Evidence on the Link Between Firm-Level and Aggregate Inventory Behavior

Scott Schuh[1]

Federal Reserve Board

November, 1996

[1]20th & Constitution Ave., N.W., Mail Stop 61, Washington, D.C. 20551 (Internet: sschuh@frb.gov).
This paper is a revised version of parts of my Ph.D. dissertation and was previously circulated under the title "Inventory Models and Aggregation." I thank my advisors, Louis Maccini, John Haltiwanger, and David Bizer, as well as Charlie Bates, Steve Blough, Ricardo Caballero, Carl Christ, Steve Davis, Mark Doms, Timothy Dunne, Steve Durlauf, Robert McGuckin, Steve Olley, Ruth Runyan, and Thomas Stoker for helpful comments. I thank the U.S. Census Bureau's M3 Branch for the data and Center for Economic Studies for financial, computing, and administrative support. Census employees Steve Andrews, Robert Bechtold, Lisa Feldman, Louise Fry, Sam Jones, Cyr Linonis, Kathy Menth, Jim Monahan, Eddie Salyers, and Robert Taylor were especially helpful. I thank Steve Baldwin, Leonard Loebach, Greg Won, and Mary Young at the U.S. Bureau of Economic Analysis for deflators and aggregate data. Mike Fratantoni and Brian Doyle provided excellent research assistance. The research in this paper was conducted while the author was a research associate at the Center for Economic Studies, U.S. Bureau of the Census. Research results and conclusions expressed are those of the author and do not necessarily indicate concurrence by the Bureau of the Census, the Center for Economic Studies, or the Federal Reserve System.

Abstract

This paper describes the finished goods inventory behavior of more than 700 U.S. manufacturing firms between 1985-93 using a new Census Bureau longitudinal data base. Three key results emerge. First, there is a broad mix of production-smoothing and production-bunching firms, with about two-fifths smoothing production. Second, firm-level inventory adjustment speeds are about an order of magnitude larger than aggregate adjustment speeds due to econometric aggregation bias. Finally, accounting for time variation in the inventory adjustment speed due to fluctuations in firm size improves the fit of a traditional aggregate inventory model by one-fifth.

JEL Classifications: E22, D21, C23, C43

Keywords: Inventory investment, production smoothing, adjustment speed, aggregation bias

Introduction

Linear-quadratic (LQ) inventory models in the spirit of Holt *et al* (1960) have not fared well empirically. As Blinder and Maccini (1991) point out, there exists a tension between microeconomic and macroeconomic views of inventory behavior. Microeconomic theory suggests that firms should use inventories as buffer stocks, allowing production to be smoothed relative to fluctuating demand. Aggregate data, however, indicate that production is more variable than sales and hence inventories appear to have an accelerator effect that generates cycles. Not surprisingly, then, econometric inventory models based on microeconomic theory and applied to aggregate data have failed miserably, producing a host of puzzling results.

Most previous attempts to reconcile this tension rest on one of three explanations. One explanation is that the microeconomic theory is wrong and the model must be modified to generate more variable production. Stockout avoidance, increasing returns to scale and (S,s) ordering policies are the leading alternatives offered in the literature. Another explanation is that the standard Department of Commerce inventory data contain measurement error that leads to faulty inference about inventory behavior. A third explanation is that econometric inference is faulty, for example with respect to the estimation of adjustment speeds and of structural (Euler equation) models. Unfortunately, however, none of these explanations has succeeded in resolving the tension.

A fourth explanation—aggregation—has been relatively unexplored. However, recent research hints that aggregation across agents may be responsible, at least in part, for the tension between microeconomic theory and macroeconomic evidence. For example, Hunt (1981), Blinder (1986b), Seitz (1993), and Lovell (1993) provide some incomplete indications that aggregation across firms may bias adjustment speed estimates downward.[1] Also, Krane (1994) and Lai (1991) show theoretically how aggregate data can exhibit production bunching even though firms smooth production. Evidence from disaggregated industries with physical units data, such as Ghali (1987), Fair (1989), and Krane and Braun (1991), provides some evidence this hypothesis may be true. Each of these studies, however, has been severely limited by a lack of broad-based firm-level data.

This paper provides new evidence on firm-level inventory behavior and examines the hypothesis that aggregation effects are responsible for the poor performance of applied LQ inventory models. The evidence comes from a new longitudinal data base developed by the Census Bu-

[1] Adjustment speeds may be biased downward for other reasons as well. Christiano and Eichenbaum (1987) cite temporal aggregation, Irvine (1988) cites misspecification, and Bivin (1989) cites product differentiation and market spillovers. This paper does not consider these interesting reasons.

reau called the M3 Longitudinal Research Database (M3LRD). As the name suggests, this data base is closely related to the Census Bureau's well-known LRD except that the M3LRD includes a more limited range of data (the Manufacturers' Shipments, Inventories and Orders survey) and primarily covers an arbitrary sample of companies rather than a probability sample of plants.[2] The M3LRD contains monthly data on shipments (henceforth, sales), stage-of-fabrication inventories, and orders for about 4,300 companies and 8,200 divisions in U.S. manufacturing during the period 1985-93. In terms of number of firms, industrial coverage, periodicity, disaggregation of inventory types, and data reliability, the M3LRD is the most comprehensive source of microeconomic data on inventory behavior available in the United States.

In this first study to exploit the M3LRD, the focus is on estimating traditional LQ inventory models at the firm level and on the relationship between the firm-level and aggregate results.[3] The investigation centers on a panel of more than 700 continuously operating firms, for which individual time-series—not cross-section—regression models are estimated. Since the LQ model primarily is suited for inventories of finished goods and since most empirical applications have focused on these inventories, this study concentrates on finished goods inventories as well.

At the microeconomic level, two key questions from the inventory literature are addressed: (1) do firms smooth production?; and (2) do traditional LQ models fit firm-level inventory data well? A central aspect of the second question is whether or not firms adjust their inventory stocks to their target levels at plausible rates. Answers to these questions about microeconomic behavior are then compared with answers obtained from aggregate data. The goal of this comparison is to determine the extent to which aggregation may adversely affect inference about the suitability of traditional inventory models.

Figure 1 summarizes the paper's two main results for all manufacturing firms. First, firms exhibit a wide array of production smoothing and nonsmoothing behavior. The upper panel plots distributions of firm-level ratios of production variance to sales variance. Ratios of less than 1.0 traditionally are interpreted as production smooothing behavior, so about two-fifths of all firms smooth production. But the median variance ratios (1.06 for companies and 1.04 for divisions) are consistent with ratios constructed from aggregate data reported in previous studies, which

[2]The name M3 originates from the alphanumeric Census label of the survey form sent to the companies.

[3]Additional studies in progress exploit the M3LRD to examine other questions and issues about microeconomic inventory behavior. Schuh (1992) tests for evidence of nonconvexity in inventory behavior by estimating an adjustment hazard model of the type developed by Caballero and Engle (1993). Maccini and Schuh (1995) examine the role of financial market conditions on stage-of-fabrication inventory behavior.

Figure 1
Summary of Main Results

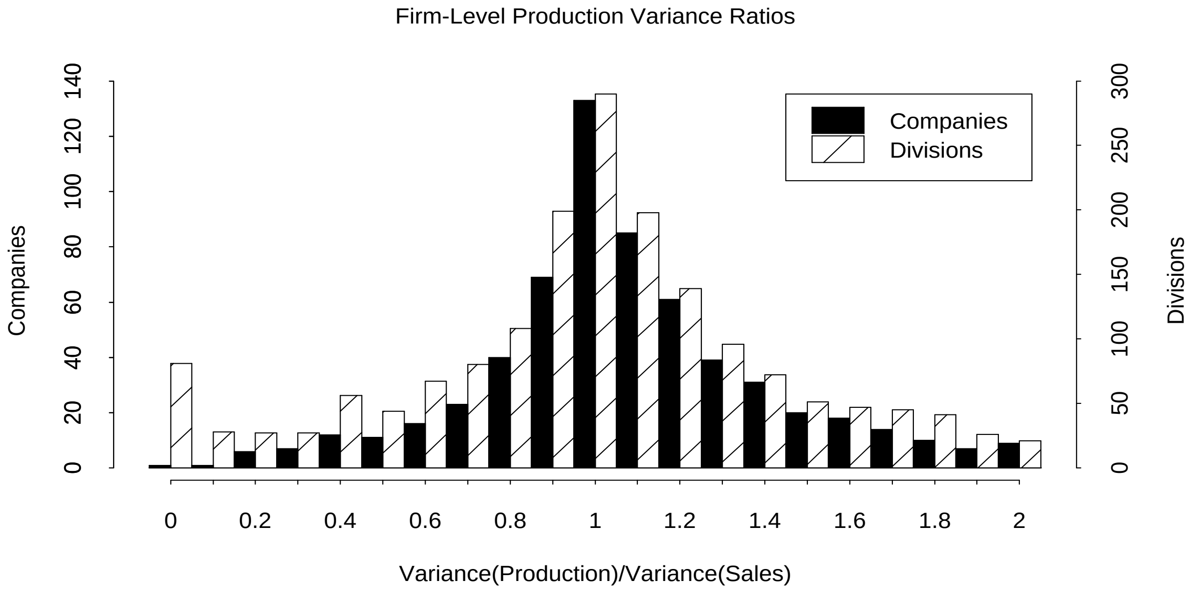

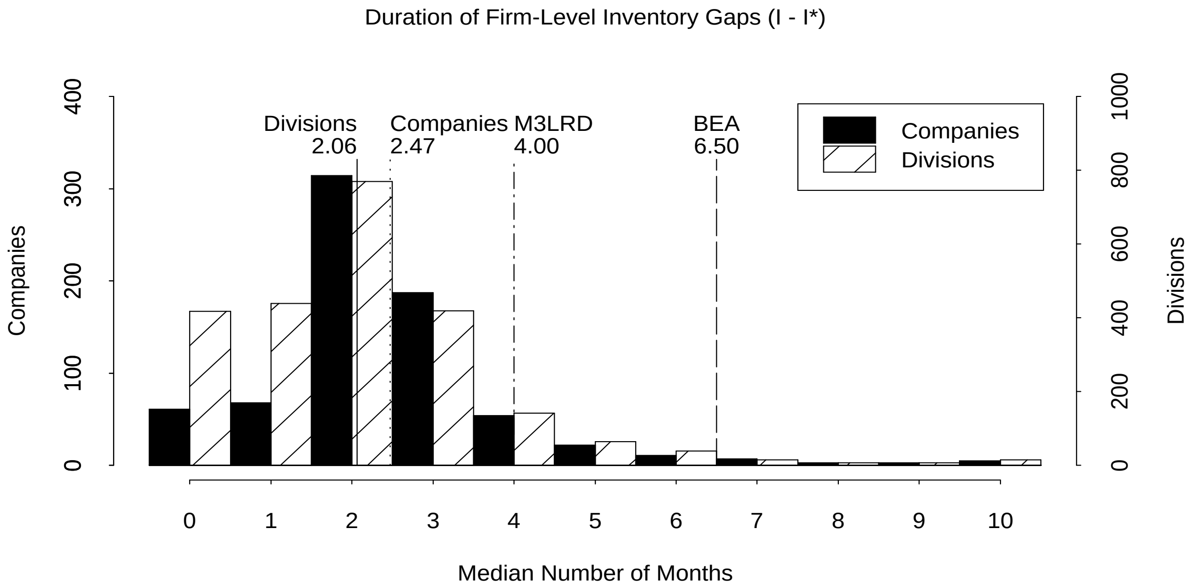

typically exceed 1.0 a bit as well. Thus, on average, a strict production smoothing inventory model is inconsistent with the firm-level data on this score as well. Nevertheless, the distributions generally are consistent with a standard LQ inventory model that incorporates both production smoothing and stockout avoidance motives. Further, evidence presented later shows that firms smooth production more at seasonal frequencies than nonseasonal frequencies.

The second result is that firms eliminate the gap between actual and target inventory stocks much more quickly than is apparent from aggregate data. The lower panel plots distributions of the duration of firm-level inventory gaps, defined as the difference between actual (I) and target (I^*) inventory stocks, where the target is derived from a typical stock adjustment inventory model. The bars represent the median number of months that each firm experiences consecutive positive or consecutive negative gaps ($I - I^*$)—that is, how long it usually takes for firms to elimate inventory gaps. The vertical lines compare means of these distributions with median durations for the aggregate M3LRD data and for the published aggregate data from the Bureau of Economic Analysis (BEA). Inventory gaps clearly persist longer on average in the two aggregate data series. The BEA data used in most previous applied inventory studies imply that the representative firm experiences inventory gaps about three times longer than the median firm. This finding seems to explain the puzzle pointed out by Feldstein and Auerbach (1976) that inventory adjustment speeds are implausibly slow.

A host of more detailed results are presented throughout the paper, but some general implications for inventory theory are worth emphasizing. First, the broad mix of production smoothing and nonsmoothing firms argues against idea that "bad" data lead to the rejection of the basic production smoothing model. Instead, it appears that inventory theory must be able to encompass both types of behavior. Second, except for more plausible adjustment speed estimates, traditional LQ inventory models fare no better econometrically with firm-level data than with aggregate data. In fact, these models only account for a small fraction of firm-level inventory variation. Third, the vast heterogeneity in variance ratios and adjustment speeds generally is not well-explained by observable firm characteristics such as industry and size. Together, these latter two implications suggest there remains considerable room for improved understanding of inventory fluctuations.

The paper proceeds as follows. Section 1 lays out the standard LQ inventory model and the econometric variants of it used in this study. Section 2 introduces and describes the new M3LRD. Section 3 presents evidence and results from the firm-level data. Section 4 presents evidence and results from the aggregate data, and compares them with the firm-level results. Section 5 summarizes and concludes.

1 The Model

The econometric models used in this study are derived from a simple version of the standard LQ model for finished goods inventories (see Blinder and Maccini (1991) or West (1993) for surveys of the inventory literature). The representative firm chooses the inventory level to minimize the present discounted value of costs subject to a production identity. Formally, the model is:

$$\min_{\{I_{t+j}\}} E_0 \left\{ \sum_{j=0}^{\infty} \beta^j \left[C_Q(Q_{t+j}) + C_I(I_{t+j},\, S_{t+j}) \right] \right\} \tag{1}$$

$$\text{s.t.} \qquad I_{t+j} = I_{t+j-1} + Q_{t+j} - S_{t+j} \tag{2}$$

$$C_Q(Q_{t+j}) = (\delta/2) Q_{t+j}^2 \tag{3}$$

$$C_I(I_{t+j};\, S_{t+j}) = (\phi/2)(I_{t+j} - \omega S_{t+j})^2 \tag{4}$$

given I_0 and the S_t process.[4] The notation is: I_t is real end-of-period finished goods inventories; S_t is real sales; Q_t is real production; β is the constant discount factor; Δ is the first difference operator; $E_t = E[.|\Omega_t]$ is the mathematical expectations operator; and Ω_t is the firm's information set, which includes all variables dated $t-1$ and earlier plus any known at t. Structural parameters δ, ϕ, and ω are assumed to be positive constants. For simplicity, production adjustment costs, $(\alpha/2)(\Delta Q_t)^2$, are assumed to be zero even though some studies — Blanchard (1983) and Fuhrer, Moore, and Schuh (1995) — have found marginally significant evidence for them in the data. Sales are assumed to be exogenous, as in most applied inventory research, so $S_t \notin \Omega_t$. The latter assumption implies that I_t is a buffer stock because the firm cannot alter its production plan during the period t.

The model's structural parameters determine the relative variability of production and sales in the model. The sign of parameter δ determines the sign of the marginal cost of production, hence the extent to which there is a production smoothing motive due to convex costs. If $\omega = 0$ and all other parameters are positive, the firm's optimal inventory decision rule will lead to production smoothing (i.e., production variance less than sales variance). But if $\omega > 0$, as advocated by Kahn

[4]Constants, linear terms and white noise shocks are omitted for simplicity here but included as necessary in the econometric work reported in sections 3 and 4.

(1987), the firm has a stockout avoidance motive and the decision rule may lead to production bunching (i.e., production variance greater than sales variance). Given $\omega > 0$, the smoothing or bunching of production depends on relative magnitudes of δ and ϕ (see Krane (1994) on this point).[5]

The model is solved in the usual manner. The Euler equation for period t is

$$E_t \{\delta Q_t + \phi(I_t - \omega S_t) - \beta\delta Q_{t+1}\} = 0 \ , \tag{5}$$

which shows that the firm chooses I_t to balance the cost of producing a unit of output today and storing it one period against the cost of producing the unit of output next period. Following Eichenbaum (1989), substitute for production and re-express equation (5) as the second-order stochastic difference equation

$$E_t \left\{(1 - \lambda L)(1 - (\lambda\beta)^{-1} L) I_{t+1}\right\} = E_t \{-S_{t+1} + (\kappa/\beta) S_t\} \tag{6}$$

where λ and $(\lambda\beta)^{-1}$ are the roots of the characteristic equation $L^2 + \psi L + \beta^{-1}$, $\psi = -[1 + \beta + (\phi/\delta)]$, and $\kappa = 1 - (\phi\omega/\delta)$. The parametric assumptions imply that $\kappa \leq 1$ and stability requires that $|\lambda| < 1$.

