

LABOR AND EMPLOYMENT RELATIONS ASSOCIATION
59TH ANNUAL MEETING
DISTINGUISHED SPEAKER ADDRESS

Chicago, Illinois
January 5, 2007

.....

Labor Markets and Monetary Policy

Thank you all so much for inviting me to be this year's Distinguished Speaker. It's fitting that I get to speak to you here in Chicago. The first time I spoke at an IRRA meeting, at the end of 1968, the meeting was held here in Chicago. My topic that day was "Multilateral Bargaining in the Public Sector," with an emphasis on public education. It was my first big academic presentation, and Arnie Weber chaired the session. My coauthor, who is in the audience today, was Ken McLennan.

Here we are in Chicago again. I'm happy to say that my background in labor economics has served me well in my current job as president of the Chicago Fed. Prices, wages, and inflation are key concerns for me now, just as they were in earlier periods in my career. In setting monetary policy, the Fed is charged with achieving both maximum sustainable growth and price stability. This is known as our dual mandate. Dual, of course, means two, but the Federal Reserve Act actually tells us to pursue three goals: price stability, maximum employment, and moderate interest rates.

Using labor market data to benchmark the economy

About two-thirds of household income comes from employment income, and employment expenses make up about two-thirds of business costs. So the analysis of labor markets plays a key role in assessing the spending and growth prospects for the economy as well as cost pressures on inflation. The published labor statistics provide a great deal of information about national labor markets and how various sectors of the economy are performing. In addition, we supplement these published data by gathering an enormous amount of nonquantitative and anecdotal information throughout the country.

One thing we do is then compare these data to various benchmarks that describe an economy performing in a manner consistent with maximum sustainable growth and price stability. For example, we ask:

- Is employment growing at a pace consistent with maximum sustainable growth?
- Is the unemployment rate roughly in line with long-run sustainable levels?
- And are wages growing relative to productivity in a way that is consistent with price stability?

An important observation is that labor markets are constantly evolving, which can cause these benchmarks to change. The demographics of the workforce are constantly changing. New labor market institutions have altered the way firms hire, organize, and pay their employees. Monetary policy could go off track if we don't recognize the impact of such structural changes on the economy and the benchmarks we use to evaluate the incoming data. So today I would like to take a more detailed look at a number of these structural changes and how they have influenced some of our important benchmarks. And as I go through them, you'll see a recurring theme: research can provide us insights into the workings of labor markets that can be very valuable in making well-informed monetary policy decisions.

Maximum sustainable growth

Let's start with maximum sustainable growth, which we also refer to as potential output growth. This is the maximum rate of real growth the economy can maintain in the long run without leading to an increase in inflationary pressures. It reflects the growth in both the number of available workers and the productivity of those workers.

In order to estimate the number of available workers, we need to understand the underlying trends in population growth and labor force participation. This can be challenging, because these trends can change over time.

For instance, the working-age population grew rapidly in the 1960s and '70s, averaging two percent per year, as the baby boomers finished school and entered the labor force. But starting in the early 1980s, working-age population growth slowed to a little more than one percent per year. And it's projected to fall a bit further over the next decade, given the size of the younger generation entering the workforce relative to the size of the retiring baby-boom group.

The labor force participation rate is also changing. From the early 1960s until around 2000, the labor force participation rate increased steadily, as a rise in the fraction of women who work outside the home offset a slow decline in the fraction of men in the labor force.

But labor force participation declined a good deal during the recession of 2001 and has not returned to its previous levels. The reasons for the decline have been the subject of some debate. Some thought it was a sign of an economy that remained weak long after the official end of the recession. This interpretation assumed that many people who couldn't find work gave up looking and withdrew from the labor force.

But we at the Chicago Fed along with many other analysts take a different view. We now think that the decline in labor force participation was mainly a manifestation of some long-run changes in demographic and cultural trends.

One of the demographic trends relates to the aging baby boomers. There's a lot of focus on them approaching the traditional retirement age of 65, but people tend to scale back their participation well before that. Indeed, more and more boomers are past the age at which people are most likely to participate in the workforce.

Furthermore, the increase in female labor force participation seems to have run its course. The women of my generation entered the labor force in greater numbers than the women of my mother's generation. And the women of my daughter's generation participate in greater numbers than the women of my generation. But these increases no longer appear to be occurring. It's likely that the women of my granddaughters' generation will participate at the same high rate as my daughter's generation.

