, s) band appears to be mitigated.

We now consider briefly impulse responses implied by the two-price models. To conserve space we consider the case of $\rho = 0.9$ only. The relevant response functions are displayed in figures 5 and 6 for z and β shocks, respectively. We have included in these figures the responses for the $(1, 0)$ model to illustrate the overall impact of (S, s) inventory policies and the price distribution mechanism on the responses of the endogenous variables.

Consider first the z impulse responses. These are qualitatively identical to the one-price responses with the exception that now average sales respond positively to a z impulse. This

is a result of the price distribution mechanism described previously. Notice that the optimal price and ordering policies of individual retailers have only a minor impact on the behavior of aggregate orders. However, the response of this variable still seems to be dominated by the assumed exogenous process for z . The overall pattern of responses seems consistent with interpreting z shocks as demand shocks.

The responses to β impulses in figure 6 are influenced by the price distribution mechanism in a manner similar to the z cases. In particular, average sales now respond positively to a β shock and the number of shoppers now always responds positively in the periods immediately following a shock, unlike in the one-price β experiments. Note that now orders clearly are more volatile than sales. The contribution of a wider (S, s) band to amplifying β shocks indicated in figures 3 and 4 is evident here also. Since the retail price falls following a β innovation, classifying β shocks as supply shocks in the two-price case seems valid.

Before turning to our simulation results, one feature of our model not illustrated by these impulse response functions is worth noting. In our description of the design of our experiments, mention was made of how saw-tooth impulse responses can emerge from this model. These patterns emerge in the two-price versions of the model with relatively high values for μ , the matching function parameter. Three factors contribute to this finding. First, as stated above, the elasticity of θ with respect to N and S is increasing in μ . Second, the price mechanism 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 mechanism 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.

4.3. Key Statistics in the Various Experiments

We now discuss results from simulations based on the specifications of the model described above. These are reported in tables 3 and 4. The layout of these tables is identical. Each table contains statistics related to sales (C), orders (Y), inventory investment (ΔI), average sales (C/N), the number of retail firms (N), the number of shoppers (S), and net entry (E). The notation σ_x denotes the standard deviation of the variable $x = C, Y, \Delta I$ relative to the standard deviation of technology and the discount rate, in tables 3 and 4, respectively. The notation $\rho(C, x)$ denotes the correlation coefficient for variables C and $x = Y, \Delta I, C/N, S, E$. Each table includes statistics based on monthly data for $\rho = 0.9$ and $\rho = 0$. In results not reported here we find the qualitative conclusions drawn from these monthly statistics translate to quarterly data. Finally, we log all variables before computing the statistics, except ΔI and E , in which case we use raw levels.

Consider table 3 first. For the most part, the results verify our findings from examining the impulse response functions. Three points are worth noting however. First, orders are more volatile than sales and inventory investment is positively correlated with sales in almost all the experiments. 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 mechanism is important for generating a positive unconditional correlation of inventories with sales.

Second, most of the variables display positive contemporaneous correlations with sales, regardless of the values taken by x^* and x' . The only exception to this rule, besides the previously mentioned inventory investment, is average sales. The unconditional correlation of average sales is only positive when the price mechanism is operational and $\rho = 0$.

Third, larger fixed costs and the implied higher values for x^* have a mixed effect on generating persistence. They tend to increase the persistence of sales fluctuations, especially in the $\rho = 0$ cases, but have little impact on the persistence of orders fluctuations. Also, they either make inventory investment more autocorrelated ($\rho = 0.9$ cases) or reduce negative autocorrelations ($\rho = 0$ cases).

Now consider table 4. The qualitative findings are identical to the case of z shocks. What stands out most here, however, is the strong association of larger values for x^* and x' with increased volatility in sales, orders and inventory investment. For example, consider the $\rho = 0.9$ case. 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. The Competitive Equilibrium as a Special Case.

The theoretical framework developed here captures what we believe are characteristic elements of previous work on (S, s) inventory policies. Of these elements, idiosyncratic risk for retailers has proven to be the most difficult to incorporate into a general equilibrium environment. We have introduced considerable microeconomic detail into the model to provide a consistent treatment of this type of risk. This allows for a rich economic structure, but it also makes it very difficult to study other interesting issues such as the interaction of inventory and physical capital investment.

The analysis of optimal inventory policies with fixed ordering costs simplifies considerably if 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. In a competitive equilibrium a 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.

The competitive equilibrium example 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 idiosyncratic risk, $\theta = \delta = 1$.³¹ 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 have performed the same experiments for the competitive economy as described above for the noncompetitive economy. With the exception of γ_0, ζ and η the parametrization of the competitive economy is exactly the same as in table 1.³² In figures 7 and 8 we display the impulse response functions for sales, orders, inventory investment, net-entry in the retail sector, the relative price of retail goods and the proportion of retailers who have sold less than half of their inventory stock.³³ The response of these variables to a persistent productivity shock is essentially the same in the competitive economy and the one-price noncompetitive economy. This is apparent from a comparison of figures 1 and 7. With one qualification the same can be said for the response of the listed variables to a persistent discount rate shock. Comparing figures 2 and 8 we see that the response of variables in the competitive economy is somewhat less persistent than in the noncompetitive economy.³⁴

In table 5 we report statistics from simulations of the competitive economy. Again the statistics are comparable to the ones reported for simulations of the one-price noncompetitive model in tables 3 and 4. We find that orders are more volatile than sales, but that inventory investment and sales can be negatively correlated.

³¹This absence of convergence is associated with fluctuations in orders, production and wages.

³²The 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 $\eta = 0$.

³³The display of variables is analogous to figures 1 through 6. We no longer display average sales per retailer, which are fixed at one, and number of shoppers, which are no longer relevant for the competitive equilibrium.

³⁴This 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 noncompetitive cases.

5. Conclusion

We have constructed a dynamic general equilibrium model to analyze the aggregate implications of (S, s) inventory policies. Key properties of our model are as follows. First, the model can replicate salient features of the business cycle. Second, it can account for observations that orders are more volatile than sales and inventory investment is positively correlated with sales. Third, the results suggest that while (S, s) inventory policies have little effect on the propagation and amplification of productivity disturbances, they contribute substantially to the amplification of discount rate shocks.

We have also shown that a competitive equilibrium version of the model closely mimics the noncompetitive economy described in the main part of the paper. 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 idiosyncratic risk for optimal price setting. We interpret this result as providing some justification for abstracting from microeconomic detail in future work in order to more fully address the aggregate implications of (S, s) inventory policies.

Currently we are using the competitive environment to study two key extensions of our analysis. First, explicit variation in the intervention points of the inventory policy over the business cycle is an important next step. Second, a comprehensive empirical evaluation of the aggregate implications of (S, s) inventory policies requires incorporating capital accumulation and sustained growth into the model.

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 consumer and producer durables, prices, employment, and cash balances. Presumably what we have learned here can be used to study these problems as well.

Appendices

A. Proofs

In this appendix we state and provide proofs for the propositions referred to in the main text.

Assumption 1. 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

$$\Delta V_x/R \leq f_1(p_x) = \frac{\Psi(p_x) - 1}{\Psi(p_x)} p_x \text{ and } p_x \geq p_l. \quad (\text{A.1})$$

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

$$\Delta V_x/R = f_2(p_x, V_{x-1}) = \frac{\theta \delta(p_x) p_x - (1 - R^{-1}) V_{x-1}}{\theta \delta(p_x) + R - 1}. \quad (\text{A.2})$$

We want to show that for any V_{x-1} there exist unique p_x and V_x which satisfy equations (A.1) and (A.2).

Suppose that $f_2(p, V) \leq f_1(p)$ for all p , then $p_x = p_l$ and V_x is determined from (A.2). Alternatively suppose that there exists a p such that $f_2(p, V) > f_1(p)$. If condition (a) from Assumption 1 applies then it follows that the function f_2 is bounded above,

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

Thus 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 (A.1) and (A.2). Since f_2 is decreasing in V , the increment ΔV_x is declining with V . \square

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 $\{\delta_x\}_{x=1}^{x^*}$, and $\theta_t \in [\underline{\theta}, \bar{\theta}]$, $0 < \underline{\theta} < \bar{\theta} < 1, \forall t \geq 0$.

