Declining Volatility in the U.S. Automobile Industry

By VALERIE A. RAMEY AND DANIEL J. VINE*

The automobile industry is a highly volatile sector of the U.S. economy. Motor vehicle production has accounted for almost 25 percent of the variance of aggregate GDP growth over the past 40 years even though gross motor vehicle output represented, on average, less than 5 percent of the *level* of aggregate GDP.[1] In the mid-1980s, however, the variance of automobile production declined drastically, falling by more than 70 percent relative to its past level. While the variance of auto sales also receded, this decline was smaller than the decline in output volatility. Moreover, the covariance of inventory investment with sales became negative in the 1980s, suggesting that inventories had begun to more actively insulate production from sales shocks. At the same time, assembly plants began to adjust average hours per worker much more often than in preceding decades, when most output volatility stemmed from changes to the number of workers attached to each plant. Interestingly, many of these changes were observed outside the motor vehicle industry as well.[2]

This paper documents these developments in the U.S. auto industry and shows how the changes observed in sales, inventories, and production in the 1980s could have stemmed from one underlying factor—a decline in the persistence of motor vehicle sales. We analyze industry-level data and micro data on production schedules from 103 assembly plants in the United States and Canada to document these developments. Using the original version of the linear-quadratic inventory model formulated by Charles C. Holt, Franco Modigliani, John F. Muth, and Herbert A. Simon (1960), we then show that a decline in the persistence of sales leads to all of the changes noted above, even in the absence of technological change.

I. Structural Change in the U.S. Automobile Industry: The Facts

A. *Production, Sales, and Inventory Variances*

The levels of monthly U.S. car and truck production, measured in physical units at seasonally adjusted annual rates, are shown in the panels of Figure 1 from January 1967 through December 2004. Sales of domestic vehicles for each of these market segments are also shown, where sales of domestic vehicles include vehicles assembled in the United States, Canada, and Mexico.[3]

As seen in Figure 1, the historical trends for production and sales in the car segment have behaved differently than the trends in the truck segment. Production and sales of cars have declined over time, while production and sales of trucks have steadily increased. As there are obvious differences in the conditional means of these two market segments, we treat them separately in most of the analysis below.

* Ramey: Department of Economics, University of California, San Diego, 9500 Gilman Drive, La Jolla, CA 92093-0508 (e-mail: vramey@ucsd.edu); Vine: Division of Research and Statistics, Board of Governors of the Federal Reserve System, 20th and C Streets, NW, Washington, DC 20051 (e-mail: Daniel.J.Vine@frb.gov). We are indebted to Edward Cho and Jillian Medeiros for outstanding research assistance and to George Hall for his Chrysler data. We have benefited from helpful comments from two anonymous referees, Nir Jaimovich, and Garey Ramey. Valerie Ramey gratefully acknowledges support from National Science Foundation grant 0213089. The analysis and conclusions set forth are those of the authors and do not indicate concurrence by the Board of Governors or the staff of the Federal Reserve System.

[1] The Data Appendix gives details of all calculations and estimates.

[2] Chang-Jin Kim and Charles R. Nelson (1999) and Margaret M. McConnell and Gabriel Perez-Quiros (2000) document a 50-percent decline in the variance of GDP growth beginning in 1984. John E. Golob (2000) shows that the covariance of inventory investment and sales for most industries switched from being positive before 1984 to being negative after 1984. James A. Kahn et al. (2002) showed that within durable goods manufacturing the variance of production fell much more than the variance of sales. Ron Hetrick (2000) documents an increase in the use of overtime hours for many industries during the 1990s expansion.

[3] The truck market segment includes vans and SUVs.

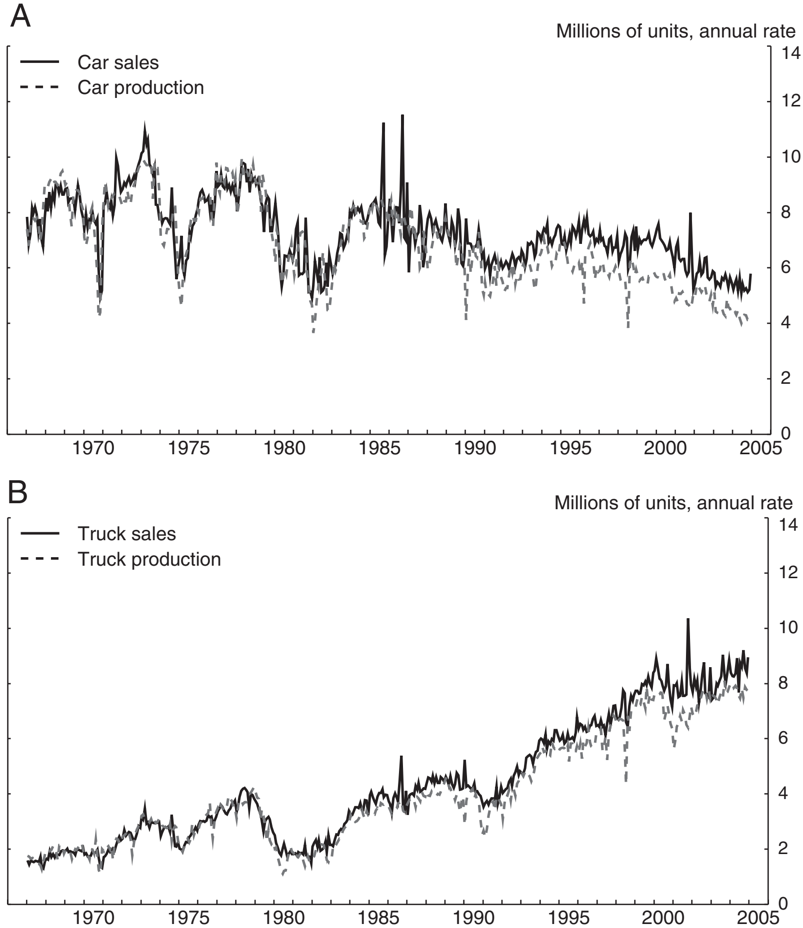

FIGURE 1. U.S. AUTOMOBILE PRODUCTION AND DOMESTIC SALES
(January 1967 through December 2004)

Notes: Production and sales are measured in millions of units at a seasonally adjusted annual rate. Domestic sales include U.S. sales of vehicles built in the United States, Canada, and Mexico.

