



OXFORD JOURNALS
OXFORD UNIVERSITY PRESS

Output Fluctuations at the Plant Level

Author(s): Timothy F. Bresnahan and Valerie A. Ramey

Source: *The Quarterly Journal of Economics*, Vol. 109, No. 3 (Aug., 1994), pp. 593-624

Published by: Oxford University Press

Stable URL: <https://www.jstor.org/stable/2118415>

Accessed: 24-04-2019 00:28 UTC

JSTOR is a not-for-profit service that helps scholars, researchers, and students discover, use, and build upon a wide range of content in a trusted digital archive. We use information technology and tools to increase productivity and facilitate new forms of scholarship. For more information about JSTOR, please contact support@jstor.org.

Your use of the JSTOR archive indicates your acceptance of the Terms & Conditions of Use, available at <https://about.jstor.org/terms>



JSTOR

Oxford University Press is collaborating with JSTOR to digitize, preserve and extend access to *The Quarterly Journal of Economics*

OUTPUT FLUCTUATIONS AT THE PLANT LEVEL*

TIMOTHY F. BRESNAHAN AND VALERIE A. RAMEY

This paper examines the short-run dynamics of manufacturing costs by detailing how plants in the U. S. automobile industry change output. Weekly data show a variety of margins on which firms adjust production. These margins, which are distinct from the usual factor demand choices, differ in their lumpiness, their adjustment costs, and their variable costs. The existence of these margins explains several empirical puzzles of output fluctuations. Using a theory of the short-run dynamic cost function, we are able to infer some of the characteristics of the underlying cost function from the dynamic behavior of the different margins.

The structure of firm-level output and adjustment costs is an important element of many economic theories. In business cycle theories, the shape of the cost function determines whether shocks to the economy are magnified or dampened and how their dynamic properties are transferred to output. In labor economics, both judgments about the effects of government labor market policies and estimates of labor productivity often depend on assumptions about the cost of adding workers relative to the cost of increasing average hours per worker. In industrial organization, estimates of the level and cyclical behavior of markups depend critically on the shape of the cost function.

For the last three decades the standard cost function specification has been convex in output as well as in the change in output or in the relevant factors of production (e.g., Holt, Modigliani, Muth, and Simon [1960]). This type of cost function is associated with an optimization problem that is easy to solve, and when coupled with a representative firm framework, has time-series implications that are roughly consistent with aggregate data (e.g., Sargent [1978]). There have, however, been many critiques of this type of cost function. For example, authors such as Alchian [1959] and Maloney and McCormick [1983] have argued that the standard cost function specification is lacking because it does not distinguish

*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, Garey Ramey, Glenn Sueyoshi, three anonymous referees, 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.

© 1994 by the President and Fellows of Harvard College and the Massachusetts Institute of Technology.

The Quarterly Journal of Economics, August 1994

between the *rate* of output and the *volume* of output. Distinguishing between the rate and volume of output is important for understanding issues such as scale economies. Furthermore, the convexity assumption has been criticized in several contexts. Rothschild [1971] argued that the economic justification for nonconvex adjustment costs is at least as good as that for convex adjustment costs. Moreover, there is a growing body of empirical work suggesting nonconvexities in the cost function. First, Hamermesh [1989] showed that at the plant level, labor force adjustment proceeds in large jumps, contrary to the convex adjustment cost model. Hamermesh concluded that the smooth adjustment observed at the aggregate level is the result of the aggregation of the plant-level nonlinear relationship. Second, Ramey's [1991] work on inventories offers additional empirical evidence against the convex cost model. She showed that in a variety of industries, the behavior of production and inventories relative to sales is consistent with nonconvex costs. Third, Bertola and Caballero [1990] and Caballero and Engel [1992] have shown that in many cases modeling macroeconomic data as the aggregation of heterogeneous agents facing nonconvexities is superior to the standard convex, representative agent framework.

This paper seeks to shed light on the nature of plant-level costs by detailing exactly how plants in the U. S. automobile industry adjust production. Changes in production in this industry are quite lumpy at the plant level. We show that the lumpiness is caused by exploitation of nonconvex operating margins. Intermittent production, with plant shutdowns and restarts at a high production level, is one significant mechanism for responding to low demand. More systematically, we examine the statistical contribution of all available operating margins to the variation of plant-level production over time. Our findings can explain both the excess volatility of production over sales and much of the short-run variation in output relative to hours.

Finally, we examine the short-run dynamics of margin choice and output determination. We estimate the duration of, and transitions between, states defined by margins of adjustment. Managers' choice among margins is also examined as a predictor of output persistence. Using a simple theory of cost and margin choice, we rank the margins on the basis of their relative static marginal costs and adjustment costs. The results add insight into the sources and consequences of nonconvexities and partial irreversibilities in production choices.

The data set used is particularly useful for studying these questions. 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 mostly 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 plants, at so high a frequency, in such detail. 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) the publicly available data for the automobile industry are very high quality. It must be noted, though, that the automobile industry is not representative of the entire economy. We believe, however, that a detailed analysis of the production dynamics in the automobile industry will provide insight into the nature of costs and output fluctuations in general. 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]). We see our work as part of a growing body of research that uses plant-level data to gain new insights into aggregate phenomena [Leonard 1987; Dunne, Roberts, and Samuelson 1989; Aizcorbe 1992; Davis and Haltiwanger 1992].

The paper proceeds as follows. Section I discusses the decision problem of a plant manager who can use various margins to vary the volume of production. Section II will present a more detailed description of the data and how it was collected. Section III characterizes each of the margins and documents their importance for the mean and variance of output. Section IV analyzes the dynamic characteristics of the margins, and their implications for the underlying cost structure. Section V concludes.

I. MULTIPLE MARGINS AND COSTS: THEORY

In this section we compare a standard cost function specification with one that we believe is more relevant for automobile assembly, and possibly for other industries as well. Using this

function, we discuss how the dynamic characteristics of the use of different margins can reveal their underlying costs.

An example of a standard output cost function is given as follows:

$$(1) \quad C(t) = C(Q, \dot{Q}, t), \text{ with } C_Q > 0, C_{QQ} > 0, C_{\dot{Q}} > 0, C_{\dot{Q}\dot{Q}} > 0,$$

where Q is the volume of production during time interval t , \dot{Q} is the change in the volume of production, and C_n denotes the partial derivative of C with respect to argument n . Typically, the unit of time is a month or a quarter, and there is no distinction made between the rate and volume of production. The function is increasing and convex in both arguments. If the problem is set up in discrete time, and if the function is quadratic, then typically optimal output will follow an ARMA process. In general, the time path of output will be smoother than its driving process (e.g., sales). Furthermore, changes in the volume of output will be gradual.

Consider the following alternative cost function, which captures a set of static and dynamic nonconvexities linked to a richer and more descriptive set of decision variables:

$$(2) \quad C(t) = F(l, h, s, t) + \alpha_l I(l) + \alpha_h I(h) + \alpha_s I(s),$$

where $I(n) = 1$ if $\dot{n} \neq 0$ and $I(n) = 0$ otherwise, for $n = l, h$, and s . l is the instantaneous rate of output, or cars per hour ("the line speed"), h is the hours worked by each shift, s is the number of shifts, and the I 's are indicator functions for changes in l, h , and s . The α 's are nonnegative parameters. The *volume* of output over a time period T is given by the identity:

$$Q(T) = \int_0^T l(t)h(t)s(t)dt.$$

We assume that plant managers minimize the expected present discounted value of costs, given by

$$(3) \quad V(\tau) = E_0 \int_0^\tau e^{-\rho t} C(t),$$

where E_0 is the expectation conditional on information in period 0, τ is the end of the model year, and ρ is the interest rate. Costs are minimized subject to targets dictated by the firm. These targets might be monthly or quarterly output volumes, or bands on inventory-sales ratios.

