Microeconomic Inventory Adjustment: Evidence From U.S. Firm-Level Data[†]

Jonathan McCarthy
Research Department
Federal Reserve Bank of New York
33 Liberty Street
New York, NY 10045
e-mail: jonathan.mccarthy@ny.frb.org

Egon Zakrajšek
Division of Monetary Affairs, MS-84
Board of Governors of the Federal Reserve System
20th Street & Constitution Avenue, NW
Washington D.C., 20551
e-mail: egon.zakrajsek@frb.gov

February 2000

[†]This is a substantially revised version of an earlier paper entitled "Microeconomic Inventory Adjustment and Aggregate Dynamics." We thank Douglas Dwyer, Mark Gertler, Jim Kahn, and Ken Kuttner for helpful comments. Benjamin Bolitzer provided expert research assistance. All remaining errors and omissions are our own responsibility. The opinions expressed in this paper do not necessarily reflect views of the Federal Reserve Bank of New York, Board of Governors, nor the Federal Reserve System.

Abstract

We examine inventory adjustment in the U.S. manufacturing sector using quarterly firm-level data over the period 1978–97. Our evidence indicates that the inventory investment process is nonlinear and asymmetric, results consistent with a nonconvex adjustment cost structure. The inventory adjustment process differs over the business cycle: for a given level of excess inventories, firms disinvest more in recessions than they do in expansions. The inventory adjustment process has changed little between the 1980s and 1990s, suggesting that recent advances in inventory control have had little effect on adjustment costs. Nevertheless, the optimal inventory-sales ratio in the durable goods sector has declined significantly during our sample period.

JEL Classification: D24, E22, E37
Keywords: inventories, (S, s) inventory policies, linear-quadratic model, business cycles

1 Introduction

The cyclical behavior of inventories has been an important feature of aggregate fluctuations in the U.S. economy. Over fifty years ago, Abramowitz (1950) showed that a plunge in inventory investment accounts for most of a contraction in output during a typical U.S. recession prior to World War II—a statistical finding that has continued to hold in the postwar data (see, for example, Blinder and Maccini 1991). This empirical regularity of aggregate output fluctuations has led macroeconomists to examine inventory investment as a potentially important channel for the propagation and the amplification of exogenous shocks to the economy.

The vast majority of macroeconomic research concerned with inventory dynamics has utilized the linear quadratic (L-Q) model of Holt, Modigliani, Muth, and Simon (1960). In the canonical L-Q model, the convex production technology at the microeconomic level, combined with the representative agent assumption, yields aggregate inventory dynamics that have proved difficult to match with the data. In particular, as documented by Ramey and West (1999), the strong procyclical fluctuations in inventory investment and the persistence of inventory movements conditional on sales are two features of the data that are difficult to reconcile with the L-Q model.

In view of the L-Q model's poor empirical performance, a number of authors have argued that the model's lack of empirical success stems from the assumption of convex production technology at the micro-level; see, for instance, Blinder (1981), Caplin (1985), and Blinder and Maccini (1991). The linear inventory decision rules implied by this assumption, although relatively easy to aggregate and estimate, do not capture potential nonlinear features of microeconomic inventory behavior. Among others, these features include (S, s)–type inventory policies, resulting from 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.[1]

Our goal in this paper is to examine the extent to which firm-level inventory adjustment in the U.S. manufacturing sector differs from the adjustment process implied by the L-Q model. We embed an inventory problem into a flexible and empirically tractable framework—the so-called generalized (S, s) approach—developed in a series of papers by Caballero, Engel, and Haltiwanger.[2] A fundamental element of this approach is the adjust-

[1]The possibility of technological nonconvexities in the context of the L-Q model, in particular increasing returns to scale, is examined by Ramey (1991), who finds industry-level evidence consistent with firms operating in a region of declining marginal costs.

[2]For a general theoretical treatment of adjustment and aggregate dynamics in (S, s) economies, see Caballero and Engel (1991). For applications of the generalized (S, s) approach to business fixed investment, see Caballero (1999), Caballero and Engel (1999), and Caballero, Engel, and Haltiwanger (1995); employment dynamics are examined by Caballero and Engel (1993) and Caballero, Engel, and Haltiwanger (1997).

ment function of the control variable in question. In our model, the inventory adjustment function relates the fraction of the deviation between the "target" and the actual level of inventories that firms close during a period to the size of that deviation. The main advantage of our approach is that the inventory adjustment function is estimated nonparametrically and, therefore, can take on a wide variety of shapes, including those implied by the L-Q model, a simple (S, s) model, and a more general nonlinear, asymmetric adjustment model.

Our analysis of inventory investment dynamics applies the generalized (S, s) framework to a long panel of high-frequency firm-level data for the U.S. manufacturing sector. Compared to other micro-level studies, our firm-level data set has several advantages for studying aggregate inventory behavior. First, our data set comprises a large portion of the U.S. manufacturing sector, averaging over 60 percent of aggregate manufacturing inventories. Second, our data cover the time period from 1978 to 1997 and thus include a number of business cycles. Third, the frequency of observation is quarterly, which is of greatest relevance for studying aggregate business cycle fluctuations.

Contrary to the predictions of the L-Q model, our results indicate that significant nonlinearities and asymmetries exist in the estimated inventory adjustment functions. The nonlinearity is consistent with the use of (S, s)–type inventory policies, while the asymmetries suggest the presence of irreversibilities in the production technology. Given the prominence of inventories in aggregate fluctuations, we examine the adjustment function at various stages of a business cycle. We find that for a given level of excess inventories, firms reduce their inventory holdings more in recessions than they do in expansions, a result consistent with the cyclical pattern of inventory investment in the aggregate data.

The long time span of our panel also enables us to examine the extent to which the inventory adjustment process may have changed over time. In particular, there has been considerable debate whether improvements in inventory control methods during the 1980s (e.g., just-in-time techniques, bar coding, etc.) may have muted the inventory cycle, translating into reduced volatility of aggregate output fluctuations; see, for instance, Filardo (1995), Allen (1995), and McConnell and Perez-Quiros (1998).

Our results paint a mixed picture. Although the rates of adjustment toward the target inventory levels are somewhat higher in the 1990s than in the 1980s, the decade-specific adjustment functions exhibit very similar shapes, indicating that the inventory adjustment process has not changed significantly during our sample period. Moreover, the cross-sectional dispersion of inventory deviations from their target levels has not declined over time, a result consistent with a stable inventory adjustment process. However, the average target inventory-sales ratio in the durable goods sector has declined significantly during our sample period, indicating that advances in inventory control have had a sizable effect on this margin of firms' inventory behavior. The relatively stable adjustment cost structure on the one hand, and the declining level of target inventory holdings relative to sales on the other, together imply that the effect of inventory control improvements on aggregate inventory fluctuations remains an open question.

The rest of the paper is organized as follows. In the next section, we discuss the issues raised by linear versus nonlinear inventory control. In Section 3, we outline our approach to the study of inventory dynamics and discuss the details behind the construction of our key state variable—the deviation between actual and target levels of inventories. In Section 4, we provide a brief description of the data and econometric issues. In Section 5, we present our main results, and Section 6 concludes.

2 Linear versus Nonlinear Inventory Control

As discussed in the introduction, a majority of empirical inventory research is based on the L-Q model. The key assumption of this model is that firms maximize profits subject to a convex production technology. This assumption implies that firms will attempt to smooth production in the face of stochastic sales. Embedding the maximization problem into a representative agent framework and using quadratic functional forms to approximate the cost structure associated with the convex production technology, the L-Q model yields a microstructure that is relatively easy to aggregate and estimate.

Although this microstructure has provided substantial insights into firm and aggregate inventory behavior, it has had difficulty explaining two fundamental and robust features of inventory investment. First, inventory investment is highly procyclical: inventories tend to be built up gradually in expansions and to be drawn down rapidly in recessions. Second, inventory movements exhibit considerable persistence, even after conditioning on sales.[3]

Within the linear-quadratic/representative-agent framework, two explanations have been advanced to reconcile these two facts with the model. The first explanation assumes that highly persistent exogenous shocks affect the cost of production. These cost shocks cause firms to bunch production—leading to procyclical inventory investment—while the built-in persistence of the cost shock process is transmitted to the inventory-sales relationship.

The second explanation posits a strong accelerator motive and high costs of adjusting production. The accelerator motive, which links current inventories to expected future sales, and positively serially correlated sales cause inventory investment to be procyclical. The high adjustment costs imply that a return to the long-run equilibrium following a

[3]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 in the Ramey and West (1999) terminology—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.

perturbation of the inventory-sales relationship will be gradual, because firms will adjust production slowly to minimize costs.

Although each explanation is economically plausible, the empirical support for either is tenuous at best. The persistent cost shock hypothesis seems to work only when these shocks are modeled as unobservable to the econometrician. Although it is plausible that the firm's cost structure could be affected by unobservable disturbances, their observable counterparts such as real unit labor costs and interest rates appear to have no appreciable effect on inventory investment. Evidence in favor of the adjustment cost hypothesis is equally unpersuasive. Estimates of adjustment cost parameters are unstable across different specifications and estimation techniques and range from negligible to economically implausible.[4]

In contrast to the L-Q vein of inventory research, a considerable operations research literature, starting with Scarf (1960), emphasizes fixed costs and other nonconvexities in the production planning problem.[5] From the macroeconomic perspective, Blinder (1981) and Blinder and Maccini (1991) have provided compelling arguments that nonconvexities in the production process may be crucial for our understanding of aggregate inventory dynamics. The optimal inventory policy in such an environment is of the (S, s)–type, implying periods of inaction when inventories are depleted, followed by periods of activity during which inventories are replenished.

Such nonlinear microeconomic inventory behavior has the potential to affect aggregate inventory dynamics and, furthermore, calls into question the representative agent assumption underlying most applied macro-inventory research. If firms follow state-dependent (S, s) policies, the time-path of the entire cross-sectional distribution of inventories has an effect on aggregate dynamics. In such an economy, a negative aggregate shock could 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.

In part because of these complications, empirical macroeconomists have been reluctant to forsake the analytically tractable L-Q framework for (S, s)–type models. Although most economists would agree that a firm's production/inventory problem will likely differ from

[4] 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. These studies 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.

Alternatively, Schuh (1996) dispenses with the representative agent assumption. Using monthly plant-level 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 L-Q model.

[5] See Hordijk and Van der Duyn Schouten (1986) for a general treatment of continuous review inventory models with fixed costs.

the assumptions underlying the L-Q model, they argue that the aggregation across goods and time may smooth the effects of any nonconvexities. As a result, the L-Q model remains a reasonable approximation to the inventory behavior of interest to macroeconomists; see, for example, Blanchard (1983) and Ramey and West (1999).

Moreover, at the macro-level, empirical evidence supportive of (S, s)–type models is limited.[6] In particular, using the theoretical aggregation results for dynamic (S, s) economies derived by Caplin (1985) and Caballero and Engel (1991), several studies have attempted to test the predictions of a simple (S, s) model using both aggregate and firm-level data for the trade sector.[7] The basic conclusion of these studies is that the time-series behavior of inventory investment in the trade sector is consistent with the (S, s) model's steady-state implications. The problem with this conclusion, however, is twofold. First, the aggregate steady-state results of Caplin (1985) are valid only under restrictive assumptions of exogenous, serially uncorrelated sales, time-invariant (S, s) bands, and no delivery lags.[8] Second, the steady-state, reduced-form relationships between inventory investment and sales are also consistent with an economically plausible parameterization of the L-Q model with stockout costs; see, for example, Blinder (1986), Kahn (1987), and Krane (1994).[9]

Therefore, if we are to shed light on whether nonconvexities matter for inventory behavior at business cycle frequencies, we must study inventory dynamics rather than steady-state behavior. Because structural (S, s)–type models must be kept simple to derive analytic results useful for empirical analysis, the 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.

3 Inventories and the Generalized (S, s) Approach

In this section, we describe the generalized (S, s) approach used to analyze firm-level inventory investment decisions. The first subsection discusses the basic elements of our model, while the second discusses the specifics in measuring the state variable of the model.

[6]At the micro-level, Hall and Rust (1999) examine nearly two years of daily transactions data from a steel wholesaler. They find the firm's inventory/order policies are consistent with a state-dependent (S, s) model. However, their data are probably too limited in scope to change the opinions of those skeptical about using these models to study aggregate inventory behavior.

