Working Paper Series

Production and Inventory Control at the General Motors Corporation During the 1920s and 1930s

Anil K. Kashyap and David W. Wilcox

Working Papers Series
Issues in Macroeconomics
Research Department
Federal Reserve Bank of Chicago
May 1992 (WP-92-10)

FEDERAL RESERVE BANK
OF CHICAGO

Production and Inventory Control at the General Motors Corporation During the 1920s and 1930s

Anil K Kashyap and David W. Wilcox*

May 1989
revised: April 1992

This paper develops a rich body of anecdotal evidence on the design and implementation of a production control system at the General Motors Corporation during the 1920s and 1930s. We evaluate that evidence by first modifying the conventional linear-quadratic model of production behavior to take account of annual shutdown, and then testing it using newly available data covering the period 1924 to 1940. On the whole, the model appears to fit the data adequately. GM appears to have been aiming to maintain a targeted level of inventory relative to expected sales, and, secondarily, to smooth production. The production control program was more successful before 1932 than after.

I. Introduction

This paper studies the development and implementation of production control methods at the General Motors Corporation (GM) during the 1920s and 1930s.[1] Especially during the early portion of this period, GM faced large seasonal fluctuations in the demand for its automobiles (much larger, in fact, than it faces today) that seemingly should have provided an ideal opportunity for production smoothing. A rich body of anecdotal evidence shows that GM's senior management thoroughly understood the costs and benefits of production smoothing, and that they implemented a program of production control that was well conceived, in principle, for the purpose of smoothing production relative to sales. We evaluate the success of the program using newly available data on units of production, inventories, and sales.

The next section of the paper reviews the design of the inventory and production procedures in place at GM during the 1920s and 1930s. Much of our information about the development of production control at GM comes from a retrospective study conducted in 1946 by F. Leslie Hayford--a GM economist who had played a leading role in the design of the system and who was later called out of retirement by GM's senior management to describe its main features and evaluate its success. We use Hayford's report together with contemporary public statements of other GM executives to document the essential objectives and operating characteristics of the GM system, and to describe the economic environment in which the company perceived itself to be operating.

Section III turns from the anecdotal record to an analytical model of inventory and production behavior. The discussion in this section focuses on the predictions of the model for the long-run comovements of production, sales, and inventories, and the implications of those comovements for the conduct of hypothesis testing. In Section IV, we modify the basic model by introducing a crude form of annual shutdown. In particular, we explore the implications of specifying that the firm can produce new cars only in selected months of the year, even though it can sell them in all months of the year. The model delivers some useful guidelines as to how we should carry out the empirical work in production data where annual shutdown is important.

Section V uses new data covering the period from 1924 to 1940 to assess empirically the new production control procedures. Much of the analysis rests on a series of charts that demonstrate five important features of the data: First, that the seasonality in sales was much greater during the 1920s and 1930s than it is today, and that GM consequently faced an ideal laboratory for the practice of production smoothing. Second, that as the model predicts, inventories, sales, and production moved together over the longer run. Third, that lumping together operating and shutdown periods can in some respects be quite misleading. Fourth, that despite the rhetoric surrounding the control policy, production seemed to be geared mainly toward maintaining a tight correspondence between inventories and near-term expected sales; production smoothing considerations appear to have influenced the behavior of production only at the margin. Finally, that the behavior of production was different after 1932 than it had been before 1932. We also report a variety of Euler equation estimates from the model. On the whole, the model estimates corroborate the graphical evidence; we interpret the exceptions to this pattern as reflecting mainly on the small-sample properties of the econometric estimates rather than on the economic content of the model.

The last section of the paper presents conclusions.

II. Production Control and Planning in General Motors: 1921-1946[2]

General Motors entered the 1920s on a tide of optimism. Believing that the post-war surge in demand would be sustained, the corporation accumulated large stocks of raw materials and made unusually large commitments to buy more (Hayford, p.2). The economy slowed sharply, however, leaving the corporation in a precarious position and leading some of its officers to push for a new program to control production more closely:

> At the beginning of [1921] material commitments were abnormal, inventories greatly over-expanded, and some $80 million had been borrowed from banks. Financial control and coordination of the Corporation's activities were imperative. (Donaldson Brown, Vice President in Charge of Finances, as quoted in Hayford, p.4)

The first step was to begin collecting better data. Early in 1921, Brown appealed to Alfred Sloan (then Operating Vice President and later President of the corporation) for help in obtaining data on stocks in the hands of dealers and distributors (Hayford, p.8-9). Later in the same year, Brown raised the possibility of also collecting data on deliveries by dealers to retail customers. Although Brown's efforts initially were resisted by others within the corporation, eventually he prevailed, and by October 1921 the divisions were reporting "production and factory sales by ten-day periods; and, as of the end of each month, orders on hand and stocks at divisions and in dealers' hands" (Hayford, p.13). Hayford further reported that "although dealers did not begin reporting their deliveries to consumers until 1925, the Central Office and the divisions were making use of 'calculated' deliveries to consumers figures long before that time" (Hayford, p.15). Indeed, GM began publishing monthly data on retail sales in July, 1924.

During late 1923 and early 1924, inventories once again piled up because the divisions had been overly optimistic about sales prospects. Whereas the 1920-21 downturn had spurred the corporation to revise its policy on purchases of materials, this second crisis caused them to act on controlling inventories of finished cars, and by mid-1924 the new control system was in place (Hayford, p.28). In the opinion of the creators of the new control program, its distinguishing characteristic was that it tied production schedules over the forecast period (current month plus three months following) to anticipated deliveries to consumers for the model year as a whole (Hayford p.2). Fluctuations in the business cycle were taken into account only insofar as they were expected to influence current-model-year sales; model changeover effectively limited the planning horizon to one year.

Under the new production control program, the divisions were required to file with the Central Office a monthly "Analysis of Production Requirements." The function of the Analysis was to record each of the elements of the inventory accumulation identity: stocks on hand at the beginning of the forecast period, projected retail deliveries during the forecast period, desired stocks at end of forecast period, and--as a residual--the "indicated maximum production required." The divisions were encouraged to give close and frequent scrutiny to their sales projections. In the words of the 1924 Procedure (the corporate handbook that detailed the mechanics of the control program), these projections "should be subject to constant consideration on the part of the division and subject to change whenever warranted by conditions of any kind" (Hayford p.43).

A key element of these monthly production statements from the divisions was the target level for stocks at the end of the forecast period. In the 1924 Procedure, this target level was defined as being "calculated to provide a sufficient number of cars or trucks to meet the requirements of Deliveries to Consumers with the plants operating on the so-called 'level production' basis." In turn, "level production" was to be calculated by assuming "'that 8.7 percent of the year's production for domestic requirements will be produced in each month of the year except December, for which month 4.3 percent is assumed, to allow for inventory taking.'" (Hayford, p.39, quoting from the Procedure of August 1924.)

Alfred Sloan summarized the objectives of the new system in the Annual Report to the stockholders for 1924:

> During 1924 the Corporation adopted a production policy as affecting stocks of finished cars which its dealers and distributors will be expected to carry. This policy is predicated upon the sale of cars to consumers as a fundamental index. Such sales are subject to seasonal fluctuations, and the merchandising policy of the Corporation requires that dealers and distributors shall accumulate stocks during seasons of relatively low retail deliveries in order to facilitate prompt deliveries in seasonal periods of heavy retail demand as well as to maintain manufacturing and distributing economies afforded by a reasonably level rate of production. The amount of such stocks varies with the seasons of the year and is based upon a careful analysis of the trend of retail demand... It is believed that the Corporation in the future will be free from the evils resulting from excess accumulation of stocks involving unnecessary storage, interest and carrying charges as well as drastic curtailment of production schedules such as have occurred at times in the past. (p.9)

In 1926, Albert Bradley--then assistant treasurer of the corporation and later chairman of the board--described the potential benefits of the control program in an address before the American Management Association:

> If this problem [the marked seasonality in demand] is solved by adjusting production to correspond closely with sales two major objections arise. Productive capacity required to take care of April's sales, for example, would be twice that needed if production were evenly distributed throughout the year, thus increasing capital investment, depreciation, taxes, and similar items and making necessary much larger earnings in order to provide an equivalent return on the investment. Labor offers a second objection to this course, for not only are workmen entitled to steady employment which might be impossible with such wide fluctuations in output, but it is unlikely that the necessary amount of labor would be available when needed for peak production and it is certain that a factory operating under such conditions would be hard pressed to keep its skilled workers, so that the quality of the product might suffer (Bradley as quoted in Automotive Industries, March 18, 1926 p.489).

Such statements notwithstanding, GM executives recognized that any move toward production smoothing would entail certain costs. Indeed, Bradley himself went on to say:

> Flattening out the rate of production and building up stocks of cars against future heavy demand also has its drawbacks. Such storage requires additional capital with its interest charge, greater insurance expense and similar items. If sales do not come up to expectation the cost of carrying stocks may be unduly increased through prolonging the storage period and further losses may result from the necessity of forced selling of the excess stock. (Automotive Industries, p.489)

We find these statements significant because they highlight several elements of cost that often feature prominently in modern-day formulations of the production scheduling problem.

The procedure as described above was kept in place with only minor changes throughout the remainder of the 1920s. However, the auto industry was particularly hard hit during the Great Depression. New car registrations fell more than 2.7 million units (over 70 percent) between 1929 and 1932 (Survey of Current Business, Annual Supplements, 1932 and 1938), compared with a 28 percent decline in real GNP over the period. GM fared little better than the rest of the industry, suffering a 66 percent reduction in retail deliveries of autos. These developments caused GM to reconsider the design of the control program:

> In view of the instability of economic conditions in the early 1930s and the resultant greater difficulty in appraising consumer demand, it was recognized that sudden and considerable changes in forecast might become necessary. (Hayford, p.87)

In May, 1932, a new Procedure was issued.

In tone, the new Procedure was similar to previous ones, the operating divisions being instructed to set production "giving consideration to the best estimate of deliveries to consumers for the complete model sales year and having regard to the extreme desirability of running at as level a rate of production as practicable" (Hayford, p.85). In practice, however, the revised program differed importantly from the ones that had been in force during the 1920s. Under the revised program, the Central Office no longer transmitted to the divisions a "Preliminary Analysis of Production Requirements" which, under the earlier program, had been a starting point for the divisional decisionmaking process. And, in recognition of the greater uncertainty about market conditions, the new Procedure specified that the divisions were required to notify the Central Office only "whenever a change in the outlook necessitates a change of ten percent or more in the total Production Schedule for the entire Forecast period" (Hayford, p.87).

Another factor that may have limited the apparent effectiveness of the production control program during the 1930s was a reduction in the amplitude of the seasonal swings in sales. In part, this reduction undoubtedly reflected the increasing mechanical reliability of newly manufactured cars, the gradual disappearance of open cars, and the increased extent of paved roads. In addition, the reduction in seasonal amplitude probably reflected the coordinated shift of the shutdown period from the turn of the year to the late summer. The Automobile Manufacturers Association proposed in 1933 "that the industry as a whole shift its new model announcement dates to the fall of the year for the purpose of attaining a greater regularity of production and employment" (Hayford, p.59). The Roosevelt Adminstration supported the plan, and between 1934 and 1935 the introduction dates were moved forward by several months (see Cooper and Haltiwanger (1991)).

Finally, labor relations became more difficult during the late 1930s. The problems reached the boiling point in 1937, when a sitdown strike by GM workers effectively closed most GM plants in late January and early February. Aside from these labor disputes, however, we found no mention (either in Hayford's report or elsewhere) of transitory fluctuations in costs.

