The Response of Prices, Sales, and Output to Temporary Changes in Demand *

Adam Copeland
Bureau of Economic Analysis†

George Hall
Brandeis University‡

September 2008

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. Using data on production, sales, and transaction prices, we estimate a dynamic profit-maximization model of the firm. We compare the results from two specifications of the model: one with convex costs and one which incorporates realistic plant-level non-convexities into the cost function. Using impulse response functions, we demonstrate that, for either specification, when an automaker is hit with a demand shock to a particular make and model, sales respond immediately and prices respond very modestly. The convex cost specification also predicts that production responds immediately. In contrast, production typically responds only after a delay in the non-convex cost specification, because plant-level non-convexites propagate the shock across time. For both specifications, after aggregating over the model year, shocks are absorbed almost entirely through adjustments in sales and production rather than prices. We examine two recent demand shocks: the Ford Explorer/Firestone tire recall of 2000, and the September 11, 2001 terrorist attacks.

Keywords: automobile pricing, inventories, revenue management, indirect inference
JEL classification: D21, D42, E22, E23, L11, L62

*We thank numerous seminar participants and three anonymous referees 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.

†U.S. Department of Commerce, Bureau of Economic Analysis, Chief Statistician BE-40, Washington, DC 20230; phone: (202) 606-9948; e-mail: adam.copeland@gmail.com ; webpage: http://www.copeland.marginalq.com

‡Department of Economics, Brandeis University, 415 South Street, Waltham, MA 02454-9110; phone: (781) 736-2242; e-mail: ghall@brandeis.edu

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). In many industries, firms have three primary margins for adjustment in the short-run, the period over which the capital stock and number of employees on the payroll is fixed. Firms can increase or decrease sales by adjusting price, raise or lower the level of production by adjusting labor inputs, or allow inventories to accumulate or de-accumulate. The relative costs of these margins determine the shape and slope of the firm's supply curve.

For the most part, the empirical analysis of firms' short-run response to demand shocks has focused on only two of these margins at a time. This restriction may generate misleading results, if in fact firms use all three margins. In this paper, we focus on the automobile industry. Not only popular discussions of the automobile industry but also formal analysis has tended to focus either on production or price adjustments, assuming the other variable is fixed. Indeed, 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.[1]

In contrast, we determine empirically how the Big Three automakers have accommodated shocks to demand explicitly taking into account all three primary margins.

We first document that automakers use all three margins. Consistent with previous work (e.g. Bresnahan and Ramey, 1994), we find that automakers frequently adjust their labor input to increase or decrease production. Further, transaction prices, net of rebates and financing incentives, fall considerably over the model year and dealer inventories are large and volatile. We then argue these margins of adjustment are interrelated, non-convex, and dynamic in nature, leading us to estimate a dynamic profit maximization model of an automaker's choice of adjustment to short-term demand fluctuations. We investigate the role of non-convexities by estimating our model with two different cost function specifications. The first is the convex cost case, which is the functional form typically used in the literature. The second is the non-convex cost case, where we explicitly model the technological and labor constraints faced by automakers.

We report two main findings. First, for either model specification, automakers only modestly respond with changes in price when faced with a demand shock to a particular vehicle. Instead, demand shocks are almost entirely absorbed by changes in sales and production. In our model simulations, we find a 10-to-1 differential between the size of the sales and price responses. Second, under the non-convex cost

[1] Jenkins H., "Why Detroit Can't Stop Haggling" *Wall Street Journal*, August 3, 2005, page A11.

specification, which fits the data better than the convex cost case, the automaker's production response are often delayed and discrete. Because of non-convexities 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 week-long shutdowns to convexify its costs. Hence, in the periods after a demand shock, the rate of production may remain unchanged. In later weeks, however, the firm modifies its level of production by discrete changes in the workweek thus smoothing its production response over time. When examining an automaker's response to a demand shock, then, an empirical analysis of only the weeks surrounding the shock will likely miss the substantial, but delayed, production response.

These results are important because production and price changes of new automobiles have observable effects on the aggregate rate of output growth and the rate of inflation. The motor vehicle sector is a sizable fraction of the economy, accounting for almost 4 percent of real GDP in the past ten years, and has a disproportionately large effect on the volatility of GDP.[2] Further, new motor vehicle prices have sizeable CPI weights of 4.7 percent.[3] We believe that understanding how automakers respond to temporary demand shocks helps in understanding firm pricing and production decisions more generally, given that motor vehicle and many other manufacturing sectors share similar characteristics.

From our reading of the literature, there was a burst of papers written on how firms response to demand shocks in the late 1960s and early 1970s.[4] As with our analysis, these papers typically found that 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 full 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.

We build upon several more recent literatures by considering how firms, in response to demand shocks, utilize the three primary margins of adjustment: price, labor inputs, and inventory. Further, we demonstrate that non-convexities in the cost of production generate a significant temporal dimension to the firm's production response to demand shocks, something missed when considering convex costs of production.

[2]Ramey and Vine (2006) document that motor vehicle production has accounted for almost 25 percent of the variance in aggregate GDP growth in the past 40 years.

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

[4]Nordhaus and Godley (1972) summarize much of this literature.

Much of the traditional inventory literature address the role of inventories on the timing and volatility of output. The bulk of this literature takes sales as given and minimizes the discounted value of expected costs.[5] We build on this literature by, first, embedding the firm's cost minimization problem within a profit maximization framework, and thus endogenizing prices. Second, we explicitly model the costs of various margins of adjustment. Given the highly non-linear cost structure of automobile production, we find this detailed modeling helps capture the within model-year dynamics of prices and production.

In operations research the study of the inventory/price tradeoff falls under the headings *revenue management* or *yield management*.[6] 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.[7] 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. These non-convexities in cost are crucial to understanding how production responses to demand shocks are propagated over the remainder of the model year. In our former paper we explained the coexistence of downward-sloped price profiles with hump-shaped sales and inventories within a deterministic model. In the current paper, we estimate a stochastic model and study how optimal policies are affected by demand disturbances.

A third literature studies the tradeoff between inventories and employment.[8] 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. In these models, however, there is no pricing decision–a potentially important margin in many manufacturing industries.

The remainder of this paper has six parts. In the first and second parts, we develop our model of an automobile assembly plant and present the data. In the third, we solve and estimate the automaker's dynamic decision problem. In the next two parts, we report impulse response functions of price, sales, and production to demand shocks and 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

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

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

[7] 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).

[8] Relevant contributions include Topel (1982), Maccini and Rossana (1984), Haltiwanger and Maccini (1988), Rossana (1990), and Galeotti, Maccini, Schiantarelli (2005), and Ramey and Vine (2006).

event represents a true demand shock. The aggregate time series of prices, sales, and production following the 9/11 attacks, however, do not accord with the expected responses from a negative demand shock. We make summary remarks in the final section.

1 The Model

The model examines an automaker selling a single product.[9] This assumption simplifies the firm's problem along two dimensions. First, we abstract away from strategic interactions between automakers. Given our focus on plant-level decisions within the model year, we believe this simplification still allows us to obtain a good approximation of automaker behavior. Second, we ignore coordination among an automaker's plants. For vehicles produced at multiple plants, this assumption may be troublesome. However, as detailed in our empirical section, we estimate our model using data on vehicles manufactured at a single plant. Both these simplifying assumptions are necessary because of computational constraints.

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, the number of vehicles to produce, q_t, and the retail price of the vehicle, p_t. For the last 52 weeks the firm makes only a pricing decision. The firm's objective is to maximize the present value of the discounted stream of profits:

$$\max_{\{p_t,q_t\}} E\left\{\sum_{t=1}^{104}\left(\frac{1}{1+r}\right)^{t-1}\{p_t s_t - h(i_t) - C(q_t)\}\right\} \tag{1}$$

where s_t is sales, $h(i_t)$ is the cost of holding i_t inventories, and $C(q_t)$ is the cost of production.

Weekly sales, s_t, depend on the vehicle's own price, p_t, the current level of inventories divided by its mean, i_t/i^{mean}, a persistent shock z_t, and a deterministic time-varying constant term μ_t. The weekly demand curves

$$\log s_t = \mu_t(1+z_t) - \eta_t^p \log p_t + \eta_t^v \log\left(\frac{i_t}{i^{mean}}\right) \tag{2}$$

take a log-log specification with η_t^p and η_t^v denoting the week t own-price elasticity and own-variety elasticity. With the variety term (i_t/i^{mean}), we seek to capture the idea that consumers are more likely to purchase a vehicle if they can find one that matches their particular tastes.[10] Within the automobile

[9] We integrate the dealership into the automaker and consider a unified pricing decision. See Blanchard (1983, page 370) for the argument for treating the manufacture and the dealer as a single entity.

[10] Womack et al (1990) emphasize the importance of providing variety, stating that a main reason automakers encourage dealerships to hold large inventories is to have "plenty of cars on hand to provide variety for the walk-in buyer" (p.171). More generally, Kahn (1987,1992) finds that inventories are productive in generating greater sales at a given price.

industry, variety means having vehicles on a dealership lot with all possible combinations of options (e.g. color, leather interior, airbags). Hence, our definition of variety translates into a measure of the number of vehicles at a dealership. Because we do not have data at the dealership level, our proxy for variety is inventories (i.e. the number of cars at dealerships) divided by the mean level of inventories for the appropriate market segment. We do not simply use the level of inventories as our measure of variety, because the number of dealerships by market segment varies. Intuitively, vehicles that appeal to buyers across the U.S. will require larger amounts of inventory to achieve the same level of variety, relative to less popular vehicles only sold in parts of the country. Mercedes-Benz, for example, only had 191 dealerships in the U.S. in 2002, while Honda had 959.[11] Dividing by the mean allows us to compare the inventory accumulation of popular vehicles such as pickups, and its resulting effect on variety, to other vehicles.[12]

While z_t is likely a function of competing vehicles' prices and inventory levels, for computational simplicity we approximate the evolution of this persistent shock using an autoregressive process:

$$z_{t+1} = \rho z_t + \omega_{t+1} \tag{3}$$

with ω distributed i.i.d $N(0, \sigma_\omega)$. 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:

$$s_t \leq i_t. \tag{4}$$

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

$$i_{t+1} = i_t + q_t - s_t \quad 0 < t \leq 52. \tag{5}$$

After 52 weeks no vehicles are produced, so inventories are simply drawn down by sales,

$$i_{t+1} = i_t - s_t \quad 52 < t \leq 104. \tag{6}$$

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

[11] Data taken from Ward's 2002 Automotive Yearbook.

[12] As noted by a referee, a dealership stock-out motive yields the same empirical prediction as our variety story. If dealerships face demand uncertainty, then sales are the minimum of consumer demand and inventories (e.g. Aguirregabiria (1999)). Aggregating over dealerships, we obtain a market-level demand function that depends on inventories.

The firm faces inventory holding costs in the form of:

$$h(i_t) = \phi_1 i_t + \phi_2 i_t^2. \tag{7}$$

Since demand for vehicles is a positive and non-diminishing function of the inventories, without a holding cost term, the firm will accumulate an unrealistic level of inventories.

We study this model of the firm under two different assumptions about its production costs.

Case I: Convex Production Costs A convex specification is the traditional model of production costs. Under this specification, we assume that each week the firm can produce q_t vehicles per week at a cost

$$C^{\text{CVX}}(q_t, g_t) \quad = \quad \gamma_1(1+g_t)q_t + \gamma_2 q_t^{\gamma_3} \tag{8}$$

where

$$g_t = \rho_g g_{t-1} + \varepsilon_t \tag{9}$$

with ε distributed i.i.d $N(0, \sigma_\varepsilon)$.

The linear, per-vehicle term, $\gamma_1(1+g_t)$ incorporates all costs (such as raw materials) that do not depend on the number of vehicles produced per week. So the disturbance g_t includes changes in input prices. If $\gamma_3 = 2$, costs are quadratic; however, since the demand curves are linear in logarithms rather than levels, the model is not linear-quadratic (LQ) even with quadratic costs. Nevertheless, given the similarities between the convex-cost specification and a traditional LQ model, we expect the implication of the two models to be qualitatively similar.

Case II: Non-Convex Production Costs As documented by Bresnahan and Ramey (1994), managers at automobile assembly plants face several important non-convexities in their production choices. In this specification, we model these non-convexities explicitly. Thus, when the firm decides how many vehicles to produce it must also decide how to organize production to minimize costs. We assume 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. Typically, plant managers increase or decrease production by altering the workweek rather than the rate of production, so we fix the number of employees per shift, n, and the line speed, LS. The firm's production function is then linear in hours:

$$q_t = D_t \times S_t \times h_t \times LS \tag{10}$$

Although this 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. Workers on the evening shift are paid 5% more than those on the day-shift. Work in excess of eight hours a day, and all Saturday work, is paid at a statutory rate of time and a half. Since the statutory rate may not equal the true shadow price of overtime (see for example Trejo, 2003), we estimate the overtime premium, ot_{prem}. 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 in a week, it must then choose D, S and h to minimize its cost of production. Given these choices, the firm's week t cost function is expressed as

$$\begin{aligned} C^{\text{nc}}(q_t,g_t) \quad = \quad & \gamma_1(1+g_t)q_t + \min_{D_t,S_t,h_t} \{ (w_1 + I(S_t=2)w_2) \times (D_t h_t n + \max[0, 0.85(40 - D_t h_t)n] \\ & + \max[0, ot_{prem} D_t (h_t - 8)n] + \max[0, ot_{prem}(D_t - 5)8n]) + \upsilon w_1 40(2 - S_t)n \}, \end{aligned} \quad (11)$$

where as in the previous case, γ_1 is the per vehicle material cost, and the cost shock, g_t, follows the autoregressive process described by (9). 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 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.

Because of these kinks, the firm minimizes average costs by operating the plant with either one or two 8 hour shifts 5 days per week, depending on the cost function's parameter values. If the plant's desired output is below this point (i.e. the firm's minimum efficient scale), the firm will minimize cost by taking a convex combination of producing at 0 and producing at its minimum efficient scale.

