

What I Learned at the Fed

*Laurence H. Meyer**

Introduction

I always judge my productivity during a particular phase of my career in terms of what I learned, not in terms of what I contributed. By that standard, I was very productive at the Fed. My talk today is about what I learned.

I first met Chuck Freedman during an earlier learning experience. Chuck was one of the first classmates I met in graduate school at MIT. My first real experience with Chuck, the student, took place on the second day of class with Charles Kindleberger during our first semester. At the end of the class, Chuck asked Professor Kindleberger if we could talk about the seven errors Chuck had found in the first two chapters of Kindleberger's textbook. I was horrified. By my naïve perspective, there were, by definition, no errors in textbooks. I had read those chapters with the sole intent of learning whatever it was that the esteemed professor wanted to teach. I was somewhat demoralized, but I recovered quickly and decided I should get to know Chuck better. Throughout our years at MIT, I leaned on Chuck more than once to explain some obscure theoretical derivation or help me understand the broader vision behind the complicated mathematics, and my learning experience was much enriched and made more enjoyable by my interaction and friendship with him.

After MIT, we took somewhat similar paths. We both started as academics, though I stayed longer, and we both ended up at central banks, though, in

* I would like to thank my former colleagues on the Federal Reserve Board of Governors and, especially, the research staff at the Board, whom I learned from and with during my tenure on the Board.

this case, Chuck stayed longer. I expect we may have had similar experiences as central bankers. I loved the opportunity to do public policy without the downside of partisan politics. Independent central banking is truly one of the great inventions of all times. But I also loved the intellectual atmosphere, the interaction with an extraordinary staff, and the opportunity to work with my colleagues on the Board and the Federal Open Market Committee (FOMC) to translate that wisdom into good policy. I trust and hope that Chuck had as rich and rewarding an experience at the Bank of Canada as I had during my more limited time at the Fed.

During the period that we were both at central banks, there were plenty of surprises, including many challenges to the models that we had brought with us from MIT and beyond, and we therefore had to continue our exercise in lifetime learning.

To be sure, all governors have a great deal of learning to do when they join the Federal Reserve because of the very wide range of responsibilities there, including bank supervision and regulation, consumer protection, and community affairs. I had a pretty steep learning curve to climb on the regulatory side and really enjoyed the investment required to fulfill my responsibilities, after a time, as oversight governor for bank supervision and regulation. But I'll focus my attention here on what I learned about how the economy works and how monetary policy should be conducted. In this case, I thought I already knew what I would need to know to be an effective central banker. I soon came to appreciate the importance of learning on the job.

1 Point of Departure: What I Thought I Knew

I came to the Board with a very structured view of how output, inflation, interest rates, and so on get determined. This reflected my long years of teaching macroeconomics, writing a textbook on this subject, my upbringing as a research assistant on the MPS (MIT-University of Pennsylvania) model project, and my experience building and using a large-scale macro-model at my forecasting firm.

One of the favourite papers that I wrote while on the Board, "Start with a Paradigm, End with a Story," is a defence of model-based forecasting and also provided an opportunity for me to lay out my view of how the macro-economy works. Here are the key features of the paradigm I described in that paper.

- (i) Output is demand determined in the short run. This is due, of course, to sticky prices. The result is that short-run output dynamics look a lot like the old-style multiplier-accelerator, IS-LM, and MPS model stories. In addition, short-run dynamics are driven predominantly by demand shocks.

- (ii) Output is supply determined in the long run. That is, output in the long run is determined by population and structural productivity, neither of which will be much, if at all, influenced by monetary policy. As a result, money is neutral in the long run. Neutrality, and the absence of money illusion, ensure that the rate of inflation, in the long run, is principally determined by monetary policy—either by money supply growth in traditional macro-models, or by the inflation objective incorporated into monetary policy rules in more modern models.
- (iii) The Phillips curve, as a model of inflation dynamics, connects the short run and the long run. That is, price flexibility is the stabilizing mechanism and the Phillips curve determines how rapidly this mechanism works. The Phillips curve also identifies the supply-side limits that, if overtaxed, result in overheating and higher inflation. The key parameters that pin down the supply-side limits are the non-accelerating-inflation rate of unemployment (NAIRU) and structural productivity growth.

In this paradigm, inflation is principally a monetary phenomenon, though the proximate source of increases in inflation is excess demand in labour and output markets. Fiscal and monetary policies in principle can both play roles in output stabilization, but in practice the role has been ceded to monetary policy, at least most of the time since the early 1980s. Fiscal policy has a role in affecting the level of potential output and the intermediate-term rate of economic growth.

In terms of monetary policy, I believed that, under an interest rate regime, there was no alternative to activist policy and that the Taylor rule summarized a well-designed strategy for monetary policy under such a regime. There is, in my view, a fundamental difference between an interest rate regime and a monetary aggregate regime. In the latter case, policy-makers could, in principle, set a monetary growth rate target and walk away. Such a “constant money-growth rule” might not yield an optimal outcome, but it would still have stabilizing properties and would generally keep you out of serious difficulties. In an interest rate regime, on the other hand, the failure to adjust interest rates in a timely fashion in response to changes in output and inflation could lead to destabilizing and open-ended errors.

The Taylor rule seemed to me to be a natural framework for monetary policy strategy. To begin with, it is consistent with the Fed’s dual mandate of promoting full employment and price stability. In addition, the prescription to adjust the real funds rate in response to deviations of output and inflation from their respective targets seemed like an efficient and obvious way for monetary policy makers to pursue their objectives. I appreciated that there would be challenges in forecasting the economy, and that the effectiveness of monetary policy would be limited by lags and uncertainty about the

structure of the economy, but still, the conduct of monetary policy didn't seem very complicated. I would leave with a much refined appreciation of the complications of conducting monetary policy, with a keener understanding of the challenges arising from uncertainties about the measurement of potential output and the NAIRU, the difficulty of identifying in real time structural changes that might be significantly affecting economic outcomes, the complications related to the zero nominal bound, the asymmetric risks associated with deflation, and the related uncertainties about output and inflation dynamics.

2 What I Learned About Macro

From virtually the moment I arrived at the Fed, I was confronted by a series of surprises that challenged the paradigm and required me to reassess the appropriate conduct of monetary policy in light of the uncertainties associated with the tension between the paradigm and recent experience. The unexpected developments included, most importantly, an acceleration in productivity, a probable decline in the NAIRU, and speculative excess. While I will focus on these surprises, I acknowledge that there were many more, including Y2K, 9/11, and the Asian financial crisis. Let me clarify with respect to Y2K that I suppose I knew there was going to be a century date change, but before joining the Board I did not appreciate that this was an issue for monetary policy makers or bank supervisors.

Focusing on the productivity acceleration, the possible decline in the NAIRU, and speculative excess, we found that one of the challenges was that these shocks did not come with labels attached. Instead, we observed, for a period, inexplicably exceptional economic performance and had to infer from the experience the shocks that were responsible. The questions I had to wrestle with were whether and how the paradigm could and should be adapted to be consistent with the unexpected developments and how monetary policy should be adapted to the higher order uncertainty about the NAIRU, productivity growth, the fundamental value of equities, and, as a result, about inflation and output dynamics. In the end, I believe that the paradigm has a basic core of validity, but it is incomplete, even for the narrow analysis of the strategy of monetary policy. In my view, the greatest flexibility had to be exercised during this period in adjusting the parameters of the old paradigm to fit the new experience and in enriching the framework where it was most seriously incomplete, specifically in the relationship between productivity and the NAIRU and with respect to the implications of speculative excess.

2.1 Productivity and the NAIRU

The major challenge from the beginning was that wages, more generally nominal labour compensation, while creeping upwards, were not rising as rapidly as the Phillips curve models were predicting. And price inflation was steady to even declining, in spite of the upward creep in labour compensation and the progressively tighter labour market.

From the beginning, these developments offered hints that the NAIRU might have declined and that productivity might have accelerated. But for a couple of years, the aggregate data for productivity gave very little hint of an acceleration. Indeed, for about a year, if anything, the data were pointing to a possible deceleration. The absence of confirmation in the aggregate data did not stop the Chairman from insisting that there was a productivity acceleration underway. He told us to ignore the data, but I remained skeptical. It was not until the summer of 1999 that revisions to the data made it clear—even to me—that there had been a sizable acceleration in productivity, and the revisions in the summer of 2000 reinforced and extended that story.