Proposition 2. Given Assumption 2, from any initial distribution \bar{g}_0 , \bar{g}_t converges to the invariant distribution (2.3),

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

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 $\tilde{\Pi}_\tau = \Pi_{\tau+(x^*-1)} \cdots \Pi_{\tau+1} \Pi_\tau$, $\tau \geq 0$. It is easy to verify that for any $\tau \geq 0$, $\tilde{\Pi}_\tau$ 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 $\tilde{\Pi}_\tau$ is always greater than some $\varepsilon > 0$. Also, for any $t \geq 0$, the invariant distribution (2.3) 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). \square

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 $C_t = \sum_x \theta_t \delta_x \bar{g}_x = \theta_t (x^* / \sum_x \delta_x^{-1})$. 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 $Y_t = \theta_{t-1} \delta_1 \bar{g}_1 = \theta_{t-1} (x^* / \sum_x \delta_x^{-1})$. 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 $(\theta_{t-1} - \theta_t) (x^* / \sum_x \delta_x^{-1})$. Thus (i), (iii) and (iv) follow trivially. From the inventory identity

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

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

Proposition 4. *Let F be as defined in (3.19) and assume that retailers use the offer price strategy (3.20) and that the household consumes the market 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 $i, j = l, h$ ’.*

Proof. Let the retailer’s price strategy be characterized by $(\bar{\lambda}, x')$. 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

$$e = \bar{\lambda} (q - \underline{g}\eta). \quad (\text{A.3})$$

with $\underline{g} = \sum_{x > x'} \left(\frac{\bar{q}_x}{N} \right)$. Analogous to the discussion of the household’s optimization problem when the idiosyncratic utility shock takes on a continuum of values in $[u_l, u_h]$, 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 $\zeta_{n_{sj}}$ 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

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

Using the definition of the expenditure function (A.3) and the equilibrium condition that $\bar{\lambda} = \lambda = w U_c / U_n$, this condition simplifies to

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

If the household consumes market 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. \square

B. Algorithms for Computing Model Solutions

In this appendix we describe our approach to solving the general equilibrium model with and without aggregate uncertainty. We also describe a way to evaluate the accuracy of our approximations and our procedure for verifying that the equilibria computed in the case of aggregate uncertainty involve optimal decision making on the part of individual agents.

B.1. An Algorithm for Computing the Steady State

The algorithm exploits the recursive structure of the steady state and can potentially be used to prove the existence of a steady state equilibrium. The crucial point of the algorithm involves finding the optimal reservation utility policy, that is the scalar λ .

Let the optimal order size x^* be given. We will first define a mapping which defines a fixed point in λ . Conditional on λ and θ the capital value of a retailer can be calculated recursively from (2.1). This defines a function $V_{x^*}(\lambda, \theta)$. We need to find a match probability $\theta \in [0, 1]$ such that there are no profits from entry, $V_{x^*}(\lambda, \theta) = \gamma_0 + \gamma_1 x^*$.³⁵ Once we have found θ , we can use the associated optimal price policy $p_x(\lambda, \theta)$ and sales probabilities $\delta_x(\lambda, \theta)$ to calculate the implied price distribution $\bar{h}_x = \bar{g}_x/N$ from (2.3). We can now find the implied average quality q and cost e of a match as functions of λ and h_x . From the first order conditions (3.12) it follows that $\lambda = e/q$. This procedure then defines a fixed point in λ . Note that, conditional on the order size x^* , the equilibrium reservation utility policy and the distribution of retailers depends on preferences only via the distribution of idiosyncratic utility shocks F .

Given the optimal reservation utility function we can now solve for the total measure of retailers in a steady state. From the definition of the invariant distribution of inventories (2.3) it follows that the measure of retailers ordering in any period \hat{g}_{x^*} is proportional to N . This determines employment in the manufacturing sector as a function of N , $(\gamma_0 + \gamma_1 x^*) \hat{g}_{x^*} = z n_y^\alpha$. From the matching function (3.8) we get search effort n_s and the success rate for search ω as functions of N . The total measure of firms is now determined by the first order condition $wU_c = \lambda U_n$.

At this stage the allocation we have constructed satisfies all conditions for a steady state but one. We have not verified that the order size x^* is indeed optimal. To find an equilibrium we have to vary the order size until it is maximizing gains from entry.

B.2. Solving the Model with Aggregate Uncertainty

We assume that preferences are as described in section 3.3. To make the description of the algorithm as concrete as possible we describe how it is applied in one particular case, namely $(x^*, x') = (4, 0)$, $\zeta < \infty$. Modifications to solve the model for other cases are straightforward and should be clear. Our description is divided into two parts. First, we describe how to formulate the model's solution in terms of the fixed point of a particular system of functional equations. In the second part we briefly review the methods we apply to find an approximate fixed point.

B.2.1. The System of Functional Equations

Suppose we have candidates for the following functions: $n_s(g, \mathbf{z})$, $\hat{g}_{x^*}(g, \mathbf{z})$, $V_1(g, \mathbf{z})$ and $V_2(g, \mathbf{z})$. Here, $g = [g_1, g_2, g_3, g_4]'$ and $\mathbf{z} = z$ or β . For given (g, \mathbf{z}) , we can use these functions to compute θ , n_y , w , and λ by using equations (3.18), (3.7), two of the three equations in (3.12) and (3.8). Using the transition rules for g and \mathbf{z} along with these same equations and the given functions, we can compute future values of these endogenous variables. (In

³⁵Note that such a θ will only exist if we have chosen λ such that $V_{x^*}(\lambda, 1) \geq \gamma_0 + \gamma_1 x^*$.

fact, with the given functions we can compute current and future values of any endogenous variable). This makes it possible to compute $\tilde{Q}(g, \mathbf{z})$. We use the fact that we can solve for these endogenous variables given the assumed functions to form the desired system of functional equations as follows.

To begin, note that free entry in the retail sector implies that, for all (g, \mathbf{z}) ,

$$V_4(g, \mathbf{z}) = \gamma_0 + \gamma_1 4. \quad (\text{B.1})$$

We also know that the value of inventories after restocking must satisfy

$$V_4(g, \mathbf{z}) = \theta[\lambda u_l + E\{\tilde{Q}V_3(g', \mathbf{z}') \mid g, \mathbf{z}\}] + (1 - \theta)E\{\tilde{Q}V_4(g', \mathbf{z}') \mid g, \mathbf{z}\}. \quad (\text{B.2})$$

Combining (B.1) and (B.2) we can solve for

$$E\{\tilde{Q}V_3(g', \mathbf{z}') \mid g, \mathbf{z}\} = [\gamma_0 + \gamma_1 4] \frac{1 - (1 - \theta)R^{-1}}{\theta} - \lambda u_l. \quad (\text{B.3})$$

Using the version of (B.2) that applies to $V_3(g, \mathbf{z})$ and (B.3) we can compute the implicit function

$$V_3(g, \mathbf{z}) = \theta[\lambda u_l + E\{\tilde{Q}V_2(g', \mathbf{z}') \mid g, \mathbf{z}\}] + (1 - \theta) \left[[\gamma_0 + \gamma_1 4] \frac{1 - (1 - \theta)R^{-1}}{\theta} - \lambda u_l \right]. \quad (\text{B.4})$$

We can compute V_3 a different way. In particular, using the implicit function given in (B.4), we have

$$\hat{V}_3(g, \mathbf{z}) = \theta[\lambda u_l + E\{\tilde{Q}V_2(g', \mathbf{z}') \mid g, \mathbf{z}\}] + (1 - \theta)E\{\tilde{Q}V_3(g', \mathbf{z}') \mid g, \mathbf{z}\}. \quad (\text{B.5})$$

These two ways of computing V_3 should equal each other if the candidate functions equal their equilibrium counterparts. Thus, the first equation in the system is

$$V_3(g, \mathbf{z}) - \hat{V}_3(g, \mathbf{z}) = 0. \quad (\text{B.6})$$

The next two equations in the system are formed using versions of (B.5) for V_2 and V_1 . In particular, we have

$$\hat{V}_2(g, \mathbf{z}) = \theta[\lambda u_l + E\{\tilde{Q}V_1(g', \mathbf{z}') \mid g, \mathbf{z}\}] + (1 - \theta)E\{\tilde{Q}V_2(g', \mathbf{z}') \mid g, \mathbf{z}\} \quad (\text{B.7})$$

and

$$\hat{V}_1(g, \mathbf{z}) = \theta \lambda u_l + (1 - \theta)E\{\tilde{Q}V_1(g', \mathbf{z}') \mid g, \mathbf{z}\}. \quad (\text{B.8})$$

The next two equations in the system are then

$$V_2(g, \mathbf{z}) - \hat{V}_2(g, \mathbf{z}) = 0, \quad (\text{B.9})$$

and

$$V_1(g, \mathbf{z}) - \hat{V}_1(g, \mathbf{z}) = 0. \quad (\text{B.10})$$

The final equation in our system involves $n_s(g, \mathbf{z})$. We solve for this function directly

because of non-linearity in the system of first order conditions characterizing the household problem. It turns out to be more efficient computationally to solve for the extra function rather than to solve a non-linear system of equations for different values of (g, \mathbf{z}) to compute the endogenous variables needed to form (B.6), (B.9) and (B.10). The final equation in our system involves the first-order condition from the household problem not imposed when solving for the endogenous variables. We use this to arrive at

$$\hat{n}_s(g, \mathbf{z}) = \frac{wn_s(g, \mathbf{z})}{\lambda [n_y + n_s(g, \mathbf{z})]^\nu}.$$

So the final equation in our system is

$$n_s(g, \mathbf{z}) - \hat{n}_s(g, \mathbf{z}) = 0. \tag{B.11}$$

In general, a solution to the model consists of the four functions $n_s(g, \mathbf{z})$, $\hat{g}_{x^*}(g, \mathbf{z})$, $V_1(g, \mathbf{z})$ and $V_2(g, \mathbf{z})$ which solve the system of functional equations (B.6), (B.9), (B.10) and (B.11). An analytical solution to this problem is not available. We find approximate solutions numerically. Our strategy for doing this is the subject of the next sub-section.