The cost function given in equation (2) may have several

properties that lead to nonconvexities. First, the adjustment cost specification implies a fixed cost of changing a margin. Second, some levels of l , h , and s may be impossible or prohibitively expensive, at least in the short run. For example, the plant may be designed, and its labor force trained, for a particular l . Similarly, operation at 1.42 shifts per day may go against prevailing labor market norms. Thus, F may not be monotonically increasing in the values of l , h , and s . Furthermore, if the h margin can take the value zero, and if adjustment costs are not too great, then the firm may use *intermittent production* [Maloney and McCormick 1983]. That is, if the plant manager wants to achieve a lower volume, rather than changing the *rate* of production, he may change the *period* of production by shutting down the plant for part of the time.

When defined, the second derivative of F with respect to each argument may be negative, zero, or positive. Aizcorbe [1992] gives evidence that the second derivative with respect to l , F_{ll} , is slightly negative in automobile assembly over the relevant range. For some of the other possible margins, such as hours per shift, there are arguments suggesting that the second derivative may be positive.

The richness of the multiple margin framework is apparent once we consider the dynamic use of each margin. In general, different margins will have different adjustment costs and different static marginal costs. As an illustration, suppose that the target volume of production changes, and the plant manager must decide whether to use margin h or s to change the volume of production. Assume that margin h is the high (static) marginal cost-low adjustment cost margin relative to margin s . That is, $F_h > F_s$, but $\alpha_h < \alpha_s$. Assuming that the plant manager minimizes the expected present discounted value of costs, the choice of margin will depend on the expected persistence of the change in target as well as the level of uncertainty associated with target changes. It is easy to see that for known temporary changes in the target, it is optimal to use the high marginal cost-low adjustment cost margin, whereas for more persistent changes in the target, it is optimal to use the low marginal cost-high adjustment cost margin. In the presence of uncertainty, the high marginal cost-low adjustment cost margin may be used more frequently since its use preserves the option value of the other margin. If the desired change in output volume is sufficiently high, then the plant manager will use both margins at the same time.

One would expect the following observations to hold if we

observed a plant over a long period of time: (1) the average duration of use of the high marginal cost-low adjustment cost margin should be less than for the low marginal cost-high adjustment cost margin; (2) on average, the change in output volume achieved by using the high marginal cost-low adjustment cost margin should be less persistent than that achieved by using the low marginal cost-high adjustment cost margin. The first observation follows directly from the example above. There are two forces leading to the second observation. First, the choice of margin reveals the manager's expectations about the persistence of the change in the output target. Second, when the manager changes the high adjustment cost margin, he is locking in an output volume that is costly to reverse. Thus, even if the manager has overestimated the persistence of the target change, the actual volume of output may stay higher for a longer period of time.

These observations concerning the relationship between the cost structure of the different margins and the observed use of the margin will be useful for our empirical work. Using these observations, we can test whether the dynamic properties of different margins are consistent with a particular cost structure. In later sections we shall characterize the dynamic properties of the different margins of output variation in the automobile industry and use them to shed light on the underlying cost structure.

II. 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*, p. 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 layoffs will be about 92,000 by the end of January. . . .

All of the information on regular hours, overtime hours, and shutdowns was taken from the *Automotive News* production articles. In a few cases we were able to detect unreported shutdowns because the production tables showed that the production of a particular model was zero for that week. The *Automotive News* production articles reported most of the changes in line speeds and shifts. Starting in the 1977 model year, *Wards' Automotive Yearbook* published the line speeds and shifts for each factory at model year start-up. For the earlier years we were able to obtain some information from *Wards' Automotive World* or directly from GM and Ford. This information was used to augment the *Automotive News* information when it was not complete. When possible, we used actual production data to infer the magnitude and date of a change.

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, and (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 [1993], 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 labeled "inventory adjustment" represents the times the company shut the plant down in order to adjust inventories. The "supply disruption" category contains shutdowns due to strikes, both on-site and off-site, 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 70 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. The data from the six matched plants show that these deviations are not significant.

In all, we study 50 assembly lines. When a plant had two lines, we treated each line as a separate plant. In three of the four cases, the second line of the plant produced specialty luxury cars, such as the Corvette or Toronado, so that it is unlikely that workers substituted between lines. 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 the 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.

III. THE IMPORTANCE OF MULTIPLE MARGINS OF OUTPUT VARIATION

A. Sources of Output Variation

We begin by identifying the margins of output variation and characterizing their apparent costs.¹ In our analysis we shall 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:

$$(4) \quad Q_{it} = (RH_{it} + OH_{it}) \times (LS_{it} - \epsilon_{it}) \times SH_{it},$$

1. By "apparent costs" we mean those costs directly spelled out in government regulations, labor contracts, or directly implied by the production technology. Apparent costs can differ from true costs because of considerations other than those included in apparent costs or if the wages set by regulations and contracts are not allocative.

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.

The decomposition in equation (4) 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. 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 or the output rate. It is the latter variables that the managers manipulate directly; it is the effect on the labor variables that generates much of the cost.

We should also note that we are excluding an additional margin in our analysis: permanent plant shutdowns or openings of new plants. We exclude this margin from the present study for three reasons. First, we interpret this margin as a *firm* production margin, rather than a *plant* production margin. Second, the focus of the present analysis is on short-run, rather than long-run output fluctuations. Permanent shutdowns are not used to achieve short-run fluctuations in output. Third, preliminary analysis indicates that all nonstationarity in aggregate industry output comes from permanent plant shutdowns and new openings, rather than from changes in output at open plants. Thus, by focusing only on plants during the time they are open, we are able to study only the stationary fluctuations in output.

Some detail about the apparent costs associated with each of the margins can be obtained from the structure of government labor laws and the labor contracts in force. Consider first overtime hours. Overtime hours are usually varied by scheduling Saturday work or by adding an hour or two to each shift; eight hours on Saturday is by far the most common form of overtime. Labor laws mandate that workers be paid a 50 percent premium for hours in excess of 40 hours per week. Thus, the apparent marginal cost of overtime hours is significant.

Variations in regular hours imply different types of costs. Regular hours are usually varied by closing the plant for a day or a week. The cost savings from closing a plant for less than a week are small because workers with one or more years of service must be provided with short workweek benefits of 80 percent of their wage rate multiplied by the difference between the number of paid hours and 40 [Bureau of Labor Statistics 1976].² If the plant is closed for an entire week, the firm does not pay the workers. The closing does, however, have implications for unemployment insurance. When a plant is closed, two types of unemployment benefits are paid: state unemployment insurance (UI) and supplemental unemployment benefits (SUB). UI benefits generally last for 26 weeks, whereas the duration of SUB benefits is determined by the seniority of the workers. The SUB fund payment structure, when combined with the state UI, is designed to give an employee an amount equal to 95 percent of weekly straight-time pay after taxes, for up to one year. The SUB fund is financed by company contributions to trust funds, at rates of several cents per hour per employee. The cost to the companies of state unemployment insurance varies, and depends on the degree of experience rating.

