

COWLES FOUNDATION FOR RESEARCH IN ECONOMICS
AT YALE UNIVERSITY

Box 2125, Yale Station
New Haven, Connecticut 06520

COWLES FOUNDATION DISCUSSION PAPER NO. 1169

Note: Cowles Foundation Discussion Papers are preliminary materials circulated to stimulate discussion and critical comment. Requests for single copies of a Paper will be filled by the Cowles Foundation within the limits of the supply. References in publications to Discussion Papers (other than mere acknowledgment by a writer that he has access to such unpublished material) should be cleared with the author to protect the tentative character of these papers.

NON-CONVEX COSTS AND CAPITAL UTILIZATION: A STUDY OF
PRODUCTION SCHEDULING AT AUTOMOBILE ASSEMBLY PLANTS

George J. Hall

December 1997

Non-Convex Costs and Capital Utilization: A Study of Production Scheduling at Automobile Assembly Plants

George J. Hall*
Yale University

November 1997

Abstract

This paper studies how managers at automobile assembly plants organize production across time. Detailed data from eleven single-source automobile assembly plants display considerable cross-plant heterogeneity. At plants which make low- and medium-selling vehicles the capital stock often sits idle, production is more variable than sales, and weeklong shutdowns are often used to vary output. In contrast, at plants which make high-selling vehicles, the capital stock rarely sits idle, production is about as variable as sales, and overtime – not weeklong shutdowns – is most frequently used to vary output. To explain this difference in production scheduling, I formulate and solve a dynamic programming model of a plant manager. The solution to the dynamic program predicts that when sales are low, non-convexities at the plant level induce the manager to bunch production at points of low average cost; thus, the manager uses less than full capital utilization on average and makes production more volatile than sales. When sales are high, the plant operates in a convex region of the cost curve. Hence the manager employs high levels of capital utilization and makes production less volatile than sales.

*Department of Economics, Yale University, P.O. Box 208264, New Haven, CT 06520-8264; e-mail: ghall@econ.yale.edu. I am grateful to the Chrysler Corporation for providing much of data used in this paper. I have benefited from discussions with participants at numerous seminars, from comments made by an anonymous referee, and from conversations with John Cochrane, Bill Dupor, Martin Eichenbaum, Lars Hansen, Anil Kashyap, John Rust, Tom Sargent, Bob Schnorbus, and Lew Segal. Pier Deganello provided able research assistance. Most of the work on this paper was done while I was on the staff at the Federal Reserve Bank of Chicago.

1 Introduction

This paper studies how managers at automobile assembly plants organize production across time. I formulate and solve a dynamic programming model that explains the production behavior observed from a new plant-level dataset. The model incorporates two non-convex margins: the adding and dropping of an additional shift and the shutting down of the plant for a week at a time. These non-convex margins play a central role in explaining much of the heterogeneity in production scheduling observed in the data. Specifically the model predicts that, when sales are low, plant managers will use primarily non-convex margins to adjust output. Thus production will be more variable than sales and the plant's capital will sit idle much of the time. In contrast, when sales are high, plant managers will use convex margins to adjust output; this production behavior is consistent with production as variable as sales and high levels of capital utilization.

I study a new database of fourteen automobile assembly plants. Eleven of these plants are the sole producers of various vehicle lines. For these eleven plants, weekly data on capital utilization and production can be accurately lined up with monthly data on employment, inventories and sales. These data display three facts that a successful model of automobile production should capture.

1. For the average plant the workweek of capital is just 66.8 hours. More striking though are the differences in capital utilization across plants. While the average workweek of capital for some plants is close to 100 hours, it is less than 15 hours at some other plants. Yet at all the plants the nominal premium for night work is modest, and the costs of having idle workers on the payroll are large. Workers on the second shift receive only about 5 percent more than workers on the first shift. Laid-off workers from these plants receive 95 percent of their straight time wage plus benefits.

Puzzling low levels of capital utilization are not unique to the auto industry. The capital stock in U.S. manufacturing industries is employed, on average, fewer than 60 hours per week (Shapiro, 1995). Shapiro argues that the true marginal premium for second shift work is closer 25 percent. Although this higher marginal shift premium partially resolves the puzzle, the question still remains: Why does the capital stock at some of these plants sit idle so much of the time?

2. The average plant makes the standard deviation of monthly production 21 percent larger than the standard deviation of sales. However, this production pattern is not uniform across all the plants. The plants that assemble the high-selling vehicle lines make production about as volatile as sales; the plants that assemble the low-selling vehicle lines tend to make production much more volatile than sales.

For a wide variety of industries, production is more volatile than sales. This fact has generated considerable attention since classic models of inventories, which assume convex short-run increasing marginal costs, imply that firms should manage inventories such that production is smoother than sales.¹ Although many explanations have been offered, there is no proposed answer to the question: Why is production more variable than sales at some plants but not at others?

3. Plant managers rarely change the number of shifts or the line speed. Managers at plants which assemble high-selling vehicles most frequently vary hours worked by using overtime. Managers at plants that assemble low- and medium-selling vehicles regularly vary hours worked by shutting down the plant for a week at a time. This production behavior is puzzling since the cost of laying off workers is high.

Previous studies, such as Bresnahan and Ramey (1994), and Matthey and Strongin (1995), have documented the frequent use of short-term layoffs and infrequent use of shift-changes and line-speed changes to vary output at manufacturing plants. But this paper attempts to explain the observed heterogeneity in production scheduling of nearly identical plants. That is, why do some automobile assembly plants – but not others – use weeklong shutdowns so frequently to vary output?

Building on the work of Hamermesh (1989), Ramey (1991), Cooper and Haltiwanger (1993), and Bresnahan and Ramey (1994), this paper argues that non-convex margins of adjustment play a key role in understanding these facts. These non-convexities arise from two sources. First, the plant faces an integer constraint on the number of shifts that can be run. Second, there are fixed costs to opening the plant each week and running a shift. Additionally, provisions in the union labor contract (i.e., the required premium for overtime and a pay floor for short-weeks) create kinks in the plant's cost function. These labor contract provisions and non-convex margins

¹See Blinder and Maccini (1991) and the citations therein.

produce large discontinuous drops in the plant's marginal cost curve. When sales are low, the plant operates in a non-convex region of its cost curve. In this region it is optimal for the plant to oscillate between periods of not producing and periods of producing a lot. This production behavior is consistent with a low average workweek of capital, production more variable than sales, and frequent plant shutdowns. However when sales are high, the plant operates in a convex region of its cost curve, so the firm wishes to smooth production and use high levels of capital utilization.

I solve a dynamic cost minimization model of an assembly plant manager who takes the sales process as given. Consequently, I do not need to make any restrictive assumptions about the market structure or the nature of demand in order to solve the model. But the large automakers do behave as if they face downward sloping demand curves for their products.² So, this model can be viewed as a sub-problem which a profit-maximizing automaker solves when choosing from a menu of prices and quantities.

The formal analysis involves solving the dynamic cost minimization model for nine of the eleven single-source plants. I use the dataset to both parameterize the model and evaluate the performance of the model. One of the advantages of modeling production at the plant level is that several of the parameters do not need to be estimated; they are simply drawn from the labor contracts. Other parameters are estimated to match moments of the employment and sales data. The results of this dynamic model demonstrate that much of the variation across plants in capital utilization and relative variability of production and sales can be attributed to the mean of the sales process.

The rest of the paper is organized as follows. The second section provides some background information on how automobile assembly plants are run. The third section presents the dataset. The fourth section develops the intuition behind the model. The fifth section presents the dynamic programming model. In the sixth section parameter values are selected, the model is solved, and moments implied by the model are compared to moments in the data. In the final section some concluding comments are made.

²See Bresnahan (1981), Blanchard and Melino (1986), and Berry, Levinsohn, and Pakes (1995) for models of the automobile industry in which both prices and quantities are endogenous.

2 Some auto industry details

Although there is some variation across plants and firms, most production decisions for automobile assembly plants are made at the monthly frequency. Once a month, there is a capacity planning meeting in which production schedules are set. At this meeting managers are presented with last month's sales and inventory numbers and a sales forecast. The managers must then set and revise their production schedule. They have five margins at their disposal.

The first margin is how many weeks the plant is scheduled to be open. The second margin is how many days per week the plant is scheduled to be open. The third margin is the scheduled number of shifts per day. The fourth is the scheduled length (in hours) of each shift. The fifth margin is the rate of output – in jobs (vehicles) per hour. This last margin is usually called the line speed. Scheduled monthly production is the product of these five margins:

$$\frac{\text{jobs}}{\text{month}} = \frac{\text{weeks open}}{\text{month}} \times \frac{\text{days open}}{\text{week}} \times \frac{\text{shifts}}{\text{day}} \times \frac{\text{hours}}{\text{shift}} \times \frac{\text{jobs}}{\text{hour}}. \quad (1)$$

The costs associated with manipulating these five margins differ. Many of these differences are due to the structure of the labor contracts these plants operate under.

Although production schedules are usually set at a monthly frequency, standard labor contracts are written with a one-week time period in mind. The average straight-time, day-shift wage at these plants about is \$18 an hour plus benefits. Workers on the second (evening) shift receive a 5 percent premium. Workers on a third (night) shift receive a 10 percent premium. Any work in excess of eight hours in a day and all Saturday work is paid at a rate of time and an half. Employees working fewer than 40 hours per week must be paid 85 percent of their hourly wage times the difference between 40 and the number of hours worked. This “short-week compensation” is in addition to the wages the worker receives for the hours s/he actually worked.