The variances of both car and truck output (relative to their respective trends) dropped sharply in the mid-1980s. According to structural break tests on the variance of detrended car production, there was a statistically significant break between February and March in 1984. For truck production, the break occurred between January and February in 1983.[4] Thus, the structural break in the variance of automobile production occurred at essentially the same time as the structural break in GDP growth volatility, which many studies place in 1984.

In order to quantify the changes in industry volatility that took place around 1984, Table 1 reports the variances of key variables in the auto industry in two sample periods—1967 through

[4] One important difference between the time-series properties of physical unit data and chained-dollar data used in other studies is that stationarity tests on the logarithm of physical unit variables reject a unit root in favor of a deterministic trend. To search for the structural break in the variance, we used data detrended with an Hodrick-Prescott (HP) filter rather than trend breaks so as not to bias the results for a particular period. In particular, we used seasonally adjusted output divided by the exponential of the HP filter trend applied to the log of output. The *p*-values were essentially zero.

TABLE 1—DECOMPOSITION OF MOTOR VEHICLE OUTPUT VOLATILITY

	Cars		Trucks	
	1967:2–1983:12	1984:2–2004:12	1967:2–1983:12	1984:2–2004:12
$Var(Y)$	3.13	0.92	9.16	1.15
$Var(S)$	2.45	1.01	6.89	0.94
$Var(\Delta I)$	0.20	0.15	0.29	0.09
$Cov(S, \Delta I)$	−0.03	−0.20	0.15	−0.08
$\frac{Var(Y)}{Var(S)}$	1.28	0.91	1.33	1.22

Notes: Monthly data are measured in seasonally adjusted physical units. Cars and trucks are measured separately. Y = production, S = sales, and ΔI = change in inventories. Production and sales were normalized by the exponential of a fitted linear trend to the log of the variable, estimated separately over each period. Inventory investment was normalized by the fitted trend in the log level of inventories. The variances and covariances in the table are 100 times the actual ones. (See Data Appendix for data sources and details.)

ADDENDUM: COMPARISON OF PRODUCTION AND SALES INCLUDING IMPORTS AND EXPORTS

	Cars	
	1967:2–1983:12	1984:2–2004:12
Var(U.S. Production + N. American Imports)	2.94	0.94
Var(U.S. Sales + Exports)	2.40	0.96
$Var(\Delta I)$	0.20	0.15
$Cov(S, \Delta I)$	−0.01	−0.18
$\frac{Var(Y)}{Var(S)}$	1.23	0.97

1983, and 1984 through 2004. The variables of interest are derived from the standard inventory identity: $Y_t = S_t + \Delta I_t$, where Y is production, S is sales, and ΔI is the change in inventories. For stationary variables, the following relationship exists between the variance of production and the variance of sales:

$$(1) \qquad Var(Y_t) = Var(S_t) + Var(\Delta I_t) + 2\ Cov(S_t, \Delta I_t).$$

The numbers in Table 1 report each term from this variance decomposition for cars and trucks separately, and, because production and sales have trends within each sample, all quantities are first normalized by their respective trends. The variance of production for cars fell by 70 percent after 1984; for trucks, it fell by 87 percent. Moreover, the variance of production for both cars and trucks fell by a larger percentage than did the variance of sales, and the covariance of inventory investment with final sales either switched from being positive to being negative or became more negative.

The variances and covariances do not form a perfect identity, in part, because domestic sales and inventories include imports from Canada and Mexico, while domestic production includes only vehicles produced in the United States. A small portion of U.S. production is also exported.[5] Changes in North American motor vehicle trade behavior, however, do not appear to have explained the structural change in industry volatility. The addendum to Table 1 shows the trade-augmented variance decomposition for cars, where U.S. production is augmented with North American imports and U.S. sales are augmented with exports. The changes in the variances after 1984 are very similar to those shown in Table 1. (Trade data for trucks are not readily available.)

[5] The variance identity also does not hold because the variables are seasonally adjusted and detrended separately.

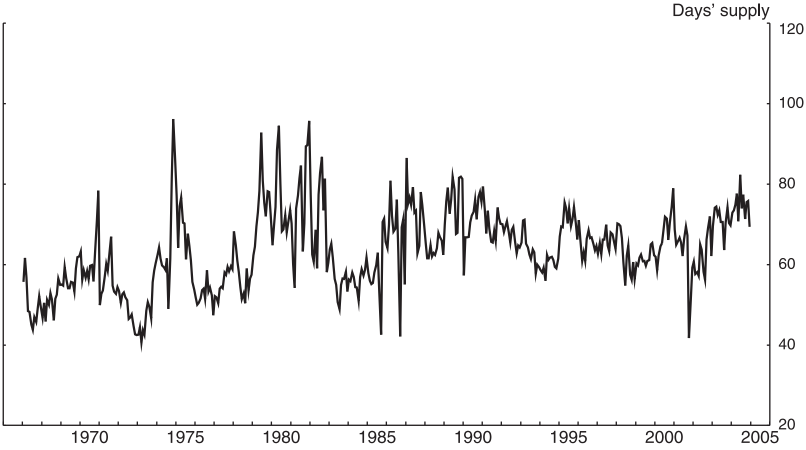

FIGURE 2. INVENTORY-TO-SALES RATIO FOR DOMESTIC CARS AND LIGHT TRUCKS
(*January 1967 through December 2004*)

Notes: Domestic unit inventories and sales include vehicles built in the United States, Canada, and Mexico. Quantities are seasonally adjusted. Seasonally adjusted days' supply is based on 307 selling days per year. Days' supply prior to January 1972 excludes light trucks.

Golob (2000) and Kahn et al. (2002) uncovered a similar decline in the volatility of aggregate chain-weighted durable goods output. Some researchers have linked this decrease in volatility to a decline in the inventory-sales ratio. As shown in Figure 2, however, the inventory-sales ratio (or "days-supply") for light vehicles shows no evidence of change after 1984.[6] Thus, the changes we have documented in the auto industry occurred for some other reason.[7]

B. *Intensive and Extensive Labor Margins*

High volatility of output in the motor vehicle industry is often associated with the ways automakers adjust work schedules and the rates of production at assembly plants.[8] Adjustments to the workweek are sometimes temporary, such as when a plant schedules overtime hours or closes for a week (called an "inventory adjustment"). At other times, changes to the workweek are more permanent, such as when a second shift is added. Automakers adjust the rate of production at a plant by raising or lowering the line speed, which has historically required a change in the number of workers on each shift.

