



Working Paper Series

Non-Convex Costs and Capital Utilization: A Study of Production and Inventories at Automobile Assembly Plants George J. Hall

Working Papers Series
Research Department
(WP-96-25)

Federal Reserve Bank of Chicago

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

George J. Hall*
Federal Reserve Bank of Chicago

November 1996

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 about as volatile than sales.

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 a second 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,

*Economic Research Department, Federal Reserve Bank of Chicago, 230 South LaSalle Street, Chicago, IL 60604; e-mail: ghall@frbchi.org. I am grateful to the Chrysler Corporation for providing much of data used in this paper; in particular I thank Van Bussmann, Mike Kuhn, Cary Greenhut and Chuck Pierce. I have benefited from conversations with John Cochrane, Mike Cubbin, Bill Dupor, Martin Eichenbaum, Lars Hansen, Anil Kashyap, Tom Sargent, Bob Schnorbus, and Lew Segal. Pier Deganello provided able research assistance. The views expressed herein are mine and not necessarily those of the Federal Reserve Bank of Chicago or the Federal Reserve System.

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.¹ This fact has generated considerable attention since classic models of inventories, which assume convex

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

short-run increasing marginal costs, imply that firms should manage inventories such that production is smoother than sales.² Although a variety of 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. So why do some plants – but not others – use weeklong shutdowns so frequently to vary output?

Building on the work of Ramey (1991), Cooper and Haltiwanger (1992), 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 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

²Three basic modifications to the classic model of inventories have been proposed. First, many authors assume firms use inventories to both smooth production and target a desired inventory-to-sales ratio (e.g., Blanchard, 1983; West, 1986; and Kashyap and Wilcox, 1993); Kahn (1987, 1992) justifies targeting an inventory-to-sales ratio by explicitly incorporating costly stock-outs. Bils and Kahn (1996) further justify targeting such a ratio by modeling sales as an increasing function of the available inventories. Second, Ramey (1991) argues that firms operate on flat or decreasing regions of their short-run marginal cost curves. Third, authors such as Blinder (1986) and Eichenbaum (1984, 1989) allow for cost shocks; thus inventories are used to smooth production costs rather than the level of production. Once one of these modifications is made to the classic model, there is no a priori reason to expect the variance of sales to exceed the variance of production. Alternatively, Fair (1989) suggests that this anomaly is just a figment of poorly constructed data.

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.³ 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 six different sales processes. The rest of the model is held fixed. 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 selected to match different features of the data. The results of these six exercises demonstrate that much of the variation across plants in capital utilization and relative variability of production and sales can be attributed to the level of the sales process.

The topic of this paper is narrow: how do managers at automobile assembly plants schedule production? Thus this paper trades generality for high data quality. Nevertheless, there are several reasons why weekly plant-level phenomenon may be of interest to economists studying the movements of aggregates at monthly or quarterly frequencies.

Precise data – particularly on capital utilization – matter. Much of the debate concerning procyclical labor productivity and increasing returns to scale centers on how to measure changes in capital utilization; but a paucity of data usually frustrates attempts to empirically evaluate models of the aggregate economy that incorporate capital utilization. So researchers must often resort to proxies for capital utilization or model capital utilization as an unobserved variable.⁴ The absence of good measures of aggregate capital utilization motivates work at the micro level.

Studies using disaggregated data and better measures of capital utilization have generated striking results. Shapiro (1993) finds that once he measures capital services by the workweek of capital, there is no evidence of short-run increasing returns to total factor inputs for the U.S. manufacturing sector. This finding is supported by Burnside, Eichenbaum, and Rebelo (1995) who measure capital utilization rates using electricity data; they also find that cyclical movements

³See 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.

⁴An example of the former is Solow (1957) who uses the employment rate as a proxy for capital utilization. Examples of the latter include Burnside and Eichenbaum (1996), Bils and Cho (1994), and Basu and Kimball (1994).

in capital utilization can explain most of the cyclical variation in labor productivity. Aizcorbe and Kozicki (1995) demonstrate that the observed procyclicality of aggregate labor productivity in the automobile industry at the monthly frequency is generated largely by reported labor hours at plants that are shutdown for a week at a time.

The absence of high-quality aggregate data on capital utilization also makes it difficult to evaluate different assumptions about the costs of increasing the fraction of time the capital stock is worked each period. In Lucas (1970) the firm faces a rising schedule of wage rates as hours worked moves from the most attractive time periods to the least attractive time periods. Similarly, in Burnside, Eichenbaum, and Rebelo (1993) and Bils and Cho (1994) the cost of increasing the workweek of capital is the increased disutility from working more hours.⁵ In Greenwood, Hercowitz, and Huffman (1988), Basu and Kimball (1994), and Burnside and Eichenbaum (1996), the primary cost of a longer workweek of capital is a more rapid depreciation of the capital stock. The current paper is silent about investment and the depreciation of capital. Nevertheless, how and when the hours worked are increased can have big effects on the costs of a longer workweek of capital. This paper documents these costs explicitly.

More generally, modeling production decisions at the plant level provides a concrete framework to study the shape and slope of the marginal cost curve. This is an important issue to macroeconomists since the shape and slope of the cost curve determine the degree to which output fluctuates in response to shocks. Several competing theories of the business cycle differ primarily in their assumptions about the shape and slope of the cost curve. For example, models of the business cycle which assume either internal or external increasing returns to scale implicitly assume that the short-run marginal cost curve is flat or downward sloping.⁶ In contrast, the real business cycle literature generally assumes a strictly upward sloping marginal cost curve.

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.

⁵In Burnside, Eichenbaum, and Rebelo (1993) variable labor effort can be reinterpreted as a variable shift length.

⁶See for example Murphy, Shleifer, and Vishny (1989), Baxter and King (1992), and Benhabib and Farmer (1994).

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.⁷ Since the SUB is a

⁷See 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.

Saturday work varies among the three-shift plants. In Canada, plants may run three seven-hour shifts on Saturday. In this case, workers are paid for eight hours at time and an half for seven hours of work. U.S. plants tend to run two nine-hour shifts on Saturdays. In this case, the firm pays the workers time and half for these nine hours of work. Thus the Canadian worker receives a slightly higher premium for Saturday work; but each shift at a U.S. plant is guaranteed every third Saturday off.

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

primarily from *Ward's Automotive Reports* and *Ward's AutoInfoBank*. *Automotive News* was used as a secondary source. *Ward's* and *Automotive News* also provided some information on 1, 3, 4, and 5. Although these sources are not entirely independent, information from multiple sources provided a filter to detect errors and inconsistencies in the data from each source.

