
BIS WORKING PAPERS

No. 63 – March 1999

**MICROECONOMIC INVENTORY ADJUSTMENT AND
AGGREGATE DYNAMICS**

by

Jonathan McCarthy and Egon Zakrajšek

**BANK FOR INTERNATIONAL SETTLEMENTS
Monetary and Economic Department
Basle, Switzerland**

BIS Working Papers are written by members of the Monetary and Economic Department of the Bank for International Settlements, and from time to time by other economists, and are published by the Bank. The papers are on subjects of topical interest and are technical in character. The views expressed in them are those of their authors and not necessarily the views of the BIS.

Copies of publications are available from:

Bank for International Settlements
Information, Press & Library Services
CH-4002 Basle, Switzerland

Fax: +41 61 / 280 91 00 and +41 61 / 280 81 00

This publication is available on the BIS website (www.bis.org).

© *Bank for International Settlements 1999.*

All rights reserved. Brief excerpts may be reproduced or translated provided the source is stated.

ISSN 1020-0959

**MICROECONOMIC INVENTORY ADJUSTMENT AND
AGGREGATE DYNAMICS**

by

Jonathan McCarthy and Egon Zakrajšek*

Abstract

We examine microeconomic and aggregate inventory dynamics in the business sector of the U.S. economy. We employ high-frequency firm-level data and use an empirically tractable model, in which the aggregate dynamics are derived explicitly from the underlying microeconomic data. Our results show that the microeconomic adjustment function in both the manufacturing and trade sectors is nonlinear and asymmetric, results consistent with firms using (S,s) -type inventory policies. There are differences in the estimated adjustment functions between the two sectors as well as the durable and nondurable goods firms within each sector. The estimated adjustment function is remarkably stable across subperiods, indicating little change in the inventory adjustment process over time. As predicted by our model, higher moments of the cross-sectional distribution of inventory deviations affect aggregate inventory dynamics.

* Both authors are affiliated with the Research Department, Federal Reserve Bank of New York. We thank Palle Andersen, Andy Caplin, Mark Gertler, John Haltiwanger, Jim Harrigan, Charlie Himmelberg, Brad Humphreys, Spence Krane, Trish Mosser, Simon Potter, Tom Sargent, Scott Schuh, Ken West, and especially Douglas Dwyer, Jim Kahn, and Ken Kuttner, as well as seminar participants at the New York Fed, the Bank for International Settlements, the Federal Reserve System Conference on Macroeconomics, the 1998 Symposium of the International Society for Inventory Research, and the 1998 European Meeting of the Econometric Society for their valuable comments. Benjamin Bolitzer provided expert research assistance. All remaining errors and omissions are our own responsibility. Most of the work on this paper was completed while the first author was visiting economist at the Bank for International Settlements, whose hospitality he greatly appreciates. Please address correspondence to Jonathan McCarthy, Research Department, Federal Reserve Bank of New York, 33 Liberty Street, New York City, NY 10045, *e-mail*: Jonathan.McCarthy@ny.frb.org. The opinions expressed in this paper do not necessarily reflect views of the Federal Reserve Bank of New York, the Federal Reserve System, nor the Bank for International Settlements.

Contents

1.	Introduction	1
2.	“Puzzles” arising from the L-Q model	3
3.	Basic methodology	8
4.	Data and econometric issues	12
	4.1 Data	12
	4.2 Measuring inventory deviations	13
5.	Results	20
	5.1 Microeconomic adjustment functions	21
	5.1.1 <i>Average adjustment</i>	21
	5.1.2 <i>Sectoral adjustment</i>	23
	5.1.3 <i>Seasonal adjustment</i>	25
	5.1.4 <i>Period adjustment</i>	26
	5.1.5 <i>Distribution of adjustments</i>	28
	5.2 The cross section of inventory deviations	30
	5.3 Aggregate dynamics and cross-sectional moments	36
6.	Conclusion	39
A.	Data appendix	41
	A.1 Selection rules	41
	A.2 Construction of variables	42
	A.3 Data summary	42
	References	45

1. Introduction

After lying dormant for much of the 1960s and 1970s, empirical research on inventories has undergone a renaissance during the last two decades. A major impetus for this renewed interest has been the long-known, although periodically overlooked, empirical observation that inventory fluctuations are highly correlated with the business cycle.¹ This observation has led macroeconomists to examine inventories as a key component in the propagation and the amplification of exogenous shocks to the economy. Relatedly, there has been considerable debate whether recent improvements in inventory control (e.g., just-in-time techniques, bar coding, etc.) have muted the inventory cycle, translating into reduced volatility of aggregate output fluctuations.²

While this effort has greatly advanced our understanding of inventory investment and its role in the business cycle fluctuations, economists continue to be perplexed by various aspects of firm and aggregate inventory behavior. As documented in a recent survey by Ramey and West (1997), the strong procyclical movements of inventories and the persistence of inventory movements conditional on sales have proven to be particularly difficult to reconcile with the predictions of canonical inventory models.

We argue in this paper that part of the problem lies in the framework that underlies most applied inventory research – the linear quadratic (L-Q) model of Holt, Modigliani, Muth, and Simon (1960). The linear dynamics implied by this model, although relatively easy to aggregate and estimate, do not capture potential nonlinear features of microeconomic inventory behavior. Among others, these would include (S,s) -type inventory policies owing to the presence of fixed or proportional costs in the production technology and asymmetries or irreversibilities induced by differing costs between drawing down and expanding inventory levels.

In addition, nonlinear microeconomic inventory behavior has the potential to affect significantly aggregate inventory dynamics and can call into question the representative agent assumption underlying most applied aggregate inventory research. For instance, suppose firms use state-dependent (S,s) rules to manage their inventory stocks. In such an economy, a negative aggregate shock may result in fewer firms reaching their trigger inventory levels, exacerbating the decline in inventory investment beyond what would be expected under a representative agent L-Q model.

Our goal in this paper is to document the prevalence of nonlinear microeconomic inventory behavior and the impact that such behavior has on the aggregate dynamics. We employ a flexible and

¹ Nearly half a century ago, Abramowitz (1950) provided the first statistical evidence, showing that a typical U.S. recession prior to World War II was characterized by intense inventory disinvestment. More recently, Blinder and Maccini (1991) have shown that this regularity continues to hold in the postwar data.

² See, for instance, Allen (1995), Filardo (1995), and McConnell and Perez-Quiros (1998).

empirically tractable framework – the so-called generalized (S,s) approach – developed initially by Caballero (1993, 1997), Caballero and Engel (1991, 1993), and Caballero, Engel, and Haltiwanger (1995, 1997) to study employment and fixed capital investment dynamics. A fundamental element in our analysis is the inventory adjustment function. This function relates the fraction of the deviation between “desired” and actual inventory levels that firms close during a period to the size of that deviation. The interaction of this adjustment function with the cross-sectional distribution of inventory deviations determines the aggregate dynamics of the model.

The advantages of our approach are twofold. First, aggregate inventory dynamics are derived explicitly from the underlying microeconomic data rather than from a representative agent framework. Second, our approach is very flexible. The adjustment function can take on a wide variety of shapes, including those implied by the partial adjustment model, a simple (S,s) model, and a more general nonlinear asymmetric adjustment model. Furthermore, if the adjustment function is constant, as in the case of a partial adjustment model, the aggregate dynamics generated within our framework are identical to the dynamics implied by the representative agent L-Q model (Rotemberg (1987) and Caballero and Engel (1993)).

We use our framework to analyze the dynamics of inventory investment in the U.S. business sector from 1981 to 1997, utilizing comprehensive, high-frequency (quarterly) firm-level data. Our results indicate that there are significant nonlinearities and asymmetries in the estimated adjustment functions in both the trade and manufacturing sectors of the economy. The nonlinearity is consistent with a use of (S,s) -type inventory policies, while the asymmetries suggest the presence of nonconvexities in the production/deliveries technology. There are sizable differences between the adjustment functions of trade and manufacturing firms as well as between nondurable and durable goods firms within each sector. In particular, durable goods firms appear to be less likely to adjust, especially when holding “excess” inventories.

In addition, we find some striking results concerning temporal variation in the inventory adjustment process. The seasonal differences found in trade firms’ adjustment functions are consistent with seasonal selling patterns and end-of-year effects. Interestingly, in both the manufacturing and trade sectors, the adjustment function appears to be remarkably stable over our sample period, indicating that the inventory adjustment process has changed little over time despite recent advances in inventory management techniques. Combined with evidence that the dispersion of inventory deviations has not declined over time, the stability of the adjustment function suggests that inventory cycles have not been muted.

One implication of our model is that the higher moments of the cross-sectional distribution of inventory deviations should have an effect on aggregate inventory dynamics. To test this hypothesis, we employ a simple parametric approach, which assumes a time-invariant polynomial approximation

to the adjustment function. This analysis indicates that the nonlinearities in the adjustment function, in conjunction with the movements in the cross-sectional density, have a significant effect on the aggregate dynamics.

The rest of this paper proceeds as follows. In the next section, we discuss the traditional L-Q model and the “puzzles” that arise in the context of the model. In section 3, we outline our approach to study inventory dynamics and discuss the implications of some specific models of inventory behavior on the shape of the inventory adjustment function. In section 4, we provide a brief description of the data and related issues and discuss the operational details behind the construction of our key state variable – the deviation between the actual and the “desired” level of inventories. In section 5, we present our main results. Section 6 concludes and discusses the implications and extensions for future research.

2. “Puzzles” arising from the L-Q model

To date, most applied inventory research has employed a version of the linear-quadratic (L-Q) model developed by Holt, Modigliani, Muth, and Simon (1960). The prototypical L-Q model assumes that firms maximize profits subject to a convex production technology, which implies that firms will attempt to smooth production in the face of stochastic sales. When incorporated into a representative agent framework, the L-Q model has provided economists with a microstructure that is relatively easy to aggregate and estimate, and that has yielded substantial insights into firm and aggregate inventory dynamics.

The key assumption of the L-Q model is a convex production technology. Accordingly, the L-Q model provides a more natural description of the behavior of manufacturers than the behavior of retailers and wholesalers. Consequently, the vast majority of recent empirical inventory research has focused on manufacturing inventories, especially inventories of finished goods – see Ramey and West (1997) for a recent comprehensive review.³ This focus has occurred despite the fact that trade (retail and wholesale) inventories now account for more than one-half of total business inventory stocks in the United States (see Figure IA).

Moreover, trade inventories play a significant role in cyclical fluctuations. The variance of trade inventory growth has accounted for about 35 percent of the variance of business inventory growth since 1959, compared to about 40 percent for manufacturing.⁴ In an analysis of more recent data,

³ The exceptions to this pattern include Trivedi (1973), Blinder (1981), Irvine (1981a, 1981b), Zakrajšek (1997), and McCarthy and Zakrajšek (1998).

⁴ Figure IB plots the growth rate of inventories for the three major components of the business sector. The data are quarterly, seasonally adjusted and for visual purposes have been smoothed by running medians.

Figure IA
Share of Trade Inventories





Figure IB
Inventory Growth Rates



Note: The shaded regions represent N.B.E.R. recessions.

Worthington (1998) identifies the retail trade sector as a particularly important component of aggregate inventory fluctuations.

Although we find this focus on manufacturing inventories disconcerting, other authors have been less concerned. Some argue that even if individual transactions are not precisely the same as those assumed in the L-Q model, the larger aggregate still could behave as if firms are solving such an optimization problem; see, for example, Blanchard (1983) and Ramey and West (1997).

Examination of the aggregate dynamics, however, point to more substantive empirical problems with the model. Despite its theoretical elegance and the relative simplicity of its empirical counterpart, the L-Q model has had difficulty explaining two fundamental and robust features of inventory investment. The first fact, noted as early as Metzler (1941), is that inventory investment is highly procyclical; inventories tend to be built up gradually in expansions and to be drawn down rapidly in recessions. Second, as discussed by Ramey and West (1997), inventory movements exhibit considerable persistence, even after conditioning on sales.⁵

The prevailing linear-quadratic-representative-agent framework has offered two competing explanations for these facts. The first assumes that there exist highly persistent shocks to the cost of production. These cost shocks lead to procyclical inventory investment because in times of low cost – that is, times of negative cost shocks – production is more efficient. So even though production technology is convex, firms will bunch up production and accumulate inventories; of course, during times of positive cost shocks, the opposite is true. The assumed persistence of the shock process implies that a cost shock that perturbs the inventory-sales relationship will take many periods to die off. This persistence, in turn, leads to the persistence in the inventory-sales relationship.

The second explanation is built around a strong accelerator effect and a high cost of adjusting production. The accelerator motive links this period's inventories to next period's expected sales. The connection reflects the economic significance of stockout (backlog) costs, which arise when sales exceed available stock. Such circumstances may lead to lost sales or at least delayed payments if orders can be backlogged. The accelerator effect and positively serially correlated sales then cause inventory movements to be procyclical. The high costs of adjusting production imply that if a shock perturbs the inventory-sales relationship, a return to long-run equilibrium will be gradual because firms will adjust production only gradually to minimize adjustment costs.

Although each explanation has some appealing aspects, the data provide limited support for both explanations. The highly persistent cost shock hypothesis seems to work only when the cost shocks are

⁵ Despite mixed evidence that inventories and sales are cointegrated at the industry- or the aggregate-level, the presumably stationary linear combinations of inventories and sales – the so-called “inventory-sales” relationship – exhibit very high first- and second-order autocorrelations, even at an annual frequency, indicating that the adjustment to long-run equilibrium takes place over many periods.

modelled as unobservable to the econometrician. While it is plausible that the firm's cost structure could be affected by persistent and unobservable disturbances, the observable counterparts to these disturbances, such as real unit labor costs and interest rates, appear to have no appreciable effect on inventories. Evidence in favor of significant adjustment costs is equally unpersuasive. Estimates of adjustment cost parameters are unstable across different specifications and estimation methods and range from negligible to economically implausibly large.⁶

One option in response to these negative results is to investigate the implications of relaxing the assumption of convex adjustment technology in the L-Q model. For instance, the operations research literature going back at least as far as the early 1960s recognized that firms are likely to face fixed costs when making production and/or order decisions (see Scarf (1960)). More recently, Blinder (1981) and Blinder and Maccini (1991) have provided compelling arguments that these fixed costs are crucial for understanding the dynamics of inventory investment, particularly in the trade sector. The presence of such fixed costs then will lead firms to adopt (S,s) -type inventory policies, which induce production/delivery bunching rather than smoothing.