So one of the lessons of the experience in the second half of the 1990s is how difficult it can be to identify important structural changes in a timely manner. Monetary policy makers cannot afford to behave like economic historians or classical econometricians. They cannot wait to make their call about a potential structural break until all the dust clears or until there is enough data to do econometric tests for structural breaks. They have to make policy in real time and make judgments and act on those judgments before the truth is fully revealed, if it ever is. Even so, policy-makers need to see at least some hints of structural change in the underlying macro data before they rely on that structural change in their interpretation of the outlook, in their forecasts, and in their perspective on the appropriate conduct of monetary policy. Another lesson is that it pays to listen to those in the trenches, rather than only to those in the ivory towers. Business leaders were championing the notion of a productivity acceleration and its implication for their cost of production.

There is a relatively higher appreciation for the value of anecdotal information in the interpretation of the outlook and in the forecasting process at the Board and FOMC, compared with other forecasters and policy-makers. It reminds me of a remark, attributed to George Stigler: “Data is just the plural of anecdote.” The Fed’s advantage is that firms will provide the Reserve Bank presidents, the Chairman, and staff with information that they would not likely share with others. Furthermore, given that the role of the presidents is to bring to FOMC meetings information about business

attitudes and plans, a comparative advantage might well have evolved over the years in both acquiring and interpreting this information.

For a while, most of the attention focused on the possibility of a decline in the NAIRU and the contribution of a coincidence of favourable supply shocks to the exceptional economic performance. In mid-1997, in response to repeated assertions by the Chairman that the productivity acceleration was already underway, the staff decided to prepare for a possible occurrence by running a set of simulations to understand the macroeconomic effects of a productivity acceleration and to learn how monetary policy should respond to such an event.

A productivity acceleration, of course, raises the rate of maximum sustainable growth. But since growth is not itself an objective of monetary policy, this did not seem to be especially relevant to monetary policy makers. They could just sit back and enjoy the faster rate of growth now sustainable at full employment. The key objective of the staff study was to understand the complicated short-run dynamics affecting aggregate demand and inflation set in motion by a productivity acceleration.

The staff found that a productivity acceleration would likely result in booms in both consumer and business spending. Consumer spending would be boosted by the wealth effect associated with higher equity prices and expectations of future faster growth in income. Business investment spending would be boosted by the increase in prospective rates of return associated with a productivity acceleration, the effect of higher equity prices on the cost of capital, and the faster pace of increase in output. A productivity acceleration would also result in a powerful swing in the federal budget from deficit to surplus, as tax revenues soared while spending was constrained by discretionary spending caps. In addition, a productivity acceleration would likely result in both a further increase in the current account deficit and an appreciation in the dollar, if the productivity acceleration were a country-specific shock. Finally, the staff found that a productivity acceleration was a powerfully disinflationary event. Because wage change was initially little affected by a productivity acceleration, labour costs fell at first, raising the share of profits in income and then leading to downward pressure on prices. Hmm, they began to think: We are seeing a productivity acceleration everywhere but in the productivity data. While these effects were likely to be temporary, the simulations suggested they could last for several years. Interestingly, at this meeting, the staff, despite the remarkable fit of the simulation results with recent experience, reaffirmed their conclusion that there had been, to date, no evidence of an acceleration in productivity.

These simulations nevertheless positioned the staff to respond more rapidly and aggressively than they otherwise might have to the first tangible signs of a productivity acceleration in the data. The staff dragged me along, grudgingly for a while, as they adjusted upwards their estimates of structural productivity growth, modestly in 1998, and more significantly in 1999 and 2000. Figure 1 shows the acceleration in productivity in the second half of the 1990s.

A second lesson, therefore, is the importance of the supply side as a source of short-run dynamics, the power of productivity shocks, in particular, and the dynamics of the economy's response to an acceleration in productivity growth, especially the effect on short-run inflation dynamics and the NAIRU.


As I was struggling with understanding the sources of the exceptional macroeconomic performance, I identified two stories as possible explanations: One I called temporary bliss and the other, permanent bliss. The temporary bliss story focused on a coincidence of favourable supply shocks—including declining oil prices after 1996, slower increase in the cost of providing health care benefits to workers, and appreciation of the dollar. The permanent bliss story focused on longer-lasting structural changes, such as a decline in the NAIRU and an increase in productivity growth. I came to appreciate that a productivity acceleration resulted in a combination of temporary and permanent bliss stories. The productivity acceleration had a permanent or at least longer-lasting effect on the rate of growth of potential output. But it also had a powerful temporary effect on aggregate demand and especially on inflation.

If I had a dime for every time I had to say faster-than-expected growth and lower-than-expected inflation during my first three years on the FOMC, I would be a wealthy man today. The core of the issue seemed to be the breakdown of the normal relationship between inflation and unemployment, i.e., the breakdown of the Phillips curve, one of the foundations of my paradigm. This is the area of the paradigm most under attack and in most serious need of repair or refinement.

Fortunately, Steve Braun had developed a model of the effect of a change in structural productivity growth on the NAIRU in his 1984 Board staff paper. His motivation was to explain the effect of the productivity deceleration in the early 1970s on inflation. This model provided one of the foundations for simulations mentioned above and became the mainstay of my interpretation of the effect of the productivity acceleration on inflation dynamics.

To motivate this analysis, let's begin with a model in which a productivity acceleration has symmetric effects on wage change and price inflation. The

Figure 1
It's productivity, stupid
 Productivity: Non-farm business sector



model consists of a wage-price Phillips curve (equation (1)) (see Table 1) and a rate of change version of the markup equation (equation (2)). In the wage-price Phillips curve, wage change depends on structural productivity growth, the unemployment rate, and expected inflation. The markup equation specifies the rate of inflation, assuming a constant markup, as the rate of wage change less the rate of structural productivity growth. The productivity term in this model appears in exactly the same way in the wage and price equation and therefore does not appear in the price inflation reduced form (equation (3)) derived by substituting the wage-price Phillips curve into the dynamic markup equation. As a result, solving for the NAIRU (equation (4)), productivity growth does not affect the NAIRU.

Now assume that the rate of wage change depends on a long-moving average of structural productivity growth (equation (1')). That is, it takes time for a change in structural productivity growth to become built into the wage bargaining process. In this case, the price inflation reduced form (equation (3')) includes a new term, the acceleration in productivity growth, the difference between the actual rate of productivity growth and its long-moving average.

Table 1
Productivity and the NAIRU

Symmetric effect on wage and price equations

$$(1) \quad w = a + q - bU + p^e$$

$$(2) \quad p = w - q$$

$$(3) \quad p = a - bU + p^e$$

$$(3a) \quad p = -b(U - U^*) + p^e$$

$$(4) \quad U^* = a/b$$

w = rate of increase in nominal labour compensation

p = inflation

p^e = expected inflation

q = trend productivity growth

q^* = moving average of trend productivity growth

U = unemployment rate

U^* = NAIRU

Asymmetric effects in wage and price equations

$$(1') \quad w = a + q^* - bU + p^e$$

$$(2) \quad p = w - q$$

$$(3') \quad p = a - [q - q^*] - bU + p^e$$

$$(3a) \quad p = -b(U - U^{**}) + p^e$$

$$(4') \quad U^{**} = U^* - (1/b)[q - q^*]$$

p = inflation

p^e = expected inflation

q = trend productivity growth

q^* = moving average of trend productivity growth

U = unemployment rate

U^* = long-run NAIRU

U^{**} = short-run NAIRU

The disinflationary impetus arises because the acceleration in productivity initially leaves wage change unaffected, but lowers unit labour costs, putting downward pressure on inflation. As a result, the economy can operate during the transition period at a lower unemployment rate without triggering higher inflation. It is as if the NAIRU had declined.

It is therefore useful to construct a short-run or effective NAIRU (equation (4')) and we find that the solution for the short-run NAIRU—the value of the unemployment rate consistent with steady inflation, taking into account the effect of the productivity acceleration—also includes the term for the productivity acceleration. As a result, a productivity acceleration temporarily lowers the short-run NAIRU. Once wage bargaining fully adjusts to the productivity acceleration—and the moving average of structural productivity converges to the prevailing rate of structural productivity growth—the short-run NAIRU converges back to the long-run NAIRU. If the unemployment rate were to remain at the trough of the short-run

NAIRU, inflation would begin to rise and continue to do so until the short-run NAIRU converged back to the long-run value.

