NBER WORKING PAPERS SERIES

OUTPUT FLUCTUATIONS AT THE PLANT LEVEL

Timothy F. Bresnahan

Valerie A. Ramey

Working Paper No. 4105

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 1992

We are very grateful to the following research assistants who helped in the construction and analysis of the data set: Mary Jo Arboleda, Subir Bose, Devajyoti Ghose, Shane Greenstein, Joao Issler, Diana Nguyen, Timothy Taylor, Michael Turner, and Wendy Wilson. We also benefited from helpful comments from Ana Aizcorbe, Frances Hammond, Joao Issler, Patrick McAllister, Garey Ramey, Glenn Sueyoshi, and participants at the 1991 NBER Summer Institute and seminars at The University of Chicago, The Wharton School, and the Federal Reserve Board. Timothy Bresnahan thanks the Sloan foundation and Valerie Ramey thanks the UCSD Committee on Research and the National Science Foundation grant SES-9022947 for financial support. This paper is part of NBER's research program in Economic Fluctuations. Any opinions expressed are those of the authors and not those of the National Bureau of Economic Research.

NBER Working Paper #4105
June 1992

OUTPUT FLUCTUATIONS AT THE PLANT LEVEL

ABSTRACT

This paper studies weekly output fluctuations from 1972 to 1983 at fifty final assembly plants in the U.S. automobile industry. The study makes use of a new data set that contains detailed information on plant operations. The main findings of the paper are: (1) Even at the simplest fabrication and assembly plant, there are a variety of margins on which production quantities are adjusted; (2) The production adjustment margins appear to have very different dynamic characteristics; and (3) The analysis of plant level data can lead to conclusions that are dramatically different from those reached using aggregated data, even though the data are driven by industry-wide shocks.

Timothy F. Bresnahan
Department of Economics
Stanford University
Stanford, CA 94305
and NBER

Valerie A. Ramey
Department of Economics, D-008
University of California, San Diego
9500 Gilman Drive
La Jolla, CA 92093-0508
and NBER

The nature of output fluctuations across time and across firms is a central theme of macroeconomic analysis. Outstanding questions include the nature of the shocks hitting the economy, the importance of returns to scale for magnifying or dampening shocks, and the role of adjustment costs in propogating those shocks across time. Most empirical studies of production dynamics have used aggregate or industry–level data on outputs, inputs and factor prices. The results, while natural for addressing macroeconomic questions, have limited informational content. Aggregation obscures fundamental facts about fluctuations at the plant level in two ways. First, and familiar, is the point that different events may be happening at different plants. That is, not only may individual plants experience different kinds of external shocks, but they may react to the same shock in different ways because they started in different "states." In this paper, we provide new evidence for the importance of individual plant heterogeneity in states and responses. Thus, this paper contributes to the growing literature that uses plant level data to analyze aggregate movements (e.g., see Davis and Haltiwanger (1989), Bartelsman and Dhrymes (1992)).

A second, and less familiar, point about aggregation concerns the focus of the analysis. Most economic models of production focus on outcomes, not processes. The outcomes include the evolution of output, employment, materials, and capital stocks over time. The production process itself, however, is a "black box" (Rosenberg (1982)). The typical macroeconomic approach is to infer the nature of the economic problem facing the manager by observing the outcomes, and to ignore completely the workings of the black box. This approach is standard even in studies based on individual plant data, where the economic analyst often treats the plant level data as if it were aggregate data and leaves the plant as a black box. Outcomes do not convey all the relevant economic information that process choices reveal. In particular, information about what the plant's managers actually decided to do and why they thought they were doing it contains information about dynamically important decisions. While data about the decision process are sometimes proprietary, they are often publically available, and thus constitute an important source of information.[1]

For this paper, we have assembled a weekly panel data set revealing plant level production decisions for 50 U.S. automobile plants over 626 weeks from 1972 to 1983. We used only public sources for the data. Entries for each plant consist of data on the hours of operation, overtime hours, the line speed, the number of shifts, the days closed, and the reasons the factory was closed. We have also gathered data on actual production by nameplate, and sporadic data on employment numbers, layoffs, and hires. We know of no other data set that covers as great a number of firms, at so high a frequency, in such detail.[2] We have chosen the automobile industry because (1) automobiles are manufactured using a fabrication and assembly process, which is the single most important technology in manufacturing; (2) the industry displays substantial cyclical volatility; and (3) fuel price shocks have had a particularly dramatic effect on the industry. We believe that a detailed analysis of the production dynamics in the automobile industry will provide insight into the nature of output fluctuations in general. Further, our use of information on the short–run decision process differentiates this work from that of others who have studied the automobile industry (e.g. Abernathy, Clark and Kantrow (1983), Blanchard (1983), Bresnahan (1981), and Ramey (1991)).

The main findings of the paper are:

---

[1]See Chew (1988) as an example of the successfully exploitation of this type of information.

[2]Ana Aizcorbe has collected monthly data for twelve assembly plants covering a later period. While her data set is not as detailed in some aspects, it does contain monthly proprietary employment statistics. See Aizcorbe (1990) and Cooper and Haltiwanger (1991) for analyses using those data.

(1) Even at the simplest fabrication and assembly plant, there are a variety of margins on which production quantities are adjusted.

(2) The production adjustment margins appear to be have very different dynamic characteristics.

(3) The analysis of plant level data can lead to conclusions that are dramatically different from those reached using aggregated data, even though the data are driven by industry–wide shocks.

The paper proceeds as follows. Section 1 will present a more detailed description of the data and how it was collected. Section 2 will analyze the margins along which production is varied and the dynamic characteristics Section 3 will discuss the implications for conclusions reached using aggregate data. Section 4 will conclude.

1. **Data Description**

The data were gathered from several automobile industry publications. The main source of data is the weekly periodical *Automotive News.* Each issue contains an article describing the production of cars during the previous week, as well as a table of production numbers by model. The article gives detailed information on which plants had overtime, both during the week and on Saturday, whether plants were closed down and for what reason, and any changes in line speeds and shifts. An example of a particularly informative article is the January 13, 1975 production article (*Automotive News*, page 22):

"Eleven plants were idle last week due to the current sales slump. Closed

> were: American Motors' Kenosha plant, Chrysler's Newark and St. Louis facilities, Ford's Chicago, Dearborn, Kansas city, Mahwah, Metuchen and San Jose car operations and General Motors' South Gate plant and the Pontiac home plant.
>
> "Chrysler, with four plants reopened for the first time since Thanksgiving, produced 7,000 cars last week...
>
> "American Motors' Kenosha plant, which has been idle since Dec. 20, will reopen today and begin production of Pacer, the new small car....
>
> "Saturday work was scheduled for the (GM) Corvette line at St. Louis....
>
> "GM's South Gate facility is idle for all of January, and the Doraville, Fairfax and Willow Run plants will be reduced to a single shift effective today...
>
> "Also, (GM) Lordstown will begin single-shift Vega-Astre operation Jan. 20 at a rate of 100 units per hour. Earlier, it had been announced that Lordstown would dip from 100 to 85 units per hour on two shifts Jan. 20. The new one-shift operation will result in the indefinite layoff of an additional 2,100 employees.
>
> "GM said indefinite hourly layofs will be about 92,000 by the end of January..."

Other sources of data that were used to identify line speed and shift changes that were not reported in the *Automotive News* articles were *Wards' Automotive Yearbook* and *Wards' Automotive World.*

Perhaps the most unique aspect of the data set is the set of reasons given for plant shutdowns. We have classified these reasons into four categories: (1) model changeovers, (2) holidays and vacations, (3) inventory adjustments, (4) supply disruptions. The first category, "model changeovers," contains the days closed due to adjustments for model changeovers. As pointed out by Cooper and Haltiwanger (1990), this category represents an important part of production volatility. The second category, "holidays," is the days closed for holidays and vacations specified in the union contracts. The category labelled "inventory adjustment" represents the times the company shut the plant down because the dealer inventories of the model produced at that plant were "excessive." The "supply disruption" category contains shutdowns due to strikes, both onsite and offsite, parts shortages, inclement weather, earthquakes, fires in the paint facility, and general machinery breakdowns.

The one difficulty with the data set is that the operations data are at the factory level, while the actual production data are at the model level. There are some seventy models during the period, with most factories producing several models and most models being produced by several factories. Therefore, matching the production data with the factory is difficult in all but a few cases. We have matched production data to six plants so far. For those plants, we analyze actual production. For the universe of plants, however, we must analyze variations in short–run *posted output*, which differs from actual output by deviations from posted line speed.

We should also mention that there are likely to be measurement errors in the data. In some cases, we knew a line speed change occurred, but did not know exactly what day it occured. In these cases, we tried to use the actual production numbers to pinpoint the date. In a few instances, we knew the number of plants that had closed down but not which ones. Again, we tried to use the actual production numbers to identify which plants closed down.

In all, we study 50 assembly lines. When a plant had two lines, we treated each line as a separate plant. Nineteen of the plants had missing values over some part of the period. The missing values occurred if there was a permanent shutdown of the plant, a conversion to light truck production, or if the plant opened during the sample period. GM plants Bowling Green, Oklahoma City, and the new Pontiac plant opened during the sample, while GM plants Fremont, Lakewood, Pontiac, Southgate, St. Louis Chevrolet, St. Louis Corvette, and the second line at Detroit closed near the end of the sample, most in 1981. Ford Los Angeles, Louisville, Mahwah, San Jose, and Twin Cities closed or converted at the end of sample, while Norfolk converted in the middle of the sample. Chrysler Hamtramck, Jefferson Avenue, and Lynch Road also closed, typically near the end of the sample.

