Labor Supply in the New Century

Labor Supply in the New Century

edited by

Katharine Bradbury, Christopher L. Foote, and Robert K. Triest

Federal Reserve Bank of Boston
Boston, Massachusetts

This edited volume is based on a conference held in June 2007 by the Federal Reserve Bank of Boston.

Conference Series No. 52

This book was set in Sabon by Sztrecska Publishing and was printed and bound in the United States of America.

Contents

Foreword

Cathy E. Minehan

When Kathy Bradbury, Bob Triest, and Chris Foote came to me last year with the proposal for the 2007 conference,"Labor Supply in the New Century," I thought their idea was inspired. Questions about the quantity and quality of the U.S. labor force have been critical to discussions of current policy setting and to longer-term concerns about the American economy's ability to continue producing the goods and services needed to maintain rising U.S. living standards.

In particular, over the past decade we first saw labor force participation rates in the United States reach a postwar peak in 2000, as the "new economy" drew in every worker it could attract. Then by 2001–2002, we witnessed labor force participation rates fall below mid-1990s levels, and wondered what could be said about a steady-state level of participation? Over the short run, how did this changing level of participation affect how one thinks about such issues as monthly employment data? Is monthly job growth of 130,000 just sufficient to absorb new entrants to the American labor force, or is it a pace that is well above equilibrium and likely over time to stress the economy's underlying resources? Is there a greater existing supply of labor just waiting to be drawn into the U.S. workforce by strengthening cyclical demand?

Is the increased participation observed among some older workers reflecting a need for more retirement income? Many of the baby boomers born between 1946 and 1964 will not have amassed the financial resources that will allow them to completely leave the workforce when they reach their mid-60s—and even if they can, many will not want to end their established careers, or may want to start new ones. Greater life expectancies mean that we need to reconsider what throughout much of

the twentieth century has been regarded as the traditional age for retirement, usually begun when people reach 62–65 years of age. What public policies and private sector human resource practices will be necessary to both shape an American workforce to take advantage of this generation's skills and expertise, and create working experiences that encourage and maintain growth in the U.S. labor supply? We know that the baby boom generation has sometimes behaved differently from its predecessors; in particular, the high labor force participation rates of baby boomer women transformed not only the U.S. labor market, but arguably also changed American society and lifestyles in very fundamental ways. Gene Steuerle, the conference's keynote speaker, contends that in the first half of the twenty-first century, older American workers will increase their labor force participation rates and have a similar impact on the U.S. workforce as women did in the second half of the twentieth century.

The quantity of labor supplied to the market depends on more than just demographics. Labor market opportunities and labor demand for certain skills also matter. And in recent decades these opportunities have increasingly favored skilled and highly educated workers relative to those with lower levels of educational attainment. Within the Federal Reserve System we have witnessed this trend firsthand. There has been a pronounced shift toward demanding more highly educated workers, particularly as paper-based retail payments instruments—checks and cash—have declined relative to electronic and card-based alternatives. Technological change has led to a sharp reduction in our employment of workers in tasks associated with check processing and check clearing, leaving our remaining jobs concentrated in areas requiring either higher degrees of educational attainment, or more service-related skills.

It is striking that the shift in labor demand toward the highly educated has been so strong that the premium that employers are willing to pay for college-educated workers and those with postcollege degrees has increased greatly, even while the workforce has, as a whole, become better educated. This, in turn, has tended to depress the labor supply of the less educated, especially among men, while at the same time providing incentives for individuals to invest in education and training.

Will the quality of the U.S. labor force be able to meet the demands of the twenty-first century? Are those demands becoming so specialized

and stratified for skills at the high end and at the low end of the wage scale, but not in the middle, that certain portions of the American labor force will be permanently left behind? In this regard, I personally found a 2006 *New York Times* article about some middle-aged American men more or less permanently separated from the workforce chilling in what it implied about the lack of hope for some in our economy. Yet as a policymaker, I am loath to accept such fate as preordained—this is a problem we can solve.

For different reasons, I am also concerned about women and their declining rates of labor force participation, particularly those who are 30 to 45 years old. Women make up more than half of all students in many of our graduate schools—and in areas such as business and engineering their numbers are growing. That's a lot of brainpower to lose if participation rates drop just as those women become experienced in their working situations. This conference addresses these issues and more.

Indeed, we are excited by this interesting collection of papers, and very distinguished group of presenters and discussants. It is of critical importance that the Federal Reserve understands how changing labor supply in the United States will affect sustainable employment growth and the growth rate of potential output which, in turn, is the key factor driving improvements in living standards. I am sure that the results of this conference will better equip us to understand the changes which are to come.

Acknowledgments

The three editors thank the many people who contributed to the Federal Reserve Bank of Boston's June 2007 Economic Conference and to this book. In particular, we thank Patricia Geagan for superb management of the conference arrangements; she was ably assisted by Tom DeCoff, Donna Dulski, Jessamyn Fleming, Nancy Gillespie, Heather Honiss, Brendan Mackoff, Marie McGinley, My Lien Phung, and Ralph Ragsdale. Suzanne Lorant and Tyler Williams put together the substantial number of exhibits (tables and figures) in the volume. We are especially grateful to Elizabeth Murry, who skillfully guided the transformation of conference manuscripts and presentations into published chapters. Krista Magnuson provided expert copyediting and proofreading, while Sally and Steve Sztrecska did their usual superb job of typesetting the manuscript. Tobiah Waldron was impressive both for his speed and thoroughness in compiling the index.

1

Introduction

U.S. Labor Supply in the Twenty-First Century

by Katharine Bradbury, Christopher L. Foote, and Robert K. Triest

The American labor force will be transformed as the twenty-first century unfolds, a change that will confront policymakers and business firms with new challenges and new opportunities. The impending slowdown of labor force growth that will accompany the retirement of the baby boom generation already is playing a central role in national debates over the future solvency of Social Security and Medicare, as well as U.S. immigration policies. But labor supply changes will be influenced by other dimensions as well. In the coming decades, American workers are likely to be, on average, older and better educated than today's labor force. The globalization of labor markets is already opening new employment opportunities for some Americans and changing the wage rates paid to others. The production technologies and personnel policies adopted by tomorrow's firms will undoubtedly reflect the numbers and types of workers available for employment.

To explore the labor-supply trends that will affect economic policymaking in the twenty-first century, the Federal Reserve Bank of Boston chose "Labor Supply in the New Century" as the theme for its 52nd Annual Economic Conference held in June 2007. In analyzing these future trends, it is helpful to consider how these changes will affect both the *quantity* of workers in the U.S. labor force and the *quality* of their skills. In terms of the policy implications, the supply of American workers is of obvious importance to the Federal Reserve System, because the size of the U.S. labor force is a direct input to the Fed's estimate of the nation's potential economic output. Moreover, in the short run, the Federal Reserve needs to understand the various ways in which the quantity of labor

supply adjusts to changes in labor demand during business cycles. But while business cycle expansions and recessions exert powerful short-run impacts on labor-market outcomes, long-run living standards are determined by the quality of skills in the aggregate labor force and the types of human and physical capital that workers can use when performing their jobs. The worker-quality dimension to labor force trends, and the impact that these trends have on capital accumulation, are therefore fundamental to policies designed to raise living standards or to expand economic opportunity throughout the population.

The conference's six papers and its keynote address by Eugene Steuerle provide a broad overview of the quantity and quality implications of labor-supply trends. The first paper, by Bruce Fallick and Jonathan Pingle, carefully documents the extent of the upcoming slowdown in labor force growth due to the aging of the population. Alicia Munnell and Steven Sass, co-authors of the second paper, investigate the likely labor-supply behavior of older workers, including the baby boomers born between 1946 and 1964, the oldest of whom are now approaching traditional retirement ages, based on the behavior of past cohorts. The third paper, by David Autor, discusses the rising income inequality in the U.S. labor market over the last 15 years, and highlights the domestic and international forces that will affect wage inequality in the future. Robert E. Hall, author of the fourth paper, discusses how domestic labor supply adjusts to fluctuations in labor demand during the typical business cycle. The fifth paper, jointly written by Dale Jorgenson, Richard Goettle, Mun Ho, Daniel Slesnick, and Peter Wilcoxen, outlines the likely U.S. trends in both labor supply and labor demand up to 2030; the authors predict that these trends point to much lower growth in future aggregate output.

I. Quantity and Quality Dimensions of Labor Force Trends

Labor Quantity and the Aging of the American Labor Force

A recurrent theme in the conference sessions is the widespread impact that the aging of the U.S. population will have on the nation's workforce. Among the conference participants, there was a clear consensus that the demographics of aging will have a quantitatively important effect on the future size of the U.S. labor force. But how large might this aging effect

actually be? And how effective will public policy be in shaping and controlling the outcomes of this effect?

Population aging will reduce the aggregate labor force participation rate, defined as the fraction of all U.S. residents at least 16 years old who are either employed or actively searching for work. As Fallick and Pingle point out, increases in the fraction of the population reaching traditional retirement ages will reduce the overall labor force participation rate. But several other factors are working to increase the age at which American workers are likely to retire. Munnell and Sass contend that benefit changes in Social Security and private pension plans will reduce the adequacy of retirement income derived from these sources. These reductions will encourage some older individuals to keep working, at least part-time, beyond traditional retirement ages. Stanford Ross explores additional institutional changes that might reduce early-retirement incentives, while several conference discussants note that the labor force participation rates of older women may increase in the future. Throughout their lives, women in the baby boom generation have had higher participation rates that did earlier cohorts of women, and this trend might continue as baby boomer women enter their 60s.

Moreover, work lives could be prolonged by recent changes in the nature of employment and the characteristics of today's older workers. The physical demands of work that older generations routinely confronted have generally decreased; jobs in offices are typically easier, from a physical standpoint, than those performed in factories or on farms. The health of older Americans has also improved, so prolonging one's working life may be feasible even in more demanding positions. Policy changes might also prompt older individuals to increase their labor supply. Ross notes that the current legal framework is favorable to taking early retirement, but that this policy orientation could change to promote greater labor force participation among older workers.

Quantity effects stemming from the aging of the population may also have more subtle effects on the labor market. Autor hypothesizes that as the elderly's share of the population increases, so too will the demand for personal services, like home aides or healthcare workers, jobs typically performed by relatively less-educated workers. As a result, population aging may help shore up the low end of the U.S. wage distribution and

provide increased incentives for labor force participation among low-skilled workers. Munnell and Sass discuss a different effect that aging will likely have on the wage structure: the increased relative supply of older workers will act to decrease the premium paid for labor market experience.

An additional effect of population aging on the labor market will operate through the nation's social insurance programs. Even if some of the policy changes currently being contemplated are enacted, Social Security and Medicare will constitute a growing share of federal expenditures and national economic output. Payroll taxes may need to be increased to help reduce long-run structural deficits in these programs, potentially exacerbating tax distortions that affect labor supply. Benefit generosity may also be further reduced, amplifying the retirement income adequacy effects discussed by Munnell and Sass. The increasing size of social insurance expenditures may tie the hands of future government policymakers, leaving little room for increased expenditures on programs serving the young. This indirect effect resulting from population aging could have an important impact on the future U.S. labor market by limiting the extent of programs, such as early childhood education, designed to increase workforce quality.

Trends in Labor Quality

While the conference consensus is that demographic shifts will undoubtedly lead to slower growth in the quantity of labor supplied, the outlook for labor quality is more encouraging. An increase in workforce quality is expected to offset, to some degree, the quantitative decline in labor force growth. The American workforce is considerably more educated today than it was twenty years ago, and education levels among American workers are expected to continue rising. Higher educational attainment is encouraged by the faster wage growth that college-educated workers have recently experienced, a development that itself reflects the ongoing shift in labor demand toward the high end of the skill distribution.

The precise size of the expected "quality offset" to lower labor-supply growth is difficult to gauge. As Autor documents, the pace of increase in college enrollments and completion rates has slowed in recent years. Since, as Autor says, "the gap in college attendance by parental income, race, and ethnicity remains large," there is room to increase the overall quality of the U.S. labor force in at least two ways. First, policies should facilitate pre-school investments in human capital, and second, policies should promote college enrollment among low- and moderate-income families. Gary Burtless, one of the paper's discussants, explores the reasons that men in the United States have lower college completion rates than men in other rich nations. He suggests that less-affluent parents have few resources available to influence their children to take a far-sighted view of the future payoffs that higher education brings. Hence, children from these families are less likely to make the investments of time and money that obtaining a college degree requires. In addition, Autor notes that immigration policy can provide an additional lever to further raise the skill and education levels of the U.S. labor force.

By incorporating labor quality into their model of the U.S. economy, Jorgenson, Goettle, Ho, Slesnick, and Wilcoxen account for changes in both the education and experience of the U.S. labor force. The authors assume that the educational composition of the U.S. population will eventually stabilize. This assumption leads labor quality to continue to rise for some time, but with the increases gradually diminishing during the next 25 years. According to this model, the quality-adjusted effective labor force continues growing more rapidly than the working-age population, but only modestly so. Even after accounting for increasing quality in the American workforce, the authors are only slightly more positive than Fallick and Pingle—who account only for quantity changes—about potential growth in the U.S. economy in the next several decades.

II. Conference Summary

Session 1: The Outlook for Labor Supply in the United States

The first session's paper provides an overview of the effects of demographic change on aggregate U.S. labor force participation, assessing shifts in age and gender mix as well as historical and possible future changes in participation rates for age by gender subgroups. Bruce Fallick and Jonathan Pingle argue that a key factor driving aggregate changes in labor force participation in coming years will be the evolution of the age distribution of the population—specifically, the movement of the

distribution's baby boom "bulge" from older working ages into their retirement years, traditionally assumed to start at age 55 and beyond. The authors note that increases in within-group participation rates can offset some of the downward pressure from population aging, but avoiding a decline in the aggregate labor force participation rate would require very substantial—and unlikely—increases in participation rates across all age groups.

The age mix of the U.S. population will change substantially during the next 35 years. For example, the Census Bureau projects that the portion of individuals aged 35 to 44 years will shrink from about 18 percent of the population to about 15 percent, while the fraction of those aged 65 years and older will rise from about 16 percent to 25 percent. Applying current (2005) participation rates by age and gender to these shifting age shares, Fallick and Pingle provide a simple forecast of the aggregate U.S. labor force participation rate that declines from about 66 percent in 2005 to about 63 percent around 2020, and to less than 60 percent by 2033 and years thereafter. The authors then compare population share projections made by the Social Security Administration with those forecast by the Census Bureau, and note that taking account of increased longevity, especially among older women, causes the Census Bureau's population projections to imply a lower aggregate U.S. labor force participation rate. Similarly, differing assumptions about future immigration also change the projected age mix of the population, as immigrants are typically concentrated in age groups with high labor force participation rates; in addition, immigrants may be more likely to be labor market participants than native-born individuals, conditional on age. Fallick and Pingle simulate the effects of various immigration assumptions, and note that even fairly substantial increases in immigration, accompanied by above-average participation by immigrants, only modestly offset future declines in aggregate labor force participation rates. They note, however, that the effects of changing immigration flows or life expectancy on the *size* of the labor force (calculated as the participation rate multiplied by population size) will be more positive than these effects on the aggregate participation rate.

With their baseline projection of a 6 percentage point drop in the aggregate participation rate over the next 30 years, Fallick and Pingle next

examine the possible effects of changes in specific age-by-gender group labor force participation rates. They compare the participation rate projections of the Bureau of Labor Statistics, the Social Security Administration, and their own cohort-based model. In terms of the gender-and-age interactions, a key question is whether older Americans, or a subset of older Americans, will begin to work longer than similar age groups did in the past. Most forecasters expect that better health, increased life expectancies, changing preferences, or changing inducements provided by government and business will lead many older Americans in coming decades to remain in the labor force longer. Such a change may have important effects on aggregate participation because this shift would occur in a fast-expanding segment of the population; indeed, Fallick and Pingle indicate that increases forecasted in participation among older women, meaning those aged 65 years and above, would offset roughly one-quarter of the projected total decline in labor force participation attributable to aging. For prime-age workers, those men and women who are 25–54 years old, the three forecasts differ substantially: the Bureau of Labor Statistics estimate is the most optimistic about future prime-age increases, while the Social Security Administration's forecast is in the middle of the pack, and the Fallick and Pingle model predicts a continuation of current participation rate trends, which reflect long-term declines in the participation rates of prime-age American men, and recent declines for prime-age American women. Teenagers, while a small fraction of the U.S. population, have contributed substantially to recent aggregate declines in labor force participation, and also to differences among the three forecasts. All in all, these forecasts of aggregate labor force participation "diverge noticeably," according to Fallick and Pingle; nonetheless, all three indicate that "likely" changes in the participation rates of various subgroups will only partially offset the aging-related declines in aggregate participation rates or even, according to the authors' own projections, possibly exacerbate these declines.

Fallick and Pingle examine potential policy changes, such as increases in Social Security's "normal retirement age" and "delayed retirement credit," which might alter the likely future path of labor force participation among older Americans. While such changes could have important effects on future participation rates, they note that policy changes for

Social Security have been gradual in the past, and are impossible to forecast in the future. Fallick and Pingle conclude their analysis by predicting that "the outlook is for slower growth in U.S. labor supply from 2007 onward than was the norm in the 1965–2000 period."

The paper's first discussant, Chinhui Juhn, takes issue with the Fallick and Pingle forecast of older Americans' future labor force participation rates. She argues that increased participation by older women, through complementarity of spouses' leisure, may have accounted for over one-third of the increase in older men's participation rate during the 1996–2006 period—and that this effect may manifest in similar ways in the future. Furthermore, Juhn notes that recent and ongoing declines in employer-provided retiree health insurance are likely to push up the participation rates of the 55–64 year old age group nearing the Medicare eligibility age, as well as those for adults aged 65 years and older. She questions Fallick and Pingle's forecast of continued declines in prime-age women's labor force participation, noting that steeper recent declines in the participation rates of never-married women and women without children provide "little evidence that the trend among married mothers—the group that fueled the increases in the earlier decades [the 1970s and 1980s]—has actually reversed and begun to decline."

Instead, Juhn argues, the pervasiveness of recent declines in participation rates among prime-age women points to time effects—specifically weak labor market conditions—not to cohort effects. She then explores what she views as a problem with the cohort-based model that underpins the Fallick and Pingle forecasts, specifically the assumption that the coefficients on time-varying variables such as cyclical factors, education, fertility, and marriage are constant over time. Juhn contends that it is likely that these coefficients are shifting, introducing error into the estimated cohort effects and hence into the forecasts. Her final comment notes that this paper's topic, the future size of the U.S. labor force, is very important because, in conjunction with labor productivity, the size of the nation's workforce affects the economy's potential growth. But to fully understand the predicted declines in the size of the labor force and these ramifications, she argues that we need to pay more attention to the *types* of labor that are forecast to shrink, and we must also sort out the degree to which current labor force trends reflect cyclical supply versus demand shifts.

Lisa Lynch, the second discussant, focuses on two aspects of the Fallick and Pingle paper: their analyses of likely changes in immigration and in labor force participation by prime-age women. She questions whether the authors' simulations of the impacts of "high" immigration—adding 200,000 immigrants annually beyond current flows—is actually so high. Lynch appeals to the Congressional Budget Office's recent estimates that proposed immigration legislation would add 180,000 workers each year for the next ten years, and she also notes that immigrants tend to participate in the labor force at above-average rates, conditional on age. In addition, Lynch contends that even though immigrants' effects on aggregate participation rates are muted by their addition to both the numerator and denominator, they do increase the overall labor supply; furthermore, immigration policy may shift the skill mix of immigrants, with corresponding effects on their contributions to economic growth.

She challenges Fallick and Pingle to ask *why* men's participation rates have been declining before assuming these declines will continue, and raises a number of questions about the authors' assumption that prime-age women's participation rates will decline in coming years. First, the fact that women's *level* of participation is not equal to that of men raises questions about why the pace of decline in women's participation should be similar to men's. Second, labor force participation is positively associated with educational attainment, and since the early 1990s women's college-enrollment rates have been rising faster than men's. In addition, Lynch argues that technological changes have contributed to increases in women's age at the birth of their first child, with implications for their labor force attachment before becoming mothers as well as likely lifetime participation patterns. Finally, like Juhn, Lynch notes that interest in future labor force participation—at least among monetary policymakers—centers on its implications for future output growth; in this context, the increased average "experience" of an older workforce should enhance labor productivity and more than proportionally add to output gains.

Session 2: The Labor Supply of Older Americans

Alicia Munnell and Steven Sass examine the labor supply of older Americans, a group they define as those aged 55 years and up, paying particular attention to men in this age group. Munnell and Sass point to several

changes in the American retirement income system that are making it necessary for people to continue labor force activity at later ages than was the norm among older cohorts in the second half of the twentieth century. Social Security's Full (Normal) Retirement Age has been gradually increasing from 65 to 67, resulting in an effective cut in benefits for workers retiring at any given age. In the coming years, rising Medicare premiums and increased taxation of Social Security income are virtually certain, and further cuts in Social Security benefits to restore the System's long-term fiscal balance are possible. Private pensions have also been transformed over the last two decades, with defined contribution plans, such as 401(k)s, largely replacing traditional defined benefit plans. Unfortunately, to date workers' defined contribution account balances have generally fallen short of the funds needed to generate the retirement income amounts typically provided by defined benefit plans. Personal non-pension savings, the third leg of the retirement income stool, have diminished in recent years and cannot make up for shortfalls in public or private pension plan funds. Munnell and Sass conclude that if people are going to maintain their living standards in retirement, then retirement ages will need to be increased; in other words, Americans will need to participate longer in the workforce than has traditionally been the case.

Until recently, the trend was toward retirement at younger ages. The increased affluence accompanying economic growth during the twentieth century made it feasible for people to spend an increasingly long period of time toward the end of their lives enjoying more leisure, a phenomenon reinforced by early retirement incentives in many private defined-benefit pension plans and by the availability of Social Security benefits starting at age 62. The trend toward older men aged 55–64 years taking early retirement ended in the mid-1980s, and since then has been partially reversed. Munnell and Sass attribute this reversal to several factors, including changes in Social Security, changes in employer-sponsored pensions, the elimination of a mandatory retirement age, shifts in employment toward less physically demanding jobs, and the decreasing availability of employer-provided retiree health insurance.

Accompanying the trend toward later retirement, Munnell and Sass document the changing employment patterns of older workers. In the United States, older workers are now less likely to remain with a long-

term employer as they approach retirement than was previously the case. The decrease in job tenure among older workers places these workers at increased risk of job displacement. Munnell and Sass show that older workers are generally at low risk of displacement, but this is because older workers are more likely than young workers to have accumulated substantial tenure on their jobs. Yet age alone does not protect workers; holding tenure constant, older workers are actually at greater risk of displacement than are young workers.

Increased life expectancy and improved health among the older U.S. working-age population, along with decreasing physical demands in the workplace, bode well for the ability of most people to extend their careers beyond retirement ages that were typical in the past. However, Munnell and Sass note that there is still a significant share of Americans in their late 50s and early 60s who would find continued participation in the workforce difficult due to poor health or disability. Other obstacles to extended work lives they discuss include many employers' resistance to part-time employment and the continued early availability of Social Security benefits at age 62. On net, however, Munnell and Sass conclude that labor force participation rates of men in their late 50s and 60s are likely to continue to increase.

In his comments, Robert Hutchens notes that the Munnell and Sass paper is very relevant to the current policy debate on how government transfer programs serving older Americans should be reformed. If older Americans can easily find decent jobs and remain physically capable of working, then solving Social Security's fiscal problems by reducing the growth of benefits below that mandated by current law would be much more palatable than it would be under an alternative scenario, where many older Americans are either unable to find work or become too frail to work as they age.

Hutchens believes that we need additional information in order to answer two important questions regarding the changing labor market for older people. First, how do the trends in labor force participation differ by educational attainment? Less-educated workers might be suffering disproportionately from changes in Social Security and private pensions, but also likely have job prospects that are less favorable than are those for individuals with more education. Second, how will employers

respond to the increased supply of older workers? Will firms be able to accommodate older workers' changing needs and allow them to stay in long-term career jobs, or will older people largely end up in a spot market for non-career jobs and lose the job-specific human capital that they had previously accumulated?

Joyce Manchester focuses her comments on two reasons for being optimistic regarding the prospects for delayed retirement and increased retirement income: trends in the early claiming of Old-Age Social Security benefits, and older women's labor force participation. Manchester presents evidence from the Social Security Administration's records showing that people who claim benefits before age 65 tend to have weaker attachment to the labor force, and lower earnings, before claiming benefits than do those who claim at 65. The full retirement age for Social Security benefits is being gradually increased from 65 to 67 under current law, and the penalty associated with claiming benefits at 62 rather than at the full retirement age is increasing. Manchester shows that this change in Social Security appears to be discouraging benefit claims at age 62, with the implication that we can expect further delays in retirement as the Social Security full retirement age continues to increase.

Turning to the issue of the labor force participation of older women, Manchester notes that women nearing retirement age increasingly do so after spending much of their adult lives in the labor force. As a result, older women are more likely to receive Social Security benefits based on their own earnings record, rather than receiving spousal benefits, and will also be more likely to have their own private pensions and health insurance than were women in earlier cohorts.

Session 3: How Structural Shifts in Labor Demand Affect Labor Supply Prospects

David Autor's paper explores the interaction of labor demand and labor supply, and its implications for the wage structure and future labor supply responses in the United States. Autor documents that the widely recognized growth in U.S. earnings inequality can be usefully divided into two stages. During the 1970s and 1980s, real wages fell at the bottom of the earnings distribution and rose moderately at the top. In contrast, during the 1990s and early 2000s there was strong growth of real wages

at the top of the distribution, and modest real growth in the bottom tier, with the middle of the wage distribution experiencing the least income growth.

Autor then turns to the question of how shifts in labor demand likely influenced these changes in the U.S. wage structure. A simple model that divides workers into high school graduates and college-educated equivalents does well in explaining changes in wages from 1963 through 1992, and implies that there was a strong shift in demand toward college-educated workers. However, a more complex model is needed to explain the data from 1993 onward, the period in which Autor argues that wage growth "polarized," with wages at the top and bottom of the distribution rising faster those in the middle.

Autor considers a number of candidate explanations for this polarization. Decreases in the real value of the minimum wage, which some studies have found to be a major source of growing income inequality in the United States, cannot explain why wages in the middle of the distribution have stagnated more than those at the bottom. Autor believes that a more promising explanation is that technological change and off-shoring have increased the demand for strong cognitive and interpersonal skills typical of highly educated professionals and managers, and decreased the demand for routine analytical and mechanical skills typical of middle-tier workers. Technological change and off-shoring have relatively little impact on low-level service jobs, which are currently difficult to automate or trade across international borders. Autor points out that the aging of the U.S. population will likely increase the demand for such services in the future, as may growing demand for services by high income households. So, continued polarization of the U.S. wage distribution seems likely.

How might labor supply respond to the recent changes in wage structure? Autor points to evidence showing that barriers to college attendance by youth from low- and moderate-income families remain substantial, and suggests that reducing these barriers would help increase the supply of highly educated workers, and attenuate further increases of inequality at the high end of the wage distribution. Liberalized immigration policies for highly educated workers would also help to reduce high-end inequality.

In his comments on Autor's paper, Jared Bernstein takes issue with the augmented version of the skill-biased technical change explanation for

growing inequality put forth by Autor. While applauding the movement away from the less-nuanced version of the skill-biased technical change story, Bernstein still has doubts about the centrality of skill bias in explaining changes in the earnings distribution. Although he does not deny the existence of complementarity between technology and skills, Bernstein cites his own research with Lawrence Mishel that raises doubts regarding whether the pace of skill-biased technical change has accelerated sufficiently over time to explain the growth of earnings inequality. Bernstein believes that a promising alternative, or perhaps complementary, explanation is change in economic policies and institutions, a hypothesis put forward in recent research by Frank Levy and Peter Temin. However, Bernstein notes that there is no "smoking gun" evidence for the centrality of either skill-biased changes in technology or changes in institutions as an explanation for growing inequality. The policy implications of the two explanations differ significantly—enhanced education and job training are the primary policy tools to address skill-biased technological change, while if institutional change is the primary source of growing inequality, a broader set of policies governing labor relations is prescribed.

Gary Burtless also notes that institutional change has played a role in the growth of inequality. Decreases in private sector union coverage and shifts in pay-setting norms have contributed to growing inequality. Most of Burtless's comment, however, focuses on the supply side of the labor market. He notes that although there were substantial increases in the educational wage premium during much of the time since 1980, there have been only relatively modest increases in college attendance and degree attainment. Men, in particular, seem to have barely responded to the increased economic rewards that accrue to postsecondary schooling. Burtless questions why educational attainment has responded so weakly to the increase in its economic return. He presents data showing that most other OECD countries have experienced substantial increases in postsecondary schooling over the past 20 years, while the United States has not. One possible explanation explored by Burtless is that many low- and moderate-income families in the United States lack the means to effectively push their children to attend college. Paying for their children's tuition and fees may seem unaffordable to such families, and they may not be fully informed about the possibility of financial aid. In many other high-income countries, in contrast, students face little or no out-of-pocket tuition expenses. This hypothesis is consistent with data showing that recent increases in college attendance have been concentrated among young adults from relatively affluent families.

Session 4: The Cyclical Sensitivity of Labor Supply

In the conference's fourth paper, Robert E. Hall analyzes the responsiveness of labor supply to business-cycle influences. He starts by noting that macroeconomists generally believe that short-run business-cycle movements in wages and employment stem primarily from changes in labor demand, not labor supply.[1] The macroeconomist's task is to learn how changes in labor demand interact with workers' labor-supply preferences to generate observed changes in wages and employment. For example, if labor supply is inelastic, workers are relatively willing to accept wage cuts in order to keep their jobs. A decline in labor demand will then bring about a sizeable drop in wages and little change in employment. On the other hand, a more elastic model of labor supply would generate large changes in employment and small changes in wages when labor demand declines. This latter, elastic pattern is closer to what we see in the macroeconomic data, but explaining it runs into a fundamental problem: most microeconomic studies find that workers are relatively unwilling to substitute work for leisure (or vice versa) when wages change. This unwillingness makes their individual labor-supply schedules inelastic, not elastic. Given the inelasticity of the labor-leisure trade-off among individual workers, it is hard to see how changes in aggregate labor demand could raise or lower employment for the entire economy.

To solve this puzzle, a number of economists are working to expand the labor-leisure trade-off assumed in traditional labor supply models. These researchers argue that a third activity—job search while unemployed—should be added to the worker's list of potential uses of time. If unemployed workers must take time to search for new jobs, then there will be a pool of potential workers that could be added to the employment ranks when labor demand increases (or subtracted from the measure of employment when demand falls). This addition would make "labor supply" for the aggregate economy more elastic than that of individual workers. For job-search considerations to be important, however, the size

of this unemployed pool of workers must be sensitive to labor demand. Total labor input must change in recessions and expansions primarily due to changes in the unemployment rate, not because more workers are choosing to participate in the labor market (the participation margin) or because workers who are working decide to spend more or less time on their jobs (the hours margin). The central goal of Hall's paper is to show that unemployment is sufficiently sensitive to labor demand for this to be the case.

Hall starts with a careful examination of the data he will use for this exercise, discussing how employment, hours, and participation are measured in the United States. He argues that the Current Population Survey's measure of employment generated by a survey of households is preferable to a measure generated by a survey of firms conducted by the Current Employment Statistics program (sometimes called the Establishment Survey). Using the household-based measure of employment is preferred because this measure is more likely to be consistent with the hours and participation measures, which also come from the household survey. Hall's next step is to obtain a direct measure of shifts in labor demand. He does this by using statistical theory to extract a common measure of labor-market cyclicality from three fundamental correlates of the health of the labor market: unemployment, average weekly hours, and real personal disposable income per capita.

Once this measure of labor demand is constructed, Hall measures its correlation with the three potential margins on which labor input can adjust over the business cycle: unemployment, participation, and hours. He finds that about 56 percent of the variability in total labor input is due to fluctuations in unemployment, with 12 percent coming from the participation margin and the remaining 32 percent coming from the hours-per-worker margin. The results suggest that unemployment is sufficiently cyclical to support the new theoretical work on job search as a formal alternative use of time in models of the business cycle. Hall concludes that: "More than half of the extra labor input in a cyclical upswing is drawn from the ranks of the unemployed. No model of the cycle in the labor market can claim any realism unless it takes this finding seriously."

Katharine Abraham begins her discussion of Hall's paper by reviewing three potentially important characteristics of his new labor-demand index. First, the index is based on common variation in the fundamental correlates (weekly hours, real personal disposable income, and unemployment). The common business-cycle variation in these correlates is assumed to be unrelated to their long-term trend movements, which are modeled as separate fourth-order functions of time. Yet detrending the data with fourth-order trends (rather than some other type of detrending procedure) may also remove useful business-cycle variation from the data. Second, using real personal disposable income as one of the fundamental variables may contaminate the measure as a pure labor-demand index, because this income variable also includes proprietors' income and government transfers, in addition to wage payments. Third and perhaps most importantly, Abraham says that Hall may want to use total hours as one of the fundamental correlates, rather than using both weekly hours and the employment rate separately. Hall's procedure allows the *common* variation in weekly hours and employment to contribute to his labor-demand index. But firms probably think about variation in hours and employment as distinctly different margins of adjustment, because hiring a new worker includes fixed costs (such as health insurance). These additional costs for new hires are not accrued when varying the hours of workers whom a firm already employs. As a result, the common variation in hours and employment does not have a clean interpretation.

Abraham then discusses how Hall's labor-demand index correlates with the three potential margins of adjustment (unemployment, per-worker hours, and participation). She notes that these margins may be correlated with the index at various time lags, which are not accounted for in Hall's analysis. Also, he finds the correlation between per-worker hours and the labor-demand index to be surprisingly large; this may reflect Hall's treatment of hours and employment as separate fundamental variables, as noted earlier. Finally, Abraham agrees that the difference between employment totals from the Bureau of Labor Statistics' separate payroll and household surveys remains a puzzle. The answer may involve a better understanding of business-cycle variation in "off-the-books" employment, or a rethinking of the population weights in the household survey.

Susanto Basu starts his comments by outlining the differences between the workhorse Keynesian model of the labor market and the newer search-based framework to which Hall is contributing. The Keynesian

view argues for pervasive "stickiness" in both prices and wages over the business cycle. Because both variables are about equally sticky, as a result there is little cyclical variation in real wages in the standard Keynesian model. The approach favored by Hall instead argues that real-wage inflexibility results from search considerations. Once a firm and worker make a good match, they are loath to destroy it simply because the real wage has not changed in exactly the way that a frictionless model of the business cycle would imply. As of yet, there is little evidence in the data to distinguish between the older Keynesian approach and the newer search-based approach. Both models not only accept the presence of involuntary unemployment (unlike more neoclassical models), but also explain this unemployment as a function of inflexible real wages. Only the explanation for real-wage rigidity differs in the two models. Eventually, distinguishing between the Keynesian and search-based frameworks in the data will likely result in a better understanding of how these models respond to specific types of shocks.

Basu then takes issue with one assumption of Hall's exercise; namely, that most of the observed short-run changes in hours and unemployment are due to shifts in labor demand rather than in labor supply. Basu pointed out that changes in government spending typically raise employment, but this cannot be because the labor-demand curve shifts to the right, since there have been no corresponding changes in the fundamental determinants of labor demand (such as the amount of capital per worker, or total factor productivity).

Session 5: Labor Supply and Labor Demand in the Long Run

The fifth paper of the conference was written by Dale W. Jorgenson, Richard J. Goettle, Mun S. Ho, Daniel T. Slesnick, and Peter J. Wilcoxen (henceforth JGHSW). The paper presents a formal model of future U.S. labor supply and demand, derived from a fully specified neoclassical growth model. At its heart, the model presented is the famous growth model developed by Nobel Prize winner Robert Solow in the late 1950s. It is meant to be a long-run characterization of an economy, so it assumes away the short-run business cycle movements that lie at the heart of the analysis in the previous paper by Hall. The neoclassical growth model focuses instead on the determinants of potential long-term growth, which

are the rate of population growth and the rate of technological progress. From these variables, a long-run rate of capital accumulation is endogenously determined. Finally, the contributions of population growth, technological progress, and capital accumulation are added together to obtain the growth rate of potential output.

This characterization of the neoclassical growth model would be familiar to most undergraduate economics students. But the JGHSW paper delves deeply into the inner workings of the neoclassical model to provide more specific predictions about where the U.S. economy is headed. For example, the authors do not merely assume that the amount of labor services available to produce output is simply proportional to the working-age population. Rather, the authors use various measures of "labor quality" (such as education and experience) to adjust the labor input according to how productive it is likely to be. Along the same lines, the model does not assume a single overall rate of technological progress for the entire economy. Instead, it uses sophisticated econometric tools to measure the rate of technological progress in a number of individual industries, and then models how these rates are likely to change in the future. Additionally, the techniques allow the authors to determine whether technological progress in any specific industry is likely to be "labor using" or "labor saving."

The central message of this bottom-up forecasting approach to predict future output is that the U.S. economy will grow much more slowly during the next quarter-century than it has since 1960. Overall GDP growth will slow from 3.2 percent in the 1960–2004 period to 1.6 percent in the 2004–2030 period. In large part, this decline is driven by well-anticipated declines in the growth of labor input. Labor services grew at an annual rate of 1.73 percent in the 1960–2004 period, but will grow at less than half this annual rate (0.74 percent per year) from 2004 to 2030. Importantly, the decline in labor input growth is less severe than the decline in the growth of the working-age population, because labor quality will continue to rise, albeit more slowly than in the past. The rate of technological progress will be essentially unchanged across the two periods, slightly less than 0.50 percent per year. "In summary," the authors write, "the potential growth of the U.S. economy will be slowing considerably between 2004 and 2030, and monetary policy will have to adapt to the new environment."

Richard Berner begins his discussion by noting that the gloomy productivity predictions in the JGHSW paper may turn out to be too pessimistic. Projecting future productivity growth is exceptionally difficult, in large part because of data limitations. As an example, official measures of productivity in the construction industry imply that it has been on a decades-long *decline*, which is hard to believe given the boom in construction that has characterized much of this period. Part of the reason that measuring productivity is so difficult is that the productivity data are based in part on compensation data, and the Bureau of Labor Statistics' procedures may not be keeping up with the changing ways that workers are rewarded in the modern American labor market.

Berner then turns to the authors' assumptions about future labor supply. He notes that labor supply among older workers appears sensitive to government policy; as an example, many older workers stay on the job just long enough until they qualify for Medicare. Changing the eligibility age for Medicare would therefore change the labor supply of older workers. Policies on immigration also affect the number of available workers in ways that the JGHSW paper cannot predict. Berner concludes by noting that the globalization of labor markets will undoubtedly affect returns that workers can expect in the U.S. labor market. Globalization will thereby influence labor-force participation decisions among American residents. Without understanding the relationship between wages and various components of globalization (such as outsourcing), it is hard to predict future labor-supply behavior. Indeed, looking around the world, we find very different labor-supply behaviors even within the small group of affluent industrialized/advanced economies; France chose to enjoy more leisure as its productivity levels rose, while the United States chose to work more.

Eric Brynjolfsson's discussion starts with two comments about the stability of the relationships on which the JGHSW paper is based. First, the model requires some predictions for future consumption patterns, which are essentially extrapolated from consumption patterns today. But some items in today's consumption bundles, which contain cell phones and iPods, were not even around 25 years ago. How can we possibly know what people will buy in 2030? A second concern about the model's stability echoes Berner's comments: it is difficult even to measure past or current productivity, let alone predict it decades from now. Official productivity figures show big changes from year to year, and even productivity levels for the same year can change dramatically as different revisions of the data are made. While applied researchers must use the productivity data currently available, the intense volatility and large revisions in these data should caution us to take productivity predictions with a grain of salt. Brynjolfsson then pointed out that the authors face difficult statistical issues that are common to many economic models. The authors make extensive use of prices (as well as quantities) to infer changes in technology, so they face classic issues of simultaneity: if a price drops, is this because supply increased, or demand declined? The authors use well-established statistical techniques to deal with this issue, and they verify that these techniques are appropriate. But some results in their paper suggest potential problems, including their finding that particular inputs appear to be used more intensively when their prices rise.

Brynjolfsson then turned to the paper's assumption that business productivity increases as soon as new machines are installed. His own research has shown that for most firms, the installation of a new information technology (IT) system is just the "tip of the iceberg" in improving productivity. After a new IT system is purchased, firms must then train their employees to use it. Occasionally, they have to rework their entire mode of operation to make the best use of the new technology.

The authors' pessimistic assumptions seem hard to square with Brynjolfsson's intuitive understanding of how tomorrow's businesses will make use of future technology. In a matter of decades, advances in computing power will allow machines to mimic and perhaps surpass the computational power of the human brain. He concludes by noting that as various computing thresholds are reached and then exceeded, technology will probably play an even bigger role in the productivity of the American economy than we can predict today.

Current and Future Challenges for Policy and Research

A wide-ranging panel discussion addressed various ways to increase the quality and versatility of the U.S. workforce as one offset to slowing growth in the quantity of prime-age labor supplied in the coming years. At the upper end of the working age distribution, however, as pointed out

repeatedly in earlier sessions, the private sector behavior of individuals and firms is highly influenced by the current rules regarding the eligibility ages for Social Security and Medicare. Thus, it seems that the government will have to drive changes in behavior by instituting changes in these programs.

Public Policy and the Labor Supply of Older Americans

Stanford Ross argues that the current legal framework in the United States is highly favorable to early retirement, and addresses the issue of whether and how the Social Security laws, U.S. tax laws, and laws governing private pensions and individual savings can be changed to provide fewer incentives to retire early and more encouragement to work longer. Ross suggests that the Social Security system was not intentionally designed to favor early retirement, but these incentives developed from "almost random" political decisions and changing circumstances. While several changes made recently have reduced the incentives to take early retirement, the changes are so gradual that potential retirees are essentially unaware that the incentives have shifted. Furthermore, once these changes to Social Security's normal retirement age and delayed retirement credit are fully in place, the regime will continue to be highly favorable to early retirement.

Ross says that while it would be possible to speed up some of the legislated transitions and change the benefit formula to provide enhancements for continued work, the major change that could make a difference would be to move the early eligibility age from 62 years to 65 years. Such a change would almost necessarily have to be part of a larger package of changes along many dimensions (such as raising minimum benefits for lower-wage workers and/or providing tax credits to employers of older workers), and probably would entail further increases in the normal retirement age. To be politically feasible, a comprehensive package would need to reflect bipartisan efforts, which is currently unrealistic.

Even without a major overhaul, however, the government might enact marginal changes that will affect retirement incentives, particularly by acting on health care. Beyond that, Ross argues that economic changes are more likely to influence the behavior of individuals and firms than are changes in laws. For example, if the economy falters and wealth prospects are diminished, resulting in workers feeling more insecure about their retirement, that could provide an impetus for working longer. On the employer side, if major labor shortages emerge, perhaps because immigration is curtailed and outsourcing is restricted, firms may need to adjust to a diminished labor supply by taking steps to employ older workers. But currently, neither employers nor aging workers seem particularly motivated to seek such changes.

The Seven Deadly Sins in Aging Policy and Research: A Cautionary List for Policymakers and Prognosticators

C. Eugene Steuerle, the conference's keynote speaker, provides a useful alternative perspective by addressing some potential inadequacies of current policy-relevant research on population aging. By casting these issues as the "seven deadly sins," Steuerle establishes a counterpoint to the rest of the conference sessions that offers a fitting end to this summary and book. Steuerle reminds us that while anticipating the future and planning for it as best we can is important, we may be eliding some considerations along the way. In other words, a little less hubris and a little more humility may be in order as we grapple with how the U.S. labor supply may unfold in the following decades.

Steuerle contends that the first "deadly sin" of aging policy research is paying "too little attention to the labor side of the aging debate." He believes that changes in retirement behavior have received too little attention as a potential solution to economic problems associated with population aging. The second sin is "policymaking without any real targets." The fundamental objectives underlying policy proposals such as preserving the current Social Security system or creating individual accounts within this system, as the Bush administration proposed in 2005, are scarcely discussed and analyzed. The third deadly sin is "limiting the debate so as to be politically correct." As an example, Steuerle cites the structure of family benefits embedded in the current Social Security framework that are at odds with contemporary social realities in the United States. Most observers would agree privately that the current system makes little sense, but would shy away from proposing bold reforms because of potential political controversy.

The fourth deadly sin is misuse of the term "aging." Steuerle notes that we persist in measuring the old-age threshold as starting at a fixed number of years, such as age 65, even though people at given ages today generally are much healthier and have longer life expectancy than did their counterparts in past generations. If instead we gauged the concept of old age as corresponding to a given remaining life expectancy, then we would have a very different perspective on the problems associated with population aging. The fifth deadly sin is "ignoring the balance sheet." As an example, Steuerle notes that calculations of the effect of changes in retirement behavior commonly extend to changes in Social Security benefits and payroll tax revenues, but rarely, if ever, extend to the effects on national output and general tax revenues. The sixth deadly sin is "assuming away arbitrary aspects of the status quo." Although status quo policies are often arbitrary or accidental, policy analysts tend to view these as having resulted from rational policymaking decisions taken when these policies were implemented, and display reluctance to recommend radical changes. A recurring idea throughout the conference has been the need for bold new thinking on these issues.

Finally, the seventh, and according to Steuerle, the most deadly sin of all is "hubris about knowing the future." In the aging field, we tend to design our policies to fit our current views at the future expense of preventing "our children from following other visions for how their society evolves." By putting in place rigid programs that promise transfer payments into the indefinite future, we essentially tie the hands of future policymakers. This runs the risk of binding future generations to policies inappropriate to their situation and values.

* * *

It is clear that the coming few decades will be accompanied by major changes in the U.S. labor supply and pose challenges to the U.S. economy, particularly in connection with the baby boom generation's transition into what traditionally have been considered retirement years. How the United States deals with the implications of these changes will help set the nation's economic and political course for the twenty-first century. This volume, consisting of papers first presented at the conference and then revised for the book, showcases some of the important considerations involved in addressing these trends, and the opportunities that may arise if we confront these challenges creatively and forthrightly.

Note

1. The primacy of labor-demand shocks as drivers of business cycles implies that employment declines in recessions because firms find it less profitable to hire employees, not because workers suddenly decide to work less. To use more formal economic language: Hall assumes that in a supply-and-demand model of the labor market, the labor-demand curve shifts along an unchanging labor supply-curve when a recession or expansion occurs.

2

The Outlook for Labor Supply in the United States

The Effect of Population Aging on Aggregate Labor Supply in the United States

Bruce Fallick and Jonathan Pingle

I. Introduction

Output growth is determined by growth in labor productivity and growth in labor input. Over the past two decades, technological developments have changed how many economists think about growth in labor productivity. However, in the coming decades, the aging of the population will change how economists think about the growth in labor input in the United States. As the oldest baby boomers born in 1946 turned 50, then 55, and then 60, an important economic change has slowly surfaced: these people have become less likely to participate in the labor force. While this shift was obscured by a labor market slump in 2002, the aging of the American population began to put downward pressure on aggregate labor supply, marking the start of what is likely to be a sharp deceleration in labor input that will last another half-century.

Beginning in the mid-1960s, movements in women's labor force participation dominated all other influences, pushing up aggregate labor force participation rates sharply over the 25 years since 1965, until leveling off after about 1990. From 1963 to 1991, the average annual growth in the U.S. labor force was 2 percent a year. In coming years, the dominant influence will likely be the evolution of the population's age distribution. The Social Security Administration projects labor force growth will slow to 0.5 percent a year by 2015 and to 0.3 percent a year by 2025[1]—a striking deceleration from the aforementioned 2 percent annual pace. Since 1995 the population bulge then comprising people in their 30s has moved on to older age groups that are associated with lower labor force participation rates. Barring an enormous change in the participation behavior

of older Americans, the shift of the population distribution away from prime working-age adults will put significant downward pressure on the nation's labor force participation rate. In fact, absent other changes, the aging of the U.S. population has the potential to undo the increases in participation rates brought about earlier by the increased entry of women into the labor force.

Estimates of labor supply growth also depend partly on uncertain population projections. Both the U.S. Census Bureau (Census Bureau) and the Social Security Administration project that growth in the population aged 16 years and over will slow to about 1 percent in 2009, from the 2005 pace of approximately 1.2 percent. However, there are subtle differences in these two projections that, combined with differing participation rate projections, act to widen the differences between various projections for aggregate labor supply growth.

As the changes in the U.S. age distribution and population growth unfold, it is unclear how the within-group trends will adapt to offset—or perhaps evolve to exacerbate—the declines induced by the aging of the population. In the last five years, in fact, within-group trends appear to have mostly exacerbated the decline in labor force participation.[2] Even abstracting from how within-group participation rates have evolved in recent years, the shifts in the age distribution have lowered the labor force participation rate by nearly 0.4 percentage points between 2002:Q4 and 2006:Q4.[3] Going forward, however, the within-group changes in participation rates are the crucial element in determining the extent to which U.S. population aging may depress aggregate labor force participation.

A population-weighted average of individual age group participation rates indicates that aging has lowered aggregate U.S. labor force participation during the 2002–2006 period. However, the level of the labor force rose 1.8 percent between December 2005 and December 2006, as the participation rate moved up from 66.0 percent to 66.4 percent. Yet in 2007:Q2, the labor force participation rate fell back to an average of 66.0 percent. Where labor supply is headed has, accordingly, become the subject of some debate. The question has spawned a small literature seeking to separate trend developments from cyclical responses, and to understand what the future course of labor supply might be.[4]

A number of uncertainties currently cloud this debate. For example, when discussing a paper on labor supply presented by the current authors and others at a recent Brookings Institution panel, members of the panel suggested that participation rate projections might be highly sensitive to assumptions about the future course of immigration, and that government projections often do not include illegal immigrants in population estimates. As will be explained in more detail below, this assertion is not quite correct, and immigration has more influence on population growth and measures such as the dependency ratio than does the labor force participation rate.

To try to answer some of these questions, this paper seeks to document a number of the features of various labor force projections made by different federal agencies, and to evaluate the importance of various assumptions to those projections in order to focus subsequent research on where the most important uncertainties lie.

The paper proceeds as follows: Section II describes the shifts underway in the nation's age distribution, and how this change relates directly to the labor force participation rate. The influence of two key assumptions underlying U.S. population projections, life expectancy and immigration, is discussed in section III. Section IV discusses within-age-sex-group trends from the Bureau of Labor Statistics, the Social Security Administration, and the predictions of a cohort-based model of participation developed by the authors. This model identifies birth cohorts' propensities to participate in the labor force, and explicitly incorporates the changing population shares.[5] Section V cautions that we do not yet know how policy may ultimately respond to the coming changes, and section VI summarizes our main points.

II. How Population Shifts Influence Labor Force Participation

The reason that U.S. population aging has the potential to drastically slow labor supply growth is because labor force participation rates decline precipitously after age 50 for both men and women. Thus as the aging baby boomers move into their 60s, and as life expectancies continue to lengthen, the rising population share of older Americans has the potential to lower the share of Americans who are working or looking for work. For example, by 2035 the share of the adult population aged 80 years or more is expected to double to approximately 15 percent, and 97 percent of this age group currently does not participate in the labor force. The

future downward pressure on participation rates is primarily a result of these two forces—the aging of the baby boomers and longer life expectancies—now pushing in the same direction, after many years in which the upward pressure of the baby boomers moving into high participation rate ages offset the downward pressure from longer life expectancies on labor supply.

Figure 2.1 shows the age profiles of U.S. labor force participation rates for men and women using 2005 annual averages for 14 age categories, and the aggregate participation rate for reference.[6] Among women, those over 55 years of age have below-average labor force participation rates. Among men, the age groups over 60 years of age have below-average labor force participation rates. Historically the age profile for men has remained relatively stable, although among women, in particular those above 20 years of age, participation rates have risen during the postwar period. These various age profiles capture an important feature of labor supply—although within-age participation rates among older age groups

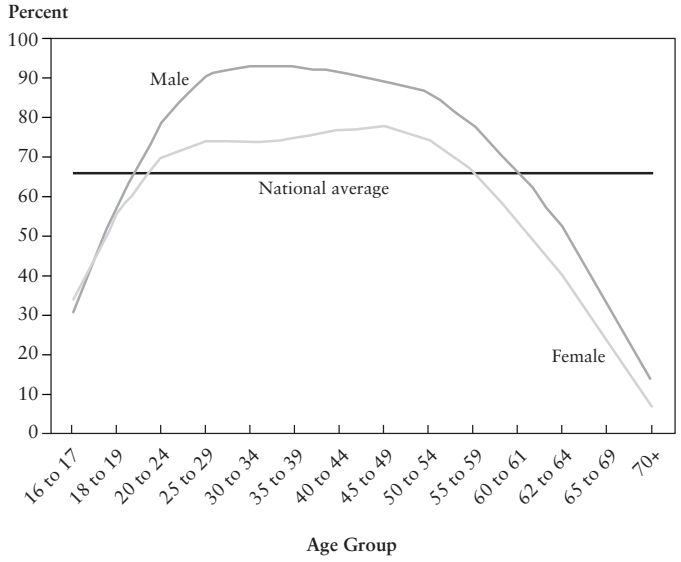

Figure 2.1
U.S. Labor Force Participation Rates by Age, 2005 (Annual Averages)
Source: U.S. Bureau of Labor Statistics.

may rise and fall, individuals in their 60s are much less likely to participate in the workforce than individuals in their 50s, and individuals in their 70s are much less likely to participate than individuals in their 60s.

In short, after age 50, labor supply is a declining function of age. Whether this lower supply arises due to failing health, disability, having adequate wealth, or retirement income, it is a feature of life-cycle labor supply unlikely to change fundamentally in the next few decades. While the first derivative of this function may change, it is quite likely to remain negative.

Figure 2.2 shows some historical and projected population shares of the age 16 and over Civilian Non-Institutional Population (CNIP), taken from the Current Population Survey, and grown out by the Census Bureau's population projections. The combination of rising life expectancy and the aging of the baby boom generation is dramatically pushing up the population share of older Americans. Around 2002, the share of individuals aged 55 years and older began to grow, while at the same time the share aged 35–44 years was shrinking. In 2006, the share of the age groups over 65 years old has begun to rise, and this increase will accelerate around 2010, with the share growing until around 2030.

Figure 2.2 highlights the fact that that the baby boomers are a bulge generation. As they leave high participation rate age groups, the cohorts behind the boomers are relatively small. In 1946 the U.S. fertility rate—defined as the number of live births per 1,000 women aged 15–44 years—leapt by nearly 20 percent, an unprecedented jump. The fertility rate rose another 10 percent in 1947. However, after peaking in 1957, fertility declined steadily and by 1966 had fallen back to levels seen during World War II. By 1973 the fertility rate had fallen to a level that at the time was the lowest recorded by the official U.S. statistics begun in 1909. The result of this fertility decline is that while the baby boomers grow older, the U.S. population in the highest labor force participation rate age groups is falling. In fact, the level of the CNIP of 30-34-year-old men, the highest participating age-sex group, has had outright declines since 2003, according to the Bureau of Labor Statistics.

The implications of these demographic shifts for future labor force participation are striking. Figure 2.3 shows the history of the actual labor force participation rate through 2006, along with a projection through

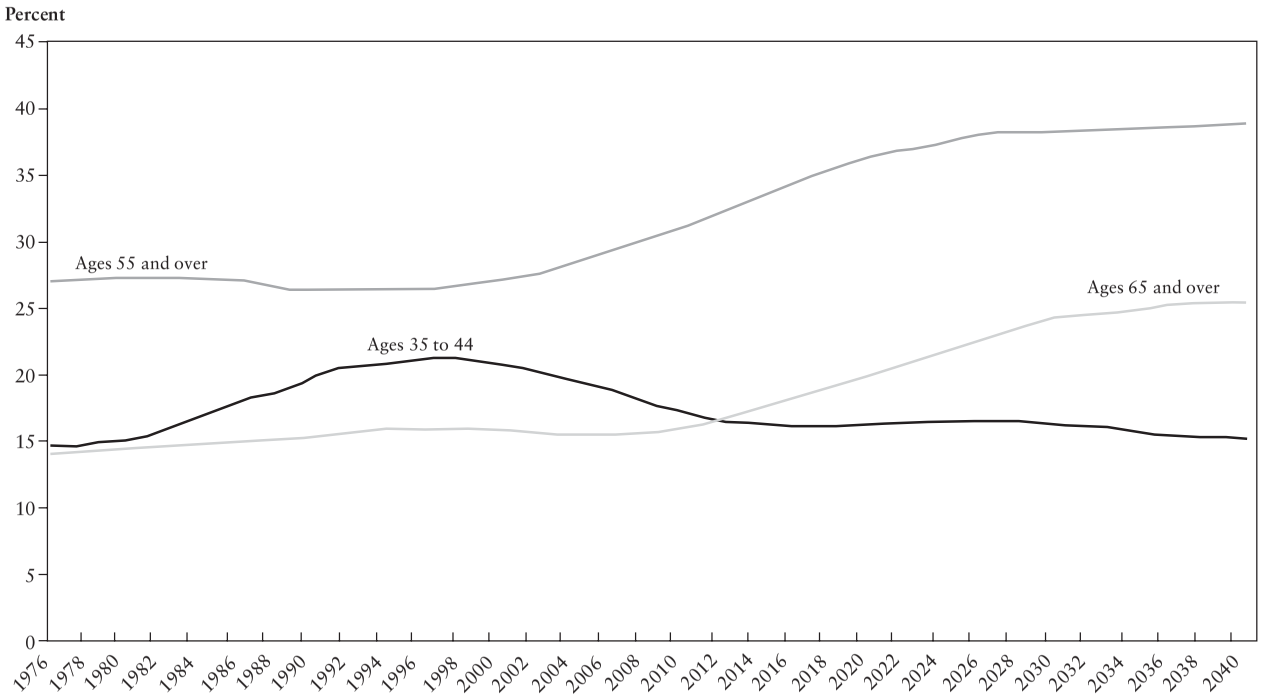

Figure 2.2
U.S. Population Shares by Age Group, 1976–2040: Actual and Projected
Source: Current Population Survey and U.S. Census Bureau.

Figure 2.3
U.S. Labor Force Participation Rates, 1950–2040: Actual and Projected
Source: U.S. Bureau of Labor Statistics and U.S. Census Bureau.
Note: Projection uses actual 2005 annual labor force participation averages within age groups, weighted by projected population shares.

2040 that uses the 2005 labor force participation rate for each age group and allows the population shares to evolve as forecast by the Census projections. Although other starting years can produce mildly different patterns, the implications are essentially the same: absent other changes, projected U.S. population aging will lower the aggregate labor force participation rate by 6 full percentage points over the next 35 years. This pace of decline dwarfs the 0.4 percentage point decline in aggregate labor force participation that shifting population shares have engendered over the past four years. In addition, using 2005 as a base year incorporates recent increases in older Americans' participation rates; the influence of aging would be a bit larger say, if we used year 2000 participation rates. In sum, the projected aging of the labor force is likely to have a sizeable influence on participation rates, with the potential to completely unwind the increases in participation attributable to the entry of more women in the workforce that began in the mid-1960s, accelerated sharply until 1990, and has since then leveled off.

The extent to which this decline shall materialize will depend on two inputs to the calculation: the projection for population shares, and the projections for within-age-group participation rates. We discuss each of these in turn.

III. Population Projections

Two government agencies are primarily responsible for U.S. population projections, and the projections' underlying assumptions overlap substantially. The primary agency is the Census Bureau, but the Social Security Administration produces independent estimates of population levels and growth. The primary difference between the two agencies' population projections is due to assumptions about net international migration, including undocumented immigration. However, in recent years important differences have arisen in the agencies' estimates for the population share of Americans over age 65.

The U.S. Census Bureau, of course, conducts the decennial census. In between these official censuses, which are essentially a benchmark, the agency produces estimates that update how many people are living in the United States, based on a variety of sources ranging from the National Center for Health Statistics to the Department of Defense; these updates

incorporate estimates from surveys such as the American Community Survey. The Census Bureau's updated estimates appear in the published reports from the Current Population Survey. Despite the Census Bureau's best efforts, these estimates can be off substantially. For example, the estimate of the national population from the 2000 decennial census was 6.8 million individuals more than the pre-decennial census estimates. Of that gap's discrepancy, 4 million was due to improved methodological reductions in undercounting, which would have raised pre-census estimates had the method been applied to the 1990 census. Only 2.8 million of the discrepancy was due to underestimated population growth, disproportionately Hispanic, which appears attributable to underestimating net international migration.[7]

While the Census Bureau's estimates are backward looking, and provide the nation's best estimate of the population at any given point in history, it also produces forward-looking projections. Based on the size of cohorts, assumptions for fertility rates, estimated death rates, and assumptions for net international migration, levels of the population at each age are projected out to 2050.[8] Similarly, the Social Security Administration produces its own projections, while using the Census Bureau's historical estimates.

However, the two agencies' assumptions have different implications for how the U.S. population is distributed across various age groups, and thus each agency's projection implies a different degree of downward pressure on the aggregate labor force participation rate. Although both the popular media and labor researchers tend to focus on how immigration may affect the future labor supply, the key difference for the U.S. labor supply outlook is the varying projections for longevity—in particular the life expectancy of older American women. Figure 2.4 shows the population shares of several age groups of men and women in the Current Population Survey estimates and in the Social Security Administration projections for 2005. The most notable difference is that the Social Security Administration predicted a lower share of women over age 65 than did the Census estimates, although the Census estimates were based on more recent source data.[9]

Figure 2.5 compares the implications of this difference for projections of the labor force participation rate. For this exercise we use the published Bureau of Labor Statistics' participation rates within age groups, so the

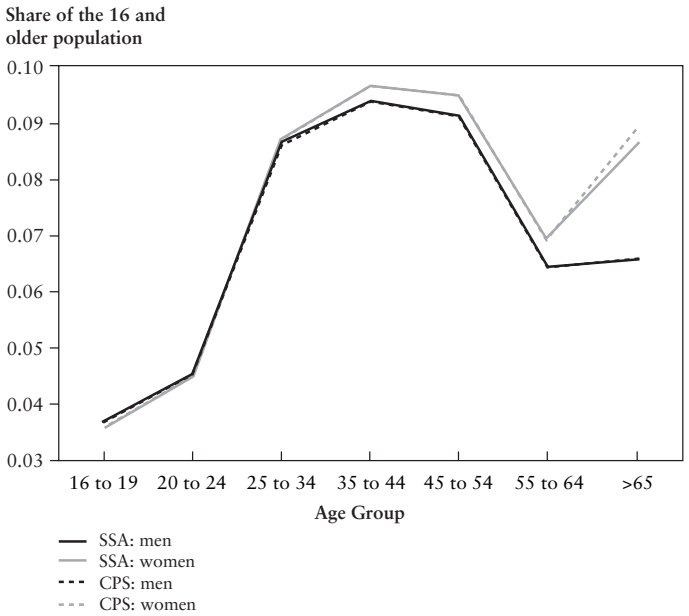

Figure 2.4
U.S. Population Distributions by Age, 2005: U.S. Census Bureau's Current Population Survey versus U.S. Social Security Administration Estimates
Source: Current Population Survey (U.S. Census Bureau) and U.S. Social Security Administration.

differences in the aggregate participation rates derive solely from differences in population shares. The solid line uses the current Census Bureau estimates of population shares in 2005 grown out by census projections for population shares going forward, while the dashed line uses the Social Security Administration's projections for the population shares. The pace of decline in the two projections is quite similar, which is not surprising, as both agencies rely mostly on current death tables for projecting future life expectancy and the immigration assumptions are relatively similar. An important part of the level gap reflects the Social Security Administration projections' smaller assumed labor market share of women over age 65, a group which has a particularly low participation rate, and the extrapolation of this difference.

This comparative exercise points to a particular downside risk to labor force participation rate projections. To the extent that most estimates

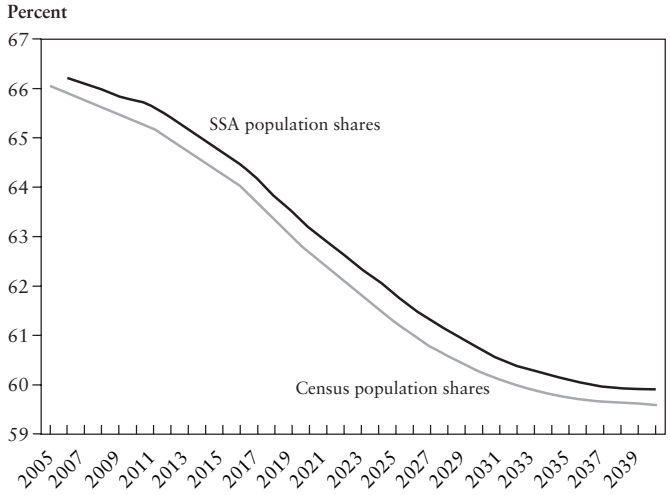

Figure 2.5
Projected U.S. Labor Force Participation Rates, 2005–2040: Census Bureau Population Shares versus Social Security Administration Estimates
Source: U.S. Bureau of Labor Statistics, U.S. Census Bureau, and U.S. Social Security Administration.
Note: Participation rates are calculated using published U.S. Bureau of Labor Statistics rates and alternative population projections.

of future life expectancy are based on current death rates, little adjustment is made for continued reductions in mortality rates.[10] To provide a sense of the magnitude of this risk, we adjusted the population growth of females aged 70 years and over from what is actually predicted by the Census Bureau projections. The change we made was to essentially simulate a shock to mortality rates for this age group. Deaths per person per year among all Americans age 65 years and over declined from nearly 0.1 in the years prior to World War II, down to nearly 0.05 by 2000.[11] For our counterfactual simulation we lower the projected path of deaths per person in the Census Bureau projections for females age 70 years and over by 0.01 in 2006 and every year thereafter. Note that by adjusting the rate of exit from the population, the population growth rate is boosted every year, which has a cumulative effect over time because the lower outflows raise the number of women aged 70 years and over who survive from year to year. The result of this counterfactual simulation is shown in Figure 2.6. In this scenario, by 2040 the downward pressure on

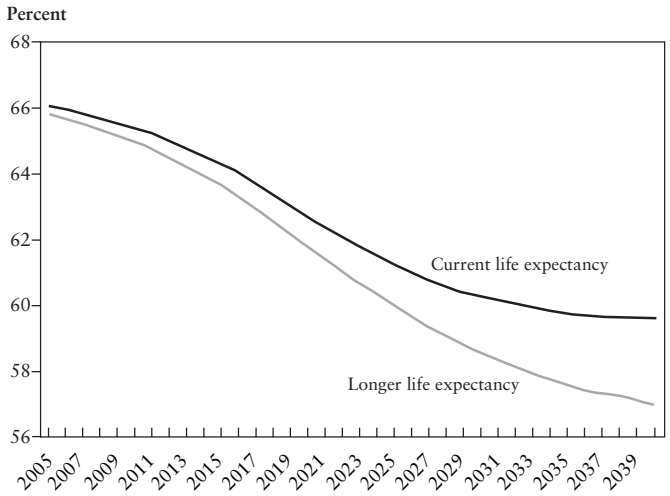

Figure 2.6
The Impact of Reduced Female Mortality Rates on Projected U.S. Labor Force Participation Rates, 2005–2040
Source: U.S. Bureau of Labor Statistics and U.S. Census Bureau.

the aggregate participation rate would be an additional 2.5 percentage points.

While continued reductions in mortality pose a downside risk to projections of the aggregate labor force participation rate, increases in immigration would likely pose an upside risk because immigrants are more likely to be in younger age groups with higher labor force participation rates, and are arguably more likely to participate even within older age groups (see Figure 2.7 for an example of the variation in participation across racial/ethnic groups).[12]

Much of the issue with immigration is one of measurement. For example, based largely on discrepancies between employment estimates from the Current Population Survey and the Current Employment Statistics program, many analysts argued that the Census Bureau has overestimated population growth, and in particular has overestimated net international migration since the 2000 census. If, indeed, the Census Bureau has overestimated net international migration, that may imply that the current level of the aggregate participation rate is overstated. And if the Census

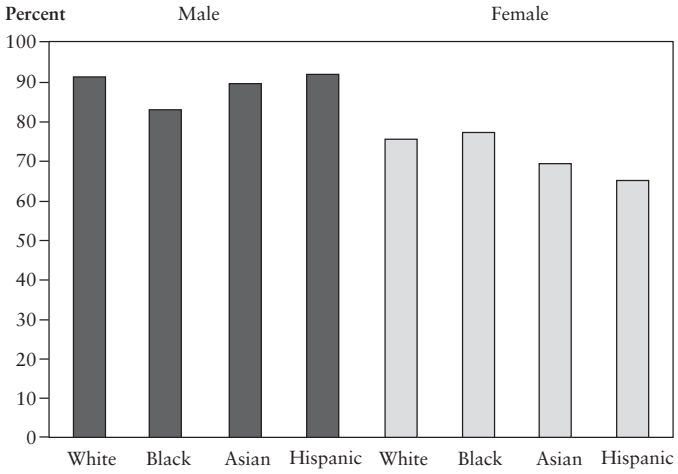

Figure 2.7
2005 U.S. Labor Force Participation Rate, Ages 25 to 54 Years, by Race and Gender
Source: U.S. Bureau of Labor Statistics.

Bureau projections for population growth also overestimate future net international migration, this may represent a downside risk to the projections for the U.S. aggregate labor force participation rate.

Undocumented immigration is the most difficult component of net international migration to measure. Accordingly, estimates of the number of undocumented immigrants vary widely, but the most recent estimates from the Pew Hispanic Center put the number of undocumented immigrants in the United States in 2006 at about 12 million—a figure consistent with the estimates in the Current Population Survey.[13] Of those 12 million undocumented immigrants, 49 percent are estimated to be male, 35 percent female, and 16 percent children. As to the total inflow of documented and undocumented immigrants, the Census Bureau estimated that total net immigration in 2005 was 1.2 million. The Social Security Administration, in its intermediate assumptions, assumes total net immigration in 2005 of 1.075 million. The Social Security Administration projects that total net immigration will decrease to 1 million annually by 2010 and to 900,000 annually by 2030. Alternatively, the Census Bureau projects total net immigration of 770,000 for 2010, and 1,161,000 for

2030, both of which are upward revisions from earlier assumptions after taking on board the data from the 2000 decennial census.

Of course, what will actually happen to U.S. net immigration 20 years from now is anyone's guess. However, compared to the influence of the aging of the population, immigration will have little influence on the aggregate added "immigrants" to the male 30-34-year-old age group—the age group with the highest labor force participation rate.[14] First, we added 100,000 of these high-participation persons each year from 2000 onward. This is represented as assumption (1) in Figure 2.8. Second, we added 200,000 per year, for an extra million (compared with an estimated stock of about 12 million) by 2006.[15] This is assumption (2) in figure 2.8. This alternative would be roughly similar to the high series nonlegal immigration assumption by the Social Security Administration, but represents a 50 percent increase over their intermediate assumption. Of course, the Social Security Administration would not likely assume that all those added immigrants would be 30-34-year-old men. In addition, we assume that even as they age, these male immigrants will maintain the

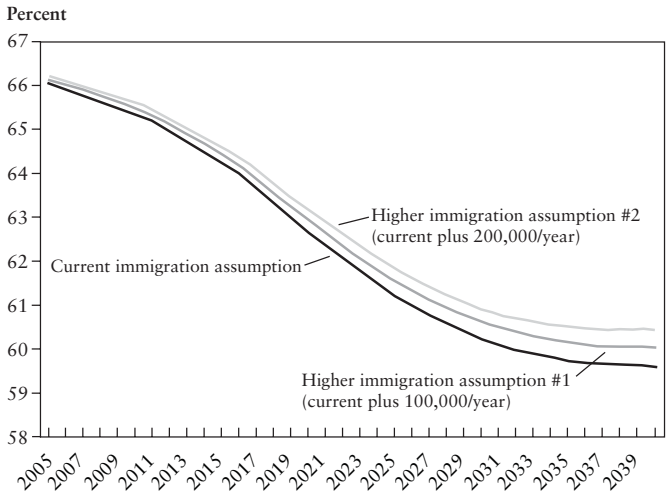

Figure 2.8
The Effect of Immigration Assumptions on Projected U.S. Labor Force Participation Rates, 2005–2040
Source: U.S. Bureau of Labor Statistics.

participation rates of current 30-34-year-old men, so our experiment represents something of an upper bound on the effect greater immigration may have on future labor force participation rates in the United States.

As Figure 2.8 shows, the accumulated extra immigration after the 2005 stepping-off point makes little difference to the aggregate U.S. labor force participation rate. And even under the extreme assumption that all these immigrants will be 30-34-year-old men, the continued higher level of immigration does relatively little to slow the pace of decline in the aggregate participation rate implied by the aging of the baby boom and longer life expectancy: immigration adds back only 1 percentage point to the 6.5-point decline in labor force participation rates due to population aging. The intuition for this muted response is that adding an immigrant to the numerator of the participation rate also adds someone to the denominator, while a retiring baby boomer removes an individual from the numerator but does not change the denominator. To completely offset the effects of aging on future labor force participation rates would require the United States to add well over 1 million 30-34-year-old males annually to population in the coming years, and for those immigrants to maintain that high 30–34-year-old participation rate even as they age.

The role of immigration leads to the question of whether one is primarily interested in the labor force participation rate, or in the overall labor supply. The participation rate is only one component of labor supply, the other two being population growth and hours worked. For example, the downside risk to the participation rate posed by increased life expectancy would largely be offset by the corresponding increase in population growth when computing labor supply. In addition, while immigration does not have a large influence on the labor force participation rate, rising population growth translates into more labor supply growth and means more for measures like the dependency ratio, which are arguably more important to policy. In our experiments, while labor force participation may continue to decline, the population aged 16 years and over is boosted by more immigration. Despite the projected declines in participation rates, immigration could increase the size of the population enough that the U.S. labor force grows. However, even the annual addition of the aforementioned 200,000 immigrants to the level of the population would currently be just a 0.1 percentage point increase in the

population growth rate, a small offset when the Social Security Administration is projecting labor supply growth to decrease to 0.5 percent a year in the next eight years.

IV. Projecting Trends in Within-Group Participation Rates

Comparing the implications for the participation rate of the Census Bureau projections and the Social Security Administration projections, as shown in Figure 2.5, indicates a certain amount of consensus among government forecasters for the direct future effects of population aging on aggregate labor force participation. Although there is uncertainty in these projections concerning primarily the outlook for U.S. immigration policy and U.S. life expectancies, the current consensus is that population aging will place substantial downward pressure on labor force participation rates, as both sets of population projections imply a decline of more than 6 percentage points over the next 30 years. How reality unfolds, of course, depends on how individuals and others respond to the changing circumstances. That is, while it seems likely that the baby boomers will age, and that life expectancies will continue to rise, what is less clear is how many people will, in the future, choose to work at any given age.

In this section of the paper we rely on three different projections of labor force participation rates. The Bureau of Labor Statistics and the Social Security Administration produce regularly updated projections of participation rates, not only for the aggregate labor supply, but within narrow age and gender categories. The Bureau of Labor Statistics' projections rely on time-series extrapolations of the recent trends within narrow demographic groups. The Social Security Administration combines model-based predictions, estimates of policy effects, and observations of the lagged behavior of birth cohorts, as well as judgmental adjustments to their projections.

