

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
Board of Governors of the Federal Reserve System

First Draft: December 2001
This Draft: January 2004

Abstract

Recent papers by Kim and Nelson (1999) and McConnell and Perez-Quiros (2000) uncover a dramatic decline in the volatility of U.S. GDP growth beginning in 1984. Determining whether the source is good luck, good policy or better inventory management has since developed into an active area of research. This paper seeks to shed light on the source of the decline in volatility by studying the microeconomic behavior of plants in the U.S. automobile industry, where the changes in volatility have mirrored those of the aggregate data. We find that changes in the relative volatility of sales and output, which have been interpreted by some as evidence of improved inventory management, are in fact the result of changes in the process driving automobile sales. Using a new dataset that tracks the production and sales of motor vehicles by assembly plant and by model, we first show that the autocorrelation of sales dropped markedly during the 1980s. A simulation of the assembly plants' cost function illustrates that the persistence of sales is a key determinant of output volatility. A comparison of the ways in which assembly plants scheduled production in the 1990s relative to the 1970s supports the intuition of the simulation. Finally, reduced form evidence suggests that changing behavior of interest rates is the force behind the change in sales persistence.

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 Garey Ramey, Jeronimo Zettelmeyer, participants in the NBER Monetary Economics group, the Hydra, Greece Workshop on Dynamic Macroeconomics, 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. Opinions expressed here are those of the authors and not those of the Board of Governors of the Federal Reserve System.

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) and Stock and Watson (2003a) show that all G-7 countries save Japan have experienced a decline in volatility in recent periods, though the timing and nature of the declines vary by country.

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 a systematically more decisive and more transparent monetary policy has been the key source of the decline in volatility of the U.S. economy (e.g. Clarida, Galí, and Gertler (2000), Boivin and Giannoni (2002) and McCarthy and Zakrajšek (2003)). 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) and Irvine and Schuh (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 difficult. First, a significant reduction in volatility has been discovered almost universally across the U.S. economy. Aggregate volatility patterns thus provide little help in narrowing the focus to a particular change in the economic environment or to the contribution

of a particular input.¹ Second, the conclusions reached from time-series analysis on aggregate data have been difficult to interpret. The reduction in GDP volatility seems to stem from a decline in the volatility of shocks, which are typically the forecast errors that result from a vector autoregression model of the U.S. economy. To further associate these forecast errors with measurable shocks, such as monetary and fiscal shocks, productivity shocks, supply shocks, etc., however, has not met with success.²

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. Not only has the automobile industry experienced a sharp decline in its volatility during the 1980s, but it is also an industry often examined by economists such as Blanchard (1983) to understand the interaction between inventories and the volatility of output. Using aggregate automobile industry data as well as a new disaggregated dataset that tracks 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 output volatility is linked to (1) a decline in the measured persistence of sales shocks and (2) plant-level non-convexities in production scheduling. In particular, sales are far less serially correlated after 1984 than they were during the 1970s and early 1980s, and interest rates appear to be a source of the change. We demonstrate that an inventory model involving non-convex 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 and to a decline in the covariance of inventory investment and sales. Finally, we show that the type of data used by other researchers may hide a change in persistence in other parts of the economy. Thus, recent evidence suggesting no change in the persistence of final sales overall may be an artifact of the data.

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

¹ Stock and Watson (2002) test 168 U.S. macroeconomic time-series and discover the pervasiveness of this volatility decline in measures of output, employment, investment and inflation. Manufacturing has been more impacted than services, however, and durable goods output volatility has experienced a particularly steep reduction.

² Stock and Watson (2002) search over several identifiable shocks in the U.S., and Stock and Watson (2003a) do so again for all G7 countries. In both cases, the behavior of observable shocks is quite different from that of the forecast errors of GDP growth.

output volatility, inventory investment and improvements in information and production technologies. Section III argues that the automobile industry is 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 the non-convex production costs typical in this industry. Section V confronts our theory with a panel of production scheduling variables across U.S. assembly plants. Section VI discusses whether the automobile industry results may be extended to the rest of the economy and Section VII concludes.

II. Overview of Volatility and Inventory Changes

Before presenting a detailed analysis of how production is scheduled in 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.

II.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 motor vehicles. All data are from the NIPA accounts of the BEA, and volatility is measured as the standard deviation of annualized growth rates. 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.³ 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 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

³ The volatility of aggregate GDP growth in the truncated early period is similar to the volatility of the extended period from 1953 – 1983 as shown in Kahn, McConnell and Perez-Quiros (2002).

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.

II.B Information Technology and Inventory Investment

Innovations in information and production technology have transformed U.S. manufacturing in many important ways in recent decades. The current question is whether technological change also underlies the moderation in GDP volatility. Kahn, McConnell and Perez-Quiros (2002), hereafter KMPQ, give perhaps the most compelling evidence that it has. They argue that the decline in GDP volatility is most closely matched in magnitude and in timing by a similar reduction within the durable goods sector.⁴ KMPQ further argue that the adoption of new machine tool technologies and inventory control systems occurred most rapidly during the 1980s within the durable goods sector. Thus, the “Information Technology Hypothesis” posits that production lines and inventory distribution systems now respond more rapidly to sales conditions than they used to, and this improvement has moderated the response of production to sales shocks.

Two of the most striking pieces of statistical evidence presented by KMPQ in favor of this 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:

⁴ 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.

$$(1) \quad \text{Var}(Y_t) = \text{Var}(S_t) + \text{Var}(\Delta I_t) + 2\text{Cov}(S_t, \Delta I_t)$$

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.⁵ 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:

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

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 $\Delta \hat{i} = \hat{y} - \hat{s}$.⁶ Then we can decompose the variance of production growth as follows:

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

Table 2 reports the results of this decomposition for all durable goods as well as motor vehicles. The 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, in an accounting sense, 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

⁵ Unit root tests on the BEA data fail to reject a unit root.

⁶ 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.

early period to negative in the post-1984 period.⁷ 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.

The connection between the “Information Technology (IT)” hypothesis and the covariance of inventories with sales is as follows: Information technology innovations, such as electronic scanning of bar codes, allow for automatic restocking based on real-time sales information and facilitate higher efficiency along the entire supply line. Elements of flexible manufacturing, such as computer numerically controlled machine tools, have additionally led to a reduction in set-up times required to produce different specifications of goods. This change in set-up 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 pro-cyclical, and hence so destabilizing.

The volatility and covariance measurements 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) the effect on production should be more pronounced than the effect 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.