III. The Basic Model

Our baseline model is similar to one proposed by Holt et al. (1960) and subsequently studied by many others.[3] We assume that the firm minimizes expected cost as given by:[4]

$$(1)\quad \min E_{t-1} \sum_{j=0}^{\infty} \beta^{j}[\alpha_0(\Delta Q_{t+j})^2 + \alpha_1(Q_{t+j})^2 + \alpha_2(H_{t+j} - \alpha_3 S_{t+j+1})^2]$$

$$\text{s.t. } H_{t+j} = H_{t+j-1} + Q_{t+j} - S_{t+j}$$

$$H_{t-1} \text{ given}$$

where β is a discount factor, H_t is the end-of-period stock of finished cars, Q_t is production, and S_t is sales. We assume that the firm makes its production decision before current-period sales are known. The three components of the cost function all are standard from the inventory literature, and are consistent with the anecdotal evidence provided above.

The first-order condition necessary for cost minimization is given by:

$$(2) \quad E_{t-1}\{\alpha_0(\beta^2\Delta Q_{t+2}-2\beta\Delta Q_{t+1}+\Delta Q_t) + \alpha_1(-\beta Q_{t+1}+Q_t) + \alpha_2(H_t-\alpha_3 S_{t+1})\} = 0.$$

In this baseline model, a similar equation holds for all periods "t".[5]

A. Identification

The parameters $\{\alpha_0, \alpha_1, \alpha_2\}$ in equation (2) are not separately identified. We achieve identification by requiring that the strong form of the Legendre-Clebsch condition be satisfied. The Legendre-Clebsch condition is an additional condition that is necessary for optimality; it states that the second derivative of the objective function with respect to the choice variable must be strictly positive (Stengel (1986) page 213). In our model, the Legendre-Clebsch condition is given by:

$$\alpha_0(1+4\beta+\beta^2) + \alpha_1(1+\beta) + \alpha_2 > 0.$$

We guarantee that this condition will be satisfied by imposing:[6]

$$(3) \quad \alpha_0(1+4\beta+\beta^2) + \alpha_1(1+\beta) + \alpha_2 = 1.$$

We also impose *a priori* the value of the discount factor β, in line with the suggestion made by Gregory, Pagan and Smith (1989), who discuss the difficulty in identifying β in this class of models, and following a long tradition in this literature.

B. Multicointegration and the importance of allowing α_3 to be non-zero

Equation (2) can be rewritten as follows:

$$(2') \quad E_{t-1}\{\alpha_0(\beta^2\Delta Q_{t+2}-2\beta\Delta Q_{t+1}+\Delta Q_t) - \alpha_1\beta(\Delta H_{t+1}+\Delta S_{t+1})$$

$$+ \alpha_1\Delta H_t - \alpha_2\alpha_3\Delta S_{t+1} + \alpha_2[H_t - (\alpha_3 - \frac{\alpha_1(1-\beta)}{\alpha_2})S_t]\} = 0.$$

This form of the equation highlights that H_t and S_t (and hence Q_t and S_t) must be cointegrated.[7] Hence, our analytical model predicts that H, Q, and

S should exhibit the type of long-run interrelationships labelled "multicointegration" by Granger and Lee (1988a, 1988b).

Equation (2′) also highlights the importance of allowing α_3 to differ from zero. As West (1986) noted, some authors have argued that α_3 should be set equal to zero. Equation (2′) makes clear, however, that the cointegrating parameter would be negative if α_3 were set equal to zero (provided α_1 and α_2 are positive), in which case a permanent increase in sales would induce a permanent *decrease* in inventories.[8] In fact, a common feature of actual data is that inventories and sales covary positively in the long run as well as in the short run.

C. Inference

Most inference in our model can be carried out using standard asymptotic distributions even if the stochastic process for sales contains a unit root.[9] The argument revolves around showing that the variables attached to the remaining parameters (after the identifying assumption has been imposed) are cointegrated. To verify that this condition obtains in our model, we drop the expectations operator from equation (2), and eliminate α_2 using equation (3) (similar arguments hold for other identifying assumptions):

$$(2'') \quad \alpha_0(\beta^2\Delta Q_{t+2}-2\beta\Delta Q_{t+1}+\Delta Q_t) + \alpha_1(-\beta Q_{t+1}+Q_t)$$
$$+ [1-\alpha_0(1+4\beta+\beta^2)-\alpha_1(1+\beta)](H_t-\alpha_3 S_{t+1}) = \xi_t$$

Then we calculate the derivative of the expectational error, ξ_t, with respect to each of the parameters remaining in the model:

$$\eta_{0t} = \frac{\partial\xi_t}{\partial\alpha_0} = (\beta^2\Delta Q_{t+2}-2\beta\Delta Q_{t+1}+\Delta Q_t) - (1+4\beta+\beta^2)(H_t-\alpha_3 S_{t+1})$$

$$\eta_{1t} = \frac{\partial \xi_t}{\partial \alpha_1} = -\beta Q_{t+1}+Q_t - (1+\beta)(H_t-\alpha_3 S_{t+1})$$

$$\eta_{3t} = \frac{\partial \xi_t}{\partial \alpha_3} = - [1-\alpha_0(1+4\beta+\beta^2)-\alpha_1(1+\beta)]S_{t+1}$$

It is easy to verify that v_t is stationary, where

$$v_t = \eta_{0t} - [\frac{1+4\beta+\beta^2}{1+\beta}]\{\eta_{1t} + [\frac{1-\beta}{1-\alpha_0(1+4\beta+\beta^2)-\alpha_1(1+\beta)}]\eta_{3t}\}.$$

Thus, η_{0t}, η_{1t}, and η_{3t} are cointegrated. Given this fact, results from Sims, Stock, and Watson (1990) and West (1988) can be applied to show that the estimated α's will be asymptotically normal, and that Hansen's (1982) J-statistic will be asymptotically chi-square.

This analysis shows that if the stochastic process followed by sales is the only source of nonstationarity, then there is no need to include time trends in the specification of the model, regardless of whether sales are trend- or difference-stationary. However, if there is some other source of a deterministic trend in the data--say, a deterministic trend in costs unrelated to production--then deterministic trends should be included in the specification of the Euler equation. Even in that case, asymptotic arguments similar to the one given above will be valid (again, see Sims, Stock, and Watson (1990) and West (1988)). Therefore, previous results derived from equations that included deterministic trends should be valid asymptotically, even if the only source of nonstationarity in fact was a stochastic trend in the sales process.[10] We carry out our estimation and hypothesis testing under the assumption that deterministic trends need not be included in the equation.

The expectational error that results from dropping the expectations operator from equation (2) is MA(1). We estimate the coefficients and their standard errors using GMM with a Newey-West (1987) covariance matrix.[11] Our RATS code and a complete dataset are available from the authors upon request.

IV. The Model With Annual Shutdown

Thus far, we have followed earlier authors in assuming that the manufacturer is producing every period. We view this assumption as particularly unsatisfactory for the automobile industry: After all, even during the 1920s and 1930s GM closed its assembly plants each year for inventory-taking and retooling. We take a small step toward reality in our theoretical model by assuming that the manufacturing plant shuts down once a year. We treat this shutdown period as exogenous to the model, as if it were, say, constitutionally imposed upon the company as a condition of incorporation.[12] We maintain the assumption that sales are made throughout the year, including during the shutdown period. We further assume that the manufacturer pays a fixed cost to open or shut a plant, and does not, in opening one up or shutting one down, bear the cost usually associated with changing production.

In most months, equation (2) remains the relevant first-order condition despite the explicit introduction of annual shutdown into the problem. When shutdown is just past or is imminent, however, a modified first-order condition is relevant. For example, in the first period following a shutdown, the relevant first-order condition is given by:

$$(4) \quad E_{t-1}\{\alpha_0(\beta^2\Delta Q_{t+2}-2\beta\Delta Q_{t+1}) + \alpha_1(-\beta Q_{t+1}+Q_t) + \alpha_2(H_t-\alpha_3 S_{t+1})\} = 0,$$

where the absence of the term in ΔQ_t reflects our assumption about the cost of re-opening a plant.

In the last period of operation before a one-month shutdown period, the following first-order condition obtains:

$$(5)\quad E_{t-1}\{\alpha_0(\beta^3\Delta Q_{t+3}-2\beta\Delta Q_{t+1}) + \alpha_1(-\beta^2 Q_{t+2}+Q_t) + \alpha_2[(H_t-\alpha_3 S_{t+1}) + \beta(H_{t+1}-\alpha_3 S_{t+2})]\} = 0;$$

while in the penultimate month of production before shutdown the following first-order condition is relevant (regardless of how long the shutdown is expected to last):

$$(6)\quad E_{t-1}\{\alpha_0(-2\beta\Delta Q_{t+1}+\Delta Q_t) + \alpha_1(-\beta Q_{t+1}+Q_t) + \alpha_2(H_t-\alpha_3 S_{t+1})\} = 0.$$

Modifications of equation (5) relevant for two- and three-month shutdowns are easily derived. Of course, we recover the condition $Q_t=0$ for months in which the plant is not operating.

Thus, the augmented model predicts that the specification of the first-order condition for any given month will depend on the orientation of that month with respect to shutdown periods, both preceding and following. Many strategies for dealing this with situation are possible; we pursue two of them here. The first involves simply dropping the shutdown-contaminated observations from our sample, and estimating the cost parameters using the remaining observations. The second involves retaining all non-shutdown observations in the sample, and applying to each observation the appropriate form of the Euler equation, enforcing that α_0, α_1, α_2, and α_3 are the same in all specifications.[13] The advantage of the second approach is that it attempts to exploit the information in the observations coming just after and just before shutdown periods, and therefore holds out the possibility of a gain in statistical efficiency. The risk in this approach, however, is that we will not be able to incorporate that additional information into our model in a valid fashion, and that in expanding the sample we may primarily be increasing the specification error in the equation rather than augmenting its statistical efficiency. One troubling indicator in this regard is that we never observe production to be zero in any month in our sample, contrary to the key assumption in our shutdown-augmented model.[14] On balance, therefore, we favor the first approach, but we implement both of them in the next section.

V. Empirical Results

We begin with five charts that highlight the main features of the data.[15] *First, monthly sales during the 1920s and 1930s were much more variable than they are today. Especially during the 1920s and early 1930s, the bulk of this variation was accounted for by regular seasonal fluctuations.* Chart 1 demonstrates these points by showing monthly domestic unit sales by all GM divisions as a percent of a twelve-month centered moving average for the periods 1924-1940 (top panel) and 1971-1987 (bottom panel). In the 1920s, the pace of sales during the busiest month of the year (generally March or April) frequently was more than three times as great as the sales pace during the slowest month (generally November or December). In the mid and late 1930s, the amplitude of the variation in sales declined somewhat, and the seasonal pattern became less regular. Even in that later period, however, the amplitude of the variation in sales was enormous compared with more recent experience: Since 1971, sales in the busiest months have rarely been even twice as great as sales in the slowest months (bottom panel).

Second, the long-run comovements between inventories, sales, and production conform to the predictions of the model. Chart 2 abstracts from short-run (especially seasonal) fluctuations by plotting twelve-month centered moving averages. The top panel shows sales and production, and leaves no doubt that those two variables are cointegrated and that the cointegrating parameter equals 1. The bottom panel shows sales against a linear transformation of inventories (where the linear transformation was chosen in line with one of our preferred sets of estimates reported below), and shows that inventories and sales also appear to have moved together over the long term. The relationship between these two variables was not as tight as the one between production and sales. However, our confidence in the validity of the underlying relationship is bolstered by the fact that all of the important deviations from the equilibrium relationship between inventories and sales have obvious economic interpretations: On the upside, inventories were unusually high during 1924 (the very episode that caused GM to institute the new program of production control), the first three years of the Great Depression, and the recession of 1938. On the downside, inventories were unusually low during 1928--at the peak of the expansion--and in the winter of 1936-1937, when labor strife culminated in the sitdown strike of 1937.