The recognition that fixed costs may play a crucial part in understanding inventory dynamics is part of a general and persuasive argument that the constraints to adjustment faced by individual firms are significantly different from the constraints implicit in the quadratic adjustment technology. At the microeconomic level, nonconvexities in adjustment technology such as irreversibilities induced by technological and/or market factors, indivisibilities and other forms of increasing returns are more likely the norm than the exception. In contrast to the L-Q model, these nonconvexities imply a pattern of microeconomic inventory investment that is highly nonlinear and asymmetric. Periods of inertia or more or less passive inventory accumulation are followed by rapid inventory investment or disinvestment – a response not only to contemporaneous shocks but also to the history of accumulated shocks.

Equally important is the fact that the nonlinear microeconomic inventory policies mean that the representative agent assumption is no longer appropriate, which greatly complicates aggregation. As discussed in greater detail in the next section, the evolution of the entire cross-sectional distribution of inventory deviations from target levels is a key to the evolution of the aggregate inventory investment. In particular, higher moments than the mean – which is a sufficient statistic in the representative agent L-Q model – may affect aggregate inventory dynamics.

⁶ As is the case in all applied work, both at the micro- and macro-level, the inventory literature is not immune to serious measurement problems in the data that undoubtedly contribute to the poor empirical performance of the L-Q model. Some studies try to mitigate this problem by using the presumably more accurate – though considerably more limited – physical product data and find somewhat greater support for the L-Q model, in particular for the production smoothing motive; see Fair (1989) and Krane and Braun (1991) for examples of this approach.

The inherent nonlinearities of the (S,s) -type models combined with complex aggregate dynamics have made these models considerably less appealing to empirical macroeconomists interested in inventory fluctuations.⁷ Despite these obstacles, some progress has been made in recent years to examine the macroeconomic implications of inventory models with (S,s) -type decision rules. Theoretical results on the aggregation of (S,s) economies by Blinder (1981), Caplin (1985), and Caballero and Engel (1991) provided the steady-state reduced-form predictions used by Mosser (1988, 1991) and Episcopos (1996), to test the basic (S,s) inventory model on the U.S. and Canadian industry-level trade data; in a recent paper, McCarthy and Zakrajšek (1998) extend the analysis to firm-level data.

Although Mosser (1988, 1991), Episcopos (1996), and McCarthy and Zakrajšek (1998) find the behavior of inventories, sales, and deliveries broadly consistent with the steady-state implications of the (S,s) model, the problem with this approach is twofold. First, the steady-state aggregation results of Blinder (1981) and Caplin (1985) are valid only under very restrictive assumptions, such as exogenous serially uncorrelated sales, time-invariant (S,s) bands, and no delivery lags.⁸ The second problem lies in the fact that the steady-state reduced-form probability relationships between inventory investment and sales are in fact consistent with an economically plausible parametrization of the standard L-Q model (see Blinder (1986), Kahn (1987), and Krane (1994)).⁹

Accordingly, if we are to shed light on the possible effects of microeconomic nonlinear inventory policies on aggregate inventory investment, we must study inventory dynamics rather than steady state behavior. Because structural (S,s) -type models must be kept relatively simple in order to derive analytic results useful for empirical analysis, studying inventory dynamics using real data almost surely would reject such models. Thus, in our approach, we will sacrifice some structural rigor to provide a tractable empirical framework, which nonetheless encompasses the possibility of nonlinear behavior at the microeconomic level.

Alternatively, Schuh (1996) dispenses with the representative agent assumption. Using plant-level monthly data, Schuh (1996) estimates the L-Q model and finds that accounting for the aggregation bias – in both the cross-sectional and the time series dimensions – moderately improves the fit of the canonical L-Q model.

⁷ An important exception is Fisher and Hornstein (1995) who embed an (S,s) inventory problem into a general equilibrium model.

⁸ In the empirical implementation, Mosser (1988, 1991) and McCarthy and Zakrajšek (1998) attempt to allow for delivery lags and serial correlation in sales and continue to find support for the simple (S,s) model.

⁹ To address this concern, McCarthy and Zakrajšek (1998) also estimate the Euler equation associated with the canonical L-Q model. They find that, although the model is overwhelmingly rejected at the industry level, estimates of the structural cost parameters are economically reasonable and statistically significant at the firm level; however, the overidentifying restrictions imposed by the model are rejected, and the parameter estimates are not stable across different asymptotically equivalent normalizations.

3. Basic methodology

In this section, we describe the basic methodology for our analysis of the relationship between the microeconomic – that is, firm-level – features of inventory investment and the implied aggregate dynamics. To model inventory investment in the manufacturing and trade sectors of the U.S. economy, we adopt a framework used by Caballero (1993, 1997), Caballero and Engel (1991, 1993), and Caballero, Engel, and Haltiwanger (1995, 1997) to examine employment and fixed investment dynamics.

The basic premise of our analysis is that underlying aggregate inventory dynamics is a population of heterogeneous firms whose inventory adjustment within a period depends upon the size of their perceived inventory shortfall or excess. In particular, we assume that firms may respond proportionally more the further away inventories are from their “target.” As noted earlier, such differences in adjustment can result from fixed or proportional costs in the production technology as postulated in (S,s) -type models.

Our model at the microeconomic level is built around a single state variable: a measure of the deviation between desired and actual log-level of inventories at the firm-level, which we label as the *inventory deviation index* z :

$$(1) \quad z_{it} \equiv \ln H_{it}^* - \ln H_{it-1}$$

In equation (1), i indexes firms, t indexes time, and H_{it} denotes real, end-of-period inventory stocks. It is important to note that z_{it} depends on the desired level of inventories, H_{it}^* , a theoretical construct, which implies that z_{it} is ultimately model-dependent.

The evolution of z_{it} over time reflects the shocks to the desired inventory level and the inventory adjustment that the firm undertakes in response to these shocks. Shocks to desired inventories can be classified into aggregate shocks – those common to all firms – and idiosyncratic shocks that are firm-specific. From equation (1), it follows that the change in a firm’s inventory deviation during period t , Δz_{it} , can be decomposed as follows:

$$(2) \quad \Delta z_{it} = \Delta \ln H_{it}^* - \Delta \ln H_{it-1} = (\eta_t + v_{it}) - \Delta \ln H_{it-1},$$

where $(\eta_t + v_{it})$ represents the decomposition of the desired inventory growth into economy-wide average desired inventory growth, η_t , and a firm-specific idiosyncratic shock, v_{it} , which, by definition, has a zero cross-sectional mean in each period.

Because we are working in discrete time, the timing convention of shocks and adjustment is important. Following Caballero, Engel, and Haltiwanger (1997), we assume that each period begins with an idiosyncratic shock v_{it} , displacing the inventory deviation that the firm had carried over from the

previous period. The result is a perturbation in the cross-section of inventory deviations that is then subjected to the aggregate shock θ_t . Finally, the period ends with firms adjusting their inventory levels.

In any given period t , the underlying heterogeneity among firms is captured by the *cross section of firms' inventory deviations*, denoted by $f(z,t)$. That is, $f(z,t)$ is the cross-sectional distribution of firms' inventory deviations immediately preceding period t 's adjustments. Therefore, the fraction of firms with inventory deviations between z and $z+dz$ in period t is approximately equal to $f(z,t)dz$.

Average inventory investment, where the average is computed over firms with a similar inventory deviation index in period t , is described by the *adjustment function*, denoted by $\gamma(z,t)$. From an operational standpoint, we consider a class of firms with similar inventory deviations prior to adjustment and calculate the fraction of the inventory deviation that is closed, on average, by firms within each of these classes. From this definition, it then follows that the average inventory investment by firms with inventory deviation z in period t is equal to $z\gamma(z,t)$.

In this framework, therefore, the three elements needed to relate firm-specific inventory investment decisions to aggregate inventory growth are the inventory deviation index z , the adjustment function $\gamma(z,t)$, and the cross-sectional density of firms' inventory deviations $f(z,t)$. Letting $\Delta \ln H_t^A$ denote the growth rate of aggregate inventories in period t , then the preceding definitions imply that

$$(3) \quad \Delta \ln H_t^A = \int z\gamma(z,t)f(z,t)dz$$

We use $\Delta \ln H_t^A$ as our measure of *aggregate inventory growth*. Note that $\Delta \ln H_t^A$ differs from the rate of growth of aggregate inventories in our sample only in that our measure does not weigh firms' inventory growth by their size at each point in time.¹⁰


Equation (3) highlights the connection between the movements in the cross-sectional distribution of inventory deviations and the growth rate of aggregate inventories. The dynamics of aggregate inventory investment are determined by the interaction between the shape of the adjustment function and the shifts in the cross-sectional density of inventory deviations induced by aggregate and idiosyncratic shocks. In general, as long as the adjustment function $\gamma(z,t)$ depends explicitly on z , aspects of $f(z,t)$ other than its mean (e.g., dispersion, skewness, and concentration) will influence aggregate dynamics.

To illustrate this interaction more concretely, let us consider some simple examples. Figure IIA displays a constant adjustment function, $\gamma(z)/\delta_0$. In this case, because the adjustment rate (solid line) is constant, a shift in the cross-sectional distribution caused by an aggregate shock (dashed lines) does


¹⁰ The time series behavior of the two measures of aggregate inventory growth is very similar; the correlation between the two series is 0.84 for the manufacturing sector and 0.92 for the trade sector.

Figure II

(a) Partial Adjustment Model



(b) (S,s) Model



not affect the adjustment rate of individual firms. The effect of the aggregate shock thus is summarized by the shift in the mean of the distribution. In fact, as shown by Rotemberg (1987) and Caballero and Engel (1993), this adjustment function generates aggregate dynamics identical to linear dynamics of the partial-adjustment model, which in turn can be obtained from a representative-agent framework with quadratic adjustment costs.

Next, suppose firms follow identical simple (S,s) policies (Figure IIB). In this case, $\gamma(z)=0$ when the inventory deviation lies in the inaction range (S,s) , and $\gamma(z)=1$ when the deviation reaches any of the trigger points – that is, firms adjust inventories to their “desired” levels when inventories move outside the (S,s) range. The effect of a positive aggregate shock on aggregate inventory investment under this policy can no longer be summarized by the shock’s effect on the mean of the distribution. Besides the effect of shifting the mean, which will have a positive impact on aggregate inventory investment, the shock increases the fraction of firms which will invest ($z>s$) and decrease the fraction of firms which will disinvest ($z<S$), which has an additional positive effect on aggregate investment. More generally, shocks which increase the dispersion of the distribution will affect aggregate inventory investment by increasing the fraction of firms outside of the inaction range.

The partial adjustment model and the (S,s) model provide two extremes of the adjustment process. The former implies smooth gradual adjustment, while the adjustment process in the latter model is lumpy and infrequent. However, even if fixed adjustment costs are prevalent among firms, adjustment functions are not likely to exhibit the extreme (S,s) shape. Even the most disaggregated economic data have been aggregated to some extent (e.g., over products, time, etc.). Hence, it seems likely that the observed adjustment process will incorporate some degree of smoothness as well as a degree of lumpiness.

The key feature of the generalized (S,s) approach is that the fraction of the inventory deviation that is closed in each period increases with respect to the (absolute) size of the gap between desired and actual inventory levels. For example, if the adjustment function is a time-invariant quadratic function $\gamma(z)=\gamma_0+\gamma_1z+\gamma_2z^2$, then equation (3) implies that

$$(4) \quad \Delta \ln H_t^A = \lambda_0 m_z^{(1)}(t) + \lambda_1 m_z^{(2)}(t) + \lambda_2 m_z^{(3)}(t),$$

where $m_z^{(k)}(t)$ denotes the k^{th} noncentral moment of the cross-sectional distribution of inventory deviations in period t . Therefore, in this case, the evolution of the standard deviation and skewness of the cross-sectional distribution of inventory deviations will have an impact on aggregate inventory investment.

4. Data and econometric issues

4.1 Data

The firm-level data used in our analysis come from the COMPUSTAT quarterly P/S/T, Full Coverage and Research data files. Two data sets are used in the analysis. The first is a panel of 3,946 manufacturing firms, and the second is a panel of 1,206 retail and wholesale trade firms. Because our measurements of inventory deviations require estimates at the two-digit SIC level, firms in SIC 29 (Petroleum refining and related industries) and SIC 55 (Automotive dealers and gasoline service stations) were excluded from the sample owing to an insufficient number of firms for reliable estimation.

Both panels span the time period 1980:Q4 to 1997:Q4 (69 quarters) and are unbalanced, with the minimum (continuous) tenure in each panel being 8 quarters. After eliminating firms with gaps in the time-series dimension or implausible entries, we were left with a total of 118,885 firm/quarter observations for the manufacturing sector and 32,666 firm/quarter observations for the trade sector. During our sample period, firms in the manufacturing panel account, on average, for almost 65 percent of aggregate manufacturing inventories, and firms in the trade panel account for about 20 percent of aggregate trade inventory stocks. The Data Appendix contains details on the exact sample selection procedure, the construction of variables as well as summary statistics of the samples.

Let us make one additional point concerning the data. If nonlinear adjustment is important in determining inventory dynamics, *plant (outlet)-level* data may be more desirable than firm-level data. Decisions and shocks at the plant level can have implications for the firm's inventory policy. For example, manufacturing plants within a single firm that produce different products may have quite different inventory policies depending upon the demand for the products that each plant produces. Similarly, retail outlets in different locations will make inventory decisions based on local demand conditions. Thus to the extent that individual plants have an independent existence and make independent decisions, the distribution of shocks and inventory deviations across plants within a firm may influence the firm's overall inventory policy. Concentrating at the firm level then may miss some information concerning microeconomic adjustment and its effect on aggregate dynamics.

Nevertheless, many aspects of inventory policies are centralized within the firm and thus are likely to depend on firm-wide conditions and shocks. Finished manufacturing inventories may go to a centralized distribution center, which enables the firm to make production and inventory decisions more efficiently. Much of the inventory for retail firms may be held in central warehouses. Furthermore, financial conditions and capital market access are firm- rather than plant-level phenomena. Hence, the well-documented sensitivity of inventory investment to movements in internal funds or net worth – caused by the limited access of many firms to external credit markets – is evidence of the role that nonconvexities may play at the firm level. Thus much of the effect that

microeconomic nonconvexities may have on aggregate inventory dynamics can be studied at the firm level. If anything, using firm-level data may understate the role of nonlinear inventory adjustment in determining aggregate inventory fluctuations.

4.2 Measuring inventory deviations

In the discussion of Section 3, equation (3) determined the dynamics of aggregate inventory investment *given* a measure of the inventory deviation index z . To construct a measure of z , we use the observation that firms typically compare their inventory levels to their sales as a gauge of the appropriate inventory level.