If the firm chooses to not operate a U.S. plant for a week, the workers are laid off. After a single waiting week each year, laid-off workers receive 95 cents on the dollar of their 40 hour pay in unemployment compensation. Of this 95 cents, state unemployment insurance (UI) pays about 60 cents. The remaining 35 cents is picked up by supplemental unemployment benefits (SUB). Firms do not pay laid-off workers directly, but laying off workers does increase the firm's experience rating and UI premiums in the future. Because of the cross-industry subsidies inherent in the UI system, firms end up paying about half of the 60 cents coming from UI.³ Since the SUB is a

³See Anderson and Meyer (1993) and Aizcorbe (1990).

negotiated benefit between the firm and the union, the firm ultimately pays all 35 cents. So, after the initial waiting week, it costs the firm about 65 percent of the 40 hour wage to lay a worker off for one week.

Unemployment insurance in Canada is slightly different. For laid-off Canadian auto workers there is a two-week waiting period each year before benefits are paid. These workers then receive 95 percent of their 40 hour wage in unemployment compensation. Government unemployment insurance pays 55 percent of a worker's full-time earnings. The remainder is picked up by SUB. Unlike the U.S., Canadian UI is not experience rated, so the firm only pays the SUB portion.

Since 1992, several North American assembly plants have started to run three seven-hour shifts per day. This allows the plant to be run 21 hours a day. Workers at these plants are paid eight hours of wages per day, Monday through Friday, for their seven hours of work. Therefore with no overtime, workers are paid a 40-hour wage for working 35 hours.

3 The data

This section describes a dataset of fourteen automobile assembly plants in the United States and Canada. The dataset contains weekly production data from the first week of 1990 to the last week of 1994 and monthly employment, sales, inventory, and production data from January 1990 through December 1994. All the assembly plants are run by the Chrysler Corporation.

For each assembly plant the following weekly data were collected: 1. the number of days the plant operated; 2. the number of days the plant was down for holidays, supply disruptions, model changeovers, or inventory adjustments; 3. the number of shifts run; 4. the hours per shift run; 5. the scheduled jobs per day (line speed); and 6. the actual production for each vehicle line produced at the plant. The Chrysler Corporation supplied data on 1, 3, 4, and 5. Data on 2 and 6 were taken from *Ward's Automotive Reports*, *Ward's AutoInfoBank*, and *Automotive News*.

For each vehicle line produced at these plants, monthly sales data were collected. Total sales by vehicle line are the sum of sales by U.S. dealers, Canadian retail sales, and exports to the rest of the world. Sales by U.S. dealers are from *Ward's*. Canadian retail sales are from the Motor Vehicle Manufacturers Association (MVMA).⁴ Exports are from the American Automobile

⁴Since *Ward's* and the MVMA aggregate the sales of the regular wheelbase minivans (Caravan and Voyager assembled at the Windsor facility) with the extended wheelbase minivans (Grand Caravan and Grand Voyager assembled at the St. Louis II facility), I use U.S. registration data provided by The Polk Company to decompose the Caravan and Voyager sales numbers.

Manufacturers Association (AAMA).

For eleven of the plants, Chrysler provided the number of paychecks written each month. At these plants a pay-period is one week. So using the weekly data described above, I was able to construct a measure of employment for each plant.

Of the fourteen plants in the sample, eleven plants are single-source plants for at least part of the time. A single-source plant is a facility that is the exclusive producer of a set of vehicle lines. By restricting myself to single-source plants, I am able to line up inventory and sales data by vehicle line to employment, production, and hours worked by plant. The assembly plants in the sample are listed in table 1. Table 1 also reports whether each plant is a single-source plant or not, and it lists the vehicle lines produced at each plant. This database is similar to the weekly database constructed by Bresnahan and Ramey (1994).⁵ In particular they identify six of the 50 plants in their sample as single-source plants; they refer to this subset as the “six matched plants.”

These data display three facts. These facts are now presented in a slightly different order than in the introduction.

Fact 3 *Managers rarely change the number of shifts or the line speed to vary output. Managers at plants that assemble low-selling vehicles most frequently vary hours worked by shutting down the plant for a week at a time. Managers at plants which assemble high-selling vehicles most frequently vary hours worked by using overtime.*

Recall from equation (1) that scheduled output is the product of five margins. Table 2 reports how often each of the five margins are used at each plant. The table reports the number of weeks each plant was open, closed, running a short-week, or running overtime. The table also reports the number of times a shift was added or dropped and the number of line speed changes. There are 261 weeks in the sample period. The plants are divided into four groups. The single-source plants are in the first three groups. The dual-source plants are in the fourth group. The single-source plants are divided up into plants which make the high-, medium-, and low-selling vehicles. Since production of the Jeep Wrangler moved from Brampton to Toledo II in 1992, these two plants are concatenated.

A plant is counted as open for the week if it is up and running at least one day during the

⁵Aizcorbe (1992), Cooper and Haltiwanger (1993), Kashyap and Wilcox (1993), and Aizcorbe and Kozicki (1995) also study plant-level data for automobile assembly plants but at the monthly frequency.

week. Otherwise it is counted as closed. If the plant is closed or open fewer than 5 days during the week, the primary reason for the downtime is reported. Following Bresnahan and Ramey (1994), every closure is classified under one of the following categories: holiday or union dictated vacation (HOL), model changeover (MC), supply disruption (SUP), inventory adjustment (IA), or long-run closure (LRUN). Columns 2 through 5 in table 2 report the number of full-week closures broken down by category. Long-run closures are not reported; a plant is classified under a long-run closure if it is closed for more than three months in a row.

Weeklong shutdowns are frequent. Consider the bottom two rows of table 2. The average plant was only open 173 weeks out of 261 total weeks; that is only 2/3 of the weeks available. Even if the long-run closures are excluded, the average plant was only open 84 percent of the available weeks (173 out of a possible 207 weeks). Thus the average plant was closed about 8 1/2 weeks each year. Weeklong shutdowns for inventory adjustment account for most of this downtime.

The averages however do not tell the whole story. Several of the plants, in particular Jefferson North, St. Louis II, and Windsor, were rarely closed for inventory adjustment (or for any other reason). The vehicles made at these plants (sport utility vehicles and minivans) have been among Chrysler's best sellers. In contrast, from 90:1 to 91:12, Bramalea was closed more weeks for inventory adjustment than it was open. During that time the slow-selling Premier and Monaco were assembled there. This is also the case for Toledo II from 90:1 through 91:6 while the Grand Wagoneer (a low seller) was assembled. Weeklong shutdowns most frequently occurred at plants which made low-selling vehicles.

Table 2 also reports the total number of weeks each plant was open for fewer than five days. This is the number of "short-weeks." In column 7, the number of short-weeks that are due to holidays is also reported. From these two columns, it is clear that almost all the short-weeks in the sample are due to holidays. Many of the remaining non-holiday short-weeks are explained by supply disruptions. Very few of these short-weeks are due to inventory adjustment. This is not surprising given the 85 percent short-week rule in the union labor contract discussed above.⁶

Column 8 reports the number weeks each plant used overtime. The average plant used overtime during 38.4 percent of the weeks in the sample. The plants which made the most extensive use of overtime (i.e., Jefferson North, St. Louis II, and Windsor) are the plants that rarely shut down for inventory adjustment. In contrast several of the plants that rarely used overtime, such

⁶See Aizcorbe (1992) and Bresnahan and Ramey (1994).

as Bramalea(90:1-91:12), St. Louis I, and Toledo II(90:1-91:6), were frequently shut down for inventory adjustment. The medium-sales plants such as Pillette Road and Toledo I used both overtime and weeklong shutdowns to vary output. In general, overtime was used frequently, and the plants which made the high-selling vehicle lines used overtime the most.

Finally, columns 9 and 10 report the number of times a shift is added or dropped and the number of times a change in the line speed is made. Changes in the number of shifts were made rarely. At all the plants, changes in the line speed occurred less frequently than weeklong shutdowns or weeks with overtime.

Fact 1 *The average plant operates only 66.8 hours of the 168 available hours in a week.*

Table 3 reports the number of shifts run and the average workweek of capital for each plant. The average workweek of capital conditioned on the plant not being under a long-run closure is presented in the far right column. The average workweek of capital conditioned on the plant being open is presented in column 4.

The three plants that were identified as frequent users of overtime and infrequent users of inventory adjustment (Jefferson North, St. Louis II, and Windsor) are plants which employed three shifts by the end of the sample. Not surprisingly these three “3-shift plants” have the longest average workweeks of capital. The plants which rarely used overtime and were often closed for inventory adjustment, Bramalea(90:1-91:12), St. Louis I, and Toledo II(90:1-91:6), all ran 1 shift and have the shortest workweeks of capital.

Shapiro (1995) states that “the workweek of capital in U.S. manufacturing averages less than 60 hours per week.” At the Chrysler plants, when the long-run closures are excluded, the average workweek of capital is 66.8 hours.⁷ This is in the ballpark of Shapiro’s statement. This finding is also consistent with other measures of capital utilization reported by Shapiro. Shapiro (1993) reports that for manufacturing plants sampled by the Census’ Survey of Plant Capacity from 1977-1988 the average workweek of capital is 80.3 hours/week. Using data from the BLS’s Industry Wage Survey, Shapiro (1995) reports that the capital stock is utilized only 11.4 hours per 24 hour day for the industries he studies.