Equation (6) is closely related to the traditional stock adjustment model of Lovell (1961), which has been used in many previous studies. The solution to equation (6) is the decision rule

$$I_t = \lambda I_{t-1} - S_t + (1 - \kappa\lambda) E_t S_t + (1 - \kappa\lambda) \sum_{i=1}^{\infty} (\lambda\beta)^i E_t S_{t+i} \tag{7}$$

which is expressed in terms of actual rather than expected inventories.[6] A simple version of the stock adjustment model is

$$\Delta I_t = \mu(I_t^* - I_{t-1}) - \pi(S_t - E_t S_t). \tag{8}$$

If the target inventory stock is $I_t^* = \nu E_t S_t$, then the model becomes

$$I_t = (1 - \mu) I_{t-1} - \pi S_t + (\mu\nu + \pi) E_t S_t. \tag{9}$$

[5]The timing of sales in the target stock specification, $C_I(\cdot)$, varies across models in the literature between S_t and S_{t+1}. Given the crude nature of the approximation it is hard to argue strenuously for one specification or the other. I find S_t more appropriate for approximating stockout costs in period t as opposed to period $t+1$, but the choice is not crucial to the econometric results.

[6]The expected inventory policy becomes the actual inventory policy as follows. Take the expectation of the inventory law of motion, equation 2: $E_t I_t - I_{t-1} = Q_t - E_t S_t$ (Q_t is known at the beginning of period t). Then rewrite the law of motion as $\Delta I_t = (Q_t - E_t S_t) + (E_t S_t - S_t)$ and substitute $E_t I_t - I_{t-1}$ for $(Q_t - E_t S_t)$ in the law of motion. Finally, substitute the solution to equation (6) for $E_t I_t$ into the law of motion and rearrange.

Equations 7 and 9 are observationally equivalent if $\pi = 1$ and either $\omega = 0$ or $E_t S_{t+i} = 0 \; \forall \; i \geq 1$. These restrictions further imply that $\lambda = (1 - \mu)$ and $(1 - \kappa\lambda) = (\mu\nu + \pi)$.[7]

A third, and related, version of this inventory model arises when the data are nonstationary. In the event that the inventory and sales data are integrated of order one (i.e., $I_t \sim$ I(1) and $S_t \sim$ I(1)), and their linear combination

$$z_t = I_t - I_t^* = I_t - AS_t \tag{10}$$

is integrated of order zero (i.e., $z_t \sim$ I(0)), then inventories and sales are cointegrated. Nickell (1985) shows that this kind of model can be written as an error correction equation:

$$\Delta I_t = \alpha - \gamma z_{t-1} + \sum_{i=1}^{L_I} \theta_{Ii} \Delta I_{t-i} + \sum_{i=1}^{L_S} \theta_{Si} \Delta S_{t-i} + \eta_t \tag{11}$$

For the model described in equations (1) through (4), $L_I = 0$; if the model includes adjustment costs due to changing production, then $L_I > 0$. Equation (11) admits any exogenous univariate sales process, which in turn determines L_S. Unlike the decision rule and stock adjustment versions of the model, the error correction model does not require specification of a sales process to estimate the parameters. In this study equation (11) is used primarily as an alternative method to estimate the adjustment speed γ, which is equivalent to μ and $(1 - \lambda)$.

Ideally, it would be preferable to obtain reliable firm-level estimates of δ, ϕ, and ω. Unfortunately, the inventory literature is filled with evidence that such estimates are not readily forthcoming. West and Wilcox (1994) and Fuhrer, Moore, and Schuh (1995) demonstrate that traditional generalized method of moments (GMM) estimation of such parameters from equations like (5) tends to produce biased and insignificant estimates. Gregory, Pagan, and Smith (1993) show that it is possible to identify the structural parameters (except, perhaps, β). However, identification requires a parametric specification of sales, which must be strictly exogenous, and definitive knowledge of the order of integration of sales (specifically, I(0) versus I(1)). Fuhrer *et al* argue that maximum likelihood estimation of inventory decision rules is a promising approach to estimating the structural parameters, but the requisite computer equipment to conduct such estimation was not available at the Census Bureau at the time of this study.

Instead, the econometric estimation focuses on using the three models to estimate the inventory speed of adjustment (i.e., $(1 - \lambda)$, μ, and γ) and addressing the well-known puzzle in the inventory

[7] In many econometric applications, the inventory stock adjustment model includes the real interest rate (cost of capital) as a determinant of I_t^*. Consequently, the real rate is included in the econometric estimation reported later in this paper.

literature that adjustment speed estimates are implausibly small. Typical adjustment speed estimates using monthly aggregate data are around 0.1, implying that the representative firm can take more than a year to adjust its actual inventory stock to the target level. Feldstein and Auerbach (1976) argued that since the entire manufacturing inventory stock amounts to less than one month of aggregate production, firms should be able to adjust their stocks much faster than is implied by estimated adjustment speeds. Maccini and Rossana (1984), Blinder (1986b), Nguyen and Andrews (1988), and Eichenbaum (1989) have attempted to resolve the puzzle with limited or no success.[8]

2 The M3 Longitudinal Research Database

2.1 Composition and Structure

The M3 Longitudinal Research Database (M3LRD) is a new panel data base containing monthly, seasonally unadjusted data on the domestic shipments, inventories, and orders of U.S. manufacturing firms (manufacturing operations only). In 1985, the Census Bureau began building the M3LRD by linking the firm-level data from the Manufacturers' Shipments, Inventories, and Orders (M3) survey. The main economic variables in the M3LRD are the dollar values of sales, stage-of-fabrication inventory stocks (finished goods, work-in-process, and materials and supplies on a current-cost basis) and new and unfilled orders. This study primarily uses two variables, the value of sales (VS), and finished goods inventories (FG), plus the production identity ($Q = VS + \Delta FG$). See the data appendix for a complete list of variables. Data are available through August 1993, and updated data become available periodically.

Two types of firm characteristics are available in the M3LRD: industry and geography. Industrial composition is determined by an industry category variable, which is based on a special Census classification system that includes about 80 categories of combined four-digit SIC manufacturing industries.[9] Census assigns a firm to the industry category in which it makes the plurality of its sales. The M3LRD includes firms from all 20 two-digit SIC manufacturing industries. Firm name and address are also included in the M3LRD, but the address is only for firms' headquarters and thus does not provide any geographic information about its production sites (plants).

The primary reporting unit in the M3LRD is the company (enterprise), which may include

[8]Eichenbaum (1989) gets the highest adjustment speed estimates, using unobservable autoregressive cost shocks. However, section 3 notes that there are some doubts about the robustness of his particular econometric technique.

[9]See Appendix B of Bureau of the Census (1987) for details about the industry categories.

one or more plants (establishments). [10] In contrast, the primary reporting unit in the LRD is the plant. Thus, although the M3LRD and LRD are drawn from the same manufacturing universe, the M3LRD is based on more aggregated reporting units and—due to different sampling schemes—includes a different mix of plants. However, Census obtains disaggregated reports for most large, diversified companies to improve data precision. Typically, the disaggregated report corresponds to a division within the company operating in relatively more homogenous industrial areas than the company as a whole. But the companies themselves determine the exact corporate entity reporting data and the entity can vary across companies. For the period 1985-93, the M3LRD contains about 8,200 disaggregated reporting units (henceforward referred to as divisions) belonging to about 4,300 companies. To maximize the degree of disaggregation, this study exploits both company-level and division-level data; for ease of exposition, both are referred to as the "firm" and the distinction is made only when necessary for clarity.[11]

2.2 Sample Selection

The M3 survey is the flagship of the Census Bureau's Current Industrial Reports program and a supplement to the Bureau's Census of Manufacturers (CM) and Annual Survey of Manufacturers (ASM) surveys.[12] The M3 survey includes "most manufacturing companies with $500 million or more in annual shipments [sales]. Selected smaller companies are included to strengthen the sample coverage in individual industry categories" (Bureau of the Census (1992), p. VII). Thus, unlike the CM and ASM, the M3 survey is not a probability sample and therefore cannot produce unbiased universe estimates.[13] Furthermore, firms may not appear continuously in the M3 survey because it is voluntary and the arbitrary sampling strategy for smaller firms changes over time. Consequently, the M3LRD does not permit accurate identification of firm entry and exit or corporate organizational dynamics (e.g., mergers, acquisitions, etc.).

Nevertheless, the M3LRD offers an opportunity to analyze the inventory behavior of hundreds

[10] A company, or enterprise, is defined as follows: "For definitions of plants (establishments) and companies (enterprises), see Bureau of the Census (1979, p. 12).

[11] This usage assumes that divisions are independent cost minimizers, which may be incorrect. Aggregating division-level inventory decision rules may not—perhaps likely will not—produce the company-level decision rule if there are important economic interactions among divisions within the company. These interactions are interesting and potentially important, but postponed for future study.

[12] For more complete details on the M3 survey and data, see Bureau of the Census (1992).

[13] Instead, Census uses the link-relative (LR) growth-rate method to obtain total manufacturing data that correspond to CM and ASM estimates. See the data appendix for details.

of manufacturing companies, and thousands of divisions, over nearly a decade-long period. From the pool of all M3 firms, this study concentrates on an M3LRD aggregation panel of 734 companies and 2,332 divisions.[14] Because the goal is to estimate firm-specific time-series econometric models, the aim in selecting the aggregation panel was to obtain a subset of M3 firms containing the highest quality data and covering the longest possible sample.

To be selected for the M3LRD aggregation panel, a company or division was required to satisfy certain criteria. The criteria included: (1) a sufficiently long sample period for all data; (2) a high percentage of reported, rather than imputed or missing, data observations; (3) a low percentage of measurement error detected from visual inspection, violation of data identities, and other miscellaneous editing procedures; and (4) limitations on outliers. In addition to the sample selection procedures, some firms' data required relatively minor editing for obvious and correctable measurement error. See the data appendix for further details.

Obviously, the arbitrary sample selection procedure may reduce the representativeness of the M3LRD, though it could cause it to be more representative. It turns out that the only significant difference between the total M3LRD and the aggregation panel is that the average firm is larger in the aggregation panel (measured by sales). The reason is that larger firms tend to report data more often and more accurately, in large part because the Census Bureau makes a more concerted effort to obtain accurate data from these firms.

2.3 Data Adjustments

Several adjustments were made to the data to make them as consistent as possible with the published aggregate M3 data. The firm-level sales and inventory data are deflated with industry-level deflators. Sales are adjusted for monthly variation in trading days. No adjustment was necessary for LIFO composition of inventories, but some inventory data were adjusted for a shift in the Census inventory reporting procedure in 1987. Sales and inventories were detrended and deseasonalized in accordance with the tradition of the literature.[15] See the data appendix for further

[14]To be clear, the 2,332 divisions do not all belong to one of the 734 companies nor, conversely, are all divisions of each of the 734 companies included in the 2,332 divisions. Each company and division enters the aggregation panel by satisfying the criteria described next.

[15]The main qualitative results in the paper are the same if the data are first-differenced data rather than detrended. The error-correction model provides evidence from an alternative way of dealing with trends. However, the issue of integration at the firm level has not been thoroughly examined yet in the literature, either theoretically or empirically, and thus the appropriate treatment is ambiguous. For the firms in the M3 aggregation panel, many data series exhibit

details about data adjustments.

2.4 General Characteristics and Statistics

Table 1 provides information about the industrial composition, size, average inventory holdings, and corporate organization of the firms in the M3 aggregation panel. The panel is almost evenly split between nondurable goods and durable goods firms in terms of number of companies and manufacturing shares, though disaggregation into divisions is more prevalent among durable goods firms (3.2 divisions per company versus 2.1).[16] Finished goods (FG) inventories are more important for nondurable goods firms, where they grew at a nearly 4 percent annual rate, than for durable goods firms, where growth was flat.

More relevant for inventory theory is the industrial disaggregation by production technology. About one-quarter of the firms in the panel operate in production-to-stock (PS) industries, i.e., they do not backlog orders. The traditional LQ inventory model is best suited to these kinds of firms. A key advantage of the M3LRD is that it permits disaggregation of the remaining production-to-order (PO) firms into those that report unfilled orders (PO-Y) and those that do not (PO-N). The latter firms, which account for almost three-fifths of the PO firms, may behave more like PS firms despite operating in PO industries.[17] Like nondurable goods firms, PS firms tend to rely relatively more on FG in terms of growth rates, shares, and inventory-to-sales ratios. PS firms also saw sales grow about 50 percent faster than PO firms. Particularly notable are the PO-N firms which, although accounting for only 6 percent of all sales, had the highest sales growth rate of any industrial category.

Despite the sampling bias toward larger firms, the panel nevertheless includes a significant number of smaller firms. About 60 percent of all firms have fewer than \$250 million annual average sales, a common cutoff for the definition of small in the literature on financial market effects on real activity. About 15 percent are just plain small, with average sales of less than \$25

positive trends, many negative, and many no trend at all. In light of this ambiguity, I elected to follow the traditional detrending procedure applied to the aggregate data.

[16] Although it is possible to examine more disaggregated industries, such as two-digit industries, there are not enough firms in all industries to produce consistent, reliable results. Schuh (1992) contains the results by two-digit industry for the interested reader.

[17] PO-N firms are identified because they did not *report* unfilled orders. Although they are assumed to produce to stock, it is possible that they hold unfilled orders but do not report them.

Table 1
M3LRD Firm Classifications and Summary Statistics

	Number of: [a]		Time Series Mean (1986-92)				
	Comp-	Div-	TM share(%) [b]		Growth Rate(%)		Level
Class	anies	isions	VS	FG	VS	FG	FG/VS
Industry [c]							
Total manufacturing (TM)	734	2332	100	100	.20	.17	.75
Nondurable goods [d] (N)	335	711	53	54	.18	.31	.76
Durable goods [e] (D)	399	1254	47	46	.24	0	.73
Production-to-stock [f] (PS)	206	473	42	42	.26	.34	.75
Production-to-order [g] (PO)	528	1492	58	58	.17	0	.74
Without backorders (PO-N)	306	349	6	6	.39	.10	.74
With backorders (PO-Y)	222	1143	53	52	.14	.03	.74
Annual average VS size, 1986-87 ($1987)							
Less than $25 million	111	498	0	0	.06	.11	.79
$25-100 million	191	689	2	2	.11	.13	.81
$100-250 million	135	438	4	4	.25	.18	.86
$250-1,000 million	159	395	13	15	.20	.21	.83
More than $1 billion	139	158	81	79	.20	.16	.73
Monthly average inventory-to-shipments ratio (FG/VS)							
0.00-0.25	101	306	12	2	.21	−.44	.12
0.25-0.50	111	326	21	10	0	.04	.36
0.50-0.75	122	269	22	19	.33	.15	.64
0.75-1.25	176	451	29	38	.26	.21	.97
More than 1.25	168	652	14	32	.24	.23	1.66
Number of establishments in the company							
One	50	65	0	0	.29	.11	.57
More than one	684	2267	100	100	.20	.17	.75
Reporting unit type							
Company	367	367	12	12	.26	.21	.75
Division	367	1965	88	88	.20	.16	.75

[a] Company counts for subclasses will not sum to the total if reporting units within companies cross over classes.

[b] Component shares may not sum to 100 due to rounding errors.

[c] The industrial aggregation identities are: (1) TM = N + D = PS + PO; (2) PO = PO-Y + PO-N.

[d] SIC industries 20-23, 26-31.

[e] SIC industries 24-25, 32-39.

[f] SIC industries 20-21, 23, 28-30.

[g] SIC industries 22, 24-27, 31-39.

million. Still, the vast bulk of sales and inventories are in the largest firms.[18] Unless meaningful differences in inventory behavior due to size occur at some size much smaller than $25 million, the M3LRD should be capable of identifying the impact of size. Middle-sized firms are the fastest growing, in terms of both sales and inventories, and have the highest inventory-to-sales ratios. Interestingly, the smallest firms have the lowest sales growth rate. For the remainder of the paper, firms are assigned to small (average annual sales less than $250 million) and large size classes.[19]

Firm-level average inventory-to-sales ratios exhibit considerable heterogeneity, running from zero to more than 1.25. A majority of of firms hold less than one-month's supply of inventories, as measured by the average monthly inventory-to-sales ratio. Interestingly, 56 companies (7.6 percent) and 270 divisions (11.6 percent) report holding no FG inventories at all. Firms in the two lowest ratio classes tend to have relatively low sales and inventory growth rates. Inventory-to-sales ratios tend to be inversely related to firm size (i.e., sales share per firm falls monotonically with ratio class).

The vast majority of firms in the data base own multiple establishments (plants)—well over 90 percent of companies and divisions. In terms of reporting units, half of the companies report company-level data and half report data by division. Companies reporting division-level data account for 84 percent of all divisions in the data base, and the average company reports data for 5.4 divisions. Single-establishment firms and company reporting unit firms account for very little of total sales and inventories. Because the traditional LQ inventory models do not capture these aspects of corporate organization, no further analysis is conducted along these dimensions.