In addition to the projections of these two agencies, we also consider the projections of a model we designed that relies on cohort effects in order to predict the future path of participation for individual birth years. Our model contains cyclical controls and variables representing education, fertility, and socioeconomic trends, as well as the features of government transfer programs, all of which are detailed in related papers we have written. The model is estimated for men and women separately, and produces a trend-cycle decomposition for labor force participation rates in 28 age-gender demographic groups that are aggregated by population weights. With some assumptions for the future path of the policy and sociodemographic variables, the model can be used to produce projections of the participation rate trend.[16] Each method is complicated enough to warrant long separate literatures, but comparing their outcomes highlights important risks to the projections for future labor supply.[17]

One important question for projecting future trends is whether people in older age groups will begin to work more than they have in the past. This change in behavior may follow from better health at older ages, more years of expected life to finance after age 65, changes in retirement preferences, or inducements to stay in the workforce provided by business or government. The popular press has questioned whether the baby boomers may keep on working past age 65, despite the fact that at younger ages, male baby boomers have tended to work *less* than the male cohorts who preceded them; see Figure 2.9.

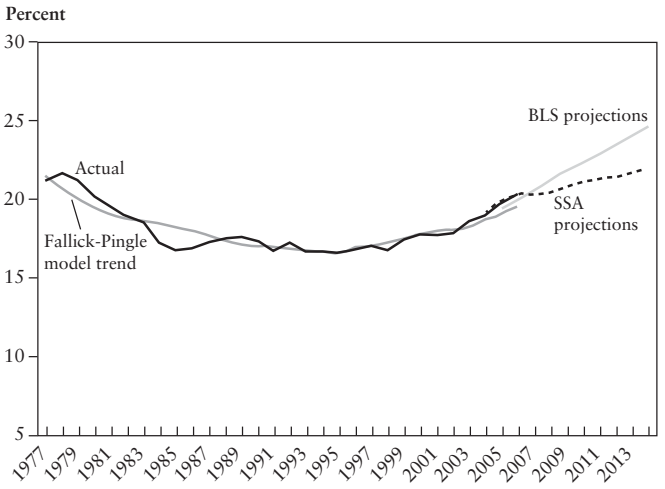

Figure 2.9
U.S. Labor Force Participation Rates for Men Aged 65 Years and Older, 1977–2017
Source: U.S. Bureau of Labor Statistics, U.S. Social Security Administration, and authors' calculations.

Among the forecasts produced by U.S. government agencies, most are in agreement that the labor force participation rates of older age groups will rise in the coming years. The participation rates for men aged 65 years and over and for women aged 65 and over are shown as the solid lines in Figures 2.9 and 2.10, respectively. The age group comprising those aged 65 years and older has increased its labor force participation markedly over the past decade. This increase, which is taking place in one of the fastest growing components of the U.S. population, has the potential to offset some of the declines in participation due to the shifts in population share.

Both Figures 2.9 and 2.10 show the historical data, the Bureau of Labor Statistics' "full-employment" projections, the Social Security Administration projection, and the projected trend from our own (Fallick-Pingle) cohort-based model. For men aged 65 years and over, all three forecasts are rising, with the Fallick-Pingle estimates roughly in line with the Social Security Administration projections. The Bureau of Labor Statistics' projections—an extrapolation of the preceding years' growth rates—point to much faster growth. In the Fallick-Pingle model the increases are attributable to a combination of longer life expectancies, higher education among the cohorts reaching these ages, and favorable Social Security incentives, specifically, the rising retirement age and the increases in the delayed retirement credit, which encourage older adults to remain in the workforce.

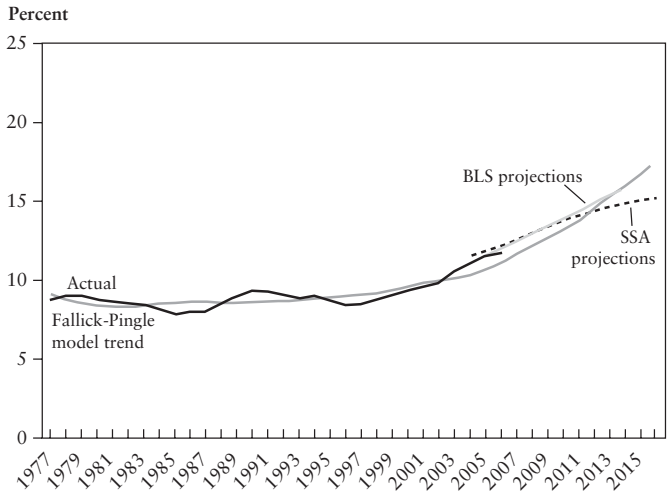

Figure 2.10
U.S. Labor Force Participation Rates for Women Aged 65 Years and Older, 1977–2017
Source: U.S. Bureau of Labor Statistics, U.S. Social Security Administration, and authors' calculations.

For women aged 65 years and over, the qualitative aspects of the forecasts are quite similar. The Fallick-Pingle trend estimates have a slightly steeper trajectory than other forecasts. In this model, these increases are largely due to female cohorts that are highly attached to the labor force—the same women who from 1970 to 1990 raised the prime-age female labor supply—filling up this older age category. All three of the forecasts expect participation rates for this demographic group will continue to rise rapidly. If we look at the rough estimates, all of the projections expect gains of nearly 0.5 percentage points per year in this group's participation rate. The share of women aged 65 years and over is expected to increase from 9 percent to 10 percent of the population over this projection horizon, to 2015. Roughly speaking, this implies this demographic group alone would add nearly 0.05 percentage points to the aggregate labor force participation rate per year. As shown in Figure 2.5, the projected pace of decline due to population aging is roughly 0.2 percentage points per year, on average, between 2005 and 2040. Thus, the increased labor force participation among older women projected by all three forecasts would offset roughly one-quarter of that decline over the period shown in the model projections.[18]

Where the various forecasts begin to diverge is among prime-age workers, those who are 25-54-years-old. This divergence is particularly important because prime-age men and women represent more than half of the population aged 16 years and over —a large weight in any aggregation. Figure 2.11 shows the historical data for prime-age men, and the forecasts for their continued labor force participation. This group's participation rate has declined steadily for nearly the entire postwar period, although the data shown begins only in 1977. Not foreseeing any reason for those declines to end, the Fallick-Pingle model extrapolates this trend forward. This estimate is roughly similar to the Social Security

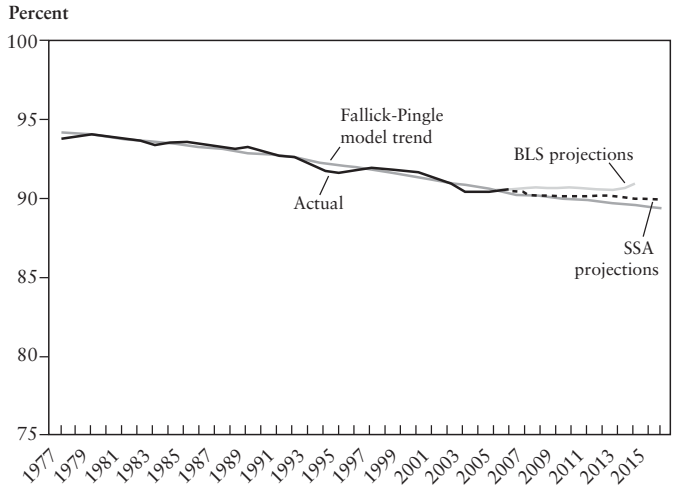

Figure 2.11
U.S. Labor Force Participation Rates for Men Aged 25 to 54 Years, 1977–2017
Source: U.S. Bureau of Labor Statistics, U.S. Social Security Administration, and authors' calculations.

Administration forecast, while in the coming years the Bureau of Labor Statistics projections expect slight increases in prime-age men's labor force participation rates. This forecast amounts to a leveling off or reversal of the data series' recent history.

As shown in Figure 2.12, more pronounced differences emerge among the forecasts for prime-age females, as the Fallick-Pingle model projects declines that the two other forecasts do not. The Fallick-Pingle model anticipates that the participation rate of prime-age women has not only leveled off, but has begun to trend somewhat like prime-age men's, with steady declines predicted in the years ahead. In contrast, the Social Security Administration predicts prime-age women's labor force participation will level off, and the Bureau of Labor Statistics projects increases that will eventually take prime-age women's participation rate to levels not seen since the business cycle peak in 2000. Given the large share of the American population represented by the age group, this divergence for

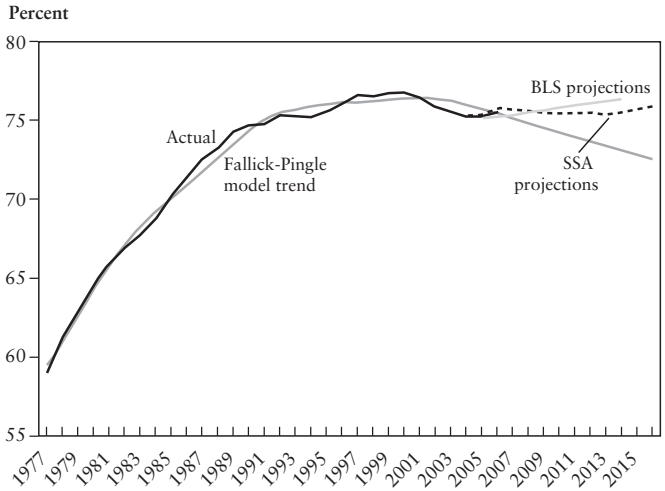

Figure 2.12
U.S. Labor Force Participation Rates for Women Aged 25 to 54 Years, 1977–2017
Source: U.S. Bureau of Labor Statistics, U.S. Social Security Administration, and authors calculations.

predicting the future labor market behavior of prime-age women accounts for a large portion of the deviations in the aggregate forecasts.[19]

The intuition behind the Fallick-Pingle estimates stems from the observed behavior of female birth cohorts over time. Figure 2.13 shows the participation rates for three age groups of women, ages 35–44 years, 45–54 years, and 55–64 years. However, instead of using the year of observation on the horizontal axis, the figure uses the birth-year cohort. As the dotted line shows, the participation of women aged 35–44 years, leveled off starting with the 1949 or 1950 birth cohort. Similarly, the participation rate of the next age group, women aged 45–54 years, also leveled off when the women born around 1949 or 1950 were in that age group (there are similar changes in slope in earlier periods which, although not shown, line up also). Thus, the pattern of a birth cohort's strong or weak labor force attachment earlier in life appears to influence that cohort's behavior later in life. For the female cohorts born in

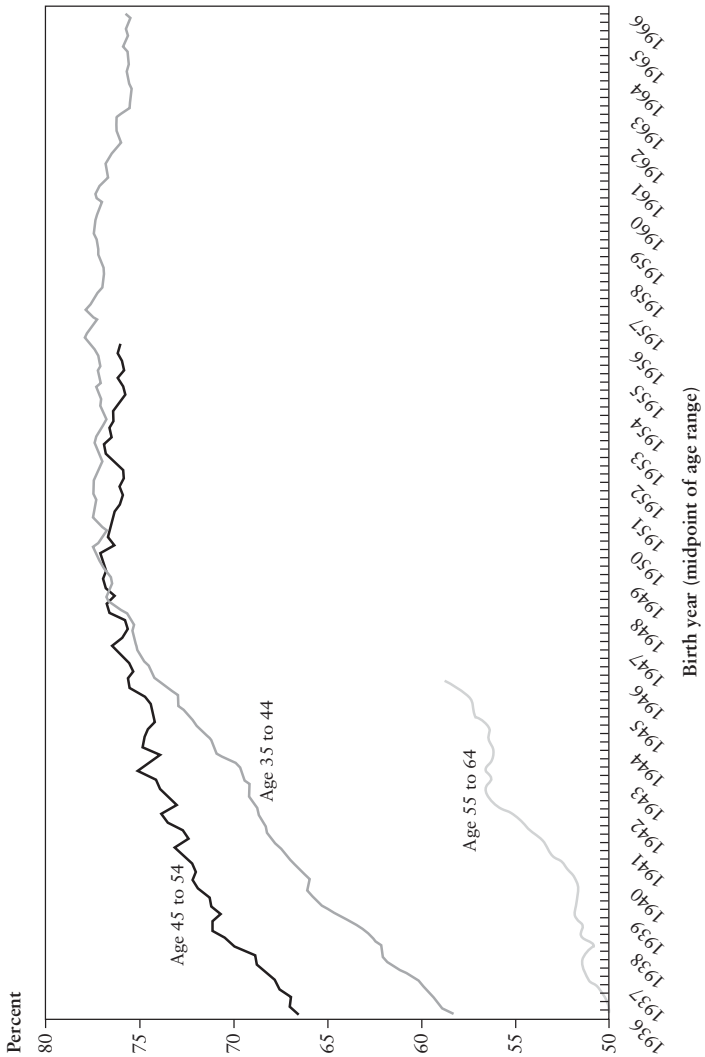

Figure 2.13
U.S. Labor Force Participation for Women Aged 35 to 64 Years, by Birth-Year Cohort and Age Group
Source: U.S. Bureau of Labor Statistics.

1957 through 1966, labor force attachment has declined from cohort to cohort, and the model carries this pattern forward. In essence, prime-age women have begun to behave similarly to prime-age men, whose labor supply has been trending down, on balance, for much of the postwar period. As of this writing, the labor force participation rate of women aged 25–54 years remains nearly 2 percentage points below the levels reached in 1997.

Although teenagers comprise only 7 percent of the U.S. population aged 16 years and over, the recent declines in their labor force participation rates have been particularly vexing to observers.[20] Figure 2.14 shows the teenage labor force participation rate, combined for males and females. Although trending down for some time, teenage labor force participation dropped precipitously following the 2001 business cycle peak, and has failed to increase even after a few years of labor market recovery. This decline is poorly explained by cyclical controls, and is not fully explained by changes in school enrollment, as a substantial amount of the decline occurred among individuals enrolled in school.[21] A full

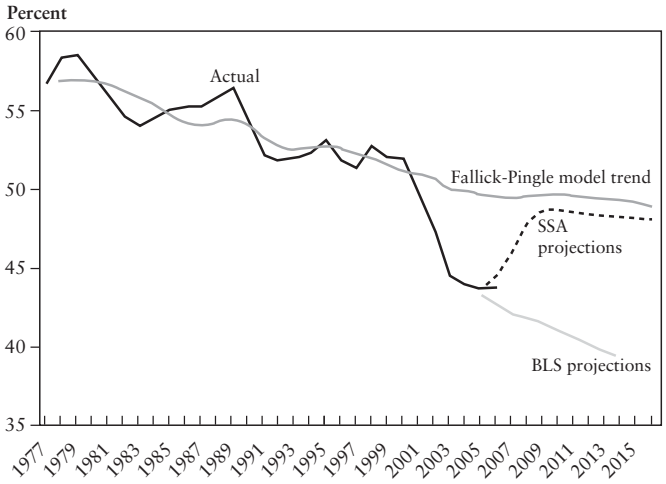

Figure 2.14
U.S. Labor Force Participation Rates for Teenagers, 1977–2017
Source: U.S. Bureau of Labor Statistics, U.S. Social Security Administration, and authors' calculations.

explanation of the labor market behavior of American teenagers remains a topic of ongoing research.

After aggregating the various components, the three forecasts for the aggregate U.S. labor force participation rate diverge noticeably—largely due to the differing assumptions made for the participation rates of prime-age workers; see Figure 2.15.

The Fallick-Pingle trend is the lowest; this forecast relies on Census Bureau weights and carries forward declines in participation rates among prime-age men and women. The Social Security Administration forecast starts with a higher level of participation among prime-age workers, reflecting the influence of an assumed smaller share of women over age 65. It also reflects a cyclical rebound that is expected to result in a 66.5 percent labor force participation rate for this group in 2007, then remains level, and then trends down beginning in 2009, although not quite at the pace predicted by the Fallick-Pingle model. This slower pace of decline in the Social Security Administration projections is in large part a function of the steep increases in teenage participation rates it has

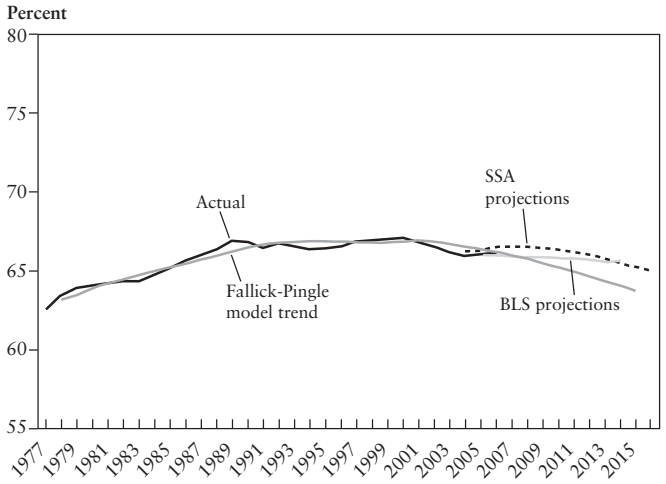

Figure 2.15
U.S. Aggregate Labor Force Participation Rates, 1977–2017
Source: U.S. Bureau of Labor Statistics, U.S. Social Security Administration, and authors' calculations.

Table 2.1
Aggregate U.S. Labor Force Participation Rates (annual averages)

Year	BLS	SSA	F/P
2005	66.0	66.28	66.37
2007	65.9	66.50	66.02
2008	65.9	66.43	65.51
2011	65.8	66.21	65.00
2013	65.6	65.76	64.40
2015	n.a.	65.25	63.79

projected. Even the Bureau of Labor Statistics estimates project aggregate labor force participation rates to edge lower in the coming years—see Table 2.1—despite robust increases forecast within nearly every age group (with the notable exception of teenagers) that will offset the downward pressures from population aging. However, as the Fallick-Pingle trend estimates show, even forecasts of strong increases in labor force participation among older age groups are not enough to offset the deeper influence of population aging in the United States. Moreover, if the labor force participation rate for prime-age workers trends down, this age group's behavior will exacerbate the declines in U.S. labor force participation rates due to population aging.

V. Policy Responses

The three projected forecasts for future U.S. labor force participation rates discussed thus far in this paper do not consider possible policy responses beyond those already enacted. As these currently stand, the Social Security, Medicare, and private pension programs embed a long list of incentives that influence the American labor supply. Large scale changes in old age programs in the United States have been rare, and policy can react with quite a lag to changing conditions. For example, the increase in U.S. women's labor supply began in earnest in the early 1960s, but federal legislation and even corporate responses to needs such as family leave, did not take hold until the early 1990s. However, within the United States,

income maintenance programs, family leave, and Social Security benefit programs have all been changed by major federal legislation in the past, and will likely be changed in the future. Such changes could have major implications for future labor supply.

Given the rising population share of older Americans, old-age programs are likely to be a particular focus of U.S. public policy in the coming years. There have been few substantive changes in these programs since the early 1980s. In 1983, the Greenspan Commission on Social Security reform suggested altering the program's incentives on when to retire and ultimately the recommendations became law. Called the National Commission on Social Security Reform, it was formed as Congress and the Reagan administration became worried that the Social Security system, facing a financial crisis, might be insolvent by the middle of that decade. That fear, combined with the projected long-run deficits expected with the retirement of the baby boomers, prompted Congressional action to encourage workers to delay receiving benefits (see Pingle 2006).

Traditionally, Americans first qualified for Social Security benefits at age 62, and then qualified for full benefits at age 65, termed the Normal Retirement Age (NRA), the age at which one qualifies for their Primary Insurance Amount (PIA). The 1983 amendments raised the NRA gradually from 65 to 67, but this shift has begun only in the few most recent years. The legislation delayed the June 1983 cost-of-living increases. The amendments also changed the Retirement Earnings Test penalty for people turning age 65 in 1990 or later—reducing the penalty levied on Social Security payments/benefits from $1 of every $2 earned by people over an earned income limit, to $1 of every $3. Subsequently, in 1996 new legislation raised the previous earned income limit, and in 2000 the Retirement Earnings Test was changed so that it only applies to people who are between 62 years of age and their NRA.

The 1983 amendments substantially changed the Delayed Retirement Credit (DRC) beginning with the cohort born in 1925 turning 65 in 1990. The DRC raises an individual's lifetime Social Security benefits for each month that he or she delays receiving their benefits after reaching their NRA—traditionally when someone turns 65. The DRC was instituted in 1972 to provide a 1 percent bonus to a person's remaining Social Security pension to compensate for each year past age 65 a person delayed

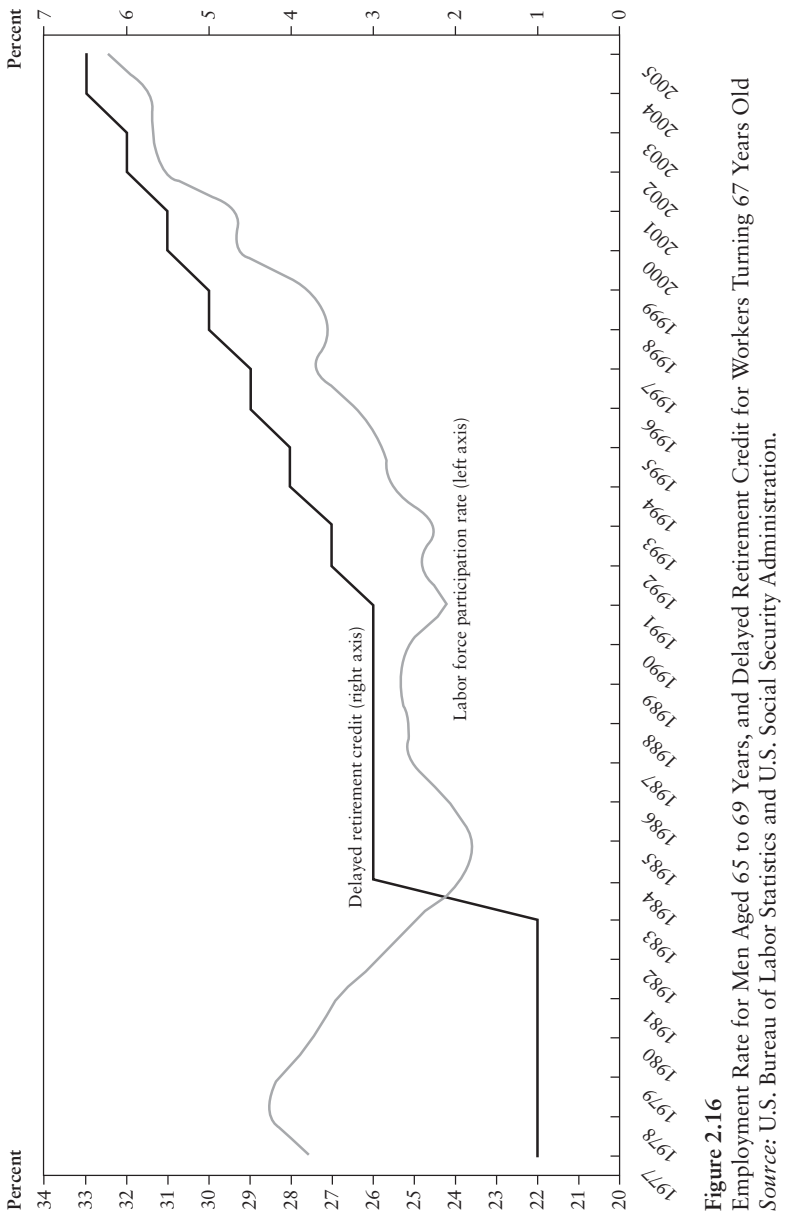

Figure 2.16
Employment Rate for Men Aged 65 to 69 Years, and Delayed Retirement Credit for Workers Turning 67 Years Old
Source: U.S. Bureau of Labor Statistics and U.S. Social Security Administration.

receiving benefits, until age 70, in order to at least partially make up over time for the present value of the deferred benefits. The credit is applied in monthly increments so that past an eligible recipient's 65th birthday, for each month during which benefits are not received, there is an upward adjustment in the lifetime amount of the monthly Social Security payment the individual eventually receives. The 1983 amendments increased the DRC from 1 percent to 3 percent. The new, higher levels of the DRC were assigned by the year an eligible worker was born. Under the new policy, the DRC rose from a 3 percent additional benefit for those born in 1924 or earlier, in half-percent increments every two birth years, until the credit reaches 8 percent for individuals born in 1943 or later. Thus, the DRC increased for workers born in 1925 or later, in effect starting in 1990 when that birth cohort turned 65.

As shown in Figure 2.16, the rising levels of the DRC correspond closely to the increases in male labor supply among men in the affected age group. This increase bucks the trend towards lower labor force participation rates among younger men. Empirical estimates of the influence of the DRC are not the specific subject of this paper, although the Fallick-Pingle projections do include the policy as a right-hand-side variable. The point is that in the years ahead, this type of policy change might influence labor supply, and as-yet-unenacted policy changes have not been accounted for in any of the forward-looking estimates shown in this paper. Outside of Social Security, as the U.S. labor force ages firms may redesign workplace policies to attract and accommodate more senior employees. However, predicting such future responses is an essentially impossible task at the present time. It took 30 years after female labor force participation had begun climbing steeply in the early 1960s before comprehensive family leave legislation was enacted in the United States. How the future course of policy may eventually evolve is anybody's guess.

VI. Conclusion

The rate at which individuals participate in the labor force declines precipitously beyond age 50. This behavioral feature of labor supply suggests that ongoing shifts in the age distribution of the U.S. population have already begun to put substantial downward pressure on the nation's aggregate labor force participation rate. Although projections of such pressure in years to come rely on current projections for the evolution of the age distribution of the population—estimates which are sensitive to assumptions about future mortality rates and immigration—different population projections have relatively little influence on the degree of downward pressure aging will exert on the aggregate participation rate. Instead, the extent to which population aging will depress aggregate participation rates in the U.S. labor force rests critically on how the participation rates within the various age groups evolve. In the end, however, substantial increases in participation rates across almost all age groups of both genders would be needed to completely offset the projected declines in labor force participation predicted over the next 30 years.

Unforeseen endogenous responses aside, the U.S. population is aging. The baby boomers are undoubtedly moving into lower participation rate age groups, relatively small younger cohorts are or will be moving into high participation rate age groups, and life expectancies are lengthening. While current labor market projections expect rising participation among older workers, a reasonable consensus is building that suggests U.S. labor supply has already slowed due to population aging and will slow further in the years ahead. The agreement among the three projections described in this paper is that the outlook is for slower growth in U.S. labor supply from 2007 onward than was the norm in the 1965–2000 period. Certainly, considering the information now in hand, the Social Security Administration's projection for labor force growth to slow to 0.5 percent a year by 2015, only seven years away, is a reasonable expectation, recognizing that there are reasonable scenarios that could lead to higher or lower growth rates. This is 25 percent of the pace observed from the three decades spanning the 1960s through the 1980s. How economists and U.S. policymakers build this general prediction into their thinking about the future performance of the macroeconomy is one of the most important research and public policy discussions on the contemporary research and policy agenda.

■ *The views expressed in this paper are those of the authors and do not represent the view of the Federal Reserve System or its staff. The authors wish to thank Stephanie Aaronson, Andrew Figura, and William*

Wascher, who co-authored a related paper on this topic. Special thanks also go to Karen Smith at the Social Security Administration for providing detailed data that underlies the agency's labor force projections, and Mitra Toosi who provided data and answered questions on the Bureau of Labor Statistics's long-run projections for labor supply. In addition, we received useful input for this line of research from Daniel Aaronson, Gary Burtless, Chinhui Juhn, Charles Fleischman, Julie Hotchkiss, Joseph Lupton, Lisa Lynch, Mark Schweitzer, Daniel Sullivan, Joyce Zickler and seminar participants at the Bureau of Labor Statistics. Leslie Carroll and Andrew Strauss provided expert assistance.

Notes

1. Social Security Administration 2007 Trustees Report.

2. See Fallick and Pingle (2007) for more discussion of the recent history.

3. This computation can be seen by using the Current Population Survey data, holding the participation rates within age groups fixed at their levels at some point in history and then allowing the actual population shares to evolve.

4. Bradbury (2005); Aaronson, Park, and Sullivan (2006); Aaronson, Rissman, and Sullivan (2004); and Toossi (2005).

5. An earlier version of this model was used in Aaronson, Fallick, Figura, Pingle, and Wascher (2006).

6. As shown in Fallick, Fleischman, and Pingle (2007), the declines do not depend on the base year chosen.

7. See West and Robinson (2005). "Understanding Factors that Contributed to the Large Error of Closure in Census 2000, a note available online at: http://paa2005.princeton.edu/download.aspx?submissionId=51262

8. The Census Bureau's population estimates are updated annually, while the Census Bureau's projections are updated about twice a decade. Therefore, the Census Bureau's projections for future population levels may not be consistent with its best current estimate of the historical population. This difference has implications for measuring labor force participation in real time, and highlights where some projection risks may lie. For example, the incoming population estimates reflected in the 2005 Current Population Survey population shares have more downward pressure on aggregate labor force participation than the share implied by the projections that the most recent Census Bureau estimates have superseded. Again, in January 2006, revisions to population estimates prompted revisions to the weights in the Current Population Survey, from which U.S. labor participation is officially measured. The resulting new population estimates, all else remaining the same, caused labor force participation rates to be revised downward by two hundredths of a percentage point.

9. The data used in this paper that refer to Social Security Administration estimates and projections refer to the data underlying the Social Security Administration 2006 Trustees Report, unless otherwise mentioned.

10. According to Ward Kinkaide in the Census Bureau's population projections branch, one reason for little extrapolation of life expectancy by birth cohort is because the data for many of cohorts born early in the century are of poor quality, thus limiting the ability to assemble a long time series of full-age-span life tables that can be used to project their evolution forward.

11. See Wilmoth (2005) for discussion of mortality rates, their history, and their role in population projections.

12. The risks to the immigration assumptions are essentially risks to the net immigration of Hispanics. As Figure 2.7 shows, Hispanic males in the United States have high labor force participation rates, but the same does not apply to Hispanic females.

13. For more discussion see Passel (2006). Jeffrey S., "Size and Characteristics of the Unauthorized Migrant Population in the U.S.," Pew Hispanic Center, 2006.

14. Here we assume that the immigrants would have the same participation rate as the general population.

15. The Pew Hispanic Center estimates that net undocumented immigration in 2006 was 500,000. The addition of 200,000 more immigrants per year would imply a large margin of error to existing estimates. However, we recognize that estimating the flow of undocumented immigrants likely has wide confidence intervals.

16. See Fallick and Pingle (2007).

17. Discussion of the Bureau of Labor Statistics projections can be found in Chapter 13 of the 2007 edition of the Bureau of Labor Statistics Handbook of Methods. This is available at http://www.bls.gov/opub/hom/homtoc.htm. The Social Security Administration projections are discussed in the Trustees Report, available on-line at http://www.ssa.gov/OACT/TR/TR07/index.html. For more information on the estimation of Social Security Administration participation rates, see the Technical Panel on Assumptions and Methods (2003).

18. This offset refers to only one aspect of the influence of this age group, as the growth in the population share of these age groups with participation rates well below average are an important part of the downward pressure on labor force participation resulting from population aging.

19. Although the data from the 2007 Trustees Report was released too late to be incorporated into this paper, one correction made to the Social Security Administration projections was to revise downward the forecasts for various groups of prime-age women. For example, the future projection for females aged 45–54 years in 2016 was revised from 74.2 percent to 73.9 percent.

20. For discussion of this decline, see Aaronson, Park, and Sullivan (2006).

21. This is easily verified using the Bureau of Labor Statistics data series from the basic monthly Current Population Survey.

References

Aaronson, Daniel, Kyung Park, and Daniel Sullivan. 2006. "The Decline in Teen Labor Force Participation." *Economic Perspectives* 30(1): 2–18.

Aaronson, Daniel, Ellen Rissman, and Daniel Sullivan. 2004. "Assessing the Jobless Recovery." *Economic Perspectives* 28(2): 2–20.

Aaronson, Stephanie, Bruce Fallick, Andrew Figura, Jonathan Pingle, and William Wascher. 2006. "The Recent Decline in the Labor Force Participation Rate and Its Implications for Potential Labor Supply." *Brookings Papers on Economic Activity* (1): 69–134.

Bureau of Labor Statistics, "Handbook of Methods," May 2007 update, available online at http://www.bls.gov/opub/hom/home.htm.

Bradbury, Katherine. 2005. "Additional Slack in the Economy: The Poor Recovery in Labor Force Participation During This Business Cycle." Federal Reserve Bank of Boston Public Policy Briefs, 05-2.

Fallick, Bruce, Charles Fleischman, and Jonathan Pingle. 2007. "The Effect of Population Aging on the Aggregate Labor Market." Paper presented at the Labor in the New Economy Conference, sponsored by the National Bureau of Economic Research Conference on Research in Income and Wealth, November 16–17.

Fallick, Bruce, and Jonathan Pingle. 2007. "A Cohort-Based Model of Labor Force Participation." Finance and Economics Discussion Series 2007-9, Federal Reserve Board, April 2007.

Passel, Jeffrey S. 2006. "Size and Characteristics of the Unauthorized Migrant Population in the U.S." Pew Hispanic Center. Available online at http://pewhispanic.org/files/reports/61.pdf.

Pingle, Jonathan F. 2006. "Social Security's Delayed Retirement Credit and the Labor Supply of Older Men." Finance and Economics Discussion Series 2006-37, Federal Reserve Board, August 2006.

Social Security Administration. 2006. "The 2006 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds." Available online at http://www.ssa.gov/OACT/TR/TR06/I_intro.html#wp1000302.

Social Security Administration. 2007. "The 2007 Annual Report of the Board of Trustees of the Federal Old-Age and Survivors Insurance and Federal Disability Insurance Trust Funds." Available online at http://www.ssa.gov/OACT/TR/TR07/I_intro.html#wp1000302.

Technical Panel on Assumptions and Methods. 2003. Report to the Social Security Advisory Board, Washington D.C. Available online at http://www.ssab.gov/2003TechnicalPanelRept.pdf.

Toossi, Mitra. 2005. "Labor Force Projections to 2014: Retiring Boomers." *Monthly Labor Review* 128(11): 25–44.

West, Kirsten K., and J. Gregory Robinson. 2005. "Understanding Factors that Contributed to the Large Error of Closure in Census 2000." Available online at http://paa2005.princeton.edu/download.aspx?submissionId=51262.

Wilmoth, John R. 2005. "Overview and Discussion of the Social Security Mortality Projections." Report to the Social Security Advisory Board, Washington, D.C., May.

Comments on "The Effect of Population Aging on Aggregate Labor Supply in the United States" by Bruce Fallick and Jonathan Pingle

Chinhui Juhn

After nearly four decades of growth, the U.S. labor force now faces the prospect of two negative blows—the retirement of the baby boom cohort and the cessation and possibly the reversal of female labor supply growth. Fallick and Pingle show that if age-specific participation rates remain at today's rates, the aging of the U.S. population could potentially offset the entire growth in the aggregate labor force participation rate we have witnessed since the mid-1960s. While there are some debatable issues regarding their exact calculation of population shares—such as projection of mortality rates and future immigrant inflows—the key question stemming from their analysis that I will focus on is: will future age-specific labor force participation rates will remain at today's levels? In particular, there are two critical groups—the older adult populations (those aged 65 years and over, and those at 55–64 years of age nearing retirement) and prime-age women between the ages of 25–54 years. Will the baby boom cohorts work longer relative to previous cohorts? Has the labor force participation rate of married U.S. women not only stopped climbing, but are there signs in the last decade that the trend has actually reversed?

On the question of older Americans, models by Fallick and Pingle and two government agencies, the Bureau of Labor Statistics (BLS) and the Social Security Administration (SSA) are in agreement and their collective outlook is fairly positive—that is, if having older Americans work longer is the goal. The authors list several factors that, going forward, seem to predict rising labor force participation rates for older groups. Among these factors are longer life expectancies and hence longer retire-

ment periods that need to be financed, better health, increases in education, and changes in the Social Security program that reduce the incentive to retire early. I would like to mention two other contingent factors. The first is the possibility that the increased labor force attachment of women is not only likely to increase participation among the women themselves at older ages, but also their husbands' participation rates, given the often-found complementarity in spouses' retirement behavior. Table 2.2 shows how for different age groups, a husband's labor force participation is related to his wife's labor force participation. The marginal effects reported in the first row of the top panel show that a wife's participation in the labor force increases her husband's participation, and this complementary/mutually reinforcing effect rises steeply with age. For men aged 55–64 years, having a wife who is in the labor force increases the husband's participation probability by 12.5 percentage points. At age 65 years and over, the marginal effect is 22.3 percentage points. While the table to some degree reflects positive assortative mating, the increasing

Table 2.2
Husband's Labor Force Participation as Function of Wife's Participation

	Husband's Age				
	25–34	35–44	45–54	55–64	65+
All Years: 1968–2006					
Wife in the LF	0.007	0.014	0.029	0.125	0.223
Number of Children <=18	0.002	0.003	0.002	0.011	0.034
Observed Probability	0.968	0.964	0.930	0.733	0.198
Recent Years: 1996–2006					
Wife in the LF	0.013	0.020	0.048	0.167	0.262
Number of Children <=18	0.004	0.008	0.006	0.023	0.063
Observed Probability	0.961	0.956	0.922	0.720	0.194

Source: March Current Population Survey 1968–2006. The sample includes married men who are 25 and older. The above reports the marginal effects from probit regressions with husband's probability of being in the labor force as the dependent variable. The regression also included husband's age, husband's and wife's education dummies, and year and state fixed effects.

positive association over the life cycle most likely reflects the complementarity of time spent at home for retired spouses. The bottom panel shows that this effect has in fact become more pronounced in the last ten years, 1996–2006, with the marginal effect for 55-64-year-old men rising to 16.7 percentage points. Of course these numbers all refer to married men. For this 55-64-year-old age category, the availability of disability benefits, together with lackluster demand for less-skilled workers, will likely continue to put downward pressure on the participation rates of less-skilled unmarried men in this age group.

Another factor not mentioned by Fallick and Pingle, but likely to become important in the next several decades, is the decline in employer-provided health insurance for retired workers. Between 1988 and 2005, the share of large firms (those having 200 or more employees) offering health insurance to their active employees and also offering health insurance to their retired employees declined precipitously from 66 percent to 33 percent. Current retirees and near-retirees may be grandfathered into employer-provided retirement health plans, so that we may not see much change in coverage rates in recent years. However, successive cohorts of retirees are surely less likely to receive such generous health benefits from their former employers. This decline in retiree health coverage could be another factor that boosts the labor force participation rate of 55-64-year-olds.

Let me now turn to prime-age women, those in the 25-54-year-old age group. The authors' forecast of these women's labor force participation differs considerably from the forecasts made by other government agencies. As illustrated in their Figure 2.12, by the year 2015 there is almost a 3 percentage point difference between the Fallick and Pingle forecast and the other forecasts. Yet the behavior of women has in the past proven to be notoriously difficult to forecast. I doubt that anyone in the 1950s or 1960s would have accurately forecasted the acceleration in married women's labor force participation during the 1970s and the 1980s. The point is, this exercise in forecasting the future is largely a guessing game. However, the authors' projections of the decline in the participation rate of prime-aged women is, in my view, overly pessimistic. First, while the rise in participation of married women has definitely come to a stop, I

do not see evidence of an actually declining trend. Second, I have some reservations about the authors' cohort-based approach.

Table 2.3 examines the labor force participation rate of prime-age women using the March Current Population Survey. The columns report three-year averages centered on the indicated years. As shown in the top row of the table, participation between 1980 and 1990 rose rapidly by nearly 10 percentage points in this decade. Since 1995, however, prime-age women's labor force participation has remained virtually constant, which suggests that the factors that lead to the rapid rise in the earlier

Table 2.3
Labor Force Participation Rates of Prime-Age Women

	Year				
	1980	1990	1995	2000	2005
Unemployment Rate	6.8	5.9	5.7	4.3	5.1
Women, Aged 25-54 Years Old	64.0	73.6	75.5	77.3	75.2
A. Marital Status					
All Married	59.3	70.6	73.6	74.4	72.6
Widowed/Divorced	74.9	79.2	79.4	82.4	79.6
Never Married	81.1	81.0	78.8	82.6	79.9
B. Mothers vs. Non-Mothers[1]					
Mothers	58.4	68.5	71.3	73.8	71.9
Non-Mothers	73.6	80.2	80.8	81.2	79.0
C. Young College Mothers vs. Never Married College Women[2]					
Young College Mothers	55.0	68.2	72.6	71.3	72.7
Never Married College Women	94.5	94.2	95.5	92.0	93.3

Source: March Current Population Survey micro data. The numbers in the columns refer to three-year averages centered on the indicated years.
[1]"Mothers" are defined as women with at least one child of their own who is aged 18 years or younger, never married, and living in the household.
[2]"Young College Mothers" refers to white married mothers aged 25–34 with a college degree while "Never Married College Women" refers to never married white women aged 40–49 with a college degree.

decades had largely played out by the mid-1990s. Their participation rate did rise from 1995 to 2000. However, the higher participation rate most likely reflected favorable cyclical conditions, as the aggregate unemployment rate in 2000 was at the lowest level since the late 1960s. During the latter half of the 1990s, welfare reform and the expansion of the Earned Income Tax Credit (EITC) also spurred strong employment growth among less-educated women. Since 2000, prime-age women's participation rate fell 2.1 percentage points. Despite articles in the popular press proclaiming that mothers are "opting out" of the U.S. labor force, participation rates fell by a smaller amount among married women than never-married women—1.8 percentage points versus 2.7 percentage points. Participation among mothers fell 1.9 percentage points, while participation fell 2.2 percentage points among non-mothers. In other words, when we take into account the weaker labor market conditions in the early 2000s by comparing the labor market behavior of married women against the behavior of non-married women, there is little evidence that the trend among married mothers—the group that fueled the increases in the earlier decades—has actually reversed and begun to decline.

Let me now turn to my reservations about the authors' cohort-based approach. Fallick and Pingle rely on a cohort-based model in which they assume a constant age profile and permanent cohort effects. This is an especially hazardous exercise for predicting women's behavior, since every birth cohort of women has a unique profile of labor force participation. The authors also use time-varying variables such as life expectancy, fertility, and the marriage rate as explanatory variables in their model. The problem is that in a reduced-form model, the coefficients are unlikely to be stable over time. To make this point, Figure 2.17 shows trends in U.S. women's labor force participation rates and marriage patterns—the percent "ever married" at age 25–29. The figure is arranged by birth cohort. Using only the earlier part of the data, we would end up with a large negative coefficient and would over-predict participation rates for the more recent birth cohorts based on marriage patterns, thereby estimating negative cohort effects. It is not clear that using information on previous cohorts leads to better predictions than are obtained by simply extrapolating from the more recent trends. To summarize, Fallick

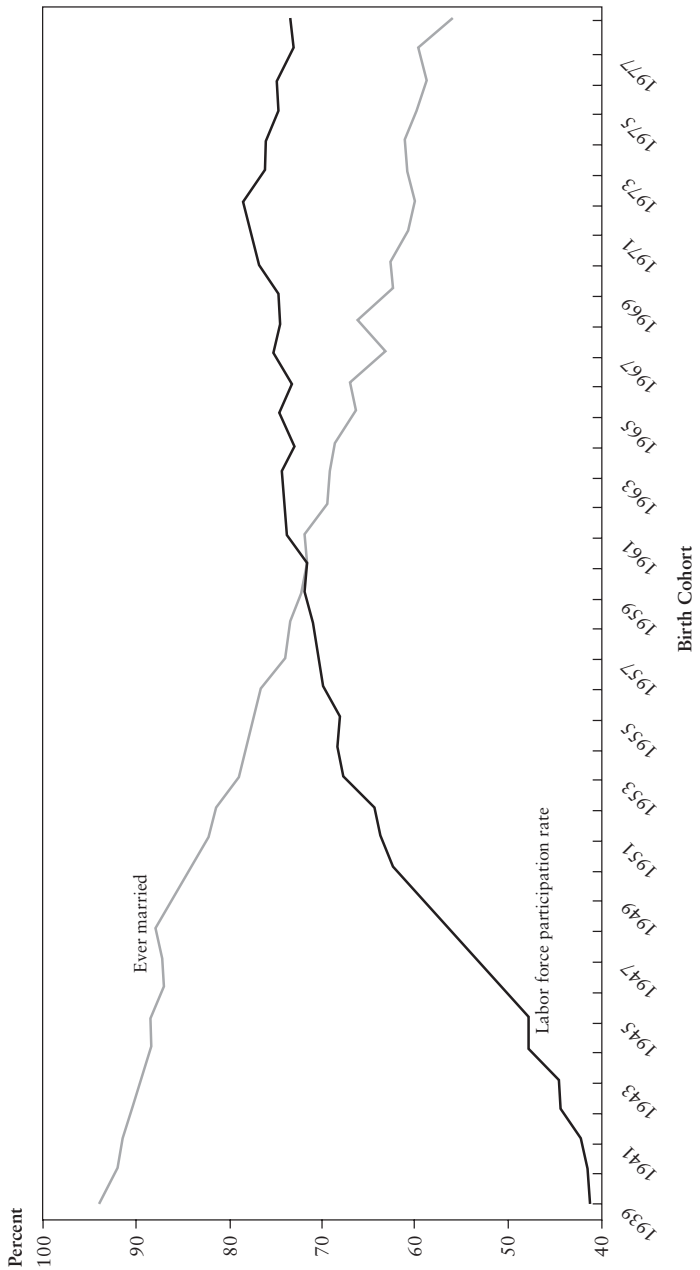

Figure 2.17
Labor Force Participation Rates and Marriage Rates for American Women Aged 25 to 29 Years, by Birth-Year Cohort
Source: Current Population Survey (U.S. Census Bureau) and U.S. Bureau of Labor Statistics.

and Pingle make an important point and demonstrate an undeniable fact that the aging of the U.S. labor force will put downward pressures on the aggregate labor force participation rate. However, in my view their projections of the within-group participation rates, particularly those of prime-aged women, are more conjectural and overly pessimistic than is warranted by a fuller examination of the data.

Finally, I would like to take a step back to look at the bigger picture. We are worried about the impact of aging and other labor force trends on potential economic growth in the United States. But in addressing this topic, we have ignored some important issues. So far we treated labor as homogenous in skill, and we have not made a distinction between shifts in demand and supply. Clearly the aging of the American labor force entails an inward shift in the aggregate labor supply curve, since older workers are less inclined to work than younger workers at the same wage. With regards to the decline in the labor force participation of prime-aged men, however, evidence points toward declining demand for their services. In view of the lackluster performance of recent wage growth, with the exception of the very top of the wage distribution, it seems not so much the lack of available working age men that is the problem. In fact, there appear to be too many men who are not very productive and employable. We may have a labor "supply" problem of sorts, but it is more of the "skills-mismatch"variety rather than simply the shortage of workers per se.

Notes

1. I draw here from a recent paper by Schirle (2007) and my own paper with Simon Potter (2007).

2. Kaisier Family Foundation and Health Research and Educational Trust (2005), also reported in Strumpf (2006).

References

Juhn, Chinhui, and Simon Potter. 2007. "Is There Still an Added Worker Effect?" Working Paper, University of Houston, June.

Schirle, Tammy. 2007. "Why Have the Labour Force Participation Rates of Older Men Increased Since the Mid-1990s?" Working Paper, Wilfrid Laurier University, Ontario, Canada, April.

Strumpf, Erin. 2006. "Employer-Sponsored Health Insurance for Early Retirees: Impacts on Retirement, Health, and Health Care." Working Paper, Harvard University, December.

Kaiser Family Foundation and Health Research and Educational Trust. 2005. "Employer Health Benefits: 2005 Annual Survey." Menlo Park, CA.

Comments on "The Effect of Population Aging on Aggregate Labor Supply in the United States" by Bruce Fallick and Jonathan Pingle

Lisa M. Lynch

Bruce Fallick and Jonathon Pingle have written a thought-provoking paper that helps set the stage for our discussion on labor supply in the twenty-first century. The paper carefully documents how—conditional on no changes in the participation behavior of workers, especially older workers—current projections of population aging will imply a decrease in the aggregate labor force participation rate in the United States of 6 full percentage points over the next 35 years. This is a huge change with significant implications for our forecasts of potential aggregate output and consequently for our conduct of monetary policy. However, before we panic too much over this large projected decrease in the aggregate labor force participation rate, there are a few caveats to keep in mind.

First, making predictions about what might happen in 35 years time, let alone 20 or even five years in the future, is a hazardous business. If 35 years ago we had used the 1972 labor force participation rates by gender and age, and had allowed the population shares to evolve as then forecast by the Census Bureau, it is doubtful that we would have come close to mapping the marked rise in labor force participation rates that the authors show in their Figure 2.3. Second, potential output is driven not only by labor force participation rates, but also by human capital and the quality of labor. Declines in labor force participation rates may not be as much of a concern in a context of rising labor productivity. Therefore, investments in education and training, along with investments in physical capital, will also play a critical role for future potential output growth.

Nevertheless, understanding the potential impact of an aging workforce is important for policymakers. The paper begins by providing a very thorough discussion of the assumptions behind the various govern-

ment projections for the U.S. population over the next 35 years. Particularly important is the discussion, as illustrated in the authors' Figure 2.8, of the impact on the projected path of labor force participation rates by adding 100,000–200,000 net new immigrants annually in the coming years to the 30-34-year-old male age group. Fallick and Pingle argue that even such a large increase to the U.S. working-age population could not reverse the underlying downward pressures on the projected labor force participation rate over the next 35 years, everything else staying constant. This is a very important simulation and one that is not that far removed from possible reforms of immigration policy in the United States. The Congressional Budget Office[1] recently released a report concluding that proposed immigration legislation in 2007 would increase the U.S. population by 1.8 million people in ten years time; this estimate translates into adding approximately new 180,000 workers a year.

Although Fallick and Pingle explain why even fairly large changes in U.S. immigration policy will have a limited impact on labor force participation rates, it is important to note that such increases would increase the overall supply of labor in the United States. In addition, their simulations assume that labor force participation rates for immigrants will be the same as for native-born American men in the age class of 30–34 years. However, we know from the Bureau of Labor Statistics that labor force participation rates for foreign-born males are much higher than for native-born males—almost 82 percent versus 72 percent for native-born workers.[2] Finally, an important component of recently proposed immigration legislation is a reconsideration of the relative priority placed on family unification versus specific professional skills for immigration policy. This potential policy change poses significant implications for future labor productivity in the United States. If we shift our immigration policy toward one that gives more weight to admitting more highly skilled people, we may not adjust the labor force participation rates significantly, but this shift would have an impact on labor productivity. So the ultimate impact on our economy of a change in immigration policy is not only on the number of people participating in the labor force, but also encompasses the human capital that is embodied in these individuals.

The paper then turns to the projections of within-age group participation rates. There are two critical parts to this discussion—what is going to happen to older workers' labor force participation rates, and what is going to happen to the labor force participation rates of men and women aged 25-54-years-old. In terms of older workers, we have seen a steady increase in the labor force participation rate of male workers over the age of 55 since the mid-1990s. I would argue, as do Munnell and Sass in their conference paper, that this increase has been fueled in part by declines in employer-provided healthcare for retirees, changes in Americans' retirement financial security due to improved health and changes in the composition of pension assets, changing occupational characteristics of workers, rising educational levels, and technologies that make working from home easier. Forecasting just how much more labor force participation rates might increase for this older age group is challenging but Fallick and Pingle, the Bureau of Labor Statistics, and the Social Security Administration all seem to agree that this participation rate will continue to rise in the near future.

However, there is much less consensus on what is going to happen to labor force participation rates of workers aged 25–54 years. Using projections from a model that relies on cohort effects to predict the future path of participation of individual birth years, Fallick and Pingle expect that labor force participation rates for prime working-age males will continue to decline, a pattern seen over the past 30 years. This prediction begs the question of why the participation rate for prime-age male workers has been falling. Without a better understanding of why labor supply for this group has been declining, it is difficult to assess whether or not extrapolating that trend forward makes sense or not. Changes in the generosity of disability insurance, incarceration rates, and school enrollment patterns may all be playing a role here. In a technical background paper associated with this conference paper, Fallick and Pingle (2007) present more details on the model they use to predict the future path of labor force participation. Specifically, in an enhanced model they show how controlling for schooling, along with age and cohort effects, is important for explaining past trends in male labor supply. However, they do not include controls for disability eligibility, changes in replacement rates associated with disability insurance, or incarceration rates. Those men in prison are not included in the numerator or denominator of the civilian labor force participation rate, so in principal changes in the U.S.

male prison population are irrelevant for male labor force projections. However, if we assume that incarceration rates are positively correlated with the pool of potential criminals, and we assume that potential criminals have a lower labor force participation rate, as suggested in work by Western, Kling, and Weiman (2001), then controlling for this variable/effect might be useful. In terms of the potential impact of enhanced disability insurance on labor force participation, we know from the work by Autor and Duggan (2003) that among males aged 25–54 years, receipts of disability insurance rose by 50 percent between 1984 and 1999. The increase was even more dramatic for those men without a high school diploma. While Fallick and Pingle have tried including the fraction of applications for disability insurance approved in each year in their model of labor force participation, they omit this variable due to its limited explanatory power. However, other proxy measures, such as the disability insurance replacement rate for low-skilled workers or the disability insurance application rate per capita, might provide more explanatory power.

The largest point of departure in the labor force projections generated by the Fallick and Pingle model versus those provided by the Bureau of Labor Statistics and the Social Security Administration has to do with the projected 25–54 year-old female labor force participation rate. Fallick and Pingle predict that the leveling off in female labor force participation that we have seen over the last couple of years will actually start trending down since, "women have begun to behave similarly to prime-age men." In large part, their claim is driven by trends in labor force participation rates that they observe for women born after 1950 when they reach the ages of 35–44 years. I want to suggest a couple of reasons why you might want to question the projections for this group that come out of Fallick and Pingle's model.

First, let me just state the obvious: since the level of women's labor force participation is not equal to men's, it is does not follow that the rate of change should exhibit the same pattern as men's. Second, Fallick and Pingle argue that more recent cohorts of women seem less attached to the labor market than past cohorts. However, in their technical background paper, the authors state that they drop the eight birth-year cohorts who entered the labor market after 1996 due to small sample sizes. I would

suggest that more recent cohorts are unlikely to be similar to previous cohorts simply on the basis of the educational choices they have been making. In Figure 2.18 we see that college enrollment rates for males and females have been trending steadily upwards since 1973, especially for women. But since the early 1990s, women have been behaving differently than men and enrolling at a higher rate in four-year colleges. The more recent cohorts of women are clearly exhibiting a greater taste for investment in higher education. This difference may actually mean that going forward, we will see a further divergence in the labor supply behavior of American men and women that reflects these different educational choices.

The role of technology is another issue that I think is quite important, but that has not been addressed in any of the papers presented here. Technology has had a profound impact on the labor supply of successive cohorts of women during the twentieth century. Improvements in maternal health and the availability of infant formula, as discussed by Albanesi and Olivetti (2007), meant that women who were born in the early part of the twentieth century spent much less time engaged in the production

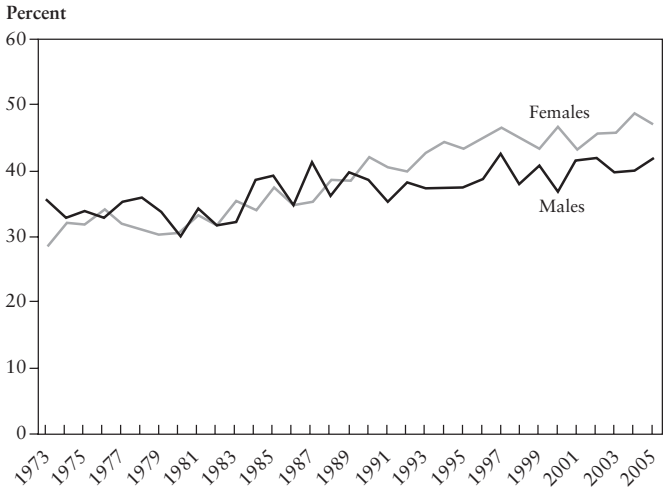

Figure 2.18
U.S. Four-Year College Enrollment Rates
Source: Digest of Education Statistics and U.S. Department of Education.

of children. This in turn allowed them to have more time to be active in the labor market. The introduction of labor-saving devices that freed women from certain household chores and the introduction of the Pill affected subsequent cohorts of women and their labor supply decisions.

Those women born after 1950 have experienced yet another technological innovation—fertility-enhancing reproductive technologies. These innovations started becoming available during the mid–1980s, when those born in 1950 were moving into the age range of 35–44 years, and insurance coverage for these treatments gradually expanded through the 1990s. These reproductive technologies have had a dramatic impact on birth rates by age of the mother, as shown in Figure 2.19. For example, if we compare the birth rates for women aged 35–39 years in 1980 and in 2005, we see that these rates have more than doubled. At the same time, teenage birth rates and birth rates for women aged 20–24 years have been declining. As women defer childbearing to later years, it is not clear whether or not the tapering off or decline in labor force participation we see in women aged 35–44 years will then reverse itself as they

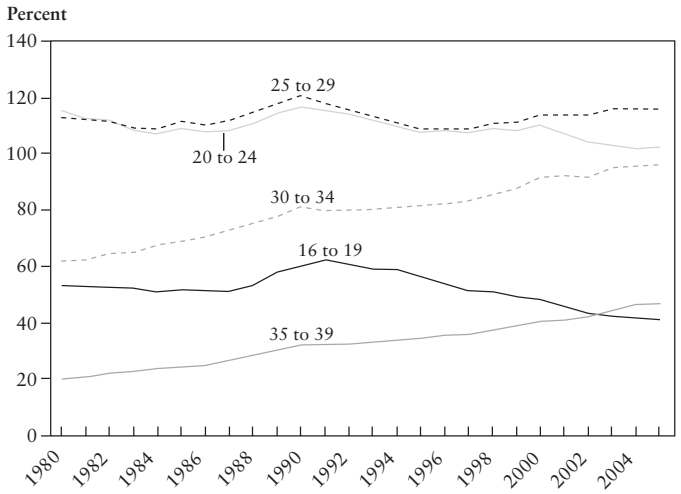

Figure 2.19
U.S. Birth Rates by Age of the Mother
Source: National Center for Health Statistics and U.S. National Vital Statistics.

get older. Fallick and Pingle try to address the impact of fertility in their extended model by including controls for fertility. However, they only do this for women aged 18–29 years. This means that they have not been able to incorporate into their projections how this technological shock may affect labor market decisions of older women. Older women who may have withdrawn from the labor market to take care of young children may exhibit very different re-entry patterns into the labor market than those who had children at a much younger age.

Finally, the aging of the U.S. working population will also have an impact on labor productivity. As workers age, they amass more experience and, depending on the amount of skills investment they have received on the job, they may also have higher human capital. These higher experience levels should, in turn, raise labor productivity. Unfortunately the failure of U.S. statistical agencies to systematically collect information on how people acquire skills in the workplace as they age limits our understanding of how declines in labor force participation may be offset by post-school investments in human capital.

In conclusion, if we have learned anything about the baby boom generation, it is that they have defied all efforts by demographers and economists to pigeonhole their labor market behavior. As they age, the boomers will continue to redefine hours of work, conditions of work, and where work is conducted. As a result, our efforts to try to model their labor market behavior going forward will constantly be challenged.

Notes

1. Congressional Budget Office, Cost Estimate. 2007. "Senate Amendment 1150 to S. 1348, the Comprehensive Immigration Reform Act of 2007." Available online at http://www.cbo.gov/ftpdocs/81xx/doc8179/SA1150_June4.pdf.

2. Bureau of Labor Statistics. 2007. "Foreign-Born Workers: Labor Force Characteristics in 2006," released April 25. Available online at http://www.bls.gov/news.release/pdf/forbrn.pdf.

References

Albanesi, Stefania, and Claudia Olivetti. 2007. "Gender Roles and Technological Progress." National Bureau of Economic Research working paper #13179, June.

Autor, David H., and Mark G. Duggan. 2003. "The Rise in the Disability Rolls and the Decline in Unemployment." *Quarterly Journal of Economics* 118(1): 157–205.

Fallick, Bruce, and Jonathan Pingle. 2007. "A Cohort-based Model of Labor Force Participation," Finance and Economics Discussion Series 2007-09, Federal Reserve Board, Washington, D.C.

Western, Bruce, Jeffrey R. Kling, and David F. Weiman. 2001. "The Labor Market Consequences of Incarceration." Princeton University Industrial Relations Section, Working Paper No. 450, January.

3

The Labor Supply of Older Americans

The Labor Supply of Older American Men

Alicia H. Munnell and Steven A. Sass

This chapter summarizes what is known about the labor supply of older American men, defined as those aged 55 years and over. The topic is of great interest because in the coming decades older individuals will comprise a much greater portion of the U.S. population, so the labor supply of older adults will have a significant impact on national output, tax revenues, and the cost of means-tested programs. Most importantly, a greater proportion of older individuals will need to remain in the workforce than is the present case, because the retirement income system is contracting and working longer is the only way for most people to ensure financial security in their old age. The paper's focus is on men, because women's work patterns are changing and increasingly reflect the work patterns of men.

Section I of this paper describes the changes to the U.S. retirement income system that will require people to work longer. Section II summarizes the long-term decline in labor force activity among older men over the course of the twentieth century, and the factors that contributed to this trend. Section III describes the recent turnaround in the labor force activity of older Americans, and the changes in Social Security and employer-provided pension plans that likely led to that reversal. In an attempt to determine whether the labor supply of older workers will continue to increase, section IV describes changes in work patterns that have emerged in the last 20 years, which have led to more labor market mobility and less job tenure among older workers, and the implications of such changes on labor supply. Section V addresses how the health of older people may influence the extent to which they can be expected to

continue in the labor force, and underscores that for 15 to 20 percent of older workers, continued employment will be impossible. Section VI discusses the remaining incentives to retire—namely, the availability of Social Security benefits at age 62 and the lack of flexible employment arrangements. Section VII concludes and estimates labor force participation rates going forward.

I. The Need for Continued Employment

As people age, earnings become dramatically less important as a source of household income, giving way primarily to income from Social Security and employer retirement income plans. Today the share of household income from earnings declines from 81 percent for those aged 55–61 years, to 57 percent for those aged 62–64, to 23 percent for those 65–69, and becomes trivial thereafter; see Figure 3.1. However, both Social Security and employer plans will replace a smaller portion of pre-retirement income in the future than is the case today. This is especially clear for

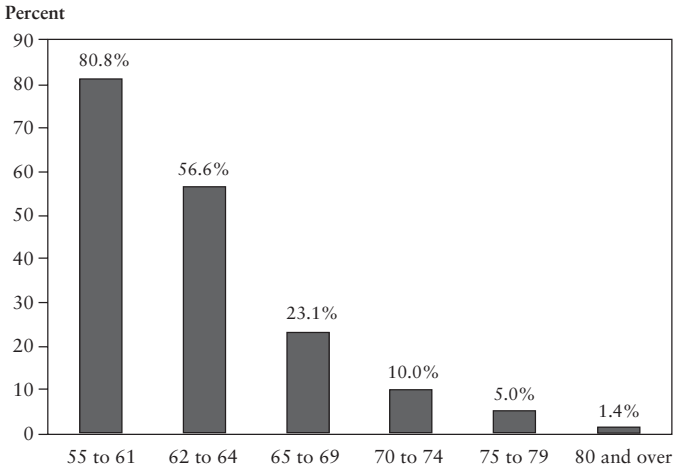

Figure 3.1
2005 Earnings as a Percent of Income, U.S. Households Aged 55 Years and Older, Middle Quintile
Source: U.S. Bureau of Labor Statistics and U.S. Census Bureau (2006), and authors' calculations.

Social Security, which is a significant source of income for most retired Americans.

The Outlook for Social Security

At any given retirement age, Social Security benefits will replace a smaller fraction of pre-retirement earnings in the future. Today, the hypothetical "medium earner" retiring at age 65 receives benefits equal to about 41 percent of his or her previous earnings. After paying the Medicare Part B premium, which is automatically deducted from Social Security benefits before the check goes in the mail, the replacement rate is 39 percent. But, under *current law,* Social Security replacement rates—benefits as a percent of pre-retirement earnings—are scheduled to decline for three reasons. First, the program's Full Retirement Age is currently in the process of moving from 65 to 67, which is equivalent to an across-the-board cut in benefits.[1] Second, Medicare Part B premiums are slated to increase sharply due to rising healthcare costs.[2] (Premiums for the new Part D drug benefit will also claim an increasing share of the monthly Social Security check.) Finally, Social Security benefits will be taxed more under the personal income tax, as the exemption amounts are not currently indexed to inflation. For the medium earner who claims benefits at age 65, these three factors will reduce the net replacement rate from 39 percent in 2002 to 30 percent in 2030; see Figure 3.2. Restoring Social Security's long-term solvency through more benefit cuts would reduce this level of support still further.

The Outlook for Private Sector Employer-Sponsored Pensions

With a diminished role for Social Security in providing retirement income, future retirees will be increasingly dependent on employer-sponsored pension plans. At any moment in time, however, less than half of the private sector workforce aged 25 to 64 years participates in an employer-sponsored retirement income plan of any type. This fraction has remained virtually unchanged since the late 1970s, and is unlikely to improve.[3] Since participation in employer-provided pension plans tends to increase with earnings, only middle- and upper-income individuals can count on receiving meaningful benefits from these plans.

The other issue is that the nature of pension coverage has changed dramatically. Twenty years ago, most American workers with pension

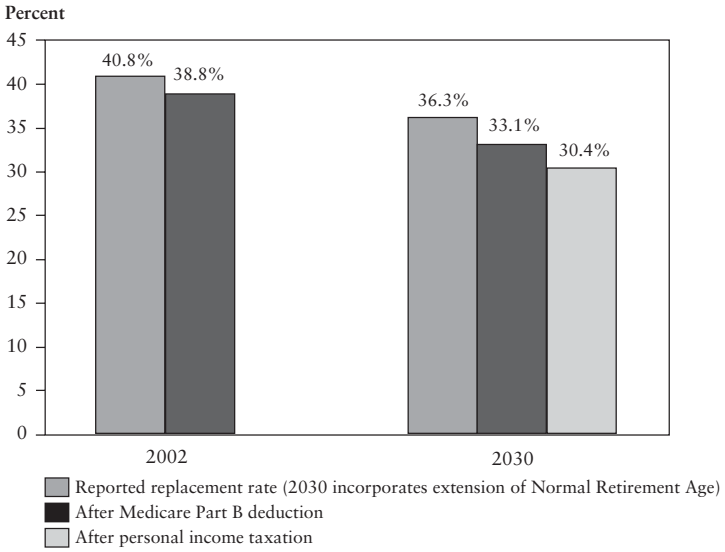

Figure 3.2
Social Security Replacement Rates for the Medium Earner, 2002 and 2030
Source: Authors' calculations based on Munnell (2003).

coverage had a traditional defined benefit plan, which pays a lifetime annuity at retirement.[4] Today the world looks very different, as depicted in Figure 3.3. Most people with an employed-sponsored pension have a defined contribution plan—typically a 401(k)—and 401(k) plans operate like savings accounts.[5] *In theory* workers may accumulate substantial pension wealth under 401(k) plans, but *in practice* they do not. For example, simulations suggest that the worker in the middle of the earnings distribution, who contributes regularly throughout his or her work life, should end up at retirement with about $300,000 in a 401(k) account and/or in an Individual Retirement Account (IRA), as most IRA assets are rolled-over balances from 401(k) plans. This $300,000, when combined with Social Security benefits, would provide an adequate retirement income. Yet reality looks quite different. The Federal Reserve Board's 2004 Survey of Consumer Finances reports that the typical individual approaching retirement had 401(k)/IRA balances of only $60,000, as shown in Figure 3.4.[6] Nor do younger cohorts seem to be on track to accumulate sufficient assets to provide an adequate retirement income. A

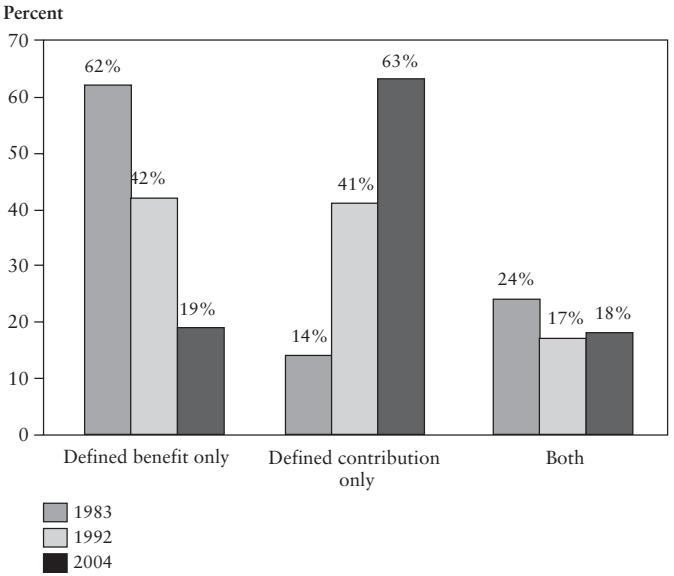

Figure 3.3
Percent of U.S. Wage and Salary Workers with Pension Coverage by Type of Plan
Source: Munnell and Sundén (2006) based on the U.S. Board of Governors of the Federal Reserve System (1983–2004).

critical factor explaining these low balances is that the entire responsibility for retirement saving has shifted from the employer to the employee, and employees make mistakes at every step along the way.[7]

Americans' Decline in Personal Saving

Given the projected decline in Social Security and the increased uncertainty surrounding employer-sponsored pensions, one might have expected to see working age adults increase their personal saving rates. This is certainly what the standard life-cycle model predicts. But a recent study of the U.S. National Income and Product Accounts (NIPA) found that saving by the working-age population has declined, and that virtually all the saving undertaken by the working-age population occurred in employer-sponsored pension plans (Munnell, Golub-Sass, and Varani 2005). In recent years, the saving rate of the working-age population outside of such plans has actually been negative; see Figure 3.5.

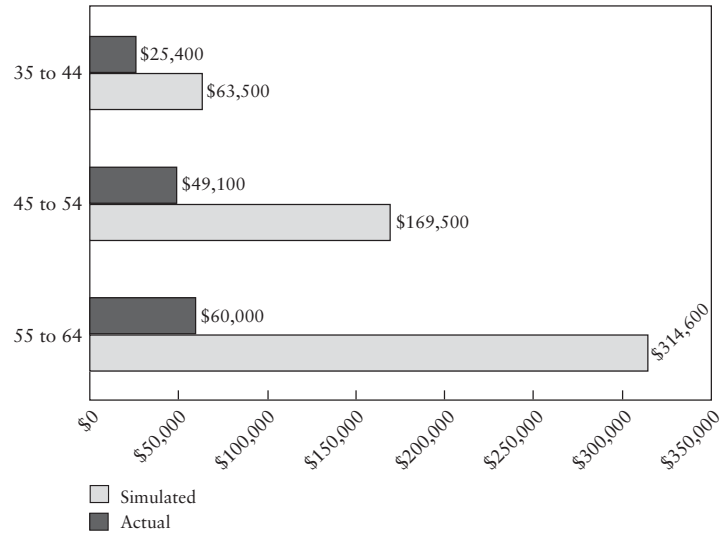

Figure 3.4
401(k)/IRA Actual and Simulated Accumulations by Age Group, in 2004
Source: Munnell and Sundén (2006).

Figure 3.5
National Income and Product Accounts Personal Saving Rate for the U.S. Working-Age Population, with and without Pensions, 1980–2003
Source: Munnell, Golub-Sass, and Varani (2005).

Thus, the outlook for retirement income for future cohorts of retirees is dismal. People are not going to be able to continue to retire at age 63 and maintain their standard of living over an increasingly long period of retirement; see Figure 3.6. Moreover, dramatically rising healthcare costs are going to erode already diminished retirement incomes.[8] Working longer is an obvious solution.[9] Each additional year in the workforce increases income directly through earnings from work and investments. It also actuarially increases Social Security benefits by 7 to 8 percent, allows retirement savings more time to accumulate investment earnings, and reduces the number of years over which those savings need to be spread. The implications are striking. As shown in Figure 3.7, a couple in the middle of the income distribution that delays retirement from 62 to 70 would reduce the assets needed to replace 80 percent of their after-tax pre-retirement income from $555,000 to $128,000.[10] Delaying retirement is clearly a powerful lever for addressing the coming decline in the nation's retirement income system. But is it realistic for most people?

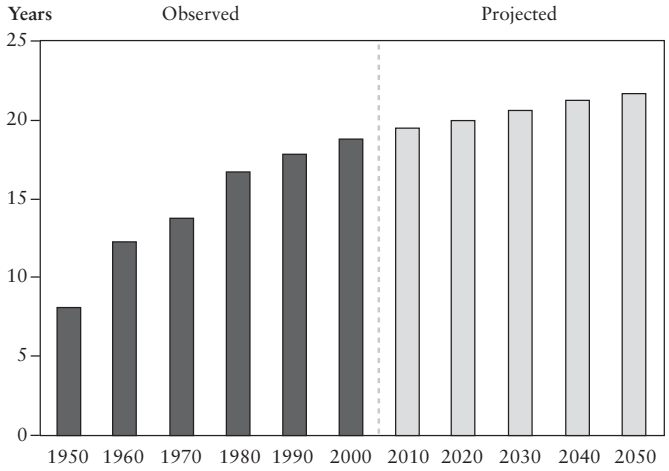

Figure 3.6
Expected Years in Retirement for American Men
Source: U.S. Bureau of Labor Statistics and U.S. Census Bureau (1962–2005), and authors' calculations based on U.S. Social Security Administration (2006).

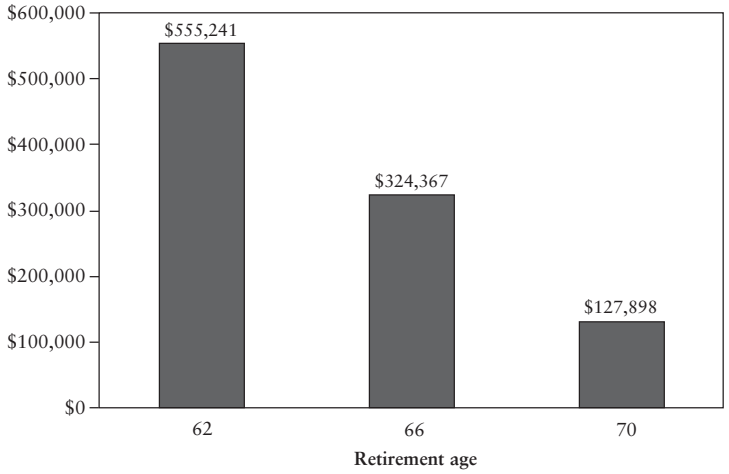

Figure 3.7
Assets Required for a Married Couple Earning $63,660 After Taxes to Maintain 80 Percent of After-Tax Pre-Retirement Income in 2007
Source: Authors' update based on Congressional Budget Office (2004).

II. The Long-term Decline in Employment Rates for Older American Men

The notion of retirement as a distinct and extended stage of life is a twentieth-century innovation. Up to the end of the nineteenth century, people generally worked as hard and as long as they could. Men in their prime put in 60 hours of work each week. And at the end of life they had only about two years of "retirement," often due to ill health. Productive capacity declined with age, as health impairments were much more prevalent and jobs much more physically demanding than is the case today. So in older age people took on less taxing jobs or worked fewer hours. But they generally stopped working only when no longer able.[11]

Beginning around the end of the nineteenth century, the percent of the older U.S. population that continued to work began to decline. This can be seen in Figure 3.8, which shows employment rates by age.[12] The employment rate among men aged 65 years and over fell from about 80 percent in 1880 to about 40 percent in 1940 to 16 percent in 1990.