[7]See Mosser (1988, 1991), Episcopos (1996), and McCarthy and Zakrajšek (1997).

[8]Mosser (1988, 1991) and McCarthy and Zakrajšek (1997) attempt to allow for delivery lags and serial correlation in sales and continue to find support for the (S, s) model.

[9]To address this concern, McCarthy and Zakrajšek (1997) also estimate an Euler equation associated with the L-Q model. They find that, although the model is rejected at the industry level, estimates of the structural cost parameters are economically reasonable and statistically significant at the firm level; however, the over-identifying restrictions imposed by the model are rejected, and the parameter estimates are not stable across different asymptotically equivalent normalizations.

3.1 Basic Elements

The premise underlying our analysis is that a firm's inventory adjustment within a period depends upon the size of its perceived inventory shortfall or excess. In particular, we allow for the possibility that firms may respond to a greater degree as inventories move further away from their target levels. If such nonlinear inventory adjustment is detected, it would be indicative of nonconvexities in the production process.

At the microeconomic level, inventory investment in our model depends upon a single state variable—the log deviation between target and actual levels of inventories, which we label as the *inventory deviation index* z:

$$z_{it} \equiv \ln H_{it}^{*} - \ln H_{it-1}. \tag{1}$$

In equation 1, i indexes firms, t indexes time, and H_{it} and H_{it}^{*} denote actual and target real end of period t inventory stocks, respectively. It is important to note that because z_{it} depends on H_{it}^{*}, a theoretical construct, it is model-dependent.

Although firms may continuously monitor their inventories and sales, observations in the model are at discrete intervals. Therefore, the timing of shocks and adjustment within a period needs to be spelled out. In our timing conventions, we follow Caballero, Engel, and Haltiwanger (1997) and assume that although firms may experience a sequence of shocks within a period, these shocks can be summarized by a single aggregate shock η_t common to all firms and a single idiosyncratic shock ν_{it} that has a zero cross-sectional mean in each period. After observing both shocks, firms determine their new target inventory levels H_{it}^{*} and thus their inventory deviation z_{it}. Firms then adjust their inventory levels accordingly, and the process is repeated.

The evolution of z_{it} over time thus reflects the shocks to a firm's target inventory level and its adjustments in response to these shocks. Using equation 1 and our within-period timing assumptions, we can decompose the change in a firm's inventory deviation index between periods $t-1$ and t, Δz_{it}, as

$$\Delta z_{it} = \Delta \ln H_{it}^{*} - \Delta \ln H_{it-1} = (\eta_t + \nu_{it}) - \Delta \ln H_{it-1}. \tag{2}$$

The next element underlying our analysis is the *cross-section of firms' inventory deviations* in period t, denoted by $f(z,t)$. This is the cross-sectional probability density of firms' inventory deviations immediately preceding inventory adjustments during period t. Therefore, the fraction of firms with inventory deviations between z and $z + dz$ in period t is approximately equal to $f(z,t)dz$.

The last element in our framework is the inventory *adjustment function*, denoted by $\Lambda(z,t)$. This function maps the inventory deviation z to the fraction of that deviation that is closed by a firm within a period. To compute it, we determine which firms have an inventory deviation close to z in period t and calculate the fraction of the deviation closed on average by those firms. It then follows that the average inventory growth rate for firms with inventory deviation z in period t is equal to $z\Lambda(z,t)$.

Much of our empirical analysis will focus on the shape of the adjustment function $\Lambda(z,t)$, because its shape provides considerable information about the nature of production technology. In particular, a constant adjustment function, $\Lambda(z,t) = \lambda_0$ for all t, is consistent with a time-invariant convex adjustment technology and generates aggregate inventory dynamics identical to those implied by a representative-agent model with quadratic adjustment costs; see, for instance, Caballero and Engel (1993). Therefore, if the estimated adjustment function differs from a constant function, this provides evidence consistent with the presence of nonconvexities in the inventory adjustment process.

To conclude this subsection, we discuss how the three elements underlying our analysis combine to relate firm-level inventory investment decisions to aggregate inventory dynamics. Letting $\Delta \ln H_t^A$ denote the aggregate growth rate of inventories in period t, the preceding definitions imply that

$$\Delta \ln H_t^A = \int z w(z,t) \Lambda(z,t) f(z,t) dz, \tag{3}$$

where $w(z,t) \geq 0$ for all t is a weighting function.

Equation 3 shows that the dynamics of aggregate inventory investment are determined by the interaction between the adjustment function and the shifts in the cross-sectional density of deviations induced by aggregate and idiosyncratic shocks. In general, as long as the adjustment function $\Lambda(z,t)$ depends explicitly on z, aspects of the cross-sectional distribution other than its mean (e.g., dispersion and skewness) will influence aggregate inventory dynamics.[10]

3.2 Measuring Inventory Deviations

A key to our analysis is the construction of a measure of the inventory deviation index z. To do so, we need to formulate and estimate a model for H_{it}^*, the unobserved target inventory level. In formulating such a model, we assume that in the absence of adjustment costs, a firm's objective is to balance the costs of a potential stockout with the costs of holding inventories. As in the standard L-Q model, this tradeoff can be captured by the following

[10] In fact, Caballero and Engel (1993) and Caballero, Engel, and Haltiwanger (1995, 1997) test this implication for aggregate investment and employment by assuming a time-invariant, polynomial approximation to the adjustment function—that is, $\Lambda(z,t) = \sum_{j=0}^{p} \lambda_j z^j$ for all t. Under such assumption, equation 3 implies that the aggregate inventory growth depends *linearly* on the first $p+1$ (weighted) moments of the cross-sectional distribution $f(z,t)$; for a criticism of such regressions, see Veracierto (1998).

quadratic cost function:

$$C(h_{it}, m_{it}, s_{it+1}) = \frac{1}{2}\left[h_{it} - (A(m_{it}) + \gamma_i E_{it} s_{it+1})\right]^2, \tag{4}$$

where h_{it} denotes the logarithm of inventories at the end of period t, s_{it+1} denotes the logarithm of real sales in period $t+1$, E_{it} is the expectations operator for firm i conditional on all the information through the end of period t, and $\gamma_i > 0$ is a firm-specific cost elasticity with respect to expected future sales. The function $A(\cdot)$ governs the significance of stockout costs relative to holding costs and depends on the vector of variables m_{it}.[11]

The function $A(\cdot)$ in equation 4 determines the optimal inventory-sales relationship. We allow this inventory-sales relationship to depend on the vector of firm-specific variables m_{it}. The source of this variation could reflect, for example, idiosyncratic and aggregate shocks to the cost of production. By affecting the opportunity costs of holding inventories, such cost shocks may induce variation in stockout costs relative to holding costs.[12] In addition, the optimal inventory-sales ratio is likely to vary because of seasonal factors. Finally, changes in inventory control techniques may induce long-run trends in the optimal inventory-sales relationship that are not captured by other variables. To allow for such variation in the inventory-sales relationship, we assume that $m_{it} = (d_t', c_{it}, c_{it+1})'$, where d_t is a k-vector of nonstochastic variables (to be specified later), c_{it} is the logarithm of the cost per unit of output in period t, and we let $A(\cdot)$ take on the following log-linear form:

$$A(d_t, c_{it}, c_{it+1}) = \alpha_i' d_t + \theta_i c_{it} + \phi_i E_{it} \Delta c_{it+1}, \tag{5}$$

where α_i is the k-vector of unknown firm-specific parameters corresponding to the vector of nonstochastic variables d_t. Firm-specific parameters $\theta_i < 0$ and $\phi_i > 0$ measure the extent to which inventories are used to buffer production from cost shocks.

Using equations 4 and 5 in a standard intertemporal cost minimization problem, it follows that the optimal log-level of inventories h_{it}^* satisfies the first-order condition,

$$h_{it}^* = \alpha_i' d_t + \gamma_i s_{it} + \theta_i c_{it} + \gamma_i E_{it} \Delta s_{it+1} + \phi_i E_{it} \Delta c_{it+1}. \tag{6}$$

Using equation 6, we can now derive our measure of the inventory deviation index z. First, let e_{it} denote the inventory deviation at the end of period t, *after* the adjustment has taken

[11] Compared to the typical accelerator term in the L-Q model, our specification of the tradeoff between inventory holding costs and stockout avoidance is somewhat nonstandard, as all variables are stated in logarithms. Hence in our specification, the firm's cost function depends on the percentage difference between the target and actual inventory levels rather than the absolute difference between the two. With firm-level data, we believe that the log-specification provides a more natural description of these costs.

[12] Specifically, higher current costs would tend to reduce the current target level of inventories given sales, whereas greater expected future growth of costs would tend to increase the current target inventories given sales; Kahn (1992) derives a version of the L-Q model that formalizes this intuition.

place. It then follows that,

$$e_{it} = h_{it}^{*} - h_{it}, \tag{7}$$

or

$$e_{it} = \alpha_i' d_t + \gamma_i s_{it} + \theta_i c_{it} + \gamma_i E_{it} \Delta s_{it+1} + \phi_i E_{it} \Delta c_{it+1} - h_{it}. \tag{8}$$

Note that from equations 1 and 7, e_{it} differs from z_{it} only in that the former incorporates the inventory adjustment during the period—that is, $e_{it} = z_{it} - \Delta h_{it}$.

In our model, the time-series properties of the ex post inventory deviation e_{it} are crucial for identifying the parameters of the model and thus the optimal inventory level. With continuous monitoring of inventories and observations at discrete intervals, the post-adjustment deviations reflect two factors. First, although we assume that all the adjustment occurs at the end of the period, the actual firm-level inventory investment decision(s) in our data may take place at times that do not coincide with the observation interval. Second, firms may not adjust inventories to their optimal level.

In this context, we assume that deviations following the adjustment do not persist indefinitely and that neither factor causes a systematic bias in the observed ex post inventory deviations. This implies that for each firm i, e_{it} is a realization from a stationary stochastic process, with an unconditional mean equal to zero, $E[e_{it}] = 0$, and finite variance, $E[e_{it}^2] = \sigma_i^2$, for all i. Therefore, e_{it} can be considered a random disturbance in the following regression:

$$h_{it} = \alpha_i' d_t + \gamma_i s_{it} + \theta_i c_{it} + \gamma_i E_{it} \Delta s_{it+1} + \phi_i E_{it} \Delta c_{it+1} - e_{it}. \tag{9}$$

To estimate equation 9, we need to specify the k-vector of nonstochastic variables d_t and parameterize the exogenous forcing processes for Δs_{it} and Δc_{it}. We assume that the deterministic part of the stockout avoidance parameter, d_t, consists of low- and high-frequency components. The low-frequency component includes linear and quadratic time trends and reflects such influences as advances in inventory monitoring technology, changes in the firm's relationship with its suppliers, and changes in product diversity over time. The high-frequency component captures movements in the stockout avoidance behavior associated with seasonal fluctuations. Thus, $d_t = (q(t)', t, t^2)'$, where $q(t) = (q_1(t), \ldots, q_4(t))'$ is a vector of quarterly indicator variables.

Finally, we assume that both the growth of sales, Δs_{it}, and the growth of costs per unit of output, Δc_{it}, can be represented by a stationary, firm-specific AR(4) process.[13] Thus, the conditional expectations of Δs_{it+1} and Δc_{it+1} in equation 9 can be replaced by four lags of these variables. Making this substitution, we obtain the regression equation used to

[13]We have experimented with the AR specifications of lag length $2, 3, \ldots, 6$ with a negligible effect on our results.

estimate the inventory deviation index z,

$$h_{it} = \alpha_i' d_t + \gamma_i s_{it} + \theta_i c_{it} + \sum_{\tau=1}^{4} \pi_{i\tau} \Delta s_{it+1-\tau} + \sum_{\tau=1}^{4} \varphi_{i\tau} \Delta c_{it+1-\tau} + u_{it}, \tag{10}$$

where $u_{it} = -e_{it}$, and the reduced-form parameters $\pi_{i\tau}$ and $\varphi_{i\tau}$, $\tau = 1, \ldots, 4$, are a combination of structural parameters γ_i and ϕ_i and the parameters of the exogenous forcing processes for Δs_{it} and Δc_{it}, respectively. (Note that in equation 10, the structural parameter ϕ_i is not identified.)

