

Interpreting the Great Moderation: Changes in the Volatility of Economic Activity at the Macro and Micro Levels

Steven J. Davis and James A. Kahn

Steven J. Davis is William H. Abbott Professor of International Business and Economics, Graduate School of Business, University of Chicago, Chicago, Illinois, and Research Associate, National Bureau of Economic Research, Cambridge, Massachusetts. James A. Kahn is Visiting Associate Professor, Stern School of Business, New York University, and Vice President, Macroeconomic and Monetary Studies Function, Federal Reserve Bank of New York, both in New York City, New York. Their e-mail addresses are <steve.davis@chicagosb.edu> and <James.Kahn@ny.frb.org>.

Abstract

We review evidence on the Great Moderation in conjunction with evidence about volatility trends at the micro level. We combine the two types of evidence to develop a tentative story for important components of the aggregate volatility decline and its consequences. The key ingredients are declines in firm-level volatility and aggregate volatility – most dramatically in the durable goods sector – but the absence of a decline in household consumption volatility and individual earnings uncertainty. Our explanation for the aggregate volatility decline stresses improved supply-chain management, particularly in the durable goods sector, and, less important, a shift in production and employment from goods to services. We provide evidence that better inventory control made a substantial contribution to declines in firm-level and aggregate volatility. Consistent with this view, if we look past the turbulent 1970s and early 1980s much of the moderation reflects a decline in high frequency (short-term) fluctuations. While these developments represent efficiency gains, they do not imply (nor is there evidence for) a reduction in economic uncertainty faced by individuals and households.

There has been a striking decline in the volatility of aggregate economic activity since the early 1980s in most advanced economies. Volatility reductions are evident for output and employment at the aggregate level and across most industrial sectors and expenditure categories. Inflation and inflation volatility have also declined dramatically over a similar time period. Previous studies offer several potential explanations for this “Great Moderation.” Some studies credit improved monetary policy for reductions in the volatility of real economic activity and inflation (for example, Clarida, et al., 2000). Others suggest that financial innovation and increased global integration play a role (Dyman et al., 2006). Still others, pointing to evidence that output volatility fell more than sales volatility, highlight the potential role of better inventory control methods (for example, Kahn et al. 2002). Another line of research (for example, Stock and Watson, 2002) stresses “good luck” in the form of smaller exogenous shocks.

These explanations are not mutually exclusive. As Bernanke (2004) remarks in his discussion of the Great Moderation: “Explanations of complicated phenomena are rarely clear cut and simple, and each ... probably contains elements of truth.” The main elements can also interact in complicated ways. Perhaps, for example, the unsuccessful monetary policy of the 1970s or the more successful policy that followed facilitated the spread of volatility-reducing financial innovations. Or perhaps sound monetary policy is easier when shocks are milder. Nonetheless, even if no single factor fully explains the phenomenon, it is useful to amass evidence for and against particular hypotheses.

We seek to address two fundamental questions about the Great Moderation: What are its causes? And does it matter for economic welfare? We consider a variety of evidence related to the Great Moderation, drawing mainly on U.S. data, and work towards a story with a few key themes. Unlike most research on the topic, ours considers volatility behavior at the micro level for clues about the sources and consequences of aggregate volatility changes. From a welfare perspective, a key issue is whether the Great Moderation led to lower consumption volatility and less economic uncertainty for individuals and households. It matters, for example, whether the Great Moderation mainly involves transient sources of income volatility or long-lasting shocks, since consumers can more readily buffer transient income movements. It also matters whether the decline in aggregate volatility translated into a decline in volatility at the micro level.

As it turns out, the micro story is complex. The average volatility of firm-level employment growth fell after the mid-1980s (Davis et al. 2006). At the individual level, several indicators point to a large decline in the risk of unwanted job loss after the early 1980s (Davis, 2008). Declines in firm-level volatility and job-loss risks fit naturally with the decline in aggregate volatility. However, when we consider household-level consumption changes, we find no evidence for a decline in volatility after 1980. The evidence on individual earnings uncertainty points to a longer term rise, not a decline. Reconciling the fuller set of facts about micro volatility trends with the aggregate volatility decline presents a significant challenge, which we meet only part way.

We begin with some facts about the Great Moderation, centered on the reduced volatility of real activity and the stabilization of inflation. Our interpretation of this evidence casts doubt on explanations for the Great Moderation that emphasize better monetary policy or financial innovation. Instead it points toward a technological story focused on the durable goods sector, and in particular on inventory investment. We then examine inventory behavior in more detail and provide additional evidence of improved supply-chain management and a specific story for how it can explain reduced volatility. Next, we consider whether evidence about volatility in micro data is consistent with our story. We present evidence on volatility in household-level consumption data and briefly discuss research on earnings and income uncertainty. In the end, we conclude that a substantial component of the Great Moderation reflects reductions in short-term volatility that have had little or no apparent impact on consumption risks at the household level.

We will not devote much space to the “good luck” hypothesis, even though a number of influential studies find support for it.¹ We interpret these findings as demonstrating the need for better models to explain the volatility reductions, or convincing evidence that measurable and plausibly exogenous, economic disturbances—oil supply disruptions, wars, or weather, for example—can explain observed declines in the volatility of economic activity. Along these lines, a recent study (Giannone, et al, 2008) argues that the econometric models that underlie the support for good luck hypothesis are too simple. When they examine more complex models with a larger

¹ See, for example, Stock and Watson (2002), Justiniano and Primiceri (2006), Ahmed et al. (2004).

number of variables, they find that the reduced volatility comes from a change in the propagation of shocks rather than in the size of the shocks.

There is also an older debate about whether economic volatility fell after World War II in comparison to the prewar era. That discussion focused on measurement issues and data quality. Romer (1986) argued that the apparent postwar stabilization was only a “figment of the data” resulting from improved statistical methods and more complete measurement in the postwar era. Our discussion will focus exclusively on postwar data where, subject to controlling for compositional changes, we believe the data permit meaningful comparisons over time.

Declines in the Volatility of Aggregate Economic Activity

An abrupt drop in the volatility of U.S. real GDP growth in the early 1980s (Figure 1) provided the initial impetus for research on the Great Moderation. Early findings of a discrete break in volatility around 1984 (Kim and Nelson, 1999; McConnell and Perez-Quiros, 2000) encouraged a focus on comparisons before and after 1984. This approach conceals the fact that many economic series did not undergo an abrupt volatility drop around 1984. Some did so much earlier, some later, and for some, volatility trended down rather than dropped.

We begin our exploration by decomposing real GDP growth and its volatility by what the National Income and Product Accounts (NIPA) call “major product” categories. Figure 2 shows annualized quarterly growth rates for the four major products that constitute GDP: nondurable goods, durable goods, services, and structures. Each sectoral growth rate is scaled by its share of nominal GDP so that we can readily assess its contribution to the aggregate growth rate in Figure 1.² These charts show that only the durable goods sector experienced a volatility change in much the same way—in terms of magnitude and timing—as GDP. Volatility of nondurables output fell after the early 1980s, but it was also lower in the 1960s before rising in the 1970s, and it was never nearly as high as the volatility of durables. Thus, the volatility decline for nondurables—such as it was—figures only modestly in the stabilization after the early 1980s.

² A small sector such as structures is very volatile, but contributes directly only modestly to overall GDP volatility.

Likewise, the service sector was never nearly as volatile as durable goods output; and moreover, its volatility fell in the early 1960s and again in the 1970s, long before the onset of the Great Moderation in the 1980s. Structures underwent a decline in output volatility similar in timing to that for overall GDP, but the size of the sector and its contribution to GDP volatility are modest.³ Thus, the drop in GDP volatility appears most closely related to developments in the durable goods sector.

Figure 3 demonstrates this point in a different way by plotting rolling five-year standard deviations of GDP growth, durables sector output growth, and, for comparison, the growth in inventory investment and residential investment.⁴ To gauge its contribution to GDP volatility, each series is scaled by its nominal share in total GDP. Inventory investment, which overlaps with durable goods output, displays nearly as large a drop in volatility as durable goods output despite having an average GDP share of less than one percent. Residential investment experienced a large drop in volatility beginning in the early 1980s, but its impact on overall GDP volatility is limited by its small share of aggregate activity (about five percent). Thus, the residential investment sector played a small role in the Great Moderation unless it had disproportionately large spillover effects on other sectors. Volatility declines in other components of GDP (not shown) are also relatively modest, and not necessarily synchronized with the Great Moderation (see Kahn, 2008).