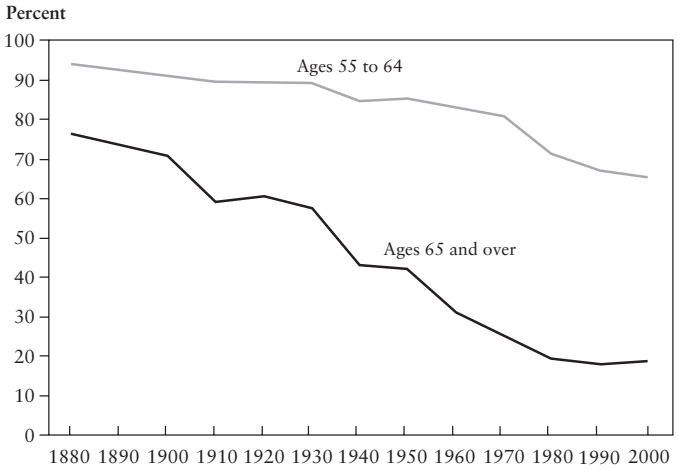

Figure 3.8
Labor Force Participation Rates of American Men Aged 55 to 64 Years, and 65 Years and Older, 1880–2000
Source: Ruggles and Sobek (2004).
Note: From 1880 to 1930, work rates are defined as reporting any gainful occupation. From 1940 to 2000, work rates are labor force participation rates, defined as working or seeking work.

Then and now, people retire for three basic reasons. Poor health may make it impossible for them to keep working. Physical strength, eyesight, hearing, and mental agility decline with age, and the incidence of contracting debilitating conditions and illnesses rises. Second, as the real or perceived productivity of older workers ebbs, employers find it unprofitable to employ them. Third, people acquire enough wealth to forgo earnings from the labor market. That is, as productivity declines and an increased incidence of ailments raises the disutility of work, older people with adequate savings can choose to quit the workforce. In terms of explaining the trend toward longer periods of retirement, increasing personal wealth and the attitudes of employers must be the primary drivers.[13] The health of older American adults has improved, not deteriorated, and would have been expected to lead to later retirement.

Economic growth has been dramatic throughout the twentieth century. Despite the Great Depression, output per hour in 1940 was 2.7 times the

level in 1880 (U.S. Bureau of Economic Analysis 1973). Workers used some of this increased affluence to reduce their labor burden. The length of the work day fell sharply between the 1880s, when the typical worker labored 10 hours a day, 6 days a week, and 1940, when the typical work schedule was 8 hours per day, 5 days per week (Costa 2000). But successful retirement requires more than rising incomes and a decision to consume more leisure. People can retire from the labor force only if they have a source of income once their earnings cease.

In theory, people could save during their working years and then tap those assets to support themselves in retirement. But this saving and investing process requires a good deal of foresight, discipline, and skill. People need to predict their earnings over their lifetime, how long they will be able to work, how much they will earn on their assets, and their life expectancy. Recent surveys suggest that even today, people are not very good at planning for retirement. Moreover, at the turn of the century most people had little reason to save for retirement since most died early, often in middle age.[14]

Instead of saving for retirement, an unexpected and substantial income stream for the elderly appeared at the end of the nineteenth century in the form of old-age pensions provided to the large number of Union Army Civil War veterans. A comprehensive study found that veterans eligible for these pensions had significantly higher retirement rates than did the American population at large (Costa 1998). It is important to note that these pensions did not require workers to retire; beneficiaries could collect these payments while remaining employed. That Union Army pensions produced an upsurge in retirements clearly illustrates the "income effect" of increased wealth on the labor supply of older workers, who often choose to consume a portion of that increased wealth in the form of more retirement.

Labor market participation rates in the United States did not return to their previous levels as the Union Army veterans died off in the early decades of the twentieth century. After a pause, the percentage of the older adult male population in the labor force continued to decline. Various analysts have argued that this reflects the growth of worker incomes (Costa 1998). But employer attitudes were also becoming more important. The U.S. workforce was rapidly shifting from self-employment, most notably as farmers, to employment in large enterprises. These organizations increasingly imposed mandatory retirement requirements on their employees, and were reluctant to hire older workers seeking employment (Moen 1987; Margo 1993).

The next big decline in the work rates of older American men, and especially of men aged 55 to 64 years, occurred after World War II, as shown in Figure 3.9. One obvious factor was the availability of Social Security benefits. Although the legislation was enacted in 1935, initially only Old Age Assistance welfare benefits were paid. Social Security's retirement benefits were not paid until 1941, and then the value of these benefits were seriously eroded by wartime and postwar inflation. The critical 1950 Social Security Amendments restored replacement rates—Social Security benefits relative to pre-retirement earnings—to 30 percent for the average earner. In the wake of the 1935 legislation, workers chose to consume a portion of their newfound Social Security wealth in the form of more retirement.

The uptick in retirement was probably also due to key features in the program design—the Social Security Retirement Earnings Test and the "take-it-or-leave-it" character of Social Security benefits. The earnings

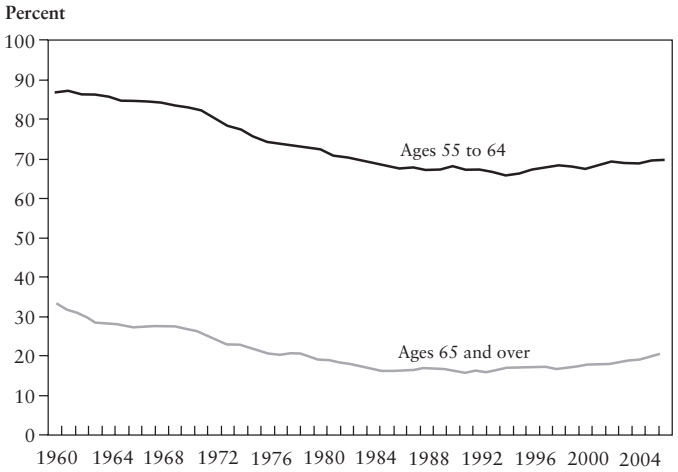

Figure 3.9
Labor Force Participation of American Men Aged 55 to 64 Years, and 65 Years and Older, 1960–2006
Source: U.S. Bureau of Labor Statistics and U.S. Census Bureau (1962–2006).

test meant that workers could not collect benefits if their earnings from work were more than a trivial sum. The "take-it-or-leave-it" character meant that a worker's benefits would not rise if he or she delayed claiming. The effective compensation of a worker who did not retire at age 65 was their compensation less their foregone Social Security benefit (and taxes and work expenses). Social Security thus decreased the value of remaining at work vis-à-vis retirement, and this "substitution effect" contributed to the decline in labor force participation. Employer pension plans had similar features and similar effects. These plans required that a worker retire in order to collect benefits, and offered no increase in benefits if a worker stayed on the job and retired at a later age.

Ultimately, Social Security's low level of earnings replacement was judged inadequate, given the widespread acceptance of retirement as a legitimate period of rest after a lifetime of work, the relative poverty of the elderly U.S. population, and the recognition that employer-provided pensions would never fill the retirement income gap. In response, Congress enacted Medicare in 1965, and in 1972 sharply increased Social Security benefits to roughly a 40 percent earnings replacement rate for the benchmark average earner.

The postwar period also saw the expansion of employer-sponsored pension plans, driven by three main considerations. First, employer-sponsored defined benefit plans had become an essential component of corporate personnel systems of large-scale organizations in the United States, so coverage grew as employment in government and corporate big business blossomed in the mid-twentieth century. Second, the special tax treatment of employer pensions became significantly more valuable in the face of mass income taxation.[15] And third, unions, which had gained powerful collective bargaining rights, made pensions a standard component of labor agreements throughout the unionized sector by the end of the 1950s.

By the early 1970s, the combination of Social Security benefits and employer-sponsored pension plans provided long-serving workers a secure and comfortable retirement income. In the wake of these developments, the labor force participation rates for men 65 and over declined from 33 percent in 1960 to 16 percent in 1985.[16]

Two factors, in addition to the sheer increase in retirement wealth created by the expansion of the retirement income system, also contributed

to a decline in the labor supply of American men in the 55–64 age group. First, many traditional employer-defined benefit plans began to offer significant subsidies for workers taking early retirement. The subsidy arises because companies offer benefits at an early retirement age, such as 55, that are not adjusted sufficiently to reflect the fact that retirees will receive benefits for 10 years longer, and begin collecting earlier, than if they retired at age 65.[17] The subsidy implicit in the less-than-actuarially fair reduction then gradually declines and disappears entirely at the plan's normal retirement age.[18] By decreasing the value of remaining at work vis-à-vis taking retirement, this produces a strong incentive to retire early.

The second factor affecting labor force participation rates for men aged 55–64 years was the availability of Social Security benefits at age 62. When in 1935 Congress established 65 as the age of eligibility for Social Security benefits, it was following precedents set internationally and by employer-sponsored plans. But in 1956, Congress lowered Social Security's Earliest Eligibility Age (EEA) for women to 62.[19] The introduction of an EEA for men followed in 1961, primarily in response to a recession that left many older male workers without employment. These early retirement benefits are actuarially adjusted, and thus involve no clear increase in retirement wealth. But numerous empirical studies, showing a spike in retirements at age 62, support the notion that the availability of benefits at 62 was an important factor in reducing the labor force participation rate of men aged 55–64 years (see Gustman and Steinmeier 1986, Rust and Phelan 1997, Burtless and Quinn 2000).[20]

III. The Recent Trend Reversal in Older Men's Labor Force Participation

The decline in the labor force activity of older American men ended in the mid-1980s. As shown in Figure 3.10, which depicts men's labor force participation rates by age for 1940, 1970, 1985, and 2005, labor force activity at each age was below that for the earlier period until 1985. The pattern then reversed, with older men's labor force participation rate in 2005 above the 1985 level for those 62 and over.[21] Observers have offered a number of explanations for this change in direction (Friedberg 2007; Burtless and Quinn 2002). We discuss some of these changes in this section.

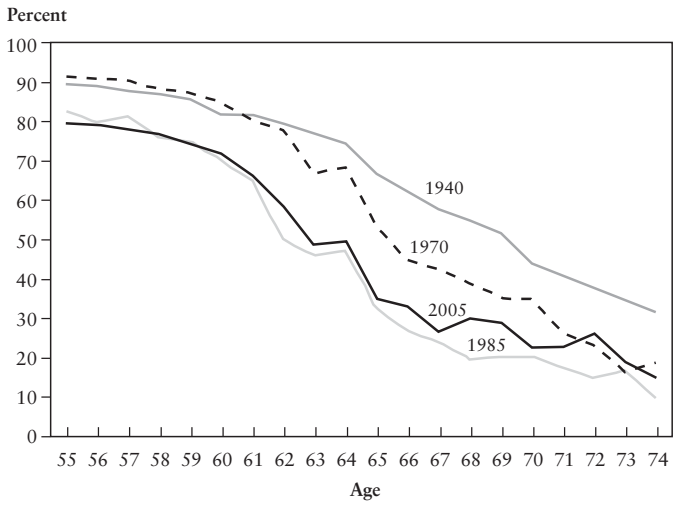

Figure 3.10
Labor Force Participation Rates of American Men Aged 55 Years, and Older, 1940–2005
Source: Authors' calculations based on the U.S. Bureau of Labor Statistics and U.S. Census Bureau (1985, 2005) and Munnell (1977).

Changes in the Social Security Program

Social Security benefits available at any given age have become less generous, and incentives for early retirement have been reduced or eliminated.

Two changes enacted in 1983 have reduced benefit amounts. First, this legislation made up to half of Social Security benefits taxable for people with earnings above a certain threshold.[22] For higher income beneficiaries, the taxable percentage was increased to 85 percent in 1994. For these higher paid workers, subjecting their Social Security benefits to taxation is equivalent to a benefit cut. Second, the 1983 legislation gradually increased the Full Retirement Age from 65 to 67, which is equivalent to an across-the-board benefit cut. Once the increase is fully phased in, for cohorts born in 1960 and later, those retiring at age 62 will receive 70 percent, as opposed to the original 80 percent, of full benefits.

The expected negative "income effect" of such benefit cuts is an increase in the labor supply of older Americans, as workers respond to this decline in wealth in part by consuming less while working, in part by consuming less in retirement, and in part by working more and "consuming" less retirement. But the labor-supply effects of these benefit cuts mainly lie in the future. The Full Retirement Age only began rising for those turning 62 in 2000, and that year the benefit reduction was small.[23] The increased taxation of benefits will also affect a much larger share of the population in the future, as the income thresholds are not indexed for inflation.

The increase in older men's labor force participation since the mid-1980s is more likely due to changes in the Social Security program that made continued work more attractive vis-à-vis retirement. The first change is the liberalization and, for some, the elimination of the earnings test. Since Social Security began as a program insuring workers and their dependents against a loss of earnings due to disability, old age, or death, the government imposed an earnings test: benefits were paid only if earnings were "lost." This test, however, encouraged workers to retire early, because it seemed like a tax. Most workers were unaware that any reduction in the amount of benefits paid out due to their continued employment triggered an increase in benefits later.[24] In recent years, Congress increased the exemption amount that workers could earn without having their benefits reduced. And, for beneficiaries older than the Full Retirement Age, it first reduced the benefit reduction for each dollar earned and then eliminated the test altogether in 2000. For those between age 62 and the Full Retirement Age, the test allows about $12,500 of earnings before reducing benefits by $1 for each $2 of earnings. Most studies suggest that the earnings test and these changes have had a substantial impact on the work effort of older people (see Friedberg 1998 and 2000; Haider and Loughran 2005; Friedberg and Webb 2006; Gustman and Steinmeier 2007), though some conclude that the test has had little effect, at least for older men (Gruber and Orszag 2003).

The Delayed Retirement Credit, which increases benefits for each year an individual postpones claiming Social Security benefits between the Full Retirement Age and age 70, has also improved older workers' incentives to remain in the labor force. When introduced in 1971, the credit increased benefits by 1 percent per year for each year of delay between the Full Retirement Age and age 72. In 1983, Delayed Retirement Credits were only granted up to age 70, but the adjustment was raised to 3

percent per year, and scheduled to increase to 8 percent per year by 2008. When fully phased in, the credit will, roughly speaking, be actuarially fair. The question then becomes what impact this increased credit for delaying claims will have on retirement decisions. Recent studies suggest that the delayed retirement credit may well have been an important factor in raising labor force participation rates among workers 65 and over (Coile and Gruber 2000; Pingle 2006).[25]

The End of Mandatory Retirement

In the early 1970s about half of all employed Americans were covered by mandatory retirement provisions that required they leave their jobs no later than a certain age, usually 65. In 1978, the earliest legal age for mandatory retirement was increased from 65 to 70. In 1986, mandatory retirement was eliminated entirely for the majority of workers. As nearly all American workers in 1986 and after were out of the labor force by age 70, however, this legislation probably had little to do with the subsequent rise in the labor supply of older workers.

Changes in Employer Pension Plans

Various changes in the structure of employer-sponsored retirement income plans have also reduced incentives to retire early. As noted earlier, in the early 1980s about 85 percent of U.S. workers with employer-sponsored pensions were covered by a defined benefit plan; by 2004 the percentage of U.S. workers with defined benefits plans had declined to 37 percent. In contrast to the early retirement incentives commonly found in defined benefit plans, 401(k)s and other defined contribution plans work like savings accounts and contain no incentives to retire at any particular age. Studies have documented that, on average, workers covered by 401(k) plans retire a year or two later than do similarly situated workers covered by a defined benefit plan (see Friedberg and Webb 2005, Munnell, Cahill, and Jivan 2003). Among recently retired workers, however, dependence on defined contribution pensions had not increased dramatically. Thus the labor supply effect of the shift from defined benefit to defined contribution plans primarily lies in the future, not in the past.[26]

Another likely change, albeit poorly documented, is a shift since the mid-1980s away from sweetened early retirement benefits in traditional defined benefit pension plans. According to one industry expert, the elimination of such early retirement incentives was a primary motive behind the conversion of a large number of pension plans, covering over 20 percent of covered workers, to cash-balance formats (see Schieber 2007). From the perspective of workers, cash-balance plans are much like defined contribution plans and neither subsidize nor penalize retirement taken at any given age.[27] In addition, many early retirement sweeteners in the past had been offered in special "one-time" windows. If the conversion to cash-balance formats does reflect a shift away from early retirement subsidies, one would expect a comparable shift away from such one-time offers.[28] The net effect could be an increase of one to two years, much like the effect of a shift from defined benefit to defined contribution plans.

The Shift to Less Physically Demanding and More Psychologically Rewarding Jobs

The nature of employment has changed dramatically in the last 20 years. As U.S. manufacturing industries have declined, the service sector has exploded. This shift, especially the expansion of knowledge-based employment, reflects the growth in jobs often thought to have significant non-pecuniary rewards, found in places such as universities and hospitals, and in occupations such as software development, management consulting, and graphic design. Even within the manufacturing sector, the composition of jobs has changed, as firms have automated or outsourced production and now employ more managers, engineers, and technicians.[29] Generally, American jobs now entail more knowledge-based activities that put less strain on older bodies, and provide more satisfaction for workers of all ages.[30] Less physical strain and more non-pecuniary rewards raises the value of remaining employed vis-à-vis taking retirement, thereby raising the supply of labor. A good portion of the increase in labor force participation since the mid-1980s, especially among workers aged 65 to 69 years, the group which saw the most dramatic gains in labor force participation, may be due to such changes.

Joint Retirement Decisions

Another factor that may be encouraging men's employment at later ages is the movement of married women into the labor force. When only the husband was working outside the home, retirement decisions could be

based on the rewards of work, the generosity of his retirement benefits, and how continued employment would affect those benefits. With wives working, retirement decisions have become more complicated. Now couples need to consider how the decision to stop working will affect the rewards and benefits of both spouses. A growing number of studies suggest that husbands and wives like to retire together.[31] Since in the United States husbands are, on average, three years older than their wives, the increased labor force participation of wives would be expected to lead to the later retirement of men.

The Decline in Post-Retirement Health Insurance

A final factor affecting the labor force participation rates for older men is related to changes in employer-provided health insurance. Among the entire working-age population, employer-provided health insurance coverage may be declining, but it is declining very slowly. In contrast, employer provision of health insurance after retirement has dropped dramatically. According to the Kaiser Family Foundation, the percent of firms with 200 or more employees offering retiree health insurance dropped in half between 1988 and 2005; see Figure 3.11. This drop dramatically changes the incentives facing workers in their late 50s and early 60s. If they stay with their employer, they will continue to receive health insurance. If they leave the workforce before age 65, when they qualify for Medicare, they will be uninsured and forced to purchase insurance on their own—a very expensive undertaking. The combination of a decline in retiree health insurance coverage with the rapid rise in healthcare costs gives workers a strong incentive to maintain their current employer-provided coverage until they qualify for Medicare at 65.

In short, a large number of factors could explain the increase in labor force participation among older male workers since the mid-1980s. The contraction of the retirement income system, which increases participation via an "income effect," is an effect that will take place mainly in the future. But substantial changes in Social Security benefits and employer-provide pension plans have raised the value of work vis-à-vis retirement, which increases labor force participation via a "substitution effect." The fact that the increase in older men's participation has occurred mainly after age 62, and especially after age 65, suggests that changes in the

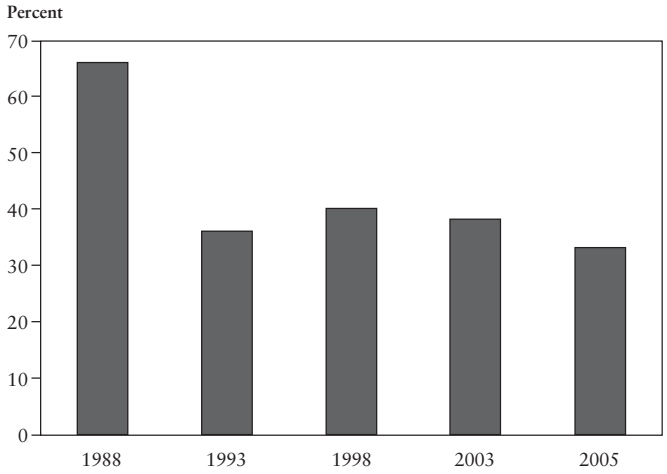

Figure 3.11
Percent of U.S. Firms with 200 or More Workers Offering Retiree Health Benefits, 1988–2005
Source: Kaiser/Health Research and Educational Trust Employer Health Benefits Survey: 2003, 2005; KPMG Survey of Employer-Sponsored Health Benefits: 1988, 1993, 1998.
http://www.kff.org/insurance/7315/sections/upload/7315Section11.pdf.

Social Security earnings test and the Delayed Retirement Credit have been quite important; refer again to Figure 3.10. On the other hand, a recent study focusing on this older segment of the workforce suggests that non-pecuniary considerations might also play an important role (see Haider and Loughran 2001). Older labor force participants tend to be among the more educated, healthiest, and wealthiest elderly Americans. Moreover, the fact that the wages earned at these older ages are lower than those of their younger counterparts, and lower than their own past earnings, suggests that money may not be the prime motivator for their continued labor force participation.[32]

The important question is whether this trend toward later retirement will continue, and whether U.S. workers will respond to the contraction of the retirement income system by remaining in the workforce longer. Boomers certainly claim that they will want to work longer, but will they follow through with their plans?[33] To provide some basis for predicting

future labor force trends, the following sections look at how career patterns have changed over time, the physical health of older workers, and the remaining incentives to retire early.[34]

IV. Patterns of Employment

The above discussion has focused on the labor force participation of older male workers. Another dimension of work patterns is the extent to which and when people change jobs over the course of their working life. This pattern is important because older workers are likely to have an easier time staying employed and enjoy higher wages if they remain with their long-term employer rather than scurrying about the labor market trying to find a new job in their late 50s and early 60s. Evidence suggests that firms are reluctant to hire older workers, and the loss of firm-specific human capital means that productivity, and hence wages, often fall when workers move to a new job (see Lahey 2006, Johnson and Kawachi 2007).

Tenure Patterns

Despite the apparent interest of older workers in remaining with their current employer, one would expect to see shorter tenures and more mobility as a result of the shift from defined benefit to defined contribution plans. The shift in employer-provided retirement plans reflects a diminished interest in career employment on the part of both firms and employees. The original purpose of defined benefit plans was to induce workers to remain with their employer until retirement, then to retire "on time" at the age specified in the plan (Sass 1997). To accomplish this goal, plans based benefits on years of service and earnings in a worker's final working years, so the value of accrued pension benefits increased rapidly as job tenures lengthened and earnings rose, and then declined as workers aged past critical benchmark ages. Workers with defined benefit plans who change jobs prematurely, even when moving to firms with identical plans and immediate vesting, receive significantly lower benefits in retirement than do workers with continuous coverage under a single plan. Both the changing mechanics of employment and changing tastes would lead one to expect more worker mobility and shorter job

tenures in a 401(k)-dominated world. These shifting incentives would be expected to affect primarily older workers, since at younger ages the pension costs of switching jobs have always been minimal.

This expectation is borne out in the median tenure data for employed males taken from the Current Population Survey and presented in Figure 3.12.[35] The results are striking in two respects. First, before 1990 the median years of job tenure is virtually flat for every male age group between age 25 and age 64. These data confirm much of the earlier work on mobility that showed very little change during the 1970s and 1980s (see Neumark 2000, Gottschalk and Moffitt 1999). Second, beginning in 1990, after a decade of 401(k) plans being in place, the median job tenure for men at older ages (55+) starts to decline. If the shift in pension coverage from defined benefit plans to defined contribution plans were to have an effect, this is where and when one would expect to find it. As noted above, pension accumulations are very small at younger ages, and never really impeded mobility among younger workers, so the shift in the

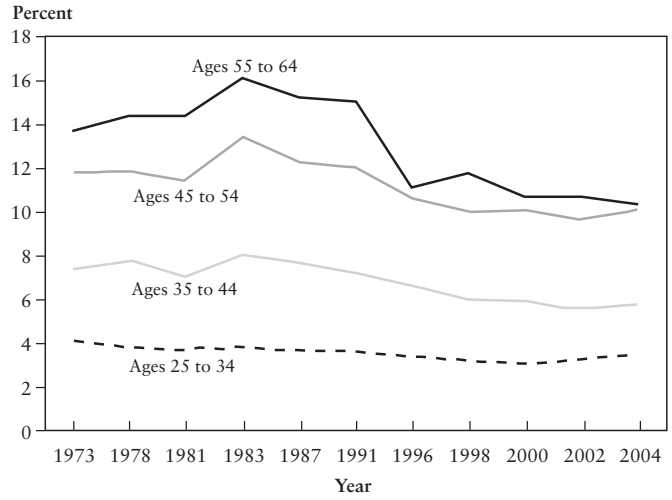

Figure 3.12
Median Years of Tenure for Employed American Men by Age, 1973–2004 (Current Population Survey data)
Source: Authors' calculations from U.S. Bureau of Labor Statistics and U.S. Census Bureau (1973–2004).

type of pension coverage would affect the mobility only of older workers.[36] Similarly, we would not expect the effect to become evident until a significant percent of older workers were covered by 401(k) plans, and this did not happen until the 1990s.

The Current Population Survey data can also be used to see how many workers remain with the employer they worked for when they were age 50.[37] The results for the years 1983 and 2004, which are shown in Figure 3.13, mirror the tenure information presented above. In the early survey, at age 60, almost 80 percent of male workers were working for the same firm as they were when they were 50 years old. By 2003, the picture changes noticeably; at age 60 less than 45 percent were working full time with their age-50 employer. In short, male workers in their 50s appear to be shifting jobs more in a pension world dominated by 401(k) plans than they did when covered by defined benefit plans. The old notion that men settle into some form of lifetime employment by middle age and stay there through retirement no longer holds for the majority of older

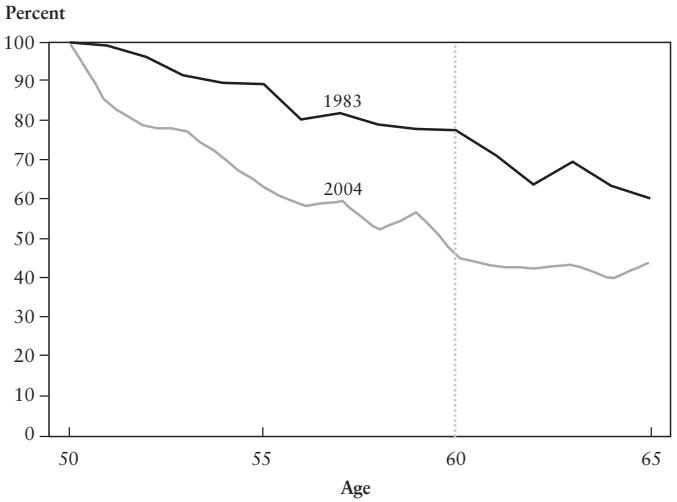

Figure 3.13
Percent of American Men Working Full-Time Who Remained in the Same Job since Age 50, 1983 and 2004
Source: Authors' calculations based on U.S. Bureau of Labor Statistics and U.S. Census Bureau (1983, 2004).

American men. One question this prompts is the extent to which this job switching at older ages is voluntary. That is, do workers move on their own volition or are they laid off from a long-held job, and forced to find a new one? One measure of layoffs is displacement rates. Have job displacement rates increased over time?

Displacement Rates

The Displaced Worker Surveys attempt to measure the number of workers who have lost their job through no fault of their own.[38] The displacement rates for older workers, while cyclical, show no discernable upward or downward trend over the period 1984–2004; see Figure 3.14.

Simple averages, however, cannot reliably indicate whether the plight of older U.S. workers is getting better or worse, because many factors are changing simultaneously. For example, the educational gap between older and younger workers has virtually disappeared, which suggests that older workers—all else remaining equal—should be less likely to be laid off. On the other hand, the shift away from career employment—defined as employment with a single firm from middle-age (at the latest) until

Figure 3.14
Job Displacement Rates by Age Group, 1984–2004
Source: Munnell, Sass, Soto, and Zhivan (2006) based on U.S. Bureau of Labor Statistics and U.S. Census Bureau (2004).

retirement—suggests older workers would be more likely to be laid off. In order to isolate the impact of age on displacement rates, it is necessary to control for the various ways in which older workers might differ from their younger counterparts. This can be done through the use of a probit regression that estimates the probability of being displaced from one's job, and includes variables for gender, marital status, race, education, industry, and full-time status as well as age.[39] Controlling for these other factors, Figure 3.15 shows the effect of age on the probability of being displaced. Being in the 50-54-year-old age group reduces the probability of being displaced by somewhere between 0 percent and 7 percent. Interestingly, the beneficial effect of increased age on job tenure appears to be declining over time.[40] Thus, the results suggest that older workers are slightly more likely to be laid off today then they were in the past.

But that is not the end of the story. Figure 3.16 reports the results for the same type of equation, but this time includes tenure variables, and

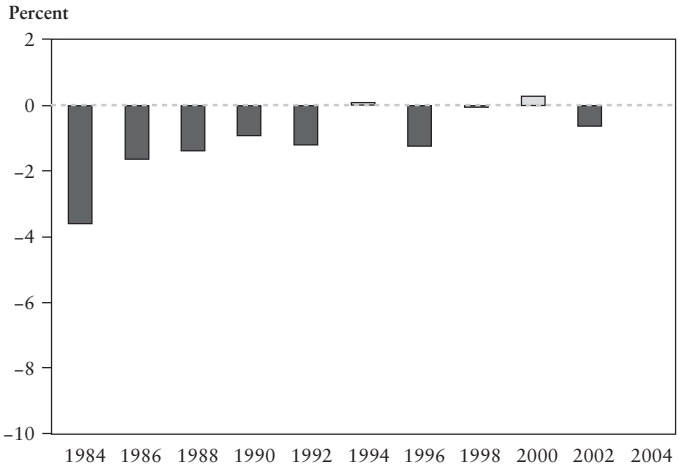

Figure 3.15
Probability of Job Displacement for U.S. Workers Aged 50 to 64 Years, Compared with the Probability for U.S. Workers Aged 20 to 49 Years, 1984–2004 (Displaced Worker Survey)
Source: Munnell, Sass, Soto, and Zhivan (2006) based on U.S. Bureau of Labor Statistics and U.S. Census Bureau (2004).
Note: Gray bars indicate results that are not statistically significant.

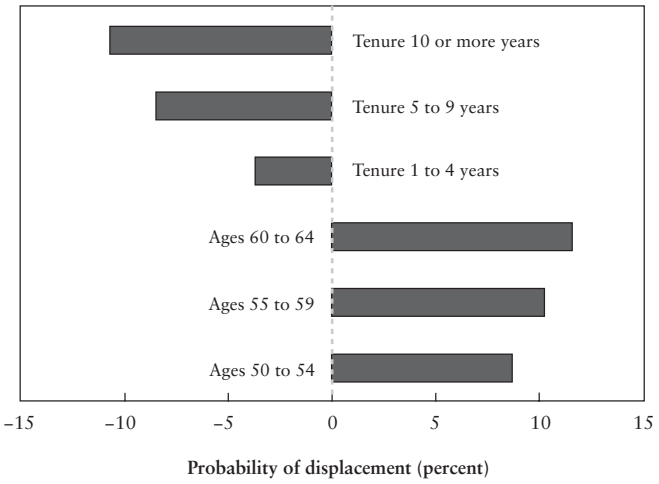

Figure 3.16
The Effect of Job Tenure and Age on the Probability of Displacement for Older U.S. Workers, 2004 (Displaced Worker Survey)
Source: Munnell, Sass, Soto, and Zhivan (2006) based on U.S. Bureau of Labor Statistics and U.S. Census Bureau (2004).
Note: All results are statistically significant.

shows that job tenure—not age—protected older workers from displacement. Holding tenure constant, older workers are actually more likely than their younger counterparts to be displaced.[41] Thus, to the extent that workers change jobs late in their careers, they are increasing their risk of eventual displacement. These older workers lose the protection afforded by long-term tenure and face the increased risk of displacement associated with age. Involuntary displacement has an extremely negative effect on the probability of older workers getting another job (Chan and Stevens 2001). This reduced probability could be the result of workers not being willing to supply their labor at the lower wages they are offered in the labor market, or of employers being unwilling to hire displaced older workers. It is very difficult to untangle the effects of labor supply and labor demand. But it appears that older workers have already experienced some increase in displacement risk, and put themselves more at risk when they change jobs. Therefore, not all of the increase in mobility among older workers appears to be voluntary.

Changes in Compensation and the Effect on Labor Supply

Two recent changes—the rapid rise in the share of older workers in the labor force and the decline of career employment—could significantly affect the compensation received by older workers, and thereby their labor supply.

The share of older workers in the U.S. labor force is increasing significantly. According to the Bureau of Labor Statistics, workers aged 55–64 years rose from 9 percent of the workforce in 1990 to 14 percent today, and are projected to exceed 18 percent in 2020, as shown in Figure 3.17.

Economic theory suggests that the age distribution of the workforce affects the wage structure, and the relative wages of older workers do appear to be inversely related to the share of older workers in the composition of the labor force. The notion here is that workers with different amounts of labor market experience are imperfect substitutes for each other. More experienced workers, who have acquired on-the-job training or simply learned by doing, generally perform different tasks and play

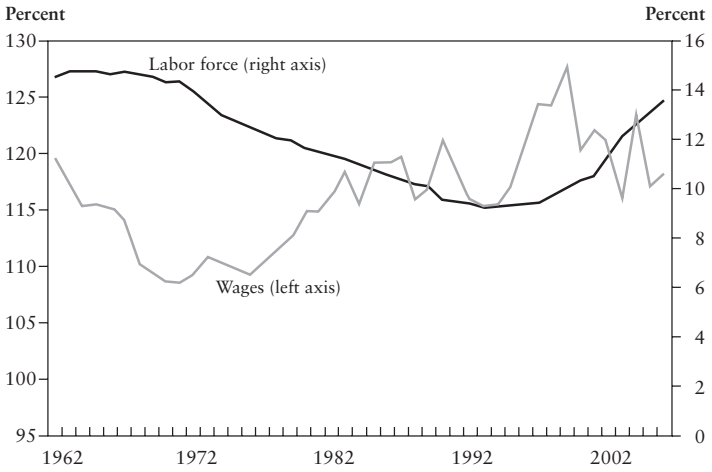

Figure 3.17
Labor Force Shares and Wages of Men Aged 55 to 64 Years as a Percent of Male U.S. Workers Aged 64 Years and Younger, 1962–2006
Source: Authors' calculations based on U.S. Bureau of Labor Statistics and U.S. Census Bureau (1962–2006) and U.S. Bureau of Labor Statistics (2007).
Note: Wages are for those who graduated from high school.

different roles within the organization. As the supply of workers with a given level of experience grows, the wages of that group will decline relative to the rest of the workforce, producing a cohort effect. The magnitude of the wage decline will depend on the extent to which workers with different degrees of experience can substitute for each other.

A number of studies have examined how relative wages have changed as the baby boom generation first entered the market and then aged. A now-famous analysis, subtitled "The Baby Boom Babies' Financial Bust," found that the wages of young white men were reduced relative to those of older white men as the baby boomer cohorts started entering the labor market (Welch 1979).[42] A recent study found that the depression of wages due to cohort crowding follows workers throughout their careers (Triest, Sapozhnikov, and Sass 2006). Thus, it seems reasonable to conclude that the increasing share of older workers in the labor force will depress their wages relative to those of younger workers.

Two further comments are required regarding the cohort crowding effect. First, the shift away from defined benefit plans has reduced the relative *compensation* of older workers even more than indicated by the decline in relative wages. Pensions in defined benefit plans are based on tenure and final salary, and become more costly to the employer as workers approach retirement; so the value of non-wage pension compensation in defined benefit plans rises rapidly at the end of workers' careers. The shift to 401(k) plans has eliminated such differential non-wage compensation received by older workers, which reinforces the finding that the increasing share of older workers in the labor market has an adverse effect on their market value. On the other hand, U.S. labor force growth in general is slated to slow. It is possible that the supply of labor may fall short of demand, thereby putting upward pressure on labor compensation, an effect that could mitigate some of the downward pressure on the compensation of older workers. On balance, however, both the experience premium and pension gains enjoyed by older workers will likely be lower in the future. As a result, work will look less desirable for older Americans relative to retirement and, as a result, they may be less willing to supply their labor.

The second labor-market change that could affect the labor supply of older men is the decline of career employment. This change, which was discussed above, is depicted clearly in Figure 3.18, which classifies the male population aged 55–64 in 1983 and in 2004 as: a) not working; b)

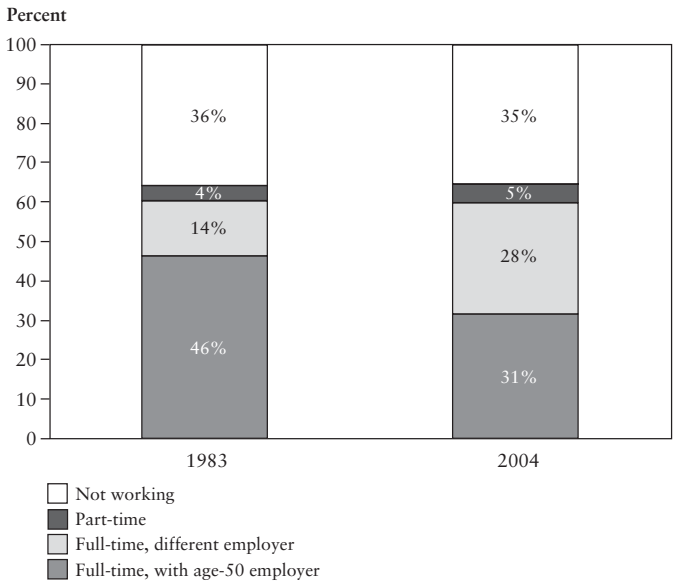

Figure 3.18
Employment Patterns of American Men Aged 55 to 64 Years, 1983 and 2004
Source: Authors' calculations based on U.S. Bureau of Labor Statistics and U.S. Census Bureau (1983, 2004).

working part-time; c) working full-time with same employer as at age 50; or d) working full-time with a different employer. The portion of this age group not working (36 percent–35 percent) or working part-time (4 percent–5 percent) was virtually identical in 1983 and 2004. But the distribution of full-time workers changed dramatically. In 1983, most full-time workers aged 55–64 years were with their age-50 employer, while in 2003 only about half of this same age group was with their same employer.[43]

This increase in mobility would be expected to impact wages. Separations from long-term employment relationships involve a loss of firm-specific human capital. Job changes also involve a loss of seniority-based protections that shield older workers from the consequences of skill erosion. Thus, a shift to a new employer would seem to suggest a fall in wage and benefit compensation. A simple comparison of wages for full-time workers who switch jobs with those who do not reveals that over the

1983–2004 period, the wages of job switchers averaged about 75 percent of those for full-time workers who remained with their age-50 employer; see Figure 3.19.

Interestingly, regardless of the reason why workers leave a long-term employer, their subsequent earnings tend to decline. A recent study (Johnson and Kawachi 2007) used the Health and Retirement Study (HRS) to explore the effect of job changes on wages, benefits, and satisfaction among workers aged 45–75 years who changed employers between 1986 and 2004. Figure 3.20 shows how those older workers leaving jobs held for more than 10 years were distributed by age and by reason for separation: retirement, layoff, voluntary quit, and "involuntary quit" (health, family reasons, personal problems, dissatisfaction with working conditions, etc.) Workers were characterized as retired if they said they left their previous job to retire. Most of the moves recorded in the HRS occurred among workers aged 51–60. Retirements accounted for about one-third,

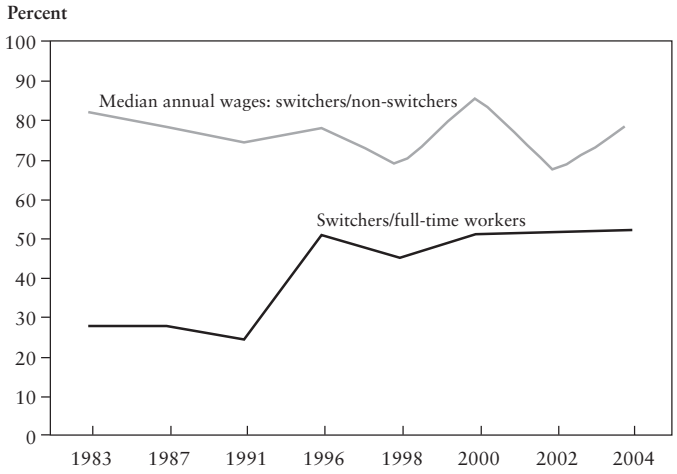

Figure 3.19
Percent of American Men Aged 55 to 64 Years Working Full-time Who Switched Jobs, and Switchers' Wages as a Percent of Non-Switchers' Wages, 1983–2004
Source: Authors' calculations based on U.S. Bureau of Labor Statistics and U.S. Census Bureau (1983–2004).
Note: A "switcher" is one who no longer works for his age-50 employer.

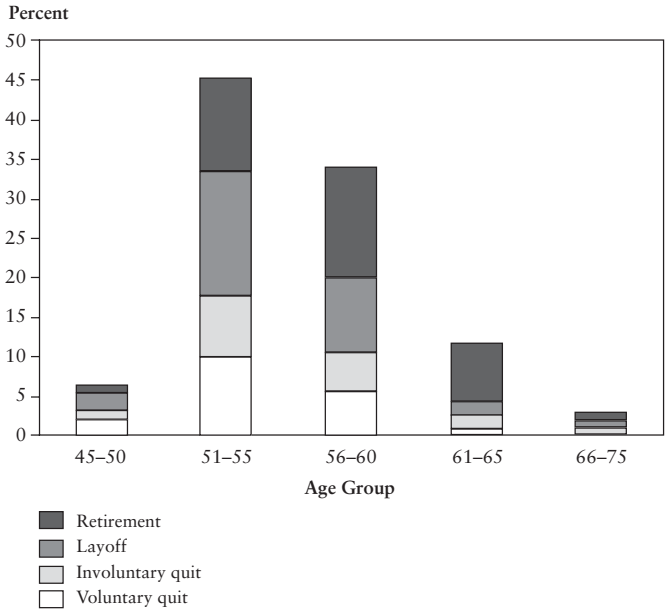

Figure 3.20
Percent of Older American Workers Who Changed Jobs, by Age and Reason for Separation, 1986–2004
Source: Authors' calculations based on Johnson and Kawachi (2007).
Note: Figure refers to those workers in the Health and Retirement Study whose former job lasted more than 10 years. "Involuntary quit" includes leaving job because of relocation, poor health and disability, family or child care responsibilities, marriage, spouse's preferences, personal problems, or dissatisfaction with work hours or length of commute.

layoffs for about one-third, and voluntary and involuntary quits for the remaining third of job separations.

Intuitively, one would think that the relationship between the reason for leaving a job and the workers' subsequent wages would be as follows: wages would fall sharply in the case of retirement, because the purpose of leaving is to work less hard. The second biggest decline would occur in the case of layoffs, because displaced workers usually face a costly search process and end up in an inferior position. A smaller decline might occur among those who quit for personal or health reasons. Finally, one might expect no decline and even an increase in wages for those who quit

voluntarily, presumably to accept "better jobs" with higher compensation and/or more non-monetary rewards. The percent losing pension and health benefits would be expected to follow a similar pattern. Figure 3.21 confirms the expected pattern, with the exception that even those who quit voluntarily suffer some drop in wages.

The conclusion that emerges from this evidence is that increased job mobility, like the effect of cohort crowding, will mean that the relative compensation of older workers will likely decline. This lower compensa-

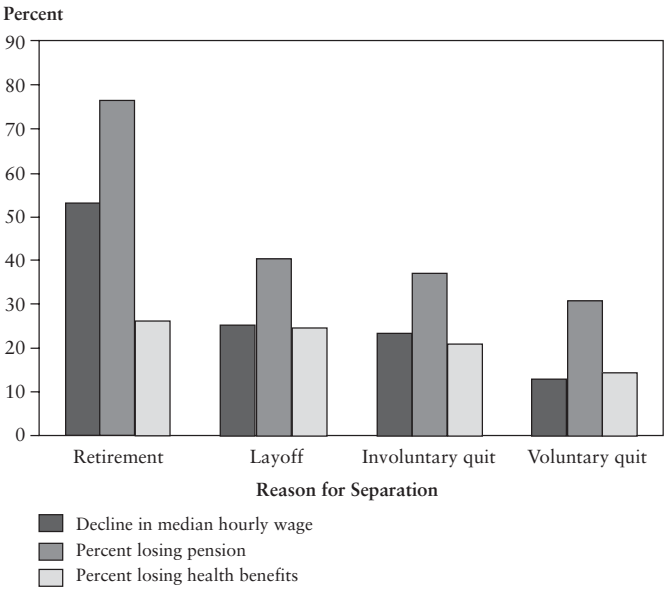

Figure 3.21
Percent Decline in Wages and Fringe Benefits Among Older American Men Who Changed Jobs after 10 Years or More with Former Employer
Source: Authors' calculations based on Johnson and Kawachi (2007).
Note: This figure shows wage and fringe benefit changes for workers whose former job lasted more than 10 years. Loss of pension encompasses those who were covered by a pension on their old job but not on their new job. Loss of health benefits encompasses those who were covered by health benefits on their old job but not on their new job."Involuntary quit" includes leaving job because of relocation, poor health and disability, family or child care responsibilities, marriage, spouse's preferences, personal problems, or dissatisfaction with work hours or length of commute.

tion can be expected to reduce the willingness of workers to supply their labor at older ages. One confounding effect of this conclusion, however, is that laid-off workers, as well as those who quit, report significant nonpecuniary gains (Johnson and Kawachi 2007). The new jobs tend to be less stressful and less physically demanding than their old ones. And more workers report that they enjoy work.

More Heterogeneity in Labor Supply at Older Ages

The shift to 401(k) plans and the increased mobility of older workers also means that in the future retirement is going to become a much messier process than it was in the past.[44] When mandatory retirement was the norm, both parties knew that the employer-employee relationship would end at a certain age. Employers also used traditional defined benefit pension plans to structure an orderly departure. No such structure exists in a 401(k) environment. Employers face the prospect of workers with declining productivity and inadequate 401(k) balances hanging onto their jobs much longer than, from the employer's standpoint, is desirable. In fact, employers in a recent survey employers indicated that they expect half of their older workers will lack the resources needed to retire at their traditional retirement age; that they expect half of these unprepared workers will want to remain on the job; and that the employers were lukewarm about retaining even half of those who will want to stay on (Munnell, Sass, and Aubry 2006). Employers will thus need new severance tools to manage an older workforce. Without such means, employers will avoid retaining or hiring older workers. The severance tool could be a "carrot," such as a generous retirement package, or a "stick," such as some form of mandatory retirement. Of course, the latter would be extremely controversial. But it is important to recognize that the shift away from employer-defined benefit plans means no mechanism exists to ease the bulk of the baby boom into retirement.

V. The Health of Older Workers

Intuitively, people's health affects their ability and desire to participate in the labor force. Poor health can make work seem very difficult and unpleasant, leading people to withdraw from the labor force. Poor health can also reduce people's productivity, leading to lower wages, and lower wages reduce the incentive to work. In the last 35 years, the impact of health on labor force activity has become a major area of research, and virtually all studies show that poor health has a negative effect on the likelihood of remaining in the labor force, and on the expected retirement age, as well as hours worked and wages received.[45] The question is the extent to which health concerns pose an obstacle to people's ability to remain in the labor force longer.

One starting point for exploring the health of older workers is to look at trends in life expectancy at age 50. Figure 3.22 shows life expectancy at age 50 for American males over the twentieth century. Interestingly, life expectancy at older ages rose very slowly at the beginning of the century and then accelerated sharply toward the end of the century. In fact, life expectancy at age 50 was not very different in 1970 than in 1900—23 years versus 21 years. After 1970, however, life expectancy for American men at age 50 took off, rising to 27 years in 2000, and is projected to increase to 30 years by 2030.

Although longer life spans generally imply improvements in health, keeping less healthy people alive could actually increase the percent of the population with disabilities. Thus, for a time, researchers referred to the "failure of success" resulting from improvements in healthcare (see Waidmann, Bound, and Schoenbaum 1995). Today, the notion of such an increase in frailty among the elderly—those aged 65 years and older—has been decisively rejected. In 2002, a technical working group examined disability trends for older Americans recorded across five major national surveys.[46] The group concluded that, when standardizing for the definition of disability, the time period, and the consistent inclusion or exclusion of the nursing home population, all five surveys showed consistent downward trends for two common disability measures—having difficulty with daily activities and requiring help with daily activities—beginning in the early to mid-1990s. The evidence remained mixed for a change in disability rates the 1980s and for the overall trend using a third measure of disability—the use of help or equipment with daily activities.

The fact that the health of older Americans has improved would lead one to conclude that the health of the older *working-age* population has also been getting better. But for a long time, such a conclusion was not

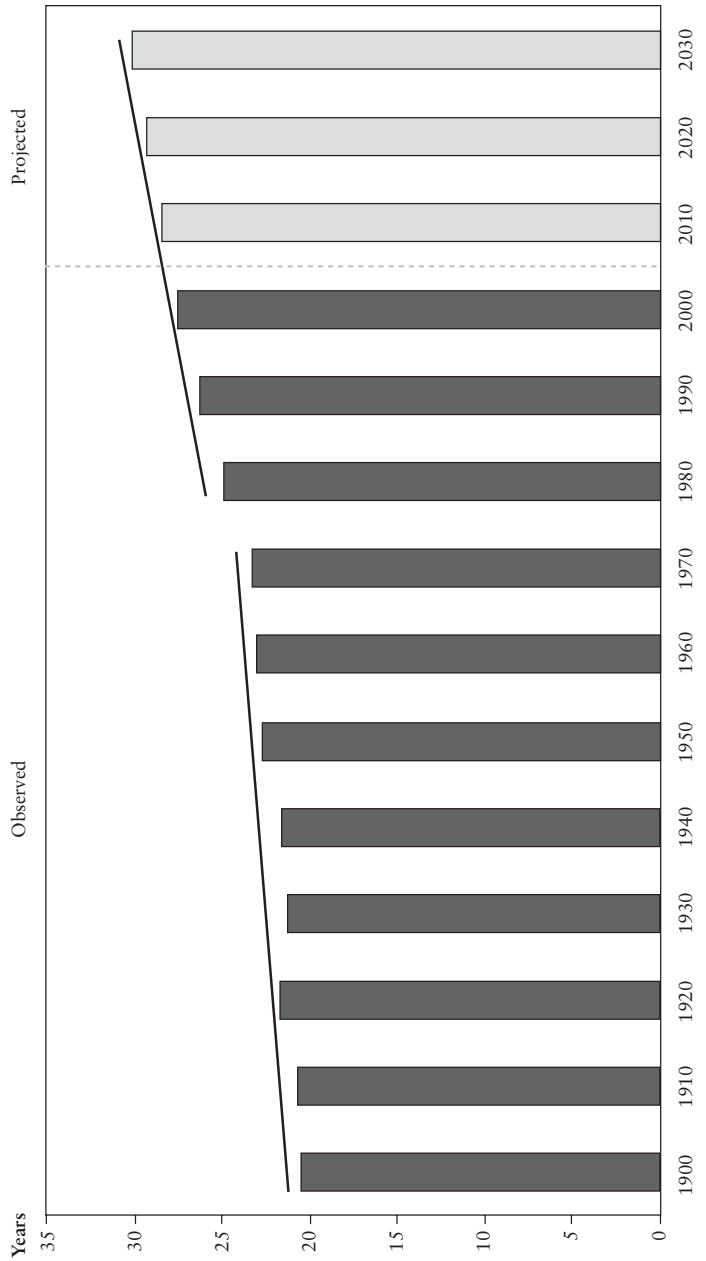

Figure 3.22
Years of Life Expectancy for American Men at Age 50, 1900–2030
Source: Munnell and Libby (2007) based on data from the U.S. Social Security Administration.

obvious. The major survey that tracked disabilities among the working-age population—the National Health Interview Survey (NHIS)—showed the percent of this population with disabilities increasing from the mid-1960s through the early 1980s; see Figure 3.23.[47] Decennial census data also showed an increase in the fraction of both working-age men and women unable to work during the 1970s. Skeptics of the increasing disability story contend that the trend during the 1970s may, at least in part, reflect social factors such as earlier detection and diagnosis of chronic diseases and greater availability of disability insurance.[48] Thus, the trend in the prevalence of disabilities during the 1970s remains controversial.

Since the early to mid-1980s, however, the health of the older U.S. working-age population has unquestionably improved. The percent of those aged 45–64 years with a disability declined through the mid-1990s, as shown in Figure 3.23. Between 1997 and 2004, responses to a similar survey question produced a more stable trend. But the general conclusion emerging from the NHIS data is one of declining disability among

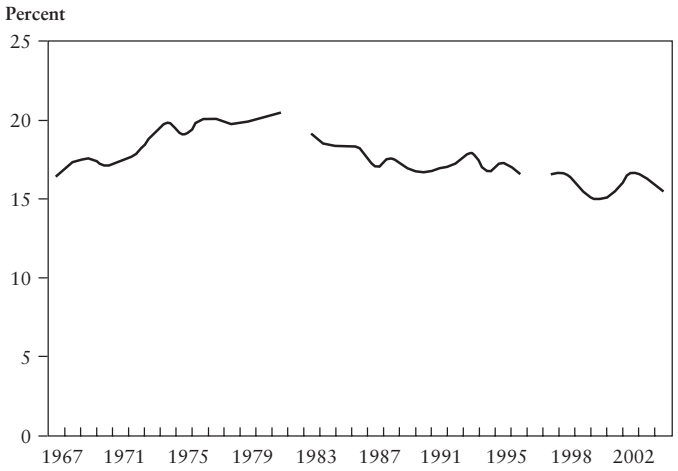

Figure 3.23
Percent of American Men Aged 45 to 64 Years with Disabilities, 1967–2004 (National Health Interview Survey)
Source: National Center for Health Statistics (1967–2004).
Note: From 2002 to 2004, the figure shows work limitations for all persons instead of males only.

older working-age individuals to a level that is at least comparable to that in the mid-1960s. Thus, the evidence suggests that the health of older workers is at least as good today as it was forty years ago. Moreover, as mentioned earlier, today's jobs are much less physically demanding. As a result, physical limitations should not inhibit the ability of the bulk of older Americans to work at least until their mid-60s.

The same data that support the possibility of continued work for the bulk of the older working population also it make clear that, despite a positive trend, 15 to 20 percent of people in their late 50s and 60s will find it virtually impossible to continue participating in the labor force. The data from the NHIS are consistent with responses from the Health and Retirement Study regarding the extent to which retirement was voluntary. As shown in Figure 3.24, 35 percent of those Americans who retired between 1992 and 2002 claimed that their retirement was involuntary, with 18 percent citing poor health as the reason for leaving the workforce. Moreover, many of those people who need to work longer—particularly low-wage workers dependent on Social Security for the bulk of their retirement income—are precisely the same individuals who have physically onerous jobs that stress their health, and who lack education, which has been shown to be important in managing their medical care. Thus, the "working longer prescription" must be administered with care, as some older Americans will simply be unable to adhere to this protocol.

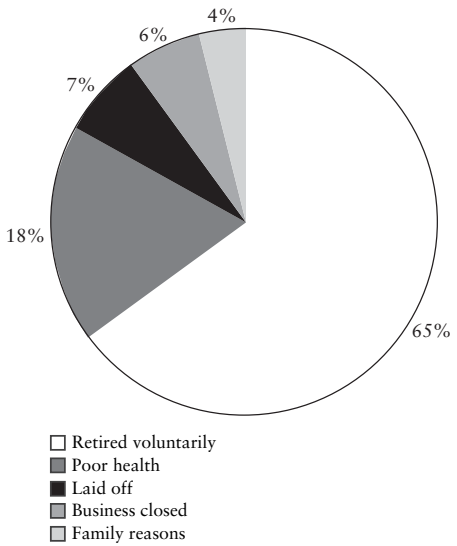

Figure 3.24
American Workers' Reasons for Retiring, 1992–2002
Source: Authors' calculations based on University of Michigan (1992–2002).

VI. Obstacles to the Labor Force Participation of Older American Workers

At least two major obstacles might hinder older workers from offering their services in the coming years and decades. The first is the availability of Social Security benefits at age 62. The second is the fact that employment seems to be an "all or nothing" full-time proposition, with relatively little room for gradually reducing hours or working part-time.

Social Security's Earliest Age of Eligibility

Social Security offers retirement benefits at 62 years of age. The early retirement benefits are actuarially reduced, and the reduction is designed to be "age-neutral." That is, two people with average life expectancy—one who claims benefits at 62, the other at 65—receive equal lifetime Social Security benefits.[49] Despite the actuarial reduction, the vast majority of American workers continue to claim Social Security benefits well before reaching age 65. In 2004, 59 percent of women and 54 percent of men claimed benefits at age 62; see Figure 3.25. The claiming of benefits coincides with the average retirement age, which is now 63 years for men and 62 years for women.[50]

Social Security's retirement age for full benefits is scheduled to increase from 65 to 67 years by 2022.[51] But under current law, the EEA remains unchanged at 62. Raising the full retirement age, however, will increase the actuarial reduction for claiming benefits at age 62 from 20 percent to 30 percent. But people's claiming behavior and retirement decisions appear more sensitive to the availability of benefits than to benefit amounts, so age 62 may well remain an important retirement benchmark for many Americans.[52]

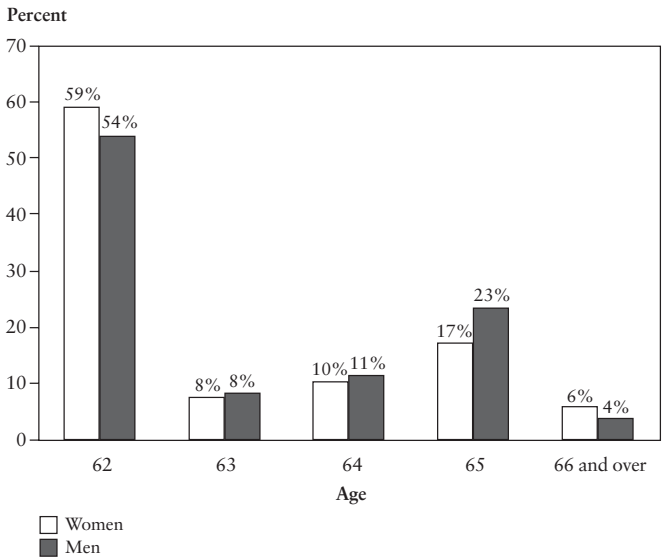

Figure 3.25
Age Distribution of Initial Receipt of Social Security Benefits, 2004
Source: U.S. Social Security Administration (2006) and authors' calculations.

Raising the EEA to 64 would likely encourage people to work longer by removing the opportunity to get benefits earlier. But this proposal is controversial. First, without instituting any other changes, raising the EEA has virtually no impact on the system's long-term finances. Any additional work effort brings in some additional payroll tax revenues, but the fact that the benefits were actuarially reduced means virtually no net savings.[53] Second, as discussed above, a significant fraction of Americans will be unable to work past age 62, either because they are in poor health, because their jobs are physically demanding, or because they have experienced job displacement later in life, and cannot find work at their age.[54] Therefore raising the EEA would inevitably involve some expansion of the disability program for older workers or some similar accommodation. Another problem is that a higher EEA would reduce lifetime Social Security wealth for those with lower-than-average life expectancies. Since African-Americans and low-wage workers have lower-than-average life expectancies, a higher EEA might be considered unfair to these groups. So raising the EEA might need to be part of a larger reform package which includes provisions that offset such losses to particular groups.

Raising the EEA, however, seems like an essential step to ensure that older adults continue to participate in the labor force. Moreover, raising this official age may not only increase the willingness of workers to supply their labor but may also enhance the willingness of employers to retain and hire older workers. A recent survey asked firms about the impact of various characteristics that affect their evaluation of older workers. A major negative factor was the perception that older workers will be on the job for only a short time. To the extent that the likely departure date can be pushed out, employers will be more willing to hire, train, and promote older workers.

Firms' Resistance to Part-Time Employment

Another hurdle to the continued employment of older workers is that they consistently report wanting to work part-time. For example, a study based on the Health and Retirement Study reports that 56 percent of respondents aged 55 to 65 years in 1996 said they would prefer to gradually reduce their hours as they age (U.S. General Accounting Office 2001). Consistent with this finding, older self-employed people tend to reduce their work hours as they approach retirement. But few older workers have part-time positions, and part-time employment does not appear to be increasing; see Figure 3.26.

Currently, part-time employment is concentrated in small business establishments and in firms in the service sector (Montgomery 1988). This remains true even after controlling for other factors that would affect labor demand, such as wages, fringe benefits, seasonal fluctuations in demand, and hiring costs. It is not exactly clear why this is the case. Large firms might avoid hiring part-time workers because such workers tend to have higher turnover rates than full-time employees (Tilly 1991). Part-time work might be more common in the service sector because it is labor intensive and is subject to large fluctuations in demand, and because employers find it is easier to manage these fluctuations by using part-time workers. While all these theories are plausible, these explanations have not been supported by rigorous empirical studies (Hutchens 2001). Without an increase in the availability of part-time employment,

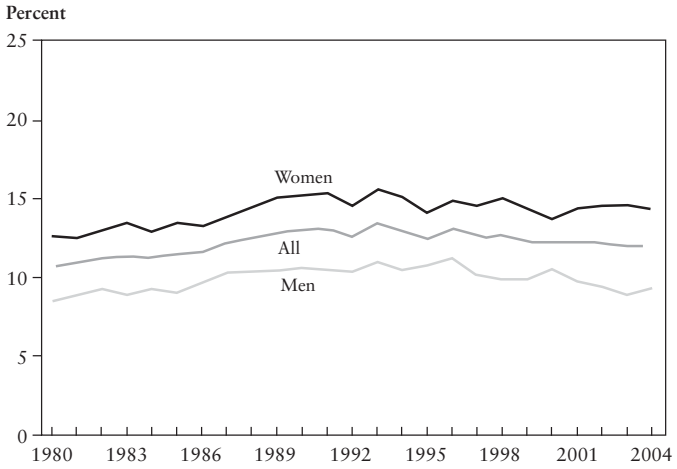

Figure 3.26
Percent of U.S. Workers Aged 55 to 70 Years Employed Part-Time, 1980–2004
Source: U.S. Bureau of Labor Statistics and U.S. Census Bureau (1981–2004).

however, many older adults may be unwilling to keep working. A recent study estimates that increased flexibility in work schedules would double the number of people entering partial retirement (Gustman and Steinmeier 2007).

In short, despite the need to build up their stock of retirement wealth, older workers may find the lack of part-time work opportunities and the availability of Social Security benefits at age 62 too tempting to pass up, and continue to retire early. And not all older people can remain in the workforce into their mid- to late-60s. Some have health problems or have been laid off and are unable to find another job, while others see continued employment as simply too onerous.

VI. Conclusion

Greater labor force participation by older U.S. workers would make an important contribution to national output, increase tax revenues, and dramatically improve retirement income security. Some indication that people might be willing to work longer comes from the fact that the century-long downward trend in the labor force participation of older men has clearly ceased, and their participation has actually been rising since the mid-1990s. The question is whether this upward trend in older men's workforce participation will continue.

Going forward, some key changes in the nation's retirement income system should encourage greater labor force participation by older workers. The share of pre-retirement earnings Social Security will replace at any claiming age is falling. Given rising longevity and the meager balances in the now dominant 401(k) accounts, the replacement income provided by employer plans, for retirement at any given age, is also likely to fall. The "income effect" of such reductions should increase labor force participation. In addition, the shift to 401(k)s and changes in the Social Security program that have essentially eliminated the subsidies for taking early retirement and penalties for taking later retirement should also raise participation rates. The "substitution effect" of these changes is to raise the cost of retirement relative to work to its actuarially appropriate level. Moreover, jobs in today's economy are less physically demanding, and today's older people are healthier than earlier cohorts.

Impediments still remain, however, to the continued employment of older workers. The most important obstacle is the availability of Social Security benefits at age 62. Even today, with the elimination of the earnings test after the Full Retirement Age and an actuarially fair Delayed Retirement Credit, the majority of U.S. workers continue to claim their benefits as soon as they are eligible to do so. Another important factor is the decline in career employment, with the majority of older workers now needing to negotiate the vagaries of the labor market if they are to work into their mid- to late-60s. Enduring extended and difficult job searches, as well as confronting the prospect of only earning low wages, may cause many older workers to simply give up and exit the workforce. Moreover, older people have a strong preference for part-time jobs and flexible work schedules, desires which, to date, many employers have been reluctant to accommodate. Finally, 15 percent to 20 percent of older people are probably not healthy enough to work beyond age 62.

What's the bottom line? Today, approximately 70 percent of American men aged 55–64 years are in the labor force, up from a low of 66 percent in the mid-1990s. Given the contraction of the retirement income system, labor force participation for this group is unlikely to start head-

ing back down. Will it continue to increase? In 1960, before men could claim Social Security benefits at age 62, before the enactment of Medicare and significant increases in Social Security replacement rates, and before employer pensions became a widespread source of retirement income, 87 percent of men aged 55–64 years were in the labor force. We are unlikely to see this high level again, given the increase in household wealth, some of which people want to spend on more leisure at the end of their work life, and the availability of Social Security benefits at age 62. Our best guess is that by 2030, without a significant change, such as an increase in Social Security's Earliest Eligibility Age, labor force participation rates for men 55–64 years may be 75 percent—up five percentage points from today's levels. This number is higher than the Bureau of Labor Statistics projection of 69 percent, as shown in Figure 3.27.

About 28 percent of American men aged 65–74 years are in the labor force today. Again, this percentage is unlikely to decline. Here, the 1960 level of almost 40 percent is a relevant benchmark for comparison. Again, some additional participation is likely to occur by 2030, but not to levels

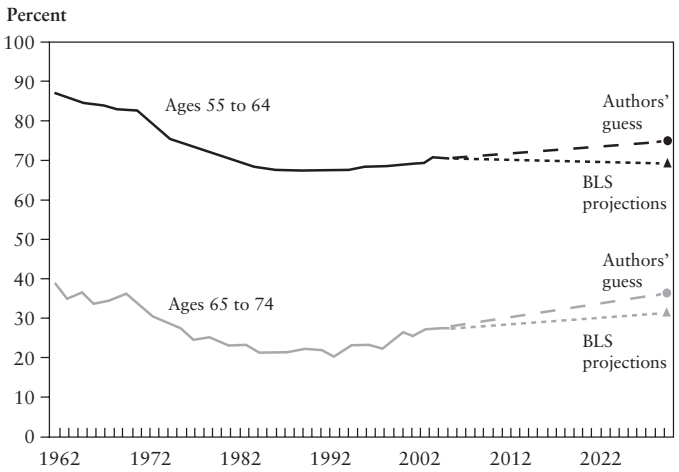

Figure 3.27
Older American Men's Actual Labor Force Participation Rates, 1962–2006, with Projection to 2030
Source: U.S. Bureau of Labor Statistics and U.S. Census Bureau (1980–2006) and U.S. Bureau of Labor Statistics (2007).

seen in the 1960s. On balance, the employment of older workers in the United States will likely rise, but fall well short of levels seen in 1960. Without changes that produce a further increase in participation rates, and raise the average retirement age to 66 or even 67 years, Americans are likely to see a significant drop in living standards in retirement.

■ *The authors would like to thank Jerilyn Libby for excellent research assistance on this project.*

1. Under legislation enacted in 1983, the increase in the Full Retirement Age began with those born in 1938, and turning 62 in 2000, and will be fully phased in for those born in 1960, and turning age 62 in 2022.

2. The premium for Medicare Part B is projected to increase from 9 percent of the average Social Security benefit in 2006 to 11 percent in 2030.

3. The pension coverage data discussed above apply only to individual workers at any given point in time. Over a lifetime and on a household, rather than an individual basis, coverage rates are somewhat higher. For households with two adults aged 55–64 years, the 2001 Survey of Consumer Finances shows that approximately 65 percent of households had some sort of pension coverage in 2001. Pension coverage is much more extensive for high-income households.

4. The annuity might provide a dollar amount per month for each year of service, say $50—so workers with 20 years of service would receive $1,000 per month at age 65. The benefit could also be a percentage of final salary for each year of service, say 1.5 percent; so workers with 20 years would receive 30 percent (20 years at 1.5 percent) of final salary for as long as they live. The employer finances these benefits by making pre-tax contributions into a pension fund; the company holds the assets in trust, directs the investments, and bears the risk. The Pension Benefit Guaranty Corporation (PBGC) insures pension benefits up to specified limits. The PBGC monthly guarantee limit in 2007 is $4,125 at age 65 years, and declines to $1,856 at age 55. Employers pay for this insurance with premiums largely determined by the plan's funding status.

5. For a 401(k), generally the employee, and often the employer, contributes a specified percentage of earnings into the account. These contributions are invested, usually at the direction of the employee, mostly in mutual funds consisting of stocks and bonds. Upon termination of employment or retirement, the worker generally receives the balance as a lump sum, albeit with the option to roll it over to an IRA.

6. This amount includes Individual Retirement Account (IRA) balances, because most of the money in IRAs is rolled over from 401(k) plans after an employer leaves a job. For further details, see Munnell and Sundén (2006).

7. In 401(k) plans, workers must decide whether or not to join the plan, how much to contribute, how to invest the assets, when to re-balance, what to do

about company stock, whether to roll over accumulations when changing jobs, and how to withdraw the money at retirement. The evidence indicates that a significant fraction of participants make serious mistakes at every step along the way. A quarter of those eligible to participate in 401(k) plans choose not to do so. Over half of those that do participate fail to diversify their investments. Many over-invest in company stock. Almost no participants re-balance their portfolios as they age or in response to market returns. Most importantly, many cash out these accounts when they change jobs, rather than rolling them over to another 401(k) or IRA, and very few annuitize these accounts at retirement to guarantee a lifetime income stream. The basic problem is that for most individuals, making their own financial decisions is difficult. Most participants lack sufficient financial experience, training, or time to figure out what to do, and the end result is often a significant shortfall in their retirement savings.

8. Penner and Johnson (2006) estimate that rising healthcare costs and the taxes required to cover these costs in retirement will require a moderate-income couple to work an additional 2.5 years, under the scenario assuming higher healthcare costs and higher tax burdens, to receive as much income in the first year of retirement—net of taxes and out-of-pocket health spending—as they would receive under the low-cost scenario of more moderate healthcare costs and future taxes.

9. In addition to addressing the financial issue, working longer appears to help individuals maintain their overall physical and mental well-being (see Calvo 2006).

10. Similarly, Butrica, Johnson, Smith, and Steuerle (2006) concluded that many people could increase their consumption by more than 25 percent at older ages simply by retiring at age 67 instead of age 62.

11. A 1570 census of the poor, in Norwich, England, thus found three widows, aged 74, 79 and 82 years, "almost past work" but still earning a small income from spinning. Estates left by the elderly in colonial America often included tools used in less strenuous trades, such as tailoring, spinning, shoemaking, and weaving. And well into the nineteenth century, about half of all 80-year-old men in America still worked (Thane 2000).

12. The Census measured the "gainful employment rate" until 1940 and then the labor force participation rate, defined as the percentage of the adult population working or actively looking for work.

13. See Graebner (1980).

14. Life expectancy at age 20 for men in 1900 was 44 years, compared to 59 years in 2000 (U.S. Social Security Administration). Also see Lee (2001) for the rapid rise in the expected length of retirement of workers entering the labor force between 1850 and 1990.

15. Favorable tax provisions had a limited effect on employer-provided pension coverage before World War II, as less than 10 percent of the adult population typically paid income tax. But the postwar growth of mass income taxation made pensions far less costly to employers and workers, and thus encouraged their spread.

16. Using evidence from the coal boom and bust, the collapse of the steel industry, and the general decline in manufacturing, Black and Liang (2005) conclude that the individual retirement decision is sensitive to prevailing economic conditions. This response most likely reflects elements of both labor supply and demand.

17. For example, suppose a person will live for 20 years and is entitled to a pension of $15,000 at age 65; lifetime benefits will equal $300,000 (20 × $15,000). To keep lifetime benefits actuarially constant, if that employee retired at 55, and was expected to live until age 85, his annual benefit should be only $10,000 per year (30 × $10,000 = $300,000). But traditional defined benefit plans typically provide far more because they use an actuarial reduction that is smaller than the full reduction. For instance, these plans might pay, say, $12,000 at age 55, which means that the worker in this example who retires at 55 would receive $360,000 (30 × $12,0000), substantially more in lifetime pension benefits than if he were to retire at 65. This exercise is actually somewhat more complicated because the employee adds to his pension if he continues to work, but the general example illustrates our main point.

18. Often, working beyond the plan's normal retirement age results in negative pension accruals. The law requires that the wage increases of those who work beyond the plan's normal retirement age be reflected in higher retirement benefits. But the law does not prevent firms from capping the years of service used to calculate benefits; nor does it require firms to provide actuarially fair adjustments for the fact that longer-working participants will receive benefits for fewer years (McGill et. al. 1996).

19. The change was made primarily to help younger widows and to allow wives, who were presumed to be two to three years younger than their husbands, to claim benefits at the same time as their husbands. Since it seemed unfair to require women workers to wait until a later age to receive benefits than mandated for non-working women, the EEA was introduced for all women. See Congressional Budget Office (1999).

20. In addition, Blau (1998) concludes that the availability of Social Security benefits is very important to the retirement decision, while changes in Social Security benefits over time have been considerably less important to this decision. On the other hand, Gruber (2000) found a sizable labor supply response to the level of disability benefits when comparing labor force participation in the Quebec system and in the rest of Canada, where disability benefits were increased.

21. For more details on recent trends, see Purcell (2005).

22. Under current law, individuals with less than $25,000 and married couples with less than $32,000 of "combined income" do not have to pay taxes on their Social Security benefits. (Combined income is adjusted gross income as reported on tax forms, plus nontaxable interest income, plus one-half of Social Security benefits.) Above those thresholds, recipients must pay taxes on either 50 or 85 percent of their benefits. Individuals must pay tax on 50 percent of their benefits if their "combined income" is between $25,000 and $34,000, and on 85 percent if above $34,000. A couple must pay tax on 50 percent of their benefits if their "combined income" is between $32,000 and $44,000, and on 85 percent if

above $44,000. (Committee on Ways and Means, U.S. House of Representatives 2000).

23. Benefits were cut a bit more than 1 percent per year until reaching a 6.7 percent cut for the cohorts turning 62 in between 2005 and 2017; the benefit cuts then resume and reach the full 13.4 percent reduction for cohorts turning 62 in 2022 and after. This full reduction will affect those cohorts born in 1960 and thereafter.

24. Prior to the introduction of early retirement, the earnings test was a tax, in that benefits lost in one year did not produce a benefit gain in later years.

25. Coile and Gruber (2000) note that in a context where workers make their retirement decisions based on the full future stream of Social Security benefits, raising the Delayed Retirement Credit could have a larger effect than raising the Full Retirement Age. Changing the Full Retirement Age has both an income effect that encourages work and a substitution effect that discourages work (via lower Social Security benefit accruals); but a change in the Delayed Retirement Credit has only a positive "substitution" effect that encourages work until age 65; after age 65, it has both an income effect (via the increase in Social Security wealth) that discourages continued work and a "substitution effect" that rewards work (via higher Social Security benefit accruals). Before age 65, their study shows that raising the Delayed Retirement Credit from 5 percent to 8 percent would increase labor force participation at age 65 by four percentage points.

