

OLIVIER BLANCHARD

Massachusetts Institute of Technology

JOHN SIMON

Reserve Bank of Australia

The Long and Large Decline in U.S. Output Volatility

SINCE THE EARLY 1980s the U.S. economy has gone through two long expansions. The first, from 1982 to 1990, lasted thirty-one quarters. The second started in 1991 and, although showing signs of faltering, has recorded its fortieth quarter as this volume goes to press and is already the longest U.S. expansion on record.

One view is that these two long expansions are simply the result of luck, of an absence of major adverse shocks over the last twenty years. We argue that more has been at work, namely, a large underlying decline in output volatility. Furthermore, we contend, this decline is not a recent development—the by-product of a “New Economy” or of Alan Greenspan’s talent. Rather it has been a steady decline over several decades, which started in the 1950s (or earlier, but lack of consistent data makes this difficult to establish), was interrupted in the 1970s and early 1980s, and returned to trend in the late 1980s and the 1990s.¹ The magnitude of the decline is substantial: the standard deviation of quarterly output growth has declined by a factor of three over the period. This is more than enough to account for the increased length of expansions.

We thank Benjamin Friedman for his comments, as well as Robert Solow and participants in the MIT macro lunch. This paper builds on John Simon’s doctoral thesis (Simon, 2000). The views expressed are our own and should not be attributed to the Reserve Bank of Australia.

1. What has happened to output volatility over a much longer time span is the subject of a well-known debate, which this paper does not revisit. See Romer (1986), Weir (1986), and Balke and Gordon (1989).

Having established this fact, we reach two other conclusions. First, the decrease in volatility can be traced to a number of proximate causes, from a decrease in the volatility of government spending early on, to a decrease in consumption and investment volatility throughout the period, to a change in the sign of the correlation between inventory investment and sales in the last decade. Second, there is a strong relationship between movements in output volatility and movements in inflation volatility. The interruption of the trend decline in output volatility in the 1970s was associated with a large increase in inflation volatility; the return to trend is associated with the decrease in inflation volatility since then.

This paper is organized as follows. We start by documenting our basic fact, namely, the secular decrease in output volatility. We then look at the stochastic process for GDP and show that this decrease in volatility can be traced primarily to a decrease in the standard deviation of output shocks, rather than to a change in the dynamics of output. Finally, we show how this decrease in the standard deviation of shocks accounts for the increased length of expansions.

We then take up the question of whether recessions are special, in a way that the formalization used earlier does not do justice to. Put another way, we ask whether what we have seen over the last twenty years is simply the absence of large shocks and nothing more. We show that this is not the case. The measured decrease in output volatility has little to do with the absence of large shocks in the recent past.

We then turn to the relationship between output volatility and inflation. We show that there is a strong relationship both between output volatility and the level of inflation, and between output volatility and inflation volatility. Both volatilities went up in the 1970s and have come down since. Correlation does not, however, imply causality. The correlation in both periods may have been due to third factors, such as supply shocks in the 1970s. This leads us to consider evidence from the other members of the Group of Seven (G-7) large industrial countries. Our motivation here is that the different timings of disinflation across these countries can help us separate out the effects of inflation from those of supply shocks. We first show that these other countries have also experienced a decline in output volatility, although with some differences in timing and in magnitude. (An interesting exception is Japan, where a decline in output volatility has been reversed since the late 1980s.) We then show that, even after controlling

for common time fixed effects, inflation volatility still appears to be strongly related to output volatility.


As a matter of accounting, the decline in output volatility can be traced either to changes in the composition of output or to changes in the variances and covariances of its underlying components. We therefore look at the components of GDP and find that, at least at the level of disaggregation at which we operate, changes in composition explain essentially none of the trend decline. The composition of GDP has changed, but the effects of the various changes have mostly cancelled each other out. We also find that, apart from a decrease in the volatility of government spending early in the postwar period, most of the decrease in output volatility can be traced to a decrease in both consumption volatility and investment volatility, and more recently to a change in inventory behavior, with inventory investment becoming more countercyclical.

We conclude by discussing the agenda for further research. In particular, it is clear that we have addressed only the proximate causes of the volatility decline. The deeper causes, from changes in financial markets to better countercyclical policy, remain to be identified.

The Decline in Output Volatility

Figure 1 shows the rolling standard deviation of quarterly real output growth (measured at a quarterly rate) since the first quarter of 1952. The measure of output is chain-weighted GDP. We use a window of twenty quarters, so that the standard deviation reported for quarter t is the estimated standard deviation over quarters $t - 19$ to t . The first available observation for chain-weighted GDP is 1947:1, and so the first observation for the standard deviation of the growth rate is 1952:1. The figure shows a clear decline in the standard deviation over time, from about 1.5 percent a quarter in the early 1950s to less than 0.5 percent in the late 1990s. This decline is not continuous, however. Volatility increases from the late 1960s to the mid-1980s, and this is followed by a sharp decline in the second half of the 1980s.

One can think of other ways of measuring volatility. One alternative is to look at the standard deviation of an output gap, for example, the difference between the level of (the logarithm of) output and a Hodrick-

Figure 1. Volatility of Output Growth, 1952–2000^a

Source: Authors' calculations based on data from Bureau of Economic Analysis, National Income and Product Accounts.
a. Twenty-quarter rolling standard deviation of quarterly real GDP growth.

Prescott-filtered series.² Another is to look at annual rather than quarterly changes in GDP. These alternatives yield very similar conclusions. The basic reason is that the standard deviation of quarterly output growth primarily reflects the high-frequency properties of the series, which are largely invariant to the detrending method.

Changes in the Output Process

The natural next step is to think about the process generating output movements over time and ask how it has changed. Does the lower volatility of output reflect a lower standard deviation of output shocks, or a change in the dynamic process through which these shocks affect output, or both?

More concretely, assume that output growth follows an autoregressive (AR) process given by

$$(1) \quad (\Delta y_t - g) = a(L)(\Delta y_{t-1} - g) + \varepsilon_t,$$

2. This is the approach taken by Taylor (2000).

where y_t denotes the logarithm of output in quarter t , Δ denotes a first difference, g is the underlying growth rate of output, ε_t is a white-noise shock with standard deviation σ_ε , and $a(L)$ is a lag polynomial.

The standard deviation of output σ_y then depends both on the standard deviation σ_ε and on the lag polynomial $a(L)$. If $a(L) = a$, for example, output growth follows a first-order AR process, and $\sigma_y = \sigma_\varepsilon / \sqrt{1 - a^2}$, so that the higher is a , the higher is the standard deviation of output.

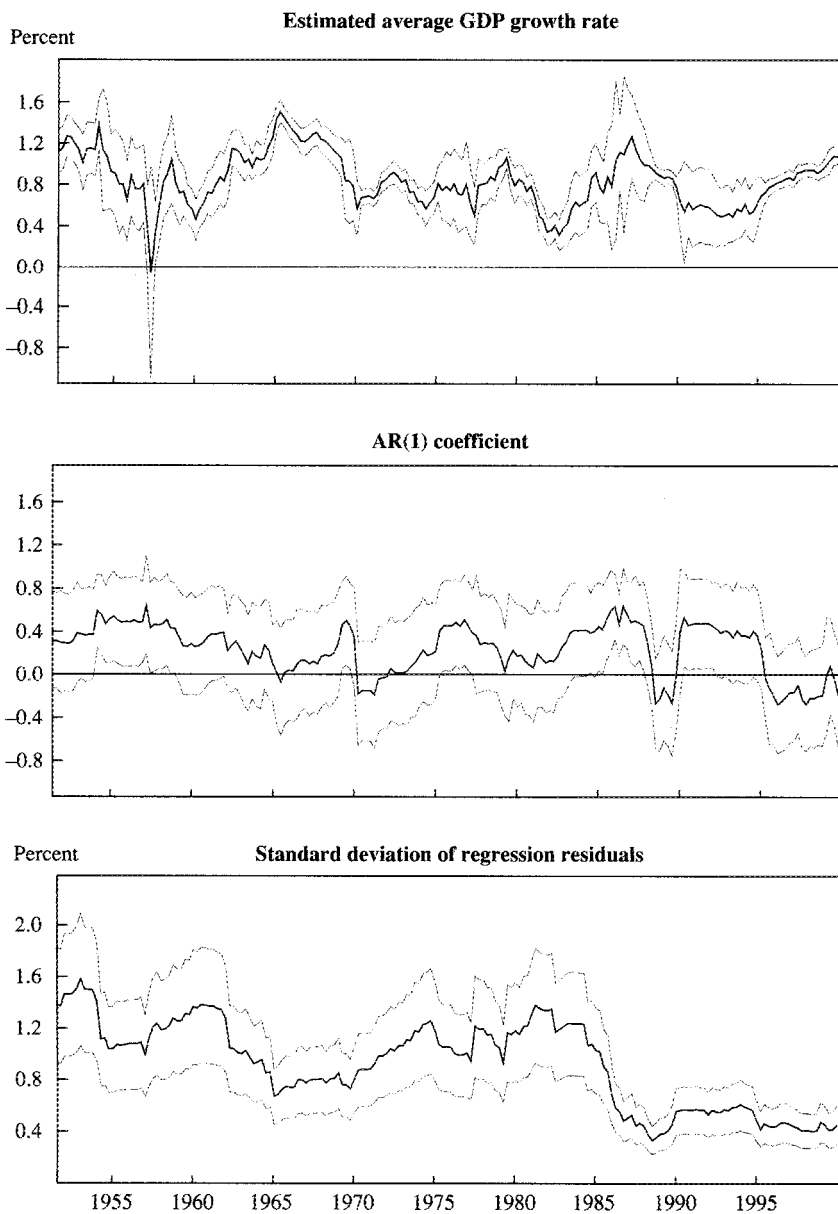
With these points in mind, we estimate equation 1 over a rolling sample from 1952:1 on, again with a window of twenty quarters. We assume the process to be first-order autoregressive, or AR(1); although this does not fully capture the dynamics of output growth, it makes for an easier interpretation of the changes in the process, and all of our conclusions extend to higher-order AR processes. The top panel of figure 2 shows the mean growth rate thus estimated, the middle panel the estimated AR(1) coefficient, and the bottom panel the estimated standard deviation of the shock. The other two lines in each panel show two-standard-deviation bands on each side of the estimate.

