Tracking the Source of the Decline in GDP Volatility: An Analysis of the Automobile Industry

by

Valerie A. Ramey
University of California, San Diego and National Bureau of Economic Research

and

Daniel J. Vine
University of California, San Diego

First Draft: December 2001
This Draft: January 2003

Abstract

Recent work by Kahn, McConnell and Perez-Quiros (2002) has argued that improvements in information technology and inventory management are the chief source of the decline in output volatility since 1984. A key piece of evidence in support of their argument is the decline in the variance of production relative to the variance of sales. This paper investigates the alternative hypothesis that declines in production volatility relative to sales stem from changes in the nature of the sales process rather than from changes in the structure of production and inventories. In particular, a small decrease in the volatility of sales can lead to a large decrease in the volatility of production if there are nonconvexities in the cost function. We confirm this intuition using simulations based on the cost function of automobile assembly plants. We then conduct an analysis of changes in production scheduling across U.S. automobile plants from the 1970s to the 1990s with a new plant-level dataset.

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 the UCSD Applied Lunch Group, John Haltiwanger and the participants at the International Society for Inventory Research seminar in Atlanta, GA. Valerie Ramey gratefully acknowledges support from National Science Foundation Grant SBR-9617437 through the National Bureau of Economic Research, and National Science Foundation grant # 0213089.

I. Introduction

Recent papers by Kim and Nelson (1999) and McConnell and Perez-Quiros (2000) uncover a dramatic decline in the volatility of the U.S. economy beginning in 1984. The volatility of GDP growth since 1984 has been 50 percent lower than it was in the post-war period before 1984. Interestingly, statistical tests point to a structural break in the first quarter of 1984 rather than to a gradual decline. The phenomenon also appears to extend beyond U.S. borders. Blanchard and Simon (2001) show that all G-7 countries save Japan have experienced a decline in volatility in recent periods. Only the UK and Canada show the sharp drop in the mid-1980s, though, with the other countries exhibiting more gradual declines or later falls.

This discovery raises an important question: Has output volatility declined in a meaningful and permanent way, or have we simply enjoyed a reprieve from the turbulence of the 1970s and early 1980s? Possible answers to this question depend on which of the three leading explanations for the decline in volatility is most accurate: (1) Good Luck, (2) Good Policy, or (3) Structural Change. The "Good Luck" hypothesis argues that the decline in volatility is a result of a fortuitous decline in the volatility of shocks hitting the economy (e.g. Ahmed, Levin, and Wilson (2000), Stock and Watson (2002)). Advocates of the "Good Policy" hypothesis argue that improved monetary policy is the key source of the decline in volatility of the U.S. economy (e.g. Boivin and Giannoni (2002) and Clarida, Galí, and Gertler (2000)). Finally, the "Structural Change" hypothesis refers to the innovations in manufacturing technology and inventory management that allow smooth production along the supply chain (e.g. Kahn, McConnell and Perez-Quiros (2002)). This potential source has recently received a lot of attention, as the decline in the volatility of aggregate output exceeds that of final sales.

Despite a rapidly developing literature in this area, answering this question in a definitive way has proven to be an elusive task. First, a significant reduction in volatility has been discovered almost universally across the U.S. economy.[1] This makes it difficult to isolate any one sector from the others, though durable goods production and residential investment appear as leading candidates. Second, the conclusions reached from time-series analysis on aggregate data have been difficult to interpret. The general consensus is that the current stability is the consequence of a reduction in the variance of the forecast errors. To further associate these

[1] Stock and Watson (2002) test 168 U.S. macroeconomic time-series and discover the pervasiveness of an overall decline in volatility.

forecast errors with measurable shocks, however, such as monetary and fiscal shocks, productivity shocks, supply shocks, etc., has achieved only limited success.

This study addresses the decline in U.S. GDP volatility in the context of decisions made at the plant-level in an industry at the forefront of this change – the U.S. automobile industry. Using aggregate automobile industry data as well as a new and highly disaggregated dataset that tracks the weekly production scheduling at automobile assembly plants between 1972 and 2001, we investigate the extent to which the decline in volatility stems from structural changes in the process governing sales versus structural changes in the process governing production. We conclude that the decline in aggregate output volatility is linked to the interaction between (1) a decline in the measured persistence of sales shocks and (2) plant-level nonconvexities in production scheduling. In particular, sales appear to be far less serially correlated after 1984 than they were during the 1970s and early 1980s. We demonstrate that an inventory model involving nonconvex costs predicts that a decline in the persistence of sales shocks leads to a decline in the variance of production relative to the variance of sales.

The organization of this paper is as follows: Section II begins with an overview of the volatility patterns for aggregate data and automobiles, and explains the connection between output volatility, inventory investment, and improvements in information and production technology. Section III argues that the automobile industry is particularly suitable for this case study, and presents evidence that the process governing automobile sales has changed noticeably in the post-1984 period. Section IV addresses how a change in sales persistence interacts with production decisions in a simulation featuring nonconvex production costs. Section V confronts our theory with the entire panel of production scheduling variables across U.S. assembly plants. Section VI concludes.

II. Overview of Volatility and Inventory Changes

Before presenting a detailed analysis of the automobile industry, we first examine how its volatility behavior compares with patterns found in aggregate data. We present summary evidence on changes in volatility and inventory behavior at the aggregate level, as well as for key disaggregates. We then discuss the connection between inventory investment and output volatility, and its relation to the "Structural Change" hypothesis.

A. Volatility Behavior in Motor Vehicle Production

Table 1 shows the volatility of output growth in the aggregate economy, as well as in the key sectors of durable goods and automobile assembly. All data are from the NIPA accounts of the BEA. We choose 1967 as the starting date because of constraints on the availability of data for the automobile sector. As the first row of the table shows, the volatility of aggregate GDP growth has declined by about 50 percent from the pre-1984 period to the post-1984 period. [2] The declines are even greater for durable goods output and motor vehicle output.

The bottom half of Table 1 shows the sectoral volatility using current dollars. We must use current dollars to subtract out durables and motor vehicles since chain-weighted numbers cannot be added. Fortunately, the story in current dollars is similar to the one in chained dollars. The volatility of GDP excluding durable goods is not only substantially less than durable good output volatility, but has also fallen by less in percentage terms. Similarly, the volatility of durable goods output excluding motor vehicles is substantially less than the volatility of motor vehicles output. Its volatility has also declined by less.

We conclude from the table that the motor vehicle industry represents an ideal case study of the decline in volatility. Its volatility behavior is similar to, but more dramatic than, the behavior of GDP and durable goods overall. Thus, understanding the decline in volatility in the automobile industry is likely to shed light on the decline in volatility overall.

B. Information Technology and Inventory Investment

Kahn, McConnell and Perez-Quiros (2002), hereafter KMPQ, give perhaps the most compelling evidence for the structural change hypothesis. They base their explanation on the fact that while there is evidence of structural breaks in volatility across many broad sectors, the evidence points to durable goods as the sector accounting for most of the decline in GDP volatility.[3] KMPQ argue that a structural change in production and inventory behavior is the

[2] The volatility of aggregate GDP growth in the truncated early period is similar to the volatility of the extended period from 1953 – 1983 shown in KMQP (2002).

[3] Kim, Nelson, and Piger (2001) present evidence of structural breaks across broad sectors. McConnell and Perez-Quiros (2000), Warnock and Warnock (2000), and Kahn, McConnell, and Perez-Quiros (2002) offer evidence that the durable goods sector plays the most important role. Stock and Watson (2002) claim residential fixed investment played an equally large role as durable goods production.

leading culprit in the decline in output volatility, and suggest that improvements in information technology are leading to better inventory management.

Two of the most striking pieces of evidence presented by KMPQ in favor of the "better inventory management hypothesis" are the differential declines in final sales versus production volatility and the changing covariance of inventories with final sales. To see how this evidence relates to inventory management, consider 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, we have the following relationship between the variance of production and the variance of sales:

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

The standard version of the production-smoothing model of inventories predicts that the variance of production should be less than the variance of sales. Thus, the covariance of inventory investment and sales should be negative. Historically, the opposite has been true.

The decomposition is somewhat more complicated in the case of nonstationary variables.[4] For comparison purposes, we follow KMPQ's procedure for decomposing the variance of production changes. To be specific, we first-difference the inventory identity and then divide by Y_{t-1} to obtain:

$$(2) \qquad \frac{Y_t - Y_{t-1}}{Y_{t-1}} = \frac{S_t - S_{t-1}}{Y_{t-1}} + \frac{\Delta^2 I_t}{Y_{t-1}}$$

For ease of notation, let $\hat{y} = \frac{Y_t - Y_{t-1}}{Y_{t-1}}$, $\hat{s} = \frac{S_t - S_{t-1}}{Y_{t-1}}$, and $\Delta\hat{i} = \hat{y} - \hat{s}$.[5] Then we can decompose the variance of production growth as follows:

---

[4] Unit root tests on the BEA data fail to reject a unit root.

[5] We define the inventory term as the residual rather than using the actual inventory data because the chain-weighted data do not add up perfectly. The variance and covariance results are similar whether we use the actual data or construct the inventory term.

(3) $$Var(\hat{y}_t) = Var(\hat{s}_t) + Var(\Delta\hat{i}_t) + 2Cov(\hat{s}_t, \Delta\hat{i}_t)$$

Table 2 reports the results of this decomposition for all durable goods as well as motor vehicles. The results for durable goods are very similar to those reported by KMPQ. The variance of production growth falls by 77 percent while the variance of final sales falls by only 37 percent. As a result of the differential fall, production becomes just as volatile as sales in the post-1984 period. The source of the differential decline is the behavior of inventories. Not only does the volatility of inventory investment fall by nearly half, but also the covariance of final sales and inventory investment turns from positive in the early period to negative in the post-1984 period.[6] Instead of contributing to the volatility of the economy as they once did, inventories now appear to stabilize the economy.

The lower panel of the table shows the results of the decomposition for motor vehicles, and the qualitative results are very similar to those for durable goods. The volatility of production falls much more than the volatility of sales, so much so that production becomes less volatile than sales in the later period. Moreover, the covariance between sales and inventory investment becomes negative.

These results suggest that an explanation of the decline in aggregate output volatility must be consistent with the following observations: (1) The source must have particularly strong effects on the durable goods sector as opposed to nondurables and services; and (2) it must have strong effects on production in that sector, but only weak effects on final sales. KMPQ argue that these observations cast doubt on the "better monetary policy" explanation since one should expect it to work mostly through sales.

KMPQ advocate the "Information Technology (IT)" hypothesis as an alternative explanation. As support, they describe the changes that have occurred since the adoption of various information technology innovations. For example, electronic scanning of bar codes allows for automatic restocking based on real-time sales information. This innovation allows for more efficient management of inventories along the entire supply line. Another innovation they discuss is flexible manufacturing. Computer numerically controlled machine tools have led to a reduction in set-up times required to produce different specifications. This change in set-up

[6] Golob (2000) first discovered this switch in the sign of the covariance term.

costs lowers optimal batch size, which varies inversely with inventory levels. Both of these innovations would be expected to work to reduce desired inventory-sales ratios. A reduction in the desired inventory-sales ratio should weaken or eliminate the tendency for inventories to be so procyclical, and hence so destabilizing.

The facts highlighted by KMPQ are consistent with the hypothesis that information technology and improved inventory management is the source of the decline in volatility. There are, however, two puzzles that cast doubt on this hypothesis. The first is why technology adoption, which usually follows an S-curve, should show up as a one-time structural break in volatility. The second concerns the inventory-sales ratio. As discussed above, we would expect the information technology innovations to reduce the inventory-sales ratio in the post-1984 period relative to the earlier period. The data do not give such a clear picture.