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).⁸ Exports are from the American Automobile Manufacturers Association (AAMA).

For ten 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 monthly measure of average employment per shift 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. 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).⁹ 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.”

Although working with single-source plants implies a smaller universe of assembly plants to be studied, it has several advantages to working with the more common dual-sourced plants. First it allows one to match inventory and sales data by vehicle line to employment, production, and hours worked data by plant. Second, it allows one to make direct comparisons between actual and scheduled production; this helps to detect errors in the data and identify unanticipated shocks to production.¹⁰ Much of the data studied in this paper are obtained directly from the

⁸Since *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.

⁹Aizcorbe (1992), Cooper and Haltiwanger (1993), and Aizcorbe and Kozicki (1995) also study plant-level data for automobile assembly plants but at the monthly frequency.

¹⁰The data for actual production are also available for the plants which assemble dual-sourced vehicle lines with the exception of Toledo III and Dodge City from 1993:7 to 1994:12. This is because the vehicle lines at the remaining dual-source plants were dual sourced with plants in Mexico. Since the production by vehicle line is decomposed by country of origin, scheduled production can be lined up with the actual production at these plants. However it is

manufacturer. Consequently this dataset contains changes in scheduled production that were not reported in *Ward's* or *Automotive News*. This suggests that these data are more accurate than data collected only from public sources.

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 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

neither possible nor sensible to decompose the sales data by country of origin for the dual-sourced vehicle lines.

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. This observation confirms a finding of Bresnahan and Ramey (1994): Automobile assembly plants rarely choose to reduce output by shortening the workweek.

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 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.

Another way to illustrate the margins used at these plants is to interpret these margins as

pseudo-states. Following Bresnahan and Ramey (1994) each week for each plant (excluding long-run closures) is classified as being in one of twelve pseudo-states. The plant can be engaged in either a 1-shift, 2-shifts, or 3-shifts operation. For each number of shifts there are four statuses: 1. Shut down for part or all of the week for inventory adjustment (IA); 2. Shut down for part or all of the week for a model changeover (MC); 3. Operating four or more overtime hours per week per shift (OT); 4. Everything else (regular hours - RH). To keep this analysis manageable, changes in line speed are ignored.

With every week classified in one of twelve pseudo-states, movements between pseudo-states are documented. Table 3 reports the pseudo-state transition probability matrix constructed from the Chrysler data. In this table, the i, j entry reports the probability that given the firm is in pseudo-state i at date t , the firm will be in pseudo-state j at date $t + 1$.¹¹ For example, one can see from the element in the third row and first column of the table that if a plant is running 1 shift with regular hours one week, then 3 percent of the time the plant will be shut down for a model changeover the next week. The diagonal elements of the table measure the persistence of each pseudo-state. The unconditional probabilities of being in any pseudo-state are reported in the bottom row.

Table 3 illustrates that no plant was shut down for inventory adjustment while running three shifts. All three of the 3-shift plants made high-selling vehicles. These three plants ran overtime weeks three times more often than regular hour weeks. This is further evidence that high-sales plants rarely used weeklong shutdowns to vary output; but these plants used overtime extensively

The probabilities reported in table 3 also confirm that changes in the number of shifts rarely occur. Overtime and inventory adjustment weeks occur often. The diagonal elements of the matrix are all greater than 50. So there is a fair amount of persistence within each pseudo-state. Nevertheless there is a substantial movement across pseudo-states; in particular, plants running one or two shifts cycle between regular hours and inventory adjustment and regular hours and overtime. These plants often alternate between running weeks with regular hours and being down for a week for inventory adjustment. The plants also alternate between running weeks with regular hour shifts and running weeks with overtime shifts. See the bold elements in the table. Plants rarely go from running overtime one week to being down for inventory adjustment the next (or

¹¹To construct table 3 I implicitly assumed that each observation is independent. Since I am working with a panel of data this is a strong assumption.

vice versa). This cyclical production behavior will be discussed in more detail below.

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

Table 4 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.¹² 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 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.

¹²If long-run closures are not excluded, the average workweek of capital is 53.1 hours. This is in line with Shapiro’s statement.

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.*

Tables 5 and 6 provide 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.

From table 5 there are two things to note. First, the weighted average of the ratio of inventories to total sales is 2.55.¹³ This moment will be used to calibrate the model presented below. Second, 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.¹⁴

More interesting are the standard deviations of production and sales presented in table 6. For all but four plants, the standard 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

¹³This ratio is weighted by the total production of each plant.

¹⁴I 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.”

more frequently, production is more volatile than sales.¹⁵ Overall the average plant (weighted by production) schedules production such that the standard deviation is 1.21 times the standard deviation of sales. This number is similar to those reported elsewhere in the literature.

Both automobile production and automobile sales have large seasonal components. To check whether the relative volatility of production and sales is simply due to seasonal variation, production and sales for each plant are regressed separately on twelve monthly dummies.¹⁶ Columns 4 and 5 of table 6 report the standard deviations of the residuals from these regressions. Columns 7 and 8 report the “standard deviations” of the seasonal components. These seasonal series are deterministic; so following Blanchard (1983) the statistic reported is:

$$\left[\frac{\sum_{i=1}^{12} (\hat{\beta}_i - \bar{\beta})^2}{11} \right]^{\frac{1}{2}},$$

where $\hat{\beta}_i$ is the coefficient on the i th monthly dummy, and $\bar{\beta}$ is the mean of the coefficients.

Because of the very short time series available for each plant, this seasonal decomposition should be interpreted with caution. With this caveat in mind, one can see that for all but three of the plants, the ratio of the standard deviations of deseasonalized production to deseasonalized sales is less than or equal to the ratio for the raw time series. Nevertheless, on average, the ratio is greater than one for the deseasonalized time series. The same holds true for the seasonal components. These results suggest that the fact the production is more volatile than sales is not due to just variations at seasonal frequencies.

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 ,

¹⁵The 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.

¹⁶Kashyap and Wilcox (1993) “seasonally adjust” their data by omitting months which include weeklong shut-downs for model changeovers and holidays. Following this strategy does not change the results in any meaningful way.

and labor, n .¹⁷ 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.¹⁸

The plant faces a standard labor contract.¹⁹ 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:

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

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

¹⁷Aizcorbe (1992) provides evidence of a close relationship between the line speed and the level of employment at automobile assembly plants.