Figure 3 also suggests that both for total GDP and especially for durables, the volatility decline in the early 1980s was an acceleration of a trend dating back to World War II. Indeed, Blanchard and Simon (2001) argue that the suddenness of the volatility drop is more apparent than real—that, in fact, large shocks in the 1970s and a deep contraction in the early 1980s obscure the longer term volatility decline that began well before the 1980s. The idea that the 1970s (really the period from 1970 to 1983) were exceptional will be a recurring theme in our discussion. But whereas Blanchard and Simon stress that volatility declines occurred across a broad range of GDP components,

³ The volatility contributions depicted in the charts are also affected by trends in sector shares over time, but the effect is very slight. The charts look virtually identical with sector shares held constant.

⁴ Our analysis will for the most part begin in 1954, so as to omit data from the Korean War period in which volatility was unusually high even for that era.

we argue that only the decline in the durable goods sector's volatility is comparable to that of overall GDP in trend, timing, and magnitude.

A related question is the extent to which the Great Moderation reflects a secular shift away from relatively volatile sectors. In particular, the rising GDP share of a low-volatility sector like services might be expected to moderate output volatility. It turns out, however, that the secular shift toward services plays a modest role in the reduction of overall output volatility. Specifically, using 1984:1 as a break point, we calculate that the volatility of quarterly GDP fell by 1.97 percentage points (standard deviations of annualized growth rates). If we reconstruct GDP growth by fixing output shares for the four major product categories at their 1959 values, we instead calculate that the volatility of quarterly GDP growth fell by 1.75 percentage points. According to this calculation, broad sectoral shifts account for about 12 percent of the long-term decline in GDP volatility.

What about shifts within major product categories? In light of Figures 2 and 3, it is natural to focus on shifts within the durable goods sector. For example, are shifts away from automobiles and toward electronic equipment partly responsible for the large decline in durables volatility? Disaggregated output data from the National Income and Product Accounts are available only annually, which throws out higher frequency volatility, and only back to 1977, which leaves only seven observations prior to the commonly used break point of 1984. A useful alternative is to construct output data from the Bureau of Economic Analysis (BEA) series on constant-dollar shipments and inventories, available monthly on a consistent basis from 1967 to 1997 for durable goods manufacturing and eleven industries at the two-digit Standard Industrial Classification (SIC) level. These data show a pattern of greatly reduced volatility after 1983, and they confirm that volatility varies greatly across the disaggregated industries. However, repeating the same type of exercise as before, we find that volatility in the durables sector fell slightly more on a fixed-weight basis (51 percent) than the 48 percent drop from 1967-83 to 1984-97 in the raw data. We conclude that long-term sectoral shifts contribute to, but are not the main force behind, the Great Moderation.⁵

⁵ Another way to cut the data is to decompose the production function into factor input and productivity components. Stiroh (2005) decomposes output growth into the growth of hours worked and labor

Another illuminating fact about the Great Moderation is that volatility fell substantially more, and earlier, for output than for final sales. In the durable goods sector, for example, comparing pre- and post-1983, the standard deviation of output growth declined from 17.8 percent to 7.7 percent, whereas the standard deviation of sales growth declined only from 10.3 to 8.4 percent.⁶ Since inventory investment makes up the difference between output and final sales, this fact implies a change in inventory behavior—either a reduction in the volatility of inventory investment or a change in the covariance between inventory investment and sales. By national income accounting conventions, the services and structures sectors do not carry inventories, so the source of the change in inventory behavior must by definition lie in the goods sector.⁷ We will return to this topic below when we discuss inventory behavior in more detail.

Finally, as discussed in the introduction, not all volatility matters equally. What is relevant for economic welfare is the impact on consumption. Households can save and borrow to buffer their consumption against income volatility, particularly if the income movements are transitory. Thus the volatility of consumption expenditure is much lower than that of GDP, especially if one removes expenditures on durables (which are really a form of investment). In fact, whereas the volatility of real GDP growth fell by 2.5 percentage points post-1983, from 4.5 percent to 2 percent, the volatility of real expenditures on non-durables and services fell by just 0.8 percentage points, from 2.1 percent to 1.3 percent.⁸ This suggests that the decline in GDP volatility may have had only a modest impact on welfare.

productivity. He finds that the volatility of both components fell, and that their covariance declined. Gali and Gambetti (2007) provide additional evidence of a major shift in the pattern of comovements among output, hours, and labor productivity. They interpret the evidence as favoring a “structural change” interpretation of the Great Moderation (rather than lower shock volatilities) in the form of more aggressively anti-inflationary monetary policy and reduced labor adjustment costs. In contrast, Arias et al (2007) argue that a decline in the volatility of the Solow residual points toward the reduction of “productivity-like” shocks in the context of a real business cycle model of fluctuations.

⁶ See McConnell and Perez-Quiros (2000), Blanchard and Simon (2001) and Kahn et al. (2002). McConnell and Perez-Quiros (2000) find evidence of a statistically significant break in the mid-1980s in durables output but not in final sales.

⁷ In structures, construction is counted as part of final output (see, for example, <http://www.census.gov/const/c30/definitions.pdf>). There is, however, evidence of a change in inventory behavior in residential construction, even though it is not treated as such in the National Income and Product Accounts. See Kahn (2000).

⁸ Constructed as a weighted average of the growth rates of expenditures on nondurables and expenditures on services; see Whelan (2000).

One possibility is that a substantial part of the output volatility decline has been of the transitory type. To examine this, we consider how output volatility is apportioned across different “frequencies”—fluctuations of different durations. We can think of fluctuations in economic activity as being composed of cycles of varying length: daily, weekly, yearly, or longer. The amplitude of fluctuations in activity between daytime and nighttime, or weekday to weekend, is considerably larger than that between business cycle peaks and troughs; yet far from harmful, these “high-frequency” fluctuations are usually seen as beneficial or benign. Even seasonal fluctuations are regarded as much less consequential for economic welfare in a modern economy, partly because of their relatively short duration and partly because of their predictable nature. Indeed, the predictable seasonal component is typically removed from economic time series via “seasonal adjustment” so as to isolate the more consequential, unpredictable part.

Fluctuations of a few years or longer, including the so-called business cycle, can have adverse welfare consequences because, first, they are much less predictable and, second, they are longer and cannot be relied on to average out. At a practical level, consumers have greater difficulty using savings to buffer expenditures against the uncertainty related to business cycles, as compared to seasonal cycles. Milton Friedman’s (1957) “permanent income” theory of consumption recognizes this point explicitly, predicting that consumers smooth spending relative to short-term or “transitory” fluctuations in income but not relative to “permanent” changes, by which he meant changes of at least several years duration.

To understand the frequency dimension of the Great Moderation, it is helpful to divide the post-Korean War era into three periods: the relatively tranquil period from 1954 to 1969, the turbulent period from 1970 to 1983, and the period from 1984 to the present. Figure 4 decomposes GDP volatility into a high frequency component (cycles lasting fewer than 12 quarters) and a “business cycle” frequency component (longer cycles) for each period.⁹ High-frequency volatility declined modestly from the first to

⁹ To obtain the decomposition in Figure 4, we start with quarterly observations on log real GDP from 1954:Q1 to 2004:Q3 and remove the low frequency component by applying a Hodrick-Prescott filter with smoothing parameter set to 1600. We then apply the band pass filter routine in the E-Views software package. (We use the Christiano and Fitzgerald, 1999, fixed-length symmetric filter with twelve leads and lags for Figure 4, but obtain similar results with other methods.) The frequency decomposition yields two

the second period, then fell sharply after 1983. This result suggests that the welfare consequences of the Great Moderation might be rather modest, because much of the Great Moderation reflects a drop in the type of short-term volatility that presents few difficulties for households in any event. Another piece of evidence pointing in this direction is the greatly diminished role of temporary layoffs in recessionary unemployment movements since the mid-1980s (Groschen and Potter, 2003). At the household level, consumption and living standards are probably less sensitive to temporary layoffs than to permanent job loss. Temporary layoffs are also naturally associated with output movements of short duration, such as inventory corrections.

