

NBER WORKING PAPER SERIES

THE RESPONSES OF PRICES, SALES,
AND OUTPUT TO TEMPORARY CHANGES

Adam Copeland
George Hall

Working Paper 11870
<http://www.nber.org/papers/w11870>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
December 2005

We thank seminar participants at Maryland, Iowa, Yale, and the Federal Reserve Bank of San Francisco for many constructive comments. In particular, we thank Lanier Benkard, Eduardo Engel, Bill Nordhaus, John Rust and Tony Smith for helpful discussions and suggestions. George Hall gratefully acknowledges financial support from the Alfred P. Sloan Foundation. The views expressed in this paper are those of the authors and do not necessarily reflect the views of the U.S. Bureau of Economic Analysis or the U.S. Department of Commerce. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

©2005 by Adam Copeland and George Hall. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Response of Prices, Sales, and Output to Temporary Changes in Demand
Adam Copeland and George Hall
NBER Working Paper No. 11870
December 2005
JEL No. D21, D42, E22, E23, L11, L62

ABSTRACT

We determine empirically how the Big Three automakers accommodate shocks to demand. They have the capability to change prices, alter labor inputs through temporary layoffs and overtime, or adjust inventories. These adjustments are interrelated, non-convex, and dynamic in nature. Combining weekly plant-level data on production schedules and output with monthly data on sales and transaction prices, we estimate a dynamic profit-maximization model of the firm. Using impulse response functions, we demonstrate that when an automaker is hit with a demand shock sales respond immediately, prices respond gradually, and production responds only after a delay. The size of the immediate sales response is linear in the size of the shock, but the delayed production response is non-convex in the size of the shock. For sufficiently large shocks the cumulative production response over the product cycle is an order of magnitude larger than the cumulative price response. We examine two recent demand shocks: the Ford Explorer/Firestone tire recall of 2000, and the September 11, 2001 terrorist attacks.

Adam Copeland
Bureau of Economic Analysis
adam.copeland@gmail.com

George Hall
Department of Economics
Yale University
P.O. Box 208268
New Haven, CT 06520-8268
and NBER
george.hall@yale.edu

In the short-run, the period over which the capital stock and the number of employees is fixed, automobile manufacturers have three primary margins to respond to temporary changes in demand. First, having some market power, they can increase or decrease expected sales by appropriately changing prices. In practice, these price changes often take the form of dealer and customer incentives. Second, they can raise or lower the level of production through adjusting labor inputs by altering the length of the work-week or engaging in temporary layoffs. Third, they may allow inventories on dealer lots to accumulate or de-accumulate. The relative costs of each of these margins determine the shape and slope of the supply curve for new automobiles.

In this paper we determine empirically how the Big Three automakers have accommodated shocks to demand. Consistent with previous work (e.g. Bresnahan and Ramey, 1994), we find that automakers frequently adjust their labor input to increase or decrease production; the average assembly plant uses temporary layoffs 6 percent of the time to reduce production and overtime 30 percent of the time to increase production. Transaction prices, net of rebates and financing incentives, fluctuate considerably, typically falling about 9 percent over the model year (see Copeland, Dunn and Hall, 2005). Dealers inventories are large and volatile; on average dealers hold about 13 weeks worth of sales in inventory and this ratio varies from 7 to 20 (25th and 75th percentiles, respectively). Because these margins of adjustment are interrelated, nonconvex, and dynamic in nature, we estimate a dynamic profit maximization model of an automaker's choice of adjustment to short-term demand fluctuations. Combining weekly plant-level data on production schedules and output with monthly data on transaction prices and sales we find that the propensity to use these alternative margins depends on: 1) the magnitude of the shock; 2) where the firm is in the product cycle; 3) the level of demand relative to the plant's minimum efficient scale; and 4) the current level of inventories.

Specifically, when an automaker is faced with a demand shock to a particular make and model, sales respond immediately. The size of this immediate response is linear in the magnitude of the shock. Prices respond gradually. The higher the initial level of demand and the lower the level of inventories, the more pronounced is the price response. The firm's production responses, on the other hand, are often delayed and discrete. Because of nonconvexities in its cost function, the firm has an incentive to operate the plant at its minimum efficient scale (MES), the rate of production that minimizes average cost. If the shock causes the firm to desire a rate of production below its MES, the firm engages in an "all on/all off" production pattern, using weeklong shutdowns to convexify its costs. In the periods immediately after a demand shock, the rate of production may remain unchanged. However, in later weeks the firm modifies its level

of production by discrete changes in the workweek thus smoothing its production response over time. For sufficiently large shocks the cumulative production response over the product cycle is between 10 and 20 times larger than the cumulative price response.

While the effects of the nonconvexity in weekly production are dampened by time aggregation, they are not eliminated. Thresholds in the production response remain and these thresholds generate small nonlinearities in the future price response. Early in the production cycle, small differences in the magnitude of a shock can generate up to 1.4 percentage point differences in level of production over the remainder of the product cycle, but only 5/100 of a percent differences in the cumulative price response. Later in the product cycle, these effects can be as large as 6 percent for cumulative production, but again only about 5/100 of a percent for cumulative prices.

How firms set prices and output in response to a demand shock is a classic issue in economics going back to at least Hall and Hitch (1939). From our reading of the literature, there was a burst of papers written on this topic in the late 1960s and early 1970s.¹ As with our analysis, these papers typically found that the demand shocks were absorbed by output changes rather than price changes. This result was sometimes interpreted as evidence of ‘sticky prices.’ While we find a small and gradual price response, prices in our model are fully flexible. Interest in firm responses to demand shocks seems to have diminished since the mid 1970s with the increased focus on supply-side shocks as the primary disturbance driving the business cycle. Nevertheless we revisit this issue because plant-level dynamics have macroeconomic implications.

The motor vehicle sector is a sizable fraction of the economy, accounting for almost 4 percent of real GDP in 2003 and 2004. Further, it has a disproportionately large effect on the volatility of GDP. In 2003 and 2004, motor vehicle output accounted for 6.3 and 4.5 percent of real GDP growth respectively.² In addition, new motor vehicle prices have CPI weights of 4.7 percent.³ Hence, production and price changes of new automobiles have observable effects on the aggregate rate of output growth and the rate of inflation. Furthermore, given the prominent role automobiles have in the manufacturing sector in the U.S. and worldwide, we believe that understanding how automakers respond to temporary demand shocks helps in understanding firm pricing and production decisions more generally.

Decisions about inventories, prices, and production are all central to firm profit maximization, but most analysis focuses on only two of these margins at a time. Much of the traditional inventory literature

¹Nordhaus and Godley (1972) summarize much of this literature.

²Computed from numbers in Selected NIPA Tables, Survey of Current Business, July 2005

³Bureau of Labor Statistics website, Table of the Relative Importance of Components in the Consumer Price Index. These are 2001-2002 weights.

address the role of inventories on the timing and volatility of output. The bulk of this literature assumes firms schedule production in the face of exogenous stochastic demand holding prices fixed to minimize the discounted value of expected costs.⁴ In these models, firms typically face both upward-sloping marginal cost curves and costs of adjustment to the level of output. To justify these assumptions, authors usually appeal to, without directly deriving, short-run diminishing marginal returns to labor and costs associated with changing the size the labor force, such as hiring and firing cost. We build on this literature by, first, embedding the firm's cost minimization problem within a profit maximization framework (and thus endogenizing prices and sales). Second, we explicitly link the cost of production to the costs of various margins of adjustment of the labor input.

In operations research the study of the inventory/price tradeoff falls under the headings *revenue management* or *yield management*.⁵ In the economics literature, work by Reagan (1982), Aguirregabiria (1999), Zettelmeyer, Scott Morton and Silva-Risso (2003), and Chan, Hall and Rust (2005) study the interaction between inventory management and pricing. These papers, along with much of the operations research literature, assume simple cost functions. In Reagan, the production function is linear, and in Zettelmeyer et. al., production (procurement) is exogenous; in Aguirregabiria and Chan, et. al. the production function is linear with a fixed set-up cost (i.e. an (S,s) framework). In the current paper, as in our previous work (Copeland, Dunn and Hall, 2005), we study the interplay of inventories and pricing in a model that explicitly incorporates realistic labor costs. In our former paper we explained the coexistence of downward-sloped prices profiles with hump-shaped sales and inventories within a deterministic model. In the current paper, we add a demand shock to study how optimal policies are affected by demand disturbances.

A third literature studies the tradeoff between inventories and employment.⁶ In this literature the link between employment and production is explicit; hence a firm that faces a change in demand can respond either by changing its labor input or allowing inventories to fluctuate. However in these models, there is no pricing decision.

Of the papers we are familiar with, the closest to ours are Haltiwanger and Maccini (1988) and Ramey and Vine (2004). Haltiwanger and Maccini use an intertemporal optimization model to characterize the

⁴Blinder and Maccini (1991a,b) and Ramey and West (1999) provide comprehensive surveys of this vast literature.

⁵This literature which started with Whiten (1955) and Karlin and Carr (1962) is reviewed by Federgruen and Heching (1999) and Elmaghraby and Keskinocak (2003).

⁶Relevant contributions include Topel (1982), Maccini and Rossana (1984), Rossana (1990), and Galeotti, Maccini, and Schiantarelli (2005).

behavior of a firm subject to random demand fluctuations. Prior to observing demand, the firm chooses the product price and the size of its work force; after the shock, the firm can vary output (and thus end-of-period inventories) by choosing the number of workers to layoff and hours work for those who work. Given the relative value to the firm of holding inventories versus laying off workers, firms are classified as being either inventory-biased or layoff-biased. Inventory-biased firms respond to relatively small declines in demand solely by accumulating inventories. Both types of firms use both margins in response to large declines in demand. In the empirical analysis of their model, Haltiwanger and Maccini (1989) take prices and new orders as exogenous to the firm, and are therefore silent on the price/inventory tradeoff.