Specifically, we assume that the desired level of inventories in period t , H_{it}^* is given by the following function of sales:

$$(5) \quad H_{it}^* = A_{it} S_{it}^{\sigma_i},$$

where S_{it} denotes real sales in period t , A_{it} is a firm-specific and time-varying parameter governing the significance of the stockout avoidance motive, and Φ_i is a firm-specific scale parameter. Given the timing of the model, we can now derive our measure of z . First, let e_{it} denote the inventory deviation index at the end of the period, *after* all the adjustments have taken place. It follows by definition that

$$(6) \quad e_{it} = \ln H_{it}^* - \ln H_{it} = h_{it}^* - h_{it} = (a_{it} + \sigma_i s_{it}) - h_{it},$$

where the lower case variable designates the logarithm of that variable. Note that from equations (1) and (6), e_{it} differs from z_{it} only in that the former incorporates the inventory adjustment during the period; that is, $e_{it} = z_{it} - h_{it}$.

The ex post inventory deviation index e_{it} is observed by the firm at the end of the period after all shocks including sales shocks have taken place. The key identifying assumption of our model is that once the adjustment has occurred, any deviation of inventories from their desired level will not persist indefinitely. This implies that the ex post inventory deviation index e_{it} is stationary and that $E[e_{it}] = 0$. From equation (6), our identifying assumption implies that e_{it} can be considered as a residual in the following regression:

$$(7) \quad \ln \left[\frac{H}{S^{\sigma_i}} \right]_{it} = a_{it} - e_{it}.$$

Equation (7) contains two unknown variables: the logarithm of the stockout avoidance constant a_{it} and the scale parameter Φ_i , which determines the desired inventory-sales relationship. We first turn to the problem of estimating Φ_i .

Although we have written the model as if Φ_i was a constant parameter, there are many economic reasons why Φ_i may vary over time as well as over individual firms. The entry into and exit from our sample could cause Φ_i to vary over time because the desired inventory-sales relationship may be markedly different when a firm is in transition. Time and/or state dependence in inventory policies and advances in inventory monitoring technology are likely to influence the importance of stockout costs over time, which in turn would cause firms to change the weighting given to sales in determining the desired inventory-sales relationship over time. It thus seems important to allow for time as well as cross-sectional variation in the estimation of Φ_i .

To achieve a reasonable compromise between cross-sectional heterogeneity and time variation and still obtain a degree of precision, we estimate the following cross-sectional regression for each period t :

$$(8) \quad h_{it} = \text{const}_{it} + \sigma_{it}s_{it} + \varepsilon_{it}.$$

We estimate the values of Φ_{it} in each period by pooling the firm-level data for each 2-digit industry. That is, we estimate equation (8) for each period by least squares, allowing both the constant and the scale parameter Φ_{it} to vary across different industries.

Least squares estimates of Φ_{it} from equation (8), however, are likely to be biased downward because of probable measurement error in the sales variable. To ascertain the severity of this bias, we also estimate the reverse regression (i.e., with s_{it} as the response variable in the regression), which because of the probable measurement error in inventories, leads to an upwardly biased estimate of Φ_{it} . From these two regressions, we obtain lower and upper bounds for the estimate of the scale parameter Φ_{it} .

Figure III displays the time path of the estimated scale parameter for each 2-digit industry. The solid line in each panel corresponds to the lower bound estimate, and the dashed line is the upper bound estimate of Φ_{it} ; the straight horizontal lines are the time *invariant* estimates of the lower and upper bounds of Φ_{it} . Note that for a majority of industries the difference between the lower and the upper bound estimates is in the order of 10 percent; the measurement error bias is the most acute in SICs 27 (Printing & Publishing), 31 (Leather & Leather Products), and 59-1 (Other Retail Durables) where the difference between the lower and the upper bound estimates exceeds 20 percent. All of our subsequent results, however, are robust to either estimate of the scale parameter Φ_{it} .¹¹

For the most part, the estimates of the scale parameter Φ_{it} are close to one, although there are industries and time periods where the parameter is substantially above or below unity. More

¹¹ This is consistent with the theoretical analysis of Caballero, Engel, and Haltiwanger (1997) who find that measurement error is more likely to conceal rather than artificially generate the nonlinear nature of the inventory adjustment process. To check the robustness of our results, we considered several alternative estimators of Φ_{it} , including the L_I and the IV estimator. However, the alternative estimates of Φ_{it} had only negligible effects on our reported results.

importantly, note that in many industries the parameter displays both low- and high-frequency (seasonal) fluctuations, indicating the need to allow the unconstrained temporal variation in this parameter.

FIGURE III
Time-Varying Estimates of the Scale Parameter

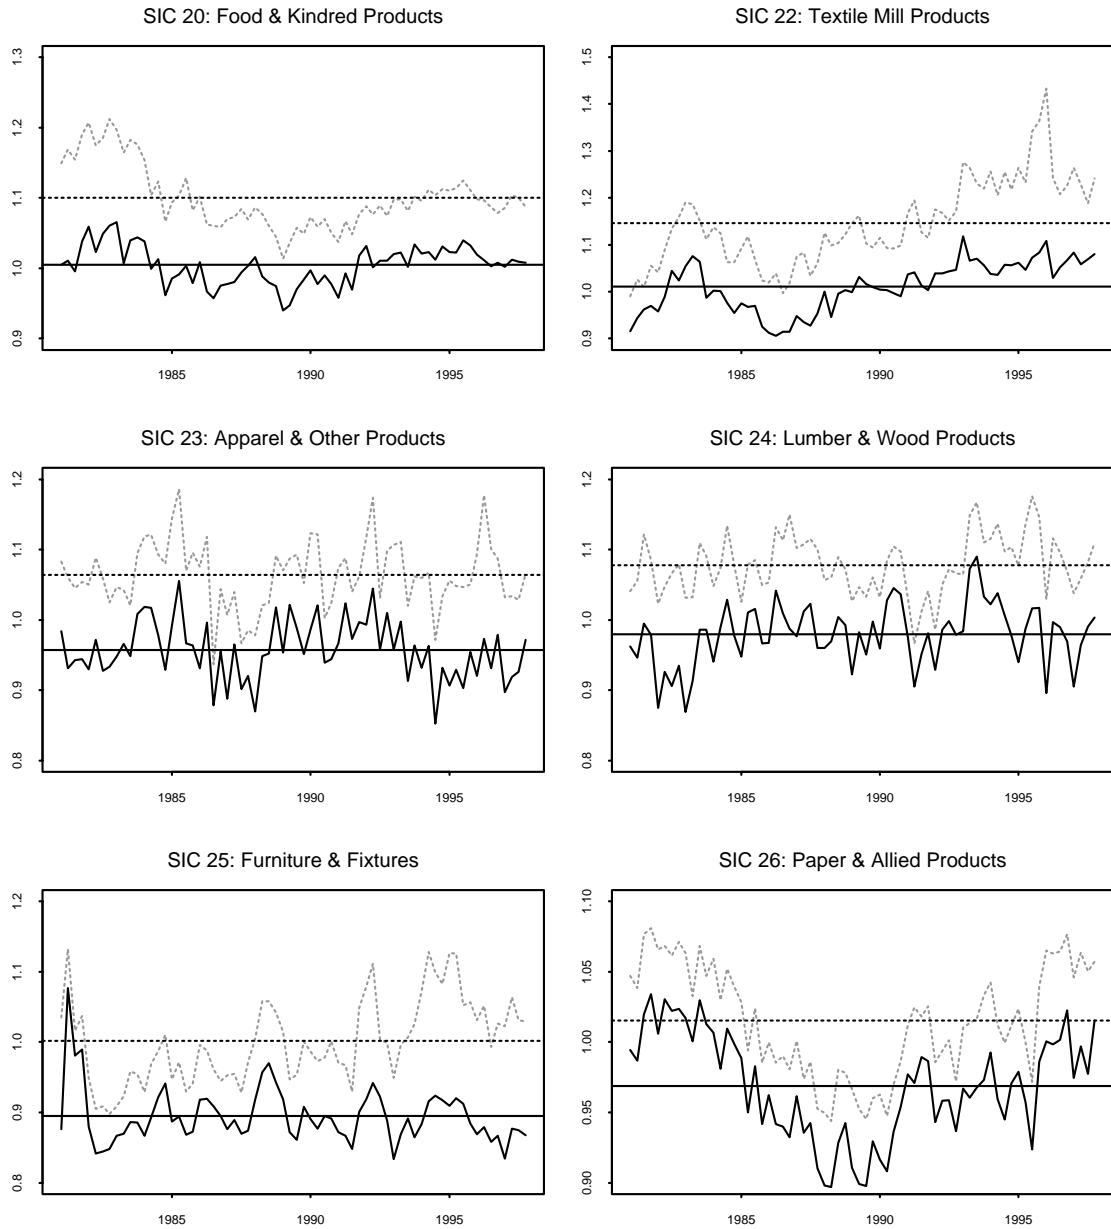


FIGURE III

Time-Varying Estimates of the Scale Parameter




FIGURE III

Time-Varying Estimates of the Scale Parameter




FIGURE III

Time-Varying Estimates of the Scale Parameter

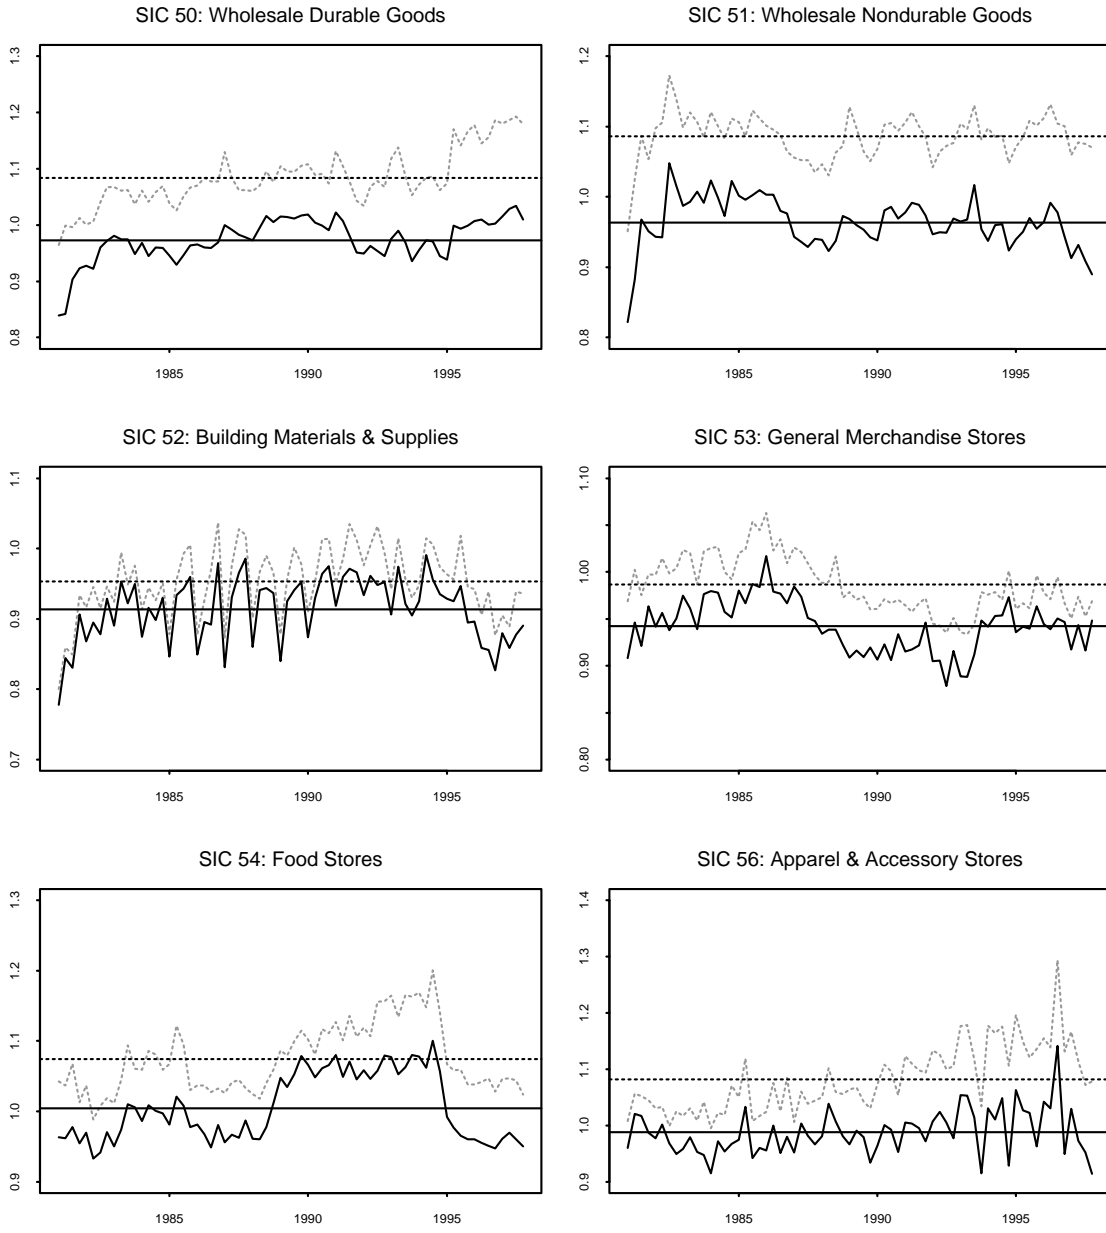
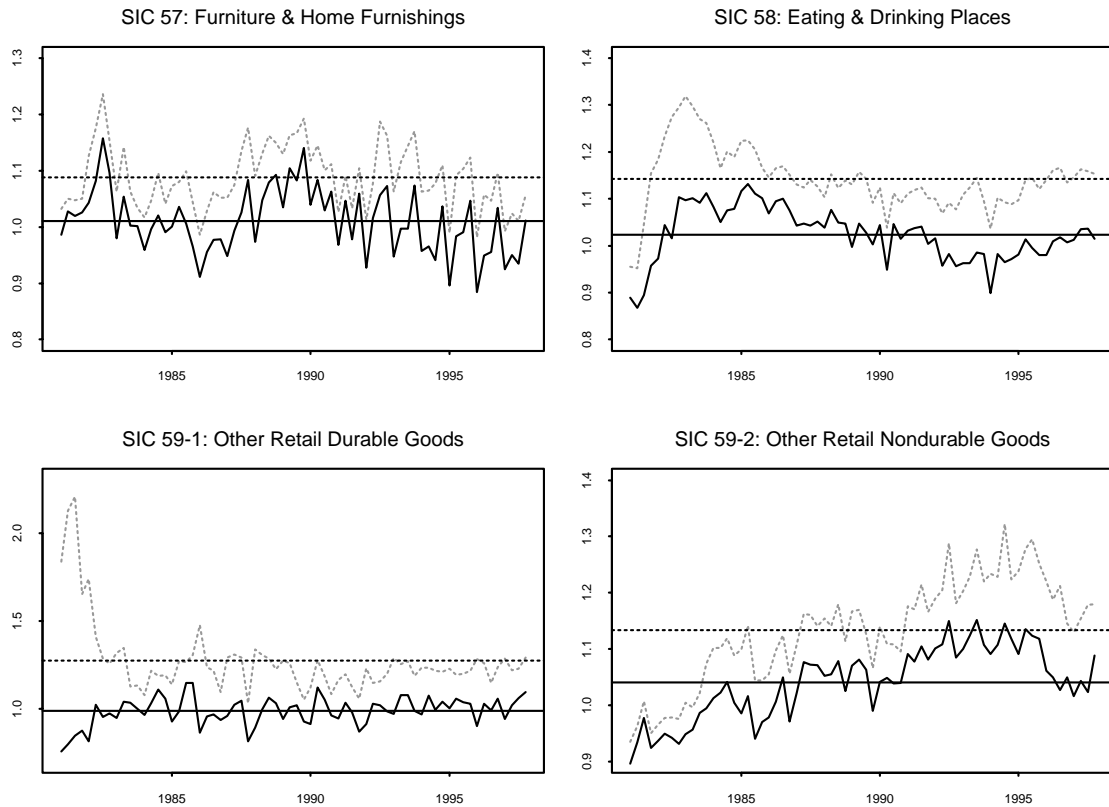


FIGURE III

Time-Varying Estimates of the Scale Parameter



We next consider the stockout avoidance parameter A_{it} . We assume that a firm-specific stockout avoidance behavior over time consists of a low- and high-frequency components. The low-frequency component reflects such influences as advances in inventory monitoring technology, changes in the firm's relationship with its suppliers, and changes in product diversity. The high-frequency component, on the other hand, is meant to capture movements in the stockout avoidance behavior associated with seasonal fluctuations.