2. Multiple Margins of Output Variation

A. Overview

In this section we will use the data set to show that there are a variety of margins on which to adjust production quantity, and that those margins have fundamentally different dynamic properties. In our analysis, we will work with a particular definition of output that relates to the margins that plants adjust. Actual production Q_{it} by factory i for week t is given by the following identity:

$$(1) \qquad Q_{it} = (RH_{it} + OH_{it}) \times (LS_{it} - \epsilon_{it}) \times SH_{it} ,$$

where

RH = regular hours, the number of hours the plant runs each shift per week for which it pays a straight–time wage to its workers.

OH = overtime hours, the number of hours the plant runs each shift per week for which it pays an overtime premium to its workers.

LS = posted line speed, i.e., potential output per hour per shift

ϵ = deviations from the posted line speed

SH = number of shifts, either one or two

Regular hours are usually varied by shutting the plant down for a day or a week. Overtime hours are usually varied by scheduling Saturday work or adding an hour or two to each shift; eight hours on Saturday is by far the most common form of overtime. A change in the line speed always involves a change in the number of workers, as does a change in the number of shifts.

The decomposition in equation (1) immediately illustrates how the focus of this paper differs from the usual approach. None of the hours or shift variables refers to hours per worker, or total employee hours. Rather, the hours and shift variables measure the workweek of the plant.[3]

What is wrong, or incomplete, in thinking of production managers as "choosing Q"? Much of the task of production scheduling at an assembly plant is making decisions about how best to meet a production plan. The production plan embodies a choice of Q already, but does not yet determine how to attain it. The circumstances affecting this choice in the short run are dynamic. Looking backward, the managers know the number of workers affiliated with the plant, the technical capabilities of the plant itself, and materials suppliers' capabilities and contractual obligations. In the short–run, these factors appear as constraints. The dynamics of the constraints are not at all the same. Asking existing workers to work overtime is an expensive way to expand Q because it requires the payment of hourly wages that are fifty percent higher. Overtime, however, has few dynamic implications. Adding workers or whole shifts, however, requires more time and involves making contractual commitments that are somewhat costly to reverse. Changing the throughput rate of operations (the "line speed" in an assembly–line based facility like the car plants we study) is difficult ("there ain't no 'go faster' switch on the wall of this factory" one manager told us) and has permanent implications. On the downside, plants face a variety of contraints on how they decrease production. For example, labor contracts specify that if the plant is in operation for part of a week, workers must be provided 80 percent of their wage multiplied by the difference between the number of hours paid and 40. Thus, shutting a plant down for less than a week does

[3]The number of workers and hours per worker do vary, of course, but those variations are the *result* of changes in the workweek of the plant, the number of shifts, or the line speed. It is the latter variables that the managers manipulate directly; it is the effect on the labor variables that generates much of the cost.

not reduce labor costs proportionately.

The same set of decisions about meeting a production plan has forward looking dynamic implications as well. To add workers or shifts is to have larger productive capabilities for the near future. To change line speeds, with all the reorganization, job redefinition, and other changes that accompany it, conditions the production scheduling decisions for years. At the other extreme, working a day of overtime costs money now, but neither increases nor decreases the plant's capabilities. One would expect, therefore, that forward looking plant managers would use different mechanisms to respond to dynamically different circumstances. Overtime, for example, might be used to respond to transitory (or possibly transitory) increases in desired output. More permanent (or apparently more permanent) increases would be met using long–run cheaper, though short–run more expensive, changes, like adding shifts.

We seek to gain a better understanding of the complexities of these margins and their differences across horizons by analyzing the components of Q given in equation (1). We begin by investigating the frequency of the use of each margin, and its importance in the overall variation in output. We then study the dynamic aspects of the manipulation of the margins.

B. Frequencies and Variance Decompositions

We analyze the importance of each margin in two ways. We first compute simple statistics that tell us how often each margin is used. We then decompose the variance of output into changes in the different margins. Periods during which a plant is permanently closed are not included in the analysis. In this paper, we do not study plant closings and openings, which we interpret as a *firm* production margin, rather than a *plant* production margin.

Table 1 quantifies how often plants use each of the margins by showing the percentage of weeks during which each margin was manipulated. The first row shows the weighted average of the universe of all plants, the second row shows the calculation for the same sample, but with closures for holidays excluded, the third row shows the weighted average of the six matched plants, and the following six rows show the data for each of the six matched plants. The weights used for the averages are based on a plant's total output for the entire sample period. Thus, plants that were open for only part of the period received less weight in this and all later calculations.

Consider first the averages for the universe of plants. The results show that on average, each plant is shut down for at least a day twenty–five percent of the weeks. Each plant is typically shut down for an entire week thirteen percent of the time or almost seven weeks per year. Overtime hours (per worker per shift) in excess of four hours per week are also frequent, occurring on average seven weeks per year. Thus the use of overtime hours is as frequent as the use of shutdowns of a week. On the other hand, both changes in shifts and in line speeds are very infrequent, occuring substantially less than once per year. The second row shows the frequency of shutdowns for reasons other than holidays. The frequency of shutdowns of at least one day drops by half, while the frequency of shutdowns of one week falls a small amount. It is clear that most non–holiday shutdowns are shutdowns of one week, and they occur ten percent of the time.

The remaining rows calculate the same frequencies (including holidays) for the six matched plants. The data are better for these plants because we were able to use the actual output data to detect overtime hours that were not reported in our sources and to detect deviations from linespeed. The average frequencies for the subset of plants are similar to those for the universe of plants. The only exception is overtime hours, which

TABLE 1
Frequency of the Use of Different Margins

	Percent of Weeks During which a Plant Experienced:				
	Shutdown $\geq$ 1 day	Shutdown of 1 week	Overtime > 4 hours per week	Change in number of shifts	Change in line speed
Weighted avg. of all plants	24.8%	12.9%	14.3%	0.61%	0.85%
Weighted avg. of all plants, holidays excluded	11.6	9.8	14.3	0.61	0.85
Weighted avg. of six matched plants	25.9	13.0	31.9	0.71	0.86
GM Bowling Green	24.4	13.3	12.6	0	0
GM Lordstown	27.2	13.7	17.9	0.96	0.96
GM Norwood	27.2	14.2	31.3	0.32	0.64
GM St. Louis	17.8	4.6	61.1	0.20	0.40
Ford St. Louis	26.5	13.3	38.8	0.64	0.48
Ford Wixom	23.3	12.0	47.9	0.96	1.44

Decomposition of Shutdowns
Percent of Days Closed Due to:

	Model Changeovers	Holidays	Inventory Adjustment	Supply Disruption
Weighted Average of all plants	33.3%	34.2%	25.2%	7.1%

are used twice as often. This increase in the calculated frequency in overtime is due, in part, to our ability to detect unreported overtime using the output data. Finally, looking down the columns, one can see that there is a good deal of heterogeneity across plants in the frequencies.

The bottom of Table 1 decomposes shutdowns, or idle time, into each of the four reasons given: model changeovers, holidays, inventory adjustment, and supply disruptions. It is important to decompose idle time in this way, because two of the categories, holidays and supply disruptions are not directly manipulated by the managers of the plant. The breakdown shows that on average a third of the idle time is due to holidays, another third is due to model changeovers, twenty–five percent is due to inventory adjustment, and seven percent is due to supply disruptions. Thus, sixty percent of the idle time is, for the most part, due to direct manipulation of that margin.

In order to give an idea of the patterns in the use of each margin, figures 1 – 6 show graphs of overtime hours, days closed for various reasons, and changes in line speeds and shifts, aggregated over all plants. The graphs illustrate the patterns brought out from the tables. In general, overtime hours and days closed, especially for inventory adjustment, are frequent events. Line speed and shift changes, on the other hand, occur relatively infrequently.

The graphs also show some interesting dynamic patterns. Note, for example, that the occurrences of days closed for inventory adjustment and overtime hours tend to come in clumps. Furthermore, as one would expect, the overtime hours tend to occur during those periods classified as booms while days closed for inventory adjustments tend to occur in those period classified as recessions. It is clear, however, that there are many periods in which some plants are using overtime while others are shutting down for inventory adjustment.