The conclusions to be drawn from figure 2 are straightforward. Neither the growth rate nor the AR(1) coefficient shows clear movement over time. The AR(1) coefficient is slightly lower at the end of the 1990s than in the rest of the sample, but the difference is not significant.³ The standard deviation of the regression residual shows the same time pattern as the standard deviation of output growth in figure 1. Indeed, if plotted on the same graph (not shown), their profiles would be nearly identical. In short, the decrease in output volatility appears to come from smaller shocks, rather than from a decrease in the persistence of the effects of these shocks on output.

Back to the Length of Expansions

Having estimated the process for output growth, we can return to the issue of the length of expansions. We proceed in two steps. First we estimate two processes, one for the period 1947 to 1981, the other for the period 1982 to 2000. We choose the split in the sample to coincide with the

3. A perhaps obvious point: changes in the estimated AR(1) coefficient for this univariate representation of output growth do not imply a change in the dynamic structure of the economy. Output movements come from many underlying shocks, each with its own dynamic effects. At different times, different shocks may dominate the (short) subsample used for estimation, leading to different estimated univariate dynamics of output.

Figure 2. Time Variation of Key Parameters, 1952–2000^a

Source: Authors' calculations based on data from Bureau of Economic Analysis, National Income and Product Accounts.
 a. Estimated from equation 1. Results are for twenty-quarter rolling regressions. Two-standard-deviation bands are shown.

peak of the cycle preceding the last two expansions. The intent of the split is to capture the major changes in the process between the start and the end of the sample.

The estimated equations for the two subperiods are

$$(\Delta y_t - 0.87) = 0.31(\Delta y_{t-1} - 0.87) + \varepsilon_t, \sigma_\varepsilon = 1.12 \text{ (for 1947-81)}$$

and

$$(\Delta y_t - 0.85) = 0.48(\Delta y_{t-1} - 0.85) + \varepsilon_t, \sigma_\varepsilon = 0.56 \text{ (for 1982-2000)}.$$

Using the first estimated equation, we generate a sequence of 100,000 observations for output growth, based on draws of the shocks from a normal distribution. Following a long tradition of using a simple rule to approximate the dating of recessions by the National Bureau of Economic Research (NBER), we define the beginning of a recession as two consecutive quarters of negative growth. Similarly, we define the beginning of an expansion as two consecutive quarters of positive growth following a recession. We then compute the mean and median lengths of expansions in the sample of 100,000 observations. Finally, we repeat the process using the second estimated equation.⁴

The results are shown in table 1. The estimates are seventeen and thirteen quarters for the mean and the median expansion length, respectively, for the first subsample; and fifty-one and thirty-five quarters for the second subsample. These means compare with actual mean expansion lengths of nineteen quarters for 1950:1 to 1981:4 (with recessions defined by the same rule as in the simulation, not by NBER dating) and thirty-six quarters for 1982:1 to 2000:4 (but with the second expansion not having ended yet). In other words, the differences between the two estimated AR processes account well for the increase in expansion length from the first to the second sample.

4. Note that the average length of expansions is a very nonlinear function of the underlying parameters of the AR process. By construction, an expansion ends when a recession starts; under our definition of a recession, this requires two consecutive quarters of negative growth. The probability of such an event depends nonlinearly on the average growth rate, the standard deviation of the residual, and the AR(1) coefficient. If, for example, the standard deviation is far below the average growth rate, small changes in the standard deviation will have little effect on the probability of a recession. If it is closer, the same small changes will have a substantial impact on the probability of a recession, and in turn on the length of expansions.

Table 1. Actual and Simulated Lengths of Expansions, 1947–2000

Quarters				
	1947–81		1982–2000	
	Mean	Median	Mean	Median
Actual	19	15	36	n.a.
Simulated ^a	17	13	51	35
Switching growth rate	17	12	55	38
Switching AR(1) coefficient	15	11	83	55
Switching volatility	99	67	14	11

Source: Authors' calculations based on data from Bureau of Economic Analysis, National Income and Product Accounts.
 a. Using parameters for each period that are estimated from actual data from that period. Additional simulations are performed switching one parameter with its opposite-period counterpart, leaving all other parameters unchanged.

To show which parameter changes are responsible for this increase, we next show the results of switching either the estimated mean growth rate, or the AR(1) coefficient, or the standard deviation of the shocks, across the two samples. Not surprisingly, given that the growth rates are nearly the same in the two samples, switching them has no effect on the length of expansions. Switching the AR(1) coefficients leads to shorter expansions in the first subsample (because the effect of a negative shock on output growth is now more persistent, making it more likely that output will decrease for two quarters in a row), and longer ones in the second. But nearly all the action comes from switching standard deviations. If the standard deviation had remained the same as in the first subsample, the mean length of expansions would now be only fourteen quarters, and the median eleven quarters. In short, the large decrease in the standard deviation of output shocks is at the root of the two long expansions the United States has recently experienced. And unless this changes, expansions are likely to be much longer in the future than they were in the past.

Are Recessions Special?

A widespread view of recessions and of output volatility holds that the estimation and the exercise we carried out in the previous section are largely tautological at best, and misleading at worst. According to that view, recessions are largely the result of infrequent large shocks—indeed, sufficiently large and identifiable that they often have names: the first and second oil shocks, the Volcker disinflation, and so on. These shocks, in this

view, dominate output volatility, and therefore there is no great mystery in the measured decline in output volatility. We simply have not had large shocks over the last two decades.⁵

To see whether this is indeed what has been going on, we explore two approaches.⁶ First, we look at what happens to our measure of volatility if we include a recession dummy. Second, we look for signs of large shocks, and associated skewness and excess kurtosis, in the relevant distributions. Both approaches yield similar conclusions. The measured decline in output volatility is not due to the absence of large shocks over the last twenty years. What it captures instead is the decline in the volatility of "routine" quarter-to-quarter changes in GDP growth.

If the decline in volatility simply reflected the absence of large negative shocks and associated recessions, excluding recessions would eliminate our finding of a decline in output volatility. Indeed, excluding recessions from the sample is clearly too strong a correction under our null hypothesis. It corresponds to eliminating large negative realizations just because they happen to be large and negative. But if a decrease in volatility remains even after this overly strong correction, it makes for convincing evidence.

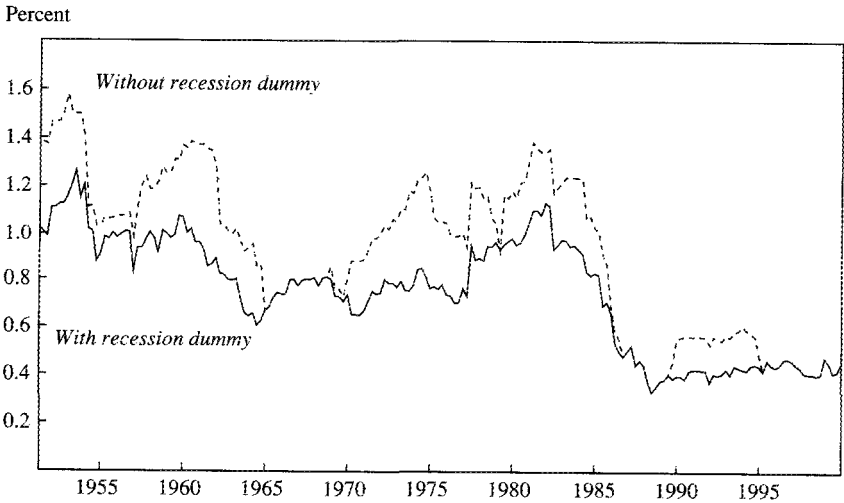
To implement this approach, we reestimate the same rolling AR regressions as before, but we allow for the presence of a dummy variable that takes the value of 1 in each quarter of an NBER-dated recession. The resulting time series for the estimated standard deviation of the residual is plotted in figure 3, together with, for ease of comparison, the standard deviation obtained without a recession dummy (from figure 2).⁷ The results are quite clear. Output volatility is indeed lower in recessions (by construction). But the general pattern is very similar, with a clear trend downward, from roughly 1.2 percent a quarter at the start of the sample to 0.4 percent at the end.

5. This view was forcefully communicated to us by the editors at the start of this project.

6. Obviously, the fact that recessions are typically associated with negative realizations of the output process is an implication of the definition of a recession, not an indication that recessions are special in any particular way. In the same way, the fact that recessions often come with unusually large negative realizations of the shocks also follows from the definition of the recession, and from the fact that the distribution of shocks, conditional on being in a recession, implies a higher probability of large negative shocks.

7. Introducing a dummy for recessions can be thought of as a way of allowing for a lower mean growth rate in recessions. In this sense, this estimation is in the spirit of the Markov switching process estimated by Hamilton (1989) for U.S. GDP.

Figure 3. Volatility of Output Shocks Including and Excluding a Recession Dummy Variable, 1952–2000^a



Source: Authors' calculations based on data from Bureau of Economic Analysis, National Income and Product Accounts.
 a. Twenty-quarter rolling standard deviation of regression residuals from estimation of equation 1. Recession dummy variables correspond to NBER-dated recessions.