These facts are consistent with the hypothesis that information technology and improved inventory management is the source of the decline in volatility. This hypothesis, however, is not so well supported by specific studies of inventory control methods and their effects on output volatility. McCarthy and Zakrajšek (2003) cite numerous articles from the operations research literature that do draw a connection between improved inventory control methods and a lower inventory-to-sales ratio at various stages of processing, but fail to find significant effects on output volatility. Similarly, Feroli (2002) models the optimal inventory-to-sales ratio as a

⁷ Golob (2000) first discovered this switch in the sign of the covariance term.

function of input prices, where inventories serve as an input to production that is substitutable with equipment and software. While declines in the relative user cost of information technology do explain the declining inventory-to-sales ratio quite well, the impact on output volatility is minute.

Patterns in U.S. inventory data also pose two puzzles for the IT 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 information technology innovations to reduce both the inventory-sales ratio and the volatility of output in a similar way if they were the source of these changes. The data do not give such a clear picture.

Figure 1 shows the ratio of nonfarm inventories to final sales since 1947. The data are in chained dollars, which best measure the real trend since the current-dollar ratio's trend is driven by relative price changes (Ramey and Vine (2003a)). Two distinct features in the graph are important. First, there was a large run-up in the inventory-sales ratio from the late 1960s to the early 1980s. Second, there has been an overall decline in the inventory-sales ratio from the early 1980s to the present. While some industries may have historically low inventory-sales ratios, the aggregate inventory-sales ratio since the early 1980s has been higher than it was during the 1950s and 1960s. 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. Thus, the behavior of the inventory-sales ratio does not line up very closely with the changes in volatility over time.

Within the stock of finished motor vehicles, there has been no trend in the inventory-to-sales ratio. 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. At the same time, however, the volatility of production for this product dropped markedly in the 1980s. 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 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? We investigate this possibility both theoretically and empirically.

III. Structural Change in the U.S. Automobile Industry

Production behavior in the U.S. automobile industry exhibits many 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 non-convex costs that sway automakers to use intermittent production in place of producing smoothly at levels that may fall 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 developed many of the 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.

III.A Overview of Patterns in Physical Unit Data for the U.S. Automobile Industry

Figures 3A and 3B depict monthly sales and production of domestic passenger cars and trucks in the U.S. from 1967:01 through 2002:09.⁸ From this point on we use unit data rather

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.

than NIPA data, as the unit data allow for direct comparison of inventories, production and sales without the constraints of chained-dollar quantities.

Several features in these plots are noteworthy. First, motor vehicle production is less volatile in the late portion of the sample than in the early portion. Following the strategy of McConnell and Perez-Quiros (2000), the variance of motor vehicle output is tested for a structural break. There is a significant reduction in volatility, and the estimated break date 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 estimates. To maintain consistency between our analysis and other studies, we chose to compare automobile industry variance estimates on either side of the generally accepted break date, which is the beginning of 1984.

A second feature visible in Figure 3 is the difference between the secular trend for passenger car sales, which have no trend, and the trend for light truck sales, which have grown steadily since 1980. Since there are obvious differences in the conditional means of these series, it is also useful to decompose the volatility of unit production into sales and inventory contributions for each class of vehicle separately.

Table 3 reports the variance decomposition for cars and trucks separately.⁹ One important difference between the time-series properties of physical unit data and NIPA data is that stationarity tests on the logarithm of physical unit variables reject a unit root in favor of a deterministic trend, with perhaps a break in trend around 1984 for trucks. Thus, variance here is based on deterministically detrended data.¹⁰ 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.

⁹ The variances and covariances do not add up because we have excluded imports and exports to and from Canada and Mexico. When these elements are included in the analysis, production volatility drops even further as the new Big Three plants in Mexico produce in a fashion that is negatively correlated with U.S. production. U.S. Exports to Canada and Mexico have little impact on the variance of sales. While import/export activity within North America raises additional interesting questions, it further augments the results obtained here and does not cause them.

¹⁰ 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

$\hat{Y}_t^T = \exp(\hat{\beta}_0 + \hat{\beta}_1 \cdot t)$ and the β 's are the estimated parameters of a regression of $\log(Y_t)$ on t .

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.

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 (see Blanchard (1983)). 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.

III.B Persistence of 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)

$$\log(\text{Sales}_t) = \alpha_0 + \alpha_1 \cdot \log(\text{Sales}_{t-1}) + \alpha_2 \cdot \text{trend}_t + \beta_0 \cdot D_t + \beta_1 \cdot D_t \cdot \log(\text{Sales}_{t-1}) + \beta_2 \cdot D_t \cdot \text{trend}_t + \varepsilon_t$$

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

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

This model allows all parameters to change in 1984:1: the coefficient on lagged 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 4 shows the results of this exercise, and the coefficient estimates indicate significant changes have occurred in the process governing sales. The constant term and the lagged sales 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 the coefficient on lagged sales, which serves as a measure of the persistence in monthly shocks to 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.

These sales data are for domestic sales of domestically produced vehicles. Thus, they include vehicles with foreign nameplates that are produced domestically. When estimating the sales process of Big Three-only vehicles, the first-order autocorrelation declines from .86 to .51 for autos, and from .84 to .58 for all light vehicles. When the definition of sales is expanded to all domestic sales of autos (including imports), this parameter estimate declines from 0.86 to 0.57. Alternatively, total sales of domestic manufacturers, including exports, have a decline in persistence from 0.86 to 0.51. Thus, including or excluding imports and exports does not change the basic results.¹¹

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.¹²

¹¹ Only the results for autos are reported in these last two cases because import or export data for light trucks is not available for the entire sample.

¹² For cars, the largest root declined from .90 to .80, the sum of AR terms declined from .90 to .79, and the half-life dropped from 4 to 2 months. For trucks, the largest root declined from .96 to .90, the sum of AR terms declined from .95 to .83, and the half-life dropped from 11 to 3 months. For a discussion on the features and limitations of each measure of persistence, see Pivetta and Reis (2002).

In summary, the aggregate motor vehicle sales process in the post-1984 period 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.

These results beg the question: why did the persistence of automobile sales decline? Ideally, one would do a structural analysis of a dynamic stochastic equilibrium model of this durable goods oligopoly to determine the source of the change. Data and space limitations prevent us from doing a full structural analysis here. Instead, we offer some reduced form evidence that is suggestive of the source of the decline.