¹⁸The 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.

¹⁹I assume the wage schedule from the labor contract is allocative.

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.

There are four free parameters in this example: k , α , $\bar{n}$, and δ . The line speed at the major assembly plants is usually between 40 and 65 vehicles per hour. Employment at plants running two shifts is usually between 2,500 and 4,000 workers. Let k be normalized to unity and $\bar{n}$ be set to 500. If it takes 1,500 workers to run a shift with a line speed of 55 vehicles per hour, then $\alpha = 0.58$. Hence the weekly wage bill for a plant that runs two 40 hour shifts with 1,500 workers per shift is about \$2.2 million.²⁰ 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:

$$n(q, S) = \left(\frac{q}{DS h k^{1-\alpha}} \right)^{\frac{1}{\alpha}} + \bar{n} \quad (2)$$

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

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

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 $(\epsilon, \bar{q})$ and $(\bar{q}, \infty)$, $TC(q)$ is still convex. The non-convexities are caused by the fixed costs associated with opening the plant and opening a second shift.

²⁰N.B. This example abstracts from benefits paid to employees.

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.

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:

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

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

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

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 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 a 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.²¹ 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 tables 5 and 6.

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 $\overline{OA}$? The answer is when:

$$(w_2 - w_1) \cdot D \cdot h \cdot n < \delta. \quad (5)$$

Plugging in reasonable numbers yields:

$$\begin{aligned} (\$18.90 - \$18.00) \cdot 5 \cdot 8 \cdot 1500 &< \delta \\ \$54,000 &< \delta. \end{aligned}$$

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 $\delta > \$54,000$. If $\delta < \$54,000$, the plant will run two shifts one week and a single shift for two weeks. One can see from table 3 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.

²¹An analogous argument can be made for the previous example.

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. In this model the manager 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. Six numerical exercises are run to illustrate the model's ability to capture the three facts listed in the introduction.

5.1 The dynamic program

Several modifications are made to the static model. First, the model is generalized to be dynamic and stochastic, and inventories are explicitly modeled. Second, the plant manager may choose to not work every (or all) employee(s) on the payroll. These modifications not only add realism, but they also allow the plant manager to bunch production. Third, adjustment costs to changing the number of employees at work and on the payroll are added to ensure that the firm pays for the unemployment compensation of the laid-off workers.

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 $\bar{n}$. Let k_t be the time t capital stock. So output produced during period t is:

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

where $0 \leq \alpha \leq 1$. Let the number of workers the plant has on its payroll at time t be $X_t n_t$. Thus X_t denotes the number of shifts of workers the plant has hired. This implies:

$$S_t n_t \leq X_t n_t. \quad (7)$$

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 both a fixed cost and a quadratic cost to changing $X_{t+1}n_{t+1}$. Workers cannot work both shifts. Workers on the payroll who do not work either shift receive unemployment compensation. This unemployment compensation is charged directly and immediately to the firm.

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

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

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 and given some costs of adjustment, the plant's time t cost function is:

$$\begin{aligned} C(t) = & (w_1 + I(S_t = 2)w_2)D_t h_t n_t + \max[0, 0.85(w_1 + I(S_t = 2)w_2)(40 - D_t h_t)n_t] \\ & + \max[0, 0.5(w_1 + I(S_t = 2)w_2)D_t(h_t - 8)n_t] + uw_1 40(X_t - S_t)n_t \\ & + \gamma_3 I(X_t n_t \neq X_{t+1} n_{t+1}) + \frac{1}{2}\gamma_4 (X_{t+1} n_{t+1} - X_t n_t)^2 + \frac{1}{2}\gamma_5 (n_{t+1} - n_t)^2 \\ & + \frac{1}{2}\gamma_1 (i_{t+1} - \gamma_2 E_t s_{t+1})^2 + \delta I(D_t > 0), \end{aligned} \quad (10)$$

where w_1 is the wage rate paid to the first-shift workers, w_2 is the wage rate paid to the second-shift workers, and u is 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 workers on the first and second shifts. 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, sixth, and seventh terms capture the costs of adjusting the line speed and the size of the payroll. The eighth term represents a convex cost to deviating from next period's desired inventory-to-sales ratio. The last (ninth) term denotes the fixed cost to opening the plant. 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:

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

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

Parameters values for the dynamic model are chosen assuming each plant faces the same cost curve. While no two plants in the dataset are identical, the plants are similar along several key dimensions. First, each plant produces common light motor vehicles for the consumer market. None of these plants produce exotic sports cars or commercial trucks. Second, all these plants have just one assembly line and, they have similar production processes; each plant has one paint shop, one body shop, one trim line and one chassis line. Third, each of these plants is run by the same company. Thus all these plants have similar corporate cultures and dictums.

Of course, I could use different parameters values for each plant in order to better match the data. Instead I am only going to vary the sales process and inventory grid. While this strategy may cause the model to miss some features of the data, it most clearly isolates the role the non-convex margins play in explaining the cross-plant heterogeneity in production scheduling.

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. As in the static example, I set the curvature parameter in the production function, α , to 0.58 and the level of overhead labor per shift, $\bar{n}$, to 500. I set the capital stock, $\bar{k}$, to 1.0. In the static example, α , $\bar{n}$ and $\bar{k}$ were chosen to match the line speed and level of employment at the “average plant.”

To check whether equation (6) as it is current parameterized is a good approximation to the true production technology, I performed the following diagnostic. I fixed α and $\bar{n}$. I then inverted equation (6) for each plant and constructed a monthly series, k_t , that reconciles the employment, hours, and production data. For each plant, table 7 reports the average line speed and the mean and standard deviation of the constructed k_t series. For a majority of the plants the mean of the k_t is within one standard deviation of 1. This is reassuring. However it appears that for the plants with slower line speeds (Bramalea, Brampton and Pillette Road), equation (6), with the assumed values for α and $\bar{n}$, is a poor approximation to the true production technology.

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 0.65.²² The average inventory-to-sales ratio reported in table 5 is 2.55. Since this ratio is for monthly sales, it is multiplied by 4.33 to obtain its weekly counterpart; so γ_2 is set to 11.0.

There is little a priori information on the remaining parameters ($\delta, \gamma_1, \gamma_3, \gamma_4$, and γ_5). Since formal estimation is not feasible, the values for these parameters are chosen to roughly match various features in the data. Perhaps the most important parameter in the model is δ , the fixed cost of opening the plant for the week. It is set to \$2.0 million. If the plant lays off two shifts of 1,500 workers, the plant is charged \$1.4 million per week in unemployment fees. So it is still \$0.6 million more expensive to open the plant than to close the plant and send the workers home. This large fixed cost generates a strong incentive for the plant manager to produce a lot on some weeks and shut the plant down on others.