The negative of the estimated residual $\hat{u}_{it}$ from equation 10 is our estimate of e_{it}. To derive an estimate of z_{it}, the state variable in our model, recall that equations 1 and 7 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},

$$\hat{z}_{it} = \hat{e}_{it} + \Delta h_{it}, \tag{11}$$

where Δh_{it} denotes the growth rate of inventories of firm i in period t.

4 Data and Econometric Methodology

4.1 Data

The data set used in our analysis come from the Standard and Poor's COMPUSTAT quarterly P/S/T, Full Coverage, and Research data files. The data set consists of a panel of 2,169 manufacturing firms and covers the time period 1978:Q1 to 1997:Q4 (80 quarters). The panel is unbalanced, with the minimum continuous tenure of 20 quarters. After eliminating firms with gaps in the time-series dimension or implausible entries, we are left with a total of 92,163 firm/quarter observations. During our sample period, the firms in the panel account, on average, for 61 percent of aggregate non-farm manufacturing inventories. The Data Appendix contains details on the exact sample selection procedure, industry composition, the construction of variables, as well as summary statistics.

Let us make one additional point concerning the data. If nonconvexities are important in determining inventory investment decisions, *plant-level* data may be preferable to firm-level data. Many decisions concerning inventory investment are made at the plant level. For example, plants within a single firm that produce different products may have inventory policies that depend on the demand for the particular products that each plant produces. Thus, to the extent that individual plants operate as independent entities within a multi-plant firm, the distribution of shocks and inventory deviations across plants within a firm may affect firm-level inventory dynamics.

Moreover, nonlinear inventory adjustment is likely to be more apparent at less aggregated levels. For example, total inventories of the firm studied by Hall and Rust (1999) exhibit considerably smoother behavior at the daily frequency than do the inventories for each individual product category. Thus, concentrating our analysis at the firm level may lose some important information concerning the nonlinear adjustment at the fundamental microeconomic level.

Nevertheless, many aspects of inventory decisions are centralized within the firm. Finished goods inventories may go to a centralized distribution center, which enables the firm to make production and inventory decisions more efficiently. 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 is evidence of the role that nonconvexities may play at the firm level. Thus, much of the effect that nonconvexities may have on inventory adjustment can be studied at the firm level. Moreover, it is interesting to examine how much aggregation to the firm level, as well as time aggregation, convexifies the nonconvexities at less aggregated levels.

4.2 Econometric Methodology

In this section, we discuss the details behind estimation of equation 10. For notational convenience, let $X_{it} = (d_t', s_{it}, c_{it}, \Delta s_{it}, \ldots, \Delta s_{it-3}, \Delta c_{it}, \ldots, \Delta c_{it-3})'$ denote the vector of all explanatory variables in equation 10, and let $\beta_i = (\alpha_i', \gamma_i, \theta_i, \pi_{i1}, \ldots, \pi_{i4}, \varphi_{i1}, \ldots, \varphi_{i4})'$ denote the corresponding vector of parameters. Observing N firms over T_i, $i = 1, \ldots, N$, periods, we can write equation 10 for firm i compactly as,

$$h_i = X_i \beta_i + u_i; \quad i = 1, \ldots, N. \tag{12}$$

Our primary concern in equation 12 is the problem of parameter heterogeneity across firms. It is widely recognized that parameter heterogeneity at the micro-level—in both cross-sectional and time-series dimensions—can have important consequences for estimation and inference. As shown by Mairesse and Griliches (1990) and Schuh (1996), for instance, evidence of slope heterogeneity from firm-level panels is pervasive.

A parsimonious specification that incorporates such parameter heterogeneity is the random coefficients model (RCM); see, for example, Hsiao (1996). In the RCM specification, all the parameters of the model are random variables and stochastic specifications are introduced to capture parameter heterogeneity. The RCM approach to parameter heterogeneity reduces the number of parameters to be estimated considerably while still allowing the coefficients to differ (randomly) across cross-sectional units and/or time.

To allow for variation in β_i across firms, we follow standard RCM assumptions that $\beta_i =$

$\beta + \xi_i$, where ξ_i is a random vector with mean equal to zero that is distributed independently of the exogenous regressors X_i and the disturbances u_i.[14] Using these assumptions and stacking observations for all N firms yields,

$$h = X\beta + Z\xi + u, \tag{13}$$

where $Z = \text{diag}[X_i]_{i=1,\dots,N}$ is a block-diagonal matrix with the matrix X_i on its i-th diagonal element, $\xi = (\xi_1', \dots, \xi_N')'$ is an unknown vector of random effects that capture parameter heterogeneity across firms, and u is an unknown random error vector. The foregoing assumptions imply that

$$E\begin{bmatrix} \xi \\ u \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}; \quad \text{and} \quad \text{Var}\begin{bmatrix} \xi \\ u \end{bmatrix} = \begin{bmatrix} G & 0 \\ 0 & R \end{bmatrix}, \tag{14}$$

where $G = \Sigma_\xi \otimes I_N$, the symbol $\otimes$ denotes the Kronecker product, and I_N is the identity matrix of dimension N.[15]

Consistent and efficient estimates of the unknown parameter vectors β and ξ in equation 13 can be obtained via Generalized Least Squares (GLS). Formally, the GLS estimators of β and ξ, denoted by $\hat{\beta}$ and $\hat{\xi}$, respectively, are obtained by minimizing

$$(h - X\beta)'V^{-1}(h - X\beta), \tag{15}$$

with respect to β, where $V = ZGZ' + R$ is the covariance matrix of h. Minimizing equation 15, of course, requires the knowledge of V and, therefore, the knowledge of G and R. Given reasonable estimates of G and R, which are denoted $\widehat{G}$ and $\widehat{R}$, respectively, we can obtain the feasible GLS (FGLS) estimates of the parameter vectors β and ξ, given by

$$\hat{\beta} = \left(X'\widehat{V}^{-1}X\right)^{-} X'\widehat{V}^{-1}h; \tag{16}$$

$$\hat{\xi} = \widehat{G}Z'\widehat{V}^{-1}\hat{\varepsilon}, \tag{17}$$

where $\hat{\varepsilon} = h - X\hat{\beta}$ and $^{-}$ denotes the generalized inverse. Using equations 16 and 17, our estimate of the ex post inventory deviation index e is then given by,

$$\hat{e} = -(h - X\hat{\beta} - Z\hat{\xi}). \tag{18}$$

Because of the unbalanced nature of out data, we use the Minimum Variance Quadratic

[14] This independence assumption seems reasonable, because any suspected systematic dependence can, in principle, be explicitly modeled and hence accounted for.

[15] Note that in equation 14, we let the covariance matrix Σ_ξ be unrestricted. Hence, we allow for the possibility that the elements of the random vector ξ_i may be correlated.

Unbiased Estimators (MIVQUE) of G and R to compute $\hat{\beta}$ and $\hat{\xi}$. Developed by Rao (1971), the non-iterative MIVQUE of variance components are an alternative to the computationally intensive ML estimators. The procedure requires no distributional assumptions other than the existence of the first four moments. Basically, the MIVQUE of a linear combination of the unknown variance components in G and R are obtained by finding a symmetric matrix W, such that $\text{Var}[h'Wh]$ is minimized subject to the conditions that $h'Wh$ is an unbiased estimator of the linear combination of the variance components and is invariant to any translation of the β parameter.[16]

To simplify computations, we also assume that $\text{Var}[u_{it}] = \sigma_i^2 = \sigma_e^2$ for all i. Hence, $R = \sigma_e^2 I_n$, where $n = \sum_{i=1}^N T_i$ is the total sample size of our unbalanced panel. Because this homoscedastic assumption is likely to be violated in practice, we compute the covariance matrix of the parameter vector $\hat{\beta}$ using the asymptotically consistent estimator described by Diggle, Liang, and Zeger (1995). This estimator is computed as follows:

$$\text{Var}[\hat{\beta}] = \left(X'\widehat{V}^{-1}X\right)^{-} \left[\sum_{i=1}^{N} X_i'\widehat{V}_i^{-1}\hat{\varepsilon}_i\hat{\varepsilon}_i'\widehat{V}_i^{-1}X_i\right] \left(X'\widehat{V}^{-1}X\right)^{-}, \tag{19}$$

where matrices with the subscript i correspond to those of the i-th firm.

Table 1 contains the FGLS estimates of the structural parameters γ and θ from equation 10 for the major industry groups in our sample. Note that across all industries, estimates of the both parameters are of the right sign: $\gamma > 0$ and $\theta < 0$. Almost uniformly, the estimates are statistically significant. Only for SIC 270 (Printing & Publishing) is the estimate of θ not statistically different from zero at the usual significance level, although it is of the right sign. In addition, there is considerable variation in the estimates of γ and θ across industries. Estimates of γ range from the low of 0.435 for SIC 320 (Stone, Clay & Glass Products) to the high of 0.925 for SIC 250 (Furniture & Fixtures), while estimates of θ range from -1.258 for SIC 340 (Fabricated Metal Products) to -0.225 for SIC 270 (Printing & Publishing).

Using the parameter estimates from Table 1, we then can calculate the time path of optimal inventories for each firm. Figure 1 presents the cross-sectional average of the ratio of optimal inventories to real quarterly sales for the nondurable and durable goods sectors

[16] As discussed by Baltagi (1995), the MIVQUE require a priori values of the variance components in G and R. Consequently, the resulting estimators are minimum variance only if these a priori values coincide with the true values. Two priors for the matrix W are typically used in practice: the identity matrix, denoted by MIVQUE(0), and the method of moments (ANOVA) estimators denoted by MIVQUE(A). We use the MIVQUE(0) estimators of G and R to obtain the FGLS estimators $\hat{\beta}$ and $\hat{\xi}$ in equations 16 and 17. Note that if one iterates on the initial values of the variance components until convergence, the MIVQUE will converge to ML estimates under normality; see Swallow and Monahan (1984) and Baltagi and Chang (1994) for a detailed discussion and relative efficiency comparisons of the various variance components estimators with unbalanced panel design.

over our sample period. For the most part, the behavior of this average ratio is quite similar to that of the sectoral inventory-sales ratios computed from the published aggregate data. The average optimal inventory-sales ratio in the durable goods sector declines fairly steadily throughout our sample period, a pattern very similar to that in the aggregate data. (The correlation between the two series is almost 0.9.) In contrast, after a decline in the early 1980s, the average ratio in the nondurable goods sector has been relatively constant. Similarly, the aggregate inventory-sales ratio in the nondurable goods sector has been relatively constant throughout our sample period; the correlation between the two series in this sector is about 0.4 after 1983.[17]

The steady decline in the optimal inventory-sales ratio for the durable goods firms suggest one way that the recent improvements in inventory control may have affected firm-level inventory behavior. The wide-spread adoption of these innovations during the 1980s presumably has enabled firms to control and monitor their inventories better than they had previously. If so, firms then face a smaller probability of encountering a stockout for a given inventory level and thus would want to carry fewer inventories relative to their sales. If these practices became sufficiently widespread throughout the sector, we would observe a decline in the average optimal inventory-sales ratio.[18]

5 Results

In this section, we present our main results. We characterize the main elements of equation 3, namely, the inventory adjustment function $\Lambda(z,t)$ and the cross-sectional distribution of inventory deviations, $f(z,t)$. Of particular interest are the shape of the inventory adjustment function and its variation across sectors and time. In addition, we examine the time path of the cross-sectional distribution of the inventory deviation index z.

5.1 Microeconomic Inventory Adjustment

To compute the inventory adjustment function $\Lambda(z,t)$, we first discretize the state space. The inventory deviation index z takes values between -0.4 and 0.4, over an equally-spaced grid with intervals of size 0.01. (This interval covers over 99 percent of the sample range

[17]The aggregate inventory-sales ratio in the nondurable goods sector does not exhibit the pronounced decline in the early 1980s that is evident in the micro data; hence, the correlation over the entire sample period is only 0.1. However, this weak correspondence between the micro and aggregate data is a result of Standard & Poor's relatively limited coverage of nondurable goods firms early in the sample. In the mid 1980s, Standard & Poor's greatly expanded its coverage, especially among smaller nondurable goods firms, and the aggregates computed from the micro data are much more closely correlated with the aggregate time series.

[18]Abernathy et al. (1999) contains a detailed case study of the apparel and textile industry, where even in the face of product proliferation, lean retailing practices have been propagated through to the manufacturers.

of z.) In each interval, we construct the value of the adjustment function by dividing the average inventory growth by z for those firms that are at z just before inventory adjustments take place.[19] In what follows, all depicted adjustment functions are smoothed by a cubic B-spline.