Likely candidates for the change in sales persistence are changes in (1) firm pricing behavior; (2) interest rates, presumably due to changes in the conduct of monetary policy; and (3) changes in aggregate income. Suppose that automobile sales depend on prices, aggregate income, interest rates, and an unobservable shock. When we estimate a univariate *AR* process on auto sales alone, the estimated autocorrelation depends on the data generating processes of these other variables as well as shocks to demand. Suppose that automobile companies began to aggressively offer incentives and price discounts in the 1980s on certain vehicles to rekindle sagging demand. If these price incentives represent a significant change of practice, are used only at key points in time and have been successful in their aim, then sales shocks in the univariate model would appear less persistent. Alternatively, if the monetary authority has become more aggressive in using interest rates to respond to underlying shifts in the economy that also affect automobile sales, then the measure of sales persistence in the univariate model would again appear to have declined.

To determine whether any of these variables have indeed led to the reduction in the autocorrelation of sales observed above, we augment Equation 4 with the current value and four lags of real automobile prices, real income, and nominal interest rates.¹³ Assuming that automobile firms do not adjust their prices to the current month's sales shock, this equation can roughly be viewed as a demand equation. The goal is to see whether including any of these other

¹³ The demand for automobiles should depend on the rental cost in theory. Rental cost variables have less explanatory power than the components (interest rates and prices) entered separately.

variables makes the estimated persistence parameter constant across the periods, thus indicating that the behavior of the extra variable was the source of the difference.

Table 5 displays the results. The first row shows the estimated change in persistence for the baseline with no other variables. For all classes of vehicles, the change in the persistence across periods is significantly negative. The second row adds real motor vehicle prices. For every class, the estimates are similar to the baseline and remain negative and significant. The third row adds real income instead, which does reduce the magnitude of the change after 1984, but the estimated change is still statistically significant. The results in the fourth row indicate that the interest rate has the largest impact on the change in the $AR(1)$ estimate. The magnitude of the change falls noticeably relative to the baseline, and 1984 no longer represents a statistically significant break date in two of the three vehicle classes. These changes become even smaller when all three variables are included. In cars and motor vehicles, the estimates are less than half what they were in the baseline case and are not significant. In the case of light trucks, the estimates are less than in the baseline case, but are still significant.

In summary, it appears that the change in persistence in sales owes in part to the behavior of interest rates and perhaps income over the two periods. Surprisingly, including prices does not have any explanatory power.

III.C Persistence of Model-Level Motor Vehicle Sales

While the decline in the propagation of aggregate motor vehicle sales shocks explains some of the change in aggregate volatility patterns, what is most important for production scheduling is the behavior of individual model sales. 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 sales have changed in the same way at the model level as at the aggregate level, we estimate the parameters for a first-order autoregression using sales data for each vehicle model. Preliminary investigations suggested that estimating the data as a panel with

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.

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 data feature that is unique to model-level analysis is the life cycle of vehicle models. Car models clearly have a life-cycle pattern, where sales build slowly after the model is introduced, and then decline as the model is phased out.¹⁵ These patterns affect the autocorrelation of sales. 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.¹⁶

For each model that exists for at least 24 months, we estimate a first-order autoregression 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 6.

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, leaving the unconditional variance of model-level sales relatively unchanged 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.

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*. Their presentation of evidence implicitly assumes there is a fixed

¹⁵ The last months of a model's life cycle often feature a spike in sales, presumably from steep markdowns at dealer clearance events. While this behavior does affect aggregate sales volatility, we wish to focus here on changes in the normal sales process.

¹⁶ Other alternatives studied were different specifications of trends, such as quadratic and cubic.

relationship between the variance of production and the variance of sales in the absence of structural change. In this section, we show how a change in the sales process alone modifies the relationship between production, inventories and sales. This is true in the standard production smoothing model when cost functions are convex (see Blanchard (1983), Ramey and Vine (2003b)), but is more likely when cost functions are non-convex. Here we examine the more realistic plant-level model with non-convex costs and lumpy production margins. 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. Using a model calibrated to the cost parameters of the U.S. automobiles industry, we show that the sales persistence changes found in the data are sufficient to explain the changing relative volatilities of production and sales.

IV.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).

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

The institutional structures in automobile manufacturing have various implications for the marginal costs of using and for the fixed costs of changing these margins. Aizcorbe (1990, 1992), for example, documents the implications for marginal cost found in the labor contract between automobile manufacturers and the United Auto Workers. Bresnahan and Ramey (1994) and Hall (2000) both study the ways in which production margin changes impact the volatility of

plant output. A conclusion that is common to these as well as other automobile industry studies is that the constraints placed on production scheduling by union rules and the high cost of retooling an assembly line make the plant's cost function non-convex over many ranges of output.

Consider first the number of hours an assembly plant works each week. Changes in regular (non-overtime) hours most often take the form of closing the plant for an entire week. This is called intermittent production. This is preferable to operating a curtailed schedule, called a short week, because plants are 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).¹⁷

The length of a workweek may be extended temporarily with overtime hours. Overtime hours take the form of one or two extra hours at the end of a regular eight-hour shift or as a Saturday shift. Employees who work more than eight hours per day or more than forty hours per week receive a 50% wage premium for the extra hours. Overtime hours are intended to be temporary and assembly plants are prevented from using overtime permanently in lieu of hiring additional workers. Frequent discontinuous spells of overtime, however, are not uncommon.

Long-term adjustments to production may involve adding or dropping the night shift. 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. 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.

¹⁷ 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.

Finally, the volume of output may also be changed directly by slowing or accelerating the rate of production. Line speed changes require a reorganization of the assembly line, which implies a period of downtime before the redesigned line is complete. Workers do not simply assemble cars faster when line speeds increase. Instead, each shift hires more workers. The UAW typically becomes involved with changes in the line speed as well.¹⁸

IV.B Cost Function Simulation with Inventories

Not surprisingly, the nature of assembly line technology and the language written into the UAW contract imply several levels of production that are either prohibitively expensive or physically impossible to attain. 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 weeklong 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 shocks is compared with the production behavior from a sales process with more transitory shocks. The conclusion is that the relationship between the volatility of production and the volatility of sales is non-linear and depends on the dynamic properties of sales.

IV.B.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 scheduled in week t , Sh_t , the number of days the plant will open in week t , D_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.