As discussed in section 4, the fixed cost to opening the 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 counter-factually that the manager will open and close the second shift rather than open and close the entire plant. But what is this fixed cost? There are some fixed costs to opening the plant: warming up the equipment, and heating the shop floor. Discussions from

²²I am assuming the plant is in the U.S. For a Canadian plant, I would set $u = 0.40$.

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. 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.

The parameter γ_1 is set to 0.0006. This value implies that the manager will change the stock of inventories freely but will never deviate dramatically from the target inventory-to-sales ratio. I set γ_3 , the fixed cost to changing the line speed, to \$65,000. The two quadratic adjustment cost terms, γ_4 and γ_5 , are set to 3.9 and 3.8, respectively. These choices for the adjustment cost parameters imply that the plant manager will not make frequent changes in line speed. Furthermore when line speed changes are made, they are made smoothly over the course of four or five weeks.

The specification of the cost function, (10), and the parameter values imply that the average cost in the industry of assembling a vehicle is about \$1,450. The average cost of assembling a vehicle (not including payments to suppliers) is about \$2,000. Equation (10) includes costs such as the convex cost of deviating from the desired inventory-to-sales ratio that never show up on the plant's books. Equation (10) also omits some important costs such as energy and maintenance costs. Nevertheless, the implied average assembly cost suggests that the estimated fixed costs and adjustment costs are not unreasonable.

Using the parameter values selected above and a specified sales process, 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 2, 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 63 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.

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 or 2. There are 41 levels of shift employment, n_t , evenly partitioning the interval [600,1800]. In order to conserve on grid points, the inventory grid is allowed to vary across exercises. The grid spaces are chosen to ensure the endpoints never bind and to keep the model computationally feasible. 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 Six exercises

In this subsection I report the results of six dynamic exercises. For each one of these exercises only the sales process and inventory grid are varied. In the first set of three exercises the plant manager faces deterministic sales processes. In the second set of three exercises, the plant manager faces stochastic sales processes.

For each one of the three deterministic exercises, the transition matrix, χ is set to a 2×2 identity matrix, and sales are set so $s_{\text{high}} = s_{\text{low}}$. The three exercises are:

1. *deterministic: high sales* Weekly sales are $s_{\text{high}} = s_{\text{low}} = 4,750$. So monthly sales are 20,568.²³ This about the mean of sales for the St. Louis II plant. The inventory grid is set from 42,750 to 57,000 in increments of 475.
2. *deterministic: medium sales* Weekly sales are $s_{\text{high}} = s_{\text{low}} = 3,500$. So monthly sales are 15,155. This about the mean of sales for the Belvidere plant. The inventory grid is set from 30,000 to 45,000 in increments of 500.
3. *deterministic: low sales* Weekly sales are $s_{\text{high}} = s_{\text{low}} = 500$. So monthly sales are 2,165. This about the mean of sales for the Bramalea plant. The inventory grid is set from 4,500 to 6,500 in increments of 125.

For each one of these deterministic exercises, the average workweek of capital and standard deviation of monthly production are computed.²⁴ These moments are reported in table 8. Con-

²³A month in this model is 4 1/3 weeks.

²⁴The standard deviation of monthly production is the sample standard deviation of 30,000 simulated “4 1/3 week” months.

sider the results from the high sales exercise. From the first row of table 8, it is clear that the plant manager always produces 4,750 cars per week with two-shifts running 8 hours each day for 5 days. Production is perfectly smooth, and the average workweek of capital is 80 hours. The plant manager sets employment (the line speed) so the plant is operating in the convex region of its total cost curve. Given the overtime premium and the short-week compensation rule, the plant manager chooses employment so that two 40 hours shifts are always run.

If the sales rate is 3,500 vehicles a week, the plant manager adopts a different strategy. Although sales are constant, the plant manager makes production variable. The plant manager runs two eight-hours shifts, five days a week for seven weeks. The plant produces 4,000 vehicles per week. So on the eighth week, the plant manager shuts down the plant and reduces inventories by 3500. This eight week cycle then repeats itself. When the plant is open, the plant operates 80 hours per week. Thus the unconditional average workweek of capital is only 70 hours. In this exercise the plant manager is using inventories to convexify the plant's cost curve. The optimal strategy is to produce some weeks and not produce others. Thus the manager makes production more variable than sales.

This cyclical production pattern is even more pronounced when sales are only 500 vehicles per week. In this case (deterministic: low sales) the plant runs one shift for six days, nine hours a day for a week. The plant makes 2,000 vehicles that week. The plant then shuts down for three weeks. This four week cycle then repeats itself. Consequently the average workweek of capital conditional on the plant being open is 54 hours. But the unconditional average workweek of capital is just $54/4 = 13.5$ hours. This low average workweek of capital matches the documented levels of capital utilization at the Bramalea and Toledo II plants. As in the previous exercise, production is variable even though sales are constant.

In the low sales example, the plant uses overtime whenever it is open. This implication is counter-factual. See table 2. In this example, employees receive three weeks of unemployment compensation for every week they actually work. Consequently it is cheaper for the plant manager to have fewer workers on the payroll and work them as much as possible (54 hours each week) than to have more employees (a higher line speed) working a 40 hour shift per week.

This counter-factual implication suggests three potential weaknesses of the model. First, the non-linear price schedule for labor may not be allocative; thus there may be additional constraints on the plant manager that limit the hours of overtime employees can work. Second, the plant

manager may not have as much flexibility as the model assumes in choosing the line speed (the number of employees). Third, equation (6) may be a poor approximation to the true production technology for plants with low line speeds. Nevertheless these three examples demonstrate that a relatively simple dynamic programming model with credible non-convex margins of adjustment can capture much of the heterogeneity in production behavior observed in the data.