Figure 1 shows the ratio of durable goods inventories to final sales of durable goods since 1947. The data are in current dollars, since it is inappropriate to compare levels of chained dollars.[7] Two distinct features in the graph are important. First, there was a large run-up in the inventory-sales ratio from the early-1970s to the early 1980s. Second, there has been an almost steady decline in the inventory-sales ratio from the early 1980s to the present. The continuing decline during the 1990s brought the inventory-sales ratio for durable goods below its level during the 1950s and 1960s. Information technology may be the reason for the recent decline; it is not clear what the source of the early 1970s run-up was. Also, KMPQ demonstrate that the volatility of GDP and durable goods was approximately equal across the periods 1953–1968 and 1969–1983, before falling in 1984. The behavior of the inventory-sales ratio does not line up very closely with the changes in volatility over time.

Within motor vehicles, there has been no trend in the inventory-to-sales ratio over time. Figure 2 depicts this ratio in terms of the number of month’s worth of sales (in physical units) in the domestic inventory stock of cars and light trucks. While this ratio shows a great deal of seasonal and business cycle variation, the average has been remarkably stable.

It is therefore interesting to explore an alternative explanation for the changes in production and sales volatility. In particular, is it possible that the decline in production

[7] Inventory-sales ratios created using chained data display a very different pattern. Their pattern suggests that inventory-sales ratios were relatively constant until 1984, and then followed an unprecedented downward trend. The large difference in the patterns across current and chained dollar data illustrates the errors one can make when comparing levels of two chain-weighted series (i.e. inventories versus final sales) in years far from the base year.

volatility relative to sales stems from changes in the nature of the sales process rather than from changes in the structure of production and inventories? This paper investigates the possibility both theoretically and empirically.

III. The U.S. Automobile Industry

Production behavior in the U.S. automobile industry exhibits many of the classic traits that characterize both seasonal and business cycle movements in U.S. manufacturing on a larger scale. As was shown in the previous section, the industry serves as an appropriate forum to investigate the changes many macroeconomists have observed in U.S. aggregate data over the 1980s. Automobile assembly also happens to be one of the most volatile components of durable goods production, a fact that past research has attributed to nonconvex costs that sway automakers to use intermittent production in place of producing at levels below full capacity.

The automobile industry has also implemented many of the technological changes showcased by KMPQ in their advocacy of the *Information Technology Hypothesis*. It has been particularly blessed with advances in assembly line technology and was one of the first industries to adopt *just-in-time* inventory management in the 1980s. Therefore, if advances in information technology have revolutionized the fundamentals of U.S. manufacturing and have delivered unprecedented economic stability as a consequence, a natural place to look for plant-level evidence is within the automobile industry.

Car production has served as a source of insight into other economic questions as well. Automobile industry data has been used to test models of inventory behavior (e.g. Blanchard (1983), Kahn (1992), Kashyap and Wilcox (1993)), and Aizcorbe (1992), Bresnahan and Ramey (1994), and Hall (2000) have documented nonconvexities in costs and increasing returns in automobile assembly. Finally, Cooper and Haltiwanger (1993a,b) have used automobile industry data to test hypotheses about industry complementarities.

A. Overview of U.S. Automobile Production

Before proceeding into an analysis of industry volatility, it is first useful to point out certain features of automobile production. From here on, we use unit data for automobiles rather than NIPA data, as the unit data are expected to have less measurement error and allow for direct comparison of inventories, production and sales without the constraints of chain-weighted dollars. Figures 3A and 3B depict monthly sales and production of U.S. passenger car and truck from 1967:01 through 2002:09.[8] These plots reveal evidence of more smoothness in the late portion of the sample relative to the early portion, though seasonal fluctuations make this difficult to see.

A second feature visible in Figure 2 is the slow decline in sales of passenger cars over the sample period, and their replacement with light trucks. In the 1970s, U.S. consumers shifted their new vehicle purchases away from large domestic sedans to more fuel-efficient cars offered by foreign manufacturers. Over time this led to a rise in the market shares belonging to foreign automakers, and this trend persisted through the 1980s. American consumers shifted their preferences yet again in the 1990s back to larger vehicles as vans, light-duty trucks and later the sports utility vehicle rapidly gained market share and overtook passenger cars.[9]

B. Testing for a Structural Break

Following the testing strategy of McConnell and Perez-Quiros (2000), the volatility of aggregate vehicle output (truck and passenger car classes combined) is tested for a structural break. The break date is assumed unknown, which necessitates special care in interpreting p-values in the presence of a nuisance parameter. The possible range of break dates include the middle 70% of the sample, and for each of these candidate dates the Wald statistic $F_n(T)$ is calculated. The timing of the structural break is determined as the date with the largest Wald statistic, and its significance judged by three test statistics - the supremum (Sup F_n), the average (Ave F_n) and the exponential (Exp F_n). The results are displayed in Table 3, where p-values are determined with tables provided by Hansen (1997).[10]

[8] Within the class of trucks, we would prefer to focus on only light trucks, since they are mostly a consumer product like cars. Unfortunately, we were only able to obtain production figures that included heavy trucks. Heavy trucks represented 22 percent of truck production in 1967 but only 7 percent in 2000 as light truck production grew.

[9] *Wards Automotive Yearbook* names 1996 as the year new light truck sales surpassed passenger car sales in the U.S. market. Light trucks often refer to vans, SUVs and light-duty pickup trucks collectively.

[10] For details on the structural break testing procedure used here, see McConnell and Perez-Quiros (2000).

The procedure reveals that U.S. motor vehicle production does show evidence of a structural break in its volatility similar to the one seen in aggregate GDP. The estimated date of the break occurs in February 1983. This roughly corresponds to the break date discovered by McConnell and Perez-Quiros (2000) for the volatility of aggregate GDP, especially considering the wide confidence intervals associated with such break date estimates. To maintain consistency with other studies of the structural break in U.S. GDP volatility, we chose to compare automobile industry variance estimates on either side of the generally accepted break date: 1984:1.

Earlier we showed the decomposition of variance based on chained dollar data from the BEA. It is useful to do the same decomposition using the unit data, where the different classes of motor vehicles may be analyzed separately. Physical unit data differs from the chain-weighted measures in that stationarity tests on the logarithm of the variables usually reject a unit root in favor of a deterministic trend, with perhaps a break in trend around 1984 for trucks. Thus, rather than difference both sides of the standard inventory identity, we divide both sides by the estimated trend.[11] Therefore, variance here is based on deterministically detrended data, and the statistical results are consistent with the patterns shown in Figures 3A and 3B.

Table 4 decomposes variance for cars and trucks separately.[12] Consider first the case of cars. The data show that both the variance of production and the variance of sales fall after 1984. Moreover, the variance of production falls by a larger percentage than sales, and the covariance of inventory investment with final sales become more counter-cyclical after 1984. These two features are consistent with the aggregate BEA data.

The results for trucks are not as clear. In both seasonally adjusted and unadjusted data, the variance of both production and sales falls in the second period. The variance of production falls by proportionally less than the variance of sales, however, so that the variance of production relative to sales is higher in the second period. On the other hand, the covariance of inventory investment with final sales does become more negative after 1984, just as in the case of cars.

---

[11] To be specific, we calculate the variances and covariances of the terms in $\frac{Y_t}{\hat{Y}_t^T} = \frac{S_t}{\hat{Y}_t^T} + \frac{\Delta I_t}{\hat{Y}_t^T}$, where $\hat{Y}_t^T = \exp(\hat{\beta}_0 + \hat{\beta}_1 t)$ and the β's are the estimated parameters of a regression of $log(Y_t)$ on a time trend.

[12] The variances and covariances do not add up because we have excluded imports and exports to and from Canada and Mexico. These are not an important part of the story.

Thus, the physical unit data give a more mixed picture than the BEA data when the variance decomposition is performed on each class of vehicles separately.

C. The Changing Behavior of Motor Vehicle Sales

The decrease in the volatility of both U.S. vehicle production and sales depicted in the tables above arises from two potential sources: (1) a reduction in the magnitude of shocks to these series, and (2) a change in the dynamic processes that propagate these shocks. Since production decisions are made in accordance with forecasts of future sales, the volatility of production depends not only on the variance of the shocks to the sales process, but also on the persistence of these shocks. Additional insight is therefore found by comparing the persistence and volatility of sales shocks between the two periods. Within the automobile industry, such an exercise reveals that the sales shocks since 1984 have been much less persistent than those prior to 1984.

Aggregate Motor Vehicle Sales

Consider the following simple univariate model of the process for monthly domestic sales data from 1967:1 through 2002:9:

$$(4) \quad Sales_t = \alpha_o + \alpha_1 \cdot Sales_{t-1} + \alpha_2 \cdot trend_t + \beta_0 \cdot D_t + \beta_1 \cdot D_t \cdot Sales_{t-1} + \beta_2 \cdot D_t \cdot trend_t + \varepsilon_t$$

where $\varepsilon_t \sim N\left(0, \sigma^2 (1 + \beta_3 \cdot D_t)\right)$

and
$$D_t = 0 \text{ for } t < 1984:1$$
$$D_t = 1 \text{ for } t \geq 1984:1$$

This model allows all parameters to change in 1984:1: the AR(1) coefficient on sales, the constant term, the trend term, and the variance of the residual. We estimate this model via maximum likelihood for cars alone, light trucks alone, and the combination of cars and light trucks, which we will call “motor vehicles.” In all cases, the regression is estimated with the logarithm of seasonally adjusted unit sales from the BEA.

Table 5 shows the results of this exercise, and the coefficient estimates indicate a significant change in the process governing sales. The constant term and the AR(1) coefficient are different across the two periods for all three aggregates. The trend (which is not significant for the entire period) changes in the case of light trucks and motor vehicles. As for the variance of the shocks, there is a significant decline in the case of light trucks, but not cars or the motor vehicle aggregate.

Of particular interest to our analysis is the change in AR(1) parameter, which serves as a measure of the persistence in monthly changes in sales. In all three cases, the first-order autocorrelation of sales falls between the early and the late periods. For the passenger car class, this parameter falls from 0.8 to 0.5, and for trucks it falls from 0.9 to 0.7. When all motor vehicles are grouped together, this estimate declines from almost 0.9 to 0.6.

To test robustness, we also estimated the model with up to 13 lags of sales and chose the optimal lag length by the BIC criterion. It chose 3 to 5 lags, depending on the aggregate used. When we estimated the expanded model, we calculated three different measures of persistence for each vehicle type and for the combined aggregate: (1) the largest autoregressive root, (2) the sum of autoregressive terms, and (3) the estimated half-life (in months) of a one-time shock to sales. We obtained very similar results with respect to changes in sales persistence, and the results of this robustness check are found in Table 6.[13]

In summary, the aggregate motor vehicle sales process in the 1990s returns to its mean much more quickly following a surprise than was previously the case in earlier decades. It is also clear that most of the change in the unconditional variance of sales described in the tables above comes from a change in the propagation mechanism for sales rather than in the variance of sales shocks.