There may also have been some decline in the long-run NAIRU in the second half of the 1990s. But the effect of the acceleration in productivity on the short-run NAIRU was, in my view, an important part of the story. Figure 2 illustrates the acceleration in structural productivity growth relative to its 10-year moving average in the second half of the 1990s, and Figure 3 shows the temporary effect of the productivity acceleration on the short-run NAIRU, based on assumed values for the parameters in the model.

In my view, the short-run NAIRU may have fallen to about 4 1/2 per cent during 1999 and 2000, from its earlier long-run value of 5 per cent in about 1995. Part of the decline was due to the productivity acceleration, but part may reflect longer-lasting structural changes in the labour market. I would put the long-run NAIRU today at 5 per cent, and I believe the short-run NAIRU has converged back towards the long-run NAIRU.

There was at least one more surprising development related to inflation dynamics. The relationship between the unemployment rate and the capacity utilization rate appears to have changed significantly in the second half of the 1990s. Before then, the two utilization rates were so highly correlated that it was impossible to get them to both enter significantly a Phillips curve, sharing the role of measures of demand balance. In the second half of the 1990s, the capacity utilization rate was consistently lower relative to the unemployment rate than would have been predicted from the earlier relationship. As a result of the change in the relationship, it now appears that both capacity utilization and unemployment rates may yield some independent information about the overall balance of aggregate demand and supply in the economy.

Recent research at Macroeconomic Advisers has shown that it is possible to get both unemployment and capacity utilization rates into a price inflation Phillips curve. One interpretation of this finding is that excess supply in the manufacturing sector puts downward pressure on prices—that is, it reduces pricing power. In this case, firms might be less willing to grant wage increases than otherwise, implying a decline in the effective NAIRU. It is not clear at this point whether the shift in the capacity utilization-unemployment rate and hence the effect on the NAIRU is a temporary or more permanent phenomenon. But this might be one of the factors that, of late, has contributed to the intensity of the disinflation in the United States, despite an unemployment rate that is relatively low for this stage of the business cycle.

Figure 2
4-quarter and 40-quarter averages of structural productivity growth
Potential productivity growth




Figure 3
Short-run and long-run NAIRUs



2.2 Speculative excess

The second star of the second half of the 1990s was speculative excess, or irrational exuberance, as Alan Greenspan termed it. At the time the Chairman coined this phrase, in December 1996, the price-earnings (p-e) ratio for the S&P 500 was about 19, well above the very long-run average of between 14 and 15, but closer to and in fact slightly below the value that some have come to accept as sustainable in low-inflation periods. But if it is not clear that there was an asset bubble at the end of 1996, the same agnosticism cannot be credibly ascribed, at least with hindsight, to the experience later in the 1990s and into early 2000. There was an equity bubble and it had a dramatic effect on economic performance during the boom period and, obviously, after the bursting of the bubble. So other lessons of this period include the difficulty of identifying an equity bubble in real time, the contribution of equity bubbles to superior economic performance during the period when the bubble is evolving, and the pain associated with the post-bubble hangover.

Unfortunately, I didn't really focus on speculative excess until late in the period and especially after the correction, when I turned to the history books for insight. For most of the time, I admit I was too fixated on the issues of productivity and the NAIRU. This failure is especially disappointing personally, since the role of speculative excess and its aftermath was a theme that my colleague at Washington University, Hy Minsky, had preached to me at every opportunity. I resisted for most of the time we were together, but what Minsky had tried so hard to teach me came back to help me better understand the dynamics of the second half of the 1990s. I now view Hy as one of my most important and influential teachers.

My failure to recognize the importance of irrational exuberance to the story of the second half of the 1990s undoubtedly reflected the fact that there was no irrational exuberance in my paradigm. And, in addition, like the productivity acceleration, it was not clear that irrational exuberance was indeed an important part of the story until later in the period, though here again, the Chairman was precocious. In any case, the speculative excess contributed to the exceptional performance of the second half of the 1990s by amplifying the consumption and investment booms that would have been generated by the productivity acceleration even in the absence of speculative excess.

During this period, many on the FOMC emphasized the difficulty in identifying equity bubbles in real time. I said this many times. In retrospect, my time would have been better spent pouring over the long regularities in equity prices and earnings. It should have been painfully clear that there was at least the presumption of a bubble by 1998, if not earlier, although I

personally believe that the bubble did not become large enough to be dangerous until 1999. But part of the reason we may have focused on the difficulty of identifying an equity bubble was that we appreciated that we really didn't know what to do if we identified one.

Another reason why it was difficult to identify the asset bubble was that this was a period of apparent structural change, with a range of new possibilities summarized by the term "new economy." New economies, by definition, don't obey the old rules. As a result, many wondered if the surge in equity prices was really more an example of rational exuberance. Indeed, Alan Greenspan has argued that the structural changes during this period could legitimately have contributed to sustainable elevated price-earnings ratios, making it more difficult to determine whether or not there was a bubble.

A simple dividend-discount model helps to make this point and appreciate its possible shortcomings. The equation in Table 2 is a simplified form of a dividend discount model. The price of equities (P) depends on the level of dividends (earnings (E), multiplied by the dividend pay-out ratio (p)) divided by the discount factor, the sum of the safe interest rate (r) and the equity premium (q), minus the expected growth rate of dividends (g).

Table 2
The dividend discount model

$$P = p E / (r + q - g)$$

P = price of equities

p = dividend pay-out ratio

E = earnings

r = safe interest rate

q = equity premium

g = growth rate of dividends

The two explanations offered for rational exuberance are the productivity acceleration and a decline in the equity premium. Faster productivity growth raises the expected growth rate of dividends and should therefore immediately translate into a higher price-earnings ratio. A long, stable expansion (following the previous long expansion and mild recession) and an apparently less inflation-prone economy could well justify a decline in the equity premium, with a similar effect. In my view, it was developments like these that made many on the FOMC, including myself, reluctant to conclude that there was an asset bubble, even when the price-earnings ratio moved above 25 and on the way to 30 plus.

This was the basic line of reasoning underpinning the book, *Dow 36,000*. Glassman and Hassett (1999) hypothesize that historical equity premiums were too high to be justified by the relative volatility of returns to equities and bonds, and that as wealth-owners were coming to appreciate this, the equity premium would decline towards zero, justifying a much higher price-earnings ratio, and indeed a DOW of 36,000. This expectation was reinforced by comparing longer-term holding period yields. For example, five-year holding period rates of return to equities almost always exceeded the five-year returns to bonds.

Such a conclusion is, in my view, excellent arithmetic and insanely bad economics. The issue here is not so much whether the equity premium had declined to zero, although I suppose we might have a different perspective about the equity premium now than we might have had in 1999 and early 2000. The key issue is the general-equilibrium effects of an increase in productivity growth or a decline in the equity premium. If you introduce a productivity acceleration or a decline in the equity premium as a shock to a general-equilibrium model, one thing you can be sure of is that the safe interest rate will not remain unchanged. That's the difference between arithmetic and economics. The safe rate would surely rise, damping, perhaps significantly or completely, the effect of a productivity acceleration or a decline in the equity premium on equity valuation.

The lesson I take from this is that I want to pay greater respect to historical regularities between equity prices and earnings and keep the emphasis on reversion to the mean, on the offsetting movements within the discount factor when there are structural changes, and maintain a strong skepticism about structural changes that might be viewed as leading to material sustained changes in the sustainable p-e ratio. One reason the historical regularities provide discipline is that they incorporate general-equilibrium effects, while not all simple models do.

I also learned a lot about bubbles and their aftermath by looking at pre-war episodes, although, unfortunately, I didn't turn to the history books until the bubble itself was history. Thinking about the current episode, I wondered whether this was simply a repeat of earlier and relatively regular historical experiences. I wondered, in particular, whether there were waves of higher and then lower rates of innovation that, in turn, contributed to waves of higher and lower productivity growth. If this were the case, I thought we should be able to show in the data a somewhat regular pattern of intervening periods of higher and then lower productivity growth. I asked Dan Sichel of the Board staff to see whether he could find such a pattern in the long span of data on U.S. output. He produced the chart depicted in Figure 4, showing intervening periods of high and low growth, each lasting about 25 years,

with low productivity periods averaging about 1 per cent and high productivity periods about 3 per cent.