While fluctuations in total worker hours are roughly proportional to the number of vehicles assembled, changes in employment and average hours depend on the margins of adjustment that are being used. Overtime hours and inventory shutdowns are intensive adjustments to hours worked, while changes to the number of shifts and to the line speed are largely extensive adjustments to the number of workers. Although inventory adjustments entail temporary layoffs, the spell typically lasts only a week or two and involves negligible adjustment costs. Inventory adjustments essentially allow the plant to reduce average hours per worker, measured on a monthly basis.

[6] Inventory data for light trucks are not available before 1972. The numbers prior to 1972 are for cars only. When inventory data for both market segments exist in the 1970s, the inventory-sales ratios for cars and light trucks were very similar.

[7] This does not imply that improvements in inventory management had no impact on the auto industry. In fact, the auto industry pioneered *just-in-time* inventories in the 1980s, and the ratio of materials and work-in-progress inventories to shipments for the entire automobile industry (SIC code 371) did fall after 1984.

[8] Plant entry and exit historically are responsible for the secular trend in output. (See Timothy F. Bresnahan and Ramey, 1994.)

TABLE 2—IMPORTANCE OF INTENSIVE AND EXTENSIVE MARGINS OF ADJUSTMENT FOR THE VARIANCE OF MONTHLY MOTOR VEHICLE OUTPUT
(*Percent of average plant-level variance attributed to use of each margin*)

Margin of adjustment	1972–1983	1990–2001
Changes in extensive margins	**64.6**	**37.3**
Shifts	33.0	5.4
Line speeds	21.1	17.8
Changes in intensive margins	**39.5**	**51.0**
Temporary closures (Inventory adjustments)	32.3	33.1
Overtime hours	9.6	24.3

Notes: Plant-level variance is calculated after holidays, supply disruptions, model changeovers, and extended closures are removed. Percent impact of each margin on output variance is calculated by comparing variance of actual production with the variance of hypothetical production if each margin (in turn) were held fixed. Contributions to variance are the weighted average among all plants operating in each period. Contributions of extensive and intensive margins do not sum to 100 because of covariance terms. The same is true for individual margins within each category. See Data Appendix for data sources.

The costs associated with adjusting each production margin have interesting implications for the volatility of output in the auto industry (see Ana M. Aizcorbe, 1992; Bresnahan and Ramey, 1994; George J. Hall, 2000). To summarize several key results, changes to extensive margins entail large adjustments costs and are therefore used when shifts in demand are perceived to be persistent. Transitory shocks to sales, on the other hand, are best accommodated through intensive margins. While overtime hours and inventory adjustments do affect marginal costs, their use incurs no adjustment costs.

Using a dataset that tracks production schedules at U.S. and Canadian assembly plants operated by the Big Three automakers (General Motors, Ford, and the Chrysler portion of DaimlerChrysler), we calculate the contributions of intensive and extensive adjustments to the variance of monthly output, and the results of this exercise are recorded in Table 2. Separate measurements are made for the periods before and after 1984.[9] Contributions to variance are measured as in Bresnahan and Ramey (1994); the variance of actual output is compared to the variance of an artificial output measure that holds each margin, in turn, constant at each plant. The difference in the variances of actual and constructed output determines the impact of each margin on the variance of plant-level output after holidays, supply disruptions, and the summer shutdowns are removed. The numbers in Table 2 are weighted averages across all plants and do not sum to 100 because of nonlinearities and covariance terms.

The contribution of extensive adjustments to the variance of output declined from 65 percent in the first period, to 37 percent in the second period. A change in the way shifts were added and pared between the two periods accounted for most of this change. The contribution to variance of adding and cutting shifts fell from 33 percent of the monthly production variance in the 1970s and 1980s, to only 5 percent during the 1990s.

Adjustments to intensive margins accordingly became more important in the latter period. The contribution of overtime hours and inventory adjustments rose from about 40 percent of plant-level variance in the early period, to 51 percent in the latter period. In particular, the contribution of overtime hours to total output variance stepped up significantly, from a contribution of less than 10 percent in the early period, to more than 24 percent in the second period. The contribution from temporary plant closures increased by a smaller amount.

C. *The Persistence of Sales*

The decline in the variance of auto sales shown in Table 1, though smaller than the decline in the variance of output, could have stemmed from a reduction in the size of sales shocks or a reduction in their persistence. To see this, assume that sales are represented by a simple AR(1) process with a first-order autocorrelation of ρ and an error-term variance of σ^2. Variance is then given by the expression $\sigma^2/(1 - \rho^2)$.

In auto industry data most of the reduction in the variance of sales that occurred in the 1980s came from a weakening in the persistence of sales as opposed to a reduction in the variance of shocks to sales. This distinction is important in the production-smoothing model, because the

[9] Our data for the more recent sample start in 1990 and not in 1984 because we did not have access to *Automotive News* for 1984–1989.

variance of production relative to sales depends a great deal on ρ but is invariant to σ^2.

To see the change in persistence for auto sales, consider the univariate model shown in equation (2):

$$\tilde{S}_t = \alpha_0 + \alpha_1 \cdot t + \alpha_2 \cdot \tilde{S}_{t-1} + D_t \cdot [\beta_0 + \beta_1 \cdot t + \beta_2 \cdot \tilde{S}_{t-1}] + \varepsilon_t. \quad (2)$$

$\tilde{S}_t$ is the logarithm of seasonally adjusted sales, and ε_t is a normally distributed error term with mean zero and variance equal to $\sigma^2 + \beta_3 \cdot D_t$. D_t is an indicator variable that equals zero from January 1967 through December 1983 and one over the rest of the sample.

The model is estimated with monthly domestic auto sales (seasonally adjusted) from January 1967 to December 2004, and it allows all parameters of the sales process to change in 1984, including the coefficient on lagged sales, the constant, the slope of the trend, and the variance of the residual. We estimate this model via maximum likelihood for cars alone, light trucks alone, and for the sum of cars and light trucks, called "light vehicles."