Figure 1: Total Weekly Overtime Hours
(Sum over all shifts at all plants)

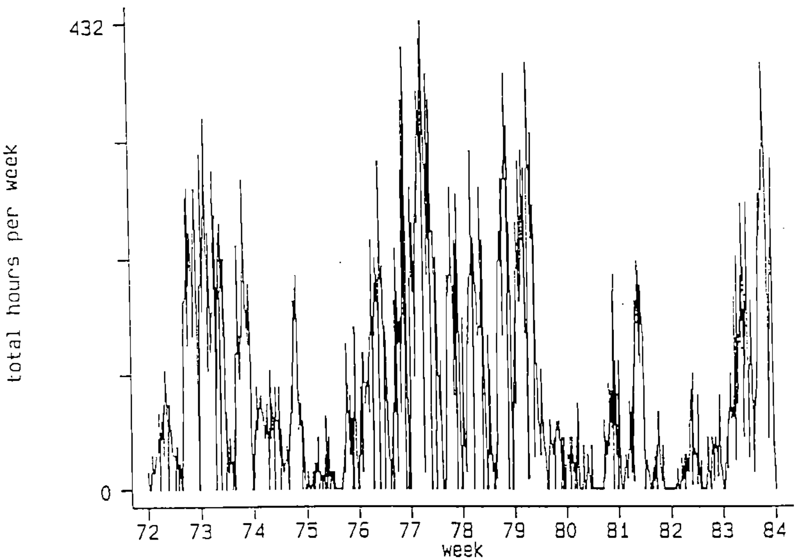

Figure 2: Total Days Closed for Inventory Adjustment
(Sum over all plants)

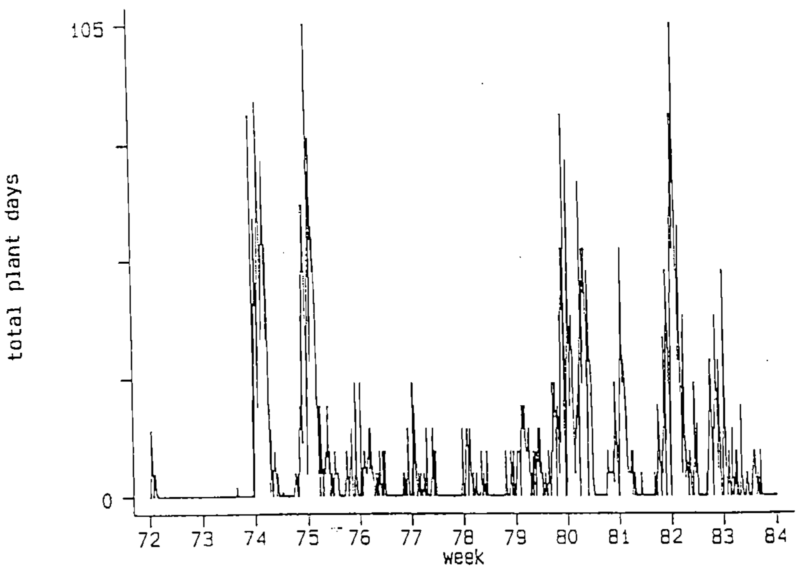

Figure 3: Total Days Closed for Model Changeover
(Sum over all plants)

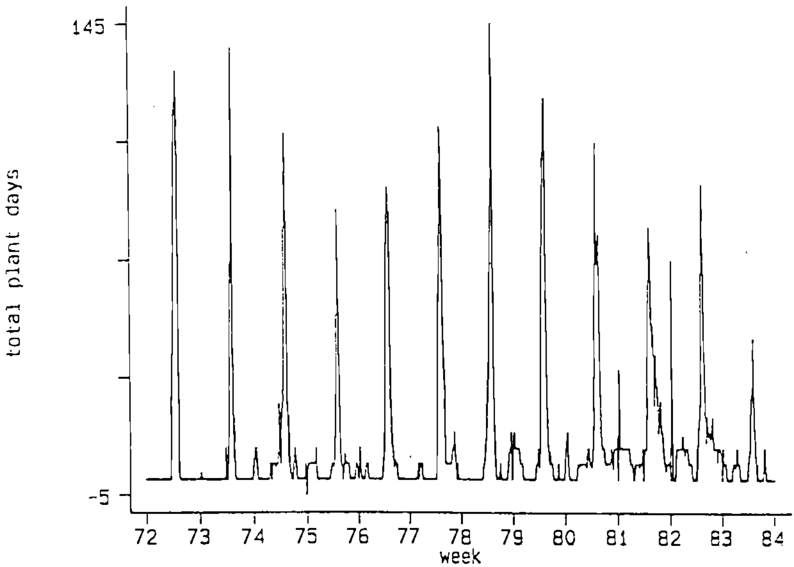

Figure 4: Total Days Closed for Supply Disruptions
(Sum over all plants)

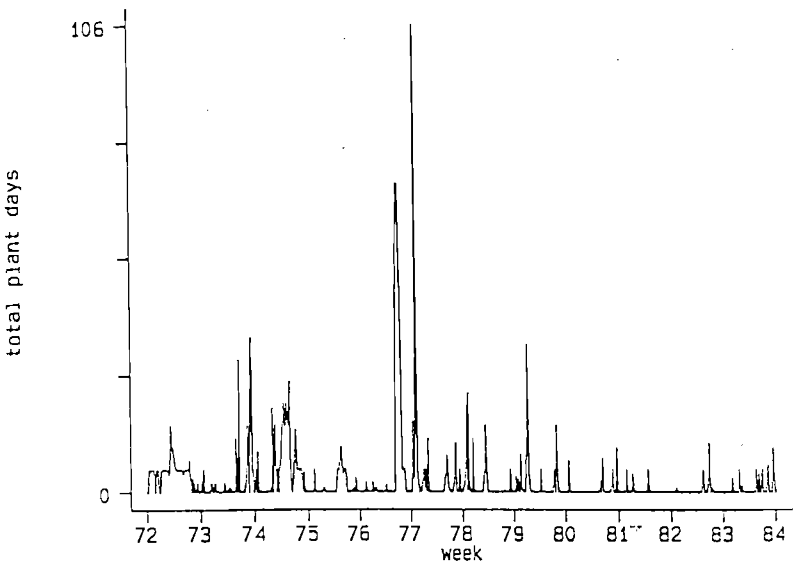

Figure 5: Line Speed Changes
(Sum over all plants)

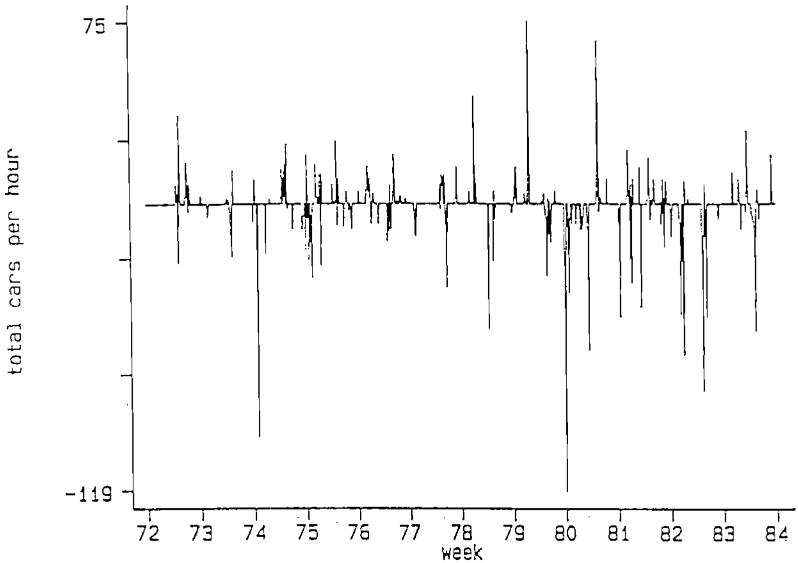

Figure 6: Shift Changes
(Sum over all plants)

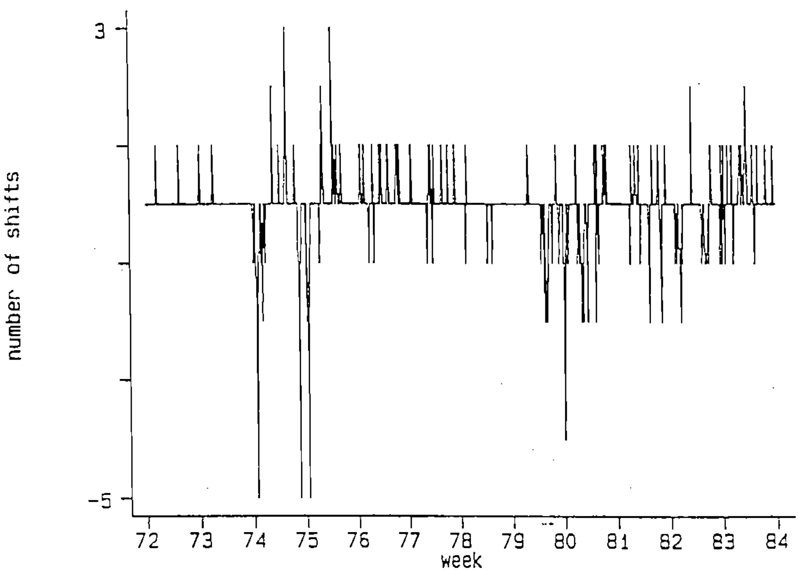

The frequencies calculated above do not reveal directly the importance of each of these margins for the short–run variation in output, because the effect of a change in each margin is different. For example, each shutdown involves a large variation in output, with the decline in weekly output ranging from 20 percent in the case of one day to 100 percent in the case of a week. Overtime hours, which often take the form of Saturday work, typically involve a twenty percent increase in output, so their impact is less than that for variations in regular hours. On the other hand, the addition of a shift doubles output. The magnitude of line speed changes varies from plant to plant.

Before decomposing the variance of output, we must discuss the measure of output we use. Recall that actual production is only publicly available at the model level. Thus, for the universe of plants we study *posted output*, which is Q in equation (1) when ϵ is equal to zero. We also present results for the six matched plants, which contain data for actual production and ϵ. The results show that in most cases ϵ is not an important source of fluctuations in output.