Figure 4 also shows that business cycle volatility jumped sharply in the middle period and then returned to approximately the same level as in the 1954-69 period. Someone lumping the first two periods together might see a decline in both types of volatility in the post-1983 era, but that approach masks highly volatile business cycle fluctuations from 1970 to 1983 compared to before and after.¹⁰ We have repeated the exercise of Figure 4 for durable goods output and obtained a similar pattern. This result reinforces the view that developments in the durable goods sector play an important role in the Great Moderation.

Inflation and the Role of Monetary Policy

In addition to the apparent break in volatility around 1983, that year is also often viewed as a turning point in monetary policy: the beginning of the Volcker-Greenspan era of low and stable inflation. A large literature, led by Clarida et al. (2000), debates the extent to which there was a change in monetary policy regime, one that had the effect of reducing both inflation and output volatility. For it to do both requires not just a more aggressive anti-inflationary policy, which according to traditional policy analysis would actually increase output volatility, but a superior policy that improves the locus of the tradeoff.

time series, one for cycles longer than 12 quarters and one for shorter cycles. Lastly, we calculate the variance of each time series within each sub period and report the results in Figure 4.

¹⁰ Ahmed et al. (2004), for example, provide evidence that the volatility decline after 1983 was essentially uniform across all frequencies. Figure 4 does not contradict their finding, but it suggests that treating the 1970-83 period separately yields a different impression. Given the shortness of the time series relative to the periodicity of business cycles, it is hard to assess the most appropriate break points. For this reason, we view the low-frequency results in Figure 4 as suggestive only.

It is worth noting that low and stable inflation is not uniquely a post-1983 phenomenon. Inflation volatility from 1954 to 1969 is almost as low as after 1983. This does not necessarily refute the “better monetary policy” hypothesis, but it requires a more complicated story. How could a low-inflation policy in one era leave output volatility high, and then, some fifteen years later be responsible for dramatically reducing output volatility? One hypothesis that can reconcile these facts goes as follows: Monetary policy can reduce inflation volatility in two ways: 1) it can choose a different point in a fixed tradeoff between inflation volatility and output volatility; or 2) it can somehow get “better” (more efficient) and reduce both. The story would be that the 1955-69 period is an example of the former, while the post-1983 period is a case where policy became more efficient.

While one can tell such a story, there is little evidence to support it. Romer and Romer (2002a, 2002b) argue, in fact, that policy in the 1950s was similar to policy in the 1990s. In addition, the “better monetary policy” story would have to account for the range of facts about the Great Moderation described in the previous section, including the long downward trend in durables output volatility, and the changes in inventory behavior implied by the small drop in sales volatility compared to output volatility.

A different argument for the “better monetary policy” hypothesis is that it is the only story compatible with a discrete drop in volatility after the early 1980s. Improved inventory management is unlikely to have emerged suddenly. And financial innovation is most plausibly a gradual process, even if there were specific discrete events such as the demise of Regulation Q interest rate ceilings in the early 1980s (see, for example, Bernanke, 2007).¹¹

Monetary policy, on the other hand, can be subject to sudden changes in “regime.” In the last twenty years, for example, many central banks around the world adopted “inflation targeting” regimes that arguably involve a discrete departure from earlier policies. In October 1979, the Federal Reserve under Volcker shifted from targeting the federal funds rate to targeting nonborrowed reserves, a regime that lasted

¹¹ Interest rate ceilings on bank and thrift deposits made credit availability to the housing sector (which was especially dependent on this source of funds) highly sensitive to market interest rates and thus to contractionary monetary policy. As argued above, however, the structures sector, and residential construction in particular, was not big enough to drive large changes in the overall volatility of GDP.

until 1983. A number of studies (for example, Clarida et al, 2000) provide evidence that the interest rate targeting regimes pre-1979 and post-1983 were fundamentally different, the former resulting in both inflation and output instability.

Nonetheless, the evidence for important policy regime shifts is not clear cut. Sims and Zha (2006) argue that changes in monetary policy regimes were relatively inconsequential and in any case do not line up very well with changes in volatility. In a series of papers, Orphanides (2002, for example) argues that the policy regime of the 1970s was not fundamentally different from earlier policies, but the Fed was hit with large structural changes (a higher “natural” unemployment rate and lower trend productivity growth) for which it had limited and imperfect information in real time.

There are, moreover, good reasons to doubt the hypothesis that advances in the conduct of monetary policy are responsible for the discrete drop in volatility post-1983. Modern research on monetary policy points to a variety of factors that influence the efficacy of monetary policy. These include the credibility of the policymaker, transparency, and commitment to rules-oriented decision making. With hindsight, the Volcker-Greenspan era looks like a discrete break with the immediate past, but it is not plausible that the Fed achieved enhanced credibility overnight. In addition, increased transparency has been an evolutionary process. The Federal Open Market Committee (FOMC) began making public its interest rate target decisions only in 1994. It did not begin releasing statements explaining its policy decisions until 1998, and the informational content of these statements continues to evolve. So the contention that there was a discrete breakthrough for monetary policy in 1983 that explains the drop in output volatility after the early 1980s does not withstand close scrutiny.

International evidence also provides grounds for skepticism about the role of monetary policy. Since the early 1980s, most industrialized countries experienced reduced output and inflation volatility, but with no clear connection to changes in monetary policy. Some countries adopted the policy of “inflation targeting,” some did not, and among those that did, there was considerable variety in its implementation. Truman (2003) finds some evidence that inflation targeting countries experienced larger declines in output volatility, but there were differences in initial conditions between adopters and non-adopters. A more robust conclusion is that volatility dropped for nearly

all industrialized countries, but the evidence is mixed regarding any connection with the timing and specifics of monetary policy changes.

The 1970s experience points to a role for monetary policy in the Great Moderation, but not in the way usually emphasized by adherents of the “better monetary policy” explanation. The evidence described earlier points to an underlying downward trend for aggregate volatility in much of the postwar era, interrupted by the turbulent 1970’s, as argued by Blanchard and Simon (2001). This reading of the evidence suggests that policy mistakes during the 1970s raised volatility for a time. Ending those mistakes with a return to policies resembling those of the 1950s and early 1960s allowed volatility to return to its underlying postwar path. According to this view, the resumption of sound monetary policy led to a large drop in volatility by removing the role of policy mistakes in the 1970s, restoring an underlying postwar trend toward lower output volatility.

To summarize, the apparent suddenness of the drop in output volatility after the early 1980s does not persuade us that new-found skill in monetary policy making drove the Great Moderation. There is evidence that the Great Moderation began before the Greenspan-Volcker era, and in any case monetary policy did not experience a once-and-for-all change in 1983. The case is stronger that monetary policy mistakes or unusual shocks in the 1970s temporarily overwhelmed the forces behind the Great Moderation, until the economic environment or policy making reverted to something resembling earlier times. A persuasive case for a more positive role for monetary policy requires 1) showing how monetary policy in the post-1983 period is distinct from the policies of both the 1970s *and* earlier; and 2) a model that traces changes in sectoral volatility patterns and inventory behavior to changes in monetary policy.

Inventory Behavior

As discussed above, there have been significant changes in the behavior of aggregate inventory investment since the early 1980s. Here we focus on the durable goods sector, where the most dramatic declines in output volatility occurred and where we have also seen evidence of a change in inventory behavior.¹² Traditionally, the

¹² We refer the reader to Bills and Kahn (2000) and Kahn et al. (2002) for discussions of inventory behavior in the nondurable goods sector.

inventory literature focuses on disaggregated data, in particular on data for two-digit manufacturing industries, but aggregate data have distinct advantages for our purposes. Disaggregated data are potentially misleading, because it is impossible to judge whether measured changes in inventory behavior reflect genuine shifts or simply result from less meaningful relocation of inventories between sectors. For example, if manufacturers shift final goods inventories downstream to wholesalers and retailers, or materials inventories upstream to their suppliers, manufacturing inventories decline relative to their shipments. Yet that decline would be largely offset by inventory increases elsewhere in the economy, and a mere re-labeling would be misinterpreted as evidence of structural change.¹³

Because output is the sum of final sales and inventory investment, we can decompose the variance of output growth rates into components for sales growth and inventory investment, plus a covariance term.¹⁴ The covariance term reflects the extent to which changes in inventory investment are positively correlated with sales growth and thereby add to output volatility. This decomposition allows us to track the contributions of sales and inventory investment to output volatility over time. Since sales volatility declines little, as mentioned above, we focus on the other two terms of the decomposition. They are shown in Figure 5, which plots the rolling five-year variance of the inventory investment term and the sales-inventory covariance term for the durable goods sector. Both terms show substantial downward trends, with the covariance term accounting for a very big drop in output volatility after the early 1980s. The results displayed in Figure 5 motivate additional investigation of inventory behavior.