Ramey and Vine (2004) use Hall's (2000) nonconvex cost-minimization model of an automobile assembly plant to argue that a decrease in the autocorrelation of sales over the last two decades has led to a reduction in the volatility of output in the U.S. automobile industry. While we are limited by our price data to study a much shorter time period than Ramey and Vine, we are able to add a pricing decision to the problem.

The remainder of this paper has three parts. In the first we present our data and describe some its key features. In the second part, we solve and estimate the automaker's dynamic decision problem. We then report impulse response functions of price, sales, and production to shocks to demand. Finally, in the third part, we examine two recent shocks to the automobile industry: the tread-separation tire recall of the Ford Explorer in 2000 and the terrorist attacks of September 11, 2001. The first event represents a true demand shock, while the aggregate time series of prices, sales and production corresponding to the second shock do not accord with the expected responses from a negative demand shock.

1 The Data

The dataset studied in this paper is a merger of two datasets that provide detailed information on sales, prices, inventories, and production of 46 models assembled at 28 single-source automobile factories run by Ford, GM, and Daimler-Chrysler (the Big Three). A single-source plant is a facility that is the exclusive producer of a vehicle line. Focusing on single-source plants, we are able to line up inventory, sales and price data by vehicle line to production and hours worked by plant.

The first dataset was obtained from Wards Communications and contains weekly production data from each assembly plant in the U.S. and Canada for the first week of 1999 through the first five weeks of 2004. For each week the plant operated, it shows: 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.⁷

We match these weekly data to a second dataset on monthly prices, sales, production and inventories by model and model year constructed in Copeland, Dunn, and Hall (2005). Foreign manufacturers are excluded because of problems measuring overseas production. The sales and production numbers in this second dataset also come from Wards Communications. The price data, however, are derived from retail transactions captured at dealerships by J.D. Power and Associates (JDPA).⁸ JDPA attempts to measure precisely the price customers pay for their vehicle, adjusting the price when a dealership under or overvalues a customer's trade-in vehicle as part of a new vehicle sale. JDPA also reports the average cash rebate and average financial package customers received from the manufacturer.

Combining the weekly production dataset with the second dataset on prices, inventories and sales provides us with a detailed picture of the Big Three's pricing and production choices. Because this paper focuses on the operation of an automobile assembly plant, we aggregate the single-source data to the plant/model-year level. This dataset includes 28 factories and has a total of 149 plant/model-year pairs. As illustrated in table 1, this output represents about 34 percent of all Big Three vehicles sold in the U.S. over our sample period.

Vehicles produced at single-source plants are like those produced at multiple plants. The mean price of single-source vehicles is only slightly above the mean price over the entire dataset. Further, with the exception of pickup trucks, single-source plants produce sizeable numbers of vehicles in all market segments.⁹ The single-source subset also is composed of roughly equal amounts from each of the Big Three, although Chrysler is over-represented.

Assembly plants usually operate at full speed (i.e. each shift works 40 hours a week), or not at all.¹⁰ This non-convexity shows up most clearly at the weekly production level, as illustrated in figure 1. This figure plots the weekly output of Chrysler's Jefferson North factory, the sole assembly plant of the Jeep

⁷We thank Dan Vine and Valerie Ramey for providing these data from 1999 to 2001. For the remaining years, the data was taken from weekly issues of *Ward's Automotive Reports* and the annual issues of *Ward's Automotive Yearbook*.

⁸The price data were constructed by Corrado, Dunn, and Otoo (2004), who obtained it from J.D. Power and Associates.

⁹Few single-source plants produce pickups mainly due to data collection and naming conventions. Unlike other market segments, a large variety of essentially different pickups tend to be grouped under one name. Ford F-series pickup trucks incorporates a variety of different vehicles (e.g. F-150, F-250, F-350, etc.), a much wider variety than those vehicles sold under model names in other categories (e.g. Ford Escort or Ford Excursion). Because the production data are collected by model name, we find that several popular pickups are produced at four or five plants.

¹⁰See for example Bresnahan and Ramey (1994) and Hall (2000). Hamermesh (1989) also reports similar results for seven large U.S. manufacturing plants of a large U.S. durable-goods producer.

	Sum of Sales (thousands)			Average Price		
	Entire Sample	Single Source	% diff	Entire Sample	Single Source	% diff
<i>Market Segment</i>						
Compact	4,706	842	18	13,648	14,776	108
Sporty	1,258	1,251	99	25,797	24,499	95
Midsize	11,272	3,927	35	19,247	19,962	104
Traditional	2,134	1,321	62	23,743	23,229	98
Luxury	2,399	1,757	73	36,554	35,757	98
Pickup	13,650	777	6	23,593	20,194	86
SUV	12,868	4,718	37	28,711	28,982	101
Vans	6,321	4,060	64	22,645	22,611	100
<i>Company</i>						
Chrysler	11,692	5,016	43	22,726	24,480	108
Ford	19,121	5,912	31	24,090	26,606	110
GM	23,794	7,726	32	23,546	23,190	98
<i>Total</i>	54,607	18,654	34	23,561	24,620	104

Table 1: Comparing Single-Source Data Relative to the Entire Sample

Grand Cherokee. The tendency for an assembly plant to shutdown completely for a week, if it shuts down at all, is clearly seen for the 2001, 2002, and 2003 model years. Over this period, the assembly plant usually produced around 5000 vehicles a week, or none at all. Of course, there are weeks when the temporary use of overtime ratcheted up production.

Shutdowns in weekly production occur for multiple reasons. Plant closures are grouped into four mutually exclusive categories: model changeovers, holidays, inventory adjustments, and supply disruptions. Model changeovers typically occur in the middle of July, and involve the re-tooling of factories so that new model-year production can start. Holidays are scattered throughout the year, with the longest single vacation occurring from December 25th to January 1st. Assembly plants are shut down for an inventory adjustments when an automaker wants to lower its level of inventories. Finally, supply disruptions are stoppages in production due to parts shortages, power outages, hurricanes, and similar events.

Over our five year sample, assembly plants shutdowns are roughly equally attributable to model changeovers, holidays, and inventory adjustments (see table 2). Supply disruptions play a minor role in explaining shutdowns, accounting for less than 5 percent of all factory shutdowns.¹¹ Table 3 displays

¹¹These numbers from single-source plants are close to the figures reported in Bresnahan and Ramey (1994), which examined a much larger set of assembly plants from 1972 to 1983.




Figure 1: Weekly Grand Cherokee Production

	Model Changeovers	Holidays	Inventory Adjustments	Supply Disruptions
Percent of days shutdown	27.3	37.4	30.7	4.6
Percent of all days	5.6	7.7	6.3	0.9

Table 2: Decomposition of shutdowns

the duration of shutdowns by type. Because of the non-convex cost structure, most plant shutdowns are either for a day or an entire week. Of all the weeks in our sample, plants were shut down for one day in the week 14.2 percent of time, while plants were shut down for an entire week 15.9 percent of the time. Shutdowns that lasted between 2 to 4 days of the week account for less than 4 percent of all weeks in our sample. Looking across the various causes for which plants stop production, we find that single day shutdowns are almost entirely attributable to holidays. Further, model changeovers and inventory adjustments, for the most part, involve a week-long shutdown.

To analyze how production varies with prices and sales, we turn to our monthly data. We first analyze the relationship between assembly plant shutdowns and prices, sales, and inventories. We estimate a probit where the dependent variable, Y , is equal to one if the assembly plant was shutdown at some point in the month. The independent variables are last month's sales and prices, beginning-of-period inventories, a model-year monthly trend term, and assembly plant fixed effects. Letting s, p, i, m be sales, prices,

	Shutdown Duration				
	One Day	Two Days	Three Days	Four Days	Entire Week
Holiday	13.5	2.3	1.1	0.0	3.4
Model Changeover	0.0	0	0	0	5.6
Inventory Adjustment	0	0	0	0.1	6.3
Supply Disruption	0.7	0.1	0.1	0.1	0.6
Total	14.2	2.4	1.2	0.2	15.9

Table 3: Frequency of shutdowns by category and duration (percent of total weeks)

	Estimate	Standard Errors	Marginal Effect
$\ln(p_{t-1})$	-1.938	0.793	-0.192
$\ln(s_{t-1})$	-0.438	0.133	-0.043
$\ln(i_t)$	0.797	0.177	0.079
m_t	-0.035	0.024	-0.003

Table 4: Estimated Probit Parameters

inventories, and the model-year trend variable respectively, we write the probit as

$$Y_t = b_1 \ln(s_{t-1}) + b_2 \ln(p_{t-1}) + b_3 \ln(i_t) + b_4 m_t + \sum_{s=1}^F \lambda_s 1_{f_t=s} + \zeta_t, \quad (1)$$

where $f_t = \{1, 2, \dots, F\}$ designates to which factory the observation belongs, $1_{x=y}$ is an indicator variable equal to one if $x = y$, and ζ is an i.i.d error term. The estimated coefficients on lagged price, sales, and inventories, shown in table 4, are significantly different from zero and have the expected signs. If prices or sales are high in the previous month (for example, if demand is strong), then the probability of the assembly plant shutting down next month decreases. Further, the higher inventory levels climb, the more likely a factory will close. Lastly, the estimate of the trend coefficient is insignificant. While these estimates accord well with theory, we are cautious in interpreting the strength of these results because the probit's explanatory power is low; the R^2 is 0.12.