So our view is that most of the decline in labor force participation that happened after 2000 was not due to a weak economy but instead to these long-run trends. Indeed, it seems very likely that it will continue to fall over the coming years.

How do these developments affect our benchmarks? Given the population growth and labor force participation rates that prevailed during the 1990s and early 2000s, the monthly job growth that was sufficient to provide jobs for the net number of entrants into the labor force was about 150,000. So this figure was an important benchmark to evaluate whether the economy was growing at a pace sustainable in the long run. The recent changes in population growth and labor force participation imply that we need to change our benchmark for employment growth to something more like 100,000 per month.

This change in the benchmark has only recently begun to be fully recognized by those who follow the economy on a day-to-day basis. And it has important implications for our view of recent labor market conditions. Job growth has averaged roughly 160,000 over the past six months. By the old standard, that would be just average, but given the current trends in the labor force, we view such growth as quite solid.

I should say, however, that while 100,000 per month appears to be the right benchmark for the next year or two, there is a lot of uncertainty about this mark in the long run. Some of this uncertainty revolves around the future participation of the baby boomers. People are living longer, healthier lives; this may allow them to work until they are older. The women now nearing retirement age are the first generation to have worked outside the home as much as they did. Perhaps they will surprise us again by working longer than we think. Moreover, wages for all workers may change in response to these trends, convincing some to work more, others to work less, than they would otherwise. Of course, the proper benchmark depends on other factors as well, including growth in household wealth, the long-run health of Social Security and Medicare, and trends in fertility, life expectancy, and especially immigration. Major changes in our immigration policy, for instance, could have a big impact on the rate at which the workforce expands and, therefore, the estimate of employment growth consistent with sustainable economic growth. Research by labor economists can be used to improve the estimates of the effects of these factors on the growth in the labor force. This can help us better anticipate changes to our employment benchmarks in the future.

As I noted previously, the growth rate of potential output also depends on labor productivity. From 1973 to 1995, productivity growth averaged 1½ percent per year, but since then it has averaged 2½ percent per year. It's now evident that a significant pick-up in productivity growth occurred in the mid-1990s, but at the time, it was not clear how long lived the change would be. Former Fed Chairman Alan Greenspan was one of the first to assert that something long lasting was in train. By recognizing the persistence of the acceleration in productivity, we were able to successfully adjust monetary policy in a way that supported economic growth at a faster pace than many at the time thought was possible without generating higher inflation.

The step up in productivity growth owed a good deal to an increase in investment in computers and other information technology. But there is still a large unexplained component. My personal opinion is that much of this has to do with innovations in management practices. Part of this is intangible, and I believe that leadership improvements have made a big difference. People matter. But more observable measures of management quality have improved as well, and recent research suggests that these measures are related to higher productivity, profitability, sales growth, and firm survivability. Some of this improvement in management quality takes the form of better human resource practices, such as setting goals for employees that are realistic, transparent, and consistent; tracking and reviewing performance; and using better incentive systems that identify and reward ability and effort. Human resource experts have long studied such practices, and lots of great work has been done already, but even more is needed to understand their importance to the performance of economy-wide productivity.

Productivity also depends on the quality of the workforce. The high education levels and greater experience of our aging population has contributed to our higher productivity rates of late. But that trend has already started to slow. It is estimated that in the late 1980s and early 1990s, improvements in worker skills were adding four-tenths of a percentage point per year to the growth of output. By the end of the 1990s, this had fallen to less than two-tenths of a percentage point, and we could soon see a decline to less than one-tenth of a percentage point, as the highly experienced workers of the baby boom generation retire in increasing numbers. Gains in education and other workforce skills by new entrants, and skill improvements by remaining workers, could allow us to avoid much of this decline. But if we consider recent education trends, I'd say we have our work cut out for us on this. College graduation rates are growing, as families have noted the very high and climbing economic returns to education. But high school completion rates have stalled, and there is a great deal of dissatisfaction with the results in our public elementary and secondary schools, especially given our already enormous investment in education. Clearly, more systemic changes are needed to achieve better educational outcomes and boost the growth in worker skills.

The slower growth of available workers and the rate at which trends in educational attainment and labor market experience add to those workers' productivity both imply slower growth in potential output. In the case of the former, the swing from an era in which labor force participation was growing around a tenth of a percent per year to one in which it is declining around two tenths percent per year implies a slowing in potential output growth. So do the slower improvements in the educational and experience composition of the labor force. Together, these suggest that the contribution of labor to potential output growth may have declined by a bit more than half a percentage point since the late 1990s.