B.2.2. A Galerkin-Chebyshev Algorithm

Our approximation strategy builds on methods emphasized by Judd (1992) (see also Christiano and Fisher (1994)). The reader is referred to these papers for a more detailed description of the methods and justifications for their use.

We seek the fixed point of a mapping of the form

$$\hat{f}(x) = H(f)(x), \tag{B.12}$$

where $f(x) = [n_s(x), \hat{g}_{x^*}(x), V_1(x), V_2(x)]'$, $x = (g, \mathbf{z})$, and H summarizes the operations implicit in forming (B.6), (B.9), (B.10) and (B.11).³⁶ Our approximation strategy begins by choosing a class of approximating functions for the elements of f . We worked with multidimensional Chebyshev polynomials. The Chebyshev polynomials in a single variable are defined recursively as follows: $T_0(y) \equiv 1$, $T_1(y) \equiv y$, and $T_i(y) = 2yT_{i-1}(y) - T_{i-2}(y)$, for $i \geq 2$. The domain and range of these polynomials is $(-1, 1)$. The exact form of our approximating functions is

$$f_{l,a}(x) = \sum_{i=1}^{21} a_{l,i} T_i^*(\varphi(x)), \quad l = 1, \dots, 4,$$

where $\varphi(x) = [\varphi_1(x_1), \varphi_2(x_2), \varphi_3(x_3), \varphi_4(x_4), \varphi_5(x_5)]'$, with x_i , $i = 1, \dots, 5$ the elements of x , and $\varphi_i(\cdot)$, $i = 1, \dots, 5$ functions that map into the interval $(-1, 1)$, and T_i^* , $i = 1, \dots, 21$

³⁶Unless otherwise indicated, the remaining notation is independent of the main text.

are a list of the elements from the set of functions

$$C_n^{(q)} = \left\{ T_{i_1}(x_1)T_{i_2}(x_2) \cdots T_{i_q}(x_q) \mid \sum_{j=1}^q i_j \leq n-1, i_1, i_2, \dots, i_q \geq 0 \right\},$$

where $q = 5$ and $n = 3$. Denote the fixed point of (B.12) by f^* . We approximate f^* by $f_a = [f_{1,a}, \dots, f_{4,a}]'$.

The Galerkin-Chebyshev procedure we adopted involves solving an *approximation problem* of the form

$$AR_l(X; a) = 0, \quad l = 1, \dots, 4. \quad (\text{B.13})$$

Here $X = [X_1, \dots, X_M]'$ is an $M \times 5$ matrix containing the elements of a grid over the state variables and $a = (a_{1,1}, \dots, a_{1,21}, \dots, a_{4,21})$ is the vector of undetermined coefficients. The term $R_l(X; a) = [r_l(X_1; a), r_l(X_2; a), \dots, r_l(X_M; a)]'$, $l = 1, \dots, 4$, is an $M \times 1$ vector of *functional equation residuals* constructed as follows

$$r_l(X_i; a) = f_{l,a}(X_i) - [H(f_a)(X_i)]_l, \quad i = 1, \dots, M,$$

where the operator $[y]_l$ picks out the l 'th element of the vector y . Finally, the term A is a $21 \times M$ matrix. The equations in (B.13) represent a system of $4 \cdot 21 = 84$ equations in the 84 undetermined coefficients summarized by the vector a . The approximation problem is to find the vector a^* that solves (B.13). An approximate solution to the model is then given by f_{a^*} .

One way to think about the approximation problem (B.13) is that it involves choosing a to set weighted averages of the functional equation residuals to zero, where the weights are determined by the elements of A and the points at which to evaluate the functional equation residuals are determined by the choice of X . What makes our approximation method a Galerkin-Chebyshev procedure is the way in which we select the grid X , and the weighting matrix A .

To construct the grid we first computed the m zeros of T_m for some $m > n - 1$. The zeros of T_m are given by

$$z_k = \cos\left(\frac{(2k-1)\pi}{2m}\right), \quad k = 1, \dots, m.$$

The grid over element i of x is then given by $\varphi_i^{-1}(z_k)$, $k = 1, \dots, m$, for $i = 1, \dots, 5$. The grid over the state space constructed using this procedure has $M = m^5$ points. These grid points make up the matrix X . The matrix A is given as follows:

$$A = \begin{bmatrix} T_1^*(\varphi(X_1)) & T_1^*(\varphi(X_2)) & \cdots & T_1^*(\varphi(X_M)) \\ T_2^*(\varphi(X_1)) & T_2^*(\varphi(X_2)) & \cdots & T_2^*(\varphi(X_M)) \\ \vdots & \vdots & \ddots & \vdots \\ T_{21}^*(\varphi(X_1)) & T_{21}^*(\varphi(X_2)) & \cdots & T_{21}^*(\varphi(X_M)) \end{bmatrix}.$$

By the orthogonality property of Chebyshev polynomials (see Christiano and Fisher (1994)), the rows of this matrix are orthogonal.

For all the versions of the model considered in section 4, including the version emphasized in the previous sub-section, we chose $m = 3$. We kept $n = 3$ throughout as well. For starting values we chose a so that the approximating functions were constants equal to the steady state values of the relevant variables. The functions φ_i were specified as follows:

$$\varphi_i(x_i) = 2 \frac{x_i - \underline{x}_i}{\bar{x}_i - \underline{x}_i} - 1, \quad i = 1, \dots, q,$$

where $(\underline{x}_i, \bar{x}_i)$ are bounds on the ergodic set of state variable i . These were chosen as (1 ± 0.025) times the steady state value of the state variable in question for all the versions of the model considered in section 4.

The algorithm is accurate and fast. We programmed in GAUSS and worked on a Gateway 486/66 machine. For the $(x^*, x') = (1, 0)$ versions of the model solutions are computed in less than a second.³⁷ For the $(x^*, x') = (2, 0)$ and $(x^*, x') = (2, 1)$ versions of the model solutions are computed in about 2 and half seconds. Computation times increase substantially in the versions of the model with $x^* = 4$, but remain under 15 minutes.

To gauge the accuracy of our approximations, we experimented with higher values for m and by adding terms to the approximating polynomials. We found that our results as summarized by the tables and figures were not changed qualitatively. In fact the quantitative implications of these perturbations in the approximation parameters are only marginal.

A measure of the accuracy of our approximations is the vector of maximum absolute values of the functional equation residuals at a^* , computed for a subset of the state space local to steady state (see Christiano and Fisher (1994)). Using the notation of section B.2.2 the exact measure is defined

$$\mathcal{A} = 100 \cdot \max(|\hat{f}_{a^*} - H(\hat{f}_{a^*})| ./ \hat{f}_{a^*}),$$

where $\hat{f}_{a^*} = [f'_{a^*}, (V_{x^*} - (\gamma_0 + \gamma_1 x^*))']'$ and $./$ is element by element division, and the max operator chooses the largest of the functional equation residuals over the specified region of the state space so that \mathcal{A} is a nonnegative scalar.

Table B. Accuracy Checks

Model (x^*, x')	Impulse of		\mathcal{A}	
	Most	Least	Most	Least
(1, 0)	$\beta, \rho = 0.9$	$z, \rho = 0$	3.2×10^{-8}	4.5×10^{-8}
(2, 0)	$\beta, \rho = 0$	$z, \rho = 0.9$	1.2×10^{-5}	1.2×10^{-5}
(2, 1)	$\beta, \rho = 0$	$z, \rho = 0.9$	1.5×10^{-5}	1.5×10^{-5}
(4, 0)	$z, \rho = 0.9$	$\beta, \rho = 0.9$	1.6×10^{-5}	1.6×10^{-5}
(4, 2)	$\beta, \rho = 0$	$z, \rho = 0.9$	1.6×10^{-5}	1.6×10^{-5}

Note: "Most" indicates the most accurate case and "Least" indicates the least accurate case.

³⁷The notation in the remainder of this sub-section is consistent with the main text.

In table B we report \mathcal{A} 's for the most accurate (smallest maximum absolute value) and least accurate approximations for the $\zeta < \infty$ cases examined in the main text. Notice that \mathcal{A} never exceeds $1.6 \times 10^{-5}\%$ of the function approximations. Differences in accuracy for a given (x^*, x') combination seem extremely small. Using this accuracy measure we find the $1/\zeta = 0$ cases are uniformly more accurate than the corresponding noncompetitive examples.