Third, the behavior of production during months we identify as "operating" months differed markedly from the behavior of production during the year as a whole. Chart 3 compares the variance of production with the variance of sales for each model year on two different bases: first using data from all months of each model year (the dashed line), and second using data from only the operating months of each model year (the solid line).[16] In more than half the years between 1925 and 1940, the variance ratio calculated over operating months was at least 30 percent lower than the variance ratio calculated over all months. Only once--in 1938--was the variance ratio for operating months higher than the variance ratio for all months.[17]

Fourth, contrary to the anecdotal evidence presented above, GM appears to have arranged its production chiefly to maintain a reasonable short-run correspondence between the level of inventories and expected sales. At the margin, production appears to have been a bit smoother than it would have been if GM had aimed to hold the inventory-sales ratio fixed to the exclusion of all other factors. Chart 4 provides evidence in support of the first part of this proposition. It plots the unfiltered levels of production and sales in the top panel and inventories and sales in the bottom panel. The bottom panel shows that inventories tracked sales not only at business-cycle frequencies, as was evident in Chart 2, but also at seasonal frequencies, suggesting that one of GM's objectives was to tie the level of inventories to the volume of sales within each model year as well as between model years. Chart 5, which shows the inventory-sales ratio (H_t/S_{t+1}) as the solid line and production as the dotted line, supports the second part of the proposition. In every year between 1925 and 1931, the inventory-sales ratio hit its seasonal low when production was at or near its high, consistent with the conclusion that production was smoother than it would have been if GM had been aiming to hold the inventory-sales ratio constant at some fixed level year-round.

Finally, the behavior of production was different after 1932 than it had been before 1932. This difference is evident in Chart 5: In some years after 1932, the correlation between the inventory-sales ratio and production appears to have been about zero; in other years (such as 1939 and 1940), it appears to have been strongly positive. Overall, it is much more difficult to make the case that production after 1932 was smoother than it would have been under a policy of fixing the inventory-sales ratio. Chart 3 provides another perspective on the change in the behavior of production in 1932: In the eight years after 1932, the variance ratio measured over operating months exceeded 1 four times; by contrast, in the eight years up to and including 1932, the variance ratio exceeded 1 only once, in 1932. We view both these pieces of evidence as suggesting that some factor or factors inhibited the apparent effectiveness of the production control program after 1932. Such factors could have been either internal or external to the firm. Two obvious candidates in this regard are the 1932 modification in the production control program and the increased difficulty of labor relations.

We now report the results of estimating the parameters appearing in the first-order conditions for cost minimization. Our goal here is to develop additional evidence either corroborating or contradicting the graphical evidence already presented. Table 1 summarizes nine sets of results, reflecting three different methods of handling the influence of annual shutdown and three different sample periods. With regard to the treatment of shutdown periods, the first method involves simply ignoring any special considerations related to the handling of the shutdown periods and applying the basic first-order condition (equation (2)) to all observations in the sample. The second method is our preferred one; it involves dropping all shutdown-contaminated observations from the sample and applying equation (2) to the remaining observations. In line with the results from our theoretical model, we treat an observation as "shutdown-contaminated" if it was either (a) a shutdown month; (b) a month that immediately followed a shutdown period; or (c) a month that preceded a shutdown period by either one or two months. The third method involves dropping only the shutdown months themselves, and selecting for each of the remaining observations in the sample the appropriate specification of the first-order condition from among equations (5)-(7) above, and the modification of (6) for multi-period shutdowns. As for sample periods, we implement each of these three methods over: the full sample period, 1925:1-1940:12; the period before any major changes were made in the original program, 1925:1-1932:5; and the remainder of the sample period, 1932:6-1940:12. All nine sets of results shown in Table 1 were derived using equation (3) as the identifying assumption.[18]

Several features of the results in Table 1 are worth highlighting briefly. First, the direct cost of *producing*, as reflected in α_1, is estimated to be negative in every specification, and significantly so in several.[19] Nonetheless, the estimated slope of the marginal cost curve, given by $\alpha_0(1+\beta)+\alpha_1$, is positive in all specifications, reflecting the positive estimated cost of changing the rate of production. Second, the cost of deviating from the target level of production also is positive in every case, but significantly so only over the early sample period. Third, the target inventory-sales ratio is positive and significant in the early sample period, but insignificantly different from zero in the late sample. (In fact, when we use either the second or third method of handling the shutdown periods, the point estimate of α_3 is negative.) Overall, the model seems to do an adequate job of describing the data, as evidenced by the fact that the residuals from these specifications look to be well-behaved. As an example, we plot the residuals from line 1b in the top panel of Chart 6. (The negative autocorrelation at the first lag evident in the residuals shown in the top panel does not contradict the assumption that they follow an MA(1) process.)

Table 2 provides additional results based on alternative assumptions about various features of the model. The first three sets of estimates reported there were derived using our preferred method of handling the shutdown periods. Lines 1 and 2 examine the ability of the data to distinguish between costs of *changing* production and direct costs of *producing* as sources of production smoothing. Like Blanchard (1983), we find that the model is about equally happy with either alternative: In both cases, the estimated slope of the marginal cost curve remains positive and highly significant. Line 3 shows the results of fixing β at 0.99 rather than 0.995; not surprisingly, this perturbation makes hardly any difference.

The last three lines in Table 2 report the results of using $\alpha_2=1$ as the identifying assumption rather than equation (3), and handling shutdown periods by the first method (that is, ignoring their presence).[20] Recall that, asymptotically, the specification of the identifying assumption should have no effect on either the nature of the decision rule implied by the parameter estimates, or the overall statistical adequacy of the model. In our sample, however, the choice of the identifying assumption turns out to matter a great deal. As shown on line 4b of Table 2, α_0 is estimated to be negative in the early sample period under the alternative identifying assumption, and the estimated slope of the marginal cost curve is negative in all three samples. Clearly, the economic implications of the alternative estimates differ dramatically from those of the baseline estimates.

In our view, however, the alternative estimates should be discounted for three reasons: First, the estimated residuals associated with those estimates are so highly autocorrelated as to suggest that the variables in the model are not cointegrated. To show this, we plot the residuals from line 4b in the bottom panel of Chart 6; the contrast with the residuals from the baseline specification is stark. The Durbin-Watson statistic for these residuals is 0.22--well below the level that would be required to reject the null of nonstationarity. A second reason for taking a skeptical view of the alternative estimates is that they fly in the face of the graphical evidence presented earlier, which pointed toward production smoothing, at least before 1932. By contrast, the baseline results reported in Table 1 were fully consistent with this evidence. Finally, we note that Krane and Braun (1991) also experimented with the $\alpha_2=1$ normalization, and found that it delivered peculiar results. For example, they report that the Legendre-Clebsch condition was violated in about a third of the industries they studied when they used $\alpha_2=1$ as their identifying assumption; by contrast, they found no such violations whey they achieved identification by setting $\alpha_0(1+\beta)+\alpha_1=1$. On balance, we view the alternative results as more of a puzzle with regard to their econometric implications--that the asymptotic approximations of the statistical distributions are very poor--than with regard to their economic implications. As a result, we focus the remainder of our discussion on the baseline econometric estimates, which we view as being both more statistically defensible, and consistent with the graphical evidence.

In particular, we now briefly revisit the questions posed earlier with respect to shutdown, the possible break in the data in 1932, and the propensity to smooth production. First, with regard to possible differences in the behavior of production during normal operating months versus shutdown periods, we find the Euler equation evidence frankly surprising. Although the relative variability of production and sales differed substantially depending on the treatment of the shutdown periods, the Euler equation estimates appear to be affected only a little. The most that can be said, in our view, is that the over-identifying restrictions are rejected less vigorously under the more cautious method of treating shutdown (lines 2a, 2b, and 2c of Table 1) than they are when shutdown is ignored. Given the substantial difference in the available number of observations under these two methods, though, we would hesitate to read too much into even this finding. Overall, while we continue to believe that the handling of shutdown can be important for some questions, it does not appear to be universally important.

On the question of whether there was a break in the data around 1932, the message from the various pieces of evidence is much more consistent. First, the autocorrelation structure of the disturbances appears to have changed between the early and late subsamples. For the residuals underlying line 2b (estimated from the pre-1932 sample), the correlogram is as follows:

lag:	1	2	3	4	5	6
correlation:	-.36	-.08	.13	-.04	.04	-.03

For the residuals underlying line 2c (estimated from the post-1932 sample), the correlogram is:

lag:	1	2	3	4	5	6
correlation:	-.43	-.17	.18	.04	.02	-.19

In both cases the asymptotic standard error is 0.13. Over the early sample period there is no evidence that would reject the view that the disturbance follow the hypothesized MA(1) process. However, the post-1932 disturbances betray greater evidence of correlation at longer lags, consistent with the view that cost shocks may have become more important after 1932 (see Eichenbaum (1989), among others, on the implications of cost shocks for models of the type we study).

Second, the estimates of α_3 derived from the later subsample are less plausible than the ones derived from the early subsample (especially the negative estimates reported on lines 2c and 3c). Third, and perhaps most importantly from an economic perspective, the decision rules implied by the early- and late-sample parameter estimates are considerably different. Below we report some simple simulations that display these differences clearly.

Finally, we return to the keystone question of whether GM was a production smoother or a production buncher. The implications of the parameter estimates for GM's propensity to smooth production are difficult to discern given the complexity of the model (especially with regard to its predictions for the dynamic response of production to a shock in sales). To illuminate the extent to which GM may have been acting as a production smoother, we conduct a simulation exercise: We transport GM to a simplified world (where there is no annual shutdown and where sales follow a random walk), solve for the decision rule in that simplified world, and then simulate GM's response to a one-unit shock to sales given a particular set of parameter values.[21] Needless to say, we are exploiting to the maximum degree the structural interpretation of our cost parameters.

The top panel of Chart 7 shows the outcome of this exercise when we use our preferred coefficients from the early period (line 2b of Table 1). The solid line represents the hypothetical sales trajectory. Sales are assumed to have been constant at 100 units per period until period 0, at which time they are assumed to have jumped permanently to 101 units. The dashed line represents the production trajectory given the parameter estimates from line 2b. Prior to period 0, production and inventories are assumed to have been at their equilibrium levels. In period 0, production remains at 100 units because the increase in sales was unanticipated; as a result, inventories decline by one unit in that period. Starting in period 1, however, production begins to follow a hump-shaped path. During much of the adjustment period, production exceeds 101 units per period, and, in this sense, production is more variable than sales. Over the entire adjustment period (including period 0), the cumulative excess of production over 101 units is sufficient not only to offset the initial inventory drawdown of 1 unit but also to build up inventories to their new higher equilibrium level.

The dotted line shows what the manufacturer's response would have been if there had been no incentive to smooth production. (We compute this path by setting $\hat{\alpha}_0$ and $\hat{\alpha}_1$ equal to zero.) In this no-smoothing world, the entire production adjustment occurs in period 1, and the manufacturer builds enough cars all at once to boost inventories to their long-run level. Clearly, production in the no-smoothing world is more variable than it is when α_0 and α_1 are at their estimated values. We conclude that, in this sense, over the early portion of our sample period, GM was a production smoother.

The bottom panel repeats the exercise using the coefficients we estimated from the late sample, shown on line 2c of Table 1. These parameters imply hardly any smoothing relative to the path that the producer would have chosen if α_0 and α_1 had been equal to zero. The simulation exercises show that seemingly slight differences in parameter values can imply important differences in economic behavior.