5.1.1 Manufacturing Inventory Adjustment

The solid line in Figure 2 shows the estimated inventory adjustment function for the U.S. manufacturing sector. The two dotted lines represent two standard deviation error bands.[20] The dashed line represents the smoothed density function of inventory deviations across all firms and quarters.

Three observations about Figure 2 stand out. First, the adjustment function in nonlinear: firms adjust inventories more as the magnitude of the deviation increases. This nonlinearity is irrespective of whether the deviation is an inventory shortage ($z > 0$) or an inventory overhang ($z < 0$). The difference between the adjustment rates for firms with large and small deviation is on the order of 15 percent and appears to be significant. The shape of the estimated adjustment function is consistent with the existence of nonconvexities in the production technology that induce firms to adopt (S, s)–type inventory policies.

Second, the inventory adjustment function is asymmetric. For small- to moderate-sized inventory deviations, 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 this asymmetry. First, because of a strong stockout avoidance motive, firms may be more willing to carry extra inventories. Second, market irreversibilities could prevent firms from reducing their excess inventories. Third, firms may be reluctant to cut output, because they find it costly not to employ their capital and labor.

The final point about Figure 2 concerns the level of the estimated adjustment rates. In contrast to the implied adjustment rate estimated from a canonical L-Q model, the adjustment rates in Figure 2 are economically plausible. According to Ramey and West (1999), typical estimates of the adjustment rate from the L-Q model on quarterly U.S. manufac-

[19]In calculating the values of 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 $\hat{e}_{it}$. All the results, however, are nearly identical if an unweighted average is used in computations.

[20]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 5,000 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.

turing data are in the 10-20 percent range, indicating large costs of adjusting production. In our model, the estimated adjustment rates vary between 60 and 80 percent per quarter, an economically plausible range given the relative size of inventories to quarterly sales for the firms in our sample.

5.1.2 Sectoral Inventory Adjustment

In this section, we investigate whether the inventory adjustment function differs between the durable and nondurable goods sectors. Sectoral differences in the inventory adjustment process may provide some insight into the nature of the nonlinearities present in the overall inventory adjustment function.

Apparent from Figure 3 is that the inventory adjustment functions for durable and nondurable goods firms exhibit a similar shape. In both sectors, the adjustment functions contain nonlinearities associated with generalized (S, s)–type inventory policies. However, there are some noticeable differences between the two sector-specific adjustment functions. First, most obviously, the adjustment rate for nondurable goods firms is greater than that for durable goods firms for all values of the inventory deviation index z (the difference ranges from 5 to 10 percent). This difference may reflect that nondurable goods manufacturers are more willing or better able to close inventory deviations, possibly owing to differing production technologies or market structures.

Second, the adjustment function for durable goods firms displays a greater asymmetry than the adjustment function for nondurable goods firms. This difference suggests that the issues of stockout avoidance and market or production irreversibilities may be more important for inventory behavior of durable goods firms than for their counterparts in the nondurable goods sector. Moreover, the combination of lower inventory adjustment rates and greater asymmetries in the durable goods sector indicates that nonconvexities in the production technology may be more prevalent in durable goods industries.[21]

5.1.3 Time Variation in Inventory Adjustment

In this section, we examine the variation of the inventory adjustment function across time. First, we compare the adjustment functions during expansions and recessions to investigate how the adjustment process may differ over the business cycle. We then compare the adjustment functions estimated over the 1980s and the 1990s to examine the extent to which technological improvements in inventory control during our sample period may have affected the inventory adjustment process.

[21] There is persuasive evidence that nonconvex production costs play an important role in the automobile industry; see Bresnahan and Ramey (1994) and Hall (1999).

Cyclical Shifts in the Adjustment Function The inventory adjustment functions for NBER-dated expansions and recessions during our sample period are displayed in Figure 4. The implied pattern of inventory investment underlying these functions is in accord with the stylized facts regarding the role of inventories in aggregate fluctuations. Most notably, the adjustment rate for small to moderate inventory overhangs ($z < 0$) during recessions is above the adjustment rate for inventory overhangs of similar size during expansions. Thus, for a given level of excess inventories, firms reduce their inventory holdings more in recessions than they do during normal economic times.

Furthermore, compared to a period of normal economic activity, target inventories are likely to be revised down in recessions. This macroeconomic state dependence could be a result of increased demand uncertainty or adverse cost or productivity shocks during periods of economic downturn. Consequently in recessions, a greater fraction of firms may find themselves holding excess inventories. The interaction between the adjustment function and the cross-sectional distribution of inventory deviations at the onset of a recession would lead to a period of rapid aggregate inventory disinvestment.

Long-Term Shifts in the Adjustment Function Next, we examine the stability of the inventory adjustment function across different subperiods of our sample. This exercise enables us to investigate the potential effect of technological improvements and innovations in inventory management practices on the inventory adjustment process.

Figure 1 of Section 4.2 shows that the estimated optimal inventory-sales ratio in durable goods industries has declined significantly during our sample period. We argued that this steady decline is consistent with the adoption of more efficient inventory control methods. In addition to this effect, the impact of inventory control innovations on the adjustment process is equally important in assessing the implications of these innovations for aggregate inventory cycles. Suppose that following the adoption of modern inventory management practices, adjustment rates have increased significantly and the adjustment function has become closer to a constant function. Such pattern would imply that the costs of adjusting inventories have declined and that nonconvexities in the cost structure have a lesser effect on firm behavior, resulting in reduced volatility of inventory investment.

Given that our data cover nearly two full decades, we compare the inventory adjustment function estimated over the 1980s with that estimated over the 1990s. In light of the above-documented cyclical shifts in the inventory adjustment process, we exclude recession quarters when we estimate the decade-specific adjustment functions. By doing so, we eliminate any differences in the inventory adjustment process that would result from business cycle-dependent shifts in the adjustment function.

The inventory adjustment functions estimated over the 1980s and 1990s are plotted in Figure 5.[22] The two functions are similar, as they both display the same nonlinear, asymmetric shape as the adjustment function estimated over the entire sample period. Nevertheless, the adjustment rates are uniformly higher during the 1990s than during the 1980s, with the largest differences occurring for inventory shortages ($z > 0$). Note also that relative to the 1980s, the adjustment function during the 1990s displays somewhat greater asymmetry between inventory investment and disinvestment.

The fact that adjustment rates are somewhat higher in the 1990s than in the 1980s is consistent with the presumed effects of improved inventory control on the inventory adjustment process. However, the shape of the adjustment function in the 1990s is also consistent with the continuing presence of nonconvexities and asymmetries in the inventory adjustment process. Combined with the relatively small increase in the rate of adjustment, this suggests that advances in inventory control methods have had little effect on the adjustment cost structure over the past two decades. On the other hand, these improvements have likely had an effect of lowering the optimal inventory holdings, particularly in the durable goods sector.

The implications of these results for aggregate inventory fluctuations are unclear. On the one hand, the recent advances in inventory control practices have done much to improve a firm's ability to continuously and accurately monitor their inventory levels, thus reducing the possibility of a "surprise" shortage. This better information, in turn, reduces the perceived probability and the expected costs of a stockout, thereby reducing the need to carry as much inventory relative to sales as before.

On the other hand, the production technology is likely to be unaffected by such inventory control innovations, and the adjustment process given the desired inventory level will remain stable over time. The decline in the optimal inventory-sales ratio implies that inventories will be lower than they would be otherwise. However, with little change apparent in the adjustment technology, it is likely that the fluctuations about the lower inventory targets may be of a similar magnitude as before, depending, ultimately, upon the distribution of inventory deviations and exogenous shocks.

5.2 Cross-Sectional Evolution of Inventory Deviations and Shocks

The cross-sectional distribution of inventory deviations, $f(z, t)$, is determined by the interaction between the aggregate and idiosyncratic shocks affecting firms' inventory targets and the inventory adjustment process. The average density—where the average is computed over all firm/quarter observations—is displayed by the dashed line in Figure 2. In

[22] We also estimated the decade-specific adjustment functions for the durable and nondurable goods sectors separately. Other than increasing noise in the data, controlling for sectoral differences has little substantive effect on our results.

this section, we consider the time variation of this distribution and the evolution of the shocks affecting inventory deviations.

5.2.1 Cross-Sectional Moments of Inventory Deviations

Figure 6 displays the time paths of the cross-sectional (weighted) mean of inventory deviations and aggregate inventory growth for the durable and nondurable goods sectors.[23] In both sectors, the cross-sectional mean of the inventory deviation index z and aggregate inventory growth move closely together; the correlation between the two series is 0.77 in the durable goods sector and 0.61 in the nondurable goods sector. The high correlation between the two series indicates that our estimates of the inventory deviations are successful in capturing the basic features of aggregate inventory investment.

Nevertheless, the correlation is far from perfect, and there are episodes—in particular, the period surrounding the economic slowdown during the 1980s—when the behavior of the two series differs markedly. This suggests that aspects other than the mean of the cross-sectional distribution of inventory deviations may affect aggregate inventory dynamics. Moreover, the use of representative agent assumption in modeling aggregate inventory investment is unlikely to capture these missing features.

The time paths of the (weighted) second, third, and fourth moments of the cross-sectional distribution are presented in Figure 7.[24] Each moment displays significant temporal variation and does not appear to be highly correlated with the other moments. Over time, the cross-sectional distribution of inventory deviations exhibits both substantial skewness and kurtosis; interestingly, these higher moments are more volatile in nondurable goods industries. The time variation of these higher moments and the imperfect correlation between the mean of the cross-sectional distribution $f(z,t)$ and aggregate inventory growth together imply that idiosyncratic shocks have an important effect on aggregate dynamics, an issue we examine further in the next section.

Finally, note that the dispersion of inventory deviations, as measured by the cross-sectional standard deviation, does not exhibit a secular decline in either sector during our sample period. A common conjecture concerning recent improvements in inventory control methods is that they make it easier for firms to align inventories more closely to their optimal levels. Thus, if advances in inventory management techniques had a significant impact on the adjustment process, we should, on average, observe a decline in the cross-sectional dispersion of inventory deviations. The fact that this decline has not occurred is consistent with the stability of the adjustment function over our sample period and indicates that

[23]The means and all other moments are weighted by last period's firm size, measured by the level of inventory stocks and have been seasonally adjusted with quarterly dummies.

[24]Excess kurtosis is defined as the difference between 3, the kurtosis of the normal distribution, and the sample kurtosis.

inventory control innovations have had little effect on the inventory adjustment process.

5.2.2 Aggregate and Idiosyncratic Shocks

In this section, we examine the evolution of the aggregate and idiosyncratic shocks. Both the time path of aggregate shocks and the time path of the cross-sectional distribution of idiosyncratic shocks are computed from equation 2, using our estimates of the inventory deviation index. The time path of estimated aggregate shocks is computed according to

$$\hat{\eta}_t = \frac{1}{N_t}\sum_{i=1}^{N_t}(\Delta\hat{z}_{it} + \Delta h_{it-1}),$$

where N_t denotes the number of firms in period t. In each period t, the estimated cross-sectional distribution of idiosyncratic shocks corresponds to the histogram of estimated ν_{it}'s, where

$$\hat{\nu}_{it} = (\Delta\hat{z}_{it} + \Delta h_{it-1}) - \hat{\eta}_t.$$

We turn first to aggregate shocks. The two panels of Figure 8 show the growth rate of aggregate inventories and the estimated time series of aggregate shocks for the durable and nondurable goods sectors. The aggregate shocks in both sectors are positively correlated with the sector-specific aggregate inventory growth rate; the correlation between the two series is 0.66 for the durable goods industries and 0.42 for the nondurable goods industries. Note that aggregate inventory investment in both sectors was subject to a sequence of large negative aggregate shocks during the recessions of the 1980s, indicating that aggregate shocks are a major factor behind economy-wide inventory fluctuations. Nevertheless, the correlation is far from perfect, suggesting that idiosyncratic shocks have a significant effect on aggregate inventory movements.

The two panels in Figure 9 show the time path of the standard deviation and skewness of the cross-sectional distribution of idiosyncratic shocks. (Recall that, by definition, the cross-sectional mean of idiosyncratic shocks is zero for all t.) In both the durable and nondurable goods sectors, the idiosyncratic shocks exhibit substantial dispersion and skewness; moreover, both moments display considerable temporal variation. In particular, note that the skewness coefficient in both sectors was relatively large and negative during the economic turmoil of the late 1970s and early 1980s. This evidence indicates that the impact of the negative aggregate shocks on inventory investment during this period was augmented by a greater than average number of negative idiosyncratic shocks, a pattern consistent with the particularly acute inventory disinvestment during the 1981-82 recession; see Kashyap, Lamont, and Stein (1994) for a case study of the 1981-82 recession.