Adding a shift entails the following types of costs. First, the premium pay for the second shift is 5 percent of day-shift rates, so the marginal cost per hour of work is slightly higher than for the first shift. Furthermore, the workers on the second shift tend to be more junior, which may have an impact on productivity. Second, adding a second shift entails adding more workers. For example, GM Fremont had 4970 hourly total workers spread between two shifts in the fall of 1981. When the second shift was eliminated in February 1982, approximately 1900 hourly workers were put on indefinite layoff. Thus, the second shift constituted 40 percent of the total number of hourly workers at the plant. Adding more workers increases the fixed costs, since many of the benefits provided by the company do not vary with the hours worked. Moreover, the cost of some benefits may increase with time since many of the benefits depend on seniority, so that the liabilities of the company to a worker increase as the worker's seniority

2. Presumably the union imposes these constraints on the firms because of the nature of state unemployment insurance benefits. Workers can only receive those benefits if they are not employed. Thus, the union prefers that the firms decrease total hours by decreasing the number employed rather than by decreasing hours per worker.

increases. On the other hand, eliminating a shift reverses some of these fixed costs. The second-shift workers who are laid off receive state unemployment insurance and supplemental unemployment benefits for a period of time that depends on the state laws and the workers' seniority.

Changing the line speed almost always involves changing the number of workers. The plant increases the line speed by reorganizing the assembly line and redefining jobs. When Lordstown increased its line speed from 70 to 80 cars on each of two shifts in July 1979, it added 185 workers. Furthermore, actual production typically remains below line speed for several weeks after a change, because it takes awhile for workers to learn the new layout [White 1971]. The ability to change the line speed varies across plants. For some plants the size of the paint facility is the binding constraint on the line speed.

The nature of these apparent costs implies a cost structure with several nonconvexities. To give an idea of how the cost function might look, Figure I shows a total labor cost curve and an average labor cost curve that incorporate some of the structure discussed above. In particular, the cost curves incorporate a 50 percent premium for overtime, a 5 percent shift differential, and the 80 percent rule for short-weeks. To highlight the effects of these rules, we assume that there is no fixed cost or adjustment cost to opening the plant or adding a shift.³ The hourly wage per shift is normalized at unity.

The cost function is clearly nonconvex in the range from 0 to 40 hours and again from 40 to 80 hours if a second shift is added. The steps at 0 and 40 hours are due to the 80 percent rule for short-weeks. Any time a shift works a positive number of hours, the plants must pay at least 80 percent of the wage times 40 hours. Thus, costs increase substantially when hours become nonzero. The only significant convexity comes from overtime hours. The average cost curve shows that the weekly hours associated with local minimum average labor costs are 40 hours and 80 hours. If we also included overhead labor and the rental cost of capital, average total costs would probably be lower at 80 hours than at 40 hours. In any case, these curves give a useful illustration of how these regulations can lead to cost curves that depart significantly from the usual textbook examples.

3. Adding fixed costs changes the shape of the cost function surprisingly little.

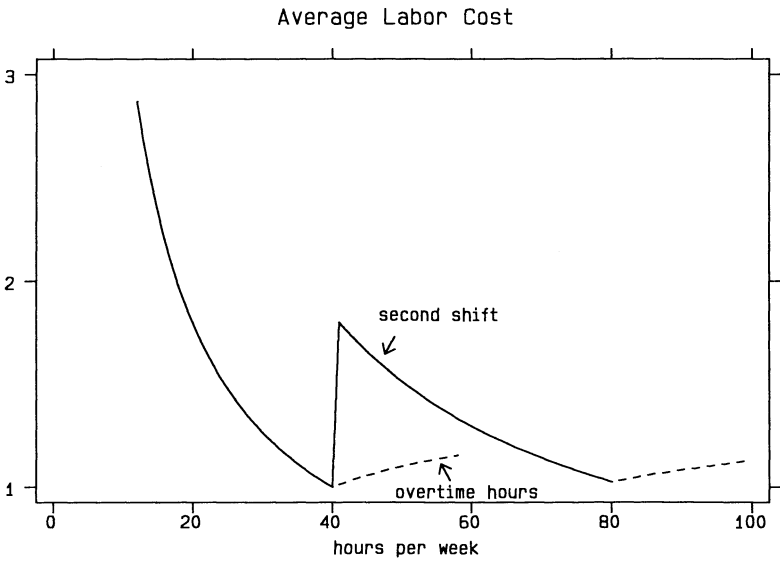
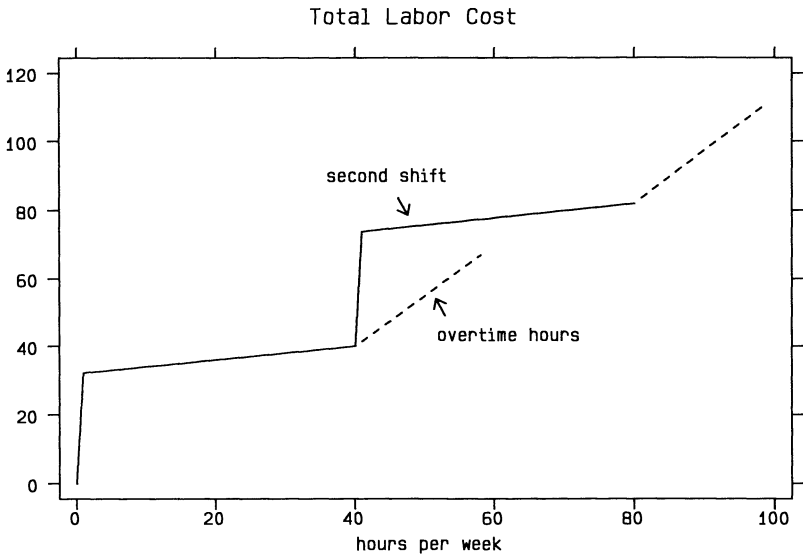


FIGURE I
The Structure of Apparent Costs at Automobile Assembly Plants

B. Frequencies and Variance Decompositions

We now analyze the importance of each margin in two ways. We first compute simple statistics that tell us how often each margin is manipulated. 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, for the reasons given in the last section.

Table I quantifies how often plants use each of the margins by showing the percentage of weeks during which each margin was manipulated. The first column shows the weighted average of the universe of all plants, the second column shows the calculation for the same sample, but with closures for holidays excluded, and the third column shows the weighted average 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 all plants. The results show that on average, each plant is shut down for at least one day in 25 percent of the weeks. Each plant is typically shut down for an entire week 13 percent of the time or almost seven weeks per year. Overtime hours (per worker per shift) of at least four hours per

TABLE I
FREQUENCY OF THE USE OF DIFFERENT MARGINS (PERCENT OF WEEKS USED)

	Weighted average of all plants	Weighted average of all plants, holidays excluded	Weighted average of six matched plants
Shutdown of at least 1 day	24.8	11.5	25.9
Shutdown of 1 week	12.9	10.3	13.0
4 or more overtime hours	14.3	14.3	32.2
Change in the number of shifts	0.6	0.6	0.7
Change in the line speed	0.8	0.8	0.9
Decomposition of shutdowns: percent of days closed by reason			
	Model changeovers	Holidays	Inventory adjustment
			Supply disruptions
Weighted average of all plants	32.5	34.7	25.5
			7.2

week are also frequent, occurring on average seven weeks per year.⁴ Thus, the use of overtime hours is as frequent as the use of week-long shutdowns. On the other hand, both changes in shifts and in line speeds are very infrequent, occurring substantially less than once per year.