Under both production-cost specifications, the firm observes ω_t and ε_t before choosing p_t and q_t. Let $V(i,z,g,t)$ be the optimal value at week t for the firm that holds inventory i and observes a demand state of z and a cost state of g. The firm's value function for weeks $t = 1,2,\ldots,52$ can be written:

$$V(i,z,g,t) = \max_{p,q} \left\{ ps(p,i,z) - h(i) - C(q,g) + \frac{1}{1+r} EV(i+q-s,z,t+1) \right\} \quad (12)$$

subject to (2), (3), (4), and (9), $h(i)$ is given by (7) and where $C(q)$ is given by (8) for the convex cost model or by (10) and (11) for the non-convex cost model. For weeks $t = 53,54,\ldots,104$ the value function becomes:

$$V(i,z,t) = \max_p \left\{ ps(p,i,z) - h(i) + \frac{1}{1+r} EV(i-s,z',t+1) \right\} \quad (13)$$

subject to (2), (3), (4), and (7). Hence the firm's pricing and production decisions are governed by the policy functions:

$$\begin{aligned} \tilde{p}_t &= \tilde{p}(i_t,z,g,t) \\ \tilde{q}_t &= \tilde{q}(i_t,z,g,t) \end{aligned} \quad (14)$$

which solve (12) and (13).

2 The Data

We draw upon two different, but related data sets. The data sets differ in their frequency and content but are consistent with one another in areas of overlap.

The first data set, constructed in Copeland, Dunn, and Hall (2005), contains monthly prices, sales, production and inventories by model and model year from 1999 to 2003. Foreign manufacturers are excluded because of problems measuring overseas production. The sales and production numbers come from Wards Communications, while the price data are derived from retail transactions captured at dealerships by J.D. Power and Associates (JDPA).[13] 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.[14] JDPA also reports the average cash rebate and average financial package customers received from the manufacturer.

This data set provides 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 consider only those vehicles produced at a single plant. We then aggregate this single-source data to the plant/model-year level. The resulting data set includes 28 factories and has a total of 149 plant/model-year pairs. This subset of vehicles represents about 34 percent of all Big Three vehicles sold in the U.S. over our sample period.

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

[14]If a customer trades in an old vehicle when purchasing a new vehicle, JDPA compares the price the customer receives on the traded-in vehicle with its wholesale price. If the wholesale price is lower (higher) than the trade-in price, then the price of the new vehicle purchased by the customer is adjusted downwards (upwards) by the difference between the wholesale and trade-in prices. In other words, JDPA adjusts the price of the new vehicle to account for instances when the customer receives a good or bad deal on the traded-in vehicle.

Vehicles produced at single-source plants are like those produced at multiple plants. The mean price of single-source vehicles is $24,910, only slightly above the mean price over all vehicles, $23,241. Further, with the exception of pickup trucks, single-source plants produce sizeable numbers of vehicles in all market segments.[15] The single-source subset also is composed of roughly equal amounts from each of the Big Three, although Chrysler is over-represented.

These single-source data reflect well our modeling assumptions of a single assembly plant producing a vehicle, and provides a complete picture of an average assembly plant's pricing and production decisions. As described in our model, the non-convex cost structure underlying vehicle production (equation 11) is a complicated function, reflecting the various technological and labor constraints faced by automakers. This detailed modeling improves the ability of the model to match the volatility of production.

To better understand what drives this volatility, we examine a second data set, also obtained from Wards Communications, which 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.[16] Since they come from the same source, the weekly production numbers in this data set are consistent with the monthly figures reported the first data set. Once again, because this paper focuses on the operation of a single automobile assembly plant, we examine only those plants which are the sole producers of a vehicle.

This detailed weekly data set provides an excellent picture of the operation of assembly plants, including the frequency with which assembly plants used different margins used to alter production. While this data set is not used to estimate our model, it does influence our cost function specification and is used to check the model's predictions of inventory shutdowns within the model year. We find that assembly plants usually operate at full speed (i.e. each shift works 40 hours a week), or not at all.[17] A clear example of this

[15]Few 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.

[16]We 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*.

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

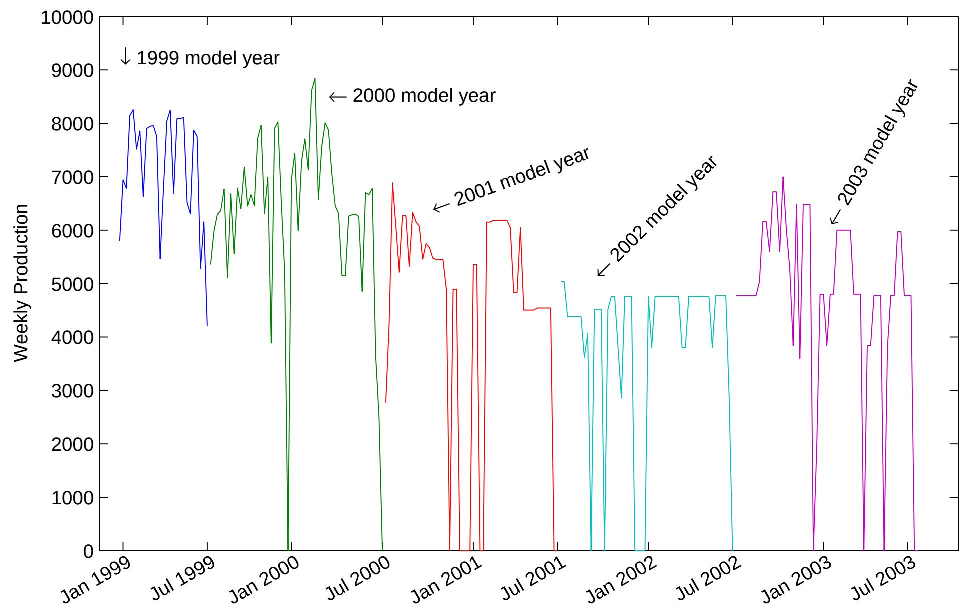

Figure 1: Weekly Grand Cherokee Production

behavior is the weekly output of Chrysler's Jefferson North factory, the sole assembly plant of the Jeep Grand Cherokee (figure 1). 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 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 plant shutdowns are roughly equally attributable to model changeovers, holidays, and inventory adjustments (see table 1). Supply disruptions play a minor role in explaining shutdowns, accounting for less than 5 percent of all factory shutdowns.[18] Table 2 displays the duration of

[18]These 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.

	Model Changeovers	Holidays	Inventory Adjustments	Supply Disruptions
Percent of days shutdown	27.2	37.5	30.8	4.6
Percent of all days	5.6	7.8	6.4	0.9

Table 1: Decomposition of shutdowns

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

shutdowns by type. 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.

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 and variety elasticities that are parameters in the market demand curves, equation (2). Second, taking these market demand curves as given we estimate the remaining parameters via indirect inference.

3.1 Estimating the Demand Elasticities

The demand elasticities are estimated using the approach described in our earlier work, Copeland, Dunn, and Hall (2005).[19] 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

[19] A 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 discrete choices of heterogeneous individuals.

The utility derived from choosing an automobile depends on the interaction between a consumer's characteristics and a product's 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 size and height), which are denoted by X; and those that are unobserved by the econometrician (such as styling or prestige), which are denoted by ξ. We 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 ℓ purchasing product j, dropping the time subscript, as

$$u_{\ell jc} = X_j\beta + \xi_j - \alpha_{\ell c} p_j + \sum_k \varphi_k \iota_{\ell k} x_{jk} + \vartheta_{\ell j}, \tag{15}$$

where p_j denotes the price of product j and $x_{jk} \in X_j$ is the kth observable characteristic of product j. The term $X_j\beta + \xi_j$, where β are parameters to be estimated, represents the utility from product j that is common to all consumers, or a mean level of utility. Included within X is a measure of variety. As mentioned earlier, our proxy for the variety of a model available to consumers is the number of that specific vehicle on dealers' lots, divided by the mean level of inventories for vehicles within the same market segment. Consumers then have a distribution of tastes over the observable characteristics. For each characteristic k, consumer ℓ has a taste $\iota_{\ell k}$, 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 $\alpha_{\ell c}$ measures a consumer's distaste for price increases in quarter $c = \{1,2,3,4\}$. Following Berry, Levinsohn, and Pakes (1999), we assume that $\alpha_{\ell c} = \frac{\alpha_c}{y_\ell}$, where α_c is a parameter to be estimated and y_ℓ 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. Finally, $\vartheta_{\ell j}$ is an i.i.d. extreme value.

Consumers choose among the $j = 1,2,\ldots,J$ 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 data set of prices and sales for the Big Three is used to estimate the model, generally following BLP's algorithm. This is the first data set we described in section 2, before we selected only single-source vehicles. Hence, it includes the full product-line offered by the Big Three from 1999 to 2003, allowing

Market Segment	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8
Compact	2.9	3.2	3.1	3.1	2.9	3.1	3.0	3.3
Full	3.5	3.7	3.7	3.6	3.5	3.6	3.7	3.4
Luxury	3.6	3.7	3.7	3.4	3.6	3.8	3.6	3.3
Midsize	3.3	3.5	3.6	3.5	3.2	3.3	3.5	3.4
Pickup	3.2	3.3	3.5	3.4	3.1	3.2	3.7	3.8
SUV	3.2	3.4	3.4	3.3	3.2	3.4	3.7	3.3
Sporty	3.5	3.9	3.7	3.4	3.5	4.1	4.0	3.3
Van	3.3	3.4	3.5	3.5	3.4	3.4	3.7	3.3
Single Source	3.4	3.6	3.6	3.4	3.4	3.6	3.7	3.4

Table 3: The Absolute Value of Own-Price Elasticities by Market Segment and Quarter

Market Segment	Q1	Q2	Q3	Q4	Q5	Q6	Q7	Q8
Compact	0.51	0.71	0.73	0.66	0.42	0.14	0.14	0.41
Full	0.52	0.76	0.81	0.81	0.45	0.08	0.28	0.50
Luxury	0.53	0.70	0.81	0.85	0.51	0.12	0.08	0.22
Midsize	0.54	0.77	0.74	0.76	0.44	0.11	0.15	0.27
Pickup	0.50	0.73	0.76	0.71	0.44	0.07	0.01	1.57
SUV	0.59	0.74	0.69	0.76	0.45	0.09	0.52	0.76
Sporty	0.41	0.61	0.79	0.65	0.66	0.17	0.08	0.42
Van	0.51	0.76	0.79	0.85	0.51	0.13	0.05	0.02
Single Source	0.49	0.70	0.78	0.76	0.54	0.14	0.09	0.40

Table 4: Own-Variety Elasticities by Market Segment and Quarter

us to accurately estimate each vehicle's own-price and variety elasticities. 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 tables 3 and 4. The own-price elasticities generated by our parameter estimates range between 2.9 and 4.1, indicating 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 a 2.9 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 and illustrates that own-price elasticities for this subset of vehicles vary little across quarters. In general, our estimated elasticities are in-line than those found in the previous literature; BLP, for example, report a range of elasticities between 3 and 6 at the model level.

Our estimates of consumer's own-variety elasticities show variety plays an important role in consumers' automobile purchasing decisions. Over the first 4 quarters of the model's product life, increases in variety significantly bolster demand. Over this period, a 1 percent increase in variety bolsters sales by roughly 0.5 to 0.8 percent. The elasticities slightly decrease in the fifth quarter before plunging downwards to about 0.1 in the sixth quarter. The estimated elasticities in the seventh and especially the eighth quarters are harder to interpret. Few models are sold for more than 6 quarters, and so these estimates are based on a small number of atypical observations.

While we compute elasticities by quarter, our model of the firm is at the weekly frequency. To construct the weekly demand curves, equation (2), we interpolate the estimated quarterly own-price and own-variety elasticities for the typical single-source vehicle to the weekly frequency using a spline. To compute the intercept terms μ_t, $t = 1, 2, ...104$, 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

Taking these demand curves as given, we turn to estimating the structural model described in section 1. 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. Because this paper focuses on explaining the dynamics of the automaker's problem at the assembly plant level, we use the monthly single-source plant-level data set on sales, prices, inventories, and production described in section 2. To capture the dynamics of the automaker's problem, we choose as moments the regression coefficients from three least-squares 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 time 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 and so control for plant-level fixed effects. Let $\hat{x}_t$ denote a variable minus its plant-level mean.

Formally, we estimate

$$\begin{array}{rcccccccccc} \hat{s}_t & = & \theta_1\hat{s}_{t-1} & + & \theta_2\hat{p}_{t-1} & + & \theta_3\hat{i}_t & + & \theta_4\hat{m}_t & + & \nu_t^s \\ \hat{p}_t & = & \theta_5\hat{s}_{t-1} & + & \theta_6\hat{p}_{t-1} & + & \theta_7\hat{i}_t & + & \theta_8\hat{m}_t & + & \nu_t^p \\ \hat{q}_t & = & \theta_9\hat{s}_{t-1} & + & \theta_{10}\hat{p}_{t-1} & + & \theta_{11}\hat{i}_t & + & \theta_{12}\hat{m}_t & + & \nu_t^q \end{array} \quad (16)$$

where ν is an i.i.d. normal error. Automakers typical produce a particular vehicle for 12 months, but through the use of inventories sell the vehicle over a 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 months of data.

In addition to the 15 regression coefficients, we augment the vector of moments with 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.[20] These last three equations provide the mean levels of sales, prices, and production at a single-source plant for the model to match. In the language of EMSM, these six regressions compose our auxiliary model. We chose this set of moments because the regression coefficients and error covariance matrix capture the dynamics of prices, sales, and production, as evidenced by their high R-squared (see table 6).

In addition to the demand curves, we fix several supply-side parameters prior to the estimation. For both production-cost specifications, we set the "scrap value" of vehicles unsold after 104 weeks, $\bar{p}_{105}$, to \$15,000. For the non-convex cost specification, we set the number of workers per shift, n, to 1,300. We set 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. The vector of the structural parameters we estimate is $\beta = \{r, \gamma_1, \gamma_2, \gamma_3, \phi_1, \phi_2, \rho_z, \sigma_\omega, \rho_g, \sigma_\varepsilon\}$ for the convex cost specification and $\beta = \{r, \gamma_1, LS, w_1, \upsilon, ot_{prem}, \phi_1, \phi_2, \rho_z, \sigma_\omega, \rho_g, \sigma_\varepsilon\}$ for the non-convex cost specification.