Of course in the data, sales are neither constant nor certain. So a second set of three exercises are run. For these three exercises, sales are stochastic. The transition matrix is set:

$$\chi = \begin{bmatrix} 0.923 & 0.077 \\ 0.077 & 0.923 \end{bmatrix}.$$

Thus the unconditional probability of each sales level is 50.0 percent. The average duration of each sales level is 13 weeks. The sales process captures roughly the seasonality in automobile sales: sales are strong in the spring and fall and weak in the summer and winter. These three exercises are:

4. *stochastic: high sales* Sales are set such that $s_{\text{high}} = 6,000$ and $s_{\text{low}} = 3,500$. So the unconditional mean and standard deviation of weekly sales are then 4,750 and 1,250, respectively. Thus the unconditional mean and standard deviation for monthly sales are 20,568 and 4,882, respectively. This roughly matches the mean and standard deviation of sales for the St. Louis II plant. The inventory grid is set from 32,000 to 68,000 in increments of 1,000.
5. *stochastic: medium sales* Sales are set such that $s_{\text{high}} = 4,200$ and $s_{\text{low}} = 2,800$. So the unconditional mean and standard deviation of weekly sales are then 3,500 and 700, respectively. Thus the unconditional mean and standard deviation for monthly sales are 15,155 and 2,733, respectively. This roughly matches the mean and standard deviation of sales for the Belvidere plant. The inventory grid is set from 27,000 to 48,000 in increments of 700.
6. *stochastic: low sales* Sales are set such that $s_{\text{high}} = 800$ and $s_{\text{low}} = 200$. So the unconditional mean and standard deviation of weekly sales are then 500 and 300, respectively. Thus the unconditional mean and standard deviation for monthly sales are 2,165 and 1,172, respectively. This roughly matches the mean and standard deviation of sales for the Bramalea plant. The inventory grid is set from 2,000 to 9,000 in increments of 200.

For each one of these stochastic exercises, the average workweek of capital and relative standard deviations of monthly production to monthly sales are computed. These moments are reported in table 8. Since the plants' production schedules are no longer perfectly periodic, tables 9, 11, and 13 report the three exercises' analogs to table 3, the pseudo-state transition probability matrix. To allow comparisons between the model and the data, selected rows and columns from table 3 for the three sets of single-source plants are reported in tables 10, 12, and 14.

To compute the model's analog to table 3, the state space is partitioned into six pseudo-states that correspond to the number of shifts and hours worked per shift. Pseudo-states labeled '1 shift' and '2 shifts' include the points in the state space in which S_t is equal to 1 or 2 respectively. If $D_t = 0$, then the number of shifts depends on the value of X_t . Pseudo-states labeled IA (for inventory adjustment) correspond to points in the state space for which $D_t < 5$. Pseudo-states labeled RH (for regular hours) include the points in the state space which satisfy $D_t = 5$ and $h_t \leq 8$. Pseudo-states labeled OT (for overtime hours) include the points such that $D_t = 5$ and $h_t > 8$ or $D_t > 5$. These sets span all the points in the state space that occur with nonzero probability.

First consider the case where sales are stochastic and high. Even though the plant is operating in a convex region of its cost curve, the manager does not set production at a constant rate. See tables 8 and 9. The manager still has an incentive to production smooth, but the manager also has an incentive to match a now-variable inventory-to-sales ratio. Consequently, production is as variable as sales. The manager makes extensive use of the overtime margin. The plant uses overtime almost 60 percent of the time and the average workweek of capital is 86 hours.

These results are consistent with the production behavior observed at the high-sales plants identified in the data. For the three high-sales plants (Jefferson North, St. Louis II, and Windsor), the average workweek of capital is about 90 hours and the standard deviations of monthly production and sales are about equal. Table 10 illustrates that these high-sales plants run overtime almost 70 percent of the time. The model does miss two things. First the high-sales plants run 3 shifts a third of the time. In the model, the plant is constrained to run either 1 or 2 shifts. The model also overestimates the use of weeklong shutdowns to adjust inventories.

For the stochastic, medium sales exercise, overtime is used only 29 percent of the time and the plant is shut down 15 percent of the time. See table 11. Reducing the level and variance of sales causes the plant to use overtime less frequently and weeklong shutdowns more frequently.

Consequently, table 8 reports that the standard deviation of production is 48 percent greater than the standard deviation of sales and the average workweek of capital is only 69.0 hours. For the average medium-sales plant, overtime is used about 23 percent of the time, and weeklong shutdowns for inventory adjustment occur about 9 percent of the time (table 12). At these plants, the standard deviation of production is about 50 percent greater than the standard deviation of sales, and the average workweek of capital is about 62 hours (tables 4 and 6). So the model does a fairly good job mimicking the production behavior at these medium-sales plants.

Finally, consider the case when sales are stochastic and low. The implied production behavior is similar to that in the deterministic, low sales exercise. The plant is shut down about 3/4 of the time, and the plant tends to use overtime when it is open. See table 13. So the model implies a low average workweek of capital. Production is considerably more (almost 3 times more) variable than sales. While this may highlight some of the weaknesses of the model already discussed, it suggests that the ratio of the standard deviations of production and sales at Toledo II may not be unreasonable. At Toledo II this ratio is 2.22.

Of course the dynamic model is too simple to match all the features of the data. 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. Furthermore, the analysis assumes that all the Chrysler plants face identical cost curves. In the six exercises I varied only the sales process and inventory grid. Clearly these plants differ along some other dimensions: size, level of technology, UAW vs. CAW, experience of the workforce. These limitations suggest some natural extensions to the analysis.

Nevertheless these six exercises illustrate that much the heterogeneity in the production behavior across plants can be explained by a simple dynamic programming model. The analysis 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 succeeds in reconciling the three facts documented in the third section. 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 often sits idle. At the same time, the dynamic model captures the fact that plants which produce high-selling vehicle lines primarily use convex margins such as overtime employment to vary output. Thus the model can also explain why production at high-sales plants varies by about as much as sales and why capital at these plants rarely sits idle.

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.²⁵

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 (1993) 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.

²⁵However the work of Cecchetti, Kashyap, and Wilcox (1994) suggests that the transportation sector may not be representative of all manufacturing.

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.
- [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] Basu, Susanto, and Miles Kimball (1994) "Cyclical Productivity with Unobserved Input Variation," manuscript, University of Michigan.
- [6] Baxter, Marianne, and Robert King (1992) "Productive Externalities and the Business Cycle," Institute for Empirical Macroeconomics discussion paper no. 53, Federal Reserve Bank of Minneapolis.
- [7] Beaulieu, J. Joseph, and Jeffery Miron (1991) "The Seasonal Cycle in U.S. Manufacturing," *Economic Letters*, Vol. 37, pp. 115-118.
- [8] Benhabib, Jess, and Roger Farmer (1994) "Indeterminacy and Increasing Returns," *Journal of Economic Theory*, Vol. 63, pp. 19-41.
- [9] Berry, Steven, James Levinsohn, and Ariel Pakes (1995) "Automobile Prices in Market Equilibrium," *Econometrica*, Vol. 63, (July), pp. 841-890.
- [10] Bills, Mark (1992) "Measuring Returns to Scale From Shift Practices in Manufacturing," manuscript, University of Chicago.