2.5 Comparison of M3LRD and BEA Aggregate data

In light of sample selection concerns about the M3 survey in general, and the M3LRD aggregation panel in particular, it is important to ascertain how closely the aggregate M3LRD aggregation panel data correspond to the published aggregate BEA data.[20] Figure 2 demonstrates that the time series

[18]Of course, the panel's size distribution is far from representative of the size distribution of firms in the universe. The point here is that the bias toward larger firms does not preclude a significant number of sufficiently small firms from entering the panel.

[19]For a more detailed investigation of the impact of size on inventory behavior, particularly with respect to financial market conditions, see Maccini and Schuh (1995).

[20]Following Reagan and Sheehan (1985), the constant dollar ($1987) BEA data were reseasonalized using the ratio of seasonally adjusted to unadjusted nominal data then deseasonalized in a manner analogous to the M3LRD data. The BEA data are based on the nominal M3 data published by Census, and thus the BEA data incorporate (potentially

Figure 2
Total Manufacturing Aggregate Data (Seasonally Adjusted)

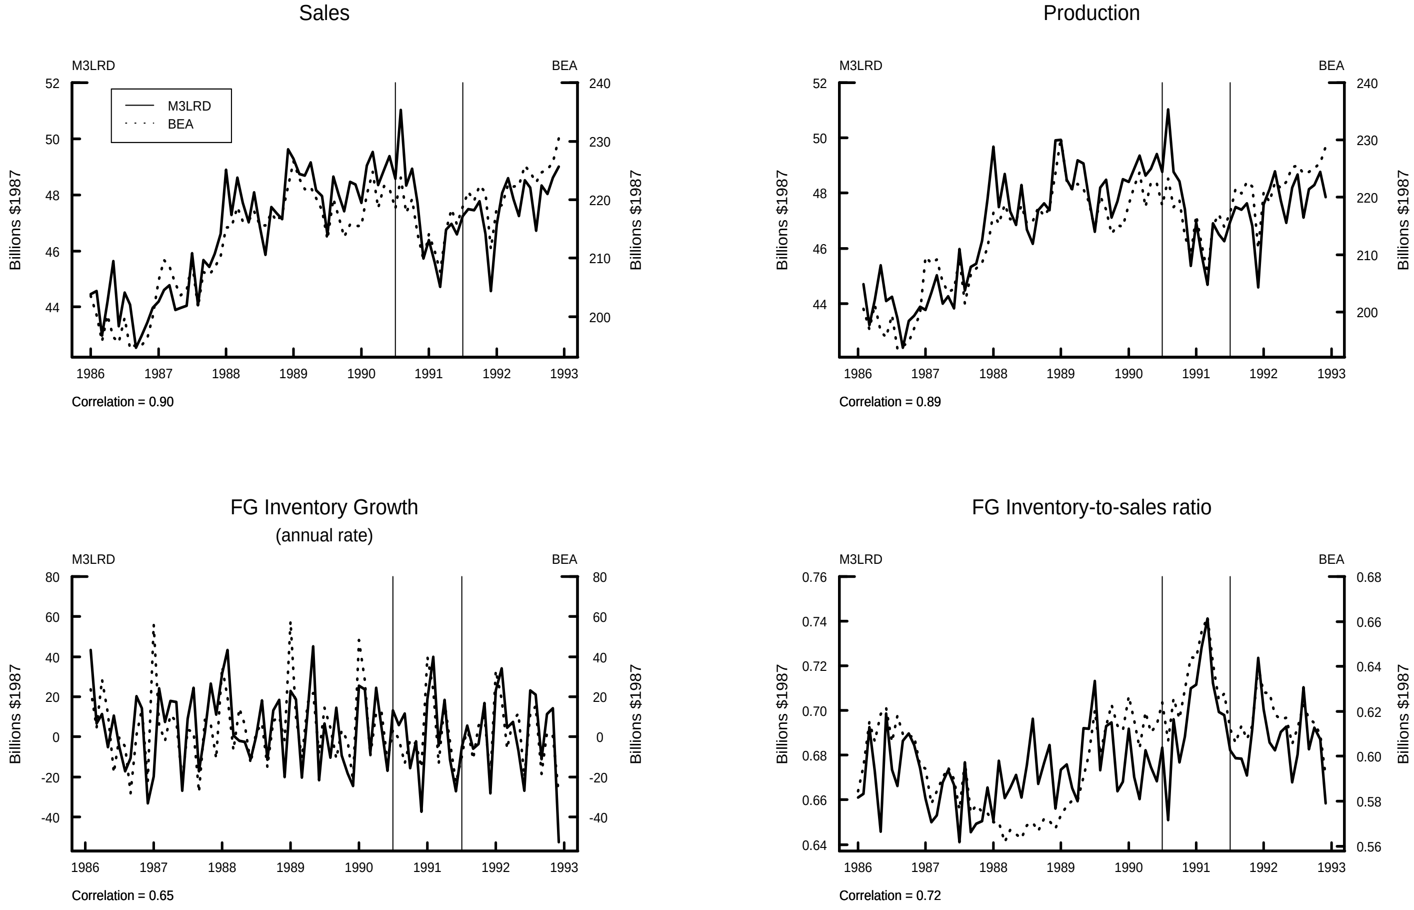

properties of the M3LRD aggregation panel data (henceforth, simply "M3LRD") and BEA data are reassuringly similar. Although the M3LRD data only account for slightly less than one-fourth of the BEA data, seasonally adjusted data from the two aggregates are highly correlated (.90 for sales and .89 for production).[21] The correlation is even stronger for seasonally unadjusted data (not shown). Correspondence between inventory series is less robust, but correlations of growth rates (.68) and inventory-to-sales ratios (.72) still significantly positive. Note that the average M3LRD inventory-to-sales ratio is slightly higher than the BEA ratio (.68 versus .60). In sum, the M3LRD aggregate data appear to be a reasonably representative subsample of the BEA data.

2.6 M3LRD Versus Other Disaggregated Inventory Data

To close this first glance at the M3LRD, it is instructive to compare the M3LRD with other disaggregated data bases containing inventory data. The comparison illustrates the advantages and disadvantages of M3LRD and, in summarizing results from other data bases, provides a useful benchmark for evaluating the results in this paper. Table 2 lists and describes the primary disaggregated inventory data bases cited in the literature. Two general impressions emerge from the table. First, no single data base exhibits all of the most desirable characteristics. Second, the M3LRD is as close as any data base—and perhaps is the closest—to exhibiting the most advantageous set of characteristics.[22]

Advantages of the M3LRD are disaggregation (of firm and inventory type) and frequency. Only the LRD provides more disaggregated data, but its data are only available annually, which is too low frequency given the extremely short-run nature of inventory decisions. Only three data bases (M3LRD, LRD, and Compustat) offer complete stage-of-fabrication disaggregation, and of these

adverse) effects of link-relative estimation of aggregate data, which permits fluctuations in sample composition, and data imputed by Census for missing observations.

[21] Aggregate VS for *all* firms in the M3LRD equals roughly half of total manufacturing VS reported by BEA. Thus, the sample selection restrictions reduce coverage by more than 50 percent. Nevertheless, the aggregate M3LRD data are broadly consistent with the BEA data. The percentage of total manufacturing VS coverage in the complete M3LRD varies from 20 to 90 percent across two-digit SIC industries. Not surprisingly, the correlations between M3LRD and BEA data aggregates at more disaggregated levels, such as two-digit industries, can be much lower.

[22] The M3LRD is a successor to the quarterly Department of Commerce Manufactures' Inventory and Sales Expectations Survey used by Hirsch and Lovell (1969) to estimate reduced-form inventory equations at the industry level and for an 83-firm subset. Some of these data were qualitative rather than quantitative. The study concluded that there is little evidence for production smoothing, but showed that smaller firms tend to have higher adjustment speeds and more flexible production plans.

Table 2
Comparison of Disaggregated Inventory Data Bases

Data Base[a]	Unit of Analysis	Coverage	Frequency[b]	Sample	Stock Type[c]	Data Type[d]
LRD	Plant	Probability sample of about 50,000-70,000 manufacturing firms	A	1972-90	SOF	$
M3LRD	Division	Arbitrary sample of about 1,700 manufacturing firms	M	1985-93	SOF	$
Ward's	Division	Domestic and transplanted foreign automakers	M	1938-95	F	P
Compustat	Company	All publically traded firms	A	1958-94[e]	SOF	$
Compustat	Company	All publically traded firms	Q	1958-94	T	$
QFR	Company	Probability sample of about 30,000 manufacturing, mining, and trade firms.	Q	1977-91	T	$
IFO	Company	Arbitrary sample of 4,000 manufacturing firms	M	1975-86	F	C
Cement	District	Arbitrary sample of 19 Portland cement industry production districts with continuous data	M	1950-60	F	P
SCB	Industry	Miscellaneous subsets of 4-digit SIC manufacturing industries	M	varies	F	P
Fair	Industry	Seven 3- and 4-digit SIC manufacturing industries	M	varies	F	P

[a]LRD=Longitudinal Research Database (Census Bureau); Ward's=Ward's Automotive Yearbook; Compustat=Standard & Poor's Compustat Services, Inc. (Securities and Exchange Commission); QFR=Quarterly Financial Report (Census Bureau); IFO=IFO Institute für Wirtschaftsforschung (Munich, Germany); Cement=Moheb Ghali (1987) data base; SCB=Survey of Current Business (Bureau of Economic Analysis); Fair=Ray Fair (1989) data base.

[b]A=annual; Q=quarterly; M=monthly.

[c]F=finished goods; SOF=all stage-of-fabrication (materials and supplies, work-in-process, and finished goods); T=total (sum of all SOF stocks).

[d]$=dollar value (current-cost); P=physical quantities; C=categorical (qualitative) responses.

[e]According to Hunt (1981), the SOF data are only available beginning in 1969 though the total inventory data go back to 1958.

only M3LRD offers it at high (monthly) frequency. Wards, the only other division-level data base, is highly specialized (auto industry) and only contains finished goods inventories.[23]

Potential disadvantages of the M3LRD are scope, sample, and data type. Unlike the LRD and QFR, the M3LRD is not based on a probability sample and, unlike the QFR and Compustat, the M3LRD covers only manufacturing. Still, the M3LRD coverage within manufacturing is broad and it includes some very small nonpublic firms that Compustat does not. Although approaching a decade, the M3LRD's sample period is on the short end. Nevertheless, new data continue to become available periodically. Finally, if critics such as Foss, Fromm, and Rottenberg (1980) and Miron and Zeldes (1989) are correct, the M3LRD's dollar-value data are subject to serious measurement error relative to physical units data.[24] But evidence reported later in this paper indicates that the dollar-value data may not be so bad after all, particularly for the purpose of quantifying the effects of aggregation within a data type. Furthermore, the M3LRD data is less subjective than the categorical-response data, which measure the opinions of firms' managers about, for example, whether inventories are too low, too high, or just right.

Two general findings in the literature pertain to the alternative disaggregated inventory data bases. First, studies using physical units data tend to find more evidence that firms smooth production. Ghali (1987), Harris (1988), Fair (1989), and Krane and Braun (1991) report considerably lower ratios of production variance to sales variance than do studies using the Commerce data—i.e., more evidence of production smoothing. This result could, however, be misleading if the narrow industries under study are more likely than others to smooth production or if aggregation across firms tends to bias the variance ratios for these industries downward.

A second general result is that disaggregated inventory data tend to produce higher adjustment speed estimates. Using BEA data, Blinder (1986b) observed that estimates for nondurable and durable goods manufacturing aggregates are generally lower than for two-digit industries within the aggregate groupings. Using categorical response data, Seitz (1993) found that the aggregate adjustment speed estimate is biased down from two-thirds to one-fourth after aggregating across

[23]To be more accurate, the Ward's data are actually available at the more disaggregated level of vehicle make (e.g., Taurus, Accord, etc.). However, product-line disaggregation in a multi-product environment does not conform well to the concept of a firm, where a centralized economic agent is assumed to be making decisions about all choice variables. Blanchard (1983) used division-level Ward's data.

[24]To investigate this issue more carefully, I compared the dollar-value and physical-units data for the motor vehicle assembly industry. M3LRD production (sales plus finished goods inventory investment) is very highly correlated with data on final production of assembled units from the Motor Vehicle Manufacturers' Association, both at the firm- and industry-level. Confidentiality laws prohibit more specific assessments and evidence on this correlation.

more than 500 German manufacturing firms. Using Compustat data, Hunt (1981) found scattered evidence in a few industries that aggregation biases adjustment speed estimates downward, but Haltiwanger and Robinson (1987) obtained extremely low estimates from pooled time-series cross-section estimation. And Harris (1988), using industry-level physical units data, did not find any systematic bias. However, none of these studies are capable of producing the breadth of quantitative evidence available from the M3LRD.

3 Firm-level Evidence and Results

This section provides evidence on firm-level production variance ratios and regression estimates of linear-quadratic inventory models for the M3LRD aggregation panel. All results include firms that reported holding no finished goods inventories. Henceforth, notation for inventories, sales, and production will be as follows: $I = FG$, $S = VS$, and $Q = VS + \Delta FG = S + \Delta I$.

3.1 Variance Ratios

Aggregate BEA manufacturing data, including total manufacturing and two-digit SIC industry-level data, suggest that firms tend to bunch production. That is, for BEA data the variance of production tends to exceed the variance of sales, expressed as $Var(Q)/Var(S) > 1$ (see Blinder and Maccini (1991), and references therein, for evidence). West's (1986) variance bounds test with BEA data also strongly rejects the hypothesis of production smoothing. On the other hand, physical units data in more disaggregated industries tend to exhibit production smoothing, with variance ratios often less than 0.9 (see Ghali (1987), Krane and Braun (1991), and Fair (1989)). There is little or no evidence, however, on firm-level variance ratios measured with data of any kind. A key contribution of this study, then, is to provide some evidence on the extent of production smoothing at the firm-level.

Figure 3 plots firm-level variance ratio distributions for all companies and divisions in the M3LRD aggregation panel.[25] The tail cells in each distribution include all ratios greater in absolute

[25] Results are presented only for $Var(Q)/Var(S)$ and not for the ratio with the value added definition of production, $Var(QSUM)/Var(S)$, where $QSUM = Q + \Delta WP$ and WP is the work-in-process inventory stock. Qualitatively, the distribution of variance ratios with $QSUM$ is the same as with Q, although the mean and variance are higher with $QSUM$. Work-in-process inventories are less precisely measured, in general, and measurement error tends to increase the variance ratios. To my knowledge, there are no theoretical models that predict whether the variance ratios should

Figure 3
Firm-Level Production Variance Ratios
(M = median, SD = standard deviation)

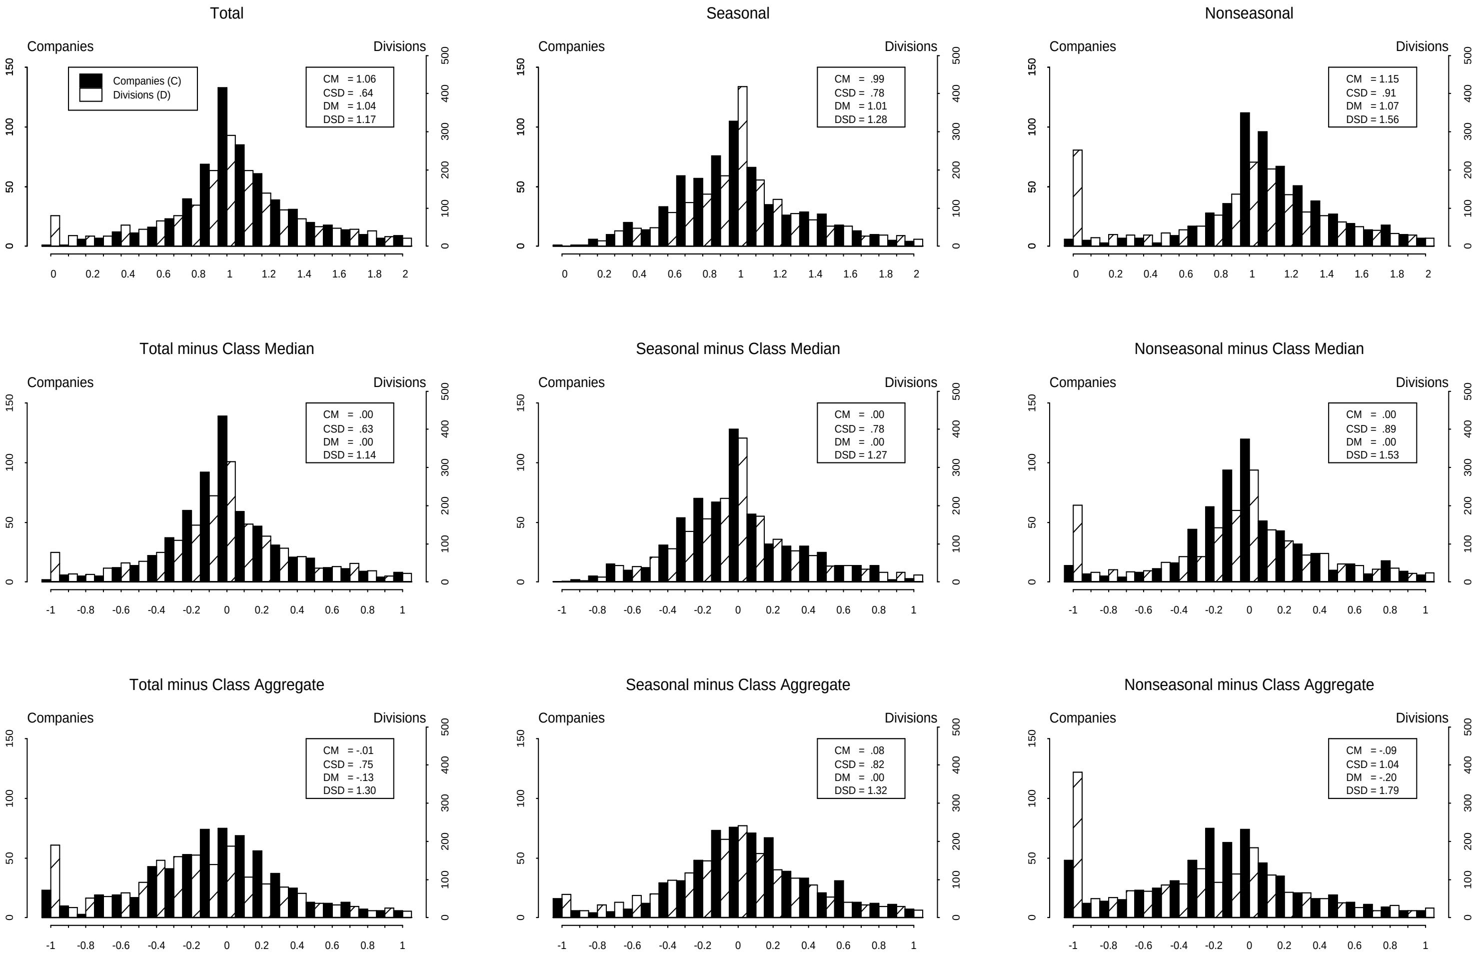

value than the endpoint values. The figure contains three columns: the total ratio (first column), and a decomposition into seasonal (second column) and nonseasonal (third column) components.[26]