The second column 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 nonholiday shutdowns are shutdowns of one week, and they occur 10 percent of the time.

The last column calculates the same frequencies (including holidays) for the six matched plants. The six matched plants are General Motors plants Bowling Green, Lordstown, Norwood, and St. Louis; and Ford plants St. Louis and Wixom. 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 line speed.⁵ The average frequencies for the subset of plants are similar to those for the universe of plants. The only exception is overtime hours, which are used twice as often. This increase in the calculated frequency of overtime is due, for the most part, to our ability to detect unreported overtime using the output data. We believe that the frequencies for overtime hours in the six matched plants are a more accurate estimate of overtime use for the entire sample as well.

The bottom of Table I 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, a quarter is due to inventory adjustment, and 7 percent is due to supply disruptions.

4. We chose four hours as the cutoff, because overtime less than four hours per week is used mostly for correcting defects in cars, rather than for increasing the volume of output.

5. These additional overtime hours and deviations from line speed are estimated. Whenever actual output exceeded posted output, the difference was ascribed to unreported overtime hours. Whenever actual output was less than posted output, the difference was attributed to deviations from line speed. Measurement error would occur if both unreported overtime hours and deviations from line speed occurred in the same week, or if the actual line speed was different from the reported line speed.

The question arises as to what percent of idle time is due to managers' direct manipulation of the volume of production. Clearly, holidays and supply disruptions can be considered exogenous changes in output volume, while the purpose of shutdowns for inventory adjustment is for the direct manipulation of volume. The case is less clear for model changeovers. Plants are shut down for model changeovers in order to retool the factory for the new model year, and thus are not directly used to manipulate the volume of production. On the other hand, we believe that many times the duration of the shutdown is used to manipulate the volume. During high demand years, model changeovers tend to be very quick, sometimes occurring over the weekend, whereas during low demand years, shutdowns for model changeover for slow-selling models appear to last considerably longer. Thus, we believe that some shutdowns for model changeover are actually shutdowns for inventory adjustment. It is difficult to distinguish individual spells, though, so we continue to separate the two categories.

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 20 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 (4) when ϵ is equal to zero. We also present results for the six matched plants, which contain data for actual production and estimates of ϵ . The results show that in most cases ϵ is not an important source of fluctuations in output.

The decomposition of the variance is not straightforward because, as equation (4) shows, posted output equals the product of the components, so the relationship is not linear. Furthermore, one cannot take logarithms because the hours components are frequently equal to zero. Therefore, we use the following method. First, for all calculations we exclude the effect of holidays, because

their contribution is not particularly interesting. For each margin we construct an artificial output measure for which the margin is held constant at some base level, and then compare this output measure to posted output. To be specific, to assess the impact of changes in regular hours due to inventory adjustment, we construct an output measure for which shutdowns due to inventory adjustment are eliminated, but for which changes in regular hours for other reasons, overtime hours, shifts, line speeds, and deviations from line speeds (in the case of the matched plants) are allowed to vary as they actually did. To construct artificial output during a shutdown, we use the number of shifts and the line speed in effect during the week before the plant shutdown; in almost every case, the plant opened up with the same number of shifts and the same line speed. Overtime hours are set to zero. We then compare that constructed output measure with posted output. We construct similar output series for the other margins. We multiplied each constructed series by a factor so that its mean was the same as the mean of posted output.

To assess the impact of a particular margin on the variance of output, we compare the variance of posted output with the variance of constructed output. The difference in the variance of posted output and constructed output as a percent of the variance of posted output gives an indication of the importance of each margin in the overall variance of output. The numbers do not add to 100 because of nonlinearities and covariance terms between the series on the different margins.

Table II shows the results of the calculations. The first column shows the results for the contribution of each margin to the

TABLE II
IMPORTANCE OF EACH MARGIN FOR THE VARIANCE OF OUTPUT
(PERCENT IMPACT ON MARGIN USE)

	All plants, weekly frequency	Six plants, weekly frequency	All plants, quarterly frequency
Inventory adjustment	28.6	17.1	20.6
Model changeover	33.5	20.2	21.4
Supply disruptions	7.5	18.0	3.0
Overtime hours	5.8	10.3	10.9
Shifts	25.1	30.1	40.2
Line speeds	10.5	8.3	19.5
Deviations from line speeds		2.4	

week-to-week variance of output in the universe of plants. Variations in regular hours for inventory adjustment and model changeover are the most important contributors to the week-to-week variation in output, followed by variations in the number of shifts. The number for inventory adjustment implies that, ignoring covariances, if the plant were not shut down for inventory adjustment, the variance of output would be 28.6 percent lower. The effect is even greater for model changeovers. Supply disruptions have a small impact. Line speed changes and overtime hours are substantially less important contributors to the weekly variance of output.

The second column shows the same calculations for the weighted average of the six matched plants. Here, shift changes are the most important contributors to weekly variance, while inventory adjustment and model changeover are somewhat less important. Supply disruptions have a greater impact on the variance in this case, equaling the effect of inventory adjustment. One reason for this result is that the plants in this sample (particularly Lordstown) had a greater than average incidence of strikes. Strikes may be more frequent in these plants precisely because they are one-car, one-plant operations. Overtime hours are twice as important in this sample, but still affect the variance by only 10 percent. Deviations from line speed are not significant for the variance.

The third column shows the results for the quarter-to-quarter variance of output for the universe of plants. For this calculation we aggregated posted output and constructed output to a quarterly frequency (thirteen-week periods) and performed the same calculations. The usefulness of examining a different frequency is illustrated by the following example.⁶ Suppose that output targets are quarterly, rather than weekly, and that the targets correspond to less output than would be produced if the plant operated for 40 hours each week of the quarter. While closing the plant temporarily might be important for the week-to-week variation in output, it might not be as important at the quarterly frequency.

The results in the table show that the conjecture is only partially true. Inventory adjustments and model changeovers still have a significant impact on the variance, but now their combined impact is equal to that of changes in shifts, which are substantially more important at the quarterly frequency than at the weekly frequency. The impact of line speed changes and overtime hours is twice as much at a quarterly frequency, but overtime hours are still

6. This example was suggested by an anonymous referee.

relatively unimportant. The contribution of overtime hours to the quarterly variance in the six matched plants, which is not shown in the table, is 12.8 percent, which is only slightly higher than for the universe of plants.

Overall, the following characterization emerges from Table II. The shifts margin and shutdowns for inventory adjustment and model changeover are the most important contributors to the variance of output. Intermittent production used expressly for inventory adjustment adds more than 20 percent to the variance of output at both the weekly and quarterly frequencies. Changes in line speeds also have a sizable impact on the variance of output at quarterly frequencies. Overtime hours contribute only 10 percent to the variance of output.

Referring back to Figure I, these results imply that most of the variance of output comes from varying hours over the nonconvex portions of the cost function, rather than from varying hours over the convex portions of the cost function. That is, movements in hours from 0 to 40 and back or from 40 to 80 and back are much more important for output volatility than variations in hours along the convex overtime hours portion of the diagram. Furthermore, these results show that almost all of the variance of production is due to margins that involve changing the number of people employed, not the average weekly hours of employed persons.