The next question was whether there were examples of bubbles in the economy during the transition to periods of higher productivity growth. A tour of the history books led to the conclusion that there were often bubbles in the innovating industries, though they did not always dominate the entire equity market.


But what was interesting was the rest of the story. The periods of innovation-induced equity bubbles were generally accompanied by a frenzy of investment in the innovating industries, resulting ultimately in a capital overhang. When the bubble burst, there was both a sharp decline in equity values and a powerful retrenchment of investment in the innovative industry. Sound familiar?

My conclusion was that in the second half of the 1990s we were in fact in a new economy, again. That is, we were repeating a pattern of pre-war experience with a wave of innovation, speculative excess, and its aftermath. I am not entirely sure how it would have helped me to make a better judgment about monetary policy, but it would at least have given me a better appreciation of the risks associated with bubbles and their aftermath. But I certainly wish I had been viewing the experience in the second half of the 1990s from this perspective while I was trying to forecast and set policy. I will return to the challenge of making monetary policy during this period.

It is also important to understand not only whether or not there was a bubble in the late 1990s, but where the bubble was. A comparison of the ratio between the NASDAQ (dominated by technology stocks) to the Wilshire 5000 (essentially the entire market) reveals a gradual and modest upward drift in this ratio until late 1998 or early 1999, when it takes off until the stock market peak and then plummets, quickly returning to near its earlier average. Technology stocks fell by 70 per cent from peak to trough (at least so far), and Internet stocks fell by 90 per cent. So this was principally a technology and especially Internet bubble. This follows the pattern of the pre-war experiences where bubbles were typically associated with transforming technologies and concentrated in the innovative industries, like electrical power, automobiles, and airplanes. This perspective is important because it raises the question of whether monetary policy could possibly be an appropriate instrument to deal with a sectoral overvaluation—a question I will return to.

While I had a lot to learn about the effects of productivity accelerations and about asset bubbles, history suggests that neither major changes in productivity growth nor asset bubbles occur very often. And not all productivity

Figure 4
Productivity growth: The long view



accelerations are associated with equity bubbles and capital overhangs. This was not the case, for example, in the previous transition to higher productivity growth early in the post-war period, although we certainly could have learned a lot from a more careful study of the effect of the productivity deceleration in the early 1970s. In any case, it is quite possible that the lessons I learned about productivity accelerations and asset bubbles, while helping put some order in the developments of the second half of the 1990s, may not be of much value until at least the next major wave of transforming technology and associated bubbles in the innovating industries—perhaps in another 45 years. By that time, I expect we may have forgotten the recent experience and will have to learn all over again about the “new economy.”

2.3 Post-bubble hangovers

While there was uncertainty both about whether there was an equity bubble and what to do if we believed there was, there was little doubt that monetary policy makers would have to move aggressively to cushion the damage if the bubble burst. On the other hand, although the technology and Internet bubble collapsed quite sharply, the overall level of the broader market indexes, such as the Wilshire, deflated more gradually and did not

immediately convey a sense of foreboding about the challenges that would lie ahead.

At the time of the bursting of the asset bubble, in March 2000, the economy was still expanding at a robust rate, and the momentum continued through the second quarter. As a result, the FOMC continued to tighten until the middle of the year, when growth began to slow. Given that the FOMC had been working hard to encourage a slowdown, it remained on hold as the economy slowed through the second half. Even by the end of the year, consensus forecast pointed to growth near 3 per cent and the unemployment rate was 3.9 per cent, just a tenth above the cyclical low of earlier in the year, and that cyclical low was the lowest unemployment rate in more than thirty years.

Still, by the end of 2000, the anecdotal stories had turned darker and the FOMC was alert to the possible need to ease. The first move was an intermeeting one in January 2001, and by the end of 2001, the FOMC had lowered the funds rate from 6 per cent to 1 per cent.

While much of the struggle to return the economy to an above-trend rate occurred after I left the Board, it was clear by late 2001 or early 2002 that the economy was struggling to unwind imbalances inherited from the boom period. These imbalances were a weight on the economy and also altered the dynamics of the economy relative to the regularities that were reflected in the data on which macro-models were estimated. As a result, those models provided less accurate guides to the strength of the economy.

The imbalances that weigh on the economy in the post-bubble period include the correction in equity prices, any capital overhang as a result of the frenzy of investment during the bubble period, and excessive debt burdens, especially for firms, that reflect the debt financing of the investment boom and the subsequent increase in leverage as a result of the decline in equity values. The correction of the overvaluation of equities led to a sharp decline in wealth and an increase in the cost of capital, undermining both consumer and investment spending. The capital overhang that appears to have developed at least in high-tech capital resulted in a sharp retrenchment of high-tech investment. Finally, the increase in balance-sheet stress encouraged firms to divert cash flow to repair balance sheets.

The weight of the imbalances not only slowed growth relative to what otherwise would have been expected, but also made traditional models, primed with data for the post-war—a period without such a bubble and post-bubble experience—less reliable for forecasting the strength of the expansion.

3 What I Learned About Monetary Policy

I had thought about and taught about the monetary policy process as well as the strategy of monetary policy before joining the Board. The basic vision of monetary policy that I brought with me to the FOMC is effectively summarized by the simple Taylor rule. That is, monetary policy makers should actively adjust the nominal funds rate in response to movements of output relative to potential and to inflation.

I learned a lot more about the process of monetary policy being on the inside—including the dynamics of FOMC meetings and the role of the Chairman in FOMC decisions—and also had to refine and adjust my views about the strategy of monetary policy, in light of the experience of this period. The greatest challenges involved how monetary policy should respond to an acceleration in productivity growth (once it became clear that it was the major shock), how the strategy of monetary policy should be adjusted in light of the higher-order uncertainty about productivity growth and the NAIRU, and whether and how monetary policy should respond to a suspicion of an evolving asset bubble. Along the way, I also had to refine my thinking on more traditional aspects of monetary policy. In each case, I would consider how the experience on the Board added complexity and richness to the simple Taylor rule story.

3.1 Productivity acceleration and measurement uncertainty

The key policy message from the simulations the staff ran in mid-1997 were that policy-makers should respond more gradually to robust growth and declining unemployment rates following a productivity acceleration than otherwise would have been appropriate, but would ultimately have to raise interest rates by more than otherwise. The initial gradualism was a reflection of the powerful disinflationary impetus of a productivity acceleration. An ultimately higher funds rate would be required because a productivity acceleration also increases prospective returns to capital and raises the economy's equilibrium real interest rate. I will return to this theme later.

My perspective on the role of uncertainty in policy making upon joining the Board was shaped by the certainty equivalence literature and the parameter or multiplier uncertainty literature. But the uncertainty that weighted most on FOMC decisions while I was at the Board was measurement uncertainty, uncertainty about the measurement of potential output and the NAIRU. These two variables—the key parameters in defining the limits to production, employment, and non-inflationary growth—had to be estimated rather than measured. The estimation in turn depends on the theoretical structure being assumed. That is, the maximum sustainable levels of output

and employment were defined as the levels consistent with stable inflation and hence were underpinned by the NAIRU and Phillips curve models.

I don't believe uncertainty about productivity growth is as important as uncertainty about the NAIRU, because monetary policy is fundamentally about responding to utilization rates, not to the rate of growth in potential output. But I learned along the way that uncertainty about productivity growth also translates into uncertainty about the output gap and the NAIRU and, therefore, carries some uncertainty about utilization rates, greatly complicating the implementation of the simple Taylor rule.


Work by the Board staff on the implications of uncertainty about the NAIRU and the output gap took two different paths, leading to different perspectives on how policy-makers should respond, but nevertheless leading perhaps to similar policy outcomes. The first approach focuses on monetary policy under uncertainty and suggests that monetary policy makers attenuate the response of the policy rate to changes in the variable about which there is heightened uncertainty. In effect, this meant attenuating the response of the funds rate to changes in the unemployment rate in the presence of heightened NAIRU uncertainty. Otherwise, by responding too aggressively to changes in a variable that was measured imprecisely, policy-makers might end up adding to the economy's instability rather than damping it.