¹³ We should note that the data we use are not entirely immune from this criticism—for example, firms could shift materials inventories offshore to foreign suppliers.

¹⁴ Although inventory investment, because it can be negative, does not have a conventionally defined growth contribution, we can define it indirectly as the difference between the growth rate of output and the growth contribution of final sales (cf. Kahn et al., 2002). Following Whelan (2000), we approximate the latter in terms of the real growth rate of sales and the nominal share of sales in output. Letting γ_{xy} denote the growth contribution of x to output y , where $x = s$ for sales and $x = i$ for inventories, define the growth contribution of inventory investment as

$$\gamma_{iy} \equiv \gamma_{yy} - \gamma_{sy}$$

where $\gamma_{sy} = \gamma_{ss} \theta_{sy}$, θ_{sy} is the nominal share of s in y (measured as the average of current and lagged shares). The growth contribution of a variable to itself is just its growth rate.

We inspected the inventory-sales ratio for the durable goods sector using a data source that captures all durable goods inventories, including those situated outside the manufacturing sector. After peaking in the years 1982-83, the ratio trends downward for nearly two decades (the solid line in Figure 6). By itself, the downward trend in the inventory-sales ratio after the early 1980s is not proof of technical progress in inventory management. It could, instead, reflect compositional shifts or movement along a fixed technological tradeoff in response to rising inventory carrying costs.¹⁵ But the timing of the break in the trend is striking. The inventory-sales ratio has also become less volatile since the early 1980s, suggesting that businesses were hit by smaller shocks, made smaller mistakes or could more readily correct inventory imbalances. Again, this pattern is not conclusive, but it is consistent with the idea of improved inventory management technologies.

Kahn et al. (2002) describe results from forecasting equations with sales and inventories that point to another important difference before and after 1983. Before 1983, sales helped to forecast inventories more than inventories helped to forecast sales. After 1983, they were similar, reflecting both a decrease in the usefulness of sales in forecasting inventories and an increase in the usefulness of inventories in forecasting sales. Moreover, the variance of sales forecast errors dropped precipitously, meaning that less of the variation in sales was unpredicted given prior sales and inventories. These results are also consistent with the idea that firms became better able to anticipate sales and adjust inventories in advance.

Stock and Watson (2002) are skeptical of the role of inventory management in the Great Moderation. But they focus on four-quarter growth rates of economic time series, a transformation that essentially filters out the higher frequency volatility associated with inventory investment. Indeed, it should be no surprise that inventory behavior appears less important in explaining the (smaller) volatility declines of four-quarter growth rates. Still, looking at four-quarter growth rates might be a reasonable thing to do if changes in high frequency volatility were economically uninteresting or played a small role in the Great Moderation, but we have argued otherwise. Stock and Watson's other grounds for

¹⁵ See Maccini et al. (2004), who argue that inventory investment is responsive to very persistent changes in real interest rates.

skepticism—the fact that sales volatility declined, and that most inventories in manufacturing are raw materials or work-in-progress—are addressed in the next section.

A Theory of Improved Inventory Control

One approach to assessing the role of improved inventory control is to be agnostic about the deeper story for technological change and look instead for changes in parameters or shocks in a dynamic empirical model of inventory behavior. McCarthy and Zakrajsek (2007) apply this approach to industry-level manufacturing data, which makes it not directly comparable to the GDP-level and sector-level data we use.¹⁶ They investigate the extent to which changes in the behavior of manufacturing inventories after 1983 reflect better inventory management versus changes in what they call the “macro structure,” which includes shock magnitudes and the effects of monetary policy. They find evidence for both types of changes.

A second approach is to apply a specific optimizing model of improved inventory control as in Kahn et al. (2002), which draws on earlier work by Kahn (1987) and Bills and Kahn (2000). In this model, firms carry finished goods inventories to avoid stockouts in the face of uncertain demand, trading off the cost of forgone profits against the cost of carrying inventories. If demand is serially correlated, mistakes in forecasting sales get magnified in production volatility, so that it exceeds the volatility of sales. To see why, suppose sales are unexpectedly low. Not only does this leave the firm with more inventories than planned, but the firm now needs lower inventories than it previously thought, because the persistence of the sales shock means lower sales than earlier forecasts for some time. Consequently, the firm cuts production by more than the current sales shock.

Given this “stockout-avoidance” mechanism, suppose that better technology gives firms better information about future demand disturbances. Then firms make smaller errors in production decisions, and the output volatility induced by correcting those errors diminishes. In addition, firms hold fewer inventories because they now face less uncertainty.

¹⁶ The disaggregated manufacturing data are subject to the measurement concerns raised earlier—both the concerns about interactions with upstream and downstream activity, and also those about basing conclusions regarding the pre-1984 era on data that are primarily from the exceptionally volatile 1970s.

This story can account for reduced production volatility in absolute terms and relative to sales volatility, but it has important drawbacks as a full explanation for the changes in inventory, output, and sales behavior that we have documented. First, depending on when the firm receives information regarding demand shocks, the story has counterfactual implications for sales volatility. To see this point, suppose that better technology yields earlier information about a future demand increase, in time to build up inventories before the demand increase arrives. Then when the demand increase actually occurs, the firm draws down the inventories it built up in anticipation. This result is consistent with the more negative correlation of inventory investment and sales, but the advance information has the counterfactual implication of greater sales volatility. Why? Because the firm is better able to accommodate fluctuating demand without stocking out.

Second, the stockout-avoidance approach does not apply so obviously or directly to the durable goods sector, much of which is best characterized as production-to-order rather than production-to-stock. Third, as pointed out by Humphries et al. (2001) and many others, most inventories, particularly in durable goods, are of materials or works in process, not final goods. Finally, while there is much anecdotal evidence concerning technology improvements that provide better information about future sales, there is no direct evidence to assist in specifying a model. And, as this discussion suggests, the details matter.

Kahn (2008) develops a variation on the “better information” story that appears to address these problems. Firms are production-to-order, and must order materials at least one period in advance. For simplicity (to avoid having to track three different stocks), materials are immediately converted into works-in-process inventories. Stockouts occur in these inventories if they are insufficient to allow the firm to meet its (stochastic) orders, in which case the order gets added to the stock of unfilled orders.¹⁷ This setup is broadly consistent with the argument of Irvine and Schuh (2005), who attribute a

¹⁷ Of course, in reality, production can involve many more stages, and many suppliers—both internal and external to the firm—at different stages. Although the formal model depicts a single firm, it could easily be applied to a vertically nonintegrated supply chain, where the two stages of production are done by different firms. As with any model, the goal is to capture the essential features of reality while eschewing unneeded complexity.

substantial volatility decline to reduced co-movement between the manufacturing and trade sectors.

The intuition for how the model works is similar to that of the stockout-avoidance motive in Kahn (1987), described above, except that stockouts occur internally when the firm runs out of works-in-process inventories, rather than externally with finished goods stocks. Once again, because orders for final goods are serially correlated, errors in forecasting demand get magnified in the volatility of production (sales plus inventory investment). The key difference is that the firm produces “to order” so that the improved information is not translated so directly into the accommodation of demand shocks and greater sales volatility. Sales are shipments of completed production, so some of the reduced production volatility lessens the volatility of shipments, potentially offsetting the effect (still present here) from improved accommodation of demand shocks.

So this model immediately addresses all but one of the objections raised above. It is specifically tailored to characteristics of the durable goods sector in that it assumes production-to-order rather than production-to-stock. It features works-in-process rather than final goods inventories. It has the property that better information about final orders has the potential to reduce the volatility of output and sales, the former more than the latter.

The remaining issue is whether and how firms have become better able to forecast orders. Here we consider survey data obtained from the Institute of Supply Management on average lead times for orders of production materials. This is imperfect evidence and is not confined to the durable goods sector, but it is striking nonetheless. The series is depicted in Figure 6, plotted against the inventory-sales ratio for the durable goods sector. While the average lead-time series does not exhibit the underlying downward trend of the volatility series, it does feature a clear drop in level post-1983 relative to earlier. It also shows some elevation in the more volatile 1970s.