C. Implications of the Results for some Empirical Puzzles

The results of the last section shed light on at least two prominent questions in economics. The first is a central question in the inventory literature: why is production more volatile than sales? The excess volatility of production is a phenomenon that occurs at many levels of aggregation, from aggregate manufacturing to the division level of the automobile industry [Blanchard 1983; Blinder 1986]. The importance of the intermittent production and shift margins gives a natural explanation for the excess volatility of production. An 80 percent increase in sales can lead to a 100 percent increase in production if the plant responds by adding an extra shift. Likewise, a moderate decrease in sales can lead to a halving of production if a second shift is eliminated. Or, if the plant uses inventory adjustments, a decrease in sales can cause the plant to jump from production at 40 hours one week to zero hours the next. Given the lumpiness of the margins we have uncovered, it is not surprising that production should be more variable than sales.

The effect of this lumpiness is particularly great at the plant level. To illustrate this effect, we compare monthly sales with monthly production at the Ford St. Louis plant during the time it produced the Mercury, which was 1972:1 to 1982:9. For the raw data the variance of production is 35 percent higher than the variance of sales. This number is similar to that obtained by Blanchard for automobile industry divisions, though somewhat greater than those obtained by Blinder for manufacturing aggregates. Several authors have argued, however, that supply shocks can make production more volatile than sales [Blinder 1986; Eichenbaum 1989]. To eliminate the effect of supply shocks, we regressed production on units lost due to supply disruptions, model changeovers, and holidays, as well as on monthly dummy variables. We also regressed sales on monthly dummy variables. Using the residuals of these regressions, the variance of production was still 30 percent higher than the variance of sales.

Of course, it is always possible that the forces leading to production volatility at the plant level are completely different from the forces leading to production volatility at the aggregate level.⁷ Simple variance formulas, however, show that the variance of aggregate production is more likely to exceed the variance of aggregate sales if production is more volatile than sales at each plant. The microeconomic phenomenon will be manifested in the aggregate data as long as the sum of the covariances of sales is not significantly greater than the sum of the covariances of production across plants.

Overall, our findings are supportive of Ramey's [1991] nonconvex cost explanation of the excess volatility of production over the cost shock explanation. The nonconvexities do not, however, take the form of wide ranges of declining marginal costs. Rather, the cost functions we presented earlier suggest that marginal costs are constant or increasing over most ranges of output. However, there are significant nonconvexities at key output values that cause the plant to bunch production at a few points of low average cost.

The second question that the results address is measurement issues in the relationship between output on the one hand and employment and hours on the other. Virtually all studies compare employment and hours measured at a point in time with output measured over an interval of time. Most employment and hours

7. Caballero [1992] uses simulations to construct examples in which asymmetries at the micro level do not lead to asymmetries at the aggregate level.

data are based on the Bureau of Labor Statistics' establishment survey or household survey, which samples the pay period including the twelfth of the month. In many manufacturing industries, such as automobiles, the length of the pay period is one week. Thus, an analysis at the monthly frequency compares production over the entire month with employment and hours during the middle week of the month.

Sims's [1974] paper on short-run increasing returns to labor is one of the few papers that have taken this temporal aggregation problem seriously. Sims showed that so long as weekly employment was smooth, the bias would not be great. Our finding that intermittent production is a common occurrence suggests that the temporal aggregation bias could be important.

To illustrate this point, Figure II shows one graph with monthly output and midmonth employment and another graph with weekly output and weekly employment for the Ford St. Louis plant. For this plant we had actual output data as well as sporadic data on employment.⁸ The employment series we construct counts those actually working, rather than those paid. The monthly data, shown in the top graph, have similar characteristics to those displayed by Hamermesh [1989], although our data are more volatile because of the period studied. Employment is relatively smooth, changing only infrequently and by large amounts. Output is considerably more variable than employment. The R^2 of a regression of monthly output volume on a constant and employment is 0.59. Thus, according to monthly data, 41 percent of the variation in output cannot be explained by employment.

The bottom part of Figure II shows weekly output and employment. At this frequency many of the output movements are matched exactly by employment movements. The R^2 of the regression of weekly output on weekly employment is 0.81. Thus, when the frequencies of employment and output are matched, only 20 percent of the variation in output cannot be explained by employment. A comparison of the two graphs shows that much of the slippage between employment and output at the monthly level is due to the measurement frequency: because of the importance of

8. In particular, from *Automotive News* articles, we knew that when the line speeds were 36 to 39 per hour, there were 1200 workers on the second shift, and that when the plant shutdown for inventory adjustment, 2800 workers on two shifts were laid off. When the line speed was 32 per hour, 2700 workers on two shifts were laid off for inventory adjustment. Information from other plants suggests that approximately 90 percent of employed workers are laid off during shutdowns. Using this information, we were able to assign employment levels to each week.

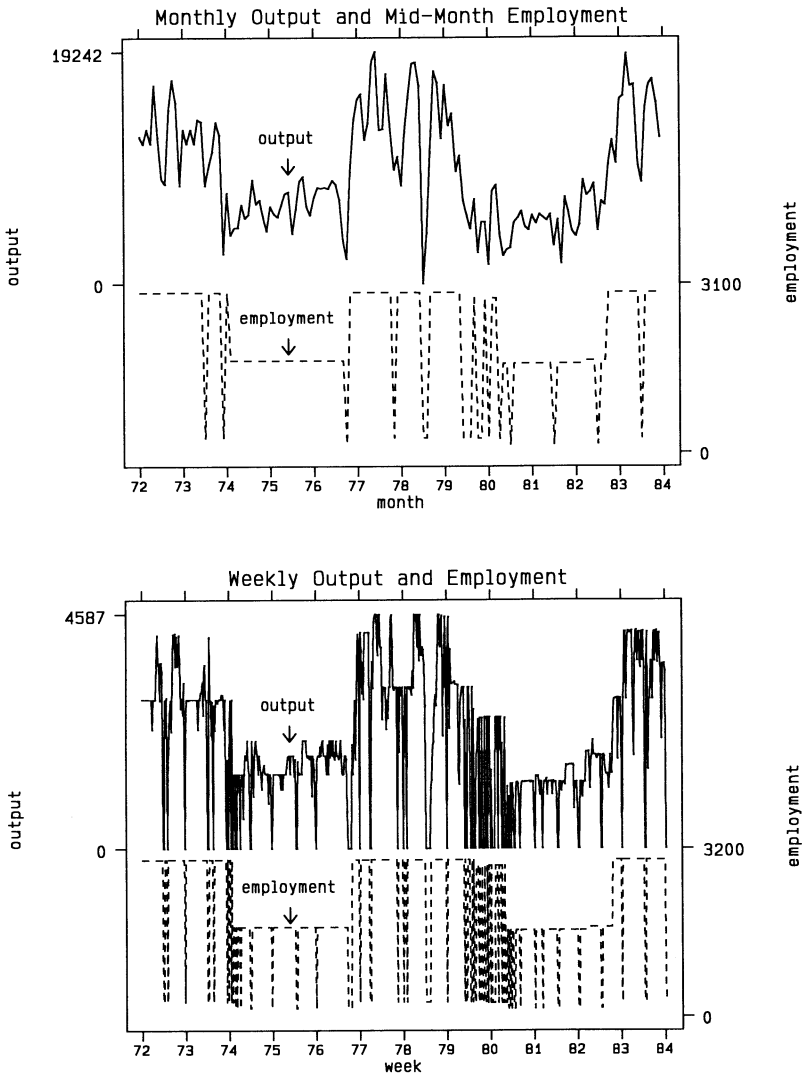


FIGURE II
Output and Employment at Ford St. Louis

intermittent production, comparing the stock of employment measured during one week with production over the entire month can lead to spurious variation in output relative to employment.