6 Conclusion

In this paper, we have examined inventory adjustment in the U.S. manufacturing sector, using a long panel of high-frequency firm-level data. Our theoretical framework is based on the generalized (S, s) methodology developed by Caballero, Engel, and Haltiwanger. The major advantage of our approach is that we are able to model explicitly the underlying microeconomic heterogeneity in the inventory adjustment process, while allowing for general nonconvexities in the production technology.

A key result of our paper is that the estimated inventory adjustment function is nonlinear and asymmetric. The nonlinearity of the adjustment function is reflected in the fact that firms with larger absolute deviations from optimal inventory levels adjust proportionally more than do firms with smaller deviations. This finding implies firm-level inventory adjustment consistent with the use of generalized (S, s)–type inventory policies, owing to the presence of nonconvexities in the production technology.

The asymmetric shape of the inventory adjustment function implies that firms with small to moderately sized inventory overhangs adjust less than do firms with similarly sized inventory shortages. This asymmetry could reflect an additional stockout avoidance motive not captured by our model. Alternatively, the asymmetry, which is particularly pronounced in the durable goods sector, could be indicative of technological and/or market irreversibilities in the inventory adjustment process.

The inventory adjustment process appears to differ significantly between recessions and expansions. The cyclical shifts in the estimated inventory adjustment function indicate that firms with excess inventories disinvest more during recessions than they do during expansions. Such macroeconomic state-dependence in (S, s) inventory policies would be consistent with reductions in the inventory target levels, owing to increased demand uncertainty or adverse cost shocks.

Outside of such cyclical shifts, the inventory adjustment process appears to have been relatively stable during our sample period. The estimated adjustment function during the 1980s is very similar to the adjustment function during the 1990s. In contrast, the average optimal inventory-sales ratio has declined significantly during our sample period, particularly in the durable goods sector. These results suggest that despite lower inventory target levels, the adoption of better inventory control methods have had little effect on the adjustment process.

Our conjecture is that recent inventory control innovations have affected relative stockout costs but not the adjustment cost structure. However, additional research into the effects of inventory control innovations on both margins of firm inventory behavior is needed to resolve this dichotomy more fully. Such research may also provide more information concerning the implications of such innovations for future aggregate inventory, and thus business cycle, fluctuations.

References

Abernathy, F. H., J. T. Dunlop, J. H. Hammond, and D. Weil (1999): *A Stitch in Time: Lean Retailing and the Transformation of Manufacturing—Lessons From the Apparel and Textile Industries.* Oxford University Press Inc., New York, NY.

Abramowitz, M. (1950): *Inventories and Business Cycles.* National Bureau of Economic Research, New York, NY.

Ahrens, H., and R. Pincus (1981): "On Two Measures of Unbalancedness in a One-Way Model and Their Relation to Efficiency," *Biometric Journal*, 23, 227–235.

Allen, D. S. (1995): "Changes in Inventory Management and the Business Cycle," *Review, Federal Reserve Bank of St. Louis*, 4, 17–33.

Baltagi, B. H. (1995): *Econometric Analysis of Panel Data.* John Wiley & Sons Ltd., West Sussex, UK.

Baltagi, B. H., and Y. J. Chang (1994): "Incomplete Panels: A Comparative Study of Alternative Estimators for the Unbalanced One-Way Error Components Regression Model," *Journal of Econometrics*, 62, 67–89.

Blanchard, O. J. (1983): "The Production and Inventory Behavior of the American Automobile Industry," *Journal of Political Economy*, 91, 365–400.

Blinder, A. S. (1981): "Retail Inventory Behavior and Business Fluctuations," *Brookings Papers on Economic Activity*, 2, 443–505.

——— (1986): "Can the Production Smoothing Model of Inventories be Saved?," *Quarterly Journal of Economics*, 101, 431–454.

Blinder, A. S., and L. J. Maccini (1991): "Taking Stock: A Critical Assessment of Recent Research on Inventories," *Journal of Economic Perspectives*, 5, 73–96.

Bresnahan, T., and V. Ramey (1994): "Output Fluctuations at the Plant Level," *Quarterly Journal of Economics*, 109, 593–624.

Caballero, R. J. (1999): "Aggregate Investment," in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, pp. 813–862. North-Holland, Elsevier, New York, NY.

Caballero, R. J., and E. M. R. A. Engel (1991): "Dynamic *(S,s)* Economies," *Econometrica*, 59, 1659–1686.

——— (1993): "Microeconomic Adjustment Hazards and Aggregate Dynamics," *Quarterly Journal of Economics*, 108, 359–383.

——— (1999): "Explaining Investment Dynamics in U.S. Manufacturing: A Generalized (S, s) Approach," *Econometrica*, 67, 783–826.

Caballero, R. J., E. M. R. A. Engel, and J. C. Haltiwanger (1995): "Plant-Level Adjustment and Aggregate Investment Dynamics," *Brookings Papers on Economic Activity*, 2, 1–54.

——— (1997): "Aggregate Employment Dynamics: Building from Microeconomic Evidence," *American Economic Review*, 87, 115–137.

Caplin, A. S. (1985): "The Variability of Aggregate Demand with *(S,s)* Inventory Policies," *Econometrica*, 53, 1395–1409.

Diggle, P. J., K.-Y. Liang, and S. L. Zeger (1995): *Analysis of Longitudinal Data.* Clarendon Press, Oxford, UK.

Episcopos, A. (1996): "Testing the (S, s) Model of Inventory Investment with Canadian Wholesale Trade Data," *Applied Economics Letters*, 3, 193–195.

Fair, R. C. (1989): "Production Smoothing Model is Alive and Well," *Journal of Monetary Economics*, 24, 353–370.

Filardo, A. J. (1995): "Recent Evidence on the Muted Inventory Cycle," *Federal Reserve Bank of Kansas City Economic Review*, 2, 27–43.

Hall, G. J. (1999): "Non-Convex Costs and Capital Utilization: A Study of Production Scheduling at Automobile Assembly Plants," Forthcoming, *Journal of Monetary Economics.*

Hall, G. J., and J. Rust (1999): "An (S, s) Model of Commodity Price Speculation," Unpublished paper, Dept. of Economics, Yale University.

Holt, C. C., F. Modigliani, J. F. Muth, and H. A. Simon (1960): *Planning Production, Inventories and Work Force.* Prentice-Hall, Englewood Cliffs, NJ.

Hordijk, A., and F. A. Van Der Duyn Schouten (1986): "On the Optimality of (s, S)–Policies in Continuous Review Inventory Models," *SIAM Journal of Applied Mathematics*, 46, 912–929.

Hsiao, C. (1996): "Random Coefficients Models," in *The Econometrics of Panel Data: A Handbook of the Theory with Applications*, ed. by L. Matyas, and P. Sevestre, pp. 77–99. Kluwer Academic Publishers, Norwell, MA.

Kahn, J. A. (1987): "Inventories and the Volatility of Production," *American Economic Review*, 77, 667–679.

——— (1992): "Why is Production More Volatile Than Sales? Theory and Evidence on the Stockout-Avoidance Motive for Inventory-Holding," *Quarterly Journal of Economics*, 152, 481–510.

KASHYAP, A., O. LAMONT, AND J. C. STEIN (1994): "Credit Conditions and the Cyclical Behavior of Inventories," *Quarterly Journal of Economics*, 109, 565–592.

KRANE, S. 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, 117–136.

KRANE, S. D., AND S. N. BRAUN (1991): "Production Smoothing Evidence from Physical Product Data," *Journal of Political Economy*, 99, 558–581.

MAIRESSE, J., AND Z. GRILICHES (1990): "Heterogeneity in Panel Data: Are There Stable Production Function?," in *Essays in Honor of Edmond Malinvaud (Empirical Economics)*, ed. by P. Champsaur, vol. 3, pp. 192–231. M.I.T. Press, Cambridge, MA.

MCCARTHY, J., AND E. ZAKRAJSEK (1997): "Trade Inventories," Unpublished paper, Federal Reserve Bank of New York.

MCCONNELL, M. M., AND G. P. 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.

MOSSER, P. C. (1988): "Empirical Tests of the *(S,s)* Model for Merchant Wholesalers," in *The Economics of Inventory Management*, ed. by A. Chikan, and M. C. Lovell, pp. 261–283. Elsevier Science Publishers, Amsterdam, The Netherlands.

——— (1991): "Trade Inventories and *(S,s)*," *Quarterly Journal of Economics*, 106, 1267–1286.

RAMEY, V. A. (1991): "Nonconvex Costs and the Behavior of Inventories," *Journal of Political Economy*, 99, 306–334.

RAMEY, V. A., AND K. D. WEST (1999): "Inventories," in *Handbook of Macroeconomics*, ed. by J. B. Taylor, and M. Woodford, pp. 863–923. North-Holland, Elsevier, New York, NY.

RAO, C. R. (1971): "Minimum Variance Quadratic Unbiased Estimation of Variance Components," *Journal of Multivariate Analysis*, 1, 445–456.

SCARF, H. (1960): "The Optimality of *(S,s)* Policies in the Dynamic Inventory Problem," in *Mathematical Methods in the Social Sciences*, ed. by K. J. Arrow, S. Karlin, and H. Scarf, pp. 196–202. Stanford University Press, Stanford, CA.

SCHUH, S. (1996): "Evidence on the Link Between Firm-Level and Aggregate Inventory Behavior," Unpublished paper, Board of Governors of the Federal Reserve System.

SWALLOW, W. H., AND J. F. MONAHAN (1984): "Monte Carlo Comparison of ANOVA, MIVQUE, REML, and ML Estimators of Variance Components," *Technometrics*, 26, 47–57.

VERACIERTO, M. (1998): "Plant-Level Irreversible Investment and Equilibrium Business Cycles," Federal Reserve Bank of Chicago Working Paper 98-1.

Table 1

FGLS Estimates of the Random Coefficients Model

Industry	Parameter γ	θ	σ_e^2	$W_{\Delta S}{}^a$	$W_{\Delta C}{}^b$	$W_{time}{}^c$	$W_{qtr}{}^d$	ω^e	*Obs.*
SIC 200	0.788	-0.699	0.023	0.01	0.01	0.14	0.01	0.78	4,769
	(0.065)	(0.104)							
SIC 220	0.768	-0.914	0.023	0.01	0.01	0.27	0.02	0.81	2,209
	(0.068)	(0.095)							
SIC 230	0.819	-1.037	0.016	0.01	0.01	0.13	0.03	0.77	1,782
	(0.097)	(0.127)							
SIC 250	0.925	-1.242	0.024	0.01	0.01	0.06	0.01	0.81	1,807
	(0.093)	(0.182)							
SIC 260	0.772	-0.468	0.020	0.01	0.02	0.56	0.01	0.78	2,666
	(0.086)	(0.154)							
SIC 270	0.868	-0.225	0.041	0.02	0.06	0.12	0.04	0.77	3,160
	(0.142)	(0.176)							
SIC 280	0.866	-0.398	0.020	0.01	0.01	0.29	0.13	0.79	5,629
	(0.055)	(0.080)							
SIC 283	0.730	-0.349	0.042	0.01	0.01	0.26	0.25	0.82	3,004
	(0.049)	(0.077)							
SIC 290	0.604	-0.506	0.038	0.01	0.10	0.10	0.01	0.79	1,837
	(0.085)	(0.179)							
SIC 300	0.775	-0.887	0.031	0.01	0.01	0.03	0.19	0.79	2,797
	(0.077)	(0.125)							
SIC 320	0.435	-0.835	0.018	0.18	0.01	0.02	0.01	0.84	1,411
	(0.099)	(0.195)							

Notes: Estimation period: 1979:Q1–1997:Q4. Dependent variable is the log-level of real end of period t inventories h_{it}. All industry-specific regressions include quarterly seasonal effects, linear and quadratic time trends (individual parameter estimates not reported) and are estimated with FGLS, using MIVQUE(0) (Minimum Variance Quadratic Unbiased Estimator) of the covariance matrix V. Heteroscedasticity-consistent asymptotic standard errors are reported in parenthesis.