26. Some researchers (see Eschtruth and Gemus 2002; Cahill, Giandrea, and Quinn 2006) suggest that those workers covered by defined contribution plans are sensitive to fluctuations in the stock market, and that the 2001 collapse of the stock market might explain why the labor force participation rate for older workers (aged 55–64 years) jumped 2 percentage points between early 2000 and 2002. This was an unprecedented increase that occurred during a recession, when labor force participation usually declines. This result would be consistent with studies by Gustman and Steinmeier (2002) and Coronado and Perozek (2003), which found that the unexpected positive shocks to wealth as a result of the stock market boom of the 1990s led to some additional retirement. Other researchers (Coile and Levine 2006) argue that few households had substantial stock holdings in this same period, and if workers were indeed so sensitive to stock market fluctuations, their labor force participation should have dropped as the market recovered, a decrease which did not happen.

27. In cash-balance plans, as in traditional defined benefit plans, the employer makes the contributions, owns the assets, selects the investments, and bears the risk. The Pension Benefit Guaranty Corporation also insures the benefits. To the employee, however, cash balance plans look very much like defined contribution plans. The employer typically contributes 4 or 5 percent of the worker's pay to a "notional" account, and provides an interest credit on the balances. Employees receive regular statements and generally withdraw the balance as a lump sum when they retire or terminate employment. Since these plans are not backloaded, employees suffer no loss in benefits as they move from job to job, and therefore these plans would not be expected to affect worker mobility. Bank of America created the first cash balance plan in 1985, and by 2003 these plans covered 22 percent of all U.S. employees and 26 percent of all assets in defined benefit plans (see Buessing and Soto 2006). Since 2003, extensive litigation has brought the expansion of cash balance plans to a virtual halt. However, the Pension Protection Act of 2006 clarified the legality of converting defined benefit plans to cash balance form, and this might prompt renewed interest among employers in converting defined benefit plans to cash balance plans.

28. Coronado and Copeland (2003) offer another perspective on the reasons for the shift to cash balance plans. They contend that these conversions occurred in competitive industries with tight labor markets, and were done largely to improve compensation for a more mobile workforce.

29. Massachusetts Office of the Governor (2001).

30. The share of men aged 55 to 60 years in a job that requires "lots of physical effort none or almost none of the time" increased from 31 percent to 38 percent between 1992 and 2002 (see Johnson 2004a).

31. Blau (1998), using the Retirement History Survey, found that among 30 to 40 percent of married couples, the spouses left the labor force within a year of each other. Hurd (1990), using the Social Security Administration's New Benefit Survey, estimated that among one-quarter of couples, the husband and wives retired within one year of each other. Johnson and Favreault (2001), looking at married couples in the 1998 wave of the Health and Retirement Study, calculated that between 22 and 40 percent of husbands and wives retired within two years of each other. These studies show that spouses tend to retire at the same time, generally because they want to spend time together. See also Johnson (2004b).

32. Indeed, a recent study (Lahey, Kim, and Newman, 2006) found that retirees who returned to work were no less financially prepared for retirement than were their counterparts who remained retired. Instead, the influential factors for returning to work were the availability of health insurance, whether or not the initial retirement decision was voluntary, and the degree of satisfaction with retirement. Maestas (2005), using the Health and Retirement Study, also concluded that financial pressures were not the reason for "un-retirement."

33. A recent study (Mermin, Johnson and Murphy 2006), using the Health and Retirement Study, reported a significant 4 percentage point increase between 1992 and 2004 in the expected probability among workers aged 51 to 56 years staying employed full-time past age 62, from 47 to 51 percent, and a similar increase in the probability of staying employed full-time past 65, from 27 to 33 percent. Controlling for other factors, self-employment, more education, and previous higher earnings increased work expectations, while defined benefit pension coverage, employer-sponsored retiree health benefits, and household wealth reduced expectations that older adults would remain in the workforce.

34. Costa (1998) cautioned researchers not to put too much emphasis on the recent uptick in labor force participation of older workers. As long as retirement

remains an attractive option and incomes continue to rise, people will want to use at least some of their increased wealth for retirement. The question is whether—even if income during people's working years continues to rise—the prospective decline in retirement income could provide the impetus for continued work.

35. The Current Population Survey (CPS) has asked respondents about job tenure since 1973. Specifically, CPS tenure supplements are available for 1973, 1978, 1981, 1983, 1987, 1991, 1996, 1998, 2000, 2002, and 2004. All data are from the Workplace Topics I (January/February) supplements, although the 1973 tenure data are from the displaced worker supplement. The job tenure question changes slightly over the period. In 1973, 1978, and 1981, the question refers to time spent working at the present job or business, while for 1983 and later the question refers to working "continuously" for the respondent's present employer. If respondents in the earlier surveys experienced temporary separations from their employer, their responses will make them look like they have more job tenure than they actually had. Since other researchers do not view this as a significant problem and make no adjustment, the raw median tenure data for employed males are presented in Figure 3.11.

36. See Allen, Clark, and McDermed (1988). Gustman and Steinmeier (1993) emphasize how small pension wealth is early in a worker's career, and argue that the main impact of defined benefit pensions would be to deter mobility for long-tenured workers.

37. Specifically, for each survey it is possible to identify those working full-time at age 55, 60, and so on who are still with the same employer they worked for at age 50. Mechanically, this exercise involves simply asking, say, the 55-year-old full-time worker, how long he has been with his current employer. If the response is five years or more, the worker is classified as working with his age-50 employer. The number working for their age-50 employer are then divided by total workers in these age groups to get the proportion in what used to be the typical pattern of employment for older workers.

38. The survey asks workers whether they lost their job for one of the following reasons: their plant or company closed down or moved; their company had insufficient work; their position or shift was eliminated; a seasonal job was completed; a self-operated business failed; other reason. These data do not include all job losses within the economy, because the survey collects and reports information on only one job loss for each individual and the distinction between layoffs and voluntary quits is not always clear. Nevertheless, this survey can be used to determine whether older workers are becoming more or less vulnerable to involuntary job displacement.

39. The analysis is limited to displacement because of plant closures, positions abolished, or slack demand for work. Using a more detailed set of 56 industry dummy variables instead of the set of private goods sector, private service sector, and public sector dummy variables had little effect on the coefficient estimates and standard errors for all other explanatory variables in the regressions.

40. As in earlier studies, women, married people, and those working full time have a low probability of being displaced, and race appears to have no impact on this probability. Private sector workers in goods-producing industries have a higher probability of being displaced than those in private sector service industries. In contrast, public sector employees have a much lower likelihood of being displaced than their private sector counterparts.

41. Over the 1996-2004 Displaced Worker Surveys, displacement rates averaged 15.9 percent for those workers with 0–1 years of tenure; 11.3 percent with those with 1–4 years; 5.5 percent for those with 5–9 years; and 4.0 percent for those with 10 or more years of tenure.

42. A study by Freeman (1979) reached similar conclusions.

43. Benitez-Silva (2002) explores the factors that lead older workers to engage in job search activities. The author finds that previous work attachment and health limitations are key factors in explaining the different job search behavior of both non-employed and employed individuals.

44. Reflecting this heterogeneity, a recent survey by Vanguard identified six different pathways to retirement. See Ameriks, Fergusson, Madamba, and Utkus (2007).

45. For a survey of the literature, see Currie and Madrian (1999); an update can be found in Deschryvere (2005).

46. See Freedman et al. (2004). The five surveys included the Health and Retirement Study (HRS), the Medicare Current Beneficiary Survey (MCBS), the National Health Interview Survey (NHIS), the National Long Term Care Survey (NLTCS), and the Supplements on Aging (SOAs).

47. The NHIS is an annual cross-sectional survey of 100,000 non-institutionalized civilians conducted by the National Center for Health Statistics. Unfortunately, the survey questions have been revised every 10 to 15 years, making it impossible to construct a series over a long period of time. Nevertheless, consistent data are available from 1967–1982, 1983–1996, and 1997–2004. For the period 1983–1996, the survey asked "Does any impairment or health problem now keep [the person] from working at a job or business? Is [the person] limited in the kind or amount of work [the person] can do because of any impairment?" A person who answers "yes" to either question is considered to have a disability that poses a work limitation.

48. See Waidmann, Bound, and Schoenbaum (1995).

49. More specifically, benefits are reduced by five-ninths of one percent for each month these are received prior to the Full Retirement Age (FRA), up to 36 months, and five-twelfths of one percent for each month thereafter. This is equivalent to a 6.67 percent reduction for the first three years prior to the FRA and 5 percent thereafter. With an FRA of age 65, a person who claims benefits at 62 years receives monthly benefits 20 percent lower than the full amount. The scheduled increase in the FRA from age 65 to 67 raises the actuarial reduction for claiming benefits at age 62 from 20 percent to 30 percent.

50. The average retirement age is defined as the age at which 50 percent of the birth-year cohort is out of the labor force.

51. The increase began with individuals born in 1938, for whom the FRA is 65 years plus two months, and increases two months per year until it reaches 66 years. Then, after a 12-year hiatus, the FRA again increases by two months per year until it reaches 67 years for individuals born in 1960 or later.

52. Studies showing that the availability of benefits has the major effect on retirement include Burtless and Moffitt (1984), Hurd (1990), and Gruber and Wise (1998). In a study of 12 countries, Gruber and Wise (2002) conclude that averaging across all countries, a reform that delayed the benefit eligibility by three years would likely reduce the proportion of men aged 56–65 years staying out of the labor force by 23 and would be closer to 36 percent in the long run.

53. However, an increase in the EEA could help set the stage for future increases in the full retirement age, one option for maintaining the solvency of the Social Security program. An EEA of 62 years makes any additional increase in the FRA highly unlikely, since a higher FRA would produce an even steeper reduction in benefits at age 62. A higher EEA, by signaling that retiring in one's early 60s is no longer economically feasible, could prepare the way to raise the FRA beyond age 67.

54. Similarly, a recent survey by Prudential Financial of a nationally representative sample of retired Americans found that 38 percent of them claimed that they had retired involuntarily.

References

Allen, Steven, Robert Clark, and Ann McDermed. 1988. "The Pension Cost of Changing Jobs." *Research on Aging* 10(4): 459–471.

Ameriks, John, Holly B. Fergusson, Anna B. Madamba, and Stephen P. Utkus. 2007. "Six Paths to Retirement." Volume 26. Vanguard Center for Retirement Research.

Benitez-Silva, Hugo. 2002. "Job Search Behavior at the End of the Life Cycle." Working Paper 2002-10. Chestnut Hill, MA: Center for Retirement Research at Boston College.

Black, Dan A., and Xiaoli Liang. 2005. "Local Labor Market Conditions and Retirement Behavior." Working Paper 2005-08. Chestnut Hill, MA: Center for Retirement Research at Boston College.

Blau, David M. 1998. "Labor Force Dynamics of Older Married Couples." *Journal of Labor Economics* 16(3): 595–629.

Buessing, Marric, and Mauricio Soto. 2006. "The State of Private Pensions: Current 5500 Data." Issue in Brief 42. Chestnut Hill, MA: Center for Retirement Research at Boston College.

Burtless, Gary, and Robert Moffitt. 1984. "The Effect of Social Security Benefits on Labor Supply of the Aged." In *Retirement and Economic Behavior*, ed. Henry Aaron and Gary Burtless, 135–174. Washington, DC: The Brookings Institution Press.

Burtless, Gary, and Joseph F. Quinn. 2000. "Retirement Trends and Policies to Encourage Work Among Older Americans." Annual Conference of the National Academy of Social Insurance. Washington, DC: National Academy for Social Insurance.

Burtless, Gary, and Joseph F. Quinn. 2002. "Is Working Longer the Answer for an Aging Workforce?" Issue in Brief #11. Chestnut Hill, MA: Center for Retirement Research at Boston College.

Butrica, Barbara A., Richard W. Johnson, Karen E. Smith, and Eugene Steuerle. 2006. "The Implicit Tax on Work at Older Ages." *National Tax Journal* 59(2): 211–234.

Cahill, Kevin E., Michael Giandrea, and Joseph Quinn. 2006. "A Micro-level Analysis of Recent Increases in Labor Force Participation among Older Men." BLS Working Paper 400. Washington, D.C.: Bureau of Labor Statistics.

Calvo, Esteban. 2006. "Does Working Longer Make People Healthier and Happier?" Work Opportunities Brief #2. Chestnut Hill, MA: Center for Retirement Research at Boston College.

Chan, Sewin, and Ann Huff Stevens. 2001. "Job Loss and Employment Patterns of Older Workers." *Journal of Labor Economics* 19(2): 484–521.

Coile, Courtney, and Jonathan Gruber. 2000. "Social Security and Retirement." Working Paper 7830. Cambridge, MA: National Bureau of Economic Research.

Coile, Courtney C., and Phillip B. Levine. 2006. "Bulls, Bears, and Retirement Behavior." *Industrial and Labor Relations Review* 59(3): 408–429.

Committee on Ways and Means, U.S. House of Representatives. 2000. 2000 Green Book. Washington, DC: U.S. Government Printing Office.

Congressional Budget Office. 1999. "Raising the Earliest Eligibility Age for Social Security Benefits." Washington, DC: U.S. Government Printing Office.

Congressional Budget Office. 2004. "Retirement Age and the Need for Saving." Economic and Budget Issue Brief. Washington, DC: Congressional Budget Office.

Coronado, Julia Lynn, and Phillip C. Copeland. 2003. "Cash Balance Pension Plan Conversions and the New Economy." Finance and Economics Discussion Series 2003-63. Washington, DC: Board of Governors of the Federal Reserve System.

Coronado, Julia Lynn, and Maria Perozek. 2003. "Wealth Effects and the Consumption of Leisure: Retirement Decisions during the Stock Market Boom of the 1990s." Finance and Economics Discussion Series 2003-20. Washington, DC: Board of Governors of the Federal Reserve System.

Costa, Dora L. 1998. *The Evolution of Retirement: An American Economic History, 1880–1990*. Chicago: University of Chicago Press.

Costa, Dora L. 2000. "The Wage and the Length of the Work Day: From the 1890s to 1991." *Journal of Labor Economics* 18(1): 156–181.

Currie, Janet, and Brigitte C. Madrian. 1999. "Health, Health Insurance, and the Labor Market." In *Handbook of Labor Economics*. Volume 3C, ed. Orley C. Ashenfelter and David Card. Amsterdam: Elsevier Science Publishers BV.

Deschryvere, Matthias. 2005. "Health and Retirement Decisions: An Update of the Literature." *ENEPRI Research Report* No. 6. Belgium: European Network of Economic Policy Research Institutes.

Eschtruth, Andrew, and Jonathan Gemus. 2002. "Are Older Workers Responding to the Bear Market?" Issue in Brief #5. Chestnut Hill, MA: Center for Retirement Research at Boston College.

Freedman, Vicki, Eileen Crimmins, Robert Schoeni, Brenda Spillman, Hakan Aykan, Ellen Kramarow, Kenneth Land, James Lubitz, Kenneth Manton, Linda Martin, Diane Shinberg, and Timothy Waidmann. 2004. "Resolving Inconsistencies in Trends in Old-Age Disability: Report from a Technical Working Group." *Demography* 41(3): 417–441.

Freeman, Richard B. 1979. "The Effect of Demographic Factors on Age-Earnings Profiles." *Journal of Human Resources* 14(3): 289–318.

Friedberg, Leora. 1998. "The Social Security Earnings Test and Labor Supply of Older Men." In *Tax Policy and the Economy*, Volume 12, ed. James M. Poterba, 121–150. Cambridge, MA: MIT Press.

Friedberg, Leora. 2000. "The Labor Supply Effects of the Social Security Earnings Test." *The Review of Economics and Statistics* 82(1): 46–63.

Friedberg, Leora. 2007. "The Recent Trend Towards Later Retirement." Work Opportunities Brief #9. Chestnut Hill, MA: Center for Retirement Research at Boston College.

Friedberg, Leora, and Anthony Webb. 2005. "Retirement and the Evolution of Pension Structure." *Journal of Human Resources* 40(2): 281–308.

Friedberg, Leora, and Anthony Webb. 2006. "Persistence in Labor Supply and the Response to the Social Security Earnings Test." Working Paper 2006-27. Chestnut Hill, MA: Center for Retirement Research at Boston College.

Gottschalk, Peter and Robert Moffitt. 1999. "Changes in Job Instability and Insecurity Using Monthly Survey Data." *Journal of Labor Economics* 17(4): S91–S126.

Graebner, William. 1980. *A History of Retirement: The Meaning and Function of an American Institution, 1885-1978*. New Haven, CT: Yale University Press.

Gruber, Jonathan. 2000. "Disability Insurance Benefits and Labor Supply." *The Journal of Political Economy* 108(6): 1162–1183.

Gruber, Jonathan, and Peter Orszag. 2003. "Does the Social Security Earnings Test Affect Labor Supply and Benefits Receipt?" *National Tax Journal* 56(4): 755–773.

Gruber, Jonathan, and David Wise. 1998. "Social Security and Retirement: An International Comparison." *American Economic Review* 88(2): 158–163.

Gruber, Jonathan, and David A. Wise. 2002. "Social Security Programs and Retirement Around the World: Micro Estimation." Working Paper 9407. Cambridge, MA: National Bureau of Economic Research.

Gustman, Alan L., and Thomas Steinmeier. 1986. "A Structural Retirement Model." *Econometrica* 54(3): 555–584.

Gustman, Alan L., and Thomas L. Steinmeier. 1993. "Pension Portability and Labor Mobility: Evidence from the Survey of Income and Program Participation." *Journal of Public Economics* 50(3): 299–323.

Gustman, Alan L., and Thomas Steinmeier. 2002. "Retirement and the Stock Market Bubble." Working Paper 9404. Cambridge, MA: National Bureau of Economic Research.

Gustman, Alan L., and Thomas Steinmeier. 2007. "Projecting Behavioral Responses to the Next Generation of Retirement Policies." Working Paper 12958. Cambridge, MA: National Bureau of Economic Research.

Haider, Steven, and David Loughran. 2001. "Elderly Labor Supply: Work or Play?" Working Paper Series 01-09. Santa Monica, CA: RAND Corp.

Hurd, Michael D. 1990. "The Joint Retirement Decision of Husbands and Wives." In *Issues in the Economics of Aging*, David A. Wise, ed., 231–254. Chicago: The University of Chicago Press.

Hutchens, Robert M. 2001. "Employer Surveys, Employer Policies, and Future Demand for Older Workers." Roundtable on the Demand for Older Workers Sponsored By the Retirement Research Consortium, The Brookings Institution. Available online at http://www.bc.edu/centers/crr/special_pubs/hutchens_ld.pdf

Johnson, Richard W. 2004a. "Trends in Job Demands among Older Workers, 1992–2002." *Monthly Labor Review* 127(7): 48–56.

Johnson, Richard W. 2004b. "Do Spouses Coordinate Their Retirement Decisions?" Issue Brief #19. Chestnut Hill, MA: Center for Retirement Research at Boston College.

Johnson, Richard W., and Janette Kawachi. 2007. "Job Changes at Older Ages: Effects on Wages, Benefits, and Other Job Attributes." Working Paper 2007-4. Chestnut Hill, MA: Center for Retirement Research at Boston College.

Johnson, Richard W., and Melissa M. Favreault. 2001. "Retiring Together or Working Alone: The Impact of Spousal Employment and Disability on Retirement Decisions." Working Paper No. 2001-01. Chestnut Hill, MA: Center for Retirement Research at Boston College.

Lahey, Joanna. 2006. "Age, Women, and Hiring: An Experimental Study." Working Paper 2006-23. Chestnut Hill, MA: Center for Retirement Research at Boston College.

Lahey, Karen E., Doseong Kim, and Melinda L. Newman. 2006. "Full Retirement? An Examination of Factors That Influence the Decision to Return to Work." *Financial Services Review* 15: 1–19.

Lee, Chulhee. 2001. "The Expected Length of Male Retirement in the United States, 1850–1990." *Journal of Population Economics* 14(4): 641–650.

Maestas, Nicole. 2005. "Back to Work: Expectations and Realizations of Work after Retirement." RAND Working Paper WR-191-1. Santa Monica, CA: RAND Corporation. Available online at http://www.rand.org/pubs/working_papers/WR196-1/

Margo, Robert A. 1993. "The Labor Force Participation of Older Americans in 1900: Further Results." *Explorations in Economic History* 30(3): 409–423.

Massachusetts Office of the Governor. 2001. Massachusetts toward a New Prosperity: Building Regional Competitiveness.

McGill, Dan M., Kyle N. Brown, John J. Haley, and Sylvester J. Schieber. 1996. *Fundamentals of Private Plans*. Pension Research Council and the University of Pennsylvania Press.

Mermin, Gordon B.T., Richard W. Johnson, and Dan Murphy. 2006. "Why Do Boomers Plan to Work So Long?" The Retirement Project: Discussion Paper 06-04. Washington, DC: The Urban Institute.

Moen, Jon R. 1987. *Essays on the Labor Force and Labor Force Participation Rates: The United States from 1860 through 1950*. Unpublished Ph.D. dissertation. The University of Chicago.

Montgomery, Mark. 1988. "On the Determinants of Employer Demand for Part-Time Workers." *The Review of Economics and Statistics* 70(1): 112–117.

Munnell, Alicia H. 1977. *The Future of Social Security*. Washington, DC: Brookings Institution.

Munnell, Alicia H. 2003. "The Declining Role of Social Security." Just the Facts #6. Chestnut Hill, MA: Center for Retirement Research at Boston College.

Munnell, Alicia H., Francesca Golub-Sass, and Andrew Varani. 2005. "How Much are Workers Saving?" *Issue Brief* No. 34. Chestnut Hill, MA: Center for Retirement Research at Boston College.

Munnell, Alicia H., Kevin E. Cahill, and Natalia Jivan. 2003. "How Has the Shift to 401(k)s Affected the Retirement Age?" Issue in Brief #13. Chestnut Hill, MA: Center for Retirement Research at Boston College.

Munnell, Alicia H., and Jerilyn Libby. 2007. "Will People Be Healthy Enough to Work Longer?" Issue Brief #7-3. Chestnut Hill, MA: Center for Retirement Research at Boston College.

Munnell, Alicia H., Steven A. Sass, and Jean-Pierre Aubry. 2006. "Employer Survey: 1 of 4 Boomers Won't Retire Because They Can't." Work Opportunities Brief #6. Chestnut Hill, MA: Center for Retirement Research at Boston College.

Munnell, Alicia H., Steven A. Sass, Mauricio Soto, and Natalia Zhivan. 2006. "Do Older Workers Face Greater Risk of Displacement?" Working Paper 2006-17. Chestnut Hill, MA: Center for Retirement Research at Boston College.

Munnell, Alicia H., and Annika Sundén. 2006. "*401(k) Plans Are Still Coming Up Short*." Issue Brief No. 44. Chestnut Hill, MA: Center for Retirement Research at Boston College.

National Center for Health Statistics. 1967–2004. *Current Estimates from the National Health Interview Survey*. Washington, DC: U.S. Centers for Disease Control and Prevention.

Neumark, David. 2000. "Changes in Job Stability and Job Security: A Collective Effect to Untangle, Reconcile, and Interpret the Evidence." In *On the Job: Is Long-Term Employment a Thing of the Past*, ed. David Neumark, 70–110. New York: Russell Sage Foundation.

Penner, Rudolph G., and Richard W. Johnson. 2006. "Health Care Costs, Taxes, and the Retirement Decision: Conceptual Issues and Illustrative Simulations." Working Paper 2006-20. Chestnut Hill, MA: Center for Retirement Research at Boston College.

Pingle, Jonathan F. 2006. "Social Security's Delayed Retirement Credit and the Labor Supply of Older Men." Finance and Economics Discussion Series 2006-37. Washington, DC: Federal Reserve Board.

Purcell, Patrick J. 2005. "Older Workers: Employment and Retirement Trends." Congressional Research Service Report.

Ruggles, Steven, and Matthew Sobek. 2004. *Integrated Public Use Microdata Series: Version 3.0*. Minneapolis, MN: Minnesota Population Center.

Rust, John and Christopher Phelan. 1997. "How Social Security and Medicare Affect Retirement Behavior in a World of Incomplete Markets." *Econometrica* 65(4): 781–831.

Sass, Steven A. 1997. *The Promise of Private Pensions: The First 100 Years*. Cambridge, MA: Harvard University Press.

Schieber, Sylvester J. 2007. "Pension Aspirations and Realizations: A Perspective on Yesterday, Today and Tomorrow." Policy paper. Watson Wyatt Worldwide. Available online at http://www.watsonwyatt.com.

Thane, Pat. 2000. *Old Age in English History: Past Experiences, Present Issues*. Oxford: Oxford University Press.

Tilly, Chris. 1991. "Reasons for the Continuing Growth of Part-Time Employment." *Monthly Labor Review* 114(3): 10–18.

Triest, Robert K., Margarita Sapozhnikov, and Steven A. Sass. 2006. "Population Aging and the Structure of Wages." Working Paper 2006-5. Chestnut Hill, MA: Center for Retirement Research at Boston College.

University of Michigan. 1992–2002. *Health and Retirement Study*. Ann Arbor, MI: Institute for Social Research.

U.S. Board of Governors of the Federal Reserve System. 1983–2004. *Survey of Consumer Finances*. Washington, DC.

U.S. Bureau of Economic Analysis. 1973. *Long Term Economic Growth, 1860–1970*. Washington, DC: U.S. Government Printing Office.

U.S. Bureau of Labor Statistics. 2007. *Projected Labor Force Data*. Available online at http://www.bls.gov/emp/emplab1.htm.

U.S. Bureau of Labor Statistics and U.S. Bureau of the Census. 1962–2006. *Current Population Survey*.

U.S. Bureau of Labor Statistics and U.S. Bureau of the Census. 2004. *Displaced Workers Survey – Supplement to the Current Population Survey*. Washington, DC: Government Printing Office.

U.S. General Accounting Office. 2001. "Older Workers: Demographic Trends Pose Challenges for Employers and Workers." GAO 02-85. Washington, DC.

U.S. Social Security Administration. *Cohort Life Tables for Males*, 1850–2000. Unpublished data.

U.S. Social Security Administration. 2006. *Annual Statistical Supplement, 2005*. Washington, DC: Government Printing Office.

Waidmann, Timothy A., John Bound, and Michael Schoenbaum. 1995. "The Illusion of Failure: Trends in the Self-Reported Health of the U.S. Elderly." *Milbank Quarterly* 73(2): 253–287.

Welch, Finis. 1979. "Effects of Cohort Size on Earnings: The Baby Boom Babies' Financial Bust." *Journal of Political Economy* 87: S65–S98.

Comments on "The Labor Supply of Older American Men" by Alicia H. Munnell and Steven A. Sass

Robert Hutchens

The paper by Alicia Munnell and Steven Sass seeks to first examine trends in labor force activity of older men in the United States, and then inquire into the economic forces that shape those trends. Of particular interest is whether the recent increase in the labor force activity of older men is likely to continue into the future.

There are at least two important reasons why we need to understand these trends. First, these trends are likely to influence government tax and transfer policies toward older Americans over the next several years. To illustrate the stakes, consider two scenarios.

> **Scenario A:** Despite reductions in real Social Security benefits, older Americans can easily find jobs and replace any lost Social Security income with wage and salary income. Reduced Social Security benefits cause no increased hardship beyond foregone leisure time.
>
> **Scenario B:** Many older Americans are unable to find jobs or are physically incapable of working. A reduction in real Social Security benefits casts large numbers of older Americans into poverty.

If scenario A is closer to the truth than scenario B, then arguably it would be easier to solve the Social Security financing problem by reducing Social Security retirement benefits—for example, by eliminating the age 62 early retirement benefit. As such, it is important that we understand how easily older Americans can increase their labor supply.

A second reason why the issues addressed in this paper are important concerns the quantity and quality of labor that will be supplied to the U.S. economy in the future. Over the next two or three decades, the baby boom generation will withdraw from the labor force, and they will be replaced

by comparatively small cohorts of young workers. At the same time, we are likely to see continued demand for skilled workers. This is at least part of the reason for the increased real wages at the upper deciles of the earnings distribution over the last three decades. An interesting question is whether older workers will help shore up the supply of skilled workers.

Munnell and Sass begin by noting that for American males over age 55 years, there has been a century-long decline in labor force activity, shown as Figure 3.8 in the paper. In a sense that decline is surprising. Over the twentieth century both life expectancy and health levels have arguably improved; one would think that with longer and healthier lives we would spend more time in the labor market. Of course, the explanation is that the wealth of the nation has continued to grow. That trend, combined with the availability of retirement income in the form of Social Security benefits and employer-provided pension plans, has made it possible for older people to withdraw from the labor force. But now comes the real surprise: despite our growing wealth, since the mid-1980s, the trend toward decreasing labor force activity has turned around and started going the other way. This is illustrated in Figures 3.8, 3.9, and 3.10 in Munnell and Sass's paper.

The paper does a nice job of discussing several explanations for this turnaround. It gives, however, special prominence to changes in employer-provided pensions (the well-documented steady decline in defined benefit plans and the rise of defined contribution pensions), and in Social Security, particularly, the partial elimination of the earnings test and the increase in the delayed retirement credit.

The paper also discusses several other labor market phenomena that occurred at roughly the same time as the turnaround in labor force participation; in particular:

1. A decline in median job tenure at older ages (see their Figures 3.12 and 3.13).
2. An increase in job displacements (roughly speaking, job losses other than voluntary quits) at older ages (see Figures 3.15 and 3.16).
3. A hypothesized decline in the compensation of older workers
4. Improved health and greater longevity of the older population.

The bottom line is that Munnell and Sass expect labor force participation of older men to continue to increase (Figure 3.27). They forecast this in large part because of the predicted contractions in the retirement income system. They quite cogently argue that under current law, Social Security replacement rates will fall over the next few decades, while employer-provided pensions are unlikely to offer a sufficient alternative source of income. Thus, Munnell and Sass predict that more older people will be working in order to augment a meager retirement income from pensions and Social Security.

I like this paper for several reasons. First, it takes a long-run perspective. This view is very useful in that it helps make clear that the turnaround in older male labor force participation in the 1980s is not particularly important in terms of its magnitude. Rather it is important because we have not seen anything like this in more than a century. This finding helps focus attention on the forces underlying the turnaround.

Second, I like the authors' agnostic view of what caused the turnaround. They give us several possible explanations, and choose their preferred explanation, changes in Social Security and pension plans. But they also make clear that the literature is not at the point where we can confidently say what actually caused the turnaround in the labor force participation of older American men.

Third, I like the way the paper brings together several different trends: in particular, employment trends, trends in job tenure, trends in displacement, and trends in health. There is an interesting breadth to the paper.

Finally, I like all the footnotes. We get 54 footnotes in 27 pages, and these are well worth spending time on. There are lots of interesting points about data and what—at least to me—are rather obscure nooks and crannies in the literature.

Let me raise two questions that struck me in reading the paper.

First, what types of older men have been increasing their labor force participation, and what types are particularly likely to do so in the future?

I wish the paper said more about the skill level of the groups that have played a major role in the post-1985 reversal in the labor force participation trends. For example, what does Figure 3.9 look like when drawn for college graduates? What about men who never went beyond a high school education? Do we see the same turnaround for different skill groups? It would be especially nice to do this analysis by deciles in the earnings distribution. But an analysis of labor force participation trends by educational category should give us the basic facts.

In particular, it would be interesting if the turnaround was occurring among men with at least a four-year college degree. Given the rapid rise in educational attainment of older cohorts, that would suggest a rapidly increasing relative demand for well-educated older workers. Perhaps such workers are presently helping to address a vital labor supply need in the U.S. economy?

The point is arguably more important regarding how these trends may play out in the future. Are the predicted increases in Figure 3.27 different by skill level? If the driving force behind the predicted increase in labor force participation is a contracting retirement system, meaning reductions in Social Security benefits and employer-provided pensions, then does that imply particularly large increases in labor force participation at lower skill and wage levels? The answer will be important because it is related to the future quality of the labor force. Looking at trends in the U.S. earnings distribution, I see little reason to think that the United States needs more low-skilled workers. Moreover, this goes back to the issue of hardship in retirement. If low-skilled older people increase their labor supply, how much can they earn to augment their income and retirement savings? Are we perhaps talking about older people who will join the ranks of the working poor?

My second question is, how will employers respond to increased labor force participation of older people?

It is not hard to believe that if the supply of older workers expands, then most will find jobs. Demand curves slope downward, wages adjust, and in the long run most older people who seek work will be able to find it. To my mind, the puzzle concerns the kind of jobs those older people will end up doing. In particular, there are two ways that this expansion in supply can occur. First, older people can delay taking retirement while remaining in a long-term career job. Second, they can leave those long-term jobs and take new jobs—presumably short-term "spot market" jobs. As a group of people from a Federal Reserve Bank can well understand, some expansions are better than others. From a social point of view, I would argue that the first type of expansion is better. If your long-term job was the highest and best use of your labor when you were 50, then that will probably still be the case at age 60 or 70. A person is probably more productive by continuing in that job than in taking a new job. Moreover, by staying in the long-term job, specific human capital is preserved as are professional and personal friendship and social support networks. Finally, by staying in the long-term job, the person is staying out of the market for new jobs; a growing literature indicates that older workers can have real difficulty finding new jobs that match their skill sets.

Of course, it will not always be the case that staying in the long-term job is desirable or even possible. My point is only that to the extent that both employers and employees benefit from workers continuing in the long-term career job, then from a social point of view that is a good thing.

There are many ways that an older worker could delay retirement while remaining in a long-term career job. Suppose the person prefers to cut back on hours. In that case, it is not necessary to change employers. Rather, a phased retirement could be arranged whereby the worker stays with the current employer but shifts to shorter hours. Another example is something like a job bank – whereby retirees are called back during periods of peak demand. Frito-Lay has done this with delivery truck drivers. All indications are, however, that such arrangements are quite rare.[1]

Is the predicted expansion in employment of older workers likely to take the form of delayed retirement while staying with the long-term career job? This paper provides good reason for doubt that the trend will play out in this manner. Looking at the older male labor market after 1980, what we learn is that (a) the fraction of men who remain in the jobs they held at age 50 decreased much more rapidly today than in the 1980s; see Figure 3.13. Moreover, (b) that change is partly because of a higher probability of displacement today than in the 1980s. Finally, (c) it is likely that this increase in job mobility is associated with reduced compensation for older workers. These reasons do not describe a labor market where people are delaying retirement by staying longer in a well-established career job.

I think this issue could be nailed down a bit more. As indicated in their Figure 3.13, one can compute the number of workers who are aged 50 years and over and who have remained at the same firm/employer since age 50. The authors do this for 1983 and 2004, but it could be done for the several years in which the Current Population Survey collected data on job tenure. Thus, one could create a version of Figure 3.9 that plots through time the fraction of men aged 55–64 years and aged 65 years and over who are both in the labor force and who have not changed employers since age 50.

Now, suppose that this new version of Figure 3.9 indicates that since 1990 there has been an increase in the fraction of men who are both in the labor force and have not changed employers since 1990. To my mind, that would be good news. It would mean that the post-1990 expansion in the older male labor force is occurring in a way that preserves specific human capital and social networks. This would indicate that people are tending to remain in the job that was the highest and best use of their labor when they were 50-years-old.

Suppose, on the other hand, that there has been a post-1990 decline in the fraction of older men who are both in the labor force and in the job they held at age 50. Given the evidence in this paper, I suspect that this scenario is what is taking place. I would view confirmation of this development as not particularly good news. While it is good for the economy to have older people working more, the expansion may well be occurring in a way that tends to not preserve specific human capital. Older workers are arguably moving into jobs that do not make full-use of their skills.

To conclude, I think Munnell and Sass have written a very interesting paper. It provides a useful perspective on several important trends in the labor market for older workers. It also continues a line of work that opens up lines of future research on the labor market for older workers. Well done.

Note

1. The Frito-Lay case is in Watson Wyatt Worldwide (2001). See Hutchens (2007) for a discussion of phased retirement.

References

Hutchens, Robert M. 2007. "Phased Retirement: Problems and Prospect." Chestnut Hill, MA: Center for Retirement Research at Boston College. Available online at http://www.bc.edu/centers/crr/wob_8.shtml.

Watson Wyatt Worldwide. 2001. *Current Practices in Phased Retirement: Transforming the End of Work*. Lake Oswego, Oregon: Watson Wyatt Worldwide.

Comments on "The Labor Supply of Older American Men" by Alicia H. Munnell and Steven A. Sass

Joyce Manchester

The paper by Alicia Munnell and Steven Sass makes the case that it is sensible for Americans to work longer in the future in order to enhance their retirement income security. Working longer also would be good for national output and tax revenues. The paper then asks the following question: based on what we know now, does it seem likely that American men will choose to work longer? Three current impediments to encouraging such a trend are identified as: 1) the early eligibility age for Social Security benefits that remains at age 62, 2) the increased mobility of workers that exposes them to the vagaries of the labor market, and 3) the significant fraction of people who are not healthy enough to work beyond age 62.

Of course, raising the early eligibility age for Social Security benefits seems like an obvious fix to some observers. But eliminating old-age benefits at 62 could lead to severe hardship for those not able to work in their early 60s, those without other sources of income, or those unlikely to live much longer. But instead of dwelling on these objections, my comments will focus on two other considerations of the question regarding whether Americans will indeed choose to work longer. The first consideration concerns the decision of when to claim old-age Social Security benefits. Specifically, what do we know about workers who choose to claim old-age Social Security benefits earlier compared with those who claim benefits later? A related issue is how recent policy changes have influenced the age at which people claim benefits, and whether those trends can help us predict future behavior. A second consideration regarding the issue of working longer involves the labor force behavior of women. At the same time that men are working less at older ages than they did in the

1960s, women's labor force participation has increased significantly. Can we say something about how their greater work experience and earnings will contribute to household retirement incomes in the future?

The Decision to Claim Old-Age Social Security Benefits

A recent analysis of data from the Social Security Administration (SSA) on the age at which people claim Old-Age and Survivors Insurance (OASI) benefits, using a 1 percent sample of administrative records, shows that:

- The most popular age to claim benefits is 62. For example, among all OASI beneficiaries who claimed benefits at ages 60 through 72 years in 2003, approximately 47 percent became entitled at age 62. The percentage is higher among women, 49 percent, than among men, 44 percent of whom claimed Social Security benefits at age 62.
- Approximately 26 percent of people who claimed benefits in 2003 were 65-years-old; at 31 percent, the percentage was higher among men than the 21 percent for women.

What do we know about people who claim old-age benefits early (starting at age 62 and prior to age 65) compared with those who claim benefits at age 65 or later? Preliminary analysis of a 1 percent sample of SSA data shows that work and earnings activities around the time of claiming vary by the claimant's age. More specifically, people who claim benefits at age 65 or later have stronger and steadier labor force activity than those who claim at age 64 or earlier. For example, the work participation rate up to 10 years prior to benefit claim and three years following benefit claim appears to be lower among those who claim benefits at age 62 than among those who claim benefits at 65; see Figure 3.28. In addition, the average annual earnings of people who claim benefits between ages 62 and 64 are lower than the average earnings of those who claim at age 65 or later. The difference can be seen as much as 15 years prior to benefit claiming as well as three years following benefit claiming; please refer to Figure 3.29. Given the different work and earnings experience of early claimants relative to late claimants, one could speculate that more emphasis on job characteristics such as flexible schedules, part-time work, and help transitioning to jobs that are not physically demanding

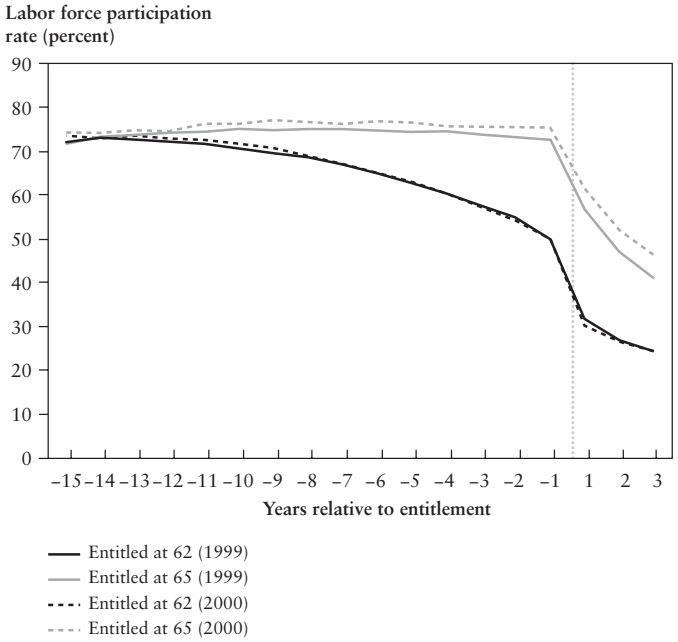

Figure 3.28
Labor Force Participation by Years Relative to Benefit Claim, By Claiming Age of Individuals Who Became Eligible for Entitlement in 1999 and 2000
Source: Tabulations by the Office of Policy, U.S. Social Security Administration, using the 1 percent sample of Social Security Administration administrative files.

could be important in encouraging older workers to remain in the labor force.

Why do so many people claim old-age benefits at age 62? Multiple reasons explain the phenomenon, but two factors that have gained attention recently are the advice commonly found in newspapers and other media, and the effective tax rate on work at older ages. Financial advisors sometimes adopt a "one-size-fits-all" attitude that for many individuals may not be the correct advice. For example, here is the punch line from one recent newspaper column: "When you turn 62, take the money and run." Unless an individual expects to die sooner than the average person in his/her cohort, that advice is usually wrong. Taking benefits early gen-

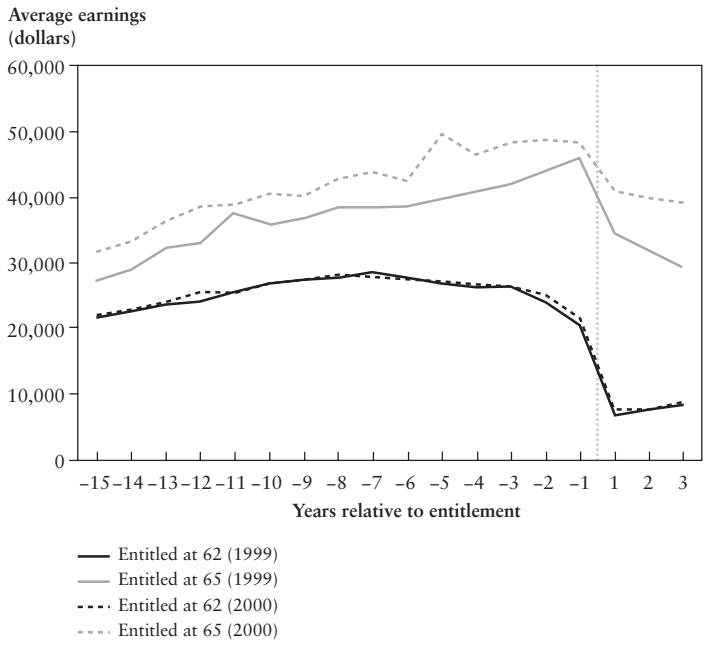

Figure 3.29
Mean Annual Earnings by Years Relative to Benefit Claim, by Claiming Age of Individuals Who Became Eligible for Entitlement in 1999 and 2000
Source: Tabulations by the Office of Policy, U.S. Social Security Administration, using the 1 percent sample of Social Security Administration administrative files.

erally means that person, and any dependents who receive Social Security benefits based on his or her earnings record, will receive lower benefits throughout their lifetimes. The reduction in benefits gets larger as the full retirement age (FRA) rises. Until 2003, the FRA was at 65 years, and a person who claimed benefits at age 62 received 80 percent of his or her full benefits. But when the FRA reaches 67 years in 2022 for individuals born in 1960 or later, the person who claims at age 62 will receive just 70 percent of the full benefit amount. Reductions also affect the widow benefit to some extent.

A second factor that may explain why so many people claim benefits at age 62 is the effective tax rate on work at older ages. Recent research shows that, for many people, work at older ages does not enhance their

Social Security benefit. A study done in the Office of Policy at the SSA using MINT, a microsimulation model developed by the SSA and the Urban Institute, examines the marginal internal rate of return (IRR) on Social Security payroll taxes from working one extra year at the end of one's work life (Reznik, Weaver, and Biggs 2007). The results show that 30 percent of men aged 62–65 years in 2005 would have faced a pure tax from working one extra year, meaning that their additional Social Security taxes paid would not lead to additional benefits. Only 21 percent of men gain from spending an extra year in the labor force at the end of their working lives.[1] Women have somewhat lower marginal internal rates of return than men because many women receive auxiliary benefits based on their spouses' earnings, or have relatively flat earnings histories. But women are somewhat less likely than men to show up as "pure tax" cases because they are more likely to have spotty work histories with years in which they have no earnings. It is still true, however, that the Social Security payroll tax represents a pure tax for 23 percent of women, meaning that 23 percent of women would see no increase in benefits if they worked an additional year. Policies that increase the payoff to work at older ages could encourage more Americans, particularly men, to stay in the workforce longer.

Do we have any evidence that raising the FRA affects the proportion of people who claim benefits early? The answer is yes. Again using a 1 percent sample of SSA administrative data, a couple of recent papers examine changes in the age at which people claim Social Security retirement benefits in response to two recent rule changes: the removal, in 2000, of the retirement earnings test at ages 65 to 69, and the gradual increase in the full retirement age (FRA), also beginning in 2000 for people who turned 62 in that year (Song and Manchester 2007, 2008). Figure 3.30 shows the cumulative distribution of benefit entitlement ages for males and females born in 1930, 1937, and 1940 who claim benefits between 62 and 65 years and 6 months.[2] The distribution of entitlement ages of the 1930 birth cohort should be relatively unaffected by both rule changes. For the 1937 birth cohort, benefit claiming at ages 63 and above took place after the elimination of the earnings test in 2000. Consequently, some members of the 1937 birth cohort clearly delayed claiming benefits between ages 62 and 65 relative to the 1930 birth cohort.

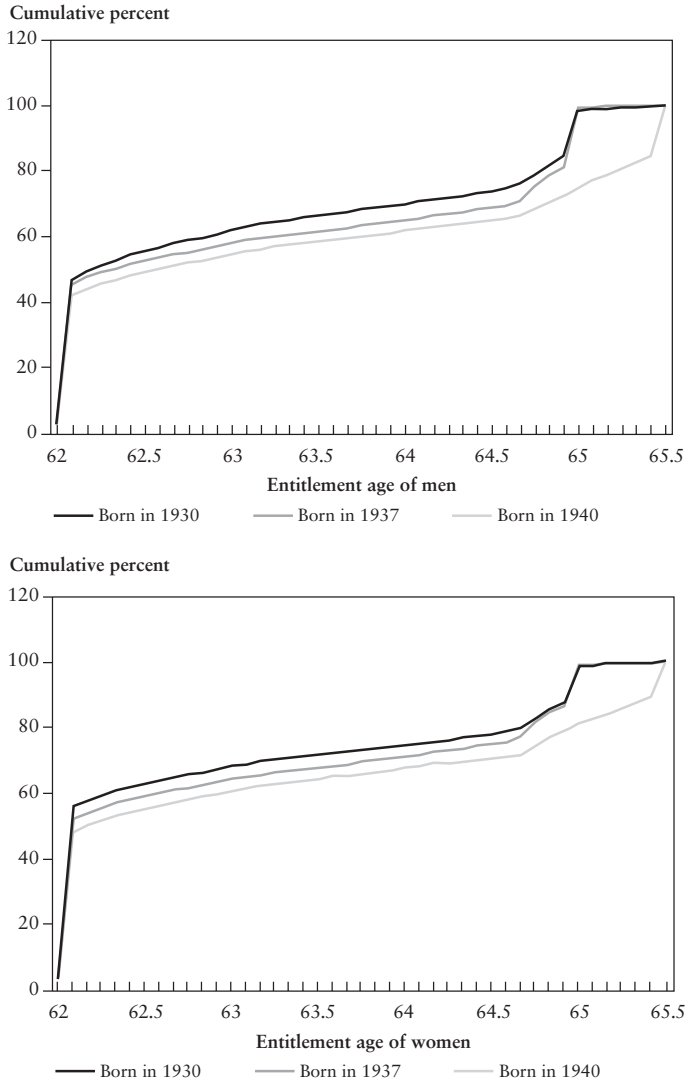

Figure 3.30
Cumulative Distribution of Benefit Entitlement Ages of Men and Women in Selected Birth-Year Cohorts Who Claimed Benefits Between the Ages of 62 Years and 65.5 Years
Source: Song and Manchester (2007).
Note: The cumulative percentages are measured among those who become entitled by age 65 and 6 months.

Further, the elimination of the earnings test in 2000 appears to accelerate benefit claims at age 65, as seen in the vertical distance at 65 years between the curves of the 1930 and 1937 birth cohorts. The 1940 cohort delayed claiming somewhat between 62 and 65 years, but the most noticeable change appears between 65 years and 65 years and 6 months. For the 1940 cohort, full Social Security benefits were not available until age 65 and 6 months, and Figure 3.30 shows that effect clearly. Women show similar, but less pronounced, responses.

It is a difficult task to sort out the economic effects of the benefit reductions from the signaling or institutional role of the FRA, but Figure 3.31 may offer some clues. The plots show the proportion of men and women who became entitled to retirement benefits, using two-month intervals between ages 62 and 65 years and 6 months among the cohorts born in 1937 through 1940. Among the 1937 birth cohort, which was not affected by the increase in the full retirement age to 65, about 42 percent of men and 49 percent of women claimed benefits at age 62, the earliest retirement age for this cohort. These percentages dropped slightly following the increase in the FRA; in the 1940 cohort, about 40 percent of men and 45 percent of women claimed benefits at age 62, the earliest eligibility age. The percentage of people who claim benefits after age 62 and up to a few months prior to the FRA stays relatively stable at about 1 percent at each two-month age increment. Benefit reductions alone affect people who retire prior to age 65, so the drop in the percentage who claim prior to age 65 largely reflects that benefit reduction.

More dramatic changes are evident at ages 65 years and above. About 18 percent of men in the 1937 cohort and 12 percent of women claimed benefits at age 65, the FRA for that cohort. As the FRA moved out by two months per year for the 1938, 1939, and 1940 cohorts, the spike at the FRA moved out as well. In the 1940 cohort, about 16 percent of men and 10 percent of women claimed benefits at 65 years and 6 months, the relevant FRA for that cohort. People who previously would have claimed benefits at age 65 but waited until their new, higher FRA are likely responding to a combination of the benefit reduction and the signaling aspect of the Social Security retirement age. It is also possible that the "full" retirement age in integrated private pension plans influences the age when individuals claim Social Security benefits.

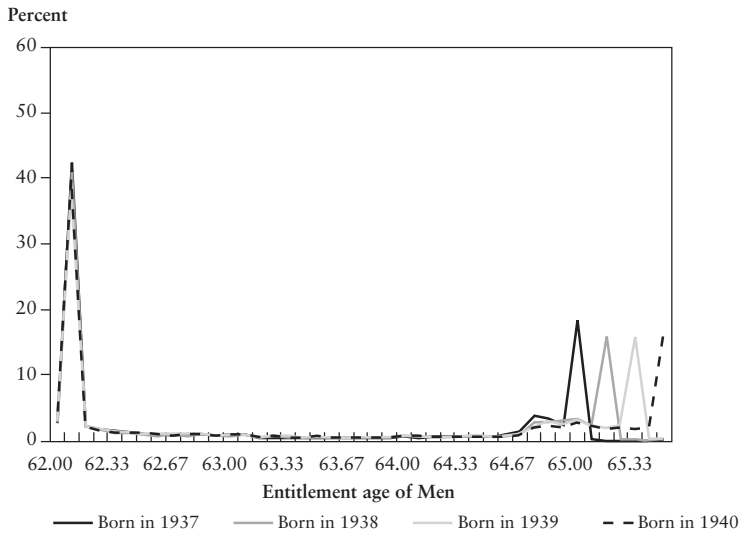

Figure 3.31
Entitlement Age Frequency Distribution by Birth-Year Cohort
Source: Song and Manchester (2007).
Note: The cumulative percentages are measured among those who become entitled by age 65 and 6 months.

Song and Manchester also conducted regression analysis of benefit entitlement status at specific ages using year- and age-specific treatment dummies over the period 2000–2005. Specifically, results show the marginal effects on the probability of entitlement claims given the elimination of the retirement earnings test for ages 65 to 69 and the increase in the FRA, separately for men and women. Both the direction and the magnitude of the estimated effects accord with our expectations. For those aged 65 years in 2000–2002, the change in the earnings test rule increases men's benefit claims by slightly more than 3 percentage points, and women's claims by slightly more than 2 percentage points. The FRA becomes the dominant rule change in 2004 and 2005, however, as the estimated marginal effect for those aged 65 turns negative; this marginal effect rises to 12.5 percentage points for men and 5.4 percentage points for women in 2005. Effects on early claimants are evident as well. Following the six-month increase in the FRA, benefit entitlement rates for men decline by 3.3, 4.4 and 5.2 percentage points at ages 62, 63, and 64 years, respectively and 2.1, 3.3, and 3.5 percentage points for women at these same three respective ages. These estimates suggest that a relatively large response occurs at age 62, and relatively small but incremental responses at ages 63 and 64. Recognizing those responses is important for policymakers who question whether people younger than the full retirement age would change their behavior; these results argue that changing the full retirement age does delay claiming by people who have not yet reached the FRA.

Elasticity estimates show the percentage change in claiming rates at a given age for a 1 percent reduction in benefit amount at that age. At 64 years, derived elasticities range from 1.3 to 1.7 for men and from 0.7 to 1.1 for women. Elasticities at age 62 are 0.8 to 1.3 for men and 0.7 to 1.2 for women. Elasticities near 1 indicate that benefit reductions do cause people to work longer. Looking forward, further benefit reductions as the FRA rises to 67 years are expected to result in longer work lives.

How Will Women's Work and Earnings Contribute to Retirement Income in the Future?

The Munnell-Sass paper points out that future retirees will face lower Social Security benefits at any given age, confront the lower likelihood of

defined benefit pensions, and contend with the vagaries of a more mobile job market. Compared to the 1960s, things look grim. At the same time, the authors see little evidence that men will increase their labor force participation much at older ages, thus compounding the pessimistic outlook. Yet there is a large and growing component of the labor force that is neglected, albeit intentionally, in the Munnell-Sass paper: women. At the same time that men's labor force participation at ages 55–64 was dropping from 87 percent in 1960 to 69 percent in 2005, women's participation rate surged from 37 percent in 1960 to 57 percent in 2005; see Figure 3.32. For men ages 65 years and above, participation dropped from 33 percent in 1960 to 20 percent in 2005. At the same time, labor force participation for women ages 65 years and above stayed approximately constant at 11 percent. Will women's increased work and earnings help to maintain American standards of living in retirement?

Some evidence on the possible contributions of women to household retirement income comes from examining married couples who are near-

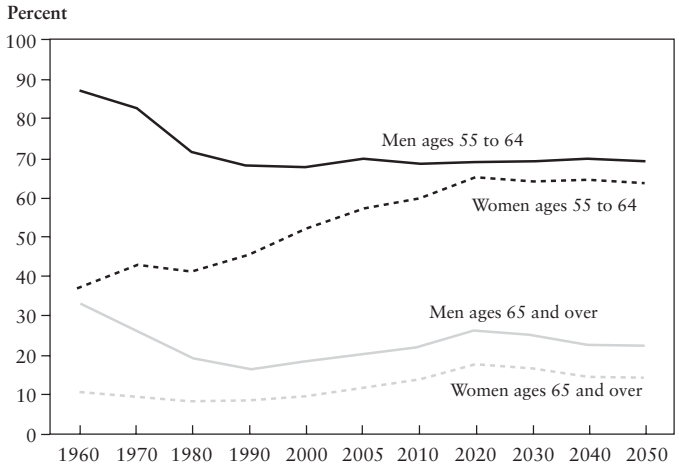

Figure 3.32
Labor Force Participation of American Men and Women, Aged 55 to 64 Years, and 65 Years and Older, 1960–2050, with Projection to 2050
Source: Author's calculations of U.S. Bureau of Labor Statistics data in Fullerton (1999) and Toosi (2006).

ing retirement. A paper by Maestas (2001) based on data from the Health and Retirement Study reports that 75 percent of men nearing retirement age in the 1990s were married, and 40 percent of them had working wives. The MINT model also shows that labor force attachment among married women is on the rise. According to the MINT model, married women in the retiree population of the late 1990s, cohorts which predate the baby boom generation born between 1946 and 1964, averaged 18 years of work experience. This contrasts rather starkly with 29 years in the labor force predicted for the early baby boomer women born in the 1946–1954 period, and 30 years for the late boomer women born between 1955 and 1964. Labor force participation is also projected to increase for non-married women.

More work experience for women translates into a greater likelihood of pension coverage for women; health insurance while they are in the labor force, and perhaps retiree health insurance as well; and receipt of their own Social Security benefits from both Old-Age Insurance and Disability Insurance. Tabulations of the Health and Retirement Study show that the percentage of women aged 55–64 years who have a pension rose from 52 percent in 1994 to 63 percent in 2004 (Iams et al. 2007). Social Security data show that the percentage of women who receive old-age benefits based on their own earnings record rather than from spousal benefits is rising as well.[3] Over the past two decades, the percentage of women with earned worker benefits increased. Current beneficiary women aged 62 to 64 years with retired worker benefits increased from 48 percent in 1984 to 56 percent in 2004, while those with disabled worker benefits doubled from 8 percent to 16 percent.

To conclude, the paper by Munnell and Sass raises a number of important issues regarding the labor supply of older American workers in the twenty-first century. But there is reason to be optimistic about their prospects, both for continuing in the labor force and for their retirement income security, once we examine the evidence to date on how older workers have responded to policy changes. Following the increase in the full retirement age, we have solid evidence that people are claiming Social Security benefits at later ages. We have identified other factors that could be changed to encourage older workers' continued labor force participa-

tion, such as job characteristics, the retirement advice and information offered, or the incentives in the tax-benefit structure of Social Security. And the other half of the potential labor force, women, have increased their contributions to household retirement incomes in recent decades, and a reversal in that trend is unlikely.

■ *These comments were written while Manchester was with the Office of Policy, Social Security Administration; she is now at the Congressional Budget Office. Any findings or opinions expressed here are those of the author and not necessarily those of the Social Security Administration or the Congressional Budget Office.*

Notes

1. In this exercise, individuals "gain" when the marginal internal rate of return is greater than 3 percent in real terms.

2. We include entitlement ages ranging from 62 to 65 years and 6 months because at the time we did the data work, the most recent data came from June 2006. The last birth cohort considered here (1940) reaches 65 years and 6 months, the FRA for that cohort, by the end of June 2006.

3. A general rule of thumb is that married women receive retired-worker-only benefits when their average lifetime earnings are more than 30 percent of their husband's earnings.

References

Fullerton, Howard N. Jr. 1999. "Labor Force Participation: 75 Years of Change, 1950–98 and 1998–2025," *Monthly Labor Review* 122(12): 3–12.

Iams, Howard M., John Phillips, Lionel Deang, and Irena Dushi. 2007. "Selected Indicators of Well-Being for People Aged 55–64: 1984, 1994, and 2004," Prepared for the Population Association of America meetings, New York City, April.

Maestas, Nicole. 2001. "Labor, Love and Leisure: Complementarity and the Timing of Retirement by Working Couples," manuscript, University of California, Berkeley.

Reznik, Gayle L., David A. Weaver, and Andrew G. Biggs. 2007. "Social Security and Marginal Returns to Work Near Retirement." Manuscript, Office of Policy, Social Security Administration.

Song, Jae G., and Joyce Manchester. 2007. "New Evidence on Earnings and Benefit Claims Following Changes in the Retirement Earnings Test in 2000," *Journal of Public Economics* 91(3–4): 669–700.

Song, Jae G., and Joyce Manchester. 2008. "Have People Delayed Claiming Retirement Benefits? Responses to Changes in Social Security Rules," *Social Security Bulletin* 67(2).

Toossi, Mitra. 2006. "A New Look at Long-Term Labor Force Projections to 2050," *Monthly Labor Review* 129(11): 19–39.

4

How Structural Shifts in Labor Demand Affect Labor Supply Prospects

Structural Demand Shifts and Potential Labor Supply Responses in the New Century

David H. Autor

I. Introduction

It is widely recognized that inequality of labor market earnings in the United States grew dramatically in recent decades. This trend may be seen in Figure 4.1, which plots the growth of real hourly wages of U.S. workers (both male and female) by earnings percentile for the years 1973 through 2005. Over the course of more than three decades, wage growth was weak to nonexistent at the bottom of the distribution, strong at the top of the distribution, and modest at the middle. While real hourly earnings of workers in the bottom 30 percent of the earnings distribution rose by no more than 10 percentage points, earnings of workers at the 90th percentile rose by more than 40 percentage points.

What is much less widely known, however, is that this smooth, monotone growth of wage inequality is a feature of a specific time period—and that this time period has passed.[1] Figure 4.2 shows that, consistent with common perceptions, the growth of wage inequality between 1973 and 1989 *was* strikingly linear in wage percentiles, with sharp falls in real wages at the bottom of the distribution and modest increases at the top.[2] Yet, starting in the late 1980s, the growth of wages "polarized," with strong, ongoing wage growth in the top of the earnings distribution, meaning at or above the 70th percentile, and modest growth in the lower tail of the distribution, defined as at or below the 30th percentile. Notably, the portion of the wage distribution that saw the least real earnings growth between 1989 and 2005 was the "middle" group, roughly the earners between the 30th and 70th percentiles of the distribution.[3] Thus, 1973 to 1989 and 1989 to 2005 represent two distinct periods of rising

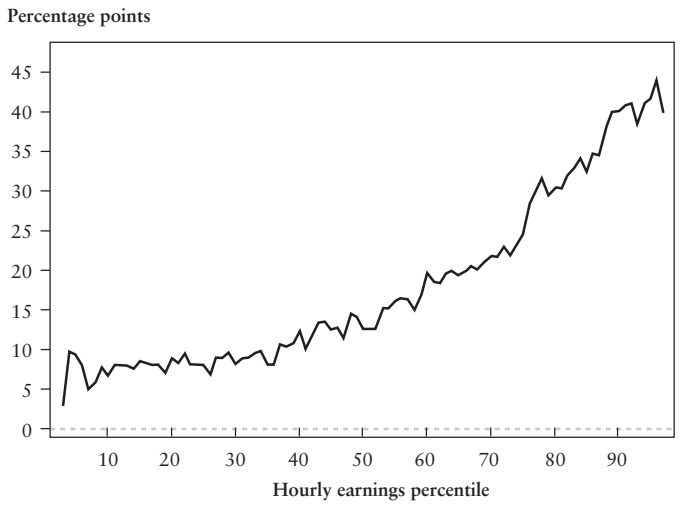

Figure 4.1
Changes in Real Log Hourly Earnings of All U.S. Workers from 1973 to 2005 by Percentile of the Hourly Earnings Distribution
Source: Current Population Survey and U.S. Census Bureau.

inequality: the first one is characterized by diverging wages throughout the distribution, and the second displays polarizing wage growth.

These two epochs are contrasted in Figure 4.3, which plots the evolution of the ratio of 90th to 50th percentile hourly earnings alongside the evolution of the ratio of 50th to 10th percentile hourly earnings.[4] The 90/50 ratio rises smoothly and secularly from 1979 to 2004. By contrast, the 50/10 ratio rises sharply from 1979 to 1987, plateaus in 1988, and then reverses course for the remainder of the time period. The divergent growth of upper- and lower-tail wage inequality in the 1980s and 1990s is also corroborated by microeconomic data on wages and total compensation from the establishment-based Employment Cost Index (see Pierce 2001). The steady growth of upper tier earnings inequality is seen in rising shares of wages paid to the top 10 and top 1 percent of U.S. earners since the late 1970s, as revealed in tax data (see Piketty and Saez 2003).

This paper evaluates the sources of the growth and then the polarization of earnings inequality in the United States, and considers these implications for the future growth of labor demand, by which we mean

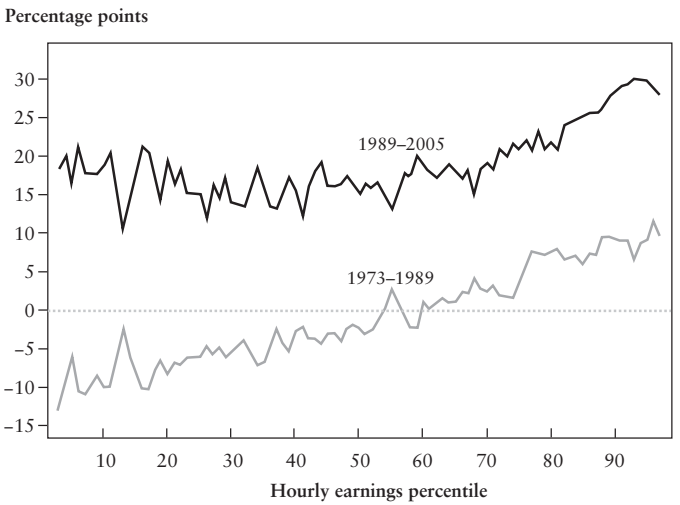

Figure 4.2
Changes in Real Log Hourly Earnings of Men and Women from 1973 to 1989 and from 1989 to 2005 by Percentile of the Hourly Earnings Distributions
Source: Current Population Survey and U.S. Census Bureau.

the demand for workers at various skill levels.[5] We begin by reviewing basic trends in earnings levels by education groups over several decades, and show how the pattern of polarization visible in Figure 4.2 is also reflected in trends in earnings by education level. We next consider whether these patterns of changing earnings by educational level can be adequately explained by canonical labor demand models of the type used by Katz and Murphy (1992); Autor, Katz, and Krueger (1998); Card and Lemieux (2001); and Acemoglu (2002), among many others. Though these models do an excellent job of explaining the evolution of U.S. income inequality to 1992, their explanatory power fares poorly thereafter, which suggests a substantial change, or structural break, in the character of labor demand over the last 15 years.

We briefly consider whether the widely discussed institutional explanations for rising U.S. wage inequality—most particularly, fluctuations in the U.S. minimum wage and the tight labor market of the 1990s—provide a sufficient alternative explanation for these same patterns. While these

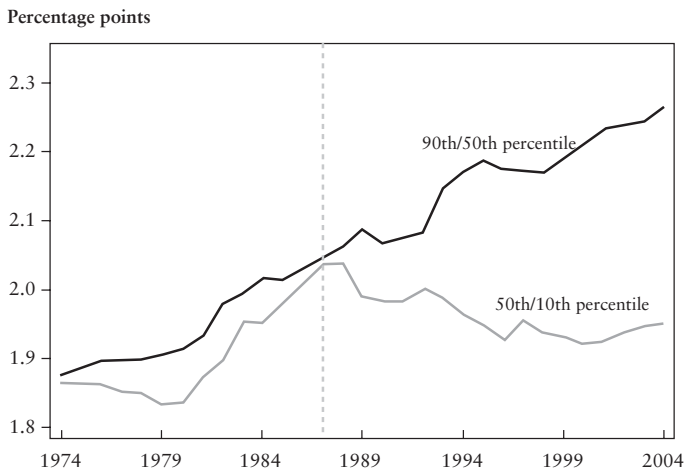

Figure 4.3
Ratios of 90th to 50th and 50th to 10th Percentile Real Hourly Earnings, 1974–2004 (Three-Year Moving Averages)
Source: Current Population Survey and U.S. Bureau of Economic Analysis.

two specific factors are likely to have contributed to rising inequality, particularly in the 1980s, neither one provides a viable explanation for the long-term secular growth of high incomes seen in the 1970s through 1990s, nor for the plateau and slight rebound of low incomes observed during the 1990s.

We next discuss how technological change and, more recently, international outsourcing, may provide a plausible, albeit still preliminary, explanation for the polarization of earnings growth. Following the conceptual model offered by Autor, Levy, and Murnane (2003), we argue that technological change (recently abetted by outsourcing) has been complementary to high-education occupations, particularly deleterious to middle-education occupations, and neither strongly complementary to nor strongly deleterious to (meaning substitutable for) low-education service occupations. A key implication of this conceptual framework is that computerization may foster a demand-driven polarization of labor market activities. Corroborating this implication, we present initial evidence that the observed polarization of earnings inequality *is* demand-driven. Drawing on this model, we speculate on the changing shape of labor demand in the United States, which in the future we argue will be characterized by rapid growth of managerial and professional occupations *and* rapid growth of comparatively low-education service employment.

The final section of the paper focuses on three sets of research and policy issues that impinge on how the changing shape of labor demand will affect employment opportunities and earnings inequality in the United States. The first set of issues considers potential supply responses in the form of human capital investment and immigration policy. A second set considers the role of labor standards and social welfare policy in shaping the quality of future jobs, particularly service jobs. The third and final set of issues considers areas of theory and measurement needing urgent attention for improving our understanding of how changes in technology and trade will affect U.S. labor demand in the ensuing decades.

II. Measuring Earnings Inequality

To summarize the basic changes in the U.S. wage structure over the last four decades, we draw on two large and representative household data sources: the March Current Population Survey (CPS) and the combined CPS May and Outgoing Rotation Group samples. The March CPS data provide reasonably comparable data on prior year's annual earnings, weeks worked, and hours worked per week for four decades. We use the March files from 1964 to 2006, which cover earnings from 1963 to 2005, to form a sample of real weekly earnings for workers ages 16 to 64 years who participate in the labor force on a full-time, full-year (FTFY) basis, defined as working 35-plus hours per week and 40-plus weeks per year. We complement the March FTFY series data with data on hourly wages of all current labor force participants using May CPS samples for 1973 through 1978 and CPS Outgoing Rotation Group samples for 1979 through 2003 (CPS May/ORG). From these sources, we construct hourly wage data for all wage and salary workers employed during the CPS sample survey reference week.[6]

We focus on two measures of relative earnings. The first is inequality in the upper and lower halves of the wage distribution, summarized by 90-

50 and 50-10 log wage gaps, which we refer to as upper- and lower-tail inequality. These trends are depicted above in Figure 4.3. The second is "between-group" inequality, which we measure using the earnings levels and earnings differentials among workers of different educational attainments.[7] Figure 4.4 displays these earnings trends for full-time, full-year workers by educational attainment for the years 1963 to 2005.[8] In this figure, the average earnings for each educational attainment level in 1963 are normalized to zero, and subsequent data points represent the logarithmic change in earnings (approximately equal to the percentage point change) since 1963. Wage levels are indexed using the Personal Consumption Expenditure deflator, and are composition-adjusted to hold constant the gender and labor market experience of workers within each educational group at their average levels over 1963 to 2005.

Figure 4.4 reveals the four major episodes in the evolution of between-group inequality in the United States. From 1963 to 1973, real wages grew strongly for all educational groups. Since growth rates were relatively comparable across educational levels (with the exception of work-

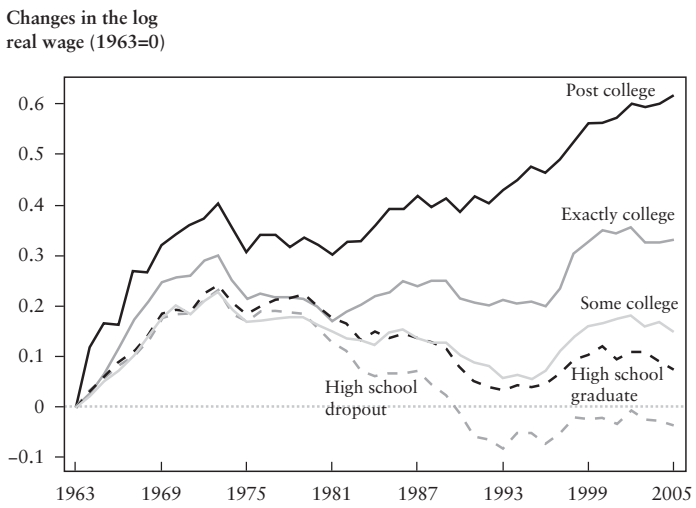

Figure 4.4
Changes in Composition-Adjusted Real Log Weekly Full-Time Wages of U.S. Men by Education, 1963–2005
Source: Current Population Survey and U.S. Bureau of Economic Analysis.

ers with a postcollege education), these sharp gains were not accompanied by a significant rise in between-group inequality. Following the 1973 oil shock, earnings levels stagnated for all educational groups, while income inequality remained largely steady. Commencing in 1979, income inequality rose rapidly even as average earnings remained stagnant. The real wages of workers with a four-year college degree or postcollege education increased significantly, while the real wages of those with a high school degree or less plummeted. Most recently, from the early 1990s forward, overall earnings levels have risen again, but this growth has been bimodal: the earnings of less educated workers (those with a high school degree or lower/less) rose modestly, the earnings of the most highly educated (those with postcollege education) rose extremely rapidly, and the earnings growth of those with some college education was comparatively weak. Thus, the polarization of *overall* earnings growth in the 1990s, as depicted in Figures 4.2 and 4.3, is reflected in a contemporaneous polarization of earnings *across* education groups.

III. Rising Inequality: The Role of Demand Shifts for College-Educated Versus Non-College-Educated Workers

To interpret the forces shaping the rise and subsequent polarization of wage inequality—and to forecast its future trajectory—it is critical to assess the degree to which shifts in labor demand are responsible for the observed patterns. In this section, we ask whether the rising wages of workers with high levels of educational attainment versus those with low levels of educational attainment can be explained by a combination of demand and supply shifts that favor more educated workers. A particularly simple and attractive formulation of this supply-demand framework posits that there are two major skill groups in the labor market, those with at least four-year college degrees and those with high school degrees. Both skill groups, termed college equivalents and high school equivalents, are in demand as employees by firms and, critically, these groups are imperfect substitutes in production. Thus, an increase in the relative supply of one group reduces its earnings relative to the other group.[9]

Figure 4.5 illustrates the intuitive appeal of this conceptual framework. In this figure, the series labeled "Log Wage Differential" plots

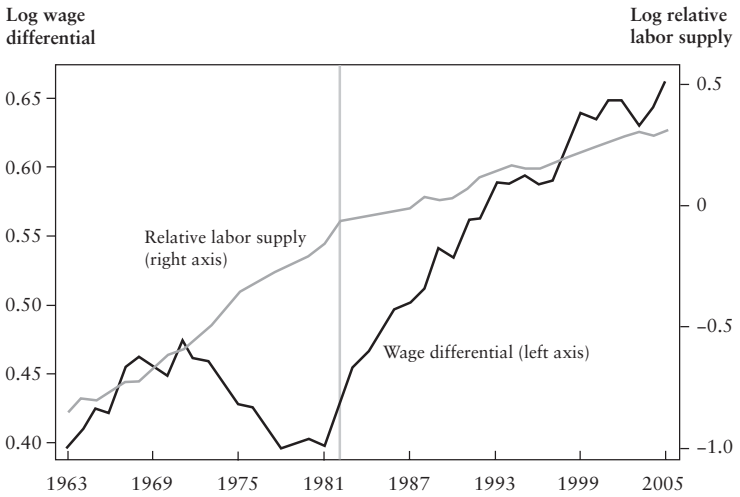

Figure 4.5
College Relative to High School Labor Supply and the College-High School Wage Differential, 1963–2005
Source: Current Population Survey.

the composition-adjusted log college/high school earnings gap for 1963 through 2005. Consistent with the more disaggregated earnings series summarized in Figure 4.4, the college/high school gap rises in the 1960s, contracts modestly in the 1970s, and then expands rapidly from 1981 forward. By 2005, the college/high school gap has attained its highest level—94 percent or 66 log points—since 1915 (see Goldin and Katz 2007). This gap is nearly double the 1963 level of 49 percent (40 log points). The second series in Figure 4.5, labeled "Log Relative Supply," depicts the evolution of the composition-adjusted supply of college-educated relative to high-school-educated workers in the same time period. This series reveals an acceleration of the growth in the relative supply of college workers in the 1970s compared to the 1960s, followed by a dramatic slowdown starting in 1982. Notably, this deceleration, caused by slowing college attainment among cohorts of youth born after 1950 (see Card and Lemieux 2000), corresponds closely with the sharp jump in the college/high school wage premium after 1981. Thus, the juxtaposition of these series suggests that fluctuations in the rate of supply growth of

college-educated workers, overlaid on secularly rising demand for college workers, may provide a reasonable summary explanation for the growth of college wage premium. Indeed, this hypothesis was famously espoused by Katz and Murphy in 1992, who found that it provided an excellent fit for trends in the college wage premium for the years 1963 to 1987, their data set's ending year.