B.3. Procedure for Verifying Optimality

The algorithm outlined in the previous section is based on a fixed (x^*, x') .³⁸ After an approximate solution has been computed it is necessary to verify that these values are optimal from the perspective of a retailer for realizations of the exogenous shock. In this section we describe our verification procedure.

First, consider the optimality of the ordering policy. We examine whether one period deviations from the $(x^*, 0)$ inventory policy in any given period of a simulation are worthwhile for a representative retailer. As described in the main text, if the capital value of inventories function is “concave” in the level of inventories then ordering at any $x > 0$ is never optimal. A direct check of the concavity of the value function once a candidate approximate model solution has been found verifies that ordering when $x > 0$ is never optimal.

To verify the optimality of the order size we check along a simulation path that the following condition is never violated:

$$V_{x^*+1}(g, \mathbf{z}) - [\gamma_0 + \gamma_1(x^* + 1)] < 0 \quad (\text{B.14})$$

For the versions of the model in which $x^* > 1$, we also check that the condition

$$V_{x^*-1}(g, \mathbf{z}) - [\gamma_0 + \gamma_1(x^* - 1)] < 0 \quad (\text{B.15})$$

is never violated. Once a candidate model solution has been found, an approximation to V_{x^*-1} is direct using the appropriate version of (B.4). An approximation to V_{x^*+1} can also be computed given an approximate model solution. For example, in the cases where $x' = 0$, this requires solving the functional equation

$$V_{x^*+1}(g, \mathbf{z}) = \theta[\lambda u_l + E\{\tilde{Q}V_{x^*}(g', \mathbf{z}') \mid g, \mathbf{z}\}] + (1 - \theta)E\{\tilde{Q}V_{x^*+1}(g', \mathbf{z}') \mid g, \mathbf{z}\}.$$

This can be accomplished by applying a version of the procedure outlined in the previous section, conditional on the candidate model solution. For the simulations underlying the results reported in section 4, conditions (B.14) and (B.15) were never violated or within the margin of error indicated by the accuracy checks in table B.

We now consider the verification procedure to ensure the optimality of an assumed x' pricing scheme. Our approach involves considering one period deviations from the assumed x' pricing scheme. We make use of the fact that the posted price is non-increasing in the inventory stock. In the $x' = 0$ versions of the model this means we need only consider whether firms with inventory level $x = 1$ have an incentive to raise their price to λu_h . To

³⁸The notation in this section follows that in the main text.

verify that this is not the case, we compute the value of this one period perturbation in the pricing strategy

$$V_1^p(g, \mathbf{z}) = (\theta/2)\lambda u_h + (1 - \theta/2)E\{\tilde{Q}V_1(g', \mathbf{z}') \mid g, \mathbf{z}\}.$$

This can be computed directly with an approximate solution in hand. Our verification procedure involves checking along a given simulation path that the condition

$$V_1^p(g, \mathbf{z}) - V_1(g, \mathbf{z}) < 0 \tag{B.16}$$

is never violated.

In the $x' > 0$ versions of the model we must check three conditions. First, we must verify that retailers with $x = 1$ do not have an incentive to raise their price to a level that ensures they make no sale in the current period. This could be optimal if the expected value of selling one unit of the good in the next period is exceptionally high. To verify that this never occurs we compute the value of this one period perturbation in the policy,

$$V_1^p(g, \mathbf{z}) = E\{\tilde{Q}V_1(g', \mathbf{z}') \mid g, \mathbf{z}\}$$

and check that the condition

$$V_1^p(g, \mathbf{z}) - V_1(g, \mathbf{z}) < 0 \tag{B.17}$$

is never violated. Second, we must verify that retailers with $x = x'$ do not have an incentive to lower their price in any period. The condition we check in this case is

$$V_{x'}^p(g, \mathbf{z}) - V_{x'}(g, \mathbf{z}) < 0, \tag{B.18}$$

where

$$V_{x'}^p(g, \mathbf{z}) = \theta[\lambda u_l + E\{\tilde{Q}V_{x'-1}(g', \mathbf{z}') \mid g, \mathbf{z}\}] + (1 - \theta)E\{\tilde{Q}V_{x'}(g', \mathbf{z}') \mid g, \mathbf{z}\}.$$

Finally, we must verify that retailers with $x = x' + 1$ do not have an incentive to raise their price in any period. The condition we check in this case is

$$V_{x'+1}^p(g, \mathbf{z}) - V_{x'+1}(g, \mathbf{z}) < 0, \tag{B.19}$$

where

$$V_{x'+1}^p(g, \mathbf{z}) = (\theta/2)[\lambda u_h + E\{\tilde{Q}V_{x'}(g', \mathbf{z}') \mid g, \mathbf{z}\}] + (1 - \theta/2)E\{\tilde{Q}V_{x'+1}(g', \mathbf{z}') \mid g, \mathbf{z}\}.$$

In none of the twenty parameterizations of the model examined in section 4 were any of (B.16), (B.17), (B.18), or (B.19) ever violated. The approximation to the violations appears good since the smallest violation over all the simulations is outside the margin of approximation error found using \mathcal{A} .

C. Specification of Idiosyncratic Preference Shocks

In section 4 we refer to a formula that can be used to compute the support of the idiosyncratic preference shocks for the various experiments involving $\zeta < \infty$. This formula is derived as follows. Ignore aggregate uncertainty and assume $p_x = \lambda u_h$ for $x \leq x'$. Then compute $v_{x'+1}$ under the two alternatives for the price. Equate these values and solve for η :

$$f(x'; \theta, \beta) = \frac{1 - \left(\frac{\beta\theta}{2-(2-\theta)\beta}\right)^{x'+1}}{3 \left[1 + \left(\frac{\beta\theta}{2-(2-\theta)\beta}\right)^{x'+1}\right]}.$$

This formula gives an exact upper bound on the value that η can take to ensure the pricing scheme summarized by (x^*, x') is optimal in a steady state where the probability of a match as seen by a retail firm equals θ . We use this formula to compute u_l and u_h in a way that ensures the optimality of a (x^*, x') pricing scheme even in the case of aggregate uncertainty.

Notice that according to the formula, η is decreasing in θ .³⁹ If we select some $\bar{\theta}$ that is much larger than any realization of θ in the aggregate uncertainty version of the model, then the assumed x' pricing scheme is likely to be optimal when $\eta = f(x'; \bar{\theta}, \beta)$ along any given simulation of a version of the aggregate uncertainty model. In practice this procedure is not enough to guarantee optimality of a given pricing scheme in our experiments. Our exact procedure is as follows. Throughout we fix $\bar{\theta} = 1$. In the $(x^*, x') = (1, 0)$, $(2, 0)$, $(4, 0)$ versions of the model we set $\eta = 2f(0; \bar{\theta}, \beta)/3$. In the $(x^*, x') = (2, 1)$ we set $\eta = [f(x'; \bar{\theta}, \beta) + f(x' - 1; \bar{\theta}, \beta)]/2$, and in the $(x^*, x') = (4, 2)$ experiments we set $\eta = 2f(x'; \bar{\theta}, \beta)/3 + f(x' - 1; \bar{\theta}, \beta)/3$. For each experiment, we check *ex post* that the hypothesised pricing scheme is in fact optimal along a given simulation. Our procedure for doing this is outlined in appendix B.

³⁹According to the formula $f(x'; \cdot)$ is increasing in x' . Thus the higher is the inventory level that is consistent with offering the high price, the more spread that is possible in the idiosyncratic shocks.

References

- [1] Aguirregabiria, Victor. 1994. A Dynamic Structural Model of Price and Inventory Decisions with Fixed Adjustment Costs. University of Western Ontario mimeo.
- [2] Bental, Benjamin and Eden, Benjamin. 1993. Inventories in a Competitive Environment. *Journal of Political Economy* 101, 863-86.
- [3] Benhabib, Jess, Rogerson, Richard and Wright, Randall. 1991. Homework in Macroeconomics: Household Production and Aggregate Fluctuations. *Journal of Political Economy* 99, 1166-87.
- [4] Bertola, Giuseppe, and Caballero, Ricardo J. 1990. Kinked Adjustment Costs and Aggregate Dynamics. In *NBER Macroeconomics Annual 1990*, eds. Olivier J. Blanchard and Stanley Fischer, 237-88. Cambridge MA: MIT Press.
- [5] 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.
- [6] 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.
- [7] Caballero, Ricardo J. 1992. A Fallacy of Composition. *American Economic Review* 82, 1279-92.
- [8] Caballero, Ricardo J. and Engel, Eduardo M.R.A. 1991. Dynamic (S, s) Economies. *Econometrica* 59, 1659-86.
- [9] Caplin, Andrew. 1985. The Variability of Aggregate Demand with (S, s) Inventory Policies. *Econometrica* 53, 1395-1410.
- [10] Christiano, Lawrence J. 1988. Why Do Inventories Fluctuate So Much? *Journal of Monetary Economics* 21, 247-80.
- [11] 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.
- [12] 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.
- [13] Greenwood, Jeremy, Hercowitz, Zvi, and Huffman, Gregory W. 1988. Investment, Capacity Utilization, and the Real Business Cycle. *American Economic Review* 78: 402-17.