The second approach focuses on the importance of continuously updating the estimate of the NAIRU (or the output gap), based on all available information, with special emphasis on prediction errors for inflation. That is, evidence that inflation was running lower than would have been expected, given the prevailing estimate of the NAIRU, would be used to justify and calibrate a downward revision to the estimate of the NAIRU. Time-varying parameter estimation proved to be precisely the technology required to efficiently carry out this updating process. During this period, the Board staff continuously updated their estimate of the short-run NAIRU based on the model I discussed earlier. In this case, the appropriate response might be to be as aggressive as otherwise to changes in the estimated output or unemployment gap, as long as the gap was constructed with appropriately updated estimates of potential output and the NAIRU. This option is, in principle, consistent with optimal policy under certainty equivalence.

Figure 5 depicts the actual funds rate relative to two specifications of the Taylor rule. The first specification is based on an estimate of the long-run NAIRU; and the second specification is based on a time-varying estimate of the short-run NAIRU, following the spirit of the model developed in Table 1.

The actual funds rate runs well below the prescription from the Taylor rule based on the long-run NAIRU during this period. One could reach either of

Figure 5
Two views of monetary policy in second half of 1990s



two conclusions from this comparison. First, the FOMC attenuated the response of the funds rate to changes in the unemployment rate during the second half of the 1990s as a result of heightened uncertainty over the NAIU. Second, the FOMC continued to be as aggressive as earlier in response to changes in the unemployment or output gap, but continuously revised downwards its estimate of the NAIU in response to better-than-expected inflation performance. This is an example of observational equivalence. I cannot empirically distinguish between these two sources for the discrepancy between the actual funds rate and prescriptions based on the long-run NAIU. But I prefer the continuous updating story, because, in my view, it better fits the process the staff was following during this period. But I also believe that the attenuation story significantly influenced the thinking of the FOMC itself. So perhaps both of these approaches played a role in shaping the policy outcome. The policy prescription that tracks the actual funds rate relatively well in Figure 5 is based on my specification of a Taylor rule with a very aggressive response of the FOMC to the unemployment rate, but based on an unemployment gap constructed using my estimate of the short-run NAIU.

3.2 Asset prices and monetary policy

I expect that monetary policy makers will continue for some time to struggle with the appropriate response to the possible emergence of asset bubbles. The experience of the 1990s suggests that asset bubbles are difficult to definitively identify in real time, especially before they become large enough so that the only decision is whether the correction should be initiated by policy-makers or left to the market. By that point, the main role of policy-makers is responding to the downside forces unleashed by a market correction.

However, looking back on this experience, perhaps with the benefit of hindsight, I am embarrassed that, at least by 1999 if not earlier, I still had any doubt at all that there was an equity bubble in the Internet and the wider technology sector. I believe that I, and perhaps some or many of my colleagues, were so focused on emphasizing how difficult it was to identify a bubble in real time that we failed to identify a bubble that, by at least mid-1999, was staring us in the face. I expect one reason we emphasized the impossibility of identifying a bubble so much is that we knew, deep down, that we had no idea about what to do if we actually identified one.

Looking back on this experience, I believe I can name three critical contributors to the equity bubble: the productivity-driven expansion, the stories about a transforming technology and a new economy, and the financial shocks of mid-1997 through 1998 (the Asian financial crisis, the Russian default and devaluation, and the implosion of LTCM).

The experience in the United States and Japan and my reading of earlier economic history suggest that asset bubbles are more likely to emerge in expansions powered by supply shocks—specifically productivity accelerations—than by expansions driven by demand shocks. In the latter case, real interest rates rise relatively early in the expansion, providing a friction that weighs against an open-ended increase in equity prices. Following a productivity acceleration, on the other hand, nominal rates may remain nearly constant despite robust growth and rising utilization rates, because the disinflationary impetus of that shock contributes to well-contained or even declining inflation despite persistent increases in utilization rates. In effect, an expansion powered by a productivity acceleration puts the economy, for a period, in the Dow 36,000 type of world—where equity markets may be responding to a perception of faster growth in dividends and perhaps also a decline in the equity premium, but without any offset (or at least without much offset) from higher real interest rates.

Such a period is also characterized by challenges to old rules—those linking unemployment and inflation and those linking equity prices and earnings. In the case of the stock market, that meant that wealth-owners were prepared to

place less attention to historical norms and dream more about future prospects. This is more likely the case when the productivity acceleration is powered by the perception (or reality) of a transforming technology, in this case information technology. Here, we would expect the bubble to be concentrated in the innovative industries, and indeed it was.

Finally, in a somewhat bizarre twist, the final blowout in the Internet and technology sectors was facilitated by the financial shocks that hit the U.S. and global economy, beginning with the currency and debt crisis in Thailand in mid-1997, continuing with the spread and intensification of the Asian financial crisis, and coming to a chaotic conclusion with the Russian default and devaluation and the implosion of LTCM. First, these shocks were expected to slow growth in the U.S. economy and therefore to substitute for tighter monetary policy than otherwise would have been appropriate. In fact, the economy didn't slow and, as a result of an incorrect forecast, we ended up with a lower path of the funds rate than otherwise might have occurred. Second, the last events were followed by an easing in monetary policy. Looking at the aggregate data for growth and utilization rates, one could not imagine why rates were lowered at this point. Of course, if one looked at forecasts (wrong again) and the extraordinary instability in the financial markets, it is easy to appreciate why policy was eased. But the net effect of these developments was that the funds rate was measurably lower going into 1999 than it otherwise would have been.

The latter episodes demonstrate the importance of changes in the forecast in precipitating changes in monetary policy. During this period, policy-makers ignored the momentum in the data and responded pre-emptively to significant discontinuities in the forecast. In principle, this is a very good idea. In practice, the forward-looking approach perhaps turned out not to be as successful during this period, because the forecasts turned out to be consistently too pessimistic.

The conventional wisdom at the Fed is that monetary policy makers should respond only indirectly to equity prices, the same way they would respond to fiscal changes and changes in exchange rates, and any other shocks to the economy. That is, policy should only respond to changes in output gaps and inflation, in the spirit of the Taylor rule. Such an indirect approach would, I believe, work well to mitigate prospects for equity bubbles during expansions driven by demand shocks, but may not be effective in the case of expansions driven by accelerations in productivity.

But that does not necessarily mean that an alternative approach would be more effective, on average, in stabilizing output and inflation during the emergence of an equity bubble. Given the difficulty in real-time identification, it is difficult to implement a direct approach—an approach

that responds, for example, to the divergence of equity prices from fair value. And it would have been even more difficult to sell a policy of wealth destruction, that is, a policy aimed at unwinding an expected bubble, than it was to tighten pre-emptively because of a concern about a decline in the unemployment rate.

Given the challenge of navigating through the bursting of a bubble, it would be prudent to think hard about steps that might be taken to mitigate the risks of a bubble. At the least, the suspicion of an evolving bubble should encourage monetary policy makers to reassess the consistency of their policy posture with their traditional objectives. An emerging bubble is like a neon sign flashing a warning—a warning that policy may be more accommodative than you think and therefore more accommodative than appropriate.

For example, one might ask whether, at the prevailing funds rate, monetary policy had, in effect, become more stimulative and, if so, whether that was desirable. There are at least two ways that monetary policy could have become more stimulative. First, to the extent that equity prices had soared and perhaps risk spreads in the capital markets had narrowed at an unchanged funds rate, financial conditions were becoming increasingly accommodative. Given the momentum in growth to begin with and the progressive tightening of labour markets already underway and of some concern, was this move to more accommodative financial conditions desirable? Another way of thinking of this might be to consider whether, if equities had not soared and risk spreads had not narrowed, monetary policy makers would have been inclined to lower the funds rate.

Second, monetary policy might have become more stimulative at the prevailing funds rate if the equilibrium real funds rate had risen relative to the unchanged funds rate. This possibility is especially relevant to the experience of the second half of the 1990s, because theory and empirical evidence suggest that the equilibrium real interest rate should rise in response to either an increase in underlying productivity growth or a decline in the equity premium. Moving the real funds rate to track such an increase in the equilibrium real rate would return the friction of rising real rates to the equity markets and reduce, if not completely avoid, the potential for an equity bubble.

This approach can be understood in terms of Wicksell's concepts of natural and market rates of interest. The equilibrium real rate in the Taylor rule is the analogue of the natural rate. If the natural rate increases relative to the market rate, financial conditions will have become expansionary and the economy will tend to grow at an above-trend rate. To avoid overheating, at some point the market rate must rise to close the gap with the natural rate.

The Board staff was experimenting with estimates of a time-varying equilibrium real rate during this period, but it was not quite ready for prime time. I will return to this later.