Motor Vehicle Sales by Model

While changes in the autocorrelation of aggregate motor vehicle sales described above help explain aggregate volatility patterns, what is most important from the perspective of plant-

[13] For a nice discussion on the features and limitations of each measure of persistence, see Pivetta and Reis (2002).

level production scheduling is the behavior of sales of individual models. To this end, we collected data on monthly model-level sales for the cars and light-trucks produced by Ford, General Motors, Chrysler and American Motors between 1965:1 and 2002:9. These companies collectively offered over 250 different models during this time period.[14]

To assess whether the individual model sales have changed in the same way as the aggregate, we estimate a simple AR(1) process on the data. Preliminary investigations suggested that estimating the data as a panel with the AR(1) coefficient constrained to be the same across all models was not consistent with the data generating process. Thus, we estimated the equation on each model separately for each sub-period, 1965 – 1983 and 1984 – 2002. While some models spanned a significant part of both sub-periods, the majority did not.

One issue that is unique to model-level analysis is the life-cycle of vehicle models. Car models clearly have a life-cycle pattern, where sales usually build slowly after the model is introduced, and then trail away as the model is phased out. These patterns can affect the AR(1) coefficient, and after case studies of various models, we decided the best specification was one in which we omitted the first six months and last six months of data for each model.[15]

For each model that existed for at least 24 months, we estimated an AR(1) regression on the logarithm of monthly sales adjusted by the BEA's seasonal factor for cars, first for the period 1965:1 – 1983:12 and then for 1984:1 – 2002:9. We summarize the collection of AR(1) estimates from each period as well as the standard deviation of the error term by taking a weighted average based on each models' contribution to total sales. The results are shown in Table 7.

The decline in the persistence of sales shocks over the periods is even more dramatic at the model level than at the aggregate level, falling from over 0.8 in the first period to below 0.3 in the second period. In contrast, the variance of the residuals rises in the second period. When both of these effects are taken together, the unconditional variance of sales at the model level has not changed much between the two periods. While assembly plants were hit by higher variance shocks after 1984 than in earlier times, these shocks also expired much more quickly in the second period.

[14] Note that we count completely redesigned cars as different models because we want to account for the effects of model introduction and ending on the sales process. Thus, the Mustang I, II and III are counted as three different models.

Overall, both the aggregate industry data and the model-level data indicate that the persistence of the shocks to sales has changed substantially during the 1980s. The next section shows how a change in persistence of sales shocks affects production in a model of the automobile industry that includes nonconvex production costs at the plant level.

IV. The Effect of Sales Persistence on Production Decisions

The crux of the IT hypothesis proposed by KMPQ rests on technological innovations in the production process that change the way production is scheduled and inventories are managed, *given a fixed sales process.* An implicit assumption in their presentation of evidence is that there should be a fixed relationship between the variance of production and sales as well as the covariance of sales with inventory investment in the absence of a structural change in the production scheduling. In this section, we show how a change in the sales process will modify the relationship between production, inventory and sales. This can be true with standard convex cost functions, but is even more likely when cost functions are nonconvex. The production smoothing model is analyzed in the appendix; here we examine the more realistic model with nonconvex costs. Specifically, we show that a decline in the persistence of sales shocks decreases the relative variance of production over sales even without IT effects on production scheduling.

[15] Other alternatives studied were different specifications of trends, such as quadratic and cubic.

A. Production Margins in Automobile Assembly Plants

In order to evaluate plant-level data, it is critical to first understand the institutional structure of the automobile industry, its labor union, and the mechanical processes involved on the assembly line. Managers of auto assembly plants have several margins at their disposal to meet production quotas, many of which involve altering the period of production as opposed to the rate of production.

Let Q_{it} represent the monthly output volume for plant i. Q_{it} is then a product of the following margins: (1) weeks in month t the plant is open; (2) days per week the plant operates; (3) the number of shifts working each day; (4) the length of each shift; and (5) the line speed in terms of jobs per hour. This is shown in Equation (5).

$$(5) \qquad Q_{it} = \frac{\text{weeks open}}{\text{month}} \times \frac{\text{days open}}{\text{week}} \times \frac{\text{shifts}}{\text{day}} \times \frac{\text{hours}}{\text{shift}} \times \frac{\text{jobs}}{\text{hour}}$$

As mentioned above, the institutional structure surrounding automobile manufacturing has various implications for the costs of using and for changing these margins. This information is found in the literature with the work of Aizcorbe (1990, 1992), who documents important implications in the labor contracts between the U.S. automobile manufacturers and the United Auto Workers, as well as in Bresnahan and Ramey (1994), and Hall (2000). A description of the production margins and the costs involved in using them can be summarized as follows:

> **Regular Hours:** Variation in regular (non-overtime) hours comes from closing the plant for either a whole or partial week. The cost to the plant of closing for a partial week is high, as the plant is required by union contract to pay short-week compensation to workers with at least one year of service. This is 85 % of a workers' regular pay for each hour less than 40 they did not work. Closing the plant for the entire week, on the other hand, entails laying workers off, in which case they receive 95% of their straight week pay through a combination of state Unemployment Insurance (UI) and Supplemental Unemployment Benefits (SUB). The state governments pay UI, and assembly plants contribute indirectly

according to their experience rating. SUBs are negotiated between the automakers and the UAW, and the plants support this fund on an employee-hour basis. Hall (2000) estimates that assembly plants pay 60 cents for each dollar distributed with UI and SUB.

Overtime Hours: Overtime hours occur either as an additional one or two hours added to regular eight-hour shifts, or in the form of an eight-hour Saturday shift. Employees who work either more than eight hours in one day, or more than five days in a week receive a 50% wage premium for the extra hours. Additionally, plants typically do not compensate for holidays by scheduling overtime, nor is a worker obligated to work overtime for more than three consecutive weeks. Thus overtime hours are intended to be temporary, and assembly plants cannot avoid hiring additional workers by using overtime permanently. Frequent discontinuous spells of overtime, however, are not uncommon.

Shifts: Most auto assembly plants operate with one or two shifts, though U.S. automakers began designing three-shift schedules in the early 1990s to increase capacity at certain facilities. The second shift pays a 5% shift premium and the third shift a 10% premium. Adding a shift involves a negotiation process with the UAW and an increase in the number of production and overhead workers on the payroll. Thus, adding a shift obliges the plant to increase their outlay of employee benefits and make additional SUB contributions. These benefits depend on the size of the payroll and not on whether these workers are actually on the job in a given week. A plant's long-run liabilities change substantially when new workers are hired.

Line Speeds: Line speed changes involve reorganizing the assembly line and redefining jobs, which imply a period of down-time before the redesigned line is complete. Workers do not simply assemble cars faster when line speeds increase. Instead, each shift hires more workers. There are natural upper-bounds on line speed changes that depend on the size of the paint facility, among other things. Line speeds will change when more or less of a car model is wanted, or when the plant switches from producing one model to another. The UAW typically becomes involved with changes in the line speed as well.[16]

B. Cost Function Simulation with Inventories

Not surprisingly, the nature of automobile assembly technology and the language written into the UAW contract imply several levels of production which are either prohibitively expensive or physically impossible to attain, and as a result it is perfectly rational for plant-level production decisions to yield output volumes that fluctuate much more than sales. Most notably, managers have the option of closing down an assembly line at week-long intervals, which is an option they exercise regularly.

This section takes the cost function for an automobile assembly plant described by Bresnahan and Ramey (1994) and Hall (2000) and investigates how properties of the sales processes feed into the cost-minimization objective function and determine production. In particular, the optimal production behavior from a sales process with persistent changes is compared with the production behavior from a sales process with relatively more transitory changes. The conclusion is that the relationship between the volatility of production and the volatility of sales is non-linear and depends upon the persistence of changes to sales.

1. The Automobile Assembly Plant Production Cost Environment

In order to minimize the discounted present value of short-run production costs while meeting vehicle sales, the plant manager schedules the workweek of the plant in week t by choosing the number of shifts (Sh_t) scheduled to report in week t, the number of days (D_t) the plant will open in week t, and the length of each shift (h_t). The line-speed (ls_t) in terms of vehicles per hour combines with the workweek variables to determine the weekly level of output as in Equation (6).

[16] Ford was negotiating with the UAW in the third quarter of 2001 in order to reduce production capacity for the Ford Explorer built at its Kentucky Truck facility, and the Ford Taurus / Mercury Sable, both built in Atlanta, GA and Chicago, IL. While Ford would prefer to pare shifts at all of these facilities, the UAW are urging instead that

(6) $$Q_t = Sh_t \times D_t \times h_t \times ls_t$$

The line-speed can be thought of as the plant's production function, as it is the flow of output made possible from employing capital (k_t) and the labor services of production workers (n_t). In this simulation we follow Hall's (2000) characterization of the line-speed as a Cobb-Douglas production function shown in Equation (7). The fact that a certain quantity of workers is necessary to achieve any positive level of output is reflected in the presence of overhead production workers ($\overline{n}_2$). The number of non-production workers ($\overline{n}_1$) employed by the plant does not affect the line-speed, but they are paid each week regardless of the plant's operating status.

(7) $$ls_t = k_t^{1-\gamma} \cdot \left(n_t - \overline{n}_2\right)^{\gamma}$$

The plant manager then solves a dynamic program built from this production identity and a series of weekly cost functions. The particular cost function used in this simulation is depicted in Equation (8).

(8)
$$\begin{aligned} c(h_t, D_t \mid Sh_t) = \;& \sum_{j=1}^{3} I(Sh_t \geq j) w_j \cdot D_t \cdot h_t \cdot n_t \\ &+ \max\left[0, 0.85 \sum_{j=1}^{3} I(Sh_t \geq j) w_j \left(40 - D_t \cdot h_t\right) n_t\right] \\ &+ \max\left[0, 0.5 \sum_{j=1}^{3} I(Sh_t \geq j) w_j \cdot D_t \left(h_t - 8\right) n_t\right] \\ &+ \max\left[0, 0.5 \sum_{j=1}^{3} I(Sh_t \geq j) w_j \left(D_t - 5\right) h_t \cdot n_t\right] \\ &+ \left[u \cdot w_1 \cdot 40 \cdot Sh_t \cdot n_t \left(D_t = 0\right)\right] + \delta \cdot I\left(D_t > 0\right) + 40 \cdot w_1 \cdot \overline{n}_1 \end{aligned}$$

The combination of production margins the plant manager chooses to obtain Q vehicles in week t will determine the value of each line in Equation (8). The first line contains the regular hours wage bill, while the second line captures the 85% short-week compensation that must be paid to

---

line speeds be reduced and the number of tag-relief workers be trimmed. (*The Wall Street Journal*, December 18, 2001) workers who spend more than 0 but less than 40 hours per week on the job. The third and fourth lines are the 50 % overtime premia charged to the plant when daily work hours exceed eight or the number of days scheduled exceeds five. The fifth line captures the costs associated with opening and closing the plant for the entire week, where the first term represents the cost of laying workers off, the second term is the fixed cost (δ) of opening the plant each week, and the third term represents the cost of salaried non-production workers. The meaning of each symbol is given below in Table 8.

In addition to the intra-period cost of scheduling a particular combination of shifts, days and hours each week, plant management must also consider the linkage between periods in minimizing total cost. The stock of inventory carried from period *t* to *t+1*, for example, is one channel through which past production decisions enter into the current environment. Equation (9) is the inventory identity used in this exercise, which simply states that the inventory level at the end of the current period, (I_t), is equal to last period's inventory plus current production, minus current sales (S_t). The stock of inventory is constrained to be above the allowable minimum, which is depicted in Equation (10).[17] Inventory holding enters the cost function in terms of its deviation from a desired level, which is determined by the target inventory-to-sales ratio $\boldsymbol{r}^*$.