The decomposition of variance is not straighforward because, as equation (1) shows, capacity equals the product of the components, so we do not have a linear relationship. Furthermore, we cannot take logs because the hours components are frequently equal to zero. Therefore, we use a Taylor series approximation to decompose the variance, and eliminate covariances by orthogonalizing the components. For the orthogonalization the variables are ordered as follows: (1) regular hours (2) overtime hours (3) shift changes (4) line speed changes. Different orderings changed the percentages by less than one percentage point. For the matched group of plants, we also included the deviation from line speed in the variance decomposition in order to assess its importance.

Table 2 shows the results of the variance decomposition. Consider first the results

TABLE 2
Variance Decomposition of Output

	Percent of Variance Explained by:				
	Regular Hours	Overtime Hours	Number of shifts	Line Speed	Deviations from line speed
Weighted avg. of all plants	70.8%	3.4%	19.5%	6.3%	—
Weighted avg. of all plants, holidays excluded	65.1	4.0	23.1	7.8	—
Weighted avg. of six matched plants	58.9	6.4	23.5	5.8	5.4
GM Bowling Green	63.8	4.4	0	0	31.8
GM Lordstown	62.4	3.0	21.1	8.0	5.4
GM Norwood	67.0	5.8	14.9	5.3	7.0
GM St. Louis	55.2	13.5	10.5	18.0	2.8
Ford St. Louis	48.8	10.2	38.1	1.2	1.6
Ford Wixom	50.5	9.0	31.2	3.2	6.0

for the universe of plants, presented in the first row. We find that variations in regular hours account for most of the variation in posted output, amounting to over seventy percent. If we multiply this number by the percent of regular hours variations caused by inventory adjustment and model changeover, which are directly controlled, we find that forty percent of the variation in production comes from the decision of managers to manipulate regular hours. Furthermore, most of those controlled changes in regular hours involve shutting the plant down for a week. Second in importance is changes in the number of shifts, which account for almost twenty percent. Overtime hours and line speed changes each account for very little of the variation.

The second row performs the same decomposition when variations due to holidays are eliminated. The importance of regular hours falls by six percentage points, and the importance of the other categories rises somewhat. The overall picture, however, is unchanged.

For the matched plants, variations in regular hours still tend to account for the bulk of variation in actual output, but the average is somewhat lower. The contribution of overtime ranges from 3 percent at Lordstown to 10 percent at GM St. Louis. The weighted average for overtime is 6 percent which is double that for the average for the universe of plants, but still low. Changes in shifts and line speeds contribute amounts similar to the universe of plants. Finally, deviations from line speed contribute a little over five percent of the variance on average. The number for Bowling Green is very atypical; the plant experienced many problems in the year after it was opened.

These results support the following interpretation. The primary way in which managers vary production is by changes in regular hours. Many (nòn holiday) related changes in regular hours involve closing the plant down for a week. Thus, managers manipulate production by varying the number of weeks its labor force works. The second most important way managers vary production is by adding or dropping a shift. We think of this as adding or dropping the second "team" of workers. Overtime hours, while used frequently, are less importance in the overall variance of output. Finally, changes in the line speed, which can be interpreted as changes in technological layout that lead to a change in the size of a "team," are relatively unimportant.

C. Dynamic Characteristics of Changes on Different Margins

We now turn to an analysis of the different dynamic implications of changes on different margins. To see whether plant managers appear to be forward–looking in their behavior, and to begin to see what their behavior might reveal about the short run dynamics of cost and of desired production, we conducted two investigations of persistence. For both analyses, we first isolated several states of the world in which a plant might be in a given week. Plants can be in either one– or two– shift operations. For each of those, there are four statuses. Plants can operate 4 or fewer (including 0) days, reporting the reason as "inventory adjustment" or "model changeover". That is status A. Status B is 0 to 4 days of operation because of holidays or supply disruptions. We distinguish between statuses A and B because status A is more "voluntary" — the events in status B are mostly unavoidable, those in status A, chosen by the managers. The regular hours status, C, is defined as operations for more than four but less than five and a half days. This definition treats small amounts of overtime as part of normal operations.[4] Status D, overtime, is defined as more than 5.5 days operations.[5]

---

[4]These definitions incorporate a few of our judgement calls that may not be obvious. For example, when a plant has a shutdown one day during the week but then works just one weekend day of overtime to make it up, we classify the plant as regular hours. More than a day of weekend overtime to make up the lost day leads to a classification in overtime status. There is an obvious conflict here between a marginal cost view of what the states should be (which would call paid overtime for a whole shift an indicator of high shadow value of the plant's capacity) vs. a planned–production view. There are so few of the

We first examine the transitions between states in order to understand the permanence of the states themselves. We use a framework with eight states, corresponding to one– or two–shift operations, each with four possible statuses A through D above. We estimated the full 8–by–8 transition probability matrix, but in Table 3 we report only the two 4–by–4 on–diagonal blocks. The transitions from 1 to 2 shift operations and back are infrequent enough that the off–diagonal blocks are basically matrices of zeroes.

The blocks on the diagonal tell a very interesting story, however. First, look down the diagonal to see which states tend to persist into the next week. The most clearly persistent state is regular hours, in either one or two shift operations. Depending on shifts, the probability of persisting in that state is around 70%. The next most persistent state is "voluntary" short days, with about a 50% probability of persistence. The overtime state is not far behind, with persistence probabilities in the forties. The exception is the "involuntary" or "supply shock" short days state, which is not very persistent at all.

The interpretation is simple. An automaker's dream life, the whole point of mass production, is persistent regular hours. That is the cheapest way to make vehicles, and the goal of the marketing, forecasting, and production planning functions is to get the plants into that state and keep them there. Even in the highly uncertain economic environment of our sample period, that is a very persistent state. On the other hand, supply/demand imbalances that lead to voluntary operations at inefficiently low (status A) or inefficiently high (status D) levels are persistent, as well. Note that movements

---

ambiguous cases that we cannot usefully give them their own separate state.

[5]The overtime data here are based only on the reports in Automotive News. As a result, they should be interpreted as plant overtime, not worker overtime. We believe that we underestimate overtime because many small events go unreported. A closer investigation of our six matched plants reveals that this is in fact the case.

Table 3
Transition Probability Matrix

Status at t / Status at t−1	Status A: Short Week IA or MC	Status B: Short Week Holid or SD	Status C: Regular Week	Status D: Overtime
1 Shift				
Short Week IA or MC	0.522	0.078	0.377	0.007
Short Week Holid or SD	0.081	0.213	0.609	0.083
Regular Week	0.059	0.151	0.728	0.052
Overtime	0.008	0.096	0.430	0.464
Probability of each state	2.8	3.8	16.6	2.2
2 Shifts				
Short Week IA or MC	0.510	0.083	0.385	0.004
Short Week Holid or SD	0.074	0.207	0.546	0.169
Regular Week	0.056	0.161	0.677	0.103
Overtime	0.004	0.096	0.488	0.412
Probability of each state	7.0	11.3	45.1	11.2

from status C to statuses A and D usually involve a change in production of at least twenty percent. Thus, although the automaker prefers to remain near forty hours a week, his manipulation of the margins leads to a great deal of volatility in actual production. The only really transitory state is low production from supply shocks — those are mostly holidays, strikes, and weather interruptions, so it is no surprise that they appear to be quite transitory. The bottom line is that output fluctuations at the monthly (or longer) frequencies appear to be mediated through persistent choices by managers.

The off–diagonal elements reinforce this view. First, consider the rows and columns of the matrices which do not involve status B. Then, the pattern is simple. Most events are stasis, staying in the same state. Most transitions are to adjacent states, with "state skipping" transitions fairly rare. This is consistent with, though it obviously does not prove, a "one–factor" model of the underlying desired production. Desired production moves slowly over time, only occasionally crossing the boundary between statuses A and C, or C and D, and causing a change in plant status. This view is incomplete, however, in that it ignores status B, the reported supply shocks. Status B can occur anywhere in a plant's transitions among statuses A, C, and D, but it is a very transitory state.

To investigate this view further, we took the full 8x8 transition probability matrix and raised it to the fourth power. If the states follow a first–order markov process, that matrix should look a lot like the transition probability matrix four weeks forward. How do they differ? Table 4 shows the diagonals of each of the matrices. We quickly noted two major differences. First, $M1^4$ underpredicts the persistence of statuses A, C, and D. It is clear that these statuses seem to be characterized by positive duration dependence. Second, the the first–order markov model overpredicts persistence in state B. Looking at the transitions, it is easy to see what is driving this. Contrast the probability of being in

Table 4
Two Measures of Monthly Persistence:
Diagonal of Transition Probability Matrix

	Status A: Short Week IA or MC	Status B: Short Week Holid or SD	Status C: Regular Week	Status D: Overtime
		1 Shift		
Given stat at week t–4 (M4)	0.306	0.117	0.674	0.377
Implied by first–order Markov Model (M1**4)	0.150	0.142	0.628	0.119
		2 Shifts		
Given stat at week t–4 (M4)	0.261	0.098	0.645	0.410
Implied by first–order Markov Model (M1**4)	0.136	0.149	0.598	0.160

status B next week, conditional on being in status B this week, under two different further conditioning events. (i) Conditional on not having been in status B last week, the probability of staying in status B is much higher than (ii) the same probability conditional on having been in status B last week. One week transits through status B are the norm, with longer stays very much less likely. Supply shocks are not only transitory at the weekly frequency, they display substantial negative duration dependence.