- [11] Bils, Mark, and Jang-Ok Cho (1994) "Cyclical Factor Utilization," *Journal of Monetary Economics*, Vol. 33, (April), pp. 319-354.
- [12] Bils, Mark, and James Kahn (1996) "What Inventory Behavior Tells Us about Business Cycles," manuscript, University of Rochester.
- [13] Blanchard, Olivier (1983) "The Production and Inventory Behavior of the American Automobile Industry," *Journal of Political Economy*, Vol. 91, (June), pp. 365-400.
- [14] 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.
- [15] Blinder, Alan (1986) "Can the Production Smoothing Model of Inventory Behavior be Saved?" *Quarterly Journal of Economics*, Vol. 101, (August), pp. 431-453.
- [16] 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.
- [17] 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.
- [18] Bresnahan, Timothy, and Valerie Ramey (1994) "Output Fluctuations at the Plant Level," *Quarterly Journal of Economics*, Vol. 109, (August), pp. 593-624.
- [19] Burnside, Craig, and Martin Eichenbaum (1996) "Factor Hoarding and the Propagation of Business Cycle Shocks," manuscript, Northwestern University.
- [20] Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo (1993) "Labor Hoarding and the Business Cycle," *Journal of Political Economy*, Vol. 101, (April), pp. 245-273.
- [21] Burnside, Craig, Martin Eichenbaum, and Sergio Rebelo (1995) "Capital Utilization and Returns to Scale," *NBER Macroeconomics Annual 1995* (Cambridge, MA: MIT Press).
- [22] 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.

- [23] Cooper, Russell, and John Haltiwanger (1992) "Macroeconomic Implications of Production Bunching: Factor Demand Linkages," *Journal of Monetary Economics*, Vol. 30, (October), pp. 107-127.
- [24] Cooper, Russell, and John Haltiwanger (1993) "The Aggregate Implications of Machine Replacement: Theory and Evidence," *American Economic Review*, Vol. 83, (June), pp. 360-382.
- [25] Eichenbaum, Martin (1984) "Rational Expectations and the Smoothing Properties of Inventories of Finished Goods," *Journal of Monetary Economics*, Vol. 14, (August), pp. 71-96.
- [26] Eichenbaum, Martin (1989) "Some Empirical Evidence on the Production Level and Production Cost Smoothing Models of Inventory Investment," *American Economic Review*, Vol. 79, (September), pp. 853-64.
- [27] Fair, Ray (1989) "The Production-Smoothing Model is Alive and Well," *Journal of Monetary Economics*, Vol. 24, (November), pp. 353-370.
- [28] Greenwood, Jeremy, Zvi Hercowitz, and Gregory Huffman (1988) "Investment, Capacity Utilization and the Real Business Cycle," *American Economic Review*, Vol. 78, (June), pp. 402-417.
- [29] Kahn, James (1987) "Inventories and the Volatility of Production," *American Economic Review*, Vol. 77, (September), pp. 667-679.
- [30] Kahn, James (1992) "Why is Production more Volatile than Sales? Theory and Evidence on the Stockout-Avoidance Motive for Inventory Holdings," *Quarterly Journal of Economics*, Vol. 107, (May), pp. 481-510.
- [31] 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.
- [32] Lucas, Robert (1970) "Capacity, Overtime, and Empirical Production Functions," *American Economic Review Papers and Proceedings*, Vol. 60, (May), pp. 23-27.
- [33] Mayshar, Joram, and Yoram Halevi (1992) "On the Time Pattern of Production: Shiftwork, Hours of Work, and Factor Employment," manuscript, Hebrew University.

- [34] Murphy, Kevin, Andrei Shleifer, and Robert Vishny (1989) “Building Blocks of Market Clearing Business Cycle Models,” *NBER Macroeconomics Annual 1989* (Cambridge, MA: MIT Press).
- [35] Ramey, Valerie (1991) “Nonconvex Costs and the Behavior of Inventories,” *Journal of Political Economy*, Vol. 99, (April), pp. 306-334.
- [36] Shapiro, Matthew (1993) “Cyclical Productivity and the Workweek of Capital,” *American Economic Review Papers and Proceedings*, Vol. 83, (May), pp. 229-233.
- [37] Shapiro, Matthew (1995) “Capital Utilization and the Marginal Premium for Work at Night,” manuscript, University of Michigan.
- [38] Solow, Robert (1957) “Technical Change and the Aggregate Production Function,” *Review of Economics and Statistics*, Vol. 39, (August), pp. 312-320.
- [39] West, Kenneth (1986) “A Variance Bounds Test of the Linear Quadratic Inventory Model,” *Journal of Political Economy* Vol. 94, (April), pp. 374-401.

Plant	Period (YR:M)	Single Source?	Vehicle Lines
Belvidere	90:1 - 93:5	yes	New Yorker Salon, Dynasty, Fifth Ave., Imperial
	93:11 - 94:12	no	Neon
Bramalea	90:1 - 91:12	yes	Monaco, Premier
	92:6 - 94:12	no	Concorde, LHS, Vision, Intrepid
Brampton	90:1 - 92:4	yes	Wrangler
Dodge City	90:1 - 93:5	yes	Ram Pickup, Dakota
	93:7 - 94:12	no	Ram Pickup, Dakota
Jefferson North	92:1 - 94:12	yes	Grand Cherokee
Newark	90:1 - 94:12	no	Acclaim, Spirit, Intrepid, LeBaron Sedan
Pillette Road	90:1 - 94:12	yes	Ram Van, Ram Wagon
St. Louis I	90:1 - 91:5	yes	Daytona, LeBaron Coupe
St. Louis II	90:1 - 94:12	yes	Grand Caravan, Grand Voyager, Town & Country
Sterling Heights	90:1 - 94:12	no	Cirrus, Daytona, Shadow, Sundance
Toledo I	90:1 - 94:12	yes	Cherokee, Commanche, Wagoneer
Toledo II	90:1 - 91:6	yes	Grand Wagoneer
	92:7 - 94:12	yes	Wrangler
Toledo III	93:9 - 94:12	no	Dakota
Windsor	90:1 - 94:12	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 Open (1)	Weeks Closed				Short-Weeks		Weeks with OT (8)	Shift Changes (9)	Line Speed Changes (10)
			HOL (2)	SUP (3)	MC (4)	IA (5)	TOTAL (6)	HOL (7)			
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
Pillette 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.