(9) $$I_t = I_{t-1} + Q_t - S_t$$

(10) $$I_t \geq 10{,}010 \text{ for all } t$$

The second channel through which the plant's history affects current decisions involves the fixed adjustment costs the plant incurred when the production schedule is changed. These costs vary according to the intrinsic nature of each production margin, as well as to the opportunity cost of production that is lost while changes are made. For example, Bresnahan and Ramey (1993) present evidence that changing the line speed or the number of shifts working

[17] In a stochastic sales setting, this *no stock-out* condition is equivalent to requiring that the inventory stock after current period production but before current period sales is large enough to accommodate the largest possible realization of sales.

entails high adjustment costs, while other margins, such as scheduling overtime hours and closing the plant for week-long intervals, involve relatively low adjustment costs.[18]

The total cost incurred in week t is a combination of $c(h_t\ ,\ D_t|\ Sh_t)$, which includes the intra-period wage bill and the fixed cost of opening the plant each week, the inventory carryover charge governed by the parameter $\boldsymbol{\alpha_I}$, and the fixed adjustment cost of changing the number of shifts working, $\boldsymbol{\alpha_{Sh}}$. The inter-temporal cost function denoted as $C(h_t\ ,\ D_t\ ,\ Sh_{t+1}|\ Sh_t)$ is in Equation (11).

$$(11) \qquad C\ (h_t, D_t, Sh_{t+1} \mid I_{t-1}, Sh_t) = c(h_t, D_t \mid Sh_t) + \alpha_{Sh} \cdot I\left(\dot{S}h \neq 0\right) + \frac{1}{2}\alpha_I \cdot \left[I_t - E_t\left(s_{t+1}\right) \cdot r^*\right]^2$$

2. Dynamic Program Simulations

In order to understand the production behavior implied by the cost minimization problem under different sales conditions, this section constructs a representative assembly plant and simulates the dynamic program the plant manager solves in making short-run production decisions. In particular, it is of interest to compare the optimal production path chosen when changes to sales are persistent with the path chosen when changes to sales are transitory. In this sense, the first simulation mimics the automobile industry environment of the 1970s, while the second closely resembles the 1990s.

Due to the prevalence of discontinuities, non-convexities and non-differentiable points in the plant's weekly wage bill (as a function of units of output), the fixed-point theorems necessary to solve the Bellman equation analytically for a time-invariant optimal policy function are not satisfied. It is precisely the influence of these troublesome points that is of interest in this exercise. As an alternative, the plant's problem is structured as a series of 156 discrete weeks (3 years) over which the plant manager must choose the workweek variables from a discrete state space. Hall (2000) conducts a very similar exercise, and many of the technical methods are similar to his paper.

To make the dynamic program tractable for numerical solution, the decision variables are limited to the number of shifts hired for the next week (Sh_{t+1}), the number of days open in the

[18] These adjustment costs are distinct from the marginal costs associated with using each production margin, such as

current week (D_t), and the hours scheduled per shift per day (h_t). The grids that define the possible values for each of these choice variables are listed in Table 9.

The choice variable grids allow the plant manager a reasonable degree of flexibility in planning the workweek of the plant, but still maintain a state-space of reasonable dimension for a grid-search solution. In particular, it is possible to schedule overtime either through opening the plant for a sixth day or by scheduling the shift length to exceed 8 hours. Inventory adjustments can take the form of either a shift reduction or a week-long plant closure. A short-week is also possible by many combinations of production margins, where the length of the short workweek can range from 2 hours through 36 hours.

The line-speed in each period is taken as exogenous and is not a choice variable in this exercise. One can think of line-speed as a long-run margin, whose optimal value is determined by the encompassing profit-maximization problem the auto manufacturer has previously solved when it designed the plant and chose the type of vehicles it would produce in the current model year. The set of decision variables in this exercise then determine the workweek of the plant, given the plant's configuration and a realized path of sales.[19]

Sales evolve according to a first-order Markov process, where the realization in any given period may take one of nine possible values. Restricting the realizations of sales to a grid of modest size is necessary if sales are to be stochastically determined state-variable in a grid-search solution algorithm. These nine grid points along with the Markov transition-probability matrix $\boldsymbol{\chi}(s' \mid s)$ are parameterized using Tauchen's (1986) procedure. The unconditional mean of sales is set to the number of vehicles produced on two shifts using regular-time hours, thus both scenarios represent a plant that has correctly matched its capacity with sales. Mismatches between the planned capacity of a plant and its realized mean sales rate are also very important in the determination of production volatility, and the implications of such occurrences are the subject of Hall (2000).

Since the evidence we have presented above indicates that the biggest change to model sales between the periods 1972 – 1983 and 1984 – 2001 has been a reduction in the AR(1)

---

the wage premium paid for overtime work, or the 2nd shift premium paid for night work.

[19] Taking sales as given in the cost minimization problem does not imply that sales are exogenous to the firm. Rather, we are using a standard micro result that allows us to focus on only the cost minimization part of the overall profit maximization problem. Automakers often use vehicle-specific incentives to boost weak sales, however it is the assembly plants' objective to keep dealers stocked with vehicles in demand, and their relationship becomes strained when the company promotes unavailable vehicles.

estimate, this exercise consists of two simulations. Simulation #1 solves the plant's cost minimization problem with a persistent monthly sales process (AR(1) = 0.85). Simulation #2 maintains all of the properties of simulation #1, but reduces the AR(1) coefficient to 0.51. In accordance with analysis of the model-level sales process presented above, the variance of the innovations is raised in Simulation #2 so that the overall variance of the sales process in unchanged between simulations.

The parameter values used throughout this exercise come from several sources, including the labor contracts between automakers and their union, parameterizations of assembly plant cost functions from previous studies (notably Hall (2000)), and from the relatively stable inventory-to-sales ratio measured in industry data. Parameters that are more difficult to discern, such as the fixed cost of changing shifts and the marginal cost of deviating from desired inventories, were chosen so that the solution to the high-persistence version of the model roughly matched the production behavior observed among assembly plants in the 1972 – 1983 period. Table 10 lists the values of cost function, production function and sales parameters for easy reference.

The sequence of decisions and the arrival of information are as follows: At the beginning of week *t,* the plant receives its sales orders for week *t*, after which the managers schedule the workweek by choosing values for Sh_t, D_t and h_t subject to the relevant constraints. The orders are then filled and the new level of inventory is carried forward into the next period.

The inter-temporal cost minimization problem is then described as follows:

$$(12) \qquad \underset{\{Sh_{t+1}, D_t, h_t\}_{t=1}^{t=T}}{\text{MIN}} V_t = E_t \sum_{t=1}^{t=T} \beta^{t-1} \cdot C\left(h_t, D_t, Sh_{t+1} \mid I_{t-1}, Sh_t\right)$$

where $C\left(h_t, D_t, Sh_{t+1} \mid I_{t-1}, Sh_t\right)$ is defined as in Equation (11). The solution is subject to the constraints:

$$Q_t \geq s_t - I_{t-1} \quad \text{for all } t \in [0, T]$$

$$I_t \geq 0 \quad \text{for all } t \in [0, T],$$

where Q_t is defined as in Equation (6), and the evolution of final sales

$$E_t\left[s_{t+1} \mid s_t\right] = \sum_{s_{t+1}} \chi(s_{t+1} \mid s_t) \cdot s_{t+1}$$

and the initial inventory level

$$I_0 = r^* \cdot E(s) = 44{,}000\,.$$

The dynamic program is then solved backwards with value functions. 1000 different paths of sales shocks are generated with a length of 3 years, and in each case the realizations of sales are constructed for both the persistent (AR(1) = .85) and the more transitory (AR(1) = .51) sales scenarios.[20] The plant solves its weekly cost minimization problem, and then the optimal paths for the workweek variables as well as for production and inventory stock is determined. These solution paths are then aggregated to a monthly frequency and their volatility properties investigated.

In the high persistence case, the average standard deviation of sales across simulations is 4725 vehicles per month, while in the low persistence case it is 5073. This change is not statistically significant within the 1000 simulations, as it was engineered not to be. The average standard deviation of the optimal production paths, however, drops from 7761 vehicles per month in the first simulation to 5529 vehicles per month in the second simulation. While the average volatility of sales actually rose by 12% in the second simulation, the volatility of production fell by 28%. Accordingly, the ratio of the standard deviation of production over sales falls from 1.7 to 1.1. These results are summarized in Table 11, which also lists the empirical 95% confidence intervals for the standard deviation point estimates across the 1000 simulations.

Figure 4 examines the fall in the ratio of the standard deviation of production over sales within each simulation. This is done by plotting the reduction in this ratio as a histogram across all 1000 simulations. On average, the decline in the volatility ratio was 0.6, and the histogram assesses how representative this behavior was across the simulations. The results indicate that the ratio of the volatility of production over sales declined in 99.4% of the simulations.

In order to assess from which sources the change in volatility behavior originates, Table 12 shows the weekly frequency with which various changes to production were made. Weeklong shutdowns for inventory adjustment, for example, occur in 11.4% of the weeks when the AR(1) parameter is set to 0.85, and that figure grows slightly to 11.5% when the AR(1) parameter is lowered to 0.51. Shift reductions, alternatively, are much more common when the persistence of sales changes is high. Shift changes occur in 4.9% of the weeks in the high

[20] The Markov process that generates weekly sales was calibrated so that, on average, monthly sales exhibited the desired first-order autocorrelation.

persistence case, and only in 0.7% of the weeks in the low persistence case. Short-weeks in both scenarios are quite rare, occurring roughly 0.5% of the time in the first scenario, and close to never in the second. The frequency of the use of overtime hours increases from 10.6% in the high persistence scenario to 18.3% in the low persistence simulation.

The conclusion of this simulation exercise is that the variance of output is highly impacted by the nature of the sales process. When changes in sales are believed to be persistent, the plant often responds by adding and paring shifts. Alternatively, when changes in sales are transitory, the plant is more likely to respond with temporary measures, such as scheduling overtime hours. Thus, if a given change in the variance of sales stems from a reduction in the persistence of the shocks to sales, this can lead to a large decline in the variance of output relative to sales. The result is reached in a simplified production model with certain nonconvex costs, but it features no changes to inventory management parameters or improvements in the flow of future sales information.

The next section leaves the simulated world behind and searches for changes in production scheduling behavior within actual assembly plants.

V. Evidence on Production Scheduling from Plant-Level Data

Our hypothesis states that the change in the nature of the sales process is decreasing the need to use the nonconvex margins that contribute so much to the volatility of production. In order to measure this change, we compare the use of the various production margins by assembly plants in the U.S. and Canada between 1972 and 2001, and determine whether there are any obvious changes in the way production is scheduled. For this purpose, a dataset has been constructed from industry trade publications that track production behavior at the pant-level on a weekly basis over the two time periods: 1972 – 1983 and 1990 – 2001. Bresnahan and Ramey (1994) collected the data covering the 50 domestic car assembly plants operating in the period 1972 – 1983, and this data set has been significantly extended to include all 103 car and light truck assembly plants operating within the two periods listed above.[21]

[21] Data for AMC car plants prior to 1983 were not available, and certain heavy-truck and specialty vehicle facilities were excluded, such as the AMC General military vehicle plant, and GMAD Truck & Coach in Pontiac, MI, which primarily produces buses.