What is the connection between shorter lead times and better information? In the model of Kahn (2008), the firm bases its forecast of future final goods orders on the history of observed final goods orders. So the longer the firm can delay materials orders, the better handle it has on how much to order. Consequently, mistakes are smaller on average, *and* the firm carries lower average inventory holdings.

Of course, what allows for shorter lead times is not modeled but taken as direct evidence of technical progress. Given that the goal of the “just-in-time” approach is greater flexibility to reduce the need to carry large stocks, this is a reasonable interpretation, but it may be only part of the story. For example, some of the increased lead times in the 1970s could be caused by Nixon-era price controls, which created shortages and frictions in materials. It would be natural for firms concerned about an inability to obtain materials in a given time frame to order farther in advance. This observation does not negate the mechanism in the model, but it points out that something other than technical progress can yield sustained movements in average lead times. Perhaps price controls, or even the high inflation of the 1970s, disrupted market signals and caused some of the rise in lead times and volatility, but this hypothesis awaits further research.

There is little doubt that vast resources have been devoted to improving what is generally referred to as “supply-chain management.” How this translates into observable behavior and data is another question entirely. As Mentzer et al. (2001) write:

“Despite the popularity of the term Supply Chain Management, both in academia and practice, there remains considerable confusion as to its meaning. Some authors define SCM in operational terms involving the flow of materials and products, some view it as a management philosophy, and some view it in terms of a management process.”

They go on to define the term as “the systemic, strategic coordination of the traditional business functions and the tactics across these business functions within a particular company and across businesses within the supply chain, for the purposes of improving the long-term performance of the individual companies and the supply chain as a whole.” This definition is vague, but it clearly encompasses a number of specifics, notably what the authors refer to as a logistics system: “the total flow of materials, from the acquisition of raw materials to delivery of finished products, to the ultimate users, as well as the related counter-flows of information that both control and record material movement.”

The strong prediction of the model is that shorter lead times give rise to a reduction in output volatility and a somewhat smaller reduction in the volatility of final sales. While the magnitudes of these declines depend on specifics such as the average

inventory-sales ratio (which is endogenous in the model and depends on price-cost markups and inventory holding costs) and the ratios of inputs to gross output at each stage of production, the qualitative results require only some degree of persistence in final goods orders.

Changes in Volatility at the Micro Level

Thus far, our discussion suggests that much of the Great Moderation involves a reduction in the volatility of durable goods output, and much of that reflects significant changes in inventory dynamics. We turn now to volatility trends at the level of firms, individuals and households and ask whether the micro evidence sheds additional light on the Great Moderation and its consequences.

Firm-Level Volatility

Figure 7 plots the average volatility of firm-level employment in the United States from 1977 to 2001. Davis et al. (2006) construct the figure using annual Census Bureau data on domestic employment for all firms in the U.S. private sector. They calculate the rolling standard deviation of employment growth rates for each firm and then compute the employment-weighted mean across firms for each year.¹⁸ As shown in the figure, the average volatility of firm-level employment fell from 49 percent in 1978 and 53 percent in 1987 to 38 percent in 2001. Thus, firm-level volatility shows a substantial long-term decline, but the timing differs somewhat from the decline in aggregate output volatility. The trend decline in firm-level employment volatility holds across one-digit industry groups and is not especially concentrated in the goods-producing sector.

Figure 7 also highlights a striking difference in the volatility behavior of publicly traded and privately held firms.¹⁹ Volatility is high and declining for privately held firms, but low and rising for publicly traded firms. This “volatility convergence” between publicly traded and privately held firms helps in understanding other recent

¹⁸ The firm-level growth and variability measures used to construct Figure 7 are defined for all firms, including entrants and exits. See equation (6) in Davis et al. (2006).

¹⁹ Here, “publicly traded” refers to firms with equity securities listed on a stock exchange or traded in over-the-counter markets. As a practical matter, given the heavy reliance on COMPUSTAT for research in this area, “publicly traded” means firms that are also in COMPUSTAT, which has reasonably comprehensive coverage of publicly traded firms since the addition of NASDAQ listings in 1973. For publicly traded firms, it is possible to compute the volatility of sales and employment; they exhibit similar trends.

findings on trends in firm-level variability. Comin and Philippon (2005) and Comin and Mulani (2006), for example, document a large upward trend in the volatility of sales among publicly traded firms. An influential paper by Campbell et al. (2001) documents a large upward trend in the volatility of firm-level equity returns for U.S. common stocks. Specifically, they find that the variance of firm-level returns in daily data more than doubles from 1962 to 1997. They also show that the trend increase in return volatility reflects a rise in the volatility of the idiosyncratic, firm-specific component. These findings for publicly traded firms led some observers to conclude that overall firm-level volatility rose sharply in recent decades, in puzzling contrast to the large drop in aggregate volatility. Figure 7, however, shows otherwise.

Why does the volatility trend among publicly traded firms depart so sharply from the overall trend? At one level, the answer is simple: Publicly traded firms account for less than one-third of private-sector employment, so there is much room for the trend among publicly traded firms to depart from the overall trend. Digging deeper reveals another, more interesting answer: There was a pronounced shift in the economic selection process governing entry into the set of publicly traded firms, and this shift greatly affected volatility trends among publicly traded firms.

To see this point, it is important to first recognize the large influx of newly listed firms in the 1980s and 1990s. Fama and French (2004) report that the number of new listings (mostly IPOs) on major U.S. stock markets jumped from 156 per year in the 1973-1979 period to 549 per year in 1980-2001. Remarkably, about 10 percent of listed firms are new in each year from 1980 to 2001. Similarly, Davis et al. (2006) report that firms newly listed in the 1980s and 1990s account for about 40 percent of employment among all publicly traded firms by the late 1990s. So the influx of new listings in the 1980s and 1990s is large in number and rather quickly accounts for a large share of activity.

In addition, Fama and French (2004) provide evidence that new listings are riskier than seasoned public firms by a variety of measures *and* that they become increasingly risky relative to seasoned firms after 1979. Based on their review of the evidence, they conclude that the upsurge of new listings explains much of the trend increase in idiosyncratic stock return volatility documented by Campbell et al. (2001). They also

suggest that there was a decline in the cost of equity that allowed weaker firms and those with more distant payoffs to issue public equity. A more recent study by Brown and Kapadia (2007) reaches an even stronger conclusion. Using a regression methodology, they find that “there is generally no significant trend in idiosyncratic risk after accounting for the year a firm lists.” They also provide other evidence that firm-specific risks in the economy as a whole did not increase, even though the volatility of firm-level equity returns rose because of an influx of successively riskier cohorts.

Davis et al. (2006) obtain similar results for firm-level employment volatility. Using a regression framework, they find that simple cohort effects for the year of first listing account for two-thirds of the volatility rise among publicly traded firms from 1978 to 2001. In contrast, firm size, age and industry effects—separately and in combination—account for little of the volatility rise among publicly traded firms.

The Risk of Job Loss

As discussed in Davis (2008), a wide variety of labor market indicators point to a secular decline in the risk of job loss. These indicators include unemployment inflows by experienced workers in the Current Population Survey (CPS), the three-year job-loss rate in the CPS Displaced Worker Survey, several measures for the gross rate of job destruction at the level of individual employers, the number of workers involved in mass layoff events, and the number of new claims for unemployment insurance benefits. All of these indicators point to a secular decline in the risk of job loss, although the extent and timing of the decline differ among the indicators. One of the most striking declines occurs in weekly new claims for unemployment insurance benefits. This measure of unwanted job loss fell from .46 percent of nonfarm employment per week in 1970-1983 to .30 percent in 1984-2007 and .27 percent since 1993.

Similarly, CPS data show a dramatic decline in unemployment inflows as a percentage of employment since the early 1980s. Davis et al. (2007) provide evidence that about half of the long-term decline in unemployment inflow rates is explained by the reduction in the gross job destruction rates and the volatility of firm-level growth rates. Hence, their study provides evidence of a direct link between the secular decline in firm-level volatility in Figure 7 and the secular decline in the incidence of job loss.