On balance, we read the bulk of the evidence as suggesting that GM was engaged in some production smoothing between 1925 and 1932. The message of the data after 1932, however, is less clear: Production appears to have been no smoother than it would have been under a rule tying inventories to near-term expected sales. On the other hand, the corporation did succeed in navigating through extremely turbulent times while avoiding a repetition of the nearly-disastrous inventory accumulation of 1923-24.

VI. Conclusion

In 1924, General Motors implemented a new production control procedure that was intended to "regularize" production and employment. In a report to the Executive and Operations Committees of the corporation, Donaldson Brown described the program as "permit[ing] an accumulation of stock during the period of the year when sales to consumers are below the average rate, and requir[ing] a liquidation of stock during the period of the year when sales to consumers are above the average rate" (Hayford, p.42). A simple plot of monthly inventories and sales suggests that, measured against Brown's yardstick, the program was a failure: During the period we study, GM usually accumulated inventory in the spring, when sales were on their seasonal upswing, and decumulated inventory in the late summer and fall, when sales were tailing off.

Further examination of the data, however, reveals considerable success in smoothing production relative to a different yardstick. A plot of production and the inventory-sales ratio shows that, during the early years of the production control program, GM drove its inventory-sales ratio down during the time of the year when production was at its highest, and allowed its inventory-sales ratio to spike up when production was at its lowest. Moreover (and to our surprise), the data contain some hints of smoothing of this type even at business-cycle frequencies: Inventories were taken down to unusually low levels relative to sales at the peak of expansion at the end of the 1920s, and were allowed to accumulate to unusually high levels, relative to sales, during the Great Depression. Formal statistical evidence, based on a modified linear-quadratic model of production behavior, largely corroborates the graphical evidence. We conclude that GM did succeed, during the 1920s and early 1930s, in making production smoother than it would have been if the corporation had only sought to stabilize the inventory-sales ratio.

In 1932, however, the corporation revised the structure of the program. We also adduce indirect evidence that cost shocks became more important. In the second half of our sample, there is little or no evidence of production smoothing, even by our more expansive definition.

We close with three summary points: First, the fact that production tracked sales does not cast doubt on the standard linear-quadratic model. Indeed, as we noted above, the model itself predicts that production and sales will be cointegrated.[22] Second, the fact that production tracked sales does not, by itself, imply that GM faced a flat or declining marginal cost curve. Indeed, our preferred parameter estimates imply that GM faced an upward-sloping marginal cost curve and yet still chose to make production mimic sales closely. Overall, our interpretation of the evidence relies heavily on the importance of the inventory-sales target as a determinant of GM's behavior, and reinforces the importance of research such as Kahn's (1992) that seeks to provide a firmer microfoundation for the observed inventory-sales targeting. Third, annual shutdown influenced GM's production behavior importantly. As Cooper and Haltiwanger (1990) note, the type of machine replacement that often motivates annual shutdown is common to many industries. We see promising possibilities in further study of the influence of machine replacement on the dynamics of production.

References

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

Blinder, Alan S. (1986) "Can the Production Smoothing Model of Inventory Behavior be Saved?" **Quarterly Journal of Economics** 101 431-454.

Cooper, Russell and John Haltiwanger (1990) "The Macroeconomic Implications of Machine Replacement: Theory and Evidence", unpublished mimeo, University of Maryland.

______________________ (1991) "Autos and the National Industrial Recovery Act: Evidence on Industry Complementarities," unpublished mimeo, University of Maryland.

Eichenbaum, Martin (1989) "Some Empirical Evidence on the Production Level and Production Cost Smoothing Models of Inventory Investment," **American Economic Review**, 79(4) 853-864.

Fair, Ray C. (1989) "The Production Smoothing Model is Alive and Well," **Journal of Monetary Economics** 24 353-370.

Forbes (1977) "1917, 1929, 1945, 1966, 1977: Sixty Years of Corporate Ups, Downs, and Outs," September, 127-60.

Granger, Clive W. J., and Tae-Hwy Lee (1988a) "Multicointegration", unpublished manuscript, University of California, San Diego.

______________________ (1988b) "Investigation of Production, Sales and Inventory Relationships Using Multicointegration and Nonsymmetric Error Correction Models", unpublished manuscript, University of California, San Diego.

Gregory, Allan W., Adrian R. Pagan and Gregor W. Smith (1989) "Estimating Euler Equations from Linear Quadratic Models", unpublished manuscript, Rochester University.

Hayford, F. Leslie (1946) "Production Control and Planning in General Motors: 1921-1946," unpublished manuscript, General Motors Corporation, Detroit.

Hansen, Lars P. (1982) "Large Sample Properties of Generalized Method of Moments Estimators," **Econometrica** 50(4) 1029-1054.

Holt, Charles F., Franco Modigliani, John F. Muth, and Herbert A. Simon (1960) **Planning, Production, Inventories and the Work Force**, Englewood Cliffs: Prentice-Hall.

Kahn, James A. (1992) "Why is Production More Volatile than Sales? Theory and Evidence on the Stockout-Avoidance Motive for Inventory Holding," forthcoming, **Quarterly Journal of Economics**.

Krane, Spencer D. (1991) "Induced Seasonality and Production Smoothing Models of Inventory Behavior," forthcoming, **Journal of Econometrics**.

____________, and Steven N. Braun (1990) "Production Smoothing Evidence From Physical-Product Data," **Journal of Political Economy** 99(3) 558-581.

Modigliani, Franco, and Franz E. Hohn (1955) "Production Planning Over Time and the Nature of the Expectation and Planning Horizon," **Econometrica** 23(5) 46-66.

___________________ and Owen H. Sauerlander (1955) "Economic Expectations and Plans of Firms in Relation to Short-Term Forecasting," in **Short-Term Economic Forecasting**, Studies in Income and Wealth, Vol. 17, National Bureau of Economic Research. Princeton, New Jersey: Princeton University Press.

Newey, Whitney K., and Kenneth D. West (1987) "A Simple, Positive Semi-definite, Heteroskedasticity and Autocorrelation Consistent Covariance Matrix," **Econometrica** 55(3) 703-708.

Ramey, Valerie A. (1991) "Non-Convex Costs and the Behavior of Inventories," **Journal of Political Economy** 99(2) 306-334.

Sims, Christopher A., James H. Stock, and Mark W. Watson (1990) "Inference in Linear Time Series Models with Some Unit Roots," **Econometrica** 58, 113-144.

Stengel, Robert F. (1986) **Stochastic Optimal Control: Theory and Application**, New York, Wiley.

West, Kenneth D. (1986) "A Variance Bounds Test of the Linear Quadratic Inventory Model," **Journal of Political Economy** 91(2) 374-401.

________________ (1988) "Asymptotic Normality when Regressors Have a Unit Root," **Econometrica** 56, 1397-1418.

Table 1
Estimated Cost Parameters: The Baseline Estimates

Sample	α_0	α_1	α_2	α_3	J-statistic
1. All months					
a. 1925:1 - 1940:12	.213	-.149	.027	1.13	32.4
	(.013)	(.041)	(.025)	(.91)	[.02]
b. 1925:1 - 1932:5	.187	-.083	.052	.715	26.0
	(.012)	(.033)	(.029)	(.34)	[.10]
c. 1932:6 - 1940:12	.252	-.294	.082	.500	25.3
	(.014)	(.049)	(.030)	(.27)	[.12]
2. Non-Shutdown-Contaminated Months					
a. 1925:1 - 1940:12	.198	-.106	.029	1.47	23.3
	(.018)	(.054)	(.021)	(.97)	[.14]
b. 1925:1 - 1932:5	.176	-.054	.056	.752	19.7
	(.014)	(.041)	(.027)	(.27)	[.19]
c. 1932:6 - 1940:12	.212	-.150	.033	-.127	16.9
	(.027)	(.088)	(.023)	(.99)	[.39]
3. All Operating Months					
a. 1925:1 - 1940:12	.215	-.143	.003	-6.03	15.3
	(.019)	(.059)	(.021)	(47.7)	[.43]
b. 1925:1 - 1932:5	.200	-.109	.025	1.56	22.6
	(.014)	(.041)	(.014)	(.85)	[.09]
c. 1932:6 - 1940:12	.194	-.091	.022	-2.27	20.7
	(.027)	(.088)	(.028)	(3.76)	[.15]

Notes:

1. Estimates of constant terms are not reported. In all regressions reported here, β was fixed at 0.995.
2. We assumed that the shutdown months were as follows: 12/25, 12/26, 11/27, 12/27, 11/28, 12/28, 11/29, 12/29, 10/30, 11/30, 10/31, 11/31, 10/32, 11/32, 11/33, 12/33, 10/34, 11/34, 12/34, 9/35, 9/36, 9/37, 8/38, 9/38, 7/39, 8/39, 9/39, 8/40, and 9/40.
3. Standard errors are in parentheses; p-values are in square brackets.

Table 2
Estimated Cost Parameters: Alternative Estimates

Sample	α_0	α_1	α_2	α_3	J-statistic
1. Non-Shutdown-Contaminated Months; $\alpha_0=0$					
1925:1 - 1932:5	0	.455	.093	.266	18.1
		(.022)	(.043)	(.24)	[.32]
2. Non-Shutdown-Contaminated Months; $\alpha_1=0$					
1925:1 - 1932:5	.159	0	.050	.688	19.9
	(.004)		(.025)	(.27)	[.22]
3. Non-Shutdown-Contaminated Months; β=.99					
1925:1 - 1932:5	.177	-.055	.057	.751	19.7
	(.014)	(.041)	(.027)	(.27)	[.19]
4. All months; $\alpha_2=1$					
a. 1925:1 - 1940:12	.174	-.747	1.0	.440	25.8
	(.060)	(.122)		(.07)	[.10]
b. 1925:1 - 1932:5	-.053	-.232	1.0	.384	20.8
	(.054)	(.105)		(.06)	[.29]
c. 1932:6 - 1940:12	.337	-1.08	1.0	.518	20.2
	(.069)	(.136)		(.09)	[.32]

Notes: See Table 1.

*University of Chicago and NBER, and Federal Reserve Board, respectively. We have benefitted from many helpful comments and suggestions, especially those from Spencer Krane, Adrian Pagan, and Mark Watson, and two anonymous referees. We also thank Maura Doyle for excellent research assistance. The views expressed are those of the authors, and not necessarily those of the Board of Governors nor of the other members of its staff. Any errors are our own.

1. As measured by total assets in 1929, General Motors was the third largest industrial firm in the United States, exceeded only by US Steel and Standard Oil of New Jersey (Forbes 1977). According to registrations data, GM had slightly less than a third of the domestic new car market in 1929.

2. This section takes its title from the report written by F. Leslie Hayford in 1946.

3. Elements of this model can be found in earlier work by Modigliani and Sauerlander (1955) and Modigliani and Hohn (1955). More recent work using this framework includes Blanchard (1983), Blinder (1986), West (1986), Eichenbaum (1989), Ramey (1991), and Krane and Braun (1991).

4. For expositional convenience, we suppress here linear terms in the cost function. In the empirical work, we allow for such terms by including an intercept in the Euler equation.

5. Equation (2) is derived by using the accumulation identity to substitute out all occurrences of Q_{t+j} in the objective function, differentiating with respect to $E_{t-1}H_t$, and then simplifying again using the accumulation identity.

6. Ramey (1991) provides the cite to Stengel (1986). She imposes $\alpha_2=1$ as

an identifying assumption, and then tests whether the Legendre-Clebsch condition is satisfied.