We then investigate what the transitions between states predict about future movements in output. Working from the definitions of the states given above we identified "events" as movements between states. For movements between states C and D we distinguished whether the firm had one or two shifts; in no other case did it make a difference. We also defined as separate events the changing of the numbers of shifts and changing the line speed. The residual event was staying in the same state with no change in shifts or line speeds. Movements from A to D and D to A were grouped in the residual category because those movements were very infrequent.

Using dummy variables for these events, we ran the following set of regressions on the pooled data set for all plants:

$$Q(t+i) - Q(t-1) = \text{ constant } + \beta\,[Q(t) - Q(t-1)] \text{x event class dummmies,}$$

for $i = 4$ weeks and $i = 13$ weeks. The estimated coefficients β reveal the permanence of the change in output resulting from a change in states. That is, the β's give the fraction of the original change in output Q that is still in effect in one month and in one quarter. We ran the regressions with and without 52 weekly dummies. We will report the results of these regressions in a descriptive sense together with a preliminary interpretation of what they might mean economically. This is for our readers' convenience, rather than because we believe that the descriptive results prove the economic case.

The results are shown in Table 5. Let us begin with the simplest result, that called "no change." In the first column, this has a coefficient of .68. Descriptively, this means that a change in quantity that does *not* lead to a change in plant status tends systematically to be partially (32%) reversed four weeks later; only 68% of any change up or down persists the whole four weeks when the managers accomodate it without a plant status change. Adding week dummies makes within–status changes look even less persistent (column 2, coefficient of .60) and looking an entire quarter ahead (columns 3 and 4) leads to smaller estimates of persistence. The economic interpretation is relatively straightforward. When the managers do not change the plant's status, they have not incurred any adjustment costs. Thus we see no endogenous forces that might have led to persistence. Why, then, is there substantial reversion? When plant status does not change, we have moved on one of the locally flat portions of the SRMC curve. If most small shocks to desired output are transitory, one would expect them to be accomodated by transitory shifts in actual behavior. An alternative theory posits somewhat more clever managers. If changes in desired output vary in their predictable permanence, then the very fact that the managers chose not to change plant status may indicate that they thought the shock was transitory.

Now contrast these results with those for adding and dropping shifts, and for increasing and decreasing line speeds.[6] Changes in quantity associated with these decisions are more persistent. A change in the number of teams or in the composition of a team certainly involves adjustment costs, so these sorts of decisions are not immediately reversed. The shift and line speed results are symmetric in both directions for horizons up

[6]Some line speed changes are due to changes in the type of car produced, and thus are due to technological changes rather than short–run changes in desired output.

Table 5
The Effect of Operations Choices on Predicted Persistence
Pooled Data Set

Dependent Variable	$Q_{t+4}-Q_{t-1}$	$Q_{t+4}-Q_{t-1}$	$Q_{t+13}-Q_{t-1}$	$Q_{t+13}-Q_{t-1}$
Constant	–48	*	–106	*
$\Delta Q(t+1)$x 52 weekly dummies		included		included
No Change	0.68 (.02)	0.60 (.02)	0.61 (.03)	0.55 (.02)
Add Shift	0.74 (.07)	0.75 (.07)	0.80 (.08)	0.81 (.08)
Drop Shift	0.80 (.08)	0.80 (.08)	0.80 (.10)	0.77 (.09)
Increase Line Speed	0.84 (.06)	0.72 (.05)	0.84 (.06)	0.71 (.06)
Decrease Line Speed	0.74 (.06)	0.70 (.05)	0.78 (.07)	0.75 (.06)
A to B	0.84 (.05)	0.90 (.05)	0.94 (.06)	0.97 (.06)
A to C	0.77 (.01)	0.74 (.01)	0.84 (.02)	0.79 (.02)
B to A	0.38 (.04)	0.25 (.04)	–0.05 (.05)	–0.02 (.05)
B to C	0.46 (.02)	0.47 (.02)	0.77 (.03)	0.62 (.02)
B to D	0.22 (.03)	0.32 (.03)	0.58 (.04)	0.54 (.03)
C to A	0.32 (.01)	0.35 (.01)	0.18 (.02)	0.26 (.02)
C to B	0.27 (.02)	0.29 (.02)	0.11 (.02)	0.29 (.02)
C to D (1 shift)	0.42 (.27)	0.69 (.25)	0.22 (.31)	0.86 (.28)
C to D (2 shifts)	0.24 (.05)	0.33 (.05)	–0.43 (.06)	–0.14 (.06)
D to B	0.44 (.04)	0.27 (.04)	0.38 (.05)	0.40 (.05)
D to C (1 shift)	0.76 (.26)	0.43 (.24)	0.44 (.30)	–0.23 (.27)
D to C (2 shifts)	0.95 (.05)	0.82 (.05)	1.34 (.06)	1.12 (.05)

A: Short week, due to inventory adjustment or model changeover
B: Short week, due to holiday or supply disruption
C: Regular week
D: Overtime

to a quarter.

The rest of the results are remarkable for their asymmetry. For example, the persistence of an output change resulting from a move from A to B is in most cases over ninety percent, while the persistence of the output change involved in a move from B to A is zero after a quarter. When two shifts are working, the movement from a regular week to overtime hours has no persistence over the quarter, while a movement from overtime hours to a regular week actually forecasts an even larger output drop one quarter ahead.

We believe that these changes are much more short–term in nature than the line speed and shift changes, and that movements between states A through D involve much lower adjustment costs. Since they are shorter–term, higher frequency events, they are involved in the complex short run dynamics of the plant's evolution. The plant transits through a complex series of state changes as shocks of various transitoriness and predictabilities hit the plant.

The various production margins appear to be quite differently dynamically. We do not see how it is possible to construct a theory of these differences without assuming that (i) managers have forecasts of the persistence of changes in desired quantity which vary over time, or (ii) the decisions themselves change the future production dynamics of the plant. While our estimates do not let us distinguish between the purely statistical versus the endogenous capability theories, we believe that both are substantially true. To add a shift or to change line speeds is, we believe, to affect one's own future decisions. Sensible managers will do that only if they think that there is some persistence to the underlying change that shifted desired output.

We draw several conclusions from the dynamic analysis. First, the knowledge of how a given change in output was attained has significant predictive power for the future path of output. This point goes back to the issue we raised in the introduction concerning the information contained in managers' decisions. Not only does the manager's method for realizing a given change in production contain information about his predictions about the future, but it also affects the path of output in the future. Second, there are three relatively "stable" states in which the plant may be. The most stable state is one where hours are around forty hours a week. Two other states that display a surprising amount of stability are the voluntary short–day state and the overtime state. Both states involve large departures of production from the forty hours a week rate. Finally, movement between the states seems to follow a process that is more complicated than first–order markov.

3. Implications for Aggregation

Finally, let us consider the implications of the behavior that we have uncovered for aggregate dynamic studies. We do this by presenting three simple examples that illuminate which types of phenomena at the plant level are, and are not, revealed in the aggregate data.

Example 1

Consider the following pair of questions: (1) How often and by how much does output deviate from "normal" capacity? and (2) When output adjusts, does it follow a smooth adjustment path or does adjustment occur in lumps? These questions are of interest because they can reveal the form of the production and adjustment cost functions that a plant faces. The form of those functions is a topic of considerable debate because they determine how shocks are propogated through the economy. We will answer each of these question in two ways, first using a pooled sample of the plant level data and second using an aggregate of our data. To answer the first question, we define normal capacity as the quantity of output the plant would produce if it ran the plant for forty hours at the given line speed and the mode of the number of shifts (equal to two for most plants). The output variable is adjusted to exclude variations that are due to holidays and model changeovers. For the aggregate answer, we aggregate normal capacity and production across plants. Figure 7 shows histograms of the ratios of output to normal capacity. The top graph shows the frequency distribution of deviations for the pooled plant level data, while the bottom graph shows the distribution for the aggregate data. The aggregate distribution looks nothing like the plant–level distribution. For example, the plant–level data shows that output is within five percent of normal capacity over sixty percent of the time, while the aggregate data show that output is within five percent of normal capacity only thirty percent of the time. On the other hand, the plant level data shows that when output does deviate from normal capacity, it deviates by twenty percent or more. Almost a quarter of the time, output is fifty percent or more away from normal capacity. The distribution of the aggregate data is more compressed. Thus, infrequent but large deviations at the plant level appear as frequent, but small deviations at the aggregate level.

Moving to the question on output adjustments, we calculated percentage changes in posted output from week to week, excluding changes involving holidays and model changeovers. We calculated these changes for the plant level data and for aggregate output.[7] Figure 8 shows a histogram of output changes. The first two graphs, which show the pooled plant–level data, display interesting patterns. The vast majority of data points on the tall spike at zero are identically equal to zero, implying that most weeks there is no change in posted output. To illuminate the distribution of points other than

[7]The denominator in the percentage change is line speed (t–1) x number of shifts (t–1) x 40. We could not use posted output in period t–1 as the denominator because it was frequently equal to zero.