Another direction, perhaps although a reinforcement or restatement of the “reassessment” direction, is that monetary policy makers should err on the side of restraint during booms when there are suspicions of an asset bubble. The analogue here is post-bubble periods where it is widely believed that central banks should err on the side of ease, especially when the risk of deflation rises. Looking back on this period, the FOMC was almost continuously on a state of alert for a possible tightening of monetary policy, from the very first meeting I attended in July 1996 through the autumn of 1997 when the Asian financial crisis had intensified and was expected to provide a substantial drag in the expansion. The lean towards tightening was almost exclusively because of concern about possible overheating. But as core inflation continued to decline, the FOMC continued to resist tightening, except for a one percentage point move in March 1997. In retrospect, I believe an additional argument could have been advanced that monetary policy should at least lean against the progressively easier financial conditions associated with soaring equity prices during this period.

Most of all, I regret that there was so little discussion of what the appropriate role of monetary policy was during a period of soaring equity prices and a suspected equity bubble. When the topic would come up, the Chairman would often argue that he had said about all there was to say on the subject in his December 1996 speech. How would we know if we were in a period of irrational exuberance in the first place and what could we do about it if we were? Most of the disciplined discussion of this topic came only after the bursting of the asset bubble, and most of it from outside the Federal Reserve system.

But this should not suggest that I hold much hope that monetary policy could do much to avoid bubbles. One of the historical lessons, repeated in the current experience, is that equity bubbles are typically concentrated in the innovating industries. Monetary policy is a blunt instrument to use to deal with an Internet bubble, or even a wider technology bubble.

3.3 Monetary policy and expectations of future policy

Before coming to the Board, I had wrestled with the practical significance of rational expectations and was content with models that specified expectations formation as adaptive. My experience at the Fed convinced me that bond markets are fiercely forward-looking and have to be modelled as such. Rational expectations also appears to be important in explaining the effect of

a productivity acceleration, specifically in terms of capturing the effect on equity valuations of forward-looking expectations of earnings growth and the effect on consumer spending of forward-looking expectations of the growth of wage income. On the other hand, I don't find rational expectations as compelling when it comes to inflation dynamics. With respect to understanding the effects of monetary policy, it is forward-looking behaviour in the bond market that is especially important.

As I observed the way the economy responded to anticipations of monetary policy actions as well as to policy actions directly, I came to the view that I have underweighted the importance of expectations of future policy.

There is, I now appreciate, a powerful synergy between monetary policy transparency and consistency, on the one hand, and monetary policy effectiveness, on the other. When markets understand how monetary policy makers are likely to respond to economic developments, they are more likely to build into current bond prices expectations of future monetary policy. As a result, long-term rates typically begin to move in anticipation of changes in monetary policy, leading rather than following movements in the funds rate. Given that one of the limits on the effectiveness of monetary policy are the long lags between policy actions and their effect on aggregate demand, this anticipatory effect on long-term rates potentially speeds the effect of monetary policy and, in principle, can make it more effective.

I also came to appreciate that the FOMC has a love-hate relationship with the bond market. The FOMC loves the bond market when it believes the bond market has gotten it right, that is, correctly built in expectations about future policy that are consistent with the Committee's expectations. On the other hand, the Committee hates it when the bond market builds in a sense of immediacy and certainty with respect to near-term actions when the FOMC is uncertain about whether or not such action will be appropriate or should occur in the very near term.

The power of transparency and the potential downside to transparency explain the struggle that the FOMC has had over the past several years with its communication through the bias and now their risk assessment in its statement. The importance of expectation about future policy, after all, puts a heavy burden on communication by the central bank. And central banks historically have not excelled in communication, originally perhaps because of a belief that secrecy contributed to the effectiveness of monetary policy and later because of concern that too much or too explicit communication could cause excessive volatility in the markets.

The role of expectations about the course of future policy has become an especially active topic of discussion after I left the Board, as part of the consideration of non-conventional policies at the zero nominal bound. Pre-commitment strategies, as part of non-conventional policies, aim to alter market expectations of the future path of short rates, lowering long-term rates relative to the low or perhaps zero policy rate. A recent paper by Eggertsson and Woodford (2003) illustrates the potential power of a history-dependent policy rule in avoiding deflation or escaping from deflation. A simple example of such a rule is a price-level rule—either a constant price level or one that rises at a target rate of inflation, say 2 per cent a year. The difference with a price-level rule compared to an inflation target is that policy-makers do not forgive their past errors. If inflation falls below 2 per cent a year, it commits the central bank to maintain a lower policy rate longer, so that inflation moves above 2 per cent long enough to restore the inflation target. In this way, a price-level rule, by altering market expectations about the future path of short-term rates, can affect long-term rates, providing some leverage for the policy authorities even if the policy rate has been driven to zero.

Board staff work had identified this property of a price-level rule in 1999. While such a formal use of a history-dependent policy rule would be an especially difficult direction for a central bank, like the Fed, that does not have an explicit numerical inflation target and one that especially prizes flexibility, this work does suggest that a central bank's communications during periods when deflation risks are mounting and the policy rate is approaching zero can play an important role in shaping expectations of future policy and contribute to the objectives of monetary policy. In my view, the recent testimony by Chairman Greenspan and the FOMC statement at the 6 May [2002] FOMC meeting were a form of an informal, pre-commitment policy aimed at lowering long-term rates relative to the funds rate. If this is what they were aiming to do, they were very successful.

3.4 The federal funds rate and financial conditions

The FOMC controls the funds rate, but spending depends on a broader set of financial conditions. I certainly understood this before I joined the Board. Models assume that a change in the funds rate will have a “historically average” effect on financial conditions, that is, that changes in the funds rate will be passed through to the broader components of financial conditions as they have usually been. During my experience on the Board, I witnessed several episodes during which the pass-throughs were markedly greater or smaller than the historical average. The difference between the pass-through

in a given episode and the historical average, in turn, will affect the degree of traction that monetary policy has on aggregate demand.

The primary reason pass-throughs may appear greater or smaller is that shocks to financial markets may augment or offset the effect of the change in the funds rate. If there is, for example, an adverse financial shock (e.g., an increase in the equity premium) at the same time as a move to an easier monetary policy, the effect on financial conditions will reflect the combination of the two developments. One could interpret the outcome as reflecting an unchanged effect of monetary policy offset by the countervailing effect of the financial shock. But the net result will be that monetary policy makers will find it more difficult to stimulate aggregate demand. As a result, they will have to move the funds rate more aggressively than otherwise, and this, in turn, can increase the risks of confronting the zero nominal bound.

Goldman Sachs and Macroeconomic Advisers, among others, have developed financial conditions indexes to summarize overall financial conditions. These indexes are weighted averages of the contributions of the key components of financial conditions—typically private short- and long-term real interest rates, the real value of household net worth, and the real exchange rate, with the weights reflecting the relative size of the effects of each of the financial variables on aggregate demand.

This distinction between reading the monetary stimulus from changes in the real funds rate versus changes in the broader set of financial conditions was, in my view, an important consideration for monetary policy makers during at least three episodes in the second half of the 1990s and immediately afterward. First, the surge in equity prices in the second half of the 1990s made financial conditions progressively more stimulative, even though the nominal funds rate was virtually stable and the real funds rate was edging slightly higher. This could have been an argument for tightening monetary policy. Indeed, the staff did raise this issue during the second half of the 1990s, but the Committee never found it persuasive enough to reach a consensus for action.

Second, during the first half of 1999 and through early 2000, the continued rise in equity prices offset some or even much of the restraint associated with the rise in the federal funds rate. Finally, during 2001 and 2002, the attempt of the Fed to stimulate the economy by lowering the funds rate was significantly frustrated by the continuing decline in equity prices and the widening of risk premiums in the capital markets. Because the lack of pass-through is, I expect, a signature of post-bubble periods, it suggests to me that monetary policy will generally have a difficult time gaining traction in a

post-bubble economy, and this suggests that fiscal policy will have an especially important role during such periods.

The Taylor rule provides a response of the funds rate to shocks, under the presumption that pass-throughs are “normal.” The potential for significant differences from normal pass-throughs is an important reason why policy-makers will sometimes need to diverge from the prescriptions of the Taylor rule. For example, during 2001 and 2002, the financial conditions story provided support for a more aggressive easing than would have been consistent with simple Taylor rules. This, I believe, is what happened.