The fact that we set the Legendre-Clebsch linear combination of the α's equal to 1 rather than any other non-zero scalar is irrelevant for the implied decision rule even in small samples. The specification of the linear combination of the parameters *is* relevant for the implied decision rule in small samples, but not asymptotically. Previous investigators have reported, however, that their empirical results were sensitive to the choice of an identifying assumption (see, among others, Krane and Braun (1991), pages 574-575, and Ramey (1991), page 322).

7. If H_t and S_t are cointegrated and S_t is I(1), then H_t must be I(1). But if H_t is I(1), then $\Delta H_t = Q_t - S_t$ must be I(0). Thus, there exists a linear combination of Q_t and S_t (namely the difference between them) that is stationary; therefore, Q_t and S_t must be cointegrated.

8. In the new steady state after a positive innovation in permanent sales, the firm will choose to incur a higher marginal stockout cost because it will be experiencing a higher marginal production cost (associated with α_1). Unless the target level of inventories is a positive function of sales, the higher marginal stockout cost will entail a lower level of inventories.

9. We are grateful to one of the anonymous referees for supplying the outline of the following argument.

10. It is an open question as to how the finite sample properties of the estimates are affected by the use of the time trends.

11. In the first step, we estimate the α's using nonlinear two-stage least squares. Then we form the cross products of the residuals from the first

stage with the instruments, using the lags=1 options in RATS to alert the program to the MA(1) error structure, and the damp=1 option to ensure the positive semi-definiteness of this matrix. Then we reestimate the α's using the inverse of the cross-products matrix as a weighting matrix, again invoking the lags=1 and damp=1 options.

12. A more complete treatment of the problem would include an explanation for the choice of a technology that requires shutdown to occur (see Cooper and Haltiwanger (1990) for a serious consideration of the machine replacement problem).

13. When we are estimating under this approach, we make room for other differences across the periods by allowing the intercept to take on a different value depending on whether the observation in question is (1) a normal operating month, (2) a month immediately following a shutdown period, (3) a month immediately preceding a shutdown period, or (4) a month preceding a shutdown period by two months.

14. We suspect that our failure to observe zero production mainly reflects that GM did not perfectly align its shutdown periods either across divisions or with the calendar months, and so was always operating at least one of its plants during at least part of a "shutdown" month. Since our model views each period as a point in time, it is silent on the issue of how production should behave during a month when GM was operating only part of its capacity part of the time.

15. The data were kindly provided to us by a member of GM staff. The data are monthly, and cover the corporation as a whole. Data at the divisional level are not available. The data for 1925 include Chevrolet trucks; data for the other years do not include Chevrolet trucks. In 1925, Chevy trucks accounted for about 5 percent of total sales and production of the

corporation as a whole. We made a simple adjustment to the level of sales and production in 1925, which preserved the inventory accounting identity between 1924 and 1926. Details of this adjustment will be provided along with the data upon request. The GM staff member cautioned that the data from the early years may be less reliable than those from the later years.

16. We were unable to develop conclusive information on the exact timing of the annual shutdowns at the individual plant level. We inferred approximate shutdown dates from Hayford and various issues of Automotive Industries, a contemporary trade journal. Whenever we had substantial doubt about whether GM was fully operational during a given month, we erred on the side of caution and flagged that month as a shutdown month. The list of months we identified as shutdown months is given in the notes to Table 1.

17. The spike in 1938 appears to reflect a colossal forecasting error. GM opened the 1938 model year with two months of extremely aggressive production, even as the economy was turning into a sharp nosedive. They then spent the rest of the model year slashing production not only in line with the weakness in sales, but also to decumulate the inventories that had piled up in the first few months of the model year.

18. The estimates in Tables 1 and 2 were obtained using an instrument set that included the following 10 variables: the contemporaneous value and the first two lags of the change in production; the first two lags each of the change in sales, the change in the log of industrial production, and the change in the log of the composite Standard and Poor's stock index; and the first lag of the difference between production and sales. In addition, the instrument set included as many seasonal dummy variables as possible (a dummy variable for December could not be included, for example, if no December observations were included in the particular sample being

estimated). This instrument set reflects our assumption that GM did not know current-period sales when it chose production. We included the contemporaneous value of production in the instrument set because it was the choice variable of the firm.

19. The statistical significance of α_0, α_1, and α_2 is difficult to interpret because in fact, of course, we are estimating the *ratio* of those parameters to the linear combination of them given in equation (3); it is difficult to know when we should be able to estimate such a ratio precisely and when we should not. Therefore, rather than relying heavily on statistical significance *per se*, we look for other clues as to the economic importance of any given parameter.

20. We obtain similar results using the second method of handling the shutdown periods.

21. This approach of simulating decision rules to assess the importance of production smoothing is similar to the approach pursued in Fair (1989).

22. In his study of the seasonal comovement of sales and production, Krane (1991) makes a similar observation in explaining the close correspondence of sales and production in some of the industries he examines.

Chart 1

Monthly Sales

as a percent of a centered twelve–month moving average

1924 – 1940

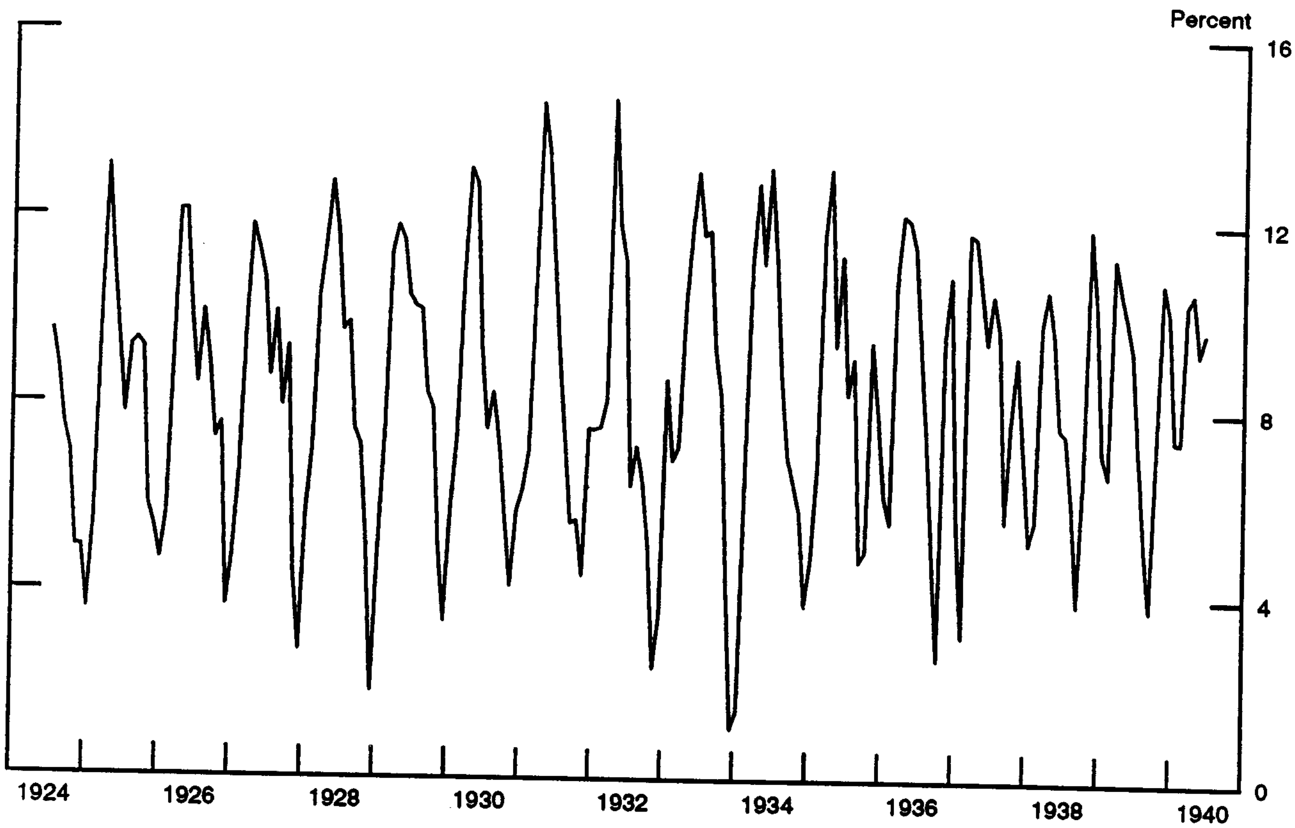

1971 – 1987

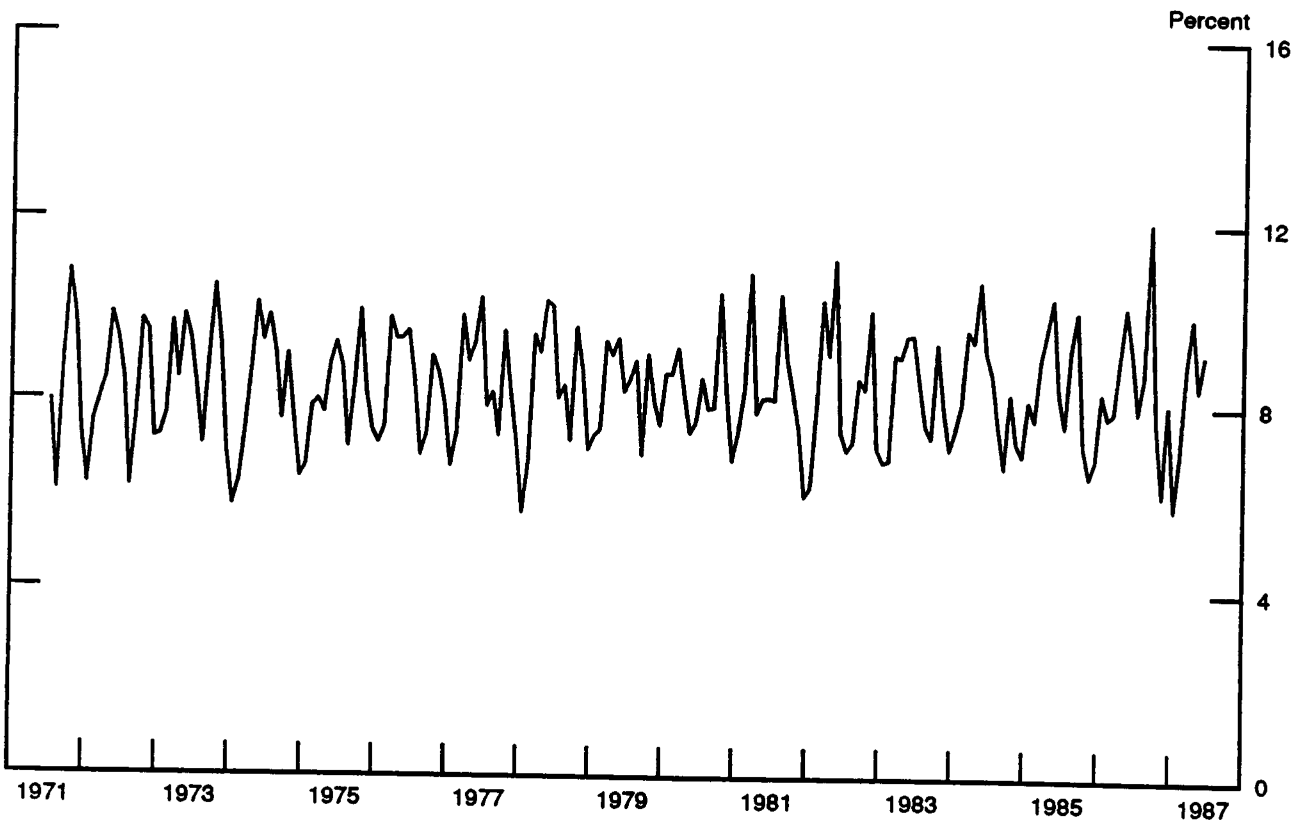

Chart 2

The Long–Run Relationships Between Production, Inventories, and Sales

(centered twelve–month moving averages)

