Chapter 13

INVENTORIES *

VALERIE A. RAMEY

University of California – San Diego

KENNETH D. WEST

University of Wisconsin

Contents

* We thank the National Science Foundation and the Abe Foundation for financial support; Clive Granger, Donald Hester, James Kahn, Anil Kashyap, Linda Kole, Spencer Krane, Scott Schuh, Michael Woodford and a seminar audience at the University of Wisconsin for helpful comments and discussions; and James Hueng and especially Stanislav Anatolyev for excellent research assistance. Email to: `kdwest@facstaff.wisc.edu`; `vramey@weber.ucsd.edu`.

Handbook of Macroeconomics, Volume 1, Edited by J.B. Taylor and M. Woodford

Abstract

We review and interpret recent work on inventories, emphasizing empirical and business cycle aspects. We begin by documenting two empirical regularities about inventories. The first is the well-known one that inventories move procyclically. The second is that inventory movements are quite persistent, even conditional on sales.

To consider explanations for the two facts, we present a linear-quadratic model. The model can rationalize the two facts in a number of ways, but two stylized explanations have the virtue of relative simplicity and support from a number of papers. Both assume that there are persistent shocks to demand for the good in question, and that marginal production cost slopes up. The first explanation assumes as well that there are highly persistent shocks to the cost of production. The second assumes that there are strong costs of adjusting production and a strong accelerator motive.

Research to date, however, has not reached a consensus on whether one of these two, or some third, alternative provides a satisfactory explanation of inventory behavior. We suggest several directions for future research that promise to improve our understanding of inventory behavior and thus of business cycles.

Keywords

JEL classification: E22, E32

Introduction

In developed countries, inventory investment typically averages less than one-half of one percent of GDP, whereas fixed investment averages 15% of GDP and consumption two-thirds. Perhaps with these fractions in mind, macroeconomists have concentrated more on the study of consumption and fixed investment than on inventories. Inventories generally do not appear as separate variables in dynamic general equilibrium models, nor in exactly identified vector autoregressive models.

It has long been known, however, that other ways of measuring the importance of inventories suggest that inventories should receive more attention, especially in business cycle research. Half a century ago, Abramowitz (1950) established that US recessions prior to World War II tended to be periods of inventory liquidations. Recent experience in the G7 countries indicates this regularity continues to hold, and not just for the USA. In six of the seven G7 countries (Japan is the exception), real GDP fell in at least one recent year. Line 2 of Table 1 shows that in five of those six countries (the United Kingdom is now the exception), inventory investment also declined during the period of declining GDP, accounting in an arithmetical sense for anywhere 12–71% of the fall in GDP. And Table 1's use of annual data may understate the inventory contribution: Table 2 indicates that for quarterly US data, the share is 49 rather than 12% for the 1990–1991 recession, with 49 a typical figure for a post-War US recession.

Such arithmetical accounting of course does not imply a causal relationship. But it does suggest that inventory movements contain valuable information about cyclical fluctuations. In this chapter, we survey and interpret recent research on inventories, emphasizing empirical and business cycle aspects. Among other points, we hope to convince the reader that inventories are a useful resource in business cycle analysis. They may be effective in identifying both the mechanisms of business cycle propagation and the sources of business cycle shocks.

Our chapter begins by documenting two facts about inventories. The first is the well-known one that inventories move procyclically. They tend to be built up in expansions, drawn down in contractions. The second, and not as widely appreciated, fact is that inventory movements are quite persistent, even conditional on sales. In many data sets, inventories and sales do not appear to be cointegrated, and the first-order autocorrelations of supposedly stationary linear combinations of inventories and sales are often around 0.9, even in annual data.

To consider explanations for the two facts, we use a linear quadratic/flexible accelerator model, which is the workhorse for empirical research on inventories. In our model, one source of persistence is from shocks to demand for the good being put in inventory – "demand" shocks. ("Demand" is in quotes because we, and the literature more generally, do not attempt to trace the ultimate source of such shocks; for example, for an intermediate good, the shocks might be driven mainly by shocks to the technology of the industry that uses the good in production.) But even if this shock has a unit root, our model yields a stationary linear combination of inventories

Table 1
Arithmetical importance of inventory change in recessions of the 1990s (annual data)[a]

Country	Canada	France	West Germany	Italy	Japan	UK	USA
(1) Peak year	1989	1992	1992	1992	n.a.	1990	1990
trough year[b]	1991	1993	1993	1993		1992	1991
(2) Peak–trough change in inventory change as percentage of peak-to-trough fall in GDP[c]	50	71	19	30	n.a.	–0.	12

[a] The figures are based on annual real data. The inventory change series is computed by deflating the annual nominal change in inventories in the National Income and Product Accounts by the GDP deflator; see the Data Appendix.

[b] The trough year was found by working backwards from the present to the last year of negative real GDP growth in the 1990s. There were no such years in Japan. The peak year is the last preceding year of positive real GDP growth.

[c] Computed by multiplying the following ratio by 100:

$$\frac{\text{inventory change in trough year} - \text{inventory change in peak year}}{\text{GDP in trough year} - \text{GDP in peak year}}.$$

By construction, the denominator of this ratio is negative. A positive entry indicates that the numerator (the change in the inventory change) was also negative. The negative entry for the United Kingdom indicates that the change in the inventory change was positive.

Table 2
Arithmetical importance of inventory changes in post-war US recessions (quarterly data)[a]

Peak quarter–trough quarter	Peak-to-trough inventory change as a percentage of peak-to-trough fall in GDP
1948:4–1949:2	130
1953:2–1954:2	41
1957:1–1958:1	21
1960:1–1960:4	122
1969:3–1970:1	127
1973:4–1975:1	59
1980:1–1980:3	45
1981:3–1982:3	29
1990:2–1991:1[b]	49

[a] The figures are based on quarterly real data. See the notes to Table 1 for additional discussion.

[b] The figure for the 1990–1991 recession differs from that for the USA in Table 1 mainly because quarterly data were used. It also differs because in this table the inventory change is measured in chain weighted 1992 dollars, whereas Table 1 uses the nominal inventory change deflated by the GDP deflator.

and sales. This stationary linear combination can be considered a linear version of the inventory–sales ratio. We call it the *inventory–sales relationship.* And our second inventory fact is that there is persistence in this relationship.

While the model is rich enough that there are many ways to make it explain the two facts, we focus on two stylized explanations that have the virtue of relative simplicity, as well as empirical support from a number of papers. Both explanations assume a upward sloping marginal production cost (a convex cost function). The first explanation also assumes that fluctuations are substantially affected by highly persistent shocks to the cost of production. Cost shocks will cause procyclical movement because times of low cost are good times to produce and build up inventory, and conversely for times of high cost. As well, when these shocks are highly persistent a cost shock that perturbs the inventory–sales relationship will take many periods to die off, and its persistence will be transmitted to the inventory–sales relationship.

The second explanation assumes that there are strong costs of adjusting production and a strong accelerator motive. The accelerator motive links today's inventories to tomorrow's expected sales, perhaps because of concerns about stockouts. Since sales are positively serially correlated, this will tend to cause inventories to grow and shrink with sales and the cycle, a point first recognized by Metzler (1941). As well, with strong costs of adjusting production, if a shock perturbs the inventory–sales relationship, return to equilibrium will be slow because firms will adjust production only very gradually.

Both explanations have some empirical support. But as is often the case in empirical work, the evidence is mixed and ambiguous. For example, the cost shock explanation works best when the shocks are modelled as unobservable; observable cost shifters, such as real wages and interest rates, seem not to affect inventories. And the literature is not unanimous on the magnitude of adjustment costs.

While the literature has not reached a consensus, it has identified mechanisms and forces that can explain basic characteristics of inventory behavior and thus of the business cycle. We are optimistic that progress can continue to be made by building on results to date. Suggested directions for future research include alternative ways of capturing the revenue effects of inventories (replacements for the accelerator), alternative cost structures and the use of price and disaggregate data.

The chapter is organized as follows. Section 1 presents some overview information on the level and distribution of inventories, using data from the G7 countries, and focussing on the USA. We supply this information largely for completeness and to provide a frame of reference; the results in this section are referenced only briefly in the sequel.

Section 2 introduces the main theme of our chapter (business cycle behavior of inventories) by discussing empirical evidence on our two facts about inventories. Procyclical movement is considered in Section 2.1, persistence in the inventory–sales relationship in Section 2.2. In these sections, we use annual data from the G7 countries and quarterly US data for illustration, and also summarize results from the literature.

Sections 3–7 develop and apply our linear quadratic/flexible accelerator model. Sections 3–5 present the model. Much of the analysis in these three sections relates to the process followed by the inventory–sales relationship, because this process has not received much direct attention in existing literature. The discussion focuses on analytical derivations, for the most part deferring intuitive discussion about how the model works to Section 6. That section aims to develop intuition by presenting impulse responses for various configurations of parameters. Section 7 reviews empirical evidence from studies using the model.

In Section 8, we discuss extensions and alternative approaches, including models that put inventories directly in production and profit functions, models with fixed costs, and the use of different data. Section 9 concludes. A Data Appendix describes data sources, and a Technical Appendix contains some technical details.

1. Sectoral and secular behavior of inventories

In this section we use basic national income and product account data from the G7 countries, and some detailed additional data from the USA, to provide a frame of reference for the discussion to come. As just noted, for the most part this is background information that will not loom large in the sequel.

Lines 1(a) and 1(b) of Table 3 present the mean and standard deviation of the real annual change in economy wide inventory stocks in the G7 countries, over the last 40 years. These were computed from the national income and product account data on

Table 3
Basic inventory statistics

	Canada	France	West Germany	Italy	Japan	UK	USA
(1) Annual NIPA change in inventories, 1956–1995 [a,b]							
(a) Mean	2.32	37.4	12.3	12.3	2.41	1.81	23.6
(b) Standard deviation	3.91	40.1	12.7	9.8	1.44	3.04	21.6
(2) Reference: 1995 GDP [c]	721	6882	2608	1351	453	584	6066
(3) 1995 Inventory level [d]	131	n.a.	411	n.a.	71	104	971

[a] The inventory change series is computed by deflating the annual nominal change in inventories in the National Income and Product accounts by the GDP deflator; see the Data Appendix. Units for all entries are billions (trillions, for Italy and Japan) of units of own currency, in 1990 prices.

[b] Sample periods are 1957–1994 for West Germany and 1960–1994 for Italy, not 1956–1995.

[c] GDP entries for Italy and Germany are for 1994, not 1995.

[d] The "level" entries for Canada, West Germany, Japan and the UK are computed by deflating the nominal end of year value by the GDP deflator; see the Data Appendix. The entry for the US is the Department of Commerce constant (chained 1992) dollar value for non-farm inventories, rescaled to a 1990 from a 1992 base with the GDP deflator.

Table 4
Sectoral distribution of US non-farm inventories [a,b]

	(1) Percent of total level, 1995	(2) Mean (s.d.) of change	(3) Mean (s.d.) of growth
Total	100	21.4	3.5
		(22.5)	(3.5)
Manufacturing	37	7.0	2.8
		(11.6)	(4.2)
Finished goods	13	2.5	3.0
		(4.4)	(4.8)
Work in process	12	2.3	2.8
		(5.9)	(6.0)
Raw materials	12	2.2	2.6
		(5.4)	(6.2)
Trade	52	12.2	4.4
		(13.4)	(4.5)
Retail	26	5.9	4.2
		(10.3)	(6.7)
Wholesale	26	6.2	4.5
		(7.3)	(4.8)
Other	11	2.2	3.1
		(5.1)	(5.8)

[a] Data are in billions of chained 1992 dollars, 1959:I–1996:IV.
[b] The inventory change differs from the US data on changes in Tables 1–3 in coverage (Tables 1–3 include changes in farm inventories), in sample period (1959–1996 here, 1956–1995 in Table 3) and in base year (1992 here and Table 2, 1990 in Tables 1 and 3).

the change in aggregate inventories. See the notes to the table and the Data Appendix for details.

Upon comparing line 1(a) to line 2, we see that in all seven countries, the average change in inventories is small, about one percent of recent GDP in Italy, well less than that in other countries. Inventory changes are, however, reasonably volatile, with the standard deviation roughly as large as the mean in all seven countries.

We have less complete data on the level (as opposed to the change) of inventory stocks. Line 3 of Table 3 indicates that in the countries for which we have been able to obtain data, total inventories are about one-sixth of GDP. This implies a monthly inventory–sales ratio of about 2, a value that will be familiar to those familiar with monthly US data.

Table 4 has a breakdown of US non-farm inventories by sector. We see in column 1 that about half of non-farm inventories are held by retailers and wholesalers (including

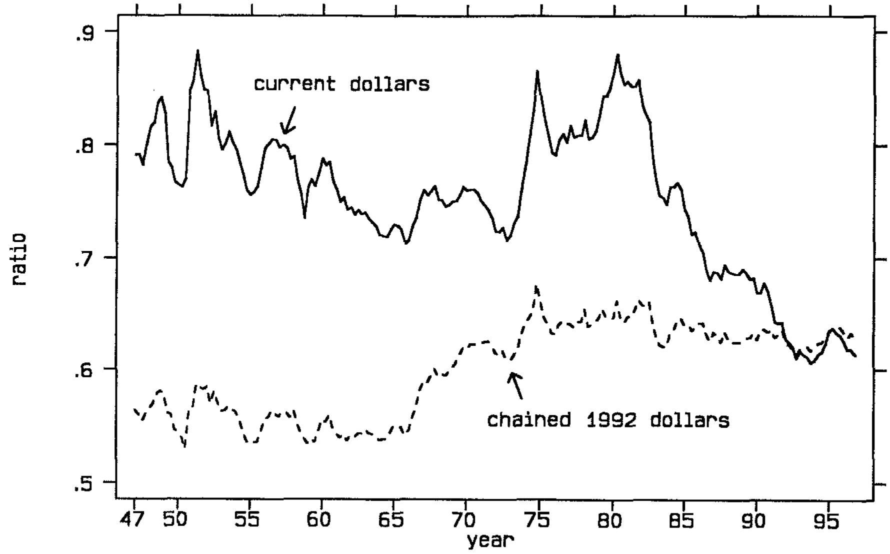

Fig. 1. Quarterly ratio of nonfarm inventories to final sales.

non-merchant wholesalers who are associated with particular manufacturers), whereas somewhat over a third are held by manufacturers. The remaining "other" category reflects holdings by a number of industries, including utilities, construction, and service companies.

Like the aggregates in Table 3, investment in each of the components is positive on average, and has standard deviations about the same size as means. This applies whether one looks at arithmetic changes (column 2) or growth rates (column 3). For future reference, it will be useful to note that manufacturers' inventories of finished goods, which have received a fair amount of attention in the inventory literature, are only 13% of total inventories, and are not particularly volatile.

Figure 1 plots the ratio of total non-farm inventories to final sales of domestic product. The dashed line uses real data (ratio of real inventories to real sales), the solid line nominal data. In the real data, the inventory series matches that in line 1 of Table 4, but over the longer sample 1947:I–1996:IV. (Table 4 uses the 1959–1996 subsample because the disaggregate breakdown is not available 1947–1958.) The real ratio shows a run-up in the late 1960s and early 1970s, followed by a period of slight secular decline. At present, the ratio is modestly above its value at the start of our sample (0.63 vs. 0.56). It will be useful to note another fact for future reference. The figure suggests considerable persistence in the inventory–sales ratio, an impression borne out by estimates of first-order autocorrelations. These are 0.98 for the sample as whole, 0.93 if the autocorrelation is computed allowing for a different mean inventory–sales ratio for the 1947:I–1973:IV and 1974:I–1996:IV subsamples.

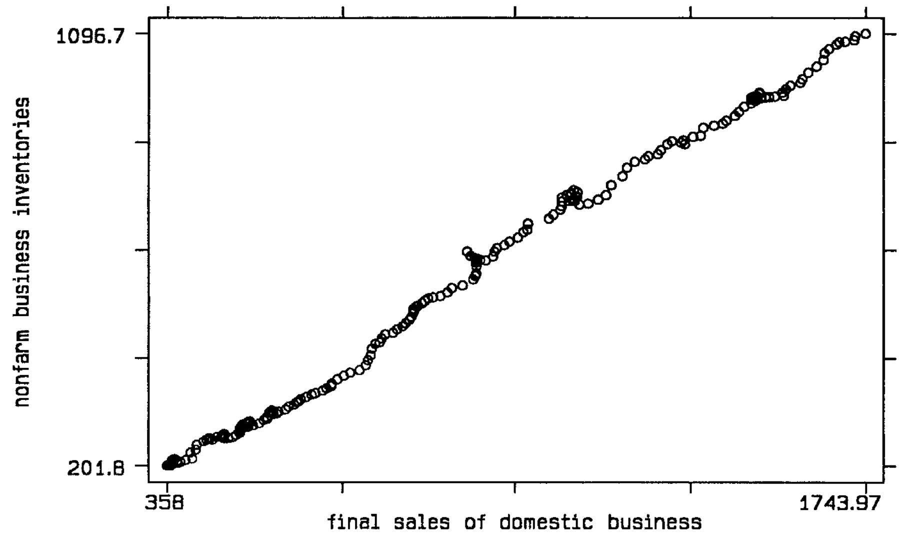

Fig. 2. Quarterly inventories and sales, 1947:1–1996:4, in billions of chained 1992 dollars.

Readers familiar with the monthly inventory–sales ratios commonly reported in the US business press may be surprised at the absence of a downward secular movement. Such monthly ratios typically rely on nominal data. The solid line in Figure 1 shows that the ratio of nominal non-farm inventories to nominal sales of domestic product indeed shows a secular decline. Evidently, the implied deflator inventories has not been rising as fast as that for final sales. We do not attempt to explain the differences between the nominal and real series. We do note, however, the nominal ratio shows persistence comparable to that of the real ratio. The estimate of the first-order autocorrelation of the ratio is 0.97 whether or not we allow a different mean inventory–sales ratio for the 1947:I–1973:IV and 1974:I–1996:IV subsamples.

To return to the secular behavior of the real series: we see from column 3 in Table 4 that the rough constancy of the overall ratio hides some heterogeneity in underlying components. In particular, raw materials, and to a lesser extent, work in progress, have been growing more slowly than the aggregate, implying a declining ratio to final sales. This fact was earlier documented by Hester (1994), who noted that possible explanations include just-in-time inventory management, outbasing of early stages of manufacturing to foreign countries, and a transitory response to transitory movements in costs.

In the sequel we do not attempt to explain secular patterns in inventory–sales ratios; see Hester (1994) for a discussion of US data, for retail as well as manufacturing, West (1992a) and Allen (1995) for discussions of Japanese data. Instead we hope that the reader will take the message away from these tables that *inventories and sales are positively related in the long run*: they tend to rise together. This is illustrated quite strikingly in Figure 2, which is a scatterplot of the inventory and sales data. A second message in the tables and the autocorrelations reported above is that while inventory movements are small relative to GDP, they are volatile and persistent. Characterizing and explaining the stochastic, and especially business cycle, behavior of inventories is the subject of the rest of this chapter.

2. Two stylized facts about inventory behavior

Our essay focuses on the business cycle aspects of inventory behavior, and is oriented around two stylized facts: (1) *inventory movements are procyclical,* (2) the inventory–sales relationship is highly *persistent* (the inventory–sales relationship is our term for a linear version of the inventory–sales ratio). These facts serve two purposes. First, they demonstrate the potential role of inventories in understanding economic fluctuations. Second, they serve as a measure by which we judge inventory models and, more generally, theories of the business cycle.

For each of the two "facts", we present illustrative evidence from annual, post-World War II data, for the G7 countries, as well as from quarterly post-War US data. We then review estimates from the literature. For the first of our stylized facts (procyclical movements), Section 2.1.1 below presents estimates, Section 2.1.2 presents the review. Sections 2.2.1 and 2.2.2 do the same for the second of our facts (persistence in the inventory–sales relationship). The remainder of this introductory subsection describes the data used in both 2.1 and 2.2.

For the G7 countries, we continue to use the aggregate (nation-wide) change in inventory stocks used in previous sections, and construct a time series of inventory levels by summing the change [1]. We measure production as GDP and sales as final sales. The quarterly US inventory data are that used in the previous section, total non-farm inventory and final sales of domestic product in chained 1992 dollars, and with sales measured at quarterly rates [2].

[1] When we summed the ΔH_t series, we initialized with $H_0 \equiv 0$. Given the linearity of our procedures, the results would be identical if we instead used the ΔH_t series to work forwards and backwards from the 1995 levels reported in Table 3. The reader should be aware that when prices are not constant, a series constructed by our procedure of summing changes typically will differ from one that values the entire level of stock at current prices. Those with access to US sources can get a feel for the differences by comparing the inventory change that figures into GDP (used in the G7 data, and in NIPA Tables 5.10 and 5.11) and the one implied by differencing the series for the level of the stock (used in our quarterly US data and NIPA Tables 5.12 and 5.13).

[2] We repeated some of our quarterly calculations using final sales of goods and structures, which differs from total final sales because it excludes final sales of services. There were no substantive changes in results.

All of these measures are linked by the identity production = sales + (inventory investment), or

$$Q_t = S_t + \Delta H_t, \tag{2.1}$$

where Q_t is production, S_t are sales, and H_t is end of period t inventories. This relationship holds by construction, with S_t being final sales.

2.1. Procyclical inventory movements

2.1.1. Illustrative evidence

Procyclicality of inventory movements can be documented in several ways. A simple indication that inventories move procyclically is a positive correlation between inventory investment and final sales. Consider the evidence in Table 5. In column 1 we see that all the point estimates of the correlation are positive, with a typical value being 0.1–0.2.