To capture the co-movements in prices, sales and production, we estimate three least square regressions of sales, price, and production. For all three regressions, the independent variables are a lag of prices, a lag of sales, beginning-of-period inventories, and a trend. Because we are interested in the dynamics of the data and not the cross section, we take out the plant-level mean of all variables. This approach allows us to control for plant-level fixed effects, without worrying about estimating fixed effects. As will be apparent later, this is a significant advantage under our indirect-inference estimation methodology. Let \hat{x}_t denote a

parameter	Sales Equation		Price Equation		Production Equation	
	estimate	std err	estimate	std err	estimate	std err
lagged price	0.172	0.035	0.787	0.006	0.787	0.072
lagged sales	0.641	0.009	0.0443	0.0051	0.376	0.028
inventories	0.0858	0.0024	-0.0076	0.0012	0.0250	0.0064
trend	-0.104	0.013	-0.0432	0.0036	-0.194	0.051
σ^2	5.33	0.10	0.691	0.009	15.39	0.44
R-squared	0.88		0.99		0.72	

Table 5: Estimated parameters

Note: std err are Newey-West standard errors with two lags.

variable minus its plant-level mean. Formally, we estimate

$$\begin{aligned}
\hat{s}_t &= \theta_1 \hat{s}_{t-1} + \theta_2 \hat{p}_{t-1} + \theta_3 \hat{i}_t + \theta_4 \hat{m}_t + v_t^s \\
\hat{p}_t &= \theta_5 \hat{s}_{t-1} + \theta_6 \hat{p}_{t-1} + \theta_7 \hat{i}_t + \theta_8 \hat{m}_t + v_t^p \\
\hat{q}_t &= \theta_9 \hat{s}_{t-1} + \theta_{10} \hat{p}_{t-1} + \theta_{11} \hat{i}_t + \theta_{12} \hat{m}_t + v_t^q
\end{aligned} \tag{2}$$

where q is production and v is an i.i.d. normal error. Automakers typically produce a particular vehicle for 12 months, but through the use of inventories sell the vehicle over longer period. The sales and price regressions are estimated using an average of 17 months of data for each vehicle, while the production regression is estimated using an average of 12 month of data. This set of regressions provides a description of the persistence of sales, prices and production as well as their covariances. The estimates of θ are reported in table 5, along with the variance of the residuals, σ^2 .

Our results show that both sales and prices are highly persistent. The estimated coefficient on lagged prices in the price equation is a high 0.79, while for the sales equation our estimate on lagged sales is 0.64. Further, we find that beginning-of-period inventories are significantly correlated with both sales and prices. Consistent with inventory-control theory, higher levels of inventories coincide with higher sales and lower prices. Finally, both sales and prices have a negative trend, suggesting a fall in demand over the model year. Turning to production, we find it is negatively correlated with last month's price and positively correlated with last month's sales. This is not surprising, given the highly persistent behavior of both sales and prices. Production is also positively correlated with inventories. This result likely reflects automakers' practice of producing enough cars in 12 months to build up inventories for 17 or more months of sales. Finally, the discreteness of the production data is clearly reflected in the large variance of the production residual, relative to the variance of the sales residual. Note that all three equations fit the data well, despite the simple specification used.

2 The Model

The model examines an automaker selling a single product. The decision period is a week. A particular model year is produced at a single plant for one year (52 weeks) and sold for two years (104 weeks). In each of the first 52 weeks, the firm must decide: (1) the number of vehicles to produce, q_t ; (2) the number of days to operate the plant, D_t , the number of shifts to run, S_t , and the number of hours per shift, h_t ; and (3) the retail price of the vehicle, p_t

Weekly sales, s_t , depend on the vehicle's own price, p_t , a persistent shock z_t , and a deterministic time-varying constant term μ_t . The weekly demand curves

$$s_t = \mu_t(1 + z_t) - \eta_t \log(p_t) \quad (3)$$

take a linear-log specification with η_t denoting the week t own-price semi-elasticity. We assume the persistent shock follows an autoregressive process:

$$z_{t+1} = \rho z_t + \omega_{t+1} \quad (4)$$

with ω distributed i.i.d $N(0, \sigma_\omega)$. This demand specification implies that above some price demand for the vehicle is zero, consistent with consumers fully substituting to other vehicles at some price. This model ignores the interaction of demand between different model years of the same model (e.g. a 1999 and 2000 Ford Escort), because previously (Copeland, Dunn, and Hall, 2005) we found these cross-price elasticities to be very small.

Unsold vehicles can be inventoried without depreciation. Let i_{t+1} be the stock of vehicles that are inventoried at the end of period t and carried over into period $t + 1$. Current production is not available for immediate sale, so sales can be made only from the beginning-of-period inventories:

$$s_t \leq i_t. \quad (5)$$

Sales cannot be backlogged. During the production year, inventories follow the standard law of motion:

$$i_{t+1} = i_t + q_t - s_t \quad 0 < t \leq 52. \quad (6)$$

After 52 weeks no vehicles are produced so inventories evolve according to

$$i_{t+1} = i_t - s_t \quad 52 < t \leq 104. \quad (7)$$

At the conclusion of week 104, any unsold vehicles are sold at a fixed price \bar{p}_{105} .

The vehicle is assembled at a single plant. Each period, the firm decides how many vehicles to produce and how to organize production to minimize costs. Plant managers increase or decrease production by altering the workweek rather than the rate of production. The plant can operate D days a week. It can run one or two shifts, S , each day, and both shifts are h hours long. We assume the number of employees per shift, n , and the line speed, LS , are fixed. So the firm's production function is linear in hours:

$$q_t = D_t \times S_t \times h_t \times LS. \quad (8)$$

Although the production function is linear, the firm faces several important non-convexities because of its labor contract. We let w_1 and w_2 denote the straight-time, day-shift and evening-shift wage rates. Any work in excess of eight hours a day, and all Saturday work, is paid at a rate of time and a half. Employees who work 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 a worker receives for the hours actually worked. If the firm chooses not to operate a plant for a week, the workers are laid off. Laid-off workers receive υ fraction of their straight-time 40-hour wage.

Such a labor contract means that if the firm decides to produce q vehicles, it must then choose how many days to operate the plant, how many shifts to run, and how many hours to run each shift in order to minimize its cost of production. Given these choices, the firm's week t cost function is expressed as

$$\begin{aligned} c(D_t, S_t, h_t | q_t) = & \gamma q_t + (w_1 + I(S_t = 2)w_2) \times (D_t h_t n + \max[0, 0.85(40 - D_t h_t)n] \\ & + \max[0, 0.5D_t(h_t - 8)n] + \max[0, 0.5(D_t - 5)8n]) + \upsilon w_1 40(2 - S_t)n, \end{aligned} \quad (9)$$

where γ is the per vehicle material cost, n is the employees per shift, and w_1 and w_2 are the hourly wage rates paid to the first-shift and second-shift workers, respectively. The first term is the per-vehicle cost. It incorporates all costs (such as materials, energy, transaction) that do not depend on the allocation of production over the week. The first term within the brackets represents the straight-time wages paid to the production workers. The subsequent terms within the brackets capture the 85 percent rule for short weeks and the statutory overtime premium. The last term is the unemployment compensation bill charged to the firm. Let $D_t = 0$ if and only if $S_t = 0$. This cost function is piecewise linear with kinks at one shift running 40 hours per week and two shifts running 40 hours per week.

To see this graphically we plot in figure 2 the labor-cost portion of (9) (i.e. $c(D_t, S_t, h_t) - \gamma q_t$) conditional on the plant running one or two shifts. The statutory 85% shortweek rule and overtime premia create kinks in each cost curve at 0 and 40 hours per shift. The total labor cost portion of the bill is the envelop




Figure 2: Labor costs as a function of weekly output conditional on running one or two shifts

of the two curves (denoted by the solid blue and green portions of the curves). Because of the kinks, the MES is determined by the unique line (the dashed red line) from the origin that intersects the cost curve only once. In the example plotted in figure 2 the firm minimizes average costs by operating the plant with two shifts for a total of 80 hours per week. If the plant's desired output is below the MES, the firm will minimize cost by taking a convex combination of producing at 0 and producing at its MES.

The firm's objective is to maximize the present value of the discounted stream of profits. For each model year the automaker's problem is to

$$\max_{p_t, q_t} E \left\{ \sum_{t=1}^{104} \left(\frac{1}{1+r} \right)^{t-1} \left\{ p_t s_t (1 - \tau(s_t/i_t)^\Psi) - \min_{D, S, h} c(D_t, S_t, h_t | q_t) \right\} \right\} \quad (10)$$

subject to (3)-(8) and where $c(D, S, h | q)$ is given by (9). The term $\tau(s_t/i_t)^\Psi$ is a "revenue tax" incurred by the automaker as a function of its sales-to-inventory ratio. This term is in the spirit of Bills and Kahn's (2000) inventory model in which inventories play a revenue-generating role. When inventories are low, it is harder for potential customers to observe and gauge a vehicle (that is, to test-drive it and view the choice set), and thus it is more costly to consummate a sale. The tax effectively disappears when the sales-to-inventory ratio is small.

We assume that the firm must choose p_t and q_t before observing z_t . Let $V(I, z_{-1}, t)$ be the optimal

value at week t for the firm that holds inventory I and observed a demand state of z_{-1} last period. The firm's value function for weeks $t = 1, 2, \dots, 52$ can be written:

$$V(i, z_{-1}, t) = \max_{p, q} \left\{ pE[s(1 - \tau(s/i)^\Psi)] - \min_{D, S, h} c(D, S, h|q) + \frac{1}{1+r} EV(i + q - s, z, t + 1) \right\} \quad (11)$$

subject to (3), (4), (5), and (8) and where $c(D, S, h|q)$ is given by (9). For weeks $t = 53, 54, \dots, 104$ the value function becomes:

$$V(i, z_{-1}, t) = \max_p \left\{ pE[s(1 - \tau(s/i)^\Psi)] + \frac{1}{1+r} EV(i - s, z, t + 1) \right\} \quad (12)$$

subject to (3), (4), and (5). Hence the firm's pricing and production decisions are governed by the policy functions:

$$\begin{aligned} p_t &= p(i_t, z_{t-1}, t) \\ q_t &= q(i_t, z_{t-1}, t) \end{aligned} \quad (13)$$

which solve (10).