Shapiro (1995) finds these low levels of capital utilization puzzling. So he asks, if second shift employees are paid only 5 percent more than their first-shift counterparts, why do more firms

⁷If long-run closures are not excluded, the average workweek of capital is 53.1 hours. This is in line with Shapiro’s statement.

not employ second shifts? He partially answers this question by providing evidence that the true marginal premium for night work substantially exceeds the nominal premium. Shapiro argues a better estimate of the shift premium is 25 percent. However the short average workweek of capital reported here is not due to the plants' failure to run second shifts – all but three plants ran more than a single shift. This short average workweek of capital is largely due to the plants being closed so much of the time. Conditional on the plants being open, the average workweek of capital is 80.0 hours.

The differences in the average workweek of capital across the plants are striking. At one extreme is Toledo II; while the Grand Wagoneer was being assembled, the Toledo II facility averaged only 12.7 hours of use per week. At the other extreme is St. Louis II; it ran, on average, almost 100 hours per week. If one thinks of 100 hours per week as a lower bound on what is possible to utilize capital, then the Toledo II facility utilized its capital only 12.7 percent of the time available. The Pillette Road facility is perhaps more representative of the sample. Pillette Road was never down for a long-run closure during the sample period but averaged only 60.4 hours of use per week. So it utilized its capital less than two-thirds of the time available. The question still remains: Why is the level of capital utilization so low at so many of the plants?

Fact 2 *For the average plant, production is more volatile than sales. For the plants that assemble the high-selling vehicle lines, production is about as volatile as sales. For the plants that assemble the medium- and low-selling vehicle lines, production is more volatile than sales.*

Table 4 provides the means and standard deviations of the monthly production, sales and inventory data for the set of single-source plants. Total sales are the sum of U.S. sales, Canadian sales, and exports to the rest of the world. Inventories are computed by a perpetual inventory method. Inventories are benchmarked so that the inventories of discontinued vehicle lines are eventually zero. Inventories for all other vehicles lines are benchmarked using December 1989 U.S. dealer inventory-to-sales ratios.

The three plants with the highest average levels of monthly production are Windsor, St. Louis II, and Jefferson North; these plants rarely closed and used overtime extensively.⁸ More interesting are the relative standard deviations of production and sales. For all but four plants, the standard

⁸I refer to these plants as the “high-sales” plants. However more vehicles were sold from the Dodge City plant than from the Jefferson North plant. But unlike Dodge City, Jefferson North began the period with zero inventories. So perhaps a better label would be “high sales + inventory accumulation.”

deviation of production is substantially greater than the standard deviation of sales. Note three of the exceptions: Jefferson North, St. Louis II, and Windsor. For the plants that rarely shut down for a week at time but use overtime extensively, production is about as volatile as sales. For the plants which shut down for inventory adjustment more frequently, production is more volatile than sales.⁹ The standard deviation of aggregate production over these eleven plants is 30 percent larger than the standard deviation of aggregate sales.

4 A static example

This section presents a simple one-period cost minimization problem of a plant manager. The static case is presented solely for pedagogical purposes. The importance of the non-convexities in the manager's problem are more easily illustrated in the static case than in the dynamic case.

Consider a plant in which the rate of production (the line speed) is Cobb-Douglas in capital, k , and labor, n . The time period is one week. The plant must produce at least q goods. The plant can operate D days. It can run one or two shifts, S , each day; both shifts are of length h . Let n employees work each shift. Workers on the first and second shifts are paid wage rates w_1 and w_2 respectively. Assume there is a fixed cost to opening the plant and it takes at least \bar{n} employees per shift to produce any output.¹⁰

The plant faces a standard labor contract.¹¹ Given this contract, the plant manager must choose how many days to operate the plant, how many shifts to run, how many hours to run each shift, and how many workers to employ on each shift, to minimize the total cost of producing q . Formally, the manager wishes to:

$$\begin{aligned} \min_{D,S,h,n} \quad & (w_1 + I(S=2)w_2)Dhn + \max[0, 0.85(w_1 + I(S=2)w_2)(40 - Dh)n] \\ & + \max[0, 0.5(w_1 + I(S=2)w_2)D(h - 8)n] + \delta \end{aligned}$$

⁹The one exception is Bramalea. When Chrysler purchased American Motors from Renault, Chrysler agreed to build a minimum number of Premiers and Monacos (using Renault parts) at Bramalea. Weak sales of these two vehicle lines forced Chrysler to offer deep discounts eventually. Consequently the volatility of sales for these two vehicle lines is large.

¹⁰The production function in this model differs from the one studied by Lucas (1970), Mayshar and Halevi (1991) and Bils (1992) in two ways. In this model, the same number of employees work each shift and the production function is generalized to allow for overhead labor. Allowing the number of employees to vary across shifts implies counter-factually that the line speed differs across shifts.

¹¹I assume the wage schedule from the labor contract is allocative.

subject to:

$$q \leq DSh(k^{1-\alpha}(n - \bar{n})^\alpha)$$

where $I(S = 2)$ is an indicator function. The parameter α is between 0 and 1. The first term in the objective function represents the straight-time wage paid to workers on both shifts. The second term captures the 85 percent rule for short-weeks, and the third term captures the overtime premium. The fourth term, δ , is a fixed cost to opening the plant. This example ignores benefits and other fixed payments to employees.

Note that production is linear in total hours worked but curved over employment. Without either the 85 percent rule for short-weeks or the requirement that at least \bar{n} employees work each shift, it would always be optimal to run both shifts since the marginal product of labor approaches infinity as $n - \bar{n}$ approaches zero. However in the presence of these fixed costs, the plant can produce low levels of output cheaper with a single shift than with two shifts.

From the discussion in section 2, it is straightforward to assign values to a subset of the parameters. The average day-shift wage at an automobile assembly plant is \$18 per hour, and evening-shift workers are paid a 5 percent premium; this implies that $w_1 = 18.00$ and $w_2 = 18.90$. Bounds can be placed on some of the manager's choice variables. The time period in this example is one week, so D can take on any integer between 0 and 7 inclusive. Most plants run either one or two shifts; so S equals 1 or 2. Hours per shift, h , is usually set between 7 and 10.

Let k be normalized to unity and \bar{n} be set to 500. I assume it takes 1,000 production workers to run a shift with a line speed of 50 vehicles per hour, so I set $\alpha = 0.63$. Set $\delta = \$100,000$. The choice of δ will be discussed in more detail below.

To illustrate the role non-convexities in the plant's cost function play in the allocation of labor, consider the following. Set D to 5 and h to 8. The manager now has two margins along which to vary output: the number of shifts and the number of employees (line speed). Conditional on the number of shifts chosen to be run, the plant manager must set employment such that:

$$n(q, S) = \left(\frac{q}{DShk^{1-\alpha}} \right)^{\frac{1}{\alpha}} + \bar{n} \tag{2}$$

in order to produce q . The cost of producing q with S shifts is then:

$$C(q, S) = (w_1 + I(S = 2)w_2)Dhn(q, S) + \delta. \tag{3}$$

The cost curves conditional on one and two shifts, $C(q, 1)$ and $C(q, 2)$ respectively, are plotted in figure 1. Both cost curves are upward sloping, convex, and cross each other once. The plant manager simply chooses to run a single shift if $C(q, 1) < C(q, 2)$ or to run two shifts if $C(q, 1) > C(q, 2)$. Hence the total cost curve for the plant, $TC(q)$, is the envelop of the two cost curves graphed in figure 1. This total cost curve is plotted in figure 2.

It is clear from figure 2 that the plant's total cost curve is non-convex. There is a kink in $TC(q)$ at the value of q such that $C(q, 1)$ is equal to $C(q, 2)$; call this value of q , \bar{q} . There is also a discontinuity between producing zero and producing ϵ . Over the subintervals (ϵ, \bar{q}) and (\bar{q}, ∞) , $TC(q)$ is still convex. The non-convexities are caused by the fixed costs associated with opening the plant and opening a second shift.

Both $C(q, 1)$ and $C(q, 2)$ individually imply standard U-shaped average cost curves. However $TC(q)$, with its kink at \bar{q} , implies a 'double-U' shaped average cost curve. See figure 3. Similarly, both $C(q, 1)$ and $C(q, 2)$ individually imply upward sloping marginal cost curves; but because of the kink in $TC(q)$ at \bar{q} , the marginal cost curve is discontinuous. See figure 4.¹²

The hours-per-shift versus the shifts-per-day margin can be studied in a similar fashion. Set D to 5 and n to 1,500. Hence the manager can now adjust the number of shifts, S , or the hours per shift, h . Conditional on the number of shifts run, the plant manager must set the hours per shift such that:

$$h(q, S) = \frac{q}{DSk^{1-\alpha}(n - \bar{n})^\alpha} \quad (4)$$

in order to produce q . So the cost of producing q goods while operating a single shift is:

$$\begin{aligned} C(q, 1) = & w_1 Dh(q, S)n + \max[0, 0.85w_1(40 - Dh(q, S))n] \\ & + \max[0, 0.5w_1D(h(q, S) - 8)n] + \delta. \end{aligned}$$

And the cost of producing q goods while operating two shifts is:

$$\begin{aligned} C(q, 2) = & (w_1 + w_2)Dh(q, S)n + \max[0, 0.85(w_1 + w_2)(40 - Dh(q, S))n] \\ & + \max[0, 0.5(w_1 + w_2)D(h(q, S) - 8)n] + \delta. \end{aligned}$$

The cost curves conditional on one and two shifts, $C(q, 1)$ and $C(q, 2)$ respectively, are plotted in figure 5. As in the previous exercise, both cost curves are upward sloping and cross each other

¹²The units on the y-axis differ across figures 1 to 4.

once. So the total cost curve for the plant, $TC(q)$, is the envelop of the two individual cost curves and is plotted in figure 6.