¹⁸ 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 line speeds be reduced and the number of tag-relief workers be trimmed. (*The Wall Street Journal*, December 18, 2001)

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

The line speed can be thought of as the plant's production function, as it is the flow of output made possible by 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, \bar{n} .

$$(7) \quad ls_t = k_t^{1-\gamma} \cdot (n_t - \bar{n})^\gamma$$

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.

$$(8) \quad c(h_t, D_t | Sh_t) = \sum_{j=1}^3 I(Sh_t \geq j) w_j \cdot D_t \cdot h_t \cdot n_t \\ + \max \left[0, 0.85 \sum_{j=1}^3 I(Sh_t \geq j) w_j (40 - D_t \cdot h_t) n_t \right] \\ + \max \left[0, 0.5 \sum_{j=1}^3 I(Sh_t \geq j) w_j \cdot D_t (h_t - 8) n_t \right] \\ + \max \left[0, 0.5 \sum_{j=1}^3 I(Sh_t \geq j) w_j (D_t - 5) h_t \cdot n_t \right] \\ + [u \cdot w_1 \cdot 40 \cdot Sh_t \cdot n_t (D_t = 0)] + \delta \cdot I(D_t > 0)$$

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 workers who spend more than 0 but less than 40 hours per week on the job. The number of shifts is chosen in the prior week from the set of 1, 2 or 3. The hourly wage for the j^{th} shift is denoted w_j , and I is an indicator variable that returns a value of 1 when the expression in parentheses is true. 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 and the second term is the fixed cost of opening the plant each week, δ .

The production schedule chosen each week depends on the following variables: (1) the level of sales in the current week, (2) expected level of sales in future weeks, (3) the plant's operating status in the prior week and (4) the level of inventory available to help meet sales in week t . The stock of inventory carried from period t to $t+1$, therefore, 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 lie within an interval depicted in Equation 10.¹⁹ 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, ω^* .²⁰

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

$$(10) \quad \underline{I} \leq I_t \leq \bar{I} \text{ for all } t ; \underline{I} \geq 0$$

The second channel through which the plant's history affects current decisions involves the fixed adjustment costs the plant incurs when the production schedule is changed. Bresnahan and Ramey (1994) present evidence that changing the line speed or the number of shifts working entails high adjustment costs, while changing other margins, such as scheduling overtime hours and closing the plant for week-long intervals, involves relatively low adjustment costs. In our exercise, changing the number of shifts entails a fixed adjustment cost, α_{Sh} .

¹⁹ 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.

²⁰ This accelerator term is common in many inventory models, and is particularly well-suited to the automobile industry. Not only do the automakers and auto trade publications track and respond to this statistic (*days' supply*) very actively, but previous automobile studies have deemed it an important part of the industry's inventory behavior as well. See Blanchard (1983) and Kashyap and Wilcox (1993).

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 α_I , and the fixed adjustment cost of changing the number of shifts working, α_{Sh} . The inter-temporal cost function denoted as $C(h_t, D_t, Sh_{t+1} | I_{t-1}, S_t, Sh_t)$ is in Equation (11).

$$(11) \quad C(h_t, D_t, Sh_{t+1} | I_{t-1}, S_t, Sh_t) = c(h_t, D_t | Sh_t) + \alpha_{Sh} \cdot I(Sh \neq 0) + \frac{1}{2} \alpha_I \cdot [I_t - E_t(S_{t+1}) \cdot \omega^*]^2$$

IV.B.2 Dynamic Program Simulations

In order to understand the production behavior implied by the cost minimization problem under different sales conditions, this section simulates the dynamic cost minimization problem 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.

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 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 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.²¹ 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 weeklong plant closure. A short week is also available through many combinations of production margins.

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.²²

Sales evolve as 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. The inventory grid consists of points compatible with the sales grid and the production possibilities, and its boundaries range from a fourteen days' supply to a ninety days' supply of the mean sales rate. The nine sales grid points along with the Markov transition-probability matrix $\chi(s' | s)$ are parameterized to mimic a desired $AR(1)$ sales process using Tauchen's (1986) procedure. The unconditional mean of sales, μ_s , 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 full-time capacity with the mean sales rate. Mismatches between the capacity of a plant and its realized sales 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 their first-order autocorrelation, this exercise consists of two simulations: Simulation #1 solves the plant's cost minimization problem with a persistent monthly sales process ($AR(1) = 0.85$), and simulation #2

²¹ Shifts may take the value of 1, 2 or 3. Days open per week are chosen from the integers 0 through 6. Hours per day are available in increments of 2 ranging from 0 through 10.

²² 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.

solves the same problem with the $AR(1)$ coefficient reduced to 0.51. These parameters match the estimated declines in persistence in the aggregate automobile data, which were less extreme than the declines in the model-level data. To determine the pure effect of a change in persistence with no overall change in unconditional variance, we raise the variance of the innovations in the second simulation so that the unconditional variance of sales is unchanged between simulations. Thus, these simulations give a lower bound on how much of the relative change in output volatility we can explain.

The parameter values used throughout this exercise come from 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 mimics the production behavior observed among assembly plants in the 1972 – 1983 period as closely as possible.^{23 24}

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 D_t and h_t as well as Sh_{t+1} 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:

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

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

²³ $\beta = .999$ (weekly discount factor); $u = 65\%$ (firm's share of unemployment compensation); $\alpha_l = 0.0024$; $\alpha_{SH} = 2.4$ weeks of regular-hours wage; $\omega^* = 60$ (days' supply); $\gamma = 0.62$; $n_1 = 364$; $n_2 = 58$; $\delta = 0.54 \cdot [40 \cdot w_1 (n_1 + \bar{n})]$ (54% of the first shift's wage bill).

²⁴ Shift changes occur too frequently in the simulation relative to actual data, but the other margins are matched quite closely.

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

$$\underline{I} \leq I_t \leq \bar{I} \quad \text{for all } t \in [0, T],$$

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

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

and the initial inventory level

$$I_0 = \omega^* \cdot \mu_S.$$

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 ($AR(I) = .85$) and the more transitory ($AR(I) = .51$) sales scenarios.²⁵ The plant solves its weekly cost minimization problem, which determines the optimal paths for the workweek variables as well as for production and inventory stock. These solution paths are then aggregated from a weekly to a monthly frequency and their volatility properties investigated.

The simulation results are summarized in Table 7, which also includes 95% confidence intervals for certain point estimates and for the changes in the point estimates between scenarios. In both the high persistence case and in the low persistence (but higher innovation variance) case, the average (unconditional) standard deviation of monthly sales across simulations is 27.0% of the mean sales rate. The average standard deviation of the optimal production path, however, changes significantly between the two scenarios. The standard deviation of production falls from 40.5% of the mean sales rate per month in the first simulation to 26.5% in the second simulation. While the average volatility of sales remains unchanged by design between the first and second simulations, the volatility of production falls. Accordingly, the ratio of the variance of production over the variance of sales falls from 2.39 to 0.96.