To explore the power of this framework for more recent trends in inequality, we re-estimate the Katz-Murphy model using earnings data extended to 2005, thus going 18 more years beyond their original work. Our illustrative conceptual framework starts with a Constant Elasticity of Substitution production function for aggregate output Q with two factors, college equivalents (c) and high school equivalents (h):

$$(1) \quad Q_t = [\alpha_t(a_t N_{ct})^{\rho} + (1 - \alpha_t)(b_t N_{ht})^{\rho}]^{1/\rho}$$

where N_{ct} and N_{ht} are the quantities employed of college equivalents (skilled labor) and high school equivalents (unskilled labor) in period t, a_t and b_t represent skilled and unskilled labor augmenting technological change, α_t is a time-varying technology parameter that can be interpreted as indexing the share of work activities allocated to skilled labor, and ρ is a time invariant production parameter. Skill-neutral technological improvements raise a_t and b_t by the same proportion. Skill-biased technological changes involve increases in a_t/b_t or α_t. The aggregate elasticity of substitution between college and high-school equivalents is given by $\sigma = 1/(1 - \rho)$.

Under the assumption that college and high school equivalents are paid their marginal revenue products, we can use equation (1) to solve for the ratio of marginal products of the two labor types, yielding a relationship between relative wages in year t, w_{ct}/w_{ht}, and relative supplies in year t, N_{ct}/N_{ht} given by

$$(2) \quad \ln(w_{ct}/w_{ht}) = \ln[\alpha_t/(1 - \alpha_t)] + \rho \ln(a_t/b_t)] - (1/\sigma)\ln(N_{ct}/N_{ht}),$$

which can be rewritten as

$$(3) \quad \ln(w_{ct}/w_{ht}) = (1/\sigma)[D_t - \ln(N_{ct}/N_{ht})],$$

where D_t indexes relative demand shifts favoring college equivalents and is measured in log quantity units. The impact of changes in relative skill supplies on relative wages depends inversely on the magnitude of

aggregate elasticity of substitution between the two skill groups. The greater is σ, the smaller the impact of shifts in relative supplies on relative wages, so the fluctuations in demand shifts (D_t) must be greater to explain any given time series of relative wages for a given time series of relative quantities. Changes in D_t can arise from (disembodied) skill-biased technological change, non-neutral changes in the relative prices or quantities of non-labor inputs, and shifts in product demand.

Following the approach of Katz and Murphy (1992), we directly estimate a version of equation (3) to explain the evolution from 1963 to 2005 of the overall log college/high school wage differential series for full-time, full-year workers from the March CPS shown in Figure 4.5. We substitute for the unobserved demand shifts D_t with a simple linear time trend. We also include an index of the log relative supply of college/high school equivalents:[10]

$$(4) \quad \ln(w_{ct}/w_{ht}) = \gamma_0 + \gamma_1 t + \gamma_2 \ln(N_{ct}/N_{ht}) + \varepsilon_t,$$

where γ_2 provides an estimate of $1/\sigma$.

Figure 4.6 plots the observed college/high school premium for years 1963 to 2005 alongside the fitted values of equation (4), generated by estimating the Katz-Murphy model for calendar years 1963 through 1987, and then extrapolating the estimates through the year 2005 based on the observed evolution of college/high school relative supply. The model implies a strong, secular growth of college/high school relative demand at the rate of about 2.6 log points annually over 1963 to 1987. Though the Katz-Murphy model is only fit to data through 1987, it does an excellent job of *forecasting* the growth of the college wage premium through 1992, thus suggesting that demand shifts favoring college-educated workers continued apace in these years. This demand growth is typically interpreted as evidence of skill-biased technological change, which refers to any introduction of a new technology, change in production methods, or change in the organization of work that increases the demand for more-skilled labor relative to less-skilled labor at fixed relative wages. Indeed, comprehensive analyses of longer time series by Autor, Katz, and Krueger (1998) and Goldin and Katz (2007) suggest that such skill-biased demand shifts have been underway for many decades—and that these shifts have accelerated in the second half of the twentieth century.

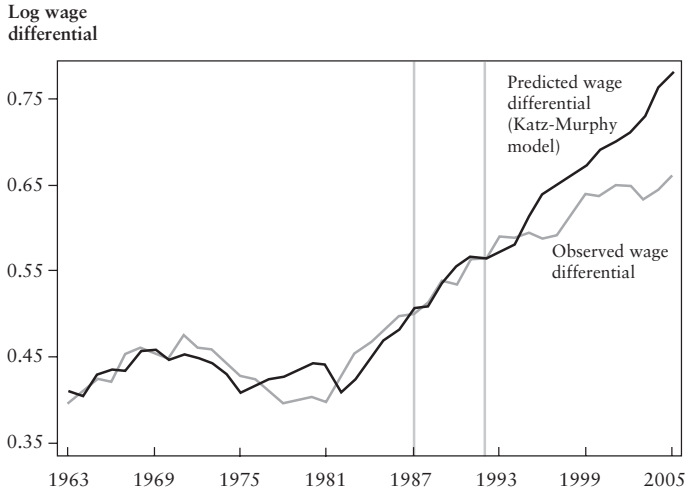

Figure 4.6
College-High School Wage Differential, 1963–2005: Observed and Predicted Values
Source: Current Population Survey.
Note: Predicted values are estimated for the years 1963–1987 and extrapolated to 2005, based on the Katz-Murphy model (Katz and Murphy 1992).

What drives these secular demand shifts? A large literature, reviewed in Katz and Autor (1999) and Katz (2000), yields two consistent findings that suggest that skill-biased technological change has played an integral role.[11] The first finding is that the relative employment of college-educated workers and non-production workers (that is, professional, managerial, and technical workers rather than line workers) has increased rapidly within detailed industries and within business establishments in the United States during the 1980s and 1990s, despite the sharp rise in the relative wages of these groups (see Dunne, Haltiwanger, and Troske 1997; Autor, Katz, and Krueger 1998). Similar patterns of within-industry increases in the proportion of skilled workers are apparent in other advanced nations (Berman, Bound, and Machin 1998; Machin and Van Reenen 1998). These findings suggest strong within-industry demand shifts favoring the more skilled, meaning more college-educated, work-

ers.[12] Second, a wealth of quantitative and case-study evidence documents a striking correlation between the adoption of computer-based technologies (and associated organizational innovations) and the increased use of college-educated labor within detailed industries, within firms, and across plants within industries (see Doms, Dunne, and Troske 1997; Autor, Levy, and Murnane 2002; Levy and Murnane 2004; Bartel, Ichniowski, and Shaw 2007).

While this simple, demand-side explanation is appealing, this story is not entirely confirmed by the data. The Katz-Murphy model accurately predicts the ongoing growth of the college wage premium between 1987 and 1992, the model substantially *overpredicts* the growth of the college wage premium going forward from 1992. This suggests, unexpectedly, that demand shifts favoring college-educated workers have *slowed* since 1992.[13] This implied slowdown in trend demand growth in the 1990s is potentially inconsistent with a simple skill-biased technical change story that appeals to the ongoing growth of computer investments, since these investments continued rapidly throughout the 1990s, particularly with the rapid diffusion of the Internet. Why has this slowdown in demand for college-educated workers occurred?

One potential explanation for this implied slowdown is the strong cyclical labor market of the expansion of the 1990s, leading to a tight labor market that may particularly boost the earnings of workers with comparatively lower levels of educational attainment. The weakening of some labor market institutions, such as the erosion of the real value of the minimum wage since the early 1980s, might also have played a role. These hypotheses are evaluated by Autor, Katz, and Kearney (2008), however, and are found lacking in explanatory power. After accounting for the role of supply shifts, the real minimum wage and prime-age male unemployment rates provide only modest additional explanatory power for the evolution of earnings inequality, and thereby reduce the extent of the estimated slowdown in trend demand growth over the last decade.[14] These cyclical and institutional factors are insufficient to resolve the puzzle posed by slowing trend-relative demand for college-educated workers in the 1990s.[15]

A closer look at the data suggests why the simple CES model with two inputs—college and high school equivalents—fails to provide an adequate explanation of the evolution of between-group wage inequality starting in the early 1990s. As shown in Figure 4.4, the real, composition-adjusted earnings of full-time, full-year workers at different levels of educational attainment polarized after 1987 in a manner consistent with the divergent trends in 90-50 and 50-10 inequality documented in Figure 4.3. In particular, the wage gap between males with a postcollege education and those with a high school education rose rapidly and monotonically from 1979 through 2005, increasing by 43.1 log points overall and 15.4, 15.7, and 12.0 points, respectively, between 1979–1988, 1988–1997, and 1997–2005.[16] By contrast, after increasing by 13.3 log points between 1979 and 1987, the wage gap between males with exactly a four-year college degree and those with a high school education rose comparatively slowly thereafter, by 4.5 and 9.0 log points, respectively, between 1988–1997 and 1997–2005. By implication, between 1988 and 2005, the earnings of postcollege educated males rose by 14.2 log points *more* than the earnings of males with only a four-year college degree.[17] Conversely, at the bottom of the wage distribution, the wage gap between high school graduates and high school dropouts increased steadily from 1979 and 1997, then flattened or reversed.

This pattern, in which wage gaps *within* college-educated and non-college-educated workers groups diverge, is inconsistent with the basic, two-factor CES model. In this model, the labor input of all college-educated worker subgroups is assumed to be perfectly substitutable up to a scalar multiple, and this substitution holds similarly for non-college-educated worker subgroups. Accordingly, the wage *ratio* of college-educated to postcollege-educated worker should be roughly constant, as should be the wage ratio of high school dropouts to high school graduates. This two-factor assumption fits the data rather well from 1963 to 1987. However, after 1987 the drastic rise in earnings of postcollege-educated workers relative to workers with only a four-year college degree, and the slightly increasing earnings of high school dropouts relative to high school graduates after 1997 represent significant departures from the model's assumptions. Fundamentally, the two-factor model does not accommodate a setting in which the wages of very high and very low-skilled workers rise relative to those of middle-educated workers—that is, the model does not accommodate a setting in which wage growth polarizes. We consider the

sources of this polarization next, after briefly considering the role of the minimum wage in greater detail.

IV. The Elusive Role of the Minimum Wage

In contrast to our conclusions above, several other studies, including Lee (1999), Card and DiNardo (2002), and Lemieux (2006b), find that fluctuations in the U.S. minimum wage play a primary role in the rise of wage inequality since 1980. The minimum wage explanation for rising wage inequality has obvious appeal. As shown by Card and DiNardo (2002), there is a striking time series relationship between the real value of the federal minimum wage and hourly wage inequality, as measured by the 90-10 log earnings ratio. This relationship is depicted in Figure 4.7. A simple regression of the 90-10 log hourly wage gap from the May/ORG CPS for the years 1973 to 2005 on the real minimum wage yields a coefficient of –0.74 and an R-squared of 0.71. Based in part on this tight correspondence, Card and DiNardo (2002) and Lemieux (2006b) argue that much of the rise in overall and residual inequality over the last two decades may be attributed to the minimum wage.[18] In a cross state analysis of the minimum wage and wage inequality for the period 1979 to 1991, Lee (1999) reaches a similar conclusion.

A potential objection to this argument is that the majority of the rise in earnings inequality over the last two decades occurred in the upper half of the earnings distribution. Since it is not plausible that a declining minimum wage could cause large increases in upper-tail earnings inequality, this observation suggests that the minimum wage is unlikely to provide a satisfying explanation for the bulk of inequality growth. Not surprisingly, as shown in the upper panel of Figure 4.7, the real minimum wage is highly correlated with lower-tail earnings inequality between 1973 and 2005; a 1 log point rise in the minimum is associated with 0.26 log point compression in lower-tail inequality. Somewhat surprisingly, the minimum wage is also highly correlated with upper-tail inequality: a 1 log point rise in the minimum is associated with a 0.48 log point compression in upper-tail inequality; see Figure 4.7, lower panel.

Autor, Katz, and Kearney (2008) explore these relationships in greater detail by estimating a set of descriptive regressions for hourly earnings inequality among three pairs of income percentiles, 90-10, 90-50, and 50-10, over the period dating from 1973 to 2005. In addition to the minimum wage measure used in Figure 4.7, Autor, Katz, and Kearney (2008) augment these models with a linear time trend, a measure of college/high school relative supply (calculated from the May/ORG CPS), the male prime-age unemployment rate (as a measure of labor market tightness), and in some specifications a post-1992 time trend, reflecting the estimated trend reduction in skill demand in the 1990s. The main finding from these models is that the strong relationship between the minimum wage and both upper- and lower-tail inequality is highly robust.

These patterns suggest that the time series correlation between minimum wages and income inequality is unlikely to provide causal estimates of minimum wage impacts. Indeed, the relationship between the minimum wage and upper-tail inequality is potential evidence of spurious causation. Although the decline in the real minimum wage during the 1980s likely contributed to the expansion of lower-tail inequality—particularly for women—the robust correlation of the minimum wage with upper-tail inequality suggests that other factors are at work.[19] One possibility is that federal minimum wage changes (or lack of changes) during these decades were partially a response to political pressures associated with changing labor market conditions and the costs a minimum wage increase would to impose on employers. This "political economy" story could help explain the coincidence of falling minimum wages and rising upper-tail inequality.[20]

V. Why Is Labor Demand Polarizing? The Possible Role of Technology and Outsourcing

Why, following the monotonic surge of earnings inequality from 1979 to 1987, did U.S. wage growth polarize, with a strong, persistent rise in inequality in the upper half of the distribution, and a slowing, then slight reversal of inequality, in the lower-half of the distribution? Based on the analysis discussed above, along with further evidence presented in Autor, Katz, and Kearney (2008), we conclude that neither standard supply-demand models nor conventional institutional explanations are sufficient to explain the evolution of U.S. income inequality since the late 1980s.

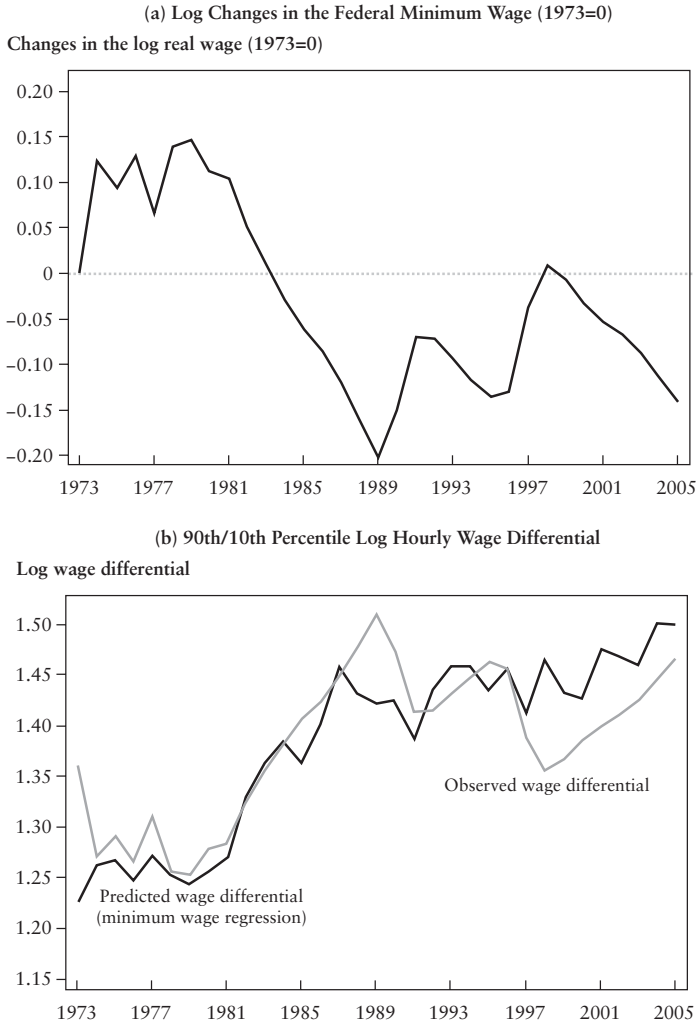

Figure 4.7
Log Hourly Wage Differentials, 1973–2005: Observed Values and Predicted Values from a Regression on the Log Real Federal Minimum Wage
Source: Current Population Survey and U.S. Bureau of Economic Analysis.
Note: Nominal minimum wages are deflated to real log values using the PCE deflator. In panel (a), the real log minimum wage measure is normalized to zero in 1973. Subsequent panels depict the observed wage gap (between the 90th and 10th percentiles, 50th and 10th percentiles, and 90th and 50th percentiles) for all hourly workers from the May and Outgoing Rotation Group Current Population Survey samples in each year plotted alongside the predicted values from separate OLS regressions of the relevant wage gap on a constant and the contemporaneous real log minimum wage.

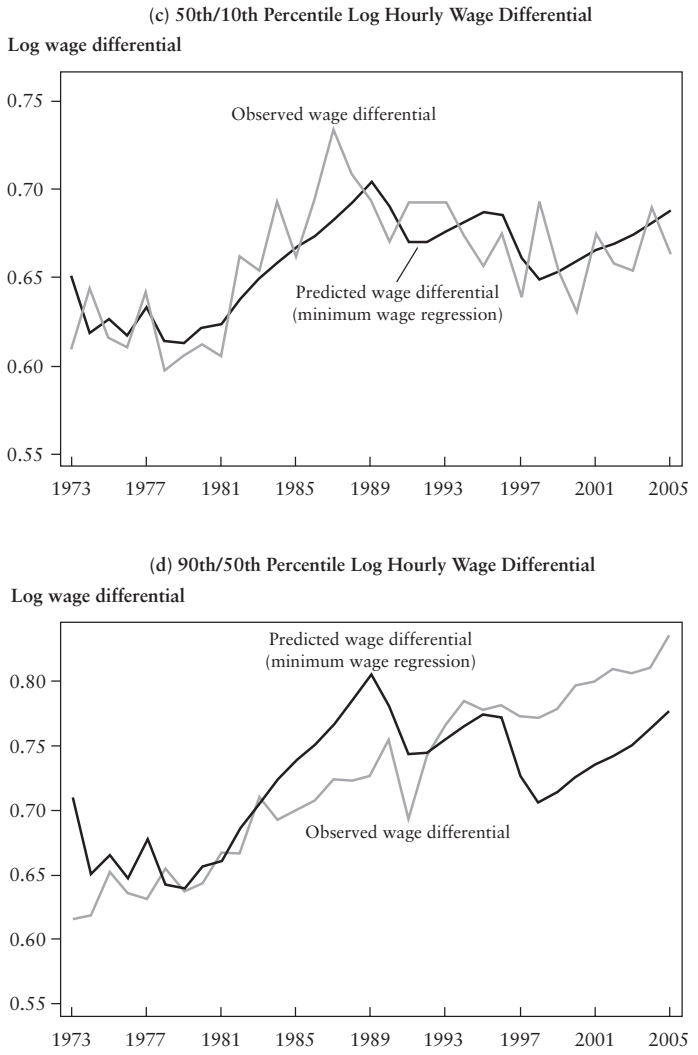

Figure 4.7 (continued)

In this section, we focus on one potentially viable hypothesis for the polarization of earnings inequality, which focuses on changing demand for job tasks, often linked to computerization and, over the longer term, outsourcing. As argued by Autor, Levy, and Murnane (2003, hereafter this article is referred to as "ALM"), and amplified by Goos and Manning (2007); Spitz-Oener (2006); Autor, Katz, and Kearney (2006); and Dustmann, Ludsteck, and Schönberg (2007), the term "skill-biased technological change" presents an inadequate description of the shifts in skill demands that were induced or abetted by the rapid price declines in computer technology over the last three decades. In the task framework proposed by ALM, computerization has non-monotone impacts on the demand for skills throughout the earnings distribution, sharply raising demand for the cognitive and interpersonal skills used by college-educated professionals and managers (termed "abstract tasks") and reducing demand for clerical and routine analytical and mechanical skills that comprised many middle-educated white collar and manufacturing production jobs (termed "routine tasks").[21] Somewhat paradoxically, computerization has probably had little direct impact on the demand for the non-routine manual skills (termed "manual tasks") used in many "low-skilled" service jobs such as health aides, security guards, hospital orderlies, janitors, and servers. Because the interpersonal and environmental adaptability demanded by these manual tasks has proven extraordinarily difficult to computerize (to date), these manual activities may in fact grow in importance as a share of labor input.[22]

The ALM framework suggests that computerization, along with complementary forces such as international outsourcing, may have raised demand for skills among higher-educated workers, depressed skill demands for middle-educated workers, and left the lower echelons of the wage distribution comparatively unscathed.[23] Goos and Manning (2007) label this process a "polarization of work," and argue that it may have contributed to a hollowing out of the wage distribution in the United Kingdom from 1975 to 2000. Spitz-Oener (2005) and Dustmann, Ludsteck, and Schönberg (2007) report a similar polarization of employment for the former West Germany during the 1979 to 1999 period.[24]

To illustrate the relevance of shifts in task demands for changes in skill demands, we link data on task intensity by occupation (information taken from the U.S. Department of Labor's online Dictionary of Occupational Titles) to data on skill level by occupation contained in the 1980 Census. In this analysis, occupational skill level is measured by the mean years of education of an occupation's workforce (weighting workers by their annual hours worked). Figure 4.8 uses a locally weighted smoothing regression to plot task intensity by occupational skill for each of the three broad task categories above: abstract, routine, and manual tasks.[25] Task intensities are measured as percentiles of the baseline distribution of job tasks in 1960. Thus, an occupation with the median intensity of "routine" task input in 1960 would receive a score of 50. This figure shows that the intensity of abstract skill input is monotonically rising in occupational skill level (reflecting more education) and, conversely, the intensity of manual task input is falling in occupational skill level. Most

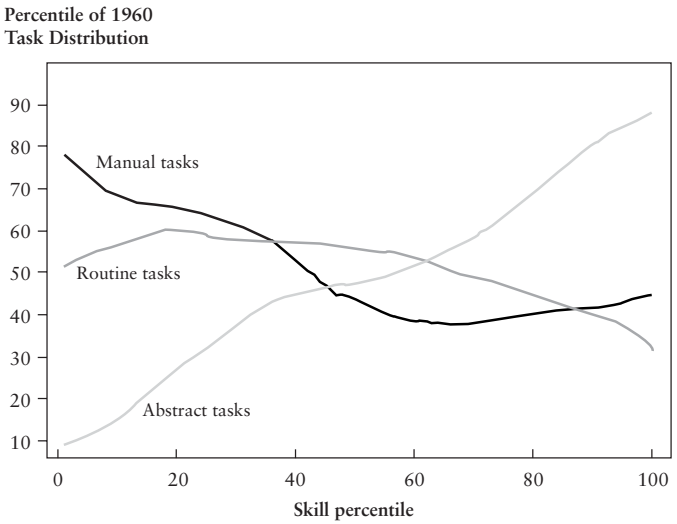

Figure 4.8
Task Intensity by Occupational Skill Percentile, Defined as Occupation's Rank (in Percentiles) in Mean Years of Education
Source: Dictionary of Occupational Titles and U.S. Census Bureau.
Note: Percentiles of the 1960 task distribution are determined using occupational task inputs defined in the Dictionary of Occupational Titles. The figure uses a locally weighted smoothing regression to calculate the plotted values.

significantly, there is a distinctly non-monotone relationship between occupational skill and routine task input. Routine task use is highest between the 20th and 60th percentiles of the skill distribution, and falls off sharply on either side of this range. This non-monotonic relationship is highly relevant because, as documented by ALM, routine task input saw the sharpest decline of all task categories over the last two decades (relative to its initial 1960 level). The substitution of information technology for routine tasks might be expected to contribute to polarization by reducing demand for middle-skill occupations relative to either high- or low-skill occupations.

An implication of the polarization hypothesis is that the twisting of the wage structure observed in recent years is, at least in part, a demand-side phenomenon, induced by rising relative demands for both high- and low-skill tasks. This implication is testable, and we provide a simple evaluation here. Following analysis for the United Kingdom in Goos and Manning (2007), we use U.S. Census data to explore how employment growth by occupation over the last two decades is related to occupational skill, as proxied by educational levels.[26] Our hypothesis is that, if the wage structure changes observed in the 1980s and 1990s are driven in substantial part by demand shifts, wage changes by earnings level and employment changes by skill level should positively covary in both decades.

To test this implication, in the upper panel of Figure 4.9 we plot the change in the share of total hours worked in the economy from 1980–1990 and 1990–2000 by occupation skill percentile, using the education-based occupational skill measure developed earlier.[27] For the decade of the 1980s, we see substantial declines in employment shares at the bottom end of the skill distribution, and observe strongly monotonic increases in employment shares as we move up the skill distribution. In contrast, employment growth in the 1990s appears to have polarized. There is rapid employment growth in highest-skill jobs (at or above the 75th percentile), a decline in the employment shares of middle-skill jobs (those at percentiles between 30 to 75), and flat or rising employment shares in the lowest-skill jobs, those in deciles one through three.

This pattern of job growth corresponds closely with the observed pattern of wage structure changes in each decade, as is shown in the lower panel of Figure 4.9. Real wage growth was essentially monotone in terms

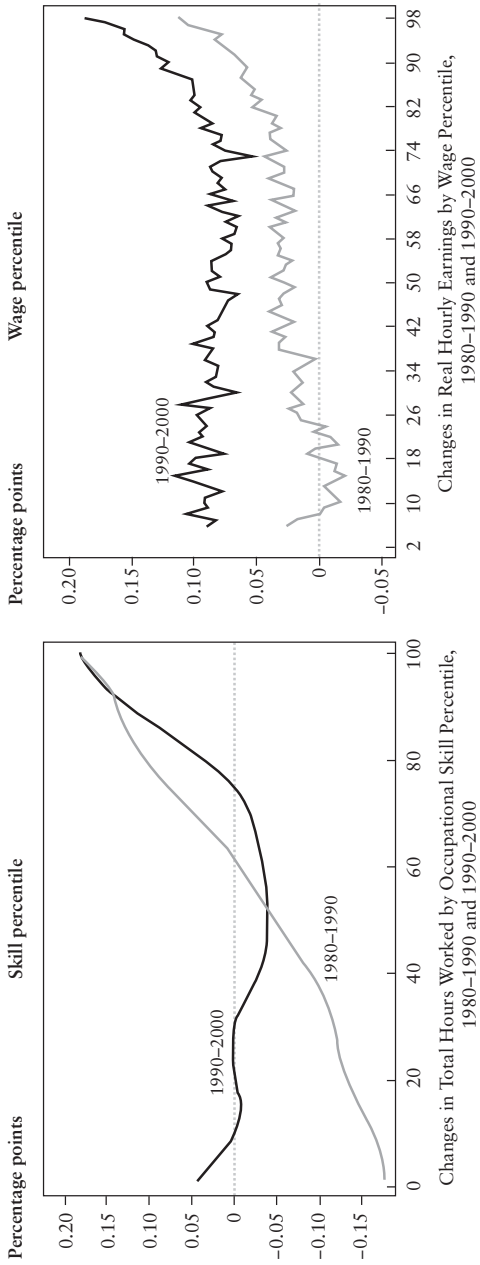

Figure 4.9
Changes in Occupational Employment Shares and Real Hourly Earnings, 1980–1990 and 1990–2000
Source: U.S. Census Bureau.
Note: Occupational skill percentile is defined as an occupation's rank (in percentile) in mean years of schooling.

of wage percentiles in the 1980s, with especially sharp wage growth above the 75th percentile and especially sharp declines below the 30th percentile. In the decade of the 1990s, however, wage growth was more U-shaped. Wage growth was stronger below the 30th percentile, and especially above the 80th percentile of the distribution, than throughout the remainder of the distribution. Thus, despite substantial differences in the evolution of inequality between the 1980s and 1990s, labor market prices and quantities (as measured by wage and skill percentiles) appear to positively covary in each decade.

To provide a slightly more rigorous assessment of this observation, we estimate a set of ordinary-least-square models of the form,

$$(5) \quad \Delta E_{p\tau} = \alpha_\tau + \beta_\tau \Delta \ln W_{p\tau} + \varepsilon_{p\tau},$$

where changes in log employment share by skill percentile are regressed on changes in log wages by wage percentile in each decade. Here, $\Delta E_{p\tau}$ represents the change in occupational log employment share at skill percentile p in decade τ, and $\Delta \ln W_{p\tau}$ is the change in real log hourly earnings at the corresponding wage percentile in the same decade.[28] Using data for the 4th through 97th percentiles of the earnings and skill distributions (thus trimming outliers at the tails), we estimate that $\beta_\tau = 300$ ($t = 3.75$) for the 1980s, and that $\beta_\tau = 2.96$ ($t = 1.90$) for the 1990s. Thus, both the monotone rise of wage inequality in the 1980s and the polarized growth of wage inequality in the 1990s are mirrored by conformal changes in employment by skill. This finding is consistent with a demand-side explanation for observed wage changes.[29]

We have further experimented with these simple models by including linear terms in wage percentiles in addition to (or instead) of estimated wage changes by percentile. For the decade of the 1980s, we find that a linear function of wage percentiles fits the observed pattern of skilled employment growth better than does the observed change in earnings by percentile. In the 1990s, by contrast, the linear term is insignificant, and the estimate of $\beta_{90\text{–}00}$ is hardly affected by its inclusion (either in magnitude or precision). These simple models do not, of course, take into account the substitutability and complementarity among various skill groups, as measured by skill percentiles, and so lack a well-grounded production function interpretation. We nevertheless view these models as suggestive evidence that labor demand shifts have favored low- and high-wage workers relative to middle-wage workers over the last fifteen years—a pattern that stands in contrast to the shifts in labor demand during the 1980s, which appear to have been monotonically rising in skill.

VI. The Jobs of the Future: Both "Lousy and Lovely" Jobs

There is no controversy to the contention that highly-educated professional and managerial jobs, meaning those jobs using abstract skills, will continue growing rapidly. Perhaps less recognized is the corollary implication to this proposition: that jobs demanding "non-routine manual" skills, meaning those skills not readily automated, and hence jobs requiring only low-levels of educational attainment, are likely to expand as well. In the memorable phrase used in Goos and Manning (2007) to describe the polarization of employment they found in the United Kingdom, we seem to be confronting a future labor market in which jobs are either "lovely" or "lousy." To provide some direct evidence on the relevance of this hypothesis, we look at the changing occupational structure of employment in the United States.

Table 4.1 shows the educational level and employment shares in six major occupational groups covering all U.S. employment categories: 1) managerial and professional specialties; 2) technicians, sales, and administrative support; 3) precision production, craft, and repair; 4) service occupations; 5) operators, fabricators, and laborers; and 6) farming, fishing, and forestry occupations. The category in which workers have the highest average educational level is managerial and professional specialty occupations, followed, at some distance, by technicians, sales, and administrative support. The four remaining categories—each averaging half the size of the first two—are demonstrably less education-intensive Whereas in the year 2000, high school dropouts made up 2.2 percent of employment in professional/managerial jobs and 6.7 of employment in technical, sales and administrative support jobs, they comprised 20-plus percent of employment in the four remaining categories.

As discussed by Autor and Dorn (2007), employment growth has not been uniform across these six categories. Figure 4.10 shows that managerial and professional specialty occupations—the highest skilled cat-

Table 4.1
Employment and Educational Statistics for Six Main Census Occupation Groups in 2000

	Employment Share	Median Hourly Wage	% High School Dropout	% No College	% Female	% Non-White	% Foreign Born
All Occupations	100.0	$13.58	12.1	39.3	42.1	21.6	14.2
Service Occupations	13.4	$9.40	21.3	55.1	51.3	30.8	19.7
All Occupations except Service Occupations	86.6	$14.42	10.7	36.8	40.8	20.1	13.3
Managerial and Professional Specialty Occupations	30.2	$19.23	2.2	11.4	46.5	16.2	11.8
Technicians, Sales, and Administrative Support	28.8	$12.50	6.7	35.0	58.8	20.8	11.6
Farming, Forestry, and Fishing Occupations	1.3	$7.50	33.0	67.2	14.9	20.6	22.3
Precision Production, Craft, and Repair Occupations	12.3	$14.40	19.9	60.4	8.6	18.7	14.3
Operators, Fabricators, and Laborers	14.0	$11.49	27.3	71.9	22.2	28.3	18.6

Source: Autor and Dorn (2007) calculated from Census IPUMS 2000 5 percent sample. All calculations are weighted by hours of annual labor supply and exclude those under age 18 or over age 65.

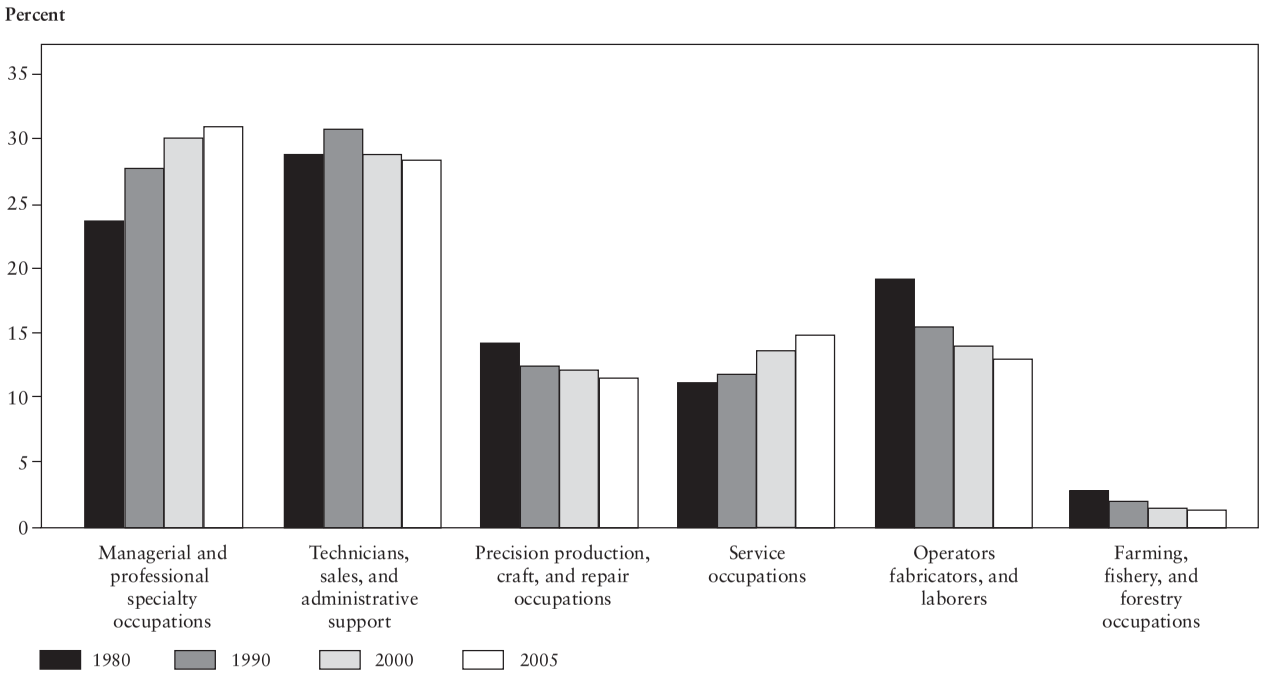

Figure 4.10
Employment Shares of Major Census Occupation Groups, 1980, 1990, 2000, and 2005
Source: U.S. Census Bureau.

egory—experienced consistent, rapid growth between 1980 and 2005, gaining 7.1 percentage points as a share of overall employment between 1980 and 2005, a 30 percent increase. By contrast, employment in the middle skill group of technical, sales and administrative support occupations showed an inverse U-shape pattern over this period, expanding in the 1980s and then contracting to below its initial 1980 level over the next 15 years; this is consistent with the growing substitution of technology for completing routine tasks. Most strikingly, employment shares in three of the four low-skill occupations fell sharply in each decade. Between 1980 and 2005, farming, forestry, and fishery occupations contracted by more than 50 percent as a share of employment, while the category comprised of operators, fabricators, and laborers contracted by 33 percent, and precision production, craft, and repair occupations contracted by 19 percent.

Standing in sharp contrast to these patterns of declining employment, however, is the experience of service occupations. Despite being among the least educated and lowest paid occupations in the U.S. economy, employment in service occupations expanded in each decade between 1980 and 2005, rising from 11.0 percent of employment in 1980 to 11.8 percent in 1990, to 13.7 percent in 2000 and to 14.9 percent in 2005. This 35 percent increase is 6 percentage points larger than the gain in employment shares of managerial and professional occupations during the same period.

What is unique about service jobs? Table 4.2 lists the major service occupations, the largest of which are: food preparation and service; health service support (a group that excludes registered nurses and other skilled medical personnel); and buildings and grounds cleaning and maintenance.[30] These are low-paying jobs; in the year 2000, 73 percent had hourly wages below the overall hourly median. From the perspective of our conceptual framework, what distinguishes these service occupations is that each is highly intensive in "non-routine manual" tasks—activities requiring interpersonal and environmental adaptability yet little in the way of formal education. These are precisely the job tasks that are difficult to automate with current technology because these jobs are non-routine *and* difficult to outsource because, in large part, the tasks involved must be produced and performed in-person in real time.

Table 4.2
Employment Statistics for Service Occupations in the Year 2000, and Growth Rates for 1980–1990 and 1990–2000

	Employment Share in 2000	Median Hourly Wage in 2000	% Female in 2000	% Non-White in 2000	% Foreign Born in 2000	% Growth 1980–90	% Growth 1990–2000
All Service Occupations	13.4	$9.40	51.3	30.8	19.7	8.5	15.8
Housekeeping, Cleaning, Laundry	0.8	$7.21	80.7	45.4	37.8	–11.0	2.5
Protective Service	2.1	$15.05	17.8	25.2	7.2	15.2	8.1
Food Preparation and Service	3.5	$7.55	53.6	31.7	25.0	8.9	11.4
Health Service Support	2.7	$9.93	73.0	31.6	15.1	5.2	64.4
Building and Grounds Cleaning and Maintenance	2.2	$9.53	19.8	32.9	25.2	10.6	–7.5
Personal Appearance	0.6	$8.65	81.4	27.4	20.5	7.1	–1.7
Child Care	0.8	$7.00	94.2	28.2	15.3	11.4	57.1
Recreation and Hospitality	0.4	$10.35	46.3	26.5	15.0	25.0	75.0
Other Personal Service	0.3	$12.02	56.2	17.3	12.3	17.2	0.0

Source: Autor and Dorn (2007), calculated from Census IPUMS 1980, 1990, and 2000 5 percent samples. All calculations are weighted by hours of annual labor supply and exclude those under age 18 or over age 65.

Employment projections from the Bureau of Labor Statistics (BLS) Employment Outlook confirm the view that low-skilled services are likely to be a major contributor to U.S. employment growth going forward. The BLS forecasts that employment in service occupations will increase by 5.3 million, or 19 percent, between 2004 and 2014.[31] The only major occupational category with greater projected growth during this time period is professional occupations, which are predicted to add 6 million jobs, a 21.2 percent increase.[32] Like all forecasts, these should be treated as tentative. Historically, the BLS has underpredicted the growing demand for professional and managerial occupations (see Bishop and Carter 1991; Freeman 2006).

It is likely that the rapid growth of service employment in the United States has multiple causes. One is the direct substitution of computerization for routine tasks, which causes the share of labor input devoted to non-routine activities to increase.[33] A second force, though of highly uncertain magnitude, is international outsourcing, which complements computerization in permitting routine tasks previously performed by domestic workers to be sourced to other locations.[34]

But these technological forces are not the only drivers of this increased demand. The aging of the U.S. population contributes to the growth of health services support occupations—and this contribution will become more important going forward. Supply-side factors may also be important. Recent work by Cortes (2006) demonstrates that influxes of low-skilled immigrants into major American cities causes the market prices of non-traded, low-skill intensive services to fall and consumption of these services to rise. Thus, the rapid growth of service employment is also partly attributable to U.S. immigration policy.

A final, relatively unstudied, factor potentially contributing to the growth of service employment is the rise of income inequality itself. Household consumption of services appears to be highly income elastic (Mazzolari and Ragusa 2007). This makes it plausible that the strong, secular rise in the earnings share of high-income households over almost three decades has increased final demand for services (see Piketty and Saez 2003; Autor, Katz, and Kearney 2008). Preliminary evidence supporting this hypothesis is offered by Autor and Dorn (2007) and Mazzolari and Ragusa (2007), who find that service employment growth in the United States has been greatest in the metropolitan areas where income inequality has increased the most. Given that the rise of high incomes shows no signs of abating, this force may stimulate additional demand for low-education, in-person services.

VII. Possible Labor Supply Responses: Human Capital Policy and Immigration Policy

Proceeding on the view that U.S. employment growth will be concentrated at the tails—in other words, in occupations requiring either high or low levels of education—how might labor supply respond? Because other papers in this volume treat this question in great detail, I offer only brief remarks on this issue, focusing on topics where policy is likely to have particular leverage.[35]

A first point of paramount economic importance is that the returns to human capital investments are currently extremely high. While some research has highlighted the fact that the college wage differential plateaued in the early 1990s, this observation needs to be placed in appropriate context. Even in the late 1990s, the college wage differential stood at a near-historic level (see Goldin and Katz 2007). And, as indicated by Figure 4.4, there was a further pickup in the pure college/high school premium after 1999. Moreover, the wage differential associated with postcollege educational returns has risen rapidly and near-continuously from 1980 to the present. Thus, postsecondary education appears to be an excellent investment.

Responding to this price signal, college enrollment of U.S. youth has risen considerably since the premium to earning a four-year college degree began its historic rise in the early 1980s. After falling slightly between 1970 and 1980, the fraction of 20 to 24 year-olds enrolled in post-secondary education rose from 35.9 percent in 1980, to 42.7 percent in 1990, to 44.7 percent in 2000, and 49.3 in 2005 (U.S. Department of Education 2007). College completion rates have not risen commensurately, however. Bound, Lovenheim, and Turner (2007) report that from 1970 forward, the share of youth obtaining the equivalent of a four-year college degree by age 23 rose only slightly for cohorts completing high school.[36] Simultaneously, the completion rate among those attending college fell by

10 percent, and the share completing a degree within four years (among degree completers) fell by 20 percent. Although some increase in the college non-completion rate is to be expected as the fraction of students enrolling in college rises, these statistics suggest that there may be room to improve the outcomes of these initial investments in a college education. Indeed, despite having led the world in high school and college completion for most of the twentieth century, U.S. young adults are now in the middle of the pack in the Organization for Economic Cooperation and Development (OECD) in terms of educational attainment (see OECD 2006; Goldin and Katz 2007).

There is ongoing debate about the degree to which financial constraints hinder college matriculation among U.S. youth. There is little doubt, however, that the gap in college attendance rates when gauged by parental income, race, and ethnicity remain large and may have potentially widened over the last 25 years (see Ellwood and Kane 2000; Heckman and Carneiro 2002).[37] Considerable evidence shows that reductions in college costs (due to tuition reductions or financial aid) greatly increase college attendance for youths from moderate income families (see Dynarski 2002; Kane 1999) and even affect the postcollege occupational choices of graduates of elite universities (see Rothstein and Rouse 2007). It is therefore abundantly clear that college costs have a substantial impact on the college-going decisions and career choices of young adults. Moreover, the economic returns to college attendance for youth from moderate income families appear to be at least as large as those for more advantaged attendees (Card 2001). Thus, there appears a solid case for reducing the financial barriers to college attendance for students from low and moderate income U.S. families.

As has been stressed by many researchers, generous college financing is not sufficient for college success. Students who do not receive adequate human capital investments early in life may gain less from later educational investments (see Heckman and Lochner 2000). Most evidence suggests large returns from early childhood educational interventions (see Currie 2001; Anderson 2007). Thus, efforts to improve college attainment among U.S. students need to commence well before students approach college-going age.[38]

In addition to fostering college attendance by U.S. residents, a critical policy lever for increasing the supply of highly educated workers for the American workforce is attracting skilled students of all nationalities to U.S. colleges and universities. In 2003, 573,000 foreign students were enrolled in U.S. institutions of higher education, an 84 percent increase from the 1980 level (U.S. Department of Education 2005, Table 408). While these numbers appear sizable, the share of foreign students attending U.S. higher educational institutions is small: 2.6 percent in 1980 and 3.4 percent in 2003.[39] Foreign student enrollment brings many of the world's most talented individuals to the United States. A substantial fraction of these students will ultimately remain in the United States and contribute to U.S. invention and economic growth. To the extent that foreign students return to their home countries following their studies, one suspects that many will maintain positive economic—as well as political and cultural—interactions with the United States. Thus, the United States maximizes the return on its leadership role in higher education by producing talent and by attracting it from all parts of the world.

Foreign students enrolled in American universities are heavily concentrated in graduate- and doctoral-level study, particularly in science, technology, engineering, and mathematics (STEM). In the year 2000, foreign students received between 25 and 55 percent of all doctorates awarded by U.S. universities in the key STEM fields of electrical engineering, physics, chemistry and biology. The growth of foreign students enrolled in U.S. graduate degree programs has raised concerns that the wages of native-born Americans are adversely impacted by the influx of foreign students (see Borjas 2006). Though this is an economically sound argument, this concern strikes me as somewhat misplaced given the high and rising earnings of highly-educated workers in the United States—particularly those with graduate degrees. If the relatively abundant supply of foreign students to STEM fields serves to buffer the ongoing growth of wage inequality in the upper half of the U.S. earnings distribution, this may arguably be viewed as an additional benefit. U.S. inequality would almost surely be higher and rising more rapidly at the top if we had to rely more heavily on home-grown talent.

Closely related to the enrollment of foreign students into U.S. universities is the entry of skilled migrants into the U.S. labor force. Though the United States has many of the world's leading universities, the majority of the world's highly educated workers are produced outside of the United

States. Freeman (2006) estimates that the U.S. fraction of Ph.D.s trained relative to total world output will have fallen from about 50 percent in the early 1970s to a projected level of 15 percent in 2010. The growing cadre of highly-educated workers produced outside the United States provides one mechanism for addressing potential skill shortages. As Freeman (2006, p. 10) observes:

> During the 1990s' rapid growth of the U.S. economy, the country greatly increased its employment of scientists and engineers. It did so despite fairly constant numbers of graduates in these fields among citizens or permanent residents and without markedly raising the salaries of these workers… The United States was able to meet increased demands for scientists and engineers without huge increases in salaries by "importing" foreign born specialists in these areas. Some of the foreign born obtained their education in the United States and remained to work in the country. But most of those with B.S. degrees and roughly half of those with higher degrees graduated overseas and came to fill jobs. If the U.S. economy demands more highly skilled workers in the period of projected slow labor force growth, it should be able to increase supplies by admitting more immigrants in areas with rising labor demand, as it did in the 1990s.

This example underscores that, should the United States ultimately face a skill shortage as the baby boom generation retires, this shortage will be a consequence of political choice as a much as demographic destiny. So long as the United States is perceived by educated citizens worldwide as a land of abundant opportunity, it will be able to attract foreign talent to meet domestic labor demand.

In recent years, U.S. immigration policy has been responsive to these demands. The H1-B Visa Program allows U.S. employers to temporarily hire skilled foreigners who have the educational equivalent of a U.S. bachelor's degree. Prior to the mid-1990s, the H1-B quota stood at 65,000 visas per year. During the "dot-com" boom, Congress increased the quota to 115,000 in 1998 and then again to 195,000 in the year 2000. The quota dropped back to 90,000 in 2004, however, and is currently coming under economic pressure. The entire quota of H1-B visas for fiscal year 2007 was exhausted within a span of less than two months. It remains to be seen whether the cap will be lifted again soon.

Over the longer term, it appears possible (though highly uncertain) that the United States will move to a skills-based immigration system. The Secure Borders, Economic Opportunity and Immigration Reform Act, which was hotly debated though ultimately rejected by Congress in 2007, would have prioritized access to U.S. visa applicants according to their educational levels, family ties, age, English language proficiency, and applicants' occupations. Of special note, priority would have been given to workers in "in-demand" occupations. While the virtues and drawbacks of such a system are too complex to adequately address here, two points deserve note. First, the notion of weighting applicants' skills in visa allocation decisions has merit. Second, accurately forecasting what skills and occupations will be "in demand" is generally not something government statistical agencies are able to predict with high reliability (see Freeman 2006). In the existing H1-B Visa Program, by contrast, employers identify and sponsor individual visa candidates. While this process is time- and resource-intensive, it does give employers a strong incentive to sponsor workers who possess particularly valuable skills. Thus, there may be efficiencies in this highly firm-level determination process. Commenting on the immigration reform act current before Congress in 2007, Lowell Sachs of Sun Microsystems opines (quoted in Broache 2007):

> The best the government can hope to do is select a pool of generically potentially qualified candidates, whereas a company knowing exactly what it needs, exactly what skills and exactly what kind of individual can best deliver is going to be far better able to make the right match… What happens if I'm interested in finding a brain surgeon and I've got a bunch of people to pick from, a pediatrician over here, a podiatrist over here, but no brain surgeon?

As this quotation highlights, it is not clear that the U.S. Congress is a better judge of the skill needs of the U.S. economy than are the U.S. employers who, under the H1-B program, hand-select individual foreign workers to meet specific skill needs.

VIII. Job Quality in the Services Occupations

There is, in my view, a solid case for meeting rising demand for professional and technical occupations, in part by importing postsecondary students and highly-trained foreign workers. The same arguments are less persuasive when applied to the demand for low-educated, in-person service workers. Unlike the earnings of four-year college graduates,

wages of high school graduates and dropouts—those most likely to perform service jobs—have fared poorly over the last three decades. Autor, Katz, and Kearney (2008) estimate that real wage growth for workers with high school diplomas and lower educational attainment levels was negative between 1979 and 1995, and only modestly positive from 1995 to 2005. Facilitating increased immigration of competing worker groups appears unlikely to improve this situation.[40] Moreover, while a case can be made that high-skilled workers generate positive human capital externalities—thus making high-skilled immigration a "public good" (see Moretti 2004a and 2004b)—this argument does not apply to low-skilled immigrants.[41] Finally, it is often argued that if the United States does not import high-skilled labor, high-skilled jobs will follow the workers to where they reside. This argument clearly is not relevant for low-skill, in-person services, since the provision of these services is primarily nontradable. In sum, rising U.S. demand for low-skilled services does not represent an economic problem that demands a policy solution. Indeed, a significant benefit of such an upward demand shift is that it is likely to increase the earnings of less-educated workers.[42]

Even given rising demand for service sector jobs, labor supply to services occupations, however, may be sufficiently elastic that wages stay low. Median real hourly wages in service jobs were $8.99 in 1980, $8.76 in 1990, and $9.40 in 2000. These hourly wage rates imply annual, full-time earnings of under $20,000 per year; but of course, many service jobs do not provide full-time, full-year earnings.[43] This income level readily exceeds the poverty threshold for the year 2000 of $17,500 for a family comprised of two adults and two dependent children. Yet $20,000 is probably an inadequate income for families to make optimal investments in childrearing and education. Echoing the concerns above regarding college attainment and early life preparedness, it appears a legitimate concern that the ongoing polarization of earnings levels among U.S. households will ultimately serve to thwart economic mobility among subsequent generations. While the impact of current economic inequality on future mobility cannot be judged until decades after the die is cast, it is clear from the current vantage point that a substantial reduction in mobility would be inimical to the U.S. ideal of meritocracy and equal opportunity.[44] Accordingly, policies that ensure access to excellent education and healthcare for all U.S. families serve in part as a precautionary investment for maintaining economic mobility in the next generation.

There are two primary means to improve the economic conditions of workers in low-skilled service jobs. One way is through transfers and other social supports. For example, the Earned Income Tax Credit has substantially raised labor force participation and earnings of single mothers (see Meyer and Rosenbaum 2001). Programs such as Medicare, Head Start, and the federal Pell grant program provide health insurance, support early childhood educational investments, and reduce the cost of postsecondary education for low-income households. Such programs could be expanded and improved to provide additional assistance to childrearing families. A significant downside risk to such social policies, however, is that these programs are vulnerable to the vicissitudes of budgetary pressures and political sentiments.[45] Transfer programs that do not create a broad constituency of middle- and upper-income beneficiaries are, over the long run, probably less likely to survive.

An alternative means to improve economic conditions of workers in low-skilled service jobs is to "professionalize" these occupations to provide better services and thus command higher wages. Occupational standards and licensing are one means to accomplish this objective. Labor unions are another. The evidence on the efficiency of such steps is decidedly mixed. Kleiner's comprehensive 2005 study of occupational licensing in the United States concludes that professional licensing has primarily served to restrict competition without improving the quality of the services provided. DiNardo and Lee (2004) find that new private sector labor unions certified in the 1980s and 1990s have had little economic impact—positive or negative—on the earnings of newly unionized employees or on the profitability of newly unionized firms. Thus, despite the intellectual appeal of improving wages and performance quality in service occupations, the specific steps to accomplish this objective are not immediately evident, at least to me.

IX. International Outsourcing: A Force of Unknown Magnitude

More than any issue discussed above, there is vast uncertainty about the degree to which international outsourcing of jobs will ultimately affect

domestic labor demand in the United States. At present, most quantitative assessments of the potential impacts of outsourcing are highly preliminary or impressionistic (see Kletzer 2006; Blinder 2007).[46] Theoretical work has also produced somewhat mixed projections on possible labor demand impacts (see Antràs, Garicano, and Rossi-Hansberg 2006a and 2006b; Grossman and Rossi-Hansberg, forthcoming). In my assessment, a safe conclusion is that outsourcing will increase the returns to "knowledge work," both by raising demand for scarce managerial and problem-solving talent, and by increasing the returns on intellectual property developed in advanced economies. Outsourcing will not directly substitute for performing in-person services. Moreover, the income gains accruing to the highly skilled might stimulate additional demand for such lower-level services occupations. Beyond this conjecture, there is little predictive certainty. The possibility appears remote that outsourcing will ultimately displace as large a share of domestic white collar work as international trade and technological change did to decrease domestic demand for blue collar manufacturing work. But then again, the possibility that manufacturing would ultimately employ less than 15 percent of the U.S. workforce in 2000, even while 42 percent of U.S. consumer spending was devoted to purchasing goods, must also have seemed remote several decades earlier (U.S. Congressional Budget Office 2004).

The profound uncertainty about the potential for the international outsourcing of jobs to affect domestic labor demand should stimulate much additional research on this topic. A key factor hindering research has been lack of measurement. Unlike trade in goods, trade in labor services is at an extremely primitive stage of measurement. A first priority for U.S. statistical agencies should be extensive data collection to assess the extent of international outsourcing and to document the nature of tasks currently being outsourced. Yet trade in services will always be more difficult to capture and quantify than trade in goods. Just as productivity measurement has become more uncertain as U.S. economic activity has moved from a concentration on manufacturing to services, tracking trade flows will become increasingly challenging as trade in services takes its place alongside trade in goods as an increasingly important source of U.S. economic activity.

X. Conclusion

Viewed from the perspective of the 1980s, the rapid, monotone rise of wage inequality appeared to presage an era of ever-increasing demand for skills, with rising incomes for the highly-educated workers and falling incomes for everyone else. Fortunately, this vision has not yet come to pass. The secular demand increases favoring more educated workers appear to have been less rapid in the 1990s and early 2000s than from the 1960s to the 1980s. Overall wage inequality continued growing from 1990 to 2005, but at a slower pace than in the 1980s. Rather than spreading continuously, wage growth polarized after 1987, with persistent increases in inequality in the upper half of the income distribution and slow or reversing inequality trends in the lower-half of the distribution.

Demand-side forces have played a key role in shaping structural changes in U.S. wages during the inequality surge of the 1980s, and the polarization that followed. In the 1980s, during which wage growth was essentially monotone in terms of skills, employment shares in the highest-educated and highest-paid occupations expanded substantially, while employment shares in the lowest-skill occupations contracted. During the subsequent decade—in which earnings growth polarized—employment shares in very low- and very high-skill occupations increased, while employment shares in moderately skilled occupations contracted. The roughly parallel movement of earnings and employment growth in each decade suggests that demand-side forces have been central to these patterns of wage changes.

Following Autor, Levy, and Murnane (2003) and Goos and Manning (2007), we argue that these patterns may in part be explained by a richer version of the skill-biased technical change hypothesis, which posits that information technology complements highly educated workers engaged in abstract tasks, substitutes for moderately educated workers performing routine tasks, and has less impact on low-skilled workers performing manual tasks. Extrapolating from these trends, we forecast (perhaps unwisely) an ongoing growth of demand for both professional and managerial jobs requiring high levels of educational attainment, and for

low-skilled in-person service jobs—tasks that are difficult to either automate or outsource, but do not require more than a high school education.

Given slowing U.S. population growth and decelerating rates of educational attainment, it is natural for the United States to look to developing and developed countries as a source of supply for future employment growth. In the case of highly educated workers, we view such efforts as sound. Attracting skilled residents to the United States, either as students or workers, is likely to raise wealth and improve the quality of life for a large number of U.S. residents. As a secondary benefit, increased skilled migration to the United States may temper the ongoing rise of upper-tail earnings inequality. These same arguments appear less compelling when applied to the immigration of low-skilled workers. Wages of low-skilled U.S. workers have been stagnant for most of the past 30 years. If labor demand is indeed rising for low-skilled, in-person services occupations, this may give a long overdue boost to earnings for these groups—a welcome development for economic mobility and social cohesion in the United States.

Though it seems banal to end a research summary with a call for further research, this bromide seems less self-serving than usual in the current context. Due to rapid economic development in Asia and improvements in computer and communications technology, international trade and outsourcing appear poised to become important determinants of U.S. domestic labor demand. Yet we have little knowledge of the scope, magnitude, speed, or even direction with which these forces will impact skill demands and earnings distributions in the United States and in other advanced economies. Devising innovative and rigorous means to measure and evaluate the impacts of these evolving forces of globalization on inequality and economic well-being constitutes a significant agenda item for further research in this field.

■ *I am grateful to Jared Bernstein and Gary Burtless for insightful comments and suggestions. I am intellectually indebted to coauthors David Dorn, Lawrence Katz, Melissa Kearney, Alan Krueger, Frank Levy, and Richard Murnane for the main themes and conclusions of this paper.*

Notes

1. This observation was, to my knowledge, first offered by Mishel, Bernstein, and Boushey (2002).

2. The Current Population Survey and Census of Populations data analyzed here do not cover the top several percentiles of the earnings distribution where the most dramatic increases in real earnings have occurred during the last three decades (see Piketty and Saez 2003). Including these top percentiles would, consistent with our discussion, reveal even greater growth at the top throughout the years studied, but this inclusion would not qualitatively change our conclusions about income inequality.

3. It bears note, however, that all percentiles of the distribution fared better in the second half of the time period (1989 through 2005) than in the first half (1973 through 1989), reflecting the considerable acceleration of U.S. productivity growth from the mid-1990s forward.

4. These series are smoothed using three-year moving averages. Thus, the data point labeled 2004 is the average of the values for 2003, 2004, and 2005.

5. I use the term "we" throughout the paper because the material in this paper draws heavily on work I performed jointly with David Dorn, Frank Levy, Lawrence Katz, Melissa Kearney, Alan Krueger, and Richard Murnane.

6. Details of the samples and data processing methods used for these data series are provided in Autor, Katz, and Kearney (2008).

7. We do not discuss inequality of earnings residuals (that is, the unexplained component of wage variance). For recent work on this topic, see Lemieux (2006b) and Autor, Katz, and Kearney (2005 and 2008).

8. For this figure, we use the full time period of 1963 to 2005 (in contrast to Figures 4.1 through 4.3) because reliable measures of average earnings levels extending back to 1963 are available from the March Current Population Survey. By contrast, trends in earnings distribution (such as the 90/50 and 50/10) are more precisely measured using the CPS May/ORG data (Lemieux 2006b), which only extend back to 1973.

9. For previous implementations of such a model, see Katz and Murphy (1992); Autor, Katz, and Krueger (1998); Katz and Autor (1999); Card and Lemieux (2001); and Acemoglu (2002), among others.

10. We use a standard measure of college/non-college relative supply calculated in "efficiency units" to adjust for changes in labor force composition by gender and experience groups.

11. Skill-biased technological change refers to any introduction of a new technology, change in production methods, or change in the organization of work that increases the demand for more-skilled labor relative to less-skilled labor at fixed relative wages.

12. Foreign outsourcing of less-skilled jobs is another possible explanation for this pattern (Feenstra and Hanson 1999). But large within-industry shifts toward more skilled workers are pervasive even in sectors with little or no observed foreign outsourcing activity. Foreign outsourcing appears likely to become increasingly important, however.

13. Less restrictive variants of this model estimated in Autor, Katz, and Kearney (2008) also imply that trend demand growth for college relative to non-college workers slowed in the early 1990s. Autor, Katz, and Krueger (1998) and Card and DiNardo (2002) reach a similar conclusion.

14. In contrast to the findings of Autor, Katz, and Kearney (2008), analyses by Bartik (2001) and Bernstein and Baker (2003) find that a low unemployment rate differentially raises earnings in low relative to high wage deciles, thus compressing wage inequality. While a resolution of these conflicting conclusions is beyond the scope of this paper, this issue merits further study.

15. The direct effects of union decline on U.S. wage inequality growth also appear to be modest. Card, Lemieux, and Riddell (2003) find that falling unionization explains about 14 percent of the growth of male wage variance from 1973 to 2001 (in models allowing for skill group differences in the impact of unions), with an even smaller union effect for the growth of female wage variance.

16. For females, earnings growth between 1988 and 2005 among postcollege-educated workers was substantially greater than for college-only-educated workers, but the pattern was reversed for the 1979–1988 period.

17. Lemieux (2006a) documents the rapid, ongoing rise in the wage return to college and postcollege education. He estimates that more than two-thirds of the rise in wage inequality between 1973 and 2005 is explained by the growing return to postsecondary education.

18. Lemieux (2006b) focuses on the tight fit between the real minimum wage and residual wage variance for men and women from 1973 to 2003. We also find greater time series explanatory power of the real minimum wage for residual wage inequality measures than for actual wage inequality measures. This is puzzling for the minimum wage hypothesis, since the minimum wage should "bite" more for actual low wage workers than for residual low wage workers.

19. Lee (1999) also noted a puzzling relationship between the "effective" state minimum wage (the log difference between the state median and the state minimum) and upper-tail inequality. Opposite to the simple time-series regressions above, Lee finds in a cross-state analysis that increases in the effective state minimum wage appear to *reduce* upper-tail inequality, both for males and for the pooled-gender distribution. This result led him to advise caution in causally attributing trends in male and pooled-gender earnings inequality to the minimum wage.

20. In a similar vein, Acemoglu, Aghion, and Violante (2001) argue that the decline in union penetration in the United States and the United Kingdom is partly explained by changing skill demands that reduced the viability of rent-sharing bargains between high and low skill workers. Furthermore, the direct effects of union decline on U.S. wage inequality growth appear to be modest. Card, Lemieux, and Riddell (2003) find that falling unionization explains about 14 percent of the growth of male wage variance from 1973 to 2001 (in models allowing for skill group differences in the impact of unions), with an even smaller union effect for the growth of female wage variance.

21. A related earlier model along these lines is developed in Juhn (1994).

22. See Levy and Murnane (2004) for numerous paradigmatic examples. The fact that computerization causes manual tasks to grow as a share of labor input may be understood as a form of Baumol's disease—that is, the tendency for advanced economies to devote an ever-rising share of resources to labor-intensive sectors characterized by slow productivity growth, such as education and health care, while sectors with rapid productivity growth (such as manufacturing or farming) ultimately require fewer resources to meet consumer demand.

23. Welch (2000) and Weinberg (2000) argue that these technical changes are particularly likely to have been favorable to demand for female labor.

24. Acemoglu (1999) offers an alternative theory of job polarization based on endogenous changes in production techniques as a response to a rise in the availability of skilled labor.

25. The task intensity data are constructed by matching Census 1980 data by occupation and gender with task measures from the Dictionary of Occupational Titles (DOT). Task intensities by occupational skill percentile are plotted using a locally weighted smoothing regression with bandwidth 0.5 (meaning, one-half of one percentile). Details on the processing and matching of DOT task measures to occupations are given in Autor, Levy, and Murnane (2003). The abstract task category we use in Figure 4.10 is the arithmetic average of ALM's "non-routine cognitive/analytic" and "non-routine cognitive/interactive" category and, similarly, our routine task category is the average of ALM's "routine manual" and "routine cognitive" categories. Our manual category is equivalent to ALM's "non-routine manual" category.

26. Autor, Katz, and Kearney (2006) present a similar analysis using Census data for changes in occupational employment and CPS May/ORG data for changes in wage levels by earnings percentile. In the present analysis, we use exclusively Census data covering the same time periods.

27. We employ a consistent set of occupation codes developed by Meyers and Osborne (2005) for Census years 1980, 1990 and 2000. We use a locally weighted smoothing regression (bandwidth 0.8 with 100 observations) to fit the relationship between decadal growth in occupational employment share and occupations' initial skill percentile in the 1980 skill distribution.

28. In contrast to the upper panel of Figure 4.9, we use raw changes in employment shares by occupational wage percentile as the dependent variable, rather than smoothed changes. If instead we were to use smoothed changes, these would not affect the point estimates by much, but would suggest a misleadingly high degree of precision in the estimation.

29. Notably, this pattern appears inconsistent with the hypothesis that a declining minimum wage played a leading role in the expansion of lower-tail inequality in the 1980s. A decline in a binding wage floor should have lead to a (modest) rise in low-wage employment rather than a sharp contraction.

30. It is critical to distinguish service *occupations*, a relatively narrow group of low-educational level occupations comprising 13.4 percent of employment in 2000 (author's calculation from Census IPUMS), from the *service sector*, a very broad category of industries ranging from healthcare to communications to real estate, and comprising 81 percent of non-farm employment in 2000 (www.bls.gov).

31. The service employment measure used by the Bureau of Labor Statistics Occupational Outlook indicates a service employment share that is several percentage points higher that our calculations (17.7 percent versus 13.4 percent). The discrepancy stems from three factors: unlike our calculations based on household data from the Census Bureau, the BLS numbers use Current Employment Statistics (CES) which, as an establishment survey, double-counts workers who hold multiple jobs; our Census-based numbers are weighted by hours of labor supply, and so part-time jobs (common in service occupations) are weighted down whereas the CES data count all jobs equally. Our Census calculations exclude workers younger than 18 years of age, whereas the CES data include workers aged 16 years and above. The service occupation in which the Census and CES data are most divergent is in Food Preparation and Service, where our data show a 3.5 percent employment share and the CES data show a 7.4 percent employment share. Despite these discrepancies in levels, we have no reason to believe that the qualitative employment trends in the Census and CES data differ.

32. The BLS category of professional occupations excludes managerial occupations, and so is more disaggregated than the Census category of professional and managerial occupations. Combined growth in professional and managerial jobs is projected at 8.2 million jobs, or 18.8 percent.

33. Though computerization appears far more complementary to abstract tasks than non-routine manual tasks, our framework implies that computerization is a *relative* complement to all non-routine tasks (meaning, relative to routine tasks).

34. Though in many respects computerization and outsourcing appear to have similar implications for the domestic organization of work (Levy and Murnane 2006), one important difference is that there is an important subset of non-routine manual tasks that are not readily computerized but can be easily outsourced—for example, call center operations or back office manual tasks, including data entry and hand-processing of bill and check images (see Autor, Levy, and Murnane 2002). However, neither outsourcing nor computerization appears a close substitute for the in-persons tasks performed by service occupations (see Blinder 2007).

35. DeLong, Goldin, and Katz (2003) provide a thoughtful, extended discussion of policies to improve U.S. human capital investment.

36. This reflects both a high non-completion rate and an increased time to degree. Thus, the share of youth obtaining the equivalent of a four-year college degree by age 28 has risen significantly more than the share of youth obtaining the degree by 23.

37. See Heckman and Krueger (2004) for a comprehensive debate.

38. See also Heckman and Krueger (2004).

39. Denominators for these calculations come from U.S. Department of Education (2005, Table 170).

40. There is heated debate about the extent to which low-skilled immigration depresses native wages (see Borjas 2003; Card 2005; Goldin and Katz 2007). Recent evidence suggests that because the jobs of many low-skilled immigrants are heavily concentrated in "manual" tasks such as cleaning, cooking, and construction, they do not directly compete with most native-born workers, including low-skilled Americans who typically have a comparative advantage in English language communication tasks (Cortes 2007; Peri and Sparber 2007).

41. Acemoglu and Angrist (2001) provide a strong test of human capital externalities and find that they are weak or nonexistent.

42. Freeman (2006, p. 20) compellingly states this case: "If firms demand more labor than workers supply due to a reduced growth of supply, should not a country that relies extensively on unfettered markets allow those markets to raise the price of labor, just as it allowed them to reduce the pay of many in recent decades?"

43. Autor and Dorn (2007) report that the median hourly wage in service jobs was between 63 and 65 percent of the median hourly wage in non-service jobs in 1970, 1980, and 1990. Accounting for differences in full-compensation (including health insurance, vacation and sick leave) among high and low-wage workers (as in Pierce 2001) would enlarge this gap.

44. Recent research by Kopczuk, Saez, and Song (2007) finds little change in mobility over the course of a career among U.S. cohorts born between 1920 and 1950. However, these data do not speak to economic mobility across generations, in particular, how likely it is that children from low-income households reach higher echelons of the earnings distribution during their careers.

45. For example, the State Children's Health Insurance Program (SCHIP), enacted in 1997, has significantly increased the health insurance coverage rate of children from low-income households (Lo Sasso and Buchmueller 2004). SCHIP is a block grant program with fixed annual funding levels, however, and SCHIP outlays have not kept pace with population increase or the rising cost of healthcare. Absent a significant policy change, the number of program beneficiaries will have to decline. The U.S. minimum wage provides another example of a politically vulnerable policy instrument for raising earnings of low-skill workers.

46. See Hsieh and Woo (2005) for a rigorous assessment of the impact of outsourcing to China on the Hong Kong labor market.

References

Acemoglu, Daron. 1999. "Changes in Unemployment and Wage Inequality: An Alternative Theory and Some Evidence." *American Economic Review* 89(5): 1259–1278.

Acemoglu, Daron. 2002. "Technical Change, Inequality and the Labor Market." *Journal of Economic Literature* 40(1): 7–72.

Acemoglu, Daron, and Joshua D. Angrist. 2001. "How Large Are Human-Capital Externalities? Evidence from Compulsory Schooling Laws." In *NBER Macroeconomics Annual 2000*, ed. Ben S. Bernanke and Kenneth Rogoff, 9–59. Cambridge, MA: The MIT Press.

Acemoglu, Daron, Philippe Aghion, and Giovanni L. Violante. 2001. "Deunionization, Technical Change and Inequality." *Carnegie-Rochester Conference Series on Public Policy* 55: 229–64.

Anderson, Michael L. 2007. "Uncovering Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects." Mimeograph, University of California Berkeley.

Antràs, Pol, Luis Garicano, and Esteban Rossi-Hansberg. 2006a. "Offshoring in a Knowledge Economy." *Quarterly Journal of Economics* 121(1): 31–77.

Antràs, Pol, Luis Garicano, and Esteban Rossi-Hansberg. 2006b. "Organization and Inequality in a Knowledge Economy." *Quarterly Journal of Economics* 121(4): 1383–1435.

Autor, David H., and David Dorn. 2007. "Inequality and Specialization: The Growth of Low-Skill Service Jobs in the United States." Mimeograph, Massachusetts Institute of Technology.

Autor, David H., Lawrence F. Katz, and Melissa S. Kearney. 2005. "Residual Wage Inequality: The Role of Composition and Prices." NBER Working Paper No. 11628.

Autor, David H., Lawrence F. Katz, and Melissa S. Kearney. 2006. "The Polarization of the U.S. Labor Market." *American Economic Review*, 96(2): 189–194.

Autor, David H., Lawrence F. Katz, and Melissa S. Kearney. 2008 (forthcoming). "Trends in U.S. Wage Inequality: Revising the Revisionists." *Review of Economics and Statistics*.

Autor, David H., Lawrence F. Katz, and Alan B. Krueger. 1998. "Computing Inequality: Have Computers Changed the Labor Market?" *Quarterly Journal of Economics* 113(4): 1169–1213.

Autor, David H., Frank Levy, and Richard J. Murnane. 2002. "Upstairs Downstairs: Computers and Skills on Two Floors of a Large Bank." *Industrial and Labor Relations Review* 55(3): 432–447.

Autor, David H., Frank Levy, and Richard J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Investigation." *Quarterly Journal of Economics* 118(4): 1279–1333.

Bartel, Ann P., Casey Ichniowski, and Kathryn L. Shaw. 2007. "How Does Information Technology Really Affect Productivity? Plant-Level Comparisons of Product Innovation, Process Improvement and Worker Skills." *Quarterly Journal of Economics* 122(4): 1721–1758.

Bartik, Timothy J. 2001. *Jobs for the Poor: Can Labor Demand Policies Help?* New York: Russell Sage Foundation.

Berman, Eli, John Bound, and Stephen Machin. 1998. "Implications of Skill-Biased Technological Change: International Evidence." *Quarterly Journal of Economics* 113(4): 1245–1279.

Bernstein, Jared, and Dean Baker. 2003. *The Benefits of Full Employment: When Markets Work for People*. Washington, DC: Economic Policy Institute.

Bishop, John H., and Shani Carter. 1991. "How Accurate are Recent BLS Occupational Projections?" *Monthly Labor Review* 114(10): 37–43.

Blinder, Alan. 2007. "How Many U.S. Jobs Might Be Offshorable?" CEPS Working Paper No. 142, Princeton University.

Borjas, George J. 2003. "The Labor Demand Curve Is Downward Sloping: Reexamining the Impact of Immigration on the Labor Market." *Quarterly Journal of Economics* 118(4): 1335–1374.

Borjas, George J. 2006. "Immigration in High-Skill Labor Markets: The Impact of Foreign Students on the Earnings of Doctorates." NBER Working Paper No. 12085.

Bound, John, Michael F. Lovenheim, and Sarah Turner. 2007. "Understanding the Increased Time to the Baccalaureate Degree." Mimeograph, University of Virginia, May.

Broache, Anne. 2007. "Tech Seeing Red over Green-Card Proposal." *CNET news.com*, June 6, 2007. Available online at http://news.cnet.com/.

Card, David. 2001. "Estimating the Return to Schooling: Progress on Some Persistent Econometric Problems." *Econometrica* 69(5): 1127–1160.

Card, David. 2005. "Is the New Immigration Really So Bad?" NBER Working Paper No. 11547.

Card, David, and John E. DiNardo. 2002. "Skill-Biased Technological Change and Rising Wage Inequality: Some Problems and Puzzles." *Journal of Labor Economics* 20(4): 733–783.

Card, David, and Thomas Lemieux. 2001. "Can Falling Supply Explain the Rising Return to College for Younger Men?" *Quarterly Journal of Economics* 116(2): 705–746.

Card, David, Thomas Lemieux, and W. Craig Riddell. 2003. "Unionization and Wage Inequality: A Comparative Study of the U.S., the U.K., and Canada." NBER Working Paper No. 9473.