Consumption Volatility and Income Uncertainty

In a simple permanent income model of consumption expenditures with stable preferences, a decline in the variability of persistent income innovations produces a decline in the average magnitude of consumption changes. The intuition is straightforward: smaller shocks to permanent income necessitate smaller adjustments to consumption. In light of this implication, we investigate whether the average magnitude of household-level consumption changes rose or fell over time. A rise indicates an increase in economic uncertainty for the average household. Conversely, if the Great Moderation, the decline in firm-level volatility, and the reduced risk of job loss led to a more tranquil economic environment for the average household, then permanent-income logic implies a decline in the average magnitude of household-level consumption adjustments.

We take a simple approach to this issue using quarterly data from the interview segment of the Consumer Expenditure Survey, which has a short panel component. To measure the magnitude of household-level consumption adjustments, we compute the absolute value of the log change in consumption expenditures for each household and then average over households.²⁰ This average value for the magnitude of household-level consumption changes is our measure of consumption volatility. To assess whether trends in consumption volatility differ between richer and poorer households, we sort households into deciles of predicted consumption before computing the consumption volatility.²¹ We sort the data this way to ensure that we do not overlook important changes in consumption volatility for certain groups, such as the lowest deciles, that we might miss by looking only at the overall average consumption volatility.

²⁰ We restrict attention to nondurable goods and services and consider six-month changes in expenditures per adult equivalent in the consumer unit. Results are similar for three-month and nine-month changes. Our measure of adult equivalents is 1.0 times the first adult plus 0.7 times each additional adult in the same consumer unit plus 0.5 times each child in the consumer unit.

²¹ We computed predicted consumption based on a regression of log expenditures per adult equivalent on sex of the household head, a quartic polynomial in the head's age, four educational attainment categories, marital status of the head, interview month, and employment status of the head and the head's spouse, if there is one. We perform this sort based on the first interview with consumption expenditures for the consumer unit

Figure 8 shows consumption volatility by decile of predicted consumption for two periods.²² There are two results. First, consumption volatility rises with the consumption level in a given time period, except at the low end of the distribution. Second, and more important for our purposes, there is no evidence for a decline in consumption volatility after the 1980s. In fact, Figure 8 points to a modest increase in household consumption volatility over most of the consumption distribution. This second result is quite surprising given the declines in aggregate and firm-level volatility and the reduced risk of job loss. Taken at face value, and seen through the lens of the permanent-income theory, the second result says that the Great Moderation failed to deliver more economic tranquility in the form of less consumption risk for the average household.

Given the surprising nature of the result in Figure 8, it would be reassuring to know that other consumption data tell a similar story. Gorbachev (2007) uses data on food expenditures in the Panel Study of Income Dynamics to estimate the volatility of household consumption after controlling for predictable variation associated with movements in real interest rates and changes in family structure. She finds that the volatility of household consumption expenditures increased over the period from 1970 to 2002. Thus, data from both the Consumer Expenditures Survey and the Panel Study of Income Dynamics point to a trend rise in consumption volatility.

Another approach to assessing trends in economic uncertainty at the level of households and individuals is to exploit income or earnings data. There is an enormous literature on earnings inequality and changes in the structure of wages, but few previous studies seek to quantify long term changes in earnings uncertainty from the vantage point of the individual. An exception is Cunha and Heckman (2007), who estimate the contribution of earnings uncertainty to the rise in earnings inequality. Their method uses data on schooling choices, in combination with data on earnings outcomes, to decompose the realized variance of earnings into a component that is predictable by individuals and an unpredictable component that is not. They estimate a substantial rise in the variance of the unforecastable component in the present value of earnings uncertainty among

²² By capturing the recessions of the early 1980s and 1990s in the first of the two periods shown in Figure 8, one might worry that we stack the deck in favor of finding a trend decline in consumption volatility. For this reason, the absence of such a trend is all the more surprising.

young workers. This result is broadly consistent with the evidence of increased consumption volatility.

A recent paper by Dynan et al. (2008) investigates the volatility of household income using data from the PSID for households whose heads are at least 25 years old and not yet retired. They find that the standard deviation of two-year percent changes in household-level income rose by a third from the early 1970s to the early 2000s, with trend increases in all major age and education groups. They trace the rise of household income volatility to a greater frequency of very large income changes. While greater income volatility does not imply greater income uncertainty, it is certainly suggestive of such a development. In this respect, the results in Dynan et al. (2008) lend credence to the evidence of rising consumption volatility.²³

Summary of Micro Volatility Evidence and Implications

The volatility of firm-level employment growth rates fell after the early to mid 1980s. The decline in average firm-level volatility is similar in magnitude to the decline in aggregate volatility, but the timing differs. Although we did not discuss it here, the volatility of state-level employment growth rates also fell after the 1980s (Carlino et al., 2007). Among publicly traded firms, the volatility in real activity and in equity returns rose sharply after the early 1980s. This rise in volatility among publicly traded firms is a striking phenomenon, but it mainly or entirely reflects shifts in the selection process governing which firms become public. Hence, considerable care is required when drawing inferences about the sources and nature of the Great Moderation from data on equity returns or from any data limited to publicly traded firms.

Declines in firm-level volatility and gross job-destruction rates are closely linked to declines in the risk of unwanted job loss, as reflected in sharply lower unemployment inflows after the early 1980s. In this respect, data on aggregate volatility, average firm-level volatility, job destruction rates and the incidence of unemployment all point to a

²³ Dynan et al. (2008) summarize many studies that investigate trends in earnings and income volatility or trends in estimated earnings uncertainty based on statistical models. Although the studies differ greatly in their particulars, none show a decline in income volatility or earnings uncertainty that even approximates the decline in aggregate volatility and the vast majority point to greater volatility or uncertainty.

much more quiescent economic environment since the early 1980s. In contrast, data on labor earnings, income and consumption expenditures do not conform to a story of greater tranquility and lower uncertainty at the household or individual level. Although there is much room for further research, the available evidence suggests at least a modest increase in individual and household economic uncertainty. Assuming that this assessment of trends in household-level consumption volatility and related measures of individual uncertainty holds up under further study, it highlights a puzzle that research on the Great Moderation has yet to confront. To wit: Why has the dramatic decline in the volatility of aggregate real activity, and the roughly coincident decline in firm-level volatility and job-loss rates, not translated into sizable reductions in earnings uncertainty and consumption volatility facing individuals and households?

We do not know the answer to this question, but we conjecture that greater flexibility in pay setting for workers played a role, possibly a major one. Greater pay flexibility is consistent with the rise in wage and earnings inequality in U.S. labor markets since 1980 and with increases in individual income volatility and earnings uncertainty. If these developments involve a rise in the variance of idiosyncratic permanent income shocks to households, then household consumption volatility also rises according to permanent income theory. Greater wage (and hours) flexibility also leads to smaller firm-level employment responses to idiosyncratic shocks and smaller aggregate responses to common shocks, because firms can respond by adjusting compensation rather than relying entirely on layoffs and hires. By the same logic, wage adjustments can substitute for unwanted job loss. So, at least in principle, greater wage flexibility offers a unified explanation for the rise in wage and earnings inequality, flat or rising volatility in household consumption, a decline in job-loss rates, and declines in firm-level and aggregate volatility measures. Sources of greater pay flexibility include the decline of real minimum wages, a diminished role for private sector unionism and collective bargaining, intensified competitive pressures that undermined rigid compensation structures, the growth of employee leasing and temp workers, and the erosion of norms that formerly restrained wage differentials and prevented wage cuts.

Concluding Remarks

We summarize the main elements of our evidence and analysis. First, macroeconomic volatility has generally trended downward in the postwar United States (and in many other countries as well), interrupted by heightened turbulence in the 1970s and early 1980s. Otherwise, however, much of the Great Moderation reflects a decline in the high frequency component of aggregate output volatility. Second, declining volatility in the durable goods sector is a major contributor to declines in U.S. aggregate output volatility. Third, changes in inventory behavior appear to play a major role in the decline of output volatility in the goods-producing sector and, hence, in the economy as a whole. Fourth, the broad decline in aggregate volatility is mirrored by declines in firm-level employment volatility and in job-loss risks for workers. Lastly, declining volatility is not evident in microeconomic data on wages, incomes, or consumption expenditures. If anything, micro evidence on consumption volatility and other measures of individual uncertainty points in the opposite direction.

From this configuration of findings, we conclude that the welfare implications of the Great Moderation are subtler than one might think. While the benefits of reduced inflation uncertainty are well understood, the benefits on the “real” side are elusive. It is likely that reduced volatility at the firm level reduces production costs, a first-order welfare benefit. But the fact that it has not coincided with reduced consumption or income volatility for individuals and households suggests that economically meaningful uncertainty—the kind that affects welfare—does not have a simple and direct connection to aggregate output volatility.