3 Estimation of the Structural Model

We estimate the structural model in two steps. First, we employ a discrete-choice methodology to estimate consumers' preferences over automobiles. We use these estimates to compute the intercepts and own-price semi-elasticities that are parameters in the market demand curves, equation (3). Second, taking these market demand curves as given we estimate the remaining parameters via indirect inference; that is, we select the structural parameters of the model such that when we estimate the set of regressions described in equation (2) using model simulations, we match as closely as possible the estimated parameters reported in table 5.

3.1 Estimating the Demand Elasticities via BLP (1995)

The demand elasticities are estimated using the approach described in our earlier work, Copeland, Dunn, and Hall (2005).¹² The demand for automobiles is modeled within a discrete-choice framework. Following Berry, Levinsohn, and Pakes (1995), henceforth BLP, we construct the demand system by aggregating over the discrete choices of heterogeneous individuals.

¹²A full description of the methodology and results are available in this earlier paper. Here, we only provide an overview of the methodology and the final results.

The utility derived from choosing an automobile depends on the interaction between consumer and product characteristics. Consumers are heterogeneous in income as well as in their tastes for certain product characteristics. We distinguish between two types of product characteristics. Those that are observed by the econometrician (such as horsepower and miles per gallon) are denoted by X . Those that are unobserved by the econometrician (such as styling or prestige) are denoted by ξ . Drawing from the nested logit literature, we also incorporate a correlation in the consumer's tastes for vehicles of the same model year. We divide vehicles into $G + 1$ mutually exclusive groups (that is, model years)— $g = 0, 1, 2, \dots, G$ —where the outside good is the sole member of group 0. We also allow households' distaste for price, denoted by α , to vary from quarter to quarter. This captures the possibility that different types of households show up to purchase a new automobile at different times of the year.

We specify the indirect utility derived from consumer i purchasing product j , dropping the time subscript, as

$$u_{ijq} = X_j \delta + \xi_j - \alpha_{iq} p_j + \sum_k \varphi_k \iota_{ik} x_{jk} + \zeta_{ig} + (1 - \kappa) \vartheta_{ij}, \quad (14)$$

where p_j denotes the price of product j and $x_{jk} \in X_j$ is the k th observable characteristic of product j . The term $X_j \delta + \xi_j$, where δ are parameters to be estimated, represents the utility from product j that is common to all consumers, or a mean level of utility. Consumers then have a distribution of tastes for each observable characteristic. For each characteristic k , consumer i has a taste ι_{ik} , which is drawn from an independently and identically distributed (i.i.d.) standard normal distribution. The parameter φ_k captures the variance in consumer tastes.

The term α_{iq} measures a consumer's distaste for price increases in quarter $q = \{1, 2, 3, 4\}$. Following Berry, Levinsohn, and Pakes (1999), we assume that $\alpha_{iq} = \frac{\alpha_q}{y_i}$, where α_q is a parameter to be estimated and y_i is a draw from the income distribution. We assume the distribution of household income is lognormal, and, for each year in our sample, we estimate its mean and variance from the Current Population Survey. The second-to-last term in equation (14) captures correlations in a consumer's tastes for products within the same group. For consumer i , the variable ζ_{ig} is common to all products in group g and has a distribution that depends upon κ . Finally, ϑ_{ij} is an i.i.d. extreme value.

Consumers choose among the $j = 1, 2, \dots, J$ automobiles in our sample and the outside good (denoted $j = 0$), which represents the choice not to buy a new automobile from the Big Three. Consumers choose the product j that maximizes utility, and market shares are obtained by aggregating over consumers.

The dataset of prices and sales for the Big Three is used to estimate the model, generally following

Market Segment	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8
Compact	8.2	9.2	7.9	8.6	7.8	8.3	8.8	8.8
Full	10.3	11.3	9.1	9.6	10.0	11.4	10.5	10.6
Luxury	9.8	11.4	8.4	9.0	9.9	11.0	9.2	9.2
Midsize	9.4	10.5	8.9	9.7	9.4	10.0	8.9	8.1
Pickup	9.7	10.8	9.1	9.0	9.8	10.8	8.9	11.6
SUV	9.7	10.7	8.6	8.6	9.9	10.8	9.6	8.6
Sporty	10.8	10.9	9.1	10.0	9.8	11.7	9.3	8.6
Van	10.2	11.6	9.7	10.1	10.2	10.9	9.6	5.8
Single Source	10.1	11.0	9.0	9.5	9.8	11.0	9.3	8.6

Table 6: The Absolute Value of Own-Price Elasticities by Market Segment and Quarter

BLP’s algorithm. We aggregate sales and prices to the quarterly frequency because of volatility in monthly sales due, in part, to intertemporal substitution. We do not estimate the model at an annual frequency because the variation in price and in the consumer’s choice set from quarter-to-quarter is a significant source of identification in the BLP framework. Lastly, we augment the data with vehicle characteristic information from Automotive News’ *Market Data Book* (various years).

The estimated elasticities that result from the discrete-choice estimation are reported in table 6. These elasticities provide the clearest picture of the values of the semi-elasticities that we use in our specified demand function. The own-price elasticities generated by our parameter estimates range between 6 and 12, an indication that manufacturers face quite elastic demand. In the first quarter a car is sold, our results imply that a 1 percent price increase for a typical compact car (roughly \$140) causes an 8.2 percent fall in sales, holding everything else equal. The average own-price elasticity for all single-source vehicles is reported in the “Single Source” row. In general, our estimated elasticities are higher than those found in the previous literature; BLP, for example, report a range of elasticities between 3 and 6. This difference is not surprising, however, because previous research estimated own-price elasticities for models, a level of aggregation higher than that of our data. Indeed, when we re-estimate the parameters from data aggregated to the model level, the implied own-price elasticities fall within the range reported in BLP.

While we compute elasticities by quarter, our model of the firm is at the weekly frequency. To construct the weekly demand curves, equation (3), we use our discrete choice model to construct quarterly own-price semi-elasticities for the typical single-source vehicle. We then interpolate these semi-elasticities to the weekly frequency using a spline. To compute the intercept terms μ_t , $t = 1, 2, \dots, 104$, we first interpolate the monthly price/quantity-sold pairs for an average single-course plant to the weekly level; we then require

each demand curve to go through the interpolated price-quantity pair for its corresponding week. This yields a set of 104 demand curves that are falling (i.e. shifting to the southwest corner) over the product cycle.

3.2 Estimating the Firm’s Decision Problem via Indirect Inference

Because of the non-convexities in the cost function, we solve for both the optimal level of output and the cost minimizing production schedule through grid search. The grids for D_t and S_t are set from 1 to 6 and from 0 to 2, respectively, in increments of 1. The plant is closed for the week whenever $S_t = 0$. The shift length, h_t , can take on values of 7, 8, 9 or 10. We allow weekly production ($D_t \times S_t \times h_t \times LS$) to take values between 0 and $120 \times LS$ in increments of LS . So there are up to 72 feasible production schedules to evaluate for each 121 possible levels of production.

We discretize the inventory grid into 26 points from 0 to 85,000. The distance between grid points increases with the level of inventories. Thus, the grid points are more densely spaced in the region where the value function has more curvature. We discretize the z grid into 11 points from -0.025 to 0.025; the grid points are more densely spaced near zero, and the distance between grid points increases the further points are away from zero. For each of the 286 (i, z) pairs, we maximize recursively the right hand side of equations (11) and (12) over each sales price and level of output. Points off the inventory and z grids are approximated using bi-linear interpolation, and all integration is done by quadrature. The sales price, p_t may take on any positive value such that quantity demanded remains positive. Finally, we impose a standard holiday schedule on production; we assume the plant is closed for days corresponding to Labor Day (1 day, week 8), Thanksgiving (2 days, week 19), Christmas/New Year’s (5 days, week 24), Martin Luther King Day (1 day, week 27), Good Friday (1 day, week 37), Memorial Day (1 day, week 46), and the July model changeover/vacation (10 days, weeks 51 and 52).

We estimate the supply-side parameters along with the demand-shock processes via indirect inference using the extended method of simulated moments (EMSM) proposed by Smith (1993). This approach selects the set of structural parameters, β , that minimizes the distance between a set of observed moments $\hat{\theta}_T$ and those generated by numerical simulations of the structural model. In this case, the moments to be matched are the 15 regression coefficients from the set of full regressions of sales, price and production, reported in table 5, the error covariance matrix of the sales and price regressions, the variance of the production regression, and three coefficients obtained from separately regressing sales, price and production

on a constant.¹³ These last three equations provide the mean levels of sales, prices and production at a single-source plant for the model to match. The full regression coefficients and error covariance matrix capture the dynamics of prices, sales, and production. Note that the price and production equations can be interpreted as empirical counterparts of the model’s decision rules for prices and production, equation (13). In the model weekly prices and production are functions of the three state variables: z_t , i_t and t . Of course, in the data we observe monthly rather than weekly prices, and we do not observe the state of demand z_t . One can interpret s_{t-1} and p_{t-1} as providing a measure of z_{t-1} .

In addition to the demand curves, we fix several supply-side parameters prior to the estimation. We set the “scrap value” of vehicles unsold after 104 weeks \bar{p}_{105} to \$19,000. We set the number of workers per shift, n to 1300. We set the overtime premium to 1.5 (i.e. “time and a half”), the second-shift premium to 1.05, (i.e. $w_2/w_1 = 1.05$), and the short week premium to 0.85 as specified in the union contracts. Let β denote the vector of the remaining structural parameters that we wish to estimate: $\beta = \{r, \rho, \sigma_\omega, \gamma, \tau, \psi, \nu, w_1, LS\}$.