Natural rate of unemployment

Employment and productivity growth are not the only benchmarks we use to assess how the real economy is performing. We also try to gauge how the level of economic activity compares to its potential at any point in time. One way we do this is by comparing measures of resource utilization, such as the unemployment rate, to benchmarks consistent with potential output. We often refer to this unemployment benchmark as the natural rate of unemployment. Such assessments also play a role in our forecasts of inflation. Here, too, we face the challenge that our resource utilization benchmarks have been changing over time.

In the mid- to late-1990s, when I had been at the Chicago Fed for only a short time, one of the big questions was whether inflation would increase given the tightening labor market. The unemployment rate had fallen from 7½ percent in 1992 to 5½ percent in early 1995. Conventional thinking at the time was that the

natural rate of unemployment was around six percent. So when unemployment went below that six percent level, a number of people worried that accelerating inflation was just around the corner. They thought that unless the Fed tightened policy, wages would start to increase at a faster rate than productivity and generate higher inflation.

Certainly the threat to price stability was something the Fed took seriously. But we realized we had to be cautious in interpreting the data. This is because there were powerful undercurrents which could affect the level of unemployment consistent with potential output.

Now it's tempting for me to cite the record since that time and simply claim vindication. The unemployment rate has averaged about five percent and core PCE inflation has averaged 1.8 percent since 1995. Most analysts have since revised their estimates of the natural unemployment rate, putting it in the range of five percent. But we live in an uncertain world, and policy makers always have incomplete information. This was particularly true back in 1995 because of the structural changes taking place. So it's important to revisit these developments, because they help shape our thinking about labor markets today and raise questions that, if answered, could help fine-tune our current estimates of the natural rate and potential output.

One reason the natural rate may have fallen is that job search and the process of matching workers with jobs has become more efficient. In the 1990s, the dissemination of fax machines and then email allowed a worker to submit a resume or application in minutes, rather than a day or two by mail. This dramatically reduced the time needed to apply for a new job. Numerous web sites now allow people to search for the jobs that meet their criteria and to almost instantaneously apply for multiple jobs; similarly, businesses can automatically search large databases of resumes and quickly screen applicants with the desired qualifications. Who knows where this is heading? For instance, I read the other day that Google has its job applicants fill out an elaborate online survey, which asks all sorts of odd questions, such as whether you have a neat or messy work space, or even whether you have ever set a record in anything. They then use a complicated algorithm to sort through the answers to determine how well the job applicant would fit into its unique culture.

Similarly, the very rapid growth of new employment practices, which some call just-in-time hiring, likely played a role in lowering the natural rate of unemployment. Thanks to a number of developments, but particularly the rise of temporary help services firms, companies now have greater flexibility to adjust their work forces on short notice. Twenty years ago, the temporary help industry was very small, representing only about one-half percent of nonfarm payroll employment. Since then, it has increased significantly and now is close to two percent of employment. At some firms, the majority of employees are temporary workers. Temps allow firms to easily vary the scale of their operations in response to actual changes in demand, rather than varying their workforce based on what they expect will happen to demand in the future. This could mean less labor hoarding during periods of slow economic activity.

These just-in-time hiring practices may even be part of the reason why the last two recessions were followed by prolonged periods of slow or nonexistent employment growth—what people have called jobless recoveries. In the early parts of a business expansion, there is typically a lot of uncertainty about its staying power. In such circumstances, firms may be reluctant to hire until they definitely need to. So the knowledge that there are readily available temp workers may have allowed firms to delay hiring longer than usual.

Furthermore, relative to the 1980s, we now see many fewer short-term layoffs of regular workers. It seems pretty likely that some of this reduction is now showing up as variation in the number of temporary work-

ers. Since temp agencies appear to be efficient at finding these workers new assignments, this trend is probably lowering overall unemployment rates.

Still, the sum of these improvements in labor market hiring and job search efficiency continues to be difficult to quantify. One reason is that the U.S. did not have good data on job vacancies until the BLS's Job Openings and Labor Turnover series, or JOLTS, began in 2001. Without historical context, we don't know how to interpret a change in the JOLTS vacancy rate, such as the rise from 2.1 to 3.2 percent that we have seen since 2003. Is this typical for an expansion in which the unemployment rate has declined by 1.5 percentage points? Or is it a slower than usual increase, which would be a sign of increased labor market efficiency? We don't know right now. As we observe these data for a longer period of time, we will be able to use them to address such issues regarding labor market efficiency as well as many other interesting questions.