The same argument applies to comparisons of output to hours

worked. Extending on the example above, a regression of monthly output at the plant on total hours worked during the middle week of the month produces an R^2 of 0.73. This would imply that 27 percent of the variation in monthly production cannot be explained by variation in hours. On the other hand, a regression of weekly output on total weekly hours has an R^2 of 0.94, so that only 6 percent of the variance in production cannot be attributed to the variation in hours. Matching up the frequencies of output and labor inputs greatly reduces the importance of variations in output relative to hours.

An important question is whether temporal aggregation leads to biases in output and hours aggregated across plants. If plants stagger their shutdowns throughout the month, then the biases might not be as important. To study this issue, we examined the relationship between output aggregated across all plants to a proxy for hours. Because we did not have employment data for all plants, we made the simplifying assumption that employment was proportional to posted output, which is the reported workweek of the plant times the line speed. We then compared the R^2 of the regression of actual monthly output on the sum of posted output across all plants, measured during the midweek of the month, with the R^2 of the regression of weekly actual production to weekly posted output. The R^2 for the monthly frequency was 0.74, while the R^2 for the weekly frequency was 0.93. Thus, the aggregate numbers are very similar to those for the one plant. The fraction of fluctuations in output not attributable to fluctuations in hours is much greater at the monthly frequency than at the weekly frequency.

The measurement errors induced by temporal aggregation do not solve the short-run increasing returns puzzle, for the measurement errors lead to biases in the direction of finding excessive *decreasing* returns. The point of our analysis is to illustrate how the importance of intermittent production can induce spurious movements in productivity, so that those studying the variance and cyclicity of production can have a better understanding of the types of measurement errors in the data. To the extent that intermittent production may be important in other industries as well, we give the following caution: when comparing a flow output with the employment and hours data, one should be very careful about the frequency of the flow data, because the use of intermittent production can lead to confusion about actual productivity and hours movements.

IV. DYNAMIC CHARACTERISTICS OF DIFFERENT MARGINS

We now turn to an analysis of the dynamic characteristics of the 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, we conducted several investigations of duration and persistence. Each of these is presented in turn.

A. Durations of Margin Use

Recall from the discussion in Section I that margins with high adjustment costs and low variable costs should be used for periods of long duration whereas margins with low adjustment costs and high variable costs should be used for shorter durations. To gain insight into the underlying cost structure of the different margins, we estimate the mean duration of the "spells" during which each margin is used.

The estimation problems are similar to those encountered in the literature on unemployment spells. In particular, some spells are not completed during the sample period. That is, the sample begins in the middle of a spell or ends before a spell terminates. Different assumptions about the underlying process can change the estimates substantially. For each case, we report several statistics. We report the average duration of completed spells, the average duration of censored spells, the estimated mean duration for all spells, assuming an underlying exponential distribution, and the estimated median, using the nonparametric Kaplan-Meier [1958] survival curve.

Table III reports the results for the margins. We show the durations of spells in which the margins were switched on as well as switched off. For example, we study both how many consecutive weeks a plant was shut down for inventory adjustment, as well as how many consecutive weeks it was *not* shut down for inventory adjustment. Consider first the ranking of the durations during which each margin was in use. By all four measures, the use of two shifts had the longest duration, with an estimated mean of over 400 weeks and a median of 262 weeks, followed by increases in line speeds, with an estimated mean of 177 weeks. On the other hand, the mean duration of a shutdown for model changeover was under 3 weeks, while the mean duration of shutdowns for inventory adjustment was 1.5 weeks. The mean duration of overtime hours was under two weeks for all plants. As the bottom of the table

TABLE III
DURATIONS OF SPELLS OF MARGIN USE (IN WEEKS)

	Mean of completed spells	Mean of censored spells	Estimated mean	Estimated median	Total number of spells
Inventory adjustment	1.5	—	1.5	1	788
No inventory adjustment	23.5	91.91	33.2	4	820
Model changeover	2.9	—	2.9	2	525
No model changeover	43.2	53.8	52.4	48	557
4 or more overtime hours	1.8	—	1.8	1	2245
Less than 4 overtime hours	9.0	33.6	10.2	2	2276
2 shifts	142.3	207.3	403.4	262	113
1 shift	57.0	168.3	89.4	43	93
Increased line speed	106.2	143.9	176.8	105	161
Decreased line speed	60.2	103.3	95.8	58	113
	Six matched plants				
4 or more overtime hours	4.2	—	4.2	2	280
Less than 4 overtime hours	6.5	13.3	7.0	2	284

shows, though, the mean duration of overtime hours was one month in the sample of six plants for which the overtime data were better.

The duration of time when the margin is not in use follows a similar ranking. Both one-shift periods and decreased line speeds have estimated mean durations near 90 weeks. Periods between model changeovers last approximately one year, as one would expect, while intervals between shutdowns for inventory adjustment last 33 weeks. The periods between the use of the overtime hours margin are the shortest, at ten or fewer weeks.

It should be noted that the estimated means often differ significantly from the medians, indicating a skewness in the duration of spells. The most dramatic example is spells with no inventory adjustment where the mean is 33 weeks but the median is just four weeks. The reason for this difference is that when a plant is shutting down for inventory adjustment, it will often close for one week, open for a week or two, and then close again, and so on for several months. Thus, there are many short spells with no inventory adjustment. On the other hand, plants often go long periods without shutdowns for inventory adjustment, so the mean duration is much longer. This pattern highlights the importance of intermittent production during downturns.

These results support the following inference about the cost

structure. Overtime hours and shutdowns for inventory adjustment and model changeovers are all low adjustment cost-high variable cost margins, while shifts and line speeds are high adjustment cost-low variable cost margins. The number of spells of inventory adjustment, model changeover, and overtime hours greatly exceeds the number of spells for shift and line speed changes, suggesting that the former margins are switched on and off much more frequently. Furthermore, when the latter margins are switched on or off, they remain in that state for a much longer period of time.

This characterization is in line with the apparent costs of each margin discussed above. There are also other aspects of the cost structure that may produce these results. For example, overtime hours have a rather complicated cost structure. For some periods and for some companies, workers who had already worked overtime hours for three weeks in a row were not obligated to work overtime hours the following week. Furthermore, most companies had the policy of not scheduling overtime hours during weeks with holidays. These two constraints have the effect of making overtime hours have less duration.

Finally, it is important to note certain peculiarities of line speed changes. There are two reasons for line speed changes: the plant wants to produce more or less of the same car, or it changes the type of car it produces. A minority of line speed changes are associated with a complete change in the model of car produced. With the others it is not so clear. Of all line speed changes, 57 percent occurred in July or August or during a time when the plant was down for model changeover. Some of these changes were probably due to changes in the specifications of the same model from one model year to the next. On the other hand, if a plant wishes to change the rate of production, it makes sense to make the change when it must reconfigure the plant for the new model year anyway. These considerations make the line speed results harder to interpret.