Specifically, we assume that A_{it} takes on the following log-linear form:

$$(9) \quad a_{it} = \gamma_i \times YEAR(t) + \theta_i \times QTR(t).$$

In equation (9), γ_i and θ_i denote fixed firm effects, $YEAR(t)$ is an indicator function that takes on the year value associated with the observation in period t , and $QTR(t)$ is an indicator function that takes on the quarter value associated with the observation in period t .¹²

Conditional on the estimate of the scale parameter Φ_{it} , we can then obtain an estimate of the ex post inventory deviations from the following regression:

$$(10) \quad \ln \left[\frac{H}{S^{\sigma_{it}}} \right]_{it} = \gamma_i \times YEAR(t) + \theta_i \times QTR(t) + u_{it}.$$

We estimate equation (10) for each 2-digit industry by weighted least squares with the weights equal to the firm's tenure in the panel.¹³ The negative of the estimated residual \hat{u}_{it} from this regression is our estimate of e_{it} . To derive an estimate of z_{it} , remember that equations (1) and (6) imply that z_{it} differs from e_{it} only in that the latter incorporates adjustment. Therefore, using our estimate of the ex post deviation e_{it} ,

$$(11) \quad z_{it} = e_{it} + \Delta h_{it},$$

where h_{it} denotes the actual growth rate of inventories of firm i in period t . These estimates of the ex ante deviations are then used in our analysis of firm inventory adjustment.

5. Results

In this section, we characterize each of the elements of equation (3). Of particular interest are the shape of the adjustment function and its sectoral, seasonal, and subsample components. Next we examine the evolution of the cross-sectional distribution of the inventory deviation index z . We measure the interaction between the shape of the adjustment function and the fluctuations in the cross-sectional density and relate its impact to aggregate dynamics.

¹² For example, in period $t = 1995:Q1$, $YEAR(t) = 1995$ and $QTR(t) = 1$. Therefore, $\gamma_i \times YEAR(t)$ denotes fixed firm-specific year effects and $\theta_i \times QTR(t)$ denotes fixed firm-specific seasonal effects.

¹³ Given our assumptions, it is clearly possible to estimate the scale parameter Φ_{it} and the ex post inventory deviation index e_{it} jointly. Because of computational feasibility, we adopt the two stage procedure, which yields consistent, although inefficient, estimates.

5.1 Microeconomic adjustment functions

To compute the average adjustment function, we first discretize the state space. The inventory deviation index z takes values between -0.5 and 0.5, over an equally-spaced grid with intervals of size 0.01.¹⁴ In each interval, we construct the adjustment function by dividing the average inventory growth of those firms that are at z just before inventory adjustments take place by z , for $z \neq 0$.¹⁵ In what follows, all depicted adjustment functions are smoothed by a cubic B-spline.

5.1.1 Average adjustment

The solid line in Figure IV shows the estimated average – over all firms and quarters – adjustment functions for the manufacturing and the trade sectors along with two standard deviation error bands.¹⁶ The dashed lines represent the smoothed average density of inventory deviations. Examining the adjustment function for the manufacturing sector in the top panel, three observations stand out. First, firms with larger absolute deviations adjust more than do firms with smaller deviations irrespective of whether the deviation is a shortage ($z > 0$) or an excess ($z < 0$). Such behavior is indicative of nonconvexities in the adjustment technology, which induce firms to adopt (S,s) -type inventory policies.

Second, the average manufacturing adjustment function displays an asymmetric adjustment process. Manufacturing firms with inventories above their desired level ($z < 0$) adjust less than firms with similar-sized inventory shortages ($z > 0$). There are several possible explanations for such asymmetric behavior. First, because of a strong stockout avoidance motive, manufacturing firms could be more willing to carry extra inventories. Second, market irreversibilities could prevent firms from reducing their excess inventory levels. Alternatively, manufacturing firms could be reluctant to cut output, because they find it costly not to employ their capital and labor.

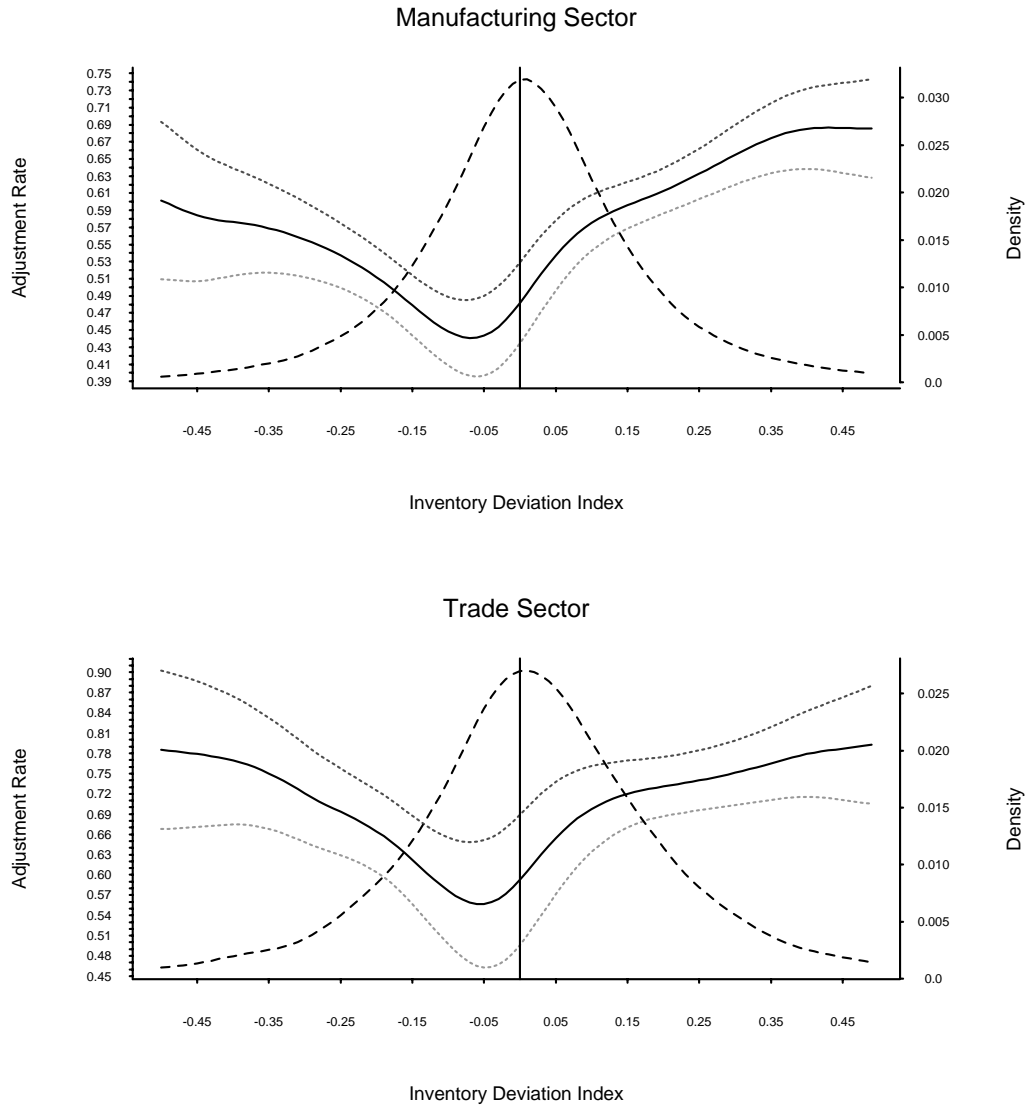
¹⁴ This interval covers over 99.0 percent of our sample.

¹⁵ In calculating the adjustment function, values of z close to zero – that is, between -0.02 and 0.02 – are excluded, because the calculation involves dividing the average adjustment rate by z . Because of the unbalanced nature of our panel and the varying degree of precision regarding our key parameter estimates, the average adjustment rate in each z -interval is computed as a weighted average of inventory growth rates, where the weights are given by the reciprocal of the firm-specific standard deviation of the estimated post-adjustment inventory deviation index e_{it} .

¹⁶ The error bands are obtained via a nonparametric bootstrap method. Specifically, from the original sample, we draw with replacement the estimated inventory deviations and the actual inventory adjustments (i.e., inventory growth rates). For each of the 1,000 bootstrap samples, we compute the adjustment function as described in the text. We then compute the standard deviation of the estimated average adjustment for each point in our z -space. Finally, the resulting ± 2 standard deviation error bands are smoothed using the same procedure as in the case of the adjustment function.

FIGURE IV

Average Aggregate Adjustment Function



Finally, note that the estimated adjustment rates are economically plausible, in contrast to the implied adjustment rate obtained from the estimation of a canonical L-Q model. According to Ramey and West (1997), typical estimates of the adjustment rate from the L-Q model on quarterly data are in the 10-20 percent range, indicating economically implausible large inventory adjustment costs.

As in the manufacturing sector, the adjustment function of trade firms is highly nonlinear (bottom panel). There are, however, several important differences. First, for a given inventory deviation, trade firms close, on average, a larger fraction of their inventory gap. Furthermore, the asymmetry of the

adjustment process is less apparent in the trade sector. Trade firms with large negative inventory deviations seem to be more willing or able to reduce their inventory overhangs.

To summarize, the adjustment functions in the manufacturing and trade sectors exhibit significant nonlinearities indicative of nonconvexities in adjustment, (S,s) -type behavior, and a possible strong stockout avoidance motive. These nonlinearities in the adjustment process imply that higher moments of the cross-sectional distribution of inventory deviations may affect aggregate inventory investment, and that the linear dynamics of the L-Q model are not sufficient to explain aggregate inventory fluctuations.

5.1.2 Sectoral adjustment

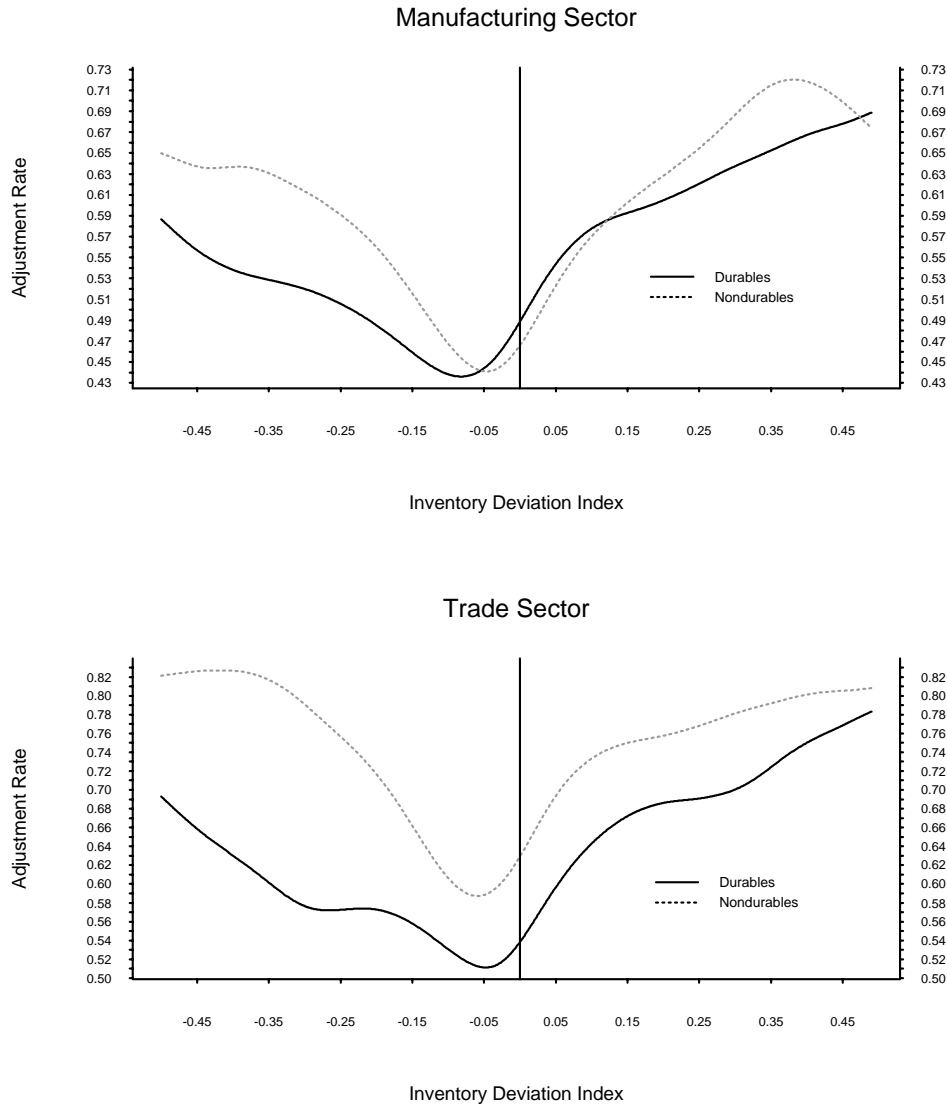
In addition to the differences in the adjustment function between manufacturing and trade firms, there remains scope for considerable heterogeneity within each sector. Furthermore, possible differences in the adjustment function within each sector may provide additional insight into the nature of the nonlinearities documented in the previous section. Figure V displays the differences in the adjustment functions for the durable and nondurable components of each sector.