The majority of the data set was collected by reading the weekly production articles in *Automotive News*, which report the following variables for all domestic assembly plants: (1) the number of regular hours the plant worked; (2) the number of scheduled overtime hours; (3) the number of shifts operating at each plant; and (4) the number of days per week the plant is closed for (a) union holidays, (b) inventory adjustments, (c) supply disruptions, and (d) model changeovers. Observations on the line speed posted on each assembly line were collected from the *Wards Automotive Yearbook.*

Table 13 examines how often each margin of production (i.e. plant closures, changes in shift length, the number of shifts working and line speed) was manipulated during the two periods. The frequency of margin use among all 103 assembly plants is summarized as a weighted average, based on each plant's contribution to total production during the period examined. Several comparisons between the periods are noteworthy. First, plants shut down at roughly the same frequency in both periods. The weeklong closures are of particular interest, as these include the inventory adjustments and model changeovers that directly relate to production decisions. The frequency of weeklong shutdowns drops from 12.4% in the early period to 11.3% in the late period, though once holidays are excluded the size of the fall is enhanced somewhat. Second, the frequency of weeks in which at least four hours of overtime are scheduled has more than doubled between the periods, rising from 14.5% to 30.3%. Finally, while changes in the line speed occurred with roughly the same frequency between the periods, changes in the number of shifts occur in 0.6% of the weeks in the early sample, and occur in only 0.1% of the weeks in the late sample. This implies that the average assembly plant either adds or pares a shift 3.75 times during the early period, but does so less than once (0.626 times) in the late period.

Table 14 isolates the plant shutdown margin and decomposes its use between the first and second periods. It shows the percent of days closed by reason across all plants that were not mothballed, on extended closure, or permanently removed from service. Thus it considers only temporary closures as opposed to exit and entry. Inventory adjustments and model changeovers each close plants for a fewer number of days in the late period than in the early period, though the increase in inventory adjustments is very minor. The number of holidays appears to have increased, and the frequency of supply disruptions, such as union strikes, parts shortages and natural disasters, is relatively unchanged.

The drop in the average downtime for model changeovers from 2.3 days per year to 1.3 days per year is particularly interesting, as this is margin through which improvements in manufacturing technology would be visible.[22] There is indeed evidence that model changeover technology has advanced over time, as the industry introduced the *week-end model changeover* in the 1970s, and the *rolling model changeover* in the 1990s. However, the primary means of managing inventories in the automobile industry, the inventory adjustment, has not changed much despite the advances in information technology.

The interpretation of these results comes with several caveats. First, the distinction between inventory adjustments, model changeovers and holidays become blurred during the winter and summer quarters. Extended Christmas holidays often mask inventory adjustments,[23] and model changeovers often take place during a summer vacation, or are much longer than the technology necessitates during periods with low demand.[24]

Finally, Table 15 shows the importance of each production margin for the variance of output. This differs from the earlier analysis of margin frequency, as some margins cause a larger change in production when they are used relative to other margins. Overtime hours, for example, typically boost weekly production by 25%, while adding a second shift will double the weekly production. To do this analysis, we construct an artificial output measure, holding each margin constant at some base level. We determine the impact of a margin on the variance of output by calculating the difference in the variance of output and constructed output. The numbers do not add to 100 because of nonlinearities and covariance terms.[25]

Table 15 displays three noticeable changes over the two periods. First, model changeovers contribute less to the variance of output during the second period. Their impact on variance falls from 31% to 21%. Second, the use of overtime hours contributes more than twice as much to the variance in the second period as it did during the first period, climbing from a

[22] These numbers are the weighted average across plants. When model changeovers do occur, they almost always occupy entire weeks.

[23] Christmas 1982 lasted until almost February 1983 in many plants!

[24] An interesting extension of this analysis could evaluate the role of model changeovers in production variance during different stages of the business cycle instead of over two discrete pieces of time as we have done. This would be similar to the Cooper and Haltiwanger (1993) study of machine replacement.

[25] See Bresnahan and Ramey (1994) for a more detailed explanation of the method.

5.8% contribution to a 13.4% contribution. Third, changes to the number of shifts at individual plants contribute half as much during the second period as the first, falling from a 24.3% contribution to 12.4% in the second period.

Thus, the two nonconvex margins that lead to so much variance of output – model changeovers and shifts – are a less important component of the variance of output in the second period than in the first. Furthermore, overtime hours, which are the classic convex margin of adjusting production, are more than twice as important during the second period than during the first.[26]

The increase in overtime hours is corroborated by other studies as well. A more intensive use of overtime hours has been observed in all of U.S. manufacturing during the 1990s and has attracted the attention of researchers at the Bureau of Labor Statistics. Hetrick (2000), for example, has concluded that U.S. manufacturing establishments rebuilt their production levels in the recovery following the 1991 recession by using overtime hours much more intensively and for a longer horizon than is typical in U.S. recoveries. This came at the expense of hiring new workers. By 1997 the level of average weekly overtime in U.S. manufacturing hit a record-setting level of 4.9 hours per worker, while the increase in the number of full-time manufacturing workers was extraordinarily weak at only 17% of the post-war recovery average. By March 1997, the number of full-time workers employed in U.S. manufacturing was still 700,000 lower than its pre-recession peak in March 1989.

In the automobile industry, the average use of overtime doubled between March 1991 and January 1998. Figure 5 charts average overtime hours per worker in motor vehicle production. In addition, the number of full-time equivalent employees lost to the rise in overtime between 1991 and 1998 reached 107,000 within *transportation equipment*, which represents one-fifth of the full-time equivalents lost to all of U.S. manufacturing.[27] This is largely consistent with evidence discovered in the plant-level production behavior evaluated above.

[26] We initially worried that the increase in overtime hours was an artifact of better reporting in *Automotive News*, which is the source of the plant level data. The BLS data on overtime use in the automobile industry show a 30 percent increase in average overtime used from the early period to the later period. Thus, we believe that the increase in overtime is a real phenomenon rather than a change in the reporting by *Automotive News*.

[27] The upside to this phenomenon is that during the 1998 financial crises in Asia, Russia and Brazil, export-sensitive industries responded largely through the elimination of overtime hours rather than employees.

VI. Conclusions

The overview of the automobile industry and analysis conducted using plant level data has highlighted several interesting facts that should serve to increase our understanding of the decline in the variance of GDP. The automobile industry experienced declines in production volatility around the same time as the rest of the economy. The declines in the automobile industry were even more dramatic than the declines overall. At least for the case of cars, the variance of production declined more than the variance of sales.

We presented evidence that the change in production volatility may be linked to changes in the sales process. We found that changes in the process driving sales appear to be an important part of the changes in the automobile industry. In contrast to the 1970s and early 1980s, a time when volatile and highly persistent movements in sales beset the automobile industry, the 1990s featured much more transient shocks to sales. We then showed how a change in the persistence of the sales shocks could lead to a proportionately larger decline in production volatility over sales volatility.

Plant-level evidence indicates that firms have responded to these changes in the sales process by reducing their use of nonconvex lumpy margins, such as shift changes, and have begun to use the classic convex margin of overtime hours much more intensively. It is likely that this induced switch is a prime cause of the sharp decline in production volatility.

The next natural step in this line of research is to determine where this change in sales originated. In our opinion, likely candidates include (1) changes in vehicle pricing, which in the case of automobiles involves consumer and dealer sales incentives at the model-level, (2) changes in aggregate demand shocks, which may be tied to monetary policy, and (3) the increasing dispersion of motor vehicle demand among a wider selection of vehicle models in the 1990s than in earlier decades.

References

Aizcorbe, A., "Experience Rating, Layoffs and Unions: A Look at U.S. Assembly Plant Layoffs," Manuscript, Bureau of Labor Statistics, 1990.

Aizcorbe, A., "Procyclical Labour Productivity, Increasing Returns to Labour and Labour Hoarding in Car Assembly Plant Employment," *Economic Journal,* 102, 1992, 860 - 873.

Ahmed, S., Levin, A. and Wilson, B. A., "Recent U.S. Macroeconomic Stability: Good Luck, Good Policies, or Good Practices?" The Board of Governors of the Federal Reserve System, January 200l.

Andrews, D. W. K., "Tests for Parameter Instability and Stuctural Change with Unknown Change Point," *Econometrica,* 61(4), 1993, 821 – 856.

Andrews, D. W. K. and Ploberger, W., "Optimal Tests when a Nuisance Parameter is Present Only Under the Alternative," *Econometrica,* 62(6), 1994, 1383 – 1414.

Automotive News, Detroit: Crain Automotive Group.

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

Blanchard, O.J. and Simon, J., "The Long and Large Decline in U.S. Output Volatility," *Brookings Papers on Economic Activity*, 2001:1, 135 – 164.

Boivin, J. and Giannoni, M.P., "Has Monetary Policy Become More Effective?" Working Paper, Columbia University, December 2002.

Bresnahan, T. F. and Ramey, V. A., "Segment Shifts and Capacity Utilization in the U.S. Automobile Industry," *American Economic Review Papers and Proceedings*, 83(2), May 1993, 213 – 218.

-----------------------, "Output Fluctuations at the Plant Level," *The Quarterly Journal of Economics,* August 1994, 593 – 624.

Caballero, R.J., Engel, E.M.R., and Haltiwanger, J., "Aggregate Employment Dynamics: Building from Microeconomic Evidence," *American Economic Review* 87(1), March 1997, 115 – 137.

Clarida, R., Gali, J., and Gertler, M., "Monetary Policy Rules and Macroeconomic Stability: Evidence and some Theory," *Quarterly Journal of Economics*, 115, February 2000, 147 – 180.

Cooper, R. and Haltiwanger, J., "The Aggregate Implications of Machine Replacement: Theory and Evidence," *American Economic Review* 83(3), June 1993, 360 – 382.

------------------, "Automobiles and the Naitonal Industrial Recovery Act: Evidence on Industry Complimentarities," *Quarterly Journal of Economics* 108(4), November 1993, 1043 – 1072.

Golob, J. E., "Post-1984 Inventories Revitalize the Production-Smoothing Model," Working

Paper, University of Missouri - Kansas City, 2000.

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

Hansen, B. E., "Approximate Asymptotic P-values for Structural Change Tests," *Journal of Business & Economic Statistics,* 15(1), 1994, 60 – 67.

Hetrick, R.L., “Analyzing the Recent upward Surge in Overtime Hours,” Monthly Labor Review, February 2000, 30 – 33.

Kahn, J.A., “Why is Production More Volatile than Sales? Theory and Evidence on the Stockout-Avoidance Motive for Inventory-Holding,” *Quarterly Journal of Economics*, 107, May 1992, 481 – 510.

Kahn, J. A., McConnell, M. M. and Perez-Quiros, G., "On the Causes of the Increased Stability of the U.S. Economy," FRNBY Economic Policy Review, May 2002, 183 – 202.

Kashyap, A.K. and Wilcox, D.W., “Production and Inventory Control at the General Motors Corporation During the 1920’s and 1930’s,” *American Economic Review* 83(3), June 1993, 383 – 401.

Kim, C. and Nelson, C. R., "Has the U.S. Economy Become More Stable? A Bayesian Approach Based on a Markov-Switching Model of the Business Cycle," The *Review of Economics and Statistics,* 81(4), November 1999, 608 – 616.

McConnell, M.M., and Perez-Quiros, G., "Output Fluctuations in the United States: What has Changed Since the Early 1980s?" *American Economic Review*, 90(5), December 2000, 1464 – 76.

Olney, M.L., “Credit as a Production Smoothing Device: The Case of Automobiles, 1913 – 1938,” *Journal of Economic History*, 49(2), June 1989, 377 - 391.

Pivetta, F. and Reis, R., “The Persistence of Inflation in the United States,” Working Paper, Harvard University, March 2002.

Ramey, V. A., "Nonconvex Costs and the Behavior of Inventories," *Journal of Political Economy,* 99, 1991, 306 – 334.