The basic strategy to estimate either model is:[21]

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 (16) as well as the least square estimates of the mean level of sales, price, and production, $\hat{\theta}_T$.

2. For a given set of parameters β, solve the structural model.

[20] We 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.

[21] Since 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.

3. Simulate the structural model for 104 weeks S times and time-aggregate each simulation to the monthly frequency to create a $24 \times S$ panel data set $y(\beta)$. For each simulation, initialize z and g using draws from their ergodic distributions.

4. Estimate the auxiliary model, using $y(\beta)$ to compute $\hat{\theta}_S^{\beta}$. Measure the distance between the vector of observed moments and the vector of simulated moments via the criterion:

$$(1+\pi^{-1})^{-1}(\hat{\theta}_T - \hat{\theta}_S^{\beta})' W_T (\hat{\theta}_T - \hat{\theta}_S^{\beta}) \tag{17}$$

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 auxiliary model. 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 the variance-covariance matrix of the observed parameters taking into account the mis-specification of the auxiliary model. The term π denotes the ratio of the simulation sample size to the data sample size.

5. Using a hill-climbing algorithm, repeat steps 2 - 4 to find the $\tilde{\beta}_T$ that minimizes (17).

We set the number of simulations S to 298, twice the number of plant/model years in our data set, thus $\pi = 2$. For both the convex cost and non-convex cost specifications, we discretize the inventory grid into 29 points from 0 to 50,000. We discretize the z grid into 7 points from -0.10 to 0.10 and the g grid into 7 points from -0.35 to 0.35. For all three grids the points are more densely spaced near zero where the value function has more curvature. For each of the 1,421 (i,z,g) triplets, we maximize recursively the right hand side of equations (12) and (13). Points off the i, z and g gridpoints are approximated using linear interpolation, and all integration is done by quadrature.

For the non-convex cost specification, 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. There are up to 72 feasible production schedules to evaluate for each 121 possible levels of production. 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 do not impose any holiday closures on the convex-cost specification.

Since log-log demand curves do not have an intercept, we fix an upper bound on the sales price, p_t. Above this price, demand for the vehicle is zero; this is consistent with consumers fully substituting to other, presumably nicer, models at some price. This upper bound never explicitly binds, but without it the firm will sell its last few vehicles for unrealistically high prices.

3.3 Empirical Results

In table 5 we report point estimates for the structural parameters for both the convex cost and non-convex cost specifications together with their estimated standard errors.[22] For both cases, the estimated parameter values are sensible. While the two specifications differ on their average production and holding costs, they yield similar predictions on the average profit per vehicle.

Under the convex cost specification, the per-vehicle linear cost, γ_1, is estimated to be \$18,679. The curvature parameter, γ_3, is estimated to be 1.95 with a standard error of 0.15, so the cost function is essentially quadratic. Over the model year, the average cost of producing a vehicle is \$19,230. With an average sales price of \$26,970, the average gross profit per vehicle is about \$7,740. The inventory-holding cost parameters, ϕ_1 and ϕ_2, imply that the average holding cost per vehicle sold is about \$2,880. Thus the average profit per vehicle net of holding costs is \$4,861, or 18% of the sales price.

Under the non-convex cost specification, the point estimate of the first-shift wage rate, w_1, at \$53.45 per hour, is reasonable if one includes benefits, but it is not particularly interesting since it can be scaled up and down by the choice of n. Our estimates of the unemployment replacement rate, υ, and the overtime premium, ot_{prem}, are of more economic interest. They are estimated to be 40.4% and 24% respectively – roughly half the statutory rates of 95% and 50% – though ot_{prem} has a rather large standard error. Nevertheless, these estimates suggest that these statutory rates are not allocative.[23] The line speed point estimate of 39.9 vehicles per hour is consistent with the observed line speeds of 30 to 70 vehicles per hour. Taken together, the estimated parameters, $\{LS, w_1, \upsilon, ot_{prem}\}$, imply an average per-vehicle labor cost of \$2,019. With a point estimate for γ_1 of \$18,087, the average per-vehicle production cost is \$20,106. While this is about \$900 more than the implied production cost from the convex model specification, the inventory-holding cost parameters imply that the average inventory holding cost per vehicle sold is about \$1,998,

[22] In the appendix, we discuss identification of the structural parameters for the non-convex specification.

[23] Trejo (2003) writes down a model of labor market equilibrium in which straight-time hourly wages adjust to neutralize the statutory overtime premium. A point estimate of less than 50 percent is consistent with partial adjustment of straight-time wages.

specification	$r^{\dagger}$	γ_1	γ_2	γ_3	LS	w_1	υ	ot_{prem}	ϕ_1	ϕ_2	ρ_z	σ_ω	ρ_g	σ_ε
convex cost	0.0182	18,679	0.219	1.95					117.1	0.00275	0.934	0.00996	0.956	0.0210
	0.0014	200	0.306	0.15					7.2	0.00026	0.013	0.00058	0.010	0.0012
non-convex cost	0.0163	18,087			39.9	53.45	0.404	0.244	65.0	0.00204	0.937	0.00979	0.936	0.0443
	0.0023	319			1.2	10.22	0.046	0.276	4.3	0.00011	0.009	0.00076	0.018	0.0112

Table 5: EMSM estimates of the structural parameters

Note: The first row for each case reports point estimates. The second row reports estimated standard errors.

† The interest rate r is reported at an annual rate.

roughly \$900 less than implied by the convex cost specification. Hence the sum of the per vehicle production and inventory holding costs are almost the same across the two specifications. Since the average sales price, \$27,189, is slightly higher under the non-convex cost specification, average profits are also slightly higher, \$5,085, or 19% of the sales price.

The real interest rate is estimated to be almost 2% at an annual rate for both specifications. These point estimates are on the low side, suggesting that some of the costs of postponing sales are being picked up by the inventory holding cost parameters.

For both specifications, the demand-side shock process, z, is estimated to be persistent with an autoregressive coefficient of 0.934 (convex cost) and 0.937 (non-convex cost). Both estimates of $\{\rho_z, \sigma_\omega\}$ imply z has a mean of zero (by assumption) and a standard deviation of 0.028. While a standard deviation of 2.8 percent may seem small, a one standard deviation movement in z results in a shift in the demand curve of typically about 400 (and up to 1,400) vehicles per week, depending on the values of μ_t and z_t.

For the supply-side shock, both point estimates of $\{\rho_g, \sigma_\varepsilon\}$ imply the g processes have mean zero ergodic distributions with standard deviations of .0716 (convex case) and 0.12 (non-convex case). In the model, the marginal cost of selling a vehicle is the shadow value of an additional unit of inventory. Since the inventory stock can be over 15 times the weekly flows of vehicles being built and sold, the model needs large and persistent shocks to the cost of production to generate significant movements in marginal cost. Consequently, the g process appears to be incorporating changes in the cost of having an additional vehicle in inventory beyond simple changes in the cost of production.

While the point estimates and average vehicle costs are similar across the two specifications, each case has different implications about the organization of production. Unlike the convex cost specification, the model with non-convex costs implies all-on or all-off production behavior, which generates time-series predictions of sales, prices and production that better fit the data.

We can see these differences in table 6, which tabulates the three sets of estimated moments: one for the observed data, a second for the convex cost specification, and a third for the non-convex cost specification. Recall that the structural parameters in table 5 minimize the difference between these regression moments from the two simulated models and their data counterparts.[24] For the non-convex cost specification all but one of the simulated moments are of the same sign and magnitude as the observed moments. The convex

[24] Our 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 that price plays a small role in absorbing increases in demand are consistent with those of Hay.

cost case replicates these moments slightly less well, getting the sign wrong on four of them.

In the data, both sales and prices are highly persistent. The estimated coefficient on lagged prices in the price equation is a high 0.83, while for the sales equation our estimate on lagged sales in 0.59. Further, 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.

We turn first to the convex cost specification. There are four moments that this specification has difficulty matching, all involving prices. First, in the sales equation, the convex cost specification estimates a negative relationship between sales and lagged price, while in the data we find a positive relationship. Second, in the price equation, the convex cost model does not generate the negative relationship between prices and inventories seen in the data. Third, in the data we find the covariance of the sales and price regression residuals is negative; under the convex cost specification, this covariance is positive. Fourth, in the production equation the convex cost specification does not generate a significantly positive relationship between production and lagged price. Because these moments capture correlations in the data, we cannot assign economic stories to these four discrepancies between the data and convex cost case. But we believe the convex cost specification's inherent inability to match the all-on and all-off behavior of production both drives the discrepancies between price and production, and pollutes the relationship between price and sales.

In contrast, the non-convex cost specification is better able to mimic the volatile production behavior in the data. Accordingly, this specification more closely matches the moments. For both the sales and price equation, the non-convex cost specification performs well, capturing all the significant relationships between the dependent and independent variables. Further, this specification matches the negative correlation between the sales and price regression residuals. Even taking realistic non-convexities into account, this specification has some difficulty matching the production equation in that it does not find a positive relationship between beginning-of-period inventories and production. Further, for both cases the R-squared statistic on the production regression is much lower compared to the statistic based upon the data. We believe this mainly because production decisions in the real world are constrained by supply chain networks and other factors outside of our model.

The estimation criterion (17) provides a test-statistic for the over-identifying restrictions of the model.[25]

[25] Readers may notice that the criterion in (17) is not scaled by the number of observations. As discussed above, 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.

	Sales Equation			Price Equation			Production Equation		
variable	observed	convex	non-convex	observed	convex	non-convex	observed	convex	non-convex
lagged price	0.191	-0.072	0.112	0.829	0.810	0.737	0.904	-0.079	0.206
	0.033	0.027	0.023	0.050	0.013	0.007	0.141	0.103	0.073
lagged sales	0.588	0.487	0.483	0.045	0.025	0.076	0.424	0.182	0.257
	0.023	0.015	0.012	0.011	0.007	0.007	0.061	0.041	0.039
inventories	0.115	0.167	0.161	-0.0087	0.0017	-0.0144	0.054	0.094	-0.016
	0.008	0.005	0.004	0.0021	0.0021	0.0022	0.020	0.021	0.016
trend	-0.055	-0.099	-0.034	-0.037	-0.042	-0.087	-0.265	-0.825	-0.694
	0.014	0.012	0.010	0.011	0.006	0.007	0.064	0.045	0.049
resid variance	4.95	3.50	2.80	0.70	0.63	0.88	15.73	21.91	20.04
	0.22	0.12	0.08	0.02	0.14	0.14	0.89	0.54	0.53
R-squared	0.88	0.83	0.86	0.99	0.69	0.68	0.67	0.10	0.12
Observations	2019	4768	4768	2019	4768	4768	1205	3278	3278

variable	observed	convex	non-convex
cov(resid sales, resid price)	-0.049	0.122	-0.041
	0.058	0.025	0.030

	Sales Equation			Price Equation			Production Equation		
variable	observed	convex	non-convex	observed	convex	non-convex	observed	convex	non-convex
constant	8.20	9.25	9.04	26.05	26.94	27.19	11.41	12.19	12.06
	0.22	0.07	0.07	0.32	0.02	0.03	0.32	0.10	0.09

Table 6: Estimated regression moments using observed data and simulated data from the convex cost and non-convex cost models

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

This statistic is distributed $\chi^2(n-k)$. In the convex case there are nine over-identifying restrictions (n-k=19-10), and the statistic is 401.5. For the non-convex case, there are seven over-identifying restrictions and the statistic is 308.5. Thus for both specifications our structural model can be overwhelmingly rejected as the true data-generating processes of the observed time series.

Nevertheless, the model, particularly the non-convex cost specification, captures much of the interesting dynamics in the data. Indeed, the model's relevance and goodness-of-fit is bolstered by the fact that it matches some key patterns in the data that are not explicitly estimated. In figure 2 we plot the the weekly paths of prices, sales, and production shutting down all the shocks (i.e. $\omega_t = 0$ and $\varepsilon_t = 0 \; \forall \; t$) for both specifications alongside corresponding trends in the data.

The simulated paths of these series are more jagged than the data. The data paths are naturally smooth since they are averages across many models and years, while the model simulation is just a single run. Some of the jaggedness in the price series, particularly in weeks greater than 60, are due to computational approximation errors. The optimal price of the vehicle is pinned down by the shadow value of an additional unit of inventory to the firm. This shadow value is the derivative of the value function with respect to inventories. Since we are linearly interpolating between grid points on the value function, there are discontinuities in this derivative.

For both specifications, the model successfully replicates the downward trend in prices coinciding with the hump-shaped pattern in sales. For the first twenty weeks in the product cycle, though the model overestimates prices and underestimates inventories and sales. Then after about week 20, the model, while still overestimating prices, overestimates inventories and sales. During the end of the production cycle, the firm wishes to build-up inventories to continue to sell once production terminates in week 50. Consequently, the model predicts that inventories peak at week 51, which is at odds with the data. Nonetheless, overall the model, with either specification, does a good job replicating the major trends in the data.