3.5 Time-varying equilibrium real rate

Like the NAIRU and the rate of structural productivity growth, the equilibrium real interest rate is potentially time-varying, not a constant. The Board staff has been working on a time-varying estimate of the equilibrium real federal funds rate. Economic theory suggests that changes in structural productivity growth, the structural budget deficit relative to GDP, and the equity premium are important drivers of the equilibrium real rate. The work at the Board, however, was able to explicitly pin down empirically only the link between productivity growth and the equilibrium real rate—with about a one-to-one effect from changes in productivity growth on the equilibrium real rate.

The concept of a time-varying equilibrium real rate is useful in interpreting developments after the productivity acceleration and during the emergence of the bubble and afterward. An increase in productivity growth, in effect, raised prospective rates of return on capital—the natural rate—making the prevailing real funds rate—the market rate—more stimulative. Monetary policy was, in this interpretation, increasingly accommodative during the second half of the 1990s, until the tightening that began in mid-1999, even though the real rate rose at least slightly and was above its historical average. The effect of the productivity acceleration on the equilibrium real rate was likely reinforced by the effect of an apparent decline in the equity premium. FOMC members began to use the distinction between market and natural rates to justify the tightening that began in mid-1999.

The concept of a time-varying equilibrium real rate also helps to understand the dynamics of the equity market correction. The bursting of the equity bubble and what followed was, in my view, precipitated, in part, by a reassessment of the profitability of owning and producing high-tech equipment, that is, by a decline in prospective rates of return to capital and hence the natural rate. As the Fed lowered the real funds rate, the market

rate, it was difficult to get traction because the equilibrium real funds rate was likely falling aggressively at the same time.

This is another important reason why the simple Taylor rule is only a point of departure and not the more definitive prescription for monetary policy. But the concept of a time-varying equilibrium real rate did not prove that useful in practice during the second half of the 1990s. First, the staff work was in the early stages and was not yet ready for “prime time.” There was simply too much uncertainty about the estimate of the equilibrium real rate and of estimate changes in the equilibrium real rate to allow it to play a more explicit role in policy decisions. Second, the staff work suggested that the equilibrium real rate had fallen to its lowest level in the post-war period in the early 1990s. The estimated equilibrium real rate increased in the second half of the 1990s—percentage point for percentage point with the increase in estimated structural productivity growth—just this brought the equilibrium real rate back to its long-run average. It did not justify a real funds rate above that suggested by the traditional Taylor rule.

One conclusion that emerges is that the Taylor rule places a lot of weight on “measurement”—really estimation—of two unobservable variables, potential output and the NAIRU. When we get these estimates “right” in the rule, it works very well. When we get them wrong, the rule does not work as well. What helps, nevertheless, is that the monetary policy response to the inflation gap dampens the damage from mistakes in the estimation of potential output and the NAIRU and hence stabilizes inflation, although, given misestimation of these variables, inflation might not settle down to its target value.

3.6 Deflation and the zero nominal bound

I hadn’t thought much about deflation or liquidity traps, except when I taught about the experience of the Great Depression. My education about real-world experience with deflation and liquidity traps began as I tried to understand developments in Japan and the challenge faced by the Bank of Japan. Most of my discussions about liquidity traps and the zero nominal bound at the Board during this period were focused on monetary policy in Japan, but we all understood that the lessons that we would draw from this experience might come in handy some day at home. Here are some of the lessons learned, all of which have been discussed in recent Board staff papers and in the paper by Governor Bernanke (2002).

First, the staff found that the forecast record during the period immediately following the bursting of the equity bubble in Japan was very poor. Forecasts were consistently too optimistic, accounting for what in retrospect looks like

an excessively slow pace of policy response. This difficulty reflects, I expect, the weight of imbalances inherited in the boom period on aggregate demand in the post-bubble period and the difficulty of traditional models in capturing this restraining effect. The implication for policy-makers, I believe, is that they should resist allowing confidence in a forecast that recovery is around the corner hold them in place when the incoming data suggest the economy is underperforming relative to expectations. Instead, they should place more emphasis, during the post-bubble period, on a timely response to the incoming data. If the data suggest the economy is struggling and underperforming relative to expectations, additional policy stimulus should be implemented in a timely fashion.

Second, the Japanese experience suggests the benefits of a positive as opposed to zero rate of inflation as a target. I came to the Fed with the view that 2 per cent CPI inflation was a sensible target, what I called price stability plus a cushion, but when I left I had a much better defence for having a cushion.

Third, related to the second lesson, it is important that the inflation target be symmetric. That is, policy-makers need to recognize that inflation can be too low as well as too high. Indeed, the experience in Japan also suggests there is, in the language of the recent Fed study, an asymmetry of risks related to inflation and deflation. Monetary policy makers can always raise rates high enough to control inflation. But they may lose control of an economy headed towards deflation if the nominal policy rate falls to zero. There may be non-conventional policy operations that could still provide incremental stimulus at the zero nominal bound, but such policies would take policy-makers into uncharted and much more uncertain territory. So monetary policy makers should act more vigorously and quickly when inflation falls below the target, compared with how they would respond when inflation rises above the target.

Fourth, as a result of the asymmetry in risks, policy-makers should err on the side of ease in post-bubble periods. That is, when the risk of deflation becomes more serious, policy-makers should put in place a degree of stimulus that would lead them to expect somewhat higher inflation than they would otherwise prefer in a couple of years. They can always clean up any excess of inflation at that point, an easier task than escaping from deflation if they fail to act in a timely and aggressive enough manner.

Fifth, monetary policy makers should not hold onto their ammunition when it is running low. Instead they should use it more aggressively. In effect, they should substitute speed and aggressiveness in their policy response for the smaller cumulative move available in the funds rate.

Finally, monetary policy makers should welcome complementary fiscal stimulus during such periods. The reliance on either monetary or fiscal policy alone would make it more likely that either would be used to excess, driving the policy rate to zero, in the case of monetary policy, or driving up deficits and the debt to income ratio to an extent that future use of fiscal policy was undermined. Of course, monetary and fiscal policy makers should try to craft fiscal solutions that provide the maximum short-run stimulus with the smallest possible adverse effect on the long-run budget position.

3.7 Non-conventional monetary policies

I had given no thought to non-conventional policies before joining the Board, but this topic became an active one for discussions in my office at the Board, focused mainly on what options we thought the Bank of Japan had and should exercise.

I came to the view that there were two potentially complementary directions for non-conventional policies. The basic principle was that once the policy rate was driven to zero, policy-makers needed to find ways to lower other interest rates relative to the policy rate, or generally raise asset prices. The first possible direction was the use of a communication strategy to alter the market expectations of the path of the policy rate in the future, lowering forward rates, and bringing down long-term interest rates and raising or at least anchoring inflation expectations to help real interest rates. The second possible direction was to conduct open-market operations in a wider range of assets, that is, to make outright purchases of longer-term government bonds, private debt, equities, sovereign debt, and so on, in the process lowering term and risk premiums in the respective markets. Since I have left the Board, the Chairman and Governor Bernanke have commented more explicitly on the options for non-conventional policies, and how these options might be exercised in the United States, if necessary.

3.8 Inflation targets and inflation targeting

One of the early surprises when I joined the Board was how the staff set the Fed's implicit inflation target in simple non-estimated specifications of the Taylor rule. The staff distributed prescriptions based on the Taylor rule to FOMC members in advance of each meeting and distributed a more detailed packet, including a variety of specifications, to Board members. For estimated versions of the rule, the staff typically estimated the implicit inflation target as part of a jumble of other terms subsumed in the constant. But when the staff ran the simple non-estimated version of the Taylor rule, it

typically used John Taylor's assumption about the inflation target. Now, I believe John has excellent judgment, but it seemed odd to ask him to set the inflation target instead of asking the FOMC. But I suppose it was not up to the staff to do so.

My education on the Fed's implicit inflation target came almost immediately, at the very first FOMC meeting I attended. That meeting featured a pre-arranged set of opening comments by two members of the FOMC on how the Committee should interpret the Fed's price-stability objective, followed by comments by the other Committee members. Janet Yellen argued in favour of a target of 2 per cent, above measured price stability and above an allowance for any inflation bias in measures of inflation. This choice was based on the now well-accepted view that inflation can be too low as well as too high. Specifically, if the average inflation rate is too low, the cyclical performance of the economy can deteriorate as a result of the inability of monetary policy makers to drive the real funds rate negative, if necessary, in the face of a particularly adverse shock. While the Chairman nevertheless made clear that he preferred an implicit target of zero, correctly measured, the consensus of the Committee was that the target should be 2 per cent, incorrectly measured.