The other approach is to actually look for signs of infrequent, large shocks. For example, if the economy is subject to two types of shocks, one frequent and small, the other infrequent and large, we would expect the distribution of output growth to exhibit either skewness or excess kurtosis, or both. It would exhibit skewness if large, infrequent shocks are typically negative, and excess kurtosis if such shocks are equally likely to be positive or negative. Other, more complex models of recessions have similar implications. Although, from the Wold representation theorem, we know that even these models are still consistent with the linear representation given by equation 1, the residuals are likely also to exhibit either skewness or excess kurtosis.⁸

8. Take, for example, the idea that serial correlation of output is greater in expansions than in recessions, clearly a nonlinear feature. Capture this by assuming that output growth is a two-state Markov process, with a high probability that output growth will remain high if it is initially high, and a low probability that output growth will remain low if initially low. It is easy to show that this Markov process will have an AR(1) representation given by equation 1, with the distribution of the residual skewed so that most residuals are small and positive, but with a long negative tail associated with recessions.

This suggests looking at the skewness and excess kurtosis of ϵ_t , the residual obtained from the estimation of equation 1. The results are shown in figure 4. Each point represents the estimate of skewness (top panel) or excess kurtosis (bottom panel) of the residual from estimation of an AR(1) process over the current and previous nineteen quarters. The two panels also show the standard 95 percent confidence bands for the hypotheses that the true measure of skewness or excess kurtosis equals zero. The two panels yield similar conclusions. Except for a brief period during the 1980 recession, there is little evidence of either significant skewness or significant excess kurtosis.⁹

We have also explored other approaches. Following Blanchard and Mark Watson,¹⁰ we estimated a specification in which the output shock is assumed to be the sum of two underlying shocks, one drawn every period from a normal distribution, the other equal to zero with probability $(1 - p)$ and drawn from a normal distribution with larger variance, with probability p . We could not reject the hypothesis that p was equal to zero, and we could not find evidence of a decrease in p over time. In other words, we could not find evidence that the decrease in output volatility has been due to a decrease in the likelihood of large shocks over time.¹¹


Output Volatility and Inflation

Having established the basic fact, we now turn to its interpretation. There are at least two ways to look at the path of output volatility in

9. Under the assumption that large shocks are indeed infrequent, the use of a short window (twenty quarters) implies that there are many subsample periods during which no large shock occurs. Those subsamples will not show evidence of skewness or excess kurtosis. But we would expect many or most recessions to be associated with measured skewness or excess kurtosis. This does not appear to be the case. Nor do we see more evidence of skewness or excess kurtosis if we use a longer window. Over the sample as a whole, there is indeed evidence of significant excess kurtosis, but this appears simply to be due to the decrease in the standard deviation over time. (The distribution of draws from a set of normal distributions with different variances will exhibit excess kurtosis.) Other evidence that skewness and excess kurtosis are not important here is that the earlier simulation results on expansion length are roughly unaffected if we draw shocks by sampling with replacement from the estimated residuals rather than from a normal distribution, as we did in our stochastic simulations earlier.

10. Blanchard and Watson (1986).

11. One hypothesis is that there are large shocks but that their effects appear over a few quarters, making them more difficult to detect. If this were the case, the results of our

Figure 4. Testing for Skewness and Excess Kurtosis, 1952–2000^a

Source: Authors' calculations based on data from Bureau of Economic Analysis, National Income and Product Accounts.
a. In the residuals from estimating equation 1. Shown with two-standard-deviation bands.

figure 1 or, equivalently, at the path of the standard deviation of the residual in figure 2, as the two are nearly identical. The first, which we have implicitly relied on until now, is to see the pattern as a trend decline, temporarily interrupted in the 1970s and early 1980s. This interpretation is shown in the top panel of figure 5, which reproduces the path of the standard deviation of output growth from figure 1 and draws in addition an estimated exponential trend over the period.

The second interpretation, which has been suggested in a number of recent papers,¹² instead invokes a step decrease some time in the early to mid-1980s. This interpretation is shown in the bottom panel of figure 5, which shows how an estimated step function can also fit the general pattern of volatility. Here the step function is drawn assuming that the step decline occurs in 1986:1. The more careful econometric work of Margaret McConnell and Gabriel Perez-Quiros, who estimate rather than assume the break date, finds a slightly earlier date, 1984:1, as the most likely break point.¹³

This second interpretation suggests looking for factors in the economic environment that changed around the mid-1980s. Plausible candidates are an improvement in the conduct of monetary policy or changes in inventory behavior.¹⁴

Under the first interpretation, however (which, we will argue, is more likely to be the right one), one needs to look for two sets of factors: those behind the underlying trend decline in volatility over the last fifty years,¹⁵ and those behind the interruption of that trend in the 1970s. Put another way, the focus shifts from what happened in the 1980s (to explain the step decline in volatility) to what happened in the 1970s (to explain the interruption of the trend for a bit more than a decade).

exercise would be very different if we were to use lower-frequency data. In fact, the results are nearly identical when using annual rather than quarterly data.

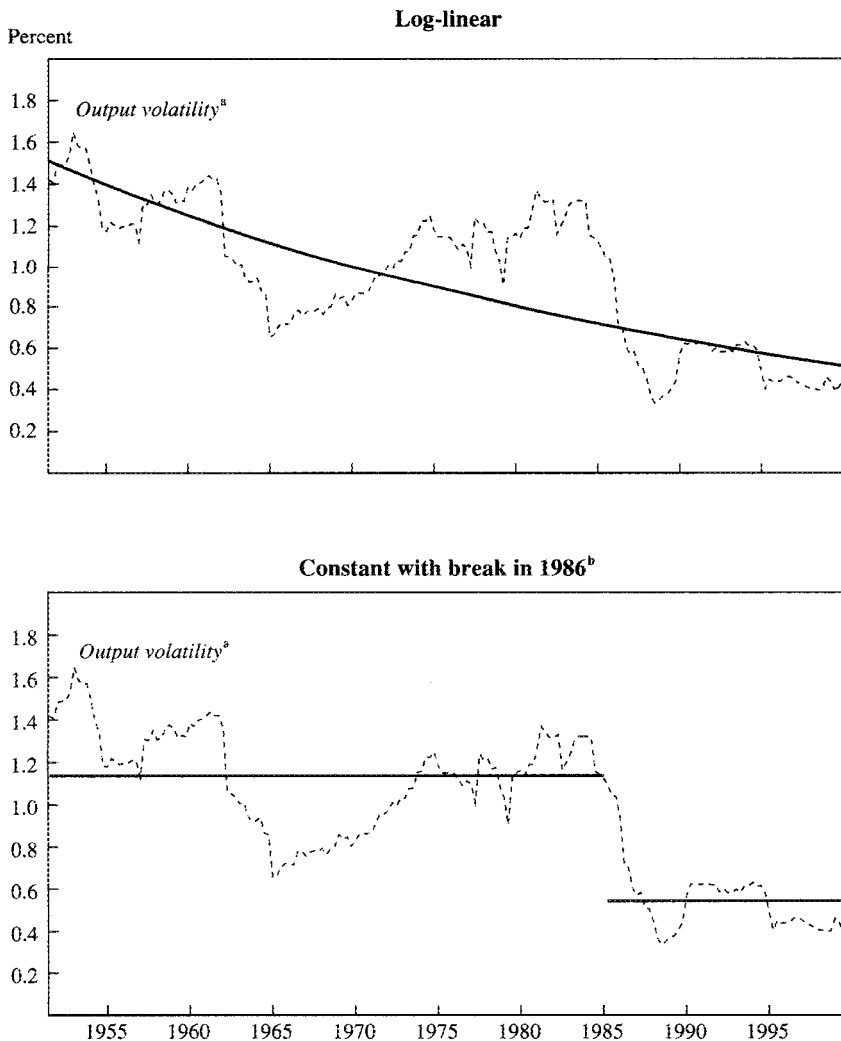
12. In particular, McConnell and Perez-Quiros (2000).

13. The difference comes from our use of a rolling window to capture volatility. A decline in 1984:1 will not necessarily show up until enough earlier observations have dropped out of the window—something that happened around 1986:1.

14. The first hypothesis is argued by, for example, Taylor (2000), and the second by Kahn, McConnell, and Perez-Quiros (2001).

15. Or, indeed, over the past century, if one takes a longer view, informed by evidence from earlier research on volatility since the late 1800s.

Figure 5. Alternative Estimates of the Trend in Output Volatility, 1952–2000



Source: Authors' calculations based on data from Bureau of Economic Analysis, National Income and Product Accounts.
 a. Twenty-quarter rolling standard deviation of quarterly real GDP growth.
 b. Means for 1952–85 and 1986–2000.

Inflation and Inflation Volatility

That the 1970s were different is not very controversial. The U.S. economy was affected by major increases in the prices of raw materials, including oil. Inflation increased, to return to a lower level only after the disinflation of the early 1980s. That these shocks, and perhaps inflation itself, may have led to more output volatility does not seem implausible.

Figure 6 shows the relationship between inflation and output growth volatility. The top panel plots output growth volatility against the twenty-quarter rolling mean of the inflation rate, with inflation measured using the GDP deflator. The bottom panel plots output volatility against inflation volatility, both constructed as twenty-quarter rolling standard deviations. All variables, including mean inflation, are measured at quarterly rates.