TABLE 3—ESTIMATES OF AGGREGATE AUTOMOBILE SALES PROCESS

Coefficient	Cars	Light trucks	Light vehicles
α_0 (constant)	0.367*	−0.0045	0.279*
	(0.136)	(0.042)	(0.117)
β_0 (Δ constant)	0.628*	0.287*	0.654*
	(0.215)	(0.086)	(0.234)
α_1 (trend)	−0.0002	0.00017	−.00005
	(0.00013)	(0.00016)	(0.00011)
β_1 (Δ trend)	−0.0003*	0.0011*	0.00050*
	(0.00017)	(0.00035)	(0.00020)
α_2 (AR(1))	0.851*	0.934*	0.886*
	(0.049)	(0.030)	(0.043)
β_2 (ΔAR(1))	−0.293*	−0.243*	−0.271*
	(0.108)	(0.076)	(0.102)
σ^2 (innov. Variance)	0.0070*	0.0089*	0.0066*
	(0.00116)	(0.0012)	(0.0011)
β_3 (Δ innov. Variance)	−0.00046	−0.0040*	−0.0013
	(0.0016)	(0.0014)	(0.0014)
Log likelihood	490.3	502.3	523.2

Notes: Estimated coefficients and standard errors (in parentheses) from equation (2). Standard errors were computed using Eicker-White methods. The sample is from February 1967 through December 2004 (455 observations).

* Significantly different from zero at the 5-percent level.

The estimates are summarized in Table 3, and it appears that the process governing sales has changed significantly between the two periods. The constant and the trend differ across the two periods for all three aggregates, reflecting the mix shift in sales between cars and light trucks that was shown in Figure 1. The parameter β_3, which measures the change in the variance of the sales shocks, shows a significant decline in 1984 for light trucks, but is unchanged for cars and the light-vehicle aggregate. The first-order autocorrelation of sales, on the other hand, fell in the late period for all three aggregates, as evidenced by the negative and significant point estimates of β_2. For cars, the first-order autocorrelation fell from 0.85 to 0.56, and for trucks it fell from 0.93 to 0.69. When all light vehicles are grouped together, this estimate declined from almost 0.9 to about 0.6.

The changes in the estimates imply that sales in the post-1984 period returned to the mean much more quickly following a surprise than was the case in earlier decades. Also, most of the change in the unconditional variance of sales described in Table 1 appears to have come from a change in the propagation of sales shocks rather than a change in the variance of sales shocks.

The persistence of sales could have declined for a number of reasons. One possibility is that the automakers began responding to shocks more aggressively with their pricing policies. Another possibility is that the types of shocks hitting the industry, such as oil price shocks or monetary shocks, became less persistent. The production-scheduling model we present in the next section does not depend on the source of the change in persistence, but only on the fact that it occurred. We take this change in persistence as given, and examine the implications for the behavior of production, inventories, average hours and employment.

II. Production Scheduling Model with Inventories and Workforce

The changes in the sales process described in the preceding section have large effects on the relationship between production, inventories, and sales. The channel through which the persistence of sales affects production volatility in the auto industry was, in fact, described by Olivier J. Blanchard (1983) before the changes uncovered here had even occurred. Blanchard concluded that the structure of costs in the auto industry was such that inventories would either stabilize or destabilize production depending on the persistence of sales. Fluctuations in auto sales were very persistent during Blanchard's sample period, 1966 through 1979, so inventories destabilized production.

In this section, we analyze the original Holt et al. (1960) model of industry costs, which distinguishes between intensive and extensive labor adjustments, and show how the persistence of sales affects the decision rules for inventories and workforce in a production-smoothing model. In the next section, we show how the persistence of sales has an impact on the variances of the decision rules.

A. *The Structure of Costs*

Consider a plant that faces a stochastic sales process and must choose the size of its workforce, N_t, and the level of output, Y_t, in order to minimize the discounted present value of production, workforce-adjustment, and inventory-holding costs. The level of workforce determines the minimum efficient scale of production in each period, denoted θN_t. The cost-minimization problem is shown in the following expression:

$$(3)\quad \min_{\{N_{t+j},Y_{t+j}\}} C_t = E_t\left\{\lim_{J\to\infty}\sum_{j=0}^{J}\beta^{t+j}\cdot\frac{1}{2}[\gamma_1 Y_{t+j} + \gamma_2(Y_{t+j}-\theta N_{t+j})^2 + \gamma_3 N_{t+j} + \gamma_4(N_{t+j}-N_{t+j-1})^2 + \alpha_1(I_{t+j-1}-\alpha_2 S_{t+j})^2]\right\},$$

where $0<\beta<1$; $\gamma_i \geq 0$ for $i = 1$ through 4; $\alpha_1 > 0$; and $\alpha_2 \geq 0$. I_t is the stock of inventories at the end of period t, and S_t is sales in period t. The minimization is subject to the inventory identity, $Y_t = I_t - I_{t-1} + S_t$, and the process governing sales, $S_t = c + \rho S_{t-1} + \varepsilon_t$, where ε_t is an i.i.d. shock to sales with mean zero and variance equal to σ_ε^2.

The plant observes S_t before it chooses employment and production in period t. While the plant does take sales as given in this cost minimization problem, this does not imply that sales are exogenous to the firm. Rather, we use a standard micro result that allows one to focus on only the cost minimization part of the overall profit maximization problem.

The terms in equation (3) summarize several key features of plant-level cost associated with scheduling production and holding inventories. The second term captures the cost per worker of scheduling overtime or short weeks, which is an intensive adjustment. To see this, consider θ the product of a normal "full-time" workweek (such as 40 hours), and let $Y_t = h_t \cdot N_t$, where h_t represents average hours per worker. The second term in equation (3) can thus be rewritten as $\gamma_2((h_t - 40)N_t)^2$.

The fourth term in equation (3) is the cost of adjusting the number of workers attached to the plant—the extensive margin. Adjustments to the workforce shift the static marginal cost curve horizontally and redefine the minimum efficient scale of production. In contrast to varying the workweek, increasing the number of workers does not lead to increasing static marginal costs, but this move does incur dynamic adjustment costs.