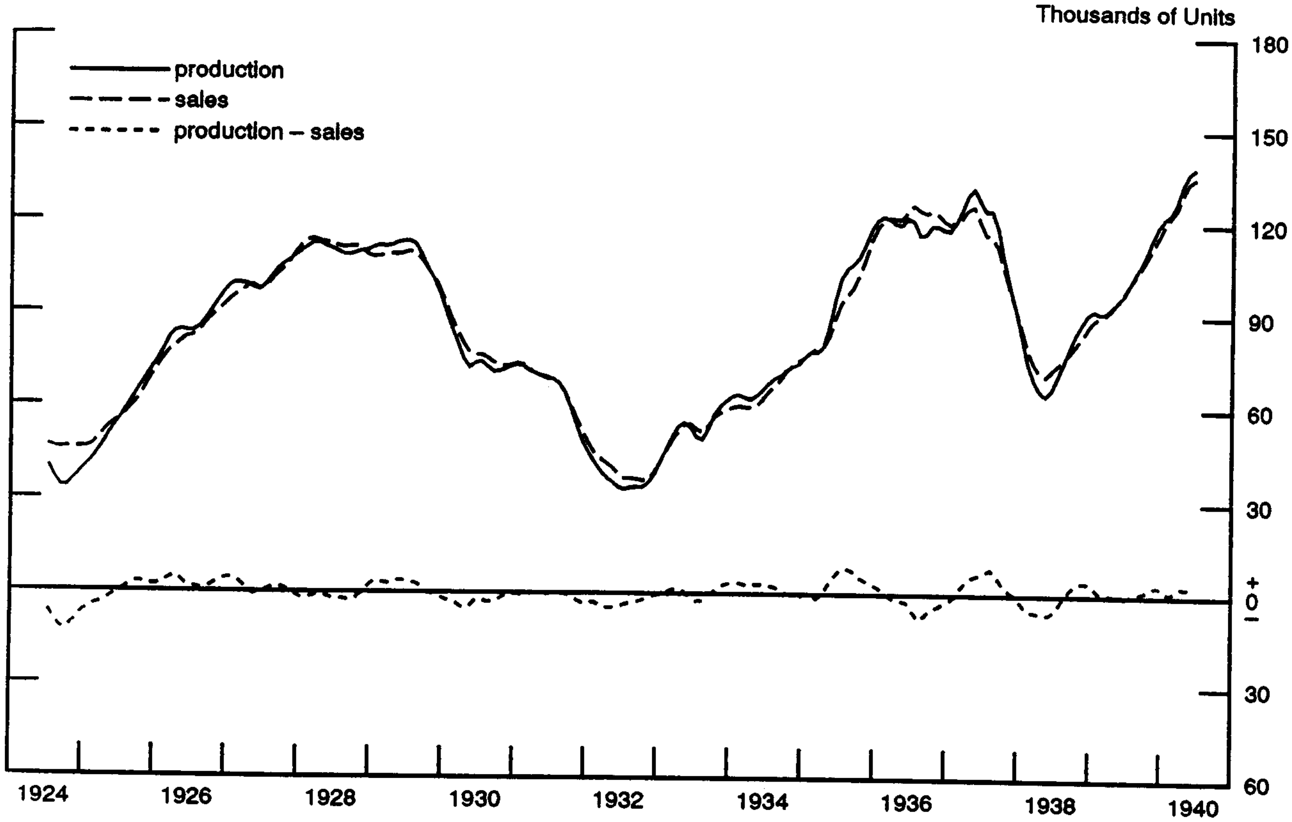

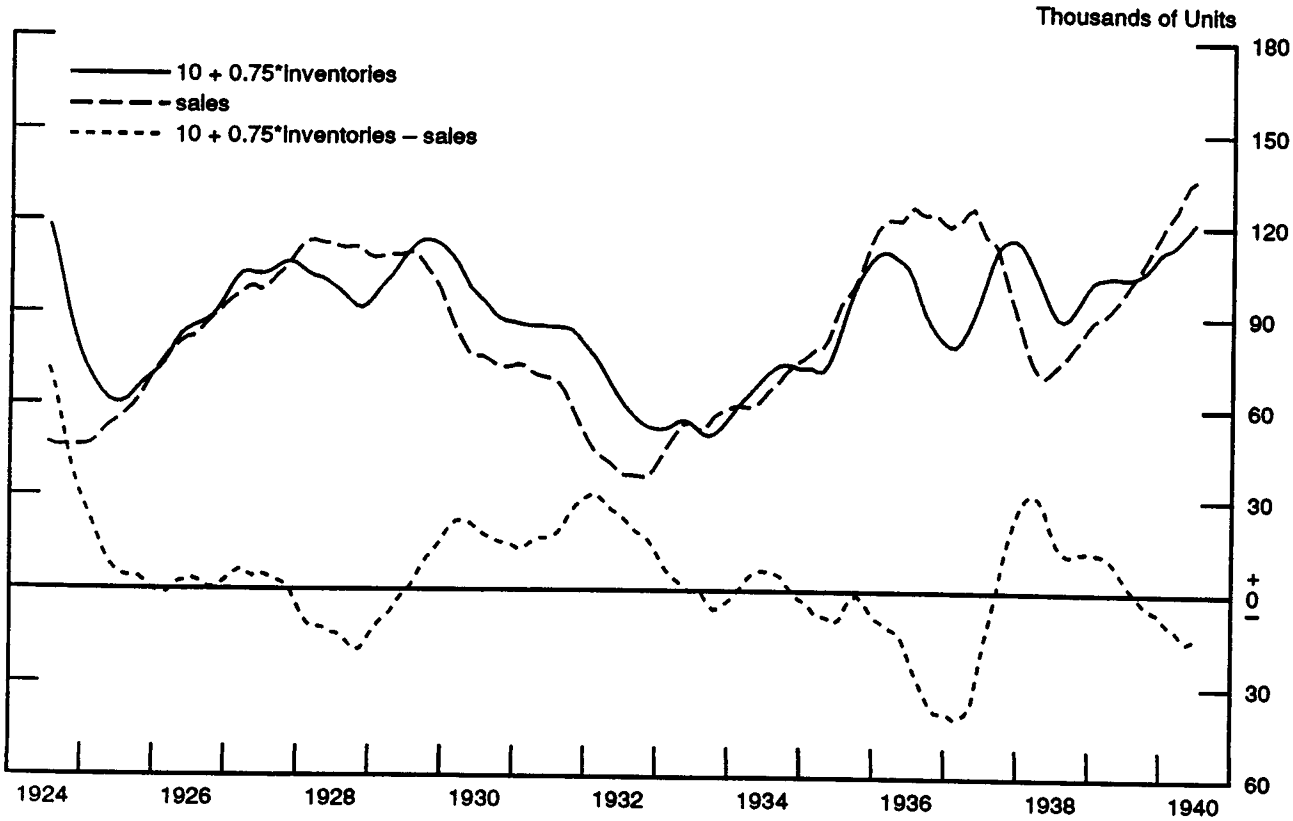

Chart 3

Ratio of Within–year Variance of Production to Within–Year Variance of Sales

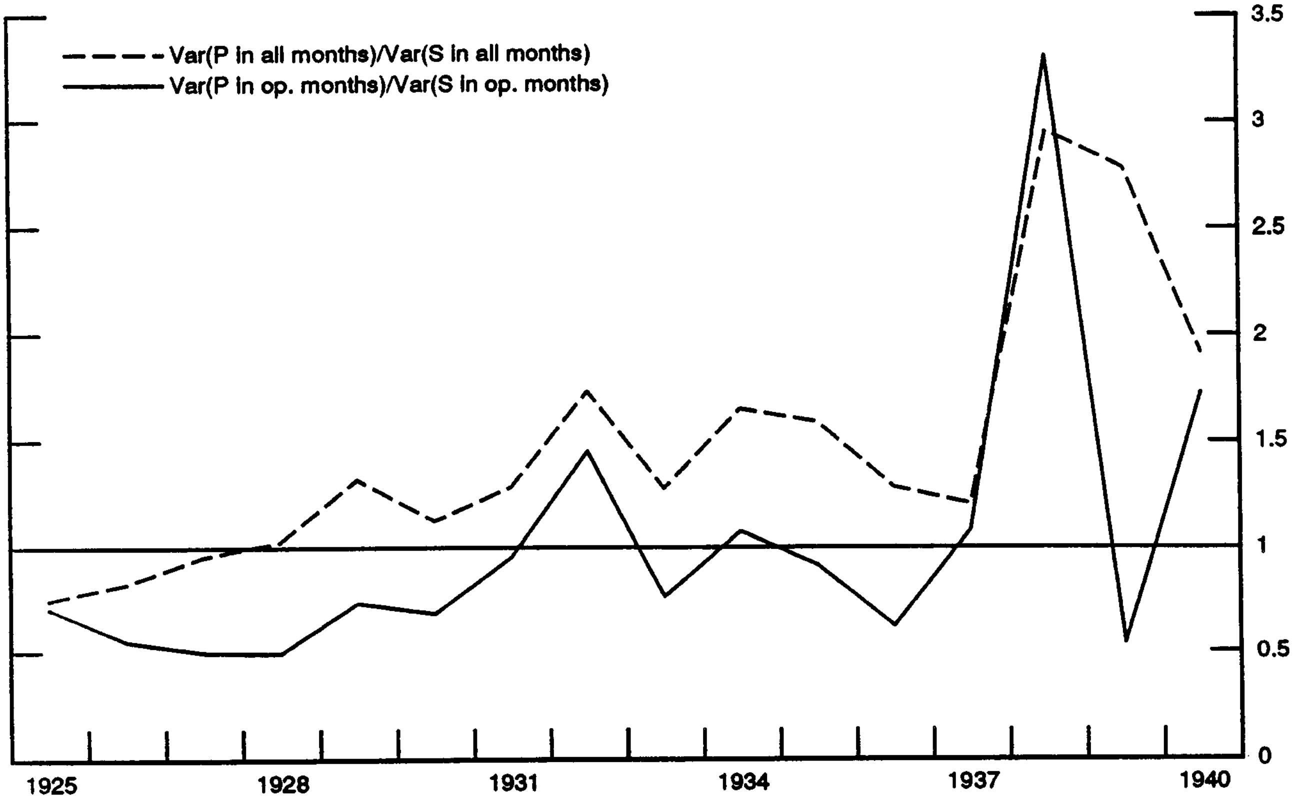

Chart 4

Production, Sales, and the Level of Inventories