- [14] Hansen, Gary. 1985. Indivisible Labor and the Business Cycle. *Journal of Monetary Economics* 16, 309-27.
- [15] 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.
- [16] Judd, Kenneth L. 1992. Projection Methods for Solving Aggregate Growth Models. *Journal of Economic Theory* 58: 410-52.
- [17] Juster, T. and Stafford, Frank. 1991. The Allocation of Time: Empirical Findings, Behavioral Models, and Problems of Measurement. *Journal of Economic Literature* 29, 471-522.
- [18] Kydland, Finn, and Prescott, Edward C. 1982. Time to Build and Aggregate Fluctuations. *Econometrica* 50, 1345-70.
- [19] Lovell, Michael C. 1993. Simulating the Inventory Cycle. *Journal of Economic Behavior and Organization* 21, 147-79.
- [20] Rogerson, Richard. 1987. An Equilibrium Model of Sectoral Reallocation. *Journal of Political Economy* 95, 824-34.
- [21] 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.
- [22] Sobel, Joel, and Takahashi, Ichiro. 1983. A multistage model of bargaining. *Review of Economic Studies* 50, 411-26.
- [23] 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.

(x^*, x')	γ_0	η	v	ζ	α	β	φ	z	γ_1	μ	ρ	σ
(1,0)	0	0.135	0.2	701.3	0.64	$1.05^{-1/12}$	0.85	1	1.175	0.25	0.9	0.01
(2,0)	0.010	0.135	0.2	699.6	0.64	$1.05^{-1/12}$	0.85	1	1.175	0.25	0.9	0.01
(2,1)	0.015	0.303	0.2	933.9	0.64	$1.05^{-1/12}$	0.85	1	1.175	0.25	0.9	0.01
(4,0)	0.046	0.135	0.2	697.1	0.64	$1.05^{-1/12}$	0.85	1	1.175	0.25	0.9	0.01
(4,2)	0.055	0.538	0.2	524.0	0.64	$1.05^{-1/12}$	0.85	1	1.175	0.25	0.9	0.01

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

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

Statistic	$(x^*, x') =$				
	(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 Manufacturing Technology Shocks.

Statistic	$\rho = 0.9$ and $(x^*, x') =$					$\rho = 0$ and $(x^*, x') =$				
	(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	1.82	1.14	1.25	0.77	0.78
σ_Y/σ_C	1.02	1.03	1.09	1.05	1.14	1.22	1.45	1.65	1.89	2.06
$\sigma_{\Delta I}$	0.10	0.23	0.24	0.34	0.34	0.30	0.72	0.48	0.61	0.59
$\rho(C, Y)$	0.99	0.97	0.97	0.93	0.94	0.98	0.73	0.95	0.53	0.69
$\rho(C, \Delta I)$	0.18	0.02	0.24	-0.04	0.15	0.59	0.06	0.73	0.01	0.28
$\rho(C, C/N)$	-0.99	-0.99	-0.13	-0.99	-0.06	-0.98	-0.92	0.64	-0.82	0.53
$\rho(C, N)$	1.00	1.00	0.99	1.00	0.98	1.00	1.00	0.86	0.99	0.87
$\rho(C, S)$	1.00	0.99	0.98	1.00	0.96	0.99	0.99	0.94	0.99	0.92
$\rho(C, E)$	0.18	0.12	0.25	0.08	0.26	0.61	0.33	0.74	0.19	0.48
$\rho(C_t, C_{t-1})$	0.93	0.97	0.97	0.99	0.99	0.25	0.76	0.31	0.92	0.77
$\rho(Y_t, Y_{t-1})$	0.90	0.94	0.92	0.95	0.94	0.03	0.23	0.03	0.10	0.17
$\rho(\Delta I_t, \Delta I_{t-1})$	0.14	0.43	0.68	0.69	0.73	-0.40	-0.25	-0.10	-0.09	-0.14

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

Table 4. Statistics Implied by Discount Rate Shocks.

Statistic	$\rho = 0.9$ and $(x^*, x') =$					$\rho = 0$ and $(x^*, x') =$				
	(1, 0)	(2, 0)	(2, 1)	(4, 0)	(4, 2)	(1, 0)	(2, 0)	(2, 1)	(4, 0)	(4, 2)
σ_C	0.19	0.72	1.07	1.41	1.66	0.12	0.20	0.21	0.17	0.17
σ_Y/σ_C	1.03	1.03	1.08	1.05	1.13	1.28	1.44	1.73	1.85	2.02
$\sigma_{\Delta I}$	0.01	0.09	0.15	0.28	0.35	0.02	0.11	0.09	0.14	0.14
$\rho(C, Y)$	0.99	0.97	0.97	0.92	0.93	0.96	0.66	0.92	0.41	0.59
$\rho(C, \Delta I)$	0.15	-0.01	0.20	-0.08	0.11	0.53	-0.05	0.66	-0.14	0.13
$\rho(C, C/N)$	-0.99	-0.97	-0.46	-0.92	-0.34	-0.96	-0.66	0.51	-0.41	0.42
$\rho(C, N)$	1.00	1.00	0.99	0.99	0.99	0.99	0.99	0.89	0.99	0.89
$\rho(C, S)$	0.99	0.99	0.98	0.98	0.97	0.98	0.90	0.93	0.82	0.88
$\rho(C, E)$	0.17	0.10	0.22	0.05	0.23	0.57	0.27	0.67	0.08	0.38
$\rho(C_t, C_{t-1})$	0.94	0.98	0.97	0.99	0.99	0.32	0.77	0.42	0.93	0.80
$\rho(Y_t, Y_{t-1})$	0.90	0.95	0.92	0.95	0.95	0.05	0.24	0.05	0.11	0.19
$\rho(\Delta I_t, \Delta I_{t-1})$	0.14	0.44	0.67	0.69	0.73	-0.40	-0.24	-0.12	-0.08	-0.14

Note: Volatility statistics are presented as a proportion of the standard deviation of the discount rate.

Table 5. Statistics Implied by the Competitive Model ($\rho = 0.9$)

Statistic	z shocks, $x^* =$			β shocks, $x^* =$		
	1	2	4	1	2	4
σ_C	2.14	2.06	1.91	0.00	0.52	1.23
σ_Y/σ_C	1.00	1.02	1.04	-	1.02	1.04
$\sigma_{\Delta I}$	0.00	0.18	0.30	0.00	0.05	0.19
$\rho(C, Y)$	1.00	0.99	0.95	-	0.99	0.95
$\rho(C, \Delta I)$	-	0.01	-0.04	-	0.01	-0.04
$\rho(C, E)$	0.23	0.14	0.10	-	0.14	0.10
$\rho(C_t, C_{t-1})$	0.90	0.96	0.98	-	0.96	0.98
$\rho(Y_t, Y_{t-1})$	0.90	0.94	0.95	-	0.94	0.95
$\rho(\Delta I_t, \Delta I_{t-1})$	-	0.22	0.63	-	0.22	0.63

Note: Volatility statistics are presented as a proportion of the standard deviation of the technology and discount rate shocks.