The figure reveals a broad mix of production smoothing and bunching behavior at the firm level. About two-fifths of all firms smooth production (ratio < 1) in total (upper left panel), though the majority do not and the central tendency is for mild production bunching. The distribution is about the same in production-to-stock industries as in production-to-order industries, as well as in more disaggregated industries. The diversity of variance ratios also supports the basic LQ inventory model of section 1, which allows for a variety of smoothing and nonsmoothing behavior by incorporating a stockout avoidance motive (model parameter $\omega > 0$).[27]

Viewing across the top row, one can see that there is more evidence of production smoothing at seasonal than nonseasonal frequencies for companies, though less so for divisions. Whereas only about two-fifths of all companies smooth production in total, about half smooth production at seasonal frequencies. Conversely, only about one-fourth smooth production at nonseasonal frequencies. For divisions, the fraction of smoothing ratios is about 45 percent for both total and seasonal ratios, and about 40 percent for nonseasonal ratios. With regard to companies, this finding supports the hypotheses and conclusions of Miron and Zeldes (1988) and Krane and Braun (1991) using more aggregate data. In particular, to the extent that seasonal fluctuations are deterministic while cyclical fluctations are stochastic, rational firms should be able to smooth through the former better than the latter. Interestingly, though, the dispersion of both the seasonal and nonseasonal distributions is higher than for the total. This finding suggests that most firms tend to exhibit substantially different seasonal (smoothing) and nonseasonal (bunching) variance ratios.

The substantial heterogeneity of firm-level variance ratios raises an important question: what accounts for dispersion in variance ratios? To answer this question, the data were divided into 40 classes based on two observable firm characteristics — industry and size (20 two-digit SIC industries and two sizes, large and small). Each firm-level ratio then was deviated around the industry-size median ratio (shown in the second row of Figure 3) and the industry-size ratio cal-

be higher or lower for $QSUM$ than for Q.

[26]The decomposition follows the method in Krane and Braun (1991). The seasonal component is obtained by running a regression with seasonal and time dummies on the production and sales data and using the fitted seasonal part. The residual from this regression is the nonseasonal component. The sample period is 1985:1 through 1993:8, depending on availability of firm-level data.

[27]Fuhrer *et al* (1995) report that the feasible parameter space for a benchmark LQ inventory model using aggregate BEA nondurable goods data produces variance ratios between about 0.5 and 1.5. It seems plausible that firm-level data easily could produce a wider range of variance ratios, as is observed.

culated with aggregated data (third row). The standard deviations of the distributions show that observable characteristics account for little of the heterogeneity in variance ratios. The median-deviation ratios have marginally smaller variance at best, and the aggregate-deviation ratios are all larger. On a size-weighted basis (not shown), it is clear that both tails of the distributions are populated primarily by smaller firms.

Lastly, it is interesting to examine the M3LRD firm-level variance ratios in the M3 industries that correspond to the four-digit industries examined by Fair (1989) and Krane and Braun (1991). At issue is whether there is evidence that firms in these industries tend to smooth production more than other manufacturing firms. The M3LRD aggregation panel contains 52 companies and 98 divisions in these industries.[28] The median total variance ratio in these industries is 1.1 for both companies and divisions; the median seasonal variance ratio is .95 for companies and 1.1 for divisions. Thus there does not appear to be a tendency for firms in these industries to smooth production in general or relative to other manufacturing firms. To the extent that M3LRD firms in these industries are representative, this result lends some support to the contention that the dollar-value M3 data are "bad" relative to the physical-units data due to measurement error.

3.2 Econometric Results

This section reports regression results from firm-level estimation of the three econometric inventory models described in section 1: the stock adjustment model, equation (9), the error correction model, equation (11), and the Euler equation model, equation (6). For each model, the econometric specification is designed to follow as closely as possible the standard practice in the inventory literature.

3.2.1 Specifications

All three econometric models share the following specifications. Regressions are run for each firm individually; these are not panel data regressions so there are no parametric restrictions between firms.[29] The sample period for each firm is the maximum number of observations available over

[28]The industries are: cigarettes and cigars (category 21A with four four-digit SIC industries), tires (category 30A with one four-digit industry); cement (category 32C with 22 four-digit industries); and copper, lead, and zinc refining (category 33C with 17 four-digit industries).

[29]In principle, this regression strategy could benefit from applying the seemingly unrelated regression (SUR) technique to the individual firm-level regressions if there exists exploitable covariance among firm residuals. Given the

the period 1985:01 to 1993:08, minus relevant leads and lags. Data are in logs, deseasonalized, and, except for the error-correction model, detrended.

The econometric stock adjustment model for firm k is

$$I_{kt} = \mu_k \nu_{0k} + (1 - \mu_k) I_{k,t-1} - \pi_k S_{kt} + (\mu_k \nu_{1k} + \pi_k) E_{kt} S_{kt} + \mu_k \nu_{2k} E_{kt} r_{it} + \epsilon_{kt} \qquad (12)$$

where ϵ_{kt} is the usual random normal OLS regression error. Equation (12) includes a real interest rate, r_{it}, defined as the difference between the nominal three-month commercial paper rate and expected industry (subscript i) output price inflation. The real rate term enters via a more general inventory target, $I_t^* = \nu 0 + \nu_1 E_t S_t + \nu_2 E_t r_t$, which is common in the inventory literature. The sign of ν_2 should be negative. Expected sales and the expected real rate are replaced in the regressions with fitted values from independently estimated AR(3) models, as is usually done.[30]

The econometric error correction model for firm k is

$$\Delta I_{kt} = \alpha_k - \gamma_k (I_{k,t-1} - \varphi_{Sk} S_{k,t-1} - \varphi_{rk} r_{k,t-1}) + \theta_{Ik} \Delta I_{k,t-1} + \theta_{Sk} \Delta S_{k,t-1} + \theta_{ri} \Delta r_{i,t-1} + \eta_t \quad (13)$$

where the industry real interest rate has been added to the cointegrating relationship. The equation also includes seasonal dummies and the inventory valuation dummy. Theoretically, equation (13) is valid only for firms for which I, S, and r are I(1) and the linear combination of these variables is I(0). Standard augmented Dickey-Fuller (ADF) unit root tests were performed on I, S, r, and their linear combination to test these conditions. Only firms that pass all the criteria at the 10 percent level were included in the error correction model estimation (288 companies and 632 divisions qualified).[31]

---

large number of total firms (especially divisions), SUR is not feasible for the entire sample. However, I explored the effects of SUR estimation on 25 companies in the stone, clay, and glass industry (SIC 32) and found little impact on either the point estimates or standard errors. The firm-level differences between the single-equation and SUR standard errors for the adjustment speed parameter had a median of zero and a 90-10 decile range of $-.03$ to .03. In general, there is very little correlation between residuals among firms, even within industries.

[30]In principle, the standard errors for equation (12) should be corrected for the two-step estimation as in Murphy and Topel (1985), for example. However, in earlier work (Schuh (1992)), I found that the adjustment made little difference. Breusch (1978) and Godfrey (1978) Lagrange multiplier tests indicated little or no serial correlation in the residuals so no corrections were made. Thus, the adjustment speed estimates are not subject to the upward bias due to multiple equilibria sometimes encountered in these models (see, for example, Blinder (1986b) and Hall and Rossana (1991)).

[31]Theoretically, the real rate probably should be I(0), but empirically one cannot reject the hypothesis that it is I(1). Because the conduct of monetary policy may induce nonstationarity in the short-term real rate, I included it. However,

The econometric Euler equation for firm k is

$$I_{k,t+1} - [\lambda_k + (\lambda_k \beta)^{-1}] I_{kt} + \beta^{-1} I_{k,t-1} + S_{k,t+1} - \kappa_k \beta^{-1} S_{kt} + c_k = \xi_{k,t+1} \tag{14}$$

where c_k is a constant and $\xi_{k,t+1}$ includes the expectational error, which is MA(1) because $S_{kt} \notin \Omega_{kt}$, plus any white noise cost shocks in the model. Discount factor β is preset to 0.995. The instrument set is $Z_{kt} = [1, I_{k,t-1}, I_{k,t-2}, S_{k,t-1}, S_{k,t-2}]$, which is consistent with the standard assumptions of exogenous and unknown sales. Parameters λ_k and κ_k are estimated with Hansen's (1982) GMM estimator using the Newey and West (1987) weight matrix with lags equal to one. The J statistic from the test of overidentifying restrictions is distributed χ^2 with three degrees of freedom. This GMM implementation is virtually identical the analogous model in Eichenbaum (1989).

For many firms, the GMM criterion function is very flat over the stable range for λ_k (i.e., 0-1). In these cases, the actual estimate of λ_k is often greater than one (unstable) and the standard errors are extremely large. To overcome this problem, I follow Eichenbaum (1989) and Durlauf and Maccini (1992) and multiply equation (14) by $(1 - \lambda_k)^{-1}$. Although in theory this transformation is legitimate because it doesn't affect the moment condition, in practice it affects small sample estimates by changing the location and increasing the curvature of the criterion function. In particular, it tends to significantly reduce estimates of λ_k (thereby increasing the adjustment speed) and their standard errors.[32]

3.3 Regression Estimates

Table 3 reports the results of the firm-level regression estimation for each of the three models at the company level (upper half) and division level (lower half). The first three columns report 10th, 50th (median), and 90th percentiles of the distributions of parameters and regression diagnostic statistics. The next two columns indicate the percentage of parameter estimates that are

the parameter estimates of the error correction model are not particularly sensitive to the inclusion of real rate except that the adjustment speed (γ) is slightly larger with the real rate (median firm-level adjustment speeds of .30 with the real rate and .25 without).

[32] I did some random experimentation with this econometric technique by transforming the equation by $(1 - \lambda_k)^{-\rho}$ for $\rho = \{.5, 1, 2, ..., 8\}$, since any positive real value of ρ is legitimate. The estimates of λ_k (the adjustment speed) were monotonically decreasing (increasing) in ρ, though at a decreasing rate. Thus it appears that this transformation permits generation of an arbitrarily large estimate of the adjustment speed. This issue requires a more systematic Monte Carlo investigation.

positive or negative and significantly different from zero at the 10 percent level. The final two columns measure the degree to which observable firm-level heterogeneity accounts for variance in the parameter estimates. "Median" indicates the percentage of variance attributable to variation in median firm-level parameter estimates across 40 industry-size classes. "Other" indicates the percentage attributable to idiosyncratic variation within industry-size classes.[33]

The overall impression conveyed by Table 3 is that the firm-level data generally do not support the linear-quadratic inventory model, regardless of the econometric methodology employed. In most respects, these firm-level estimates are consistent with prior results from regressions with aggregate data. On one hand, the adjustment speed estimate is highly significant and the correct sign and magnitude. But on the other hand, most parameter estimates are insignificantly different from zero, the wrong sign, or both. Specifically:

- **Stock adjustment model**—The median buffer stock parameter (π) is zero instead of one. Roughly half of the estimates of π, ν_1, ν_2 are the wrong sign, and estimates of all three are about equally likely to be significantly negative as significantly positive.

- **Error correction model**—The lagged difference coefficients (θ's) are about evenly distributed around zero. The estimates are significantly different from zero less than half the time, but about twice as likely to be significantly positive rather than negative. In this unrestricted form, the coefficients don't have much of a structural interpretation other than a significant θ_I indicating the presence of adjustment costs.

- **Euler equation model**—The κ parameter is less than 1.0 for most firms, as predicted, but not for all firms. In addition, less than two-fifths of the κ estimates are significantly different from zero.

Despite these generally unsupportive results, the table reveals three notable differences between the firm-level results and previous aggregate results:

1. **Adjustment speeds**—Firm-level adjustment speed estimates are considerably larger than most estimates reported from regressions with monthly aggregate data. Figure 4 plots the

[33] As in the standard variance formula, there is a covariance-type term but it is generally small and can be inferred from the table. The actual formulas are as follows. Let $\overline{\beta}$ denote the median parameter estimate for all firms and N the total number of firms. Subscripts are as follows: i denotes industries, j denotes sizes, and k denotes firms. Then the "median" column is $(N-1)^{-1}\sum_{i=1}^{20}\sum_{j=1}^{2}(\overline{\beta}_{ij}-\overline{\beta})^2$, and the "other" column is $(N-1)^{-1}\sum_{k=1}^{N}(\beta_{ijk}-\overline{\beta}_{ij})^2$.

Table 3
Firm-Level Inventory Model Parameter Estimates

Model [a]		Percentiles			% Significant		% Variance Accounted for by:	
		P10	P50	P90	+	−	Median	Other
Companies:								
SA	μ	.13	.32	.72	99	0	1	100
	π	−.30	−.01	.51	27	29	0	97
	ν_1	−2.20	.20	3.08	46	36	0	98
	ν_2	−.03	.00	.03	33	33	0	100
	R^2	0	.02	.09				
	SC(1)	.23	.78	1.00				
ECM	γ	.12	.30	.63	94	0	4	88
	θ_I	−.25	−.03	.21	14	26	3	91
	θ_S	−.28	−.02	.24	11	23	0	98
	θ_r	−.02	0	.02	13	15	1	92
	R^2	.10	.23	.44				
	SC(1)	.28	.81	1.00				
Euler	$1-\lambda$	.31	.48	.69	89	1	0	94
	κ	−1.23	.47	1.54	35	4	0	99
	χ^2	.17	.72	.97				
Divisions:								
SA	μ	.13	.40	1.03	93	4	0	99
	π	−.51	−.02	.86	24	32	0	94
	ν_1	−3.28	.09	3.90	40	37	0	95
	ν_2	−.03	0	.03	28	32	0	100
	R^2	0	.02	.27				
	SC(1)	.26	.78	1.00				
ECM	γ	.11	.32	.79	93	0	2	95
	θ_I	−.25	−.03	.27	13	27	1	95
	θ_S	−.30	−.03	.24	9	26	0	97
	θ_r	−.02	0	.03	11	17	0	97
	R^2	.11	.26	.50				
	SC(1)	.26	.81	1.00				
Euler	$1-\lambda$	.33	.50	.70	86	0	0	96
	κ	−1.35	.50	1.70	33	4	0	98
	χ^2	.17	.72	.97				

[a] SA denotes stock adjustment, ECM denotes error correction model, and Euler denotes the Euler equation model. SC(1) and χ^2 statistics report the probability values for tests of first-order serial correlation and overidentifying restrictions, respectively. See text for more details.