The correlation between sales and inventory investment is related to the relative variances of production and sales. As in Table 5, let "var" denote variance, "cov" covariance. Since (2.1) implies $\mathrm{var}(Q) = \mathrm{var}(S) + \mathrm{var}(\Delta H) + 2\,\mathrm{cov}(S, \Delta H)$, it follows from the positive correlation in column 1 that $\mathrm{var}(Q) > \mathrm{var}(S)$ (column 2). Other indications of procyclical behavior include two variance ratios robust to the possible presence of unit autoregressive roots. The column 3 estimates indicate that $\mathrm{var}(\Delta Q_t)/\mathrm{var}(\Delta S_t) > 1$, the column 4 estimates that $E(Q_t^2 - S_t^2) > 0$. [$E(Q_t^2 - S_t^2)$ is essentially an estimate of $\mathrm{var}(Q) - \mathrm{var}(S)$ robust to the presence of unit autoregressive roots; see the Technical Appendix.]

To illustrate the pattern of correlation over different short-run horizons, we present impulse response functions. The responses are based on a bivariate VAR in the level of inventories and sales for the quarterly US data, including eight quarterly lags, a time trend, and breaks in the constant and trend at 1974. In accordance with this section's aim of presenting relatively unstructured evidence, we present responses to a one standard deviation shock to the VAR disturbances themselves, and not to orthogonalized shocks. Figure 3 shows the responses of inventories and sales to a disturbance to the sales equation, Figure 4 the responses to a disturbance to the inventory equation. To prevent confusion, we note that on the horizontal axis, we plot the pre-shock (period -1) values of the variables; the shock occurs in period 0.

Figures 3 and 4 both show a positive comovement of inventories and sales. In Figure 3, by construction the contemporaneous (period 0) response of inventories is zero. But the 7 billion (approximately) dollar rise in sales in period 0 is followed in the next quarter by a 1.5 billion dollar increase in inventories. Inventories continue to rise for the next five quarters, even after sales turn down. Both series smoothly decline together.

Figure 4 shows that after a 3 billion dollar shock to inventories, sales rise by nearly 2 billion dollars. Both inventories and sales subsequently show some wiggles.

Table 5
Relative variability of output and final sales[a–e]

Country	Period	(1) corr($S, \Delta H$)	(2) var(Q)/var(S)	(3) var(ΔQ)/ var(ΔS)	(4) $1+[E(Q^2-S^2)/$ var(ΔS)]
Canada	1956–1995	0.14	1.16	1.53	1.41
	1974–1995	0.17	1.21	1.55	1.24
France	1956–1995	0.17	1.36	1.65	1.68
	1974–1995	0.32	1.63	2.09	1.41
West Germany	1957–1994	0.12	1.10	1.36	1.01
	1974–1994	0.13	1.08	1.27	1.03
Italy	1960–1994	0.13	1.30	1.81	1.12
	1974–1994	0.11	1.27	1.83	1.08
Japan	1956–1995	0.23	1.07	1.10	1.30
	1974–1995	0.51	1.15	1.08	1.12
UK	1956–1995	0.28	1.21	1.52	1.10
	1974–1995	0.26	1.17	1.38	1.04
USA	1956–1995	0.26	1.19	1.48	1.12
	1974–1995	0.25	1.21	1.50	0.98
USA	1947:I–1996:IV	0.30	1.26	1.39	1.41
	1974:I–1996:IV	0.14	1.13	1.40	1.48

[a] "var" denotes variance, "corr" correlation, Q = output, S = final sales, ΔH = change in inventories. The variables are linked by the identity $Q=S+\Delta H$.
[b] In all but the last row, data are annual and real (1990 prices), with Q = real GDP, S = real final sales, ΔH = real change in aggregate inventories. In the last row the data are quarterly and real (1992 prices), with S = final sales of domestic business goods and structures, ΔH = change in non-farm inventories, and $Q \equiv S+\Delta H$. See the text and Data Appendix for sources.
[c] In columns 1 and 2, Q and S were linearly detrended, with the full sample estimates allowing a shift in the constant and trend term in 1974 (1974:I in the last row); ΔH was defined as the difference between detrended Q and S. In columns 3 and 4, ΔQ and ΔS were simply demeaned, again with the full sample estimates allowing a shift in the mean in 1974 (1974:I).
[d] In column 4, the term $E(Q^2-S^2)$ essentially is the difference between the variance of Q and the variance of S, computed in a fashion that allows for unit autoregressive roots in Q and S. See the Technical Appendix for further details.
[e] The post-1973 sample, as well as the post-1973 shifts in the full sample estimates, were included to allow for the general slowdown in economic activity.

This shock appears to have more persistent effects than does the sales shock, with inventories still over 2 billion dollars above their initial level after six years.

The important point is that both sets of impulse response functions offer the same picture of procyclical inventories as the statistics in Table 5. Thus, inventories seem to amplify rather than mute movements in production.

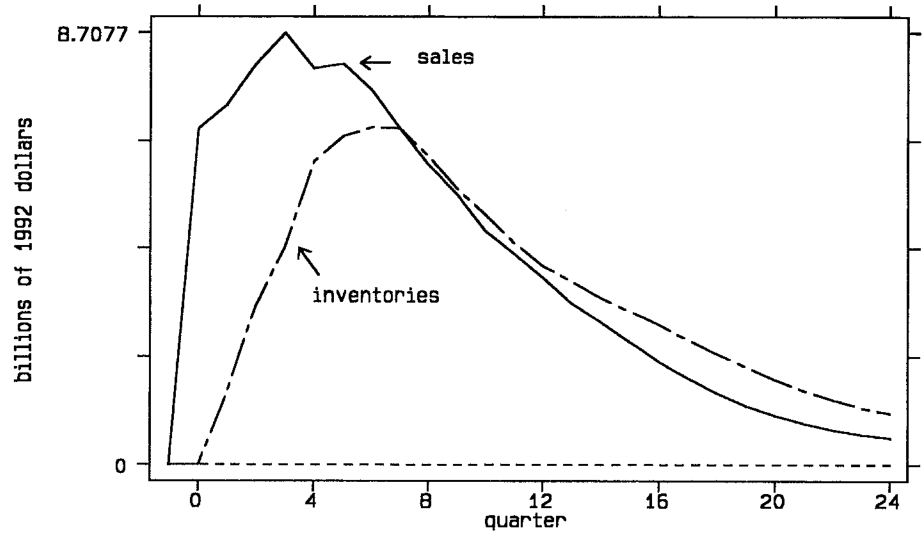

Fig. 3. Response to sales equation shock, quarterly VAR.

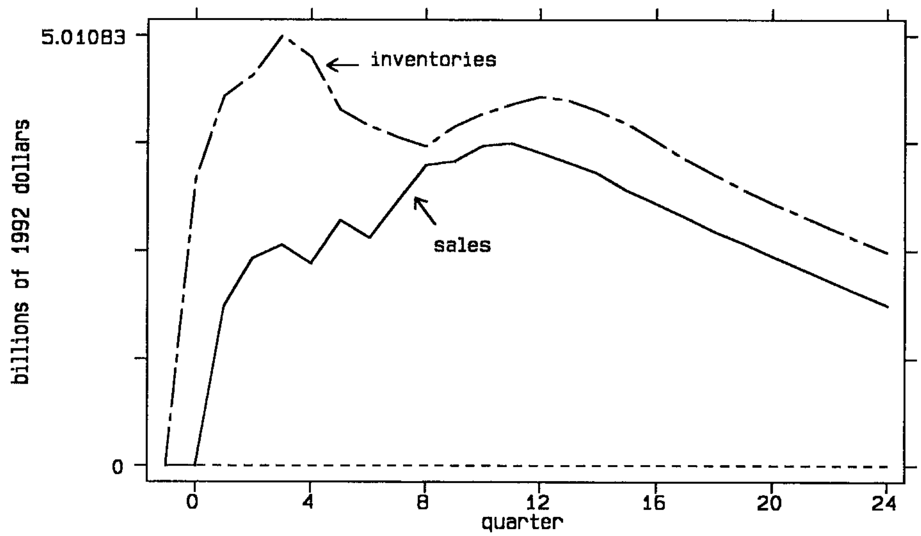

Fig. 4. Response to inventory equation shock, quarterly VAR.

2.1.2. A survey of results

As many readers no doubt are aware, similar findings have been reported for many though not all data sets. A brief summary of estimates of variance inequalities in studies using aggregate data: Fukuda and Teruyama (1988) report comparable results for industrialized countries, but also conclude that by contrast in less developed countries GDP tends to be smoother than final sales[3]. Beaulieu and Miron (1992) report that in manufacturing in some industrialized countries, seasonals in production are no less variable than those in sales.

For US quarterly economy-wide or monthly two-digit manufacturing data, the pattern is as pronounced as in Table 5. This applies first of all to demeaned or detrended data such as is reported in Table 5. For example, Blinder (1986a) reports $\mathrm{var}(Q)/\mathrm{var}(S) > 1$ for 18–20 two-digit manufacturing industries. It also applies to deterministic seasonals [West (1986), Miron and Zeldes (1988), Cecchetti et al. (1997)]: in US manufacturing, seasonal variation in production tends to be larger than seasonal variation in sales[4].

Finally, for both deterministic and stochastic terms, studies that have taken sampling error into account sometimes but not always find it quite unlikely that sampling error alone accounts for the lack of evidence production smoothing [West (1986, 1990b)].

Does *aggregation* substantially account for the lack of production smoothing evident in these studies, either because of measurement error in aggregate data sets, or heterogeneity across firms? Probably not. It has been argued persuasively that disaggregate data measured in physical units are more accurate than the aggregate data used in most studies [see Fair (1989), Krane and Braun (1991), and Ramey (1991)][5]. But as summarized below, studies with disaggregate data still find production more variable than sales in many cases.

For evaluation of the effects of firm heterogeneity, analytical arguments are not particularly helpful. If $\mathrm{var}(Q) > \mathrm{var}(S)$ for a single firm, the inequality may be shown analytically to apply to an aggregate of firms if each individual firm solves a linear quadratic problem such as the one presented below with identical parameters, regardless of the correlation of demand shocks across firms [West (1983)]. But if the cost functions are different for different firms, analytical results appear not to be available. Lai (1991) and Krane (1994) show by example that aggregate variance ratios might differ substantially from individual firm ratios. And even if we assume identical parameters across firms, which as just noted implies that $\mathrm{var}(Q)/\mathrm{var}(S) \leqslant 1$ in the aggregate, this aggregate ratio may be larger (or smaller) than that of individual firms, with the direction of the bias depending on the correlation across firms of demand shocks.

So to consider possible biases from aggregation we must turn from analytical to empirical studies of data disaggregated to the firm or perhaps physical product level.

[3] So far as we know, there have been no systematic attempts to explain this finding. It is possible that measurement error plays a large role.

[4] Carpenter and Levy (1998) report a related finding: for manufacturing, the spectra of inventory investment and production show very high coherence at seasonal frequencies.

[5] And the G7 data that we use for illustration are among the worst measured. In some countries, the change in inventories apparently is constructed at least initially as the difference between product and income estimates of GDP, and thus includes the statistical discrepancy. See West (1990a).

The disaggregate picture is broadly similar though less striking than the aggregate one, at least in the relatively well-studied USA. Production smoothing is markedly absent in the automobile industry [Blanchard (1983), Kashyap and Wilcox (1993)]. But Krane and Braun (1991) found production is less variable than sales in about two thirds of a set of 38 physical products. Finally, Schuh (1996) found deseasonalized production less variable than deseasonalized sales in only one fourth of 700 manufacturing firms; deterministic seasonals in production were less variable than those in sales in about half the firms. For the deseasonalized data, Schuh reports that the median ratio of production to sales variance was about 1.1, which is consistent with the US figure reported in column 2 of Table 5.

2.2. Persistent movements in the inventory–sales relationship

Our second stylized fact is that the inventory–sales relationship is highly persistent. While there may be a steady state linear relationship between inventories and sales, movement towards that steady state is very slow. This characteristic may be more recognizable to inventory experts if it is stated as a "slow speed of adjustment"; a link between persistence in the inventory–sales relationship and the speed of adjustment is demonstrated formally in Sections 4 and 5 below. Section 2.2.1 illustrates this with the annual G7 and quarterly US data used above, Section 2.2.2 documents it with citations to various papers.

2.2.1. Illustrative evidence

To illustrate persistence, we use a standard technique described in the Technical Appendix to attempt to find a stationary relationship between the levels of inventories and sales. This yields $H_t - \widehat{\theta} S_t$ for a parameter $\widehat{\theta}$ estimated from the data. Those familiar with the literature on cointegration will recognize $H_t - \widehat{\theta} S_t$ as the (estimated) error-correction term if inventories and sales have unit autoregressive roots and are cointegrated. We refer to $H_t - \widehat{\theta} S_t$ as the *inventory–sales relationship*, since it is a generalized linear version of the inventory–sales ratio. More precise terminology for this variable is "deviation from the long-run inventory–sales relationship." In our view, the disadvantages of the length of this term outweigh the advantages of being more precise.

After obtaining $\widehat{\theta}$, we estimate the first two autocorrelations in the putatively stationary variable $H_t - \widehat{\theta} S_t$, to gauge the extent of persistence in the relationship. We emphasize that our aim is merely to document quickly evidence of high persistence, not to work towards a complete time series model, nor even to endorse the use of unit roots as a modeling device; among other tasks, the latter would require testing for unit roots in H_t and S_t, perhaps allowing for a one-time shift in mean in 1974, and so on. Rather, this is a way of organizing facts and theories about inventories that has some rough plausibility for a wide range of data.

Table 6
Persistence in stationary linear combinations of inventories and sales[a,b]

Country	Period	(1) Estimate of cointegrating parameter $\widehat{\theta}$	(2) First two autocorrelations of $H_t - \widehat{\theta}S_t$	
Canada	1956–1995	0.16	0.92	0.82
France	1956–1995	0.31	0.95	0.89
West Germany	1957–1994	0.27	0.93	0.82
Italy	1960–1994	0.45	0.88	0.80
Japan	1956–1995	0.22	0.97	0.91
UK	1956–1995	0.20	0.95	0.87
USA	1956–1995	0.24	0.88	0.82
USA	1947:I–1996:IV	0.68	0.94	0.88
USA	1947:I–1996:IV, H and S in logs	1.12	0.95	0.90

[a] See notes to Table 5 for description of data and variable definitions.
[b] In the inventory–sales relationship $H_t - \widehat{\theta}S_t$, the estimate $\widehat{\theta}$ is obtained with the Stock and Watson (1993) procedure to estimate a cointegrating vector. Details are in the Technical Appendix. If the inventory and sales series have unit autoregressive roots, the adjective "stationary" used in the title to this table applies only if the inventory and sales series in fact are cointegrated, and then only asymptotically.

Table 6 presents results. The Technical Appendix describes computational details. Columns 1 and 2 present the estimates of $\widehat{\theta}$ and the first two autocorrelations of $H_t - \widehat{\theta}S_t$. In column 1, the estimates of $\widehat{\theta}$ are all positive, suggesting that inventories and sales move together positively in the long run[6]. But even if this is the case, the result that we wish to emphasize is in column 2: all of the first-order autocorrelations are above 0.8. There is of course a downward bias in estimates of autocorrelations near 1, and an additional bias imparted by time aggregation (recall that the data are annual or quarterly rather than monthly). The column 2 estimates thus suggest that mean reversion takes place quite slowly[7].

[6] The technique used is not invariant to normalization. For the annual data, we re-estimated with S_t on the left instead of H_t. Positive estimates of $(1/\theta)$ resulted. On the other hand, if time trends are included, negative estimates of θ result for 4 of the 7 annual data sets.
[7] Consistent with the high autocorrelation, the null of no cointegration cannot be rejected at the 5% level in three of the seven annual data sets (Canada, Japan, USA), nor in the quarterly US data. This suggests that for these three countries an appropriate model might be one in which there is no steady-state linear relation between inventories and sales. As we shall see, the model to be presented rationalizes such lack of a steady state with unit root cost shocks. When we quickly investigated the extent to which a one-time change in regime in 1974 could account for the persistence, we got mixed results. With post-1974 annual data, the autocorrelations fell; what is perhaps notable is that in one case (the USA) the fall was dramatic, to 0.28 (vs. the 0.88 reported in Table 6). Next, using the quarterly data, we re-estimated using

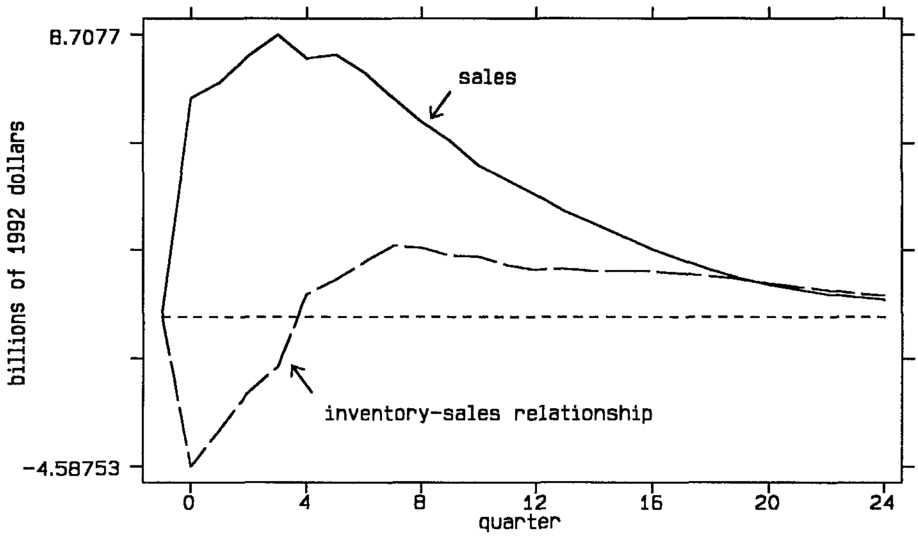

Fig. 5. Response to sales equation shock, quarterly VAR.

We note that this generalization applies when we use logs instead of levels of inventories and sales (last row of the table). We noted above that the inventory–sales ratio displays high autocorrelations. The upshot of the last row is that transforming to logs, and allowing $\theta \neq 1$ – that is, considering not $\log(H_t/S_t)$ but $\log(H_t/S_t^{\theta})$ for a freely estimated θ – still suggests considerable persistence[8].

To illustrate the persistence from a different perspective, Figures 5 and 6 depict the response of sales and the inventory–sales relationship to shocks, using the same VAR estimated for Figures 3 and 4. The sales responses are identical to those in Figures

the whole sample but allowing for different means in the 1947:I–1973:IV and 1974:I–1996:IV sample. This reduced the estimate only slightly, to 0.89 and 0.94 (vs. the 0.94 and 0.95 values in Table 6). Rossana (1998) carefully investigates the question of stability in two-digit US data. He finds evidence of instability, and lack of cointegration within subsamples.

[8] Now seems an appropriate time to comment on the use of logs versus levels of the data, in response to comments by two of the readers of an earlier version of this paper. Much empirical work in inventories, and a distinct majority of work using intertemporal dynamic models, has relied on levels rather than logarithms of the variables in question. We follow that convention in most of this paper. Working in levels has the advantage of preserving the identity that links inventories, production and sales [see Equation (2.1)]. In addition, inventory investment, as reported by the National Income and Product Accounts, is defined as the change in the level of inventories, and is frequently negative, which further discourages the use of logarithms. In a few cases, such as in the last row of Table 6, we do use logarithms. That little turns on levels versus logs is suggested by the similarity of the results in the last two rows in Table 6, and, more generally of the review of the literature that we are about to present: flexible accelerator studies typically use data in logs, while linear quadratic studies typically use levels. Both find great persistence.

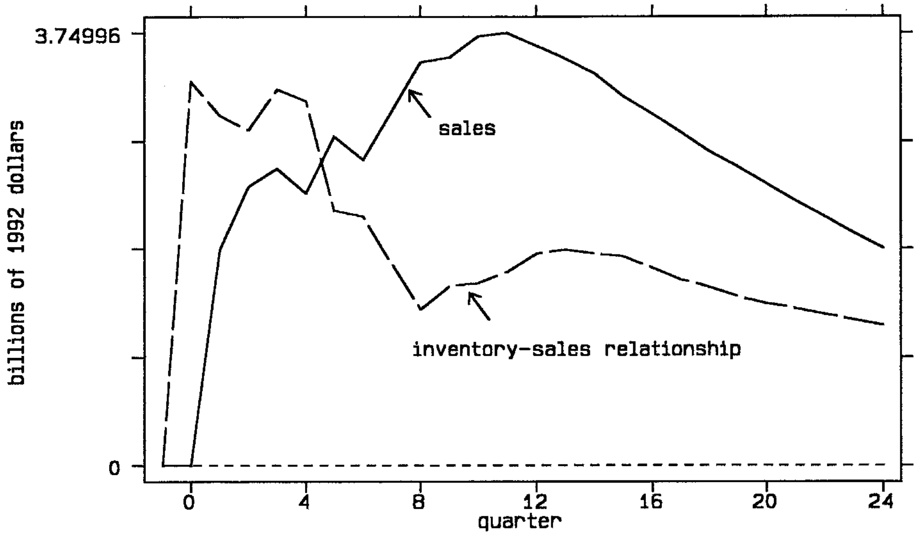

Fig. 6. Response to inventory equation shock, quarterly VAR.

3 and 4. The inventory–sales relationship is computed using the estimated value of $\widehat{\theta} = 0.68$ from Table 6 to form a linear combination of the impulse response functions for sales and inventories.

In response to a positive one standard deviation shock to the sales equation, inventories rise, but less rapidly than do sales (see Figure 3). This induces the fall in the inventory–sales relationship depicted in Figure 5. Eventually, inventories build up more rapidly, leading to overshooting and an increase relative to the initial position. Since both inventories and sales are close to their pre-shock values after six years (see Figure 3), the inventory relationship returns as well, as Figure 5 shows. By contrast, a shock to the inventory equation leads to a complicated pattern displaying more persistence. The inventory–sales relationship rises initially by construction, and then declines erratically for eight quarters before briefly rising again and then slowly decaying. At the end of six years, a return to the pre-shock value is not evident: these unrestricted VAR estimates suggest great persistence in the inventory–sales relationship.