[a]Probability value for the Wald test of the null hypothesis that the coefficients on $\Delta s_{it}, \ldots, \Delta s_{it-3}$ are jointly equal to zero.

[b]Probability value for the Wald test of the null hypothesis that the coefficients on $\Delta c_{it}, \ldots, \Delta c_{it-3}$ are jointly equal to zero.

[c]Probability value for the Wald test of the null hypothesis that the linear and quadratic time trends are jointly equal to zero.

[d]Probability value for the Wald test of the null hypothesis that the quarterly seasonal effects are jointly equal to zero.

[e]A measure of panel unbalancedness given by Ahrens and Pincus (1981): $\omega = N/\overline{T}\sum(\frac{1}{T_i})$, where $\overline{T} = \sum T_i/N$. Note that $0 \leq \omega \leq 1$, with $\omega = 1$ when the panel is balanced.

Table 1 (continued)

FGLS Estimates of the Random Coefficients Model

Industry	Parameter γ	θ	σ_e^2	$W_{\Delta S}$ [a]	$W_{\Delta C}$ [b]	W_{time} [c]	W_{qtr} [d]	ω^e	*Obs.*
SIC 330	0.605	-0.803	0.026	0.01	0.01	0.29	0.01	0.76	3,289
	(0.081)	(0.130)							
SIC 340	0.770	-1.258	0.011	0.01	0.01	0.13	0.59	0.78	4,204
	(0.055)	(0.121)							
SIC 350	0.736	-0.811	0.029	0.01	0.01	0.24	0.01	0.81	7,402
	(0.034)	(0.075)							
SIC 357	0.804	-0.525	0.059	0.01	0.01	0.39	0.01	0.81	4,361
	(0.046)	(0.076)							
SIC 360	0.820	-0.826	0.022	0.01	0.02	0.15	0.01	0.83	4,444
	(0.053)	(0.076)							
SIC 366	0.690	-0.889	0.038	0.02	0.06	0.83	0.09	0.81	3,379
	(0.066)	(0.123)							
SIC 367	0.621	-0.568	0.036	0.01	0.01	0.09	0.58	0.80	4,677
	(0.049)	(0.089)							
SIC 370	0.729	-0.929	0.031	0.01	0.01	0.03	0.01	0.79	4,259
	(0.064)	(0.133)							
SIC 380	0.652	-0.936	0.026	0.01	0.01	0.41	0.17	0.81	1,673
	(0.080)	(0.116)							
SIC 382	0.714	-0.584	0.027	0.01	0.01	0.33	0.60	0.81	4,973
	(0.049)	(0.072)							
SIC 384	0.695	-0.452	0.049	0.01	0.01	0.83	0.15	0.81	3,628
	(0.054)	(0.081)							
SIC 390	0.683	-1.081	0.032	0.01	0.01	0.35	0.02	0.79	3,958
	(0.056)	(0.106)							

Notes: Estimation period: 1979:Q1–1997:Q4. Dependent variable is the log-level of real end of period t inventories h_{it}. All industry-specific regressions include quarterly seasonal effects, linear and quadratic time trends (individual parameter estimates not reported) and are estimated with FGLS, using MIVQUE(0) (Minimum Variance Quadratic Unbiased Estimator) of the covariance matrix V. Heteroscedasticity-consistent asymptotic standard errors are reported in parenthesis.

[a]Probability value for the Wald test of the null hypothesis that the coefficients on $\Delta s_{it}, \ldots, \Delta s_{it-3}$ are jointly equal to zero.

[b]Probability value for the Wald test of the null hypothesis that the coefficients on $\Delta c_{it}, \ldots, \Delta c_{it-3}$ are jointly equal to zero.

[c]Probability value for the Wald test of the null hypothesis that the linear and quadratic time trends are jointly equal to zero.

[d]Probability value for the Wald test of the null hypothesis that the quarterly seasonal effects are jointly equal to zero.

[e]A measure of panel unbalancedness given by Ahrens and Pincus (1981): $\omega = N/\overline{T}\sum(\frac{1}{T_i})$, where $\overline{T} = \sum T_i/N$. Note that $0 \leq \omega \leq 1$, with $\omega = 1$ when the panel is balanced.

Figure 1: Average Optimal Inventory-Sales Ratio

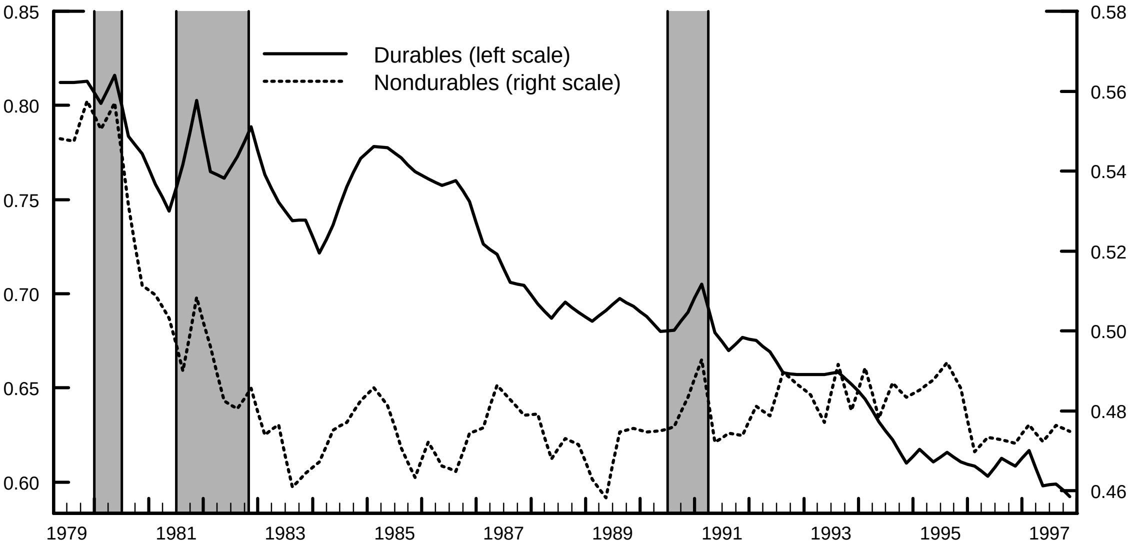

Seasonally adjusted data. Shaded regions indicate NBER-dated recessions.