The last term in equation (3) captures the trade-off between inventory-holding costs and stock-out costs, which depends on the level of sales. For industries that produce to stock, such as motor vehicles, this is a standard way to obtain an industry equilibrium in which inventories are nonzero.

The modern production-smoothing model of inventory behavior is a simplified version of this original Holt et al. model. The models used by Blanchard (1983) and those surveyed by Ramey and Kenneth D. West (1999), for example, do not distinguish between the intensive and extensive margins of labor input, and thus are special cases of the Holt et al. model in which $\theta = 0$ and output equals (or is in fixed proportion to) workforce. All increases in production imply rising marginal costs in these models, and there is no distinction between boosting the workweek and hiring more workers. This distinction is very important in the auto industry, however.

The main reason that the special case of the original inventory model became dominant is that it has only one endogenous state variable—the level of lagged inventories. In contrast, the Holt et al. model has an additional endogenous state variable—the level of lagged workforce. Going from one state variable to two state variables makes the system significantly more difficult to solve analytically.

B. *Solution: Optimal Production Scheduling*

For simplicity, we set γ_1, γ_3, and c equal to zero, since these linear terms affect only the means of the variables in the solution and not the dynamics. The first-order conditions with respect to workforce and inventories in the current period are written respectively as equations (4) and (5):

$$(4)\qquad \gamma_4(N_t - N_{t-1}) - \theta\gamma_2(Y_t - \theta N_t) = E_t\{\beta\gamma_4(N_{t+1} - N_t)\};$$

$$(5)\qquad E_t\{\gamma_2(Y_t - \theta N_t) + \beta\alpha_1(I_t - \alpha_2 S_{t+1})\} = E_t\{\beta\gamma_2(Y_{t+1} - \theta N_{t+1})\}.$$

First-order condition (4) states that employment, given some level of output, is optimized when the cost of adding one more worker this period, less the savings in current-period production costs, equals the discounted cost of adjusting workforce by one less worker next period. First-order condition (5) is analogous to the first-order condition obtained in the simple model without workforce, and it states that the cost of producing one more unit in the current period and storing it in inventory equals the discounted savings of producing one less unit next period. When workforce and output decisions are optimized simultaneously, the solution path is homogeneous of degree zero in α_1, γ_2, and γ_4.

The decision rules for N_t and I_t, assuming rational expectations, depend on two state variables—N_{t-1}, I_{t-1}—and on S_t, as shown in equation 6:[10]

$$(6)\qquad \begin{bmatrix} N_t \\ I_t \end{bmatrix} = \mathbf{C} \cdot \begin{bmatrix} N_{t-1} \\ I_{t-1} \end{bmatrix} + \mathbf{d} \cdot S_t.$$

The persistence of sales, ρ, affects only the coefficients in **d** and not those in **C**.

III. The Effects of Sales Persistence on the Variances of Output, Inventories, and Workforce

Expressions for the variances of output, inventories, and workforce and the covariance between inventory investment and sales are readily derived from the decision rules. Because the decision rules from the general model are rather complicated functions of the parameters, it is first useful to analyze two special cases. The results are then shown to hold up in the general model as well.

A. *Special Case 1: Simplified Modern Production-Smoothing Model*

Consider first a special case of the Holt et al. model that sets $\theta = 0$ and $\gamma_4 = 0$. This is a simpler version of the model used by Blanchard (1983), in which $\theta = 0$ and $Y = N$. For a particular parameterization of his model, he showed that the production response to a sales shock was greater if the persistence of sales was greater.

The optimal decision rule for production in this simplified model, assuming rational expectations, is given by:

$$(7)\qquad Y_t = -(1 - \lambda)I_{t-1} + \phi S_t,$$

where $\lambda = \frac{1}{2}\left\{\frac{1}{\beta} + 1 + \frac{\alpha_1}{\gamma_2} - \sqrt{\left[\frac{1}{\beta} + 1 + \frac{\alpha_1}{\gamma_2}\right]^2 - \frac{4}{\beta}}\right\}$

and $\phi = \dfrac{1 - \lambda + \beta\lambda\rho\dfrac{\alpha_1\alpha_2}{\gamma_2}}{1 - \beta\lambda\rho}.$

As long as α_1, α_2, and γ_2 are nonnegative, λ will lie between zero and unity, and ϕ will be positive. While λ depends on neither α_2 nor ρ, ϕ is increasing in both of these parameters.

The relative variances of production and sales and the covariance of sales and inventory investment in the simplified model are given in equations (8) and (9):

$$(8)\qquad \frac{\sigma_Y^2}{\sigma_S^2} = 1 + \frac{2(1-\rho)(\phi - 1)}{1 - \lambda\rho} + \frac{2(\phi - 1)^2(1 - \rho + \rho\lambda^2)}{(1 + \lambda)(1 - \lambda\rho)};$$

[10] See the Mathematical Appendix (available at http://www.e-aer.org/data/dec06/20040244_app.pdf) for the full solution.

$$(9) \qquad \sigma_{S,\Delta I} = \frac{(\phi - 1)\sigma_{\varepsilon S}^2}{(1 - \lambda\rho)(1 + \rho)}.$$

The value of ϕ relative to unity is an important determinant of the relationship between production, sales, and inventories, and ϕ is an increasing function in ρ. If $\phi > 1$, which is more likely when shocks to sales are persistent, the covariance is positive and the variance of production is greater than the variance of sales. The intuition is as follows: when ρ is high, the firm anticipates that sales will remain elevated for a longer time following a positive shock to sales, and it should raise production more in order to prevent its inventory-sales ratio from dipping too low for an extended period. If $\phi < 1$, on the other hand, the covariance of sales and inventory investment is negative, and the variance of production is potentially, but not necessarily, less than the variance of sales.

A decline in ρ alone thus could turn the sign of the covariance between sales and inventory investment from positive to negative, and it could also reduce the variance of output below that of sales. Neither measure, however, is monotonically increasing in ρ for all possible parameter values. In order to determine sufficient conditions under which these measures are strictly increasing in ρ, we numerically investigated the parameter space for these functions, with β preset to 0.997 (annual discount rate of 4 percent). We searched over values of α_2 from 0 to 5 months, values of α_1/γ_2 from 0.001 to 5, and values of ρ from 0.01 to 0.99.