Cortes, Patricia. 2006. "The Effect of Low-Skilled Immigration on US Prices: Evidence from CPI Data." Mimeograph, University of Chicago Graduate School of Business.

Currie, Janet. 2001. "Early Childhood Intervention Programs: What Do We Know?" *Journal of Economic Perspectives* 15(2), 212–238.

DeLong, J. Bradford, Claudia Goldin, and Lawrence F. Katz. 2003. "Sustaining U.S. Economic Growth." In *Agenda for the Nation*, ed. Henry J. Aaron, James M. Lindsay, and Pietro S. Nivola, 17–60. Washington, DC: The Brookings Institution Press.

DiNardo, John, and David S. Lee. 2004. "Economic Impacts of New Unionization on Private Sector Employers: 1984-2001." *Quarterly Journal of Economics* 119(4): 1383–1441.

Doms, Mark, Timothy Dunne, and Kenneth R. Troske. 1997. "Workers, Wages, and Technology." *Quarterly Journal of Economics* 112(1): 253–290.

Dunne, Timothy, John Haltiwanger, and Kenneth R. Troske. 1997. "Technology and Jobs: Secular Changes and Cyclical Dynamics." *Carnegie-Rochester Conference Series on Public Policy* 46: 107–178.

Dustmann, Christian, Johannes Ludsteck, and Uta Schönberg. 2007. "Revisiting the German Wage Structure." Mimeograph, University of Rochester, February.

Dynarski, Susan. 2002. "The Behavioral and Distributional Implications of Aid for College." *American Economic Review* 92(2): 279–285.

Ellwood, David T., and Thomas Kane. 2000. "Who is Getting a College Education? Family Background and the Growing Gaps in Enrollment." In *Securing the Future*, ed. S. Danziger and J. Waldfogel, 47–83. New York: Russell Sage Foundation.

Feenstra, Robert C., and Gordon H. Hanson. 1999. "The Impact of Outsourcing and High-Technology Capital on Wages: Estimates for the United States, 1979–1990." *Quarterly Journal of Economics* 114(3): 907–940.

Freeman, Richard B. 2006. "Is A Great Labor Shortage Coming? Replacement Demand in the Global Economy." NBER Working Paper No. 12541.

Goldin, Claudia, and Lawrence F. Katz. 2007. "The Race between Education and Technology: The Evolution of U.S. Educational Wage Differentials, 1890 to 2005." NBER Working Paper No. 12984.

Goos, Maarten, and Alan Manning. 2007. "Lousy and Lovely Jobs: The Rising Polarization of Work in Britain." *Review of Economics and Statistics* 89(1): 118–133.

Grossman, Gene M., and Esteban Rossi-Hansberg. Forthcoming. "Trading Tasks: A Simple Theory of Offshoring." *American Economic Review*.

Heckman, James J., and Pedro Carneiro. 2002. "The Evidence on Credit Constraints in Post-Secondary Schooling." *Economic Journal* 112(482): 705–734.

Heckman James J., and Alan B. Krueger. 2004. *Inequality in America: What Role for Human Capital Policies*? Cambridge, MA: The MIT Press.

Heckman, James J., and Lance Lochner. 2000. "Rethinking Education and Training Policy: Understanding the Sources of Skill Formation in a Modern Economy." In *Securing the Future*, ed. S. Danziger and J. Waldfogel, 47–83. New York: Russell Sage Foundation.

Hsieh, Chang-Tai, and Keong T. Woo. 2005. "The Impact of Outsourcing to China on Hong Kong's Labor Market." *American Economic Review* 95(5): 1673–1687.

Juhn, Chinhui. 1994. "Wage Inequality and Industrial Change: Evidence from Five Decades." NBER Working Paper No. 4684.

Kane, Thomas J. 1999. *The Price of Admission*. Washington, DC: Brookings Institution Press and Russell Sage.

Katz, Lawrence F. 2000. "Technological Change, Computerization, and the Wage Structure." In *Understanding the Digital Economy*, ed. E. Brynjolfsson and B. Kahin, 217–244. Cambridge, MA: The MIT Press.

Katz, Lawrence F., and David H. Autor. 1999. "Changes in the Wage Structure and Earnings Inequality." In *Handbook of Labor Economics*, Volume 3, ed. O. Ashenfelter and D. Card, 1463–1555. Amsterdam: Elsevier-North Holland.

Katz, Lawrence F., and Kevin M. Murphy. 1992. "Changes in Relative Wages, 1963–87: Supply and Demand Factors." *Quarterly Journal of Economics* 107(1): 35–78.

Kleiner, Morris M. 2005. *Licensing Occupations: Ensuring Quality or Restricting Competition?* Kalamazoo, MI: W.E. Upjohn Institute.

Kletzer, Lori G. 2006. "The Scope of Tradable Services and the Task Content of Offshorable Services Jobs." Mimeograph, University of California Santa Cruz.

Kopczuk, Wojciech, Emmanuel Saez, and Jae Song. 2007. "Uncovering the American Dream: Inequality and Mobility in Social Security Earnings Data since 1937." Mimeograph, University of California Berkeley.

Lee, David S. 1999. "Wage Inequality in the U.S. during the 1980s: Rising Dispersion or Falling Minimum Wage?" *Quarterly Journal of Economics* 114(3): 977–1023.

Lemieux, Thomas. 2006a."Post-Secondary Education and Increasing Wage Inequality." NBER Working Paper No. 12077.

Lemieux, Thomas. 2006b. "Increased Residual Wage Inequality: Composition Effects, Noisy Data, or Rising Demand for Skill." *American Economic Review* 96(3): 461–498.

Levy, Frank, and Richard J. Murnane. 2004. *The New Division of Labor*. Princeton, NJ: Princeton University Press.

Levy, Frank, and Richard J. Murnane. 2006. "How Computerized Work and Globalization Shape Human Skill Demands." Mimeograph, Massachusetts Institute of Technology.

Lo Sasso, Anthony T., and Thomas C. Buchmueller. 2004. "The Effect of the State Children's Health Insurance Program on Health Insurance Coverage." *Journal of Health Economics* 23(5): 1059–1082.

Machin, Stephen, and John Van Reenen. 1998. "Technology and Changes in Skill Structure: Evidence from Seven OECD Countries." *Quarterly Journal of Economics* 113(4): 1215–1244.

Mazzolari, Francesca, and Giuseppe Ragusa. 2007. "Spillovers from High-Skill Consumption to Low-Skill Labor Markets."IZA Discussion Paper 3048. Bonn: Institute for the Study of Labor.

Meyer, Bruce D., and Dan T. Rosenbaum. 2001. "Welfare, the Earned Income Tax Credit, and the Labor Supply of Single Mothers." *Quarterly Journal of Economics* 116(3): 1063–1114.

Meyers, Peter B., and Anastasiya M. Osborne. 2005. "Proposed Category System for 1960–2000 Census Occupations." U.S. Bureau of Labor Statistics Working Paper No. 383.

Mishel, Lawrence, Jared Bernstein, and Heather Boushey. 2002. *The State of Working America: 2002–03*. Ithaca, NY: ILR Press.

Moretti, Enrico. 2004a. "Estimating the Social Return to Higher Education: Evidence from Longitudinal and Repeated Cross-Sectional Data." *Journal of Econometrics* 121(1–2): 175–212.

Moretti, Enrico. 2004b. "Workers' Education, Spillovers, and Productivity: Evidence from Plant-Level Production Functions." *American Economic Review* 94(3): 656–690.

Organization for Economic Cooperation and Development. 2006. *Education at a Glance: OECD Indicators, 2006*. Paris: OECD.

Peri, Giovanni, and Chad Sparber. 2007. "Comparative Advantage and the Gains from Immigration." Mimeograph, University of California Davis.

Pierce, Brooks. 2001. "Compensation Inequality." *Quarterly Journal of Economics* 116: 1493–1525.

Piketty, Thomas, and Emmanuel Saez. 2003. "Income Inequality in the United States, 1913–1998." *Quarterly Journal of Economics* 118(1): 1–39.

Rothstein, Jesse, and Cecilia Elena Rouse. 2007. "Constrained After College: Student Loans and Early Career Occupational Choices." NBER Working Paper No. 13117.

Spitz-Oener, Alexandra. 2006. "Technical Change, Job Tasks and Rising Educational Demands: Looking Outside the Wage Structure." *Journal of Labor Economics* 24(2): 235–270.

U.S. Department of Education, National Center for Education Statistics. 2005. "Digest of Education Statistics." Available online at http://nces.ed.gov/.

U.S. Department of Education, National Center for Education Statistics. 2007. "The Condition of Education." Available online at http://nces.ed.gov/programs/coe/.

U.S. Congressional Budget Office. 2004. "What Accounts for the Decline in Manufacturing Employment?" Available online at http://www.cbo.gov/ftpdoc.cfm?index=5078&type=0.

Weinberg, Bruce. 2000. "Computer Use and the Demand for Female Workers." *Industrial and Labor Relations Review* 53: 209–308.

Welch, Finis. 2000. "Growth in Women's Relative Wages and in Inequality Among Men: One Phenomenon or Two?" *American Economic Review* 90(2): 444–449.

Comments on "Structural Demand Shifts and Potential Labor Supply Responses in the New Century" by David H. Autor

Jared Bernstein

Introduction

David Autor has written an insightful and admirable paper, which in important ways updates economists' thinking about one of the more important questions in labor economics: what is the relationship between skill demands, technological change, and wage inequality?

I'll stress the insights from the paper below, but what's admirable about its contribution? Well, as J.M. Keynes famously said when confronted with accusations regarding shifts in his analysis: "When my information changes, I change my opinion. What do you do, sir?" A healthy debate regarding the impact of skill-biased technological change has simmered among economists for a few years now, and some of us believed that the trends in the wage data did not support conventional wisdom. With this paper, building on his earlier work, Autor agrees with this contrarian assessment, and his more nuanced view of the issue is both gratifying and interesting.[1]

Below, I offer some objections to the plasticity with which Autor treats the concept of skill-biased technological change. One has the sense that he and his fellow authors remained wedded to this concept and loathe to let it go, so they've morphed the definition in ways that may strike some as stretching the concept to the breaking point. Also, in the course of this review, I question a) whether the original construction of skill-biased technological change provided an adequate description of the trends observed in wage inequality, and b) whether Autor's new interpretation is sufficient to explain more recent relative wage trends, such as those that prevailed in the latter 1990s or over the 2000s. But for those of us who

have been skeptical of skill-biased technological change as the dominant explanatory factor in the growth of wage inequality in recent decades, Autor's evolving views are a welcome point of departure.

Motivation: That Pesky Trend in the 50/10 Wage Percentiles

As Autor shows in his Figure 4.2, the character of hourly wage inequality in the United States has changed over time. In the 1970s and 1980s, the wage distribution was fanning out largely monotonically: as your wage level went up, so did your wage growth. The top income groups pulled away from the middle; the middle income groups pulled away from the groups at the bottom of the wage distribution. Over the 1990s, however, growth rates were much more comparable at the middle and the bottom of the distribution. In fact, depending on which end points you choose, there was even some compression of the 50/10 ratio since the mid-to-late 1980s, especially for men.

Larry Mishel and I viewed this change as quite important, and thus focus on it in various editions of *The State of Working America*, in part because the compression of the 50/10 ratio raises questions about the economist's most popular answer for the "why" of growing wage inequality: skill-biased technological change. Among others (see Howell and Wolff 1992, Handel 2000, and Card and Dinardo 2002), we wondered why the skill demands of the workplace in the 1990s, a period of very deep, if not downright bubbly, technological capital deepening, would be more complementary to low-wage workers than had been the case in earlier years.

It seemed to us then, and still does now, that if skill-biased technological change were the main perpetrator of higher wage inequality, relative wage trends should look more like Autor's Figure 4.2, bottom line, than this figure's top line. Autor, in an earlier paper written with Levy and Murnane (2001), analyzes occupations from the perspective of skill content and tasks, and asks whether those tasks are complementary or substitutable to computer technology. In that paper, they raise the hypothesis that the skill content of recent technological change, particularly regarding the impact of computers on skill demands, may not generate the monotonic trends in relative wages that a basic skill-biased technological change hypothesis predicts.

With this paper, as with Autor, Katz, and Kearney (2006), Autor further develops a different version of skill-biased technological change, which is, as he puts it, "a richer version of the skill-biased technical change hypothesis, which posits that information technology complements highly educated workers engaged in abstract tasks, substitutes for moderately educated workers performing routine tasks, and has less impact on low-skilled workers performing manual tasks."

Two Questions

Autor's new view raises two questions. First, does the theory that Autor and his co-authors have been developing comport with the data in a lasting manner? That is, given a) his earlier research findings suggesting nonmonotonic demand impacts, and b) the fact that the relative wage data trends failed to support the old view, Autor came up with an interesting and plausible explanation that fits the part of the data for which the old view did not. But if the new view fails to fit the further evolution of wage inequality, Autor and his coauthors will be forced to invent an even "richer version of the skill-biased technological change hypothesis." Going down that path can lead to "ad hockery" and, while spinning of lots of interesting analysis, an enriched explanation is unlikely to provide analysts and policymakers with a reliable sense as to the determinants of wage inequality's growth.

If that does indeed turn out to be the case, as I suspect it might (supporting evidence to follow), many would be compelled to conclude that skill-biased technological change, an obviously important theory in ways I describe below, is insufficient to explain the path of wage inequality. This is the second big question that this paper raises: is skill-biased technological change the right horse to bet on in the race to explain the evolution of wage inequality in our economy, or for that matter, any economy?

Let us first examine this second question, regarding the usefulness of skill-biased technological change as a framework for explaining and understanding wage inequality. All economists recognize that technology tends to boost labor demand for those workers whose skills are complementary to that technology (and, conversely, reduce demand for the technology's substitutes). And the long-term reality of capital skill complementarity is also an obvious feature of our labor market. It is the

main reason why we can double the share of four-year college graduates, as we have done over the past 30 years, yet maintain their unemployment rate at frictional levels (typically at around 2 percent, per year).

But whether technology and human skills are complements is not the question being asked in this literature. The question asked is whether the extent of this complementarity has increased in a manner that would explain the pattern of wage inequality observed over the past few decades. As Mishel and I view it in a series of papers from the latter 1990s (some co-authored with John Schmitt), the key question was whether skill-biased technological change had accelerated to such a degree that it could explain the acceleration in wage inequality.[2]

In statistical analysis, we tested for the acceleration of skill-biased technological change by regressing changes in wage inequality across industries on measures of capital investment associated with technology, and allowing the coefficients on those variables to change over time. We found no consistent evidence to support the notion that skill-biased technological change and its impact on wage inequality had accelerated.

Interestingly, we found evidence against accelerating skill demands in work by economists closely tied to the skill-biased technological change story. Katz, in various papers [most recent in work with Goldin (2008)], uses a simple CES production function model of the labor market to create indexes of labor supply and demand by skill level, stretching back many decades. Note that these models examine just one dimension of wage inequality: that existing between college-educated and non-college-educated workers. So-called residual wage inequality—the part not explained by the usual regressors in human capital models—is not part of this analysis (analysis like that in Autor's Figure 4.2 captures both residual and "between group" aspects of wage inequality).

The Katz/Goldin model assumes that changes in relative wages (skilled relative to unskilled workers, or college-educated to noncollege-educated workers) are a function of shifting demand and supplies of different types of labor, the degree to which relative supply changes effect relative wages (the substitution elasticity),[3] and technological progress.

For our purposes in evaluating the utility of skill-biased technological change and the case for accelerating wage inequality, the relevant outputs from the model are changes in relative demand for skills across decades.

If unmet skill demands are accelerating, meaning these are unmet by increasing skill supplies, the result will be accelerating relative wages, the pace of which is partly determined in the model by the degree to which relative skill-changes map onto relative wages through the substitution elasticity. All else equal, a larger substitution elasticity will diminish the growth of relative wages, because employers can more easily substitute away from more expensive workers.

For points I turn to next, it is important to recognize that the demand index is a residual. This is important for our alternative story, since we and others argue that there's more than demand embedded in that term. In fact, any nonsupply factor—tighter job markets or lower minimum wages—that affects relative wages gets subsumed under demand here, a critique we return to below. The top panel in Table 4.3 just reprints Goldin and Katz's table, showing that that relative wages grew most quickly in the 1980s, in fact more quickly than in any decade since 1940. The second column, relative supply growth, shows the 1970s was a decade when the share of college-educated workers grew quickly, leading to a compression in the relative wage. Since then, relative skill supplies have increased more slowly, especially in the 2000s.

But our focus is on the demand column, shown in panel B. Some authors simply interpret the positive values in this column as evidence of skill-biased technological change driving up relative demand, and thus see this as the key determinant of the between-education group part of wage inequality. But our model argues that unless these decadal demand indexes are accelerating, these are insufficient to explain the acceleration in wage inequality over these years. It is the second derivative that matters, not the first.

And, in fact, the demand index shows quite a significant deceleration over much—but not all—of the last few decades. This deceleration seems really quite revealing. According to this simple but plausible and widely accepted model of the labor market, the rate of change in relative demand for skills was slower in the 1990s than the 1980s, and slower still in the 2000s, at least through 2005.

In fact, compared to the 1980s, demand for skills in the 2000s (through 2005) grew 3.6 log points per year more slowly. This decline seems particularly remarkable when we consider the dissemination of computer-

Table 4.3
The College Wage Premium, Supply and Demand: 1940–2005 (Annual Percent Changes)

Period	Relative Wage	Relative Supply	Relative Demand
1940–50	–1.86	2.35	–0.69
1950–60	0.83	2.91	4.28
1960–70	0.69	2.55	3.69
1970–80	–0.74	4.99	3.77
1980–90	1.51	2.53	5.01
1990–2000	0.58	2.03	2.98
2000–05	0.34	0.89	1.42
Panel B: Deceleration in Relative Demand			
1960s over 50s	–0.59		
1970s over 60s	0.08		
1980s over 70s	1.24		
1990s over 80s	–2.03		
2000s over 90s	–1.56		

Source: Author's analysis of Goldin and Katz (2007), Table 1, using their preferred subsition elasticity of 1.64 (implying that a 10 percent increase in the relative supply of college graduates lowers the college premium by 6.1 percent (1/.61 = 1.64)).

ization and information technology in general since the 1980s. Surely, the spread of information technology has accelerated. The literature explaining the post-1995 acceleration in productivity growth is, in fact, quite clear on this point: the spread of computers and information technology explains most of productivity's acceleration.[4] In this climate, if skill-biased technological change were truly a critical determinant of relative wage trends by education level, would we not expect to see a sharp acceleration of relative demand? Instead, the model yields the opposite finding.

Other authors, such as Handel, Howell, and Wolfe (see various citations in references) had similar findings, noting that the timing of information technology's capital deepening did not match the changes in wage

inequality such that skill-biased technological change would be a likely determinant. Card and DiNardo (2002) derived various tests of the skill-biased technological change (SBTC) hypothesis, and they too found little supporting evidence: "we conclude that the SBTC hypothesis is not very helpful in understanding the myriad shifts in the structure of wages that have occurred over the past three decades."

So if skill-biased technological change is not the reason, what has been the main driver of wage inequality over these years? Our work has found that there is no smoking gun, no single factor that explains more than half of the growth in wage inequality. Instead, there appear to be many perpetrators, including high unemployment, the sharp expansion of unbalanced international trade, the decline in the real value of the minimum wage, and the decline in unions. Many of these factors fit within an "institutional" framework, and are very compellingly discussed in Levy and Temin's recent paper (2007), an analysis all the more fascinating since Levy's arguments used to be squarely in the skill-biased technological change camp. Along with Temin, Levy now writes:

> [T]he current trend toward greater inequality in America is primarily the result of a change in economic policy that took place in the late 1970s and early 1980s. The stability in income equality where wages rose with national productivity for a generation after the Second World War was the result of policies that began in the Great Depression with the New Deal and were amplified by both public and private actions after the war. This stability was not the result of a natural economy; it was the result of policies designed to promote it (41–42).

As Levy and Temin stress, it's not the case that skill-biased technogical change is not important, it's that a) wage differentials are often moved by changes in institutional arrangements favoring one class of worker over another, and b) skill-biased technological change has been a fairly smooth, ongoing dynamic in our economy, consistently driving up employers' skill demands. But the institutional context within which skill-biased technological change occurs is highly determinative of relative wage outcomes.

In the interest of honest, full disclosure, there's no smoking gun for the institutional explanations either. Solid research casts doubt on skill-biased technological change and much work, including our analysis in many editions of *The State of Working America*, confirms the important influences of institutional forces on both absolute and relative wages. But

Autor and his co-authors would be well within their rights to point out that our evidence is also limited. We can point to various studies showing that declining union membership explains 10–20 percent of the increase in wage inequality, with international trade and the impact of the minimum wage each also contributing about the same percentage. But, as Autor notes (see his note 14) there is not enough convincing work on these causes either.[5]

In the next section, I explore the first question raised above: is Autor's newer model likely to offer a more reliable interpretation of the factors behind rising wage inequality?

As in Autor, Katz, and Kearney (2006), Autor's analysis is based on the shifts in employment by occupation, along with the educational composition of the occupations over time. Occupations constitute a legitimate source of variation in this type of work, since it is widely assumed that occupational demands necessarily embody technological change.

Autor's key figure supporting this analysis is Figure 4.9, panel A, which contrasts the monotonic pattern of occupation skill demands of the 1980s with the more U-shaped pattern of the 1990s. The significant accomplishment here bears remark: Autor has crafted a model which explains the changing pattern of wage inequality over the 1980–2000 period. Whether it is legitimate to call it skill-biased technological change is another matter, as I discussed above and elaborate upon in the conclusion. But this is laudable, interesting work that advances our understanding of the labor market.

But is it built to last? Is Autor's model telling us something important about the causal factors behind wage inequality—providing actionable intelligence, as the saying goes—that we can count on continuing in a way that should inform policy? Or is this model simply mapping relative wage trends onto occupations in a way that will change again when the relative wage trends shift again?

To test this, I derived a simple (too simple?), "poor man's AKK" (that's Autor, Katz, Kearney, where this model first appeared) model of occupational employment, weighted by hours and wages.[6] Sophistication-wise, my method is a mere shadow of Autor's, but my results for the 1980s and 1990s roughly match his; see Figure 4.11 for my 1980s version, which, due to a coding change, only goes from 1984–1989. My Figure 4.12

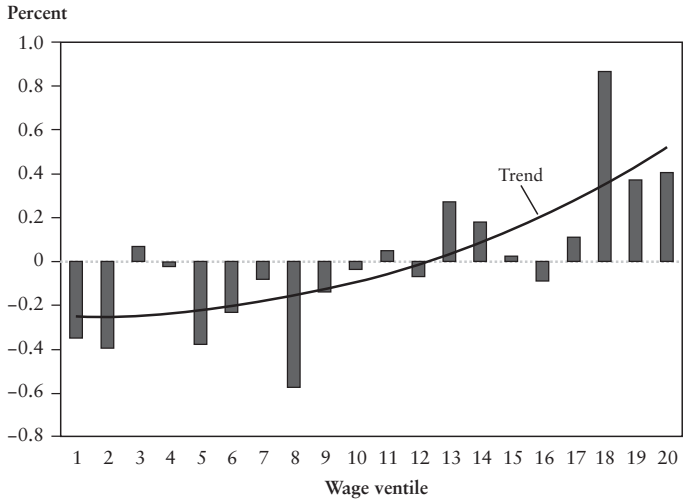

Figure 4.11
Change in Occupational Employment Shares by Wage Ventile, 1984–1989
Source: Current Population Survey.
Note: Calculation method loosely based on Autor, Katz, and Kearney (2006).

roughly reflects the polarization the Autor finds, but my results only find evidence for this in the years 1989–1995. That's because much work at Economic Policy Institute has found that the latter 1990s were a very unique period for wage inequality, as the first full employment job market in decades gave a lift to U.S. workers' bargaining power.

Given that important dynamic, I plotted the latter 1990s separately in Figure 4.13, and these seem to revert back to the earlier, more monotonic pattern. This is quite an interesting finding, in that we know relative wages were compressing between middle- and low-wage workers over this period, yet the model shows a pattern of relative demand much like that of the 1980s.

Finally, I plot the 2000s through 2006 in Figure 4.14. This decade (so far) takes on more polarizating characteristics, although one could also view this change as showing weak or negative skill demands throughout most of the skill distribution, except at the very top, possibly due to the uniquely weak job creation that prevailed over this period.

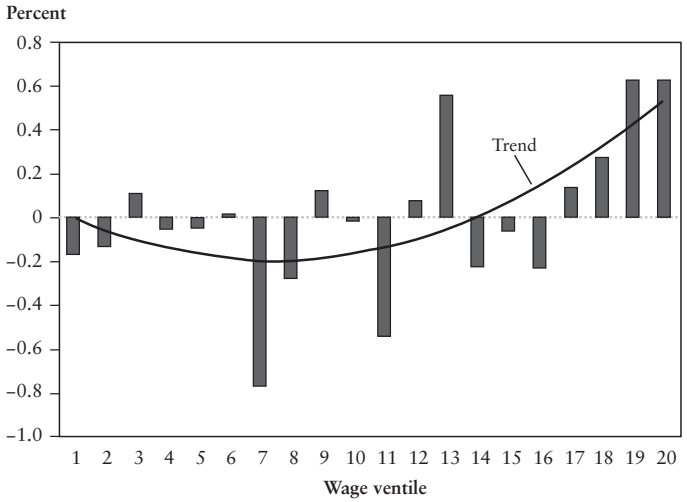

Figure 4.12
Change in Occupational Employment Shares by Wage Ventile, 1989–1995
Source: Current Population Survey.
Note: Calculation method loosely based on Autor, Katz, and Kearney (2006).

Based on the differences in our methods, this exercise by no means establishes a claim that Autor is wrong. But if his model's method really yields quite different results regarding the nature of skill demands over relatively short time periods, we need to wonder what it is really telling us. If the full employment of the latter 1990s raised the bargaining power of less-skilled workers enough to move these curves around, as my work suggests, then we do not want to conflate that change with a shift in the nature of skill-biased technological change.

I thus encourage Autor to apply his method to these other time periods, including the 2000s, to see if and how the results change over different time periods.

A second and final critique of Autor's method relates to the underlying elasticities implied by the work. The movements in the occupational index in a given decade are actually very small compared to the movements in relative wages. The 50/10 ratio falls by about .10 percentage points over the 1990–2000 period. Autor's Figure 4.9 reveals that the demand at the median fell about 0.0005 points, while demand at the low

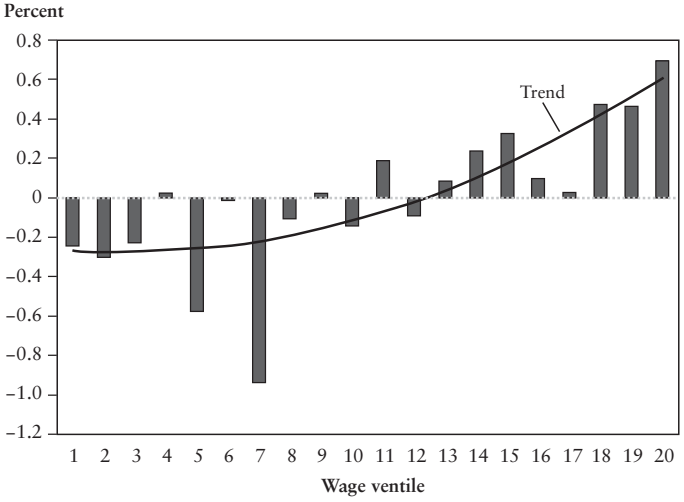

Figure 4.13
Change in Occupational Employment Shares by Wage Ventile, 1995–2000
Source: Current Population Survey.
Note: Calculation method loosely based on Autor, Katz, and Kearney (2006).

end rose about the same absolute amount. It seems unlikely that such a small shift in the occupational shares could really be the major factor driving relative wage trends of this magnitude.

As Autor has pointed out to me in a private conversation, the Autor, Katz, and Kearney model does not measure underlying shifts in supply and demand, along with relevant elasticities, as does the Katz model featured in Table 1. Instead, the model by Autor, Katz, and Kearney is driven by, in Autor's words, "realized prices and quantities," meaning the occupational and wage outcomes of an implicit labor demand and supply model that is a lot more complicated than the simple, two-skill Katz version. Underlying this implicit model is a rich set of elasticities regarding the ease with which employers can substitute different types of workers over a broad range of tasks—tasks which, unlike levels of educational attainment, are not necessarily exogenous.[7]

All of which leads one to conclude that Autor, Katz, and Kearney's model may be telling a small part of the story of shifts in relative skill/labor demand. Other factors must be playing an important role as well,

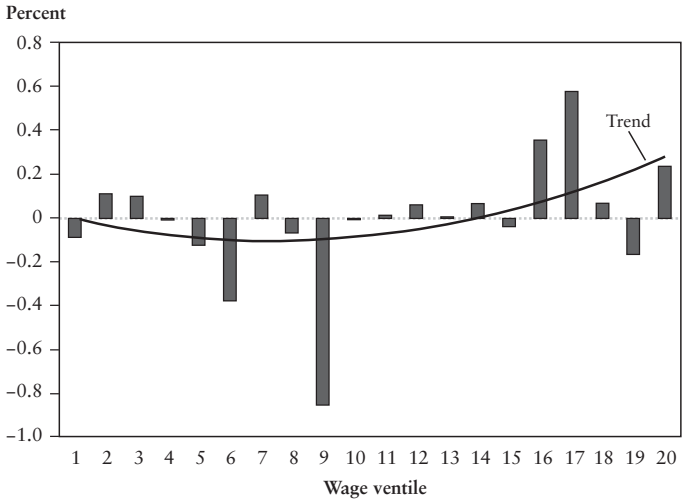

Figure 4.14
Change in Occupational Employment Shares by Wage Ventile, 2000–2006
Source: Current Population Survey.
Note: Calculation method loosely based on Autor, Katz, and Kearney (2006).

a point that I suspect Autor, an economist with an open mind and a demonstrable interest in a wide range of causes, would agree with.

Conclusion

Though I have not pulled any punches in expressing the difficulty that I and others have had in assigning a primary role to skill-biased technological change in the wage inequality debate, I still find that this paper represents a useful contribution to that debate. Autor does what good economists should do: when the old model stops working, you build a new one. One can raise concerns, as I do, as to whether the new one has any more explanatory power than the old one, outside of a specific time period, perhaps a relatively short period—he says the 1990s; I say the early 1990s—wherein it fits the data. In this regard, I hope and expect Autor to keep testing whether his polarization model continues to fit the trends in relative wages.

One final policy point: I don't think Autor (or Autor, Katz, and Kearney, for that matter) should refer to the pattern of relative demand they discover in the 1990s as skill-biased technological change. For that label to apply, technology must be monotonically biased toward skilled workers and against unskilled ones. In their new co-authored work, skill is biased toward the high end, and—perhaps to a lesser extent, but still—toward the low end.

Words matter. And the hurly-burly world of Washington, DC economic policy—the world I travel in—sustains little in the way of subtle discourse. Among DC policy makers, skill-biased technological change translates into the notion that employers' skill demands continue to shift hard against non-college-educated workers. As they envision it, the shift is absolutely monotonic; there's little room, I fear, for the notion that "skill bias" as Autor now understands it doesn't really mean skill bias as they understand it.

Autor could explain that the concept of skill-biased technological change now means a bias in favor of high- and low-end workers and a bias against middle-wage workers. That's what he's trying to establish in this paper, and he makes the case. But why insist on calling it skill bias? The policy implications are significant. If you believe in the traditional skill-biased technological change story, you're prone to think exclusively in terms of education and job training. If you instead think in terms of Autor's concept of skill-biased technological change, along with those critical skill-enhancement policies, you also recognize that demand for less-skilled jobs is strong and will remain so. And that leads economic policymakers to worry about the quality of the jobs available to the American workforce, not just the quality of the labor supply.

As I noted, Keynes would have liked this paper—a very fine compliment indeed. I urge Autor and his co-authors to continue testing this model as new trends form to see if it consistently describes the evolving patterns in relative wages. I also urge them to keep looking for other explanations that may have less to do with skill biases and technology, and more to do with bargaining power, full employment, and institutional forces within which a full and complete concept of skill-biased technological change plays out.

Notes

1. These new views on the impact of skill-biased technological change on wage inequality appeared in an earlier paper of which Autor was a co-author. See Autor, Katz, and Kearney (2006).

2. See, for example, Mishel and Bernstein (1998).

3. That is, how would an increase in the supply of less-skilled workers affect the demand for more-skilled workers? If substitution between these two groups is a simple matter for employers, then the falling relative price of less-skilled workers would induce employers to use more of them. But if skill requirements are such that employers cannot easily substitute less-skilled for more skilled workers, then the decline in relative price will have less impact of changes in employment.

4. Oliner and Sichel (2000).

5. Autor, Katz, and Kearney are very much aware of these institutional forces and test the impacts of unemployment and minimum wages. As Autor notes (note 14), they don't find much, but their analysis is fairly cursory, as they include both the minimum wage and the national unemployment rate (male, prime age) in a reduced-form, time series model with relative wages by education as the dependent variable. More detailed work, such as that cited by Autor in the aforementioned note, taps both geographical variation and examines the impact of tight labor markets on wages at various deciles, and these analyses reveal a greater impact.

6. Using Current Population Survey data, I calculated hours-weighted occupation employment shares (using three-digit codes) and rank them by their average wage level over the time period in question. I then calculate the changes in these shares between the two periods. Next, I find the ventile cutpoints in the cumulative distribution of these changes, and plot them, along with a polynomial curve.

7. Unlike quasi-fixed education levels, tasks that workers undertake can change quickly in response to relatively prices (i.e., the wages associated with the task). In the model underlying Autor's recent work (meaning Autor,Katz, and Kearney, 2006), workers reallocate time from routine to non-routine tasks as the price of routine tasks declines, so such changes are endogenous to price changes.

References

Autor, David H., Lawrence F. Katz, and Melissa S. Kearney. 2006. "The Polarization of the U.S. Labor Market." *American Economic Review* 96(2): 189–194.

Autor, David H., Frank Levy, and Richard Murnane. 2001. "The Skill Content of Recent Technological Change: An Empirical Exploration," NBER Working Paper No. 8337.

Card, David, and John E. DiNardo. 2002. "Skill-Based Technological Change And Rising Wage Inequality: Some Problems And Puzzles." *Journal of Labor Economics* 20(4): 733–783.

Goldin, Claudia, and Lawrence F. Katz. 2007. "The Race between Education and Technology: The Evolution of U.S. Educational Wage Differentials, 1890 to 2005." NBER Working Paper No. 12984.

Handel, Michael J. 2000. "Is There a Skills Crisis? Trends in Job Skill Requirements, Technology, and Wage Inequality in the United States." Jerome Levy Economics Institute Working Paper 295.

Howell, David. 1994. "The Skills Myth." *American Prospect* 5(18). Available online at http://www.prospect.org/cs/articles?article=the_skills_myth

Howell, David, and Edward Wolff. 1991. "Trends in the Growth and Distribution of Skills in the U.S. Workplace, 1960-1985." *Industrial and Labor Relations Review* 44(3): 486–502.

Howell, David, and Edward Wolff. 1992. "Technical Change and the Demand for Skills by U.S. Industries." *Cambridge Journal of Economics* 16(2): 127–146.

Levy, Frank, and Peter Temin. 2007. "Inequality and Institutions in 20th Century America." Mimeograph, Department of Urban Studies and Planning and Department of Economics, Massachusetts Institute of Technology.

Mishel, Lawrence, and Jared Bernstein. 1998. "Technology and the Wage Structure: Had Technology's Impact Accelerated Since the 1970s?" *Research in Labor Economics* 17: 305–355.

Oliner, Stephan D., and Daniel E. Sichel. 2000. "The Resurgence of Growth in the Late 1990s: Is Information Technology the Story?" *Journal of Economic Perspectives* 14(4): 3–22.

Comments on "Structural Demand Shifts and Potential Labor Supply Responses in the New Century" by David H. Autor

Gary Burtless

David Autor has written a clear, judicious, and even-handed survey of recent evidence on shifts in demand for U.S. labor, particularly the shifts that affect American employers' demand for skills. Since Autor is himself a major contributor to this literature, it is hardly surprising he has done such a fine job.

Almost everyone who regularly reads a daily paper or subscribes to a business magazine recognizes that earnings inequality in the United States has increased without interruption over the past three decades. Household income disparities have risen, too, and the growth in earnings inequality is an important contributor to this rise.

Autor's paper lucidly explains a crucial fact about the increase in earned income inequality. Earnings disparities did not actually rise uniformly and steadily over time. There have been major changes in the pattern of change in inequality. From the late 1970s up to the early 1990s, earnings inequality increased in almost every dimension. Workers earning the lowest wages saw their earnings fall in relation to median wages. Workers earning wages close to the middle of the distribution saw their earnings fall in comparison to wages at the top. Contrary to a common impression, however, this basic pattern did not continue uninterrupted up to the present day. It came to an end sometime around 1990. Since then earnings inequality has risen at the top—the earners with the highest wages are still pulling away from earners in the middle—but the bottom is not falling further behind the middle. Autor's Figures 4.2 and 4.3 show a sharp break in the pattern of increasing inequality around 1990. Now low-wage workers are pulling to closer workers in the middle, and have closed part of the earnings gap that opened up in the 1980s.

The shift in the pattern of growing wage inequality naturally has implications for how we think about the demand-side changes that may be contributing to wider earnings inequality. Theories that account for the trends in the 1980s may not do a good job of accounting for developments after 1990. Of course, some theories that helped explain why labor incomes at the very top of the distribution rose so much in the 1980s might do a reasonable job in accounting for wage trends after 1990. However, theories that explained the 1980s decline in relative wages at bottom of the distribution may have to be modified to reflect the fact that wages at the bottom have partly recovered since 1990.

Autor and a variety of his coauthors offer a plausible demand-side theory that accounts for the earnings pattern we have seen since the early 1990s. In recent years advances in information technology and communications methods have significantly increased the demand for the cognitive, decision-making, and interpersonal skills of managers and professionals who are adept at performing abstract, non-routine tasks. The same technical advances have reduced the relative demand for routine clerical, analytical, and mechanical tasks that can now be performed more cheaply with the assistance of inexpensive machines, such as personal computers. Technical advance has been less successful in reducing the need for people who perform some of the economy's least-well paid tasks, many of which require the on-the-spot presence of a manual worker. Autor notes that a variety of low-skill, low-pay, service sector occupations fit this description—health aides, security guards, hospital orderlies, cleaners, and servers. The result is a surge in relative demand for very highly skilled workers who can perform abstract, non-routine tasks, comparative stability in the demand for workers with the lowest skills, and a decline in the relative demand for workers with a middle range of skills.

I do not have any basic disagreement with this theory, which seems to me quite plausible. In addition to the technical and globalization trends that Autor emphasizes, however, I think there has also been a change in wage-setting practices inside of private firms. In the private sector as a whole, the change has occurred partly because unions represent a shrinking percentage of the American workforce. Unions now exercise direct and indirect influence over wages in a narrower and narrower slice of the private labor market. Since unions tended to equalize the wages of workers across skill categories, the reduced influence of unions has tended to weaken the bargaining power of many workers who perform routine, repetitive tasks. Shifts in pay-setting norms within large companies and innovations in executive compensation arrangements have also helped fuel wage growth at the very top of the pay structure. As already noted, this phenomenon has continued up to the present day, and it may continue in the future.

In the remainder of my comments I want to focus on a deep mystery regarding the supply-side response to the demand-side developments Autor describes. In particular, why has the response been so sluggish and small? And why has it been particularly small among American men? The average payoff to post-high-school formal education has risen, but the schooling and skill attainments of U.S. workers, especially men, have increased relatively little, both absolutely and in comparison with trends in other rich industrialized countries.

As noted in Autor's paper, one of the most important contributors to the growth in U.S. wage inequality has been the growing premium that people derive from earning a formal degree after high school. Figure 4.15 shows my own estimates of the earnings premium received by workers for completing a four-year college degree and earning a postcollege degree. The chart displays estimates of the log earnings difference between high school graduates and two groups of workers with higher levels of school attainment, college graduates and workers with at least one postcollege degree. To measure the premium for four-year college degrees and postcollege degrees, I regressed the logarithm of workers' annual labor earnings on age and educational attainment for years between 1968 and 2005 using the Census Bureau's March Current Population Survey files. In order to reduce the sampling variability of the displayed results, estimates shown in Figure 4.15 reflect the centered average of regression coefficients for five successive calendar years. The estimation sample includes full-time, year-round workers between 25 and 64 years old who have a valid report of their annual labor income, including both wages and net self-employment earnings.[1]

The top and bottom panels of Figure 4.15 show sizeable increases in the earnings premium enjoyed by college and postcollege degree holders during much of the period after 1980. Female degree holders saw their

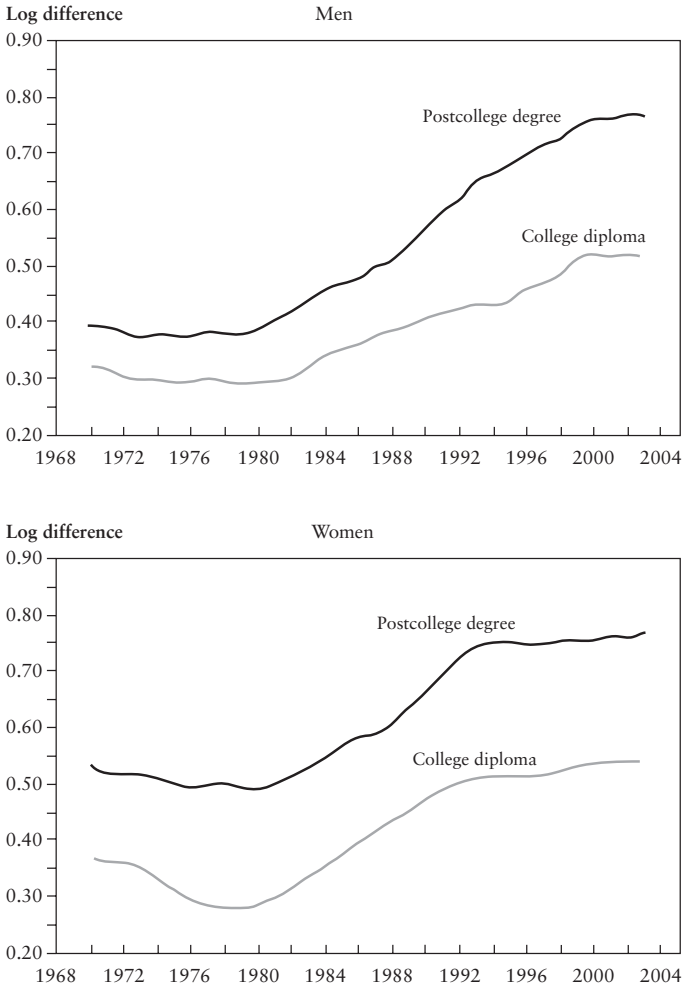

Figure 4.15
U.S. Trends in Earnings Premium for a Four-Year College Degree versus a High School Diploma, 1968–2005
Source: Author's tabulations of 1969–2006 March Current Population Survey files.
Note: The figure shows trends in the log pay differential between earnings of the indicated educational group and U.S. workers of the same sex who have a high school diploma. To reduce sampling variability, the chart shows 5-year centered moving averages of the regression coefficients.

educational pay premium rise substantially between 1980 and 1992, but their gains since 1992 have been very small; see Figure 4.15's lower panel. The increase in the educational pay premium persisted over more years for male degree holders, but their gains appear to have slowed or stopped in the late 1990s, as shown in the top panel of Figure 4.15. Since four-year college completion has become more common among working-age Americans, the rise in the payoff to advanced schooling has occurred against a backdrop of an increasing relative supply of well-educated workers. This development leads many labor economists to infer that the rising earnings premium for higher education must have signaled a rise in the relative demand for highly educated workers. The fact that the earnings premium for college and postcollege degrees stopped increasing in the middle or late 1990s suggests that some of the factors pushing up relative demand for highly educated workers slowed or the availability of college-educated workers increased.

Figure 4.16 shows estimates of the earnings penalty suffered by U.S. workers if they have failed to complete secondary school. The estimates were obtained using the same method and with the same sample described above. These results show that the pay differential between high school dropouts and graduates widened between 1980 and the late 1990s, but the differential has not widened much in recent years.

Now we come to the puzzle. Why was the supply-side response to these relative wage changes so sluggish and small? Figure 4.17 shows the trend in college completion rates in the U.S. population between the ages 25 and 34 years. This group is comprised of young adults who have just attained an age where we should expect that they have completed their college education. The broken line in the chart indicates college completion rates among women; the solid line shows the same trends among men. The financial reward for completing a four-year college degree rose steadily and strongly from 1979 to 2000, but college completion rates rose relatively slowly for women in the 1980s and actually declined among 25–34 year-old men in the same decade. College completion rates for young men and women improved in the 1990s. Nonetheless, the college completion rate for men is about the same in 2006 as it was in the late 1970s, when the pay premium for a four-year college degree was considerably smaller than it is today; see Figure 4.15. Young

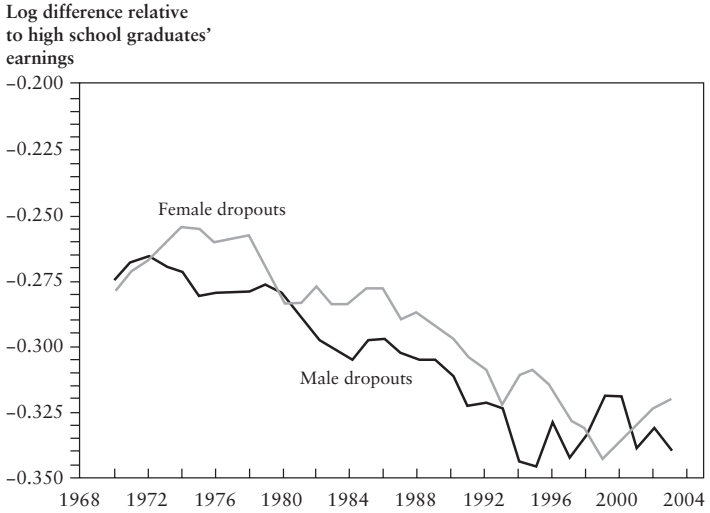

Figure 4.16
U.S. Trends in the Earnings Penalty for Failure to Complete High School, 1968–2005
Source: Author's tabulations of 1969–2006 March Current Population Survey files.
Note: The figure shows trends in the log pay differential between earnings of U.S. high school dropouts and workers of the same sex who have a high school diploma. The sample consists of full-time, year-round workers. To reduce sampling variability, the chart shows 5-year centered moving averages of the regression coefficients.

women apparently pay closer attention to wage trends described in the daily papers and business magazines. Today, college completion rates for women are appreciably higher than in the late 1970s. Their completion rates are, in fact, currently higher than young men's college completion rates. We cannot reject the hypothesis that young women are somewhat smarter or more forward-looking than young men. Even among women the trends shown in Figure 4.17 represent a bit of a mystery. The pay premium for a four-year college degree and a postcollege degree rose more strongly between 1980 and 1992 than it did in earlier or later years, as shown in the lower panel of Figure 4.15. Yet the trend toward higher college completion rates actually decelerated in that period. Even young women appear slow in responding to price signals in the job market.

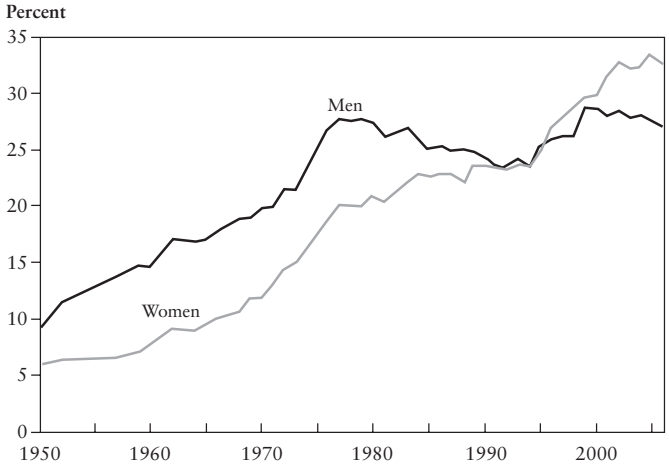

Figure 4.17
College Completion Rates of Americans Aged 25 to 34 Years, 1950–2006
Source: U.S. Census Bureau.

What about college completion rates among prime-age workers? Figure 4.18 shows the rates of four-year college completion among 35–54 year-olds in the five and a half decades after 1950. Prime-age males reached a peak rate of college completion in the late 1980s. There has been essentially no rise in prime-age men's college completion since that time. Of course, college attainment in this population mainly reflects educational decisions that were made at least a decade, and sometimes up to three decades, earlier. Very few 45-year-old men think seriously about enrolling in college. This middle-aged group is not a population segment where we would expect to see an instantaneous response to a bigger college pay premium. The college completion rate among prime-age women has risen much more steadily than the comparable rate among men. This probably reflects the fact that a much bigger percentage of women now expects to earn a large fraction of their families' total income over the course of their careers, raising the importance of obtaining a good educational credential that commands a premium in the labor market.

The educational decisions that should respond fastest to changes in wage signals are those made by adolescents and young adults between 16 and 24 years of age. In most states 16 years is the oldest age for which

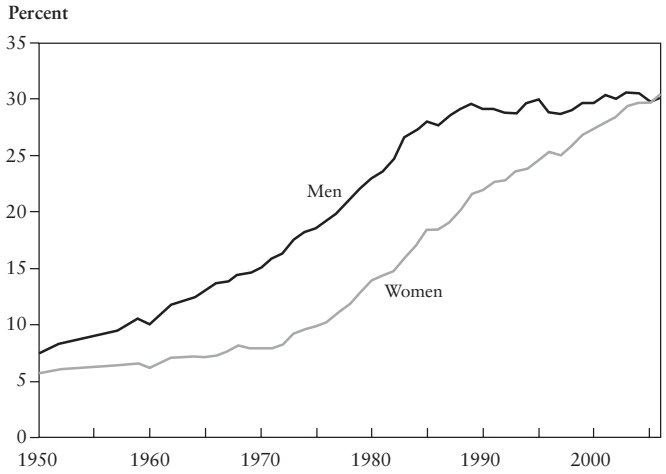

Figure 4.18
College Completion Rates of Americans Aged 35 to 54 Years, 1950–2006
Source: U.S. Census Bureau.

school attendance among youth is compulsory. Few people between 16 and 24 years of age have to enroll in school if they do not wish to attend. Figure 4.19 shows trends in the number of years that young people are enrolled in school between the ages of 16 and 24 years. The Census Bureau conducts an annual household survey in October asking about school enrollment. The estimates displayed in Figure 4.19 show the total number of Octobers between ages 16 and 24 years that members of successive birth cohorts have spent as enrollees in secondary school, college, or university. Someone enrolled in each October would have been enrolled for a total of nine years. The cumulative number of enrollment years rose from 2.9 years for 24-year-old men in 1955 to 4.5 years for 24-year-old men attaining in 1971. For 24-year-old men in the first half of the 1980s, the cumulative number of enrollment years fell below 4.0 years. From 1986 to 1998 the number of enrollment years increased, but for male cohorts attaining age 24 in years after 1998, the number of enrollment years has stagnated or declined slightly. In 2005 24-year-old young men accumulated only slightly more years of school enrollment than 24-year old men in the early 1970s. These enrollment trends seem

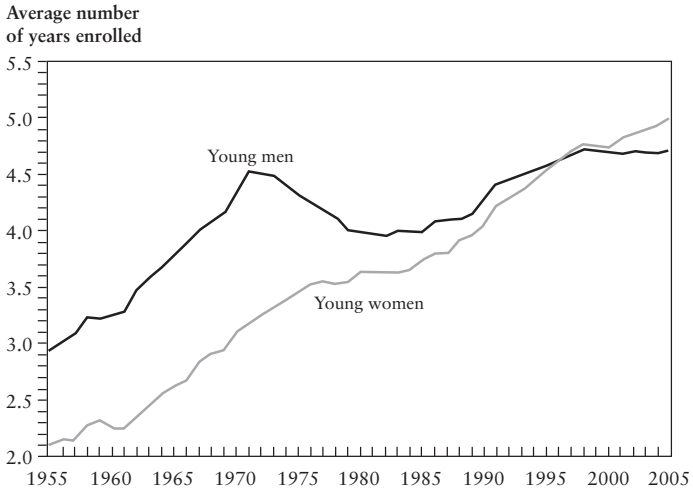

Figure 4.19
U.S. Trends in Educational Enrollment for Americans Aged 16 to 24 Years, 1955–2006
Source: U.S. Census Bureau.

puzzling in view of the fact that the male pay premium for completing a four-year college degree increased from 28 percent in 1980 to 51 percent in the late 1990s. The pay premium for obtaining a postcollege degree increased from 39 percent to 76 percent over the same period. Tuition, one component of college costs, increased over that period. However, another big component of enrollment costs fell— the opportunity cost of being enrolled in school. When the real wages of male high school dropouts and high school graduates declined, the foregone earnings of young men enrolled in school also declined.

The years young U.S. women spend enrolled in school have increased more steadily over time. Since 1997, 24-year-old women have accumulated more years of schooling than 24-year-old men. Female enrollment rates, however, do not show a clear pattern of response to the rise in the pay premium commanded by a four-year college degree. Women's enrollment rates rose rapidly between the 1960s and early 1980s, when the female college premium shrank, and these rates did not rise any faster when the college pay premium soared after the early 1980s.

It is worth comparing trends in school attainment in the United States with educational trends in other rich countries. Tertiary school completion rates have been calculated by the Organization of Economic Cooperation and Development (OECD) for a number of OECD-member countries. Tertiary schooling is defined as educational attainment that goes beyond a secondary degree but falls short of four-year college completion. Figure 4.20 contains results for 22 OECD countries. Each square indicates the tertiary completion rate among people who were between 45 and 54 years old in 2003. Each triangle indicates the rate of tertiary completion among people who were between 25 and 34 years old in the same year. By comparing the two tertiary completion rates, we have a rough indicator of the trend in tertiary completion in each country.

This comparison suggests the United States has seen little trend change in its tertiary completion rate over the past 20 years. Note that this estimate combines the completion rates of both men and women. There are a couple of other countries where tertiary completion rates have also been stagnant, notably New Zealand and Germany. By and large, however, most OECD countries have seen increases in their tertiary completion rate. In many countries, the gains have been substantial. Note that in the older age group, adults between 45 and 54 years of age, the United States has the highest tertiary completion rate of any of the 22 countries. In the younger age group, people between 25 and 34 years of age, the United States ranks only eighth out of the 22 countries for tertiary completion. Among these countries, the United States may have experienced the biggest increase in the gross financial payoff to obtaining a postsecondary educational degree. The wage premium for attainment of a postsecondary degree, meaning a four-year college degree or a postcollege degree, was typically larger in the United States than it was in other industrial countries in 1970s, and the premium increased more in the United States than it did in most other countries after 1979. In the face of bigger increases in the college pay premium, why have tertiary completion rates remained stagnant in the United States?

One explanation is the country's high rate of immigration. Although immigrants have a college graduation rate that is close to that of native-born Americans, immigrants also have exceptionally low rates of high school completion. They have much lower high school completion rates

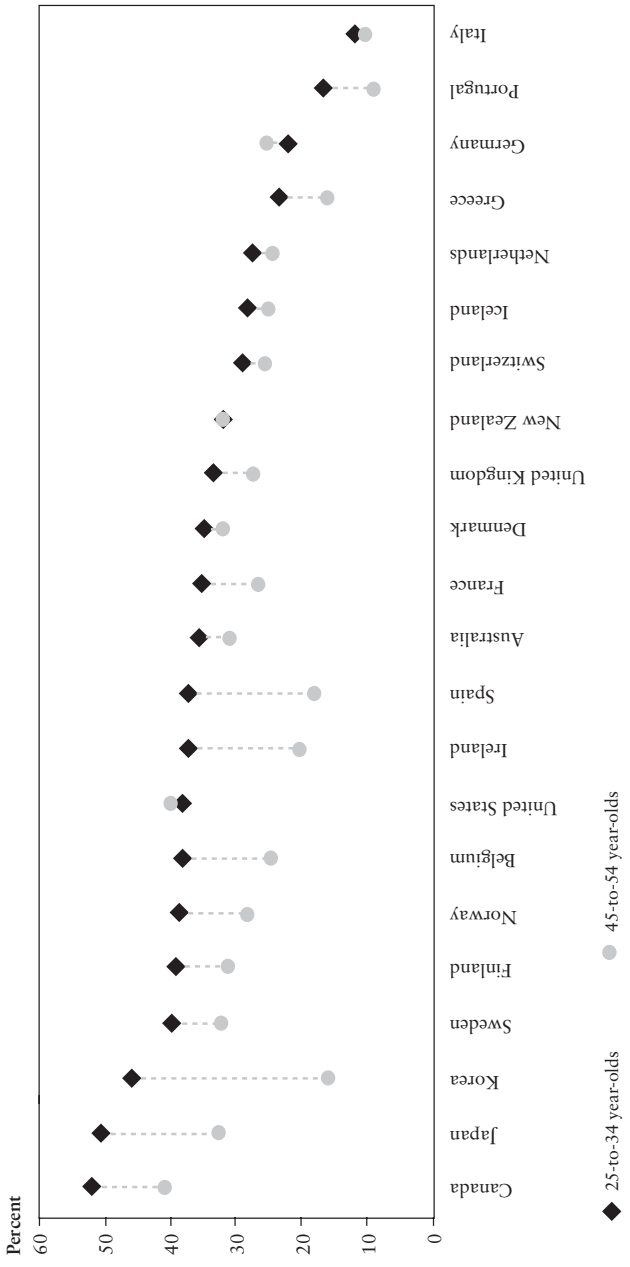

Figure 4.20
Trends in Tertiary School Completion Rates in 22 OECD Countries, 2003
Source: OECD, *Education at a Glance, 2005* (Paris: OECD, 2005). Chart A1.3.
Note: For Iceland, the Netherlands, and Italy, data are from 2002 rather than 2003. Tertiary school completion is defined as "... the population that has attained tertiary-type B education or tertiary-type A and advanced research programmes."

than native-born residents who are the same age, and many children of immigrants fail to enroll in or to complete college. Since immigrants and immigrants' children are a rising percentage of the U.S. resident population, the trends shown in Figures 4.17 through 4.20 show slower educational progress than we would see if U.S. immigration rates were lower or if admitted immigrants had more schooling.

Another explanation for Americans' slow educational progress is that school institutions have captured part of the increase in the educational pay premium. This has been accomplished by charging higher prices, mainly for tuition and fees. Most economists recognize that this cost increase has been partly or entirely offset by the declining opportunity cost of attending college. As young people's wages, especially young men's wages, have declined, their cost of delaying entry into the workforce has fallen. Nonetheless, a stubbornly high percentage of young American men has failed to attend or complete college.

Another possible explanation is a rise in the perceived risk in the payoff of attending a postsecondary institution. The college pay premium has increased on average, but so too has the variance around the average graduate's pay. There may be a bigger risk that workers with just one or two years of college will earn a lower wage than the average wage earned by a high school graduate. For instance, a high-school graduate who works as a plumber may well earn more than a community-college graduate who might work as a store clerk. However, it would require a very strong degree of risk aversion for far-sighted workers to remain out of college as a result of the increased uncertainty of college graduates' pay.

Why are college enrollment rates rising in other rich countries but remaining relatively constant in the United States? One possibility is that U.S. teenagers and young adults do not know how to perform the benefit-cost calculations that would inform them of the financial advantages of college attendance. Compulsory schooling laws provide some protection against the shortsightedness of 15-year-olds. State laws oblige 15-year-olds to attend school. Among people between 16 and 24 years of age, the main protection against short-sighted decisions is the influence of a rational and far-sighted parent. Upper middle-class and middle-class parents have many resources with which to bribe their youngsters. Some

parents pay for the full cost of a child's tuition and college room and board, and others pay a substantial part of the charges that are not covered by financial aid. A great deal of evidence suggests these parental bribes are successful in persuading middle-income and affluent children to attend college. Recent increases in college enrollment have been concentrated among young adults in the middle class and especially in the upper middle class.[2] In contrast, poorer parents have fewer resources to influence their children's secondary and postsecondary educational decisions. Many low-income students cannot borrow enough money or work enough hours to pay their college bills while simultaneously maintaining an acceptable grade point average. For students who do not particularly enjoy schoolwork, the decision about whether to attend college may depend on whether their parents bribe them to attend. When returns to higher education rise, far-sighted parents will want their children to attend college. But while affluent parents can influence their children's decisions by offering to pay, poor parents cannot. As a result, affluent students respond to changes in the long-term benefits of obtaining a college education, while poor students respond mainly to changes in the short-term costs of attending college.

An explanation for the divergence between educational attainment trends in the United States and other industrial countries may be that U.S. postsecondary institutions charge much higher fees than counterpart institutions in other rich countries. For that reason, the short-term cost of attending college may loom larger in the decisions made by adolescents and their parents. Most rich countries impose low charges on the students who qualify for admission to college, and a few routinely provide generous subsidies to cover the living expenses of enrolled students. It is of course possible in the United States for low- and moderate-income students to obtain generous financial aid or to enroll in low-cost public institutions, at least for the first two years of college. But many eligible students may not apply to college if they do not realize how much aid is available. Others do not apply because the short-term costs of attending seem large in relation to the distant and uncertain income gains they may achieve as a college graduate. While adolescents and young adults may be equally short-sighted and ill-informed in all industrial countries, the United States is unusual in imposing such high and erratic direct costs on

students who may enroll in college. The evidence in Autor's paper and in Figures 4.17 through 4.20 strongly suggests that the supply response to bigger college pay premiums in the United States has been too sluggish and too small to bring the college premium back to its level in the 1960s and 1970s.

Notes

1. The basic time trends are similar if I perform regressions based on all workers who earn at least $1 per year in wages or net self-employment income. In order to offset the effects of year-to-year changes in the Census Bureau's top-coding procedures, I top-coded earnings in every year using a simple and uniform procedure. Reported earnings that exceeded the 97th percentile of male earnings in a given year were recoded to the 97th percentile value for male earners, and earnings reports that exceeded the 99th percentile of female earnings in a given year were recoded to the 99th percentile value for female earners. This procedure means that the estimated education premiums do not capture the full earnings advantage enjoyed by well-educated earners in the top 3 percent of the male earnings distribution and in the top 1 percent of the female distribution. Thus, the estimates almost certainly understate the increase in the education premium, especially for men.

2. David Ellwood and Thomas Kane, "Who Is Getting A College Education? Family Background and the Growing Gaps in Enrollment," in Sheldon Danziger and Jane Waldfogel, eds., *Securing the Future* (New York: Russell Sage, 2000).

5

The Cyclical Sensitivity of Labor Supply

Cyclical Movements along the Labor Supply Function

Robert E. Hall

I. Introduction

A consensus in macroeconomics holds that the observed higher-frequency movements in employment and hours of work are movements along a labor-supply function caused by shifts of the labor demand function. Recent theoretical thinking has extended this view to include fluctuations in unemployment, so that macroeconomists can speak coherently of movements along an unemployment function caused by shifts in labor demand.

I develop an empirical framework for measuring the movements along the labor supply function and for measuring shifts of labor supply. I review data sources for the U.S. economy and conclude that the household survey is the only source of data that supports a clean set of measures of hours and employment. While recognizing the discrepancy between short-run movements of employment from the household survey and the employer payroll survey, at this point in my study I am unable to make any further contribution to reconciling the puzzle of the higher amplitude of employment fluctuations in the employer survey.

The measurement framework presented in this paper rests on the inference of an underlying single unobserved variable that determines labor supply. At the paper's end I discuss how this single driving force is related to modern macroeconomic labor-market theory. Here I use an econometric model with a latent variable to infer the unobserved variable, which turns out to move closely with unemployment. This variable has a high correlation with weekly hours as well, though there is much more

noise in the measurement of hours from the household survey than in the employer payroll survey.

II. Labor Supply and Labor Demand

Figures 5.1 through 5.3 describe four different views of changes that occur in the aggregate labor market when labor demand shifts outward. The shift may be the result of improved aggregate productivity, declines in the prices of inputs other than labor, or a favorable shift in the terms of trade. The horizontal axis is total labor input measured in hours per year. The vertical axis is the hourly real wage.

Figure 5.1 shows the standard neoclassical view of the labor market, which holds that labor supply is fairly inelastic. The labor market clears at all times at the intersection of supply and demand. A large outward shift in demand raises labor input by a small amount and the wage by a substantial amount. As a theory of fluctuations, the neoclassical view fails in describing both dimensions of demand and supply, as cyclical fluctuations in hours are large but small for wages.

Figure 5.2 shows two views with the same properties but very different rationalizations. In the real business cycle model, labor supply is highly

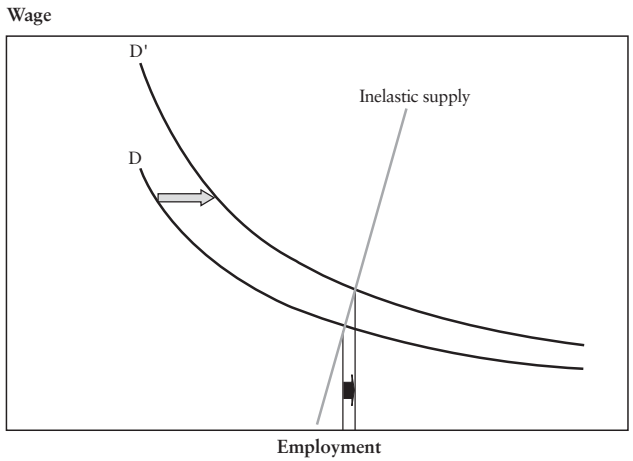

Figure 5.1
Demand Shift with Inelastic Supply

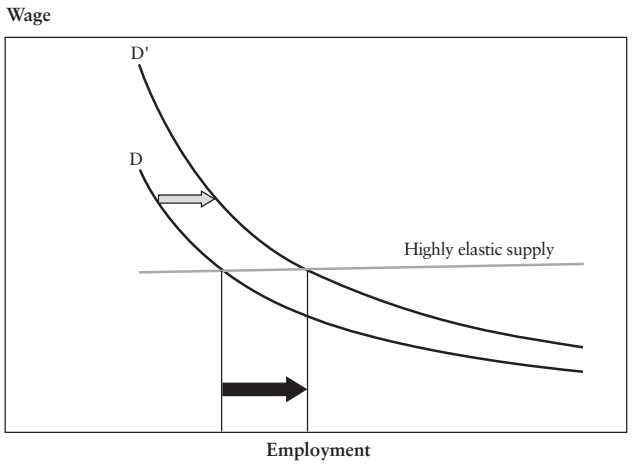

Figure 5.2
Demand Shift with Elastic Supply or Rigid Wage

elastic—the aggregate labor-supply schedule is essentially flat. Real business cycle theorists, notably Rogerson (1988), have provided analytical foundations for elastic supply and have addressed the important question of why studies of labor supply estimated at the level of individuals find relatively low elasticities—see Rogerson and Wallenius (2007). Whether microeconomic and macroeconomic estimates of labor-supply elasticities can be reconciled is a lively topic of debate today.

The other interpretation of Figure 5.2 takes the middle horizontal line to express rigidity of the real wage. It is not a standard supply schedule derived from the choices of workers about participation and hours, but the operation of a system of employment governance in which employers choose a level of labor input given a fixed real wage rate. Theoretical rationalizations of this system of governance have not fared well in recent years, after early enthusiasm about the possibility that contracts made under asymmetric information might take this form. The rigid real wage model carries with it an explanation of unemployment—it is the horizontal distance between the actual level of employment and the labor-supply curve of Figure 5.1. This rigid wage model is a gap theory of disequilibrium unemployment. Yet little theoretical work has been done in this framework in recent years, especially in the American context.

Figure 5.3 illustrates the theory of labor-market fluctuations underlying the measurement work discussed in this paper. As mentioned earlier, labor supply has its inelastic neoclassical form. Absent frictions in the labor market, shifts in labor demand would cause small changes in hours and large changes in hourly real wages. But the model used in this paper embodies an economic equilibrium view of unemployment derived from an explicit consideration of frictions. Unemployment is not a gap but is the result of the interaction of search and matching frictions and compensation determination. The search and matching elements are from Mortensen and Pissarides (1994). As Shimer (2005) demonstrated, search and matching frictions are not enough to explain cyclical fluctuations in unemployment. Shimer's paper set off an enthusiastic investigation of many different modifications of the Mortensen-Pissarides model. It is too early to say which will emerge as the leading explanation.

The Mortensen-Pissarides model describes physical frictions (such as the search and hiring process on the part of firms and workers) in the labor market but not wage frictions. Wages clear the market in a sense that they are the result of an unrestricted voluntary bargain between employers and workers. The simplest way to alter the Mortensen-Pissarides model in a way that makes it consistent with Figure 5.3 is to introduce what

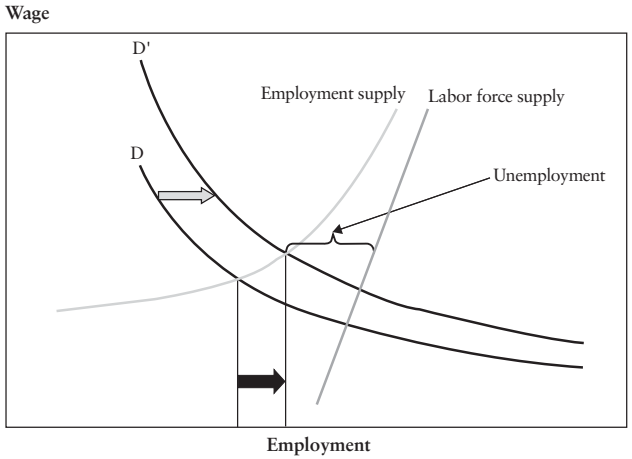

Figure 5.3
Demand Shift with Equilibrium Unemployment

I call equilibrium wage stickiness (see Hall 2005a). With this form of wage rigidity, the extended Mortensen-Pissarides model implies that an outward shift of labor demand, as it tries to push the wage up, will also reduce unemployment substantially. The result, as Figure 5.3 shows, is an increase in labor input that is much larger than the movement along the labor supply schedule because of the added effect of drawing people out of unemployment and putting them to work.

The line of thought expressed in Figure 5.3 embodies a full economic treatment of an individual's three possible activities related to the job market—remaining completely out of the labor force (specializing in nonmarket activity), looking for work, and being employed. In that sense, it is a natural extension of modeling from two activities, as illustrated in the first two figures, to three activities. But it is important to understand that the unemployment curve shown in Figure 5.3 is not just an expression of individual choice about how much time to spend looking for work. Rather, it is the equilibrium of the search and matching process together with wage determination. Individual choices about search and job acceptance are only one component of that equilibrium. For further discussion, see Hall (2008).

III. Earlier Work on Cyclical Fluctuations in the Labor Market

All of the earlier macroeconomic research that I have located so far takes unemployment, employment, or output as the measure of the business cycle in the labor market. I am not aware of work that infers an unobserved index.

Participation

Tella (1964) was an early and influential investigator of higher-frequency movements in aggregate labor-force participation. He considered the relation between the participation rate and the employment-to-population ratio, focusing on higher frequencies by using first differences and finding coefficients of 0.40 for men and 0.62 for women. These figures are substantially higher than those found in later work and in this paper, probably because participation is one of the components of the right-hand variable, and because he used data from 1948 through 1962, when

women's labor force participation was lower than it has been in subsequent decades.

Wachter (1977) found that for men, participation increased for all age groups in tight labor markets with low unemployment, though the effects are small except for teenagers and those over 65 years of age. For women, he found similar results for all but the older groups, for whom participation declines in tight markets.

Hours

Raisian (1983) studied the cyclical variation of hours per week and weeks per year as a function of experience in the labor market, using data from the Panel Study of Income Dynamics. He found that the elasticity of hours per week with respect to the employment rate (1 minus the unemployment rate) was 0.30, and that the elasticity of weeks per year was 1.14. The latter figure implies an elasticity of participation of 0.14.

Cho and Cooley (1994) took as a stylized fact of the U.S. business cycle that one-quarter of the variation in total hours of work is in hours per worker, and the remaining three-quarters is in workers per member of the working-age population. These are approximately the relative standard deviations of Hodrick-Prescott filtered hours per worker and employment.

IV. Framework and Data

The objective of this paper, which is part of a larger research agenda, is to develop a conceptual framework and corresponding data in which the three dimensions of labor supply—participation, employment rate, and hours—play roles derived from the macroeconomic theory of labor supply and unemployment.

The modern theory that provides the logical starting point for this paper's measurement framework runs as follows: Individuals have preferences defined over hours spent at home, hours of search, and hours of work. Each period (months are used in this paper) they choose an allocation of hours out of a set of available choices. Hours spent looking for work and hours spent at home are not restricted, but hours spent at work depend on the jobs available that period—workers do not have unilateral choice over jobs or the hours of jobs.

The macroeconomic theory of unemployment that has emerged from Mortensen and Pissarides's pioneering work focuses on the interacting behavior of job-seekers and employers. Hall (2007) gives an extended discussion of a generalization of their model. The job-creation efforts of employers control the unemployment rate. Employers respond to the job-creation incentive defined by the gap between the marginal product of labor (that is, labor demand) and the wage they expect to pay a newly hired worker. Wage flexibility is a key issue. If an increase in labor demand results in an equal increase in the wage, job creation and thus unemployment remain unchanged. The stickier the wage, the stronger is the decline in unemployment in response to an increase in labor demand. Some bargaining models imply sticky wages—see Hall and Milgrom (2008). Another source of wage stickiness is efficiency wages—see Alexopolous (2004).

Individuals' and firms' choices map into the three observable labor market activities for individuals. The Current Population Survey (CPS) uses certain important conventions in assigning individuals to activities. Although the CPS is a monthly survey, it uses a combination of time periods in the assignment process. The first convention is that work trumps any other activity, in the sense that a person who worked even one hour in the week before the survey is counted as employed, notwithstanding any other time spent at home or in a job search. The second convention is that a person not recorded as employed is recorded as unemployed if the person was not participating in the workforce in the previous week, but made any of a variety of designated types of efforts to find a job in the preceding four weeks. Those who fail to meet the criteria for employment or unemployment are counted as out of the labor force.

The recent launch of the American Time Use Survey (ATUS) will provide a far more complete view of the allocation of household time. The new survey focuses on measuring all uses of time rather than assigning individuals to categories based on partial measures. However, the size of the ATUS sample is not large enough to support good national estimates of monthly labor-market status.

Flinn and Heckman (1983) make the reasonable proposal that the unemployed should be taken to be non-working individuals who have a probability, above a designated threshold, of finding work in the coming

period. The CPS definition of unemployment appears to implement a rough approximation to the Flinn-Heckman definition. Along with Flinn-Heckman, the CPS definition does not classify people as unemployed if these individuals have decided that no job realistically likely to become available would be superior to engaging in non-work activities. The CPS has a separate category for these people, often called discouraged workers.

The home activities that occupy all individuals, employed or not, include home production as well as leisure. As the ATUS shows, these activities include shopping, cooking, and caring for others, together with sleep and pure leisure, such as watching television or socializing.

V. Measuring Employment

The Bureau of Labor Statistics (BLS) runs two independent surveys aimed at determining a seemingly simple concept: the number of people at work in the United States at a given moment. In addition to the CPS count of employment, the BLS surveys employers about the number of workers on their payrolls. Almost from the beginning of the household survey, economists noted cyclical discrepancies between the two surveys—the payroll measure of employment rises faster in expansions and falls faster in recessions than does the CPS measure. The report of the President's Committee to Appraise Employment and Unemployment Statistics commented extensively on the issue in 1962. Economists affiliated with the party in power find reasons to praise the CPS measurement during recessions—especially the most recent one—while others cite the payroll survey as the more accurate description of the ravages of the downturn.