²⁵ The Markov process that generates weekly sales was calibrated so that, on average, monthly sales exhibited the desired first-order autocorrelation.

Our simulations also match the facts concerning the change in the covariance of inventory investment and sales. As Table 7 shows, when sales have a persistence parameter of 0.85 the correlation between inventory investment and sales is 0.3. In contrast, when sales have a persistence parameter of 0.51, the correlation becomes negative. The intuition is the same as in the production smoothing case. If a sales shock is thought to be very persistent, the firm wants to increase production dramatically in order to keep maintain its desired inventory-sales ratio, since it knows sales are likely to stay high for awhile. In contrast, if the shock is more transitory, the firm is willing to allow a deviation from the desired inventory-sales ratio since it expects the deviation to be short-lived.

In order to assess from which sources the change in volatility behavior originates, Table 8 shows the weekly frequency with which various changes to production were made. Weeklong shutdowns for inventory adjustment, for example, occur in 10.5% of the weeks when a shock to sales has a persistent effect, and that figure declines slightly to 9.0% when the persistence is reduced. Shift reductions, alternatively, are much more common when the persistence of sales changes is high. Shift changes occur in 3.8% of the weeks in the high persistence case, and only in 0.03% of the weeks in the low persistence case. Short-weeks in both scenarios are non-existent. The frequency of the use of overtime hours increases from 14.3% in the high persistence scenario to 20.3% in the low persistence simulation.

The conclusion of this simulation exercise is that the variance of output and the covariance of inventory investment and sales are 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, and smaller, 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 non-convex costs, but it features no changes to the cost function parameters, inventory targets or improvements in the flow of future sales information.

The next section investigates changes in production scheduling that have actually taken place at domestic assembly plants. The changes actually observed line up quite well with the model's predictions.

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 non-convex 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 plant 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.^{26 27}

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 works; (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*.²⁸

Table 9 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

²⁶ The period 1984 to 1989 was excluded only because we did not have access to *Automotive News* from that period when we were collecting the data.

²⁷ 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.

²⁸ See Bresnahan and Ramey (1994) for more detail about how data is extracted from the weekly production articles.

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 occur with roughly the same frequency in both 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 10 isolates the plant shutdown margin in the early and in the late 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 decrease in inventory adjustments is very minor. The number of holidays has 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 the margin through which improvements in manufacturing technology would be visible.²⁹ There is indeed evidence that model changeover technology has advanced over time, as the industry introduced the *weekend 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,³⁰

²⁹ These numbers are the weighted average across plants. When model changeovers do occur, they almost always occupy entire weeks.

³⁰ Christmas 1982 lasted until almost February 1983 in many plants!

and model changeovers often take place during a summer vacation or are much longer than the technology necessitates during periods with low demand.³¹

Finally, Table 11 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, boost weekly production by up to 25%, while adding a second shift doubles 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 actual output and constructed output. The numbers do not add to 100 because of nonlinearities and covariance terms.³²

Table 11 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 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 they do during the first period, falling in contribution from 24.3% to 12.4%.

Thus, the two non-convex 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. The increase in overtime hours is corroborated by other data sources as well. The BLS data on overtime use in the automobile industry show a significant increase in overtime used from the early period to the later period.

³¹ 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.

³² See Bresnahan and Ramey (1994) for a more detailed explanation of the method.

VI. Relevance for the Aggregate Economy

The persistence of sales has declined significantly in the post-1984 period in the automobile industry, and this decline in persistence can explain the observed changes in relative volatilities and covariances. The question arises as to whether this phenomenon is unique to the automobile industry or whether it also occurred in other sectors of the economy that experienced declines in output volatility. At first glance it might seem that our story of declining persistence is unique to the automobile industry. Blanchard and Simon (2001), Ahmed, Levin and Wilson (2002), and Stock and Watson (2002) are among those who have tested for a break in the autoregressive parameters and found none.

We believe that the literature's failure to find a structural break may owe to data limitations. Our data differs from the data used in the literature in three important ways: (i) its level of disaggregation; (ii) the higher frequency of its observation; and (iii) its measurement in physical units rather than in chained-dollars. These choices are important to our story. Since our theory addresses the way in which plants schedule production in response to sales, it is the change in persistence of plant-level sales, not aggregate sales, which matter most directly. Second, plant-level decisions typically occur at the weekly or monthly frequency, so the appropriate sales data would also be high frequency. Third, production volatility most likely depends on the properties of sales in *physical units* as opposed to the properties of their *value*. These alternative measures of sales can have very different persistence properties. Sales measured in chained-dollar reflect the real value of a good along with its physical-unit quantities. As shown below for automobiles, the real value of a car is a time-series with its own very high persistence.

Table 12 shows how these three data features (the frequency of observation, the level aggregation and the unit of account) affect the tests for a change in persistence. We estimate the model given in Equation 4, which estimates the change in the first-order autocorrelation as well as in other elements of the conditional mean and variance in 1984, for motor vehicle sales data constructed at different frequencies and with units of account. The first row replicates Table 6 and measures the change in the average attributes of the physical unit sales process at the model level. These data show the largest decline in persistence of sales.³³ The next three rows replicate

³³ Because we used the weighted average of estimates over all models, standard errors are not reported.

changes in the monthly physical unit sales process measured in Table 4 for cars, light trucks, and their aggregate. In all three cases there is a significant change in persistence, and in two of three cases there is no significant change in the variance of the error term.

Consider next the lower panel of Table 12, which inter-temporally aggregates the physical unit data from the upper panel to a quarterly frequency using averages. In all cases the changes in persistence parameter decline by 50% in magnitude and are no longer significant. The changes in the variance of the innovations, on the other hand, now become significant in each of these series. Monthly data imply a change in the autocorrelation of sales and no change in the innovation variance, whereas the quarterly data imply just the opposite!

The final row of Table 12 tests the final sales of domestic product of motor vehicles for changes in persistence and innovation variance. This series is most similar to the types of variables others have used in more aggregated studies, and it is observed quarterly in units of chained 1996 dollars.³⁴ The estimated decline in persistence is even smaller in this case than it was when measured with quarterly physical units in the preceding line, and the change is not significant. The variance of the innovations, however, appears to have increased significantly in the second period. Thus, this variable gives answers that are even farther from those obtained using our preferred data.