2.2.2. *A survey of results*

Table 6 and Figures 5 and 6 suggest that there is little mean reversion in economy-wide inventory–sales relationships. Congruent evidence comes from two sources, unstructured tests such as in Table 6, and structural estimates of what is called a "speed of adjustment".

Consider first the unstructured tests. Using quarterly economy-wide data, West (1990b, 1992b) cannot reject the null of no cointegration for the USA and for

Japan. Using monthly US data, Granger and Lee (1989) find only mild evidence for cointegration in monthly two-digit manufacturing and trade series; fewer than a third of their 27 data sets reject the null of no cointegration at the 5% level, and the median first-order autocorrelation of the putatively stationary linear combination of inventories and sales is about 0.9 [9]. Rossana (1993, 1998) uses similar data and allows a vector of cost variables W_t to enter the cointegrating relationship. He uses regression techniques to search for a combination $H_t - \theta S_t - \alpha' W_t$ that is stationary. W_t is defined to include real wages, real materials prices, nominal interest rates and inflation. In the end, however, he finds mixed evidence of cointegration across H_t, S_t and W_t. (He does not report autocorrelations of stationary linear combinations.)

In addition, a large literature has used structural inventory models to estimate the "speed of adjustment". We shall describe such models below. For the moment, what is relevant is that under conditions described below, a slow speed of adjustment is equivalent to persistence in the inventory–sales relationship.

Let ρ be the largest autoregressive in $H_t - \theta S_t$, or $H_t - \theta S_t - \alpha' W_t$, with the latter variable the relevant one if, as in the Rossana (1993, 1998) papers cited above, a vector of cost variables is entered into the cointegrating relationship. The large empirical literature has estimated inventory equations, with results implying that $\widehat{\rho}$ is near 1. A typical value in quarterly data is around 0.8–0.95.

The following are some examples. Using quarterly economy-wide data for some industrialized countries, Wilkinson (1989) found $\widehat{\rho} \sim 0.75–1.0$, with W_t including one or more of: sales shocks, inflation, wages, raw materials prices and capacity utilization. Similar results have repeatedly been found using monthly or quarterly US manufacturing data. Examples include Maccini and Rossana (1981), Blinder (1986b), and Haltiwanger and Maccini (1989). In these papers, W_t included one or more of: factor prices such as wages, raw materials prices and real and nominal interest rates, other factors of production such as labor, and sales expectational errors.

Does *aggregation* across heterogeneous firms substantially account for this high serial correlation? We are not aware of analytical arguments establishing general conditions under which aggregate estimates of ρ will be higher than typical firm-specific estimates, though no doubt such arguments could be constructed. Simulations in Lovell (1993) and a limited amount of empirical evidence suggests that aggregation does impart a bias. Schuh (1996) reports that for monthly data for 700 manufacturing firms, the median estimate of ρ is about 0.6 (W_t includes a real interest rate); for quarterly data for some publicly owned firms, and with W_t including measures of credit conditions, $\widehat{\rho}$ is reported to be about 0.6–0.8 by Carpenter et al. (1994, 1998), although a somewhat higher value of about 0.9 is found by Kashyap et al. (1994) [10].

[9] The autocorrelation was computed from the Durbin–Watson statistic reported in the last column of Table II in Granger and Lee, using $2(1 - \widehat{\rho}) = \text{d.w.}$

[10] A small literature has discussed how estimation of ρ is affected when the decision interval for firms is smaller than the sampling interval of the data [Christiano and Eichenbaum (1989), Jorda (1997)]. Such time aggregation does not appear capable of explaining the high persistence.

Of course, tests for cointegration have notoriously low power. And there is heterogeneity of estimates of ρ implied by the flexible accelerator literature. But these results suggest that there is little evidence that mean reversion of $H_t - \theta S_t$ towards its mean takes place rapidly, and considerable evidence of persistence. To interpret such persistence, as well as the procyclical behavior discussed in the previous section, we now present a standard inventory model.

3. Linear quadratic models

3.1. Introduction

The linear quadratic inventory model dates back to Holt et al. (1960). Although this book is written in an operations research style that suggests that its main aim was to provide practical advice to managers, it may still be profitably reviewed by economists interested in inventory behavior. In fact, Holt et al. (1960) develop models more general than those in many of the applications we review here: they allow for stockout costs, order backlogs, and stochastic variation in costs.

Early uses of a linear quadratic model to interpret macro data include Childs (1967) and Belsley (1969). Tools developed in the 1970s and 1980s to estimate and interpret decision rules and first-order conditions from rational expectations models [e.g., Hansen and Sargent (1980), Hansen and Singleton (1982)] were subsequently used by many authors to estimate one or another version of the model. Some papers used aggregate data (sometimes economy wide, sometimes at the two-digit level); some used data at the level of individual firms or physical products.

It is this more recent literature that we review in this and the next four sections. The focus of this literature, and of our discussion, is on how production parameters and constraints influence the intertemporal interaction between production, sales and inventories. While different papers of course vary in the details of the model, there is sufficient commonality that the model presented in Section 3.2 may fairly be described as representative. Section 3.3 derives a first-order condition. Section 3.4 discusses whether the model is applicable to a wide range of inventories.

Table 7 lists notation, and may be a useful reference in this and subsequent sections.

3.2. A model

We assume that a firm maximizes the present value of future cash flows. In macro applications this will be a representative firm. In the formal statement of the model about to be given, we omit constant, linear and trend terms for notational simplicity. Earlier work provides tediously detailed treatment of such terms, allowing for trends that are either arithmetic [West (1983)] or geometric [working paper versions of West (1988, 1990b)].

Table 7
Variable and parameter definitions[a]

A. Definitions of basic variables

H_t	inventories at end of period t
$H_t - H_t^*$	stationary linear combination of H_t, S_t and W_t when U_{dt} and W_t have unit autoregressive roots, but u_{ct} does not; $H_t - H_t^* = H_t - (\theta S_t + \alpha' W_t)$
$H_t - \theta S_t$	inventory–sales relationship
Q_t	production in period t
S_t	sales in period t
U_{ct}	cost shifters, $U_{ct} = \tilde{\alpha}' W_t + u_{ct}$
u_{ct}	cost shock, unobservable to the economist
U_{dt}	demand shock, unobservable to the economist
W_t	vector of observable cost shifters

B. Definitions of basic parameters

Mnemonic	Description	Section[b]
a_0	cost of changing production	3.2
a_1	cost of production	3.2
a_2	inventory holding cost	3.2
a_3	accelerator term for inventories	3.2
b	discount factor	3.2
g	inverse of slope of demand curve	4.2
α	$\alpha = -(ba_2)^{-1}(1-b)\tilde{\alpha}$, coefficient vector on W_t in H_t^*	3.3
$\tilde{\alpha}$	coefficient vector on W_t in U_{ct}	3.2
ϕ_c	AR(1) coefficient of u_{ct}, $u_{ct} = \phi_c u_{ct-1} + e_{ct}$	4.2
ϕ_d	AR(1) coefficient of U_{dt}, $U_{dt} = \phi_d U_{dt-1} + e_{dt}$	4.2
Φ_w	AR(1) matrix of coefficients of W_t, $W_t = \Phi_w W_{t-1} + e_{wt}$	4.2
θ	$= a_3 - [a_1(1-b)/ba_2]$, coefficient on S_t in H_t^*	3.3
π_H	an autoregressive root in $H_t - H_t^*$ when $a_0 = 0$	4.2
π_1, π_2	two autoregressive roots in $H_t - H_t^*$ when $a_0 \neq 0$	4.3

[a] All variables and parameters are scalars, with the exceptions of e_{wt}, W_t, α, $\tilde{\alpha}$ and Φ_w. The variables in panel A are introduced in Section 3.2.
[b] Section where parameter is introduced.

We use the following notation: P_t is real price (say, ratio of output price to the wage), S_t real sales, Q_t real production, H_t real end of period inventories, C_t real period costs, b a discount factor, $0 \leqslant b < 1$, and E_t mathematical expectations conditional on information known at time t, assumed equivalent to linear projections. The objective function is

$$\max \lim_{T\to\infty} E_t \sum_{j=0}^{T} b^j (P_{t+j}S_{t+j} - C_{t+j})$$
$$\text{subject to } \begin{cases} C_t = 0.5a_0\Delta Q_t^2 + 0.5a_1Q_t^2 + 0.5a_2(H_{t-1} - a_3S_t)^2 + U_{ct}Q_t, \\ Q_t = S_t + H_t - H_{t-1}. \end{cases} \tag{3.1}$$

The scalar U_{ct} is a cost shock, and, as discussed below, may depend on both observable and unobservable variables. The term $P_{t+j}S_{t+j}$ is revenue. The analysis in this section does not depend on specification of demand or market structure, so we defer discussion of $P_{t+j}S_{t+j}$ until the next section.

The cost function C_t allows two possible roles for inventories. One is a *production smoothing* role, in which inventories facilitate intertemporal allocation of production. (N.B.: in much of the inventory literature, the phrase "production smoothing" references smoothing from demand shocks; we use it to reference smoothing from cost shocks as well.) This role is reflected in the terms in ΔQ_t^2 and Q_t^2. The second role is a *revenue* role, in which inventories allow a firm to satisfy demand that cannot be backlogged. This role is reflected in the "a_3S_t" term in $(H_{t-1} - a_3S_t)^2$; $(H_{t-1} - a_3S_t)^2$ induces an *accelerator* motive. In our discussion of these terms, we assume for the moment that a_1 and a_2 are positive, a_0 and a_3 nonnegative.

The first production smoothing term, $a_0\Delta Q_t^2$, captures increasing costs of changing production and of production. This represents, for example, hiring and firing costs. Not all authors include this term, and, as discussed in Section 7, some empirical tests find estimates of a_0 insignificantly different from zero.

The second production smoothing term, $a_1Q_t^2$, reflects costs of production. It can be interpreted as the second order term in a quadratic approximation to an arbitrary convex cost function associated with a decreasing returns to scale technology. In data with trends, this approximation would likely be around a growth path. (Recall that constant and trend terms are omitted for notational simplicity.)

The accelerator term $a_2(H_{t-1} - a_3S_t)^2$ embodies inventory holding and backlog costs. Consider first when $a_3 = 0$, so that the term becomes $a_2H_{t-1}^2$. Then this can be interpreted as the second order term in a quadratic approximation to an arbitrary convex inventory holding cost function. When $a_3 \neq 0$, the term is intended to reflect backlog (stockout) and batch as well as inventory holding costs, and thus captures a revenue-related motive for holding inventories. Stockout costs arise when sales exceed the stock on hand, perhaps entailing lost sales, perhaps entailing delayed payment if orders instead are backlogged. Batch costs vary inversely with the stock of inventories, since fewer production runs and larger lot sizes imply larger inventory levels on average.

Ceteris paribus, the higher the stock of inventories, the less likely is a stockout and the lower are stockout costs. As well, higher stocks result when the number of batches falls, since lot sizes rise. On the other hand, higher stocks entail higher inventory holding costs. This quadratic term approximates the tradeoff between the two costs, with a_3 rising as stockout costs rise relative to backlog costs. Holt et al.'s (1960) formal derivation of this time invariant approximation to this inherently nonlinear, and time-varying, cost is presented in Section 8. Note in any case that in many applications, inventories are strictly positive and large relative to sales in all time periods, perhaps in some data sets because of aggregation over heterogeneous firms. So in such data, careful treatment of nonlinearity may have limited empirical payoff. Section 8 below discusses alternative approaches.

The final term in the cost function is $U_{ct}Q_t$. This captures exogenous stochastic variation in costs. (We omit terms of the form cost shock $\times\, \Delta Q_t$ and cost shock $\times\, H_t$ to avoid needless algebraic complications.) In some applications $U_{ct} \equiv 0$ and this shock is absent [West (1986)]; in others it is not observed by the economist and follows an exogenous process [Eichenbaum (1989)]; in still others it has both observable and unobservable components [Ramey (1991)]. To cover all three cases, we write

$$U_{ct} = \widetilde{\alpha}' W_t + u_{ct}. \tag{3.2}$$

In (3.2), W_t is a vector of observable components; u_{ct} is unobservable to the economist and follows an exogenous process. In one or another paper, W_t includes variables such as real wages, materials prices and real and nominal interest rates. If there are no such components, $W_t \equiv 0$; if cost shocks are entirely absent, $u_{ct} \equiv 0$ as well. For convenience we refer to u_{ct} as a "shock", even though it may be serially correlated and partially or fully observable.

3.3. *A first-order condition*

An optimizing firm will not expect to increase cash flow by producing one more unit this period, putting the unit in inventory, and decreasing production by one unit next period, all the while holding revenue constant. Formally, upon differentiating the objective function (3.1) with respect to H_t we obtain [11]

$$E_t[a_0(\Delta Q_t - 2b\Delta Q_{t+1} + b^2 \Delta Q_{t+2}) + a_1(Q_t - bQ_{t+1}) + ba_2(H_t - a_3 S_{t+1}) + U_{ct} - bU_{ct+1}] = 0. \tag{3.3}$$

For discussion of the second of our stylized facts (persistence in the inventory–sales relationship), it will be helpful to note a low frequency implication of Equation (3.3), namely, that inventories and sales are cointegrated if S_t is $I(1)$ and U_{ct} is $I(0)$. [Here,

[11] The first-order condition derived assuming that $[\partial(p_{t+j}S_{t+j})/\partial H_t] = 0$. This assumption is not particularly appealing, since $a_3 > 0$ is motivated in part by stockout costs, and presumably increases in H_t will decrease stockouts and increase revenues (increase $p_t S_t$). As noted above, our modeling of stockout costs is crude but given data constraints in some applications the gains from more sophisticated treatments perhaps are small. See Section 8 below for alternative treatments.

we use standard time series notation: a variable that is integrated of order 0 – $I(0)$, for short – is one that is covariance stationary, with a spectrum that is finite and strictly positive at all frequencies. An "integrated" variable – $I(1)$, for short – is one whose arithmetic difference is $I(0)$. $I(1)$ variables are sometimes called "difference stationary", or as having unit autoregressive roots (a term that we used in Section 2).]

To see the cointegration result, write the term in brackets in Equation (3.3) as x_{t+2}, so that (3.3) is $E_t x_{t+2}=0$. If we replace expectations with realizations, we can write

$$x_{t+2} = \eta_{t+2},\ \eta_{t+2} \equiv (x_{t+2} - E_t x_{t+2}) \sim I(0). \tag{3.4}$$

Note that even if some or all of the variables that comprise x_{t+2} are integrated, η_{t+2} will still be stationary. Kashyap and Wilcox (1993) observe that x_{t+2} may be rewritten

$$\begin{aligned} x_{t+2} &= a_0(\Delta Q_t - 2b\Delta Q_{t+1} + b^2\Delta Q_{t+2}) - ba_1(\Delta H_{t+1} + \Delta S_{t+1}) + a_1\Delta H_t \\ &\quad - a_2a_3\Delta S_{t+1} - ba_2\theta\Delta S_{t+1} + ba_2(H_t - \theta S_t) + U_{ct} - bE_tU_{ct+1}, \\ \theta &\equiv a_3 - \frac{a_1(1-b)}{ba_2}. \end{aligned} \tag{3.5}$$

Suppose $U_{ct}\sim I(0)$ but $H_t, S_t \sim I(1)$. Then $\Delta H_t \sim I(0)$ and $\Delta Q_t \sim I(0)$, and Equation (3.4) (stationarity of x_{t+2}) requires $H_t - \theta S_t \sim I(0)$: inventories and sales are cointegrated with cointegrating parameter θ.

We focus on $S_t \sim I(1)$ for concreteness and empirical relevance, but it is worth noting that θ is still economically interpretable in other cases. If S_t and U_{ct} are both $I(0)$, then $\theta = EH_t/ES_t$ is the inventory sales ratio evaluated at steady state values of S_t and H_t. (Recall that we have omitted deterministic terms for notational simplicity.) If S_t and H_t have deterministic drift ($E\Delta S_t \neq 0$) – which is consistent with U_{ct} either $I(0)$ or $I(1)$, and S_t either $I(1)$ or $I(0)$ around trend[12] – then $\theta = E\Delta H_t/E\Delta S_t$.

Suppose there are observable cost shifters. If $U_{ct} = \tilde{\alpha}' W_t + u_{ct}$ with $W_t \sim I(0)$ or $W_t \sim I(1)$ [see Equation (3.2)] then from similar logic $u_{ct} \sim I(0) \Rightarrow H_t - \theta S_t + (ba_2)^{-1}(1-b)\tilde{\alpha}' W_t \equiv H_t - \theta S_t - \alpha' W_t \equiv H_t - H_t^* \sim I(0)$: after controlling for possibly nonstationary observable cost shifters, inventories and sales are cointegrated. (Once again, $\theta S_t + \alpha' W_t$ remains economically interpretable under other conditions on presence or absence of unit roots.)

We summarize the preceding two paragraphs as follows: if $S_t \sim I(1)$,

$$\begin{aligned} & H_t - H_t^* \sim I(0), \\ & U_{ct} \sim I(0) \quad \Rightarrow \quad H_t^* \equiv \theta S_t, \quad \theta \equiv a_3 - a_1\frac{1-b}{ba_2}, \\ & u_{ct} \sim I(0) \quad \Rightarrow \quad H_t^* \equiv \theta S_t + \alpha' W_t, \quad \alpha \equiv -\frac{1-b}{ba_2}\tilde{\alpha}. \end{aligned} \tag{3.6}$$

Observe that this result does not require parametrization of the demand curve, or specification of market structure. Rossana (1995, 1998) assumes exogenous revenue

[12] A variable x_t is $I(0)$ around trend if $Ex_t = m_0 + m_1 t$ for some $m_1 \neq 0$ and $x_t - Ex_t = I(0)$. Such variables are sometimes called trend stationary.

and uses the resulting decision rule to provide a complementary proof that $U_{ct} \sim I(0)$ $[u_{ct} \sim I(0)]$ implies cointegration between H_t and S_t (between H_t, S_t and W_t).

3.4. Whose inventories?

The preceding description of the model suggests that manufacturers' inventories of finished goods are a natural area for applying the model. Indeed, a large fraction of the inventory literature focuses on manufacturers' finished goods inventories, often in six two-digit industries that are known as "production to stock".[13] To some, in fact, the model is not a particularly attractive one for studying any other types of inventories [e.g., Blinder and Maccini (1991)].

In connection with our discussion of Table 4's sectoral breakdown of US inventories, however, we noted that manufacturers' inventories of finished goods are a small and not particularly volatile component of total inventories. If, indeed, the model is applicable only to finished goods inventories, then it has limited relevance for aggregate inventory fluctuations. We conclude this section by briefly noting that this arguably constitutes an unadvisedly narrow reading of the model.

First, works in progress inventories apparently function as a buffer in many industries, particularly ones that are production to order [West (1988)], and thus fit naturally into the model. More importantly, however, it is possible that an invisible hand causes a large aggregate to solve an optimization problem such as (3.1), using transactions not explicitly modeled in Equation (3.1) to do so. This point was first made by Blanchard (1983), who combined data from the production and retail sectors of the auto industry, modelling explicitly the transactions between the two sectors, and showed that the industry as a whole solved an optimization problem such as (3.1). This is an illustration of the equivalence between a social planning and a decentralized equilibrium, which is well known to occur under general circumstances.

We are therefore sympathetic to the view that there has been too strong a focus on manufacturers' inventories of finished goods. But we feel the implication may be not that the model is of limited relevance, but that it, or related models that capture similar forces, should be applied more widely.

4. Decision rule

4.1. Introduction

Section 4.2 briefly reviews alternative treatments of demand, and solves for a decision rule under a simple specification. Section 4.3 derives the process followed by the

[13] "Production to stock" industries are ones that typically sell finished goods off the shelf. By contrast, "production to order" industries are ones that maintain order backlogs, often deferring final assembly until orders are in hand. See Abramowitz (1950) and Belsley (1969).

inventory–sales relationship, and may be skipped without loss of continuity. Section 4.4 summarizes the implications of Section 4.3 for persistence in the inventory–sales relationship. We present a detailed treatment of the inventory–sales relationship because this process has not received much direct attention in existing literature. Discussions of procyclicality in subsequent sections will cite analytical results in existing literature.

4.2. Derivation of decision rule

Derivation of a decision rule requires specification of revenue because current and expected sales and output appear in the first-order conditions. We begin by reviewing some alternatives. We then work through in detail a specification that is relatively tractable.

In some applications, sales is taken as exogenous, with cost minimization (i.e., $\min E_t \sum_{j=0}^{\infty} b^j C_{t+j}$) the objective[14]. Examples include Holt et al. (1960), Belsley (1969) and Blanchard (1983). The use of this assumption in industry-wide data is valid only if the demand curve facing the industry is vertical.

In other applications, an industry equilibrium is analyzed, and a linear demand curve is specified [e.g., Eichenbaum (1984)]. We write such a demand curve in inverse form as

$$P_t = -gS_t + g\widetilde{U}_{\mathrm{d}t} \equiv -gS_t + U_{\mathrm{d}t}. \tag{4.1}$$

In (4.1), $\widetilde{U}_{\mathrm{d}t}$ and $U_{\mathrm{d}t}$ are stochastic. The demand curve is written in this form so that exogenous sales is a special case of Equation (4.1), implemented by letting the parameter $g \to \infty$ and specifying $\widetilde{U}_{\mathrm{d}t}$ as exogenous; with $g \to \infty$, $S_t = \widetilde{U}_{\mathrm{d}t}$. In practice, one needs to allow for serial correlation in $\widetilde{U}_{\mathrm{d}t}$. In principle one might want to rationalize such serial correlation with (say) costs of adjustment on the part of purchasers, or with observable shifters of the demand curve [West (1992b)]. But since the model focuses on production, and, moreover, is typically not used to study the effects of a hypothetical intervention or change in regime, taking such serial correlation as exogenous is a useful simplification that will be maintained here.