Figure 7

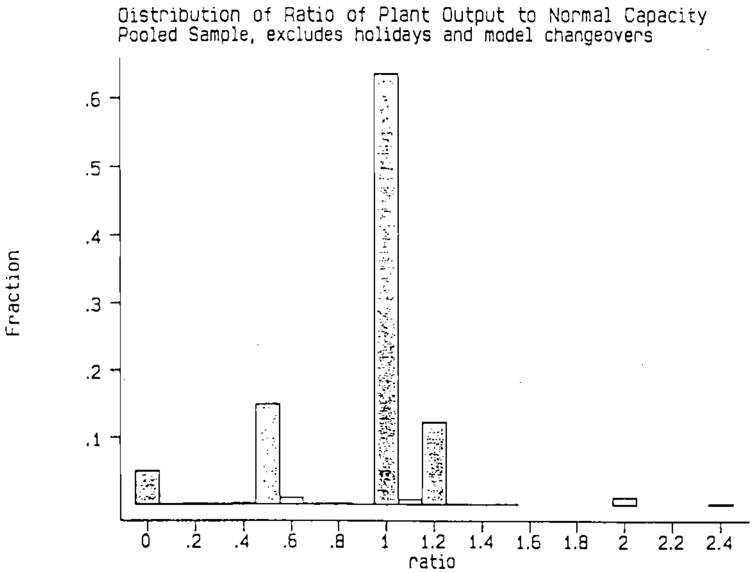

Oistribution of Ratio of Plant Output to Normal Capacity
Pooled Sample, excludes holidays and model changeovers
Fraction
.6
.5
.4
.3
.2
.1
0
.2
.4
.6
.8
1
1.2
1.4
1.6
1.8
2
2.2
2.4
ratio

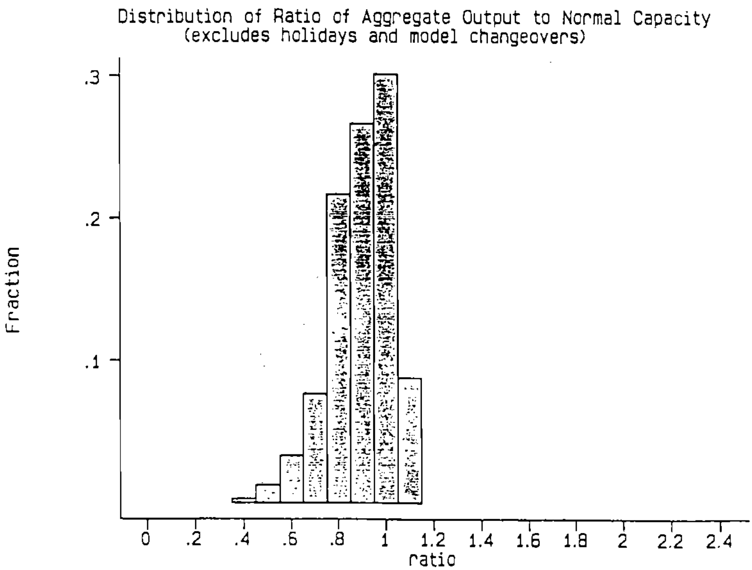

Distribution of Ratio of Aggregate Output to Normal Capacity
(excludes holidays and model changeovers)
Fraction
.3
.2
.1
0
.2
.4
.6
.8
1
1.2
1.4
1.6
1.8
2
2.2
2.4
ratio

Figure 8
Percentage Changes in Weekly Output
(holidays and model changeovers excluded)

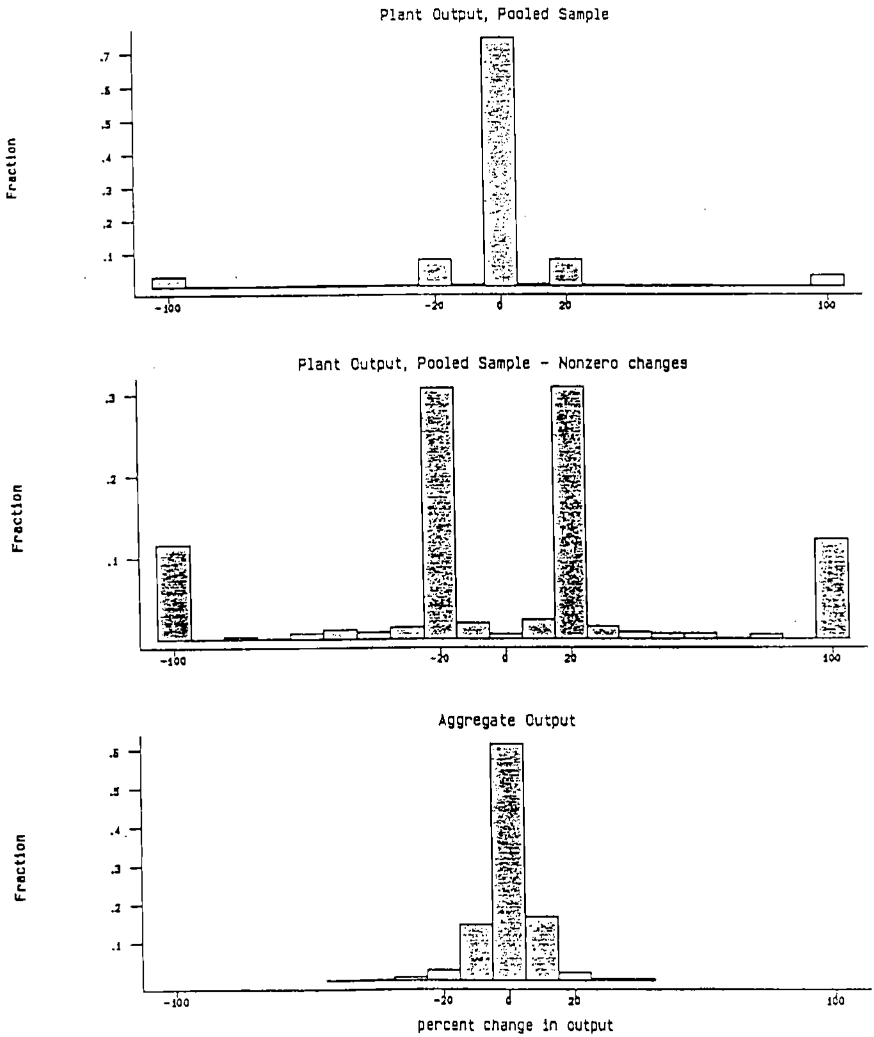

at zero, the middle graph shows the histogram of the nonzero changes in output. Most changes are approximately twenty percent, but close to a quarter of the output changes are 100 percent in either direction.[8] In contrast, the bottom graph shows that most of the changes in the aggregate data are very small, with the majority being much less than twenty percent in absolute value. The average (absolute) nonzero change is 43 percent for the plant data and only 5.7 percent for the aggregate data. This calculation makes the same point for output that Hamermesh (1989) made for employment: plant level dynamics suggest nonconvex costs of adjustment while aggregate dynamics suggest convex costs of adjustment. Thus, adjustment cost parameters estimated on aggregate data do not bear a direct relationship to the adjustment cost parameters for the individual plant. An implication is that the high convex cost of adjusting production found by Blanchard (1983) and Ramey (1991) for the automobile industry is not reflective of plant level costs.

Example 2

The histograms in figures 7 and 8 clearly show important plant–level events which are not reflected in aggregate data. To put this in perspective, we undertake a simple statistical investigation of aggregate and idiosyncratic shocks to plant posted output in an analysis of variance framework. This will reveal the extent to which output is correlated across plants, which is a distinct question from the amount of information lost in aggregation.

The plants are of different capacities; line speed can be as low as 15 in a luxury or specialty car plant or as high as 100 in a compact car plant. This will likely make the

[8]There were a few percentage changes that were greater than 100 in absolute value in the plant level data, but these were excluded from the histogram because they amounted to less than 0.1 percent of the nonzero changes.

plant–level data heteroskedastic. For the analysis of variance, then, we divide posted output by planned, one–shift capacity, which is line speed times 40 hours. Our dependent variable at plant i in week t is R_{it}, where:

$$R_{it} = Q_{it} / (LS_{it} \times 40).$$

The denominator measures how many units a plant would make at its line speed during one shift regular operations, so that R_{it} measures percentage deviations from one–shift normal operations. (See equation (1).)

The simplest analysis of variance model is

$$R_{it} = u_t + \epsilon_{it}$$

where u_t and ϵ_{it} are independent with standard deviations σ_u and σ_ϵ, respectively. Of course, this "model" is merely an interpretation of the variances and covariances of all the plant–level outputs. It does serve, however, to offer a descriptive summary of output movements which is closely related to the question of whether plant level events are aggregate phenomena: the larger is σ_u relative to σ_ϵ, the more important the correlated (industry–wide) portion of the individual plant shock to output.

In our data, the estimated variances are $\sigma_u^2 = 860$ and $\sigma_\epsilon^2 = 4154$. These figures have two different implications. First, let τ_p be the fraction of the individual plant variance accounted for by the common factor u, i.e., $\tau_p = \sigma_u^2/(\sigma_u^2 + \sigma_\epsilon^2)$. In our data, this is 0.172, which corresponds to a 17.2% r–squared in a regression of plant–level R_{it} on time dummies.