Ramey, V. A. and West, K. D., "Inventories," in *Handbook of Macroeconomics,* eds. John D. Taylor and Micheal Woodford, Elsevier Science Press.

Stock, J.H. and Watson, M.W., “Has the Business Cycle Changed and Why?” *NBER Working Paper Series* #9127, September 2002.

Tauchen, G., “Finite State Markov-Chain Approximations to Univariate and Vector Autoregressions,” *Economics Letters*, 20(1986), 177 – 181.

Wards Automotive Yearbook, Detroit: Wards Communications.

Appendix A: Production Decisions in a Standard Production-smoothing Model

In a standard production-smoothing model of inventories, a change in the first-order autocorrelation of sales has an impact on the relative variances of production and sales. Consider the problem of a firm or plant that seeks to minimize production and inventory costs given a particular process for sales:

Minimize $V = E_0 \sum_{t=0}^{\infty} \beta^t \left[\frac{1}{2}\alpha_1 Y_t^2 + \frac{1}{2}\alpha_2 (I_t - \alpha_3 S_{t+1})^2 \right]$ subject to $Y_t = S_t + I_t - I_{t-1}$

where E_0 denotes the expectation conditional on information in period 0, β is a discount factor between 0 and 1, Y_t is production during period t, I_t is the stock of inventories at the end of period t, and S_t is sales during period t. The firm schedules production given a process for sales.[28] In doing so, the firm weighs two different types of costs: increasing marginal costs of production and the cost of allowing inventories to deviate from a desired ratio relative to sales. For more detail on the motivation for this cost function, see Ramey and West (1999).

Suppose that sales are given by an AR(1) process:

(A1) $S_t = \rho S_{t-1} + \varepsilon_t, \qquad 0 < \rho < 1, \qquad \varepsilon_t \ i.i.d$

For this sales process, the optimal rule for production is given by:

(A2) $Y_t = -(1-\lambda) I_{t-1} + \phi S_t$

where $\lambda = \frac{1}{2}\left[1 + \frac{1}{\beta} + \frac{\alpha_2}{\alpha_1 \beta} - \sqrt{\left(1 + \frac{1}{\beta} + \frac{\alpha_2}{\alpha_1 \beta}\right)^2 - \frac{4}{\beta}} \right]$

and $\phi = \dfrac{1 - \lambda + \beta\rho\lambda \frac{\alpha_2 \alpha_3}{\alpha_1}}{1 - \beta\rho\lambda}$

[28] Taking sales as given in the cost minimization problem does not imply that sales are exogenous to the firm. Rather, we are using a standard micro result that allows us to focus on only the cost minimization part of the overall profit maximization problem.

As long as the α's are nonnegative, λ will be positive and less than unity, and ϕ will be positive. It is only the ratio of α_2/α_1, which measures the relative cost of deviating from desired inventories to the slope of the marginal cost curve that matters rather than the levels of each parameter.

Note that while λ depends on neither α_3 (which gives the firm's preferred inventory-sales ratio) nor ρ (the persistence parameter for sales shocks), ϕ is increasing in both of these parameters. If a firm prefers to maintain a higher ratio of inventories to sales (higher α_3), then a given increase in sales will lead the firm to produce more. How much it responds depends, of course, on the cost of deviating from the desired ratio relative to the cost of changing production.

Why does ρ increase the response of production to an increase in current sales? If ρ is high, then the firm anticipates that sales will be above normal for a long time, so it will raise production in order not to allow inventories to deviate from their desired level. On the other hand, if ρ is low, the sales shock is thought to be temporary, so there is no point in building up inventories in anticipation of high future sales. In the numerator of ϕ, ρ and α_3 multiply each other, so there is also an interaction effect.

This model shows that the relative variances of production and sales and the covariance of sales with inventory investment are not independent of the parameters of the sales process. Using the optimal rule for production and inventories, the inventory identity and the process for sales, one can show that:

(A3) $$\frac{Var(Y)}{Var(S)} = 1 - \frac{2(1-\rho)(1-\phi)}{1-\lambda\rho} + \frac{2(1-\phi)^2[1-\rho(1-\lambda^2)]}{(1+\lambda)(1-\lambda\rho)}$$

and

(A4) $$Cov(S, \Delta I) = \frac{-(1-\phi)\sigma_\varepsilon^2}{(1-\lambda\rho)(1+\rho)}.$$

As these formulas show, both of these variables are complicated functions of the persistence of sales. (Recall that the ϕ parameter also depends on ρ.)

It is difficult to sign the derivatives of these functions with respect to ρ, so we studied several simulations based on the automobile industry. The average inventory-sales ratio has been quite stable at 2.5 months since 1965, so we preset $\alpha_3 = 2.5$. We also set $\beta = 0.997$ as the monthly discount rate. We then studied how the variance of production relative to sales varied with changes in ρ. We found that most parameter values imply that an increase in ρ leads to an increase in the variance of production relative to the variance of sales. In particular, for $\alpha_2/\alpha_1 < 1$ and $\rho \leq 0.85$, $Var(Y)/Var(S)$ is monotonically increasing in the value of ρ. After $\rho > 0.85$, the ratio falls a little. For $\alpha_2/\alpha_1 = 10$, $Var(Y)/Var(S)$ is increasing in ρ for $\rho < 0.66$. Thus, as long as ρ is not too high and the cost of deviating from inventory-sales ratios relative to the slope of marginal cost is not too high, an increase in the persistence of sales shocks raises $Var(Y)$ relative to the $Var(S)$.

The results are even clearer for the covariance between sales and inventory investment. For every parameter combination we studied, the covariance was monotonically increasing in the value of ρ.

To illustrate how a change in ρ might explain the changes in variances and covariances we discussed in previous sections, suppose $\alpha_2/\alpha_1 = 0.5$. According to the formulas above, if ρ falls from 0.8 to 0.3 (as we saw in the model-level data), $Var(Y)/Var(S)$ falls from 1.83 to 0.73 and $Cov(S,\Delta I)$ falls from 0.61 to –0.17. Thus, the change in sales persistence alone, with no change in the "structure" of production scheduling, is enough to make the variance of production fall more than the variance of sales and for the covariance to change signs.

Appendix B: U.S. Domestic Assembly Plants and Vehicle Models

Domestic Car and Light-Truck Assembly Plants

Chrysler and American Motors

Belvidere, IL

Bramalea, Ont.

Detroit, MI (Conner Ave.)

Detroit, MI (Jefferson Ave.)

Detroit, MI (Jefferson North)

Detroit, MI (Lynch Rd)

Detroit, MI (New Mac Ave.)

Hamtramck, MI

Kenosha, WI (East)

Kenosha, WI (West)

Missouri Truck (St. Louis)

Newark, DE

St. Louis, MO #1

St. Louis, MO #2

Sterling Heights, MI

Toledo, OH (Parkway Ave.)

Toledo II, OH

Warren, MI (Dodge City)

Windsor, Ont. (Pillette Road)

Windsor, Ont. (Tecumseh Road)

Ford Motor Company

Atlanta, GA

Avon Lake, OH

Chicago, IL

Dearborn, MI

Edison, NJ

Kansas City, MO

Kansas City, MO

Kentucky Truck (Louisville)

Loraine, OH

Los Angeles, CA

Louisville, KY

Mahwah, NJ

Michigan Truck (Wayne, MI)

Norfolk, VA

Norfolk, VA

Oakville, Ont line 1

Oakville, Ont line 2

Ontario truck (Oakville, Ont.)

San Jose, CA

St. Louis, MO

St. Thomas, Ont.

Twin Cities, MN

Wayne, MI

Wixom, MI

General Motors

Arlington, TX

Baltimore, MD

Bowling Green, KY

Detroit, MI #1

Detroit, MI #2

Doraville, GA

Fairfax, KS

Flint, MI #1

Flint, MI #2

Flint, MI #4 (Buick City)

Janesville, WI

Lake Orion Twnshp, MI

Lakewood, GA

Lansing, MI A / M

Lansing, MI B / C

Lansing, MI Grand River

Lansing, MI Reatte Craft Ctre.

Leeds, MO

Linden, NJ

Lordstown, OH

Pontiac, MI #8

Pontiac, MI [Central]

Pontiac, MI [East]

Pontiac, MI [West]

Scarborough, Ont.

Shreveport, LA

Southgate

Spring Hill, TN

St. Louis, MO Chevrolet

St. Louis, MO Corvette

Flint, MI #40	Moraine, OH	Ste. Therese, Que
Fort Wayne, IN	Norwood, OH	Tarrytown, NY
Framingham, MA	Oklahoma City, OK	Van Nuys, CA
Fremont, CA	Oshawa, Ont.	Wentzville, MO
Hamtramck Allante line	Pontiac, MI #1	Willow Run, MI
Hamtramck, MI	Pontiac, MI #2	Wilmington, DE

Domestic Car and Light Truck Model Lines

Chrysler / American Motors

Acclaim	Chrysler E Class	Eagle	Lancer	PT Cruiser	Talon
Alliance	Chrysler LHS	Encore	Laser	Ram Charger	Town & Country
Ambassador	Cirrus	Fifth Avenue	Laser - Plymouth	Ram Pickup	Trail Duster
Aries	Comanche	Fury	Lebaron	Ram Van	Turismo
Aspen	Concorde	Grand Cherokee	Liberty	Rampage	Valiant
Avenger	Cordoba	Grand Fury	Matador	Reliant	Viper
Barracuda	Dakota	Grand Wagoneer	Mirada	Royal Monaco	Vision
Breeze	Dart	Grand Wagoneer ZJ	Monaco	Scamp	Volare
Caravan	Daytona	Gremlin	New Yorker	Sebring	Voyager
Caravelle	Diplomat	GTS	Omni	Shadow	Voyager - Chrysler
Challenger	Dodge 024	Horizon	Pacer	Spirit	Wagoneer XJ
Charger	Dodge 400	Hornet	Phantom	Sport	Wrangler
Cherokee XJ	Dodge 600	Imperial	Plymouth Neon	St. Regis	
Chrysler	Dodge Neon	Intrepid	Premier	Stratus	
Chrysler 300M	Durango	Javelin	Prowler Chrysler	Summit	
	Dynasty	Jeep Pickup	Prowler Plym	Sundance	

Ford Motor Company

Aerostar	Cougar	F Series Pickup	LTD	Mustang	Topaz
Blackwood	Crown Victoria	Fairmont	LTD II	Mystique	Torino
Bobcat	Econoline	Falcon	Lynx	Navigator	Tracer
Bronco	Elite	Focus	Mark IV	Pinto	Versailles
Bronco II	Escape	Ford	Marquis	Probe	Villager
Capri	Escort	Granada	Maverick	Ranger	Windstar
Club Wagon	Excursion	Grand Marquis	Mercury	Sable	Zephyr
Comet	EXP	Lincoln	Monarch	Taurus	
Continental	Expedition	Lincoln LS	Montego	Tempo	
Contour	Explorer	LN7	Mountaineer	Thunderbird	

General Motors

Pontiac 6000	Fiero	Sportvan	CTS	Oldsmobile 88	Sonoma
Astro	Firebird	Starfire	Custom Cruiser	Oldsmobile 98	Suburban
Buick	Firenza	Sunbird	El Camino	Parisienne	Suburban
Cadillac	Grand Am	Pontiac T1000	Envoy	Park Avenue	Sunfire
Calais	Grand Prix	Toronado	Escalade	Prizm	Supreme
Camaro	J2000	Vega	Escalade EXT	Reatta	Tahoe
Cavalier	Lesabre	Achieva	EV1	Rendezvous	Tracker
Celebrity	LeMans	Alero	Express Van	Road Master	Trailblazer
Century	Monte Carlo	Allante	Fleetwood Brougham	S 10 Pickup	Trans Sport
Chevelle	Monza	Aurora	Fleetwood Deville	S 15 Pickup	Vandura/Rally

Chevrolet	Nova	Avalanche	Impala	S Blazer	Venture
Chevette	Oldsmobile	Aztek	Intrigue	Safari Pontiac	Vibe
Ciera	Omega	Blazer	Ion	Safari Van	Vue
Cimarron	Phoenix	Bonneville	Jimmy	Saturn LS	Yukon
Corvette	Pontiac	Bravada	S Jimmy	Saturn S	Yukon XL
Cutlass	Regal	Caballero	Lumina	Savana	
Deville	Riviera	Caprice	Lumina Van	Sierra	
Electra	Seville	Chevy CK Pickup	Malibu	Silhouette	
Eldorado	Skyhawk	Citation	Metro	Silverado	
	Skylark	Corsica Beretta	Montana	Somerset	

Table 1: The Standard Deviation of Output Growth*
(Annualized growth rates)

	1967:1 - 1983:4	1984:1 - 2002:3	Percent Change in volatility
Chained $1996			
GDP	4.5	2.2	-52
Durable Goods	17.1	8.0	-53
Motor Vehicles	52.2	21.0	-60
Current $			
GDP	4.4	2.3	-50
GDP excluding Durable Goods	3.4	2.1	-37
Durable Goods	15.7	8.1	-49
Durable Goods (Excluding Motor Vehicles)	13.5	8.1	-40
Motor Vehicles	51.4	20.6	-60

* Based on quarterly NIPA data.