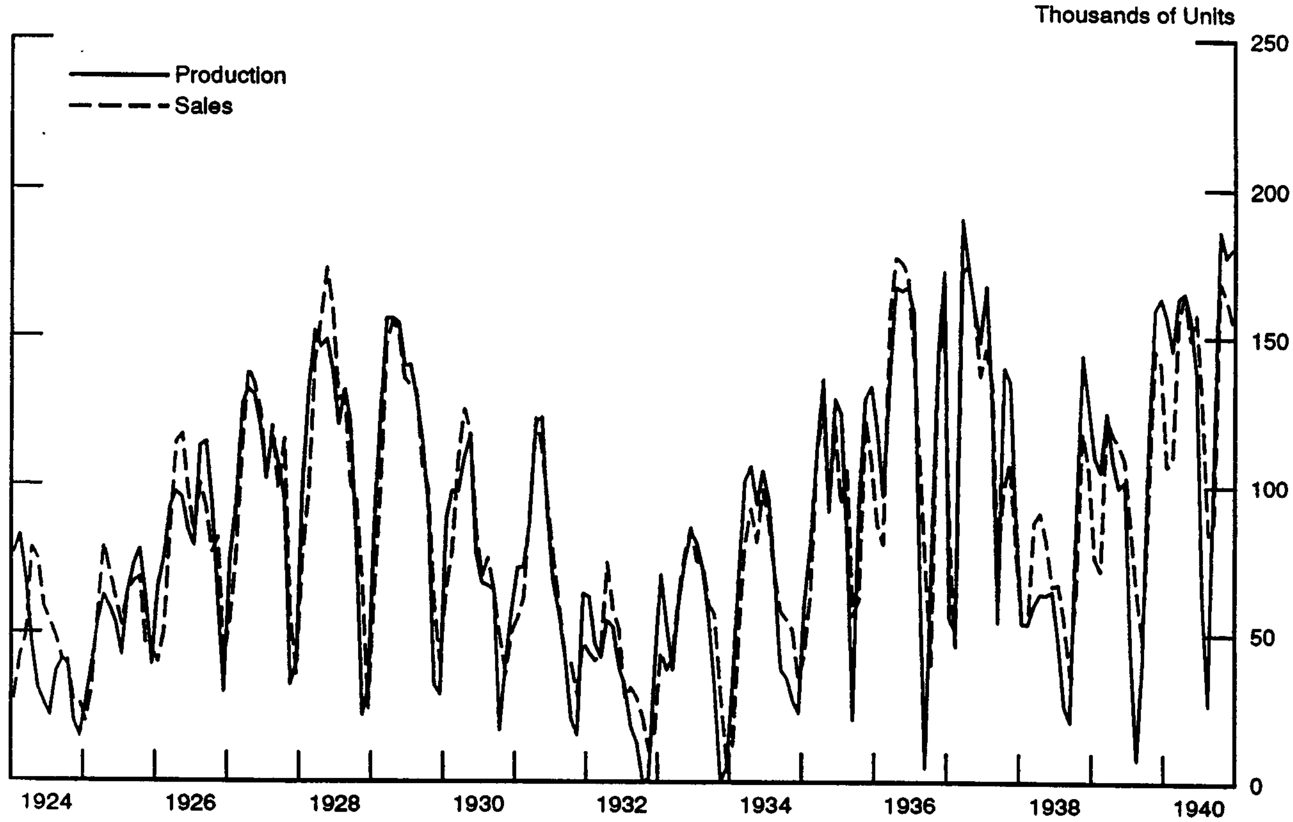

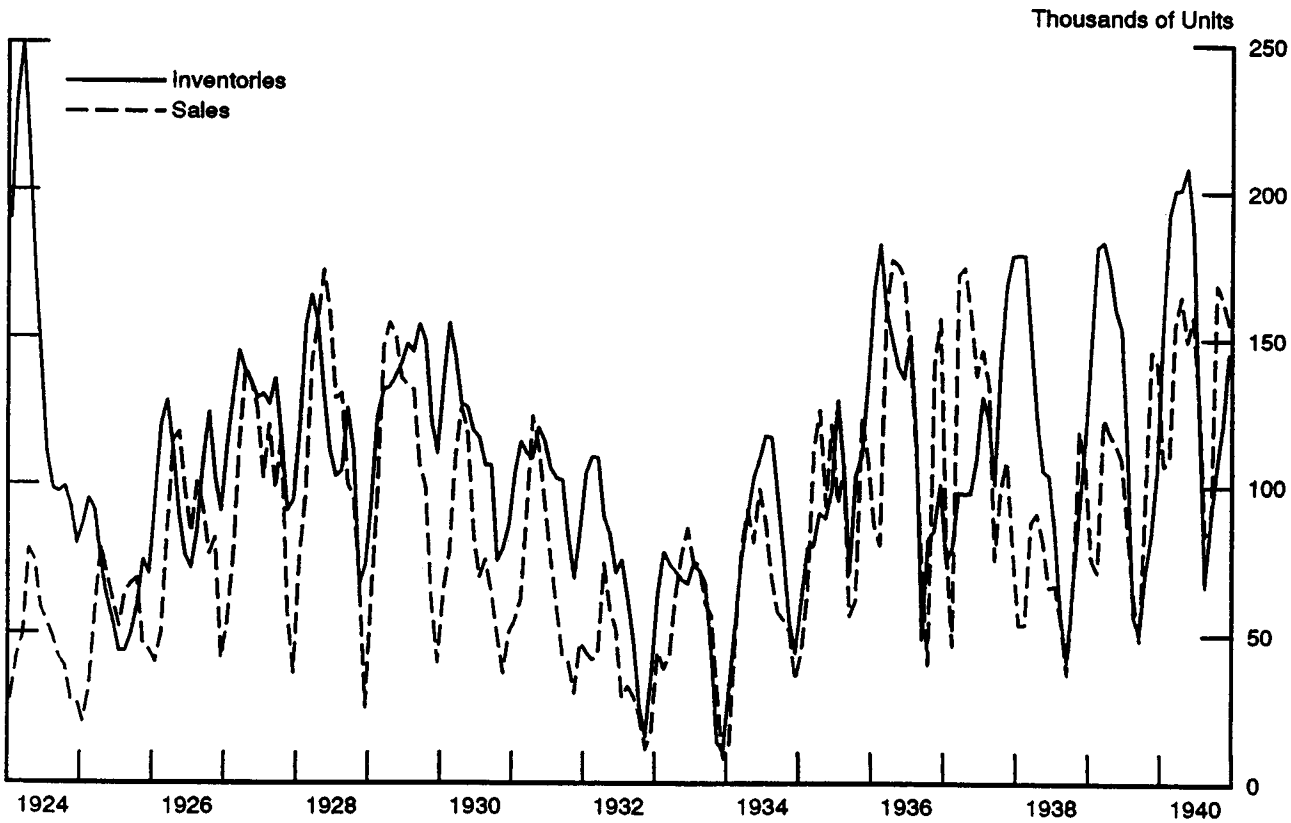

Chart 5

The Inventory–Sales Ratio and Production

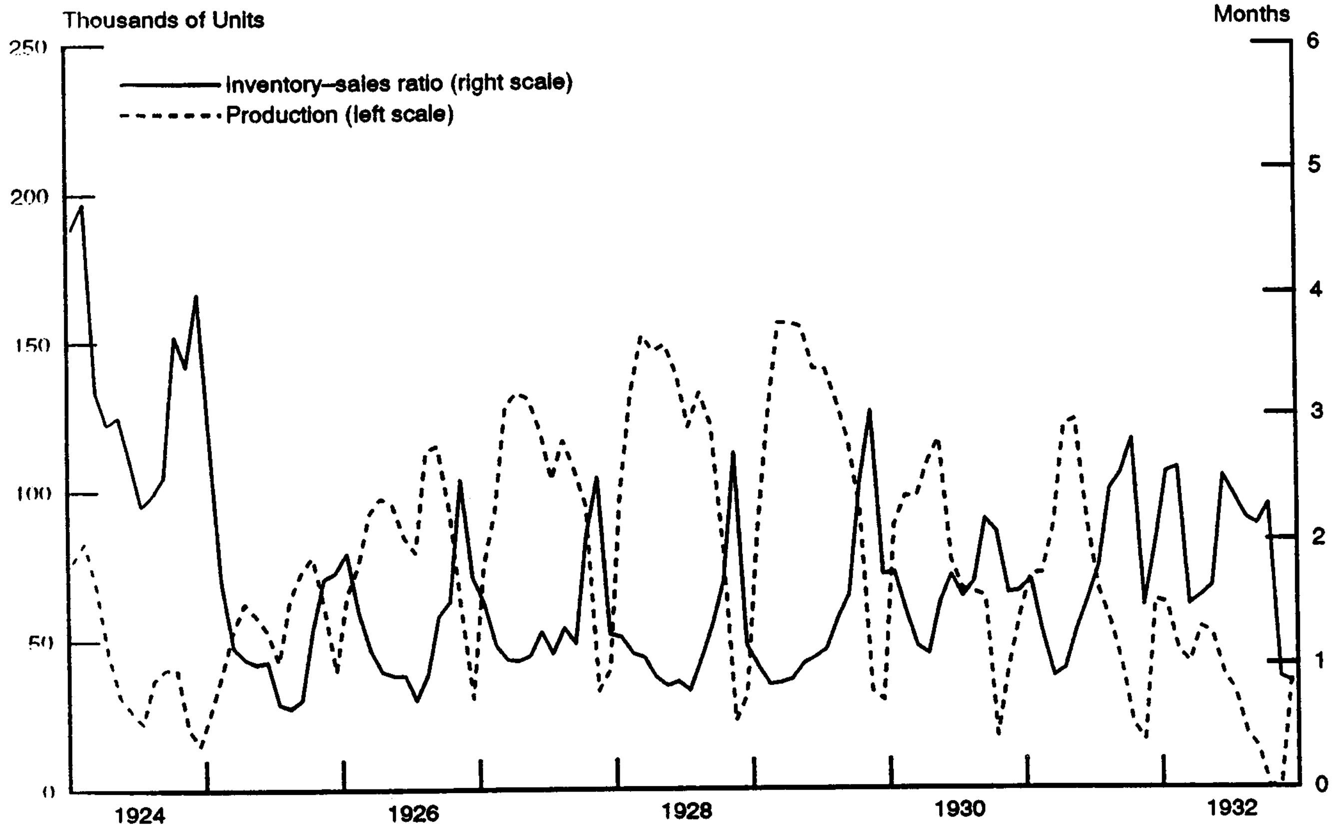

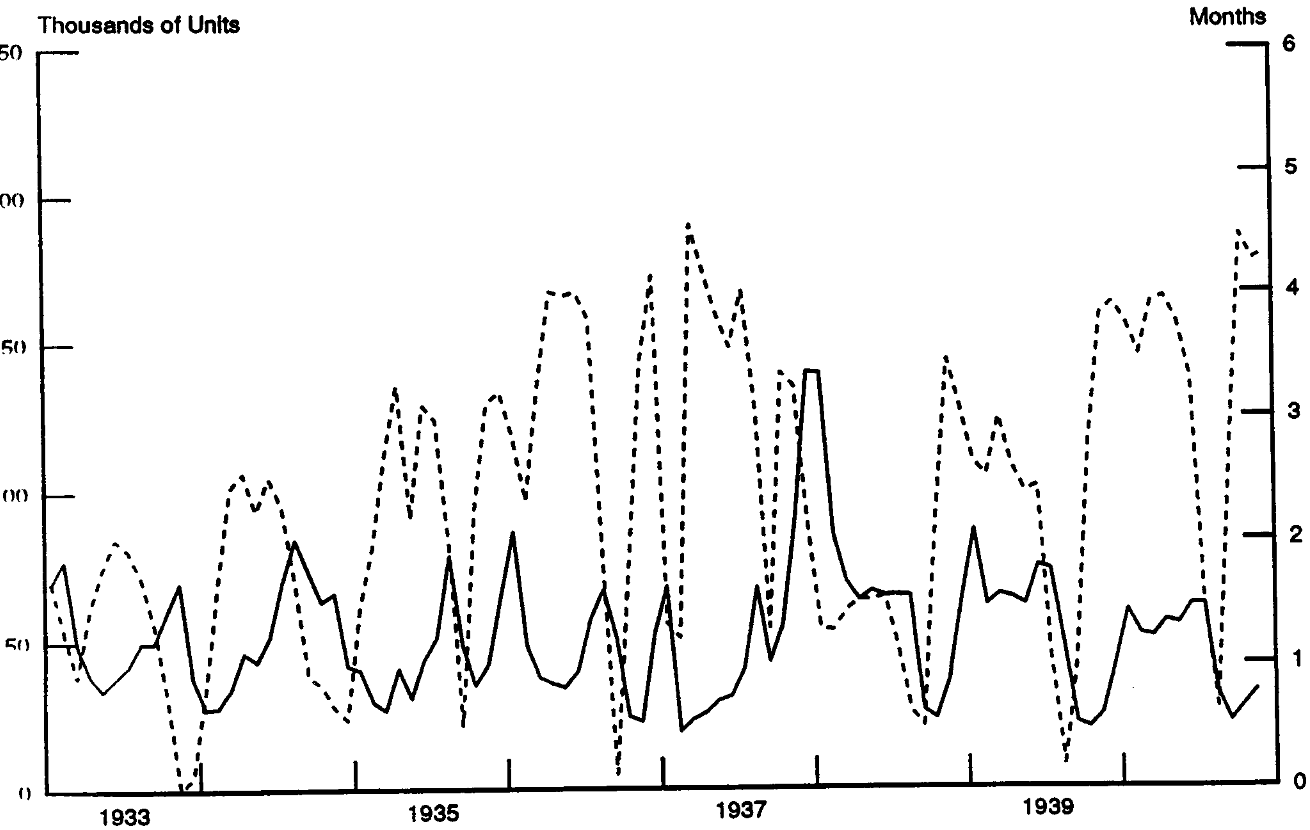

Chart 6

Residuals from the Basic Euler Equation Estimated Over All Months