The production graphs in figure 2 plot the weekly baseline time paths for production under the two specifications. Under the non-convex cost assumption, the plant operates two 60-hour shifts (full capacity) for the first three weeks, two 48-hour shifts (Saturday overtime) for the next four weeks, and then (with the exception of holidays) runs two 40-hours shifts per week for the remainder of the product cycle. This pattern generates the negative monthly time trend in the full production regression reported in table 6. Production is predicted to be more volatile than we observe in the data. The variance of the residual for the production regression is one-third higher than the variance we see in the data. Overall the plant in the non-convex cost specification runs overtime 36.7 percent of the time (versus 30 percent in the data) and is

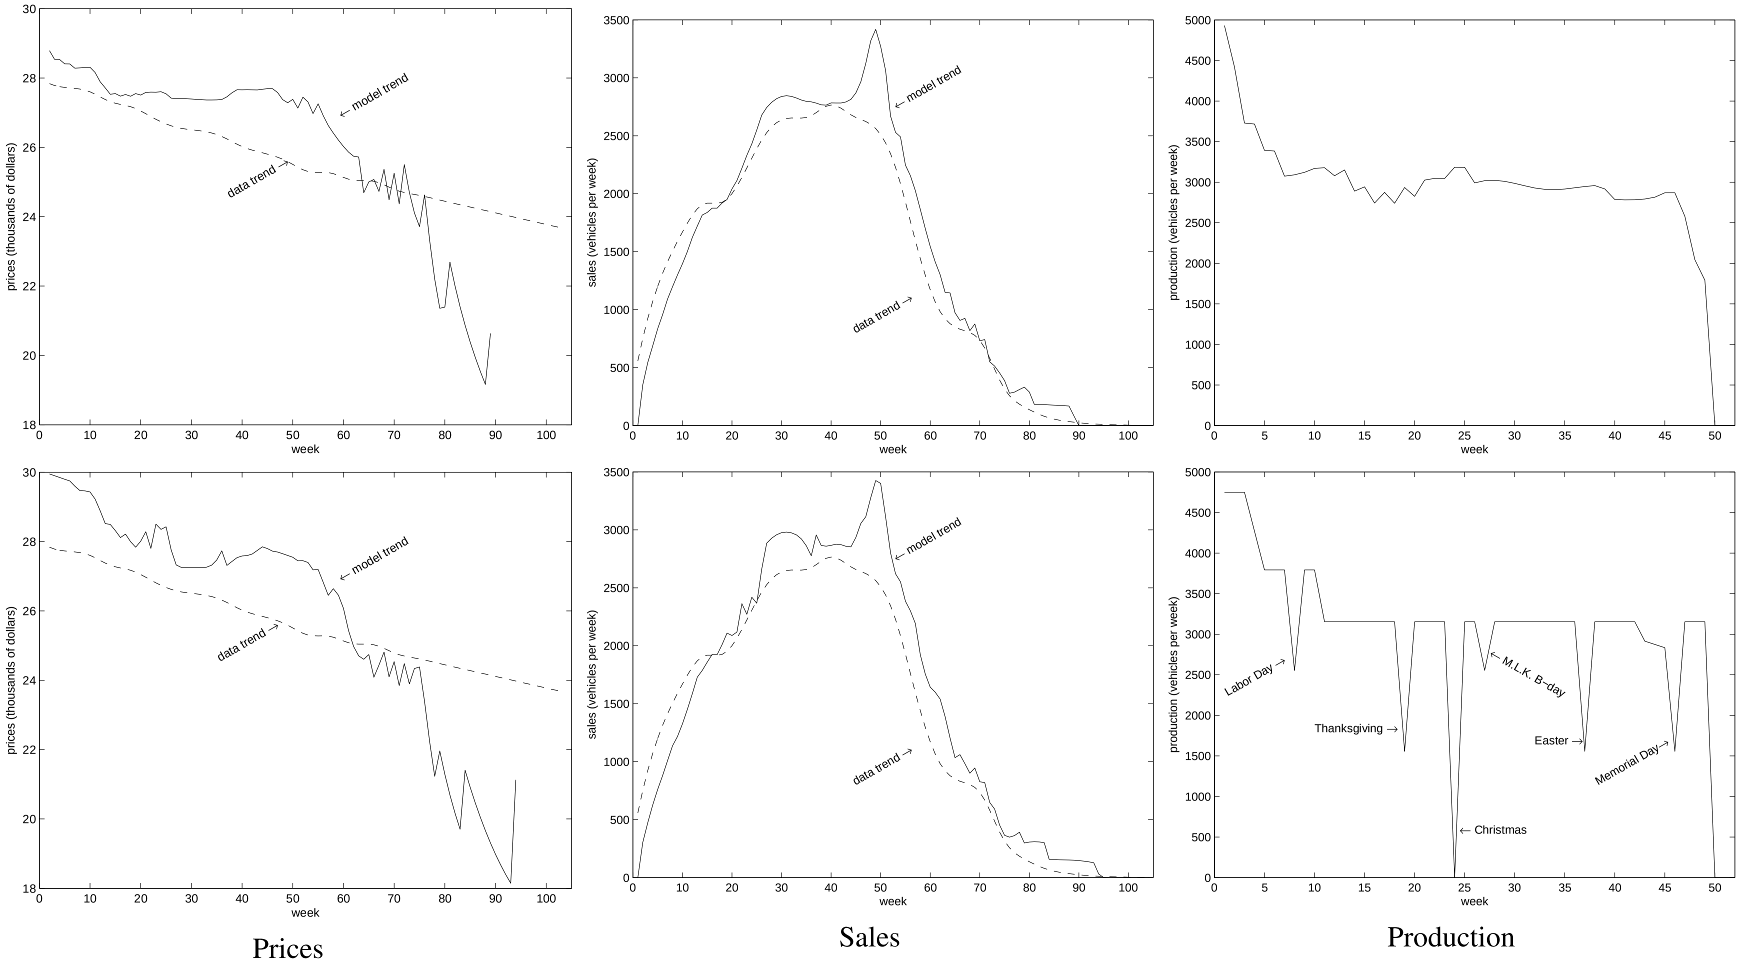

Figure 2: Baseline Time Paths of Prices, Sales, Production, and Inventories for the Convex Model (Top Panel) and Non-Convex Model (Bottom Panel).

Notes: The dashed lines in price and sales graphs are the price and sales trends from the data.
In all six figures the solid line is a simulation of the model with all innovations set to zero (i.e. $\omega_t = 0$ and $\varepsilon_t = 0 \, \forall \, t$).

shut down for inventory adjustments 10.7 percent of the time (compared to 6.4 percent in the data). The hump-shaped pattern of inventories is similar to that observed in the data, and the model generates the right level of inventories. Specifically, the non-convex cost model predicts an average inventory-to-sales ratio of 68 days of supply with a standard deviation of 16. For the single-source models in our data, this average ratio is 70 with a standard deviation of 28.

The convex cost specification, by construction, is silent about shift changes, overtime, and inventory adjustments. It too, however, captures the downward time-trend in production and generates a hump-shaped pattern of inventories. Further, the convex cost case predicts an average inventory-to-sales ratio of 64 days of supply with a standard deviation of 15.

The own-variety elasticity term in the demand curves, equation (2), plays a critical role in generating the time paths for these three series. During the first weeks of the production cycle, inventories are naturally low and thus demand is depressed. In order to increase demand in the future, the automaker needs to accumulate inventories. Hence, early on, the automaker sets prices high, dampening sales and producing at "full" capacity allowing the inventory stock to rise. Once inventories reach about 35,000, the benefits of additional inventories are offset by the quadratic holding cost term, equation (7), 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.

As a last check on the model's goodness-of-fit, we measure its propensity to use weeklong shutdowns to adjust production. We accomplish this by estimating a probit model of inventory shutdowns on prices, sales and inventories for both the data and 298 simulations from the non-convex cost specification. As mentioned earlier, the convex cost case is silent on issues regarding shutdowns and other margins of adjustment in production. Let the dependent variable, Y, be equal to one if the assembly plant was shutdown for inventory adjustment at some point in the month.[26] Because price, sales, and inventories all have particular shapes over the model year, we want to de-trend these variables before analyzing their relationship with plant shutdowns; thus we regress price, sales, and beginning-of-period inventories on a quadratic model-year trend. Denoting $\{\tilde{p}, \tilde{s}, \tilde{i}\}$ as the price, sales, and inventory residuals from these regressions, we

[26] We use the weekly production data to determine if a plant closed down for inventory adjustment at some point in the month. There is little information lost by considering inventory shutdowns as a binary event. Over 70 percent of the inventory shutdowns in our data are one week long, with almost all of the remaining instances lasting two weeks.

variable	Data				Non-convex Model			
	Probit 1		Probit 2		Probit 1		Probit 2	
lagged price	-0.082	(0.039)	-0.165	(0.133)	-0.046	(0.029)	-0.030	(0.043)
twice lagged price			0.103	(0.131)			0.121	(0.043)
lagged sales	-0.116	(0.023)	-0.106	(0.033)	-0.085	(0.015)	-0.219	(0.024)
twice lagged sales			-0.058	(0.033)			0.114	(0.023)
inventories	0.034	(0.006)	-0.059	(0.019)	0.0003	(0.006)	-0.077	(0.001)
lagged inventories			0.108	(0.022)			0.138	(0.011)
trend	-0.078	(0.020)	-0.116	(0.029)	-0.014	(0.009)	0.037	(0.012)
R-squared	0.146		0.188		0.258		0.376	
Observations	1,057		909		3,278		2,980	

Table 7: Estimated Probit Explaining Inventory Shutdowns

Note: Standard errors are in parenthesis. The dependent variable is an indicator function equal to one if the plant is shut down anytime during the month to adjust its inventory.

estimate two probit models: one with only one-period lags and the other with one and two-period lags,

$$Pr(Y_t = 1) = \Phi\left(\eta_1 \tilde{p}_{t-1} + \eta_2 \tilde{s}_{t-1} + \eta_3 \tilde{i}_t + \eta_5 m_t + \sum_k I_{f_t=k} \kappa_k\right), \tag{18}$$

$$Pr(Y_t = 1) = \Phi\left(\eta_1 \tilde{p}_{t-1} + \eta_2 \tilde{p}_{t-2} + \eta_3 \tilde{s}_{t-1} + \eta_4 \tilde{s}_{t-2} + \eta_5 \tilde{i}_t + \eta_6 \tilde{i}_{t-1} + \eta_7 m_t + \sum_k I_{f_t=k} \kappa_k\right), \tag{19}$$

where Φ is the cdf of the normal distribution, m is a model-year trend, f identifies a plant, and $I_{x=y}$ is an indicator function equal to 1 if x equals y. This last term captures plant-level fixed effects. The estimated coefficients are shown in table 7. With only one-period lags, all coefficient estimates using actual data are statistically significant and have the expected sign. If prices or sales are high in the previous months, indicating strong demand, then the probability of the assembly plant shutting down in the current month decreases. Higher beginning-of-period inventories increase the probability of shutting down, and, everything else equal, plants are less likely to shutdown later in the model year. Turning to the second probit with one and two-period lags, the results are less clear. The coefficients on the two price lags are no longer statistically significant and have opposite signs. But the sales lags still have a significant and negative effect. Further, while current beginning-of-period inventories are now negatively correlated with shutdowns, the lagged inventories have a stronger, positive correlation. 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 for the two models are between 0.15 and 0.19.

The estimated profit coefficients using simulated data generated by the non-convex cost specification demonstrate similar patterns. For the probit model with one-period lags, higher prices and sales last period are associated with fewer plant shutdowns in the current period; shutdowns are also less likely later in the model year. However unlike what we see in the data, the coefficient on current inventories is effectively zero. For the probit model with one and two-period lags, the estimated coefficients on the simulated data match up well with those estimated on the data, except for the trend and two-period lag on sales.

The cumulation of all these results demonstrate two points. First, the model, under either specification, fits the data well. Second, the non-convex cost specification replicates an automaker's adjustment of production margins, allowing it to better fit the data compared to the convex cost case. In particular, the non-convex cost specification does well in capturing firms' propensities to use week-long inventory shutdowns.

4 Dynamics and Conditional Responses

This section examines how the firm under both cost specifications responds to persistent demand shocks to a particular make and model. The firm's decision rules (equation 14) are non-linear functions of the four state variables. In particular, for the non-convex specification there are threshold levels of inventories below which the firm wishes to operate "all on" (e.g. two 40-hour shifts per week) and above which it will operate "all off" (e.g. an inventory shutdown). Since prices are a function of the shadow value of inventories, there are discrete jumps at these thresholds in the pricing rule as well. Thus, we want to measure how the firm responds to shocks in disparate regions of the state space. We report the responses of sales, prices, and production to innovations in 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 which the firm is likely to inhabit.

To vary the initial conditions of z and i, we consider three alternatives: 1) no shocks in the weeks prior to the innovation, 2) a series of positive shocks in the weeks prior to the innovation; and 3) a series of negative shocks in the weeks prior to the innovation. More precisely, in the first alternative, we shut down all the shocks except for a single innovation to z at week t^*; that is we set

$$\omega_t = \begin{cases} \Lambda\sigma_\omega & \text{if } t = t^*, \text{ where } \Lambda = \{-1, 0, 1\} \\ 0 & \text{otherwise.} \end{cases} \tag{20}$$

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

$$\omega_t = \begin{cases} \Lambda\sigma_\omega & \text{if } t = t^*, \text{ where } \Lambda = \{-1, 0, 1, \} \\ \frac{1}{4}\sigma_\omega & \text{if } t^* - 10 < t < t^* \\ 0 & \text{otherwise.} \end{cases} \tag{21}$$

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

$$\omega_t = \begin{cases} \Lambda\sigma_\omega & \text{if } t = t*, \text{ where } \Lambda = \{-1,0,1\} \\ -\frac{1}{4}\sigma_\omega & \text{if } t^* - 10 < t < t^* \\ 0 & \text{otherwise.} \end{cases} \tag{22}$$

In the top panel of figure (3) we plot impulse response functions for prices, sales, and production to a negative one-standard deviation shock to z during week 14 (month 4) under the convex-cost specification. The lines plotted in these three graphs are the percent differences between the response for $\Lambda = -1$ and the response for $\Lambda = 0$. In each case, the time paths have been aggregated to the monthly frequency. To determine whether the response to the shock "washes out" over time, we plot in the lower panel of figure 3 the sum of the response over time. We repeat this exercise in figures 4 to 6, reporting the responses to a positive one-standard deviation shock to z (i.e $\Lambda = 1$) as well as the responses to negative and positive shocks in the non-convex cost specification.