Second, let τ_A be the fraction of the variance in average plant R_{it} (or in the aggregate) accounted for by the common shocks. The variance of the average R_{it} is

$$\sigma_A^2 = \text{var}\,(\sum_{i=1}^{I} R_{it}/I) = \sigma_u^2 + \sigma_\epsilon^2/I.$$

The fraction of this variance explained by the aggregate shocks, u, is:

$$\tau_A = \frac{\sigma_u^2}{\sigma_u^2 + \sigma_\epsilon^2/I}$$

In our data, I=50 and this fraction is .912.[9,10] Thus, 91.2% of the variance in aggregate (or average) capacity utilization is explained by the common shocks, even though they explain only 17% of the plant level variance. As you can see by examining the formulas for τ, this is simply a law of large numbers result; the plant errors drop out in the aggregate. An exactly analogous calculation leads to the fraction of aggregate output, Q_t, which is explained by the aggregate shock:[11]

$$\tau_{Q(t)} = \frac{\Sigma_i \Sigma_j LS_{it}\ LS_{jt}\ \sigma_u^2}{\Sigma_i \Sigma_j LS_{it}\ LS_{jt} \sigma_u^2 + \Sigma_i LS_{it}^2\ \sigma_\epsilon^2}$$

---

[9]Of course, some of the high covariance is seasonal. When we adjust the data for holidays and model changovers, estimated τ_A is 0.88. Thus, nonseasonal industry–wide shocks dominate nonseasonal movements in the aggregates.

[10]It is not clear that I should equal 50, because 50 plants were never open at the same time. The fraction is similar, though, for I = 40.

[11]This calculation assumes that line speeds may be taken to be exogenous in the run in which the shocks occur. This assumption is an innocent one for the week–to–week fluctuations we study. Line speed changes are typically either substantial (and infrequent) changes in plant process or changes in the type of product the plant produces. Below, we will show evidence on the limited supply substitutability (in the relevant run) of plants making vehicles in different segments.

In a hypothetical week in which all 50 plants were operating at their sample average line speed, this would be 89%, since $\Sigma_i \Sigma_j \, LS_i \, LS_j \; / \; \Sigma_i \, LS_i^2 = 40$. The same law of large numbers point is at work here, of course.

The simple statistical relationship among τ_A, τ_P, and τ_Q contains an economic phenomenon. Low τ_P is a familiar fact: plant level data are extremely noisy (e.g. see Griliches and Mairesse (1983), Clark and Hayes (1986)). High τ_A and τ_Q is another familiar fact: aggregate data are driven by the nonlocal, systematic (economy– or industry–wide) part of plant–level shocks. There is no tension between these two statements, as the law of large numbers works against the importance of plant–level shocks in the aggregate.

Our example, shows, however, that high τ_A or τ_Q is no evidence for the utility of aggregate data in drawing inferences about "representative plant" phenomena. We exhibit an industry in which the aggregate data are driven by the common shocks, and yet in which the lumpiness of the micro adjustments is completely washed out in the aggregates. The right condition for the utility of aggregate data in drawing behavioral inferences is high τ_P. Unfortunately, this condition is very rarely met.

Example 3

In this example, we will conduct a comparison of the histories of two plants and discuss the broader implications of that comparison. One of the points we will highlight is how shocks that are considered to be macroeconomic shocks can have very different consequences for plants in a narrowly defined industry.

Figure 9 shows the behavior of adjusted output and short–run capacity for Ford St. Louis, which produced intermediate and full sized cars and GM's Lordstown, which produced compact and subcompact cars. Output is adjusted by removing the variance

Figure 9

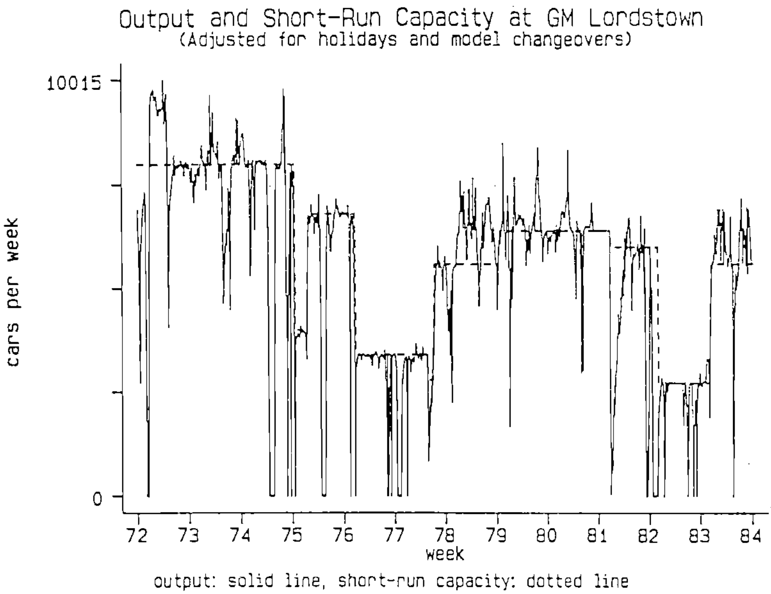
Output and Short-Run Capacity at GM Lordstown
(Adjusted for holidays and model changeovers)
10015
0
cars per week
72
73
74
75
76
77
78
79
80
81
82
83
84
week
output: solid line, short-run capacity: dotted line

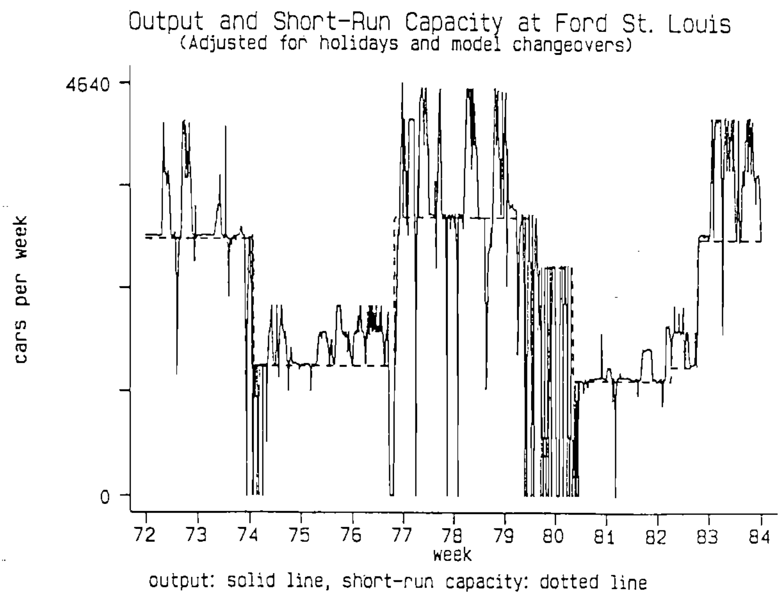
Output and Short-Run Capacity at Ford St. Louis
(Adjusted for holidays and model changeovers)
4640
0
cars per week
72
73
74
75
76
77
78
79
80
81
82
83
84
week
output: solid line, short-run capacity: dotted line

caused by holidays and model changeovers; short–run capacity is line speed times the number of shifts times 40 hours. During 1972 and 1973, output at both St. Louis and Lordstown was generally above short–run capacity, with heavy use of overtime; the deviations below capacity tended to be due to parts shortages or strikes. The first oil shock hit at the end of 1973, shifting demand from large cars to small cars. At the beginning of 1974, St. Louis began closing down for periods of inventory adjustment, and followed with the elimination of its second shift in February. Lordstown, on the other hand, maintained high production (except for strikes in August 1974) until November 1974, when it began to close down for inventory adjustment. On the other hand, in the first half of 1975 when the economy begain recovering and queues at gasoline stations dwindled, Lordstown reduced its line speed while St. Louis began scheduling overtime. In 1976, St. Louis increased its output further by adding a second shift and scheduling heavy overtime, while Lordstown eliminated its second shift. Thus, output at Lordstown declined during the recovery from the recession. Both plants operated at high levels in 1978 and 1979. During 1980 in the aftermath of the second oil shock, however, St. Louis cut its output substantially, while Lordstown did not undertake significant cuts until the beginning of 1982. Both plants kept output low through 1982, and then increased it in 1983.

These graphs show the important differences across plants producing different size cars. They also give insight into how oil shocks affect the economy. In the case of automobiles, the oil shocks served to shift demand back and forth across size classes, leaving firms with mismatched capacities. After the first oil shock, a number of plants were idled for long periods of time in order to convert them to small car production. When these plants came on line, though, demand had already shifted back to large cars, and the plants were significantly underused.

To see whether this "capacity mismatch" story was systematic, we undertook another analysis of variance of R_{it} (defined in example 2). Each plant's products were assigned to one of five automotive market segments, s(it).[12] The analysis of variance was extended to include segment effects and segment*week effects:

$$R_{it} = u_t + v_s + w_{st} + \epsilon_{it}$$

where the four error components are once again assumed, for descriptive purposes, to be orthogonal. We use this model for two purposes. First, as a descriptive statistical matter, can the segment*week effects (w_{st}) be excluded? They cannot: The F(2500,24013) is 1.91, which is highly significant.

The second purpose of the model is to see the generality of the point about the price of fuel and demand. The mean predicted segment*week effects from the ANOVA are plotted in Figure 10.[13] This shows time on the horizontal axis and the estimated value of the average of w_{st} for the two small–car segments and for the three large–car segments on the vertical axis. Note that the segments move both together and apart. The fuel price crises mark periods of capacity shortage in the subcompact and compact segments, and of severe excess capacity elsewhere in the industry. Other time periods, for example that in between the two fuel–price crises, show the reverse tendency. Thus the simple demand story we told about the two specific plants is general to the entire industry.