The covariance of sales and inventory investment is monotonically increasing in ρ for virtually all parameter values. An increase in ρ always leads to an increase in the covariance of sales and inventory investment, as long as the penalty for deviating from desired inventories is not too small relative to the marginal cost of production, or as long as $\alpha_1/\gamma_2 > 0.05$. For $\alpha_2 = 2.5$, which coincides with the average inventory-sales ratio in the automobile industry (stated in months), $\alpha_1/\gamma_2 > 0.027$ guarantees that the covariance is increasing in ρ.

Most parameter values also imply that an increase in ρ leads to a rise in the variance of production relative to the variance of sales. When $\rho < 0.6$, σ_Y^2/σ_S^2 is monotonically increasing in ρ for all values of α_2 and α_1/γ_2 within the ranges explored. For $\alpha_2 = 2.5$, as shown in Figure 3, the

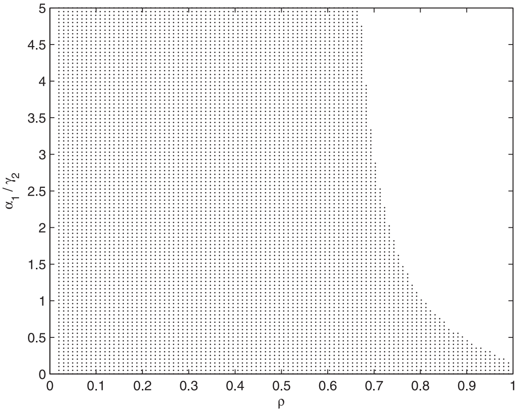

FIGURE 3. PARAMETER REGION FOR WHICH $\frac{\partial(Var(Y_t)/Var(S_t))}{\partial\rho} > 0$

Notes: Parameters other than ρ and α_1/γ_2 are held fixed. β is set to 0.997 (4-percent annual discount rate) and α_2 is set to 2.5 months.

ranges of values for α_1/γ_2 and ρ over which the derivative with respect to ρ is not positive is confined to high values of ρ when α_1/γ_2 exceeds about 0.1, a threshold that is higher than the level consistent with the historical volatilities of U.S. production and sales of autos, as we show below.

B. *Special Case 2: Model without Inventories*

Consider next a special case in which plants cannot hold inventories but are able to choose workforce (i.e., $Y_t = S_t$ and $\alpha_1 = \alpha_2 = 0$). The stochastic Euler equation for this problem is given by

$$(10) \qquad E_t\{\beta\gamma_4 N_{t+1} - [(1 + \beta)\gamma_4 + \gamma_2\theta^2]N_t + \gamma_4 N_{t-1}\} = -\gamma_2\theta S_t.$$

The rational expectations solution to equation (10) that satisfies the transversality condition is:

$$(11) \qquad N_t = \lambda N_{t-1} + \frac{\gamma_2\theta}{\gamma_4}\frac{\lambda}{1 - \beta\lambda\rho}S_t,$$

where $\lambda = \frac{1}{2}\left\{\frac{1}{\beta} + 1 + \frac{\gamma_2\theta^2}{\beta\gamma_4} - \sqrt{\left[\frac{1}{\beta} + 1 + \frac{\gamma_2\theta^2}{\beta\gamma_4}\right]^2 - \frac{4}{\beta}}\right\}.$

As long as all parameters are positive, λ will be between zero and unity. The variance of workforce relative to the variance of sales (which in this simple case is equal to the variance of production) is given by

$$(12)\quad \frac{\sigma_N^2}{\sigma_S^2} = \frac{\sigma_N^2}{\sigma_Y^2} = \frac{1}{1-\lambda^2}\left[\frac{\gamma_2\theta}{\gamma_4}\frac{\lambda}{1-\beta\lambda\rho}\right]^2\left[1 + 2\,\frac{\lambda}{1-\lambda\rho}\right].$$

This expression is always increasing in ρ, and the intuition is as follows: Because changes to workforce entail adjustment costs while changes to average hours do not, workforce will account for a larger part of output volatility when the adjustment costs pay off—when sales are persistent and the new workforce is expected to yield lower static marginal costs well into the future.

C. *General Model*

In the general model, plants choose both output and workforce in each period. While closed-form solutions to the model do exist, the effects of individual parameters on the decision rules are difficult to see in these complicated expressions. Fortunately, the intuition developed in the special cases is largely unchanged when inventories and workforce are optimized jointly.

The variance of output relative to sales, the covariance between sales and inventory investment, and the variance of workforce relative to output are plotted against values of ρ in the top row of Figure 4. The solid lines represent these measures of volatility for a baseline set of parameters, chosen so that the variance of output relative to sales and the variance of workforce relative to output in the model are consistent with the empirical counterparts in industry data prior to 1984.[11] We do not attempt to estimate the parameters from this model because, while having the advantage of being simple and intuitive, the model does not capture the important nonconvexities in the automobile industry.[12] The other (not solid) lines plot these measures for alternative parameterizations of the model and will be discussed below.

The relationship between the variance of production and ρ is pictured in the left-most graph. Production is less volatile than sales for low values of ρ, though the ratio is not monotonically increasing in ρ for all values of ρ. As shown in the middle graph, the covariance between sales and inventory investment increases monotonically in ρ with this set of parameter values. The variance of workforce relative to output, shown in the right-most graph, is generally increasing in ρ, but again, not universally. The ratio increases in ρ as long as ρ is sufficiently large.

The impact of ρ on the variances is consistent with the magnitude of the changes observed in the data. For monthly sales of domestic cars, recall that ρ declined from 0.85 to 0.56. According to the baseline parameterization of the model, this would lead the ratio of the output-to-sales variance to decline from well over unity to around 0.7, even greater than the change shown in Table 1. The persistence of truck sales declined from 0.93 to 0.69, and the model predicts a smaller decline in the variance ratio, which also matches the pattern in the data.

The decline in the covariance between inventory investment and sales that occurs in the model when the persistence of sales is reduced is qualitatively consistent with the change measured in industry data. This simple convex approximation to the industry cost function does not, however, replicate the exact magnitude of the change very well. Covariance declines from 0.28 to -0.20 in the model, while in the industry data it declined from -0.03 in the early period, to -0.20 in the late period.