Finally, Christiano and Eichenbaum (1989) and West (1990b), building on Sargent (1979, ch. XVI) derive the linear demand curve (4.1) in general equilibrium. Both papers assume a representative consumer whose per period utility is quadratic in S_t and linear in leisure. The disturbance $U_{\mathrm{d}t}$ is a shock to the consumer's utility. There is no capital; the only means of storage is inventories. See the cited papers for detail.

[14] To prevent confusion, we note that the first-order condition (3.3) also results if one assumes cost minimization. So if one aims to use condition (3.3) to estimate model parameters, one can motivate the equation by reference to cost minimization without taking a stand on the how revenue is determined (apart from the caveat stated in Footnote 11).

Here, we do not derive Equation (4.1) in general equilibrium but take (4.1) as given. We do not attempt to trace the demand shock back to preferences or other primitive sources. We therefore caution the reader that despite the label "demand", U_{dt} should not be thought of as literally a nominal or monetary shock, since it (like all our variables) is real. As well, one can imagine scenarios in which U_{dt} reflects forces typically thought of as supply side. If the good in question is an intermediate one, for example, one can imagine that shocks to the technology of the industry that produces the good dominate the movement of U_{dt}.

Whatever the interpretation of U_{dt}, we derive an industry equilibrium assuming a representative firm. Even so, to obtain a decision rule, we must be specific about market structure and the structure of the demand and cost shocks. We assume here that the market is perfectly competitive, and normalize the number of firms to one. If the firm is a monopolist, the reduced form is identical, but with a certain parameter being the slope of the marginal revenue curve rather than the slope of the demand curve.

We assume that W_t, U_{dt} and u_{ct} follow exogenous AR processes, possibly with unit autoregressive roots. By "exogenous" we mean "predictions of W_t, U_{dt} and u_{ct} conditional on lagged Ws, U_ds and u_cs are identical to those conditional on lagged Ws, U_ds, u_cs, and industry-wide Hs and Ss: W_t, U_{dt} and u_{ct} are not Granger-caused by industry-wide H_t or S_t". (Of course in general equilibrium, such exogeneity of W_t is doubtful.)

Finally, for notational simplicity, and to make contact with the literature on the "speed of adjustment" (see the next section), we tentatively assume that

$$a_0 = 0. \tag{4.2}$$

This assumption is arguably not a good one empirically, and we will relax it below.

Under the assumption of perfect competition, there are two equivalent methods for deriving the decision rules for inventories and sales. The first method, which studies the decentralized optimization problem, derives the individual firm's first-order conditions and then incorporates those into the industry equilibrium. The second method, which uses a social planning approach, derives the first-order conditions for the social planner problem and obtains decision rules for those. Both methods yield identical answers.

We exposit here the decentralized method, and present in the Technical Appendix the social planner approach. With $a_0 = 0$, the first-order condition for sales S_t for the representative firm is

$$P_t - E_t[a_1 Q_t - a_2 a_3(H_{t-1} - a_3 S_t) + U_{ct}] = 0. \tag{4.3}$$

In the absence of inventories, this would simply tell our competitive firm to set marginal revenue P_t equal to marginal cost $a_1 Q_t + U_{ct}$. The additional term $a_2 a_3(H_{t\,1} - a_3 S_t)$ is the effect on inventory holding costs of an additional unit produced for sale.

Upon using $P_t = -gS_t + U_{dt}$ in Equation (4.3) and $Q_t = S_t + \Delta H_t$ in Equation (4.3) and the inventory first-order condition (3.3), we obtain a pair of linear stochastic difference equations in H_t and S_t. This two equation system is solved in the Technical Appendix. The resulting decision rule is

$$H_t = \pi_H H_{t-1} + \text{ distributed lag on } u_{ct}, U_{dt} \text{ and } W_t, \tag{4.4a}$$

$$S_t = \pi_S H_{t-1} + \text{ a different distributed lag on } u_{ct}, U_{dt} \text{ and } W_t. \tag{4.4b}$$

In (4.4a), π_H is the root to a certain quadratic equation, $|\pi_H| < 1$. Both π_H and π_S depend on b, g, a_1, a_2 and a_3. (Note two differences from the relatively well-understood case of exogenous sales. Even when $a_1, a_2 > 0$, if $g < \infty$: (a) it is in principle possible to have $\pi_H \leqslant 0$, and (b) the accelerator coefficient a_3 affects π_H.) The distributed lag coefficients on u_{ct}, U_{dt} and W_t depend on b, g, a_1, a_2 and a_3 as well as the autoregressive parameters governing the evolution of the u_{ct}, U_{dt} and W_t. In the empirically relevant case of $\pi_H > 0$, π_H increases with marginal production costs a_1 and decreases with marginal inventory holding costs a_2. The signs of $\partial\pi_H/\partial a_3$ and $\partial\pi_H/\partial g$ are ambiguous.

The solution when revenue is exogenous ($g \to \infty$) is obtained by replacing U_{dt} with $g\widetilde{U}_{dt}$ [see Equation (4.1)] and letting $g \to \infty$. In this case, $\pi_S = 0$, $S_t = \widetilde{U}_{dt}$ and the solution (4.4) may be written in the familiar form

$$H_t = \pi_H H_{t-1} + \text{ distributed lag on } S_t \text{ and on measures of cost}, \tag{4.5a}$$

$$S_t \sim \text{ exogenous autoregressive process}. \tag{4.5b}$$

On the other hand, when revenue is endogenous, $\pi_S \neq 0$ and we see in Equation (4.4b) that inventories Granger-cause sales. The intuition is that forward looking firms adjust inventories in part in response to expected future conditions. Thus industry-wide stocks signal future market conditions, including sales. This signalling ability is reflected in Equation (4.4b).

These same results can be obtained directly from the social planner problem that maximizes consumer surplus plus producer surplus, which is equal to the area between the inverse demand and supply curves. See the Technical Appendix.

In empirical application, matching the data might require allowing shocks with rich dynamics. Such dynamics may even be required to identify all the parameters of the model. Blanchard (1983), for example, assumes that the demand shock follows an AR(4). For expositional ease, however, we assume through the remainder of this section that all exogenous variables – U_{dt}, u_{ct}, W_t – follow first-order autoregressive processes (possibly with unit roots). Specifically, assume that

$$\begin{aligned}
&E_{t-1}W_t = \Phi_w W_{t-1}, \qquad W_t = \Phi_w W_{t-1} + e_{wt}, \qquad E_{t-1}e_{wt} = 0,\\
&|I - \Phi_w z| = 0 \Rightarrow |z| \geqslant 1,\\
&E_{t-1}u_{ct} = \phi_c u_{ct-1}, \qquad u_{ct} = \phi_c u_{ct-1} + e_{ct}, \qquad E_{t-1}e_{ct} = 0, \qquad |\phi_c| \leqslant 1,\\
&E_{t-1}U_{dt} = \phi_d U_{dt-1}, \qquad U_{dt} = \phi_d U_{dt-1} + e_{dt}, \qquad E_{t-1}e_{dt} = 0, \qquad |\phi_d| \leqslant 1.
\end{aligned} \tag{4.6}$$

(Given the growth in the number of symbols, it may help to remind the reader that Table 7 summarizes notation.) The Technical Appendix shows that the distributed lags in Equation (4.4) are all first order, and Equation (4.4) is

$$H_t = \pi_H H_{t-1} + f'_{H\mathrm{w}} W_t + f_{H\mathrm{c}} u_{\mathrm{c}t} + f_{H\mathrm{d}} U_{\mathrm{d}t}, \tag{4.7a}$$

$$S_t = \pi_S H_{t-1} + f'_{S\mathrm{w}} W_t + f_{S\mathrm{c}} u_{\mathrm{c}t} + f_{S\mathrm{d}} U_{\mathrm{d}t}. \tag{4.7b}$$

See the Technical Appendix for explicit formulas for the "f"s in terms of b, g, a_1, a_2, a_3, Φ_{w}, ϕ_{c} and ϕ_{d}. Of course, if $\tilde{\alpha} = 0$ so that $U_{\mathrm{c}t} = u_{\mathrm{c}t}$, then $f_{H\mathrm{w}} = f_{S\mathrm{w}} = 0$.

4.3. Persistence in the inventory–sales relationship

To analyze the second of our stylized facts (persistence in the inventory–sales relationship), we now further assume that the demand shock and the observable cost shifters follow random walks:

$$\phi_{\mathrm{d}} = 1, U_{\mathrm{d}t} = U_{\mathrm{d}t-1} + e_{\mathrm{d}t}, \Phi_{\mathrm{w}} = I, W_t = W_{t-1} + e_{\mathrm{w}t}. \tag{4.8}$$

Recall that if $U_{\mathrm{c}t}$ is stationary, $H_t - \theta S_t$ is stationary as well, where the cointegrating parameter θ is defined in Equation (3.5).

When there are no observable cost shifters ($\tilde{\alpha} = 0 \Rightarrow f_{H\mathrm{w}} = f_{S\mathrm{w}} = 0$), tedious manipulation of Equation (4.7) yields

$$H_t - H_t^* \equiv H_t - \theta S_t = \pi_H (H_{t-1} - \theta S_{t-1}) + m_{0\mathrm{c}} u_{\mathrm{c}t} + m_{1\mathrm{c}} u_{\mathrm{c}t-1} + m_{0\mathrm{d}} e_{\mathrm{d}t}, \tag{4.9}$$

where $m_{0\mathrm{c}}$, $m_{1\mathrm{c}}$ and $m_{0\mathrm{d}}$ depend on θ, $f_{H\mathrm{c}}$, $f_{H\mathrm{d}}$, $f_{S\mathrm{c}}$ and $f_{S\mathrm{d}}$ (see the Technical Appendix). Let "L" be the lag operator. Since $(1 - \phi_{\mathrm{c}} L) u_{\mathrm{c}t} = e_{\mathrm{c}t}$, it follows from Equation (4.8) that

$$\begin{aligned} &(1 - \pi_H L)(1 - \phi_{\mathrm{c}} L)(H_t - H_t^*) = v_t, \\ &v_t = m_{0\mathrm{c}} e_{\mathrm{c}t} + m_{1\mathrm{c}} e_{\mathrm{c}t-1} + m_{0\mathrm{d}} e_{\mathrm{d}t} - \phi_{\mathrm{c}} m_{0\mathrm{d}} e_{\mathrm{d}t-1} \sim \mathrm{MA}(1). \end{aligned} \tag{4.10}$$

Thus, $H_t - H_t^* \sim \mathrm{ARMA}(2, 1)$ with autoregressive roots π_H and ϕ_{c}. (This presumes that the moving average root in v_t does not cancel an autoregressive root in $H_t - H_t^*$, which generally will not happen.) Note that the innovation $e_{\mathrm{d}t}$, rather than the shock $U_{\mathrm{d}t}$, appears in Equation (4.9) and thus in Equation (4.10). With $\phi_{\mathrm{d}} \neq 1$, however, the right hand side of Equation (4.10) would include a linear combination of $U_{\mathrm{d}t}$ and $U_{\mathrm{d}t-1}$ that would not reduce to a linear function of $e_{\mathrm{d}t}$, and ϕ_{d} would also be one of the autoregressive roots of $H_t - H_t^*$. In this case, if $\phi_{\mathrm{d}} \sim 1$, then $H_t - H_t^*$ would also have a moving average root that would approximately cancel the autoregressive root of ϕ_{d}.

Similarly, when there are observable cost shifters ($\alpha \neq 0$), it may be shown that Equations (4.6) and (4.7) imply

$$\begin{aligned} &H_t - H_t^* \equiv H_t - \theta S_t - \alpha' W_t = \pi_H (H_{t-1} - \theta S_{t-1} - \alpha' W_{t-1}) + \text{disturbance}, \\ &\quad \text{disturbance} = m'_{0\mathrm{w}} e_{\mathrm{w}t} + m_{0\mathrm{c}} u_{\mathrm{c}t} + m_{1\mathrm{c}} u_{\mathrm{c}t-1} + m_{0\mathrm{d}} e_{\mathrm{d}t}. \end{aligned} \tag{4.11}$$

Once again, persistence in $H_t - H_t^*$ is induced by π_H and ϕ_{c}.

We close this subsection by re-introducing costs of adjusting production a_0. Suppose

$$a_0 \neq 0. \tag{4.12}$$

It is well known that when revenue is exogenous ($g \to \infty$), costs of adjusting production put additional persistence in inventories [Belsley (1969), Blanchard (1983)]: in this case Equation (4.5a) becomes

$$H_t = \pi_{H1} H_{t-1} + \pi_{H2} H_{t-2} + \text{distributed lag on } S_t \text{ and on measures of cost,} \tag{4.13}$$

with $\pi_{H2} \neq 0$. Unsurprisingly, inventory decisions now depend on $Q_{t-1} = S_{t-1} + H_{t-1} - H_{t-2}$ and thus on H_{t-2}, even after taking into account H_{t-1} and the sales process.

As one might expect, the presence of costs of adjusting production has a similar effect even when sales and revenue are endogenous, and on the inventory–sales relationship as well as inventories. The Technical Appendix shows that $a_0 \neq 0$ puts an additional autoregressive root in $H_t - H_t^*$, which now follows an ARMA(3,2) process. One autoregressive root is ϕ_c. We let π_1 and π_2 denote the two additional (possibly complex) roots. These are functions of b, a_0, a_1, a_2, a_3 and g. Intuition, which is supported by the simulation results reported below, suggests that increases in a_0 increase the magnitude of these roots.

4.4. Summary on persistence in the inventory–sales relationship

We summarize the preceding subsection as follows: assume the shocks follow the AR(1) processes given in Equation (4.6), with the additional restriction (4.8) that the demand shock and observable cost shifters follow random walks. Then

$$\begin{aligned} &a_0 = 0 \Rightarrow H_t - H_t^* \equiv H_t - \theta S_t - \alpha' W_t \sim \text{ARMA}(2, 1), \\ &\quad \text{with AR roots } \pi_H \text{ and } \phi_c. \end{aligned} \tag{4.14}$$

The root π_H is a function of b, g and the a_i, but not the autoregressive parameters of the shocks, and is increasing in the marginal production costs a_1. In addition,

$$\begin{aligned} &a_0 \neq 0 \Rightarrow H_t - H_t^* \equiv H_t - \theta S_t - \alpha' W_t \sim \text{ARMA}(3, 2), \\ &\quad \text{with AR roots } \pi_1, \pi_2 \text{ and } \phi_c; \\ &\text{if} \quad \phi_c = 0, \quad H_t - H_t^* \sim \text{ARMA}(2, 1) \\ &\quad \text{with AR roots } \pi_1 \text{ and } \pi_2. \end{aligned} \tag{4.15}$$

The roots π_1 and π_2 are functions of b, g and the a_i, but not the autoregressive parameters of the shocks; both analytical manipulations of the formulas in the

Technical Appendix and simulations reported in Section 6 indicate that the modulus of the larger of the roots increases with a_0 and a_1 [15].

Thus the persistence documented in Section 2.2 above follows if there are sharply increasing production costs (a_0 and/or a_1 are sufficiently large) and/or serially correlated cost shocks. In addition, it is important to observe that qualitatively similar reduced forms are implied by the following two scenarios: (1) serially correlated cost shocks with no costs of adjusting production, and (2) serially uncorrelated cost shocks and sharply increasing costs of adjusting production. We shall return to this point below.

Of course persistence may also follow if we put different dynamics into the shocks W_t, u_{ct} and U_{dt}.

5. The flexible accelerator model

We now derive (4.10)–(4.11) from another optimization problem. This optimization problem is one that underlies empirical work motivated by the *flexible accelerator* model. In this model, pioneered by Lovell (1961), firms solve a static one period problem, balancing costs of adjusting inventories against costs of having inventories deviate from their frictionless target level H_t^*. Specifically, the firm chooses H_t to minimize

$$0.5(H_t - H_t^*)^2 + 0.5v(H_t - H_{t-1})^2 + u_t H_t. \tag{5.1}$$

In (5.1), $v > 0$ is the weight of the second cost relative to the first, and u_t is an exogenous unobservable disturbance [16]. The first-order condition is then

$$H_t - H_{t-1} = [1/(1+v)](H_t^* - H_{t-1}) - [1/(1+v)]u_t. \tag{5.2}$$

The coefficient $1/(1+v)$ is the fraction of the gap between target and initial inventories closed within a period. If v is big (cost of adjusting inventories is big), the fraction of

[15] Under the present set of assumptions, then, the parameter called "ρ" in Section 2.2 is $\max\{\pi_H, \phi_c\}$ if $a_0 = 0$, $\max\{|\pi_1|, |\pi_2|, \phi_c\}$ if $a_0 \neq 0$.

[16] H_t and S_t are sometimes measured in logs [e.g., Maccini and Rossana (1981, 1984)], and the variable u_t is sometimes split into a component linearly dependent on the period t surprise in sales and a component unobservable to the economist [e.g., Lovell (1961), Blinder (1986b)]. We slur over differences between regressions in levels and logs, which in practice are small (see Footnote 8), and omit a sales surprise term in the inventory regression, which in practice has little effect on the coefficients that are central to our discussion.

the gap expected to be closed is, on average, small. To make this equation operational, target inventories H_t^* must be specified. Let

$$H_t^* = \theta S_t + \alpha' W_t. \tag{5.3}$$

Here, W_t is a vector of observable cost shifters [as in Section 2.2.2 and Equation (3.2)]. Notation has been chosen because of link about to be established with θ and $\alpha' W_t$ as defined earlier. Suppose

$$S_t = S_{t-1} + e_{\mathrm{d}t}, \qquad W_t = W_{t-1} + e_{\mathrm{w}t}, \tag{5.4}$$

$E_{t-1}e_{\mathrm{d}t} = 0$, $E_{t-1}e_{\mathrm{w}t} = 0$. (In practice, $E_{t-1}S_t$ is usually approximated as a linear function of a number of lags of S, the actual number dependent on the data, and similarly for W_t [e.g., Maccini and Rossana (1984)]. The single lag assumed here is again for simplicity.) Then with straightforward algebra, the first-order condition (5.2) implies

$$\begin{aligned} &H_t - \theta S_t - \alpha' W_t = \pi_H (H_{t-1} - \theta S_{t-1} - \alpha' W_{t-1}) + \text{disturbance}, \\ &\quad \pi_H = [\nu/(1+\nu)], \quad \text{disturbance} = [1/(1+\nu)](\theta e_{\mathrm{d}t} + \alpha' e_{\mathrm{w}t} - u_t), \end{aligned} \tag{5.5}$$

which is in the same form as Equation (4.11).

We have thus established that in the simple parameterization of this section, in which sales follows an exogenous random walk, high serial correlation in a stationary linear combination of inventories and sales is the same phenomenon as slow speed of adjustment of inventories towards a target level.

6. Dynamic responses

To develop intuition about how the model works, and what the two stylized facts suggest about model parameters and sources of shocks, this section presents some impulse responses. Specifically, we present the industry equilibrium response of (1) H_t, S_t and Q_t, or (2) H_t, S_t and $H_t - \theta S_t$, to a shock to $U_{\mathrm{d}t}$ or $u_{\mathrm{c}t}$, for various parameter sets, with no observable cost shifters ($\alpha = W_t = 0$). While the parameter values we use are at least broadly consistent with one or another study, we choose them not because we view one or more of them as particularly compelling, but because they are useful in expositing the model.

Table 8 lists the parameter sets. It may be shown that *the solution depends only on relative values of* g, a_0, a_1 and a_2; multiplying these 4 parameters by any nonzero constant leaves the solution unchanged. [This is evident from the first-order conditions (3.3), (B.4) and (B.5): doubling all these parameters leaves the first-order conditions unchanged, apart from a rescaling of the shocks.]

Our choice of $a_2 \equiv 1$ is simply a normalization. We fix $g = 1$ in part because some of the properties documented below can be shown either to be invariant to g [see West (1986, 1990b) on procyclicality of inventories] in part because a small amount of

Table 8
Parameter sets[a]

(1) Mnemonic	(2) g	(3) a_0	(4) a_1	(5) a_2	(6) a_3	(7) ϕ_c	(8) ϕ_d
A	1	0	1	1	0	n.a.[b]	0.7
B	1	0	1	1	1	n.a.	0.7
C	1	0	−0.1	1	0	n.a.	0.7
A'	1	0	1	1	0	n.a.	1
A''	1	0	1	1	0	0.7	n.a.
D	1	3	1	1	0	n.a.	1
E	1	3	1	1	1	n.a.	1

[a] See Table 7 for parameter definitions. The behavior of the model depends only on the scale of the parameters g, a_0, a_1 and a_2; doubling all these leaves behavior unchanged. The discount factor b is set to 0.99 in all experiments.
[b] "n.a." means that the autoregressive parameter is irrelevant for the impulse responses plotted in Figures 7–13: the response is for a shock to the other variable, whose AR(1) parameter is given.

experimentation indicated little sensitivity to g. To prevent possible confusion, we note explicitly that the parameter a_3 is identified in absolute terms and not just relative to other parameters. Throughout, we set the discount factor $b = 0.99$, and interpret the time period as quarterly. To facilitate discussion, in the graphs we set vertical tick marks labelled "−1", "1", "2", and so on, but this (or any other) choice of units to measure the response is arbitrary. [In actual application, the units used would of course be monetary (e.g., billions of 1992 dollars in the impulse responses in Section 2 above)].

The production smoothing aspect of the model is most clearly evident when shocks are mean reverting. We therefore begin with three parameter sets illustrating the response to an innovation in a stationary AR(1) demand shock U_{dt}, with AR parameter $\phi_d = 0.7$. Since $\phi_d = 0.7$ is probably far enough from unity to make the notion of cointegration between H_t and S_t unappealing, we plot the responses of Q_t, S_t and H_t but not those of $H_t - \theta S_t$[17].