In figure 6, the total cost curve is not differentiable at four points. First, the 85 percent short-week rule and the fixed cost to opening the plant cause a discontinuity at zero. Second, the required overtime premium causes kinks at points A and C . Finally there is kink at the point where $C(q, 1) = C(q, 2)$. Call this point B . Let the origin be denoted by O . As in the previous example, kinks in the total cost curve cause discontinuities in the marginal cost curve and multiple local minima in the average cost curve. See figures 7 and 8.

These non-convexities can be exploited to lower the plant's costs. From figure 6 - 8 one can see for any value of π between 0 and 1,

$$\pi TC(O) + (1 - \pi)TC(q(C)) \leq TC(\pi O + (1 - \pi)q(C)).$$

Thus, a plant manager who must produce q such that $O < q < q(C)$ would ideally like to take a linear combination of producing O and producing $q(C)$. Following such a strategy would lower the plant's total cost and make production more volatile than sales.¹³ If this is possible, the plant manager would never produce in the region $0 < q < q(C)$. The manager's incentive to exploit Jensen's inequality motivates the need to model the manager's problem as a dynamic problem and introduce inventories.

If the manager must produce q such that $q > q(C)$, then the plant operates on a convex portion of the cost curve. Indeed the marginal cost curve is flat in this region (holding employment fixed). See figure 8. In such a region there is no incentive to make production more volatile than sales. From this intuition, it is not surprising that the assembly plants which produced the most vehicles per month (Jefferson North, St. Louis II, and Windsor) use overtime extensively and rarely shut down for inventory adjustment. They are also the plants for which the standard deviation of production is about equal to the standard deviation of sales. See table 4.

From looking at figure 6 it is not obvious that the line segment \overline{OC} convexifies $TC(q)$. It is not clear that the point A is above the line segment \overline{OC} . So when is $TC(q)$ convexified by the single line segment \overline{OC} ? And when is $TC(q)$ convexified by the two line segments \overline{OA} and \overline{AB} ? This is equivalent to asking: under what conditions is slope of \overline{OC} less than the slope of \overline{OA} ?

¹³An analogous argument can be made for the previous example.

The answer is when:

$$(w_2 - w_1) \cdot D \cdot h \cdot n < \delta. \tag{5}$$

Plugging in reasonable numbers yields:

$$\begin{aligned}
 (\$18.90 - \$18.00) \cdot 5 \cdot 8 \cdot 1000 &< \delta \\
 \$36,000 &< \delta.
 \end{aligned}$$

So this simple static model does imply some restrictions on the data. Consider a simple multi-period problem with no costs of holding inventories. Suppose a plant must produce four shifts worth of output in three weeks. The manager will choose to operate two shifts for two weeks and close down for the third week if $\delta > \$36,000$. If $\delta < \$36,000$, the plant will run two shifts one week and a single shift for two weeks. One can see from table 2 that shift changes rarely occur, but plants are often completely shutdown for a week at a time. This suggests that the fixed cost to opening the plant, δ , is large.

5 The dynamic model

The above discussion appeals to the plant manager’s ability to exploit Jensen’s inequality without formally discussing a multiperiod model. This section formulates a dynamic programming model of an automobile assembly plant. As in the static example, the manager in the dynamic model controls the plant’s labor allocation (and thus production) to minimize the expected discounted cost of production subject to technological constraints and the nonlinear price schedule for labor.

5.1 The dynamic program

Consider a plant which produces q_t output at time t . As in the static model, the plant manager has four margins along which to adjust output each period: the number of days the plants is open, the number of shifts run, the length of each shift, and rate of output per unit of time. Let D_t denote the number of days the plant is open. Let S_t denote the number of shifts that are run. Let h_t denote the number of hours each shift runs. Finally let n_t denote the number of employees who work each shift.

Line speed is Cobb-Douglas in capital and the number of employees at work in excess of \bar{n}_2 . Thus \bar{n}_2 is the number of overhead workers it takes to run a shift. Let k_t be the time t capital

stock. So output produced during period t is:

$$q_t = D_t S_t h_t [k_t^{1-\alpha} (n_t - \bar{n}_2)^\alpha] \quad (6)$$

where $0 \leq \alpha \leq 1$.

The total number of workers the plant has on its payroll at time t is $X_t n_t + \bar{n}_1$. Let \bar{n}_1 denote the number of non-production workers (e.g. engineers, administrative personnel) at the plant. Non-production workers are paid a fixed wage each period and are never laid off. Let X_t denotes the number of shifts of production workers the plant has hired. So $X_t n_t$ are the total number of production workers hired. Individual production workers can only work one shift. Production workers on the payroll who do not work either shift receive unemployment compensation. This unemployment compensation is charged directly and immediately to the firm.

I impose the following restriction:

$$S_t n_t + \bar{n}_1 \leq X_t n_t + \bar{n}_1. \quad (7)$$

In words, the total number of employees working must be less than or equal to the number of employees on the payroll. Each period the manager chooses the number of workers to have on the payroll next period. There is a fixed cost to changing $X_{t+1} n_{t+1}$. This fixed cost is set large enough so that once a line speed is chosen, the manager will never wish to change it. This fixed cost also ensures that the firm pays unemployment compensation to employees who do not work either shift.

The plant faces sales each period of s_t . Assume s_t takes on one of three discrete values and evolves according to a first-order Markov chain,

$$\chi(s, s') = \text{Prob}\{s_{t+1} = s', s_t = s\} \text{ for } s, s' \in S = \{s_{\text{high}}, s_{\text{medium}}, s_{\text{low}}\}.$$

Unsold output can be inventoried without depreciation. Let i_{t+1} be the stock of finished goods inventoried at the end of period t carried over into period $t + 1$. Feasibility then requires that:

$$q_t + i_t \geq s_t + i_{t+1}. \quad (8)$$

Inventories cannot be negative:

$$i_{t+1} \geq 0. \quad (9)$$

Assuming the plant's labor contract is of the form described in section 2, the plant's time t cost function is:

$$\begin{aligned}
C(t) = & (I(S_t \geq 1)w_1 + I(S_t \geq 2)w_2 + I(S_t = 3)w_3)D_t h_t n_t \\
& + \max[0, 0.85(I(S_t \geq 1)w_1 + I(S_t \geq 2)w_2 + I(S_t = 3))](40 - D_t h_t) n_t \\
& + \max[0, 0.5(I(S_t \geq 1)w_1 + I(S_t \geq 2)w_2 + I(S_t = 3))D_t (h_t - 8)n_t] \\
& + u w_1 40(X_t - S_t)n_t + \gamma I(X_t n_t \neq X_{t+1} n_{t+1}) + \delta I(D_t > 0) + 40 w_1 \bar{n}_1,
\end{aligned} \tag{10}$$

where w_1 , w_2 , and w_3 are the hourly wage rates paid to the first-shift, second-shift and third-shift workers, respectively. I let u denote the fraction of the 40-hour day-shift wage charged to the firm per idle employee. So the first term represents the straight time wages paid to the production workers. The second and third terms capture the 85 percent rule for short-weeks and the required overtime premium, respectively. The fourth term is the unemployment compensation bill charged to the firm. The fifth term is a fixed cost of adjusting the line speed and the size of the payroll. The sixth term denotes the fixed cost to opening the plant. The last term (seventh) are the wages paid to the plant's non-productive workers. This last term is a constant and has no effect on the manager's allocation of labor. Recall that $I(\cdot)$ are indicator functions. To simplify the notation, assume $D_t = 0$ if and only if $S_t = 0$.

The plant manager's problem is to minimize the present value of the discounted stream of costs given a constant real risk free interest rate, r . Assume the stock of capital, k_t , is fixed at \bar{k} for all t . The manager's problem is then to choose a set of stochastic processes $\{X_{t+1}, i_{t+1}, n_{t+1}, D_t, S_t, h_t\}_{t=0}^{\infty}$ to minimize:

$$E \sum_{t=0}^{\infty} \left(\frac{1}{1+r}\right)^t C(t) \tag{11}$$

subject to (6) - (9) and given $\{X_0, i_0, n_0\}$.