We first turn to the manufacturing sector. The qualitative shape of the adjustment function for durable and nondurable goods firms is very similar – in both sectors, the adjustment function exhibits significant nonlinearities associated with (S,s) -type inventory policies. However, when firms carry excess inventories ($z < 0$), nondurable goods manufacturers adjust more than durable goods manufacturers, except when the inventory overhang is very small. In the case of an inventory shortage ($z > 0$), the adjustment rate of firms in the nondurable goods sector appears to be somewhat higher than that of durable goods firms. These differences may reflect the fact that nondurable goods manufacturers are more willing or better able to close inventory deviations, possibly owing to differing production technologies or market structures.

The differences between the adjustment functions of the durable and nondurable goods firms in the trade sector are more pronounced. For the entire range of the inventory deviation index z , nondurable goods firms adjust considerably more than firms in the durable goods sector. Note also that the adjustment function of durable goods firms exhibits greater asymmetry: durable goods firms with inventories above their desired level ($z < 0$) adjust less than firms with similar-sized inventory shortages ($z > 0$). This asymmetry does not seem to be present in the nondurable goods sector to the same extent. As noted earlier, the difference in the adjustment function between durable and nondurable goods trade firms could reflect the differing nature of the goods sold – for example, market irreversibilities – the relative significance of the stockout avoidance motive between the two sectors, or technological differences in the delivery process.

FIGURE V

Average Sectoral Adjustment Function



In summary, there are significant differences between the adjustment functions for nondurable and durable goods firms in both the manufacturing and the trade sectors. Each subsector provides considerable evidence of (S,s) -type behavior, while the durable goods subsectors provide greater evidence of asymmetries in inventory adjustment. In both manufacturing and trade, nondurable goods firms appear to be better able or more willing to adjust, especially when it comes to decreasing their inventory stocks. These nonlinearities indicate that the higher moments of the cross-sectional distribution of inventory deviations should matter for the aggregate dynamics in each subsector.

5.1.3 *Seasonal adjustment*

We now turn our attention to the temporal variation in the adjustment function. Undoubtedly, variables other than the inventory deviation have an impact on firms' inventory adjustments. Examining the differences in the adjustment function over time provides a gauge for evaluating the effects of these unmodelled variables. Because our data are quarterly and not seasonally adjusted, we proceed in two steps. First, we examine seasonal fluctuations in the adjustment function. In the next subsection, we look at differences over different subperiods of our sample.

The seasonal patterns in the adjustment function are displayed in Figure VI. In this figure, the fourth quarter adjustment function refers to the average inventory adjustment between the end of the third and fourth quarters. By considering quarterly variation in the adjustment function, we are computing the average adjustment using relatively few observations, especially for the trade sector. Although our panels contain a fair amount of observations, especially for the manufacturing sector, our data sets are unbalanced and considerably smaller than the LRD panel used by Caballero et al.¹⁷ Therefore, the seasonal variation in the adjustment function is likely to be imprecisely estimated and caution must be exercised when interpreting time variation at a quarterly frequency.

Given this caveat, seasonal variations in the manufacturing sector seem to be only minor. The adjustment functions for the second, third, and fourth quarters are essentially parallel shifts of each other, and have similar shape as the overall average adjustment function. The first quarter adjustment function – which captures average inventory adjustment between the end of the fourth and first quarters – exhibits the greatest qualitative differences. Despite these differences, the estimated adjustment function is nonlinear and asymmetric. A puzzling feature of the first quarter adjustment function, especially in light of all our other results, is that the adjustment rate is relatively constant over a considerable range of positive inventory deviations (i.e., inventory shortages).


In the trade sector, the seasonal fluctuations in the adjustment function are more pronounced. In particular, note that the fourth quarter adjustment function indicates that trade firms with excess inventories disinvest more than they do in other quarters. In fact, firms with excess inventories are more likely to adjust than firms with shortages in the fourth quarter, a feature not observed in any of the other quarterly adjustment functions. This may reflect the sell-off of inventory during the Christmas and year-end selling season.

Despite our caveats, seasonal variations in the adjustment function for trade firms appear to be substantial. Besides year-end effects, the differences are consistent with the effect of the Christmas selling season on trade firms.

¹⁷ In their analysis of business fixed investment and employment, Caballero et al. have a *balanced* panel of 10,000 individual plants in the manufacturing sector over the time period 1972 to 1980.

FIGURE VI

Average Seasonal Adjustment Function

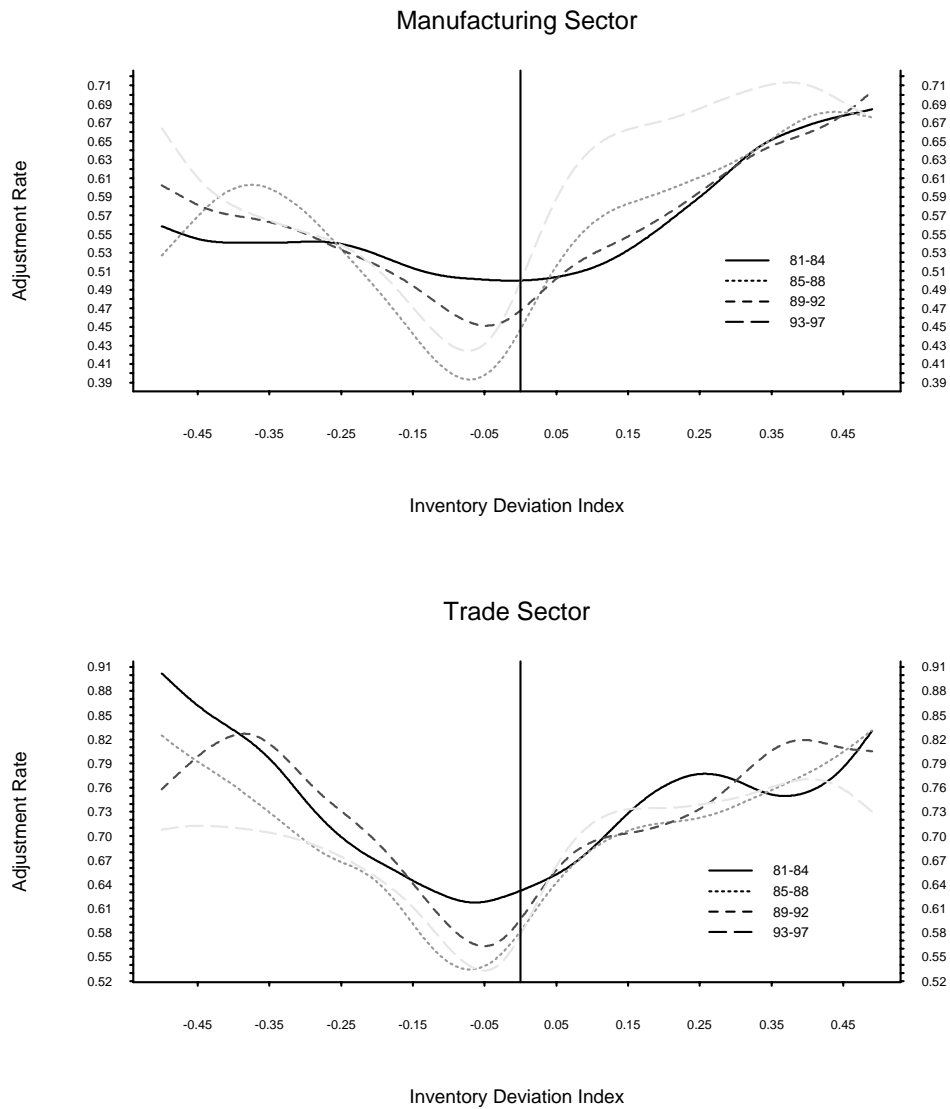


5.1.4 Period adjustment

In this section, we examine the stability of the adjustment function across different subperiods of our sample. This analysis allows us to investigate the extent to which technological improvements in inventory control may have affected the adjustment process. Furthermore, it provides an idea of how low-frequency movements in the adjustment function affect aggregate inventory dynamics.

FIGURE VII

Average Time-Varying Adjustment Function



Given that our data span a long period of time, we ideally would like to examine yearly variations in the adjustment function. Unfortunately, the number of observations in each year is not sufficient to provide a reliable estimate of the adjustment function. Hence, we compute the average adjustment function for four different subperiods of our sample: 1981-1984, 1985-1988, 1989-1992, and 1993-1997. Even with this grouping, the estimates of the adjustment function across different subperiods are likely to be imprecise, much as in the case of seasonal adjustment.

Despite this caveat, Figure VII indicates that the adjustment function has been remarkably stable across the four subperiods. In each sector, the adjustment function displays a similar nonlinear shape – consistent with the use of (S,s) -type inventory policies – in each of these periods. The most pronounced time variation in the adjustment function occurs during the latest subperiod (1993-1997) in the manufacturing sector, where the adjustment rates associated with inventory shortages appear somewhat higher.

The relative time invariance of the adjustment function indicates that the inventory adjustment process has not changed significantly over our sample period. Furthermore, the stability of the adjustment function suggests that the aggregate inventory fluctuations are driven primarily by movements in the cross-sectional distribution of inventory deviations rather than by shifts in the adjustment function. This result is in contrast to the popular view, which argues that advances in inventory control practices have reduced the volatility of inventory investment and its contribution to the overall output fluctuations.¹⁸

5.1.5 *Distribution of adjustments*

Figures VIII and IX show the distribution of adjustment rates conditional on different ranges of inventory deviations for the manufacturing (Figure VIII) and trade (Figure IX) sectors. The horizontal axis in each panel represents the fraction of the inventory deviation closed by a firm in any given quarter. We use data pooled over all firms and quarters to generate these conditional distributions. For example, the bar in the bottom panel of Figure VIII at the value 0.1 on the horizontal axis represents the fraction of firm/quarter observations with an inventory deviation index in the interval $[0.1,0.3]$ that closed 10 percent of their inventory gap. The upper panel in each figure corresponds to a situation where firms have “large” excessive holdings of inventories ($z \in [-0.5,-0.3]$); the middle panel corresponds to “large” inventory shortages ($z \in [0.3,0.5]$), and the bottom panel corresponds to a situation of a “normal” inventory shortage ($z \in [0.1,0.3]$).

There are a couple of noteworthy features in the conditional distribution of adjustments. First, even though there is a wide dispersion in the distribution of adjustment rates – in particular, there is a fraction of firms which overshoot their target ($7(z) > 1.0$) – most firms have adjustment rates lying between 0 and 1.0, which is the range consistent with the model. This occurs despite the potential distortions induced by entry into and exit from our sample.

Second, a sizable fraction of firms exhibits no adjustment, indicating substantial inaction in the quarterly adjustment that is consistent with the use of (S,s) -type inventory policies. In particular, the

¹⁸ Our results are consistent with Allen (1995) and Filardo (1995) who find that the net effect of improved inventory control methods on the business cycle is negligible.

mode of the conditional distribution for manufacturing firms with large inventory overhangs is at zero (top panel), which suggests that a large fraction of these firms does not reduce their inventory holdings. This result is particularly at odds with the partial-adjustment model, because for most firms inventories are a relatively small fraction of sales at the quarterly frequency, and we would expect these firms to reduce their excess inventory levels. Consequently, the evidence of the inertia in the inventory adjustment process is indicative of nonconvexities in the production technology, which results in infrequent production adjustment.