The basic strategy to estimate the model is:¹⁴

1. Use the data to compute estimates of the coefficients and the variance-covariance matrix of the residuals for the set of regressions stated in equation (2) as well as the least square estimates of the mean level of sales, price, and production, $\hat{\theta}_T$. Our estimates of $\hat{\theta}_T$ are reported in table 8.
2. For a given set of parameters β , solve the structural model.
3. Simulate the model for 104 weeks S times and aggregate each simulation to the monthly frequency to create a $24 \times S$ panel dataset $y(\beta)$. For each simulation, initialize z using a draw from its ergodic distribution.
4. Estimate the set of full regressions and the three least squares means listed in the first step, using $y(\beta)$ to compute $\hat{\theta}_S^\beta$. Measure the distance between the vector of observed parameters and the vector of simulated parameters via the criterion:

$$(1 + \pi^{-1})^{-1}(\hat{\theta}_T - \hat{\theta}_S^\beta)'W_T(\hat{\theta}_T - \hat{\theta}_S^\beta) \quad (15)$$

¹³We do not attempt to match the covariances between the production residuals and the price and sales residuals because of the different sample sizes. Recall that the production regression used the 12 months of data that a plant produced a vehicle, while the sales and price regressions used the 17 months of data for which a typical vehicle was sold.

¹⁴Since we follow Smith’s (1993) methodology rather closely we describe it only generally and refer readers to Smith’s paper for a complete description of the derivations and asymptotics.

r^\dagger	γ	τ	ψ	ρ	σ_ω	υ	w_1	LS
0.0140	22,270	0.423	0.907	0.954	0.000772	0.697	58.92	40.18
0.0107	114	0.080	0.140	0.005	0.000061	0.185	7.21	0.80

Table 7: EMSM estimates of the structural parameters

Note: The first row of the table reports point estimates. The second row of the table reports estimated standard errors.

\dagger The interest rate r is reported at an annual rate.

where the weighting matrix, $W_T \equiv A_T(\theta_T)B_T(\theta_T)^{-1}A_T(\theta_T)$. $A_T(\theta_T)$ and $B_T(\theta_T)$ are the Hessian of the likelihood function and the information matrix, respectively, for the set of full regressions described earlier and the three OLS regressions of sales, price and production on a constant. We compute these matrices numerically. We compute $B_T(\theta_T)$ using the Newey-West (1987) estimator with two lags. Since $-A_T(\theta_T) \approx B_T(\theta_T)$ the weighting matrix is the inverse of variance-covariance matrix of the observed parameters taking into account the model mis-specification. The term π denotes the ratio of the simulation sample size to the data sample size.

- Using a hill-climbing algorithm, repeat steps 2 - 4 to find the $\tilde{\beta}_T$ that minimizes (15).

We set the number of simulations S to 298, twice the number of plant/model years in our dataset. Thus $\pi = 2$.

In table 7 we report point estimates for the structural parameters together with estimated standard errors. The estimated parameter values are sensible. The interest rate, which is the sole cost of holding inventories, is estimated to be 1.4 percent at an annual rate with a standard error of 1.1. Even though this represents a real rate of interest, it still seems low. The estimate may be picking up some implicit value to holding inventories. The per-vehicle material cost, γ is estimated to be \$22,270. Because of the non-convexities in the cost function, the marginal cost of production is a discontinuous, non-monotonic function. However, when the plant operates two 40-hour-shifts per week, the average cost of production is \$24,176. In this case the per-vehicle labor cost is \$1906. The average profit per vehicle is about \$1804. The revenue-tax parameters, τ (the tax rate) and ψ (the curvature parameter) are estimated to be 0.42 and 0.91 respectively. Hence, the revenue tax is a concave function of the sales-to-inventory ratio.

The z process is estimated to be quite persistent with an auto-regressive coefficient of 0.954 and a standard deviation of the innovations (σ_ω) of 0.000772. This implies the ergodic distribution of z has a mean of zero (by assumption) and a standard deviation of 0.0026. While a standard deviation of 3/10 of one-percent may seem small, a one standard deviation drop in z results in a downward parallel shift in the

demand curve of 350 to 780 vehicles per week (depending on the value of μ_t).

The point estimate of the first-shift wage rate w_1 is reasonable, \$58.92 per hour, if one accounts for both benefits and wages but it is not particularly interesting since it can be scaled up and down by our choice of n . Our estimate of the unemployment replacement rate, υ , is of more economic interest. It is estimated to be 0.697. In the U.S., if an assembly plant closes 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 covered 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 future UI premiums. Aizcorbe (1990) and Anderson and Meyer (1993) report that due to the cross-industry subsidies inherent in the UI system, firms end up paying about half of the 60 cents coming from UI. Since the SUB is a 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. This is quite close to (i.e. well within one standard error) our estimate of 0.697.

Finally our point estimate of the line speed of 40.18 vehicles per hour is consistent with observed line speeds at assembly plants. In our dataset plant line speed vary between 30 to 70 vehicles per hour.

The estimation criterion (15) provides a test-statistic for the over-identifying restrictions of the model.¹⁵ This statistic is distributed $\chi^2(n - k)$. In our case there are ten over-identifying restrictions ($n-k=19-9$), and the statistic is 746.7. Thus our structural model can be overwhelmingly rejected as the true data-generating process of the observed time series. Nevertheless, the model captures much of the interesting dynamics in the data.

Table 8 tabulates two sets of estimated parameters, one set for the observed data and one set for the simulated data given the estimated structural parameters in table 7.¹⁶ Eight of the 19 simulated moments are within a single standard error of their corresponding observed moments, and almost all of the simulated moments are of the same sign and magnitude as the observed moments. The model implies a high level of auto-correlation in both sales and prices. The model also implies the higher inventories are associated

¹⁵Readers may notice that the criterion in (15) is not scaled by the number of observations. As discussed in section 1, the number of observations differ across the sales, price and production regressions. The individual elements of the $A_T(\theta_T)$ and $B_T(\theta_T)$ matrices are scaled appropriately to take this in account, so we do not 'pull a T out to the front' of the expression.

¹⁶Our use of a set of regressions of sales, price and production to evaluate the fit of the firm's decision problem is quite similar in spirit to the analysis by Hay (1970). Hay calibrated a linear-quadratic model of the firm, took first order-conditions and compared informally the SUR of prices, production, and inventories implied by his model to two SURs estimated using data on the lumber and paper industries. Our findings are consistent with those of Hay.

parameter	Sales Equation		Price Equation		Production Equation	
	observed	simulated	observed	simulated	observed	simulated
lagged price	0.172	-0.0146	0.787	0.865	0.787	1.762
	0.035	0.0209	0.006	0.045	0.072	0.254
lagged sales	0.641	0.644	0.0443	0.0311	0.376	0.127
	0.009	0.009	0.0051	0.0109	0.028	0.020
inventories	0.0858	0.0794	-0.0076	-0.0053	0.0250	0.0473
	0.0024	0.0028	0.0012	0.0015	0.0064	0.0084
trend	-0.104	-0.236	-0.0432	-0.0349	-0.194	-0.181
	0.013	0.007	0.0036	0.0141	0.051	0.062
resid variance	5.33	1.60	0.691	0.619	15.39	4.95
	0.10	0.04	0.009	0.252	0.44	0.20
R-squared	0.88	0.85	0.99	0.68	0.72	0.29
	parameter		observed	simulated		
	cov(resid sales, resid price)		-0.0741	0.0568		
			0.0503	0.0135		
parameter	Sales Equation		Price Equation		Production Equation	
	observed	simulated	observed	simulated	observed	simulated
constant	8.952	8.959	26.036	25.970	12.600	12.298
	0.202	0.049	0.223	0.015	0.332	0.045

Table 8: Estimated parameters using observed and simulated data

Note: The top and bottom numbers in each cell are, respectively, the point estimate and standard errors.

with higher sales, lower prices and higher production. In both the model and the data higher lagged prices are associated higher production.

The three moments the model fails to match by wide margins are the variance of residuals in the sales and production regressions, and the covariance between the sales and price residuals. To check how important these moments are to the estimates of the parameter values, we re-estimated the model dropping the four second moments. The value of the estimation criterion falls from 746.7 to 165.5, but the point estimates and the standard errors of the structural parameters are not interestingly different. The estimated wage rate rises to 64.2, and the line speed falls to 37.1; the other parameter are essentially unchanged. So these second moments are not playing an essential role in determining the point estimates. To better match these second moments, we tried adding an idiosyncratic multiplicative shock to the production function (8) and an idiosyncratic additive shock the demand curve (3). In each case, these additional shocks were estimated to have tiny variances and thus negligible impacts on the dynamics of the model. Further these additional shocks had very little impact on the value of the criterion (15).

The model's relevance is further bolstered by the fact that it matches some key patterns in the data

that are not explicitly estimated. Figures 3 to 6 plot the the weekly paths of prices, sales, production and inventories shutting down all the shocks (i.e. $\omega_t = 0 \forall t$). Figures 3 and 4 also plot the price and sales trends observed in the data. The model's baseline price and sales paths track the trends in the data reasonably well. The model successfully replicates the downward track in prices coinciding with the hump-shaped pattern in sales. The model also nails the averages. In the data the mean price is \$26,036 and the mean weekly sales rate is 8,952 vehicles; in the model, these averages are \$25,970 and 8,959. Early in the product cycle, the model slightly overestimates prices and underestimates sales. Later in the product cycle (after about week 40) the model underestimates prices and overestimates sales. Nonetheless, overall it does a good job mimicking the levels and basic shapes of these two series.

Figures 5 and 6 plot the weekly baseline time paths for production and inventories. The plant operates at near full capacity (two 60-hour shifts per week) for the first four weeks in the production cycle and then (with exception of holidays) runs two 40-hours shifts per week for the remainder of the product cycle. This pattern generate the negative monthly time trend in the full production regression reported in table 8. Adding shocks breaks up this pattern somewhat, but production is still predicted to be smoother in our model than we observe in the data. The variance of the residual for the full production regression for the model is one-third the variance we see in the data. Overall the plant in the model runs overtime 8 percent of time (versus 30 percent in the data) and is shut down for inventory adjustments 2.5 percent of the time (compared to 6 percent in the data). The hump-shaped pattern of inventories is similar to that observed in the data, and the model generates about the right level of inventories. Specifically, the model predicts an average inventory-to-sales ratio of 75 days of supply with a standard deviation of 25. For the single-source models in our data, this average ratio is 79 with a standard deviation of 51.