Table 2: Decomposition of Durable and Motor Vehicle Volatility
(Chained 1996 dollars)

	1967:1 – 1983:4	**1984:1 – 2002:3**
Durable Goods		
$Var(\hat{y})$	18.3	4.3
$Var(\hat{s})$	6.7	4.2
$Var(\Delta\hat{i})$	8.4	4.5
$Cov(\hat{s},\Delta\hat{i})$	1.6	-2.2
$\frac{Var(\hat{y})}{Var(\hat{s})}$	2.7	1.0
Motor Vehicles		
$Var(\hat{y})$	214.7	27.2
$Var(\hat{s})$	104.4	40.0
$Var(\Delta\hat{i})$	88.6	45.5
$Cov(\hat{s},\Delta\hat{i})$	10.8	-29.2
$\frac{Var(\hat{y})}{Var(\hat{s})}$	2.1	0.6

$\hat{y}=\frac{Y_t-Y_{t-1}}{Y_{t-1}}$, $\hat{s}=\frac{S_t-S_{t-1}}{Y_{t-1}}$, $\Delta\hat{i}=\hat{y}-\hat{s}$, where Y = production, S = sales, and I = inventories.

Table 3: Structural Break Test Results on the Conditional Variance of Monthly Unit Vehicle Builds

	Sup F	Exp F	Ave F	Break Date
U.S. Vehicle Production (p-value)	9.24 (0.00)	3.06 (0.08)	5.17 (0.02)	February, 1983

Table 4: Motor Vehicle Volatility

A. Cars

	Not Seasonally Adjusted		Seasonally Adjusted	
	1967:1-1983:12	1984:1-2002:8	1967:1-1983:12	1984:1-2002:8
$Var(Y)$	6.31	2.98	3.86	1.24
$Var(S)$	3.94	3.12	2.81	1.78
$Var(\Delta I)$	2.75	3.32	1.42	1.33
$Cov(S,\Delta I)$	-0.02	-1.13	-0.09	-0.44
$\frac{Var(Y)}{Var(S)}$	1.59	0.96	1.38	0.70

B. Trucks

	Not Seasonally Adjusted		Seasonally Adjusted	
	1967:1-1983:12	1984:1-2002:8	1967:1-1983:12	1984:1-2002:8
$Var(Y)$	10.47	3.14	9.35	1.51
$Var(S)$	9.39	2.10	8.43	1.23
$Var(\Delta I)$	2.79	2.84	1.91	1.39
$Cov(S,\Delta I)$	-0.12	-0.62	0.18	-0.24
$\frac{Var(Y)}{Var(S)}$	1.11	1.49	1.11	1.22

All variables were normalized by the exponential of a fitted trend to log production.

Data were seasonally adjusted using dummy variables since the government's seasonal adjustment factors for production do not extend back very far.

Table 5: Estimates of Aggregate Automobile Sales Process
(Standard errors in parenthesis)

Coefficient on:	**Cars**	**Light Trucks**	**Motor Vehicles**
Constant	0.334** (.116)	0.0348 (.0191)	0.275** (.107)
Constant $\cdot D_t$	0.712** (.246)	0.271** (.094)	0.633** (.255)
AR(1)	0.848** (.050)	0.934** (.031)	0.884** (.044)
AR(1) $\cdot D_t$	-.339** (.125)	-.232** (.085)	-.263** (.112)
Trend	-.0002 (.00014)	0.00016 (.00016)	-.00007 (.0001)
Trend $\cdot D_t$	-.0002 (.0002)	0.0011** (.0004)	0.00056** (.00022)
σ^2	0.0070** (.00117)	0.0090** (.0012)	0.0066** (.0011)
$\sigma^2 \cdot D_t$	-.00035 (.0016)	-.0042** (.0014)	-.120 (.142)
Log likelihood	458.386	471.4	488.0

- Standard errors were computed using Eicker-White methods.
- * denotes significant at the 5 % level, ** denotes significant at the 1 % level.
- N = 426
- $D_t = 0$ for $t \leq 1983{:}12$; $D_t = 1$ for $t \geq 1984{:}1$

Table 6: Alternative Measures of Sales Persistence: U.S. Domestic Cars, Trucks and Motor Vehicles 1967:1 – 2002:9

	Cars		**Trucks**		**Motor Vehicles**	
Optimal Lags	5		3		4	
	1967:1 – 1983:12	1984:1 – 2002:9	1967:1 – 1983:12	1984:1 – 2002:9	1967:1 – 1983:12	1984:1 – 2002:9
Largest Autoregressive Root	.900	.797	.960	.900	.940	.920
Sum of AR Terms	.904	.790	.947	.832	.915	.825
Half-life (No. of Months)	4	2	11	3	4	2

Optimal lag length chosen by BIC over entire sample period

Table 7: Estimates of the First Order Autocorrelation and the Standard Deviation of the Residuals from U.S. Domestic Car and Light-Truck Sales*

(Weighted Average across Models)

Estimates	1965:1 – 1983:12	1984:1 – 2002:9
AR(1) coefficient	0.845	0.273
Standard deviation of residuals	0.237	0.418

- Estimates from first-order autoregressions on monthly sales in log physical units for each vehicle model

Table 8: Weekly Cost Function Variables and their Meaning in Assembly Plant Simulation

Variable	Meaning
Q_t	Production in week *t* in physical units
Ls_t	Line-speed in units per hour
n_t	Number of production workers per shift
$\overline{n}_1$	Number of non-production workers at the plant
$\overline{n}_2$	Number of overhead production workers per shift
I	Indicator variable (= *1* if expression is true)
I_t	Inventory carried from week *t* to *t+1*
Workweek of the Plant	
Sh_t	Shifts working in week *t*
D_t	Days per week
h_t	Hours per shift per day
Parameters	
u	Percentage of wage paid to laid off workers
w_1 , w_2	Wage per hour for 1st and 2nd shift workers
δ	Fixed cost of opening plant for in week *t*
r^*	Target inventory-to-sales ratio
α_I	Cost per unit of deviating from desired inventory
α_{Sh}	Fixed cost of adjusting the number of shifts working

Table 9: Choice Variable and State Variable Grids in Assembly Plant Simulation

State Variable at *t*	Choice Variable at *t*	Allowable Values
Sh_t	Sh_{t+1}	{ 1 , 2 , 3 }
--	D_t	{ 0 , 1 , 2 , 3 , 4 , 5 , 6 }
--	h_t	{ 0 , 2 , 4 , 6 , 8 , 10 }
ls_t	--	{ 55 }
s_t	--	9 grid points equi-spaced between 990 and 7810
I_{t-1}	--	910 grid points between 10,010 and 60,010 at 55 unit intervals

Table 10: Parameter Values Used in Assembly Plant Simulation

Parameter	Value
Cost Function Parameters	
β	0.999 per week (*= 5% annual discount rate*)
w_1	\$18.00/hr
w_2	\$18.90/hour
w_3	\$19.80/hour
u	65%
$\overline{n}_1$	364 workers (non-production)
α_I	0.002
δ	\$400,000/week open
r^*	60 days-supply
Production Function Parameter	
$\overline{n}_2$	658 workers (overhead)
γ	0.62
Monthly Sales Process Parameters	
AR(1) in simulation #1	0.85
AR(1) in simulation #2	0.51

Table 11: Average Standard Deviations of Monthly Production and Sales over 1000 Simulations

Sales Path	*σ(S)*	*σ(Q)*	*σ(Q)/σ(S)*	*Ave. % Drop in σ(S)*	*Ave. % Drop in σ(Q)*
AR(1) = .85	**4725** (2752 , 7219)	**7761** (5552 , 10049)	1.7	--	--
AR(1) = .51	**5073** (3797 , 6487)	**5529** (3853 , 7225)	1.1	-12.5%	27.7%

Table 12: Frequency of Production Behavior with High and Low Persistence Sales over 1000 Simulations

Sales Path	*Regular Hour Weeks*	*Inventory Adjustment Weeks*	*Weeks with Shift Changes*	*Short Weeks*	*Overtime Hours Weeks*
AR(1) = .848	77.4%	11.4%	4.9%	0.5%	10.6%
AR(1) = .509	70.3%	11.5%	0.7%	~ 0%	18.3%

Table 13: Frequency of the Use of Different Margins

(Percent of Weeks Used)

	1972 – 1983		**1990 – 2001**	
	Weighted average of all plants	Weighted average of all plants, **holidays excluded**	Weighted average of all plants	Weighted average of all plants, **holidays excluded**
Shutdown of at least 1 day	24.5	7.9	23.6	10.9
Shutdown of 1 week	12.4	9.3	11.3	6.1
4 or more overtime hours	14.5	14.5	30.3	30.3
Change in the number of shifts	0.6	0.6	0.1	0.1
Change in the line speed	0.9	0.9	1.0	1.0

Table 14: Percent of Days Closed, by Reason

Reasons for closure	1972 – 1983	1990 - 2001
Inventory adjustment	4.1	3.7
Model changeover	4.4	2.5
Supply disruptions	1.2	1.1
Holidays	5.6	7.6

All percentages are calculated using the sum of days during which a plant exists and is not on permanent or extended shutdown as the denominator.