FIGURE VIII
Distribution of Estimated Adjustment Rates
Manufacturing Sector

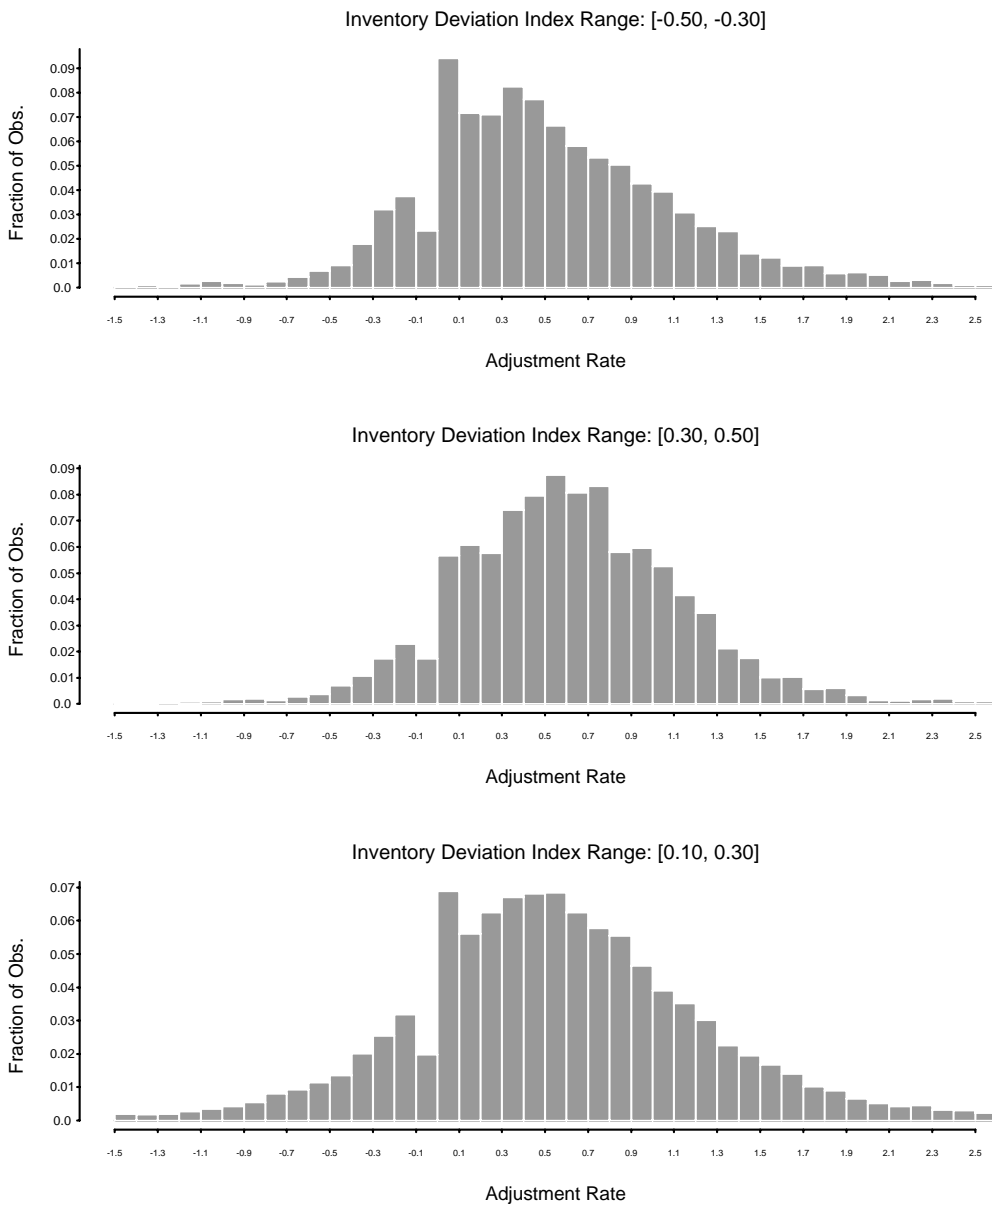
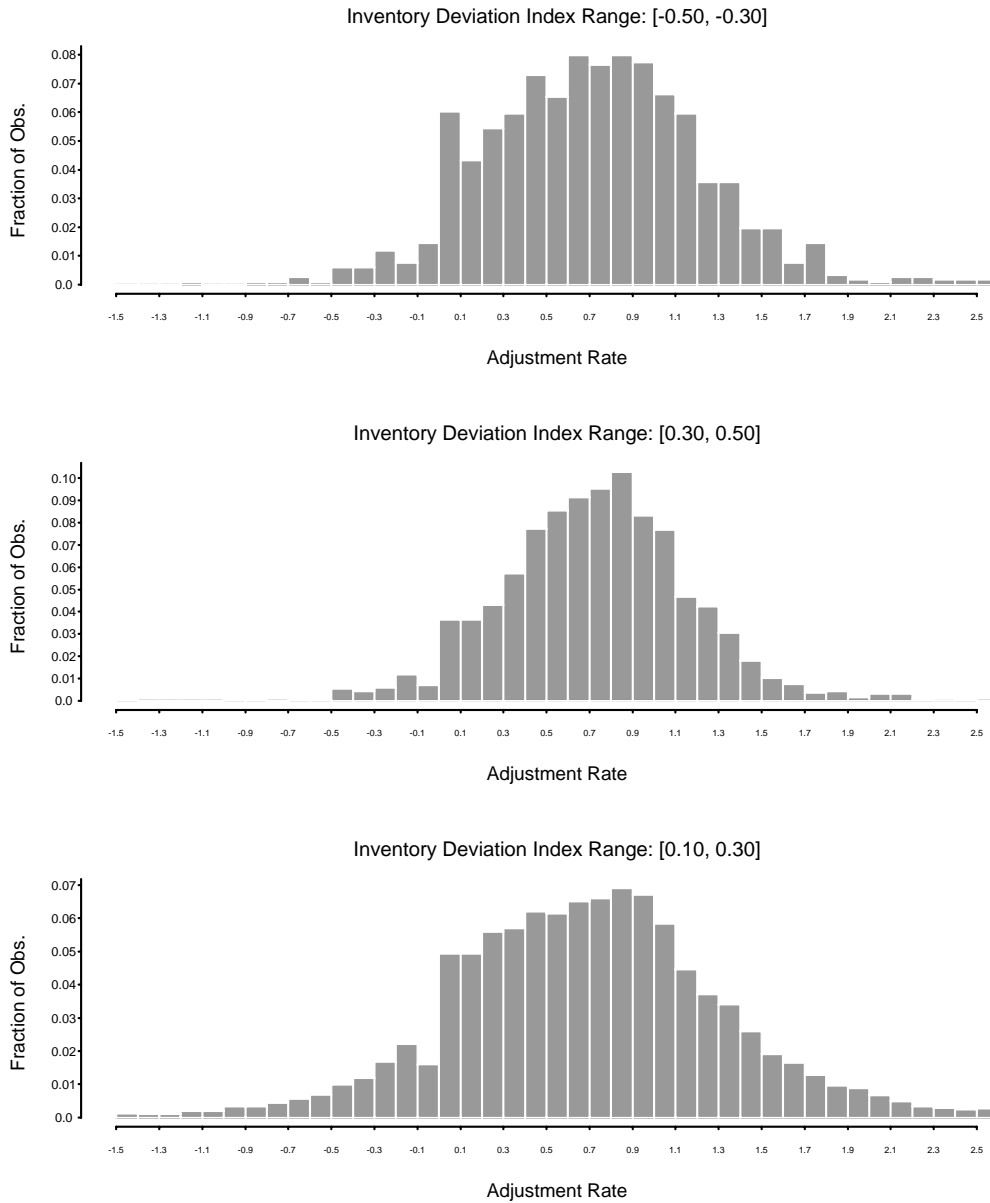


FIGURE IX
 Distribution of Estimated Adjustment Rates
 Trade Sector



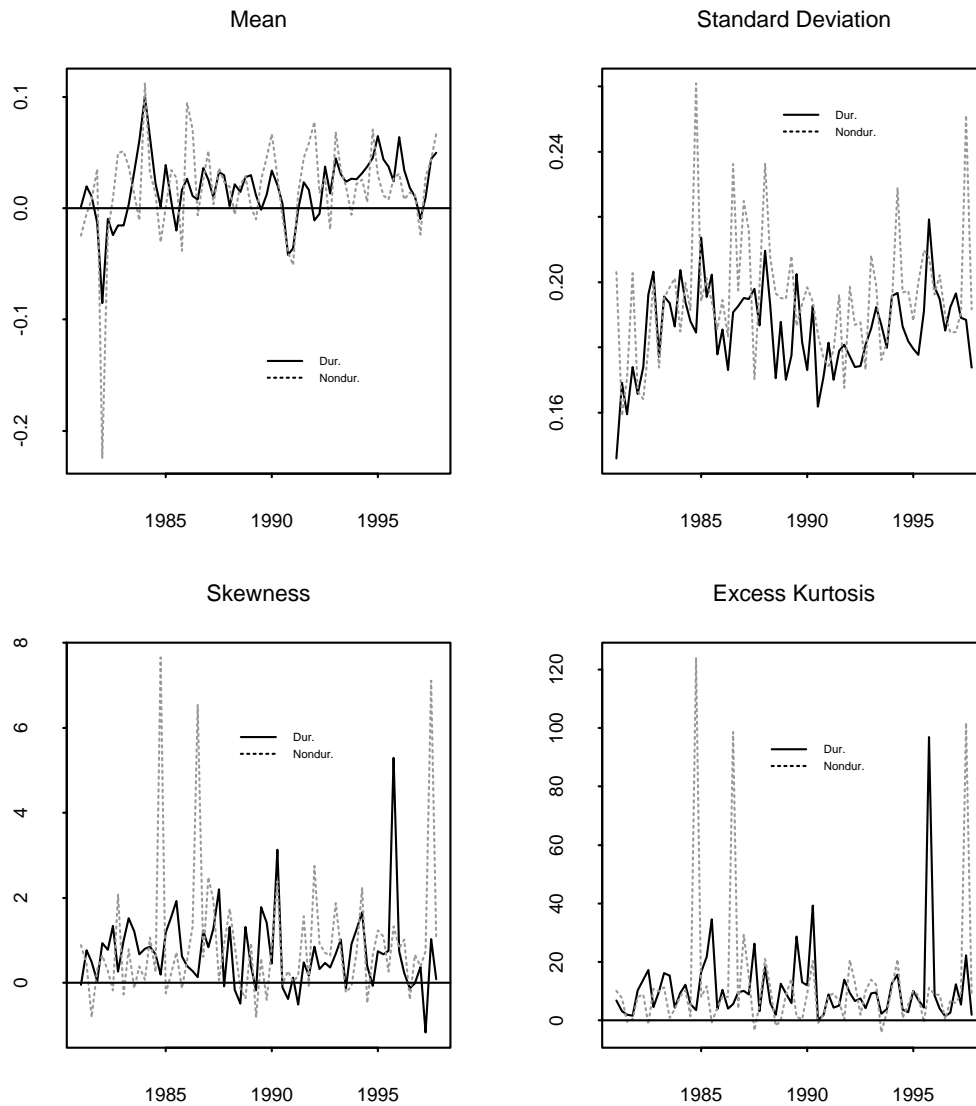
5.2 The cross section of inventory deviations

The cross section of inventory deviations results from the interaction between aggregate and idiosyncratic shocks and the microeconomic adjustment process. The density function of this cross section corresponds to the histogram of deviations in each period. The average densities – where the average is computed over all firm/quarter observations in the manufacturing and the trade sectors – are displayed by the dashed lines in Figure IV. In this section, we consider the time variation of these distributions and the evolution of the shocks affecting inventory deviations.

FIGURE X

Cross-Sectional Moments of Inventory Deviations

Manufacturing Sector



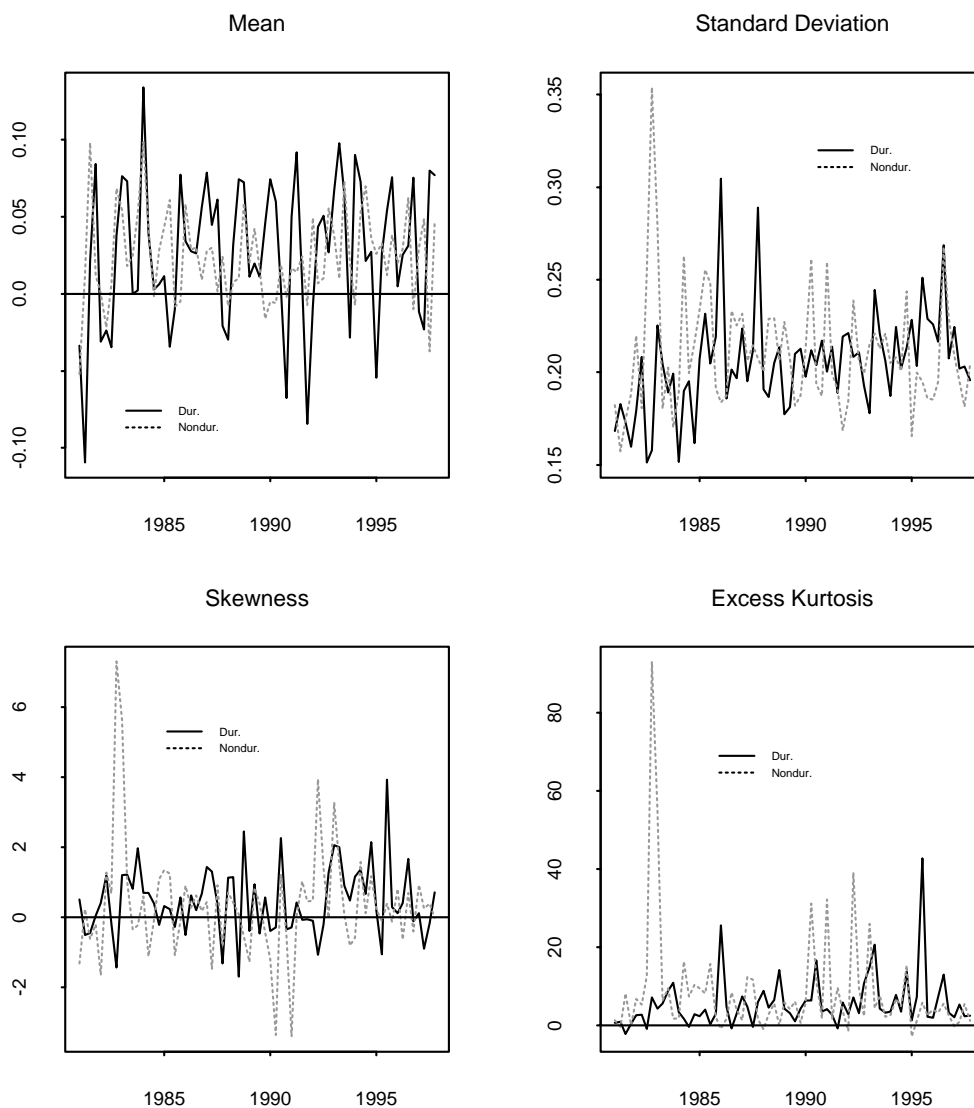
The time paths of the first four moments of the cross-sectional distribution of the inventory deviations are presented in Figures X and XI for the primary subcomponents of the manufacturing and trade sectors, respectively.¹⁹ Each moment displays substantial temporal variation and does not appear to be correlated with other moments. In particular, the implied cross-sectional distributions of inventory

¹⁹ All moments are unweighted and have been seasonally adjusted with quarterly dummies; excess kurtosis is defined as the difference between 3, the kurtosis of the normal distribution, and the sample kurtosis.

deviations exhibit both substantial skewness and kurtosis; interestingly, these higher moments are more volatile in nondurable goods industries – for example, note the skewness and kurtosis in the manufacturing sector.

FIGURE XI
Cross-Sectional Moments of Inventory Deviations

Trade Sector



The visual evidence of the time variation in the cross-sectional density of inventory deviations is confirmed by formal statistical evidence. Figure XII presents a histogram of significance levels from a sequence of nonparametric tests for the equality of cross-sectional distributions over time in

Figure XII
 Stability Tests of the Cross-sectional Distribution








FIGURE XIII

Moments of Inventory Deviations and Shocks

Manufacturing Sector




manufacturing and trade sectors. Specifically we compute the probability value for the Kolmogorov-Smirnov test statistic (see DeGroot (1975)) of the null hypothesis that $f(z,t) = f(z,t+1)$, against the two-sided alternative that $f(z,t) \neq f(z,t+1)$, for all t . Note that for most of our sample period, we can reject the null hypothesis that the cross-sectional distribution of inventory deviations is stable across adjacent periods at any of the usual significance levels.

Finally, note that the dispersion of inventory deviations – as measured by the standard deviation of the cross-sectional distribution – does not exhibit a secular decline during our sample period. Such a

FIGURE XIV

Moments of Inventory Deviations and Shocks

Trade Sector



decline would be consistent with advances in inventory management techniques, which, presumably, allow firms to align inventories more closely to their desired levels.