Figure 4
Firm-Level Inventory Adjustment Speed Estimates

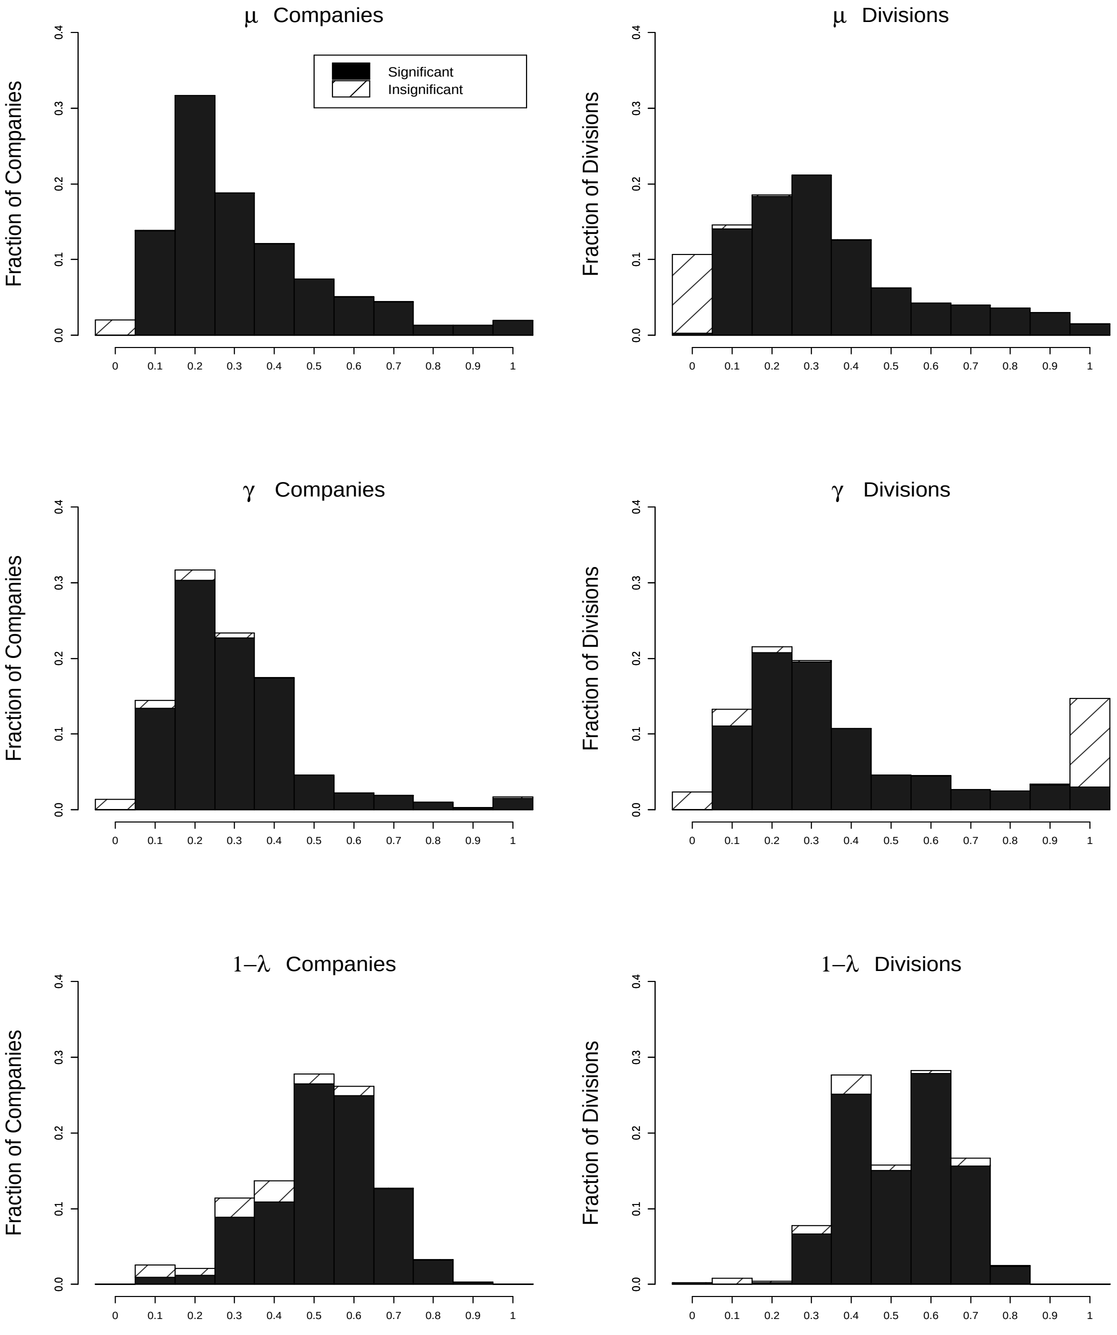

weighted distributions of adjustment speed estimates for further inspection. The median firm-level adjustment speeds for the stock adjustment and error correction models run from .30 to .40. Compare these with estimates in Blinder (1986b, Table 1) of .06 and .14 for durable and nondurable goods industries, respectively. The median firm-level adjustment speeds for the Euler equation model are .48 and .50; Eichenbaum (1989, Table 1) reports estimates of .30 and .39 for nondurable goods industries.[34] The adjustment speed estimates are slightly higher for divisions than companies in all three models, supporting Blinder's assessment from two-digit industries that "aggregation seems to bias the estimated speed of adjustment downward" (p. 359). Also, for μ and γ, disaggregation from companies to divisions leads the adjustment speed distributions to move from a skewed (e.g., χ^2) shape closer to a uniform shape.

2. **Serial correlation**—Firm-level residuals exhibit very little serial correlation, unlike residuals from regressions using aggregate data. The table reports p-values for: (1) tests of the null hypothesis of no serial correlation against the alternative of first-order serial correlation (denoted by SC(1)); and (2) tests of overidentifying restrictions (denoted by χ^2). There is virtually no evidence of serial correlation in the stock adjustment and error correction models. In contrast, the inventory literature has devoted considerable space to debating how the treatment of serial correlation in regressions with aggregate data can radically affect regression results. Similarly, there is virtually no evidence against the overidentifying restrictions in the Euler equation model. In contrast, most GMM estimates of Euler equations produce highly autocorrelated residuals and rejections of the restrictions (see, for example, Eichenbaum (1989) and Kashyap and Wilcox (1993)).

3. **Model fit**—Firm-level models do not account for much of the variation in the inventory data. The median R^2 in the stock adjustment model is close to zero; the median in the error correction model is considerably higher, but still only about one-fourth. Given that the

[34]Two robustness issues regarding the firm-level Euler equation estimates are worth noting. First, if there is no buffer stock motive (S_t is known and in the instrument set) then the adjustment speed distributions have the same central tendencies. However, the distributions are considerably less disperse and the overidentifying restrictions are rejected more frequently. Second, when the parametric transformation $(1-\lambda_k)^{-1}$ is not made to equation (14) the distribution of λ_k changes markedly — the median estimate is 0.63 and about one-fourth of all estimates are greater than or equal to 1.0 (i.e., the adjustment speeds are much smaller). These issues apply equally to the firm-level and aggregate data.

stock adjustment model is approximately a restricted version of the error correction model, this result weighs against the restrictions of the stock adjustment model. In contrast, regressions with aggregate data usually generate R^2 of one-half or more or more for both types of models. It seems reasonable to infer, however, that this result is directly related to the other two.

The table also portrays considerable heterogeneity in the firm-level results, but the heterogeneity is essentially totally unexplained. For example, the 90-10 decile band of adjustment speed estimates encompasses most of the stable region. Likewise, all other estimates exhibit a broad mix of negative and positive values, many (but not most) of which are significantly different from zero. What can explain this heterogeneity? The last two columns of the table indicate that substantially less than 10 percent of the dispersion, and usually none of it, can be attributed to observable characteristics (industry and size). Instead, virtually all the dispersion is idiosyncratic.

Combining these findings, one is left with the following depiction of microeconomic inventory behavior. Although important theoretically, standard accelerator (expected demand) and cost of capital (real interest rate) effects explain very little of the variation in firm-level inventory investment. Observable firm characteristics cannot explain the idiosyncratic component of firm-level inventory investment. In addition, the idiosyncratic component of firm-level inventory investment is not very persistent, which is consistent with the finding that firms do not experience inventory gaps for abnormally long periods of time.

Several factors could potentially explain the large and transitory idiosyncratic component of firm-level inventory investment. Measurement error is an obvious candidate, given the heterogeneity and lack of persistence. But this explanation is, of course, unsatisfying. Misspecification due to missing variables is an obvious economic explanation, though it is not immediately clear what firm-specific variables are missing. Some low-frequency factors, such as inventory-saving capital and production techniques, are possibilities, but these cannot account for the high-frequency volatility of the idiosyncratic component. Also, because the inventory target depends on expected variables, it is possible that simple univariate time-series projection methods do not adequately account for all information available to firms which may lead them to very different sales plans, for example. Finally, at the firm level there may be functional nonlinearities or discontinuities at work, such as an (S,s) rule. In any case, the results point to much room for improvement and further research.

4 Aggregate Evidence and Results

This section provides evidence on aggregate production variance ratios and regression estimates of linear-quadratic inventory models for a balanced-panel subset of the M3LRD aggregation panel (613 companies and 903 divisions). The balanced panel includes all firms from the M3LRD aggregation panel with continuous data over the period 1986:01 through 1992:12. A balanced panel is necessary for examining the link between firm-level and aggregate inventory behavior because neither the data nor the model framework adequately account for firm entry and exit.

In this section, three aggregate concepts arise. First is MW, which pertains to the weighted firm-level mean value where the weight accounts for variation in average firm size. Second is MA, which pertains to the aggregate data (simple summation across all firms) obtained from the balanced-panel subset of the M3LRD aggregation panel. Third is BEA, which pertains to the aggregate data published by the Bureau of Economic Analysis (the data used in most previous inventory studies).

4.1 Representative Agent Issues

Before turning to the aggregate results, it is necessary to address the representative agent assumption. In particular, the question arises: can aggregate inventory behavior be modeled adequately using a representative agent hypothesis? This hypothesis has come under increasing fire recently. For example, Geweke (1985) and Kirman (1992), argue strenuously that there is no guarantee that a representative agent exists even in theory, much less in practice. In the inventory literature—as in most of macroeconomics—relatively little attention has been given to the representative agent hypothesis. But the M3LRD provides an opportunity to assess existing assumptions and to compile evidence on the hypothesis.

The remainder of the paper investigates this issue by adopting Theil's (1954, 1971) approach to aggregation analysis. The Theil approach begins with the assumption that there exists a true microeconomic model for each of the N firms in the economy. Then, using the simple linear aggregator function, one can derive the implied "true" aggregate model. The representative agent hypothesis, then, is that the aggregate model accurately reflects the true aggregate of the underlying microeconomic relations.

It turns out that linearity of the microeconomic relations is critical to having a hope of evaluating the representative agent hypothesis, so the linear-quadratic inventory model qualifies. However, even with linear microeconomic relations, one of the following two conditions is necessary for the existence of a representative agent model: (1) all data (e.g., inventories and sales) are the same across firms; or (2) all micro parameters are the same across firms. West (1986, p. 383), for example, adopts the latter approach by assuming that the structural parameters of his inventory model are identical across firms. Clearly, however, previous sections have provided ample evidence against both homogeneity assumptions by documenting tremendous heterogeneity in the data.

Ramey (1989, p. 345) adopts another approach to the aggregation problem by assuming: (1) firm-specific components of observed and unobserved variables are independent across firms and time; (2) firms face identical prices, wages, and interest rates; and (3) cross-sectional output variance is constant. The extensive unexplained heterogeneity in the firm-level data and the lack of correlation among firm-level data series might be interpreted as evidence in favor of (1). However, Abbott (1987) and Davis and Haltiwanger (1991) provide strong evidence against (2) for manufacturing prices and wages, respectively.

In addition, Figure 5 shows evidence against Ramey's point (3). The two panels in the first column plot the time-series of the first four cross-section moments of detrended firm-level production. Clearly, the cross-section mean of production can vary more than 10 percent over the business cycle (1990-91), and even more for the other moments. Note, especially, that the cross-section variance of production rises substantially and persistently during the 1990-91 recession much the same way job reallocation (a measure of dispersion) does, as reported by Davis and Haltiwanger (1992). The two panels in the second column demonstrate the same principle for the moments of the equilibrium error from the cointegrating equation.

4.2 Variance Ratios

This section provides evidence on the relationship between firm-level and aggregate production variance ratios. At issue is whether the production variance properties of the aggregate data are representative of the firm-level evidence discussed previously. Put another way: is it possible that the firm-level data exhibit more production smoothing than is observed in the aggregate data?

The first four columns of Table 4 report the evidence on four aggregate measures of total, seasonal, and nonseasonal production variance ratios (i.e., $Var(Q)/Var(S)$). Comparing the first two columns with the next two, it is clear that the firm-level data do not reveal more production smoothing behavior than the aggregate data. The MA ratios (column 3) generally are about the

Figure 5
Cross-Section Moments in the M3LRD Aggregation Panel

Mean and Standard Deviation of Firm-Level Production

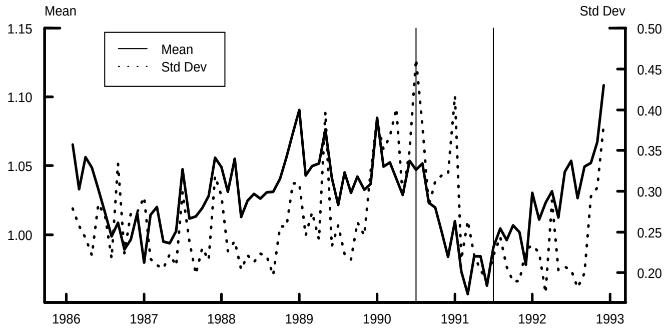

Mean and Standard Deviation of the Cointegrating Equilibrium Error

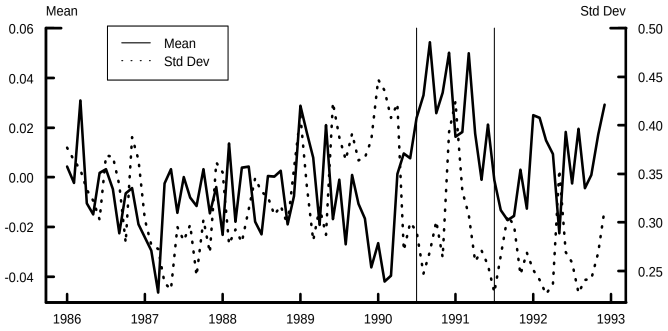

Skew and Kurtosis of Firm-Level Production

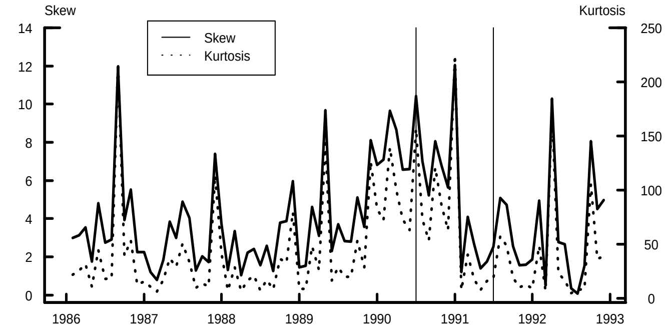

Skew and Kurtosis of the Cointegrating Equilibrium Error

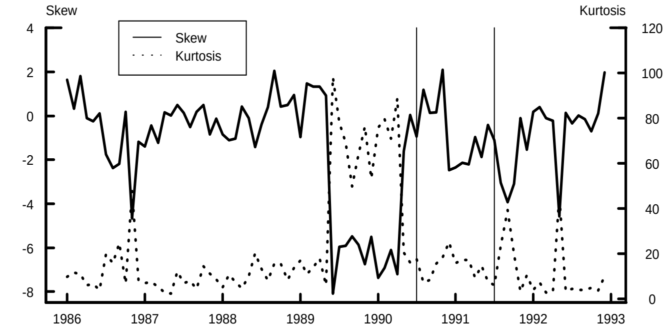

Table 4
Aggregate Variances of Production and Sales

	Variance Ratios				Variance Sums			
					Companies		Divisions	
Industry	Median	MW	MA	BEA	S	Q	S	Q
Total:								
TM	1.05	1.39	1.08	1.43	64.7	77.5	120.5	153.2
N	1.04	1.54	1.18	0.76	24.5	33.3	32.8	37.7
D	1.07	1.26	1.02	2.50	40.2	44.3	53.2	76.7
PS	1.08	1.76	1.26	0.66	13.9	21.1	26.4	29.7
PO	1.04	1.25	1.04	2.16	50.9	56.4	59.6	84.7
PO-N	1.05	1.28	1.17		24.8	30.9	8.7	17.0
PO-Y	1.03	1.22	0.91		26.0	25.5	50.8	67.7
S	1.04	1.39	1.02		44.4	49.6	83.0	100.1
L	1.08	1.37	1.22		20.4	27.9	37.5	53.0
Seasonal:								
TM	1.00	1.44	0.91	0.91	27.4	25.7	48.9	48.4
N	1.00	1.79	0.98	0.66	11.6	11.7	13.5	12.2
D	1.00	1.15	0.86	1.30	15.8	14.0	17.2	20.3
PS	1.00	2.30	1.33	0.64	5.1	6.9	9.7	8.6
PO	1.00	1.13	0.82	1.18	22.3	18.8	20.9	24.0
PO-N	1.00	1.14	0.94		9.2	8.6	4.2	5.5
PO-Y	0.99	1.11	0.73		13.1	10.1	16.7	18.5
S	1.00	1.50	0.83		21.3	18.5	35.1	32.2
L	0.99	1.33	1.16		6.0	7.1	13.8	16.2
Nonseasonal:								
TM	1.15	1.51	1.22	3.91	37.4	51.8	73.7	105.7
N	1.12	1.60	1.38	1.38	12.9	21.5	20.5	26.2
D	1.17	1.44	1.13	8.08	24.5	30.2	36.9	56.6
PS	1.18	1.63	1.22	0.99	8.8	14.1	17.9	21.9
PO	1.15	1.47	1.22	5.94	28.6	37.6	39.5	61.0
PO-N	1.17	1.51	1.31		15.7	22.3	4.7	11.4
PO-Y	1.10	1.43	1.10		12.9	15.4	34.9	49.5
S	1.12	1.53	1.20		23.0	31.0	48.9	68.0
L	1.21	1.49	1.25		14.4	20.8	24.8	37.7

NOTES: The table reports the variance ratios, $\mathrm{Var}(Q)/\mathrm{Var}(S)$, and variances of Q, estimated over the period 1986:01 to 1992:12. Column headings indicate variance ratios constructed with the following data: Median = firm-level median from balanced panel; MW — weighted sum of firm-level ratios; MA — M3LRD balanced panel aggregate data; B — BEA reseasonalized aggregate data. See the text for definitions of total, seasonal, and nonseasonal.

same magnitude, or only slightly larger, than the median firm-level ratios (column 1). But the MA ratios are notably smaller than the MW ratios (column2). Apparently, firms with higher variance ratios are disproportionately large. For all manufacturing, the BEA data exhibit considerably more production bunching, but the cross-sectional evidence is disparate (perhaps due to sample selection or small sample problems). The evidence on total ratios versus the seasonal and nonseasonal components is essentially the same as at the firm level. In sum, it appears that aggregation actually tends to *reduce* the degree of production bunching in the aggregate data.