[12]The five segments are subcompact, compact, midsize, fullsize, and luxury. Some plants changed the segments they produced in. In the analysis of variance, they are classified at each time according to their current products.

[13]These are not calculated as the coefficients one would obtain for segment*week dummies in a regression. Instead, they are the predicted means in each segment and each week. Thus, they need not sum to zero.

Figure 10: Small and Large Car Deviations by Week
(four-week moving average)

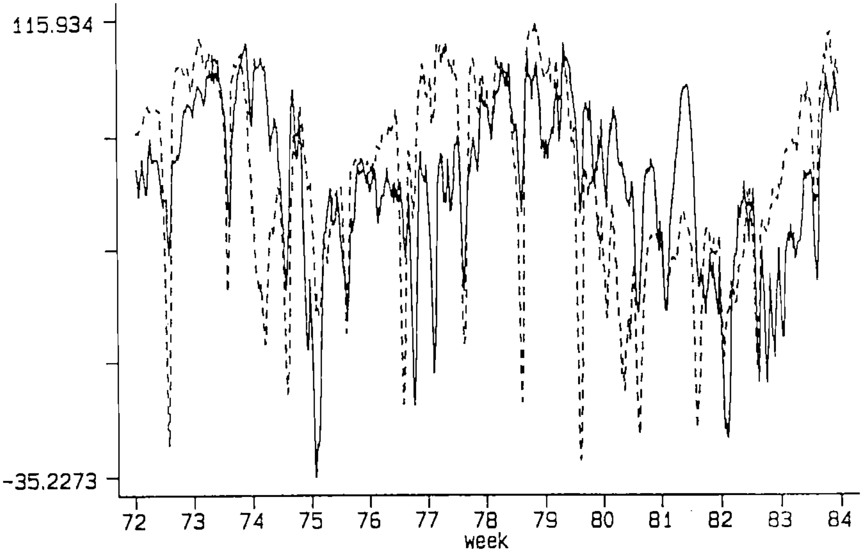

small cars: solid line, large cars: dotted line

What is going on here is the interplay between firms' fixed competitive capabilities and highly variable demand. Firms must commit not only to the aggregate capacity they have, but also to their product line and, to a very considerable extent, to the capacity for each product. In a run of a year or more, capacity is largely fixed, and not very fungible between products in different market segments. On the demand side, rapid changes in the price of fuel in our period altered the composition of the demand for automobiles as well as its level. Thus firms found themselves with perennial mismatches between the composition of their capacity and consumers' desired purchases. The resulting loss in productivity compared to a more predictable or stable demand environment is considerable.[14]

The three examples have shown how answers provided by aggregate data can lead to incomplete or different conclusions from those provided by the plant level data. Some of our discussion is related to an important new literature arguing that the smooth behavior of some aggregates is the result of very lumpy behavior at the individual level. As discussed above, Hamermesh (1989) has shown using plant level data that employment fluctuations are very erratic, and are best described by nonconvex costs of adjustment. Further, Blinder (1981), Caplin (1987), Caballero (1991) and others have investigated the implications for aggregate data of individuals who follow S–s type of rules. Several of our conclusions are similar to their conclusions.

In the previous section, we demonstrated that the various margins for increasing or decreasing output in the short run are different predictors of the future. We have not distinguished between the two distinct theories of: (i) heterogeneous, mostly correct,

[14]This conclusion is similar to the point made by Ramey and Ramey (1991) that planning mistakes can have negative consequences for economic performance.

expectations about future economic circumstances at different plants or (ii) endogenous state dependence resulting from the plant level choices. What we want to observe here is that both theories can be problematic for analysis at the aggregate level.

Let us begin with theory (ii), and our extension of Hamermesh's story. Suppose that, in the short run, the supply curve for output is smoothly rising not because any plant has a smoothly rising SRMC, but because there is a smooth distribution function of the heterogeneous costs of a lumpy change. This is Hamermesh's story of employment adjustment, and in our data it corresponds to the adjustments associated with shutdowns, shift changes or line speed changes. What is wrong with doing analysis on the aggregate supply curve in these circumstances? Hamermesh's point is that the resulting inferences reveal nothing about the plants' SRMC if they are interpreted in a representative plant framework.

The story quickly gets worse with multiple margins. In particular, consider our plants with dynamically distinct margins. Suppose that in some period, aggregate production has increased. We want to predict future movements of aggregate quantity and shipments based on this fact. Yet the same increase in quantity could come from (a) an added shift at one–fifth of the plants or (b) an added day of overtime at all plants or (c) anything in between. Estimation a model of aggregate quantity dynamics without knowing which of these events occurred can lead to seriously biased estimates of the parameters of interest. There are no stable aggregate dynamics in such a system; the analyst has gone beyond mislabelling the macro phenomenon as a micro one to estimating a relationship that has no relationship to the underlying dynamic production technology at either the plant or industry level.

The problem is not all that different with expectational heterogeneity. Suppose now a variety of plants all want to increase quantity, but with different expectations about the permanence of the increased run–rates. The event in which one plant thinks its demand is up by 300% is not at all the same as the event in which each of three plants thinks its demand is up by 100%. The shock to aggregate desired quantity is exactly the same (desired quantity after the shock is three times that before) but the visible resulting event is not at all the same.

4. Conclusions

Our investigation of weekly production dynamics at the plant level has uncovered characteristics of output volatility that are not apparent in aggregate data. Furthermore, by analyzing operational aspects of production we have been able to provide a more complete picture of output fluctuations.

Our most important overall finding is that adjusting production is a more complicated process than simply "changing Q." Of the multiple margins used by the managers of an automobile assembly plant, varying regular hours by shutting the plant down for a week is one of the most important. Second most important is adding or dropping a shift. These margins are very different dynamically. How managers chose to adjust output contains information about the permanence of the output change. Probably because the different output adjustment margins involve different amounts and lumpiness and irrevesibility, the recent history of output alone (as opposed to the margins used to adjust it) is an incomplete summary of the plant's state. The transitions of plants through high and low output states is not a process that can be readily characterized by standard time series methods.

We also examined the relationship between plant–level events and aggregate ones. A key analytical distinction is that (i) the usefulness of aggregate data for drawing conclusions about costs, demand, or expectations is not the same as (ii) dominance of aggregate data by shocks that are common across all plants. In an industry where (ii) is quite true, we found two kinds of economically important departures from (i). Aggregation hid heterogeneous shocks, as when the same change in the price of fuel shifted some plants' demand up, others' down. Heterogeneous responses are lost as well, as when some managers choose to make a lumpy adjustment, others to wait or temporize. The predictive importance of agggregate shocks is logically unrelated to the analytical value of aggregate data.

Our analysis has been descriptive, but it serves to illustrate the value of operations data in economic production analysis. By analyzing not only the usual output data, but also the "black box" data on the managers decisions, we feel that we have been able to paint a more accurate picture of output fluctuations.

References

Abernathy, William J., Kim B. Clark, and Alan M. Kantrow, *Industrial Renaissance: producing a Competitive Future for America*, New York: Basic Books, 1983.

Aizcorbe, Ana, "Procycical Labor Productivity, Increasing Returns to Labor and Labor Hoarding in U.S. Auto Assembly Plant Employment," Bureau of Labor Statistics Working Paper No. 203, March 1990.

Automotive News, various issues from 1972 to 1983.

Bartlesman, Eric and Phoebus Dhrymes, "Productivity Dynamics: U.S. Manufacturing Plants, 1972–1986," February 1992 Center for Economic Studies working paper, U.S. Bureau of the Census.

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

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

Bresnahan, Timothy, "Departures from Marginal–Cost Pricing in the American Automobile Industry," *Journal of Econometrics* 17 (November 1981): 201–227.

Caballero, Ricardo, "Durable Goods: An Explanation for their Slow Adjustment," Columbia University mimeo, 1990a.

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

Chew, Bruce, "No–Nonsense Guide to Measuring Productivity," *Harvard Business Review 66* (January/February 1988): 110–118.

Clark, Kim B. and Robert H. Hayes, "Why Some Factories are More Productive than Others," *Harvard Business Review* (September/October 1986): 66–66.

Cooper, Russell and John Haltiwanger, "The Aggregate Implications of Machine Replacement: Theory and Evidence," May 1991 manuscript.

Davis, Steve and John Haltiwanger, "Gross Job Creation and Destruction: Microeconomic Evidence and Macroeconomic Implications," in *NBER Macroeconomics Annual* 1989.

Griliches, Zvi and J. Mairesse, "Comparing Productivity Growth: An Exploration of French and U.S. Industrial and Firms Data," *European Economic Review* 21 (March/April 1983): 89–119.

Hamermesh, Daniel "Labor Demand and the Structure of Adjustment Costs," *American*

Economic Review 79 (September 1989): 674–689.

Ramey, Valerie, "Nonconvex Costs and the Behavior of Inventories," *Journal of Political Economy* 99 (April 1991): 306–334.

Ramey, Garey and Valerie Ramey, "Technology Commitment and the Cost of Economic Fluctuations, " NBER working paper 3755, June 1991.

Rosenberg, Nathan, *Inside the Black Box, Technology and Economics*, Cambridge: Cambridge University Press, 1982.

Wards' Automotive World, various issues.

Wards' Automotive Yearbook, various years.