This minimization problem is split into an intra-period problem and an inter-period problem. The intra-period problem is as follows. For each realization of $\{X_t, i_t, n_t, s_t, X_{t+1}, i_{t+1}, n_{t+1}\}$ the firm chooses the feasible set, $\{D_t, S_t, h_t\}$, that minimizes (10). Let:

$$\mathcal{C}(X_t, i_t, n_t, s_t, X_{t+1}, i_{t+1}, n_{t+1}) = \min_{D_t, S_t, h_t} C(t) \quad \text{subject to (6), (7) and (8).}$$

The inter-period problem is then solved by dynamic programming. Let $V(X, i, n, s)$ be the optimal value function for the plant that has X shifts of n employees on the payroll, carries inventories i

into the period, and faces sales s . Thus, the plant's Bellman equation can be written:

$$V(X, i, n, s) = \min_{X', i', n'} \left\{ \mathcal{C}(X, i, n, s, X', i', n') + \frac{1}{1+r} \sum_{s'} \chi(s, s') V(X', i', n', s') \right\} \quad (12)$$

subject to (9). The solution to this Bellman equation yields time invariant decision rules.

5.2 Parameter values

The time period in the dynamic model is one week. The interest rate r is set such that $(1+r)^{-1}$ equals 0.999; this corresponds to a 5 percent annual rate. I set the capital stock, \bar{k} , to 1.0.

I estimated the parameters, α , \bar{n}_1 , and \bar{n}_2 with non-linear least squares using the data on line speed and employment for each plant. The line speed data is weekly and the employment data is the average number of paychecks written each month. After talking with Chrysler, I made the following assumptions: during weeks when the plant is open and during holiday weeks, $S_{it}n_{it} + \bar{n}_{1i}$ workers are paid; during inventory adjustment weeks, supply disruption weeks and long term closure weeks, \bar{n}_{1i} employees are paid; during model changeover weeks, $\bar{n}_{1i} + \bar{n}_{3i}$ employees are paid. So \bar{n}_{3i} represents the number of employees, above and beyond \bar{n}_{1i} it takes to perform a model changeover at plant i . These assumption imply:

$$E_{it} = \frac{1}{WK_t} [(OP_{it} + HL_{it})(S_{it}n_{it} + \bar{n}_{1i}) + (IA_{it} + SD_{it} + LT_{it})\bar{n}_{1i} + (MC_{it})(\bar{n}_{1i} + \bar{n}_{3i})]$$

where

- E_{it} = average number of paychecks written during month t at plant i ,
- S_{it} = number of shifts plant i ran during month t ,
- WK_t = number of weeks (pay-periods) during month t ,
- OP_{it} = number of open weeks during month t at plant i ,
- HL_{it} = number of holiday weeks during month t at plant i ,
- IA_{it} = number of inventory adjustment weeks during month t at plant i ,
- SD_{it} = number of supply disruption weeks during month t at plant i ,
- LT_{it} = number of long-term closure weeks during month t at plant i .

Using equation (6) to eliminate n_{it} yields:

$$E_{it} = \bar{n}_{1i} + \frac{OP_{it} + HL_{it}}{WK_t} \bar{n}_{2i} + \frac{MC_{it}}{WK_t} \bar{n}_{3i} + \frac{(OP_{it} + HL_{it})S_{it}}{WK_t} \left(\frac{LS_{it}}{\bar{k}^{1-\alpha_i}} \right)^{\frac{1}{\alpha_i}} \quad (13)$$

where LS_{it} is the average line-speed for the month at plant i . I estimated \bar{n}_{1i} , \bar{n}_{2i} , \bar{n}_{3i} , and α_i , using equation (13) for all but four of the plants. The point estimates are presented in table 5; standard errors are reported in parentheses.

Four of the plants in the sample underwent some form of major investment during the time period I studied: Belvidere, Bramalea, and Sterling Heights switched vehicle lines; and Dodge City was down for nine weeks in 1993 for a major re-tooling. See table 1. For these four plants, I allowed \bar{n}_{1i} to vary across the two sub-periods. Specifically, I estimated

$$E_{it} = \bar{n}_{1i} + \bar{n}_{4i}DUM_{it} + \frac{OP_{it} + HL_{it}}{WK_t}\bar{n}_{2i} + \frac{MC_{it}}{WK_t}\bar{n}_{3i} + \frac{(OP_{it} + HL_{it})S_{it}}{WK_t} \left(\frac{LS_{it}}{k^{1-\alpha_i}} \right)^{\frac{1}{\alpha_i}}$$

where DUM_{it} is equal to zero during the first sub-period, and is equal to one during the first sub-period. The results are also presented in table 5.

The employment data was incomplete. I did not have employment data for the Toledo plants. For six of the plants, seven months of employment data during 1990 are missing. Furthermore the sample period is short. So some caution is in order when interpreting any single point estimate. Nevertheless, most of the point estimates seem reasonable.

For all but two of the plants, the point estimates of the curvature parameter, α , are between 0.5 and 1.0. And most of the point estimates of α are within two standard errors of 0.64, the usual estimate of labor's share. On average, the point estimates of \bar{n}_1 and \bar{n}_4 imply that about one-third of the workers at these plants are non-production workers. The point estimates of \bar{n}_2 imply that at the average plant about one-half of the production workers on a shift are overhead workers. Of course there is considerable variation in the parameter estimates across the plants.

I estimated the sales processes for each of the single-source plants by assuming that weekly sales follow an AR(1). I estimated the AR(1) parameters by maximum likelihood. Since the sales data are monthly, I assumed weekly sales were a latent variable and used the Kalman filter to construct the likelihood function. For each plant, I used Tauchen's (1986) method to compute weekly three-state Markov chains whose sample paths approximate those of the estimated AR(1) processes. For each Markov chain, the grid width was chosen to match the standard deviation of actual sales process. The grid points were then rounded to make them compatible with the inventory grid. To conserve on space, the estimated Markov chains are not reported but are available from the author.

Following the discussion in the second section, wage rates are set as: $w_1 = \$18.00$ per hour,

$w_2 = \$18.90$ per hour. The per idle employee fee for unemployment compensation, u , is set to 0.65 for the U.S. plants. For the Canadian plants, I set $u = 0.40$.

The parameter γ is set to large enough so that changes in employment per shift (line speed) and changes in the number of shifts hired are rarely made; but γ is set small enough so that the terminal level of employment does not effect the decision rules. I allow γ to vary across the plants between \$1.0 million and \$3.0 million.

There is one remaining free parameter, δ , the fixed cost of opening the plant for the week. As discussed in section 4, the fixed cost to opening a two-shift plant each week must be large. If the firm is operating in the non-convex region of its cost curve and the fixed cost is small, then the model will predict that the manager will open and close the second shift rather than open and close the entire plant. So I set δ to \$1.0 million.

But what is this fixed cost, δ ? There are some fixed costs to opening the plant: warming up the equipment, and heating the shop floor. Discussions from industry sources indicate that it is considerably easier to control many of these costs, particularly energy costs, by shutting down for a week at a time rather than sending a single shift home. Additionally, managers usually encourage salaried workers to take vacation when the plant is shutdown. Thus the firm can avoid having key workers on vacation when the plant is running.

But there may be other factors besides the fixed costs that influence the manager's decision whether to shut down the plant or just lay off a single shift. The union contract dictates a strict hierarchy concerning who gets laid off before whom. By laying the entire work force off, the firm treats all the workers equally – thus saving the firm the cost of figuring out who to lay off and who to not.¹⁴ More generally, if the workers face diminishing marginal utility in leisure, then the workers and the firm may prefer a complete one-week shutdown over the firm sending the second shift home for two weeks. While these other factors are credible, the model assumes workers are homogeneous and is silent on worker preferences.

Using the parameter values selected above, the intra-period problem is solved via grid search. The grids for D_t and S_t are set from 0 to 6 and from 0 to 3, respectively, in increments of 1. The plant is closed for the week whenever $S_t = 0$ or $D_t = 0$. Recall, $S_t = 0$ if and only if $D_t = 0$. The shift length, h_t , can take on values of 7, 8 or 9. So there are 84 grid points to evaluate for each $\{X_t, i_t, n_t, s_t, X_{t+1}, i_{t+1}, n_{t+1}\}$ sept-tuple.

¹⁴See Aizcorbe's (1990) discussion of the UAW contract with Ford Motor Company.

To make the inter-period problem a finite state, discounted dynamic program, the state space is discretized. The number of shifts of workers on the payroll, X_t , can take on values of 1, 2, or 3. For each plant, there are 26 levels of shift employment. The employment grid for each plant is set so that the line speed can vary between 0 and 100 vehicles per hour. In order to conserve on grid points, the inventory grid is also allowed to vary across plants. For each plant, inventories can take on 31 points from 0 to $2 \times s_{high}$. The inter-period problem is solved by iterating on the Bellman equation, (12). Once the Bellman equation is solved, the transition matrix and the invariant probability distribution for the state space are computed. The state space is checked to be ergodic. Using the invariant probability distribution and the decision rules, a wide variety of population moments can be computed.

5.3 Results for nine plants

In this subsection, I use the production and sales parameters estimated above to solve the model plant-by-plant. A set of the model's predictions for each plant is reported in table 6. The corresponding moments in the data are also reported.

The first column of table 6 reports the model's prediction for the ratio of the monthly standard deviation of production to the monthly standard deviation of sales.¹⁵ The second and third columns report the model's predictions for the average workweek of capital conditional on the plant being open and unconditionally. The predicted number of shifts run and line speed are reported in columns 4 and 5, respectively. Columns 6 and 7 report the unconditional probabilities that the plant is closed for an inventory adjustment or the plant is running overtime.