Similarly, the ratio of the variance of workforce to output in the model declines from about 0.65 to 0.50 when ρ falls. In the plant-level data, the ratio fell from 0.65 in the early period, to 0.4

[11] Specifically, α_2 is set to 2.5, which is the average inventory-sales ratio (in months), θ is normalized to 1, ρ is estimated from the first-order autocorrelation of sales to be 0.85, and β is preset to 0.997 (4-percent annual rate). With these parameters in place, $\alpha_1/\gamma_2 = 0.085$ and $\gamma_4/\gamma_2 = 1.65$ together yield decision rules in which $\sigma_Y^2/\sigma_S^2 = 1.28$ and $\sigma_{\theta N}^2/\sigma_Y^2 = 0.647$, values which match their empirical counterparts for cars in Tables 1 and 2.

[12] Ramey and Vine (2004) analyzed a model of industry costs that accounted for these nonconvexities. A further complication for estimation is that the worker-output ratio (θ) is not constant in the data, creating time-varying coefficients.

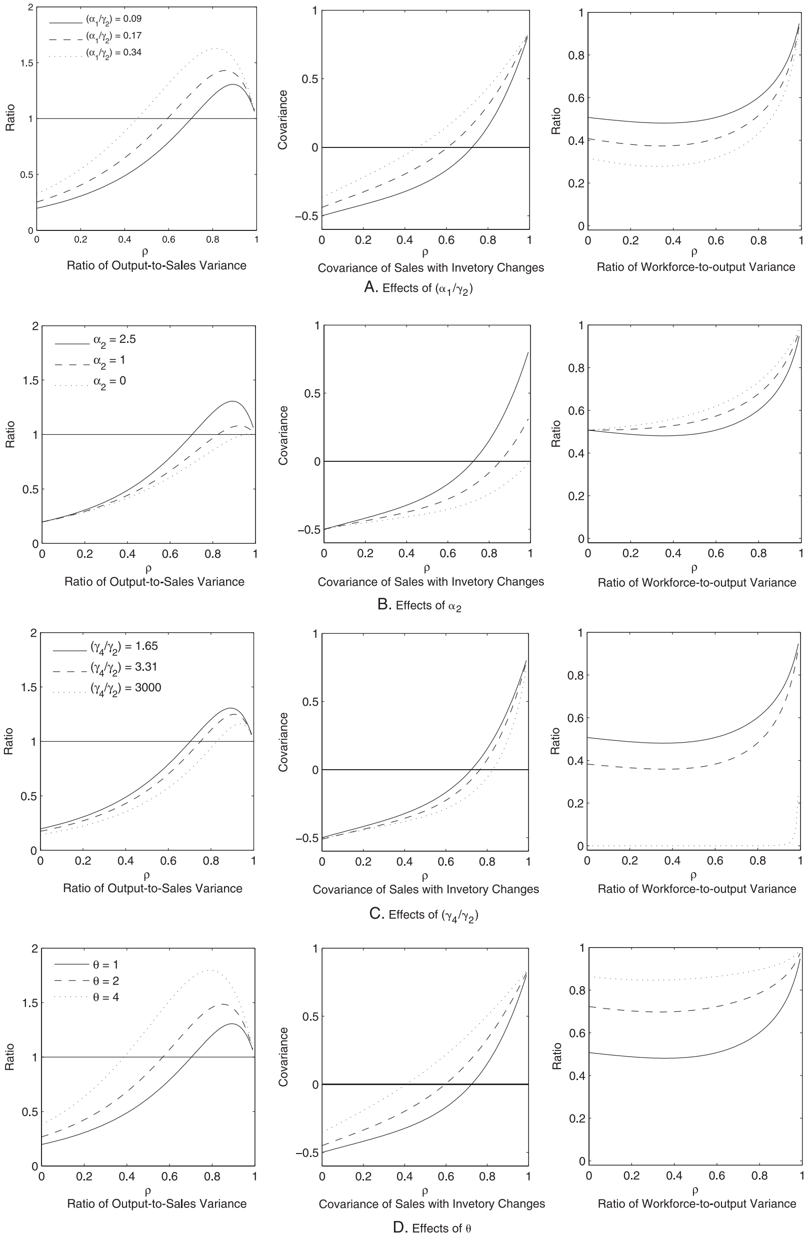

FIGURE 4. VARIANCE OF SALES, OUTPUT, AND EMPLOYMENT AS A FUNCTION OF ρ

Notes: Parameters not subject to variation in each graph are fixed at their benchmark level: $\alpha_1/\gamma_2 = 0.09$, $\gamma_4/\gamma_2 = 1.65$, $\theta = 1$, $\alpha_2 = 2.5$, $\beta = 0.997$. The variance of sales innovations is fixed so that the variance of sales equals 2.45 when $\rho = 0.85$. This implies that covariance is measured on the same scale as cars in Table 1.

in the late period, a decline that was more pronounced than the model predicts.

D. *The Effects of Other Model Parameters on Variance*

Could changes to parameters other than ρ also be consistent with the changes measured in the auto industry data around 1984? To investigate this, the effects of various parameter values for α_1/γ_2, α_2, γ_4/γ_2, and θ on the key volatility measures (as a function of ρ) are shown in the panels of Figure 4. Alternative parameterizations appear as the dashed and dotted lines in each graph.

The effects of raising α_1/γ_2 on the volatility measures are shown in panel A. When the penalty for deviating from the desired inventory-sales ratio is higher, the contribution of workforce adjustments to output variance moves down. The reason is that average hours per worker can adjust quickly following a sales shock, so this margin is used more intensively when maintaining inventories becomes relatively more important. The variance of output relative to the variance of sales, however, moves up when α_1/γ_2 is higher, and the correspondence between the covariance of sales with inventory investment and ρ also shifts up.

Panel B describes the effects of lowering the target inventory-sales ratio, α_2, a change that has been observed in some industries (though not in inventories of finished autos). The variance of output declines relative to the variance of sales, and inventory investment becomes less procyclical. The contribution of workforce to output variance, however, increases for all values of ρ between zero and one.

Increases in the cost of adjusting workforce, γ_4/γ_2, shown in panel C, do move the key volatility measures in the desired directions. However, the sensitivities of changes in the variance of output relative to sales and the covariance of inventory investment with sales to changes in γ_4/γ_2 are very low. Even if γ_4/γ_2 is raised to 3,000, the effects on the variance of output relative to sales and on the covariance are small.