This discussion was not expressly directed at moving towards an explicit inflation target, just a sharing of individual views of Committee members about their respective interpretations of the price-stability mandate. But this discussion does suggest that there was, at least at that point, a very strong consensus about the numerical rate of inflation that would have been selected as the target if the Committee were to move in that direction. And it also provides direct evidence on the Fed's implicit inflation target.

However, before you conclude that you now know the Fed's implicit target, even at that moment, consider the discussion that followed once the Chairman summarized the consensus as favouring a 2 per cent inflation target. The Chairman asked what measure of inflation the 2 per cent target applied to. Some members admitted they did not have an explicit measure in mind, but most were thinking of the CPI or core CPI. The Chairman remarked that the target should be defined in terms of the best measure of consumer price inflation and that was, in his judgment, the perfectly competitive equilibrium (PCE). The transcripts for this meeting are now publicly available and can be found on the Web site for the Federal Reserve Board of Governors.

This choice between the CPI and PCE is important, because the average difference between these two measures has been more than one-half of a percentage point over the past five years. President McTeer also made the very important point that the appropriate numerical target, when there is

more than one price index to choose from, should depend on the choice of the target. We can therefore conclude from the discussion at the June 1996 meeting that the implicit target for inflation, at least at that time, was 2 per cent for some index, but we are not sure whether in the final analysis that means 2 per cent for the CPI (and hence 1 1/2 per cent for the PCE) or 2 per cent for the PCE (and hence 2 1/2 per cent for the CPI).

The increase in the number of countries opting for formal inflation-targeting regimes during the time I was on the Board forced me to consider more carefully whether or not the Fed should adopt an explicit inflation target and, if so, what type of inflation-targeting regime would be a good fit with our mandate and experience.

I came to believe that the design of an inflation-targeting regime involved two sets of choices. First, should the mandate be hierarchical, assigning a priority to achieving the inflation target, or should it be a dual mandate, giving equal importance to the full employment and price-stability objectives? So-called inflation-targeting countries opted for the hierarchical mandate, except in the case of Australia, while the United States operates under a dual mandate. The second choice is between implicit targets, as in the case of the Fed, or explicit targets, as in the case of formal inflation-targeting regimes.

My conclusion was that the best regime for the United States was one that reaffirmed the commitment to the dual mandate but introduced an explicit inflation target. The Fed prides itself on a flexible and pragmatic approach to monetary policy. Inflation targeting has been described by Governor Bernanke and others as a form of constrained decision making. So the challenge is how to preserve an important degree of flexibility while at the same time constraining decisions! In addition, there is, in my judgment, a strong internal view that the purpose of any refinement in operating procedures is to ensure continuity of monetary policy, as practised in the Greenspan years, rather than to change the strategy. Finally, I see no chance at all that the U.S. Congress would allow the Fed to diminish its responsibility for pursuing full employment in the conduct of monetary policy. Hence, the only possible way and also in my view the best way for the United States would be to adopt an explicit numerical inflation target in the context of its dual mandate.

3.9 Money and monetary policy

One last surprise when I joined the Board was the limited attention given to the monetary aggregates in monetary policy decisions. One could attend several FOMC meetings without a mention of the money supply, at least by

a member of the Committee. That should not lead you to believe that the FOMC and the Board staff never paid attention to the monetary aggregates. They did so until the early 1990s, when the link between money growth and economic activity deteriorated very sharply. So the staff's "indifference" to the monetary aggregates was empirically driven.

The main opportunities for discussion of the monetary aggregates during the first several years of my tenure were during the semi-annual, two-day FOMC meetings at which the Committee was required by law to update monitoring ranges for a number of the monetary aggregates. This turned out to be a surprisingly mechanical exercise, with changes virtually never made in the ranges. The main objective of this exercise was to avoid drawing any attention to the monetary aggregates, specifically to avoid any hint that they were now playing a more important role in monetary policy decision-making. While there was evidence in the early part of the second half of the 1990s that the relationship might be becoming somewhat more reliable again, it was widely agreed within the FOMC that such a conclusion was premature. As a result, the best course was to continue to downplay the role of the monetary aggregates.

I came to the view that it was time to give the monetary aggregates at least a limited secondary role. I liked the notion of a reference value for money growth, the rate of growth of the monetary aggregates that, on average, would be consistent with the estimated rate of growth in potential output and the Fed's implicit inflation target. Of course, the fact that the Fed did not have an official inflation target made it more challenging to set such a reference value, but, given the discussion above, that was not an insurmountable obstacle. The ECB had made a reference value for money growth the "first pillar" of its monetary policy strategy, according money growth far more prominence than I was comfortable with. Nevertheless, I thought it deserved a promotion from its virtually neglected role at the Fed.

I had the staff develop a monetary aggregate packet for me that tracked the growth of M2 relative to a reference value that I set, based on a 1 per cent inflation target for the chain-weighted measure of the GDP price index, consistent with a 2 per cent target for the CPI measure of inflation. I would meet with the staff each FOMC cycle to assess developments in the monetary aggregates and look for any insights that might reinforce or contradict what other indicators were suggesting. While I believe this is a useful exercise, I have to admit that this discussion never provided insights that altered my view about the appropriate course of monetary policy. Indeed, the staff confided that they would prefer that the Committee did not elevate the role of the monetary aggregates. If the Committee did so, they expected they would be spending most of their energy explaining shifts in

the demand for money, as they had earlier and, in light of such shifts, advising the Committee to disregard various instances when money growth moved outside the monitoring range.

Conclusion

When I look back on my tenure at the Board and the FOMC, what I miss the most was not the opportunity to participate in FOMC meetings, but the extraordinary intellectual climate at the Board. My office was often filled with staff, brainstorming on various issues related to the productivity acceleration, the NAIRU, equity valuation, and the strategy of monetary policy. I am much richer today, intellectually I mean, as a result of that experience.

Bibliography

- Bernanke, B.S. 2002. "Deflation: Making Sure 'It' Doesn't Happen Here." Remarks by Governor Ben S. Bernanke before the National Economists Club, Washington, DC, 21 November. <<http://www.federalreserve.gov/boarddocs/speeches/2002/20021121/default.htm>>.
- Braun, S.N. 1984. "Productivity and the NAIRU (and Other Phillips Curve Issues)." National Income Section, Working Paper No. 34. Board of Governors of the Federal Reserve System, June.
- Eggertsson, G. and M. Woodford. 2003. "The Zero-Bound on Interest Rates and Optimal Monetary Policy." Presented at the Brookings Panel on Economic Activity, Washington, DC, March. <www.princeton.edu/~woodford/>.
- Glassman, J.K. and K.A. Hassett. 1999. *Dow 36,000: The New Strategy for Profiting from the Coming Rise in the Stock Market*. New York: Random House.
- Greenspan, A. 1996. "The Challenges of Central Banking in a Democratic Society." Remarks by Chairman Alan Greenspan at the Annual Dinner and Francis Boyer Lecture of The American Enterprise Institute for Public Policy Research, Washington, DC, 5 December. <<http://www.federalreserve.gov/boarddocs/speeches/1996/19961205.htm>>.
- Malkiel, B. 2002. *A Random Walk Down Wall Street: The Best Investment Advice for the New Century*. Seventh edition. New York: W.W. Norton & Company.
- Meyer, L.H. 2000. "The Economic Outlook and the Challenges Facing Monetary Policy." Remarks at the Century Club Breakfast Series, Washington University, St. Louis, MO, 19 October. <<http://www.federalreserve.gov/boarddocs/speeches/2000/20001019.htm>>.

Meyer, L.H. 2001. "Inflation Targets and Inflation Targeting." Remarks at the University of California at San Diego Economics Roundtable, San Diego, CA, 17 July. <<http://www.federalreserve.gov/boarddocs/speeches/2001/20010717/default.htm>>.

———. 2001. "What Happened to the New Economy?" Remarks before the New York Association for Business Economics and The Downtown Economists, New York, NY, 6 June. <<http://www.federalreserve.gov/boarddocs/speeches/2001/20010606/default.htm>>.