For four plants (Jefferson North, St. Louis II, Windsor, and Brampton), the solution to the dynamic programming problem implies implausibly high line speeds (over 80 vehicles/hour). Therefore at these plants I constrained the feasible line speeds to be less than or equal to the plant's average line speed. For each of these four plants, the constrained solution implies running the line at the fastest feasible speed. For the remaining five plants, the predicted line speeds are in the interior of the feasible set.

Given the estimated parameter values, the model replicates the three fact described above. The model predicts that at the high-sales plants, production is less variable than sales, the average workweek of capital is over 85 hours, and overtime is frequently employed. For the medium- and

¹⁵In the model, a month is 13/3 weeks.

low-sales plants, the model predicts that production is more variable than sales, the plants use lower levels of capital utilization, and weeklong shutdown are frequent. The model also correctly predicts the number of shifts used at each plant.

Of course the dynamic model is too simple to match all the features of the data.

1. At all but two of the plants, the model under-predicts the standard deviation of monthly production. In the data, weeklong shutdowns tend to be bunched together; the average duration of any type of weeklong closure (except supply disruptions) is greater than one week. In the model, the duration of a shutdown at most plants is one week. Since I aggregate each plant's output to the monthly frequency, the effect of these single-week shutdowns tends to wash out. The duration of the shutdowns is longer in the data, so their influence is not dampened as much by time aggregation.

Adding a desired inventory-to-sales ratio target to the plant's cost function (equation 10) can increase the implied duration of the weeklong shutdowns at these plants; this in turn implies an increase in the standard deviation of monthly production. Furthermore adding an inventory-to-sales ratio target to the model can generate the inventory accumulation observed in the data when sales increase. The work of Blanchard (1983) and Kashyap and Wilcox (1993), as well as the auto industry's great interest in days-supply inventory data, suggest that automakers target such a ratio. Equation (10) does not include such a target to isolate the effect the non-convex margins play in production scheduling.

2. With exception of Belvidere, the model over-predicts the use of overtime at the low- and medium-sales plants. It may be that the non-linear price schedule for labor is not allocative; there may be additional constraints on the plant manager that limit the hours of overtime employees can work. Or it could be that production is not linear in hours; for example, as the shift length increases, mistakes may increase as workers tire.
3. At five of the plants (Jefferson North, St. Louis II, Windsor, Brampton, and Bramalea), the model dramatically over-predicts the line speed. This suggests that equation (6) may be a poor approximation to the true production technology; given the estimated parameters there is too little curvature in employment to rationalize the observed line speed at these five plants.

4. In the model, the only reason the plant ever shortens the workweek is to reduce inventories. This assumption causes the model to ignore other states identified in the data for which the plant might be shut down for all or part of the week. Thus the model is silent about holidays, model changeovers, and supply disruptions.

These limitations suggest some natural extensions to the analysis.

Nevertheless the analysis illustrates that much of the heterogeneity in the production behavior across the nine plants can be explained by a simple dynamic programming model. Of course, since I allow several of the parameters to vary across the plants, it is not clear how much of this cross-plant variation is due to difference in the parameter estimates and how much is due to the differences in the mean of the sales processes. So in this following subsection, I resolve the model holding fixed all the parameters and varying just the sales rate.

5.4 Results assuming a deterministic sales process

In this subsection, the sales process is deterministic; the transition matrix, χ , is the scalar 1. I set α , n_1 , and n_2 to the point estimates for the Belvidere plant. I set δ to \$1.0 million, and u to 0.65. The weekly sales rate varies from 200 to 7000 in increments of 200. The employment and inventory grids are fixed throughout this exercise.

I solve the dynamic model at each sales rate. For each sales rate, I compute the average workweek of capital, the standard deviation of monthly production, and the total cost of production. The total cost for a given sales rate, $TC(s)$, is the sum of the value function at each state, $V(X, i, n|s)$, weighted by the unconditional probability of each state, $\lambda(X, i, n|s)$. I multiply $TC(s)$ by $(1 - \beta)$ to make the units compatible with the static example. I trace out a “long-run marginal cost” curve for the plant by computing the one-sided derivative of the total cost curve. More precisely,

$$MC_{lr}(s) = \frac{(1 - \beta) \sum_{X,i,n} \lambda(X, i, n|s + 200) V(X, i, n|s + 200) - (1 - \beta) \sum_{X,i,n} \lambda(X, i, n|s) V(X, i, n|s)}{200}$$

where $MC_{lr}(s)$ denotes the long-run marginal cost at sales rate, s .

To compute a “short-run” cost curve, I fix the line speed, n , and number of shifts hired, X , at their optimal levels for a sales rate of 3600 vehicles per week (the average rate at Belvidere). Thus the plant manager can only manipulate “short-run” margins: i_t , D_t , S_t , and h_t . Bresnahan and Ramey (1994) provide evidence that line speed and shift changes are associated with permanent

changes in output while changes in the shift length and week-long shutdowns are associated with temporary changes in output.¹⁶ I then repeat the above exercise.¹⁷

Consider the long-run analysis displayed in figures 9 - 12. Figure 9 illustrates that the plant's cost total curve has both a concave and a convex region. Consequently the marginal cost curve, plotted in figure 10, is U-shaped: it is downward sloping when sales are less than 3200 vehicles per week; it is essentially flat in the region, $3200 < \text{sales rate} < 4400$; and it becomes upward sloping when sales are greater than 4400 vehicles per week. Thus the Belvidere plant operates in a region of increasing long-run marginal costs only when sales are high.

The concave region in the total cost curve occurs even though the manager has the ability to manipulate inventories to exploit some of the non-convexities in the cost minimization problem. Two factors imply this concavity. First the production function, equation (6), does not exhibit constant returns-to-scale; it takes at least \bar{n}_2 overhead workers to run a shift. Second the unemployment insurance provision and the 85% short-week rule make it relatively expensive for the plant to operate at low levels. The marginal savings of laying off a worker for one week is just $40(1 - u)w_1$; the marginal savings of reducing a worker's workweek by one hour is just $(1 - .85)w_1$. The combination of the minimum number of workers needed to produce and the high costs associated with idling these workers imply a downward sloping marginal cost curve at low levels of output. It is reassuring that the Belvidere plant sold on average about 3600 vehicles a week, the nadir of the long-run marginal cost curve.¹⁸

The model predicts that at low- and medium-sales levels, the plant manager primarily changes the frequency of weeklong shutdowns, a non-convex margin, to vary output. If the sales rate is below 5200, the optimal strategy is for the plant to produce for several weeks and build up an inventory stock equal to one week of sales; the plant then shuts down for a week and inventories fall to zero. Consequently, the plant manager chooses to make production volatile – despite the fact that the sales rate is constant. See figure 12. Furthermore the optimal strategy implies that capital often sits idle for a week at a time; note that in figure 11 the difference between the

¹⁶Even though the changes in the sales rate are permanent, I refer to the computed cost curves as “short-run” cost curves since I fix the “long-run” margins.

¹⁷It would be interesting to study the simultaneous movements of price and marginal cost as functions of the sales rate. Thus one could study the relative size and cyclical movements of mark-ups. However, I was unable to obtain transaction price data for any of the vehicles produced at these plants.

¹⁸Chrysler executives told me that when sales fall, they first adjust the price of vehicle (e.g. rebates, dealer-incentives) to try to increase sales; if demand is not sufficiently elastic, they then adjust output. Such a strategy is consistent with a U-shaped marginal cost curve.

unconditional average workweek on capital and the workweek of capital conditional on the plant being open (the vertical difference between the dotted and solid lines) is quite large.

The dynamic model also captures the fact that when sales are high the plant manager primarily manipulates the shift length, a convex margin, to vary output. When sales are greater than 5200 vehicles per week, the plant is never closed for a week at a time; as figure 11 shows, the conditional and unconditional average workweeks of capital are equal (and over 80 hours) when sales are high. Overtime is used extensively. Moreover the implied time series on production is relatively smooth; figure 12 illustrates that the standard deviation of monthly production falls when the sales rate rises above 5200.

When I fix the two long-run margins, the total cost curve becomes piece-wise linear.¹⁹ See figure 13.²⁰ When sales are below 4600, the plant manager changes hours worked (and thus output) by changing the number of weeks the plant operates. When sales are above 4600 the plant manager changes hours worked by changing the shift length. Hence the marginal cost curve, plotted in figure 14, is not U-shaped; it is flat with one discontinuous jump. The average workweek of capital becomes just a linear function of sales (figure 15). Figure 16 illustrates that the standard deviation of monthly production falls when the plant varies output using overtime rather than weeklong shutdowns.

A simple dynamic programming model with credible non-convex margins of adjustment can capture much of the heterogeneity in the production behavior observed across a set of automobile assembly plants. The model attributes the differences across plants in capital utilization and relative volatility of production and sales to differences in the level of sales. High sales imply that the plant is operating in a convex region of the cost curve, while low and medium sales imply that the plant is operating in a non-convex region of the cost curve.

The model captures the fact that plants with low and medium sales often use weeklong shutdowns, a non-convex margin, to vary output. Thus the model can explain why production at these plants is more volatile than sales and why capital at these plants is often idle. At the same time, the dynamic model captures the fact that plants which produce high-selling vehicle lines primarily use convex margins of adjustment such as overtime employment to vary output; therefore the model can explain why production at high-sales plants varies by about as much as

¹⁹Recall the production function, equation (6), is linear in hours worked.