Figures XIII and XIV plot the evolution of the aggregate and idiosyncratic shocks computed from equation (2) using the path of the estimated inventory deviations. The time path of aggregate shocks is computed according to $\eta_t = \frac{1}{N_t} \sum_{i=1}^{N_t} (\Delta z_{it} + \Delta h_{it-1})$, where N_t denotes the number of firms in period t .

In each period, the distribution of idiosyncratic shocks corresponds to the histogram of estimated \langle_{it} s, which are obtained from $\langle_{it} = (z_{it+1} - h_{it-1}) - 0_t$. The left and right top panels of each figure show the time

path of the cross-sectional mean of the inventory deviation index z (solid line) and the estimated time path of the aggregate shock 0_t (dashed line) for the primary subcomponents of the two sectors. The left and right bottom panels show the temporal variation in the standard deviation (left panel) and skewness (right panel) of the idiosyncratic shocks in each of the primary subcomponents of the two sectors.

Note that the time path of aggregate shocks in the manufacturing sector is closely correlated with movements in the mean of the cross-sectional distribution of inventory deviations; the correlation between the two series in the durables goods component of the manufacturing sector is 0.85, and the correlation in the nondurable goods sector is 0.83. In the trade sector, however, the same correlations are only 0.66 for the durables and 0.67 for the nondurable goods sector. Consequently, fluctuations in the cross-sectional density of inventory deviations in the trade sector could be driven to a greater extent by idiosyncratic shocks. Note that in both the manufacturing and trade sectors, the cross-sectional distribution of idiosyncratic shocks exhibits substantial dispersion and skewness. Moreover, both moments display considerable temporal variation.

5.3 Aggregate dynamics and cross-sectional moments

In this section, we consider the impact of nonlinear microeconomic adjustment on the aggregate inventory dynamics. We use a simple parametric approach to analyze whether higher moments of the cross-sectional density affect aggregate dynamics. The density functions pictured in Figures IV indicate that for a significant fraction of our firm/quarter observations, inventories are within $\forall 25$ percent of their desired level. More importantly, the distribution of inventory deviations is concentrated in the interval over which the adjustment function is particularly nonlinear. This concentration suggests that – with the proper specification of the aggregate adjustment function – cross-sectional moments other than the mean should affect aggregate dynamics.

Following Caballero and Engel (1993), we test this hypothesis as follows. Suppose that the adjustment function can be approximated by a second-degree polynomial, $\gamma(z)=\delta_0+\delta_1z+\delta_2z^2$. From equations (3) and (4) it follows that the aggregate dynamics will depend on the first three noncentral moments of the cross-sectional density $f(z,t)$. Expansion of these noncentral moments in terms of mean and central moments implies,

$$(12) \quad \Delta \ln H_t^A = \lambda_0 \mu_z(t) + \lambda_1 \left[\sigma_z^2(t) + \mu_z^2(t) \right] + \lambda_2 \left[\sigma_z^3(t) \gamma_z(t) + 3\sigma_z^2(t) \mu_z(t) + \mu_z^3(t) \right]$$

where μ_z denotes the mean of the cross-sectional density $f(z,t)$, σ_z the standard deviation, and γ_z the skewness coefficient. The constant adjustment function implicit in the partial-adjustment model implies that $\delta_1=\delta_2=0$, while a model with a nonconstant adjustment function implies nonzero coefficients on the higher moments.

Estimates of equation (12) as well as the partial-adjustment model are presented in Tables 1 and 2 for the manufacturing and trade sectors respectively. The first two columns in each table report the results for the durable goods component and the last two columns for the nondurable goods sector. The equations for durable and nondurable goods in each sector are estimated jointly in a SUR framework using least squares over the period 1981Q1–97Q4.²⁰

Turning first to the manufacturing sector, the estimates of the constant adjustment rate in the partial adjustment model are presented in the first and third columns of Table 1. Both estimates are statistically highly significant and indicate quarterly adjustment rates of 0.632 in the durable goods manufacturing and 0.339 in the nondurable goods sector. Although the adjustment rate in the durables sector is not unreasonable, the adjustment rate in the nondurables sector appears to be too low, especially compared to the adjustment function we previously estimated for this sector (Figure V).

Table 1
Aggregate dynamics and cross-sectional moments

$\Lambda(z)$	Mfg. durables		Mfg. Nondurables	
	P.A.M. ^a	Quadratic ^b	P.A.M.	Quadratic
λ_0	0.632 (0.049)	0.523 (0.052)	0.339 (0.048)	0.386 (0.047)
λ_1	–	0.138 (0.152)	–	0.404 (0.202)
λ_2	–	0.545 (0.170)	–	0.015 (0.097)
<i>SSE</i>	0.013	0.012	0.029	0.028
<i>p</i> -value ^c	–	0.000	–	0.041
<i>DW</i> (1) ^d	0.145	0.237	0.087	0.119
<i>DW</i> (2)	0.329	0.220	0.078	0.028
<i>DW</i> (3)	0.193	0.176	0.097	0.121
<i>DW</i> (4)	0.586	0.567	0.045	0.263
<i>S</i> - <i>W</i> ^e	0.006	0.001	0.653	0.003

Notes: Estimation period: 1981Q1-97Q4 ($T=68$). All data are seasonally unadjusted. All specifications include a constant term (not reported). The dependent variable is the growth rate of (real) average inventories. The sector-specific equations are estimated jointly in a SUR framework using OLS. Asymptotic standard errors are reported in parentheses.

^a Partial-adjustment model: $\Lambda(z) = \lambda_0$.

^b Quadratic adjustment function: $\Lambda(z) = \lambda_0 + \lambda_1 z + \lambda_2 z^2$.

^c Probability value for the Wald test of the null hypothesis that $\lambda_1 = \lambda_2 = 0$.

^d $DW(p)$ indicates the probability value for the Durbin-Watson test of the p th-order serial correlation.

^e Probability value for the Shapiro-Wilk test of the normality of the residuals.

Estimates of the quadratic adjustment function in the second and fourth columns provide additional evidence that the adjustment function is not constant. In the durable goods sector, the null hypothesis

²⁰ Because the regressors are derived from an auxiliary econometric model, they are measured with sampling error – the so-called “generated regressor” problem. Consequently, although the parameter estimates are consistent, hypothesis tests based on the estimated covariance matrix are biased.

of $\delta_1=\delta_2=0$ is rejected at any of the usual significance levels. In particular, the quadratic term, δ_2 , is highly significant, which is indicative of (S,s) -type behavior in this sector. In the nondurable goods sector, the null hypothesis of a constant adjustment function is also rejected, although the rejection is not as strong as that in the durable goods component. In this case, the linear term, δ_1 , is statistically significant but the quadratic term, δ_2 , is not statistically different from zero. Furthermore, the implied shape of the quadratic adjustment function for the nondurable goods sector indicates a likely misspecification.

The results for the trade sector are presented in Table 2. For the partial adjustment model, the estimated constant adjustment rate is 0.245 in the durable goods sector and 0.730 for nondurables – both estimates are statistically significant. The implied adjustment rate in the durable goods sector appears to be unreasonably low, compared both to economic intuition and our nonparametric estimates of the adjustment function from the firm-level data (Figure V).

Table 2
Aggregate dynamics and cross-sectional moments

$\Lambda(z)$	Trade durables		Trade nondurables	
	P.A.M. ^a	Quadratic ^b	P.A.M.	Quadratic
λ_0	0.245 (0.065)	0.105 (0.085)	0.730 (0.058)	0.748 (0.059)
λ_1	–	0.205 (0.296)	–	–0.197 (0.209)
λ_2	–	0.640 (0.255)	–	–0.037 (0.073)
<i>SSE</i>	0.059	0.052	0.058	0.054
<i>p-value</i> ^c	–	0.033	–	0.253
<i>DW(1)</i> ^d	0.203	0.118	0.937	0.950
<i>DW(2)</i>	0.263	0.155	0.709	0.496
<i>DW(3)</i>	0.870	0.934	0.970	0.969
<i>DW(4)</i>	0.000	0.003	0.000	0.000
<i>S-W</i> ^e	0.794	0.001	0.915	0.994

Notes: Estimation period: 1981Q1-97Q4 ($T=68$). All data are seasonally unadjusted. All specifications include a constant term (not reported). The dependent variable is the growth rate of (real) average inventories. The sector-specific equations are estimated jointly in a SUR framework using OLS. Asymptotic standard errors are reported in parentheses.

^a Partial-adjustment model: $\Lambda(z) = \lambda_0$.

^b Quadratic adjustment function: $\Lambda(z) = \lambda_0 + \lambda_1 z + \lambda_2 z^2$.

^c Probability value for the Wald test of the null hypothesis that $\lambda_1 = \lambda_2 = 0$.

^d $DW(p)$ indicates the probability value for the Durbin-Watson test of the p th-order serial correlation.

^e Probability value for the Shapiro-Wilk test of the normality of the residuals.

However, the parametric estimates of the quadratic adjustment function in the trade sector provide only limited evidence that the adjustment function is nonlinear. In the durable goods sector, the null hypothesis of a constant adjustment function is rejected at better than a 5-percent significance level. As in the manufacturing durables sector, it is the quadratic term, δ_2 , that is most significant, indicating

(S,s) -type behavior in the sector; the implied adjustment rates for this sector, however, remain rather low. In the nondurable goods sector, the higher moment terms have little effect and the constant adjustment function – consistent with the partial adjustment model – cannot be rejected.²¹

Overall, our parametric approach provides some support for the hypothesis that the higher cross-sectional moments of the inventory deviation distribution influence the dynamics of average inventory investment, especially for the durable goods industries. Nevertheless, it appears that the influence of these higher moments on average inventory growth is more complicated than that implied by a time-invariant second-order polynomial approximation of the adjustment function.

6. Conclusion

In this paper, we have examined aggregate inventory dynamics by building directly on the microeconomic evidence. Our empirical methodology employs high-frequency firm-level data and the generalized (S,s) framework developed Caballero (1993, 1997), Caballero and Engel (1991, 1993) and Caballero, Engel, and Haltiwanger (1995, 1997). The combination of these two factors allows us to model explicitly the underlying microeconomic heterogeneity and provides us with an empirically tractable aggregation methodology. One contribution of our paper is a unified approach in modelling the behavior of manufacturing and trade inventories, providing a framework that is consistent with numerous empirical regularities of the data.

The key finding of our paper is that the estimated adjustment function is nonlinear and asymmetric, implying a pattern of microeconomic inventory adjustment consistent with the use of (S,s) -type inventory policies. The nonlinearity of the adjustment function is reflected in the fact that firms with larger deviations from “desired” inventory levels adjust proportionally more than do firms with smaller deviations – behavior consistent with the presence of nonconvexities in the adjustment technology. The asymmetry, especially evident in the manufacturing sector, reflects the fact that firms with excessive inventory holdings adjust less than firms with similarly sized inventory shortages. In part, these asymmetries are consistent with a presence of a stockout avoidance motive. Furthermore, we find sizable differences between the adjustment functions of manufacturing and trade firms as well as between those of firms in the durable and nondurable goods industries in each sector.

The seasonal variations in the adjustment function display interesting features, notably in the trade sector, where there is evidence consistent with the Christmas selling season and year-end effects. More importantly, we find that the adjustment function is remarkably stable across different subperiods of

²¹ Note that the first moment of the cross-sectional distribution, $\mu_x(t)$, includes a period t aggregate shock and hence uses more information than the typical aggregate time series partial adjustment specification.

our sample. This result suggests that the inventory adjustment process has changed little during our sample period, despite a popular view that recent developments in inventory management techniques have allowed firms to adjust faster and more completely. The time invariance of the adjustment function – combined with our evidence that shows that the dispersion of the inventory deviations has not declined – indicates that the inventory cycle has yet to be tamed. This further implies that movements in the cross-sectional distribution of inventory deviations are likely to play a key role in determining aggregate business cycle fluctuations.

In particular, the nonlinearities and asymmetries in the inventory adjustment function would imply that the higher cross-sectional moments of inventory deviations affect aggregate dynamics. In contrast to the implied dynamics of the partial-adjustment model, a simple parametric approach employed in this paper suggests that the higher moments of the cross-sectional distribution of inventory deviations have some explanatory power for aggregate inventory investment, especially for the durable goods industries.

Overall, our results strongly support the applicability of the generalized (S,s) framework in connecting microeconomic inventory policies to aggregate dynamics. However, much work on the aggregate implications of our model remains to be done. A limitation of our model is that it is built around a single state variable, the firm's inventory deviation from its desired level. Although this state variable may be a sufficient statistic for many aspects of the firm's inventory policy, it is probably insufficient to capture all relevant considerations. In particular, the financial condition of firms has been shown to have significant effects on cyclical inventory fluctuations in all sectors; see, for example, Gertler, and Gilchrist (1994), Fazzari, Hubbard, and Petersen (1997), and Zakrajšek (1997). An index of the firm-level financial capacity, therefore, would appear to be a particularly promising candidate to be included in the model.

A. Data appendix

This section describes the selection rules used to construct our firm-level panels and the construction of all variables in the analysis. The data for our paper come from the P/S/T, Full Coverage, and Research COMPUSTAT data files. As we discussed in the main text, firms in SIC 29 (Petroleum refining and related industries) and SIC 55 (Automotive dealers and gasoline service stations) were excluded from the analysis, because these two industries contained an insufficient number of observations for reliable estimation.

The COMPUSTAT data files are compiled in a fiscal-year format. The fiscal quarters in the data are aligned with calendar quarters as follows:

1. If the firm's fiscal year ends in the same month as a calendar quarter, the adjustment is straightforward, as the fiscal quarters are relabeled to correspond to calendar quarters.
2. If the firm's fiscal-year end does not coincide with the end of a calendar quarter, the data are adjusted so that the majority of the fiscal quarter is placed into the appropriate calendar quarter.

A.1 Selection rules

We selected all firms with positive total inventories, positive net sales, positive total assets, and with at least eight continuous quarters of data between 1980Q4–97Q4. To avoid results that are driven by a small number of extreme observations, two criteria were used to eliminate firms with substantial outliers or obvious errors:

1. If a firm's estimate of (real) output/deliveries from the accounting identity $Y/S+\Delta H$ was negative at any point during a firm's tenure in the sample, a firm was eliminated in its entirety.²²
2. If a firm's inventory-sales ratio was above (below) the 99.0 (1.0) percentile of the distribution in any period during the firm's tenure in the panel, the firm was eliminated in its entirety.

As a consequence of these selection rules, 798 firms were eliminated from the manufacturing panel, and 216 firms were eliminated from the trade panel.²³

²² Note that the accounting identity, $Y \equiv S + \Delta H$, holds only for finished goods inventories. Because the quarterly COMPUSTAT data report only the dollar value of total inventory stocks, this relationship does not hold exactly in our data. Nevertheless, the results reported in this paper are virtually identical when we include firms that violate this "accounting identity."