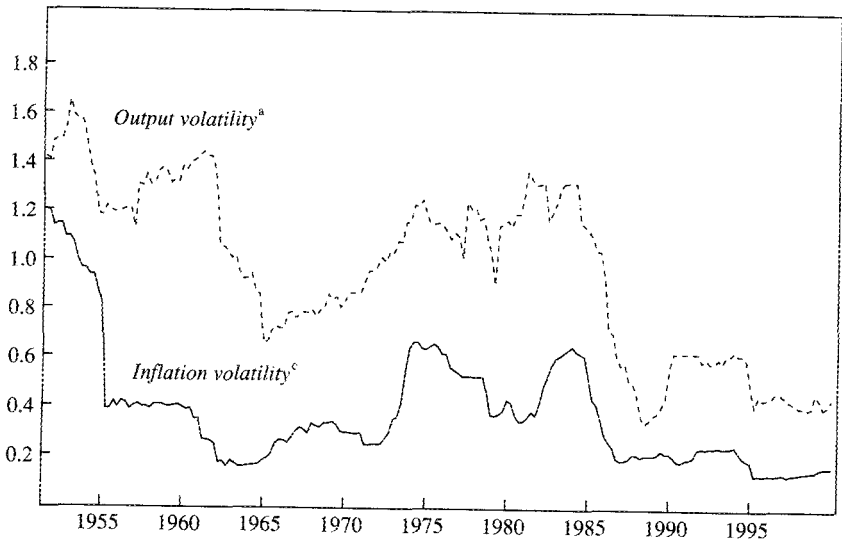
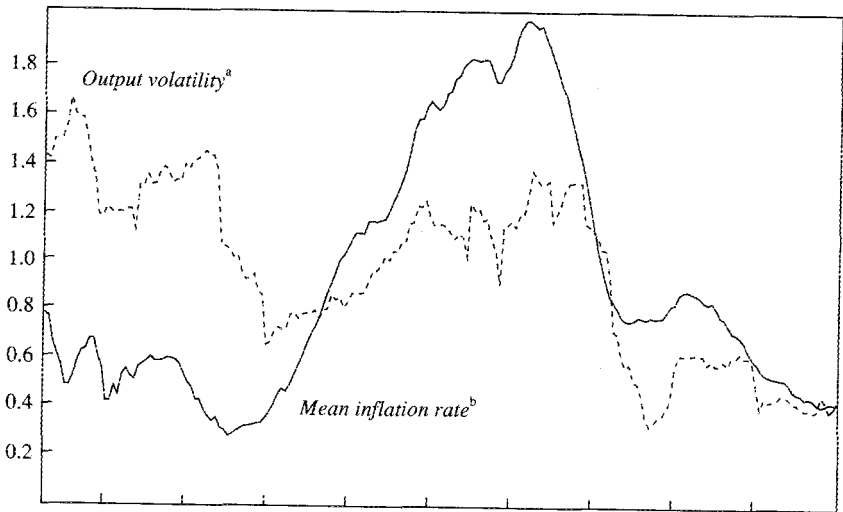
The temporary increase in output volatility in the 1970s and early 1980s is clearly correlated with the temporary increase in the level of inflation. Output volatility seems, however, more strongly related to the volatility than to the level of inflation. Simple regressions of output volatility on the level of inflation, inflation variability, and an exponential time trend show all three factors to be significant, with the level and the variability of inflation playing roughly similar roles quantitatively, and the negative time trend remaining important and statistically significant.¹⁶

Correlation between inflation and output volatility, however, does not imply causality from inflation to output volatility. At least one plausible alternative is that the correlation reflects a common dependence of inflation and output volatility on third factors, such as the supply shocks of the 1970s. Here international evidence can help. First, and obviously, it can tell us whether the patterns observed in the United States are representative of what happened to output volatility, and to the relation between output and inflation volatility, elsewhere. But also, if we are willing to assume that the supply shocks of the 1970s were largely common across

16. One worry is that measurement noise in the decomposition of nominal GDP will create a spurious positive correlation between output and inflation volatility. But the results are very similar if we use the consumer price index, where the issue is likely to be less important. Another problem with regressions of this kind is the use of moving averages for standard deviations and means on both the left- and the right-hand sides. Estimation of a potentially more appropriate GARCH (generalized autoregressive conditional heteroskedasticity) model for output growth, allowing the variance of output shocks to be a function of inflation volatility, the inflation level, and a time trend, yields very similar results.

Figure 6. Relationship between Output Volatility and Inflation, 1952–2000

Percent



Source: Authors' calculations based on data from Bureau of Economic Analysis, National Income and Product Accounts.

a. Twenty-quarter rolling standard deviation of quarterly real GDP growth.

b. Twenty-quarter rolling mean of the quarterly inflation rate, as measured by the GDP deflator.

c. Twenty-quarter rolling standard deviation of the quarterly inflation rate, as measured by the GDP deflator.

countries, then cross-country evidence gives us a way of controlling for the presence of these shocks, by treating them as fixed effects in a cross-country panel regression. In other words, such a regression can help us establish the relationship between output and inflation volatility, controlling for supply shocks. With this in mind, we now turn to the evidence from the G-7 countries.


A Look at the Other Group of Seven Countries

Figure 7 shows the path of output volatility for all the G-7 countries. For clarity's sake, we have grouped the countries into three panels: the United States, the United Kingdom, and Canada (top panel); France, Germany, and Italy (middle panel); and Japan (bottom panel). In each case the measure of output volatility is again the five-year rolling standard deviation of output growth, using a window of twenty quarters. Because the data we have start only in 1960 (in 1982 for Italy), the different measures are available only from 1965:1 on (1987:1 for Italy), resulting in a shorter sample than the one used for the United States above.

The top and the middle panels show that output growth volatility in six of the G-7 countries has followed a roughly similar path over the period. In all of these countries the standard deviation of output growth has declined: whereas in the early 1960s it ranged from about 1.5 percent in Germany to a little below 1.0 percent in the United States, by the late 1990s it was around 0.5 percent in all six countries. One of the striking characteristics of these two panels is indeed how similar the standard deviation of output growth is across these countries today. The general decline and convergence suggest the presence of common, long-lasting forces across countries. Looking more closely, however, there are also clear differences across countries, especially in timing. After the general increase of the early 1970s, the decrease in volatility took place earlier in Germany, and later in Canada.

The only G-7 country where the pattern is clearly different is Japan. After falling from the early 1960s to the late 1980s, the standard deviation of Japanese output growth rose in the 1990s and is now higher than it was at the start of the sample. To the extent that this increase largely coincides with the long Japanese slump of the 1990s, it is tempting to search in that direction for an explanation. For example, a decrease in liquid asset holdings by both firms and consumers may have led to stronger

Figure 7. Volatility of Output Growth in the G-7 Countries, 1965–2000^a



Source: Authors' calculations based on data from Datastream.

a. Twenty-quarter rolling standard deviation of quarterly real GDP growth. Data are through 2000:3.

effects of cash flows on consumption and investment, leading to stronger multiplier effects of shocks on output. The zero floor on nominal interest rates may have constrained monetary policy responses. We do not explore these hypotheses further in this paper, but we find the coincidence of the long slump and the increase in volatility intriguing and potentially useful in learning what has happened in other countries.

Leaving Japan aside, we return to the relationship between output volatility and inflation in the other six countries. To do so, we run the following panel regression:


$$(2) \quad \sigma_{yit} = \beta_i + \beta_t + a_1 \bar{\pi}_t + a_2 \sigma_{\pi t} + \varepsilon_{it}$$

where i indexes countries and t time, so that the β_i s are country fixed effects, the β_t s are time fixed effects, σ_y and σ_π are rolling standard deviations of output and inflation, and $\bar{\pi}$ is a rolling mean of the inflation rate.

If the effects of the supply shocks of the 1970s on output volatility were indeed common across countries, this specification will give us the relationship between output volatility and inflation volatility, controlling for supply shocks. The assumption is probably too strong, however: the effects of supply shocks were different across countries, and these differences may well have been associated with different paths through time of both the level and the volatility of inflation. In this case the coefficient on inflation will still pick up some of the effects of supply shocks. Nevertheless, even in this case, this cross-country specification is an improvement over the U.S. regression (in which we could not introduce fixed time effects) presented earlier. (Note that, even under the assumption of common supply shocks, this specification does not resolve other potential identification problems. One of these is the possibility that the relationship between output volatility and inflation volatility reflects causality from output to inflation volatility, through the response of monetary policy, or a dependence on other, third causes, such as an improvement in the conduct of monetary policy leading to both lower output volatility and lower inflation volatility.)

Estimation yields coefficients of $a_1 = -0.02$, with a t statistic of -0.7 , and $a_2 = 0.67$, with a t statistic of 13.7 . Thus it is inflation volatility, rather than the level of inflation, that appears to matter, and to matter strongly. The best way to summarize the implications of the regression is through the three panels of figure 8. The top panel shows the actual and the fitted values of output volatility for the United States (dropping $\bar{\pi}_t$ from the panel regression; nothing is changed by this). The conclusion to be drawn is that the panel specification does a good job of fitting the U.S. data.

Figure 8. Volatility of Output Growth in the United States and Its Main Explanatory Variables, 1965–2000^a



Source: Authors' calculations based on data from Datastream.
 a. Estimated from equation 2 using panel data for the G-7, excluding Japan.
 b. Twenty-quarter rolling standard deviation of quarterly real GDP growth.

The other two panels show how much of the variation derives from movements in inflation, and how much is due to the common time components. The second panel shows the fitted value of output volatility (repeated from the first panel), together with the inflation component, $a_2\sigma_{\pi_{it}}$. This panel makes clear that the increase in inflation volatility accounts for the reversal in trend from the early 1970s to the early 1980s. The third panel shows the fitted value of output volatility (again repeated from the first panel), together with the common time component, β_t . This suggests a steady underlying trend decrease from 1960 on.

In short, this decomposition suggests a trend decrease in output volatility, temporarily interrupted by an increase in inflation volatility. Under that interpretation, the sharp decline in output volatility in the 1980s appears to be associated with a sharp decline in inflation volatility at that time. To get a better understanding of both the trend decrease and the temporary reversal in output volatility, the next section goes one level down and looks at trends in the individual components of GDP.

A Disaggregated Look

In his 1960 presidential address to the American Economic Association, Arthur Burns argued that a trend decline in output volatility was indeed under way.¹⁷ Composition effects (including the steady shift to services), improvements in capital markets, the increasing ability of consumers to smooth consumption in the face of variations in income, the increase in the income tax, and stronger automatic stabilizers all led and, Burns argued, would continue to lead to more economic stability.