Recently, Krane (1994) and Lai (1991) have hypothesized that production bunching may be observed in the aggregate data even though all (or most) firms smooth production. Their argument is clear from the aggregation formula for the difference between production variance and sales variance:

$$Var(Q) - Var(S) = \sum_{k=1}^{K} Var(q_k) - \sum_{k=1}^{K} Var(s_k) + 2\left[\sum_{i \neq j}\sum_{j \neq i} Cov(q_i, q_j) - Cov(s_i, s_j)\right] \quad (15)$$

where lowercase letters pertain to firm-level data and uppercase to aggregate data. Suppose each firm smooths production, i.e., that the sum of firms' production variance is less than the sum of firms' sales variance. Then aggregate production variance would exceed aggregate sales variance only if the sum of covariance among firms' production sufficiently exceeded the sum of covariance among firms' sales. Whether this covariance condition is true depends largely on the underlying parameters of the cost functions $C_Q(\cdot)$ and $C_I(\cdot)$. Krane shows that the condition is true when marginal stockout costs (a function of ϕ and ω) sufficiently dominate marginal holding costs (which implicitly enter the model through constants). Lai shows that the condition is true when shocks to firms' market shares are sufficiently more variable than common (aggregate) shocks.

The remaining four columns of Table 4 report evidence on this hypothesis for companies and divisions in the M3LRD balanced aggregation panel. The data speak strongly against the Krane and Lai hypotheses: for the vast majority of industries for all types of ratios, the sum of production variance exceeds the sum of sales variance. Thus, there is not a significant difference between the covariances among firm production and firm sales to generate a misleading aggregate variance ratio. In fact, the correlation among firm-level production and firm-level sales variables is generally close to zero.

4.3 Econometric Results and Aggregation Bias

This section provides evidence on the relationship between firm-level and aggregate parameter estimates from the three inventory econometric models. At issue is whether the aggregate parameter estimates are representative of the firm-level parameter estimates. That is, we want to determine whether or not there is evidence in favor of the representative agent hypothesis for the standard inventory models.

For linear models like the inventory models, Theil's (1954) analogy principle provides a framework for analyzing the effects of aggregation on the parameters of macroeconometric models. For a general linear model, the framework is as follows. Let uppercase variables denote aggregates (simple linear sums) and lowercase variables denote firm-level variables. The firm-level models are

$$y_{kt} = x_{kt}\beta_k + \epsilon_{kt} \quad k = 1, ..., K \tag{16}$$

where x_{kt} and β_k are J-dimensional vectors of predetermined variables and time-invariant firm-level parameters, respectively. Thus the aggregate model is

$$Y_t \equiv \sum_{k=1}^{K} y_{kt} = \sum_{k=1}^{K} \left(x_{kt}\beta_k + \epsilon_{kt}\right) \ . \tag{17}$$

Relation (17) implies that

$$Y_t = \sum_{k=1}^{K} \left[\sum_{j=1}^{J} \left(w_{jkt}\beta_{jk}\right) X_{jt} + \epsilon_{kt}\right] \tag{18}$$

where $w_{jkt} = (x_{jkt}/X_{jt})$, and X_{jt} and x_{jkt} are (T $\times$ 1) matrices from the jth columns of X_t and x_{kt}. Therefore each aggregate parameter, $\beta_{jt} = \sum_{k=1}^{K} w_{jkt}\beta_{jk}$, is the size-weighted sum of firm-level parameters. Although the firm-level parameters β_{jk} are time-invariant (by assumption), the "true" aggregate parameters β_{jt} are time-varying.

The "true" aggregate parameters become time-invariant only if one (or both) of two conditions is true: 1) all of the weights are time-invariant (i.e., $w_{jkt} = w_{jk}$), which implies that firm size is time-invariant; or 2) the firm-level parameters are the same for all firms (i.e., $\beta_{jk} = \beta_j \ \forall \ k$). If one of these restrictions holds, then the aggregate model will be "representative" of the underlying firm-level models. Otherwise, a fixed-parameter (FP) aggregate model will be misspecified relative to the "true" time-varying-parameter (TVP) aggregate model.

On one hand, the FP aggregate model may provide misleading information about firm-level behavior if the regression estimates of a fixed aggregate parameter (β) are not representative of the underlying firm-level parameters (β_k). A simple measure of the "true" fixed aggregate parameter is a weighted sum (average) of the firm-level parameters: $\overline{\beta} = \sum_{k=1}^{K} \overline{w}_k \beta_k$, where $\overline{w}_k$ is the time series mean of the true weight w_{kt}. The difference between these two fixed aggregate parameters, $\beta - \overline{\beta}$, is a measure of econometric aggregation bias.[35] Past research has shown that aggregation bias can be as large as an order of magnitude (ten times) or more—a finding confirmed in this paper for the inventory model.[36]

On the other hand, however, $\overline{\beta}$ is not necessarily the value of β that minimizes the sum of squared residuals in an aggregate OLS regression. As Grunfeld and Griliches (1960) point out, the FP aggregate model may very well be preferred to the firm-level models for two related reasons. First, the FP aggregate model is likely to obtain a much higher R^2 than are the firm-level models—a finding also confirmed in this paper. Second, the forecasting performance of the FP aggregate model may exceed that of the collective forecasts of the firm-level models.

In light of substantial econometric aggregation bias but poor fit of firm-level inventory models, the question arises as to whether the disaggregated data offer any improvement over FP aggregate models in explaining inventory behavior. To answer the question, this paper compares the FP aggregate model with the TVP model, which allows for heterogeneity in firm-level fixed parameters and time series fluctuation in the size distribution of firms. At issue is whether the TVP aggregate model can fit the aggregate data substantially better than the FP aggregate model.[37]

The remainder of this section shows that severe aggregation bias arises in standard aggregate inventory models. Because the inventory models do not perform well, at either the firm level (as shown earlier) or at the aggregate level (as shown in the literature), the focus is restricted to the one parameter that is estimated precisely—the adjustment speed. In particular, Table 5 reports adjustment speed estimates (μ, γ, and $(1-\lambda)$) for: (1) the weighted sum of firm-level parameters (MW column); (2) the FP aggregate model using aggregate M3LRD data (MA column); and (3)

[35] See Theil (1954, 1972) for a derivation of the aggregation bias and a demonstration that $E[\beta] \neq \overline{\beta}$.

[36] See Boot and de Wit (1960) and Sasaki (1978) for examples with capital investment stock adjustment models, Gupta (1971) and Lee, Pesaran, and Pierse (1991) for examples with labor demand equations, and Bils (1985) for an example with a wage equation.

[37] A TVP model fits the aggregate data better essentially by definition, of course, because a FP model is just a special case of a TVP model. In some sense, it might be more appropriate to compare the aggregation TVP model with other TVP models such as those obtaining time variation from model nonlinearities, variation in cross-section distributions, etc. But given that the firm-level models are actually FP models, this exercise provides relevant information on the extent to which accounting for firm-level size variation can account for unexplained variation in aggregate inventory investment.

Table 5
Aggregate Inventory Adjustment Speed Estimates

	μ				γ				$1-\lambda$			
Industry	DMW	CMW	MA	B	DMW	CMW	MA	B	DMW	CMW	MA	B
TM	.45*	.37*	.27*	.05	.37*	.32*	.18*	.02	.50	.48*	.56*	.35
	(.15)	(.08)	(.08)	(.05)	(.11)	(.06)	(.06)	(.02)	(.23)	(.22)	(.17)	(.14)
N	.43	.40	.38*	.22	.41*	.36*	.11*	.02	.50	.49*	.56*	.35
	(.16)	(.08)	(.10)	(.06)	(.13)	(.06)	(.07)	(.03)	(.23)	(.21)	(.11)	(.36)
D	.50*	.34*	.23*	.06	.35*	.29*	.26*	.05	.50	.47	.48*	.35
	(.14)	(.07)	(.07)	(.04)	(.10)	(.06)	(.07)	(.03)	(.23)	(.22)	(.21)	(.13)
PS	.40	.39	.40*	.19	.42*	.35*	.12*	.04	.50	.50	.55*	.34
	(.16)	(.08)	(.10)	(.06)	(.14)	(.07)	(.07)	(.03)	(.24)	(.23)	(.13)	(.38)
PO	.50*	.36*	.32*	.06	.35*	.30*	.27*	.04	.50	.47	.44*	.35
	(.14)	(.08)	(.08)	(.04)	(.10)	(.06)	(.07)	(.03)	(.23)	(.21)	(.23)	(.13)
PO-N	.50*	.36*	.33		.33*	.29*	.28		.52	.47	.49	
	(.15)	(.08)	(.08)		(.07)	(.06)	(.07)		(.22)	(.20)	(.19)	
PO-Y	.50*	.35*	.55		.37*	.32*	.15		.49	.47*	.42	
	(.14)	(.08)	(.19)		(.10)	(.05)	(.05)		(.23)	(.22)	(.11)	
S	.55*	.38*	.26		.41*	.34*	.01		.51	.48*	.38	
	(.25)	(.08)	(.09)		(.12)	(.06)	(.03)		(.25)	(.20)	(.27)	
L	.43*	.35*	.29		.36*	.29*	.20		.49	.47*	.60	
	(.12)	(.08)	(.08)		(.11)	(.06)	(.06)		(.22)	(.23)	(.18)	

NOTES: Regressions are estimated over the period 1986:01 to 1992:12, less appropriate lags. Asymptotic standard errors appear in parentheses. Columns indicate parameter estimates associated with the following data: DMW — weighted sum of division-level parameters; CMW — weighted sum of company-level parameters; MA — aggregate data from the M3LRD balanced aggregation panel; B — BEA reseasonalized aggregate data (not available for PO-N, PO-Y, S, and L industry categories). All data were detrended and deseasonalized using seasonal dummies. A * indicates that the parameter estimate is significantly different from the estimate in the column to the right at the 10 percent level. See the text for more details.

the FP aggregate model using aggregate Commerce Department data (BEA).[38] The BEA estimate cannot be compared strictly with the MW and MA estimates because the M3LRD aggregation panel represents only a subset of the BEA data. However, the BEA estimate gives some indication of the effects of higher degrees of aggregation and the effects of the sampling scheme used to produce the data (see the data appendix on this latter point).

Table 5 illustrates the most striking result of the paper: aggregation tends to bias adjustment speeds estimates downward. For the stock adjustment and error correction models, the adjustment speed estimates (μ and γ) in nearly all industry classes decline monotonically as the level of aggregation increases from divisions (DMW) to companies (CMW) to aggregate (MA and BEA).[39] Further, the bias is most often statistically significant and large. For total manufacturing, the division-level weighted mean adjustment speed is 67 percent and 105 percent larger than the aggregate (MA) estimate in the two models, respectively. Aggregation bias is evident in the Euler equation adjustment speed estimates $(1 - \lambda)$, but considerably smaller and less significant. However, given the sensitivity of the Euler equation estimation technique to parametric transformations, it is unclear how seriously to take the results.

Two other important characteristics stand out in Table 5. First, there is some variation in the extent of aggregation bias across types of firms. Aggregation bias tends to be smaller in production-to-stock (PS) and nondurable goods industries—for which the standard inventory models are more applicable—than in production-to-order (PO) and durable goods industries. Perhaps aggregation in these latter industries cancels out some misspecification biases such as failing to include unfilled orders in the model.

A second important characteristic is that the BEA adjustment speed estimates are markedly lower than any other estimates. For most industry categories, the decline in the BEA estimate relative to the MA estimate is greater than the decline due to explicit aggregation bias (MW-MA). Apparently, the additional aggregation and the sampling methodology together reduce adjustment speeds even more. Thus the BEA adjustment speeds, which are comparable to most estimates

[38]The econometric specifications for the aggregate models are identical to those described for the firm-level models in section 3.

[39]The MA estimates of γ are not strictly comparable with the DMW and CMW estimates. The reason is that the MW estimates include only the firms which exhibited cointegration between inventories and sales, while the MA estimate is for all companies in the M3LRD balanced aggregation panel. MA-type estimates from data aggregating only over the subset of firms with cointegrated inventories and sales are not uniformly smaller than the DMW and CMW estimates. I chose to report the estimates this way because I believe that the smaller sample adversely affects the aggregate more than it does the central tendency of the firm-level parameters.

reported in the inventory literature, are the *least* representative of the underlying firm-level adjustment speeds.[40]

What are the economic implications of this finding that aggregation biases adjustment speeds downward? The result seems to resolve the long-standing puzzle in the literature that adjustment speeds are implausibly slow (Feldstein and Auerbach (1976)). At the aggregate level, the BEA data imply that the representative firm reduces its inventory gap by a small fraction each month. But the firm-level M3LRD data reveal that the average firm actually eliminates close to half of its gap—a much more plausible rate of adjustment. In fact, with even more disaggregated data (e.g., at the plant level) we should expect to find even higher adjustment speeds—perhaps well above one-half. More generally, this finding provides evidence that persistence in aggregate data may not be due to adjustment costs, learning, "stickiness", incomplete markets, and other theoretical hypotheses advanced in many macroeconomic models.

The aggregation bias also suggests that there may be econometric gains to exploiting disaggregated data in aggregate models. Clearly, it is not theoretically or computationally feasible to build and maintain separate models for each agent in the economy. But it is possible to investigate whether the heterogeneity in firm-level parameters and time-variation in firm size can be exploited to improve the aggregate inventory model.

Figure 6 demonstrates how time variation in the "true" aggregate adjustment speed parameter can impact the aggregate inventory model. The upper panel plots the time series of "true" aggregate adjustment speeds for each of the three inventory models. For example, the stock adjustment model estimate (μ) of the "true" aggregate adjustment speed (based on divisions) has a time series range of about .015 (or a little more than 3 percent).

Although the time variation is small, the impact on the fit of the stock adjustment model is large. The lower panel shows the residuals from the FP stock adjustment model using MA data and the TVP model with μ varying only.[41] The figure shows that the TVP model fits the data better than the FP model: the R^2 rises from .74 in the FP model to .88 in the TVP model—an improvement

[40] In most other respects, the aggregate MA and BEA estimates conform well with the estimates typically reported in the literature. For example, the R^2 estimates for the stock adjustment and error correction models are much higher than for the firm-level regressions, typically well above .5. Estimates other than the adjustment speed parameters are often insignificant or the wrong sign or both. Surprisingly, though, there is little evidence of serial correlation in the aggregate regressions for this sample period.

[41] For the most part, allowing other parameters to vary doesn't change the results much largely because the other parameters are insignificant. However, allowing parameters that are occasionally the wrong sign to vary, especially the real interest rate parameter ν_2, can actually worsen the fit of the model.