There are two main points to take away from these four figures. First under the convex cost specification, all three series, price, sales, and production, respond immediately and relatively smoothly to the shock. In contrast, under the non-convex cost specification, prices and sales respond in the months immediately following the innovation but production tends to respond months later. Because automobiles typically are built-to-stock rather than built-to-order, production does not need to respond simultaneously with prices and sales.[27] Since under the non-convex cost case production may not immediately adjust, more of the shock is transmitted to prices than in the convex cost specification. Second, under both specifications, the price responses are quite small. The magnitude of the sales response is over 15 times larger than the price responses for the convex cost specification and over 8 times larger for the non-convex cost specification. While we estimate demand to be quite elastic, with own-price elasticities around 3, we get more than a 10-to-1 differential in the magnitude of the sales and price responses. Under both specifications almost the entire shock is ultimately absorbed through changes in sales and production.

In figures 3 and 4 we see that under the convex cost specification the firm adjusts all three margins at impact. For both positive and negative shocks, the marginal response of sales and production are largest in the month right after the shock.[28] Prices respond very little (only about 6/10 of one percent or about $150) in the month after the shock and quickly return to the baseline path. This modest response in the price

[27] In this paper, we simply assert the usual assumption that automobiles are built-to-stock. Empirically, the most compelling evidence for this assumption is the large days-supply of vehicles held by automakers. Further, Womack et al (1990, page 174) describe how US automakers in the 1980s heavily pushed to eliminate special orders with the aim to improve efficiency in their factories and supply chains.

[28] We suspect the same thing can be said about prices as well. The large price responses after month 15, when few inventories are held and few vehicles are sold, appear to be largely due to approximation error.

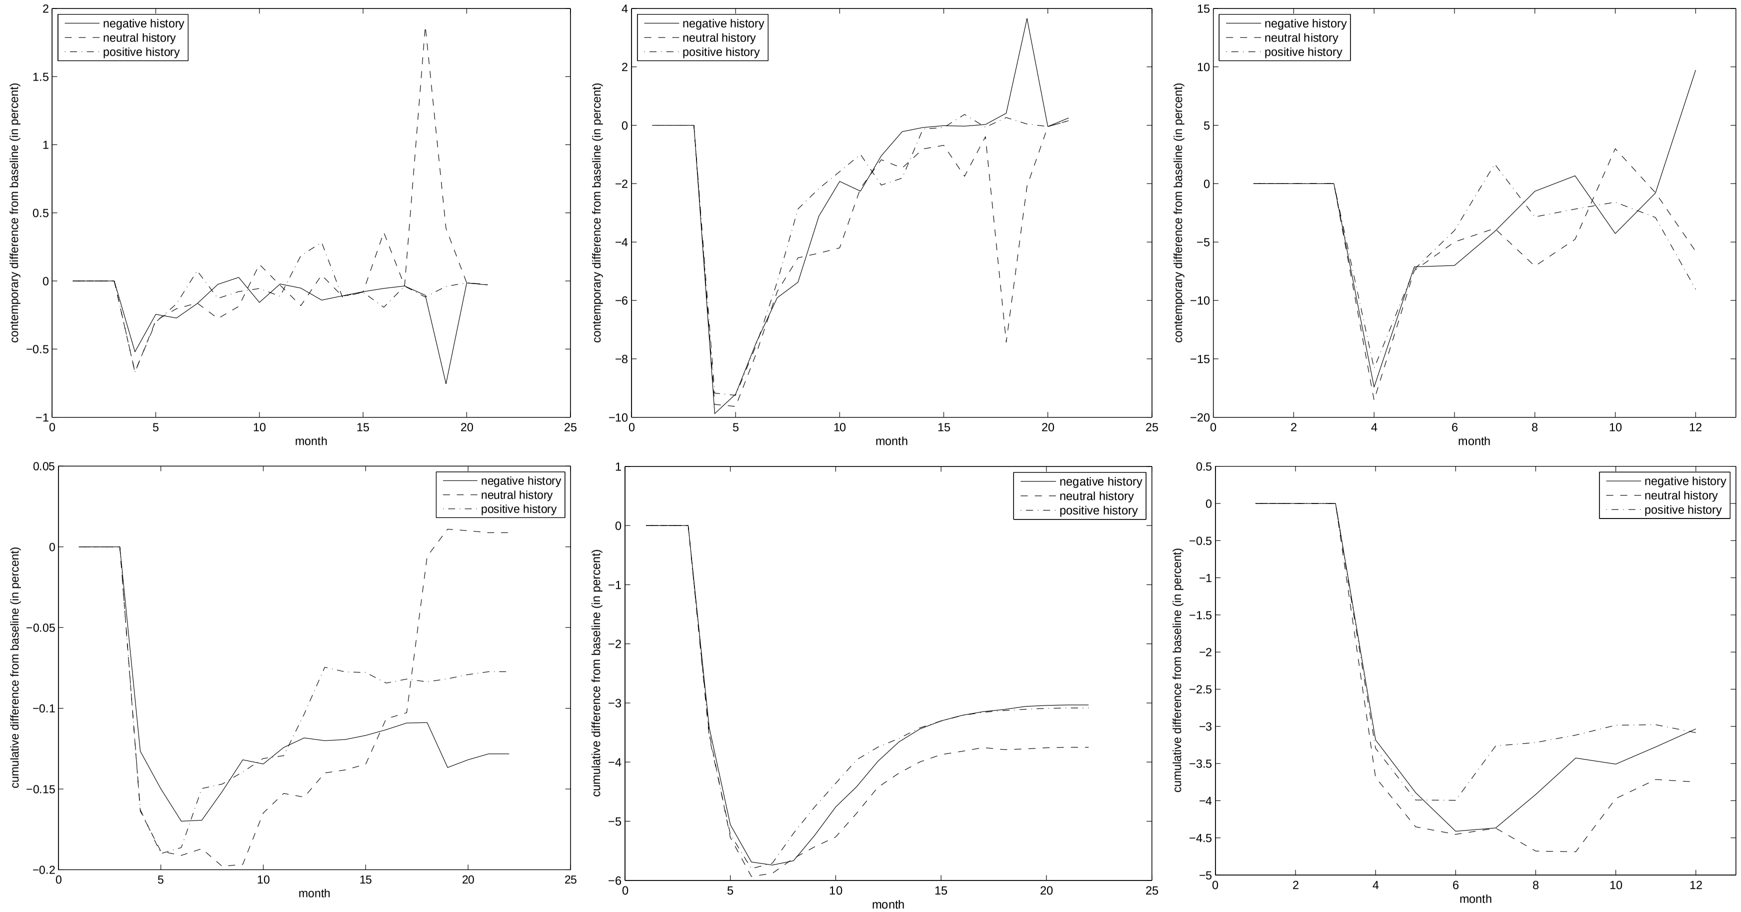

Price Response

Sales Response

Production Response

Figure 3: Contemporary (Top Panel) and Cumulative (Bottom Panel) Responses of Prices, Sales, and Production to a One Standard-Deviation Negative Innovation to z in the Convex Model at Week 14 (Month 4).

Notes: The responses have been time-aggregated to the monthly frequency.
In the top panel, each line plots the contemporary percent difference between the time path of the variable with $\Lambda = -1$ and the time path with $\Lambda = 0$; that is, $100 \times \left(\log(x_t^{\Lambda=-1}) - \log(x_t^{\Lambda=0})\right)$ for $x = p, s,$ or q.
In the bottom panel, each line plots the cumulative percent difference between the time path of the variable with $\Lambda = -1$ and the time path with $\Lambda = 0$; that is, $100 \times \left(\log\left(\sum_{j=0}^{t} x_j^{\Lambda=-1}\right) - \log\left(\sum_{j=0}^{t} x_j^{\Lambda=0}\right)\right)$ for $x = p, s,$ or q.
The solid line is the response of the variables under the negative history case (i.e. $\omega_{14} = -\sigma_\omega$; $\omega_t = -\sigma_\omega/4$ for $t = 4, 5, ..., 13$; and $\omega_t = 0$ for $t > 14$).
The dashed line is the response of the three variables under the neutral history case (i.e. $\omega_{14} = -\sigma_\omega$; $\omega_t = 0$ otherwise).
The dot-dashed line is the response of the three variables under the positive history case (i.e. $\omega_{14} = -\sigma_\omega$; $\omega_t = \sigma_\omega/4$ for $t = 4, 5, ..., 13$; and $\omega_t = 0$ for $t > 14$).

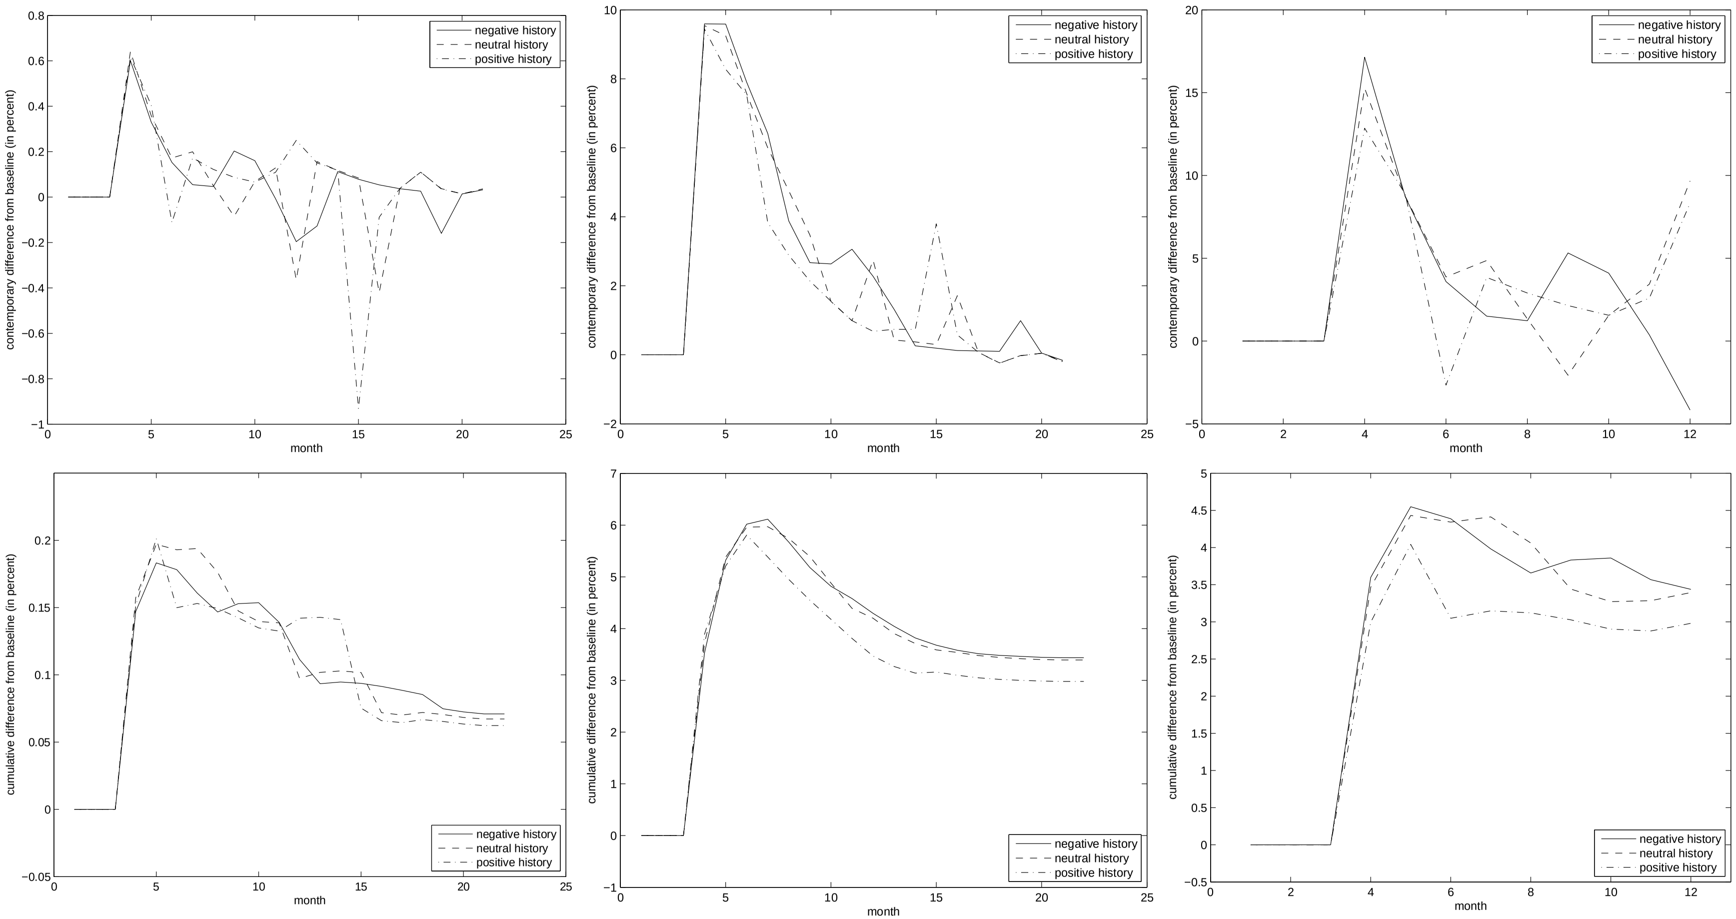

Price Response

Sales Response

Production Response

Figure 4: Contemporary (Top Panel) and Cumulative (Bottom Panel) Responses of Prices, Sales, and Production to a One Standard-Deviation Positive Innovation to z in the Convex Model at Week 14 (Month 4).

Notes: The responses have been time-aggregated to the monthly frequency.
In the top panel, each line plots the contemporary percent difference between the time path of the variable with $\Lambda = 1$ and the time path with $\Lambda = 0$; that is, $100 \times \left(\log(x_t^{\Lambda=1}) - \log(x_t^{\Lambda=0})\right)$ for $x = p, s,$ or q.
In the bottom panel, each line plots the cumulative percent difference between the time path of the variable with $\Lambda = 1$ and the time path with $\Lambda = 0$; that is, $100 \times \left(\log\left(\sum_{j=0}^{t} x_j^{\Lambda=1}\right) - \log\left(\sum_{j=0}^{t} x_j^{\Lambda=0}\right)\right)$ for $x = p, s,$ or q.
The solid line is the response of the variables under the negative history case (i.e. $\omega_{14} = \sigma_\omega$; $\omega_t = -\sigma_\omega/4$ for $t = 4, 5, ..., 13$; and $\omega_t = 0$ for $t > 14$).
The dashed line is the response of the three variables under the neutral history case (i.e. $\omega_{14} = \sigma_\omega$; $\omega_t = 0$ otherwise).
The dot-dashed line is the response of the three variables under the positive history case (i.e. $\omega_{14} = \sigma_\omega$; $\omega_t = \sigma_\omega/4$ for $t = 4, 5, ..., 13$; and $\omega_t = 0$ for $t > 14$).