B. Transitions between States

In order to gain more insight into how plants use different margins and into the underlying cost structure, we now investigate the temporal ordering of the use of margins and the different combinations of margins used. That is, thinking of these combinations of margins as "states," we wish to know, for example, how often plants are in states in which they use one shift with overtime

TABLE IV
TRANSITION PROBABILITY MATRIX (WEEKLY FREQUENCY)

State at t	1shift <i>IA</i>	1shift <i>MC</i>	1shift <i>RH</i>	1shift <i>OH</i>	2shift <i>IA</i>	2shift <i>MC</i>	2shift <i>RH</i>	2shift <i>OH</i>
State at $t - 1$								
1 shift <i>IA</i>	35	1	64	0	0	0	0	0
1 shift <i>MC</i>	2	69	26	1	0	1	2	0
1 shift <i>RH</i>	4	2	86	7	0	0	1	0
1 shift <i>OH</i>	0	0	51	48	0	0	0	0
2 shifts <i>IA</i>	1	0	3	0	30	1	66	0
2 shifts <i>MC</i>	0	0	1	0	1	66	32	1
2 shifts <i>RH</i>	0	0	0	0	3	2	81	13
2 shifts <i>OH</i>	0	0	0	0	0	0	57	43
Probability of each state	1.4	1.4	20	2.5	2.8	4.1	55	12
Probability of each state in six matched plants	1.1	1.6	17	7.2	2.0	2.5	38	31

Abbreviations: *IA* = plant shut down for inventory adjustment for at least part of the week. *MC* = plant shut down for model changeover for at least part of the week. *RG* = regular hours; plant not shut down for inventory adjustment or model changeover and not working four or more of overtime hours per week per shift. *OH* = overtime hours greater than or equal to four per week per shift.

hours versus two shifts with regular hours. We also want to analyze how plants progress between states.

We first define several states of the world in which a plant might be in a given week. For this analysis we ignore changes in line speeds, because it is difficult to define a small number of states that allow for variations in line speeds. Plants can be in either one- or two-shift operations. For each of those there are four statuses: (1) shut down for at least part of the week for inventory adjustment; (2) shut down for at least part of the week for model changeover; (3) not shut down for inventory adjustment or model changeover, and operating less than four overtime hours per week per shift; (4) operating four or more overtime hours per week per shift.⁹ Thus, there are eight possible states.

We first examine the transitions between states in order to understand the permanence of the states themselves. Table IV reports the transition probability matrix for week-to-week move-

9. There was only one case of a plant being categorized in more than one state. During one week, Ford Los Angeles was down for most of the week for model changeover. When it came back up, it used overtime. We classified it as a model changeover.

ments across states. Consider first the elements on the diagonal. The most clearly persistent states are those with regular hours, and either one or two shifts. The probability of staying in those states is over 80 percent. The bottom of the table shows that the probability of being in either regular hours state is 75 percent for all the plants, and 55 percent for the six matched plants. The second most persistent states are those with model changeovers, followed by overtime hours, and finally inventory adjustment with a persistence over 30 percent.

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, operations at inefficiently low (inventory adjustment) or inefficiently high (overtime hours) levels are fairly persistent, as well. Note that movements to inventory adjustment or overtime hours states usually involve a change in production of at least 20 percent. Thus, although the automaker prefers to remain near 40 hours a week, his manipulation of the margins leads to a great deal of volatility in actual production. Model changeover states are also quite persistent. There are two reasons for this. First, for conversions to significantly different models, the time to retool is significant. Second, when demand for a particular type of car is low, firms sometimes extend the downtime for model changeovers.

Consider now transitions between states. The entries off the block diagonals support the frequencies reported above in Table I: changing the number of shifts is a very low probability event for any given week. A shift is most likely to be added if the plant was down for model changeover the week before; a shift is most likely to be dropped if a plant was down for inventory adjustment the week before.

Now consider transitions between states that do not involve a change in the number of shifts. The probability of going from a shutdown for inventory adjustment in one week to overtime the next and vice versa is zero. This behavior supports the posited cost structure and the assumption that plant managers are rational. A plant down for model changeover is most likely to stay in that state or to move to regular hours operation.

The probabilities at the bottom of the table show how often the

plant is in each state. The numbers for regular hours and overtime hours for the six matched plants are probably more accurate because of the better data on overtime hours. The most likely state is the one with two shifts and regular hours. The plant is not very likely to be in a state with one shift and overtime hours operation. These results support the following interpretation. Overtime hours have a high marginal cost, and are only used in two situations: first, when the change in demand is thought to be temporary, or when the firm is waiting for more information; and second, when two shifts are already working and the plant is still below target. Typically, if demand stays high, the company will start producing the particular model at another plant, if possible.

C. Predictions about Future Movements in Output

Finally, we investigate what the transitions between states predict about future movements in output. Recall from the theoretical discussion that we would expect the use of high adjustment cost margins to signal persistent changes in output, and the use of low adjustment cost margins to signal transitory changes in output. Working from the definitions of the states given above, we identified "events" as movements between states. For movements between regular hours and overtime hours, 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.

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

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

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.

The results are shown in Table V. Let us begin with the simplest result, that called "no change." In the first column this has a coefficient of 0.40. Descriptively, this means that a change in quantity *not* achieved by changing the plant status tends to be substantially reversed four weeks later; only 40 percent of any

TABLE V
MARGIN CHOICES AS PREDICTORS OF PERSISTENCE (POOLED DATA SET)

Dependent variable	$Q(4) - Q(-1)$		$Q(13) - Q(-1)$	
	Coefficient	Standard error	Coefficient	Standard error
Constant	-51	11	-82	13
No change	0.40	0.01	0.42	0.01
Add shift	0.74	0.07	0.80	0.09
Drop shift	0.80	0.09	0.80	0.10
Increased line speed	0.84	0.06	0.83	0.06
Decrease line speed	0.74	0.06	0.78	0.07
IA to MC	0.75	0.57	0.94	0.66
IA to RH	0.71	0.02	0.82	0.02
IA to OH	0.25	0.57	0.74	0.41
MC to IA	1.14	0.46	1.25	0.46
MC to RH	0.93	0.02	0.88	0.03
MC to OH	0.91	0.17	0.84	0.19
RH to IA	0.37	0.02	0.24	0.02
RH to MC	0.25	0.02	0.03	0.03
RH to OH (1 shift)	0.50	0.14	0.74	0.17
RH to OH (2 shifts)	0.41	0.03	0.40	0.03
OH to IA	0.96	0.16	0.59	0.19
OH to MC	0.26	0.14	0.31	0.16
OH to RH (1 shift)	0.56	0.16	0.35	0.18
OH to RH (2 shifts)	0.64	0.03	0.81	0.04

IA = shut down for inventory adjustment; MC = shut down for model changeover; RH = regular hours; OH = overtime hours greater than or equal to four.

change up or down persists the whole four weeks when the managers accommodate it without a plant status change. The same is true for a quarter ahead. The economic interpretation is relatively straightforward. These types of output changes result from small variations in overtime hours, or from variations in hours due to supply disruptions, such as the weather. 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 short-run marginal cost curve. If most small shocks to desired output are transitory, one would expect them to be accommodated by transitory shifts in actual behavior.

Now contrast these results with those for adding and dropping shifts, and for increasing and decreasing line speeds. Changes in

quantity associated with these decisions are more persistent, with 80 percent of the change remaining after a quarter. The shift and line speed results are symmetric in both directions for horizons up to a quarter.