state at $t + 1$	1 shift MC	1 shift IA	1 shift RH	1 shift OT	2 shifts MC	2 shifts IA	2 shifts RH	2 shifts OT	3 shifts MC	3 shifts IA	3 shifts RH	3 shifts OT
state at t												
1 shift MC	52	5	38	5	0	0	0	0	0	0	0	0
1 shift IA	1	66	31	2	0	0	0	0	0	0	0	0
1 shift RH	3	18	59	19	0	0	1	0	0	0	0	0
1 shift OT	1	2	23	74	0	0	0	1	0	0	0	0
2 shifts MC	1	0	0	0	68	0	30	2	0	0	0	0
2 shifts IA	0	0	0	0	0	44	54	1	0	0	0	0
2 shifts RH	0	0	0	0	3	8	75	14	0	0	0	0
2 shifts OT	0	0	0	0	1	0	19	79	0	0	0	0
3 shifts MC	0	0	0	0	0	0	0	0	29	0	57	14
3 shifts IA	—	—	—	—	—	—	—	—	—	—	—	—
3 shifts RH	0	0	0	0	0	0	0	0	4	0	29	67
3 shifts OT	0	0	0	0	0	0	0	0	2	0	22	76
unconditional probability of pseudo-state	0.7	4.3	7.3	6.0	4.6	5.6	37.4	26.5	0.2	0.0	1.8	5.4

Table 3: Complete Pseudo-State Transition Probability Matrix: Chrysler Sample

This table reports the probabilities of week-to-week movements across pseudo-states for the sample of Chrysler assembly plants. IA = inventory adjustment, plant shut down for inventory adjustment for at least part of the week. MC = model changeover. RH = regular hours. OT = overtime hours, 4 hours or more of overtime per shift.

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 4: Average Workweek of Capital (in hours/week)

Plant	Period	Production	U.S. Sales	Canadian Sales	Exports	Total Sales	Inventories	Inventories Total Sales
Jefferson North	92:1 - 94:12	17490	14880	998	980	16858	28196	1.72
St. Louis II†	90:1 - 94:12	20669	18268	2300	357	20924	41214	2.17
Windsor†	90:1 - 94:12	24193	19702	2674	854	23230	65406	2.95
Belvidere	90:1 - 93:5	14600	13902	1394	279	15575	32644	2.10
Brampton/Toledo II	90:1 - 94:12	5581	4778	421	590	5790	11281	2.23
Dodge City	90:1 - 93:5	15851	15561	1627	52	17240	61707	3.69
Toledo I	90:1 - 94:12	12643	10655	747	2061	13464	29425	2.13
Pillette Road	90:1 - 94:12	6567	6141	381	82	6603	18850	3.05
Bramalea	90:1 - 91:12	1783	2087	89	2	2178	6781	3.96
St. Louis I	90:1 - 91:5	6414	6903	535	461	7898	27838	3.77
Toledo II	90:1 - 91:6	445	456	10	60	527	1865	3.90
weighted average		13089				13340		2.55

Table 5: Basic Monthly Statistics: Means

† U.S. registration data provided by The Polk Company was used to decompose U.S. and Canadian sales of the Dodge Caravan and Plymouth Voyager.

Plant	Period	Raw Time Series			Deseasonalized Series			Seasonal Component		
		$\sigma_{\text{production}}$	σ_{sales}	$\frac{\sigma_{\text{production}}}{\sigma_{\text{sales}}}$	$\sigma_{\text{production}}$	σ_{sales}	$\frac{\sigma_{\text{production}}}{\sigma_{\text{sales}}}$	$\sigma_{\text{production}}$	σ_{sales}	$\frac{\sigma_{\text{production}}}{\sigma_{\text{sales}}}$
		(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Jefferson North	92:1 - 94:12	6992	7729	0.90	4943	5511	0.90	4151	3851	1.08
St. Louis II	90:1 - 94:12	5499	5208	1.06	4937	4670	1.06	3185	3105	1.03
Windsor	90:1 - 94:12	6778	6343	1.07	4542	4499	1.01	5695	5129	1.11
Belvidere	90:1 - 93:5	5479	3496	1.57	4265	3290	1.30	3917	2676	1.46
Brampton/Toledo II	90:1 - 94:12	1962	1599	1.22	1598	1034	1.54	1293	1377	0.94
Dodge City	90:1 - 93:5	4991	3502	1.43	3499	3233	1.08	4010	2749	1.46
Pillette Road	90:1 - 94:12	2707	1848	1.47	1707	962	1.77	2265	1750	1.29
Toledo I	90:1 - 94:12	3943	2236	1.76	2684	1964	1.37	3232	1709	1.89
Bramalea	90:1 - 91:12	1253	1172	1.07	2340	3002	0.78	747	935	0.80
St. Louis I	90:1 - 91:5	3208	2261	1.42	1678	1176	1.43	2828	1785	1.58
Toledo II	90:1 - 91:6	373	168	2.22	1940	1268	1.53	317	131	2.41
weighted average		4220	3500	1.21			1.16			1.26

Table 6: Basic Monthly Statistics: Standard Deviation of Production and Total Sales




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	Period (YR:M)	Average Line Speed	mean(k_t)	std(k_t)
Belvidere	90:1-93:5	58.3	1.15	0.25
	93:11-94:12	63.5	1.19	0.11
Bramalea	90:1-91:12	32.1	0.64	0.26
	92:6-94:12	61.4	1.49	0.12
Brampton	90:1-92:4	19.2	11.6	9.9
Dodge City	90:1-94:12	55.2	1.05	0.14
Jefferson North	92:1-94:12	50.7	1.05	0.15
Newark	90:1-94:12	58.0	0.93	0.18
Pillette Road	90:1-94:12	25.3	0.50	0.09
St. Louis I	90:1-91:5	58.8		
St. Louis II	90:1-94:12	49.9	0.87	0.21
Sterling Heights	90:1-94:12	55.6	1.47	0.19
Toledo I	90:1-94:12	45.9		
Toledo II	90:1-91:6	8.1		
	92:7-94:12	36.0		
Toledo III	93:9-94:12	30.3		
Windsor	90:1-94:12	65.0	0.95	0.10

Table 7: Average Line Speeds (vehicles/hour) and Mean and Standard Deviation of Constructed k_t Series

Due to incomplete employment data, I am unable to construct a k_t series for St. Louis I and the three Toledo plants.