The revenue tax in equation (10) plays a central role in generating the time paths for these four series. During the first weeks of the production cycle inventories are naturally low, and thus it is costly to sell a lot of vehicles. In order to reduce this tax in the future, the automaker needs to accumulate inventories. Hence, early on, the automaker sets prices high to dampen sales and produces at a high capacity allowing the inventory stock to rise. Once inventories have reached a sufficient level, the tax effectively disappears and the automaker lowers prices in order to stimulate sales. Further exacerbating this fall in prices, demand for the vehicle decreases as the product cycle progresses.

To measure the model's propensity to use weeklong shutdowns to adjust production, we estimated the probit model described in equation (1) on the set of 298 model simulations (allowing for shocks to ω). The estimated coefficients and standard errors are reported in table 9 along with the estimates from the




Figure 3: Prices




Figure 4: Sales




Figure 5: Production




Figure 6: Inventories

Baseline Time Paths of Prices, Sales, Production, and Inventories

Notes: The red dashed lines in figures 3 and 4 are the price and sales trends from the data.

In all four figures the solid black line is a simulation of the model with all innovations set to zero (i.e. $\omega_t = 0 \forall t$).

parameter	Estimated Coefficients	
	observed	simulated
$\ln(p_{t-1})$	-1.938 (0.793)	-38.32 (4.10)
$\ln(s_{t-1})$	-0.438 (0.133)	-0.528 (0.65)
$\ln(i_t)$	0.797 (0.177)	23.786 (2.55)
m_t	-0.035 (0.024)	2.63 (0.32)
R^2	0.12	0.91

Table 9: Estimated probit parameters using observed and simulated data

Note: Standard errors are in parenthesis.

data first reported in table 4. The first thing to note is that signs on the coefficients of price, sales, and inventories are the same as in the data: the likelihood of a shutdown increases as the level of inventory increases; it decreases as prices and sales rise. However, the three of the four model's coefficients are an order or two magnitude higher than we see in the data. Perhaps this reflects considerable more noise in the data than in the simulations. This possibility is supported by the R^2 on the simulation probit of .87 while the R^2 of the data probit is .12.

4 Dynamics and Conditional Responses

This section examines how the firm responds to persistent demand shocks. The firm's optimal response to these shocks is non-linear and history dependent. Therefore we report the responses of sales, prices and production to negative innovations to z conditioning on three distinct histories. These distinct realizations of prior shocks push the level of inventories, i , and the state of demand, z , into different regions of the state space. Thus, we can measure how the firm responds to shocks of various magnitudes at different points in the state space.

To implement this, we look at negative innovations of varying magnitudes both early and late in the production cycle: weeks 12 and 35.¹⁷ To vary the initial conditions of z and i , we study cases in which the firm faces: 1) no shocks in weeks prior to the innovation, 2) a series of positive shocks in the weeks prior to the innovation; and 3) a series on negative shocks in the weeks prior to the innovation. More precisely,

¹⁷While the firm does not respond symmetrically to positive and negative innovations, the results are qualitatively similar. To conserve on space, we only report results for negative innovations.

in the first case, we shut down all the shocks except for a single innovation to z at week t^* ; that is we set

$$\omega_t = \begin{cases} -\Phi\sigma_\omega & \text{if } t = t^*, \text{ where } \Phi = 0, \frac{1}{2}, 1 \text{ or } 2 \\ 0 & \text{otherwise.} \end{cases} \quad (16)$$

We refer to this first case as the *neutral history* case. In the second, or *positive history*, case we set

$$\omega_t = \begin{cases} -\Phi\sigma_\omega & \text{if } t = t^*, \text{ where } \Phi = 0, \frac{1}{2}, 1, \text{ or } 2 \\ \frac{1}{3}\sigma_\omega & \text{if } t^* - 10 < t < t^* \\ 0 & \text{otherwise.} \end{cases} \quad (17)$$

In the third case, or *negative history*, case we set

$$\omega_t = \begin{cases} -\Phi\sigma_\omega & \text{if } t = t^*, \text{ where } \Phi = 0, \frac{1}{2}, 1, \text{ or } 2 \\ -\frac{1}{3}\sigma_\omega & \text{if } t^* - 10 < t < t^* \\ 0 & \text{otherwise.} \end{cases} \quad (18)$$

In figures 7 to 10 we plot impulse response functions for prices, sales, production, and inventories to a two-standard deviation shock to z during week 12 (month 3). The lines plotted in the four figures are the differences in levels between the response for $\Phi = 2$ and the response for $\Phi = 0$. In each case, the time paths have been aggregated to the monthly frequency.

The main point to take away from these graphs is that prices and sales respond in months immediately following the innovation (which occurs in month 3) but production responds months later. Since automobiles are built-to-stock rather than built-to-order, production does not need to respond simultaneously with prices and sales.

At impact both prices and sales fall. The marginal response of sales is largest in the month right after the shock. The response diminishes over time, and after 5 months or so (month 8), the effects of the shock on sales are largely gone. For the positive and neutral history cases, sales are higher in than without the shock after month eight, even though our demand specification does not have any role for ‘pent-up demand’. Prices, on the other hand, fall only very little in the month after the shock but continue to decline for several months. For all neutral and positive histories, prices never return to their no-shock baseline levels. This persistence in the price response is not due to ‘sticky prices.’ There are no price rigidities in the model. Instead these persistent lower prices are due to the desire to boost sales to make up for lost sales earlier in the model year (i.e ‘pent-up supply’). Further the price response is quite modest; a two-standard deviation shock to demand causes only a \$200 decline in prices on a \$26,000 vehicle.

The production response is completely different. Examination of figure 9 shows that for all three histories output does not respond to the shock for seven months. The output response occurs in months 11 and 12 after the sales response has largely died out. This propagation occurs even though there are




Figure 7: Price Response




Figure 8: Sales Response




Figure 9: Production Response




Figure 10: Inventory Response.

Response of Prices, Sales, Production, and Inventories to a Two Standard-Deviation Negative Innovation to z at Week 12 (Month 3).

Notes: The responses have been time-aggregated to the monthly frequency.

Each line plots the difference between the time path of the variable with $\Phi = 2$ and the time path with $\Phi = 0$.

The black solid line is the response of the variables under the negative history case (i.e. $\omega_{12} = -2\sigma_\omega$; $\omega_t = -\sigma_\omega/3$ for $t = 2, 3, \dots, 11$; and $\omega_t = 0$ for $t > 12$).

The blue dashed line is the response of the four variables under the neutral history case (i.e. $\omega_{12} = -2\sigma_\omega$; $\omega_t = 0$ otherwise).

The red dot-dashed line is the response of the four variables under the positive history case (i.e. $\omega_{12} = -2\sigma_\omega$; $\omega_t = \sigma_\omega/3$ for $t = 2, 3, \dots, 11$; and $\omega_t = 0$ for $t > 12$).

no adjustment costs in the model. Because of the nonconvexities in the firm's cost function, the firm wishes to operate the plant at its minimum efficient scale. In this case, the firm minimizes average cost by running two 40-hour shifts per week producing 3175 vehicles per week. Below the MES the firm can only convexify its cost function over time via temporary shutdowns; therefore the nonconvexities can induce a lag between the price and production responses. Further because higher inventories reduce the revenue tax, the firm prefers to postpone shutdowns until the end of the product cycle. Finally, because of this lag between the two responses, inventories rise in the months immediately after the innovation.

Despite the importance of the nonconvexities in the timing of the production response, the firm can convexify its costs over time. Thus it is not obvious that the optimal longer-term responses are nonlinear. To address this issue, we summarize the impulse response for 18 different innovations in tables 10. For each innovation (i.e. history, week, and Φ), we report the percentage difference between the realization of the model with $\Phi = 1/2, 1$ or 2 and the realization of the model with $\Phi = 0$.

Some basic patterns emerge from table 10. First the marginal price and sales response at impact (columns (2) and (3)) are proportional to the size of the shock. In each case the magnitude of the marginal response with $\Phi = 2$ is twice the magnitude with $\Phi = 1$ and four times that of $\Phi = 1/2$. Also in each case, the sales responses at impact are considerably larger than the price responses. This is partially due to the size of the elasticities estimated in section 3.1. These elasticities measure the sales response for a particular make and model to a change in price holding the demand for all other vehicles constant. If we want to interpret these demand shocks as aggregate demand shocks, the sizes of our elasticities represent upper bounds.

Holding Φ fixed, the price responses are largest (and thus the sales responses are smallest) for the positive history paths. The prices responses are smallest (and thus the sales responses are largest) for the negative history paths. While this effect is small (the price responses for the first four months following the shock in figure 7 look almost identical), the intuition is straightforward. The lower the level of inventories, the higher is the marginal value of an additional unit of inventory to the firm. Since a series of positive shocks to z reduces the level of inventories, it raises the shadow value of inventories to the firm leading the firm to charge higher prices and dampen sales.