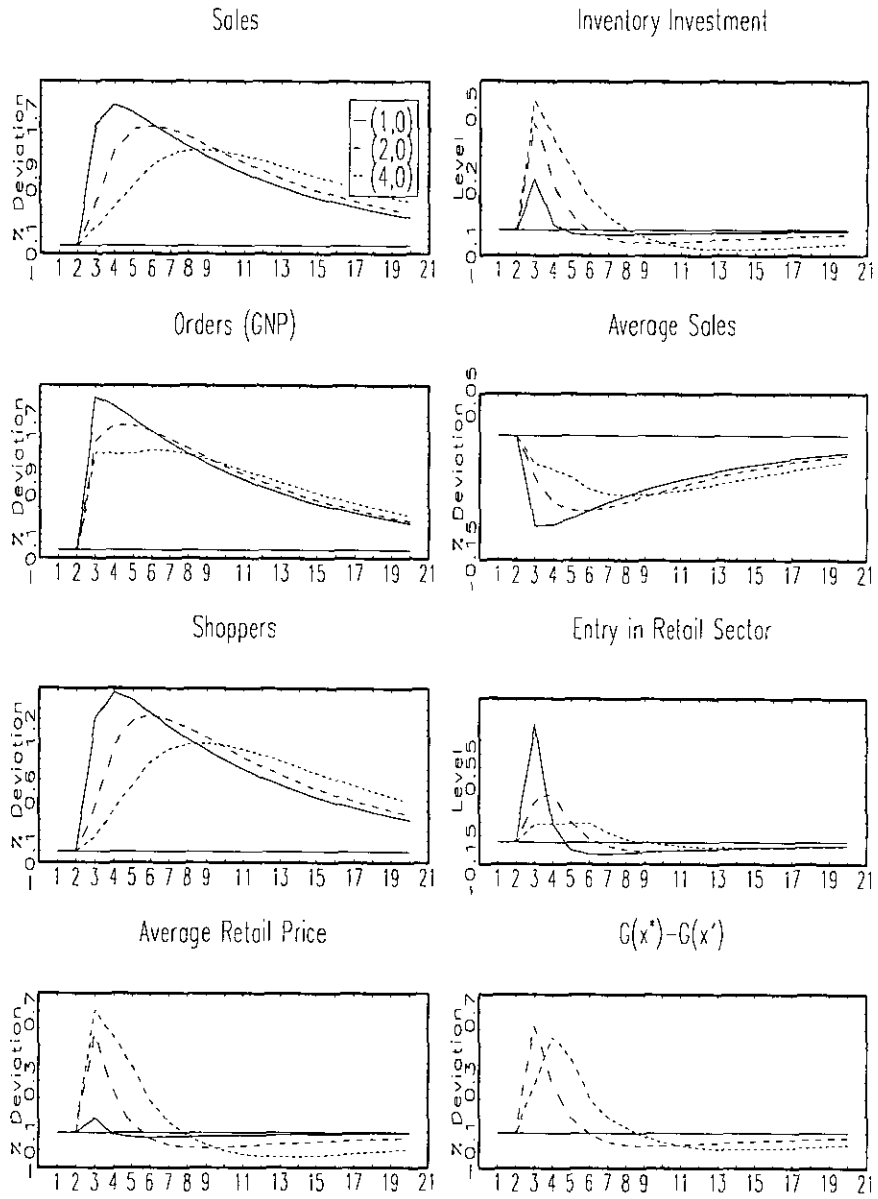


Figure 1: One-Price Impulse Responses to Technology Shocks, $\rho = 0.9$

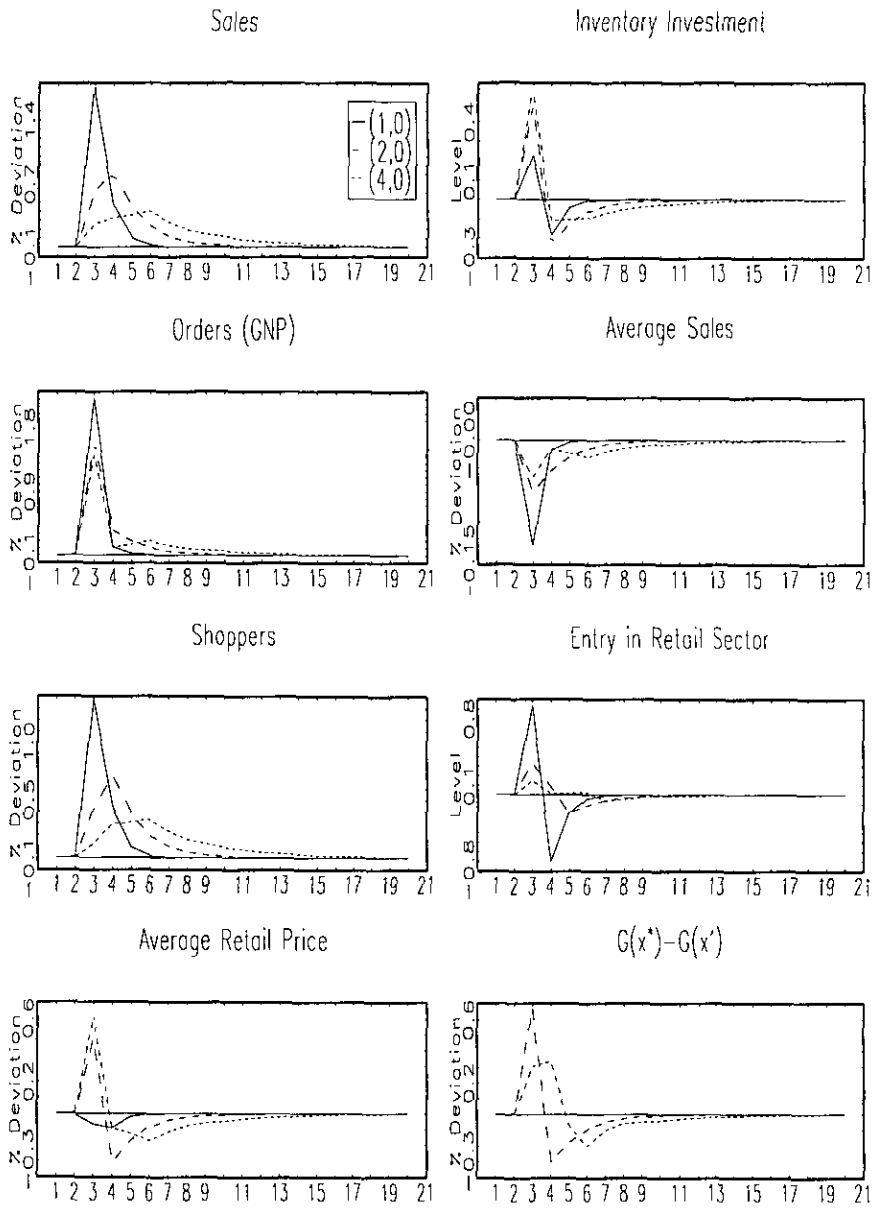


Figure 2: One-Price Impulse Responses to Technology Shocks, $\rho = 0$

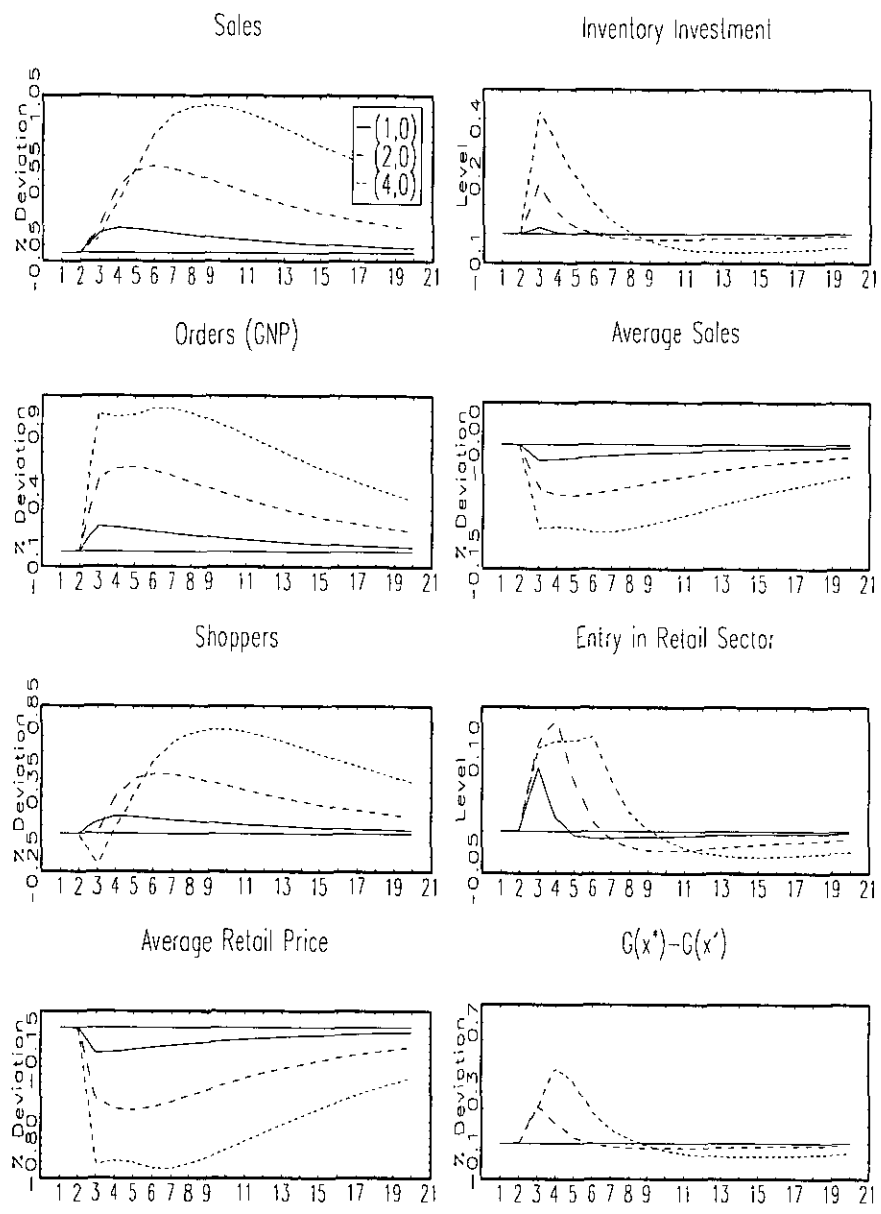


Figure 3: One-Price Impulse Responses to Discount Rate Shocks, $\rho = 0.9$