He was clearly right about the trend. Was he right about the channels? Here we make a first pass at the answer. From a statistical accounting point of view, one can think of the volatility of output as depending on three sets of factors: the volatility of its components, their covariation, and their relative weights. We look at each in turn.

Volatility of Output Components

Take the standard decomposition of GDP by type of purchase and type of purchaser: consumption, investment, government spending, net exports,

17. Burns (1960).

and inventory investment. Let each of these components, in real terms, be denoted by X_i , so that

$$(3) \quad Y_t = \sum_i X_{it},$$

where Y_t is real GDP, measured quarterly. For each component we consider two measures of volatility.

The first is the same as for GDP earlier, namely, the rolling standard deviation of the rate of growth of X_{it} , which we denote by $\sigma_{x_{it}}$. This measure makes little sense, however, for inventories and net exports, which change sign and are frequently close to zero. Thus we construct the rolling standard deviation of growth for consumption, investment, and government spending only.¹⁸

The second measures the volatility of a variable commonly called the "growth contribution" of each component, which adjusts for the share of the component in GDP. A very volatile component may have little effect on overall output volatility if it accounts for a small share of GDP. The variable is defined as $\Delta X_{it}/Y_{t-1}$, and our measure of volatility is once again the rolling standard deviation. Note that this measure is well defined for all components of GDP regardless of whether they change sign or are close to zero. Note also that the variable can be rewritten as $(\Delta X_{it}/X_{it-1})(X_{it-1}/Y_{t-1})$, so that, if the share is stable at high frequency, the standard deviation will be roughly equal to the share of the component of GDP times the standard deviation of that component's growth rate. For both volatility measures, the window we use to compute standard deviations is again twenty quarters. Rates of change are quarterly, not annual, rates.

The results are plotted in figures 9 and 10. The figures show that the two measures move together at high frequency, reflecting the stability of the shares.

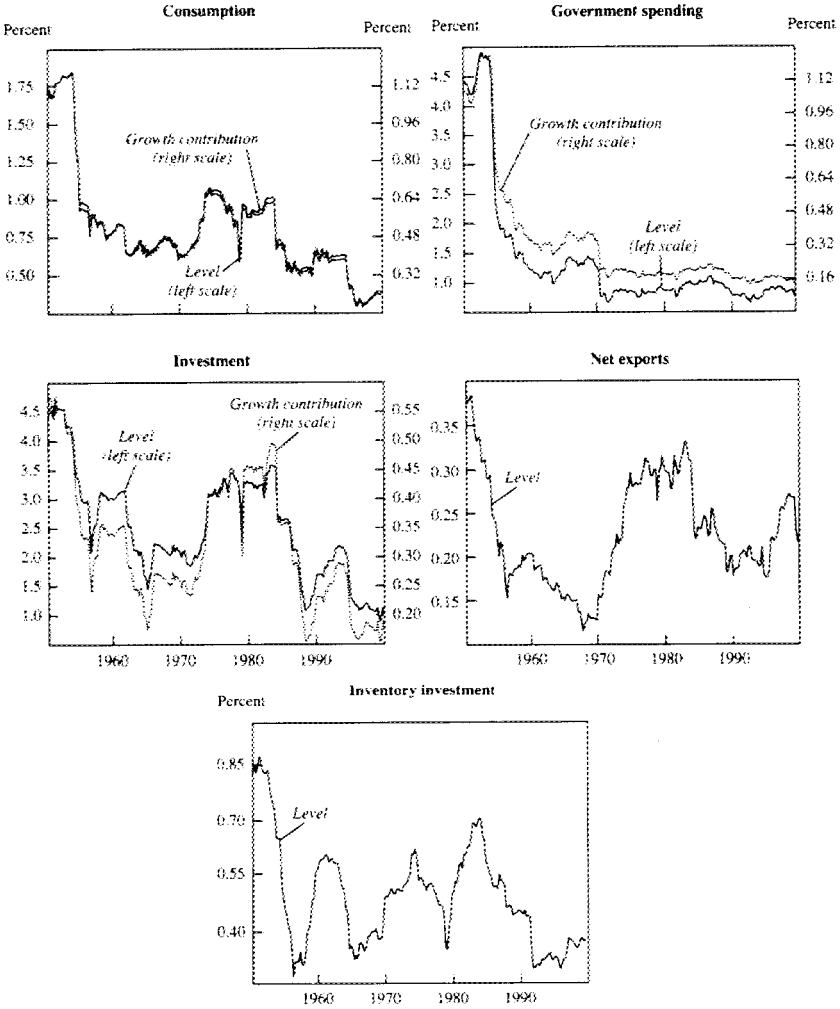
From figure 9 we draw the following conclusions:

—The volatility of government spending (and of fiscal policy in general) was very high during the Korean War. It fell rapidly in the 1950s and has remained low ever since.

—There is no clear trend in the volatility of net exports or in the volatility of inventory investment, although the latter was low in the 1990s. (This,

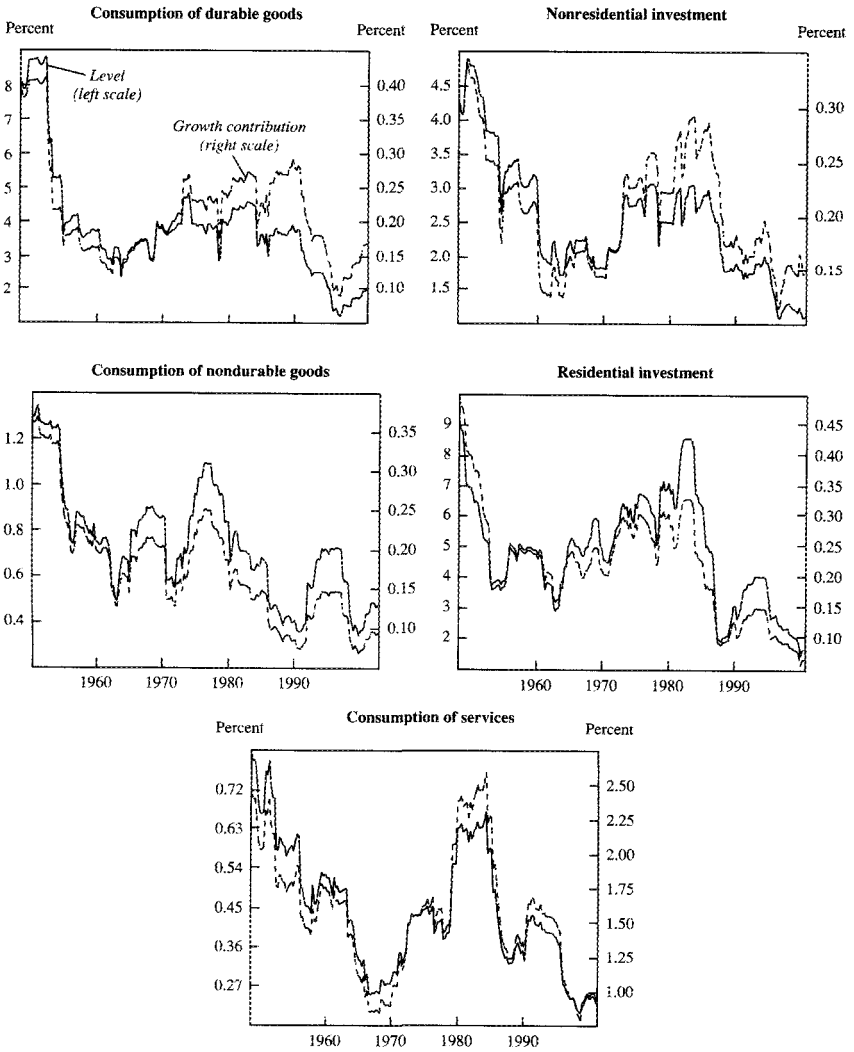
18. In parallel with our exploration of GDP, we have estimated AR processes for each component. Although we do not present the results here, the general conclusion is the same as for GDP. For the most part, the decrease in volatility comes from a decrease in the volatility of the shocks rather than from a change in dynamics.

Figure 9. Volatilities of Components of GDP, 1952–2000^a



Source: Authors' calculations based on data from Bureau of Economic Analysis, National Income and Product Accounts.
a. Twenty-quarter rolling standard deviation of quarterly growth and growth contribution.

Figure 10. Volatilities of Components of Consumption and Investment, 1952–2000^a



Source: Authors' calculations based on data from Bureau of Economic Analysis, National Income and Product Accounts.
 a. Twenty-quarter rolling standard deviations of quarterly growth and growth contribution.

together with the change in correlation reported below, suggests a recent change in the behavior of inventory investment.)

—Most of the trend decrease in output volatility can be traced to a decrease in the volatility of consumption and investment. After a large decrease in the 1950s, consumption volatility has continued to decrease, from about 0.75 percent in the 1960s to 0.30 percent in the late 1990s. The decrease in investment volatility has been more limited. The standard deviation of our second measure is nearly the same in the late 1990s as in the 1960s.

Given that much of the action comes from consumption and investment, figure 10 goes one step further in the disaggregation, to trace the volatility of the components of consumption and investment. Relative declines in the volatility of all three components of consumption—spending on durables, on nondurables, and on services—are roughly similar. Their timing is slightly different, however, with much of the trend reversal in consumption in the 1970s and the early 1980s coming from consumption of services.

We think these are slightly surprising findings. One might have expected improvements in financial markets to lead consumers to choose a smoother consumption path, thus leading to smoother consumption of services and nondurables. But one would also have expected an improved ability to borrow and lend to lead to a stronger stock-flow adjustment for purchases of durables, and thus potentially to more volatility of durables purchases. This, however, does not seem to be the case.

The two series for investment volatility exhibit quite different patterns. Nonresidential investment shows a steady decline and a limited increase in the 1970s. Residential investment volatility shows a steady increase from the 1950s to the mid-1980s and a sharp decrease after that. The latter coincides with the elimination of interest rate ceilings on savings and loan institutions (the end of Regulation Q), making it a plausible candidate explanation.