In all 18 cases production (column (4)) does not respond in the weeks immediately after the shock. This however does not imply there is no production response. In columns (5)-(7) of table 10 we report the marginal cumulative response of prices, sales and production over the remainder of the product cycle. Unlike the price and sales response at impact, the cumulative price and sales responses are not proportional

Φ	Response at Impact			Cumulative Response over Cycle				
	Price	Sales	Product	Price	Sales	Product	Wages	Profits
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
<i>Negative History: Week 12 Shock</i>								
1/2	-0.004	-4.75	0.00	0.00	-1.10	-1.39	-0.47	-0.27
1	-0.008	-9.48	0.00	-0.04	-1.10	-1.39	-0.47	-0.56
2	-0.016	-18.98	0.00	-0.09	-2.18	-2.74	-0.88	-1.18
<i>Neutral History: Week 12 Shock</i>								
1/2	-0.004	-3.92	0.00	-0.04	-0.00	0.00	0.00	-0.22
1	-0.009	-7.85	0.00	-0.08	-0.00	0.00	0.00	-0.45
2	-0.019	-15.64	0.00	-0.12	-1.11	-1.37	-0.47	-0.96
<i>Positive History: Week 12 Shock</i>								
1/2	-0.005	-3.40	0.00	-0.04	-0.00	0.00	0.00	-0.13
1	-0.009	-6.80	0.00	-0.08	-0.00	0.00	0.00	-0.29
2	-0.020	-13.55	0.00	-0.15	-0.01	0.00	0.00	-0.65
<i>Negative History: Week 35 Shock</i>								
1/2	-0.004	-4.54	0.00	-0.01	-1.71	-3.38	-0.96	-0.19
1	-0.007	-9.09	0.00	-0.08	-1.71	-3.38	-0.96	-0.37
2	-0.015	-18.17	0.00	-0.15	-3.47	-6.85	-2.07	-0.77
<i>Neutral History: Week 35 Shock</i>								
1/2	-0.004	-3.88	0.00	-0.07	-0.01	0.00	0.00	-0.14
1	-0.008	-7.77	0.00	-0.13	-0.01	0.00	0.00	-0.30
2	-0.016	-13.22	0.00	-0.20	-1.77	-3.24	-1.08	-0.62
<i>Positive History: Week 35 Shock</i>								
1/2	-0.004	-3.32	0.00	-0.06	-0.01	0.00	0.00	-0.08
1	-0.009	-6.65	0.00	-0.13	-0.01	0.00	0.00	-0.17
2	-0.018	-13.30	0.00	-0.26	-0.02	0.00	0.00	-0.39

Table 10: Conditional Responses to ω Innovations in Weeks 12 or 35

Note: For each panel, columns (2), (3), and (4) report the percent differences in prices, sales, and production one week after the shock between the $\Phi = 1/2, 1, 2$ cases and the $\Phi = 0$ baseline.

Columns (5), (6), (7), (8), and (9) report the percent differences in the sum of prices, sales, production, wages and profits for the remainder of the product cycle between the $\Phi = 1/2, 1, 2$ cases and the $\Phi = 0$ baseline.

to size of the shock. This non-proportionate result follows from the nonconvexities in the cost function. When the firm operates below its minimum efficient scale, it will tend to use more non-convex margins to adjust output (e.g. weeklong shutdowns); but when it operates above its MES it will tend to use more convex margins of adjustment (e.g. overtime). Hence we would expect the nonconvexities to matter more (and the responses to be more nonlinear) under the negative and neutral history cases and less under the positive history cases.

It is often argued that firms respond to small shocks by changing prices and large shocks by changing production. For example, Haltiwanger and Maccini (1988) find that a fraction of the firms in their model only adjust production in the face of large shocks. This view is generally consistent our results. In four of the six cases (shock week and histories) there is no production response to shocks of $\Phi = 1/2$ and $\Phi = 1$; the shock is absorbed by prices alone. But in four of the six cases with $\Phi = 2$ size shocks the firm uses both margins, prices and production, to absorb the shock.

Finally, we can also examine the response of the discounted streams of wages for the workers and profits for the firm. For both the week 12 shock and week 35 shock, the profit stream falls roughly one-for-one (perhaps a little more than one-for-one) with the size of the shock. It appears that the owners have enough margins at their disposal to dampen the effects of the demand shocks on profits. Also wages fall by only a third of the fall in output.¹⁸

The relationship between the size of the price and production responses and the size of the shock to z merits additional analysis. Figures 11 and 12 report the cumulative relative responses of prices and production under the negative history case as functions of the size of the shock.¹⁹ In both figures, the production decline is a step function in the size of the shock. Between each threshold, prices decline linearly in the size of the shock. Because the production response is fixed between the thresholds and demand is falling linearly in the size in the size of the shock, the firm lowers prices so that supply equals demand. The discrete drops in production thus generate discrete increases in prices.

Globally, the relationships between the size of the shock and the size of the production and price responses are linear. In week 12, the size of the marginal cumulative production and price responses for $\Phi = 2$ are twice the size of the responses for $\Phi = 1$. However important local nonconvexities still survive time-aggregation. In both weeks 12 and 35, a small negative shock to ω (less than one-half of a standard

¹⁸When computing wages, we assume workers on temporary layoffs receive 95 percent of their straight-time 40-hour wage, the statutory replacement rate. We interpret the difference between .95 and our point estimate of υ as representing cross-subsidies from other firms inherent in the UI system.

¹⁹These curves correspond to columns (5) and (7) in table 10.




Figure 11: Week 12 Shock: The Marginal Cumulative Responses of Prices and Production

Note: This figure plots the magnitudes of the marginal cumulative responses of prices and production as functions of the magnitude of a negative shock to z during week 12.




Figure 12: Week 35 Shock: The Marginal Cumulative Responses of Prices and Production

Note: This figure plots the magnitudes of the marginal cumulative responses of prices and production as functions of the magnitude of a negative shock to z during week 35.

deviation) causes output to fall but prices to rise over the remainder of the product cycle. A small increase in the size of the shock during week 12 that pushes the firm over a threshold results in a cumulative production decline of about 1.4 percent; for week 35, this production decline can be as large as 6 percent. In sum, although the nonconvexities make the production decision almost discrete at the weekly frequency (either all-on or all-off), the firm has enough margins to vary the workweek of capital over the remainder of the production cycle so that its costs are quasi-linear. Hence its production and price responses aggregated over time become globally linear to the size of the shock.

5 A demand shock that was and a demand shock that was not

Our model and data set can be used to explain automakers' reactions to two recent events. One is when the Ford Explorer tire-tread separation problems became public during 2000. The second, is the terrorist attacks of September 11, 2001.

5.1 The Firestone/Ford Explorer Tire Recall of 2000

On August 9, 2000, Ford and Firestone issued the second largest tire recall in history, recalling more than 6.5 million tires because of tire tread-separation problems. Tires on several models were recalled, but the majority were mounted as original equipment on the Ford Explorer, a highly popular SUV. Even before the recall, bad publicity surrounding the Explorer had begun to snowball as law firm web sites and television news shows attributed 46 deaths to the tires. Sales of new Explorers fell, while sales for other SUVs rose, as concerns about the Explorer's safety prompted consumers to switch to other models. This episode provides an example of a demand shock to a single make and model.

Figures 13 through 16 show the difference between Ford Explorer's monthly sales, prices, production, and inventories in 2000 and the average monthly sales, prices, production, and inventories for Ford Explorers in all other years in our sample (1999, 2001, 2002, 2003). At the beginning of 2000, prices, sales and production of the Ford Explorer were above their benchmark averages, likely driven by the robust economic growth at that time. By the end of the first quarter, however, sales and prices started to fall relative to their averages, a trend that continued throughout the year. Looking at the scales of the price and sales paths (figures 13 and 14) we see that the relative magnitudes of the responses (roughly 10 to 1 in sales to prices) are consistent with the responses reported from our model in figures 7 and 8.

Ford Explorer production did not immediately react to the fall in consumer demand. Rather, it continued above the benchmark average throughout the first half of 2000, before finally declining in the second




Figure 13: Prices




Figure 14: Sales




Figure 15: Production




Figure 16: Inventories

The Monthly Path of Prices, Sales, Production, and Inventories for the Ford Explorer During the Year 2000
 In each graph the solid blue line displays the difference between the monthly series during 2000 and the average monthly series Ford Explorers in all other years in our sample (1999, 2001, 2002, 2003).

half. In addition to reacting to declining demand, Ford Explorer production was halted for three weeks in August to increase the supply of new tires available for the tire recall. Explorer inventories remained at or below its average through the first half of 2000, before exploding upward in June, July, and August. The slowdown in September production helped bring inventories down, but they still remained high at the end of 2000. Note that inventories and prices are negatively correlated with a correlation coefficient of -0.46.

The Ford Explorer time series of sales, prices, production and inventories in 2000 are generally in line with our model's predictions. As the public began to learn of the Ford Explorer's tread-separation problems in the spring of 2000, consumer demand fell. Similar to the impulse-response graphs generated by our model, Ford initially responded to this fall in demand by lowering price and maintaining production. Then in the latter half of 2000, Ford reacted to the slump in demand for Ford Explorers by cutting production, and bringing inventories back to its historical average.

5.2 Post-September 11, 2001

The tradeoff between automobile price and production was discussed prominently in the popular press during September and October of 2001. In the days immediately following the terrorist attacks of September 11, auto sales fell by one-third and Standard & Poor's reported, "Industry demand is now expected to be exceptionally weak for the next two quarters, at least, and the likelihood of any improvement beyond that time is highly uncertain."²⁰ Ford Motor Company then announced it was cutting third quarter output by 12 percent. This decision was subtly criticized as being detrimental to the macroeconomy during a time of war. Dieter Zetsche, head of Daimler Chrysler AG's Chrysler group stated, "I think it is our responsibility to try to do whatever we can to contribute to stability. Not to overreact ... not to try to pre-empt shortfalls on the demand side with production cuts." GM North American President Ron Zarrella added "... GM has a responsibility to help stimulate the economy by encouraging Americans to purchase vehicles, to support our dealers and suppliers, and to keep our plants operating and our employees working."²¹ After a September 19 meeting in Detroit of Commerce Secretary Donald Evans and Labor Secretary Elaine Chao with top auto executives and union officials, General Motors reaffirmed its existing production schedules and introduced zero percent financing incentives under its "Keep America Rolling" campaign. Ford, Chrysler, and several foreign automakers soon matched these discounts.