Of course, more research and better data may alter our perception of the facts. Consumption at the individual and household level is poorly measured, and it may be that noisy data or idiosyncratic risks swamp the effect of lower aggregate volatility in micro data. But, at this juncture, the weight of the evidence points toward a different conclusion: The Great Moderation brought few benefits in the form of lower consumption volatility or reduced economic uncertainty for individuals and households.

Acknowledgements

Doug Elmendorf, Nick Bloom, and other participants at the San Francisco Fed's CSIP conference offered helpful comments, as did Charles Steindel. The editors of this journal – Jim Hines, Andrei Shleifer, Jeremy Stein and Timothy Taylor – provided many constructive comments that led to significant improvements in the paper. Parts of the paper also owe a significant debt to discussions with Meg McConnell. Responsibility for any flaws in the final product remains with the authors. The views expressed in this article are those of the authors and do not necessarily reflect the position of the Federal Reserve Bank of New York or the Federal Reserve System.

References

- Ahmed, Shaghil, Andrew Levin, and Beth Anne Wilson, 2004. "Recent U.S. Macroeconomic Stability: Good Policies, Good Practices, or Good Luck?" *Review of Economics and Statistics* 86, 824-832.
- Arias, Andre, Gary Hansen, and Lee Ohanian, 2007, "Why Have Business Cycle Fluctuations Become Less Volatile?" *Economic Theory*, 32, 43-58.
- Bernanke, Ben, 2003. "'Constrained Discretion' and Monetary Policy," before the Money Marketeers of New York University, New York, New York, February 3.
- Bernanke, Ben, 2004. "The Great Moderation." At the meetings of the Eastern Economic Association, Washington, D.C., February 20.
- Bernanke, Ben, 2007. "Housing, Housing Finance, and Monetary Policy," at the Federal Reserve Bank of Kansas City's Economic Symposium, Jackson Hole, Wyoming, August 31.
- Bils, Mark., and James A. Kahn, 2000, "What Inventory Behavior Tells Us about Business Cycles," *American Economic Review* 90, 458-481.
- Blanchard, Olivier, and John Simon, 2001. "The Long and Large Decline in U.S. Output Volatility," *Brookings Papers on Economic Activity*, 1, pp. 135-64.
- Brown, Gregory and Nishad Kapadia, 2007, "Firm-Specific Risk and Equity Market Development," *Journal of Financial Economics*, 84, pp. 358-388.
- Campbell, John Y., Martin Lettau, Burton Malkiel and Y. Xu, 2001, "Have Individual Stocks Become More Volatile? An Exploration of Idiosyncratic Risk," *Journal of Finance*, 56, no. 1, 1-43.
- Carlino, Gerald, Robert DeFina and Keith Sill, 2007 "The Long and Large Decline in State Employment Growth Volatility," Working Paper no. 07-11, Research Department, Federal Reserve Bank of Philadelphia.
- Christiano, L. J. and T.J. Fitzgerald, 2003, "The Band Pass Filter," *International Economic Review* 44(2), 435-465.
- Clarida, Richard, Jordi Gali, and Mark Gertler, 2000. "Monetary Policy Rules and Macroeconomic Stability: Evidence and Some Theory," *Quarterly Journal of Economics*, 115, pp. 147-80.

- Comin, Diego and Sunil Mulani, 2006, "Diverging Trends in Aggregate and Firm Volatility," *Review of Economics and Statistics*, 88, no. 2, 374-383.
- Comin, Diego, and Thomas Philippon. 2005. "The Rise in Firm-Level Volatility: Causes and Consequences." *NBER Macroeconomics Annual 2005*, M. Gertler and K. Orgoff, eds. MIT Press.
- Cunha, Flavio and James J. Heckman, 2007, "The Evolution of Inequality, Heterogeneity and Uncertainty in Labor Earnings in the U.S. Economy," NBER working paper no. 13526.
- Davis, Steven J., 2008, "The Decline of Job Loss and Why It Matters," *American Economic Review: Papers and Proceedings*, 98, no. 2.
- Davis, Steven J., R. Jason Faberman, John Haltiwanger, Ron Jarmin and Javier Miranda, 2007, "Business Volatility, Job Destruction and Unemployment," working paper, October.
- Davis, Steven, J., John Haltiwanger, Ron Jarmin and Javier Miranda, "Volatility and Dispersion in Business Growth Rates: Publicly Traded versus Privately Held Firms," *NBER Macroeconomics Annual 2006*, Daron Acemoglu, Kenneth Rogoff, and Michael Woodford, eds. MIT Press.
- Dynan, Karen E. Douglas W. Elmendorf, and Daniel E. Sichel, 2006, "Can Financial Innovation Help Explain the Reduced Volatility of Economic Activity," *Journal of Monetary Economics* 53, 2006, pp. 123-150.
- Dynan, Karen E. Douglas W. Elmendorf, and Daniel E. Sichel, 2008, "The Evolution of Household Income Volatility," The Brookings Institution. Downloadable at http://www.brookings.edu/papers/2008/02_us_economics_elmendorf.aspx.
- Fama, Eugene F. and Kenneth R. French, 2004, "New Lists: Fundamentals and Survival Rates," *Journal of Financial Economics*, 73, no. 2, 229-269.
- Friedman, M., *A Theory of the Consumption Function*, National Bureau of Economic Research, Princeton, N.J., 1957.
- Gali, Jordi, and Gambetti, Luca, 2007, "On the Sources of the Great Moderation," Working Paper no. 1041, Universitat Pompeu Fabra.
- Giannone, D., M. Lenza, and L. Reichlin, 2008, "Explaining the Great Moderation: It is not the Shocks," European Central Bank, Working Paper Series No. 865, February.
- Gorbachev, Olga, 2007, "Did Household Consumption Become More Volatile?" working paper, available at <http://www.columbia.edu/~ocg2001/JMP.pdf>.

- Groshen, Erica, and Simon Potter, 2003, "Has Structural Change Contributed to a Jobless Recovery?" *Current Issues in Economics and Finance* 9 (August), Federal Reserve Bank of New York.
- Humphries, Brad, Louis J. Maccini, and Scott Schuh, 2001, "Input and Output Inventories," *Journal of Monetary Economics* 47, 347-375.
- Irvine, F. Owen, and Scott Schuh, 2005, "The Roles of Comovement and Inventory Investment in the Reduction of Output Volatility," Federal Reserve Bank of Boston Working Paper 05-9.
- Justiniano, A., and G. Primiceri, 2006, "The Time-Varying Volatility of Macroeconomic Fluctuations," NBER Working Paper No. 12022, February 2006.
- Kahn, James A., 1987, "Inventories and the Volatility of Production," *American Economic Review* 77 (September), 667-679.
- Kahn, James A., 2000, "Explaining the Gap between New Home Sales and Inventories," Federal Reserve Bank of New York, *Current Issues in Economics and Finance*, 6 (May).
- Kahn, James A., 2008, "Durable Goods Inventories and the Great Moderation," Federal Reserve Bank of New York, Staff Report No. 3xx.
- Kahn, James A., Margaret McConnell, and Gabriel Perez-Quiros, 2002. "On the Causes of the Increased Stability of the U.S. Economy," Federal Reserve Bank of New York, *Economic Policy Review*, 8, pp. 183-202.
- Kim, Chang-Jin, and Charles Nelson, 1999. "Has the U.S. Economy Become More Stable? A Bayesian Approach Based on a Markov-Switching Model of the Business Cycle," *Review of Economics and Statistics*, 81, pp. 608-16.
- Kim, Chang-Jin, Charles Nelson, and Jeremy Piger, 2003. "The Less Volatile U.S. Economy: A Bayesian Investigation of Timing, Breadth, and Potential Explanations," *Journal of Business and Economic Statistics*, forthcoming.
- Maccini, L.J., B. J. Moore and H. Schaller, 2004, "The Interest Rate, Learning, and Inventory Investment," *The American Economic Review* 94(5), December 2004, 1303-1327.
- McCarthy, Jonathan, and Egon Zakrajsek, 2007. "Inventory Dynamics and Business Cycles: What Has Changed?" *Journal of Money, Credit and Banking*, Vol. 39, No. 2-3 (March-April), 591-613.