Figure 5.4 compares employment counts from the two sources. It shows the raw ratio of the payroll count to the household count together with its higher-frequency component. The latter comprises the residuals from a regression of the ratio on a fourth-order polynomial in time. The payroll count rose irregularly from 82 percent of the CPS level in 1959 to 97 percent at its maximum at the end of the 1990s and then fell to its current level of 94 percent.

The higher-frequency component of the ratio comparing the household survey with the employer survey is conspicuously correlated with

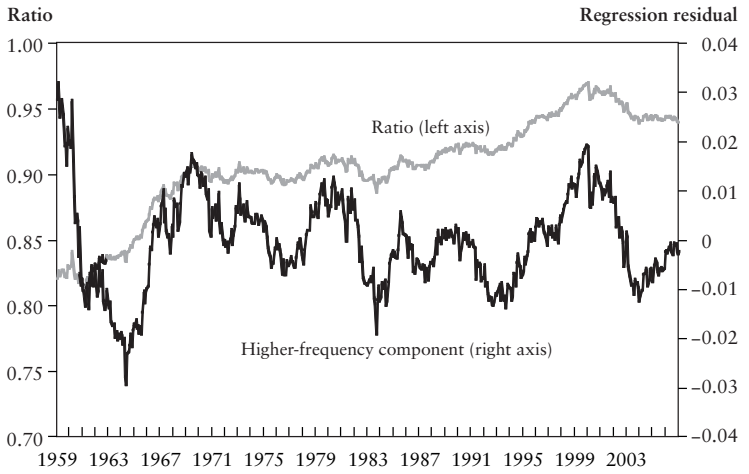

Figure 5.4
Ratio of Payroll Employment to Current Population Survey Employment
Source: U.S. Bureau of Labor Statistics, U.S. Census Bureau, and author's calculations.
Note: The higher-frequency component comprises the residuals from a regression of the ratio on a fourth-order polynomial in time.

the business cycle. In each recession, the payroll count falls by 1 to 3 percent of the CPS count. The decline was particularly large in the most recent recession. It was large in the worst postwar recession that took place in 1981–1982, but not as large in the other deep postwar recession that occurred between 1973 and 1975.

The cyclical discrepancy between these two measures remains almost entirely unexplained. Table 5.1 shows a dissection of the conceptual differences between the two employment measures based on Bowler and Morisi (2006). The top line is the percentage shortfall of the payroll count from the CPS count. During the expansion years, 1994 through 2000, the shortfall shrank and then expanded during the recession and following years, 2000 through 2004. The column headed *Cycle* is the percentage growth from 2000 through 2004 plus two-thirds of the shrinkage from 1994 though 2004. This figure is zero if the figures to the left grow linearly with time and is positive if the figures to the left fall during the expansion and rise during the contraction, as the payroll shortfall plainly does.

Table 5.1
Components of CPS Employment Related to Conceptual Differences from Payroll Data

	1994	2000	2004	Cycle
Shortfall of payroll jobs	7.13	3.73	5.58	4.12
Components from CPS				
Agriculture	2.77	2.47	1.60	–0.67
Non-ag self employed	7.32	6.40	6.80	1.00
Non-ag unpaid family workers	0.11	0.08	0.06	0.01
Private household workers	0.78	0.66	0.56	–0.01
Unpaid absence	1.62	1.47	1.38	0.01
Multiple jobholders	–5.51	–5.20	–5.07	–0.07
Total components	7.09	5.88	5.33	0.27

The entries collectively labeled "Components from CPS" report components of CPS employment measures that are conceptually different from the payroll data, stated as percentages of the total CPS employment count. The business-cycle measure is given in the right column for each adjustment. A positive cycle measure means that the component helps explain the pro-cyclical discrepancy between the employer payroll measurement and CPS household counts.

The first of the conceptual differences between the two employment measures is that the CPS one includes the self-employed and wage earners in agriculture, whereas the payroll data exclude agricultural employment. The cycle measure is negative for this component—the strong labor market of 2000 resulted in an upward deflection in agricultural employment. This phenomenon only deepens the mystery of the cyclical discrepancy, as it would make the CPS measurement by itself more cyclical than the payroll data.

The second adjustment shows an important source of the cyclical discrepancy—self-employment—which declined sharply as a fraction of CPS employment during the 1994–2000 expansion, and rose a bit during the 2001 recession and its aftermath. The payroll data exclude the self-employed.

The four remaining CPS components shown in Table 5.1 account for trivial percentages of the cyclical movements. Unpaid family workers and private household workers, included in the CPS but excluded from the payroll data, comprise tiny fractions of total employment and have no cyclical component. People who have jobs but are not currently being paid—individuals counted in the CPS household measurement but not in the payroll data—make no contribution to the cycle. And second jobs—counted twice in the payroll count of jobs but only once in the CPS count of employed people—make a small contribution in the wrong direction to explain the discrepancy between the two employment estimates.

Notice that the total CPS components almost perfectly match the payroll counts in the years 1994 and 2004, but result in an excess of payroll employment in the peak year, 2000.

According to Table 5.1, the cyclical discrepancy in employment counts between the two surveys is almost completely a mystery. The table covers all but one of the important conceptual differences between the surveys, the length of the reference period. In the CPS measurement, a person who worked one hour or more in the week before the survey counts as employed. The payroll survey counts the number of people on an employer's payroll at any time during the pay period that includes the twelfth of the month. My impression is that pay periods are generally two weeks, half a month, or a month.

Explaining the relationship between the length of the pay period and the overstatement of monthly snapshot of unemployment by the payroll data is simple: the overstatement is the weekly rate of new hires times the number of weeks in the pay period. Hall (2005b) discusses evidence on cyclical variation in the new hire rate. The Job Openings and Labor Turnover Survey (JOLTS) measures the new hire rate directly and shows little variation in the only business cycle that has occurred since it was launched in 2000. The separation rate is an excellent proxy for the new hire rate—the two measures differ only by the rate of change of employment, which at all times is insignificant in comparison to the levels of new hires and separations. The CPS has measured total separations since 1994, so it too includes only the most recent cycle. Figure 2.4 in Hall (2005b) shows that the monthly separation rate fell by about half a percent from the strong labor market of 2000 to the weak market of 2003. The weekly rate thus fell by a little over one-tenth of one percent. Even if the pay period is monthly, or four weeks, cyclical variations in the

overstatement of employment caused by longer pay periods are tiny in relation to the observed discrepancy in the cyclical behaviors of CPS and payroll employment measurements.

Absent an understanding of the source of the extra cyclical movements of the payroll employment data, it is not possible to use the data in the three-activity framework normally used in research on labor-market dynamics. The difficulty is that the individual fractions of the population engaging in the three possible activities—out of the labor force, unemployed, and working—must sum to one. The payroll survey provides no measure of the first two activities. One would have to adjust the fractions from the CPS household survey for calculating the fractions of the population that are out of the labor force and are unemployed to satisfy adding up to one. There is no basis for making the fraction of those out of the labor force and the fraction of those unemployed more countercyclical than is reported in the CPS, but these adjustments would be needed to incorporate the payroll employment data.

VI. Data on Hours

The CPS asks the household respondent (often not the actual worker whose activity is reported), a question like, "So, for last week, how many hours did the individual actually work at her or his job?" (the computer tailors the question to the individual worker). The survey procedure gathers hours separately by job for multiple-job holders. The respondent decides what constitutes an hour of work—whether it includes breaks, setup time, and the like.

The CPS measure of hours drops dramatically at random, such as when a holiday falls in the reference week. The choice of the reference week as the one including the twelfth of the month dodges Thanksgiving, Christmas, and several other major holidays, but cannot exclude every holiday celebrated in the United States. Monthly plots of hours show these drops.

The BLS also provides a comprehensive measure of hours based primarily on the payroll data, extended to agriculture and self-employment with CPS data; see the Major Sector Productivity and Costs (MSPC) index at http://bls.gov/lpc/home.htm. The payroll survey determines hours paid per job from employers. The MSPC index restates the results on the basis of hours worked rather than hours paid, using another survey that collects both. The MSPC index also uses CPS hours for workers not covered by the payroll survey. Although the MSPC measure of hours is mainly hours per job rather than hours per worker, there is so little cyclical variation in jobs per worker that the distinction is unimportant for the study of cyclical phenomena.

Figure 5.5 compares the two sources of data on hours. Hours as measured by the MSPC fell by about 10 percent relative to CPS hours from 1959 to 2005. I am not aware of any discussion or explanation of this behavior. As in Figure 5.4, I also show the higher-frequency component. It is relatively small and not conspicuously cyclical. Apart from the differing trend, there seems no important discrepancy between the measures.

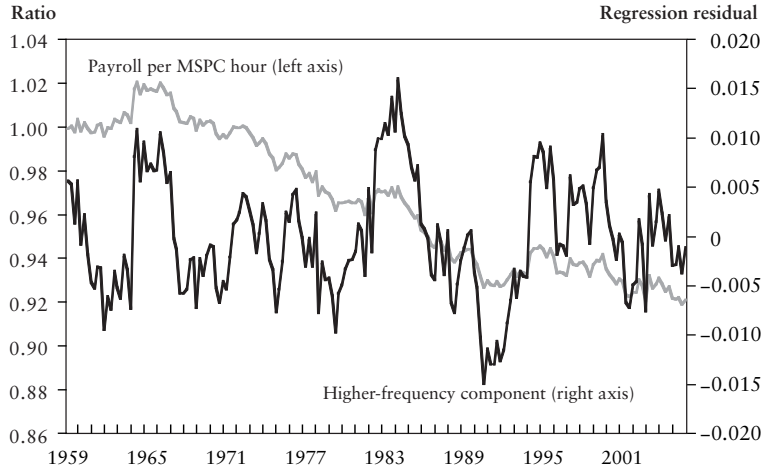

Figure 5.5
Ratio of Major Sector Productivity and Costs Index Hours to Current Population Survey Hours
Source: U.S. Bureau of Labor Statistics, U.S. Census Bureau, and author's calculations.
Note: The higher-frequency component comprises the residuals from a regression of the ratio on a fourth-order polynomial in time.

VII. The Single Driving Force of Movements along the Labor Supply Function

The consensus of modern macroeconomics is that shifts of labor supply are not a significant driving force of the business cycle. Rather, productivity shocks, oil shocks and other shifts in the terms of trade, and changes in other factor prices move workers along their labor supply schedules. Hours of work reflect variations over time in the current payoff to work relative to the value of other activities. Choices about whether to participate in the labor market also reflect a similar choice. Both of these derive from perfectly standard models of labor supply.

A more recent extension, deriving from the work of Mortensen and Pissarides (1994), of this consensus view has developed a model incorporating job search, the third use of an individual's time, that responds to the same factors affecting labor supply. Hall (2007) shows how unemployment fits into a model of labor-market fluctuations. That paper derives two indexes that jointly capture the driving forces of labor-market fluctuations. One index describes the overall well-being of households, based on expectations of future earnings. The other describes the current state of the labor market. The two indexes are highly correlated, so it is a reasonable approximation to treat the labor market as having a single driving force, the approach taken in this paper.

The important point that derives from this line of thought is that a single force drives all three key measures—participation, unemployment, and hours of work. This single force is the current position of the labor demand function in relation to its typical level trend. When labor demand is unusually strong, labor force participation rises, unemployment falls, and hours of work rise. The rest of this paper will derive a measure of the single driving force from the multiple indicators and measure the relative cyclical sensitivities of participation, unemployment, and hours.

The model underlying this work—and the conclusion about a single driving force—does not necessarily rest on any ideas of the kind usually labeled Keynesian. In fact, all of the paper's conclusions, except the magnitude of the fluctuations, will hold in a neoclassical, real-business-cycle model, extended only in the direction of Mortensen-Pissarides. Although the easiest way to explain the observed amplitude of the responses of labor-market variables to the driving force is with sticky wages, it is an open and very interesting question whether other mechanisms may be involved as well, or if any wage or price stickiness is required to explain the labor supply response.

To derive a measure of the single driving force, I use three monthly measures that track the business cycle. Two measures are from the labor market: unemployment and hours. To put unemployment in a form that makes it interchangeable (except for sign) with employment per participant in a log-additive framework, I measure unemployment as the negative of the log of the employment rate. The third measure in the cyclical model is real disposable personal income per capita (see National Income and Product Accounts, Table 2.6).

The following econometric setup enables the measurement of the common driving force, z_t:

$$(1) \quad y_t = \gamma_y z_t + \tau_y(t) + \varepsilon_{y,t}$$

$$(2) \quad u_t = \gamma_u z_t + \tau_u(t) + \varepsilon_{u,t}$$

$$(3) \quad h_t = \gamma_h z_t + \tau_h(t) + \varepsilon_{h,t}.$$

Here y_t is log real income, u_t is the unemployment rate, and h_t is the log of weekly hours of work. The γs are the loading factors of the observed variables on the unobserved driving force, z_t. These factors are interpreted as elasticities of the component with respect to the cyclical driving force. The $\tau(t)$ functions capture slower-moving non-cyclical determinants of the observed variables and the εs are the idiosyncratic higher-frequency movements not associated with the cyclical driving force z_t—the εs are assumed to be uncorrelated with z_t. I assume that z_t, whose units are arbitrary, has a variance of one. I also assume that γ_u is negative, so z_t is procyclical.

I specify the $\tau(t)$ functions as fourth-order polynomials in time. I also include seasonal dummies for hours because the data are not seasonally adjusted. The model has two sets of moment conditions. The first are standard regression conditions—orthogonality of the time variables in the τ functions with the disturbances. The regression part—like all regressions—has the same number of moment conditions and unknown parameters, and is exactly identified.

The second set of moment conditions describes the latent-variable structure of the disturbances. This part of the model has six observed moments: three variances of the ε disturbances, V_y, V_u, and V_h, and three covariances, $C_{u,y}$, $C_{u,h}$, and $C_{h,y}$. It has six unknown parameters, γ_y, γ_u, γ_h, σ_y, σ_u, and σ_h, where the last three are the standard deviations of the idiosyncratic components. The latent-variable model is exactly identified and has the following moment conditions:

$$(4) \quad C_{u,y} = \gamma_u \gamma_y,$$

$$(5) \quad C_{u,h} = \gamma_u \gamma_h,$$

$$(6) \quad C_{h,y} = \gamma_h \gamma_y,$$

$$(7) \quad \sigma_y^2 = V_y - \gamma_y^2,$$

$$(8) \quad \sigma_u^2 = V_u - \gamma_u^2,$$

$$(9) \quad \sigma_h^2 = V_h - \gamma_h^2.$$

The overall model is exactly identified. Its moment conditions are block-triangular—I can solve for the regression parameters first, and then derive the latent-variable parameters. The first step is to estimate regressions of the three variables on the components making up the τ functions (powers of t and seasonal dummies). I denote the residuals from these regressions as $\hat{y}_t$ and similarly for u and h. The variances and covariances in the moment conditions for the latent-variable model then refer to the hatted residuals.

From the moment conditions, I derive

$$(10) \quad \gamma_y = \sqrt{\frac{C_{u,y} C_{h,y}}{C_{u,h}}},$$

with the square root taken as positive. The remaining parameters come directly from the moment conditions. Notice that the model imposes a condition on the signs of the covariances—the expression under the square root is non-negative. In addition, the implied values of the squared values of three σ parameters must be non-negative.

To infer the values of the single driving force z_t, I use the projection of z on the observed variables; that is, the fitted values of the regression of z on those variables. The regression coefficients are the inverse of the covariance matrix of the variables (observed), multiplying the vector of covariances of z and the variables. The covariances are just the estimated parameters γ, because the variance of z is one.

Table 5.2 shows the results of these calculations. The top panel shows the variances and covariances of the residuals from the preliminary regressions. The unemployment rate is in percent and real income and hours in 100 times their natural logs. Hours and unemployment have about the same variances but the variance of real income, around its lower-frequency trend, is quite a bit higher. The covariances of the three variables are as expected—unemployment is countercyclical and income and hours are procyclical.

The first line in the lower panel of Table 5.2 shows the loading coefficients, γ, for the three variables. Unemployment has a loading coefficient on the cyclical driving force of just under 1. The next line shows that unemployment has a fairly low idiosyncratic movement—the variance of its non-cyclical higher-frequency movements is only 0.22. Real income loads on the cyclical component with an elasticity of 1.39 and has an idiosyncratic variance of 2.27, about half its total variance of 4.20. Hours load on the cyclical driving force with an elasticity of 0.56, leaving a large idiosyncratic variance of 1.32 out of its total variance of 1.64.

The a coefficients for extracting the implied time series for the driving force z show that the optimal inference places a large negative coefficient on unemployment and smaller positive coefficients on real income and hours. Figure 5.6 shows the index $\hat{z}_t$.

Table 5.2
Inference of Cyclical Driving Force from Data on Unemployment, Real Income, and Hours

	Unemployment	Real income	Hours
Moments			–0.54
Unemployment	1.14	–1.33	0.78
Real income		4.20	1.64
Hours			
Parameters			
Loading on z, γ	–0.96	1.39	0.56
Variance, σ^2	0.22	2.27	1.32
Coefficients for z, a	–0.696	0.097	0.068

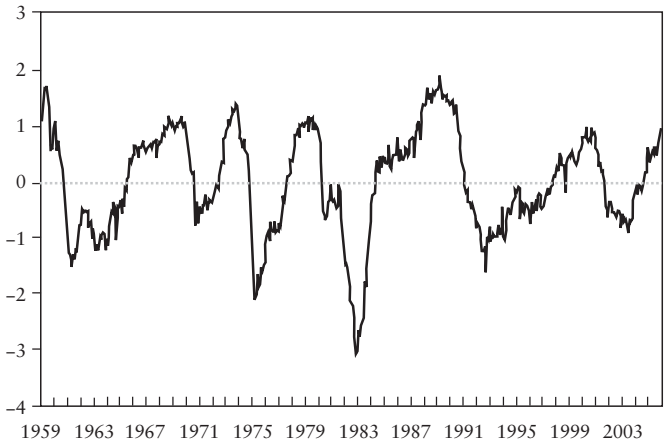

Figure 5.6
Index of the Single Driving Force
Source: U.S. Bureau of Labor Statistics, U.S. Census Bureau, and author's calculations.

VIII. Cyclical Sensitivity of Participation, Unemployment, and Hours

I am now equipped to answer the basic question of the cyclical sensitivities of participation, unemployment, and hours. Table 5.3 shows the loading factors for the three dimensions of work effort on the driving force, z. For employment, the coefficient is the positive value of the one shown in Table 5.2 and for hours, it is the value shown there. For participation, not included in the earlier model, I show the coefficient of the regression of log of the CPS participation rate on the inferred measure, $\tilde{z}$; the regression also includes the fourth-order polynomial in t as in the earlier regressions. For all three components, I measure the standard error from that type of regression. The total loading shown at the bottom is just the sum of the loadings of the three components.

The first line of Table 5.3 shows the small but statistically unambiguous cycle in participation. Recall that the units of the cyclical driving forces are standard deviations of cyclical movements in the labor market. A one standard deviation tightening of the labor market raises participation by 0.2 percent. Because the aggregate level of participation is around 60 percent, this increase in the driving force increases labor force participation by about 0.12 percentage points. The response of participation is 11.6 percent of the total response of labor input.

Employment, shown in the second line, is a bit more than half of the total cyclical variation. A tightening of the market by one standard deviation raises employment and lowers unemployment by just under one percentage point.

Weekly hours, shown in the third line, account for a third of total cyclical variation in labor input.

Tables 5.2 and 5.3 deal with labor measures per person. Table 5.4 considers the employment count, the product of population, participation, and the employment rate. I will not consider the employment count in the rest of the paper, but it does permit a further consideration of the difference between the CPS and payroll data, as the latter take the form of employment counts only, without the breakdown into population, participation, and the employment rate.

Table 5.4 needs to put population on the same footing as the other measures, as the higher-frequency component obtained as residuals from

Table 5.3
Loading Coefficients for the Three Dimensions of Work on the Cyclical Driving Force

	Loading	Standard error	Percent of total
Participation	0.197	(0.008)	11.6
Employment	0.957	(0.008)	56.5
Hours	0.534	(0.099)	31.5
Total	1.696	(0.075)	

Table 5.4
Cyclical Loading Coefficients for Number of Employed Workers

	Loading	Standard error
Population	–0.147	(0.018)
Participation	0.197	(0.008)
Employment rate	0.958	(0.008)
CPS employment	1.006	(0.029)
Payroll employment	1.512	(0.048)

the regression of the log of population on a fourth-order polynomial in time. Then, to reconcile the CPS measures, including population, with the payroll measure in the measurement framework used in the paper, it needs to measure the loading of population on the cyclical driving force. The population loading component is –0.147, with a standard error of 0.018. How can population be countercyclical? Obviously population does not respond to the forces that cause the business cycle, but population swings could be a contributor to the cycle. This hypothesis seems to be part of the explanation—unusually low population growth during World War II led to a tighter labor market in the late 1960s, before the baby boomers entered, and unusually high population growth led to a slacker market in the period containing the weakest labor market, 1973 through 1983, as the boomers started work. Another part of the explanation is discontinuous increases in the population estimates used in the CPS at the beginning of 1990 and 2000, both near cyclical peaks.

The loading of the CPS employment count on the cyclical driving force is very close to 1. The loading is only slightly higher than the loading for employment per participant in Table 5.3, because the negative effect of population offsets the positive role of participation.

By contrast, the loading of the log of payroll employment on the cyclical driving force is much higher, at 1.512. The stronger cycle in payroll employment shows through prominently in the framework of the cyclical driving force, even though the driving force is derived completely independently of the payroll data.

Although higher-frequency changes in the working-age population are shifts of labor supply rather than movements along a labor-supply function, the movements in participation, unemployment, and hours considered in this paper are movements along their respective functions. The discovery that population movements are part of the driving force of those movements is quite consistent with the overall framework of this paper.

IX. Cyclical Responses by Demographic Groups

Table 5.5 breaks down the responses shown in Table 5.3 by age and sex, to the extent that the data are available from the BLS. Long historical tabulations of the data are incomplete, though the important features of the differences among demographic groups are visible and in accord with prior beliefs. The hours data in Table 5.5 begin in June 1976.

Table 5.5
Loading Coefficients for Participation, Unemployment, and Hours by Age and Sex

	Sex	Age	Loading	Standard error
Participation	Male	16 to 24	0.339	(0.071)
		25 to 54	0.146	(0.016)
		55+	0.349	(0.065)
	Female	16 to 24	0.679	(0.095)
		25 to 54	0.147	(0.047)
Unemployment	Male	16 to 24	–2.283	(0.029)
		25 to 54	–1.081	(0.018)
	Female	16 to 24	–1.537	(0.028)
		25 to 54	–0.867	(0.013)
Hours	Male	16 to 19	2.273	(0.105)
		20 to 24	1.179	(0.057)
		25+	0.699	(0.057)
	Female	16 to 19	2.125	(0.107)
		20 to 24	1.135	(0.066)
		25+	0.573	(0.070)

Table 5.5 confirms that the participation elasticity is higher for younger workers (those under 25 years of age) and for older workers (those over 54 years of age), and among younger workers is higher for women than for men. The more elastic groups contain a larger fraction of people who are close to the margin between choosing to participate in the labor force and choosing to specialize in non-work activities, primarily activities at home and attending school. Unemployment among men and among younger workers is more sensitive to the driving force.

The elasticities of hours with respect to the cyclical driving force are slightly lower for women than for men. For both sexes, the response of hours is much higher for the younger workers.

X. A More Complete View of Driving Forces

Hall (2008) develops a model of labor supply and unemployment derived from the underlying principles of labor supply and a generalization of the Mortensen-Pissarides search-and-matching view of unemployment. In the Hall model, hours of work and unemployment are linked to the rest of the economy through two variables. One is the marginal product of labor, which conveys the demand for labor. The other is the marginal utility of consumption, which conveys the long-run well-being of workers. Because the marginal product of labor is the primary determinant of well-being and because the marginal product of labor tends to evolve as a random walk, the marginal product of labor and the marginal utility of consumption are quite highly correlated. In this paper, I make the approximation that the correlation is so high that these can be treated as a single variable, which I call the driving force, z.

The analysis of some issues of the response of labor-market variables requires the two-variable framework. In particular, all macroeconomic models agree that an expansion in government purchases of goods and services tightens the labor market, raises hours of work, and decreases unemployment. The increased government spending does not change the demand for labor as measured by the marginal product of labor. Instead, it depresses long-term well-being because any increase in government purchases must be paid for, sooner or later, by reduced consumption. Hours of work increase because of the negative wealth effect in labor supply—higher government purchases decrease wealth and raises hours. Hall (2007) shows that the same effect operates on unemployment—lower wealth results in lower unemployment.

In the framework with two variables, the consensus I noted in the introduction might be phrased more precisely as "a consensus in macroeconomics holds that the observed higher-frequency movements in employment and hours of work are movements along labor-supply and employment functions caused by changes in fundamentals acting through the marginal product of labor and the marginal utility of consumption." The consensus viewpoint rules out shifts of the labor-supply and employment functions as important sources of fluctuations.

XI. Interpretation

I have not tested the consensus viewpoint that shifts in labor demand account for most of the cyclical variation in labor input. But it holds up well provisionally in the analysis of this paper. First, all three components of the labor demand function—participation, the employment rate for participants, and hours per week of workers—respond positively to my measure of cyclical shifts in labor demand. Because these shifts are transitory, these involve mostly substitution effects. Basic labor-supply theory shows that the substitution effect in participation and in hours per worker should be positive. The extended Mortensen-Pissarides theory requires the substitution effect for the employment rate to be positive as well.

More than half of the extra labor input in a cyclical upswing is drawn from the ranks of the unemployed. No model of the cycle in the labor market can claim any realism unless it takes this finding seriously. It is inappropriate to lump those assigned by the CPS to unemployment together with those workers found to be out of the labor force, because the unemployed are much more likely to be employed a month later. The unemployed are truly different from other people who are not working because they generally wind up working within a few months.

Research trying to explain the high cyclical elasticity of unemployment has made exciting advances in the past few years, but a great deal remains to be done.

■ *Prepared for the 52nd Economic Conference of the Federal Reserve Bank of Boston, Cape Cod, June 18 to 20, 2007. This research is part of the Economic Fluctuations and Growth program of the National Bureau of Economic Research. Data and programs are available from the author—Google "Robert Hall".*

References

Alexopoulos, Michelle. 2004. "Unemployment and the Business Cycle." *Journal of Monetary Economics* 51(2): 277–298.

Bowler, Mary, and Teresa L. Morisi. 2006. "Understanding the Employment Measures from the CPS and CES Surveys." *Monthly Labor Review* 129(2): 23–38.

Cho, Jank-Ok, and Thomas F. Cooley. 1994. "Employment and Hours Over the Business Cycle." *Journal of Economic Dynamics and Control* 18(2): 411–432.

Flinn, Christopher J., and James J. Heckman. 1983. "Are Unemployment and Out of the Labor Force Behaviorally Distinct Labor Force States?" *Journal of Labor Economics* 1(1): 28–42.

Hall, Robert E. 2005. "Employment Fluctuations with Equilibrium Wage Stickiness." *American Economic Review* 95(1): 50–65.

Hall, Robert E. 2005. "Job Loss, Job Finding, and Unemployment in the U.S. Economy over the Past Fifty Years." In *NBER Macroeconomics Annual,* Volume 20, ed. Mark Gertler and Kenneth Rogoff, 101–137. Cambridge, MA: The MIT Press.

Hall, Robert E. 2008. "Sources and Mechanisms of Cyclical Fluctuations in the Labor Market." Available at http://www.stanford.edu/~rehall/SM022208.pdf.

Hall, Robert E., and Paul R. Milgrom. 2008. "The Limited Influence of Unemployment on the Wage Bargain." *American Economic Review*, forthcoming.

Mortensen, Dale T., and Christopher Pissarides. 1994. "Job Creation and Job Destruction in the Theory of Unemployment." *Review of Economic Studies* 61(208): 397–415.

Raisian, John. 1983. "Contracts, Job Experience, and Cyclical Labor Market Adjustments." *Journal of Labor Economics* 1(2): 152–170.

Rogerson, Richard. 1988. "Indivisible Labor, Lotteries and Equilibrium." *Journal of Monetary Economics* 21(1): 3–16.

Rogerson, Richard, and Johanna Wallenius. 2007. "Micro and Macro Elasticities in a Life Cycle Model With Taxes." NBER Working Paper 13017.

Shimer, Robert. 2005. "The Cyclical Behavior of Equilibrium Unemployment and Vacancies." *American Economic Review* 95(1): 24–49.

Tella, Alfred. 1964. "The Relation of Labor Force to Employment." *Industrial and Labor Relations Review* 17(3): 454–469.

Wachter, Michael L. 1977. "Intermediate Swings in Labor-Force Participation." *Brookings Papers on Economic Activity*(2): 545–576.

Comments on "Cyclical Movements along the Labor Supply Function" by Robert E. Hall

Katharine G. Abraham

I very much enjoyed reading Hall's paper. From an empirical point of view, the paper does two things. First, it develops an innovative measure of high frequency shifts in labor demand. Second, given this monthly index of labor demand, it decomposes the total response of labor input to shifts in labor demand into the pieces represented by changes in labor force participation, changes in the share of people in the labor force who are employed, and changes in weekly hours. I will discuss these topics in turn.

As conceptualized in the paper, the index of labor demand is a latent variable that drives observable outcomes in the labor market. The behavior of this driving force is inferred from de-trended data series on the log of real personal disposable income per capita, the unemployment rate (measured as the negative of the log of the employment rate), and the log of weekly hours of work. Three regression equations are specified, each of which relates one of the observed variables just mentioned to the unobserved cyclical driving force (z_t), a fourth-order polynomial in time and, in the hours equation, a set of seasonal dummies. The values of the z_t variable are inferred using a set of moment conditions. These implied values are somewhat surprising, in the sense that they do not conform especially well to what most people think they know about the relative severity of recessions over the last few decades. In particular, the behavior of this index of labor demand is virtually the same in the recession of the early 1990s as in the recession of the mid-1970s, whereas conventional measures show the mid-1970s recession to have been much more severe than the recession of the early 1990s. This leads me to wonder what it

is about the way in which the labor demand index is constructed that produces this result. One observation I would make is that the index of labor demand in essence reflects the behavior of the residuals of the three observable variables—personal income per capita, unemployment, and weekly hours—from a quadratic trend. The quadratic trend may pick up not only the longer-term influences on the observable variables it is intended to capture, but also movements that more properly could be considered cyclical.

A second observation relates to the use of personal disposable income in the system of equations from which the labor demand index is derived. Given the purpose for which the system is estimated, it would seem that the dependent variables in the individual equations should be outcomes that depend rather directly on the demand for labor. Much of personal disposable income can be categorized as labor income, but personal disposable income also includes the part of proprietor's income that represents returns to capital as well as a substantial amount of government transfer payments. I wonder whether it might not have been better to construct the system using a measure of employee compensation as a dependent variable in place of the measure of personal disposable income.

The most serious question I have about the system of equations Hall used to estimate the labor demand index, however, is the presence of separate equations for employment and hours. Thinking about how adjustments to changes in the demand for labor occur, employment and hours can be viewed as substitutes for one another: if the demand for labor falls, firms may reduce employment, cut workers's hours, or a combination of both. The point here is that employment and hours need not respond in the same way to changes in the demand for labor. With the estimating equations as specified, however, it is the co-movements in these variables' responses that will be reflected in the index of labor demand. This then leads me to wonder whether it might not have been better to use a measure of total hours, rather than separate measures of employment and weekly hours, in the estimating equations.

Let me turn now to the second part of the paper, which looks at the responsiveness of participation, employment, and weekly hours of work to changes in the index of labor demand. The equations used to estimate these responses are specified in a rather parsimonious fashion. For starters, the model makes no allowance for lags in the response of any of the three labor input variables to changes in labor demand. It also makes no allowance for changes over time in the way in which the different labor input variables respond to changes in labor demand. As an aside, I might also mention that one needs to be careful in characterizing the estimates obtained from the model. The paper is very careful in its choice of words—talking about labor input rather than labor supply—but this is a conference on labor supply and, to the extent that decisions about hiring and hours are made by the employer rather than by the worker, it should be emphasized that the outcomes that are observed cannot be interpreted as labor supply responses, but rather can be traced to labor demand.

The main findings in this part of the paper decompose the response of labor input to changes in labor demand into shares due to changes in labor force participation, changes in the (un)employment rate, and changes in hours of work. Changes in labor force participation are relatively unimportant. A suggestion to be made here is that, in carrying out this sort of decomposition exercise, there could be value in thinking about where the boundary between being in and being out of the labor force is drawn. In the official data series that underlie the estimates reported in the paper, people are considered to be in the labor force if they are either counted as working or as being unemployed, with the unemployed including people who want to work, are available for work, and have actively looked for work in the last four weeks. But this boundary could be drawn differently, with either more or fewer people counted as belonging to the labor force and, at least in principle, this could affect the results obtained. The Bureau of Labor Statistics produces several alternate measures that draw the boundary somewhat differently – for example, changing the position of the boundary to include people who searched for work in the past year but not in the past four weeks as belonging to the labor force. In practice, the cyclical behavior of these alternate measures appears to be very similar to the cyclical behavior of the official unemployment rate, which makes me think this is probably not a major issue for the analysis. Still, this seems like something that might merit a bit more careful exploration.

Most of the response of labor input to changes in the index of labor demand is attributable to changes in the (un)employment rate (the share

of the labor force that is (un)employed), and changes in weekly hours. At least to my eyes, in Hall's paper, changes in weekly hours account for a large share of this response. My reading of the message to be taken from the literature on the dynamics of labor demand is that, in the United States, adjustment to changing labor demand conditions tends to occur through changes in employment rather than through changes in weekly hours. In the results reported here, however, fluctuations in weekly hours account for about a third of the overall response. The framework here is rather different than the framework used in the labor economics literature with which I am more familiar, but I nonetheless find this paper's results a bit surprising. In the remainder of my comments, I would like to consider why hours adjustment appears to be so important here.

In part, I suspect, these findings may reflect the way that the labor demand index is constructed. As I have already noted, by construction, the labor demand index reflects only the common movements in income, employment, and hours. I worry that the approach adopted in this paper builds in an association between weekly hours worked and the labor demand index that would not be there if, for example, a measure of total hours had been used in place of the separate employment and hours measures.

A second issue, discussed at some length in the paper, is that the measure of employment derived from the Current Population Survey (CPS) on which the paper rests is much less cyclical than the measure derived from the monthly employer payroll survey. One might suspect that the different behavior of the two measures could be explained by their different conceptual underpinnings—for example, the inclusion of the self-employed in the CPS measure but not in the payroll survey measure—but accounting for these conceptual differences in fact does little to make the two series more comparable. If the payroll survey does a better job of measuring employment than the household survey, then, properly measured, the responsiveness of hours is relatively less important than implied by the estimates reported here.

Which measure of employment—the CPS measure or the payroll survey measure —is more believable? To answer this question, we need first to understand why the cyclical behavior of the two employment measures has been so different. A variety of explanations have been proposed, but

at this point there is no obvious "smoking gun." Potential explanations that seem to me to merit further investigation include possible problems with the classification of people as self-employed versus wage-and-salary workers in the CPS, cyclical movements in the amount of "off-the-books" employment recorded in the CPS that do not show up in the payroll survey, and problems with the CPS population controls related to difficulties in accounting for immigration—but I cannot tell you that any of these potential explanations will end up offering an answer. In the meantime, the uncertainty about which employment series should be believed makes it difficult to know how much confidence to place in this paper's estimates.

In summary, Hall's paper offers a creative and interesting approach to the measurement of fluctuations in labor demand and their effects on the labor market. At this point, I have a number of questions about the results obtained. My hope would be that Hall's further work on this issue will provide a better basis for assessing the robustness of the conclusions reported here.

Comments on "Cyclical Movements along the Labor Supply Function" by Robert E. Hall

Susanto Basu

Bob Hall is famous for working in different areas of macroeconomics, making important contributions in each one, and then forsaking his current area of research to work on other issues. But there is one topic to which he has returned again and again since his earliest days as an academic, understanding the behavior of unemployment and hours worked over the business cycle. The last few years have seen Hall produce a burst of papers that collectively deepen our understanding of this central issue in macroeconomics and labor economics. The present paper is another step in this important research program.

Let us review the facts that Hall takes as his starting point. Over the business cycle, we see large changes in employment and total hours worked, with relatively small changes in real wages. If one adopts the perspective that Hall takes in this paper, namely that business cycles are due to shifts in labor demand along a stable labor supply curve, these facts suggest that the labor supply curve must be quite flat. However, microeconomic estimates of labor supply elasticities suggest that the elasticity of hours worked with respect to wages for continuously employed workers is small—that is, the labor supply curve for employed workers is steep, not flat.

But of course, changes in hours worked by the employed account for only a small fraction (perhaps 15 percent) of the cyclical variation in total hours worked. The vast majority of the variation in labor supply over the business cycle comes from changes in the number of people employed. Together, these facts suggest that the underpinnings of "macro labor supply"—the supply of total hours worked—are quite different from those

assumptions typically used to explain micro labor supply, which involve utility-maximizing choices made by individuals in a frictionless, neoclassical framework.

Hall thus moves to a framework where the determinants of micro and macro labor supply are different. In the paper that derives formally the framework explained heuristically in this paper, Hall (2008) shows how a variant of the popular Mortensen-Pissarides (MP) search-matching model yields a labor supply function of the form

$$(1) \quad L = H(\lambda, w)N(\lambda, w)$$

where L is total hours worked, H is hours per worker, N is the number of workers, w is the real wage, and λ is the marginal utility of wealth (assumed equal for everyone due to perfect consumption insurance). The supply of hours for employed workers takes the standard Frisch form: it is a function of the real wage and (expected) lifetime wealth. Hall's striking achievement is to show that the number of workers available for employment is also a function of the same two variables. In the same paper, which I recommend highly to all who are interested in this topic, Hall shows that one can use a mix of calibration and estimation to infer the shapes of the H and N functions from household studies, plus aggregate data on consumption, employment, and hours worked. The MP model fits the facts well once one modifies the model as in Hall (2005), by assuming that the real wage is constant as long as it is within the bargaining range that is efficient for workers and firms.

While Hall's interpretation of the facts is insightful and consistent with labor market paradigms that are currently in vogue, it is useful to ask whether there are alternative models of the labor market that can also explain the basic stylized facts. And if there are indeed alternative theoretical explanations, how might one use data to discriminate empirically among these competing models?

Models based on fixed costs of going to work that are incurred by workers, such as Hansen (1985) and Rogerson (1988), try to explain the difference between micro and aggregate labor supply in a neoclassical framework. But Mulligan (2001) shows that the strong implications of these models come from the assumption in both papers that the fixed cost of going to work is identical across all workers. Once Mulligan allows for a distribution of this parameter, he finds that there are few implications of fixed costs per se, and none that are significant for labor supply over the business cycle.

The early Keynesian interpretation of the cyclical facts that motivate Hall also differentiates between micro and macro labor supply. But where Hall modifies the neoclassical model by introducing the information friction central to the MP model, the Keynesian story dismisses the neoclassical paradigm altogether, at least for the purpose of understanding short-run aggregate fluctuations in the labor market. In the Keynesian framework, workers agree to work as much as employers demand at a pre-set nominal wage. As in the model that Hall develops, the effective labor supply curve in the (Old) Keynesian model is indeed flat—the pre-set wage is independent of employment—and fluctuations in employment and hours are determined by labor demand. If prices are approximately as sticky as wages, then real wages change relatively little over the business cycle, matching what the data show.

One might object that the Keynesian framework does not determine the split of total hours worked between overall employment and hours per worker. But a small change in the set-up that introduces fixed per-worker costs incurred by the firm—for example, benefits like health insurance, the value of which often is independent of the number of hours worked—would remedy this problem.

The more fundamental problem that Hall sees in the Keynesian framework is that its predictions are not derived from optimal decision-making by workers and firms. In a paper that has been central to Hall's thinking on these issues, Barro (1977) pointed out that the contracts assumed in sticky-wage models are not Pareto-efficient—that is, both firms and workers would gain by renegotiating to reduce wages in a downturn instead of cutting employment. Hall has adopted Barro's view that even if observed nominal (or real) wages are sticky, the decisions on employment and hours are made in accordance with a long-term implicit contract between workers and firms in which wages are fully flexible, with the observed spot wages being "installment payments" of the agreed-upon total lifetime wages due workers.

A different set of models, based on efficiency-wage considerations, also has implications for the difference between micro and macro labor supply. As an example, consider the efficiency-wage model of Shapiro and Stiglitz (1984). Unlike the Keynesian model, the Shapiro-Stiglitz model assumes optimizing behavior by firms and workers. In fact, the famous labor-market diagram summarizing the predictions of the model bears a striking resemblance to Figure 5.3 of Hall's paper. Both models predict equilibrium unemployment. Both display steep supply functions for total hours in the absence of frictions, but a relatively flat effective labor supply curve in the environment with frictions. In Hall's model the friction is an informational one—people do not automatically know where to find a good match for their skills, and must invest in job search. In the Shapiro-Stiglitz model, the friction is imperfect monitoring—workers and firms contract over the number of hours spent on the job, but the firm cannot contract over how hard the worker works each hour, since effort is not observed perfectly. Is one friction clearly more important than the other? The answer is not obvious to me. But macro labor research in recent years has focused almost exclusively on the MP model and its variants, and ignored the Shapiro-Stiglitz model.

Which of these four classes of models best explains the data? One would hope that such questions would be settled by confronting these competing models with data in a systematic fashion. Unfortunately, few such efforts have been made, and most are tests of a single model against an unspecified alternative. The failure to test the models against one another is due partly to the fact that these models explain different features of the data. A major strength of the Mortensen-Pissarides search-matching framework, for example, is its ability to explain data on worker flows and job vacancies, but the other models have little or nothing to say about such issues. On the other hand, all the models of the labor market discussed above are able to match the key stylized facts of the data—in most cases, they were created to do so! Unfortunately, beyond matching these facts, the models make surprisingly few empirical predictions, and the predictions they make are often ambiguous and subject to varying interpretations. A good example is the literature on efficiency-wage models and inter-industry wage differentials. Efficiency-wage models do predict that workers in some industries might be paid higher wages than identical workers in other industries, but it is impossible to rule out the possibility that workers in different industries who appear identical to the econometrician actually differ in their labor market characteristics.

Thus, in practice the choice between the matching-based framework that Hall advocates and the other models is guided as much by aesthetic considerations (what makes a good model?) and intuitive ones (what are the basic institutions and frictions in the labor market?) as by any formal empirical testing.

These are my comments on Hall's current research agenda in labor economics, of which the current paper is a part. Now let me turn to the novel contribution of the conference paper at hand. Relative to Hall (2008), the contribution in this paper is to derive an index of labor demand, and study how each component of labor input responds to labor demand. Hall is very clear in noting that his exercise is possible only under the assumption that the labor supply curve is stable over the business cycle. In terms of equation, the assumption is that changes in λ happen at low frequencies, but are not relevant for business cycles.

Unfortunately, this assumption does not hold for an important category of shocks, namely shocks to government expenditure. The empirical literature on these shocks, summarized in Perotti (2007), shows that these increase output and hours worked at high frequencies. Even in the post-World War II period, there have been large, exogenous changes in U.S. government expenditure, usually associated with national security crises, notably the Korean and Vietnam Wars. But in the neoclassical labor supply framework in which Hall operates, spending shocks increase output and hours worked by raising λ, that is, by making consumers feel poorer which results in them working (and producing) more.

Another way of making this point is to note that in a neoclassical framework with distortionary taxes, the labor demand curve can be expressed as:

$$(2) \quad w = (1 - \tau) A F_L(K, AL),$$

where τ is the (labor income) tax rate, F is a production function, K is capital, A is the state of technology, and the L subscript denotes the marginal

product of labor. Since technology does not change in response to a fiscal shock and capital, a state variable, is essentially unchanged as well, the labor demand curve in (w, L) space shifts only if the tax rate changes. But one would think that current tax rates would rise in response to an increase in government spending, since the government requires more revenue, and indeed the empirical work confirms this conjecture. In this situation, then the labor demand curve must shift inward. So if output and hours must rise to match the empirical evidence, then an outward shift in labor supply must be responsible for more than 100 percent of the increase in hours worked after an increase in government purchases.

In some non-neoclassical equilibrium models with imperfect competition (for example, Rotemberg and Woodford 1992), labor demand can shift out due to a fall in the mark-up of price over marginal cost. Even so, the mark-up changes in response to a change in output, and output would not change unless the labor supply curve shifts outward.

To confirm that government purchase shocks do indeed have their expected effects on unemployment and hours per worker, I took the preferred series of government purchase shocks from Perotti (2007). (Roberto Perotti kindly supplied the data.) Figure 5.7 shows the effects of government purchase shocks on employment and hours per worker. As one would surmise, an unexpected increase in government purchases increases weekly hours and lowers unemployment. Some of the fluctuations are quite large, on the order of a half-percentage point change in the unemployment rate. Thus, under the maintained hypothesis that Hall's framework provides the correct interpretation, the evidence shows that it is not safe to assume that changes in λ are negligible at high frequencies.

In sum, Bob Hall is revolutionizing our understanding of one of the most perplexing issues in macroeconomics, the behavior of labor supply over the course of the business cycle. I remain an interested observer. Since I am unconvinced that a lack of information is the most important friction preventing the labor market from functioning smoothly over the business cycle, I am somewhat skeptical that this research program will attain all its promised objectives, but I remain hopeful. However, I am fairly sure that this particular paper will not contribute in any important way to the attainment of the desired outcome.

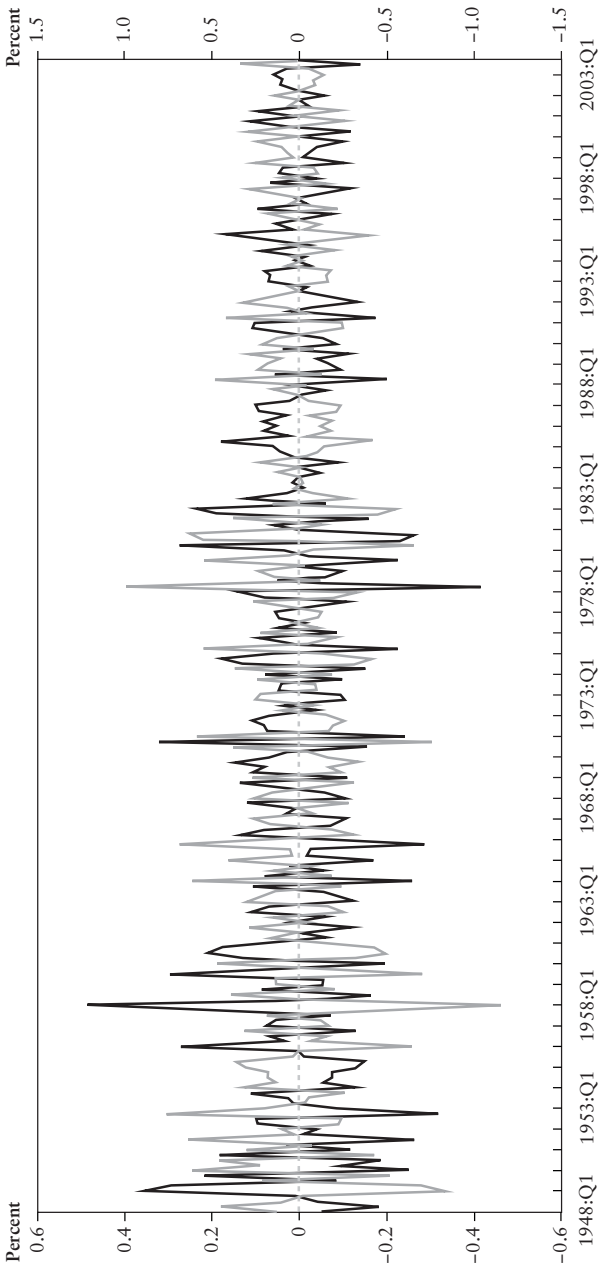

Figure 5.7
Changes in Unemployment and Weekly Hours Due to Shocks in Government Spending
Source: Blanchard and Perotti (2002).

References

Barro, Robert J. 1977. "Long-Term Contracting, Sticky Prices, and Monetary Policy." *Journal of Monetary Economics* 3(3): 305–316.

Blanchard, Olivier, and Robert Perotti. 2002. "An Empirical Characterization of the Dynamic Effects of Changes in Government Spending and Taxes on Output." *The Quarterly Journal of Economics* 117(4): 1329–1368.

Hall, Robert E. 2005. "Employment Fluctuations with Equilibrium Wage Stickiness." *American Economic Review* 95(1): 50–65.

Hall, Robert E. 2008. "Sources and Mechanisms of Cyclical Fluctuations in the Labor Market." Available at http://www.stanford.edu/~rehall/SM022208.pdf.

Hansen, Gary D. 1985. "Indivisible Labor and the Business Cycle." *Journal of Monetary Economics* 16(3): 309–327.

Mulligan, Casey B. 2001. "Aggregate Implications of Indivisible Labor." Advances in Macroeconomics 1(1). Available at: http://www.bepress.com/bejm/advances/vol1/iss1/art4.

Perotti, Roberto. 2007. "In Search of the Transmission Mechanism of Fiscal Policy." NBER Working Paper 13143.

Rogerson, Richard. 1988. "Indivisible Labor, Lotteries and Equilibrium," *Journal of Monetary Economics* 21(1): 3–16.

Rotemberg, Julio J., and Michael Woodford. 1992. "Oligopolistic Pricing and the Effects of Aggregate Demand on Economic Activity." *Journal of Political Economy* 100(6): 1153–1207.

Shapiro, Carl, and Joseph E. Stiglitz. 1984. "Equilibrium Unemployment as a Worker Discipline Device." *American Economic Review* 74(3): 433–444.

6

Labor Supply and Labor Demand in the Long Run

U.S. Labor Supply and Demand in the Long Run

Dale W. Jorgenson, Richard J. Goettle, Mun S. Ho, Daniel T. Slesnick, and Peter J. Wilcoxen

I. Introduction

In this paper we model U.S. labor supply and demand in considerable detail in order to capture the enormous heterogeneity of the labor force and its evolution over the next 25 years. We represent labor supplies for a large number of demographic groups as responses to prices of leisure and consumption goods and services. The price of leisure is an after-tax wage rate, while the final prices of goods and services reflect the supply prices of the industries that produce them. By including demographic characteristics among the determinants of household preferences, we incorporate the expected demographic transition into our long-run projections of the U.S. labor market.

The U.S. population will be growing older over the next quarter-century, and elderly households have very different patterns of labor supply and consumption compared to their younger counterparts. Our projections for the period spanning 2004 to 2030 thus incorporate the expected fall in the supply of labor per capita. These changes in labor supply patterns are the consequence of population aging, rather than wage and income effects. Despite the anticipated aging of the U.S. population, moderate population growth will provide growing supplies of labor well into the twenty-first century. Improvements in the quality of U.S. labor input, defined as increased average levels of educational attainment and experience, will also continue for some time, but will gradually disappear over the next quarter-century.

We represent labor demand for each of 35 industrial sectors of the U.S. economy as a response to the prices of productive inputs—labor, capital,

and intermediate goods and services. In addition, labor demand is driven by changes in technology. Technical change generates productivity growth within each industry. Rates of productivity growth differ widely among industries, ranging from the blistering pace of advance in computers and electronic components to the gradual decline in construction and petroleum refining. In addition, changes in technology may be skill-biased. Labor-saving technical change reduces labor demand for given input prices, while labor-using technological change increases labor demand. Over the next 25 years, productivity growth for the U.S. economy as a whole will be below long-term historical averages. However, productivity growth in information technology equipment and software will continue to outpace productivity growth in the rest of the economy. The output of the U.S. economy will continue to shift toward industries with high rates of productivity growth. Labor input biases of technical change are substantial in many industries. Labor-using, rather than labor-saving, biases predominate. Labor-using technical change will continue to be a stimulus to the growth of labor demand, and differences in the biases for different industries will play an important role in the reallocation of labor.

We incorporate the determinants of long-term labor supply and demand into a model of U.S. economic growth. We refer to this model as the IGEM,[1] which stands for the Inter-temporal General Equilibrium Model. Markets for labor, capital, and the aggregate output of the economy equilibrate through the price system at each point of time. In the labor market, for example, wage rates determine the labor supplied by the current population and the labor quantity demanded by employers in the many sectors of the economy. In the IGEM model and in the U.S. economy, year-to-year changes in the level of economic activity are primarily the consequence of the accumulation of capital. However, over the next quarter-century the driving forces of economic growth are projected to be demography and technology—as encapsulated in the neoclassical theory of economic growth.

In the IGEM, capital formation is determined by the equilibration of saving and investment. We model household saving at the level of the individual household. Consumption, labor supply, and saving for each household are chosen to maximize a utility function, defined as the stream of future consumption of goods and leisure, subject to an inter-

temporal budget constraint. The forward-looking character of savings decisions allows changes in future prices and rates of return to affect the current labor supply. The availability of capital input in the U.S. economy is the consequence of past investment. This backward-looking feature of capital accumulation links current markets of capital input to past investment decisions.

II. A Long-Run Model of the U.S. Economy and the U.S. Labor Market

Our household model generates demand for a detailed list of personal consumption expenditures given in Table 6.1. Household preferences are structured in a nested, or tiered, manner. At the top tier, utility is a function of non-durable goods, capital services, consumer services, and leisure. Lower tiers allocate non-durable goods to specific categories, like food and clothing, and consumer services to transportation, finance, and other services. Household consumption patterns for goods and leisure are derived from the Consumer Expenditure Survey (CEX).[2] The items in Table 6.1 are based on the consumption categories in the National Income and Product Accounts (NIPAs).[3] These items are linked to the supplying industries listed in Table 6.2.

As the owner of the economy's wealth, the household sector makes a second contribution to the demand side of the economy through the demand for investment goods. Household sector savings are allocated between domestic and foreign investment, and the domestic portion is distributed among investments in assets such as building structures, capital equipment, consumer durables, and inventories. Capital stocks and capital services are derived primarily from the Fixed Asset Accounts of the Bureau of Economic Analysis,[4] which include information on investment by 60 asset categories. Data on labor input by industry are derived from detailed demographic and wage data in the annual Current Population Surveys and the decennial Censuses of Population, as described by Jorgenson, Ho, and Stiroh (2005).

We separate the production sector in the IGEM into 35 individual industries. The complete list is given in Table 6.2, together with the value of each industry's output in 2000 and the corresponding Standard Industrial Classification codes. Each industry produces output from labor,

Table 6.1
Personal Consumption Expenditures and leisure, IGEM categories, 2000.
Leisure, IGEM Categories, 2000

	IGEM categories	Billions of dollars	Category
1	Food	568.6	3
2	Meals	376.5	4
3	Meals-Employees	9.9	5,6
4	Shoes	46.3	12
5	Clothing	267.4	14,15,16
6	Gasoline	164.4	75
7	Coal	0.2	40
8	Fuel oil	17.9	40
9	Tobacco	72.2	7
10	Cleaning supplies	115.8	21,34
11	Furnishings	38.3	33
12	Drugs	156.3	45
13	Toys	62.7	89
14	Stationery	23.4	35
15	Imports (travel)	3.3	111
16	Reading	51.7	88,95
17	Rental	247.4	25,27
18	Electricity	101.5	37
19	Gas	40.8	38
20	Water	48.8	39
21	Communications	130.6	41
22	Domestic service	16	42
23	Other household	48.5	43
24	Own transportation	210.8	74,76,77
25	Transportation	56.9	79,80,82,83,84,85
26	Medical Services	921.3	47,48,49,51,55
27	Health Insurance	70.6	56
28	Personal services	76.2	17,19,22
29	Financial services	517.7	61,62,63,64
30	Other services	114.8	65,66,67
31	Recreation	255.5	94,97,98,99,100,101,102,103
32	Education and Welfare	354.1	105,106,107,108
33	Foreign Travel	80.9	110
34	Owner maintenance	90	authors' imputation
35	Durables flow	1394.4	authors' imputation
	Leisure	13786.3	authors' imputation

Source: U.S. Bureau of Economic Analysis and Consumer Expenditure Survey (U.S. Bureau of Labor Statistics).
Note: National Income and Product Accounts Personal Consumption Expenditure category refers to the line number in Table 2.4 of Survey of Current Business 2002.

Table 6.2
Industry Output and Value Added, 2000

Code	Industry Name	Output	Value-Added	SIC
1	Agriculture	388994	195781	01-02, 07-09
2	Metal Mining	15603	7167	10
3	Coal Mining	23081	14175	11-12
4	Petroleum and Gas	136651	72669	13
5	Nonmetallic Mining	18894	10619	14
6	Construction	995279	419200	15-17
7	Food Products	487587	156127	20
8	Tobacco Products	35853	10108	21
9	Textile Mill Products	61629	21811	22
10	Apparel and Textiles	84273	62899	23
11	Lumber and Wood	115974	43305	24
12	Furniture and Fixtures	87965	39619	25
13	Paper Products	175955	72942	26
14	Printing and Publishing	233523	137723	27
15	Chemical Products	422655	183438	28
16	Petroleum Refining	235145	26422	29
17	Rubber and Plastic	170270	77459	30
18	Leather Products	10616	4028	31
19	Stone, Clay, and Glass	111040	53522	32
20	Primary Metals	191627	59691	33
21	Fabricated Metals	279540	125540	34
22	Industrial Machinery and Equipment	472251	193646	35
23	Electronic and Electric Equipment	433257	195913	36
24	Motor Vehicles	427709	83072	371
25	Other Transportation Equipment	186241	87121	372-379
26	Instruments	183293	104351	38
27	Miscellaneous Manufacturing	52715	21889	39
28	Transport and Warehouse	553535	263335	40-47
29	Communications	430330	231027	48
30	Electric Utilities	245950	166618	491, %493
31	Gas Utilities	81196	26421	492, %493, 496
32	Trade	1965715	1187180	50-59
33	FIRE	2009429	1240039	60-67
34	Services	3455269	2197343	70-87, 494-495
35	Government Enterprises	256268	167722	
36	Private Households	1394410	1394410	88
38	General Government	1194160	1194160	

Source: U.S. Bureau of Economic Analysis and U.S. Census Bureau.
Note: All figures in millions of current dollars. % indicates part of an SIC code.

capital, and intermediate inputs, using a technology that allows for substitution among these inputs. Although technology can be represented by means of a production function, we find it much more convenient to use a dual approach, based on a price function that gives each sector's output price as a function of its input prices. Technologies are structured in a nested or tiered manner, with intermediate inputs divided between energy and materials; both energy and materials are further subdivided among inputs that correspond to the 35 commodity groups produced by the 35 industries.

Our representation of the technology embedded in each sector includes its respective rate and biases of technical change. The rate of technical change captures improvements in productivity or growth in output per unit of input. The biases of technical change correspond to increases or decreases in the shares of inputs in the value of output, holding input prices constant. The evolution of patterns of production reflects both price-induced substitution among inputs and the impact of changes in technology. We project the historical patterns of technical change represented in our database in order to incorporate future changes in technology into the demand for inputs of labor, capital, and intermediate goods and services.

The production of each commodity by one or more of the 35 U.S. domestic industries is augmented by imports of that commodity from the rest of the world to generate the U.S. domestic supply of goods and services. This supply is allocated to U.S. industries as an intermediate input and to final demand for consumption by U.S. households and governments; investments by U.S. businesses, households, and governments; and net exports. Since imports are not perfect substitutes for commodities produced domestically, we also explicitly model the substitution between imports and domestic production. The rest of the world absorbs exports from the United States, and the net flow of resources in each period is governed by an exogenously specified current account deficit.

The final sector explicitly considered in our model is the government sector, which taxes, spends, and makes transfer payments. Public consumption of goods and services is one component of final demand, while public sector borrowing is one of the uses of private savings. The flow of goods and factors among the four sectors of the U.S. economy—household, business firms, government, and net exports to the rest of the world—is illustrated in Figure 6.1. Prices adjust to equate the supply from domestic and foreign producers to the demand from households, investors, government, and exports in each period.

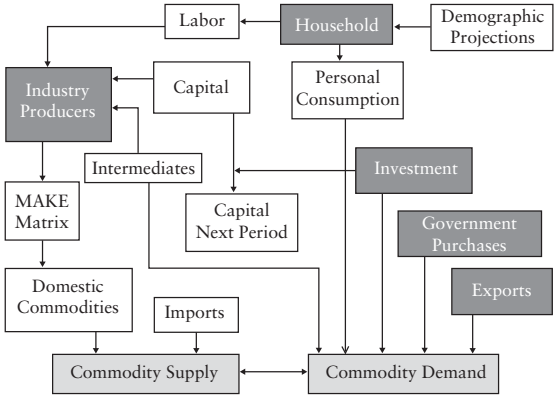

Figure 6.1
Flow of Goods and Factors in the IGEM

Our model of the U.S. economy is implemented econometrically. Parameters describing the behavior of producers and consumers are estimated statistically from a data set that we have constructed specifically for this purpose. These data are based on a new system of national accounts that integrates the wealth accounts with the National Income and Product Accounts.[5] The capital accounts include investment goods, capital services, capital stocks, and the corresponding prices. These data are described in detail by Jorgenson, Ho, and Stiroh (2005). Similar data have recently been released for members of the European Union by the EU KLEMS project.[6]

III. Exogenous Variables in the Projections

Our model of the U.S. economy simulates the future growth and structure of the economy over the intermediate term of 25 years. Of course, our model's time path of outcomes is conditional on projections of exogenous variables. Among the most important of these variables are the

total population, the time endowment of the working-age population, the overall government deficit, the current account deficit, world price levels, and U.S. government tax policies. Many of these variables are developed from published sources, "official" and otherwise. In addition, we project the evolution of technology in each of the 35 industries that make up the model's production sector. These variables are projected from the historical data set that underlies the production model and its estimation.

The key exogenous variables that describe the growth and composition of the U.S. population are population projections by sex and individual year of age from the U.S. Census Bureau.[7] During the sample period the U.S. population is allocated to educational attainment categories using data from the Current Population Survey[8] in a way that is parallel to our calculation of labor input. Each adult is given a time endowment of 14 hours a day to be used for work and leisure. The number of hours for each sex-age-education category is weighted by labor compensation rates and aggregated to form the national time endowment presented in Figure 6.2.

Our projections use Census Bureau forecasts by sex and age. We assume that the educational attainment of those aged 35 or younger will be the same in the projected period as in the last year of the historical sample period; that is, a person who becomes 22 years old in 2014 will have the same chance of having a bachelor's degree as a person in 2004. Those aged 55 years and over carry their educational attainment with them as they age; that is, the educational distribution of 70-year-olds in 2014 is the same as that of 60-year-olds in 2004. Those between 35 and 55 years of age have a complex adjustment that is a mixture of these two assumptions to allow a smooth improvement of educational attainment that is consistent with the observed profile in 2004. The result of these calculations, shown in Figure 6.2, is that the U.S. population is expected to grow at just under 1 percent per year through 2030, reaching a level slightly in excess of 365 million inhabitants. The gradually slowing improvement in the average level of educational attainment implies that the time endowment grows at a modestly faster rate of around 1 percent through 2030.

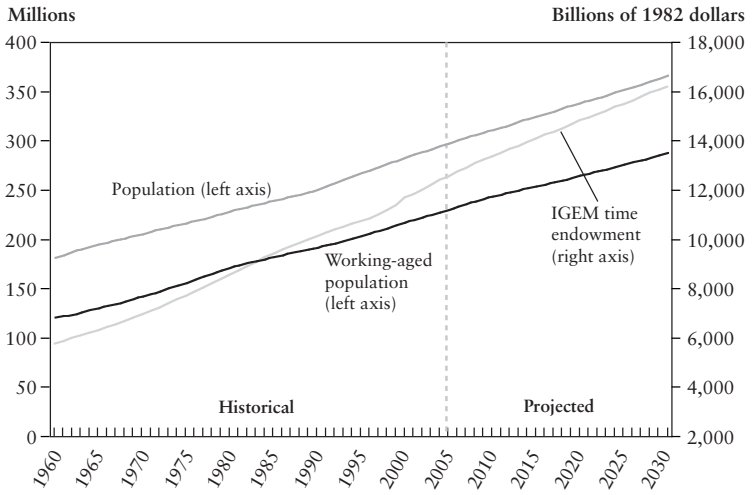

Figure 6.2
Population and Household Time Endowment for the United States
Source: Consumer Expenditure Survey of U.S. Bureau of Labor Statistics and the Current Population Survey from the U.S. Census.

We project productivity growth for each of the 35 industries, using the state-space approach of Jin and Jorgenson (2007). To illustrate this approach, Figure 6.3 gives historical data for the period 1960–2004, based on the estimates of Jorgenson, Ho, Samuels, and Stiroh (2007). These data update and revise the estimates of Jorgenson, Ho, and Stiroh (2005). Figure 6.4 presents projections of productivity growth for the period 2004–2030, using the state-space approach. Positive productivity growth reduces output prices, relative to costs of inputs, while negative growth raises output prices relative to costs.

For 2004–2030 our baseline projections reveal steadily improving productivity in 30 of the 35 sectors in the IGEM. Electrical machinery, which contains electronic components such as the semiconductor devices used in computers and telecommunications equipment, leads the list in projected productivity growth. Although this industry's projected productivity growth rate exceeds 3 percent, this represents a slight reduction in the rate of productivity growth of just under 4 percent for the

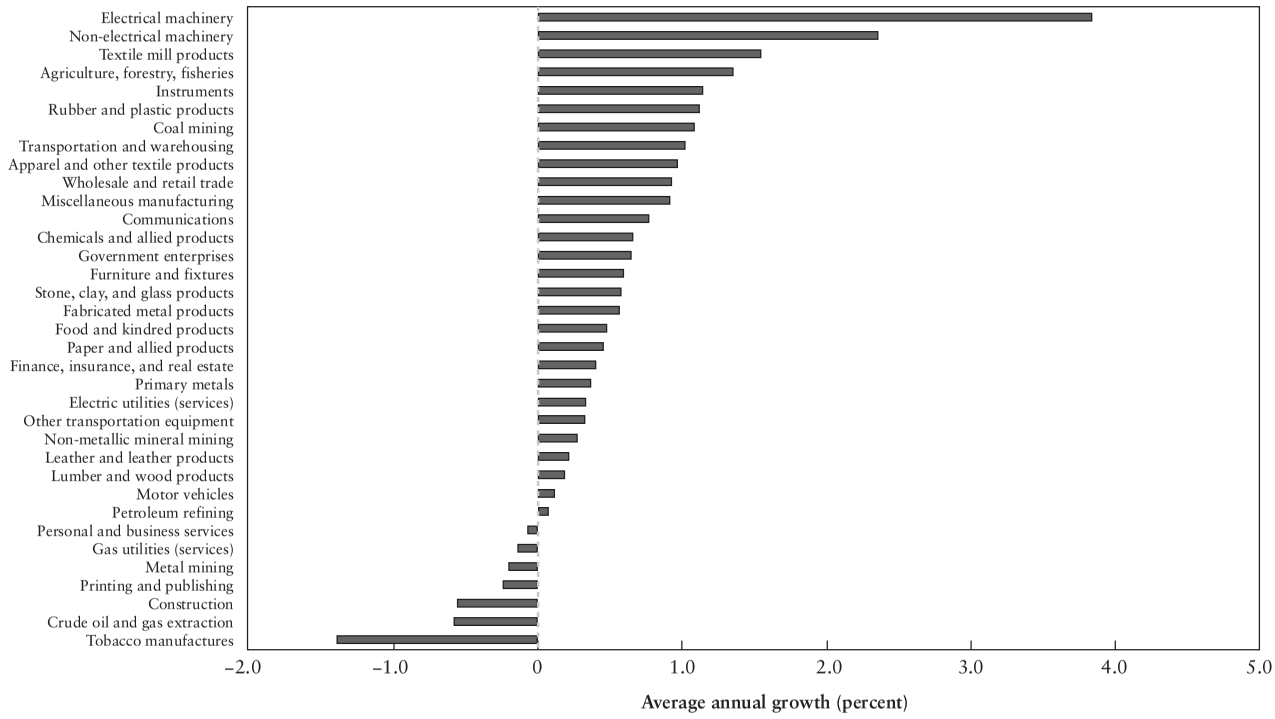

Figure 6.3
Growth in U.S. Total Factor Productivity, 1960–2004
Source: U.S. Bureau of Economic Analysis, National Income and Product Accounts.

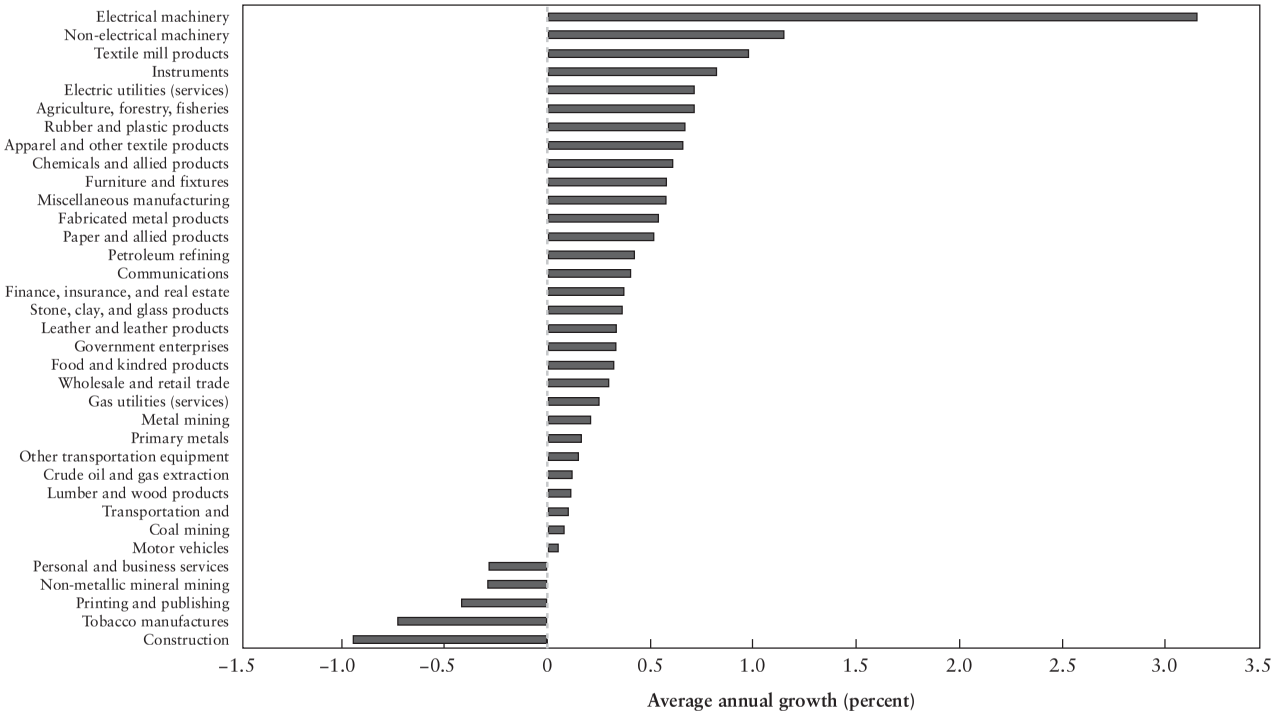

Figure 6.4
Projected Growth in U.S. Total Factor Productivity, 2004–2030
Source: U.S. Bureau of Economic Analysis, National Income and Product Accounts.

historical period 1960–2004. Non-electrical machinery, including computers, has the second highest rate of productivity growth in both the historical period and the projection period, but the projected growth rate between 2004 and 2030 is considerably lower than the historical rate.

Below we show that the overall rates of productivity growth projected for the U.S. economy are substantially below those attained for the historical period 1960–2004. It is also important to recognize productivity losses as well as productivity gains at the industry level. There are several sectors with negative projected productivity growth, including the very large construction industry and the relatively small tobacco industry. Both industries also have declining productivity during the 1960–2004 sample period.

Projections of the input biases are accomplished in a similar manner to the projections obtained for productivity. Figure 6.5 gives historical data for the period 1960–2004, while Figure 6.6 gives our projections for the period 2004–2030. Recall that the definition of skill-biased technical change is the effect of changes in technology on the share of labor input in the value of industry output, holding prices of labor input, as well as capital, energy, and materials inputs, constant. It is important to keep in mind that we have fitted and projected biases of technical change for capital, energy, and materials inputs, as well as labor input, but these are not presented in this paper due to space considerations.

During the historical sample period of 1960–2004, technical change is predominantly labor-using rather than labor-saving. Metal mining, a relatively small industry, has a very large labor-using bias of technical change, while coal mining has a large labor-saving bias. Biases of technical change differ substantially among industries, and both labor-using and labor-saving changes occur with some frequency. It is important to project rates of technical change to determine the growth rate of individual industries and the economy as a whole. However, it is also important to project biases of technical change in order to capture the impact of changes in technology on the distribution of labor input among sectors.

Two other important assumptions that determine the shape of the economy are the government and trade deficits. Our projection of the government deficit follows the forecasts of the Congressional Budget Office for the next 10 years, and then is set on course to a zero balance by 2030.[9]

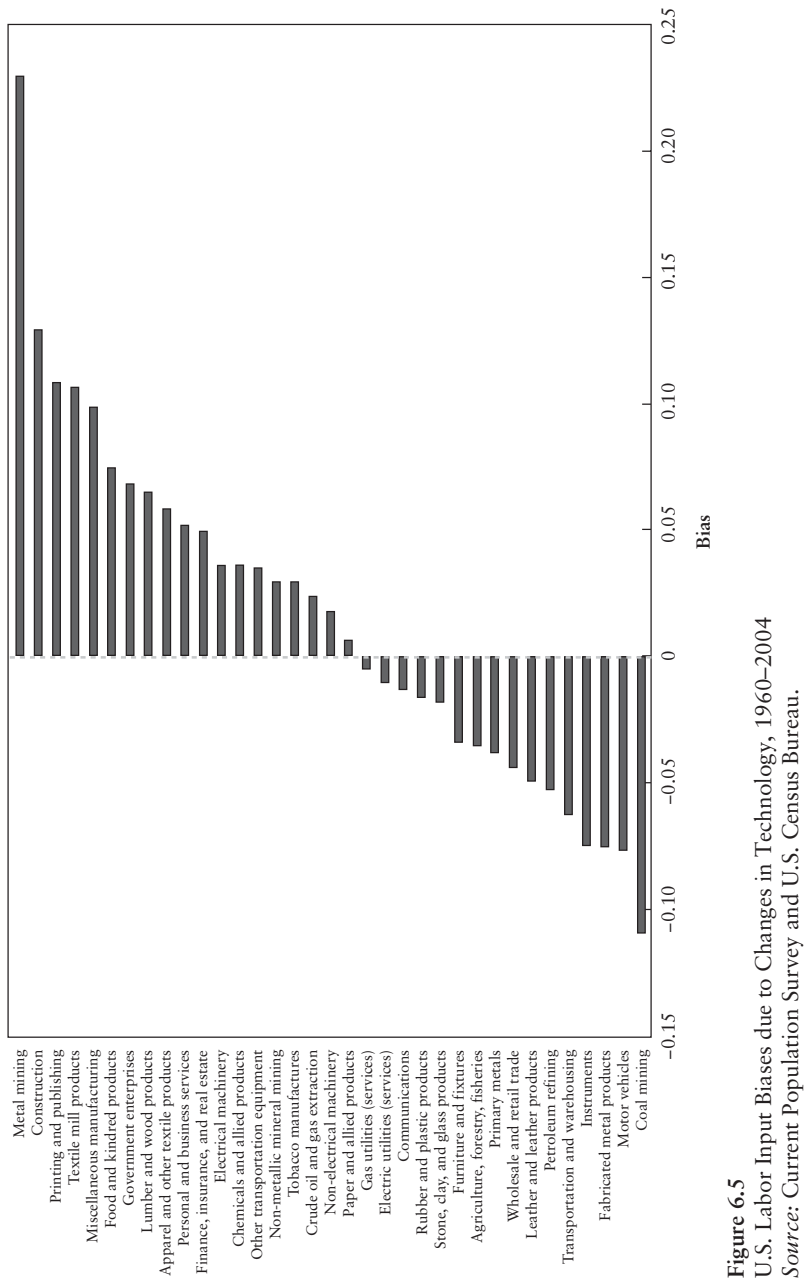

Figure 6.5
U.S. Labor Input Biases due to Changes in Technology, 1960–2004
Source: Current Population Survey and U.S. Census Bureau.