²⁰Krebs M.L. "Driving Through and Altered Landscape," *New York Times*, September 23, 2001, section 12, page 1.

²¹Both quotations are from White, G.L. and J.B. White "GM Unveils Interest-Free Offer on All U.S. Model" *Wall Street Journal*, September 20, 2001, page A3. For more quotations on the patriotism of price cuts, see Burton T.M. and J.T. Hallinan "Is It Unpatriotic to Lay Off Workers When the Nation Faces a Crisis?" *Wall Street Journal*, October 2, 2001, page B1.




Figure 17: Industry Price Response




Figure 18: Industry Sales Response




Figure 19: Industry Production Response




Figure 20: Industry Inventory Response.

The Aggregate Time Paths of Prices, Sales, Production, and Inventories During Late 2001 and Early 2002.

In each graph the solid blue line displays the difference between the monthly series during 2001 and the average monthly series for all other years in our sample .




Figure 21: Value of Cash Rebates and Financing Incentives

Patriotism as well as long-term public relations considerations no doubt played key roles in these decisions during the emotional weeks after 9/11; nevertheless we would not expect the automakers to throw profit maximization out the window. To analyze the industry response to the terrorist attacks, we graph the difference between prices, sales, production and inventories levels for every month from June of 2001 through February of 2002 against the average price, sales, production and inventory level for all remaining months in our sample. The first, and a surprising, fact illustrated in figures 17 through 20 is the increase in relative prices from September to October. This is not an artifact of the normalization. Average prices rose 3.7 percent un-normalized. Perhaps even more surprising, this price increase corresponds with a massive sales increase of 42 percent (un-normalized). These price and sales responses are inconsistent with a persistent drop in demand.

Despite the desires voiced by executives to maintain high levels of production, September production was quite a bit lower than the average. This drop in production was largely due to parts disruptions related to increased border security arising after September 11th. October production remained low, however, largely because of a number of inventory shutdowns. Using weekly production data for single source plants, during September and October of 2001, shutdowns for inventory adjustment accounted for 8.7 percent of all production days. This is almost three times as large as the average 3.0 percent of production days factories closed for inventory adjustment during the months of September and October in 1999, 2000,

2002, and 2003.

The conventional wisdom that automakers heavily slashed prices on their vehicles after 9/11 is not confirmed by our data. Despite the zero-percent financing incentives introduced in late September, the average price of new vehicles net of incentives and rebates rose slightly. Part of the explanation lies in the mix of incentives that customers received. In figure 21, we plot the time paths of the average value per vehicle of financing incentives and cash rebates. Automakers increased financial incentives modestly in late 2001. Nonetheless, this increase was more than offset by the drop in cash rebates.

Why did demand not fall, but actually rise during the Autumn of 2001? Some consumers may have been motivated to buy a new car out of patriotism.²² But it appears to us that the zero-interest financing, while not reducing prices, reduced the need for consumers to haggle and search across dealership to find the best deal. Zero-percent financing is an easily understood pricing arrangement and eliminates at least one dimension that car dealers can price discriminate across consumers. It simplifies the buying process much like the “employee discount pricing” programs in the Summer of 2005. It appears that consumers prefer simplified pricing; they were eager to buy and even paid more to avoid more complicated haggling. While the solution to the firm’s decision problem formulated in this paper provides insights into the timing and relative magnitudes of price and production responses, it is silent on the value of price discrimination and opaque pricing to the firm.

6 Conclusion

One often reads statements such as:

With its labor costs fixed because of employment guarantees and large pension and retiree health costs, Detroit can’t adjust supply to meet demand – so it must rely on price adjustments alone.²³

Not only popular discussions of the automobile industry but formal analysis as well has tended to focus either on production adjustments or price adjustments, assuming the other variable is fixed. This paper has shown the Big Three automakers utilize both price and production adjustments concurrently even in the short run. In contrast to the statement above, we find that for sufficiently large demand shocks, most of the adjustments over time are made to supply rather than prices.

²²For example, Freeman S., “September Auto Sales Showed Resilience” *Wall Street Journal*, October 3, 2001, page A2 quotes a consumer who is bought a new PT Cruiser to “do his part” for the economy.

²³Jenkins H., “Why Detroit Can’t Stop Hagglng” *Wall Street Journal*, August 3, 2005, page A11.

Our model suggests that the use of inventories along with the nonconvexities present in the automaker's cost function causes production adjustments to be propagated throughout the model year even though prices and sales move immediately. Thus an observer with a static supply-and-demand model in mind could be misled to believe the supply curve is vertical. This propagation occurs even though there are no adjustment costs to varying the workweek of capital over time. These nonconvexities make the weekly production decision nearly discrete (either all on or all off); but over the course of several months automakers have sufficient margins to dampen the effect of these nonconvexities considerably. As a consequence, production and price responses aggregated over time are near linear in the size of the shock.

References

- [1] Aguirregabiria V., 1999. The dynamics of markups and inventories in retailing firms. *Review of Economic Studies* 66, 275-308.
- [2] Aizcorbe A., 1990. Experience rating, layoffs and unions: a look at U.S. auto assembly plant layoffs. manuscript, Bureau of Labor Statistics.
- [3] Anderson P., Meyer B., 1993. Unemployment insurance in the United States: layoff incentives and cross-subsidies. *Journal of Labor Economics* 11, S70-S95.
- [4] Berry S., Levinsohn J., Pakes A., 1995. Automobile prices in market equilibrium. *Econometrica* 63, 841-890.
- [5] Berry S., Levinsohn J., Pakes A., 1999. Voluntary export restraints on automobiles: Evaluating a trade policy. *American Economic Review* 89, 400-430.
- [6] Bilts M., Kahn J., 2000. What inventory behavior tells us about business cycles. *American Economic Review* 90, 458-481.
- [7] Blinder A., Maccini L.J., 1991a. The resurgence of inventory research: what have we learned? *Journal of Economic Surveys* 5, 291-328.
- [8] Blinder A., Maccini L.J., 1991b. Taking stock: a critical assessment of recent research on inventories. *Journal of Economic Perspectives* 5, 73-96.
- [9] Bresnahan T., Ramey V., 1994. Output fluctuations at the plant level. *Quarterly Journal of Economics* 109, 593-624.
- [10] Chan H., Hall G., Rust J., 2005. Price discrimination in the steel market. manuscript, University of Maryland.
- [11] Copeland A., Dunn W., Hall G., 2005. Prices, production, and inventories over the automotive model year. NBER working paper 11257.
- [12] Corrado C., Dunn W., Otoo M., 2004. An initial look at incentives and prices for motor vehicles: What has been happening in recent years? manuscript, Board of Governors of the Federal Reserve.

- [13] Elmaghraby W., Keskinocak P. 2003. Dynamic pricing in the presence of inventory considerations: Research overview, current practices, and future directions. *Management Science* 49, 1287-1309.
- [14] Federgruen A., Heching A., 1999. Combined pricing and inventory control under uncertainty. *Operations Research* 47, 454-475.
- [15] Galeotti M., Maccini L.J., Schiantarelli F., 2005. Inventories, employment, and hours. *Journal of Monetary Economics* 52, 575-600.
- [16] Hall G., 2000. Non-convex costs and capital utilization: A study of production scheduling at automobile assembly plants. *Journal of Monetary Economics* 45, 681-716.
- [17] Hall R., 1939. Price theory and business behavior. *Oxford Economic Papers* 2, 12-45.
- [18] Hay G., 1970. Production, price and inventory theory. *American Economic Review* 60, 531-545.
- [19] Haltiwanger J.C., Maccini L.J., 1988. A model of inventory and layoff behavior under uncertainty. *Economic Journal* 98, 731-745.
- [20] Haltiwanger J.C., Maccini L.J., 1989. Inventories, orders, temporary and permanent layoffs: an econometric analysis. *Carnegie-Rochester Series on Public Policy* 30, 301-366.
- [21] Hamermesh D.S., 1989. Labor demand and the structure of adjustment costs. *American Economic Review* 79, 674-689.
- [22] Karlin S., Carr R., 1962. Prices and optimal inventory policy. in *Studies in Applied Probability and Management Science*, K. Arrow, S. Karlin, H. Scarf (eds.), Stanford, CA: Stanford University Press.
- [23] Maccini L.J., Rossana, R.J., 1984. Joint production, quasi-fixed factors of production, and investment in finished goods inventories. *Journal of Money, Credit, and Banking* 16, 218-236.
- [24] Newey W., West K., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance-matrix. *Econometrica* 55, 703-708.
- [25] Nordhaus W., Godley W., 1972 Pricing over the trade cycle. *Economic Journal* 82, 853-882.
- [26] Ramey V., Vine D., 2004. Tracking the source of the decline in GDP volatility: An analysis of the automobile industry. NBER working paper 10384.

- [27] Ramey V., West K.D., 1999. Inventories. in Handbook of Macroeconomics, J.B. Taylor, M. Woodford (eds.), Amsterdam: North-Holland.
- [28] Reagan P., 1982. Inventory and price behavior. *Review of Economic Studies* 49, 137-142.
- [29] Rossana R.J., 1990. Interrelated demands for buffer stocks and productive inputs: estimates for two-digit manufacturing industries. *Review of Economics and Statistics* LXXII, 19-29.
- [30] Smith A.A., 1993. Estimating nonlinear time-series models using simulated vector autoregressions. *Journal of Applied Econometrics* 1993, S63-S84.
- [31] Topel R.H., 1982. Inventories, layoffs, and the short-run demand for labor. *American Economic Review* 72, 769-787.
- [32] Whiten T., 1955. Inventory control and price theory. *Management Science* 2, 61-80.
- [33] Zettelmeyer F., Scott Morton F., Silva-Risso J., 2003. Inventory fluctuations and price discrimination: The determinants of price variation in car retailing. manuscript, Yale School of Management.