- McConnell, Margaret, and Gabriel Perez-Quiros, 2000. "Output Fluctuations in the United States: What Has Changed since the Early 1980s?" *American Economic Review*, 90, pp. 1464-76.
- Mentzer, John, W. DeWitt, J. Keebler, S. Min, N. Mix, C. Smith, and Z. Zacharia, 2001, "Defining Supply Chain Management" *Journal of Business Logistics*.
- Nelson, Edward, 2004. "The Great Inflation of the Seventies: What Really Happened?" Federal Reserve Bank of St. Louis, working paper (January).
- Orphanides, Athanasios, 2002 "Monetary-Policy Rules and the Great Inflation," *The American Economic Review*, 92 (May), Papers and Proceedings, pp. 115-120.
- Orphanides, Athanasios, 2003. "The Quest for Prosperity without Inflation," *Journal of Monetary Economics*, 50 (April), pp. 633-63.
- Orphanides, Athanasios, and John C. Williams, 2003a. "Imperfect Knowledge, Inflation Expectations, and Monetary Policy," forthcoming in Ben Bernanke and Michael Woodford, eds., *Inflation Targeting*, Chicago: University of Chicago Press for NBER.
- Rogoff, Kenneth, 2006, "The Impact of Globalization on Monetary Policy," prepared for the Federal Reserve Bank of Kansas City symposium, "The New Economic Geography: Effects and Policy Implications," Jackson Hole, Wyoming, August 24-26.
- Romer, Christina, 1986, "Is the Stabilization of the Postwar Economy a Figment of the Data?" *American Economic Review* 76 (June), 314-334.
- Romer, Christina, and David Romer, 2002a. "A Rehabilitation of Monetary Policy in the 1950's," *American Economic Review (Papers and Proceedings)* 92 (May), pp. 121-127.
- Romer, Christina, and David Romer, 2002b. "The Evolution of Economic Understanding and Postwar Stabilization Policy," *Rethinking Stabilization Policy*, Federal Reserve Bank of Kansas City, pp. 11-78.
- Sims, Christopher, and Tao Zha, 2006, "Were There Regime Switches in U.S. Monetary Policy?" *American Economic Review*, 96 (March), pp. 54-81.
- Stiroh, Kevin, 2005, "Volatility Accounting: A Production View of Increased Economic Stability," working paper, Federal Reserve Bank of New York.
- Stock, James H. and Mark W. Watson, 2002, "Has the Business Cycle Changed and Why?" in Mark Gertler and Kenneth Rogoff, eds., *NBER Macroeconomics Annual 2002*. Cambridge, MA: MIT Press, 159-218.

Truman, Edwin M., 2003. *Inflation Targeting in the World Economy*. Institute for International Economics, Washington, D.C.

Warnock, M.V. Caddac, and Francis Warnock, 2000. "The Declining Volatility of U.S. Employment: Was Arthur Burns Right?" Board of Governors of the Federal Reserve System, International Finance Discussion Paper no. 677 (August).

Whelan, Karl, 2000, "A Guide to the Use of Chain Aggregated NIPA Data," Division of Research and Statistics Discussion Paper, Federal Reserve Board.


Figure 1: GDP Growth, 1947-2007



Quarterly, Annual Rate. Source: NIPA


Note: Shaded periods represent NBER-designated recessions

Figure 2: GDP Growth Contributions by Major Product, 1947-2007



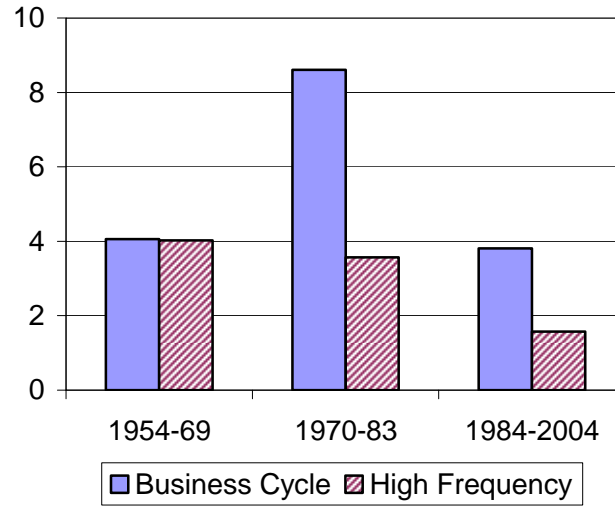
Quarterly, Annual Rate. Source: NIPA
Note: Growth contributions are growth rates scaled by the product category's share in total GDP. Shaded periods represent NBER-designated recessions

Figure 3: Volatility* over Time in Key Categories




*5-year rolling standard deviations of quarterly annualized growth rates
**Inventory investment growth contributions were computed by subtracting fixed investment's contribution from total investment's.

Figure 4: GDP Volatility by Frequency



Note: Units are standard deviations of annualized growth rates.

Figure 5: Contributions to Durables Output Variability



5-year rolling variances of growth contributions

Figure 6: Inventory-Sales Ratios in Durable Goods in Relation to Lead Times for Materials Orders in Manufacturing





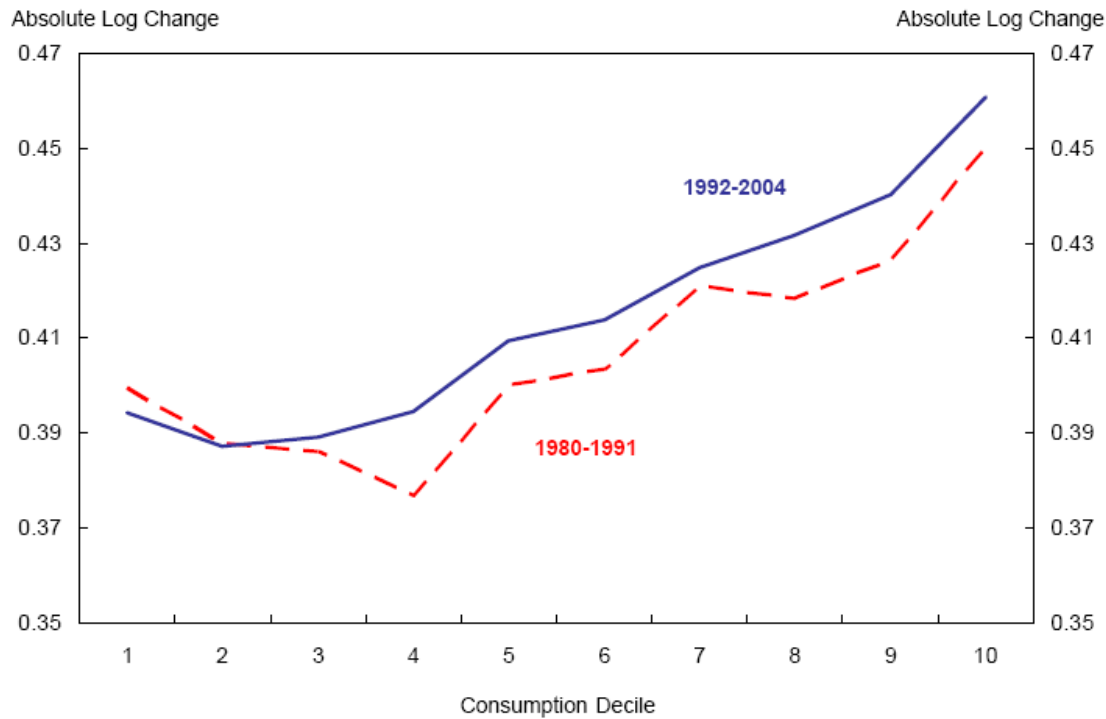
Figure 7: Volatility in Firm-Level Employment Growth Rates, Overall and by Ownership Status, U.S. Private Sector, 1977 to 2001



Notes: Firm-level volatility calculated as a ten-year weighted moving average of growth rates, inclusive of entry and exit and with a degrees-of-freedom correction. See equation (6) in Davis et al. (2006). Average volatility across firms computed on an employment-weighted basis.

Source: Calculations on the Longitudinal Business Database by Davis et al. (2006).

Figure 8: Household-Level Consumption Volatility by Deciles of Predicted Consumption



Note: We compute the absolute value of six-month log changes in expenditures on nondurable goods and services per adult equivalent in each household. Averaging the absolute changes by time period and decile yields the reported measure of consumption volatility.