A.2 Construction of variables

- **Inventories:** The COMPUSTAT data report the book value of total inventories. Because the firm-level COMPUSTAT data provide limited and incomplete information on the inventory accounting practices, we assumed that all inventory stocks are evaluated using the FIFO method; namely, once a finished good is placed on shelves, it is given a price tag that remains on the item regardless of what subsequently happens to the price of newly produced goods.²⁴ This implies that the replacement value of inventory stocks equals their book value. To convert nominal reported value of inventories to real terms, inventory stocks were deflated by the sector-specific (i.e., durable and nondurable) implicit (1992=100) inventory deflator.
- **Net Sales:** To construct a real measure of sales, the reported nominal value of sales was deflated by the sector-specific (i.e., durable and nondurable) implicit (1992=100) sales deflator.
- All other variables were deflated by the chain-weighted (1992=100) GDP deflator.

A.3 Data summary

Tables A.1-A.4 provide information on the two-digit SIC composition and summary statistics for the two data sets used in our analysis. Although each two-digit industry has at least 1,100 observations, the coverage of the manufacturing sector is considerably more extensive, especially in SICs 28, 34, 35, and 38. This is due to the fact that, relative to the trade sector, a greater fraction of manufacturing firms is publicly traded and thus included in the COMPUSTAT data files.

Because all firms in the sample are publicly traded, most of them are relatively large. In the manufacturing sector, the median firm size, measured by total assets, is almost \$79 million, while in the trade sector, the median firm size is almost \$115 million. The distributions of most variables display considerable skewness – the means of inventories, sales, (approximate) output, and assets are much greater than the medians. The distribution of the inventory-sales ratio, on the other hand, is considerably more symmetric. Also note that even after excluding outliers there remains a great deal of heterogeneity in inventory investment and in the growth of inventories and sales.

Finally, as mentioned in the text, the panels are unbalanced, with firms entering and exiting the data sets. In the manufacturing sector, the average tenure is 30 quarters and the median is 25 quarters. For the trade sector, the average tenure in the panel is 27 quarters and the median is 21 quarters.

²³ Over 3/4 of eliminated firms were deleted because of the second selection criterion. Visual inspection of the eliminated firms revealed severe anomalies and likely errors in their reported data.

²⁴ It has been noted that firms are most likely to use LIFO accounting to value their inventory stocks – an accounting method that creates a wedge between the reported book value and the replacement value of inventory stocks – during times of high inflation. Given that our samples spans the time period of low and stable inflation, the assumption that firms use FIFO accounting seems reasonable.

Table A.1

Industry composition: manufacturing sector

Industry classification	No. of firms	Observations
SIC 20: Food & Kindred Prod.	212	6,494
SIC 22: Textile Mill Prod.	107	2,976
SIC 23: Apparel & Other Prod.	108	2,820
SIC 24: Lumber & Wood Prod.	69	2,230
SIC 25: Furniture & Fixtures	74	2,404
SIC 26: Paper & Allied Prod.	111	3,735
SIC 27: Printing & Publishing	116	3,876
SIC 28: Chemicals & Allied Prod.	397	12,497
SIC 30: Rubber & Plastic Prod.	153	4,396
SIC 31: Leather & Leather Prod.	39	1,374
SIC 32: Stone, Clay & Glass Prod.	79	2,079
SIC 33: Primary Metal Industries	161	4,941
SIC 34: Fabricated Metal Prod.	182	5,614
SIC 35: Industrial Machinery	654	19,206
SIC 36: Electronic Equip.	636	19,293
SIC 37: Transportation Equip.	198	5,929
SIC 38: Instruments & Related Prod.	527	15,752
SIC 39: Misc. Durables	123	3,269

Table A.2

Summary statistics: manufacturing sector

Variable	Mean	Std. Dev.	Median	Minimum	Maximum
Inventories	134.20	519.40	16.00	0.003	12414.90
Net Sales	248.20	1147.60	25.50	0.005	39335.00
Gross Output	253.10	1164.00	26.30	0.001	38280.70
Total Assets	1013.90	6243.20	78.80	0.107	259303.00
Inv/Sales Ratio	0.74	0.46	0.64	0.070	5.20
Inv. Investment	1.07	49.10	0.08	-2444.70	3077.60
Inv. Growth Rate (%)	2.00	18.00	1.30	-242.50	451.00
Sales Growth Rate (%)	1.90	24.10	2.11	-331.10	477.00
No. of Firms	3,946				
Tenure (Quarters)	30.1	18.6	25.0	8.0	69.0
Observations	118,885				

Notes: Sample period: 1984:Q4-1997:Q4. All variables are in millions of 1992 dollars.

Table A.3

Industry composition: trade sector

Industry classification	No. of firms	Observations
SIC 50: Wholesale Durables	296	8,000
SIC 51: Wholesale Nondurables	151	4,052
SIC 52: Building Materials	38	1,145
SIC 53: General Merch. Stores	99	3,179
SIC 54: Food Stores	98	3,082
SIC 56: Apparel & Accessory Stores	84	2,533
SIC 57: Furniture & Home Furnishings	80	2,152
SIC 58: Eating & Drinking Places	171	4,035
SIC 59-1: Misc. Retail Durables	72	1,805
SIC 59-2: Misc. Retail Nondurables	117	2,683

Table A.4

Summary statistics: trade sector

Variable	Mean	Std. Dev.	Median	Minimum	Maximum
Inventories	165.30	599.70	28.70	0.007	18369.00
Net Sales	341.40	1133.70	61.00	0.017	34571.50
Gross Output	351.70	1155.00	63.90	0.013	31885.70
Total Assets	681.90	3686.70	114.60	0.047	108274.70
Inv/Sales Ratio	0.57	0.43	0.50	0.020	3.01
Inv. Investment	3.38	90.20	0.11	-3065.60	3110.90
Inv. Growth Rate (%)	2.55	21.80	2.03	-331.20	529.80
Sales Growth Rate (%)	2.30	25.70	2.31	-203.40	331.60
No. of Firms	1,206				
Tenure (Quarters)	27.1	17.6	21.0	8.0	69.0
Observations	32,666				

Notes: Sample period: 1984:Q4-1997:Q4. All variables are in millions of 1992 dollars.

References

- Abramowitz, Moses (1950): *Inventories and Business Cycles*. New York: National Bureau of Economic Research.
- Allen, Donald S (1995): "Changes in Inventory Management and the Business Cycle". *Review, Federal Reserve Bank of St. Louis*, 4, pp. 17-33.
- Blanchard, Olivier J (1983): "The Production and Inventory Behavior of the American Automobile Industry". *Journal of Political Economy*, 91, pp. 365-400.
- Blinder, Alan S (1981): "Retail Inventory Behavior and Business Fluctuations". *Brookings Papers on Economic Activity*, 2, pp. 443-505.
- Blinder, Alan S (1986): "Can the Production Smoothing Model of Inventories be Saved?" *Quarterly Journal of Economics*, 101, pp. 431-54.
- Blinder, Alan S and Louis J Maccini (1991): "Taking Stock: A Critical Assessment of Recent Research on Inventories". *Journal of Economic Perspectives*, 5, pp. 73-96.
- Caballero, Ricardo J (1993): "Durable Goods: An Explanation for Their Slow Adjustment". *Journal of Political Economy*, 101, pp. 351-84.
- Caballero, Ricardo J (1997): "Aggregate Investment". Forthcoming in the *Handbook of Macroeconomics*.
- Caballero, Ricardo J and Eduardo M R A Engel (1991): "Dynamic (S,s) Economies". *Econometrica*, 59, pp. 1659-86.
- Caballero, Ricardo J and Eduardo M R A Engel (1993): "Microeconomic Adjustment Hazards and Aggregate Dynamics". *Quarterly Journal of Economics*, 108, pp. 359-83.
- Caballero, Ricardo J, Eduardo M R A Engel and John C Haltiwanger (1995): "Plant-Level Adjustment and Aggregate Investment Dynamics". *Brookings Papers on Economic Activity*, 2, pp. 1-54.
- Caballero, Ricardo J, Eduardo M R A Engel and John C Haltiwanger (1997): "Aggregate Employment Dynamics: Building from Microeconomic Evidence". *American Economic Review*, 87, pp. 115-37.
- Caplin, Andrew S (1985): "The Variability of Aggregate Demand with (S,s) Inventory Policies". *Econometrica*, 53, pp. 1395-409.
- Carpenter, Robert E, Steven M Fazzari and Bruce C Petersen (1997): "Financing Constraints and Inventory Investment: A Comparative Study with High-Frequency Panel Data". Unpublished paper, Department of Economics, Washington University.
- DeGroot, Morris H (1975): *Probability and Statistics*. Reading, Massachusetts: Addison-Wesley.
- Episcopos, Athanasios (1996): "Testing the (S,s) Model of Inventory Investment with Canadian Wholesale Trade Data". *Applied Economics Letters*, 3, pp. 193-5.
- Fair, Ray C (1989): "Production Smoothing Model is Alive and Well". *Journal of Monetary Economics*, 24, pp. 353-70.
- Filardo, Andrew J (1995): "Recent Evidence on the Muted Inventory Cycle". *Federal Reserve Bank of Kansas City Economic Review*, 2, pp. 27-43.
- Fisher, Jonas D M and Andreas Hornstein (1995): " (S,s) Inventory Policies in General Equilibrium". *Research Report*, 9514, Department of Economics, The University of Western Ontario.
- Gertler, Mark and Simon Gilchrist (1994): "Monetary Policy, Business Cycles and the Behavior of Small Manufacturing Firms". *Quarterly Journal of Economics*, 109, pp. 309-40.
- Holt, Charles C, Franco Modigliani, John F Muth and Herbert A Simon (1960): *Planning Production, Inventories and Work Force*. Englewood Cliffs, NJ: Prentice-Hall.

- Irvine, F Owen (1981a): "The Influence of Capital Costs on Inventory Investment: Time-Series Evidence for a Department Store". *Quarterly Review of Economics and Business*, 21, pp. 25-44.
- Irvine, F Owen (1981b): "Retail Inventory Investment and the Cost of Capital". *American Economic Review*, 71, pp. 633-48.
- Kahn, James A (1987): "Inventories and the Volatility of Production". *American Economic Review*, 77, pp. 667-79.
- Krane, Spencer D (1994): "The Distinction Between Inventory Holding and Stockout Costs: Implications for Target Inventories, Asymmetric Adjustment, and the Effect of Aggregation on Production Smoothing". *International Economic Review*, 35, pp. 117-36.
- Krane, Spencer D and Steven N Braun (1991): "Production Smoothing Evidence from Physical Product Data". *Journal of Political Economy*, 99, pp. 558-81.
- McCarthy, Jonathan and Egon Zakrajšek (1998): "Trade Inventories". Federal Reserve Bank of New York, *Staff Report*, No. 53.
- McConnell, Margaret M and Gabriel Perez-Quiros (1998): "Output Fluctuations in the United States: What has Changed Since the Early 1980s?" Federal Reserve Bank of New York, *Staff Report*, No. 41.
- Metzler, Lloyd A (1941): "The Nature and Stability of Inventory Cycles". *Review of Economics and Statistics*, 23, pp. 113-29.
- Mosser, Patricia C (1988): "Empirical Test of the (S,s) Model for Merchant Wholesalers," in A Chikan and M C Lovell (eds.) *The Economics of Inventory Management*, Amsterdam, The Netherlands: Elsevier Science Publishers, pp. 261-83.
- Mosser, Patricia C (1991): "Trade Inventories and (S,s) ". *Quarterly Journal of Economics*, 106, pp. 1267-86.
- Ramey, Valerie A and Kenneth D West (1997): "Inventories". National Bureau of Economic Research Working Paper No. 6315. Forthcoming in the *Handbook of Macroeconomics*.
- Rotemberg, Julio J (1987): "The New Keynesian Microfoundations," in O J Blanchard and S Fischer (eds.) *NBER Macroeconomics Annual*, Cambridge, MA: The MIT Press, pp. 69-116.
- Scarf, Herbert (1960): "The Optimality of (S,s) Policies in the Dynamic Inventory Problem", in K J Arrow, S Karlin and H Scarf (eds.) *Mathematical Methods in the Social Sciences*, Stanford, CA: Stanford University Press, 1960, pp. 196-202.
- Schuh, Scott (1996): "Evidence on the Link Between Firm-Level and Aggregate Inventory Behavior". Unpublished paper, Board of Governors of the Federal Reserve System.
- Trivedi, Peter K (1973): "Retail Inventory Investment Behavior". *Journal of Econometrics*, 1, pp. 61-80.
- Worthington, Paula R (1998): "The Increasing Importance of Retailers' Inventories". *Federal Reserve Bank of Chicago Economic Perspectives*, 3, pp. 2-12.
- Zakrajšek, Egon (1997): "Retail Inventories, Internal Finance, and Aggregate Fluctuations: Evidence from U.S. Firm-Level Data". Federal Reserve Bank of New York, *Research Paper*, No. 9722.

Recent BIS Working Papers

No.	Title	Author
47 September 1997	Financial asset prices and monetary policy: theory and evidence	Frank Smets
48 September 1997	Some multi-country evidence on the effects of real exchange rates on output	Steven B Kamin and Marc Klau
49 September 1997	Why does the yield curve predict economic activity? Dissecting the evidence for Germany and the United States	Frank Smets and Kostas Tsatsaronis
50 November 1997	The euro and the dollar	Robert N McCauley
51 November 1997	Forecast errors and financial developments	Palle S Andersen
52 January 1998	Inflation and disinflation in Iceland	Palle S Andersen and Már Guðmundsson
53 March 1998	Exchange rate regimes and inflation and output in Sub-Saharan countries	Marc Klau
54 June 1998	The coming transformation of continental European banking?	William R White
55 June 1998	Spread overreaction in international bond markets	Gregory D Sutton
56 June 1998	Commercial banks in the securities business: a review	João A C Santos
57 July 1998	One-step prediction of financial time-series	Srichander Ramaswamy
58 November 1998	The importance of bank seniority for relationship lending	Stanley D Longhofer and João A C Santos
59 November 1998	Portfolio selection using fuzzy decision theory	Srichander Ramaswamy
60 November 1998	Output gap uncertainty: does it matter for the Taylor rule?	Frank Smets
61 November 1998	Foreign direct investment and employment in the industrial countries	P S Andersen and P Hainaut
62 March 1999	The pricing of bank lending and borrowing: evidence from the federal funds market	Craig H Furfine