Correlations of Output Components

The standard deviation of output depends not only on the standard deviations of its components, but also on their correlations. To show what has happened in this regard, we construct the correlation of each component with final sales (GDP minus inventory investment) or, more specifically,

the correlation of $\Delta X_{it}/Y_{t-1}$ with $\Delta S_t/Y_{t-1}$, where S_t is final sales. Again we use a window of twenty quarters.

These results are shown in figure 11. The correlations change over time (as we would expect if the subsamples are dominated by different shocks, with different implications for the correlation between each component and final sales). But except for one series, they do not show clear trends. The exception is the correlation between inventory investment and sales. Until the mid-1980s, inventory investment tended to move with sales, leading to a higher variance of production than of sales—a fact studied at length in the research on inventory behavior. Since the mid-1980s, inventory investment has become countercyclical, leading to a decline in the variance of output relative to sales. This fact, which Kahn, McConnell, and Perez-Quiros have examined, must have come from a change in the inventory management methods of firms. It is clearly one of the factors behind the decrease in output volatility in the 1980s, although only the last in a long series of structural changes.¹⁹

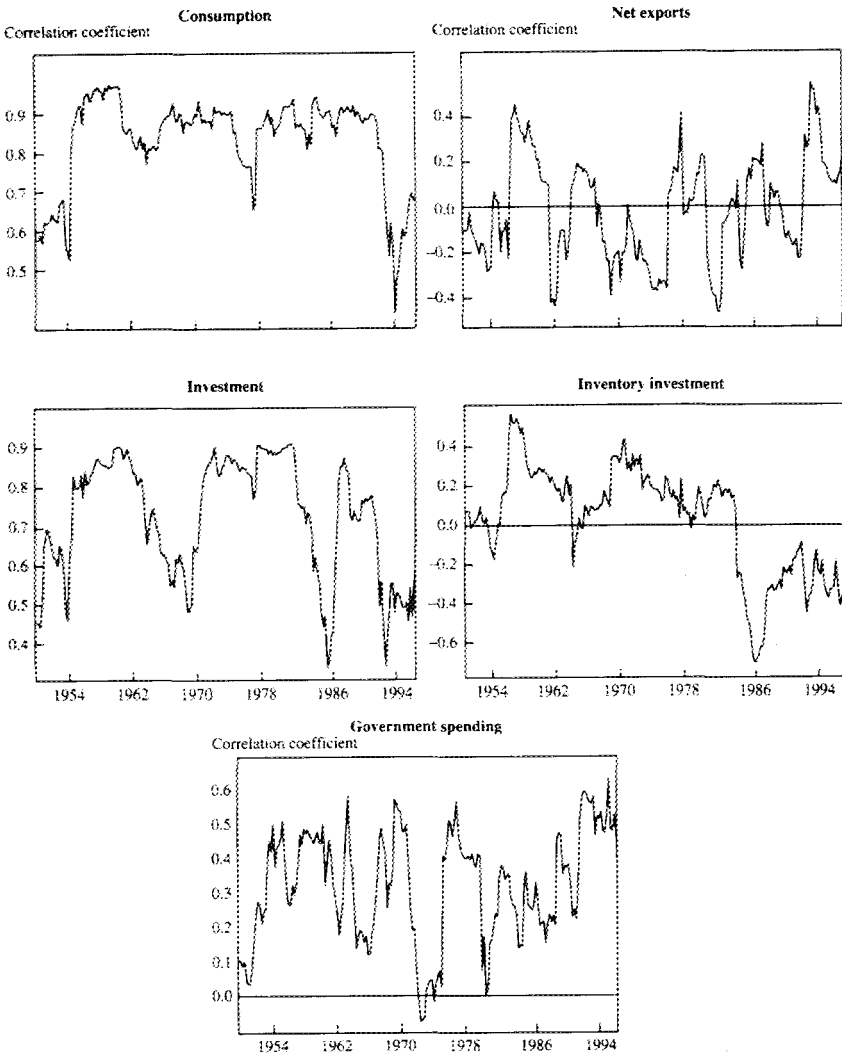
Composition Effects

The composition of GDP has changed substantially over the last fifty years. The three main changes, at the level of disaggregation examined here, are an increase in the share of (high-volatility) fixed nonresidential investment, from 9.4 percent of GDP in 1950 to 13.7 percent in 2000; a decrease in the share of nondurables consumption, from 33.4 percent to 20.2 percent; and a mirror increase in the share of (low-volatility) consumption of services, from 21.7 percent to 39.4 percent.

To characterize the effects of composition on output volatility, a simple approach is to compute the volatility of a counterfactual series for output growth, using 1947 shares rather than current shares to weight the components. Specifically, we construct the counterfactual output growth series as follows. Write the rate of growth of output as

19. There is a puzzle here as well. The change in correlation roughly coincides with the introduction of just-in-time inventory management methods, which have led to lower inventory-to-sales ratios. It is not clear, however, why they should have led the correlation to change from positive to negative. Better tracking and forecasting of sales, and the ability to maintain a stable inventory-to-sales ratio, should lead to more procyclical, not less procyclical, inventory investment.

Figure 11. Correlations between Final Sales Growth and the Growth of Its Components, 1952–2000^a



Source: Authors' calculations based on data from Bureau of Economic Analysis, National Income and Product Accounts.
a. Twenty-quarter rolling correlations.

$$(4) \quad (\Delta Y_t / Y_{t-1}) = \sum_i (\Delta X_{it} / Y_{t-1}) + \sum_j (\Delta X_{jt} / Y_{t-1}),$$

where the terms in the first sum are the terms that are always positive, and the terms in the second sum are the terms that change sign in the sample (net exports and inventory investment). We can rewrite this expression as

$$(5) \quad (\Delta Y_t / Y_{t-1}) = \sum_i \alpha_{it-1} (\Delta X_{it} / X_{it-1}) + \sum_j (\Delta X_{jt} / Y_{t-1}),$$

where α_{it-1} is the share of component i at time $t - 1$. Once again, constructing $\Delta X_{jt}/X_{jt-1}$ does not make much sense for inventories and net exports, as they are frequently around zero. Consequently, we treat them separately. We then construct the counterfactual series for output growth as

$$(6) \quad (\Delta Y_t / Y_{t-1})^* = \sum_i \alpha_{i:1947} (\Delta X_{it} / X_{it-1}) + \sum_j (\Delta X_{jt} / Y_{t-1}),$$

where $\alpha_{i:1947}$ is the 1947 share of component i . We then construct rolling standard deviations for actual and counterfactual output series. There is no need to present a figure, as the two series are nearly indistinguishable. The different changes in composition nearly offset each other, and the final values are within 0.05 percent of each other. Composition effects therefore have little to do with the general pattern of output volatility over the last fifty years.

Conclusions

We have documented the long and large decline in output volatility over the last half-century. We have shown that this phenomenon does not have one, but many proximate causes. Among them are a steady decrease in investment volatility and, even more so, of consumption volatility; a decrease in the volatility of government spending early on; and a change from procyclical to countercyclical inventory investment in the 1990s. Many questions remain, however.

The first set of questions concerns the deeper causes of the decrease in volatility, from the role of policy (especially monetary policy) to the role of structural changes (especially changes in financial markets). Our findings suggest that monetary policy has played a complex role. On the one hand, the trend decrease in output volatility from 1950 on does not lend

much support to the idea that what we are seeing is primarily the result of a dramatic recent improvement in the conduct of monetary policy—that is, of a Greenspan effect. On the other, the dramatic decrease in output volatility in the mid-1980s can be interpreted in two ways, both related to monetary policy.

The first is that this decrease was indeed the result of smarter counter-cyclical monetary policy, leading to better output stabilization from the 1980s on. This explanation runs into a puzzle, however. Given the lags in the effects of monetary policy on output, one would have expected better monetary policy to show up primarily as shorter-lived effects of shocks on output, and thus as a decrease in the AR(1) coefficient in the univariate AR representation. There is no evidence that this has been the case.²⁰ The other interpretation is that the decrease in output volatility was associated with—and may have been largely caused by—the decrease in inflation volatility that occurred around the same time. But even this second interpretation implies a role for monetary policy. The increased stability of inflation is likely to have been due, in large part, to better monetary policy.

Our findings also suggest a role for improvements in financial markets in reducing consumption and investment volatility. But here again the argument is not straightforward. On theoretical grounds, it is not obvious that more efficient financial markets should lead to lower consumption volatility. Although, for given interest rates, they plausibly lead to a decrease in the volatility of consumption services and nondurables, they also allow consumers to adjust faster toward their desired stock of durables, leading, other things equal, to more volatility of spending on consumer durables. The same argument applies to investment. The evidence is, however, of a decrease in volatility in all components of consumption and investment.

The issue of the relative roles of monetary policy and financial market improvements in reducing output volatility is a fascinating one. In that respect, developments in Japan in the 1990s are both intriguing and incon-

20. A more sophisticated argument is that, despite the lags in monetary policy, better policy might have reduced the variance of measured output shocks, leading agents to expect shorter-lived effects of the underlying shocks on GDP, and thus to react less to these shocks in the first place. But even in this case, better policy should be reflected both in a smaller variance of measured output shocks and in shorter-lasting effects of shocks on output—and thus in a decrease in the AR(1) coefficient.

clusive. As we saw, output volatility increased substantially in Japan in the 1990s. But was that increase due to monetary policy or to changes in financial markets (or to something else)? The answer is far from obvious. Monetary policy, both current and anticipated, was clearly limited by the constraint that interest rates be nonnegative—the liquidity trap. And because of the problems Japanese banks were facing, intermediation was clearly disrupted. Only a more disaggregated examination will help attribute blame.