Figure 6
Time-Varying Aggregate Parameter Model

'True' Time-Varying Aggregate Adjustment Speed Parameters

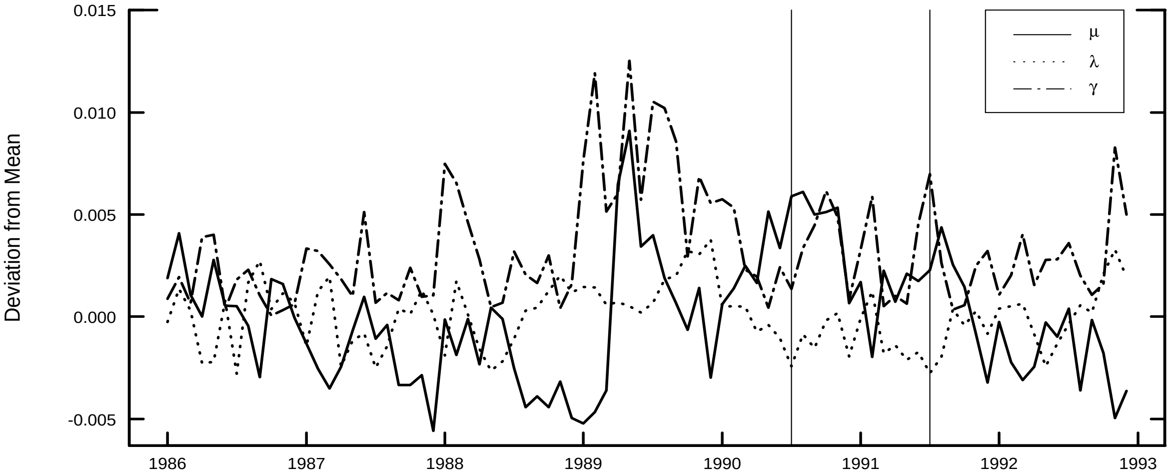

Inventory Stock Adjustment Model Residuals

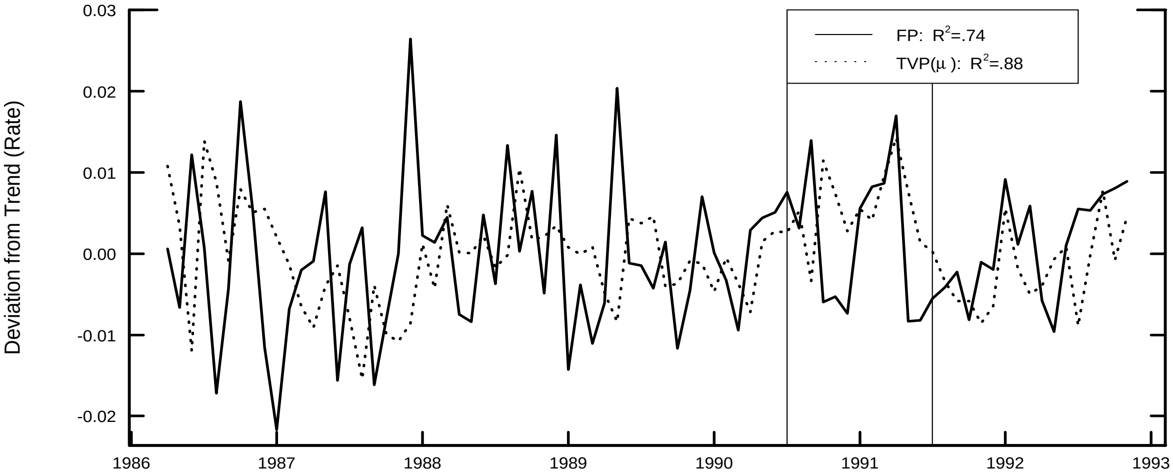

of 19 percent. Note that this improvement is about the same magnitude as the improvement in fit reported by Caballero, Engel, and Haltiwanger (1994) for an aggregate employment model that incorporates cross-section moments to account for firm-level nonlinearities.

5 Summary and Conclusions

This paper informs us of three new facts about inventory behavior. First, there is a broad mix of production smoothing and production bunching bunching firms. Second, firm-level inventory adjustment speeds are much higher than has been inferred previously from aggregate adjustment speed estimates. Finally, although inventory models generally perform no better with firm-level data than with aggregate data, accounting for variation in the size distribution of firms can dramatically improve the fit of aggregate inventory models.

What are the implications of these new facts? Following in a summary list:

1. **Adjustment speeds**—Given that firm-level adjustment speeds appear to be more plausible than previously thought, a new puzzle arises: how do we reconcile slow aggregate adjustment speeds with fast firm-level adjustment speeds? One possible answer is that the appropriate periodicity of interpretation is different for aggregate and firm-level adjustment speeds. Firm-level adjustment speeds probably should be interpreted at calendar frequency. With the average M3LRD inventory stock equal to about three-quarters of a month of sales, it should not take a rational, optimizing firm many months to eliminate an inventory gap. But aggregate adjustment speeds probably should be interpreted at business cycle frequency, where the typical gap lasts about as long as the recession. Thus, the aggregate adjustment speed reveals how fast the aggregate economy eliminates the aggregate inventory gap. Slow stock adjustment arises naturally from the heterogeneity of firm-level adjustment speeds (some of which are quite slow) if demand shocks hit all firms simultaneously. But slow stock adjustment can be exacerbated if there is variation in the timing of the demand shock across firms as well.

2. **Persistence**—The finding on adjustment speed bias may pertain to other macroeconomic variables as well. To explain the extreme persistence observed in most real macroeconomic data, economists have resorted to explanations such as adjustment costs, price "stickiness" and wage contracts, imperfect or incomplete markets, limited information, less-than-rational

expectation formation, strategic behavior, and fixed costs. But the results in this paper support: (1) Trivedi's (1985) contention that aggregate persistence in the form of long distributed lags is simply a consequence of aggregation across firms with short distributed lags; and (2) Stoker's (1986) finding that "distributional effects can be mistakenly interpreted as evidence of dynamic effects". Thus, a problem of interpretation arises. Consider an aggregate inventory model that includes a fixed-parameter cost of changing output to account for persistence in the data. Typically, this term is interpreted as some sort of physical adjustment costs. But the results in this paper suggest that the interpretation of this specification is incorrect on two grounds: (1) it is unrepresentative of firm-level behavior; and (2) it misses time variation in persistence that occurs due to fluctuations in firm size.

3. **Convexity and aggregation**—The results in this paper show that accounting for aggregation over convex firm-level behavior can produce dramatic improvements in the fit of aggregate models equal to the improvements obtained by aggregating over nonconvex firm-level behavior. Of course, this finding does not rule out the existence or importance of nonconvex behavior. But it does indicate that disaggregating only to the firm level, where behavior is relatively convex, rather than to the plant level, where behavior is relatively nonconvex, can yield significant gains for macroeconomic analysis.

4. **Room for improvement**—The poor overall performance of the standard inventory models at the firm-level suggests that there is still much room for improvement in developing applied inventory models. While it is important to understand why aggregate adjustment speeds are so slow, this study seems to rule out aggregation bias as the reason for the overall poor fit of applied inventory models.

5. **Production smoothing**—The heterogeneity in production variance ratios suggest that future inventory modeling efforts should allow flexibly for both production smoothing and bunching behavior at the firm level. The M3LRD data demonstrate that there is a substantial fraction of firms with smoothing behavior. While there may be some measurement problems in this data—perhaps even enough to cause the average variance ratio to exceed one instead of being less than one—there is too much heterogeneity to force an inventory model to necessarily produce one behavior or the other.

6. **Disaggregated data**—This study demonstrates the value of working with disaggregated data by showing two things. First, it shows that the data reveal how an important feature of firm-

level behavior (the inventory adjustment speed) had been misinferred from past research using aggregate data. Second, it shows that the data provide an opportunity to account for heterogeneity and time variation in size that can markedly improve a simple aggregate model relatively easily. These achievements motivate expanding the availability of disaggregated data.

A Data Appendix

This data appendix provides more detailed information about the Census Bureau's M3LRD and how it was used in the study reported in this paper.

A.1 Contents

Table A.1 lists the main variables in the M3LRD and their definitions. There are three main categories of variables: information, economic, and reporting status. The information variables pertain primarily to the name, address, and Census identification numbers of each firm's central administrative office, but these names and addresses are not necessarily associated with physical production sites. Two other information variables tell about the firm's accounting period and, relatedly, it's monthly trading-day variation in the economic data. Finally, there is a sample weight for the limited number of firms that were involved in historical testing of random sampling techniques (most firms have a weight of 1.0).

In addition to the two main variables used in this study (VS and FG), the M3LRD includes economic variables containing information on other stage-of-fabrication inventories—materials and supplies and work-in-process—and on new and unfilled orders. The total inventory and new orders identities provide an opportunity to check for data measurement error, as does the production identity (to some extent). The percentage of nonzero values of the LIFO variables is less than 3 percent.

The last group of variables are reporting statuses. These useful variables provide information about M3LRD data reliability by indicating whether the data are reported by the firm or not. If so, they show how the data were reported; if not, they show how Census "filled in" the data. There is one reporting status variable associated with each economic variable and with each version of the data (described next).

A.2 Data Versions

The M3LRD contains three data versions: REPORTED, FINAL, and BENCH. Each version corresponds to a different stage of the data collection, editing, and publication process. REPORTED is data reported by firms in the M3 survey and included in the preliminary published M3 report. FINAL is REPORTED data edited by Census analysts in preparation for final publication, plus late

Table A.1
Variables in the M3 Longitudinal Research Data Base

	Information Variables
ADDRESS	Central office address of reporting unit or company
DATE	Month and year
ID	Census reporting unit identification number indicating company affiliation
IND	M3 industry category
NAME	Name of reporting unit's company
OID	Old ID from preceding parent company (if applicable)
RP	Accounting reporting period
STATE	State associated with ADDRESS
TDF	Trading day factor
WT	Sample weight
	Economic Variables
FG	Finished goods inventories
LIFO	Portion of TI valued on a LIFO basis
LRES	LIFO adjustment or reserve
MS	Materials and supplies inventories
NOD	New orders, derived (VS+ ΔUO)
NOR	New orders, reported
Q	Production (VS + ΔFG)
TI	Total inventories (MS+WP+FG)
UO	Unfilled orders
VS	Value of shipments
WP	Work-in-process inventories
	Reporting Statuses for Economic Variables [a]
A	Correction of reported data by Census analyst
C,K	Correction of reported data by company
D	Derived by computer edit
E	Analyst estimate with no reported data available
G	Estimated by company
H	Estimated from company's historical data
I,M	Imputed from industry's historical data
P	Reported data received by phone
R	Reported data received by survey

[a] One reporting status variable is associated with each of the 10 economic variables. These reporting statuses are designated RFI, RLIFO, etc.

reported data. Analyst edits include imputation, deletion of non-manufacturing data, reclassification of stage-of-fabrication (SOF) inventory data, etc. BENCH is FINAL data benchmarked by Census analysts to the Annual Survey of Manufactures (ASM) data, plus other late reported data.

Data versions may differ significantly, but only for certain rare observations. The differences, which occur only in the FINAL and BENCH versions, signify time-series structural breaks in the data caused by corporate reorganization (e.g. a merger), reporting procedure changes (e.g. redefinition of company divisions), and data redefinitions (e.g. changes in a firm's assessment of what is work-in-process inventory versus finished goods inventory). In the event of a structural break, Census creates two different time series: one that reflects the proper month-to-month growth rate of the data based on the old structure and one on the new structure. Consequently, there are two versions of each of the FINAL and BENCH data series: FINAL1, FINAL2, BENCH1, and BENCH2 data series (1 and 2 denote the series on a pre- and post-structural break basis, respectively).

Because considerable skill, time, and knowledge about specific industries and firms are required to edit the data versions, I chose to work with the BENCH data. It would be interesting to re-estimate the inventory models using the REPORTED data to discover what affect, if any, the Census editing process has on the parameter estimates. However, it turns out that there is very little difference among the versions anyway. Less than .5 percent of the REPORTED, FINAL, and BENCH data differ within any particular month, and less than .02 percent of the BENCH1 and BENCH2 data differ. Where the BENCH versions differed, the BENCH1 version was selected because BENCH1 represents the actual data value.

A.3 Sample selection and restriction

The entire M3LRD contains more than 222,000 monthly observations for more than 8,200 divisions of more than 4,300 companies. A subset of the M3LRD sample was selected for this study that would provide reliable data with which to test the LQ inventory model. The following criteria were used to select and restrict the sample:

- **Sample size**—Only firms with sufficiently long continuous data samples were included. Sales (VS) data were required to be nonzero for 90 out of 104 possible months. This criterion ensures sufficient observations to maintain acceptable degrees of freedom for econometric estimation.

- **Reported data**—Only firms with high percentages of reported, rather than imputed, data were included. For each firm and each economic variable I calculated the time-series percentage of monthly observations of "reported" data as those observations with reporting statuses A, C, G, K, P, and R. All other observations were considered not reported, or imputed. Sales and total inventory data (VS and TI) were required to be reported for 90 percent or more of all time-series observations. Stage-of-fabrication (SOF) inventory data (FG, WP, and MS) were required to be reported for 70 percent or more of all time series observations because the SOF data have lower average rates of reported data. This criterion reduces the extent of measurement error in the data attributable to imputation.

- **Data identities**—Only firms with limited violation of data identities were included. The two main identities are total inventories (TI) and derived new orders (NOD). Firms with data violating these identities for a majority of observations were eliminated.[42] (Some of these identities were used to correct data as well; see below.) This criterion also limits the extent of measurement error.

- **Size distribution**—Only firms that fit smoothly into the size distribution were included. Because the M3 sample is unrepresentative, and because sample selection procedures further reduce the sample, the possibility arises of very large outlier firms dominating the results. Consequently, the largest firm in the sample was prohibited from having annual average sales more than twice as large as the next largest firm. This criterion prevents the skewness of the size distribution from distorting the results given the nonrepresentativeness of the sample.[43]

[42]Using identities to judge the data can be problematic because the data are book-value (dollar) data. For example, although the total inventory identity holds in terms of physical quantities, the equation may not hold in dollar terms if accounting or tax provisions allow companies to delay or speed up the recording of inventory investment. Nevertheless, most identities are satisfied by the reported data — in fact, Census forces the new orders identity — so the identities seem to be a reasonable quality check.

[43]To be more concrete, the composition of the M3LRD aggregation panel is such that it is possible to have one firm with more sales than the sum of the sales of all other firms in the panel. Consequently, the behavior of one firm can largely determine the outcome of the aggregate M3LRD analysis. Because there is only one firm, there is little chance that it is exactly representative of all very large firms. Thus, although it would be preferable to include very large firms, in small samples selection bias can be severe.

A.4 Data Editing and Imputation

Even after passing these sample selection criteria, some firms still required some editing and imputation of a small fraction of their data observations. Obvious data errors, such as random rounding errors, were cleaned via visual inspection of time series plots of each data item. Data imputed by the Census Bureau using industry average growth rates were replaced with data imputed using firm-specific time-series methods described below. Other editing and imputation is due to the occurrence of zero, missing, and identity-violating data.

Sales (VS) data are considered by Census to be the most reliable data variable and were addressed first. The need for editing and imputation was minimal because the vast majority of sales data is nonzero, nonmissing, and reported. In the few cases where they are not, the data were imputed using the SAS software PROC EXPAND function, which fits a cubic spline to the firm-level data.

Total inventory (TI) data are considered the second most reliable data and were addressed next. For the relatively few observations where the total inventory data required imputation, fitted values of total inventory were obtained from the regression

$$TI_{kt} = \beta_k VS_{kt} + \sum_{i=1}^{12} \gamma_{ki} SD_{ikt} + \delta_{k1} T + \delta_{k2} T^2 + \epsilon_{kt} \tag{19}$$

where k indexes firms, SD indicates monthly seasonal dummies, T is a linear time trend, and only reported VS observations were included. Use of this regression was predicated on the theory that there is a stable, long-run relationship between sales and total inventories.

Stage-of-fabrication (SOF) inventory data are considered notably less reliable than the sales and total inventory data and were addressed last. Some firms' SOF data also exhibit occasional compositional changes apparently related to changes in definition. For example, firms may change their assessment of what is a material inventory versus a work-in-process inventory. Where obvious definitional changes could be identified, the stocks were redefined to be consistent over time (the stock with data for the majority of observations was selected). Then, the remaining zero, missing, imputed, and identity-violating data were obtained from the fitted values of the regression

$$I_{kt} = \beta_{Ik} TI_{kt} + \gamma VS_{kt} + \sum_{i=1}^{11} \gamma_{ki} SD_{ikt} + \delta_{k1} T + \delta_{k2} T^2 + \epsilon_{Ikt} \tag{20}$$

where $I = \{MS, WP, FG\}$ and only reported SOF data observations where the TI identity is satisfied were included. The β_I parameters represent the average proportion of I in TI, conditional on sales, and sum to one.

A.5 Trading Day Adjustments

The sales (VS) data were trading-day adjusted using the Census adjustment factor. Trading-day adjustment is necessary because accounting data collected on a monthly calendar basis are not comparable between months throughout the year for two reasons. First, the number of calendar and work days vary across months. Second, accounting periods vary across firms. Since time periods in the LQ inventory model implicitly are assumed to be identical, the data months must be standardized. Three types of accounting periods are reported in the M3LRD data: 4-4-5 periods (two 4-week months and one 5-week month per quarter); 13-4 periods (13 4-week months per year), and calendar month periods.

Although many different trading day adjustment techniques have been used in empirical work, the Census Bureau has used some form of the following technique since the 1960s. Weekdays receive a production weight of one and weekend days a weight of one-half; thus, there are six production days per week. A standard number of monthly trading days is obtained from the following formula:

$$Trading\ days/month = \frac{6 \times (365.25/7)}{12} \approx 26.089\ . \tag{21}$$

Consequently, the trading day factor for 4-4-5 reporters in month t is

$$TDF_t(4-4-5) \approx \frac{26.089}{6X} \tag{22}$$

for $X = \{4, 5\}$. A similar factor is used for the 13-4 reporters. The factor for calendar month (M) reporters is

$$TDF_t(M) \approx \frac{26.089}{WD_t + 0.5WE_t} \tag{23}$$

where WD_t is the number of weekdays and WE_t the number of weekend days in month t.