Another possible reason why inflation stayed under control in the late 1990s, despite what seemed at the time to be very tight labor markets, is that perhaps "worker anxiety" had caused employees to accept lower compensation packages. According to this theory, workers became more concerned about their job security in response to increased reports of layoffs and stories about jobs being eliminated. Former Chairman Greenspan speculated a number of times that this was a factor behind the more modest than expected wage pressures. These theories, which relate to the balance of bargaining power between labor and management, are hard to test. But certainly, the experts in this audience should have some further insights.

This issue may be pertinent again now. Up until a year or so ago, the increase in unit labor costs was running noticeably lower than in previous economic expansions. Some have speculated that this could be explained by an increase in worries about outsourcing, especially overseas, which has kept workers in a weak bargaining position.

Demographic changes also are a factor in the decline in the natural rate of unemployment. As the baby boom generation acquired more working experience, their ability to find jobs improved. It happens to every generation. Experienced workers have more employable skills, know more about what jobs match their skills, and have built a broader network to help them conduct a job search. Since 1960, the unemployment rate for 25-34 year olds has been about two percentage points higher than the unemployment rate for 45-54 year olds, on average. In the late 1980s, about 15 percent of the labor force was 45-54. This share has since trended up and now appears to be peaking at 23 percent. In turn, the natural rate of unemployment for the entire economy should be lower because of the relatively higher employability of these workers, who now constitute a larger portion of the workforce. Other changes in the workforce, such as higher average education levels, increased incarceration rates, and greater use of the disability insurance program, may also have brought the natural rate down.

Wage growth

Of course, in addressing our price stability mandate, we look carefully at data on labor costs. We monitor compensation trends and compare them against benchmarks. In particular, if wages rise faster than productivity for a prolonged period, the resulting persistent increases in unit labor costs can fuel price inflation. The words "prolonged" and "persistent" are important here. Real wages and productivity track each other closely over the long term but not necessarily over shorter periods. In other words, swings in the share of national income going to labor can last for a number of years.

One challenge is understanding what the various different wage measures are telling us about this dynamic process. Over the last several years, changes in compensation practices likely have caused these measures to send conflicting signals about labor costs. For example, the total compensation series, from the BLS's Employment Cost Index, has grown by roughly 3½ percent per year over the last decade. Over the same time, the compensation per hour figure, based on the National Income and Product Accounts and reported in the BLS's productivity report, has grown by just over 4½ percent per year. But in the decade and a half prior to 1996, each grew at nearly identical rates of around four percent per year.

One reason why the two series may have diverged is that pay practices have changed. The NIPA-based measure includes stock-option realizations and lump-sum bonuses, while the ECI total compensation measure does not. As you know, these and other forms of variable pay now have a more prominent role in compensation. Indeed, recent studies suggest that stock-option realizations account for about one-quarter to one-half of a percentage point per year of the growth in compensation per hour during the late 1990s. Stock-option realizations also add quite a bit of year-to-year variability in compensation that is tied to swings in the equity markets. Such movement might have little to do with labor market conditions at the time the options were granted.

Because variable pay is easier to change from year to year, it can reduce the degree of downward nominal wage rigidity in the economy. This may lessen inflation persistence by reducing the tendency for current wage increases to become locked in to firms' labor costs. Flexible wages also strengthen the tie between individuals' performance and pay. This can enhance aggregate economic performance by enabling labor markets to achieve a more efficient alignment of labor resources with their most valuable use.

The effects of new compensation practices are a fertile area for monetary-policy-relevant labor-market research. Knowing how such changes may have altered the relationships among current labor market conditions, compensation, and firms' pricing decisions would be very useful for improving our inflation forecasts. Those forecasts, of course, are key tools in the Federal Reserve's efforts to achieve its price stability mandate.

Conclusion

So to conclude, the changes we have seen in demographics and labor market institutions over the years mean that we must carefully and continually examine the assumptions that underlie the job-market benchmarks we use to gauge the performance of the economy. The achievement of our goals of sustainable growth and price stability is at stake. So, any help that you can give us in better understanding the dynamics behind employment, productivity, or wages is greatly appreciated. Thank you.