²⁰Given the fixed line speed, the plant can produce at most 6000 vehicles in one week.

sales and why capital at these plants rarely sits idle. Thus, the model succeeds in reconciling the three facts documented in the third section.

6 Concluding remarks

The paper focuses on understanding the high-frequency production behavior of a small set of automobile assembly plants. Thus this paper trades generality for precise data. But the non-convexities identified in this paper are not unique to automobile assembly plants. Managers at most manufacturing plants that produce-to-stock face these same non-convex margins: how many shifts to run and whether to open or close the plant each week. Thus the results of this paper may apply to other industries.²¹

It is unclear whether the important role non-convexities play at the plant level do not just wash out at the aggregate level. However there is evidence that production decisions are not independent across plants and firms. Automobile assembly plants are just one component of a large network of suppliers and dealers. The work of Beaulieu and Miron (1991) and Cooper and Haltiwanger (1992) provide evidence that in the presence of strategic complementarities, multiple firms synchronize output. These papers suggest that the dramatic high frequency variations in output observed at the plant level may not be completely smoothed out by modest aggregation.

References

- [1] Aizcorbe, Ana (1990) “Experience Rating, Layoffs and Unions: A Look at U.S. Auto Assembly Plant Layoffs,” manuscript, Bureau of Labor Statistics.
- [2] Aizcorbe, Ana (1992) “Procyclical Labour Productivity, Increasing Returns to Labour and Labour Hoarding in Car Assembly Plant Employment,” *Economic Journal*, Vol. 102, (July), pp. 860-873.
- [3] Aizcorbe, Ana, and Sharon Kozicki (1995) “The Comovement of Output and Labor Productivity in Aggregate Data for Auto Assembly Plants,” working paper 95-33, Board of Governors of the Federal Reserve System.

²¹However the work of Cecchetti, Kashyap, and Wilcox (1994) suggests that the transportation sector may not be representative of all manufacturing.

- [4] Anderson, Patricia, and Bruce Meyer (1993) "Unemployment Insurance in the United States: Layoff Incentives and Cross-Subsidies," *Journal of Labor Economics*, Vol. 11 (January), pp. S70-S95.
- [5] Beaulieu, J. Joseph, and Jeffrey Miron (1991) "The Seasonal Cycle in U.S. Manufacturing," *Economic Letters*, Vol. 37, pp. 115-118.
- [6] Berry, Steven, James Levinsohn, and Ariel Pakes (1995) "Automobile Prices in Market Equilibrium," *Econometrica*, Vol. 63, (July), pp. 841-890.
- [7] Bils, Mark (1992) "Measuring Returns to Scale From Shift Practices in Manufacturing," manuscript, University of Chicago.
- [8] Blanchard, Olivier (1983) "The Production and Inventory Behavior of the American Automobile Industry," *Journal of Political Economy*, Vol. 91, (June), pp. 365-400.
- [9] Blanchard, Olivier, and Angelo Melino (1986) "The Cyclical Behavior of Prices and Quantities: The Case of the Automobile Market," *Journal of Monetary Economics*, Vol. 17, (May), pp. 379-407.
- [10] Blinder, Alan, and Louis Maccini (1991) "Taking Stock: A Critical Assessment of Recent Research on Inventories," *Journal of Economic Perspectives*, Vol. 5, (Winter), pp.73-96.
- [11] Bresnahan, Timothy (1981) "Departures from Marginal-Cost Pricing in the American Automobile Industry: Estimates for 1977-78," *Journal of Econometrics*, Vol. 17, (November), pp. 201-227.
- [12] Bresnahan, Timothy, and Valerie Ramey (1994) "Output Fluctuations at the Plant Level," *Quarterly Journal of Economics*, Vol. 109, (August), pp. 593-624.
- [13] Cecchetti, Stephen, Anil Kashyap, and David Wilcox (1994) "Do Firms Smooth the Seasonal in Production in a Boom? Theory and Evidence," manuscript, University of Chicago.
- [14] Cooper, Russell, and John Haltiwanger (1992) "Macroeconomic Implications of Production Bunching: Factor Demand Linkages," *Journal of Monetary Economics*, Vol. 30, (October), pp. 107-127.

- [15] Cooper, Russell, and John Haltiwanger (1993) "The Aggregate Implications of Machine Replacement: Theory and Evidence," *American Economic Review*, Vol. 83, (June), pp. 360-382.
- [16] Hamermesh, Daniel (1989) "Labor Demand and the Structure of Adjustment Costs," *American Economic Review*, Vol. 79, (September), pp. 674-689.
- [17] Kashyap, Anil, and David Wilcox (1993) "Production and Inventory Control at the General Motors Corporation During the 1920's and 1930's," *American Economic Review*, Vol. 83, (June), pp. 383-401.
- [18] Lucas, Robert (1970) "Capacity, Overtime, and Empirical Production Functions," *American Economic Review Papers and Proceedings*, Vol. 60, (May), pp. 23-27.
- [19] Matthey, Joe and Steve Strongin (1995) "Factor Utilization and Margins for Adjusting Output: Evidence from Manufacturing Plants," working paper 95-12, Board of Governors of the Federal Reserve System.
- [20] Mayshar, Joram, and Yoram Halevi (1992) "On the Time Pattern of Production: Shiftwork, Hours of Work, and Factor Employment," manuscript, Hebrew University.
- [21] Ramey, Valerie (1991) "Nonconvex Costs and the Behavior of Inventories," *Journal of Political Economy*, Vol. 99, (April), pp. 306-334.
- [22] Shapiro, Matthew (1993) "Cyclical Productivity and the Workweek of Capital," *American Economic Review Papers and Proceedings*, Vol. 83, (May), pp. 229-233.
- [23] Shapiro, Matthew (1995) "Capital Utilization and the Marginal Premium for Work at Night," manuscript, University of Michigan.
- [24] Tauchen, George (1986) "Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions," *Economic Letters*, Vol. 20, pp. 177-181.

Plant	Period (YR:M)	U.S. or Canada	Single Source?	Vehicle Lines
Belvidere	90:1 - 93:5	U.S.	yes	New Yorker Salon, Dynasty, Fifth Ave., Imperial
	93:11 - 94:12		no	Neon
Bramalea	90:1 - 91:12	Canada	yes	Monaco, Premier
	92:6 - 94:12		no	Concorde, LHS, Vision, Intrepid
Brampton	90:1 - 92:4	Canada	yes	Wrangler
Dodge City	90:1 - 93:5	U.S.	yes	Ram Pickup, Dakota
	93:7 - 94:12		no	Ram Pickup, Dakota
Jefferson North	92:1 - 94:12	U.S.	yes	Grand Cherokee
Newark	90:1 - 94:12	U.S.	no	Acclaim, Spirit, Intrepid, LeBaron Sedan
Pillette Road	90:1 - 94:12	Canada	yes	Ram Van, Ram Wagon
St. Louis I	90:1 - 91:5	U.S.	yes	Daytona, LeBaron Coupe
St. Louis II	90:1 - 94:12	U.S.	yes	Grand Caravan, Grand Voyager, Town & Country
Sterling Heights	90:1 - 94:3	U.S.	no	Daytona, Shadow, Sundance
	94:8 - 94:12		no	Cirrus
Toledo I	90:1 - 94:12	U.S.	yes	Cherokee, Commanche, Wagoneer
Toledo II	90:1 - 91:6	U.S.	yes	Grand Wagoneer
	92:7 - 94:12		yes	Wrangler
Toledo III	93:9 - 94:12	U.S.	no	Dakota
Windsor	90:1 - 94:12	Canada	yes [†]	Caravan, Voyager

Table 1: Assembly Plants and Their Vehicle Lines

[†] The Eurostar plant in Austria produced a version of the Voyager beginning in the fourth quarter of 1991 solely for the European market.

Plant	Period (YR:M)	Weeks	Weeks Closed			Short-Weeks		Weeks	Shift	Line Speed	
		Open (1)	HOL (2)	SUP (3)	MC (4)	IA (5)	TOTAL (6)	HOL (7)	with OT (8)	Changes (9)	Changes (10)
Jefferson North	92:1-94:12	154	3	0	4	1	20	20	139	2	15
St. Louis II	90:1-94:12	248	5	0	8	0	27	22	183	1	9
Windsor	90:1-94:12	243	5	0	12	1	25	22	164	1	8
Belvidere	90:1-93:5	141	4	0	12	20	29	26	31	1	6
Brampton/Toledo II	90:1-94:12	219	5	0	9	16	35	33	97	0	10
Dodge City	90:1-93:5	153	4	0	7	12	19	18	53	2	5
Pilette Road	90:1-94:12	202	4	1	24	30	43	39	36	0	17
Toledo I	90:1-94:12	217	6	2	15	21	34	31	58	0	16
Bramalea	90:1-91:12	39	1	0	5	56	12	9	0	0	1
St. Louis I	90:1-91:5	53	1	0	2	17	10	9	1	0	2
Toledo II	90:1-91:6	27	1	1	4	44	8	5	2	0	0
Belvidere	93:11-94:12	56	2	1	1	0	10	9	20	1	3
Bramalea	92:6-94:12	124	3	0	4	0	17	14	85	0	5
Dodge City	93:7-94:12	73	4	0	1	0	7	6	63	2	5
Newark	90:1-94:12	220	4	0	14	23	39	34	90	0	5
Sterling Heights	90:1-94:12	193	5	0	8	33	32	29	56	2	6
Toledo III	93:9-94:12	63	2	0	2	0	8	6	37	0	2
average plant		173.2	4.1	0.4	10.1	19.6	26.8	23.7	79.6	0.8	8.6
percentage of non-LRUN weeks in each state		83.6	2.0	0.2	4.9	9.4	12.9	11.4	38.4		

Table 2: Margins Used by Each Plant

This table reports the number of weeks each plant is open, closed, running a short-week or running overtime. The plant closures are decomposed into five categories: HOL = holiday/vacation; SUP = supply disruption; MC = model changeover; IA = inventory adjustment; LRUN = closed for more than 3 months. The long-run closures are not reported. The data are weekly from 1990:1 to 1994:53; so the total number of weeks in the sample is 261.