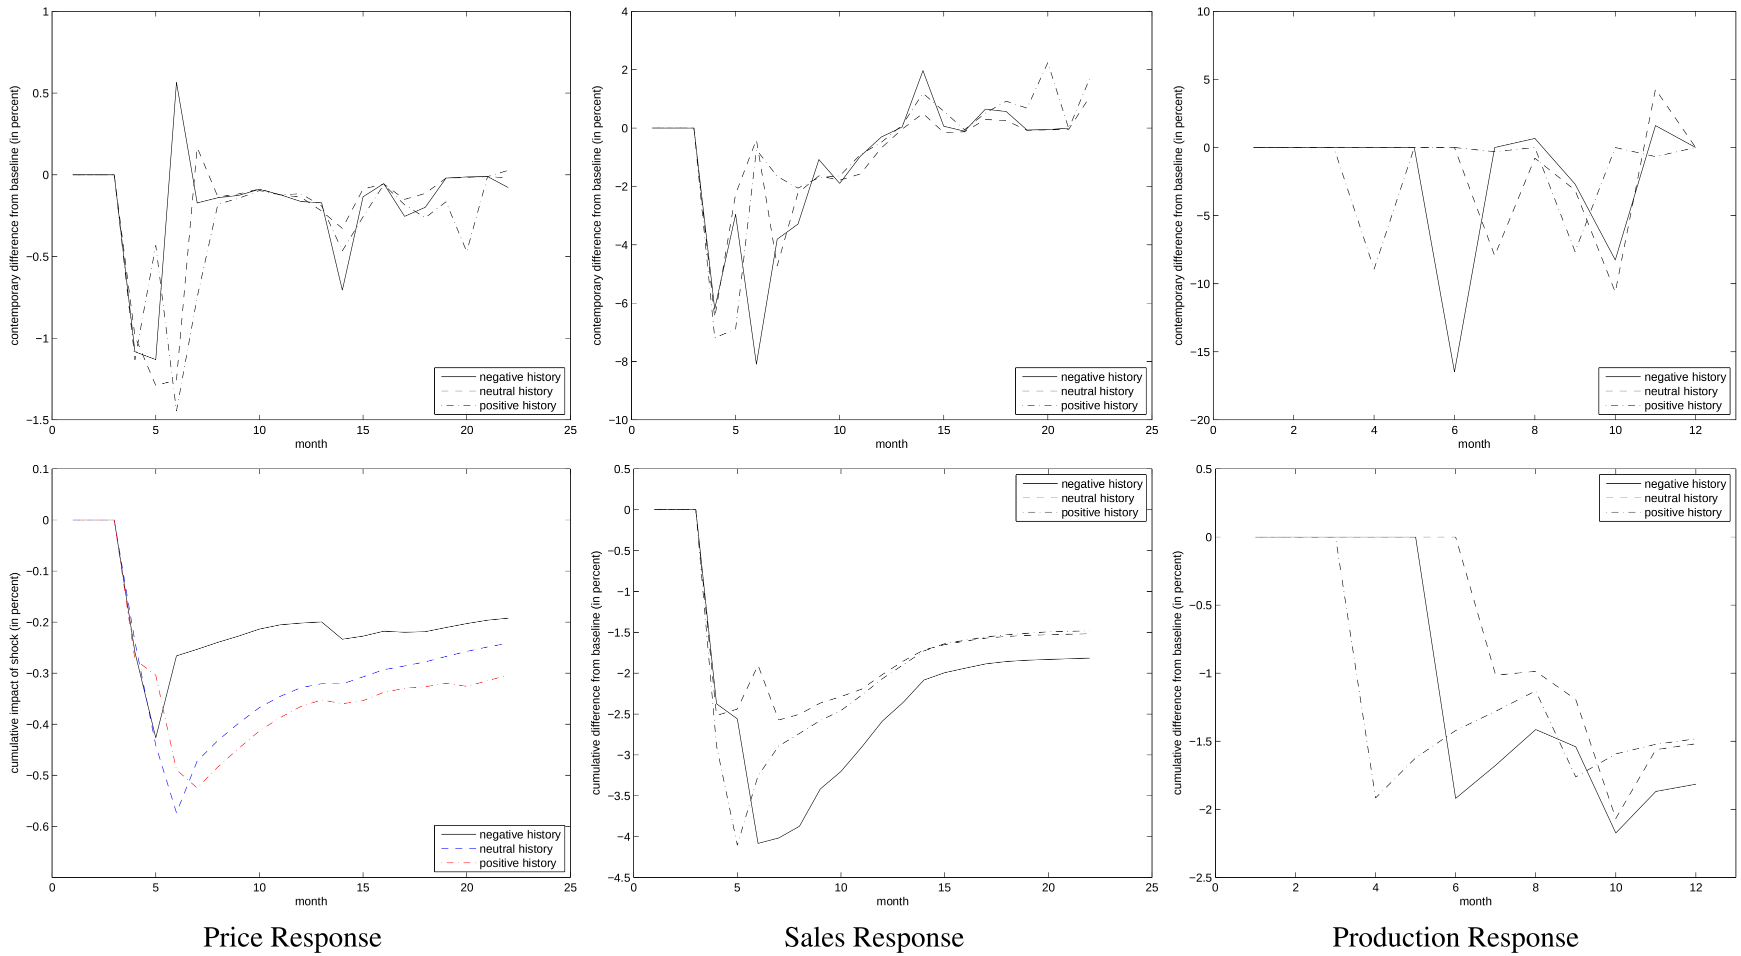

Figure 5: Contemporary (Top Panel) and Cumulative (Bottom Panel) Responses of Prices, Sales, and Production to a One Standard-Deviation Negative Innovation to z in the Non-Convex Model at Week 14 (Month 4).

Notes: The responses have been time-aggregated to the monthly frequency.
In the top panel, each line plots the contemporary percent difference between the time path of the variable with $\Lambda = -1$ and the time path with $\Lambda = 0$; that is, $100 \times \left(\log(x_t^{\Lambda=-1}) - \log(x_t^{\Lambda=0})\right)$ for $x = p, s,$ or q.
In the bottom panel, each line plots the cumulative percent difference between the time path of the variable with $\Lambda = -1$ and the time path with $\Lambda = 0$; that is, $100 \times \left(\log\left(\sum_{j=0}^{t} x_j^{\Lambda=-1}\right) - \log\left(\sum_{j=0}^{t} x_j^{\Lambda=0}\right)\right)$ for $x = p, s,$ or q.
The solid line is the response of the variables under the negative history case (i.e. $\omega_{14} = -\sigma_\omega$; $\omega_t = -\sigma_\omega/4$ for $t = 4, 5, ..., 13$; and $\omega_t = 0$ for $t > 14$).
The dashed line is the response of the three variables under the neutral history case (i.e. $\omega_{14} = -\sigma_\omega$; $\omega_t = 0$ otherwise).
The dot-dashed line is the response of the three variables under the positive history case (i.e. $\omega_{14} = -\sigma_\omega$; $\omega_t = \sigma_\omega/4$ for $t = 4, 5, ..., 13$; and $\omega_t = 0$ for $t > 14$).

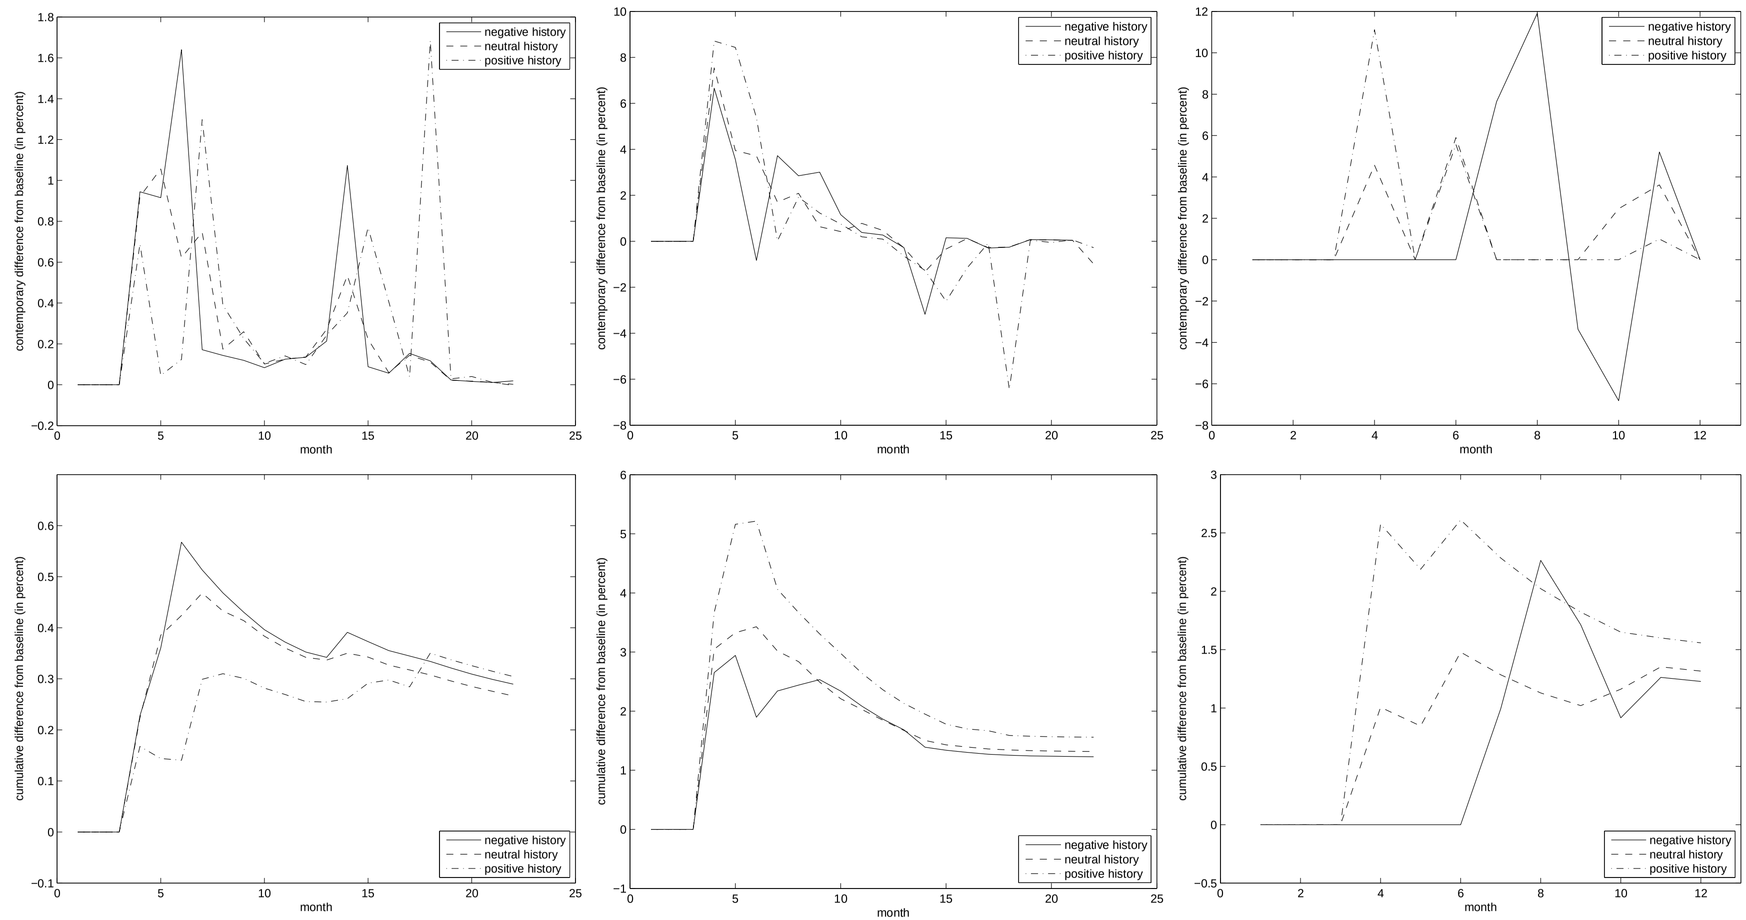

Price Response

Sales Response

Production Response

Figure 6: Contemporary (Top Panel) and Cumulative (Bottom Panel) Responses of Prices, Sales, and Production to a One Standard-Deviation Positive Innovation to z in the Non-Convex Model at Week 14 (Month 4).

Notes: The responses have been time-aggregated to the monthly frequency.
In the top panel, each line plots the contemporary percent difference between the time path of the variable with $\Lambda = 1$ and the time path with $\Lambda = 0$; that is, $100 \times \left(\log(x_t^{\Lambda=1}) - \log(x_t^{\Lambda=0})\right)$ for $x = p, s,$ or q.
In the bottom panel, each line plots the cumulative percent difference between the time path of the variable with $\Lambda = 1$ and the time path with $\Lambda = 0$; that is, $100 \times \left(\log\left(\sum_{j=0}^{t} x_j^{\Lambda=1}\right) - \log\left(\sum_{j=0}^{t} x_j^{\Lambda=0}\right)\right)$ for $x = p, s,$ or q.
The solid line is the response of the variables under the negative history case (i.e. $\omega_{14} = \sigma_\omega$; $\omega_t = -\sigma_\omega/4$ for $t = 4, 5, \ldots, 13$; and $\omega_t = 0$ for $t > 14$).
The dashed line is the response of the three variables under the neutral history case (i.e. $\omega_{14} = \sigma_\omega$; $\omega_t = 0$ otherwise).
The dot-dashed line is the response of the three variables under the positive history case (i.e. $\omega_{14} = \sigma_\omega$; $\omega_t = \sigma_\omega/4$ for $t = 4, 5, \ldots, 13$; and $\omega_t = 0$ for $t > 14$).

is not due to ‘sticky prices,’ because there are no price rigidities in the model. Instead, almost the entire shock is absorbed through quantities rather than prices. Looking at the lower panels in figures 3 and 4, we see that a one-percent shock in the fourth month has a 3 to 4 percent impact on total sales and output over the entire product cycle.