Figure 2: Inventory Adjustment Function

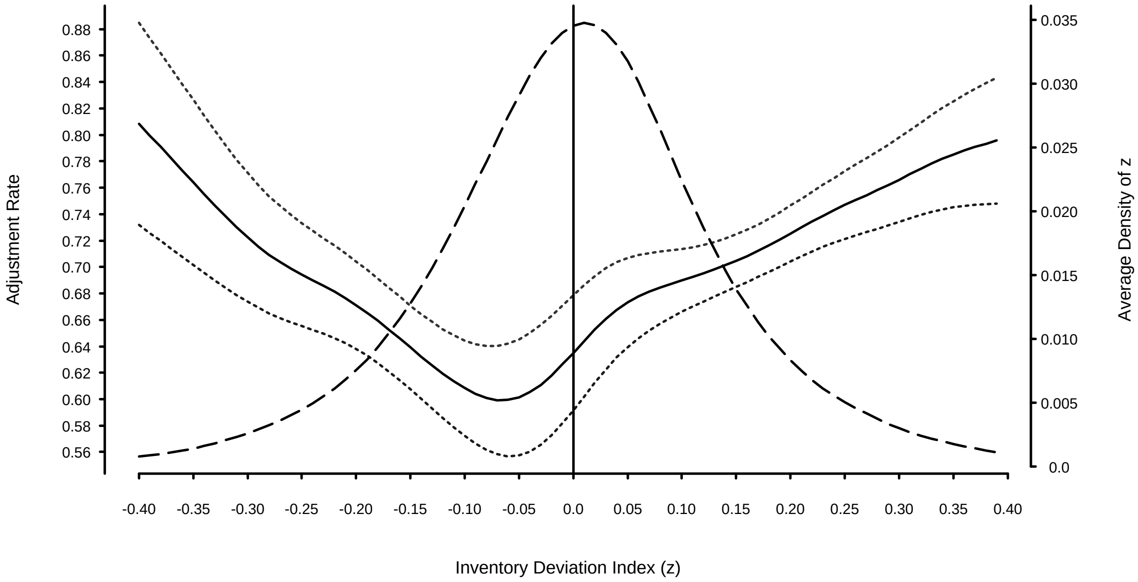

Figure 3: Sector-Specific Inventory Adjustment Functions

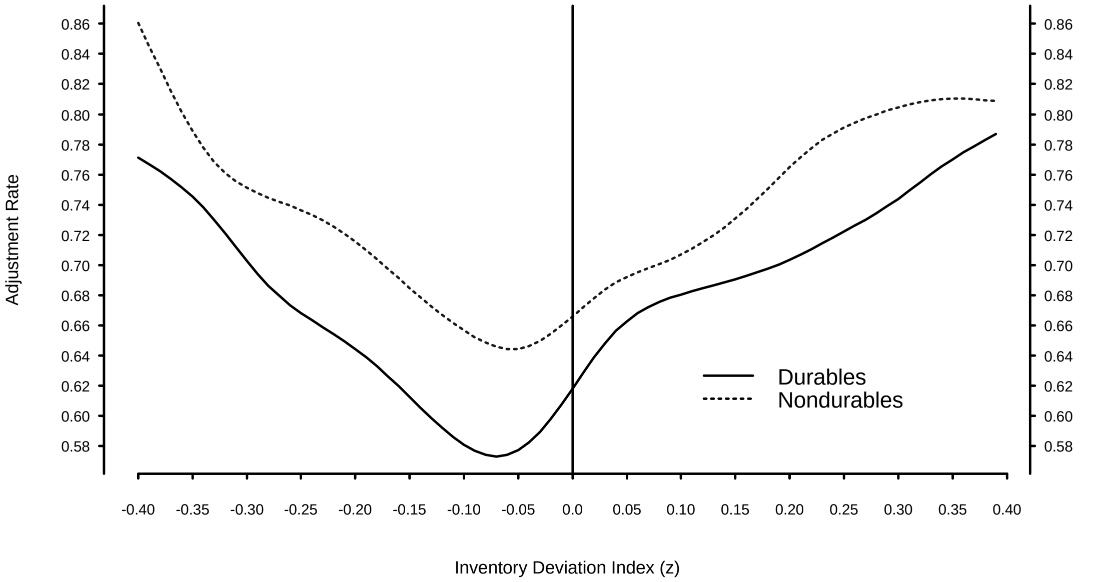

Figure 4: Cyclical Shifts in the Adjustment Function

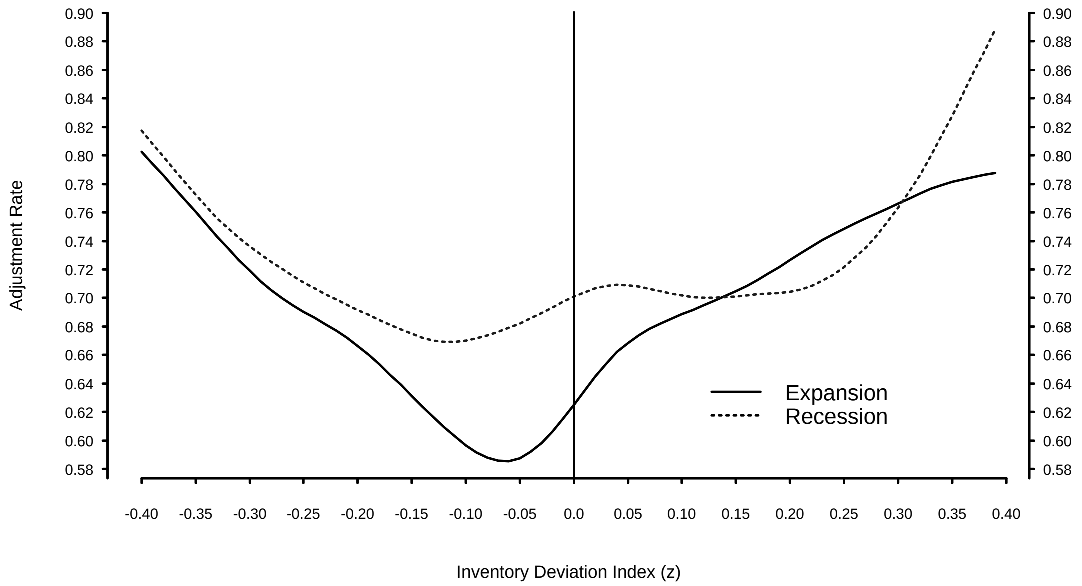

Figure 5: Decade-Specific Inventory Adjustment Functions

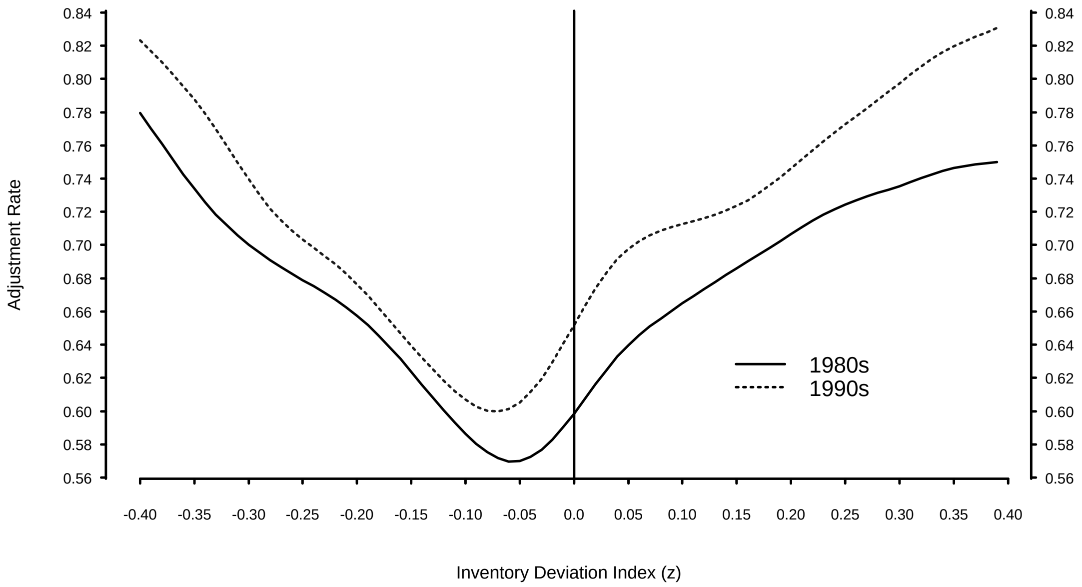

Figure 6: Aggregate Inventory Growth and Inventory Deviation Index

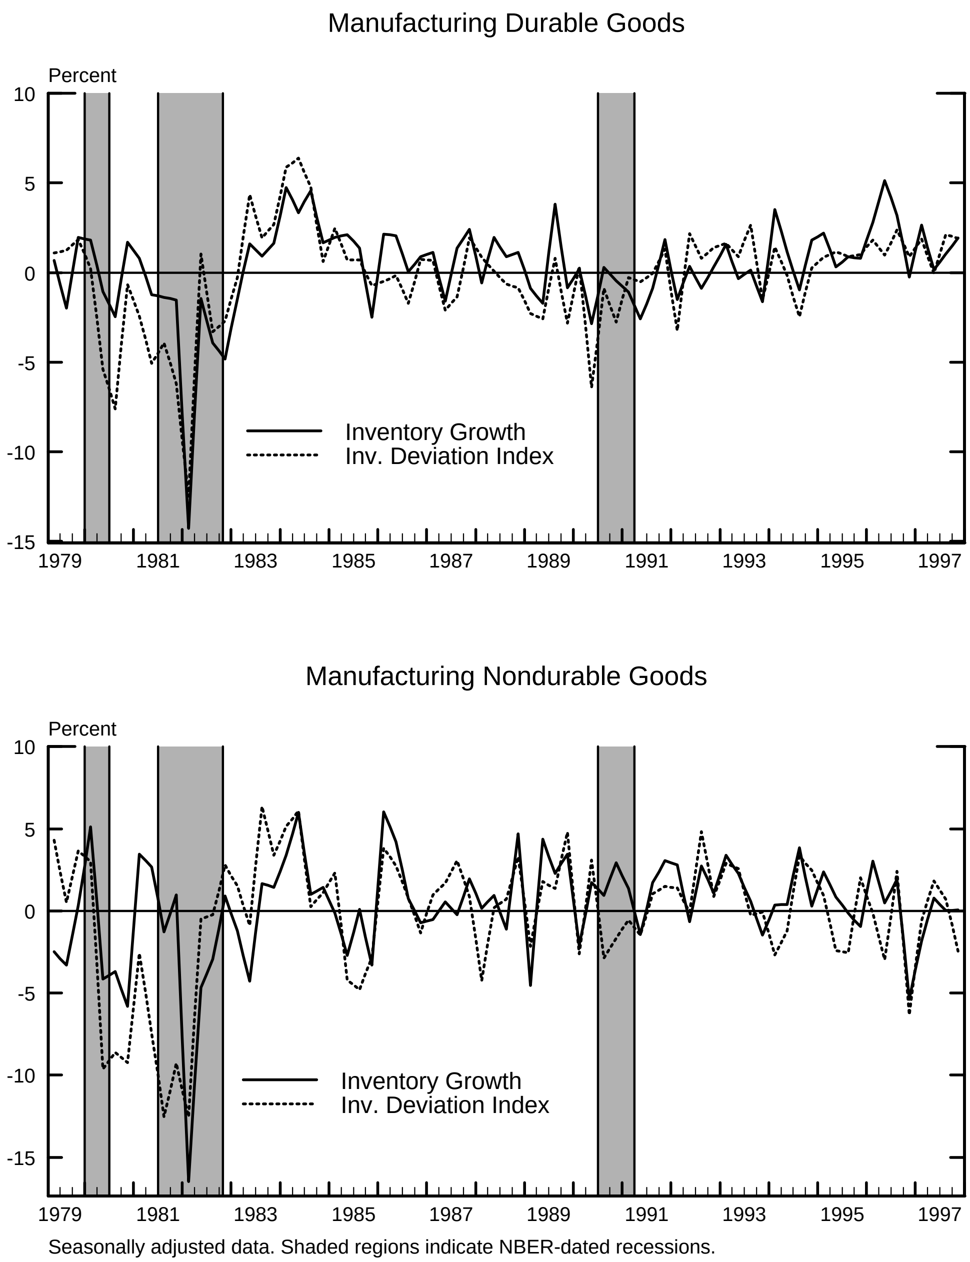

Seasonally adjusted data. Shaded regions indicate NBER-dated recessions.

Figure 7: Cross-Sectional Moments of Inventory Deviations

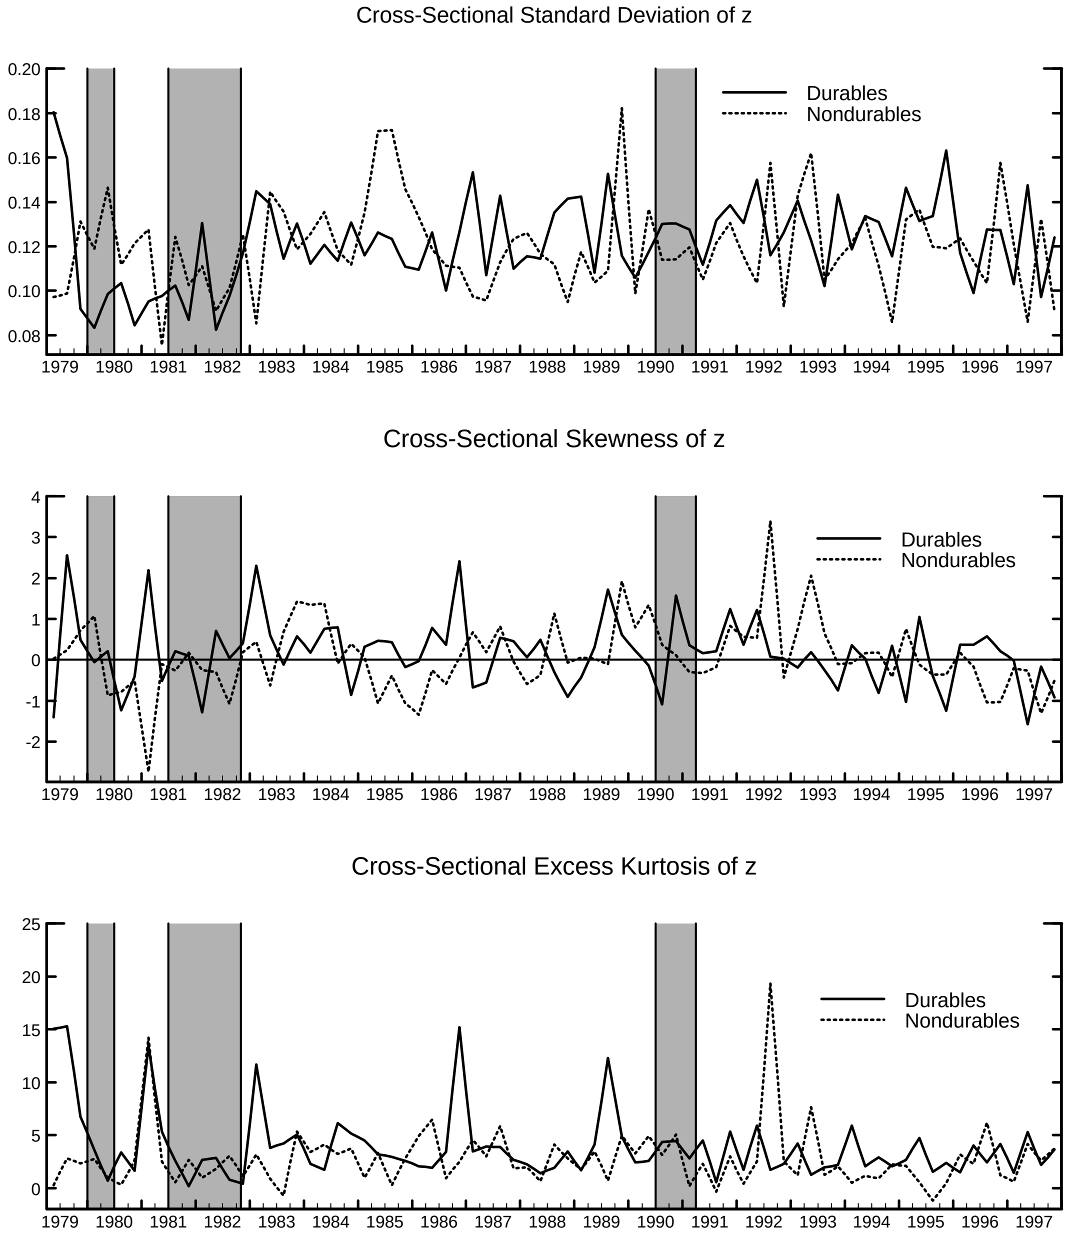

Seasonally adjusted data. Shaded regions indicate NBER-dated recessions.