The chained data contains another element that may be masking the change in persistence: the behavior of average real expenditures for cars. The chained data is the number of physical units multiplied by chained dollar average expenditure per unit. As the addendum at the bottom of Table 12 shows, average real expenditures per car show no break in persistence or variance. Average real unit expenditures, however, have a high persistence (first-order autocorrelation is 0.92), so their effect on chained dollar total expenditures likely hide any persistence changes in the unit data.³⁵

These exercises suggest that testing for change in sales persistence using standard chain-weighted quarterly data is not revealing. While the decline in macroeconomic volatility was

³⁴ This variable includes exports. As noted in an earlier section, though, including exports in sales does not change the estimates noticeably when physical unit data are used.

³⁵ The case for using physical unit data rather than chained value data can also be made by analogy to studies of the labor market. Economists always study employment variables measured in physical units – either hours or number of workers. The stylized facts about labor markets would probably change significantly if economists only studied the real value of total hours (i.e. the product of real wages and hours) because the high persistence of real wages would hide the time series properties of the hours variables.

originally observed in quarterly U.S. GDP growth data, it is possible that the true source of this change is detectable only at higher frequencies and with less aggregation. Thus, the correct data to use in order to verify that the results given here do or do not extend to other industries are high frequency physical unit data, which we intend to pursue as future research.

At this point, we offer indirect evidence supportive of the relevance of our hypothesis for the economy overall. Both the simulations and plant-level data show that when sales persistence declined in the automobile industry, firms were more likely to vary hours per worker (e.g. overtime hours) than the number of workers (e.g. the number of shifts). Varying the hours per worker typically has low adjustment cost but high marginal cost, whereas varying the number of workers has high adjustment cost but lower marginal cost. Thus, firms will use hours per worker to respond to transitory shocks and will use employment to respond to more persistent shocks.

To determine whether production decisions in the overall economy appear to have changed in the same way as in the automobile industry, we analyze the change in the variances of the intensive margin (hours per worker) and the extensive margin (number of workers) over the two periods. The variables used are average weekly hours per worker and the number of employees in the nonfarm private sector (divided by the population aged 16 and over) at the monthly frequency. The product of the two variables is total hours per capita worked in the nonfarm private sector. We study the logarithms of the variables over the period for which all of the data are available, which is 1964 to the present.

There are substantial trends in the separate series, which are not the object of our analysis. Thus, we detrend the data using two different methods. The first method assumes unit roots in the series and studies the variances of the growth rates. The second method assumes a linear deterministic trend (that differs across the two periods) and studies the variance of the deviations from the trend.

As Table 13 shows, both methods give the same qualitative results. The variance of total hours fell substantially from the pre-1984 period to the post-1984 period. Fluctuations in the average hours per employee accounted for 48 percent of the variance of total hours in the early period but 65 percent of the variance in the later period. Meanwhile, changes in employment became a smaller part of the overall variance in the second period. The qualitative results are similar for the linearly detrended data.

Thus, the aggregate labor data display patterns similar to those found in the automobile industry. The use of margins of adjustment is entirely consistent with the theory that the sales shocks faced by individual firms in the early period had a higher persistence than those faced by firms in the later period.

VII. 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 in 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 non-convex 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 brief investigation of the source of the change in the sales process suggests that monetary policy could be the key. This interpretation is consistent with Clarida Gali and Gertler (2000), Boivin and Giannoni (2003), and McCarthy and Zakrajšek (2003), all of who find changes in the estimated monetary policy function at the Federal Reserve to underlie declines in the persistence of demand shocks and reduced output volatility.

Finally, an analysis of the data suggests that previous failures to find changes in the persistence of sales at the aggregate level may be due to temporal aggregation and other

problems with chain-weighted data. Evidence from the aggregate labor market suggests that firms have varied their labor input in a way that is consistent with sales shocks becoming more transitory.

VIII. Tables and Figures

Table 1: The Standard Deviation of Output Growth *
(Quarterly data, 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 <i>(Excluding durable goods)</i>	3.4	2.1	-37
Durable Goods	15.7	8.1	-49
Durable Goods <i>(Excluding motor vehicles)</i>	13.5	8.1	-40
Motor Vehicles	51.4	20.6	-60

*Based on quarterly NIPA data.

Table 2: Decomposition of Durable Goods and Motor Vehicle Output Volatility

(Quarterly data, 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

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

$I = \text{inventories.}$

Table 3: Decomposition of Motor Vehicle Output Volatility
(Physical units, cars and trucks separately)

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 4: Estimates of Aggregate Automobile Sales Process

Coefficients (and standard errors) from the regression:

$$Sales_t = \alpha_0 + \alpha_1 \cdot Sales_{t-1} + \alpha_2 \cdot trend_t + \beta_0 \cdot D_t + \beta_1 \cdot D_t \cdot Sales_{t-1} + \beta_2 \cdot D_t \cdot trend_t + \varepsilon_t$$

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

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

Coefficient	Cars	Light Trucks	Motor Vehicles
α_0 (constant)	0.334** (.116)	0.0348 (.0191)	0.275** (.107)
β_0 (Δ constant)	0.712** (.246)	0.271** (.094)	0.633** (.255)
α_1 (AR(1))	0.848** (.050)	0.934** (.031)	0.884** (.044)
β_1 (Δ AR(1))	-.339** (.125)	-.232** (.085)	-.263** (.112)
α_2 (trend)	-.0002 (.00014)	0.00016 (.00016)	-.00007 (.0001)
β_2 (Δ trend)	-.0002 (.0002)	0.0011** (.0004)	0.00056** (.00022)
σ^2 (innov. Variance)	0.0070** (.00117)	0.0090** (.0012)	0.0066** (.0011)
β_3 (% Δ innov. Variance)	-.00035 (.0016)	-.0042** (.0014)	-.0012 (.0014)
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.

Actual estimation is from 1967:4 – 2002:9, $N = 426$ $D_t = 0$ for $t \leq 1983:12$; $D_t = 1$ for $t \geq 1984:1$

Table 5: Estimated Change in Persistence of Sales: Explanations

Change in Late Period (1967 – 1983) from Early Period (1984 – 2002)

	Other variables included	Vehicle Category		
		Cars	Light Trucks	Motor Vehicles
Baseline	--	-.339** (.125)	-.232** (.085)	-.263** (.112)
	Real price of motor vehicles	-0.336** (.122)	-0.252** (0.082)	-0.270** (0.108)
	Income	-0.295** (0.129)	-0.225** (0.084)	-0.233** (0.111)
	Federal funds rate	-0.231* (0.135)	-0.184** (0.089)	-0.162 (0.121)
	Prices, income, and federal funds rate	-0.156 (0.145)	-0.169* (0.091)	-0.102 (0.128)

426 observations. The estimate reported is the difference in the $AR(1)$ coefficient on sales between 1967:1 – 1983:12 and 1984:1 – 2002:9. ** indicates significant at the 5% level, * indicates significant at the 10% level.