Parameter set A illustrates the production smoothing model. Figure 7 presents the response to a demand shock. As may be seen, when there is a positive innovation to demand, sales of course rise. But part of the increase in sales is met by drawing down inventories, thereby buffering production from the demand shock. As sales return to

[17] Naturally, even though we do not include the plots here we did examine them ourselves. As it turned out, $H_t - \theta S_t$ showed persistence. From Equation (B.11) in the Technical Appendix, we see that $H_t - \theta S_t$ has an autoregressive root of ϕ_d that is cancelled by a moving average root only when $\phi_d \to 1$. This autoregressive root apparently explains the persistence. In our view such persistence is not particularly interesting: in a stationary model, $\theta \equiv a_3 - [a_1(1-b)/ba_2]$ is not the parameter corresponding to a projection of H_t onto S_t, and thus $H_t - \theta S_t$ does not correspond to the quantity displaying persistence in, for example, Table 6.

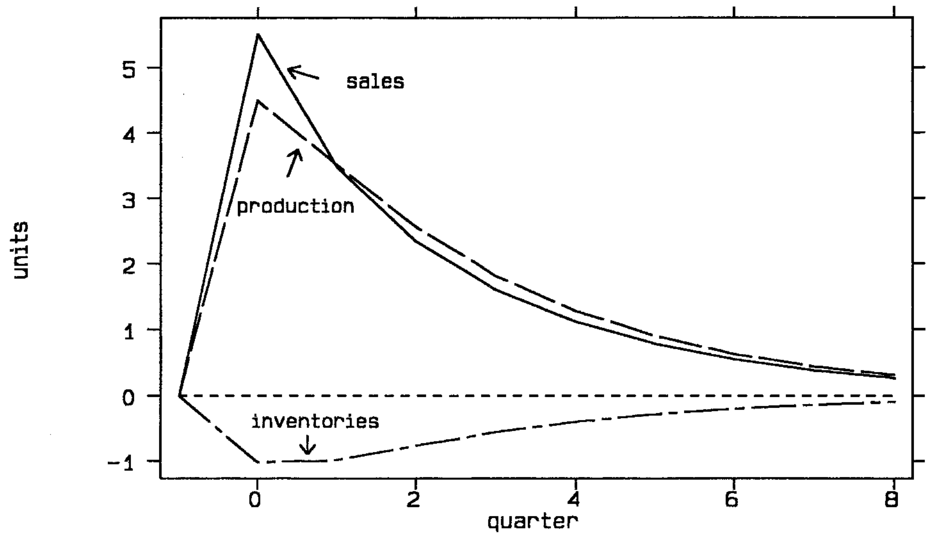

Fig. 7. Response to a stationary demand shock; parameter set A.

the steady state, inventories are gradually built back up. It may be seen in the graph that production is smooth relative to sales.

The intuition is straightforward: given increasing marginal costs ($a_1 > 0$), it is cheaper to produce at a steady rate than to produce sometimes at a high rate, sometimes at a low rate. So the increased demand is met partly with inventories, and production is smoothed relative to sales. (Note that this logic applies even for a competitive firm that can sell as much as it wants at the prevailing market price.) Inventory movements are countercyclical, in the sense that they covary negatively with sales. It may be shown analytically that such countercyclical behavior will obtain when $a_1 > 0$, $a_3 = 0$ and there are no cost shocks [West (1986) for a stationary model, working paper version of West (1990b) for a model with unit roots].

One can obtain procyclical movements when costs are convex if the accelerator term is operative ($a_3 > 0$) and is sufficiently strong to offset the production smoothing motive. In Figure 8, which shows results when $a_3 = 1$ rather than $a_3 = 0$, inventories initially rise along with sales when there is a positive innovation to the stationary demand shock. So production rises even more than does sales, and is more variable. All three variables then fall smoothly back towards the steady state.

Some algebra may help with intuition: if $a_0 = a_1 = 0$, and $U_{ct} = 0$, the first-order condition for Equation (3.1) is simply $H_t = a_3 E_t S_{t+1}$. Thus inventories will covary positively with expected sales, and thus with sales themselves since S_t is positively serially correlated in equilibrium. With $a_0 \neq 0$, $a_1 \neq 0$, inventory movements will reflect a balance of accelerator and production smoothing motives. If the

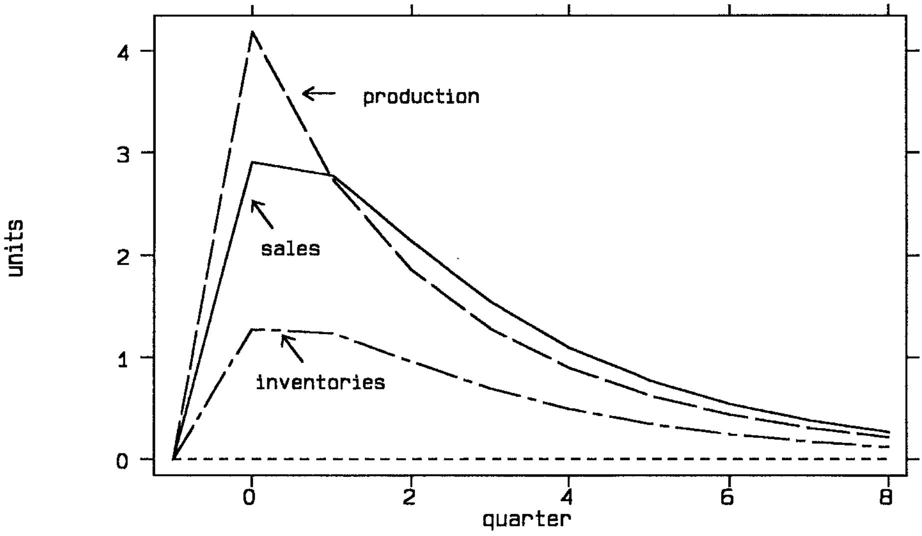

Fig. 8. Response to a stationary demand shock; parameter set B.

accelerator motive dominates, as it does in this parameter set, inventories will move procyclically [18].

Another way to obtain procyclical movements in response to demand shocks, is with nonconvex production costs [Ramey (1991)]. Parameter set C captures this with a small negative value for a_1. [The linear quadratic problem will still be well-posed, and lead to an internal solution, as long as the nonconvexity is not too marked; in the present context, this essentially demands that a_2 and g be sufficiently large relative to $|a_1|$. See Ramey (1991).] We see in Figure 9 that a positive innovation to demand causes inventories to rise (though by a small amount – an artifact of our choice of parameters): with $a_1 < 0$ it is cheaper to bunch rather than smooth production. Thus, firms build up inventories when sales are high. If there is a cost of changing production ($a_0 \neq 0$), marginal production costs are $(1+b)a_0 + a_1$. Ramey (1991) has noted that $a_1 < 0$ may induce a tendency to bunch production even if $(1+b)a_0 + a_1 > 0$ [see West (1990b) for a particular set of parameters for which this happens].

We now turn to parameter sets with a unit root in the demand shock ($\phi_d = 1$). Figure 10 plots the response of inventories, sales and the inventory–sales relationship

[18] Recall that the accelerator term is motivated in part by stockout costs. Kahn (1987) rigorously shows that when nonnegativity constraints are imposed, demand uncertainty (which implies uncertainty about whether a stockout will occur) will lead to procyclical movements if demand is serially correlated.

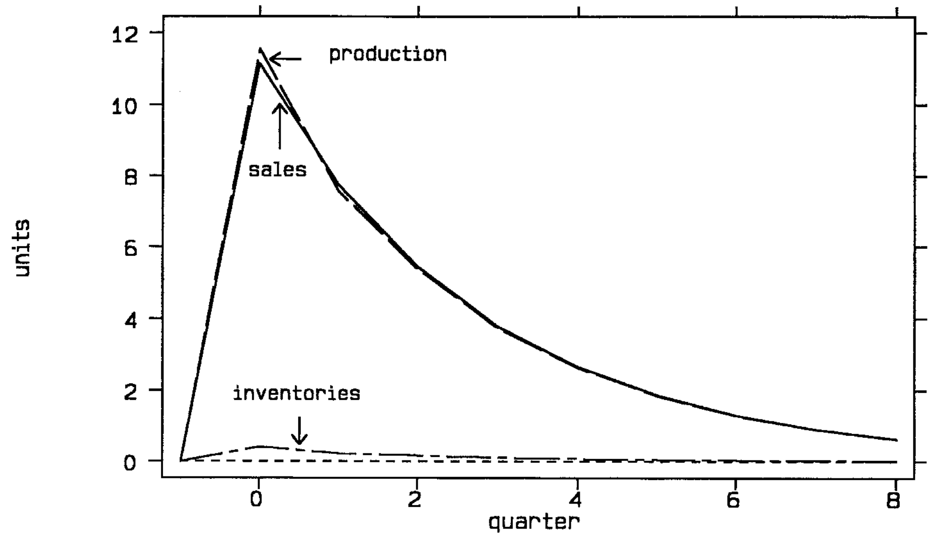

Fig. 9. Response to a stationary demand shock; parameter set C.

$H_t - \theta S_t$, with the technology parameters matching those in parameter set A. As with a stationary shock, firms draw down inventory stocks when demand increases (though the fall is slight in our figure). They will replenish stocks in response to a negative shock (not depicted). Thus inventories buffer production. With our choice of parameters, the transition to the new steady state is quite rapid[19]. In parameter sets B and C, a demand shock with a unit root leads to procyclical inventory movements (not depicted): with or without a unit root in the demand shock, inventories buffer production.

Figure 10 also plots the response of the inventory sales relationship $H_t - \theta S_t$. To understand the pattern it exhibits, some mechanics may be helpful. Since $a_3 = 0$, $\theta < 0$ and $(-\theta) > 0$ [see Equation (3.5)]. In Figure 10 we see that the response of H_t is negative but small in absolute value, that of S_t positive and relatively large; in the end, the net response of $(-\theta)S_t > 0$ is greater than that of $H_t < 0$, and $H_t - \theta S_t$ increases in response to a demand shock. The inventory–sales relationship has little persistence, however; it has a first-order autocorrelation coefficient of a little under 0.3.

When we computed impulse responses for parameter sets B and C with a unit root demand shock (not depicted), the sign of the initial response of $H_t - \theta S_t$ happened to be negative for parameter set B, positive for parameter set C: the sign of the initial response to a demand shock is sensitive to exact parameter values. A characteristic

[19] To prevent confusion: sales and revenue *are* endogenous in this experiment ($g < \infty$). Although S_t looks like a random walk in the figure, in fact ΔS_t does have a little bit of serial correlation, and is Granger-caused by inventories.

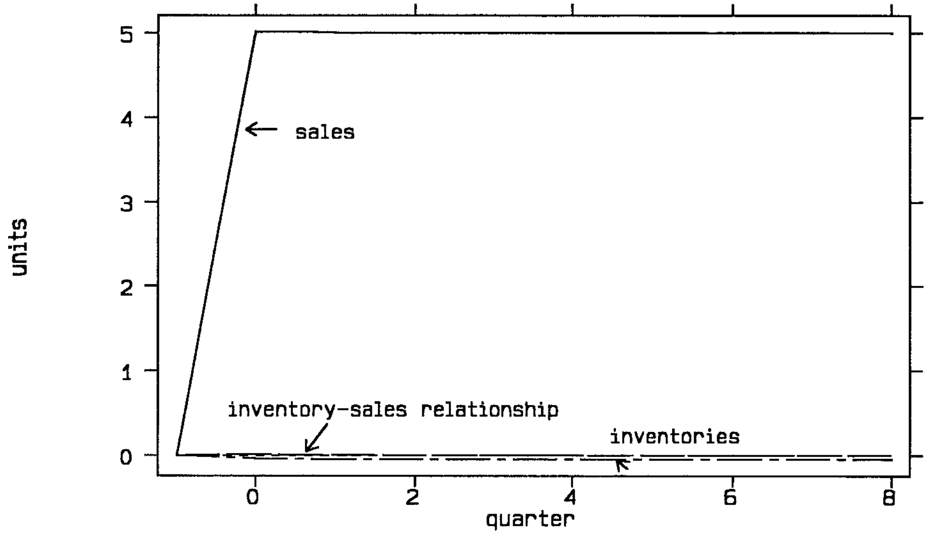

Fig. 10. Response to a unit root demand shock; parameter set A′.

of all three parameter sets, however, was rapid mean reversion in $H_t - \theta S_t$. The autoregressive root π_H [see Equation (4.9)] was 0.270 (parameter set A), 0.269 (B) and −0.148 (C).

The negative sign of π_H in parameter set C perhaps deserves a word of mention. Recall that $a_1 < 0$ in this parameter set, and that downward sloping marginal costs induces production bunching. So bunching generates a negative autocorrelation, since high activity in one period tends to be followed by low activity in the next.

In parameter set A'', we continue to use the same technology parameters, but now plot responses to a cost rather than demand shock. The cost shock is stationary, with an AR parameter of $\phi_c = 0.7$. We see in Figure 11 that inventories move procyclically in response to a negative cost shock: the shock causes both inventories and sales to rise, and makes production (not depicted) more variable than sales. (Recall that we are studying industry equilibrium, and sales of course change as costs change.)

The intuition once again is straightforward: a firm with a convex cost function will use periods of low cost to produce a lot and to build up inventory stocks (as in the figure), and use periods of high cost (not depicted) to produce little and instead sell out of inventory stocks that have already been built up. In this case, inventories serve to buffer production from cost shocks. (Reminder: the phrase "production smoothing" is conventionally used only to describe smoothing from demand shocks but not, as in the present paragraph, smoothing from cost shocks.) Inventory movements are procyclical. It may be shown analytically that such procyclical behavior will obtain when $a_1 > 0$,

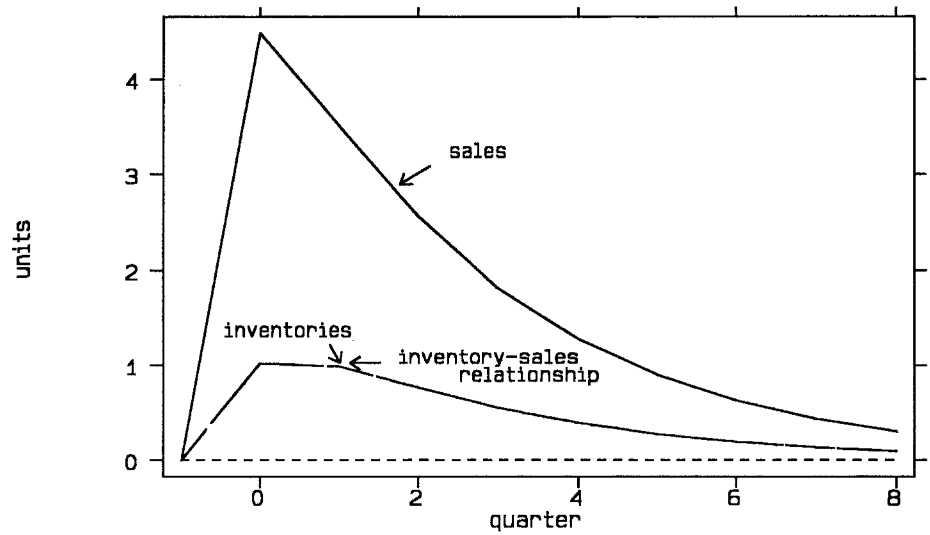

Fig. 11. Response to stationary cost shock; parameter set A″.

$a_3 = 0$ and there are no demand shocks [see the working paper version of West (1990b)]. We see that there is now persistence in $H_t - \theta S_t$[20].

We showed above that when $a_0 = 0$, persistence in $H_t - \theta S_t$ results even in the absence of cost shocks if production costs a_1 are sufficiently large relative to inventory costs a_2. When sales are exogenous ($g \to \infty$), it is straightforward to establish analytically that persistence also increases as a_0 increases relative to a_1 and a_2. We use the final two parameter sets to document via simulation that large values of a_0 also cause persistence when sales are endogenous. Parameter set D varies from A' only by introducing $a_0 = 3$ instead of $a_0 = 0$. Figure 12 plots responses of H_t, S_t and $H_t - \theta S_t$ to a positive innovation in a unit root demand disturbance. From the perspective of buffering, there are no new results: sales increase and inventories are drawn down to buffer production. But now there is persistence in the inventory–sales relationship $H_t - \theta S_t$, which has an autocorrelation a little above 0.7. The intuition is that if costs of adjusting production are large, and if there is a (say) jump in demand, firms only very gradually increase production to build the stock up.

Finally, parameter set E varies from D by allowing the accelerator motive, setting $a_3 = 1$ (as in parameter set B). In Figure 13 we see that inventories now move procyclically in response to demand shocks, and that there continues to be persistence in the inventory–sales relationship.

[20] Of course, for the computed response of $H_t - \theta S_t$ to give us insight into the empirical behavior of $H_t - \theta S_t$, there must be unit roots or near unit roots in H_t and S_t, and hence this experiment and parameter set cannot provide the whole explanation for the two stylized facts.

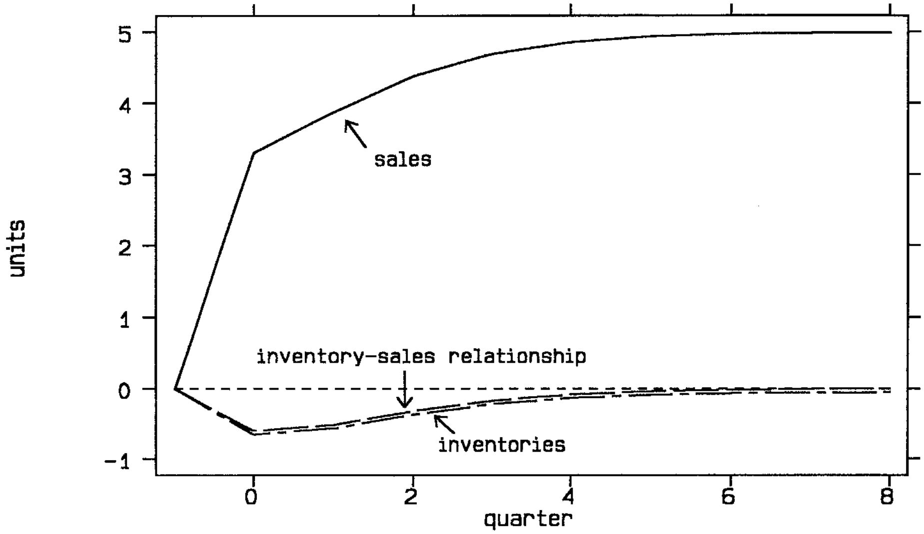

Fig. 12. Response to a unit root demand shock; parameter set D.

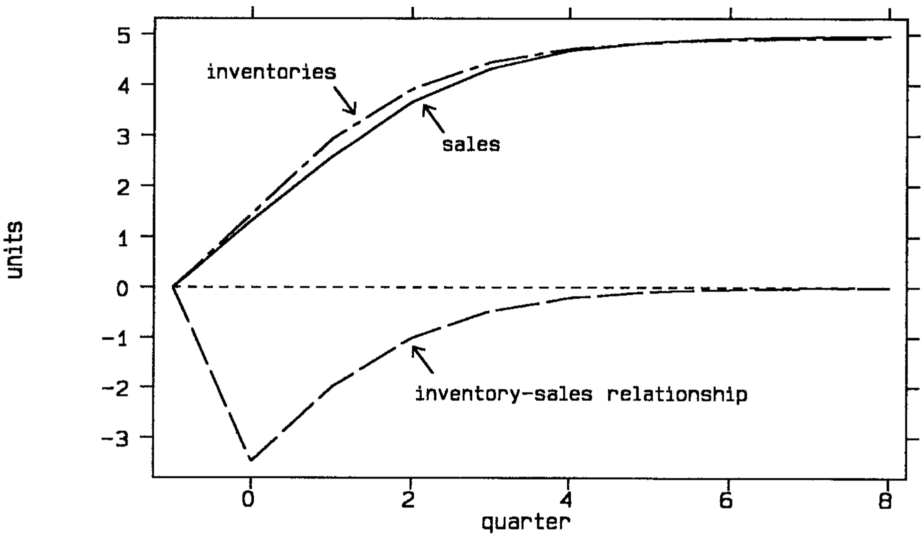

Fig. 13. Response to a unit root demand shock; parameter set E.

Table 9 summarizes the discussion in this section. Evidently, we must consider empirical evidence on: the accelerator motive, the relative values of a_0, a_1 and a_2, and cost shocks.

Table 9
Possible explanations of stylized facts[a]

	Inventories move procyclically	Persistence in inventory–sales relationship[b]
In response to demand shocks:		
(1) Marginal production costs rise very rapidly relative to marginal inventory holding costs (a_2 small relative to a_1 and/or a_0)	no	yes
(2) Declining marginal production costs[c] ($a_1 < 0$)	yes	no
(3) Strong accelerator motive ($a_2 a_3$ large relative to a_0 and a_1)	yes	no
In response to cost shocks;		
(4) Increasing marginal production costs	yes	yes

[a] The parameters are defined in Table 7.
[b] The implications for persistence in the inventory–sales relationship apply if there is a unit root in the demand shock [explanations (1)–(3)] and persistence in the cost shock [explanation (4)].
[c] The reference to "declining marginal production costs" abstracts from costs of changing production (from a_0). With $a_0 \neq 0$, marginal production cost is $a_0(1+b)+a_1$. But even if $a_0(1+b)+a_1 > 0$, inventory movements may be procyclical if $a_1 < 0$.

7. Empirical evidence

7.1. Introduction

The analytical results in West (1986, 1990b) and Section 4 and the simulations in Section 6 suggest at least two different ways of rationalizing the procyclicality of inventory movements and the persistence of the inventory–sales relationship. One is a demand-driven model with rapidly increasing marginal production costs (marginal production costs a_0 and/or a_1 are large relative to marginal inventory holding costs a_2), together with a strong accelerator motive ($a_2 a_3$ large relative to a_0 and a_1). The second is a cost-driven model, with increasing marginal production costs; such a model may or may not have a role for the accelerator. For simplicity we somewhat loosely refer to these as our *demand-driven* and our *cost-driven* explanations. We do so with some reservations: please recall that our demand shock U_{dt} may in some data basically reflect supply side forces.