Figure 8: Aggregate Inventory Growth and Aggregate Shocks

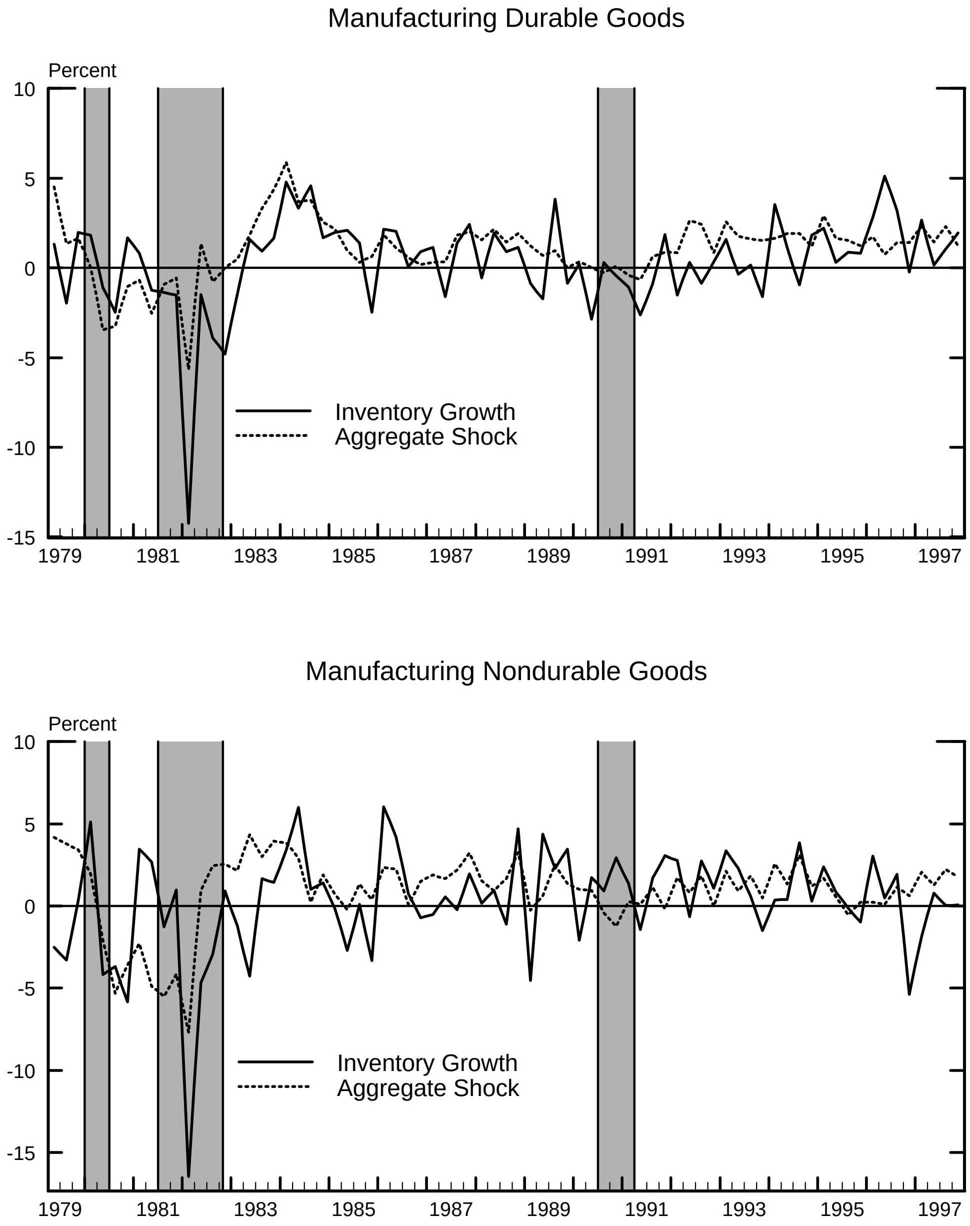

Seasonally adjusted data. Shaded regions indicate NBER-dated recessions.

Figure 9: Cross-Sectional Moments of Idiosyncratic Shocks

Standard Deviation of Idiosyncratic Shocks

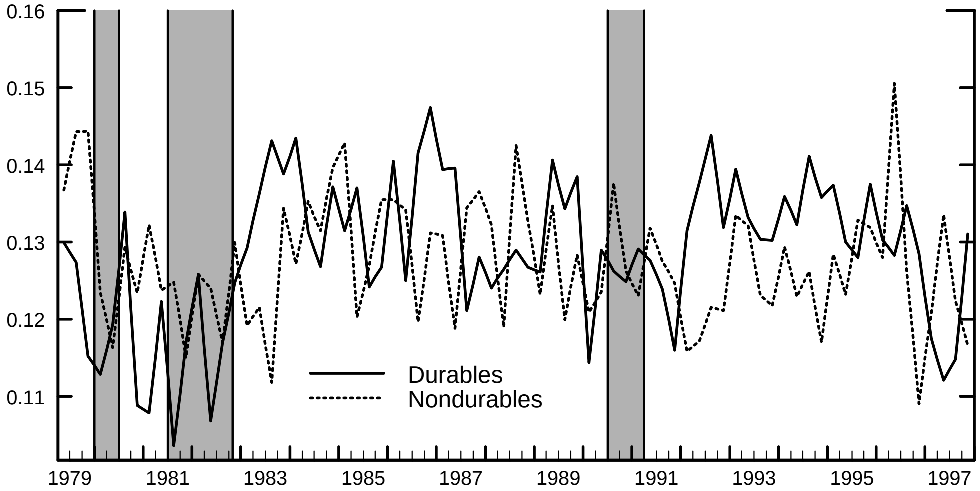

Skewness of Idiosyncratic Shocks

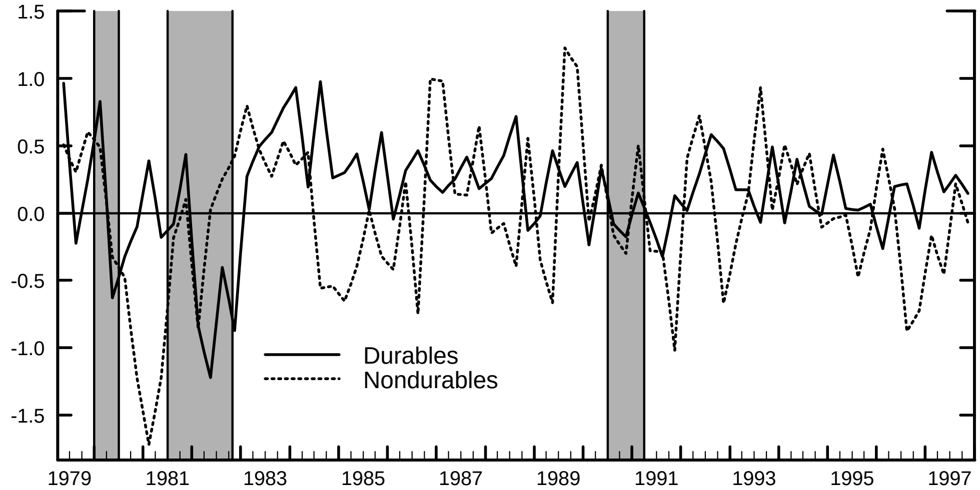

Seasonally adjusted data. Shaded regions indicate NBER-dated recessions.

A Data Appendix

This section describes the selection rules used to construct our firm-level panel and the construction of the variables used in the analysis. The data for our paper come from the quarterly P/S/T, Full Coverage, and Research COMPUSTAT data files. The firm-level COMPUSTAT data 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 20 continuous quarters of data between 1978Q1 and 1997Q4. To avoid results that are driven by a small number of extreme observations, three criteria were used to eliminate firms with substantial outliers or obvious errors:

1. If a firm's estimate of (real) gross output from the accounting identity $Y \equiv S + \Delta H$ was negative at any point during a firm's tenure in the sample, a firm was eliminated in its entirety.
2. If a firm's quarterly growth rate of real inventories was above (below) the 99.5 (0.5) percentile of the distribution in any period during the firm's tenure in the panel, the firm was eliminated in its entirety.
3. If a firm's quarterly growth rate of real sales was above (below) the 99.5 (0.5) 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, 911 firms were eliminated from the original panel.[25] Table A.1 provides a detailed industry breakdown of our panel. All of our industry groups contain more than 1,500 firm/quarter observations. The sparsest industry in our data set is SIC 290 (Petroleum & Coal Products), which contains only 40 firms. On the other hand, SIC 350 (Industrial Machinery & Equipment) contains 189 firms. Finally, as mentioned in the text, the panels are unbalanced, with firms entering and exiting the data set. The lowest average industry-specific tenure in the panel is almost 38 quarters (SIC 357: Computers & Office Equipment), and the highest is almost 51 quarters (SIC 290: Petroleum & Coal Products).

[25] Over 3/4 of eliminated firms were deleted because of the second and third selection criteria. Visual inspection of the eliminated firms revealed severe anomalies and likely errors in their reported data.

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. This implies that the replacement value of inventory stocks equals their book value. To convert the reported nominal value of inventories to real terms, inventory stocks were deflated by the sector-specific (i.e., durable and nondurable) implicit (1992=100) inventory deflator. The inventory stocks are measured as of the end of the period.

- **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.

- **Gross Output**: An estimate of a firm's (real) gross output in period t, Y_t, was obtained from the accounting identity $Y_t \equiv S_t + \Delta H_t$, where S_t denotes real final sales in period t, and ΔH_t denotes real inventory investment in period t.[26]

- **Cost per Unit of Output**: A real cost per unit of output was constructed by converting the nominal Cost of Goods Sold to real terms using the implicit (1992=100) GDP deflator. The ratio of real total costs to real gross output is our measure of the average cost per unit of output.

- All other variables were deflated by the implicit (1992=100) GDP deflator.

Table A.2 provides summary statistics for the key variables used in our analysis. Because all firms in the sample are publicly traded, most of them are relatively large. The median firm size, measured by total assets, is $125 million. The distributions of most variables display considerable skewness—the means of inventories, sales, (gross) 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.

[26]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."

Table A.1

Industry Composition

Industry Classification	# of Firms	*Avg.* T_i	*Med.* T_i	*Obs.*
SIC 200: Food & Kindred Prod.[a]	127	42.6	38.0	5,404
SIC 220: Textile Mill Prod.	67	38.0	30.0	2,544
SIC 230: Apparel & Other Prod.	52	39.3	30.0	2,042
SIC 250: Furniture & Fixtures	48	42.7	44.0	2,047
SIC 260: Paper & Allied Prod.	62	48.0	44.5	2,976
SIC 270: Printing & Publishing	75	47.1	40.0	3,535
SIC 280: Chemical & Allied Prod.[b]	128	49.0	46.0	6,269
SIC 283: Drugs	81	42.1	36.0	3,409
SIC 290: Petroleum & Coal Prod.	40	50.9	43.5	2,037
SIC 300: Rubber & Misc. Plastic Prod.	81	39.5	34.0	3,202
SIC 320: Stone, Clay & Glass Prod.	43	37.8	34.0	1,626
SIC 330: Primary Metal Industries	84	44.2	37.5	3,709
SIC 340: Fabricated Metal Prod.	113	42.2	38.0	4,769
SIC 350: Industrial Machinery & Equip.[c]	189	44.2	39.0	8,347
SIC 357: Computers & Office Equip.	134	37.5	34.0	5,031
SIC 360: Electronic & Other Electric Equip.[d]	121	41.7	39.0	5,049
SIC 366: Communications Equip	95	50.6	37.0	3,854
SIC 367: Electronic Components	125	42.4	38.0	5,302
SIC 370: Transportation Equip.	112	43.0	37.0	4,819
SIC 380: Instruments & Related Prod.[e]	50	38.5	33.5	1,923
SIC 382: Measuring & Controlling Prod.	131	43.0	39.0	5,628
SIC 384: Medical Instruments	108	38.6	34.5	4,168
SIC 390: Misc. Manufacturing Industries[f]	103	43.4	40.0	4,473

[a]Includes SIC 21 (Tobacco and Tobacco Products).
[b]Excludes SIC 283 (Drugs).
[c]Excludes SIC 357 (Computers and Office Equipment).
[d]Excludes SIC 366 (Communications Equipment) and SIC 367 (Electronic Components).
[e]Excludes SIC 382 (Measuring and Controlling Products) and SIC 384 (Medical Instruments).
[f]Includes SIC 24 (Lumber & Wood Products) and SIC 31 (Leather & Leather Products).

Table A.2

Summary Statistics for All Industries

Variable	*Mean*	*Std.Dev.*	*Median*	*Minimum*	*Maximum*
Inventories	191.4	661.5	24.5	0.039	12,660.6
Net Sales	403.5	1699.7	39.6	0.012	39,335.0
Gross Output[a]	408.0	1,708.9	40.4	0.010	38,280.7
Total Assets	1,588.5	7,955.1	125.0	0.235	259,303.0
Inv. Investment	0.84	58.8	0.09	-2,444.7	3,077.6
Inv/Sales Ratio	0.73	0.46	0.64	0.01	15.4
Average Cost[b]	0.73	0.44	0.72	0.25	49.4
Inv. Growth Rate (%)	1.31	13.3	1.04	-72.1	82.5
Sales Growth Rate (%)	1.49	17.9	1.80	-112.3	114.5
# of Firms	2,169				
Observations	92,163				

Notes: Sample period: 1978:Q1–1997:Q4. All variables are in millions of 1992 dollars.

[a]Real gross output in period t, Y_t, was estimated from the accounting identity $Y_t \equiv S_t + \Delta H_t$, where S_t denotes (real) final sales in period t, and ΔH_t denotes (real) inventory investment in period t.

[b]Real average cost is defined as the ratio of (real) cost of goods sold to (real) gross output.