All specifications allow for breaks in the constant term, trend term, and variance of the error term between 1983:12 and 1984:1.

Table 6: Estimates of the Model-level Domestic Car and Light-Truck Sales Process: Persistence and Standard Deviation of Sales Shocks*

(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 7: Assembly Plant Simulation Results: Average Standard Deviation of Monthly Production and Sales over 1000 Simulations

Sales Path	$\frac{\sigma_S}{\mu_S}$	$\frac{\sigma_Q}{\mu_S}$	$\frac{\sigma_Q^2}{\sigma_S^2}$	$\rho_{S,\Delta I}$	Ave. % Change in σ_S/μ_S	Ave. % Change in σ_Q/μ_S
$\rho = .85$	27.0% (26.5 , 27.6)	40.5% (39.9 , 41.2)	2.39 (2.34 , 2.43)	0.34 (.334 , .351)	--	--
$\rho = .51$	27.0% (26.7 , 27.3)	26.5% (26.2 , 26.9)	0.96 (.951 , .976)	-0.26 (-.269 , -.246)	0% (-.51 , .42)	- 14.0% (-14.6 , -13.5)

Table 8: Assembly Plant Simulation Results: Frequency of Production Behavior with High and Low Persistence Sales over 1000 Simulations

Sales Path	Regular Hour Weeks	Inventory Adjustment Weeks	Overtime Hours Weeks	Short Weeks	Weeks with Shift Changes
$\rho = .85$	75.2%	10.5%	14.3%	0%	3.8%
$\rho = .51$	70.7%	9.0%	20.3%	0%	0.03%

Table 9: Frequency of Use of Different Margins at Big Three Assembly Plants

(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 10: Percent of Days Closed at Big Three Assembly Plants 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 11: Importance of Each Margin for the Weekly Variance of Output at Big Three Assembly Plants

(Percent impact of margin use)

Margin	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

Table 12: Effects of Using Different Types of Data on Changes in Persistence and Conditional Variance Estimates of Motor Vehicle Sales

Changes from 1967:I-1983:IV to 1984:I-2002:III

Frequency	Units of measurement	Aggregation level	Change in $AR(1)$ parameter	Change in variance of innovations
Monthly	Physical units	Model-level	-0.572	0.119
Monthly	Physical units	Cars	-0.339** (.125)	-0.00035 (.0016)
Monthly	Physical units	Light trucks	-0.232** (.085)	-0.0042** (.0014)
Monthly	Physical units	Cars & light trucks	-0.263** (.112)	-0.0012 (.0014)
Quarterly	Physical units	Cars	-0.179 (.165)	-0.0077** (.0035)
Quarterly	Physical units	Light trucks	-0.183 (.115)	-0.010** (.0028)
Quarterly	Physical units	Cars & light trucks	-0.126 (.145)	-0.0077** (.0033)
Quarterly	Chained data	All motor vehicles	-0.074 (.123)	-0.0047** (.0023)

Addendum:

$AR(1)$ results for average expenditures per car (deflated by BEA deflator):

- No evidence of break in $AR(1)$ parameter or variance of the innovation
- Full sample $AR(1)$ parameter is 0.916 with standard error of 0.0285.

Table 13: Aggregate Evidence on the Hours Margin Use
Nonfarm private sector

	Variances of growth rates		Variances of deviations from deterministic trends	
	1964:1-1983:12	1984:1-2002:9	1964:1-1983:12	1984:1-2002:9
Total hours	48.9	22.2	8.2	5.3
Employment	14.7	5.2	7.2	4.0
Average weekly hours	23.7	14.4	0.4	0.4
Covariance	5.2	1.3	0.3	0.5

Data are from the BLS. Total hours and employment are in per capita terms.

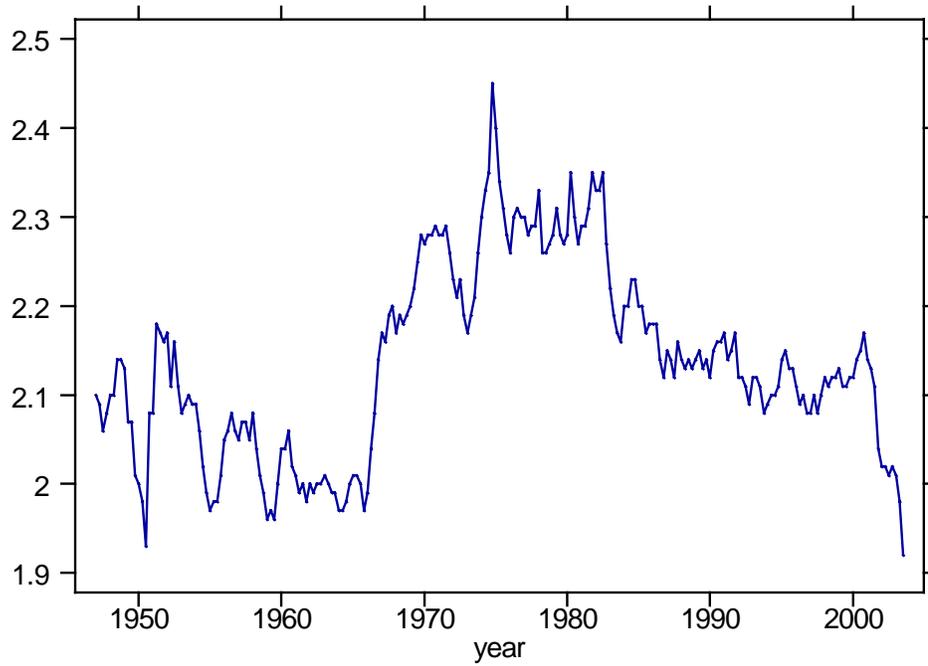


Figure 1: Nonfarm Inventory to Final Sales Ratio (Chained dollars)