These two do not exhaust the possibilities, and many economists (including us) would expect both cost and demand shocks to be important over samples of reasonable length. Our own work, for example, has emphasized the possibility of declining marginal production costs [Ramey (1991)]. In combination with highly persistent cost shocks, both procyclicality of inventories and persistence of the inventory–sales relationship may result. And West (1990b) finds both stylized facts explicable with a model with strong costs of adjusting production and a substantial role for both cost and demand shocks, but with no accelerator.

But our demand-driven and cost-driven explanations have the virtue of simplicity, and both have support from a number of papers: as summarized in this section and the next, most aggregate studies, and the limited microeconomic evidence available, do not point to declining marginal cost, and do find a role for the accelerator. In citing such support we do not cast a wide net but instead selectively cite representative papers. In addition, after some introductory remarks on papers using the flexible accelerator model (Section 7.2), we focus on papers that explicitly use the linear quadratic model, for ease of exposition.

Section 7.2 reviews parameter estimates from the linear quadratic literature, Section 7.3 discusses sources of shocks, and Section 7.4 provides an interpretation. We remind the reader that the behavior of inventories depends only on the relative values of g, a_0, a_1 and a_2. All statements referencing "large" values of one of these parameters should be understood to mean "large relative to another parameter or linear combination of parameters". The normalization involved will be clear from the context.

7.2. Magnitude of cost parameters

Our discussion will focus on estimates of the linear quadratic model. We begin, however, with a brief discussion of results from less structured studies, including those using the flexible accelerator. We record two results.

The first is that in flexible accelerator studies, actual or expected sales is generally found to be an important determinant of inventory movements, with a positive relationship between the two series. See, for example, Maccini and Rossana (1981, 1984) or Blinder (1986b). In terms of the model in Section 5, a positive relationship may be interpreted as $\theta > 0$, where θ is the coefficient on sales in the expression for target inventories [see Equation (5.3)]. As well, in direct estimation of a cointegrating parameter, Granger and Lee (1989) do obtain $\widehat{\theta} > 0$ in all 27 of their US two-digit manufacturing and trade series.

To interpret this with the linear quadratic model, recall that under certain conditions, the decision rule from the flexible accelerator model (5.1) can be mapped into that of the linear quadratic model (3.1). Under those conditions, $\theta = a_3 - [a_1(1-b)/(ba_2)]$ [see Equations (3.5), (4.11) and (5.5)]. Thus $a_3 > 0$ is necessary for the cointegrating parameter θ to be positive, as noted by Kashyap and Wilcox (1993). [This holds even when U_{ct} is present (although cointegration requires $U_{ct} \sim I(0)$).] Thus here and in linear quadratic studies (see below) there is support for a nontrivial role for the accelerator motive – a result that may be unsurprising or reassuring to some, but in any event is not particularly helpful in discriminating between our two candidate explanations.

The second result from the flexible accelerator literature concerns the structure of production costs. As discussed in Section 2, this literature has found large autoregressive roots in $H_t - H_t^*$, which implies slow adjustment of H_t towards H_t^*. In quarterly data, a typical estimate of the root is around 0.8–0.9, implying that about

10–20% of the gap between actual and target inventories is closed in a quarter. Ever since Carlson and Wehrs (1974) and Feldstein and Auerbach (1976), many observers have found such estimated speeds puzzling and perhaps not well-rationalized by the flexible accelerator model. One reason is that even the largest quarterly movements in inventories amount to only a few days production. This suggests to Feldstein and Auerbach (1976, p. 376) and others that costs of adjusting inventories [v, in the notation of Equation (5.1)] cannot be very large[21].

To interpret this second result with the linear quadratic model, recall that we set $a_0 \equiv 0$ when we established a mapping from the flexible accelerator to the linear quadratic model (3.1). With $a_0 \equiv 0$, an arbitrarily slow speed of adjustment results when a_1 is arbitrarily large. It is not clear to us how large a value of a_1 is implausibly large. But we take from the flexible accelerator literature the message that many find this simplest version of the model unappealing [see Blinder and Maccini (1991) for a recent statement].

Accordingly we consider the other sources of persistence isolated above: costs of adjustment (discussed in this subsection), and serial correlated cost variables (discussed in the next subsection). To focus the discussion of costs of adjustment, we highlight estimates from some recent linear quadratic studies using two-digit manufacturing data from the USA. Different studies present estimates of a_0, a_1 and a_2 relative to different parameters or linear combinations of parameters. To display results from various studies in consistent form, we restate published estimates of a_0, a_1 and a_2 relative to a common linear combination of the published estimates of those parameters. This linear combination is

$$c \equiv (1 + 4b + b^2)\, a_0 + (1 + b)\, a_1 + b a_2, \tag{7.1}$$

with $b \equiv 0.99$. Here, "c" is the second derivative of the objective function (3.1) with respect to H_t; the Legendre–Clebsch condition states that $c > 0$ is a necessary condition for an optimal solution. [See Stengel (1986, p. 213) or Kollintzas (1989, p. 11).] Note that the estimates we discuss will therefore not be comparable to those used in the simulations in the previous section and in Table 8, and often are not as easily interpreted as those expressed relative to a single parameter. We nonetheless use this normalization since studies sometimes report negative estimates of a_0, a_1 or a_2, which can make interpretation of estimates relative to one of those parameters problematic.

Most authors examine more than one specification. Table 10 presents results for a specification that seemed to be preferred by the author(s). For the preferred specification, columns 2–6 present the median point estimate of a_0/c, a_1/c, $[(1+b)a_0 + a_1]/c$

[21] The logic apparently is that it should be easy to make inventory movements rapid if firms are beginning from a starting point in which current movements are small relative to production. But small inventory movements seem to be exactly what one would associate with slow adjustment speeds, if costs of adjustment determine both the size of movements and the adjustment speeds; if, instead, the slow adjustment speeds were accompanied by large movements in inventories, there would be a puzzling contrast between regression results and basic data characteristics.

Table 10
Median point estimates of model parameters[a–d]

(1) Reference[e]	(2) a_0/c	(3) a_1/c	(4) $[(1+b)a_0+a_1]/c$	(5) a_2/c	(6) a_3	(7) Number of industries
Models with serially correlated cost variables:						
(1) Durlauf and Maccini (1995)	0.	0.43	0.43	0.15	0.55	5
(2) Eichenbaum (1989)	0.	0.21	0.21	0.58	1.15	7
(3) Kollintzas (1995)	−0.16	0.83	0.64	−0.09	1.14	6
(4) Ramey (1991)	0.15	−0.63	−0.43	1.69	0.40	6
Models without serially correlated cost variables:						
(5) Fuhrer, Moore and Schuh (1995)	0.13	0.12	0.38	0.00	0.67	1
(6) West (1986)	0.05	0.34	0.44	0.01	1.12	10

[a] In the column definitions, $c \equiv (1+4b+b^2)a_0 + (1+b)a_1 + ba_2$, $b \equiv 0.995$. Note that the magnitudes in columns 2–4 are therefore not comparable to those in columns 4–6 in Table 8.

[b] Different papers expressed point estimates relative to different linear combinations of parameters. For each paper, the reported point estimates were restated relative to c. The Legendre–Clebsch condition states that $c > 0$ is a necessary condition for an interior solution of the optimization problem. The table reports the median of the restated estimates. When $a_0 \equiv 0$ (lines 1 and 2), or when there is only one industry (line 5), the column 4 entry for marginal production cost is by construction equal to: $(1+b)$ times column 2, plus column 3.

[c] All the studies used two-digit manufacturing data from the USA. The exact data, sample period, specification and estimation technique vary from paper to paper.

[d] Most papers present more than one set of results. We chose the specification that seemed to be favored by the author(s).

[e] Sources by reference: (1) Table 7 (p. 85), entries labelled "Table 3"; (2) Table 2 (p. 861); (3) Tables 1–6 (pp. 77–80), columns labelled "random walk"; (4) Table 1 (p. 323), excluding autos; (5) Table 4 (p. 128), entry labelled "FIML-endogenous sales"; (6) Table 4 (p. 391).

(=marginal production cost, taking into account costs of adjusting production), a_2/c and a_3. The median is computed across the datasets considered by the author; the number of datasets is given in column 7.

A skim of the table suggests a broad consensus on a_3 (column 6). As well, there is relatively little disagreement on the sign of the slope of marginal production costs (column 4); with the exception of Ramey (1991), the studies find an upward slope to marginal production cost. There is, however, some variation in the extent to which the cost of adjustment a_0 contributes to this upward slope. Consistent with the demand-driven explanation, Fuhrer et al. (1995) (line 5) and to a lesser extent West (1986) (line 6) find that a_0 contributes to the upward slope.

Some studies with other datasets have found an even stronger role for the cost of adjustment a_0, with a_0 positive and significant but with estimates of the production cost a_1 negative [consistent with Ramey (1991)], or economically or statistically indistinguishable from zero. For example, Kashyap and Wilcox's (1993) study of the automobile industry in the 1920s and 1930s yielded median estimates of parameters as follows:

$$\begin{array}{ccccc} a_0/c & a_1/c & [(1+b)a_0+a_1]/c & a_2/c & a_3 \\ 0.20 & -0.11 & 0.29 & 0.03 & 0.72 \end{array} \tag{7.2}$$

Similar results are reported for the modern automobile industry by Blanchard (1983) and Ramey (1991), and for US aggregate inventories by West (1990b).

On the other hand, we see in lines 1 and 2 that the preferred specifications in the Eichenbaum (1989) and Durlauf and Maccini (1995) set the cost of adjustment to zero. In these two papers, the estimates of a_1 tended to be positive but perhaps not so large as to imply a speed of adjustment that Feldstein and Auerbach (1976) would find implausibly slow. In part these papers set a_0 to zero – because in a setup similar to that of Kollintzas (1995) in line 3, negative and insignificant point estimates of a_0 tended to result.

Rounding out the cost-driven story requires finding substantial persistence from stochastic variation in costs. This is discussed in the next subsection.

7.3. Shocks

There is much circumstantial evidence that serially correlated cost shifters have important effects on inventory behavior. In particular, the data often seem happy with specifications in which the unobservable disturbance u_{ct} is highly autocorrelated [e.g., Eichenbaum (1989), West (1990b), Ramey (1991)].

One's confidence that this unobservable disturbance really reflects stochastic variation in production costs would be increased if inventories could be shown to respond aggressively to observable measures of costs. Unfortunately, this appears not to be so. In practice, factor prices and interest rates usually are insignificant (in both economic and statistical terms), and sometimes have effects opposite of the theoretical predictions. For statistical significance, Table 11 shows a selection of results using cost variables, from studies of two-digit manufacturing in the USA, and now including flexible accelerator as well as linear quadratic studies.

It may be seen in columns 1–4 that a finding of a statistically significant effect of observable measures of costs is rare: only 2 entries are "y"s, indicating that in only two of the 21 studies did significance at the 5% level characterize at least three-fourths of the coefficients estimated in a given study. 11 entries are "n"s, indicating that in these 11 studies fewer than one-fourth of the coefficients were significant. On the other hand, in column 5 it may be seen that for the unobservable disturbance, three of the 6 entries are "y"s, and that two of these "y"s are for studies that also included some observable measures of costs (lines 6 and 8); none of the 6 entries are "n"s.

7.4. Interpretation

We showed in Section 4 that the demand-driven and cost-driven explanations put two large autoregressive roots in the inventory sales relationship $H_t - \theta S_t$; in fact,

Table 11
Statistical significance of cost variables[a–c]

Reference[d]	Wage	Materials prices	Energy prices	Interest rate	Unobservable shock
(1) Blinder (1986b)	?	?		n	?
(2) Durlauf and Maccini (1995)	?	n	n		
(3) Eichenbaum (1989)					y
(4) Kollintzas (1995)					?
(5) Maccini and Rossana (1981)	y	?		n	?
(6) Maccini and Rossana (1984)	n	y		n	y
(7) Miron and Zeldes (1988)	n	?	n	n	
(8) Ramey (1991)	n	?	n		y
(9) Rossana (1990)	?	?		?	

[a] This table is an updated version of a table in West (1995).
[b] All the studies used two-digit manufacturing data from the USA. The exact data, sample period, specification and estimation technique vary from paper to paper.
[c] A "y" entry indicates that the coefficient on the variable in a given column was significantly different from zero at the 5% level in at least three-fourths of the datasets in a given study, a "n" that it was significant in at most one-fourth of the datasets, a "?" that it was significant in more than one-fourth but fewer than three-fourths of the datasets. A blank indicates that the variable was not examined.
[d] Sources by reference: (1) Table 1 (pp. 360–61); (2) Table 3, inst. set. 4; (3) Table 2 (p. 861); (4) Tables 1–6 (pp. 77–80), columns labelled "HP filter" and "quadratic trend"; (5) Table 1 (p. 20); (6) Table 3 (p. 231) and discussion on p. 227; (7) Table II (p. 892); (8) Table 1 (p. 323); (9) Tables 3 and 4 (pp. 26–27), with the cost of capital variable "ce" and "cp" interpreted as interest rate variables.

under certain conditions, both imply ARMA(2, 1) processes of $H_t - \theta S_t$. Specifically, this happens when $\phi_c = 0$ in the demand-driven explanation [see Equations (4.14) and (4.15)].

That the similar ARMA structures might allow both models to fit a given body of data is illustrated by Kollintzas (1995). Kollintzas' results in line 3 of Table 10 were for a specification with a random walk ($\phi_c = 1$) cost shock (i.e., Kollintzas differenced the first-order condition before estimating). Among other specifications, Kollintzas allowed for an i.i.d. unobservable cost shock. In the specification with i.i.d. cost shocks, the median estimates of the parameters were:

$$\begin{array}{ccccc} a_0/c & a_1/c & [(1+b)a_0 + a_1]/c & a_2/c & a_3 \\ 0.03 & 0.42 & 0.47 & 0.01 & 2.51 \end{array}. \tag{7.3}$$

While the estimate of marginal production costs was not wildly different (0.47 with i.i.d. shocks vs. 0.64 with random walk shocks [line 3, column 4 of Table 10)], the median estimate of a_0 was higher. In fact, the estimate of a_0 was higher in 5 of the 6 datasets[22].

An interpretation is that a model that fits his data would imply two autoregressive roots. When serial correlation in the cost shock was suppressed, the positive values for a_0 rationalized a second autoregressive root; when serial correlation in the cost shock was imposed, large (or even positive) values of a_0 would imply an autoregressive structure too elaborate for the data, and accordingly the regression yielded diminished values of a_0.

Discriminating between the two explanations thus means distinguishing between costs of adjustment [when $a_0 \neq 0$ and the serial correlation of the cost shock is zero ($\phi_c = 0$)] and exogenous serial correlation (when $a_0 = 0$ and ϕ_c is near one). In principle this may be done, using either cross-equation restrictions, or additional variables such those in the W_t vector. But in both inventory and non-inventory contexts this has proved difficult [e.g., Blinder (1986b), McManus et al. (1994), Surekha and Ghali (1997)].

And in any case, our discussion so far perhaps has understated the extent of conflict across empirical results. There is a range of estimates of most parameters (including some wrong-signed or otherwise implausible ones), and we have pushed papers into one of just two camps in the interest of summarizing a complex set of results: while in principle it may be possible to pin down important macroeconomic parameters and sources of shocks by simply estimating linear inventory models with aggregate data, this tantalizing idea has not proved true in practice so far.

The conflict across papers, or the range of estimates, may be no worse than in empirical work in other areas. For example, those familiar with the real business cycle literature will probably not be surprised that it is difficult to find observable counterparts to unobservable cost shocks. And Lovell (1994) shows that the estimated speed of adjustment of H_t towards H_t^* is in fact no slower than those of some other variables. As well, part of the conflict across papers no doubt results from econometric problems related to sample size or estimation technique [West and Wilcox (1994, 1996), Fuhrer et al. (1995)]. Finally, it may be that careful analysis would reveal that seemingly disparate conclusions in fact result mainly from the use of different sample periods, datasets, and observable cost shifters ("W_t", in the notation of the previous sections).

But pointing out (perhaps unfairly!) that other literatures have similar problems will not advance our knowledge about inventories. Nor, most likely, will sharp estimates be produced by even the most refined econometric technique, at least when applied to familiar data. We therefore suggest some alternative approaches.

[22] Such statements potentially are sensitive to how the parameters are expressed (relative to "c", as in Table 10, or some other linear combination of parameters). But in this case the statement applies not only with respect to the normalization we have used, but also with respect to the normalization used by Kollintzas, which was relative to $a_0(1+b)+a_1$.

8. Directions for future research

8.1. Introduction

In this section we offer what we believe to be fruitful directions for future research. Sections 8.2 and 8.3 describe alternative modelling strategies. Some of our suggestions are based on alternatives to the basic linear quadratic production smoothing model; others extend the basic model. All build on the insights delivered by the basic model: that procyclical movements result when inventories facilitate sales (a force captured in the basic model with the accelerator term), and that the shape of production costs influences both the character of cyclical movements and the persistence of the inventory–sales relationship. In addition, all seem intuitively capable of helping explain either or both of our stylized facts (procyclicality of inventories, persistence of the inventory–sales relationship), although, of course, research to date involving these suggestions has its share of blemishes (e.g., wrong-signed parameter estimates). Finally, Section 8.4 describes how the use of different data may help understand inventory behavior.

8.2. Inventories in production and revenue functions

The potential importance of the accelerator term (a_3) in explaining both the business cycle and long-run behavior of inventories suggests that the relationship between inventories and sales deserves more study. Consider first Holt et al.'s (1960) original motivation for this formulation as $a_2(H_t - a_3S_{t+1})^2$. As discussed on pp. 56–57 of their book, their initial model of optimal inventory holdings used lot-size formulas, where the optimal batch size, the number of batches and optimal inventory levels all increase with the square root of the sales rate. They used two approximations to capture the costs and benefits associated with inventory holdings. First, they approximated the square-root relationship with a linear relationship between inventories and sales (e.g., $H_t^{\sim} = a_3S_{t+1}$). Second, they approximated all costs and benefits associated with inventories with a quadratic in which costs rise with the square of the deviation of inventories from the optimal level. This generates the accelerator term $a_2(H_t - H_t^{\sim})^2$.

While this tractable formulation provides a plausible mechanism for procyclical inventory movements, there are two potential problems with it. First, the approximations may be inadequate. As we will discuss below, the approximations used imply that the cost of a marginal reduction in inventories is linear in the stock of inventories, whereas at least one paper found significant convexities. Second, inventories may directly affect revenue in a way that is not well captured by including the accelerator term in the cost function.

One alternative strand of the literature has modelled inventories as factors of production, or considered interrelationships between inventories and other factors of production. Christiano (1988), Ramey (1989), Galeotti et al. (1997) and Humphreys et al. (1997) are examples. Ramey (1989) argues that since inventories at all stages of production facilitate production and shipments, they should be considered factors of production. She includes three stages of inventories in the production function, and estimates factor demand relationships. This approach obviously has the potential to make inventories move procyclically, since factor usage fluctuates with output. The results from the linear quadratic model suggest if costs of adjustment are allowed, persistence would result as well.

A second line of research has considered the revenue role of inventories. Kahn (1987, 1992) develops a theory of a stockout avoidance motive for holding inventories and tests some of its implications using automobile industry data. Kahn argues that demand uncertainty and a nonnegativity constraint on inventories can explain several important patterns in the data. Bils and Kahn (1996) extend this line of research by assuming that the demand function is a (nonlinear) function of the stock of goods available for sale. They apply the model to two-digit manufacturing data with mixed success. Rotemberg and Saloner (1989) offer another potential role for inventories in revenue, arguing that inventories may be accumulated to deter deviations from an implicitly collusive arrangement between firms. The nature of the equilibrium implies that inventories will be high when demand is high. They show empirically that the correlation between inventories and sales is higher in concentrated industries, as predicted by their model.

A third line of research studies uses more general functional forms for the relationship between inventories and sales. [The work by Kahn (1992) and Bils and Kahn (1996) also fits into this category.] Pindyck (1994) studies the convenience yield of inventories for three commodities. Augmenting the usual production, sales and inventory data with futures prices, he provides evidence that the marginal convenience yield is very convex, increasing sharply as inventories approach zero. This indicates that the approximation embodied in the basic Holt et al. (1960) model may miss some important aspects of inventory behavior.

8.3. *Models with fixed costs*

We next consider arguments and evidence that a key shortcoming of the linear quadratic model is that it fails to account for fixed costs facing firms. Blinder (1981), Caplin (1985), Mosser (1988), Blinder and Maccini (1991), and Fisher and Hornstein (1996) all argue that fixed costs of ordering may be very important for understanding the behavior of retail and wholesale inventories as well as manufacturers' materials and supplies. In some environments the aggregation argument presented in Section 3.4 will not apply, and research to date has shown that under certain conditions such fixed costs may lead to (S, s) type of decision rules. In their review article, Blinder and Maccini (1991) recommend that future inventory research concentrate on the (S, s) model. This will require resolution of difficult problems of aggregation, perhaps partly through the use of simulations [Lovell (1996)]. While the results look suggestive at the level of a single-product firm, the implications for a multi-product firm, let alone for an industry or economy, have been harder to obtain because of difficulties of aggregating (S, s) rules.

The studies by Bresnahan and Ramey (1994) and Hall (1996) show that fixed costs are an important determinant of production costs in the automobile industry. Bresnahan and Ramey follow some fifty assembly plants on a weekly basis from 1972 to 1983, and uncover important lumpiness in the margins for varying production. Hall studies fourteen assembly plants from 1990 to 1994. Both studies isolate two important methods for varying production, which appear to involve some sort of nonconvex costs. First, they find that complete shut-down of a plant for a week at a time is an important method for temporarily decreasing output rates. Second, the adding and dropping of extra shifts (each of which doubles or halves production) are an important source of output variability, and appear to involve fixed costs and lumpiness of production levels. Thus, costs in the automobile industry deviate from the linear quadratic production smoothing model in two important ways. First, there appear to be fixed costs of adjusting production, not convex costs as postulated in the production smoothing model. Second, the lumpiness of the margins, accompanied by the fixed costs, leads to a nonconvex cost function. It is important to point out that the nonconvexity is due not to declining marginal costs [as Ramey (1991) originally posited], but rather to the existence of large fixed costs at key points in the cost curve.