Baseline Identifying Assumption: The Legendre–Clebsch Condition

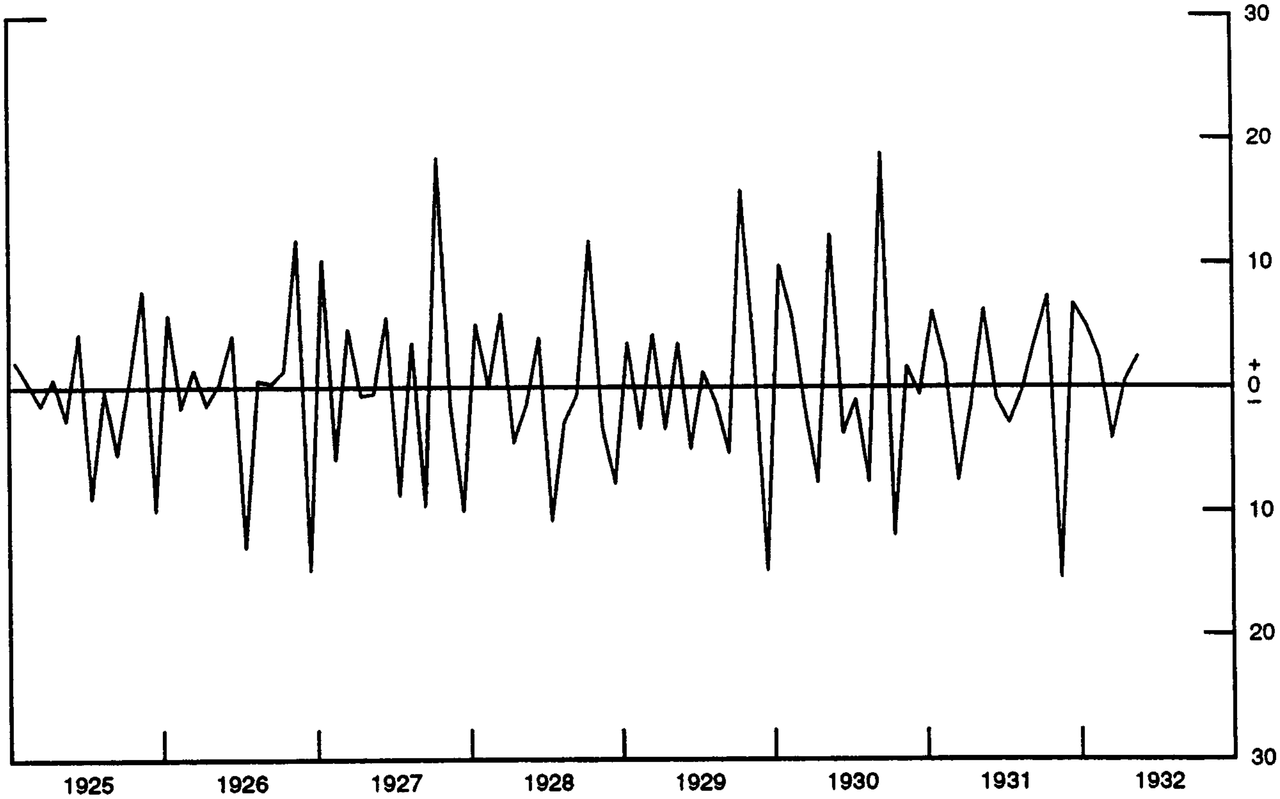

Alternative Identifying Assumption: a2=1

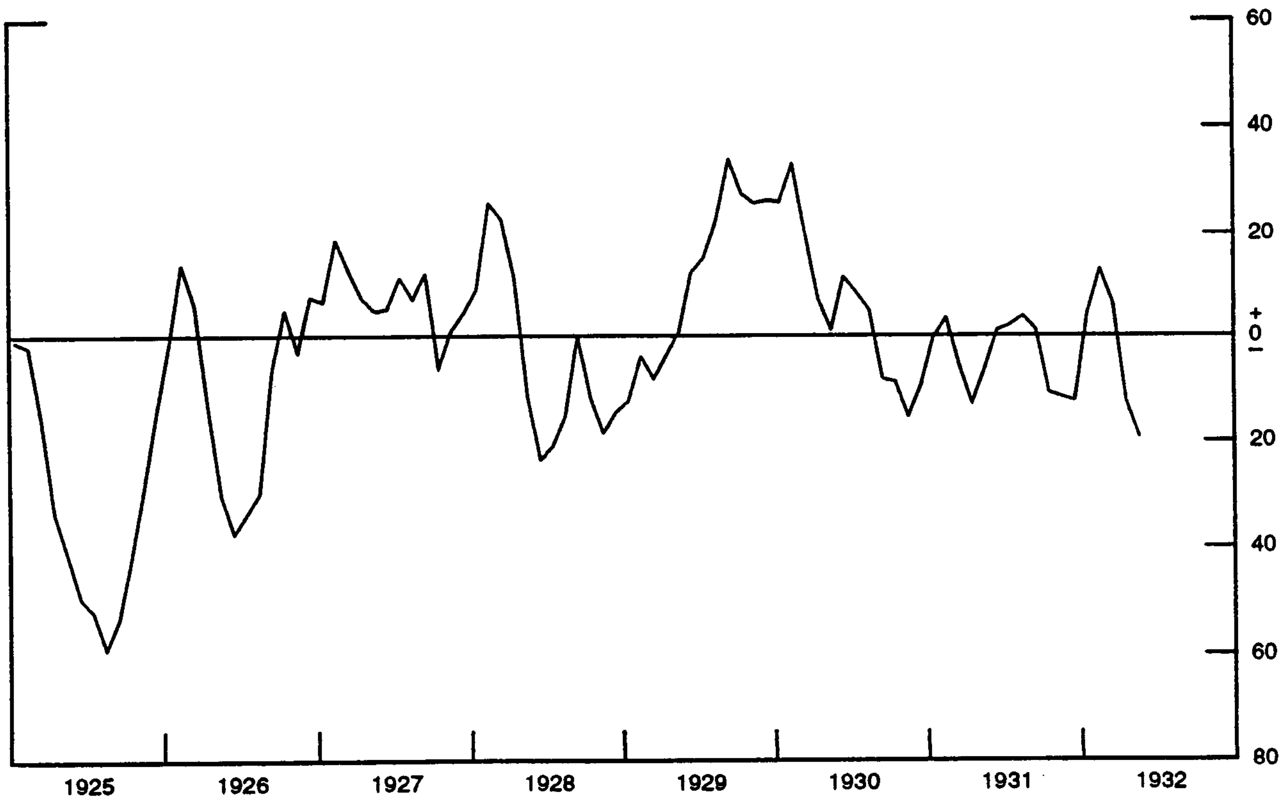

Chart 7

Impulse Response Functions

Early Sample Period

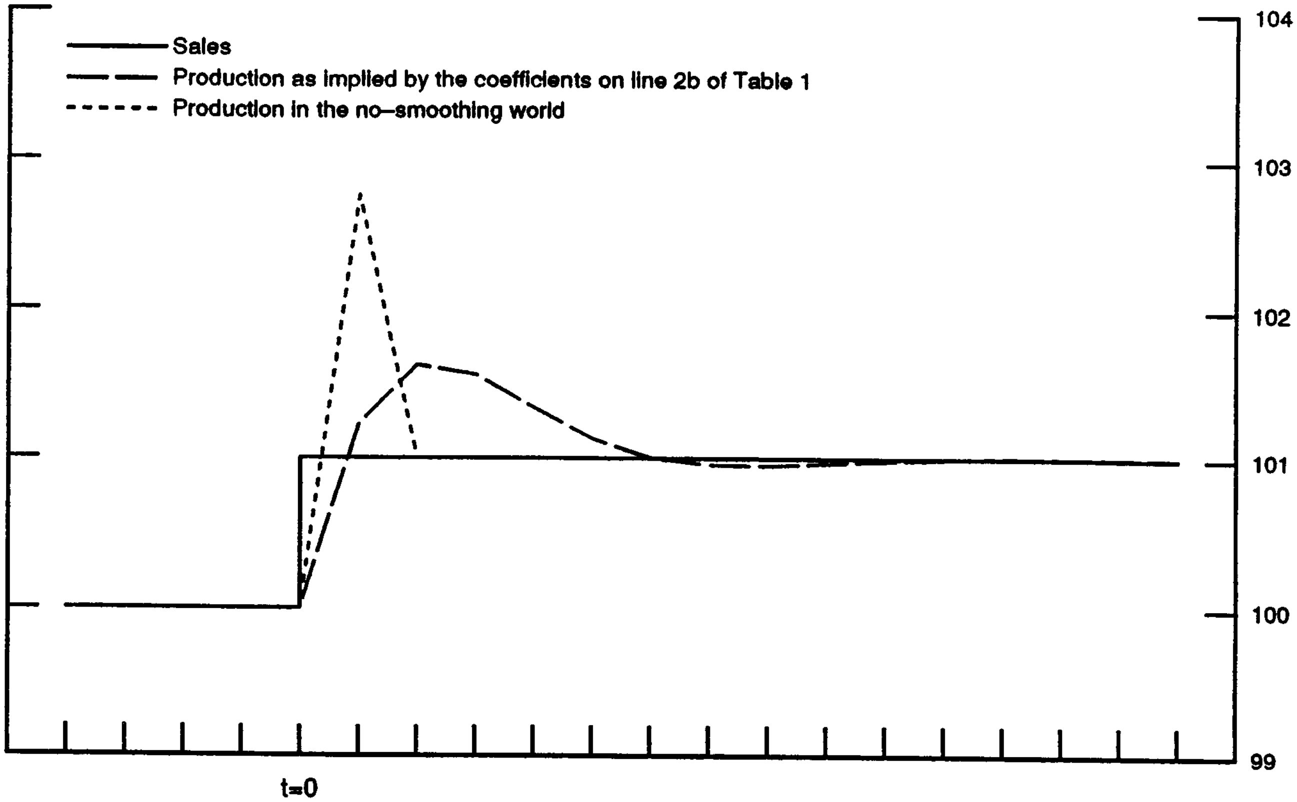

Late Sample Period