Plant	Period (YR:M)	# Shifts Run	Conditional On Open	Conditional On Not LRUN
Jefferson North	92:1-94:12	1,2,3	89.5	85.1
St. Louis II	90:1-94:12	2,3	104.4	99.2
Windsor	90:1-94:12	2,3	94.4	87.9
Belvidere	90:1-93:5	1,2	73.7	58.7
Brampton/Toledo II	90:1-94:12	1,2	59.5	52.3
Dodge City	90:1-93:5	2	81.6	71.0
Pillette Road	90:1-94:12	2	78.0	60.4
Toledo I	90:1-94:12	2	80.7	67.1
Bramalea	90:1-91:12	1	36.3	14.0
St. Louis I	90:1-91:5	1	38.2	27.7
Toledo II	90:1-91:6	1	33.8	12.7
Belvidere	93:11-94:12	1,2	80.4	75.0
Bramalea	92:6-94:12	2	80.3	76.0
Dodge City	93:7-94:12	1,2	92.5	88.9
Newark	90:1-94:12	2	83.2	70.2
Sterling Heights	90:1-94:12	1,2	80.4	64.9
Toledo III	93:9-94:12	1	43.3	40.7
average plant			74.5	63.1
Weighted average			80.0	66.8

Table 3: Average Workweek of Capital (in hours/week)

Plant	Period	Production	Total Sales	Inventories	$\frac{\text{Inventories}}{\text{Total Sales}}$	$\sigma_{\text{production}}$	σ_{sales}	$\frac{\sigma_{\text{production}}}{\sigma_{\text{sales}}}$
Jefferson North	92:1 - 94:12	17490	16858	28196	1.72	6992	7729	0.90
St. Louis II	90:1 - 94:12	20669	20924	41214	2.17	5499	5208	1.06
Windsor	90:1 - 94:12	24193	23230	65406	2.95	6778	6343	1.07
Belvidere	90:1 - 93:5	14600	15575	32644	2.10	5479	3496	1.57
Brampton/Toledo II	90:1 - 94:12	5581	5790	11281	2.23	1962	1599	1.23
Dodge City	90:1 - 93:5	15851	17240	61707	3.69	4991	3502	1.43
Toledo I	90:1 - 94:12	12643	13464	29425	2.13	3943	2236	1.76
Pillette Road	90:1 - 94:12	6567	6603	18850	3.05	2707	1848	1.47
Bramalea	90:1 - 91:12	1783	2178	6781	3.96	1253	1172	1.07
St. Louis I	90:1 - 91:5	6414	7898	27838	3.77	3208	2261	1.42
Toledo II	90:1 - 91:6	445	527	1865	3.90	373	168	2.22
Aggregate	90:1 - 94:12	103826	104420	266265	2.55	22970	17638	1.30

Table 4: Monthly Statistics: Means and Standard Deviations




Figure 1: Cost conditional on running one shift, $C(q, 1)$, and running two shifts, $C(q, 2)$ holding hours per shift fixed.




Figure 2: Total cost allowing either one or two shifts to run, $TC(q)$, holding hours per shift fixed.




Figure 3: Average cost curve allowing either one or two shifts to run, holding hours per shift fixed.




Figure 4: Marginal cost curve allowing either one or two shifts to run, holding hours per shift fixed.




Figure 5: Cost conditional on running one shift, $C(q, 1)$, and running two shifts, $C(q, 2)$, holding employment fixed




Figure 6: Total cost curve allowing either one or two shifts to run, $TC(q)$, holding employment fixed




Figure 7: Average cost curve allowing either one or two shifts to run, holding employment fixed




Figure 8: Marginal cost curve allowing either one or two shifts to run, holding employment fixed

Plant	time period	no. of usable obs.	α	\bar{n}_1	\bar{n}_2	\bar{n}_3	\bar{n}_4	R^2
Belvidere	90:1-94:12	53	0.628 (0.020)	567 (159)	532 (154)	268 (369)	316 (123)	0.76
Bramalea	90:1-94:12	60	0.707 (0.039)	196 (51)	670 (113)	20 (199)	599 (71)	0.93
Brampton	90:1-92:4	27	1.0 [†]	563 (79)	201 (44)	184 (203)		0.48
Dodge City	90:1-94:12	53	0.616 (0.029)	1123 (228)	158 (247)	0 [‡]	1096 (191)	0.41
Jefferson North	92:1-93:5	49	0.620 (0.017)	364 (54)	658 (98)	1588 (348)		0.97
Newark	90:1-94:12	53	0.702 (0.062)	1405 (309)	596 (191)	268 (420)		0.47
Pillette Road	90:1-94:12	60	0.522 (0.018)	703 (140)	80 (137)	254 (170)		0.67
St. Louis I	90:1-91:5	10	0.767 (0.590)	894 (301)	409 (781)	0 [‡]		0.33
St. Louis II	90:1-94:12	53	0.915 (0.358)	970 (231)	1005 (145)	2067 (624)		0.76
Sterling Heights	90:1-94:12	53	1.0 [†]	708 (195)	772 (113)	253 (563)	283 (195)	0.53
Windsor	90:1-94:12	60	0.663 (0.020)	465 (187)	1121 (96)	1664 (288)		0.90

Table 5: Parameter values for the production function

[†] To avoid estimates of α_i greater than 1.0 for Brampton and Sterling Heights, I fixed α_i to be 1.0 prior to estimation.

[‡] To avoid negative estimates of \bar{n}_{2i} for Dodge City and St Louis I, I fixed \bar{n}_{3i} to be zero prior to estimation.

Plant		$\frac{\sigma_{\text{mon. prod.}}}{\sigma_{\text{mon. sales}}}$	Ave. Workweek of Capital cond. on open	unconditional	# of Shifts	average line speed	uncon prob of IA week	uncon. prob of OT week
HIGH-SALES PLANTS								
Jefferson North	<i>model</i>	1.00	89.5	89.3	1,2,3	54.0 [†]	0.5	98.0
	<i>data</i>	0.90	89.5	85.1	1,2,3	53.7	1.0	88.1
St. Louis II	<i>model</i>	0.67	98.8	98.8	2,3	53.4 [†]	0.0	81.0
	<i>data</i>	1.06	104.4	99.2	2,3	54.0	0.0	71.5
Windsor	<i>model</i>	0.89	95.5	88.0	2,3	65.1 [†]	7.8	48.2
	<i>data</i>	1.07	94.4	87.9	2,3	65.0	0.4	63.5
MEDIUM-SALES PLANTS								
Belvidere	<i>model</i>	1.10	80.0	62.3	2	57.1	22.1	0.3
	<i>data</i>	1.57	73.7	58.7	1,2	58.3	11.9	17.0
Brampton	<i>model</i>	1.12	108.0	66.8	2	19.0 [†]	38.1	61.9
	<i>data</i>	1.27	78.8	65.1	2	19.0	12.5	13.3
Dodge City	<i>model</i>	1.08	90.3	76.0	2	50.1	15.8	84.1
	<i>data</i>	1.43	81.6	71.0	2	55.2	7.4	21.1
Pillette Road	<i>model</i>	1.11	108.0	70.2	2	23.8	35.0	65.0
	<i>data</i>	1.47	78.0	60.4	2	25.3	11.5	13.5
LOW-SALES PLANTS								
Branalea	<i>model</i>	1.37	49.9	7.5	1	66.9	85.0	15.0
	<i>data</i>	1.07	36.3	14.0	1	32.1	58.0	3.0
St. Louis I	<i>model</i>	1.13	54.0	27.0	1	65.0	50.0	50.0
	<i>data</i>	1.42	38.2	27.7	1	58.8	23.6	1.4

Table 6: Results for the nine plants

† To avoid implausible line speeds for Jefferson North, St. Louis II, Windsor, and Brampton, I constrained the line speeds for these plants to be less than or equal to the average line speed observed.




Figure 9: The long-run total cost curve.




Figure 10: The long-run marginal cost curve.




Figure 11: Unconditional average workweek of capital (solid line) and the average workweek of capital conditional of the plant being open (dotted line).




Figure 12: Standard deviation of monthly output as a function of the weekly sales rate.




Figure 13: The short-run total cost curve.




Figure 14: The short-run marginal cost curve.




Figure 15: Unconditional average workweek of capital (solid line) and the average workweek of capital conditional of the plant open (dotted line).




Figure 16: Standard deviation of monthly output as a function of the weekly sales rate.