Exercise	Mean sales/week	Mean sales/month	$\sigma_{\text{mon. sales}}$	$\sigma_{\text{mon. prod.}}$	$\frac{\sigma_{\text{mon. prod.}}}{\sigma_{\text{mon. sales}}}$	Ave. Workweek of Capital cond. on open unconditional
deterministic: high	4750	20,568	0	0		80.0
deterministic: medium	3500	15,155	0	2904		80.0
deterministic: low	500	2,165	0	1014		54.0
stochastic: high	4750	20,568	4906	4969	1.01	86.0
stochastic: medium	3500	15,155	2747	4056	1.48	81.8
stochastic: low	500	2,165	1177	3401	2.89	47.0
						13.0

Table 8: Results From the Six Exercises

The reported standard deviations of monthly production and sales are the sample standard deviations from 30,000 simulated “4 1/3 week” months.

state at $t + 1$ state at t	2 shift IA	2 shift RH	2 shift OT	3 shifts IA	3 shifts RH	3 shifts OT
2 shift IA	1	4	95	0	0	0
2 shift RH	7	85	8	0	0	0
2 shift OT	4	9	86	0	0	0
3 shifts IA	—	—	—	—	—	—
3 shifts RH	—	—	—	—	—	—
3 shifts OT	—	—	—	—	—	—
unconditional probability of pseudo-state	5.1	36.6	58.3	0.0	0.0	0.0

Table 9: Pseudo-State Transition Probability Matrix: Stochastic, High Sales

This table reports the probabilities of period-to-period movements across pseudo-states in the model. IA = inventory adjustment, plant shut down for at least part of the week. RH = regular hours. OT = overtime hours, 4 hours or more of overtime per shift.

state at $t + 1$ state at t	2 shift IA	2 shift RH	2 shift OT	3 shifts IA	3 shifts RH	3 shifts OT
2 shift IA	0	100	0	0	0	0
2 shift RH	1	56	36	0	0	0
2 shift OT	0	11	88	0	0	1
3 shifts IA	—	—	—	—	—	—
3 shifts RH	0	0	0	0	29	67
3 shifts OT	0	0	0	0	22	76
unconditional probability of pseudo-state	0.1	14.0	45.9	0.0	7.6	23.1

Table 10: Partial Pseudo-State Transition Probability Matrix: High-Sales Chrysler Plants

This table reports the probabilities of week-to-week movements across pseudo-states for the high-sales Chrysler assembly plants. Since this table is just selected rows and columns from the complete transition matrix, the rows do not necessarily sum to 100. IA = inventory adjustment, plant shut down for inventory adjustment for at least part of the week. RH = regular hours. OT = overtime hours, 4 hours or more of overtime per shift.

state at $t + 1$ state at t	1 shift IA	1 shift RH	1 shift OT	2 shifts IA	2 shifts RH	2 shifts OT
1 shift IA	–	–	–	–	–	–
1 shift RH	–	–	–	–	–	–
1 shift OT	–	–	–	–	–	–
2 shifts IA	0	0	0	0	7	93
2 shifts RH	0	0	0	9	90	2
2 shifts OT	0	0	0	37	16	46
unconditional probability of pseudo-state	0.0	0.0	0.0	15.7	55.5	28.8

Table 11: Pseudo-State Transition Probability Matrix: Stochastic, Medium Sales

This table reports the probabilities of period-to-period movements across pseudo-states in the model. IA = inventory adjustment, plant shut down for at least part of the week. RH = regular hours. OT = overtime hours, 4 hours or more of overtime per shift.

state at $t + 1$ state at t	1 shift IA	1 shift RH	1 shift OT	2 shifts IA	2 shifts RH	2 shifts OT
1 shift IA	0	0	100	0	0	0
1 shift RH	0	51	44	0	0	0
1 shift OT	1	20	79	0	0	0
2 shifts IA	0	0	0	41	57	2
2 shifts RH	0	0	0	8	77	12
2 shifts OT	0	0	0	2	47	50
unconditional probability of pseudo-state	0.1	3.8	8.5	9.0	58.1	14.2

Table 12: Partial Pseudo-State Transition Probability Matrix: Medium-Sales Chrysler Plants

This table reports the probabilities of week-to-week movements across pseudo-states for the medium-sales Chrysler assembly plants. Since this table is just selected rows and columns from the complete transition matrix, the rows do not necessarily sum to 100. IA = inventory adjustment, plant shut down for inventory adjustment for at least part of the week. RH = regular hours. OT = overtime hours, 4 hours or more of overtime per shift.

state at $t + 1$ state at t	1 shift IA	1 shift RH	1 shift OT	2 shifts IA	2 shifts RH	2 shifts OT
1 shift IA	73	13	14	0	0	0
1 shift RH	99	1	0	0	0	0
1 shift OT	56	1	44	0	0	0
2 shifts IA	—	—	—	—	—	—
2 shifts RH	—	—	—	—	—	—
2 shifts OT	—	—	—	—	—	—
unconditional probability of pseudo-state	72.4	9.9	17.7	0.0	0.0	0.0

Table 13: Psuedo-State Transition Probability Matrix: Stochastic, Low Sales

This table reports the probabilities of period-to-period movements across pseudo-states in the model. IA = inventory adjustment, plant shut down for at least part of the week. RH = regular hours. OT = overtime hours, 4 hours or more of overtime per shift.

state at $t + 1$ state at t	1 shift IA	1 shift RH	1 shift OT	2 shifts IA	2 shifts RH	2 shifts OT
1 shift IA	68	30	1	0	0	0
1 shift RH	32	65	1	0	0	0
1 shift OT	25	25	50	0	0	0
2 shifts IA	—	—	—	—	—	—
2 shifts RH	—	—	—	—	—	—
2 shifts OT	—	—	—	—	—	—
unconditional probability of pseudo-state	48.4	45.6	1.6	0.0	0.0	0.0

Table 14: Partial Pseudo-State Transition Probability Matrix: Low-Sales Chrysler Plants

This table reports the probabilities of week-to-week movements across pseudo-states for the low-sales Chrysler assembly plants. Since this table is just selected rows and columns from the complete transition matrix, the rows do not necessarily sum to 100. IA = inventory adjustment, plant shut down for inventory adjustment for at least part of the week. RH = regular hours. OT = overtime hours, 4 hours or more of overtime per shift.