For the non-convex cost specification, the response to a shock in z is quite different. Examination of figures 5 and 6 show that output may not respond to the shock for several months. In both the positive and negative shock cases, much of the output response occurs in months 6 to 12 after the sales and price responses have largely died out. This propagation occurs even though there are no adjustment costs in the model. Because of the non-convexities 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 3,150 vehicles per week. Below the MES the firm can only convexify its cost function over time via temporary shutdowns; therefore the non-convexities can induce a lag between the price and production responses. Further because higher inventories stimulate sales, the firm prefers to postpone reductions in production until later in the product cycle.

5 A Demand Shock That Was and a Demand Shock That Was Not

Our model and data set can be used to understand 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 7 through 10 show the percent 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,

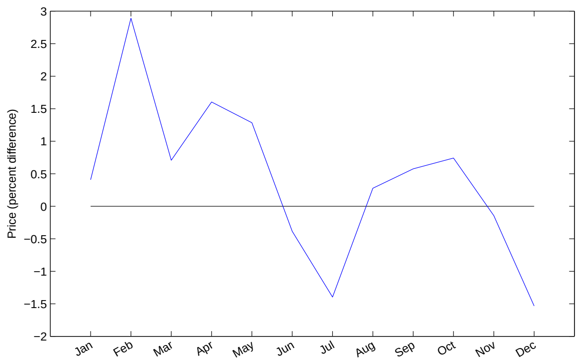

Figure 7: Prices

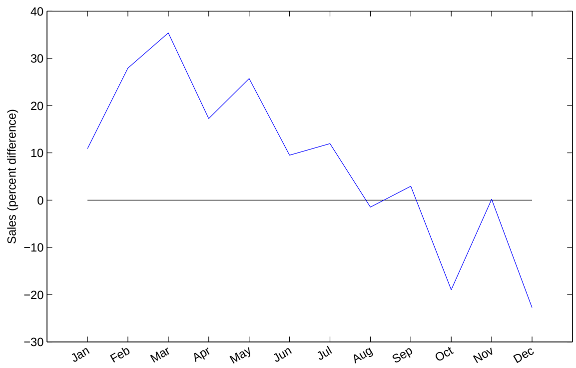

Figure 8: Sales

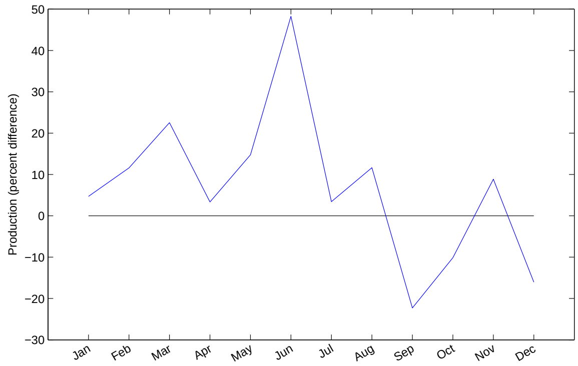

Figure 9: Production

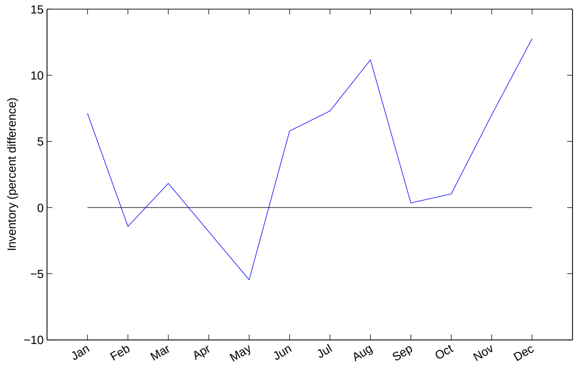

Figure 10: Inventories

The Monthly Path of Prices, Sales, Production, and Inventories for the Ford Explorer During the Year 2000

Each graph displays the percent difference between the monthly series during 2000 and the average monthly series for all other years in our sample.

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 7 and 8) we see that the relative magnitudes of the responses (over 10 to 1 in sales to prices) are consistent with the responses reported for either specification of the model.

In line with the non-convex cost specification, 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 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 non-convex cost model's predictions (see figure 5). 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 only modestly lowering the 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 prices 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."[29] 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 Zetesche, 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 "...

[29] Krebs M.L. "Driving Through and Altered Landscape," *New York Times*, September 23, 2001, section 12, page1.

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."[30] 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 it 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.

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 percent difference between prices, sales, production and inventories levels for every month from June of 2001 through February of 2002 and the average price, sales, production and inventory level for all remaining months in our sample. The first, and a surprising, fact illustrated in figures 11 through 14 is the increase of 6 percent in relative prices from September to November. This is not an artifact of the normalization; prices rose 3.7 percent un-normalized. Perhaps even more surprising, this price increase corresponds with a massive sales increase of over 40 percent. 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, weeklong 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 that 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 15, we plot the time paths of the average value per vehicle of financing incentives and cash rebates. Automakers increased financial incentives modestly in

[30] Both 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.

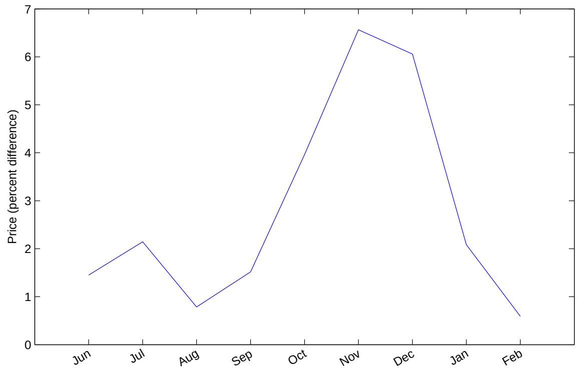

Figure 11: Industry Price Response

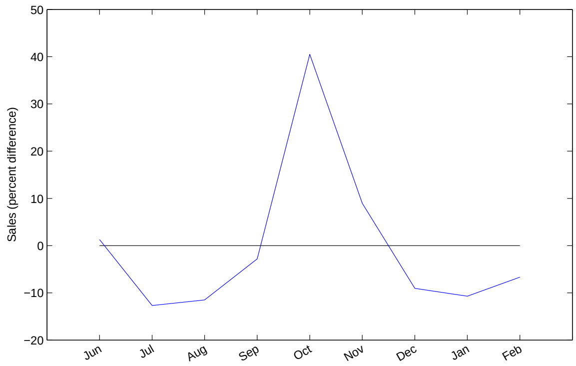

Figure 12: Industry Sales Response

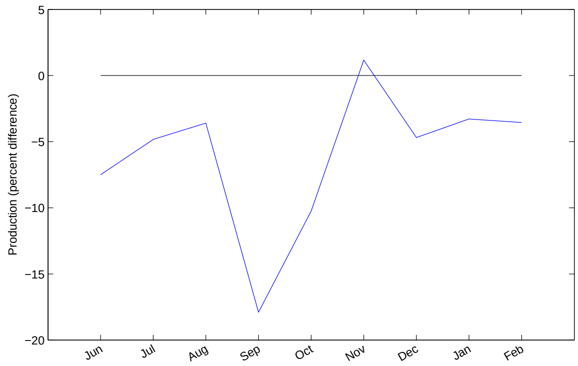

Figure 13: Industry Production Response

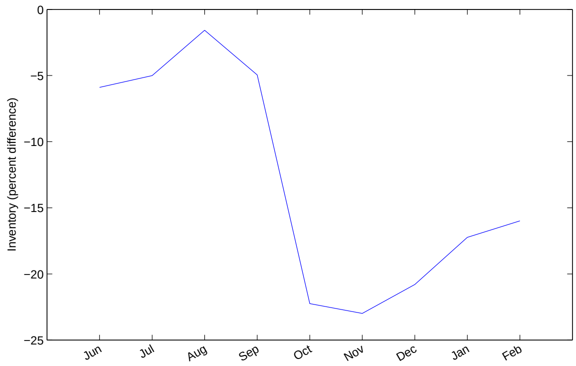

Figure 14: Industry Inventory Response.

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

Each graph displays the percent difference between the monthly series during 2001 and the average monthly series for all other years in our sample.

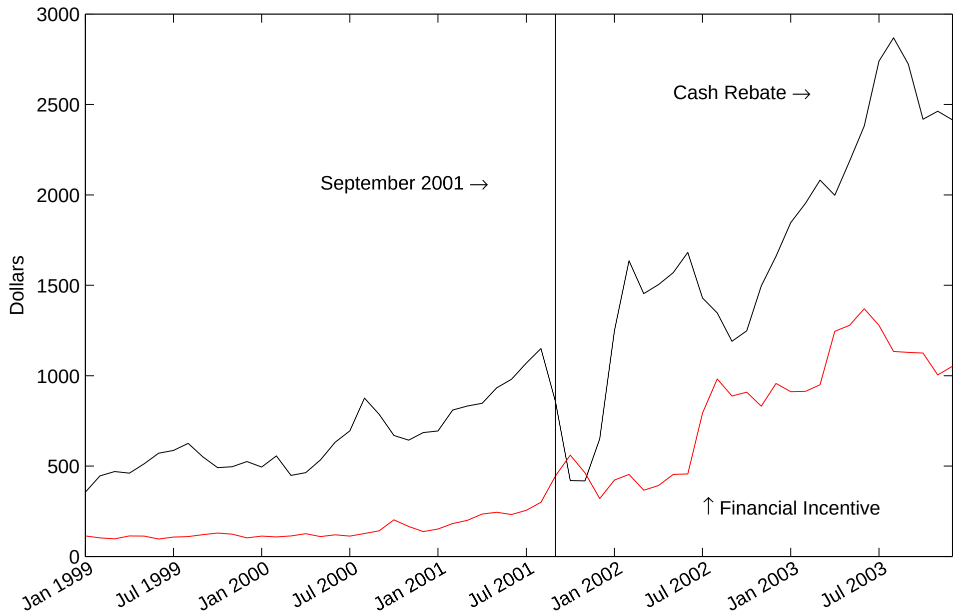

Figure 15: Value of Cash Rebates and Financing Incentives

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.[31] 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. Further, neither of our specifications can reconcile rising prices with simultaneous production cuts in response to a demand shock.

[31] For 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.

6 Conclusion

In this paper, we present a model in which an automaker can use all three primary margins of adjustment when responding to a short-run demand shock. This is important for motor vehicle production, because we find that automakers steadily reduce prices throughout the model year and frequently adjust labor inputs and inventory stocks. In analyzing an automaker's response to temporary demand shocks, we show that non-convexities in the firm's cost structure induce delayed production responses. Thus an observer with a static supply-and-demand model in mind could be misled to believe the supply curve is vertical. Contrary to industry wisdom[32], we find that prices only respond modestly despite the absence of any price rigidities. Unexpectedly, these shocks are almost entirely absorbed by changes in sales and production.

Our model suggests that the use of inventories along with the non-convexities present in the automaker's cost function causes production adjustments to be propagated throughout the model year even though prices and sales move immediately. This propagation occurs even though there are no adjustment costs to varying the workweek of capital over time. These non-convexities 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 non-convexities.

[32] See the quote in the second paragraph of this paper.

7 Appendix: Identification of the Structural Parameters

The non-convex model is too complex for us to provide analytical results on identification. Instead we perform two exercises. First, we plot concentrated slices of the criterion function parameter-by-parameter. Compared to the standard errors reported in table 5, these graphs provide a more detailed representation of the slope and shape of the criterion function. Second, we report the effect of an increase in each structural parameter on select moments in the auxiliary model. This exercise illustrates how each parameter is identified by tracing how increases in each structural parameter are detected by the auxiliary model through changes in the simulated sales, price, and production time series. We conclude by discussing why we are confident that combinations of the structural parameters are not un-identified.

Consider first Figure A1. In this figure we plot the criterion function for different values of each of the twelve parameters holding the remaining eleven fixed at their estimated values. Perhaps the most striking feature of the plots is the jaggedness of the criterion function along most dimensions. The source of this jaggedness appears to be largely due to the linear interpolation of the value function. For the firm, the marginal cost of selling an additional vehicle is the derivative of the value function with respect to inventories. Linear interpolation creates discrete jumps in this derivative.[33] These discrete jumps translate into cliffs in the criterion function.

For the four parameters governing the shock processes ρ_z, σ_ω, ρ_g and σ_ε and two of the production parameters, γ_1 and LS, the concentrated criterion function is clearly U-shaped and the minimum is easily recognizable. This suggests that individually these parameters are well identified. For the remaining six parameters, the plots of the criterion are dominated by sharp spikes and dips. However, despite this local jaggedness, the more "global" curvature of the criterion function is apparent as one move further away from the minima. Take for example the plot of the criterion for different values of ot_{prem}. While there are many values between 0.1 and 0.25 that yield similar minima of the criterion function, values outside this region fit the data less well. Of particular interest, values around the statutory rate of 0.5 generate values of the criterion function above 370, clearly above the minima of 308.5 found at $ot_{prem} = 0.244$.

Given the many local minima, it is reasonable to wonder if the results reported in tables 5 and 6 are for a local rather than the global minima of our criterion. While we can not prove that no other minima exists we found it reassuring that when we initialized the estimation procedure with different starting values, our derivative-based minimization routine (specifically MATLAB's `fmincon.m` routine) consistently converged to parameter values in the same region of the parameter space. We then performed grid searches, such as the one displayed in figure A1, to search for other nearby minima.