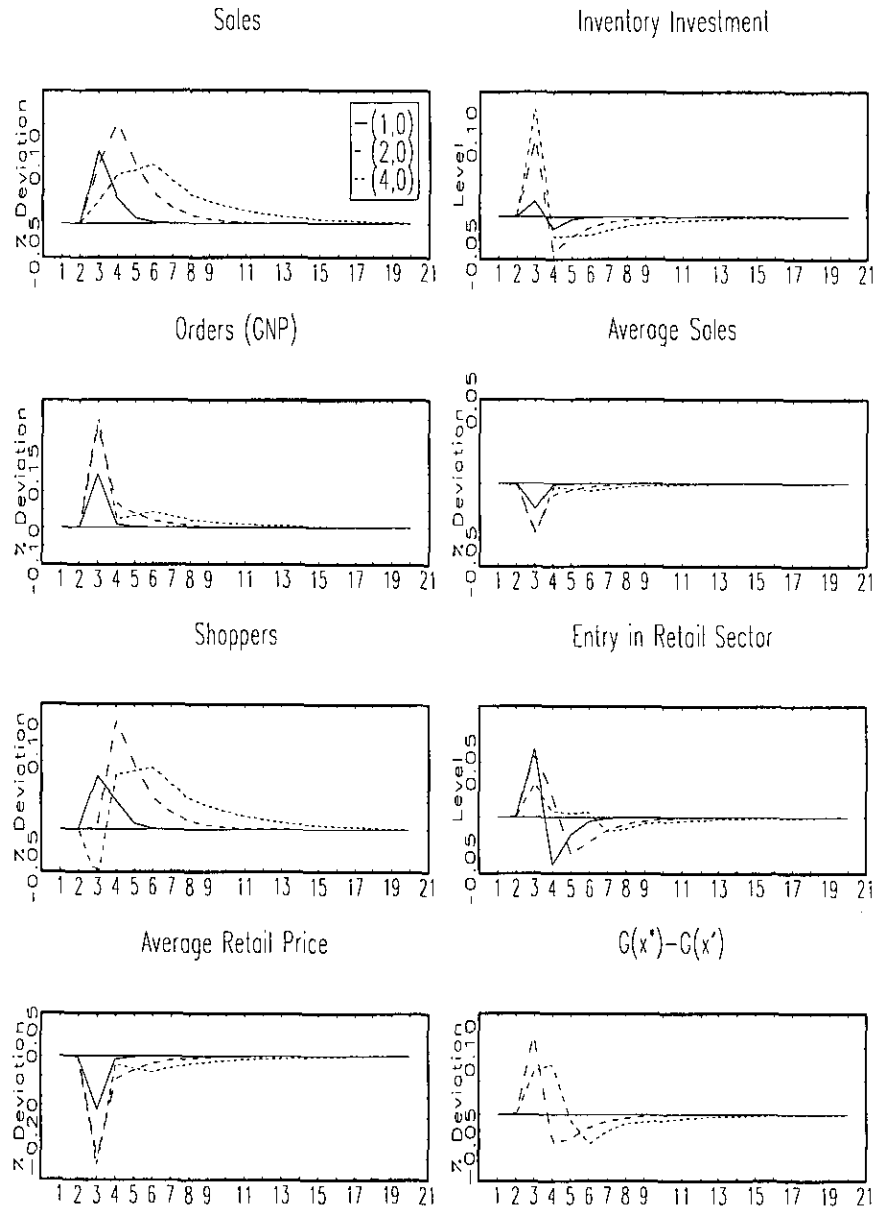


Figure 4: One-Price Impulse Responses to Discount Rate Shocks, $\rho = 0$

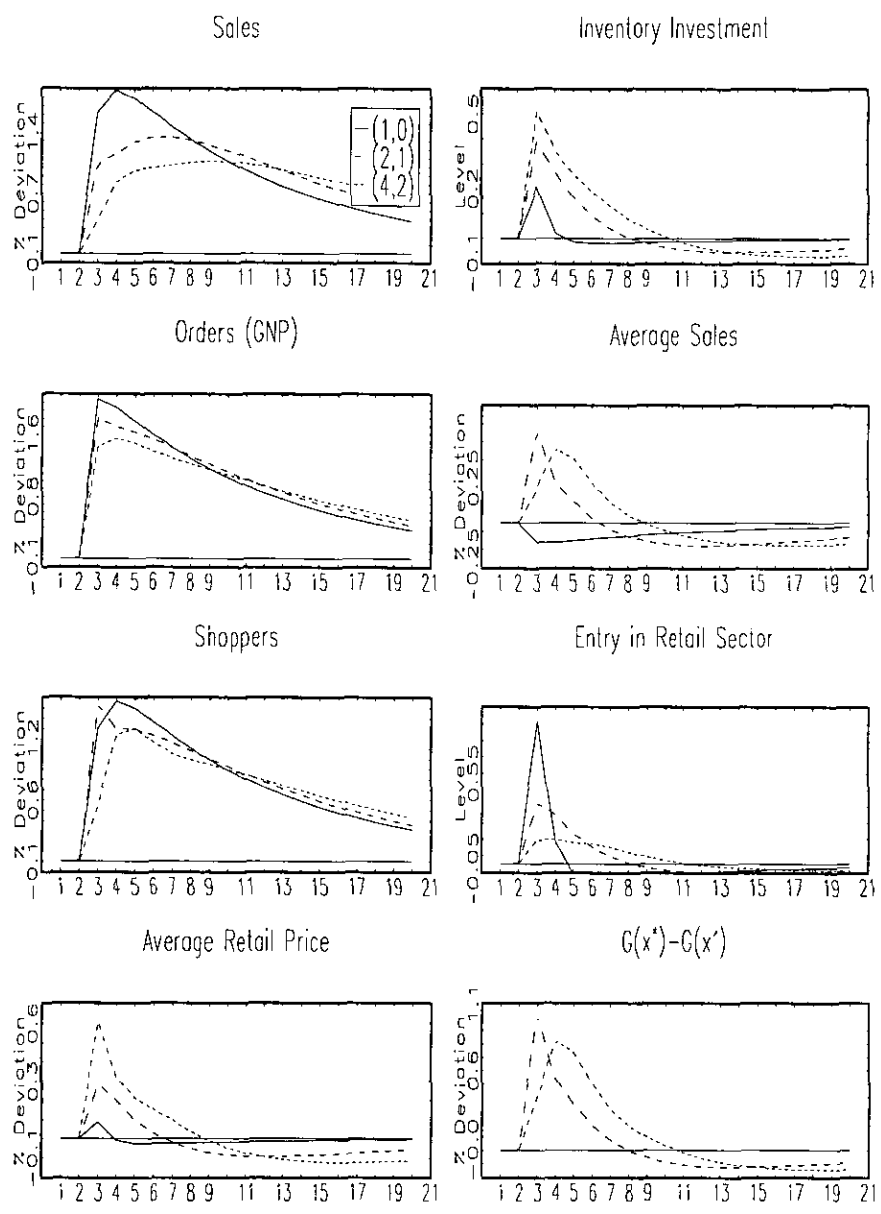


Figure 5: Two-Price Impulse Responses to Technology Shocks, $\rho = 0.9$

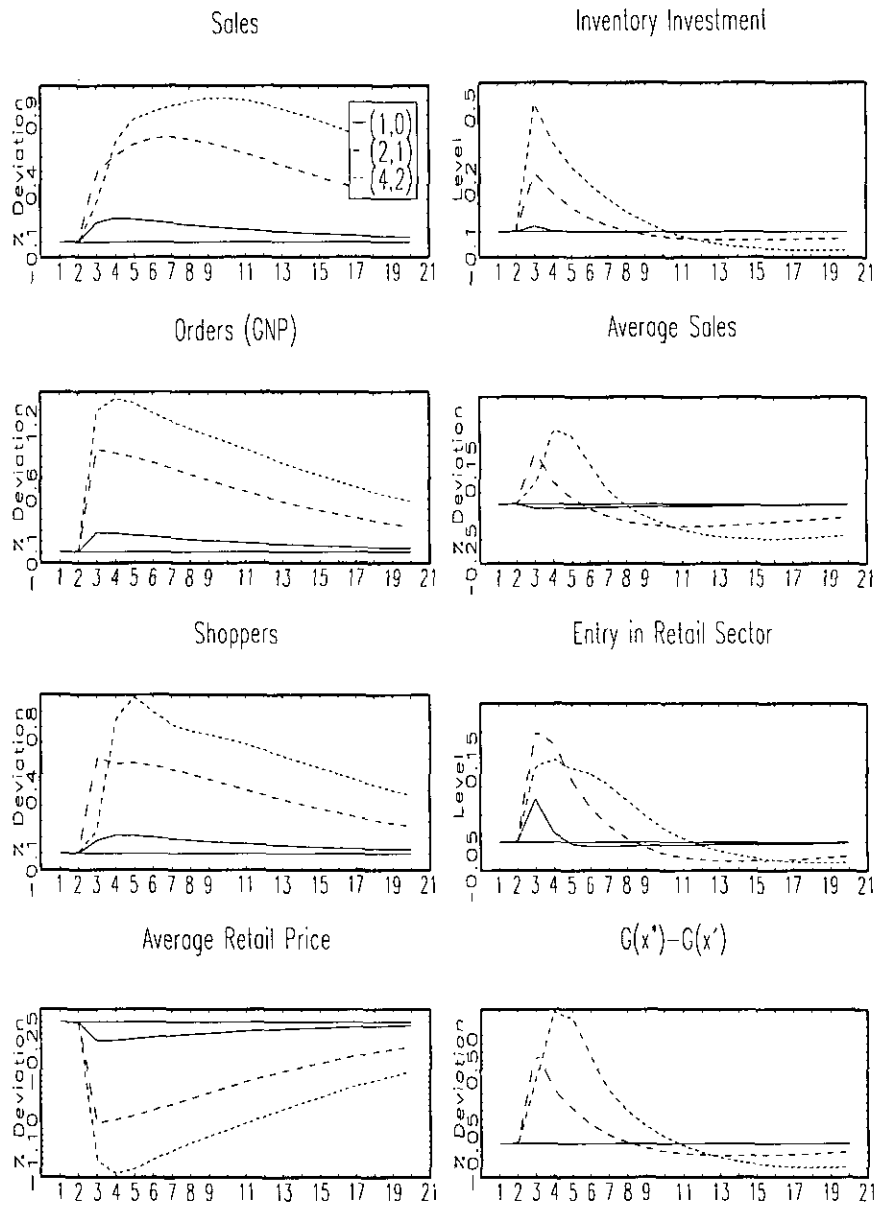