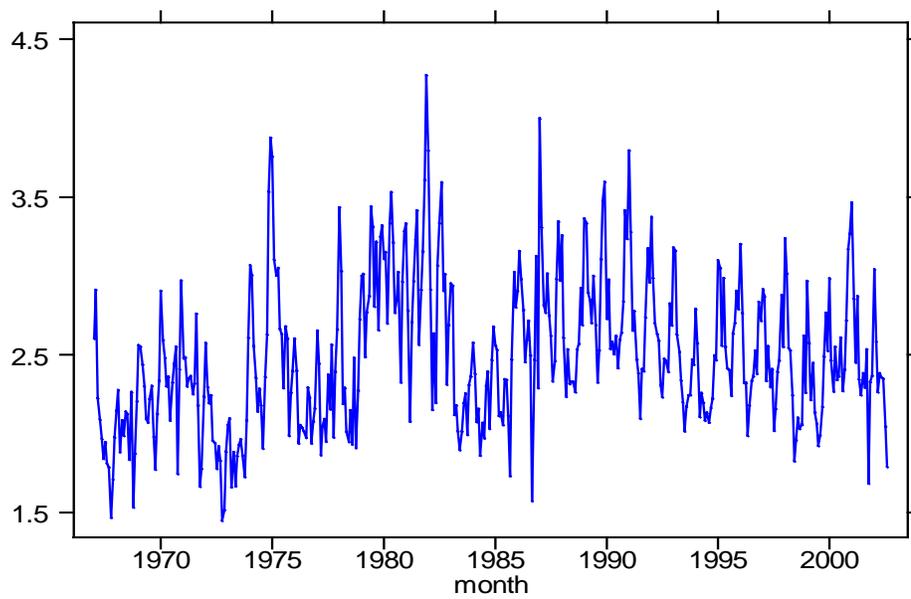
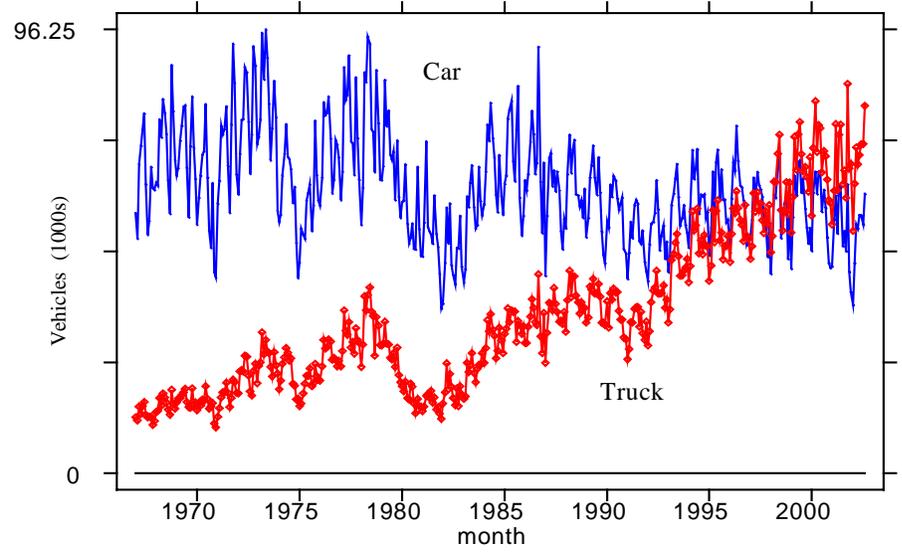


Figure 2: Inventory to Sales Ratio for U.S. Domestic Cars and Trucks
(In Month's of physical units)

U.S. Monthly Car and Light-Truck Sales



U.S. Monthly Car and Light-Truck Production

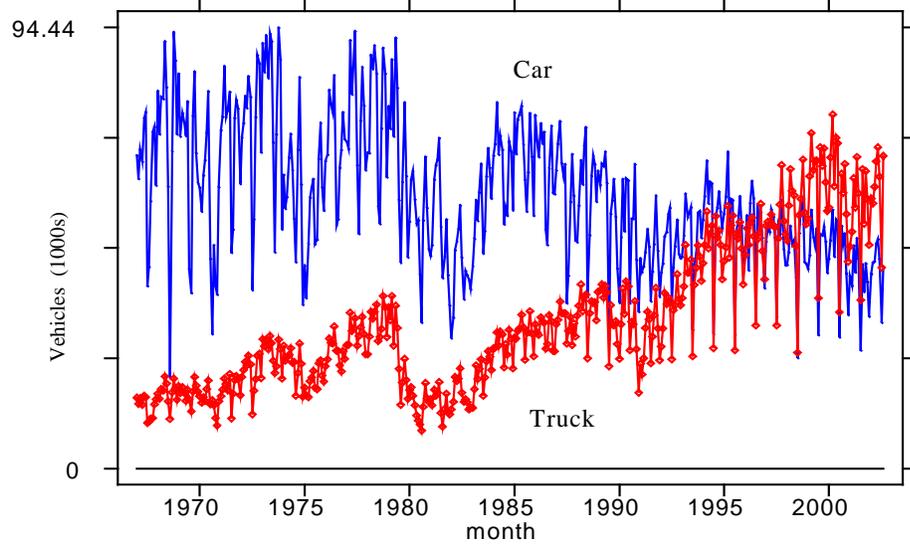


Figure 3: U.S. Monthly Automobile Production and Sales
(Physical Units)

IX. 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 2001.

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?" *NBER Working Paper Series #9459*, January 2003.

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.

Clarida, R., Galí, 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 National Industrial Recovery Act: Evidence on Industry Complementarities," *Quarterly Journal of Economics* 108(4), November 1993, 1043 – 1072.

Feroli, M., "An Equilibrium Model of Inventories with Investment-Specific Technical Change," Working paper, New York University, November 2002.

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.

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

Irvine, O. and Schuh, S., "Inventory Investment and Output Volatility," Working paper, Federal Reserve Bank of Boston, December 2002.

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," *FRBNY 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.

McCarthy, J. and Zakrajšek, E., "Inventory Dynamics and Business Cycles: What has Changed?" *Finance and Economics Discussion Series*, Federal Reserve Board of Governors, June 2003.

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 Vine, D. J., "Why Do Real and Nominal Inventory-Sales Ratios Have Different Trends? 2003a, manuscript.

Ramey, V. A. and Vine, D. J., "The Effect of Sales Persistence on Volatility in a Production Smoothing Model," 2003b, manuscript.

Ramey, V. A. and West, K. D., "Inventories," in *Handbook of Macroeconomics*, eds. John D. Taylor and Michael 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.

-----, "Understanding Changes in International Business Cycle Dynamics," *NBER Working Paper Series* #9859, July 2003a.

-----, "Has the Business Cycle Changed? Evidence and Explanations," August 2003b.

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

Wards Automotive Yearbook, Detroit: Wards Communications.