A second set of issues concerns the implications of our findings. We feel reasonably confident in predicting from our results that the increase in the length of expansions is here to stay. (This, however, is not a prediction that the United States will not go through a recession in the near future, nor do we claim that the New Economy has eliminated the business cycle.) The decrease in output volatility appears sufficiently steady and broad based that a major reversal appears unlikely. This implies a much smaller likelihood of recessions.

Lower output volatility suggests lower risk, and thus changes in risk premiums, in precautionary saving, and so on. Interestingly, however, the decrease in output volatility has not been reflected in a parallel decrease in asset price volatility. As others have documented, there is little evidence of a trend in the volatility of the Dow Jones Industrial Average.²¹ Ultimately, of course, what matters is not aggregate risk but the risk borne by individuals. We do not know what has happened to the volatility of idiosyncratic shocks during the period. We intend to explore all these avenues in the near future.

21. This is not necessarily a puzzle. If we think of the better use of monetary policy as one of the factors behind the decrease in output volatility, stronger stabilization efforts may require sharper movements in interest rates, and thus potentially stronger movements in asset prices. There is, however, little evidence of increased volatility in real interest rates, let alone nominal interest rates.

Comment and Discussion

Benjamin M. Friedman: The goal of Olivier Blanchard and John Simon's paper is to study the business cycle, focusing in particular on aggregate real output in the United States. Their principal finding is that the shocks affecting output growth have become less volatile in recent years. The standard error of the autoregression that serves as their central vehicle of analysis, estimated on a rolling-sample basis, has declined from about 1 percent a quarter (4 percent annualized), on average from the 1950s through the mid-1980s, to roughly $\frac{1}{2}$ percent a quarter since then. Because the average growth rate of real output for the entire sample, from 1947 through 2000, is 0.9 percent a quarter, this sharp decline in the absolute magnitude of the estimated shocks obviously means that lately there have been fewer quarters when the measured growth rate has been less than zero. To the extent that movements of real GDP as measured by the U.S. Department of Commerce capture the broadly based fluctuations in nonfinancial economic activity that the NBER business cycle dating process emphasizes, there have therefore been fewer recessions.

The immediate question is what to make of this finding. In their introduction the authors say that their paper will "take up the question of whether recessions are special . . . whether what we have seen over the last twenty years is simply the absence of large shocks and nothing more." They preview their conclusion as follows: ". . . this is not the case. The measured decrease in output volatility has little to do with the absence of large shocks in the recent past." Later on they summarize their results by stating, "The measured decline in output volatility is not due to the absence of large shocks over the past twenty years. What it captures

instead is the decline in the volatility of 'routine' quarter-to-quarter changes in GDP growth."

To be sure, one should not read the results of Blanchard and Simon's paper to say that the infrequency of recessions in the United States over the last two decades has reflected the absence of large shocks *and nothing more*. But the absence of large shocks certainly *is* a major part of what the authors have to report. Inspection of the estimated residuals from the authors' equation (which Olivier Blanchard was kind enough to provide to me) shows that, without exception, every one of the nine official NBER recessions in the United States since World War II has involved at least one negative residual larger (in absolute value) than the 0.9 percent mean quarterly growth rate of real output. And in the past there were plenty of large residuals other than just during recession episodes. Including both positives and negatives, from 1947 through 1984 their results show fifty-eight estimated residuals larger than 0.9 percent.

By contrast, since 1985 only four of their estimated quarterly residuals have exceeded 0.9 percent; two of these were positive and two negative. And one of the two large negatives occurred during the lone recession the United States has experienced during these years, that of 1990–91. (The other was in the first quarter of 1993.) Blanchard and Simon's basic point is that output has become less volatile, and this is surely true. But there is no way to duck the fact that, in their estimated autoregression, the absence of recessions is very much associated with the absence of large negative shocks.

Yet another way of looking at the data, which the authors do not discuss, leads to the same conclusion. It is conventional to estimate regressions for real output using quarterly data, presumably because the Commerce Department reports GDP this way. But no economic theory of which I am aware guarantees that the calendar quarter is the right level of time aggregation for investigating the kinds of shocks that matter for business cycles. A single large shock spread out over more than a single measured quarter would look in quarterly observations like a series of smaller shocks occurring over a sequence of consecutive quarters.

Blanchard and Simon's autoregression exhibits no meaningful serial correlation. (The Durbin-Watson statistic is a comforting 2.05.) Nonetheless, from time to time there are runs of consecutive estimated negative residuals, and these runs are also very much part of the story of recessions. Of the nine recorded postwar recessions, seven have involved three or

more consecutive negative residuals, and six of those seven have involved four or more negative residuals in succession. (The only two exceptions each involved a “near run,” in which four out of five consecutive residuals were negative.) Moreover, the identification of a recession with a run of negative Blanchard-Simon residuals is not only almost necessary, but almost sufficient as well. Only once in their sample spanning more than 200 quarterly observations does their regression exhibit three successive negative residuals outside of an NBER recession period. And that episode was during 1989, which numerous analyses have indicated was unusual, and in particular recession-like, in a variety of ways.¹

In sum, whether the matter is to be framed in terms of the absence of large shocks or more generically in terms of reduced volatility, the result from this part of the paper is clear: the time-series behavior of output in the United States has changed (or, to anticipate the discussion below, has been changing), and that change has a lot to do with the recent infrequency of recessions. The broader question to be put to this part of Blanchard and Simon’s analysis, however, is, What is learned from viewing the time series of real output in this way—that is, through the lens of their first-order autoregression? The regression, estimated over the full sample, has an adjusted R^2 of 0.11. The coefficient on the lagged dependent variable is 0.34. (These figures are not reported in the paper. I am again grateful to Olivier Blanchard for providing them.) Hence the fitted value of real output growth for each quarter is simply the 0.9 percent mean adjusted by one-third of whatever was the difference between the measured growth rate and this mean in the previous quarter. All else is “shocks.” To put the question in more quantitative terms, What do we learn from excluding this 11 percent of the variation of real output growth, and focusing our attention on the remaining 89 percent, that we did not already know from looking at real output growth itself—as, indeed, Blanchard and Simon do in much of the rest of the paper?

The most interesting part of Blanchard and Simon’s paper, and the part that I think offers the greatest promise for future research, is the demonstration that the decline in the volatility of output growth (focusing now on the raw series, not the estimated autoregression residuals) roughly lines up with a decline in the volatility of price inflation. In particular, the match-up with output volatility is distinctly better for the volatility of inflation

1. See, for example, Perry and Shultze (1993).

than for the mean of inflation. Taking the research in this direction opens room to investigate (although the authors do not do so) how familiar, systematic forces like monetary and fiscal policies, and understandable categories of shocks like energy price movements and other supply shocks, can enter the central story of business cycles.

For example, one longstanding view of how recessions come about in the postwar U.S. economy is that the Federal Reserve makes them happen, either deliberately or by overshooting the mark, but in either case through the application of tight monetary policy. The presumed motivation for that tight policy is, of course, the need to prevent inflation from developing or to slow an inflation that has already begun. Hence recessions should be observed following periods in which monetary and fiscal policies have stimulated the economy beyond its productive capacity, or in which exogenous movements in consumer spending or business investment (due to a surge of confidence, for example) have done the same, or perhaps both. A separate but highly similar view is that recessions follow after an adverse shock, such as a rise in oil prices, has reduced the economy's ability to produce outputs from inputs at any given cost, and the central bank chooses not to allow rising prices to absorb the entire backshift in the aggregate supply curve.

The rough correspondence shown in the lower panel of the authors' figure 6 is consistent with either of these accounts of how recessions occur, both of which assign a central role to monetary policy reacting to either the reality or the anticipation of inflation. The correspondence is also consistent with a quite different view, however: that recessions come and go for reasons entirely independent of inflation, and that it is the fluctuation of output, sometimes above the economy's productive capacity and sometimes below, that then causes inflation to move up or down.

As Blanchard and Simon rightly point out, in this context correlation does not necessarily imply causality. Their response to this problem is to turn to evidence from the other G-7 economies, using a panel regression (omitting Japan) of moving-average output volatility on moving-average inflation volatility and the moving-average inflation mean, allowing for time effects. The result mirrors the finding for the United States alone that is evident from figure 6: output volatility is systematically related to inflation volatility, not to the inflation mean.

But the panel regression resolves the *causality* question no better than a single equation for the United States would—nor, for that matter, any

better than does simple inspection of figure 6. The authors' results are consistent with the interpretation that shocks, some common across all six countries and some not, create inflation; that central banks respond with tight monetary policies that slow output growth; and that these episodes of tight monetary policy in turn render output more volatile (perhaps in a way that involves out-and-out recessions, perhaps not). But the results are also consistent with the alternative account whereby movements in output relative to capacity trigger movements in inflation, so that independent forces that sometimes render output more volatile then cause inflation to be more volatile as well.

The authors' aim in turning to the panel regression is to control for any such independent forces—for example, oil price shocks—that act in common across the six included economies. It is instructive that the results hold up when they do so. But controlling for common supply effects (and common demand effects, too) is not the same as resolving the causality issue. On the assumption that the authors are right that what is involved here is inflation volatility and not inflation itself, the panel regression is no more informative than a single-country regression on the question of whether we should be regressing output volatility on inflation volatility or vice versa.

Finally, the analysis at the end of the paper, focusing on changes in the time-series behavior of individual components of aggregate spending (now again for the United States only) is instructive in its own right. But it, too, cannot resolve the more interesting question of whether output volatility is driving inflation volatility or vice versa. At first blush, the finding that all major components of U.S. GDP have shown less volatility in recent years lends credence to the story that places Federal Reserve (over)reaction to inflation at the heart of the matter. Otherwise, why would *all* of the components have become less volatile over the last fifty years? Exogenous shocks to spending, such as changes in consumer confidence or in the "animal spirits" of entrepreneurs, would more likely have affected different components of spending differently.