Next we examine the effect of increasing each structural parameter one-by-one on the individual coefficients in the auxiliary model. This exercise, reported in table A1, illustrates how each of the parameters have different effects and are therefore identified. To keep the presentation concise, we report the effect on only 8 of the 19 auxiliary moments: the five sales regression coefficients and the three constants.

First consider the five production-cost parameters, γ_1, LS, w_1, υ, and ot_{prem}. As one would expect, increases in costs of producing (i.e. γ_1, w_1 and ot_{prem}) lower sales and production and increase prices. Increases in the line speed and unemployment premium increase sales and production and lower prices. Each

[33] We experimented with shaped-preserving splines but found that using them increased the computation time considerably and that they were less robust than linear interpolation.

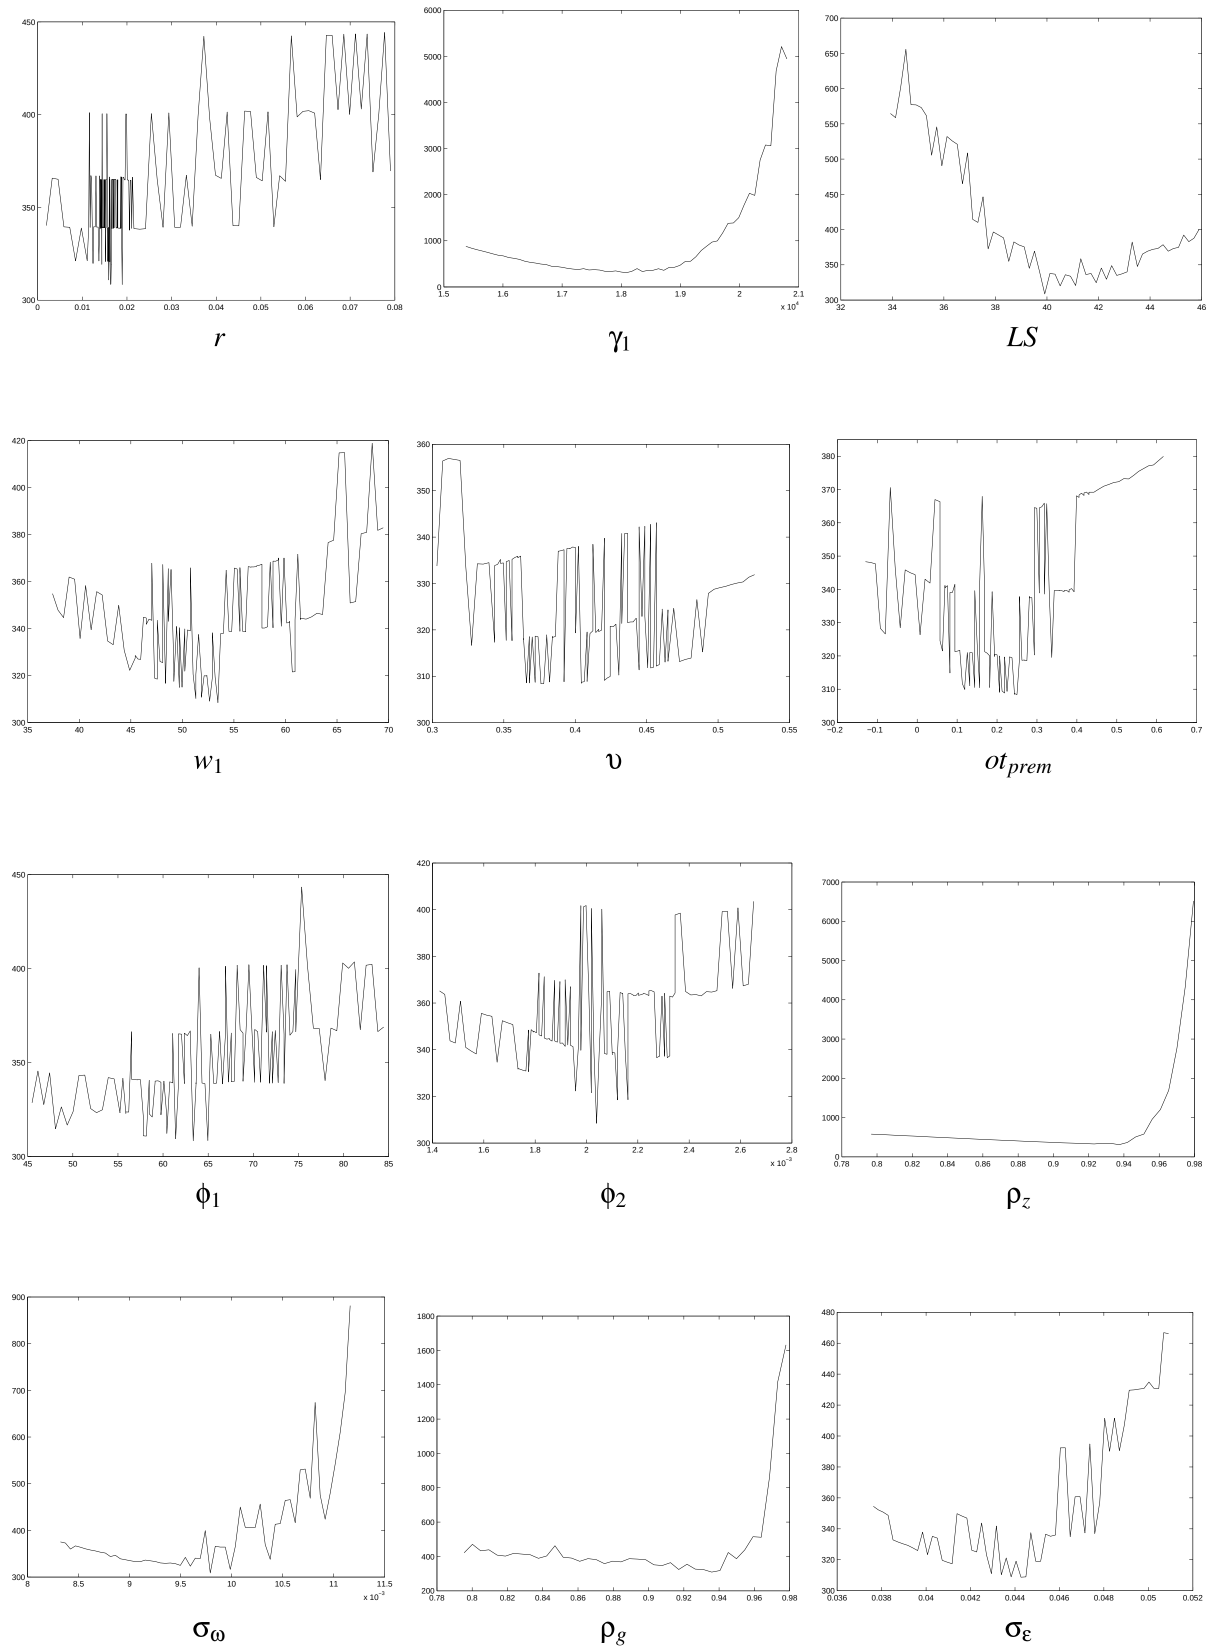

Figure A1: Concentrated Slices of the Criterion Function Parameter-By-Parameter

	Sales Equation					Constants		
parameter	lag p	lag s	inv	trend	var(res)	Sales	Price	Prod
r	–	–	+	–	≈ 0	+	–	≈ 0
γ_1	++	+	– –	++	–	– –	++	– –
LS	–	+	–	– –	+	++	– –	++
w_1	+	+	– –	–	–	–	+	–
υ	–	+	–	–	–	+	–	+
ot_{prem}	+	+	–	+	–	–	+	–
ϕ_1	+	–	+	–	–	–	–	–
ϕ_2	+	–	++	–	–	–	≈ 0	–
ρ_z	+	++	– –	– –	+	–	+	–
σ_ω	+	+	–	+	+	–	+	≈ 0
ρ_g	– –	+	++	– –	+	–	– –	–
σ_ε	–	–	++	+	+	–	–	–

Table A1: Effect of An Increase In Each Structural Parameter on Select Moments in the Auxiliary Model

A "+" denotes that increasing a parameter results in an increase in the moment. A "–" denotes a decrease. A "++" or a "– –" denotes a large increase or decrease.

of these five parameters also have differential effects on the sales regression parameters. These effects, along with the additional differential effects in the price and production regressions, further contribute to the identification of these parameters.

Since vehicles must be held in inventory before they can be sold, an increase in either inventory holding cost parameter, ϕ_1 or ϕ_2, reduces the quantity produced and sold. Also since these parameters directly effect the marginal value of an additional unit of inventory, an increase in their values reduces prices and increases the sensitivity of sales to current inventories. These two parameters can be separately identified by the differential response to average price and the coefficient on inventories in the sales regression. The interest rate, r, also represents a cost of holding inventories. Like increases in ϕ_1 and ϕ_2, an increase in r decreases the shadow value of inventories thus lowering the average price and increasing the sensitivity of sales to inventories. But unlike (ϕ_1, ϕ_2), increases in r have almost no effect on production; instead it moves sales forward in the product cycle. With more vehicles being sold in the first 17 months of the product cycle, the constant term on the sales regression rises. Hence, r can be identified separately from the two holding cost parameters.

Now consider the four shock-process parameters, ρ_z, σ_ω, ρ_g and σ_ε. Increases in the persistence and variance of the demand-side shocks, (ρ_z, σ_ω), raise the importance of shifts in the demand curve on the simulated price and sales data. Hence the correlations of sales with lagged sales and prices increase and the correlation of sales with current inventories decreases. In contrast, increases in ρ_g and σ_ε raise the importance of shifts in the marginal cost curve, increasing the correlation between sales and current inventories and decreasing the correlation between sales and lagged prices. However increases in the persistence of the supply-side shock increase the serial correlation of sales. These differential responses (which are also picked up in the price regression, though not reported in table A1) allow us to identify the supply and demand disturbances.

Finally neither the plots displayed in figure A1 nor the results shown in table A1 show that linear combinations of the parameters are unidentified. The best way to address this concern would be to run a monte carlo experiment using the structural model to repeatedly create synthetic datasets, and then re-estimating the structural model employing these synthetic datasets to determine if the original parameters are recovered. Unfortunately, the non-convex cost model as described in the text takes several days to estimate making such an exercise infeasible. Nevertheless, over the course of conducting this research, the non-convex model was estimated many dozen, perhaps hundreds, of times as we learned more about the model and experimented with different functional forms, different solution and approximation methods, and different specifications of the auxiliary model. At no time did we find that two or more parameters would move together into unexpected regions of the parameter space. Furthermore, when we estimated the model using different starting values, our estimation method repeatedly returned to the same region of the parameter space. Had we found evidence of under-identification, we would have either fixed a parameter or changed the specification of our auxiliary model. Consequently we are confident there are not unidentified combinations of the parameters.

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] Bils M., Kahn J., 2000. What inventory behavior tells us about business cycles. American Economic Review 90, 458-481.

[7] Blanchard O., 1983. The production and inventory behavior of the American automobile Industry. Journal of Political Economy 91, 365-400.

[8] Blinder A., Maccini L.J., 1991a. The resurgence of inventory research: what have we learned? Journal of Economic Surveys 5, 291-328.

[9] Blinder A., Maccini L.J., 1991b. Taking stock: a critical assessment of recent research on inventories. Journal of Economic Perspectives 5, 73-96.

[10] Bresnahan T., Ramey V., 1994. Output fluctuations at the plant level. Quarterly Journal of Economics 109, 593-624.

[11] Chan H., Hall G., Rust J., 2005. Price discrimination in the steel market. manuscript, University of Maryland.

[12] Copeland A., Dunn W., Hall G., 2005. Prices, production, and inventories over the automotive model year. NBER working paper 11257.

[13] 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.

[14] 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.

[15] Federgruen A., Heching A.,1999. Combined pricing and inventory control under uncertainty. Operations Research 47, 454-475.

[16] Galeotti M., Maccini L.J., Schiantarelli F., 2005. Inventories, employment, and hours. Journal of Monetary Economics 52, 575-600.

[17] 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.

[18] Hall R., Hitch C.J., 1939. Price theory and business behavior. Oxford Economic Papers 2, 12-45.

[19] Hay G., 1970. Production, price and inventory theory. American Economic Review 60, 531-545.

[20] Haltiwanger J.C., Maccini L.J., 1988. A model of inventory and layoff behavior under uncertainty. Economic Journal 98, 731-745.

[21] 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.

[22] Hamermesh D.S., 1989. Labor demand and the structure of adjustment costs. American Economic Review 79, 674-689.

[23] Kahn J., 1987. Inventories and the volatility of production. American Economic Review 77(4), 667-79.

[24] Kahn J., 1992. Why is production more volatile than sales? Theory and Evidence on the Stockout-Avoidance Motive for Inventory-Holding. Quarterly Journal of Economics 109(3), 565-92.

[25] 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.

[26] 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.

[27] Newey W., West K., 1987. A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance-matrix. Econometrica 55, 703-708.

[28] Nordhaus W., Godley W., 1972 Pricing over the trade cycle. Economic Journal 82, 853-882.

[29] Ramey V., Vine D., 2006. Declining volatility in the U.S. automobile industry. American Economic Review 96, 1876-1889.

[30] Ramey V., West K.D., 1999. Inventories. in Handbook of Macroeconomics, J.B. Taylor, M. Woodford (eds.), Amsterdam: North-Holland.

[31] Reagan P., 1982. Inventory and price behavior. Review of Economic Studies 49, 137-142.

[32] 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.

[33] Smith A.A., 1993. Estimating nonlinear time-series models using simulated vector autoregressions. Journal of Applied Econometrics 1993, S63-S84.

[34] Topel R.H., 1982. Inventories, layoffs, and the short-run demand for labor. American Economic Review 72, 769-787.

[35] Trejo S., 2003. Does the statutory overtime premium discourage long workweeks? Industrial and Labor Relations Review 56, 530-551.

[36] Whiten T., 1955. Inventory control and price theory. Management Science 2, 61-80.

[37] Womack J.P., Jones D.T., Roos D., 1990. The Machine That Changed The World. Rawson Associates, New York, USA.

[38] 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.