Thus, both the (S, s) literature and the limited amount of factory evidence available suggests that fixed costs may be very important. Furthermore, the types of fixed costs highlighted can potentially explain both the procyclicality of inventories and the persistence of the inventory–sales relationship. For example, the lumpiness of the shifts margin in the automobile industry can explain why an increase in sales might lead to a more than proportional increase in production. Also, the importance of fixed adjustment costs can explain why significant deviations in the inventory–sales relationship are allowed to occur before production responds.

It is not yet clear, however, how general are the results from the automobile industry. And more generally it is not well understood how and whether fixed costs at the plant or firm level translate into industry- or economy-wide behavior. Thus, the role of these types of fixed costs in explain aggregate inventory fluctuations remains an important topic for future study.

8.4. The value of more data

Finally, we discuss how the addition of more data may help narrow the estimates obtained from the linear quadratic model, as well as shed light on the unobserved cost shocks. We will argue that there are several available sources of data that have the potential to clear up ambiguities.

One possible explanation for the range of estimates obtained from the production smoothing model is the data are not sufficient for distinguishing the relative values of the parameters. One way to glean more information from macroeconomic data is to use information contained in prices, something done in a handful of papers including Eichenbaum (1984), Blanchard and Melino (1986) and Bils and Kahn (1996). Pindyck's (1994) results using futures prices provides additional evidence of the information contained in prices.

A second possible use of new data is to measure the stochastic variation in cost. As Table 11 indicates, a number of authors have experimented with several observable cost shifters, but generally do not find effects. Another possible source of cost shocks that has been studied in a few papers is credit conditions. We remarked in Section 2.2 that Kashyap et al. (1994) and Carpenter et al. (1994, 1998) still find persistence in the inventory–sales relationship after including measures of credit conditions. But they also regularly find that credit conditions affect inventory holding behavior of small firms, across various specifications. If these credit conditions are serially correlated (which they are likely to be), and if small firms are important enough to substantially affect industry- and economy-wide aggregates, credit conditions may ultimately help explain our two stylized facts.

Finally, we advocate more plant and firm-level studies, although gathering such data requires substantial work. Schuh (1996), for example, uses panel data from the M3LRD to calibrate biases from aggregation. And consider Holt et al.'s (1960) study of six firms ranging from a paint producer to an ice cream maker and Kashyap and Wilcox (1993) and Bresnahan and Ramey's (1994) studies of the automobile industry. They use not only firm-level data on production, inventories and sales, but also company reports and industry press, which provide valuable insights into the cost structure facing firms. For example, Bresnahan and Ramey (1994) were able to categorize the cause of every plant shutdown using information from *Automobile News*, which chronicled drops in demand and cost shocks such as strikes and model year change-overs.

9. Conclusions

We conclude by briefly reiterating several points we have made in this chapter. We began by asserting that inventories are a useful resource in business cycle research. The theoretical dependence of the comovements of sales, production, and inventories on important parameters such as the slope of marginal costs, and on the nature of the underlying shocks, indicates that inventory models can in principle be used to identify these important macroeconomic features. The two stylized facts we highlight – the procyclicality of inventories and the persistence of the inventory–sales relationship – are intimately linked to other aspects of business cycle fluctuations. Thus, inventory movements have valuable business cycle information.

To consider explanations for the two facts, we presented a linear quadratic model. We showed that the model can rationalize the two facts in a number of ways, but focused on two stylized explanations have the virtue of relative simplicity and support from a number of papers. Both assume that there are persistent shocks to demand for the good in question, and that marginal production cost slopes up. The first explanation assumes as well that there are highly persistent shocks to the cost of production. The second assumes that there are strong costs of adjusting production and a strong accelerator motive.

Our review of the empirical evidence, however, indicates that the range of estimates of key parameters and of the relative importance of cost versus demand shocks is too wide to allow us to endorse one of the two or some third explanation. But while the literature has not reached a consensus, it has identified mechanisms and forces that can explain basic characteristics of inventory behavior. We believe that several research strategies, and use of different data, promise to continue to improve our understanding of inventory movements and therefore of business cycle fluctuations.

Appendix A. Data Appendix

Data sources for annual G7 data: all data on inventory changes were obtained from *International Financial Statistics*, mostly from the 1996 CD-ROM. From the CD-ROM, we obtained nominal and real GDP and the nominal change in aggregate inventories. The GDP deflator was used to convert the inventory change from nominal to real. For the Canada, France, the UK and the USA, 1955 data were available to compute ΔQ and ΔS in 1956. For all other countries an observation was lost in computing the initial ΔQ and ΔS. Additional sources were used for West Germany and Italy.

West Germany: (a) 1957–1978: the IFS data used in West (1990a), rebenchmarked to a 1990 from a 1980 base, and output measured with GNP instead of GDP. (b) 1979–1994: in both the CD-ROM and in recent hardcopy versions of IFS, the figures on the annual change looked suspicious: they were uniformly positive and large relative to 1957–1958, bore no obvious connection to the figures on the levels reported in the Statisches Bundesamt publication cited below, and in recent years bore no obvious connection to the average of the reported quarterly figures. So for 1979–1990, we used the annual change reported in the hardcopy IFS, obtaining a given year's data from the April issue three years later (e.g., the 1990 figure came from the April 1993 issue of IFS). (For 1980 we used the May 1983 issue, because the April 1983 issue was missing.) For 1991–1994, we used the average of the quarterly figures from the April 1995 hardcopy version of IFS.

Italy: 1993 and 1994 real GDP came from OECD *Economic Surveys, Italy, 1996*, rebenchmarked to a 1990 from a 1985 base. We checked the US data against the Department of Commerce's 1992 chain-weighted NIPA data, and while there were notable differences, overall the two perhaps were tolerably close: the correlation between inventory investment as constructed here and the Department of Commerce measure was 0.96.

Data sources for non-US data on inventory levels: Canada: private communication from Statistics Canada gave a nominal 1995:IV inventory figure for all nonfinancial industries of 140.8 billion Canadian dollars, which we deflated with the GDP deflator. West Germany: Statisches Bundesamt, *Volkswirtschaftliche Gesamtrechnungen*,

Table 3.2.9. Agriculture (line 2) was subtracted from total (line 1), and the result was deflated by the GDP deflator. Japan: Economic Planning Agency, *Annual Report on National Accounts, 1997*, table on "Closing Stocks". The nominal figure for total stocks was deflated by the GDP deflator. United Kingdom: Office for National Statistics, *United Kingdom National Accounts: The Blue Book, 1996*, Table 15.1. Agriculture (series DHIE) and government (AAAD) were subtracted from total (DHHY), and the result was deflated by the GDP deflator.

Data sources for sectoral distribution of US inventories: broad sectoral categories were obtained from Citibase, and manufacturing inventories by stage of processing were obtained from the BEA. The stage of processing inventories were converted from monthly to quarterly data by sampling the last month of the quarter.

Appendix B. Technical Appendix

This appendix discusses the following: (1) solution of the model[23]; (2) Computation of $E(Q^2-S^2)$ in Table 4; (3) Estimation of $\widehat{\theta}$ in Table 5; (4) the social planning approach to derivation of the first-order conditions.

B.1. Solution of the model

We assume throughout that $a_1, a_2, g>0$ and $a_0, a_3 \geqslant 0$. See Ramey (1991) for solutions when $a_1<0$.

We begin by working through in detail the solution discussed in Section 4, when $a_0=0$ and the forcing variables follow first-order autoregressions. For simplicity, for the most part we set $\widetilde{\alpha} \equiv 0$ as well. Thus $U_{ct}=u_{ct}$ [see Equation (3.2)], $E_{t-1}u_{ct}=\phi_c u_{ct-1}$ and $E_{t-1}U_{dt}=\phi_d U_{dt-1}$. To insure a unique stable solution, we assume that either (B.1a) or (B.1b) holds:

$$g > a_2 a_3(1-a_3), \tag{B.1a}$$

$$a_2 a_3(1-a_3) > g > -\frac{2(1+b^{-1})\,a_1 a_2(a_3-0.5)(a_3-b(1+b)^{-1})}{a_2+2a_1(1+b^{-1})}. \tag{B.1b}$$

Note that the right-hand inequality in (B.1b) follows if a_3 falls outside $(b(1+b)^{-1}, 0.5)$, a narrow range when $b \sim 1$. There will also be a stable solution when $a_2 a_3(1-a_3)=g$. But to allow us to divide by $g-a_2 a_3(1-a_3)$ at certain stages in the derivation, we rule this out for conciseness.

[23] We thank Stanislav Anatolyev for assistance in the preparation of this part of the Technical Appendix.

When $a_0 = 0$, differentiating the objective function (3.1) with respect to S_t gives

$$P_t - E_t[a_1 Q_t - a_2 a_3 (H_{t-1} - a_3 S_t) + U_{ct}] = 0. \tag{B.2}$$

Use $P_t = -gS_t + U_{dt}$, $Q_t = S_t + \Delta H_t$, and our tentative assumption that $U_{ct} = u_{ct}$. (B.2) becomes

$$-a_1 H_t - (a_1 + a_2 a_3^2 + g)S_t + (a_1 + a_2 a_3)H_{t-1} - u_{ct} + U_{dt} = 0. \tag{B.3}$$

$$\Rightarrow$$

$$S_t = -\frac{a_1}{d}H_t + \frac{a_1 + a_2 a_3}{d}H_{t-1} - \frac{1}{d}u_{ct} + \frac{1}{d}U_{dt}, \qquad d \equiv (a_1 + a_2 a_3^2 + g). \tag{B.4}$$

Use (B.4) and (B.4) led one period to substitute out for S_t and S_{t+1} in H_t's first-order condition (3.3) (with $a_0 \equiv 0$). After some rearrangement, the result may be written

$$\begin{aligned}
0 &= bE_t H_{t+1} - (1 + b + m)H_t + H_{t-1} + g_{Hc}u_{ct} + g_{Hd}U_{dt} \\
&\equiv bE_t H_{t+1} - \eta H_t + H_{t-1} + g_{Hc}u_{ct} + g_{Hd}U_{dt}, \\
m &\equiv \frac{a_2[b(a_1 + g) + a_1 a_3(1 - b)]}{a_1[g + a_2 a_3(a_3 - 1)]}, \\
g_{Hc} &\equiv -\frac{g + a_2 a_3^2 - b\phi_c[g + a_2 a_3(a_3 - 1)]}{a_1[g + a_2 a_3(a_3 - 1)]}, \\
g_{Hd} &\equiv -\frac{a_1 - b\phi_d(a_1 + a_2 a_3)}{a_1[g + a_2 a_3(a_3 - 1)]}.
\end{aligned} \tag{B.5}$$

It can be shown that inequality (B.1) guarantees that there is exactly one root less than one to the polynomial

$$bx^2 - \eta x + 1 = 0. \tag{B.6}$$

Call this root π_H, where

$$\pi_H \begin{cases} = 0.5b^{-1}[\eta - (\eta^2 - 4b)^{1/2}] & \text{if} \quad \eta > 0, \\ = 0.5b^{-1}[\eta + (\eta^2 - 4b)^{1/2}] & \text{if} \quad \eta < 0. \end{cases} \tag{B.7}$$

Using techniques from Hansen and Sargent (1980) it follows that the solution to problem (B.5) is

$$\begin{aligned}
&H_t = \pi_H H_{t-1} + f_{Hc}u_{ct} + f_{Hd}U_{dt}, \\
&f_{Hc} \equiv [\pi_H/(1 - b\pi_H\phi_c)]g_{Hc}, \qquad f_{Hd} \equiv [\pi_H/(1 - b\pi_H\phi_d)]g_{Hd}.
\end{aligned} \tag{B.8}$$

Upon substituting Equation (B.8) into Equation (B.4) and rearranging, we obtain

$$\begin{aligned}
&S_t = \pi_S H_{t-1} + f_{Sc}u_{ct} + f_{Sd}U_{dt}, \qquad \pi_S \equiv \frac{a_1(1 - \pi_H) + a_2 a_3}{a_1 + a_2 a_3^2 + g}, \\
&f_{Sc} \equiv -\frac{1 + a_1 f_{Hc}}{a_1 + a_2 a_3^2 + g}, \qquad f_{Sd} \equiv \frac{1 - a_1 f_{Hd}}{a_1 + a_2 a_3^2 + g}.
\end{aligned} \tag{B.9}$$

Let L be the lag operator, "adj" the adjoint of a matrix. From Equations (B.7) and (B.8), a representation for the bivariate $(H_t, H_t - \theta S_t)' \equiv Y_t$ process is

$$\begin{aligned}
&Y_t = AY_{t-1} + BU_t,\\
&A \equiv \begin{vmatrix} \pi_H & 0 \\ \pi_H - \theta\pi_S & 0 \end{vmatrix}, \qquad B_1 \equiv \begin{vmatrix} f_{H\mathrm{c}} & f_{H\mathrm{d}} \\ f_{H\mathrm{c}} - \theta f_{S\mathrm{c}} & f_{H\mathrm{d}} - \theta f_{S\mathrm{d}} \end{vmatrix}\\
\Rightarrow \quad &Y_t = (I - AL)^{-1}BU_t = \frac{\mathrm{adj}(I - AL)}{|I - AL|}BU_t\\
\Rightarrow \quad &|I - AL|\, Y_t = \mathrm{adj}(I - AL)\, BU_t.
\end{aligned} \tag{B.10}$$

This may be used to solve for the univariate process for $H_t - \theta S_t$, which is

$$\begin{aligned}
(1 - \pi_H L)(H_t - \theta S_t) = {} & (f_{H\mathrm{c}} - \theta f_{S\mathrm{c}})\, u_{\mathrm{c}t} + \theta(\pi_H f_{S\mathrm{c}} - \pi_S f_{H\mathrm{c}}) u_{\mathrm{c}t-1}\\
& + (f_{H\mathrm{d}} - \theta f_{S\mathrm{d}})\, U_{\mathrm{d}t} + \theta(\pi_H f_{S\mathrm{d}} - \pi_S f_{H\mathrm{d}}) U_{\mathrm{d}t-1}.
\end{aligned} \tag{B.11}$$

Suppose that $U_{\mathrm{d}t}$ follows a random walk, so that $\phi_{\mathrm{d}} = 1$ and $(f_{H\mathrm{d}} - \theta f_{S\mathrm{d}})\, U_{\mathrm{d}t} = (f_{H\mathrm{d}} - \theta f_{S\mathrm{d}})(U_{\mathrm{d}t-1} + e_{\mathrm{d}t})$. Upon using the definition of θ in Equation (3.5), and in light of the quadratic equation used to obtain π_H, tedious manipulations reveal that $(f_{H\mathrm{d}} - \theta f_{S\mathrm{d}}) + \theta(\pi_H f_{S\mathrm{d}} - \pi_S f_{H\mathrm{d}}) = 0$. It follows that

$$\begin{aligned}
&\phi_{\mathrm{d}} = 1 \quad \Rightarrow\\
&(1 - \pi_H L)(H_t - \theta S_t) = (f_{H\mathrm{c}} - \theta f_{S\mathrm{c}})\, u_{\mathrm{c}t} + \theta(\pi_H f_{S\mathrm{c}} - \pi_S f_{H\mathrm{c}})\, u_{\mathrm{c}t-1} + (f_{H\mathrm{d}} - \theta f_{S\mathrm{d}})\, e_{\mathrm{d}t}\\
&\qquad\qquad \equiv m_{0\mathrm{c}} u_{\mathrm{c}t} + m_{1\mathrm{c}} u_{\mathrm{c}t-1} + m_{0\mathrm{d}} e_{\mathrm{d}t}\\
&\Rightarrow\\
&(1 - \pi_H L)(1 - \phi_{\mathrm{c}} L)(H_t - \theta S_t) = v_t \equiv m_{0\mathrm{c}} e_{\mathrm{c}t} + m_{1\mathrm{c}} e_{\mathrm{c}t-1} + m_{0\mathrm{d}} e_{\mathrm{d}t} - \phi_{\mathrm{c}} m_{0\mathrm{d}} e_{\mathrm{d}t-1}\\
&\quad \sim \mathrm{MA}(1).
\end{aligned} \tag{B.12}$$

Thus, when $\phi_{\mathrm{d}} = 1$, $H_t - \theta S_t \sim \mathrm{ARMA}(2, 1)$ with autoregressive roots π_H and ϕ_{c}.

Now suppose that $\tilde{\alpha} \neq 0$, so that $U_{\mathrm{c}t} = \tilde{\alpha}' W_t + u_{\mathrm{c}t}$, with $E_{t-1} W_t = \Phi_{\mathrm{w}} W_{t-1}$. Algebra similar to that used above may be used to conclude that the first-order condition (B.5) and the decision rules (B.8) and (B.9) become

$$\begin{aligned}
&0 = E_t(bH_{t+1} - \eta H_t + H_{t-1} + g'_{H\mathrm{w}} W_t + g_{H\mathrm{c}} u_{\mathrm{c}t} + g_{H\mathrm{d}} U_{\mathrm{d}t}),\\
&H_t = \pi_H H_{t-1} + f'_{H\mathrm{w}} W_t + f_{H\mathrm{c}} u_{\mathrm{c}t} + f_{H\mathrm{d}} U_{\mathrm{d}t},\\
&S_t = \pi_S H_{t-1} + f'_{S\mathrm{w}} W_t + f_{S\mathrm{c}} u_{\mathrm{c}t} + f_{S\mathrm{d}} U_{\mathrm{d}t},\\
&g_{H\mathrm{w}} \equiv -\frac{(g + a_2 a_3^2)\boldsymbol{I} - b[g + a_2 a_3(a_3 - 1)]\Phi'_{\mathrm{w}}}{a_1[g + a_2 a_3(a_3 - 1)]}\tilde{\alpha},\\
&f_{H\mathrm{w}} \equiv \pi_H(\boldsymbol{I} - b\pi_H \Phi'_{\mathrm{w}})^{-1} g_{H\mathrm{w}}, \qquad f_{S\mathrm{w}} \equiv -\frac{\tilde{\alpha} + a_1 f_{H\mathrm{w}}}{a_1 + a_2 a_3^2 + g},
\end{aligned} \tag{B.13}$$

with the other parameters unchanged. It follows that when $\phi_{\mathrm{d}}=1$

$$\begin{aligned}
&(1-\pi_H L)(H_t-\theta S_t)\\
&\quad=(f_{H\mathrm{w}}-\theta f_{S\mathrm{w}})'W_t+\theta(\pi_H f_{S\mathrm{w}}-\pi_S f_{H\mathrm{w}})'W_{t-1}\\
&\qquad+(f_{H\mathrm{c}}-\theta f_{S\mathrm{c}})u_{\mathrm{c}t}+\theta(\pi_H f_{S\mathrm{c}}-\pi_S f_{H\mathrm{c}})u_{\mathrm{c}t-1}+(f_{H\mathrm{d}}-\theta f_{S\mathrm{d}})e_{\mathrm{d}t}.\\
&\Rightarrow\\
&(1-\pi_H L)(H_t-\theta S_t-\alpha' W_t)\\
&\quad=(f_{H\mathrm{w}}-\theta f_{S\mathrm{w}}-\alpha)'W_t+[\theta(\pi_H f_{S\mathrm{w}}-\pi_{\mathrm{s}} f_{H\mathrm{w}})+\pi_H\alpha]'W_{t-1}\\
&\qquad+(f_{H\mathrm{c}}-\theta f_{S\mathrm{c}})\,u_{\mathrm{c}t}+\theta(\pi_H f_{S\mathrm{c}}-\pi_S f_{H\mathrm{c}})\,u_{\mathrm{c}t-1}+(f_{H\mathrm{d}}-\theta f_{S\mathrm{d}})\,e_{\mathrm{d}t}\\
&\quad\equiv m_{0\mathrm{W}}'W_t+[\theta(\pi_H f_{S\mathrm{w}}-\pi_S f_{H\mathrm{w}})+\pi_H\alpha]'W_{t-1}+m_{0\mathrm{c}}u_{\mathrm{c}t}+m_{1\mathrm{c}}u_{\mathrm{c}t-1}+m_{0\mathrm{d}}e_{\mathrm{d}t}.
\end{aligned}\tag{B.14}$$

When $\Phi_{\mathrm{w}}=\boldsymbol{I}$ and $W_t=W_{t-1}+e_{\mathrm{w}t}$, it may be shown that $m_{0\mathrm{w}}+\theta(\pi_H f_{S\mathrm{w}}-\pi_S f_{H\mathrm{w}})+\pi_H\alpha=0$. Then (B.14) implies

$$(1-\pi_H L)(H_t-\theta S_t-\alpha' W_t)=m_{0\mathrm{w}}'e_{\mathrm{w}t}+m_{0\mathrm{c}}u_{\mathrm{c}t}+m_{1\mathrm{c}}u_{\mathrm{c}t-1}+m_{0\mathrm{d}}e_{\mathrm{d}t}.\tag{B.15}$$

Now allow $a_0\neq 0$, as well as arbitrary autoregressive processes for W_t, $u_{\mathrm{c}t}$, and $U_{\mathrm{d}t}$. When $a_0\neq 0$, the first-order condition for S_t (B.2) becomes

$$-gS_t+U_{\mathrm{d}t}-E_t[(a_0\Delta Q_t-ba_0\Delta Q_{t+1})+a_1Q_t-a_2a_3(H_{t-1}-a_3S_t)+U_{\mathrm{c}t}]=0,\tag{B.16}$$

where $P_t=-gS_t+U_{\mathrm{d}t}$ has been used to substitute out for P_t. The solution is most concisely derived if one uses $S_t=Q_t-\Delta H_t$ to remove the S_t from the first-order condition (equivalently, if one makes Q_t and H_t rather than S_t and H_t the choice variables). After so doing, Equations (B.16) and (3.3) may be written

$$E_t[bA_1'X_{t+1}+A_0X_t+A_1X_{t-1}+B_0U_t+B_1U_{t+1}]=0, X_t\equiv(H_t,\ Q_t)'.\tag{B.17}$$

The (2×1) vector U_t is $(U_{\mathrm{d}t},U_{\mathrm{c}t})'$. The matrices A_1, A_0 and B_0 and B_1 are (2×2), with

$$\begin{aligned}
A_0&\equiv\begin{vmatrix}(1+b)g+a_2a_3^2+ba_2(1-a_3)^2 & -(g+a_2a_3^2)\\ -(g+a_2a_3^2) & g+a_1+a_2a_3^2+(1+b)a_0\end{vmatrix},\\
A_1&\equiv\begin{vmatrix}-[g-a_2a_3(1-a_3)] & 0\\ [g-a_2a_3(1-a_3)] & -a_0\end{vmatrix},\qquad B_0\equiv\begin{vmatrix}0 & 1\\ 1 & -1\end{vmatrix},\qquad B_1\equiv\begin{vmatrix}0 & -b\\ 0 & 0\end{vmatrix}.
\end{aligned}\tag{B.18}$$

Suppose that $V_t\equiv(u_{\mathrm{c}t},U_{\mathrm{d}t},W_t')'\sim\mathrm{AR}(p)$ (possibly with unit autoregressive roots), $E_{t-1}V_t=\Phi_1\mathrm{V}_{t-1}+\cdots+\Phi_pV_{t-p}$. (There may be many zeros in the Φ_i if there is lots more dynamics in say W_t than in either $u_{\mathrm{c}t}$ or $U_{\mathrm{d}t}$). Note for the future that

$$U_t=\begin{vmatrix}1 & 0 & \tilde{\alpha}'\\ 0 & 1 & 0\end{vmatrix}V_t\equiv CV_t.\tag{B.19}$$

Guess a solution of the form

$$X_t = RX_{t-1} + G_0V_t + \cdots + G_{p-1}V_{t-p+1} \equiv RX_{t-1} + G(L)V_t. \tag{B.20}$$

In Equation (B.20), the G_i are $2\times(2+\text{dimension of } W_t)$. In Equation (A.17), use (B.20) led once to substitute out for E_tX_{t+1}, and then substitute out for X_t using (B.20). For condition (B.17) to hold, we must have

$$(bA_1'R^2 + A_0R + A_1)X_{t-1} = 0, \tag{B.21a}$$

$$bA_1'\{RG(L)V_t + [E_tG(L)V_{t+1}]\} + A_0G(L)V_t + B_0CV_t + B_1CE_tV_{t+1} = 0. \tag{B.21b}$$

Equation (B.21a) requires $bA_1'R^2+A_0R+A_1=0$. Given model parameters, and thus knowledge of A_1 and A_0, this equation may be used to solve for a stable matrix R. (The matrix equation will have multiple solutions, just as does the scalar quadratic Equation (B.5). Restrictions similar to those in (B.1) will insure that there is a unique stable solution.) It may help to note that the solution matches the earlier one if we reimpose the assumption that there are no costs of adjusting production. With $a_0=0$, the second column of A_1 is zero, from which it follows that the second column of R is zero. Further, $R(1,1)\equiv r_{11}=\pi_H$, $r_{21}=\pi_H+\pi_S-1$ for π_H and π_S defined in Equations (B.7) and (B.9).