Table 15: Importance of Each Margin for the Weekly Variance of Output

(Percent impact of margin use)

	1972 - 1983	1990 – 2001
Inventory adjustment	28.7	30.5
Model changeover	31.1	21.1
Supply disruption	7.7	11.6
Overtime hours	5.8	13.4
Shifts	24.3	12.4
Line speeds	11.7	9.2

Figure 1: Durable Goods Inventory-Sales Ratio

(Current dollars)

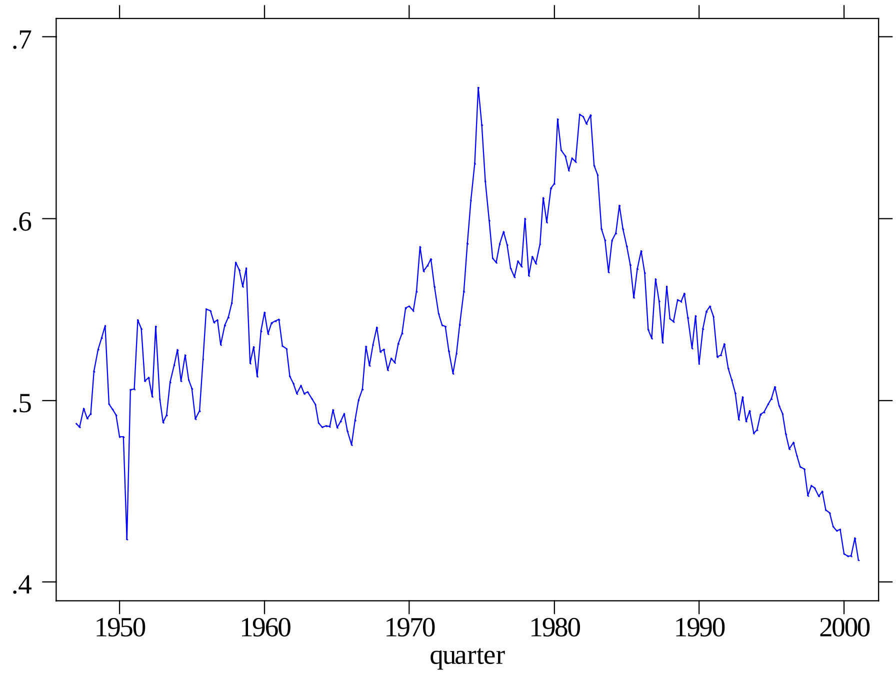

Figure 2: Inventory to Sales Ratio for U.S. Domestic Cars and Trucks

(In physical units)

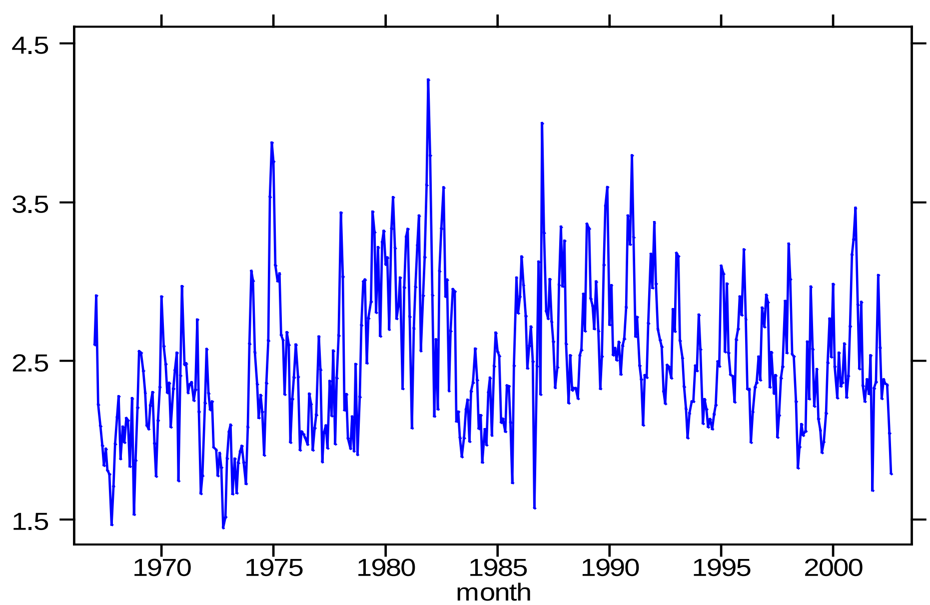

Figure 3A: U.S. Monthly Car and Light-Truck Sales

(In Physical Units)

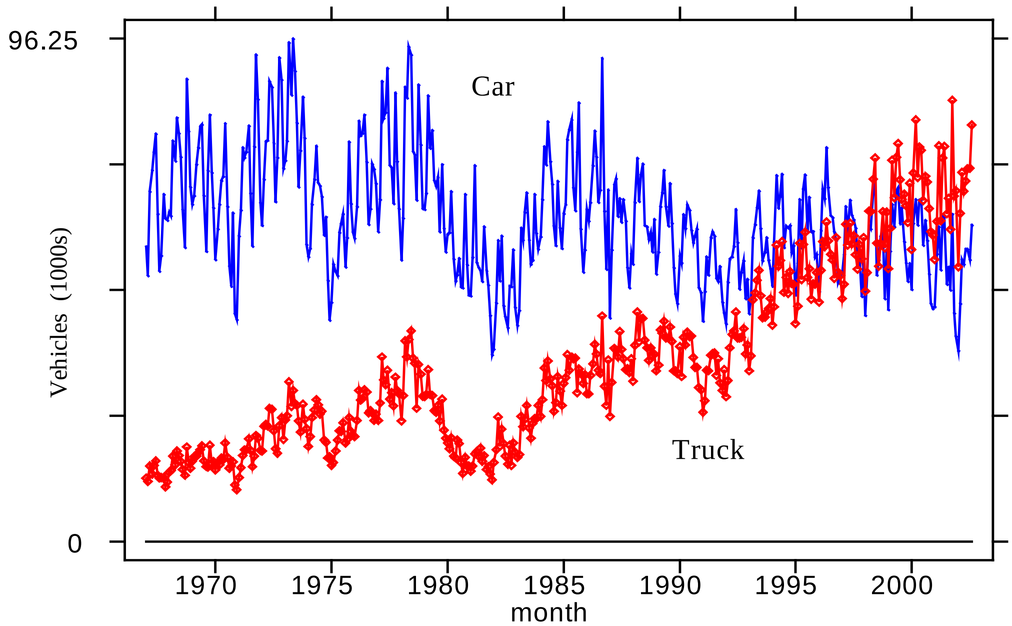

Figure 3B: U.S. Car and Light-Truck Monthly Production

(In Physical Units)

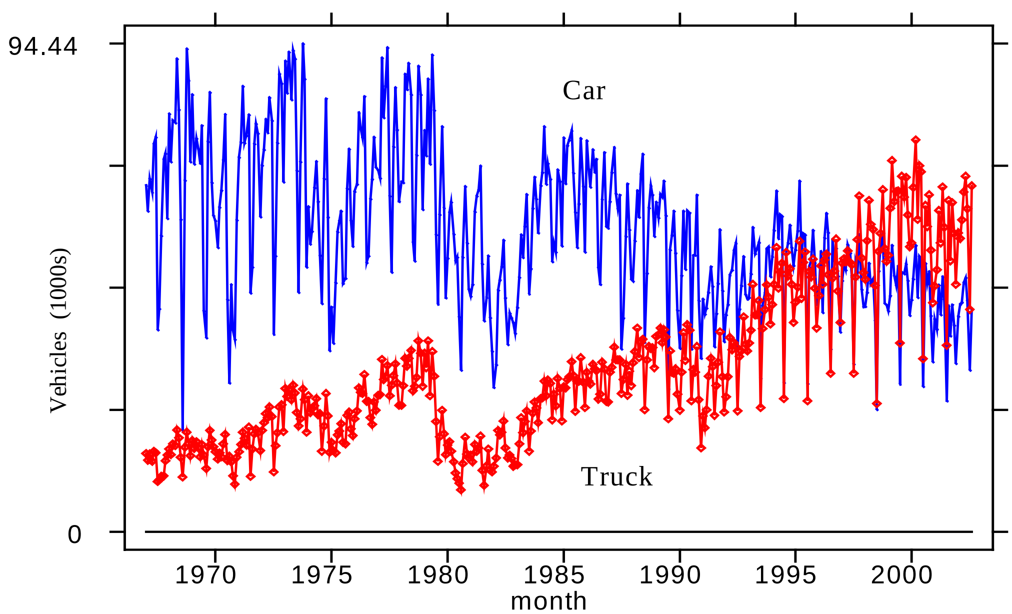

Figure 4: Density of the Drop in the Production Volatility over Sales Volatility Ratio between High and Low Persistence Sales Scenarios

(1000 Simulations, Volatility ratio measured in standard deviations)

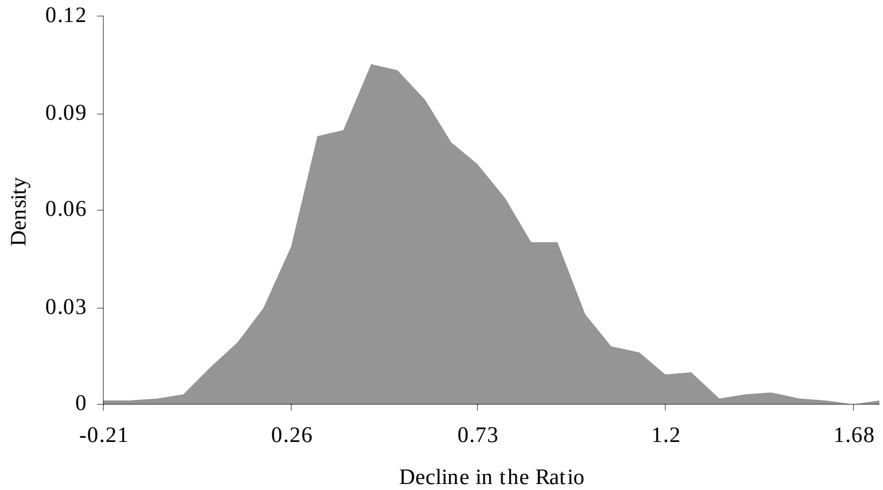

Figure 5: Average Weekly Overtime Hours per Worker: Motor Vehicle Manufacturing*

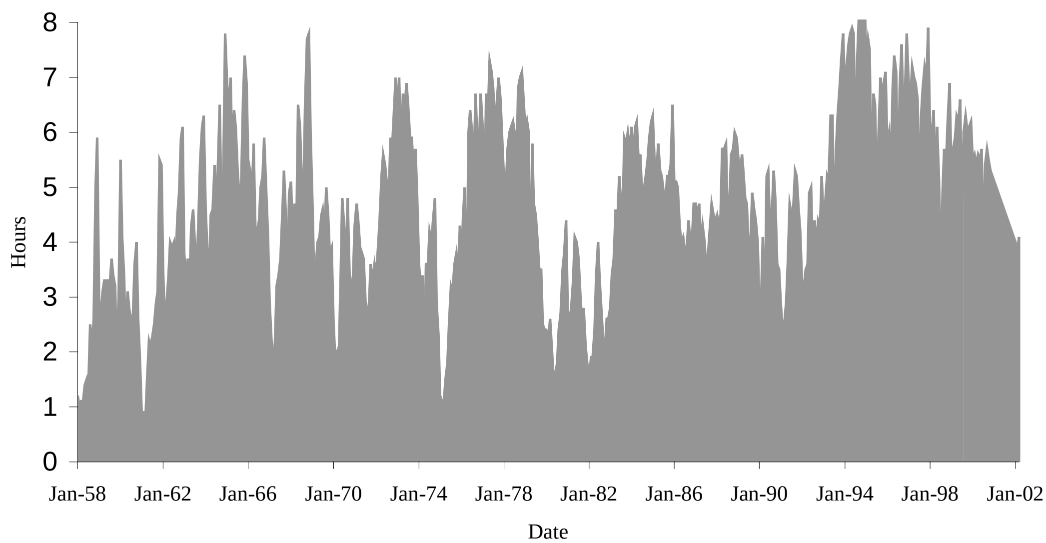

*Not seasonally adjusted