On closer inspection, however, the timing of the decline in volatility does turn out to vary from one component of GDP to another. Moreover, Blanchard and Simon themselves identify some, but far from all, of the structural changes in the U.S. economy that plausibly have accounted for the most salient reductions in output volatility for reasons other than the response of monetary policy to inflation. For example, the sharp decline in

the volatility of homebuilding after the early 1980s nicely corresponds to the removal of Regulation Q interest ceilings (which they mention) and the development of the secondary market for home mortgages (which they do not). The steep decline in the volatility of consumer durables purchases after the Korean War similarly corresponds to the removal of Regulation W controls on consumer financing (which they also do not mention). Perhaps the most interesting single element of what the authors find is the dramatic change in inventory behavior, in the mid-1980s, from what used to be production *unsmoothing*—a phenomenon that runs counter to standard economic theory—to production smoothing. As they rightly point out, this development in particular has potentially important implications in the context of many other discussions not pursued here.

But this component-by-component dissection of real GDP growth also makes the pattern that Blanchard and Simon see in the data and choose to emphasize throughout the paper—a trend decline in output volatility throughout their half-century-long sample period, “temporarily interrupted in the 1970s and 1980s”—appear less satisfactory as a comprehensive description of what has happened. The volatility of spending on durable consumption and that of both residential and nonresidential investment initially declined but then began to increase again starting around 1965. Volatility of services consumption began a similar reversal in either 1967 or 1971, depending on how one reads the authors’ figure 10, reached a peak in 1975, then declined to an even lower level through 1990, and finally rose again through the mid-1990s. To be explicit, in each of these cases the midsample reversal predated the first OPEC oil price increase and the other familiar supply shocks of the 1970s. Of the five components of GDP shown in figure 10, only for nondurable consumption does the movement of volatility over time look much like a declining trend throughout, with an interruption beginning in the 1970s and peaking in the mid-1980s.

The subject of this paper is important. Within our lifetimes the death of the business cycle has been foretold as often as the coming of the messiah (perhaps because, in so many people’s minds, the two are identical). Every run of a few good years produces much talk, much of it from the business community and much of it fatuous, to the effect that new techniques of business management, or improved understanding on the part of policymakers, has relegated business cycles to the realm of historians. Perhaps some day this will come to pass. In the meanwhile, as the cir-

cumstances in which this panel met remind us, all it takes is a whiff of slowdown—companies that fail to meet overly optimistic earnings projections, or a down stock market, or an easing in the pace of industrial production or housing starts or retail sales—and all such talk is readily forgotten.

The more important questions, however—those that form the proper basis for economic inquiry, as in this paper—persist. The world does not always remain the same. Carefully documenting how it has changed, or is changing, is an important task for economists. So is seeking to understand the origins of those changes. The paper by Blanchard and Simon, especially in its highlighting of the relationship between declining output volatility and declining inflation volatility both in the United States and elsewhere, points toward a useful direction in which to look.

General discussion: Several participants discussed how much of the credit for decreased output volatility should be ascribed to better luck and how much to policy. Gregory Mankiw suggested that the behavior of food and energy prices during the 1970s was a piece of bad luck, which turned into good luck in the 1990s. He believed the volatility of food and energy prices relative to those of other goods, a measure of supply shocks, had decreased in the past ten or fifteen years and was responsible for some of the increase in stability. On the other hand, he acknowledged, a number of studies have found that interest rates are much more responsive to inflation than they were in the past, suggesting that better monetary policy should also get part of the credit.

Alan Blinder agreed that improved Federal Reserve behavior is a candidate for explaining the dramatic reduction in volatility in the second half of the 1980s. But, he suggested, this is not a fully satisfactory way to view the performance of policy. One could argue that then-Fed Chairman Paul Volcker's crackdown on inflation in the early 1980s was important in creating a more stable environment later in the decade, but it was also a large shock to output, increasing volatility in the short run. Blinder also found it hard to understand the precise timing of the abrupt decline in volatility.

Robert Gordon commented that the drop in volatility was likely to be in part an artifact of using rolling regressions to estimate the output process. A large fluctuation in output has a large effect on estimated volatility at the time it occurs and an opposite effect of the same magnitude when it drops out of the sample five years later. The decline in volatility looks like a

reflection of the dramatic inventory swing between the fourth quarter of 1982 and the first quarter of 1983. Hence not too much should be made of the precise timing of changes in volatility. William Brainard suggested that this problem could be avoided by using an explicitly time-varying parameter model rather than rolling regressions.

Gordon observed further that the suggestion of a causal relation running from inflation volatility to output volatility had a taste of *déjà vu*, recalling a famous *New York Times* headline of the mid-1970s: "Inflation triggers recession." We have learned that not all price shocks are adverse: some are benign, decreasing rather than raising volatility. In recent years we have experienced falling real import prices, an acceleration in the rate of decline of computer prices (and higher productivity growth), and, for a while, falling real energy and medical care prices. These have created an environment in which the Federal Reserve could refrain from its normal *réponse* to rapid output growth, which in past circumstances would have been to create a spike in short-term interest rates, bringing on the next recession.

Several panelists wondered how sensitive the authors' results were to particular observations or assumptions about the underlying processes. Gordon would have preferred to omit the Korean War episode, because consumption and government spending during that period were too peculiar (and highly correlated) to be of general significance. George Perry questioned the authors' decision to divide their sample period into halves and assume a constant rate of growth for each. He noted that many observers thought the underlying growth rate had varied significantly within each of the two subperiods. Since the probability of a recession and the expected length of an expansion are quite sensitive to this rate of growth, such a refinement might have resulted in a quite different picture. For example, it seems plausible that the growth rate at the end of the 1960s was relatively high; hence, according to the model, falling into a recession in 1970 was an extremely low probability event.

Edmund Phelps found anomalous the paper's conclusion that volatility is much lower than it used to be, given that we recently experienced one of the most powerful booms of the last 100 years. Robert Hall agreed: he found it hard to say we have been in a period of low volatility when recent five- and ten-year forecast errors have been huge. Hall concluded that the focus of the paper's analysis was on movements in real GDP of too high frequency, and he thought it would be more appropriate to look at the

medium-frequency movements. He also cautioned against identifying a decline in volatility on the basis of a small sample: after all, from 1875 to 1929 volatility had decreased before shooting up during the 1929 crisis. Robert Shimer suggested that inflation and output variability may have common determinants (such as the 1970s oil shock) and that it would have been more appropriate to treat both as dependent variables rather than try to explain one with the other.

Gordon and Hall expressed surprise at the insignificant effect of changes in the composition of output over a period that has seen a massive shift from factory work to desk jobs. Phelps would have liked to have seen an analysis of changes in the volatility of the employment (or the unemployment) rate, both for its intrinsic interest and as a possible explanation of changes in output volatility. He conjectured that employment has become a more sluggish variable, in part because of changes in the composition of output, in part because of institutional changes, and in part because of developments that we need to understand better, such as the fact that quit rates have decreased. Perry was interested in the results on inventory, where it appeared the correlation with output has reversed. If, as is widely believed, firms are getting much better at anticipating sales and adjusting production quickly, one might have expected the correlation to have become more positive.

Susan Collins thought that a more detailed look at the components of output (investment, inventory, and so forth) in European countries would be informative. She believed that a quite different picture would emerge from that in the United States in the aggregate.

References

- Balke, Nathan S., and Robert J. Gordon. 1989. "The Estimation of Prewar Gross National Product: Methodology and New Evidence." *Journal of Political Economy* 97(1): 38–92.
- Blanchard, Olivier, and Mark Watson. 1986. "Are Business Cycles All Alike?" In *The American Business Cycle: Continuity and Change*, edited by Robert J. Gordon. University of Chicago Press.
- Burns, Arthur F. 1960. "Progress Towards Economic Stability." *American Economic Review* 50(1): 1–19.
- Hamilton, James D. 1989. "A New Approach to the Economic Analysis of Non-stationary Time Series and the Business Cycle." *Econometrica* 57(2): 357–84.
- Kahn, James, Margaret M. McConnell, and Gabriel Perez-Quiros. 2001. "The Reduced Volatility of the U.S. Economy: Policy or Progress?" Unpublished paper. Federal Reserve Bank of New York (March).
- McConnell, Margaret M., and Gabriel Perez-Quiros. 2000. "Output Fluctuations in the United States: What Has Changed Since the Early 1980's?" *American Economic Review* 90(5): 1464–76.
- Perry, George L., and Charles L. Schultze. 1993. "Was This Recession Different? Are They All Different?" *BPEA*, 1:1993, 145–95.
- Romer, Christina D. 1986. "Is the Stabilization of the Postwar Economy a Figment of the Data?" *American Economic Review* 76(3): 314–34.
- Simon, John A. 2000. "Essays in Empirical Macroeconomics." Ph.D. dissertation, Massachusetts Institute of Technology.
- Taylor, John B. 2000. "Remarks for the Panel Discussion on 'Recent Changes in Trend and Cycle.'" Paper prepared for the conference on Structural Change and Monetary Policy, sponsored by the Federal Reserve Bank of San Francisco and the Stanford Institute for Economic Policy Research, March 3–4.
- Weir, David R. 1986. "The Reliability of Historical Macroeconomic Data for Comparing Cyclical Stability." *Journal of Economic History* 46(2): 323–65.

A vertical bar on the left side of the page, consisting of a series of horizontal segments in shades of yellow and orange, with a small red diamond at the top.

COPYRIGHT INFORMATION

TITLE: The long and large decline of U.S. output volatility
SOURCE: Brookings Papers on Economic Activity no1 2001
WN: 0100104142003

The magazine publisher is the copyright holder of this article and it is reproduced with permission. Further reproduction of this article in violation of the copyright is prohibited.

Copyright 1982-2001 The H.W. Wilson Company. All rights reserved.