Figure 6: Two-Price Impulse Responses to Discount Rate Shocks, $\rho = 0.9$

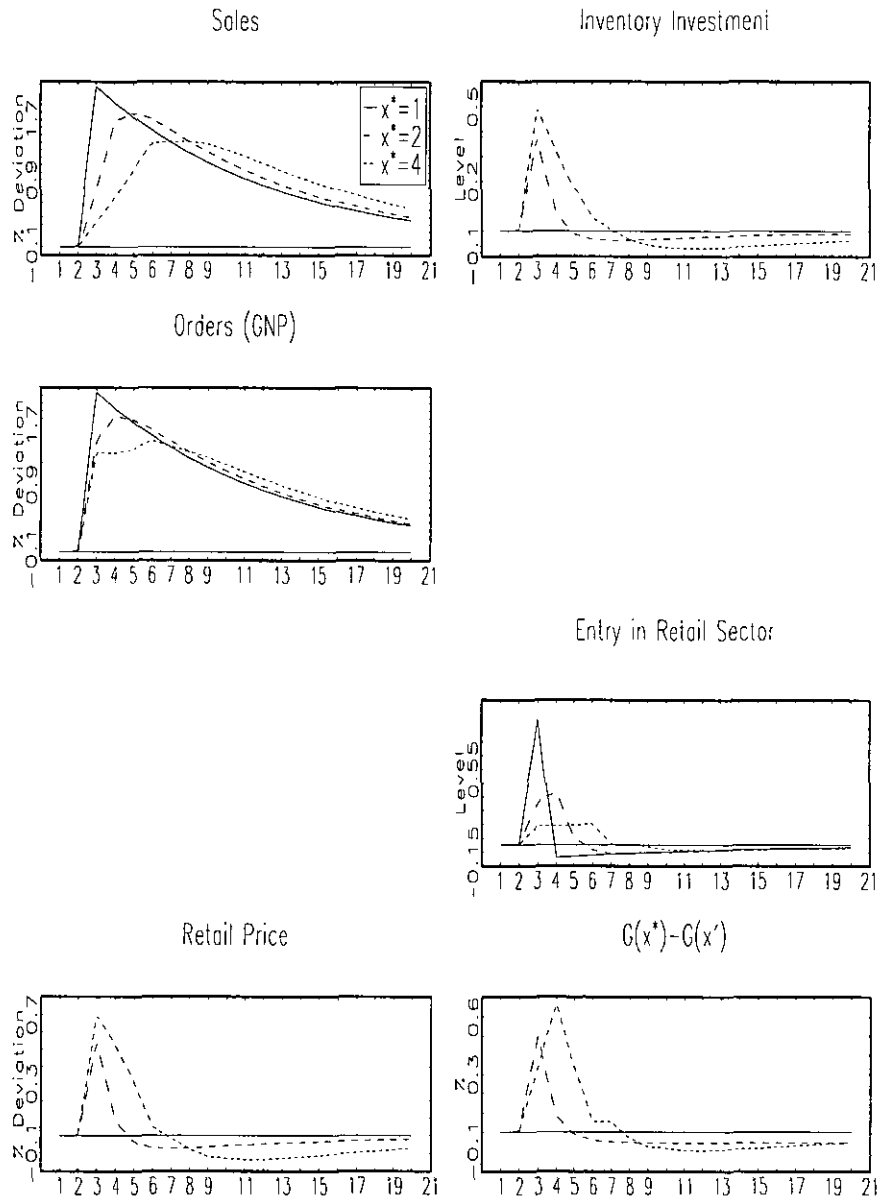


Figure 7: Responses to Technology Shocks in the Competitive Model, $\rho = 0.9$

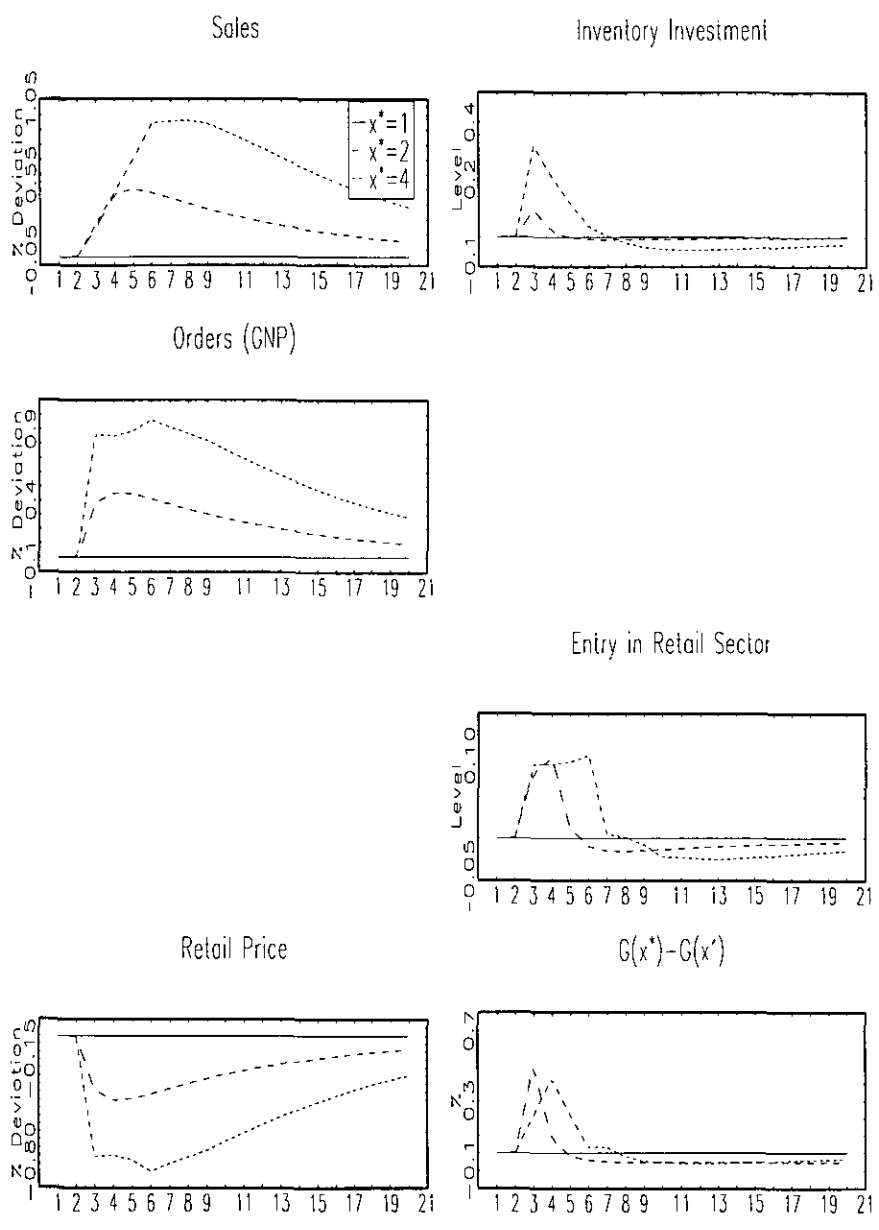


Figure 8: Responses to Discount Rate Shocks in the Competitive Model, $\rho = 0.9$