The trading day factor is applied multiplicatively to flow variables, such as VS, but not to stock variables, such as the various inventory types. Although the LQ inventory model is specified in terms of inventory stocks, it implicitly determines the change in inventories (inventory investment), which is also a flow. Consequently, it is logical to adjust the inventory *stocks* so that inventory *investment* (a flow) is trading-day adjusted in a manner consistent with the trading-day adjustment of sales. Trading-day adjustment factors for stock variables, denoted TDF^*, can be derived from the trading-day factor for flow variables, TDF, using the equation

$$TDF^*_{It} = TDF_{It} + (TDF^*_{I,t-1} - TDF_{It})(I_{t-1}/I_t) \tag{24}$$

where $I = \{MS, WP, FG, TI\}$ denotes the inventory type and $t > 0$. By assumption, $TDF^* = 1.0$ at $t = 0$. Experimentation with stock trading-day factors revealed they had no substantive or systematic effect on the construction of aggregate data, on the basic time-series patterns of the data, or on firm-level or aggregate regression results, and thus they were not used.

A.6 Other Data Adjustments

A number of other adjustments were made to the raw M3LRD data:

- **Deflation:** Firm-level price data were not available, so all nominal data were deflated at the firm level with highly disaggregated fixed-weight industry-level price deflators (1987=100) from the Bureau of Economic Analysis (BEA), as described in Hinrichs and Eckman (1981). Sales (VS) deflators are at the M3 industry category level and all inventory deflators are at the two-digit SIC industry level. Although the deflators are not firm-specific, the breadth of product diversification in many of the large M3 companies increases the suitability of industry-level deflators.

- **LIFO inventory adjustment:** No adjustment was made for differences in inventory valuation methodology, specifically differences between LIFO and non-LIFO methods. Since January of 1987, Census has collected all nominal inventory data on a current-cost, or pre-LIFO, basis, which eliminates the need for a LIFO adjustment. During 1985-86, the extent of LIFO inventory accounting reported in the M3LRD is less than 3 percent of all observations during that period. Further, the usefulness of LIFO data is limited because most firms using LIFO usually only make the appropriate calculations annually or quarterly rather than monthly.

- **Detrending and deseasonalizing:** Because the M3LRD contains time-series data, nonstationarity is an issue. Following inventory literature tradition, e.g. Blinder (1986a), the data were detrended and deseasonalized at the firm level with the following regression:

$$log(X_{kt}) = \sum_{i=1}^{12} \beta_{ki} SD_{ikt} + \gamma_{1k} T + \gamma_{2k} T^2 + \delta D87_t + \epsilon_{kt} \tag{25}$$

  where $X = \{VS, TI\}$ and $D87$ is a dummy variable for potential intercept shift due to valuation changes (one in months from January, 1987, forward and zero prior to that). The $D87$ adjustment is only made where δ is positive and significant (because the switch from

LIFO to current cost will increase the nominal value of inventories), and where the firm experience abnormally large growth rates in all inventory stocks (but not sales) in early 1987.

A.7 M3 Sampling Methodology and Sample Selection

The Census Bureau conducts three main economic surveys of manufacturing: (1) the quinquennial Census of Manufactures (CM), which covers the universe of manufacturing plants and companies; (2) the Annual Survey of Manufactures (ASM), which covers a probability sample of plants and companies; and (3) the monthly Manufacturers' Shipments, Inventories, and Orders (M3) survey, which covers a non-probability sample of about 1,700 companies. Thus, the CM and ASM surveys produce unbiased estimates of manufacturing data levels, but the M3 survey does not.

Census designs the M3 sample as follows. It attempts to include all companies with 1,000 or more employees plus about 60 percent of an arbitrarily selected sample of companies with 100 to 1,000 employees. The 1,000 employee cutoff roughly corresponds to \$500 million in annual sales and pertains to about 550 companies. The original M3 sampling frame, established in 1962, included a sample of companies with fewer than 100 employees. Census dropped the small companies in 1963 because the response rate was too low. Since then, resampling occurs occasionally to improve coverage in certain industries but Census makes no systematic effort to update the sample. See Bureau of the Census (1992) for more details about the survey.

Without probability-based sampling weights, the M3 data levels cannot provide unbiased estimates of the universe. However, there is some chance that the growth rate of the M3 data is an unbiased estimate of the universe growth rate. Thus, the Census Bureau constructs a link-relative (L) growth-rate estimator of the universe inventory stock, $I = \sum_{k=1}^{K} I_{kt}$, which is

$$I_{my}^{L} = \prod_{m=1}^{t} \left[\frac{I_{my}^{M}}{I_{m-1,y}^{M}} \right] I_{y-1}^{A} \tag{26}$$

where subscript m denotes month and y denotes year; I_{my}^{M} is the aggregate of all M3 companies; and I_{y}^{A} is the (unbiased) aggregate inventory stock for year u from the ASM.[44] Unfortunately, even this aggregate estimator is not representative because large firms tend to have much lower average growth rates than small firms (see Dunne, Roberts, and Samuelson (1989), for example). Wakim

[44]The Department of Labor, Bureau of Labor Statistics, uses a similar technique to construct monthly employment estimates from its establishment survey.

(1986) found some evidence that the M3 data tended to underestimate the ASM growth rates, primarily because the M3 data was not representative for about half of the industries examined. However, the Bureau's M3 Branch later found that the M3 level estimate for total manufacturing actually overestimated the ASM estimates, which themselves badly understated the CM levels (in 1982).[45]

A second M3 sample selection shortcoming taken up in this study is nonresponse. Because the M3 survey is voluntary, important differences may arise between responding and nonresponding firms. For example, if economically distressed firms stop reporting data (to cut costs) but prosperous firms do not, then the M3 link-relative growth rate would likely be biased. Furthermore, firms may temporarily stop reporting but not be dropped from the sample. For nonreporting firms, Census imputes their data using industry-level growth rates. To the extent that there is considerable heterogeneity among firms, random imputation with industry growth rates is likely to impart bias in the firm-level data as well as in the link-relative estimator.[46] To avoid this potential bias, this study uses a panel of M3 firms with only minimally imputed data and re-imputes the data with firm-level statistical models.

[45] Although the ASM produces, in principle, unbiased estimates of the universe, it too has difficulties with biases entering between censuses (see Davis, Haltiwanger, and Schuh (1990) for a description of this benchmarking problem). In any case, this study cannot correct for, or study the effects of, this sample selection problem.

[46] Only firms with reported data for at least two of the three most recent months are allowed to enter into the link-relative calculation.

References

[1] Thomas A. Abbott. *Producer Price Dispersion and the Analysis of Production*. PhD thesis, Harvard University, Cambridge, MA, 1987.

[2] David Bivin. Disaggregation and the speed of adjustment in inventory models. Unpublished paper, Indiana University-Purdue University at Indianapolis, 1989.

[3] Olivier J. Blanchard. The production and inventory behavior of the american automobile industry. *Journal of Political Economy*, 91(3):365–400, June 1983.

[4] Alan S. Blinder. Can the production smoothing model of inventory behavior be saved? *Quarterly Journal of Economics*, 101(3):431–453, August 1986a.

[5] Alan S. Blinder. More on the speed of adjustment in inventory models. *Journal of Money Credit and Banking*, 18(3):431–453, August 1986b.

[6] Alan S. Blinder and Louis J. Maccini. Taking stock: A critical assessment of recent research on inventories. *Journal of Economic Perspectives*, 5(1):73–96, Winter 1991.

[7] J. C. G. Boot and G. M. de Wit. Investment demand: An empirical contribution to the aggregation problem. *International Economic Review*, 1(1):3–30, January 1960.

[8] T. S. Breusch. Testing for autocorrelation in dynamic linear models. *Australian Economic Papers*, 17:334–355, 1978.

[9] Ricardo J. Caballero and Eduardo M. R. A. Engel. Microeconomic adjustment hazards and aggregate dynamics. *Quarterly Journal of Economics*, 108(2):359–383, May 1993.

[10] Ricardo J. Caballero, Eduardo M.R.A. Engel, and John Haltiwanger. Aggregate employment dynamics: Building from microeconomic evidence. Unpublished paper, May 1994.

[11] Lawrence J. Christiano and Martin S. Eichenbaum. Temporal aggregation and structural inference in macroeconomics. *Carnegie-Rochester Conference Series on Public Policy*, 26(1):63–130, Spring 1987.

[12] Steven J. Davis and John Haltiwanger. Wage dispersion between and within U.S. manufacturing plants 1963-86. *Brookings Papers: Microeconomics*, pages 115–180, 1991.

[13] Steven J. Davis, John C. Haltiwanger, and Scott Schuh. Published versus sample statistics from the asm: Implications for the lrd. In *1990 Proceedings of the Business and Economic Statistics Section*, pages 52–61. American Statistical Association.

[14] Economic Surveys Division. The standard statistical establishment program. Technical report, U.S. Department of Commerce, Bureau of the Census, January 1979. Technical Paper 44.

[15] Timothy Dunne, Mark Roberts, and Larry Samuelson. The growth and failure of u.s. manufacturing plants. *Quarterly Journal of Economics*, 104(4):671–698, November 1989a.

[16] Steven N. Durlauf and Louis J. Maccini. Measuring noise in inventory models. *Journal of Monetary Economics*, 36(1):65–89, December 1995.

[17] Martin Eichenbaum. Some empirical evidence on the production level and production cost smoothing models of inventory investment. *American Economic Review*, 79(4):853–864, September 1989.

[18] Ray C. Fair. The production smoothing model is alive and well. *Journal of Monetary Economics*, 24(3):353–370, May 1989.

[19] Martin Feldstein and Alan Auerbach. Inventory behavior in durable goods manufacturing: The target-adjustment model. *Brookings Papers on Economic Activity*, (2):351–396, 1976.

[20] Murray F. Foss, Gary Fromm, and Irving Rottenberg. Measurement of business inventories. Technical report, U.S. Bureau of the Census, 1980. Economic Research Report 3.

[21] Jeffrey C. Fuhrer, George R. Moore, and Scott D. Schuh. Estimating the linear-quadratic inventory model: Maximum likelihood versus generalized method of moments. *Journal of Monetary Economics*, 35:115–157, 1995.

[22] John Geweke. Macroeconomic modeling and the theory of the representative agent. *American Economic Review*, 75(3):206–210, May 1985.

[23] Moheb A. Ghali. Seasonality, aggregation, and the testing of the production smoothing hypothesis. *American Economic Review*, 77(3):464–469, June 1987.

[24] L. G. Godfrey. Testing against general autoregressive and moving average errors when the regressors include lagged dependent variables. *Econometrica*, 46(6):1,293–1,302, November 1978.

[25] Allan W. Gregory, Adrian Pagan, and Gregor W. Smith. *Estimating Linear Quadratic Models with Integrated Processes*. Blackwell, 1993.

[26] Yehuda Grunfeld and Zvi Griliches. Is aggregation necessarily bad? *The Review of Economics and Statistics*, 42(1):1–13, February 1960.

[27] Kanhya L. Gupta. Aggregation bias in linear economic models. *International Economic Review*, 12(3):293–305, June 1971.

[28] Alastair Hall and Robert J. Rossana. Estimating the speed of adjustment in partial adjustment models. *Journal of Business and Economic Statistics*, 1991.

[29] John C. Haltiwanger and Marc S. Robinson. The effect of taxes on inventories. UCLA Working Paper 182, January 1987.

[30] Lars Peter Hansen. Large sample properties of generalized method of moments estimators. *Econometrica*, 50(4):1029–1054, July 1982.

[31] Ethan S. Harris. A reexamination of the inventory buffer effect with disaggregated data. Federal Reserve Bank of New York Research Paper No. 8817, July 1988.

[32] John C. Hinrichs and Anthony D. Eckman. Constant-dollar manufacturing inventories. *Survey of Current Business*, 61(9):16–23, September 1981.

[33] Albert A. Hirsch and Michael C. Lovell. *Sales Anticipations and Inventory Behavior*. Wiley, New York, 1969.

[34] Charles C. Holt, Franco Modigliani, John F. Muth, and Herbert A. Simon. *Planning Production, Inventories, and Work Force*. Prentice-Hall, Inc., Englewood Cliffs, N.J., 1960.

[35] Paul T. Hunt. *An Investigation of Finished-Goods Inventories Using Micro Data*. PhD thesis, Stanford University, Palo Alto, CA, 1981.

[36] James A. Kahn. Inventories and the volatility of production. *American Economic Review*, 77(4):667–679, September 1987.

[37] Anil K. Kashyap and David W. Wilcox. Production and inventory control at the general motors corporation during the 1920's and 1930's. *American Economic Review*, 83(3):383–401, June 1993.

[38] Alan P. Kirman. Whom or what does the representative individual represent? *Journal of Economic Perspectives*, 6(2):117–136, Spring 1992.

[39] Spencer D. Krane. 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(1):117–136, February 1994.

[40] Spencer D. Krane and Steven N. Braun. Production smoothing evidence from physical-product data. *Journal of Political Economy*, 99(3):558–581, June 1991.

[41] Kon S. Lai. Aggregation and testing of the production smoothing hypothesis. *International Economic Review*, 32(2):391–403, May 1991.

[42] Kevin C Lee, M. Hashem Pesaran, and Richard G. Pierse. *Disaggregation in Economic Modelling*, chapter Aggregation Bias in Labour Demand Equations for the U.K. Economy, pages 113–49. Routledge, London and New York, 1990.

[43] Michael Lovell. Manufacturers' inventories, sales expectations, and the acceleration principle. *Econometrica*, 29(3):293–314, July 1961.

[44] Michael C. Lovell. Simulating the inventory cycle. *Journal of Economic Behavior and Organization*, 21(2):147–79, June 1993.

[45] Louis J. Maccini and Robert J. Rossana. Joint production, quasi-fixed factors of production, and investment in finished goods inventories. *Journal of Money Credit and Banking*, 16(2):218–236, May 1984.

[46] Louis J. Maccini and Scott Schuh. Credit market conditions and inventory investment. Unpublished paper, Johns Hopkins University and the Federal Reserve Board, January 1995.

[47] Jeffrey A. Miron and Stephen P. Zeldes. Seasonality, cost shocks, and the production smoothing model of inventories. *Econometrica*, 56(4):877–908, July 1988.

[48] Jeffrey A. Miron and Stephen P. Zeldes. Production, sales, and the change in inventories: An identity that doesn't add up. *Journal of Monetary Economics*, 24:31–51, July 1989.

[49] Kevin M. Murphy and Robert H. Topel. Estimation and inference in two-step econometric models. *Journal of Business and Economic Statistics*, 3:370–379, October 1985.

[50] Whitney K. Newey and Kenneth D. West. A simple, positive definite, heteroscedasticity and autocorrelation consistent covariance matrix. *Econometrica*, 55(3):703–708, May 1987.

[51] Sang V. Nguyen and Stephen H. Andrews. Stage-of-fabrication inventory behavior: A general target-adjustment model. *Applied Economics*, 21(1):175–192, February 1988.

[52] Stephen Nickell. Error correction, partial adjustment and all that: An expository note. *Oxford Bulletin of Economics and Statistics*, 1985.

[53] U.S. Bureau of the Census. Manufactures' Shipments, Inventories, and Orders: 1982-92. Technical report, Current Industrial Report M3-1(92), 1993.

[54] Valerie A. Ramey. Inventories as factors of production and economic fluctutations. *American Economic Review*, 79(3):338–354, June 1989.

[55] Patricia Reagan and Dennis P. Sheehan. The stylized facts about the behavior of manufacturers' inventories and backorders over the business cycle. *Journal of Monetary Economics*, 15:217–246, March 1985.

[56] Komei Sasaki. An empirical analysis of linear aggregation problems: The case of investment behavior in japanese firms. *Journal of Econometrics*, 7(3):313–331, June 1978.

[57] Scott Schuh. *Aggregation Effects in Production Smoothing and Other Linear Quadratic Inventory Models*. PhD thesis, Johns Hopkins University, Baltimore, MD, 1992.

[58] Helmut Seitz. Still more on the speed of adjustment in inventory models: A lesson in aggregation. *Empirical Economics*, 18(1):103–27, 1993.

[59] Henri Theil. *Linear Aggregation of Economic Relations*. North-Holland Publishing Company, Amsterdam, 1954.

[60] Henri Theil. *Principles of Econometrics*. John Wiley & Sons, Inc, 1971.

[61] P.K. Trivedi. Distributed lags, aggregation and compounding: Some econometric implications. *The Review of Economic Studies*, 52(1):19–35, 1985.

[62] Anne Wakim. Evaluation of coverage and response in the manufacturers' shipments, inventories, and orders survey. *American Statistical Association Proceedings of the Business and Economic Statistics Section*, pages 523–528, 1986.

[63] Kenneth D. West. A variance bounds test of the linear quadratic inventory model. *Journal of Political Economy*, 94(2):374–401, April 1986.

[64] Kenneth D. West. Inventory models. NBER Technical Working Paper No. 143, September 1993.

[65] Kenneth D. West and David W. Wilcox. Some evidence on finite sample behavior of an instrumental variables estimate of the linear quadratic inventory model. NBER Technical Working Paper No. 139, December 1994.