Lastly, the graphs in panel D show that a higher value for θ boosts the variance of output relative to sales and the covariance of inventory investment with sales. This change also raises the variance contribution of workforce substantially for all values of ρ less than one, because every dollar paid in workforce adjustment costs moves the minimum efficient scale of production by a larger magnitude. This, in turn, makes workforce a more cost-effective margin of adjustment.

To summarize, changes in only two parameters, γ_4/γ_2 and ρ, are capable of simultaneously moving all three key variance measures in the right direction. The effects of changes in γ_4/γ_2 on two of the key variances, however, are much smaller than the effects of changes in ρ on the variances. Thus, among these parameters, only changes in ρ are likely to reproduce the changes in the data that were documented earlier.

IV. Conclusion

This paper has documented significant changes in the behavior of production, sales, and inventories in the automobile industry. The variance of production has fallen more than 70 percent since 1984, whereas the variance of sales has fallen by less. The covariance of inventory investment and sales has also become more negative. Moreover, plants are now more likely to vary average hours per worker than the number of workers employed.

These changes in production scheduling occurred together with a significant decline in the persistence of sales shocks. Our theoretical analysis of the original Holt et al. (1960) production-smoothing model showed that a reduction in sales persistence, all else equal, lowers the volatility of output relative to sales, lowers the covariance of inventory investment and sales, and reduces the portion of output volatility that stems from employment changes. Changes in the other model parameters cannot replicate all three of these changes.

Many of these developments in production volatility have been documented in other industries as well. Our analysis suggests that it would be fruitful to determine whether the behavior of sales changed in these industries as well. Determining the source of the decline in persistence is also an important area for future research.

Data Appendix

The contribution of the variance of motor vehicle production to GDP is calculated by comparing the variances of the growth rates of the following two variables: chain-weighted total GDP and chain-weighted GDP less motor vehicles. Both of these variables are available from Table 1.2.3 from the National Income and Product Accounts. From 1967 through 2004, the variance of GDP growth is 11.30 and the variance of the growth of GDP less motor vehicles 8.65. We used nominal GDP figures to compare levels.

Figure 1 and Table 1—Car and truck sales are seasonally adjusted by the Bureau of Economic Analysis (BEA) and are in millions of units at an annual rate. We would have preferred to limit our analysis to light vehicles, but light-truck production is not distinguishable from total truck production prior to 1977, and data for light-truck inventories are not available before 1972. Thus, Figure 1 and Table 1 include all trucks. All production data are seasonally adjusted by the Federal Reserve Board. Car inventories are seasonally adjusted by the BEA, though we had to use our own seasonal adjustment method for trucks. For truck inventory investment, we regressed the unadjusted inventory investment on the difference between seasonally adjusted and unadjusted production and sales of trucks.

Figure 2—Car sales, car inventories, and light-truck sales are available on a seasonally adjusted basis from the BEA. The level of light truck inventories was seasonally adjusted with X12ARIMA.

Table 2—The dataset was constructed from industry trade publications in part by Bresnahan and Ramey (1994), who collected the data covering the 50 domestic car assembly plants operating during the 1972–1983 period, and by Ramey and Vine (2004), who extended it to include all 103 car and light truck assembly plants operating in the 1972–1983 and 1990–2001 periods. The data were collected by reading the weekly production articles in *Automotive News*, which report the following variables for all North American assembly plants: (a) the number of regular hours the plant works; (b) the number of scheduled overtime hours; (c) the number of shifts operating; and (d) the number of days per week the plant is closed for (i) union holidays, (ii) inventory adjustments, (iii) supply disruptions, and (iv) model changeovers. Observations on the line speed posted on each assembly line were collected from the *Wards Automotive Yearbook*.

REFERENCES

Aizcorbe, Ana M. 1992. "Procyclical Labour Productivity, Increasing Returns to Labour and Labour Hoarding in Car Assembly Plant Employment." *Economic Journal*, 102(413): 860–73.

Automotive News. Various issues, 1972–2001. Detroit: Crain Automotive Group.

Blanchard, Olivier J. 1983. "The Production and Inventory Behavior of the American Automobile Industry." *Journal of Political Economy*, 91(3): 365–400.

Bresnahan, Timothy F., and Valerie A. Ramey. 1994. "Output Fluctuations at the Plant Level." *Quarterly Journal of Economics*, 109(3): 593–624.

Golob, John E. 2000. "Post-1984 Inventories Revitalize the Production-Smoothing Model." Unpublished.

Hall, George J. 2000. "Non-Convex Costs and Capital Utilization: A Study of Production Scheduling at Automobile Assembly Plants." *Journal of Monetary Economics*, 45(3): 681–716.

Hetrick, Ron L. 2000. "Analyzing the Recent Upward Surge in Overtime Hours." *Monthly Labor Review*, 123(2): 30–33.

Holt, Charles C., Franco Modigliani, John F. Muth, and Herbert A. Simon. 1960. *Planning Production, Inventories and Workforce*. Englewood Cliffs, NJ: Prentice-Hall.

Kahn, James A., Margaret M. McConnell, and Gabriel Perez-Quiros. 2002. "On the Causes of the Increased Stability of the U.S. Economy." *Federal Reserve Bank of New York Economic Policy Review*, 8(1): 183–202.

Kim, Chang-Jin, and Charles R. Nelson. 1999. "Has the U.S. Economy Become More Stable? A Bayesian Approach Based on a

Markov-Switching Model of the Business Cycle." *Review of Economics and Statistics*, 81(4): 608–16.

McConnell, Margaret M., and Gabriel Perez-Quiros. 2000. "Output Fluctuations in the United States: What Has Changed since the Early 1980's?" *American Economic Review*, 90(5): 1464–76.

Ramey, Valerie A., and Daniel J. Vine. 2004. "Tracking the Source of the Decline in GDP Volatility: An Analysis of the Automobile Industry." National Bureau of Economic Research Working Paper 10384.

Ramey, Valerie A., and Kenneth D. West. 1999. "Inventories." In *Handbook of Macroeconomics.* Vol. 1B, ed. John B. Taylor and Michael Woodford, 863–923. Amsterdam: Elsevier Science, North-Holland.