Whether or not $a_0=0$, once R is recovered from Equation (B.21a), Equation (B.21b) may then be used to solve for the G_i. For example, suppose that W_t and the shocks are first-order autoregressive, so that $p=1$: Φ_1 is block diagonal with ϕ_c in its $(1,1)$ element, ϕ_d in its $(2,2)$ element and Φ_w in the block in its lower right hand corner. Then Equation (B.21b) implies

$$bA_1'(RG_0 + G_0\Phi_1) + A_0G_0 + B_0C + B_1C\Phi_1 = 0, \tag{B.22}$$

which can be used to solve linearly for G_0 in terms of the model parameters and Φ_1.

Return to the solution (B.20). Let $R=[r_{ij}]$. Transform from a solution in $(H_t,Q_t)'$ to one in $(H_t,S_t)'\equiv Z_t$, using $Q_t=S_t+\Delta H_t$. The result is

$$\begin{gathered}Z_t = \Pi_1Z_{t-1} + \Pi_2Z_{t-2} + F_0V_t + \cdots + F_{p-1}V_{t-p+1},\\ \Pi_1 \equiv \begin{vmatrix} r_{11}+r_{12} & r_{12} \\ 1+r_{22}+r_{21}-(r_{11}+r_{12}) & r_{22}-r_{12}\end{vmatrix}, \qquad \Pi_2 \equiv \begin{vmatrix} -r_{12} & 0 \\ -(r_{22}-r_{12}) & 0\end{vmatrix},\\ F_i = MG_i, \qquad M = \begin{vmatrix} 1 & 0 \\ -1 & 1\end{vmatrix}.\end{gathered} \tag{B.23}$$

Thus S_{t-2} does not appear in the reduced form, and the first column of Π_2 is the negative of the second column of Π_1. To repeat an earlier point: if $a_0=0$, then $r_{12}=r_{22}=0$, $r_{11}=\pi_H$, and $1+r_{21}-r_{11}=\pi_S$.

Finally, suppose we use Equation (B.23) to derive the autoregressive process for $(H_t, H_t-\theta S_t)'$, and then solve for the univariate process for $H_t-\theta S_t$, using the method that led to Equation (B.11). When $a_0 \neq 0$, the relevant determinant [the analogue of $|I - AL|$ in Equation (B.10)] is a second-order polynomial in the lag operator. Specifically, for Π_1 defined in Equation (B.23), let $\Pi_1 \equiv [\pi_{ij}]$. Then this second order lag polynomial is

$$1 - (\pi_{11} + \pi_{12})L + (\pi_{11}\pi_{22} + \pi_{12} - \pi_{21}\pi_{12})L^2. \tag{B.24}$$

The roots to this polynomial are called π_1 and π_2. Equation (B.22) may be used to show that the moving average component of $H_t - \theta S_t$ is first order when $\phi_c = 0$ and $\phi_d = 1$.

B.2. Computation of $E(Q^2 - S^2)$

The computation of $E(Q^2 - S^2)$ follows West (1988). In the stationary case, and ignoring deterministic terms, $E(Q_t^2 - S_t^2)\,\mathrm{var}(\Delta H_t) = 2ES_t\Delta H_t = 2\,\mathrm{cov}(S_t, \Delta H_t) = 2\,\mathrm{cov}[(\Delta S_t + \Delta S_{t-1} + \Delta S_{t-2} + \cdots), \Delta H_t] = 2[\mathrm{cov}(\Delta S_t, \Delta H_t) + \mathrm{cov}(\Delta S_{t-1}, \Delta H_t) + \cdots]$. It may be shown that even in the case $S \sim I(1)$: (1) under regularity conditions, $\sum_{j=0}^{\infty} \mathrm{cov}(\Delta S_{t-j}, \Delta H_t)$ is finite and may be consistently estimated; (2) the implications of the model for the sign of $E(Q^2 - S^2)$ are the same as those for $\mathrm{var}(Q) - \mathrm{var}(S)$ when S_t (and H_t) are $I(0)$ [see the working paper version of West (1990b)]. Let $\widehat{\gamma}_j$ be an estimate of $\mathrm{cov}(\Delta S_{t-j}, \Delta H_t)$, which we computed as $\widehat{\gamma}_j = T^{-1} \sum_{t=j+1}^{T} \Delta s_{t-j}\Delta h_t$, where Δs and Δh are residuals produced by demeaning or detrending. When $S_t \sim I(1)$, consistent nonparametric estimation of $\sum_{j=0}^{\infty} \mathrm{cov}(\Delta S_{t-j}, \Delta H_t)$ involves computing $\sum_{j=1}^{m} \widehat{\gamma}_j$ and letting $m \to \infty$ as $T \to \infty$. In the computations in Table 4, we let the number of lags m be 5 for the full samples, 4 for post-1973 samples.

B.3. Estimation of $\widehat{\theta}$

The cointegrating parameters were estimated using dynamic OLS, as advocated by Stock and Watson (1993). We regressed inventories on a constant and sales (with no deterministic trend), and included current plus up to two years (or eight quarters) of leads and lags of the change in sales. We successively eliminated variables with t-statistics less than 1.65 in absolute value. Heteroskedasticity and autocorrelation consistent standard errors were computed using a Bartlett kernel with two years (or eight quarters) of lags.

B.4. Social planning derivation of the model's first-order conditions

The area under the inverse demand curve is just the integral of the inverse demand (4.1), and is given by

$$\int_0^{\infty} (U_{\mathrm{d}t} - gX_t)\,\mathrm{d}X_t = U_{\mathrm{d}t}S_t - 0.5gS_t^2. \tag{B.25}$$

The area under the supply curve is equal to the cost function presented as part of Equation (3.1). Thus, the competitive equilibrium solution is equivalent to the solution to the following social planner problem that maximizes the difference between these two functions:

$$\mathrm{Max}_{H,S}\ V = E_t \sum_{j=0}^{\infty} b^j \left[U_{\mathrm{d}t+j} S_{t+j} - 0.5 g S_{t+j}^2 - 0.5 a_0 \Delta Q_{j+j}^2 - 0.5 a_1 Q_{t+j}^2 - 0.5 a_2 (H_{t+j-1} - a_3 S_{t+j})^2 - U_{\mathrm{c}t+j} Q_{t+j}\right] \tag{B.26}$$

subject to the inventory identity that $Q_{t+j} = S_{t+j} + \Delta H_{t+j}$. The first-order condition for inventories is identical to the one obtained from the firm-level problem. The first-order condition for sales is identical to the one obtained when the industry demand curve is substituted into the firm-level first-order condition for sales.

References

Abramowitz, M. (1950), Inventories and Business Cycles (National Bureau of Economic Research, New York).

Allen, D.S. (1995), "Changes in inventory management and the business cycle", Federal Reserve Bank of St. Louis Review 77:17–26.

Beaulieu, J.J., and J.A. Miron (1992), "A cross country comparison of seasonal cycles and business cycles", The Economic Journal 102:772–788.

Belsley, D. (1969), Industry Production Behavior: The Order–Stock Distinction (North-Holland, Amsterdam).

Bils, M., and J.A. Kahn (1996), "What inventory behavior tells us about business cycles", manuscript (University of Rochester).

Blanchard, O.J. (1983), "The production and inventory behavior of the American automobile industry", Journal of Political Economy 91:365–400.

Blanchard, O.J., and A. Melino (1986), "The cyclical behavior of prices and quantities: the case of the automobile market", Journal of Monetary Economics 17:379–408.

Blinder, A.S. (1981), "Retail inventory behavior and business fluctuations", Brookings Papers on Economic Activity 1981(2):443–505.

Blinder, A.S. (1986a), "Can the production smoothing model of inventory behavior be saved?", Quarterly Journal of Economics 101:431–453.

Blinder, A.S. (1986b), "More on the speed of adjustment in inventory models", Journal of Money, Credit and Banking 18:355–365.

Blinder, A.S., and L.J. Maccini (1991), "Taking stock: a critical assessment of recent research on inventories", Journal of Economic Perspectives 5:73–96.

Bresnahan, T.F., and V.A. Ramey (1994), "Output fluctuations at the plant level", Quarterly Journal of Economics 109:593–624.

Caplin, A.S. (1985), "The variability of aggregate demand with (S, s) inventory policies", Econometrica 53:1395–1409.

Carlson, J.A., and W. Wehrs (1974), "Aggregate inventory behavior", in: G. Horwich and P. Samuelson, eds., Trade, Stability and Growth: Essays in Honor of Lloyd A. Metzler (Academic Press, New York).

Carpenter, R.E., and D. Levy (1998), "Seasonal cycles, business cycles and the comovement of inventory", Journal of Money, Credit and Banking 30:331–346.

Carpenter, R.E., S.M. Fazzari and B.C. Petersen (1994), "Inventory investment, internal-finance fluctuations, and the business cycle", Brookings Papers on Economic Activity 1994(2):75–138.

Carpenter, R.E., S.M. Fazzari and B.C. Petersen (1998), "Financing constraints and inventory investment: a comparative study using high frequency data", Review of Economics and Statistics 80:513–519.

Cecchetti, S.G., A.K. Kashyap and D.W. Wilcox (1997), "Interactions between the seasonal and business cycles in production and inventories", American Economic Review 87:884–892.

Childs, G.D. (1967), Inventories and Unfilled Orders (North-Holland, Amsterdam).

Christiano, L.J. (1988), "Why does inventory investment fluctuate so much?", Journal of Monetary Economics 21:247–280.

Christiano, L.J., and M.S. Eichenbaum (1989), "Temporal aggregation and the stock adjustment model of inventories", in: T. Kollintzas, ed., The Rational Expectations Inventory Model (Springer-Verlag, New York) 70–109.

Durlauf, S.N., and L.J. Maccini (1995), "Measuring noise in inventory models", Journal of Monetary Economics 36:65–90.

Eichenbaum, M. (1984), "Rational expectations and the smoothing properties of finished goods", Journal of Monetary Economics 14:71–96.

Eichenbaum, M. (1989), "Some empirical evidence on the production level and production cost smoothing models of inventory investment", American Economic Review 79:853–864.

Fair, R.C. (1989), "The production smoothing model is alive and well", Journal of Monetary Economics 24:353–370.

Feldstein, M.S., and A.J. Auerbach (1976), "Inventory behavior in durable goods manufacturing: the target adjustment model", Brookings Papers on Economic Activity 1976(2):351–396.

Fisher, J.D.M., and A. Hornstein (1996), "(S, s) inventory policies in general equilibrium", working paper WP-96-24 (Federal Reserve Bank of Chicago).

Fuhrer, J.C., G.R. Moore and S. Schuh (1995), "Estimating the linear-quadratic inventory model: maximum likelihood versus generalized method of moments", Journal of Monetary Economics 35: 115–157.

Fukuda, S.-i., and H. Teruyama (1988), "Some international evidence on inventory fluctuations", Economics Letters 28:225–230.

Galeotti, M., L. Guiso, B. Sack and F. Schiantarelli (1997), "Inventories, production smoothing and the shape of the cost function", manuscript (Università degli Studi di Bergamo).

Granger, C.W.J., and T.H. Lee (1989), "Investigation of production, sales and inventory relationships using multicointegration and non-symmetric error correction models", Journal of Applied Econometrics 4:S145–S159.

Hall, G. (1996), "Non-convex costs and capital utilization: a study of production and inventories at automobile assembly plants", manuscript (Federal Reserve Bank of Chicago).

Haltiwanger, J.C., and L.J. Maccini (1989), "Inventories, orders, temporary and permanent layoffs: an econometric analysis", Carnegie-Rochester Series on Public Policy 30:301–366.

Hansen, L.P., and T.J. Sargent (1980), "Formulating and estimating dynamic linear rational expectations models", Journal of Economic Dynamics and Control 2:7–46.

Hansen, L.P., and K.J. Singleton (1982), "Generalized instrumental variables estimation of nonlinear rational expectations models", Econometrica 50:1269–1286.

Hester, D.A. (1994), "Changing relations between inventories and bank loans", in: R. Fiorito, ed., Inventory Cycles and Monetary Policy (Springer-Verlag, Berlin) 125–147.

Holt, C.C., F. Modigliani, J.F. Muth and H.A. Simon (1960), Planning Production, Inventories and Work Force, (Prentice Hall, Englewood Cliffs, NJ).

Humphreys, B.R., L.J. Maccini and S. Schuh (1997), "Input and output inventories", manuscript (Federal Reserve Bank of Boston).

Jorda, O. (1997), "Random-time aggregation in partial adjustment models", manuscript (University of California at Davis).

Kahn, J.A. (1987), "Inventories and the volatility of production", American Economic Review 77: 667–680.

Kahn, J.A. (1992), "Why is production more volatile than sales? Theory and evidence on the stockout-avoidance motive for inventory holding", Quarterly Journal of Economics CVII:481–510.

Kashyap, A.K., and D.W. Wilcox (1993), "Production and inventory control at the General Motors Corporation during the 1920s and 1930s", American Economic Review 83:383–401.

Kashyap, A.K., O.A. Lamont and J.C. Stein (1994), "Credit conditions and the cyclical behavior of inventories", Quarterly Journal of Economics CIX:565–592.

Kollintzas, T. (1989), "The linear rational expectations equilibrium inventory model: an introduction", in: T. Kollintzas, ed., The Rational Expectations Equilibrium Inventory Model (Springer-Verlag, New York) 1–32.

Kollintzas, T. (1995), "A generalized variance bounds test with an application to the Holt et al. inventory model", Journal of Economic Dynamics and Control 19:59–89.

Krane, S.D. (1994), "The distinction between inventory holding and stockout costs: implications for target inventories, asymmetric adjustment, and the effect of aggregation on production smoothing", International Economic Review 35:117–136.

Krane, S.D., and S.N. Braun (1991), "Production smoothing evidence from physical product data", Journal of Political Economy 99:558–581.

Lai, K.S. (1991), "Aggregation and testing of the production smoothing hypothesis", International Economic Review 32:391–403.

Lovell, M.C. (1961), "Manufacturers' inventories, sales expectations, and the acceleration principle", Econometrica 29:293–314.

Lovell, M.C. (1993), "Simulating the inventory cycle", Journal of Economic Behavior and Organization 21:147–179.

Lovell, M.C. (1994), "Researching inventories: why haven't we learned more?", International Journal of Production Economics 35:33–41.

Lovell, M.C. (1996), "Macroeconomic implications of *S, s* versus accelerator finished goods inventory strategies", International Journal of Production Economics 45:55–64.

Maccini, L.J., and R.J. Rossana (1981), "Investment in finished goods inventories: an analysis of adjustment speeds", American Economic Review 71:17–22.

Maccini, L.J., and R.J. Rossana (1984), "Joint production, quasi-fixed factors of production, and investment in finished goods inventories", Journal of Money, Credit and Banking 16:218–236.

McManus, D.A., J.C. Nankervis and N.E. Savin (1994), "Multiple optima and asymptotic approximations in the partial adjustment model", Journal of Econometrics 62:91–128.

Metzler, L.A. (1941), "The nature and stability of inventory cycles", The Review of Economics and Statistics XXIII:113–129.

Miron, J.A., and S.P. Zeldes (1988), "Seasonality, cost shocks and the production smoothing model of inventories", Econometrica 56:877–908.

Mosser, P.C. (1988), "Empirical tests of the (*S, s*) model for merchant wholesalers", in: A. Chikan and M. Lovell, eds., The Economics of Inventory Management (Elsevier, Amsterdam) 261–284.

Pindyck, R.S. (1994), "Inventories and the short-run dynamics of commodity prices", The RAND Journal of Economics 25:141–159.

Ramey, V.A. (1989), "Inventories as factors of production and economic fluctuations", American Economic Review 79:338–354.

Ramey, V.A. (1991), "Nonconvex costs and the behavior of inventories", Journal of Political Economy 99:306–334.

Rossana, R.J. (1990), "Interrelated demands for buffer stocks and productive inputs: estimation for two digit manufacturing industries", Review of Economics and Statistics 72:19–29.

Rossana, R.J. (1993), "The long-run implications of the production smoothing model of inventories: an empirical test", Journal of Applied Econometrics 8:295–306.

Rossana, R.J. (1995), "Technology shocks and cointegration in quadratic models of the firm", International Economic Review 36:5–17.

Rossana, R.J. (1998), "Structural instability and the production smoothing model of inventories", Journal of Business and Economic Statistics 16:206–215.

Rotemberg, J.J., and G. Saloner (1989), "The cyclical behavior of strategic inventories", Quarterly Journal of Economics 104:73–98.

Sargent, T.J. (1979), Macroeconomic Theory (Academic Press, New York).

Schuh, S. (1996), "Evidence on the link between firm-level and aggregate inventory behavior", Finance and Economics Discussion Paper No. 1996-46 (Federal Reserve Board of Governors).

Stengel, R.F. (1986), Stochastic Optimal Control: Theory and Applications (Wiley, New York).

Stock, J.H., and M.W. Watson (1993), "A simple estimator of cointegrating vectors in higher order integrated systems", Econometrica 61:783–820.

Surekha, K., and M. Ghali (1997), "The speed of adjustment and production smoothing", manuscript (Western Washington University).

West, K.D. (1983), "Inventory models and backlog costs", Ph.D. Thesis (Massachusetts Institute of Technology).

West, K.D. (1986), "A variance bounds test of the linear quadratic inventory model", Journal of Political Economy 94:374–401.

West, K.D. (1988), "Order backlogs and production smoothing", in: A. Chikan and M. Lovell, eds., The Economics of Inventory Management (Elsevier, Amsterdam)305–318.

West, K.D. (1990a), "Evidence from seven countries on whether inventories smooth aggregate output", Engineering Costs and Production Economics 19:85–90.

West, K.D. (1990b), "The sources of fluctuations in aggregate inventories and GNP", Quarterly Journal of Economics CV:939–972.

West, K.D. (1992a), "A comparison of the behavior of US and Japanese inventories", International Journal of Production Economics 26:115–222.

West, K.D. (1992b), "Sources of cycles in Japan, 1975–1987", Journal of the Japanese and International Economies 6:71–98.

West, K.D. (1995), "Inventory models", in: M. Pesaran and M. Wickens, eds., Handbook of Applied Econometrics, vol. I, Macroeconometrics (Basil Blackwell, Oxford) 188–220.

West, K.D., and D.W. Wilcox (1994), "Some evidence on the finite sample behavior of an instrumental variables estimator of the linear quadratic inventory model", in: R. Fiorito, ed., Inventory Cycles and Monetary Policy, (Springer-Verlag, Berlin) 253–82.

West, K.D., and D.W. Wilcox (1996), "A comparison of alternative instrumental variables estimators of a dynamic linear model", Journal of Business and Economic Statistics 14:281–293.

Wilkinson, M. (1989), "Aggregate inventory behavior in large European economies", European Economic Review 33:181–194.