On the other hand, the coefficient estimates imply that all movements to model changeovers lead to very transitory changes in output. The only large coefficient is the *IA* to *MC* coefficient, which is not estimated precisely. On the other hand, movements from model changeover to any other state are very persistent.

Consider now movements between the regular hours and overtime hours states. The transitory movements are regular hours to overtime hours when there are two shifts and overtime hours to regular hours when there is one shift. The other two combinations are much more persistent at the quarterly frequency. Seventy-four percent of the change persists one quarter later when there is a movement from regular hours to overtime hours and only one shift. There is a simple explanation for this persistence. In almost 10 percent of the cases, a movement from regular hours to overtime hours is followed one quarter later by the addition of a second shift. Since the addition of the second shift leads to an increase in output that is more than five times as much as that brought about by overtime hours, the average output change appears to be very persistent. The other persistent case is overtime hours to regular hours when there are two shifts, which simply reflects the fact that regular hours are a persistent state.

Altogether, the regression results support the same ranking of costs implied by our duration analysis. The persistence of output changes brought about by line speed and shifts changes relative to those achieved by overtime hours and shutdowns supports the proposition that the adjustment cost of changing shifts and line speeds is substantially higher than the adjustment cost of using overtime hours and shutdowns. The former are used to achieve persistent changes in output whereas the latter are used to achieve temporary changes in output.

V. CONCLUSIONS

Our investigation of weekly production at 50 automobile assembly plants has uncovered characteristics of output volatility that are not apparent in aggregate data. Our most important overall finding is that adjusting production is a more complicated process than simply “changing *Q*” or choosing the mix of capital

and labor. Of the multiple margins used by plant managers, adding or dropping a shift and varying regular hours by shutting the plant down for a week are the most important. Overtime hours are substantially less important.

Our findings can also explain why production is often more volatile than sales. The margins most often used, shifts and shutdowns, are very lumpy and result in dramatic variations in the volume of production. It is easy to see how a moderate change in sales can lead to a larger change in production because of the underlying cost structure.

The importance of intermittent production has further implications for comparisons of output volume to labor input measured at a point in time. In particular, the use of intermittent production can explain why the correlation between the monthly volume of output and midmonth employment and hours is relatively low. We find that weekly employment and hours can explain a significantly greater fraction of the variation in weekly output.

The dynamic characteristics of the margins reveal other aspects of the cost structure. The fact that shifts and line speeds are changed infrequently, that the duration of use is long, and that the consequent change in output is persistent, relative to overtime hours and shutdowns, supports an underlying cost structure in which shifts and line speeds are low marginal cost-high adjustment cost margins whereas overtime hours and shutdowns are high marginal cost-low adjustment cost margins. The fact that overtime hours do not contribute significantly to the variance of output suggests that their true marginal cost is indeed high.

STANFORD UNIVERSITY AND THE NATIONAL BUREAU OF ECONOMIC RESEARCH
UNIVERSITY OF CALIFORNIA, SAN DIEGO AND THE NATIONAL BUREAU OF ECONOMIC RESEARCH

REFERENCES

- Abernathy, William J., Kim B. Clark, and Alan M. Kantrow, *Industrial Renaissance: Producing a Competitive Future for America* (New York, NY: Basic Books, 1983).
- Aizcorbe, Ana, "Procyclical Labour Productivity, Increasing Returns to Labour and Labour Hoarding in U. S. Auto Assembly Plant Employment," *Economic Journal*, CII (1992), 860-73.
- Alchian, Armen, "Costs and Outputs," in *The Allocation of Economic Resources*, Moses Abramovitz et al., eds. (Stanford, CA: Stanford University Press, 1958).
- Automotive News*, various issues from 1972 to 1983.
- Bertola, Giuseppe, and Ricardo J. Caballero, "Kinked Adjustment Costs and Aggregate Dynamics," *NBER Macroeconomics Annual 1990*, 237-88.
- Blanchard, Olivier, "The Production and Inventory Behavior of the American Automobile Industry," *Journal of Political Economy*, XCI (1983), 365-400.

- Blinder, Alan, "Can the Production Smoothing Model of Inventory Behavior be Saved?" *Quarterly Journal of Economics*, CI (1986), 431–53.
- Bresnahan, Timothy, "Departures from Marginal-Cost Pricing in the American Automobile Industry," *Journal of Econometrics*, XVII (1981), 201–27.
- Bureau of Labor Statistics, *Industry Wage Survey: Motor Vehicles and Parts, 1973–1974*, Bulletin 1912, 1976.
- Caballero, Ricardo J., "A Fallacy of Composition," *American Economic Review*, LXXXII (1992), 1279–92.
- Caballero, Ricardo J., and Eduardo Engel, "Microeconomic Adjustment Hazards and Aggregate Dynamics," MIT manuscript, 1992.
- Cooper, Russell, and John Haltiwanger, "The Aggregate Implications of Machine Replacement: Theory and Evidence," *American Economic Review*, LXXXIII (1993), 360–82.
- Davis, Stephen J., and John Haltiwanger, "Gross Job Creation, Gross Job Destruction, and Employment Reallocation," *Quarterly Journal of Economics*, CVII (1992), 819–64.
- Dunne, Timothy, Mark Roberts, and Larry Samuelson, "Plant Turnover and Gross Employment Flows in the U. S. Manufacturing Sector," *Journal of Labor Economics*, VII (1989), 48–71.
- Eichenbaum, Martin, "Some Empirical Evidence on the Production Level and Production Cost Smoothing Models of Inventory Investment," *American Economic Review*, LXXIX (1989), 853–64.
- Hamermesh, Daniel, "Labor Demand and the Structure of Adjustment Costs," *American Economic Review*, LXXIX (1989), 674–89.
- Holt, Charles C., Franco Modigliani, John F. Muth, and Herbert A. Simon, *Planning Production, Inventories, and Workforce* (Englewood Cliffs, NJ: Prentice-Hall, Inc., 1960).
- Kaplan, E. L., and P. Meier, "Nonparametric Estimation from Incomplete Observations," *Journal of the American Statistical Association*, LIII (1958), 457–81.
- Leonard, Jonathan S., "In the Wrong Place at the Wrong Time: The Extent of Frictional and Structural Unemployment," in *Unemployment and the Structure of Labor Markets*, Kevin Lang and Jonathan S. Leonard, eds. (New York, NY: Basil Blackwell, Inc., 1987).
- Maloney, Michael T., and Robert E. McCormick, "A Theory of Cost and Intermittent Production," *Journal of Business*, LVI (1983), 139–54.
- Ramey, Valerie, "Nonconvex Costs and the Behavior of Inventories," *Journal of Political Economy*, XCI (1991), 306–34.
- Rothschild, Michael, "On the Cost of Adjustment," *Quarterly Journal of Economics*, LXXXV (1971), 605–22.
- Sargent, Thomas J., "Estimation of Dynamic Labor Demand Schedules under Rational Expectations," *Journal of Political Economy*, LXXXVI (1978), 1009–44.
- Sims, Christopher, "Output and Labor Input in Manufacturing," *Brookings Papers on Economic Activity* (1974), 695–736.
- Wards' Automotive World*, various issues.
- Wards' Automotive Yearbook*, various years.
- White, Lawrence J., *The Automobile Industry since 1945* (Cambridge, MA: Harvard University Press, 1971).