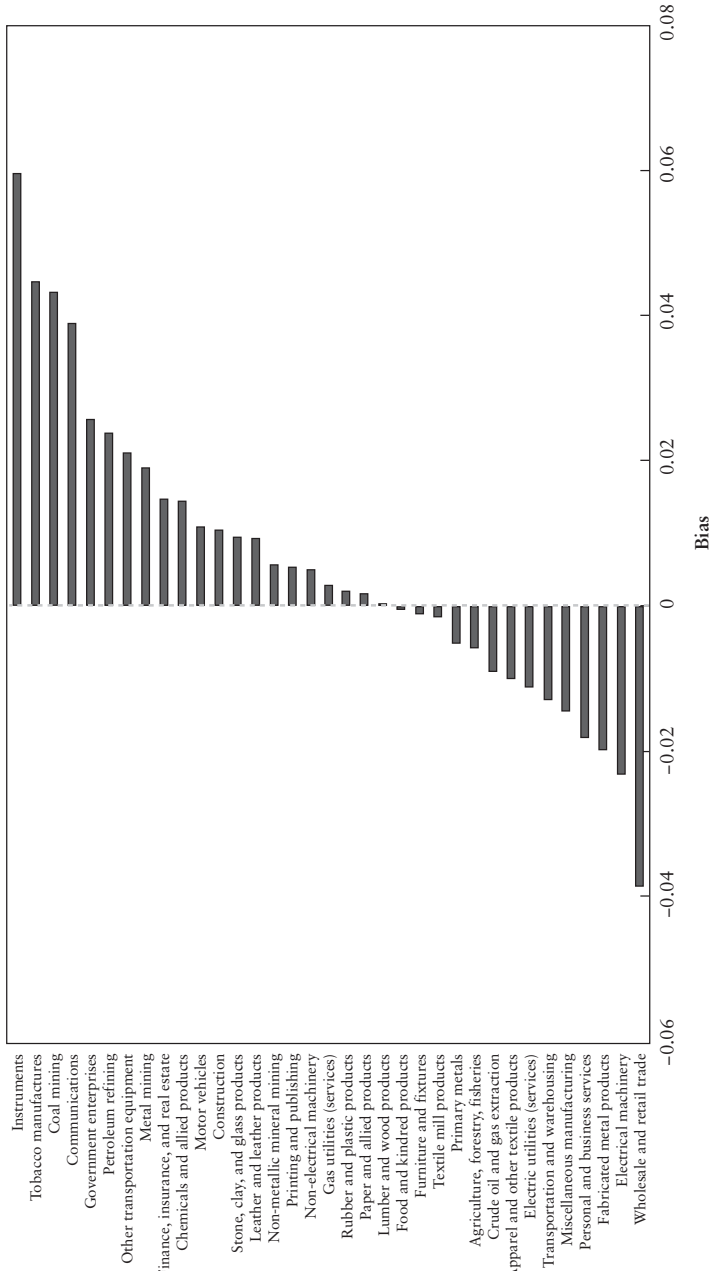

Figure 6.6
Projected U.S. Labor Input Biases due to Changes in Technology, 2004–2030
Source: Current Population Survey and U.S. Census Bureau.

The current account deficit is assumed to shrink steadily, relative to the GDP, so that it also reaches a sustainable balance by 2030. These simplifying assumptions allow the simulation to produce a smooth time path. The government and current account deficits are determinants of long-run growth to the extent that these deficits influence capital formation, but are substantially less important than the exogenous demographic and technology variables we have described.

IV. Projection of U.S. Economic Growth

Our baseline path for the U.S. economy generates a labor force participation rate, defined as the ratio of labor input to the time endowment. We have used this to extrapolate the ratio of hours worked to discretionary hours available from the working age population. The participation rate presented in Figure 6.7 reached a peak in 2000, before the shallow recession of 2001 and the "jobless" recovery that followed. The historical

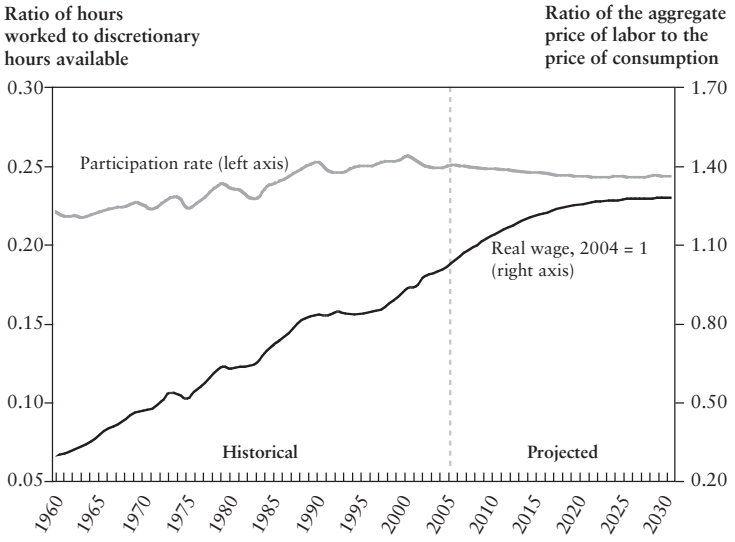

Figure 6.7
U.S. Labor Participation Rates and Real Wages, 1960–2004, with Projection to 2030
Source: Current Population Survey and U.S. Census Bureau.

data from 1960 to 1990 show substantial gains in labor force participation. No such gains in participation are in prospect for the next quarter-century. At the same time, projections beginning in 2004 do not suggest a large decline in labor force participation.

It is important to keep in mind that the rate of population growth will be declining throughout the projection period of 2004–2030. The U.S. working-age population will be growing at a very similar rate to the population as a whole during our projection period. During the 1960–2004 historical period, the working-age population grew considerably more rapidly than the U.S. population as a whole. Finally, the time endowment, which adjusts the population for changes in composition by educational attainment and labor market experience, will continue to grow more rapidly than the working-age population. However, changes in composition will gradually disappear as average levels of education and experience stabilize.

Real wages, defined as the ratio of the price of labor input to the price of consumption goods and services, are also presented in Figure 6.7. Contrary to historical trends often described in the popular business press, real wages have risen steadily throughout the postwar period with especially rapid growth rates during the period 1995–2004. Our projections of real wages rise steadily during the period 2004–2030, but at a decreasing rate. This declining rate of increase mimics the historical data from 1973–1995, prior to the U.S. growth resurgence that began around 1995 and continued into the 2000–2004 period. A slowdown in the growth rate of real wages will occur despite the continuation of historical productivity trends summarized in section III.

We next turn to the sources of U.S. economic growth during the historical and projection periods. Figure 6.8 presents historical data on the sources of U.S. economic growth during 1960–2004 recently compiled by Jorgenson, Ho, Samuels, and Stiroh (2007). The overall rate of growth is an impressive 3.34 percent per year. The most important source of growth is capital input, which contributes 1.70 percent or well over half of growth during the historical period. The next most important source of growth is labor input, which contributes 0.95 percent per year. These contributions are the growth rates of capital and labor inputs, each

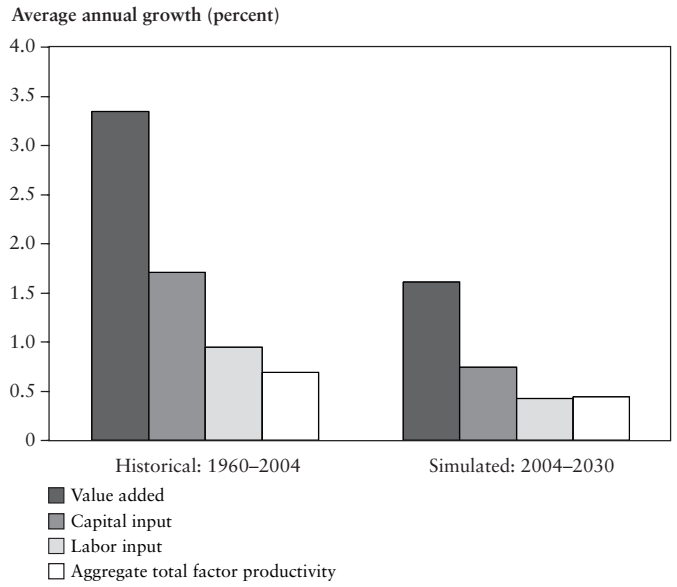

Figure 6.8
Sources of U.S. Economic Growth
Source: Current Population Survey, Fixed Asset Accounts and National Product Accounts of U.S. Bureau of Economic Analysis, and U.S. Census Bureau.

weighted by the corresponding share in the value of output. Total factor productivity growth contributes 0.69 percent per year or slightly more than 20 percent of growth during the historical period.

We project that the growth of the U.S. economy during the 2004–2030 period will be only 1.61 percent per year. The contribution of capital input will remain the most important source of growth at 0.74 percent per year. The growth of total factor productivity will decline very slightly to 0.44 percent per year, and will outstrip the sharply lower contribution of labor input of 0.42 percent. While the contributions of capital and labor inputs will still greatly predominate among the sources of U.S. economic growth, the relative importance of total factor productivity growth will jump substantially. This reflects the strength of the projected productivity trends described in section III.

We conclude our discussion of projected U.S. economic growth with a description of the growth of output and labor input at the industry level. Figure 6.9 presents growth rates of labor input for each of the 35 industries in the IGEM during the historical period 1960–2004. Slightly less than half the industries experienced an increase in labor input, led by personal and business services. However, many industries experienced sharp declines in labor input, led by leather and leather products, apparel and textile products, and gas utilities. The growth rate of labor input overall was 1.64 percent per year.

We have projected a substantial slowdown in the growth rate of labor input for the projected 2004–2030 period to 0.70 percent per year. Figure 6.10 provides a breakdown by industries. Positive growth in labor input predominates in the projections. Relatively small sectors with low projected productivity growth like tobacco and petroleum refining will show substantial increases in labor input. As widely anticipated, the large service sectors like finance, insurance, and real estate, will greatly predominate in the growth of labor input. Primary metals and metal mining will continue to release labor input to a future U.S. economy that is increasingly constrained by the slow growth of the labor supply.

Labor input biases are an important component of changes in demand for labor input. Labor-using technical change results in an increase in the share of labor input, holding prices of labor, capital, energy, and materials inputs constant. This effect dominates in our projections, as well as in the sample period. The share of labor input in instruments will increase by 0.06 during the projection period 2004–2030, reversing a similar decline in the share of labor input during the sample period 1960–2004. Metal mining, a small sector that had a large labor-using bias of technical change during the historical sample period, has a smaller labor-using bias during the projection period. Biases of technical change are an important component of labor input demand, along with the steady rise in the price of labor input relative to other inputs.

Growth in industry output completes our picture of future U.S. economic growth.

Figure 6.11 gives historical data on output growth for the period 1960–2004. Economic growth during the period 1960–2004 differed widely among industries, with a relatively narrow range of industries

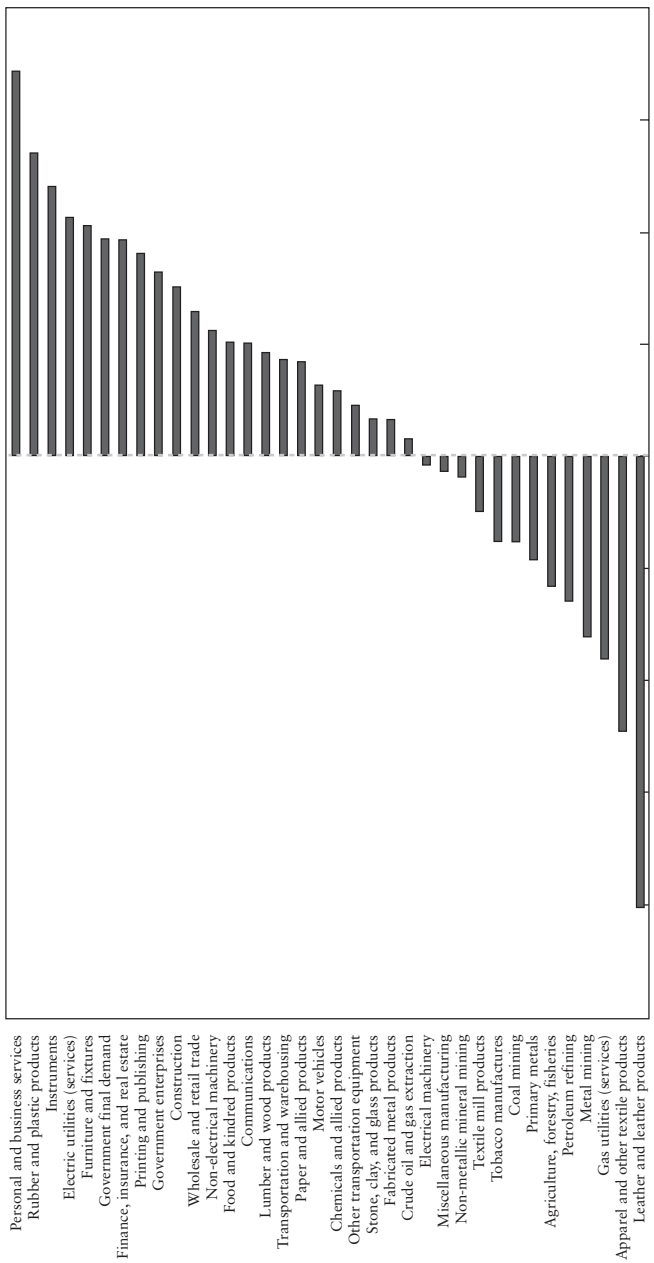

Figure 6.9
Growth in U.S. Labor Output, 1960–2004
Source: Current Population Survey and U.S. Census Bureau.

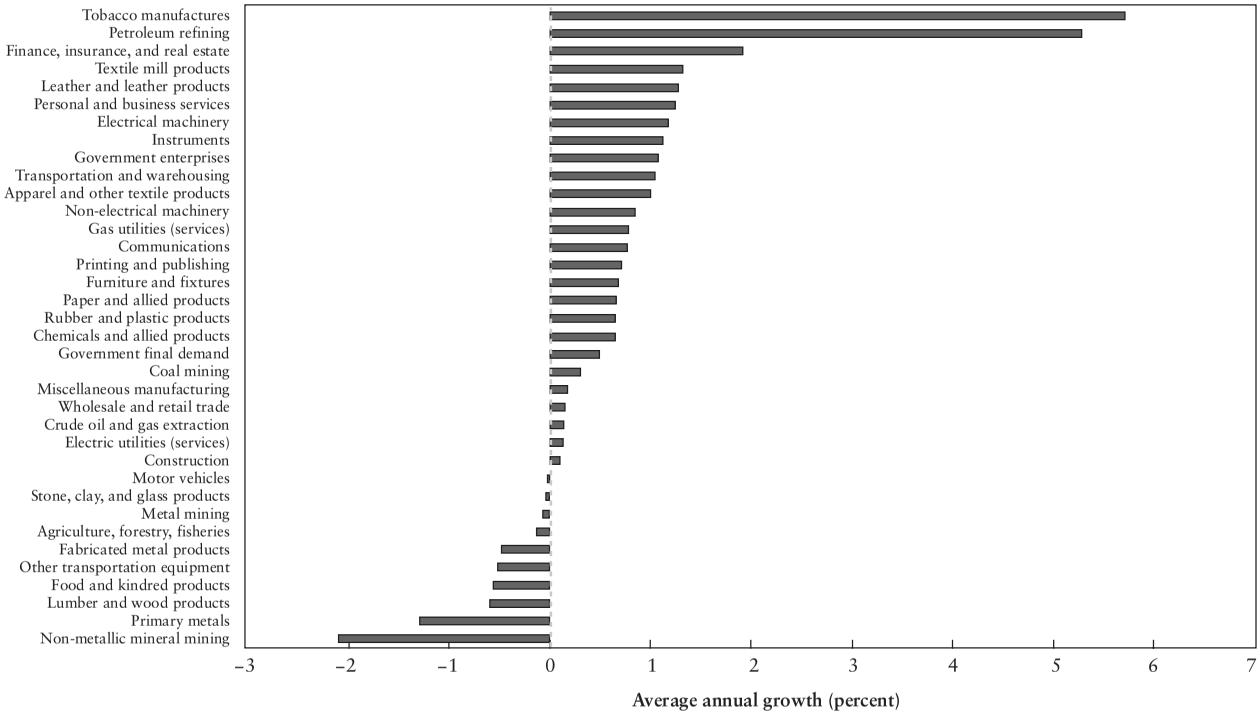

Figure 6.10
Projected Growth in U.S. Labor Output, 2004–2030
Source: Current Population Survey and U.S. Census Bureau.

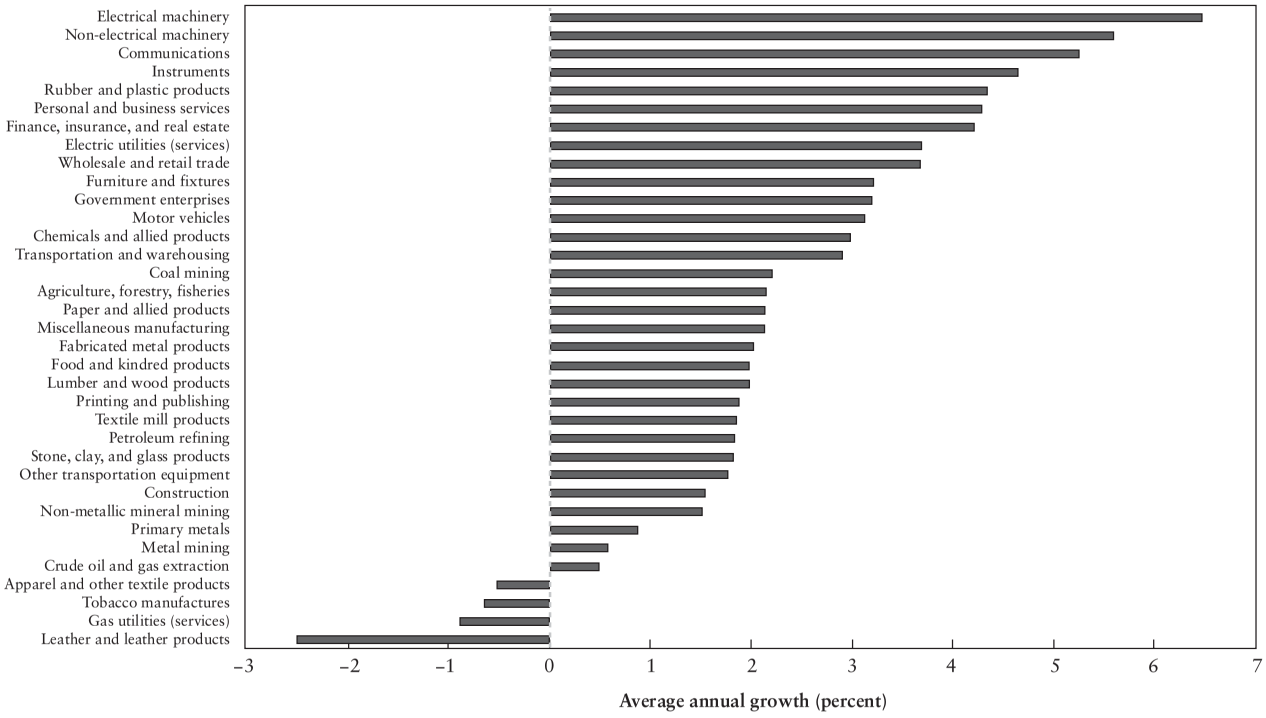

Figure 6.11
Growth in U.S. Domestic Output, 1960–2004
Source: U.S. Bureau of Economic Analysis, National Income and Product Accounts.

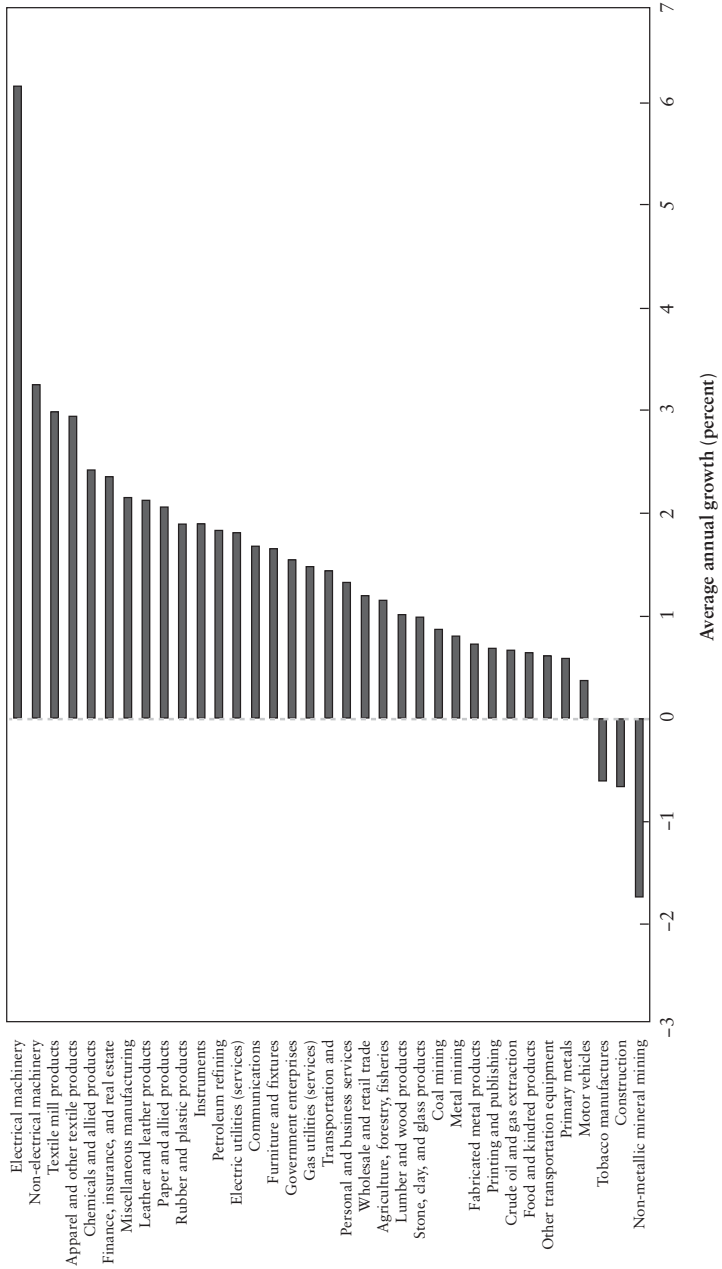

Figure 6.12
Projected Growth in U.S. Domestic Output, 2004–2030
Source: U.S. Bureau of Economic Analysis, National Income and Product Accounts.

exceeding the economy-wide average of 3.22 percent. As expected, the rapidly growing sectors were led by electrical machinery, including electronic components. Substantial growth also took place in three industries, non-electrical machinery, which contains computers; communications, the largest consuming sector for information technology equipment and software; and instruments, another major consumer. Only three industries experienced declining output growth—leather and leather products, gas utilities, and tobacco manufactures.

Figure 6.12 gives U.S. economic growth during the projection period 2004–2030.

Again, growth rates will differ substantially among industries, with electrical machinery exhibiting growth at the very rapid pace of more than 6 percent per year, comparable to the historical period of 1960–2004. Most of the remaining industries, including non-electrical machinery, one of the stars of the historical period, will scale back growth in the projection period. The relatively small leather industry will reverse the negative growth of the historical period and exceed the economy-wide average of 1.50 percent.

V. Summary and Conclusions

Our first and most important conclusion is that future supply and demand for labor in the U.S. economy will be driven by demography and technology. The supply side of the labor market will be dominated by the slowdown in the growth of the U.S. working-age population, partly offset by continuing increases in the quality of labor input due to rising average levels of educational attainment and experience. From 1960 to 1990 the participation rate of the working-age population increased fairly steadily as more women joined men as participants in the labor market. No such increases in labor force participation are in prospect for our projection period of 2004–2030.

The widely discussed aging of the labor force is reflected in the slowing growth of the working-age population, relative to the total U.S. population. The working-age population will continue to expand more rapidly than the population as a whole, and labor force participation rates will decline very slowly. However, the slowdown in the growth of the time

endowment will reduce the growth rate of the U.S. economy very substantially. This will be reinforced by the decline in investment and growth of capital input that will accompany the slow growth of labor supply. It is important to keep in mind that in the neoclassical theory of economic growth embodied in the IGEM, the growth of capital input is endogenous and is equal to the growth of output in the long run.[10]

Finally, future productivity growth will remain robust, despite waves of technological pessimism that sometimes accompany cyclical downturns. Rapid changes in technology will continue to be concentrated in the industries that produce information technology equipment and software, led by electrical machinery, the industry that includes electronic components like semiconductors. This industry has had very rapid growth of total factor productivity or output per unit of unit, throughout the historical period 1960–2004. We project that this will continue for the next quarter-century, although the specific form of the underlying changes in technology will undergo the same dramatic evolution as in the recent past.

At the level of individual industries, the demand for labor depends not only on the growth of output and the substitution of capital input for labor input, but also on the character of technical change. We have emphasized the wide variations in rates of productivity growth among industries. However, labor demand at the industry level is also strongly affected by biases of technical change. We have focused attention in labor-saving and labor-using biases for each of the 35 industries in the IGEM. We have assessed the importance of these biases during the historical period 1960–2004, and projected these biases for the 2004–2030 projected period. Part of the growth of labor input in industries like instruments, tobacco, coal mining, and communications will be due to ongoing labor-using biases.

In summary, the potential growth of the U.S. economy will be slowing considerably between 2004 and 2030, and monetary policy will have to adapt to the new environment. The changes we have projected embody many features of the future labor market that are well known to economists and monetary policymakers—slowing population growth, particularly for the working-age population, and declining growth in labor quality. We have quantified these factors by relying on official population projections from the Census Bureau and our own estimates of labor quality growth. This data has enabled us to characterize the future growth of labor supply with some precision.

The future growth of the U.S. economy depends on the contribution of labor input, that is, the growth rate of labor input multiplied by the labor share of output. However, future growth also depends on the rate of growth of total factor productivity and the contribution of capital input. In the neoclassical theory of growth embodied in the IGEM, the contribution of capital input, the growth rate of capital input multiplied by the capital share, is endogenous. To a reasonable approximation, growth rates of output and capital input must converge in the long run. The only component of the sources of growth not yet accounted for is productivity growth.

We have projected future productivity growth on the basis of the historical data on productivity growth constructed by Jorgenson, Ho, Samuels, and Stiroh (2007). We have augmented this description of future changes in technology at the level of individual industries by estimating and projecting labor-saving and labor-using biases of technical change. This enables us to conclude that future productivity growth during the next quarter-century will be substantially less than productivity growth during our historical period of 1960–2004. This completes our analysis of labor demand and its distribution by industry.

Economists and policymakers, especially in the Federal Reserve System, have made important contributions to our present understanding of the role of technology in the evolution of labor demand and the growth of the U.S. economy.[11] The remaining challenge will be to build the new understanding of technology and the sources of economic growth into the framework for the conduct of monetary policy. This new policy framework can be erected on the solid foundation provided by projections of future demographic change. The new framework will be an important addition to the Federal Reserve's highly successful policy structure for understanding and mitigating the impact of the business cycle.

Notes

1. Detailed information about earlier versions of the IGEM and a survey of applications are available in Jorgenson (1998).

2. See http://www.bls.gov/cex/home.htm. Detailed documentation for the CEX is available at http://www.bls.gov/cex/home.htm#publications.

3. See http://www.bea.gov/national/index.htm. Detailed documentation for the NIPAs is available at http://www.bea.gov/methodologies/index.htm.

4. See http://www.bea.gov/national/index.htm#fixed. Detailed documentation for the Fixed Assets Accounts is available at http://www.bea.gov/methodologies/index.htm.

5. See Jorgenson and Landefeld (2006).

6. See http://www.euklems.net/. This data set was released on March 15, 2007, and is described in "Use IT or Lose It," *The Economist*, May 19–25, 2007, p. 82.

7. See: http://www.census.gov/popest/estimates.php. Historical data are taken from http://www.census.gov/popest/archives/. These population data are revised to match the latest censuses (e.g., 1981 data is revised to be consistent with the 1990 Census).

8. See http://www.census.gov/cps/.

9. See www.cbo.gov/showdoc.cfm.

10. Jorgenson, Ho, and Stiroh (2008) have pointed out the implications of this fact for growth in an intermediate run of ten years.

11. An excellent summary of this research is provided by Oliner, Sichel, and Stiroh (2007). The implications for monetary policy are discussed by Chairman Ben Bernanke in his August 31, 2006, speech on "Productivity," available at http://www.federalreserve.gov/BOARDDOCS/Speeches/2006/20060831/default.htm.

References

Jin, Hui, and Dale W. Jorgenson. 2007. "Econometric Modeling of Technical Change." Presented at the Center for Business and Economic Research. Copenhagen: Copenhagen Business School. Available online at http://post.economics.harvard.edu/faculty/jorgenson/files/EconometricModelingTechChangehuijin_draft070203.pdf.

Jorgenson, Dale W. 1998. *Energy, the Environment, and Economic Growth*. Cambridge, MA: The MIT Press.

Jorgenson, Dale W., Mun S. Ho, Jon D. Samuels, and Kevin J. Stiroh. 2007. "Industry Origins of the American Productivity Resurgence." *Economic Systems Research* 19(3): 229–252.

Jorgenson, Dale W., Mun S. Ho, and Kevin J. Stiroh. 2005. *Information Technology and the American Growth Resurgence*. Cambridge, MA: The MIT Press.

Jorgenson, Dale W., Mun S. Ho, and Kevin J. Stiroh. 2008. "A Retrospective Look at the U.S. Productivity Growth Resurgence." *Journal of Economic Perspectives* 22(1): 3–24.

Jorgenson, Dale W., and J. Steven Landefeld. 2006. "Blueprint for an Expanded and Integrated U.S. National Accounts: Review, Assessment, and Next Steps," in *A New Architecture for the U.S. National Accounts*, ed. Dale W. Jorgenson, J. Steven Landefeld, and William D. Nordhaus, 13–112. Chicago: University of Chicago Press.

Oliner, Stephen D., Daniel E. Sichel, and Kevin J. Stiroh. 2007. "Explaining a Productivity Decade." *Brookings Economic Papers* (1): 81–137.

Comments on "U.S. Labor Supply and Demand in the Long Run" by Dale W. Jorgenson et al.

Richard Berner

Dale Jorgenson has given us an important paper, describing and using a long-term model of the U.S. economy that can inform our judgment about potential growth and the factors behind these predictions. The news is not good. Over the 2004–2030 period that Jorgenson uses for this projection, potential U.S. economic growth plummets to just over 1.5 percent per year.

There are two factors at work behind this forecast:

1. Jorgenson projects slower growth in the U.S. labor supply, resulting from the now-familiar combination of slowing population growth and the reduced labor force participation that accompanies an aging population.

2. Jorgenson is a self-proclaimed productivity and technology optimist. But the projected pace of productivity growth is slower between 2004 and 2030 than in the 1990s, when information technology posted very impressive gains. In part, this slower predicted growth is because in many key industries, the bias in technical change is labor-using, not labor-saving.

We now know about the consequences of the demographic transition if cohort participation rates stay on current trends. Growth optimists were hoping that high rates of productivity growth would bail us out. Dale's work in this paper argues that this scenario won't take place.

Nonetheless, is there any hope for aging societies and their economic prospects?

Let's first consider productivity. I am concerned that over the next 25 years, productivity growth in the United States may slow, but I'm

not sure that Dale's estimates conclusively prove the case. Much of my skepticism about those estimates revolves around the poor quality of the output data available for certain economic sectors, and for overall and sectoral compensation; my doubt is not a criticism of Dale's model or econometrics. The data indicate that construction productivity has been declining almost monotonically, and especially since the 1990s real estate bust. I find this hard to believe. Likewise, the Bureau of Labor Statistics' programs have simply not kept up with the changing structure of worker compensation. Finally, because the estimates of biases in technical change are based on factor prices, including wages and salaries, I think we should take them with a grain of salt.

Turning next to labor supply, I think the debate about whether labor force participation for older cohorts will increase or decline in the future is still a wide open question. It is clear that only heroic increases in the labor force participation of older American adults will offset population aging. But several factors may influence just such a change, including the fact that the next wave of retirees will come from the baby boom generation, and this group has a habit of upending expectations and rewriting the rules. There are three traditional legs to the retirement saving stool: 1) employer-sponsored pension plans, whether defined benefit or defined contribution plans like 401(k)s; 2) Social Security benefits; and 3) other personal savings. Working longer, in my view, is the fourth critical leg, and future policy changes in Social Security and other retirement saving incentives do influence labor force participation. There is a fifth leg: access to health insurance. Many older adults stay in the workforce to get healthcare coverage, and retire once they turn 65 years old and become eligible for Medicare. Thus any changes to health care financing or Medicare have the potential to trigger significant change in labor force behavior.

In particular, the United States—and advanced countries generally—might also look abroad for help in filling gaps in their labor supply. Such relief could come either from flows of immigration that might alter the nation's demographic profile, or from higher-return investments that will provide more income for retirement. Fallick and Pingle's paper in this volume, and the work by Ralph C. Bryant and John F. Helliwell presented at Jackson Hole in 2004, both suggest that neither increased immigration nor

increased investment returns will be a complete panacea, but clearly each one can help improve the situation. Yet in the post-9/11 world, barriers to immigration are higher and could rise further still. Dale's rudimentary and exogenous treatment of the "rest of the world" is probably worth enhancing to analyze those questions.

More broadly, because Jorgenson's analysis focuses on the longer run, it seems critical to model the most important development in global labor markets of the past decade, namely the emergence of key Asian and other economies that are possessed of a rapidly expanding labor force, strong productivity growth, and high saving rates. As David Autor notes in his conference paper, we really do not know what impact offshoring has on the United States, and other industrial economy labor markets. Much more work is needed in this area if we are to accurately predict future labor flows.

At the present time, however, we can say some things. I think that this positive supply shock has been disinflationary—both through new sources of labor supply and the offshoring of some jobs. Of course, other factors, most importantly monetary policy, not just in the United States but around the world, have also been disinflationary. So we don't know the contributions of each factor. But the existence of such new sources of supply has put some pressure on both labor compensation and employment in both goods-producing and service-producing industries. This is because global connectivity means that workers engaged in a broad array of occupations can work effectively with both customers and colleagues several time zones removed. I know from personal experience that such arrangements work well, and that global companies are reckoning headcount in terms of worldwide numbers. Just as with U.S. labor markets, the question in global labor markets is whether these will have their own demographic transition that could reverse, or at least partially unwind, the favorable supply shock of the last 10–15 years.

I think it is quite likely that the currently favorable supply conditions in global labor markets will eventually experience some type of reversal, but it will take a long time for this shift to play out. To find out how long this might take, it is worth looking at overseas labor markets to compare and contrast our own experience. In these remarks, I'll look briefly at Three Ps: the population overseas, participation rates, and productivity.

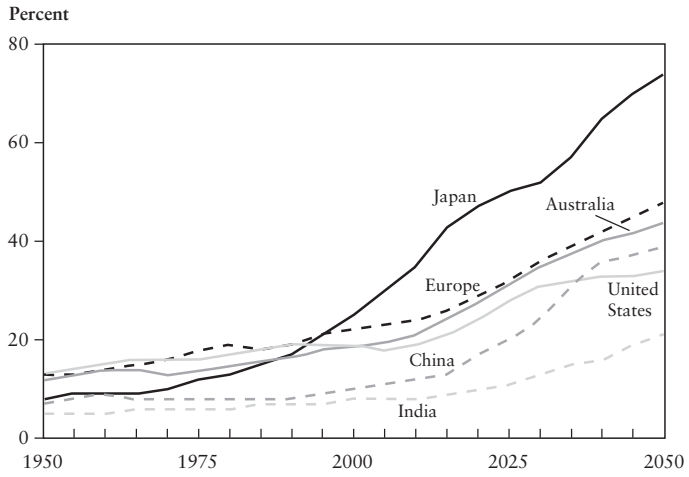

Figure 6.13
Rising Dependency Rates: Ratio of Individuals Over 64 Years of Age to Individuals Aged 15 to 64 Years
Source: Australian Government, Intergenerational Report, 2007; Australian Treasury projections and United Nations 2006 Revision Population Database, medium variant projections; Morgan Stanley Research.

As you see in Figure 6.13, dependency ratios are rising across the board, but what may happen in China and India is especially important for how the global labor supply may play out in the next few decades. China's dependency ratio, partly as a result of its one-child policy, likely will rise and eclipse our own around 2037. Labor force participation there may slow as the population ages, and China is developing a new pension system that may somewhat change labor force participation. In contrast, India has had no one-child policy, and its demographic transition toward an older population is a long way off.

Figure 6.14 compares labor force participation and productivity across countries. The isoquants define countries with the same GDP per capita. Note that in this figure, the hours worked data per capita is for the entire population, not just the workforce. French per capita GDP is below that in the United States, primarily because the French choose to work fewer hours per week. Yet productivity in France is actually very similar to U.S. productivity. Figure 6.15 documents the differences in the French and

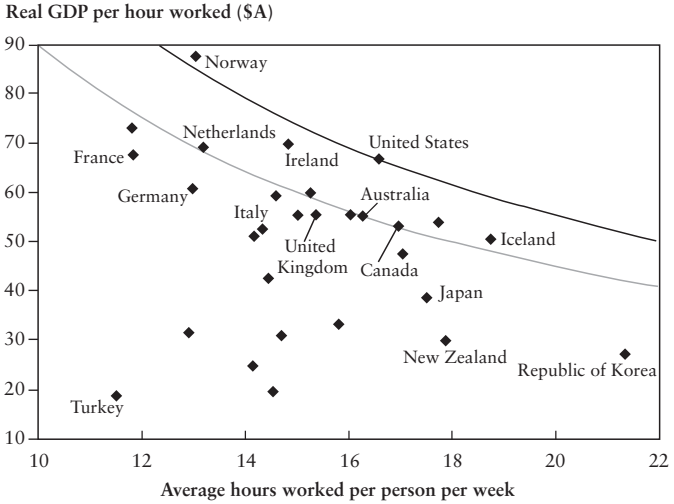

Figure 6.14
Labor Force Participation and Productivity: A Cross-Country Comparison
Source: Morgan Stanley Research; OECD Productivity Database, September 2006.
Note: Average hours worked per person are calculated across the whole population, not just the labor force. Thus, the horizontal axis combines the population and participation components of the 3Ps. Countries on the same contour line have the same GDP per head.

American workweek more clearly. This difference famously reflects cultural differences, but it also reflects labor market regulation. The policy message is that deregulation of labor markets abroad may reverberate in U.S. labor markets. Figure 6.16 shows how strong is the incentive to outsource in Asia and elsewhere based solely on wage differentials, although these data reflect the Bureau of Labor Statistics's definition of compensation on an hourly basis. U.S. compensation includes the fixed cost of employer-provided healthcare, so defining it on an hourly basis may not be entirely accurate, but this won't change the relative position here. The differences are still quite huge. Dale's paper offers a rich menu for future research. Below I identify two key areas for further study that flows from the analysis presented in this paper.

1. The need to assess the impact policy changes may have on labor force participation and aggregate U.S. output. We know from Munnell

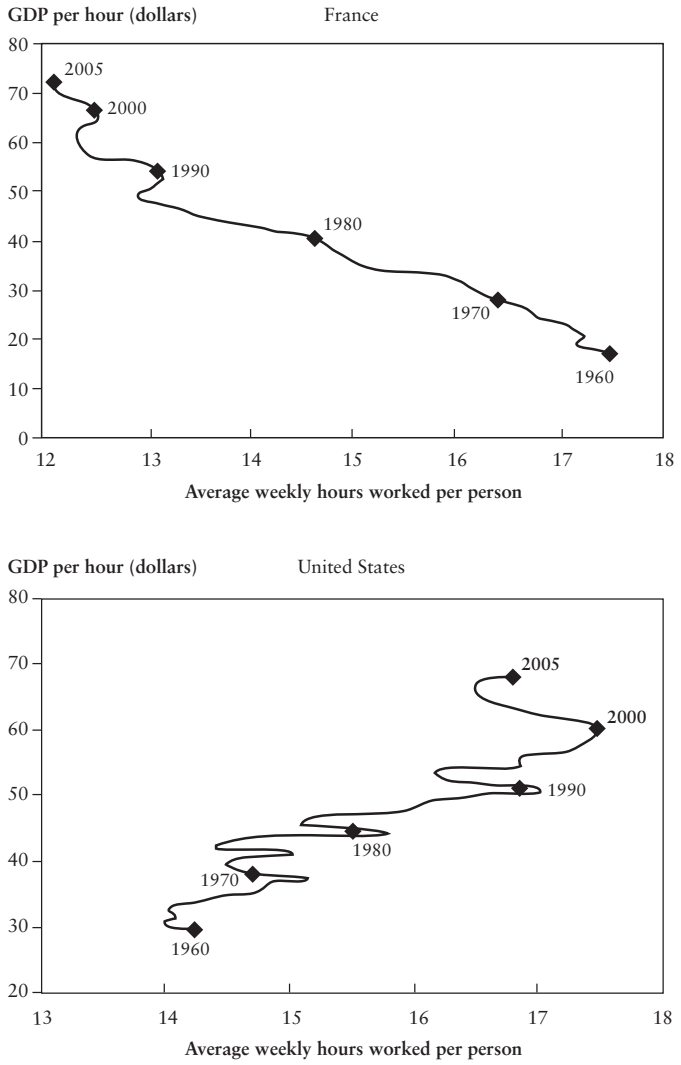

Figure 6.15
Productivity Differences Between the French and American Workweeks, 1960–2005
Source: Australian Federal Budget, 2007–2008, Statement No. 4, http://www.budget.gov.au/2007-08/bp1/download/bp1_bst4.pdf; Groningen Growth and Development Centre, Total Economy Database, January 2007; Morgan Stanley Research.

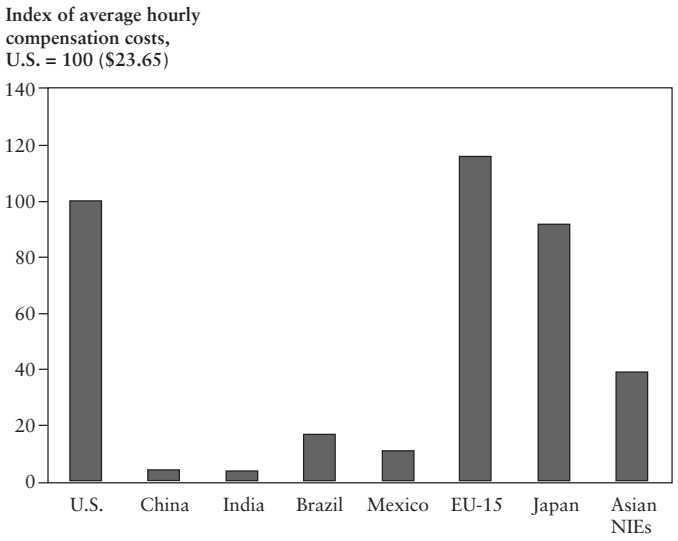

Figure 6.16
Global Wage Differentials for Manufacturing Workers, 2005
Source: Morgan Stanley Research, U.S. Bureau of Labor Statistics.
Note: Asian NIEs (newly industrialized economies) include Hong Kong, Korea, Singapore, and Taiwan. Numbers for China and India are estimates.

and Sass's paper and from Gene Steuerle's work that changes in Social Security retirement ages and benefits, as well as the start of Medicare, induced important shifts in retirement decisions and in labor force participation among older workers. Will changes in the financing of healthcare, such as proposals to tax employer-provided health benefits or proposals among the 2008 presidential candidates to finance U.S. healthcare outside of the workplace, induce significant changes in U.S. labor force participation? Dale tells me that he is doing some work for the Department of Health and Human Services' Centers for Medicare and Medicaid Services along this line, which I applaud. Immigration and tax policies are also worth analyzing.

2. The globalization of product and labor markets has been disinflationary for much of the last 10–15 years. I suspect that this situation will reverse, and changing demographics abroad, as well as the emergence of giants like China and India, may play a role in that reversal. Over time, if these demographic transitions occur abroad, the disinfla-

tionary effect may reverse, and so an analysis of changing demographics abroad may play a role in that analysis.

To conclude, I have a story that knits together thoughts about labor supply, incentives, cultural differences, and the risks of going global. It's especially appropriate that a guy from Wall Street tells this anecdote.

An American investment banker was at the pier of a small coastal Greek village when a small boat with just one fisherman docked. Inside the small boat were several large yellow fin tuna. The American complimented the fisherman on the quality of his fish and asked, "How long does it take to catch them?" The fisherman replied: "Only a little while." The American then asked why didn't he stay out longer and catch more fish? The Greek said he had enough to support his family's immediate needs. The American then asked, "But what do you do with the rest of your time?" The Greek fisherman said, "I sleep late, fish a little, play with my children, take siesta with my wife, Maria, stroll into the village each evening where I sip wine and play cards with my friends, I have a full and busy life." The American scoffed, "I am a Harvard M.B.A. and I could help you. You should spend more time fishing and with the proceeds, buy a bigger boat. With the proceeds from the bigger boat you could buy several boats, and eventually you would have a fleet of fishing boats. Instead of selling your catch to a middleman you would sell directly to the processor, eventually opening your own cannery. You would control the product, processing, and distribution. You would need to leave this small coastal fishing village and move to Athens, then London, and eventually New York where you will run your expanding enterprise." The Greek fisherman asked, "But, how long will this all take?" To which the American replied, "15–25 years." "But what then?" The American laughed and said that's the best part. "When the time is right you would announce an initial public offering, sell your company stock to the public, and become very rich—you would make millions." "Millions ... then what?" The American said, "Then you would retire. Move to a small coastal fishing village where you would sleep late, fish a little, play with your kids, take a siesta with your wife, stroll to the village in the evenings where you could sip wine and play cards with your friends."

Comments on "U.S. Labor Supply and Demand in the Long Run" by Dale W. Jorgenson et al.

Erik Brynjolfsson

It's been said that productivity isn't everything, but in the long run, it's almost everything. And there are few people, if any, who have made a greater contribution to our understanding of productivity growth than Dale Jorgenson, so it's nice to have a few minutes of fame sharing the podium with him.

It's fascinating to discuss U.S. labor supply and demand in the long run. We get to look way out into the future and speculate on what might be, which is always fun. Unlike Bob Hall's business cycle weather predictions, no one can really know for sure whether or not these estimates by Dale and his co-authors are going to come true for quite some time. But with that opening disclaimer, let me try to add some value to the discussion of the issues surrounding future labor supply and demand.

Briefly, here's a summary of some of the key takeaways from Dale and his co-authors' work in this paper, and elsewhere, a lot of which is methodological. The Inter-temporal General Equilibrium Model (IGEM) that he has developed is very impressive, with its inclusion of thirty-five industrial sectors in its number of inputs. As Dale mentioned, in the long run economic growth is really driven by demographics and technological change, his projections of which over the 2004–2030 period are being examined here. Quite a lot of effort went into that model and long-run projection, and he has harnessed an enormous amount of data that isn't really very visible in the particular paper here at hand. You have to dive into a number of his other papers to see the enormous effort and technical detail that went into this model by Dale and a whole host of co-authors and other researchers, including related papers with Hui Jin and Kevin Stiroh. I think that soon there will have to be a separate sector added

to the U.S. economy to account for all the output from Dale's research team.

In this paper, Dale goes ahead and estimates long-run U.S. labor demand and supply, using data he's constructed based on the National Income and Production Accounts (NIPA), and calibrates these relationships between the different sectors and the inputs, particularly on the price side. Dale then makes his projections based on certain exogenous variables, some of which we can pin down pretty well, like what will be happening to population growth and time endowment over time. Those variables are pretty much baked in already. But the projections also need to make some assumptions about what's going to happen to the federal government deficit (I'm glad to see that Dale thinks it is going to decline and go away), the current account deficit, and some of the other most important macroeconomic variables moving forward. Dale's collective research has really made a lot of contributions to predicting what will be happening with education, which really translates into labor quality and capital quality. When Dale gets his Nobel citation, I'm sure it's going to prominently mention the contributions he has made to our understanding that it is not just the quantity of capital and labor that matters, but it is the quality of these inputs that counts, which thanks to Dale are now important factors in everyone's model of the aggregate economy.

Dale then projects the evolution to the economy that these technology changes enable going forward. Again, you need to go to another paper of his to really understand and appreciate some of the novelty of what's gone into this model and its projections. Dale just summarizes this research in a few sentences in this conference paper, but there is a tremendous amount of work that went into this model. His IGEM is a really dramatic extension of Bob Solow's work—as Dale mentioned, 50 years ago Bob developed the basic model of technological change that has really become the workhorse model of growth for economists. But if you want to understand input substitution in the various relationships between labor and capital, you can't just have unbiased technical change in the way Bob Solow's model includes it. You have to look at the bias of technical change for each of the different sectors, and that's what Dale's contribution has done. His IGEM is a very important innovation, and is

really going to be valuable for helping to address these questions about future labor supply and demand in the long run.

The key result, as Dale mentioned, is that we are going to have a significant growth slowdown in a lot of the key input factors. Population growth, labor force participation, and educational attainment are all slowing down. It's hard to argue with any of these projections—you can quibble here and there on the margin, but the broad picture seems largely set in stone. Dale also sees a dramatic decline in investment growth and significant multifactor productivity slowdowns. We can argue about whether that's significant or not, as in an early version of this paper he did not project much of a change. In this more recent version, the numbers have shown more of a decline, so maybe some of the text needs to match up with the numbers a little bit more on that score. But I think that the overall picture and the big takeaway is much slower growth is in the future for the U.S. economy. If you put all three together, slower labor supply growth, dramatically less investment, and significantly lower multifactor productivity growth, you see that overall growth is less than half of what it was in the past, and that's a function of all three of those factors, labor, capital, and productivity growth, interacting in his model. I'll call this "the Great Slowdown," as depicted in Figure 6.8 in Dale's paper.

So that's the summary and background; now for some comments. I'd like to be somewhat of an optimist, like Dale used to be in some of his earlier papers. So in thinking about what's going forward, I'll try to raise some questions that challenge these projections. Of course, it is much easier to be the critic on these sorts of things.

Stability of the Parameters

One question I'll raise is about the stability of the parameters used in the IGEM. Every forward-looking projection has to be estimated from the data we actually have in hand now. Since we do not have data on future years, we need to extrapolate from the historical data to project the future. This a natural thing to sort of do; for instance, to get future consumption patterns, you need to look at what people consumed in the

past and are consuming now in order to predict what people like my sons will be buying in the year 2030, when they will be the prime-age consumers. What my sons are consuming now—iPods, Xboxes, cell phones, and lots and lots of hours of instant text messaging—pretty much accounts for all of their disposable income and all of their disposable time, as far as I can tell. It's an interesting market basket, because none of those items existed 25 years ago. The point is that right now I cannot predict exactly what my sons are going to be consuming in the year 2030—obviously some of it will be housing and food, and maybe some of the discretionary consumer goods will be in keeping with the electronic technologies they are consuming today, but one would suspect that there is going to be a different consumption pattern and set of relationships. Now if you aggregate these predictions enough, maybe you can lump them all into some broad categories, but it is not so obvious that all these fine-grained parameters are going to stay the same, and we need to remember this. While it is important to try and predict the future in order to make better policy choices today, we have to also recognize that to some degree this is an exercise in uncertainty. Who would have predicted the radio in 1900, television in 1925, video cassette recorders in 1965, or iPods in 1985? Later this week, I'm going out to Cisco Systems to see this new product they call TelePresence, a virtual "in-person" communications system that they think is going to revolutionize the world. A TelePresence unit costs $300,000 and enables people to have meetings across great distances, which cuts down on travel and saves time. While it costs $300,000 now, you can be pretty sure that in the next 10 or 20 years it is going to cost a small fraction of that and will have similar or better capabilities. Is a TelePresence unit going to be a big part of the consumption bundle in the future, much like cell phones or personal computers are today? I would bet it may be a standard part of the consumer market basket in 20 years, even though it's not today, and it could substitute for or complement other things today.

Even aside from the uncertaintly about the future substitution relationships among specific components, the unfortunate reality is that the key variables moved around quite a bit in the past. Take productivity growth, which has bounced around quite a bit over the last couple decades. Dale referred to the pessimism people have with the latest downturn, but over

periods of 10 or 20 years we have seen a doubling or a tripling of key productivity numbers. Surprisingly, in some cases the numbers change even for the exact same year. That happens because of revisions to earlier data. In one of his previous papers, Dale pointed out that for the year 1996, the data were revised to show that productivity growth moved from 0.8 to 2.7 percent. So you've got to take these past numbers with some pretty broad confidence bounds, and some of them with a grain of salt. But again, I'm being the critic here, which is a very easy task. Yet it is very difficult to come up with an alternative approach to the one Dale offers here, as you need to work with the data you have in hand. But I think that bearing in mind how these numbers can change just makes you want to be a little cautious about how precise you are about the predictions made going forward.

Price Identification

Another key issue is price identification, which is critical to this type of analysis. Estimating price effects at the intersection of supply and demand is the canonical example of problems introduced from simultaneity. It's a very tough question to sort out, rather like the chicken-and-the-egg issue of which factor really determines the outcome of the other one. Dale does a fantastic job of using instrumental variables to address these issues of simultaneity—he works out these details in some of the companion papers to the IGEM that he pointed me to. While his instruments pass all the statistical tests, I can't help but feel this nagging sense that I'm not fully understanding the model, because some of the correlations just seem a little unexpected. For instance, Dale mentions that increased price tends to be correlated with more technological change, such that this technology input is used more and more. That seems pretty counterintuitive to me—why would higher prices increase demand?—although that's what the data say. It seems plausible that instead there's some causality going the other direction, with higher demand leading to higher costs, as for instance when greater demand drives up wages in a technologically advancing sector. I'm not sure whether the technical change is leading to more demand for the input, which drives the price increases, but maybe there are some explanations that he can help us with. A related issue on

prices is that demand and supply tend to be much more elastic in the long run than in the short run. This raises some questions about which numbers are being used and how these are projected into the future.

Adjustment Costs and Organizational Capital

Another issue is how you deal with adjustment costs and delays and the role of intangible capital. From my understanding of the model, it appears that the instant capital investments are made, these have effects on output and productivity, even before the capital is installed. Now it is difficult to reconcile this assumption with certain categories of intangible capital investment like software and information technology. I've done a lot of work looking at enterprise resource planning (ERP) systems. For instance, Scientific-Atlanta, which is part of Cisco Systems, purchased a large ERP system in late 1994, spent a lot of money and time, and didn't go live with it until 1997; even then, the firm really did not realize its full impact for several years after; see Figure 6.17. This kind of delay partly reflects the complementary of these unmeasured investments in

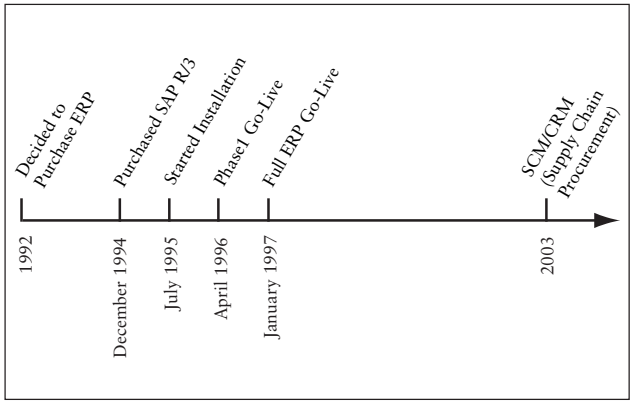

Figure 6.17
Scientific-Atlanta Enterprise Resource Planning (ERP) Timeline
Source: Aral, Sinan, Erik Brynjolfsson, and D. J. Wu. 2006. "Which Came First, IT or Productivity? The Virtuous Cycle of Investment and Use in Enterprise Systems." Working Paper. Cambridge, MA: MIT Center for Digital Business.

intangible capital, which I have become very convinced is a huge factor in the economy even though we don't measure this kind of intangible investment very well.

How do you measure things like organizational capital? I visited a Dell computer factory a few years ago which had doubled its output in response to increased demand. But the way they accomplished this was not by building a new factory next to the first one. Instead, Dell installed some software to redesign their business processes, and by establishing electronic links to their suppliers and their customers, eliminated some of the work-in-process inventory. This set of changes allowed them to produce twice the output with the same bricks and mortar. So had they really built a second factory? Actually, yes. Dell did build a second factory, but it was made out of software and business processes, not bricks and mortar. They installed the organizational capital that was producing real output in the form of physical computers, which resulted in real market value and real revenues. So in that sense, I think installing these new business processes constituted a real capital investment, but unfortunately this investment is not something that shows up in the conventional GDP

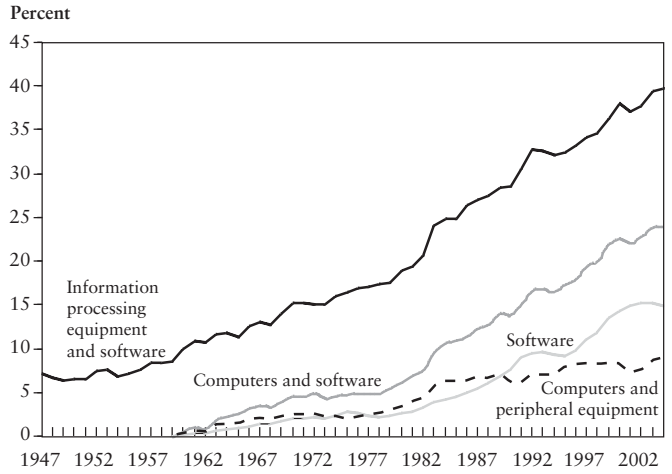

Figure 6.18
Investment Share of Information Technology (Percent of Private Nonresidential Fixed Investments)
Source: U.S. Bureau of Economic Analysis.

accounts as adding to the nation's capital stock. This is not just an isolated story—you can look at many different examples from different firms, and see a similar pattern where hardware is just the visible tip of the iceberg that involves much larger investments like process engineering and human capital investment. These intangible capital investments actually account for the bulk of this capital investment, which I estimate is on the order of $2 trillion of information technology-related and computer- related intangible assets that as a nation we have built up in recent years.

Pessimism on Productivity

In this paper, Dale and his team are very pessimistic on future U.S. productivity gains given some of the other numbers generated in his model's projections. I got the sense that the role of multifactor productivity is grinding to a stop in their model. They don't provide exact numbers for the year 2030, when the projections end, but over the 2004–2030 period productivity is slowing down markedly.

But has productivity growth in the United States has been slowing down historically? If anything, there has been an upsurge in the last decade or so, which in large part reflects gains from information technology. It may be that to calibrate their neoclassical model they need to make some kinds of steady-state assumptions that eliminate the role of exogenous productivity growth. But I'm not sure exactly how this works through their model, and that is something I would be interested in hearing more about in Dale's response. This model's predictions of declining future productivity growth are certainly much more pessimistic than the projections made by a lot of other people who have looked forward, albeit with less elaborate and less sophisticated models. If you look at the bottom line here, for a ten-year period going forward, one of Dale's earlier papers with Min Ho and Kevin Stiroh estimated multifactor productivity growth at about 0.91 percent, which is a bit more than double what it is in the current paper. To square it with the slowdown Dale now anticipates for the next 30 years, maybe all of the projected slowdown in the new paper happens *after* the next 10 years.

Similarly, the paper sees real wages as kind of hitting the ceiling, as you can see in Figure 6.7 of Dale's paper. Real wages rise, but at an ever-

declining rate over the 2004–2030 period. Again, from the paper I get the sense this is because of the assumption that long-run capital growth has to equal output growth over time, so you're going to have this asymptote. But I think this prediction may depend on where that asymptote is located. Are we approaching this point in the twenty-first century or sometime in the thirty-first century? I'm not sure, and in my opinion that obviously is going to have a big effect on when we're going to start seeing this kind of tailing off.

To extrapolate these trends that drive the model, we need to speculate on what's going to happen going forward with information technology and labor demand. A big question is whether or not this trend of productivity gains from advances in information technology is going to continue for the next 25 years. Dale predicts that productivity growth in information technology equipment and software will be slower than it was historically over the 1960–2004 period, especially in the computer sector (which is called "non-electrical machinery," ironically). Understanding the "why" behind the trend, as Lisa Lynch mentioned the other day and like the work that David Autor presented yesterday, is especially important when applied to these potential future outcomes.

Today, we have more and more computer applications that are "intelligent," from computer chess grand masters to software agents recommending books to different kinds of robots. Interestingly, in many cases, these applications were made possible simply because of increases computational power. This is useful to know, because improvements in computer power, not just microprocessors and memory, but also in hard drives and other components, are highly predictable. Moore's Law, the doubling of processor power every 18 months, has held for nearly 40 years. Computer scientists and engineers are confident that it will continue for at least another decade, and most think longer than that.

If information technology does continue to post gains as seen in the past, we will start hitting some key thresholds. For example, a lot of the reason that today's computers can't do tasks like using vision to recognize objects as well as a two-year-old human is because our brains are simply much more powerful at the necessary raw computations than computers. The vision centers of our brains have billions of processors that are actually arranged in a very simple way. When engineers try to simulate

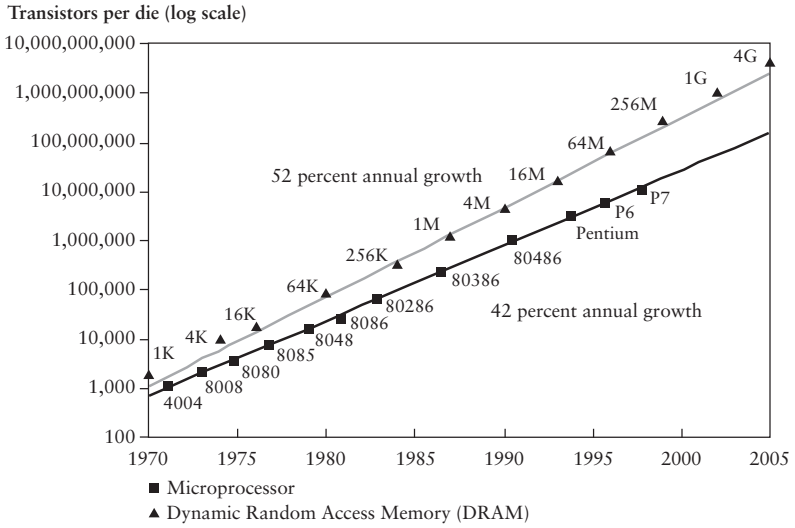

Figure 6.19
Moore's Law and Computer Investments, 1970–2005
Source: Grove, A. S. 1990. "The Future of the Computer Industry." *California Management Review* 33(1): 148–60, and company data. Trend lines are authors' estimates.
Note: P6 and P7 microprocessors and 256M, 1G, and 4G DRAMs are estimated by Intel and the Semiconductor Industry Association.

that ability in computers, it turns out that today's machines do not have the same level of computational power as the human brain. But that's not going to be true if these information technology trends continue for another decade or two. By then, we may actually have a level of computational power that by today's standard falls into the category of things that only humans can do. That will result in a fundamental change in the way the economy works.

Assuming that within the next three decades we will get to this point is a more optimistic scenario than the one Dale presented. In addition, it is quite possible not only that the underlying trends in productivity will continue, but also that the factor share of information technology will continue to grow. Historically, information technology *has* become a bigger and bigger share of the economy, despite the rapidly falling prices of this technology. If this sector continues to take up a bigger share of the economy, then we are going to have not only rapid productivity growth

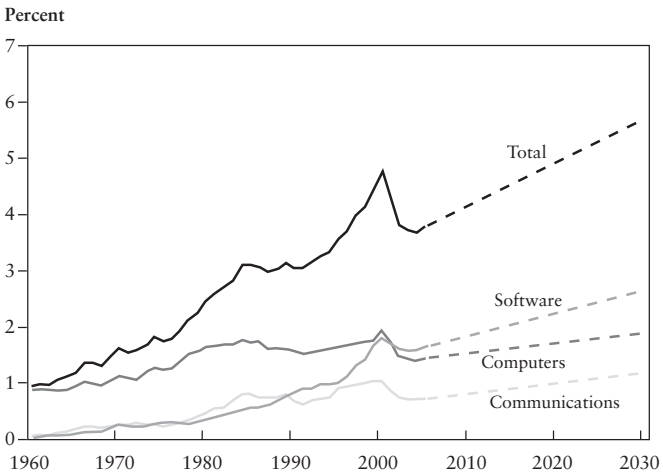

Figure 6.20
Value Added as a Nominal Percentage Share of GDP, 1960–2006, with Projection to 2030
Source: U.S. Bureau of Economic Analysis.

in the related sectors, but this productivity gain is going to have a bigger overall weighted average affect on the economy. This effect will tend to raise the average productivity growth of the economy as a whole, above even its current growth rate.

Let me just close with a comment made by Todd Thompson, the former chief financial officer of Citigroup. When I asked him about whether he was going to be able to cut his information technology budget because of rapidly falling prices, he said, "Oh, I hope not. I'm already spending billions of dollars on technology, and I want to spend more." And I said, why? You don't have to pay as much any more for your computers as you used to, doesn't that mean you can reduce their share of the budget? Thompson replied that Citigroup wants to try and get as much of the rest of the firm's labor force over onto the technology bandwagon, riding down the gains from Moore's law. I think there are a lot of other executives out there who are actively trying to be creative about shifting labor to areas where it can be augmented or replaced with computer power. If they succeed, then instead of winding down, we can expect to maintain or even increase the productivity growth rates we've seen historically.

7

Current and Future Challenges for Policy and Research

Public Policy and the Labor Supply of Older Americans

Stanford G. Ross

Most of the papers prepared for this conference make clear the desirable economic effects if older Americans worked longer and spent fewer years in retirement. Despite a small upturn in labor market participation by older workers in recent years, there is substantial room for significantly greater movement in this direction. The public policy framework is a major determinant of when Americans decide to retire.[1] Both workers and employers take some account of the rules related to retirement that are present in the Social Security laws, tax laws, and regulations governing private pension plans and individual retirement savings. This paper addresses the issue of whether the current set of laws can be changed to provide fewer incentives to retire early and offer more encouragement to work longer.[2]

I. The Current Policy Regime Favors Early Retirement

The current legal and institutional framework in the United States is highly favorable to early retirement. Public policy does far less than it could to encourage older workers to stay in the labor force and firms to employ older workers. Social Security basically allows workers to retire at age 62 (which about 55 percent of Americans currently choose to do) and the vast majority (about 75 percent) do retire under Social Security before the normal retirement age of 65 years. About another 20 percent retire at age 65, and at present very few workers (about 5 percent) work beyond the normal retirement age. Moreover, Medicare benefits are provided at age 65 (and, in the case of disability, at earlier ages). Pensions and other private savings often are available at early ages, which offer

further inducement to retire earlier rather than "later." Indeed, the shift from defined benefit to defined contribution plans makes what assets are accumulated available at even earlier ages.

The Social Security system was not intentionally designed to favor early retirement. The system currently in place has resulted from a historical evolution of almost random political decisions and changing economic and societal circumstances. In 1935, when Social Security was enacted, the retirement age of 65 years was adopted without a great deal of debate or rationale, and only a small percentage of workers at that time survived beyond that age. Since the mid-1930s, life expectancies have grown steadily but the normal retirement age has only been adjusted once. In 1983, largely to help restore Social Security's financial solvency, the normal retirement age was raised to 67 years, albeit with a very long transition over 40 years. The change to age 67, which will affect birth cohorts born in 1960 and later, did not fully reflect the growth in life expectancy that had taken place over recent decades, nor was there any attempt to index the retirement age to take account of future expected gains in longevity.

In 1956, early retirement at age 62 was provided for women (and the benefit formula enhanced), but this was at a time when there were relatively few women in the labor force. Men were allowed to retire at age 62 in 1961 as a response to arguments for equitable treatment and the recession in the late 1950s. These were piecemeal benefit enhancements provided without a full appreciation of the possible long-term consequences.

Significantly, the early retirement age of 62 years was not changed in 1983, when the normal retirement age was raised to age 67. Although the possibility was considered, the increase was not seen as relevant to the system's solvency, the motivating consideration of the legislation, nor worth the political cost. In the case of early retirement at age 62, an actuarially fair reduction of the normal benefit was provided, which reduced somewhat the incentive for early retirement, as did the amendments to the retirement test and the addition of the delayed retirement credit. The labor market implications of leaving the early retirement age unchanged were not seriously addressed.

There have been recent changes in the current policy regime, making it somewhat less favorable to take early retirement. As the normal retirement age rises from 65 to 67, the benefits for early retirees are reduced.

For example, when fully phased in, the benefit at age 62 will be reduced from 80 percent of the normal benefit to 70 percent. The retirement earnings test has been liberalized so that after the normal retirement age individuals can work and retain their earnings without penalty. The law increasingly provides an actuarially fair delayed retirement credit for workers who choose to work after the normal retirement age. When fully phased in, an additional 8 percent will be provided for each additional year of work up to five years. Thus, the normal benefit can be increased 40 percent by working five years longer.

All of these provisions are in the process of being phased in, however, and, to a great extent, constitute "stealth" changes because of the lack of publicity about what is happening over the long transitions. This gradual implementation means that there are no dramatic changes at any given moment to influence behavior in a significant way. For example, if the benefit at age 70–72 years was highlighted as the "full" benefit attainable, it could be pointed out that the early benefit at age 62 is only 50 percent of that amount (rather than 70–80 percent of the normal benefit). Nonetheless, in the final analysis, even after all the presently legislated changes are fully operative, the Social Security regime will continue to be highly favorable to early retirement.

II. Possible Directions for Legislative Changes to Encourage Later Retirement

It would be possible to speed up the transitions to previously enacted later retirement ages. It would also be possible to speed up the delayed retirement credit changes. The benefit formula could also be changed to provide enhancements for continued labor force participation (wages are now indexed only to 60 years of age, and work beyond this age rarely improves a person's ultimate benefit entitlement). But the major change that could make a large difference would be to move the early retirement eligibility age from 62 to 65 years, albeit with an appropriately timed transition phase.

A change in the eligibility age for early retirement would almost necessarily have to be part of a larger package of adjustments, probably including further changes in the normal retirement age, perhaps to 70 years of age (or even beyond) to provide an appropriate structure. Other actions

might include easing the requirements for disability benefits for workers aged 62 to 65 years, and providing a significant minimum benefit for lower-wage workers and others who would be adversely affected by the changes in the early retirement age and the normal retirement age.

It would also be possible to provide tax credits for employers to hire older workers, perhaps through a remission of the employer's share of the Social Security taxes that would otherwise be imposed. Further, a tax benefit could be provided to older workers by remitting their share of Social Security taxes. Income tax credits and allowances could also be provided to employers and workers. In other words, given the political will to change the laws to encourage later retirement, it would be entirely possible to provide greatly enhanced incentives for older workers to remain in the active labor force and for employers to employ these seasoned workers.

III. Why the Prospects are Dim for Major Legislative Changes

What is the likelihood of enacting legislative developments that establish a public policy framework that is more responsive to the realities of the circumstances facing older Americans in the future and confronting the national economy?

To review the history, the Social Security program was established in the 1930s, reconstituted in the 1950s, and expanded in the 1960s and 1970s. Congress adjusted the system in the early 1980s. Yet since the disability reforms enacted in 1980 and major old-age and survivor reforms enacted in 1983, there have not been further changes with the same degree of policy significance. Despite the Social Security financing issues stemming from projected long-term deficits, which first became very apparent in the early 1990s, the political conditions for significant Social Security change have not been present for almost 25 years. The program has largely remained static while the U.S. economy and society have changed dramatically.

Since 1983, the closest we have come to a major Social Security bill was during President Clinton's second term, when he held a series of policy forums across the country and began to assemble popular support for a major reform package. These changes would have involved restoring long-term financial stability to the traditional Social Security program while introducing the concept of an individual account system, possibly as an add-on to the traditional system with some subsidization of contributions for lower-income workers. The federal budget surpluses then projected from the vantage point of the late 1990s could have been used to help finance these changes. However, once Clinton's personal difficulties began to emerge and he became reliant on the more liberal House Democrats during the impeachment proceedings, the congressional support for such a package collapsed. Liberals would not support the addition of an individual account system and conservatives would not support changes needed to ensure the long-term financing of the traditional system. The opportunity to modernize the program and reset it for the twenty-first century was lost.

During his first term President Bush put a great deal of energy into seeking Social Security reform that would have introduced an individual account system. However, many viewed his approach as substantially undermining the traditional Social Security system, and his proposals never received the popular acceptance and broad bipartisan support in Congress that would be required for such a major overhaul. His adamant persistence after the battle was lost produced adamant opposition, with the result that individual account possibilities may have been doomed for the foreseeable future. This highly partisan experience has clearly postponed any significant opportunity for major Social Security reform until the next presidential administration, at the earliest.

It is important to understand that the problems of the current Social Security system go well beyond the issues of early retirement incentives and long-term solvency. The system has not been modernized and adapted to current societal conditions—it is a mid-twentieth century framework that does not mirror the realities of the twenty-first century. For example, family benefits for spouses and dependents need to be reconsidered in the light of the greater participation of women in the workforce, increases in the rate of divorce, and the greater diversity of family patterns. The program's entire structure and its administration need to be thoughtfully reviewed. In short, a comprehensive reform package designed to make the system sound now and in the future is in order—a challenge for some future president and Congress.

The political sensitivity of addressing the early retirement issue is particularly intense, as the reform efforts of the late Senator Daniel Patrick Moynihan (D-NY) reveal. Moynihan was the leading champion of Social Security during a long and distinguished Senate career, which culminated in his chairing the Finance Committee, which has jurisdiction over Social Security and Medicare. He developed a comprehensive reform bill during the late 1990s as he approached retirement and reflected on how to set the system on the right course for the future. The bill retained 62 years as the early retirement age, and eliminated the retirement earnings test at this age, even as it substantially raised the normal retirement age and provided for increases in Social Security taxes to restore the system's long-term solvency. Moynihan's rationale for retaining the early retirement age at 62 years was to maximize choice for individual workers, although it was apparent that based on a later normal retirement age, with the larger actuarial reduction of the early benefit, many workers taking early retirement would receive lower benefits, and that if some lived long enough, that shortfall could possibly lead to providing Supplemental Security Income (SSI) (means-tested benefits) in greater amounts.

Moynihan's bill, however, would have introduced an individual account system within the Social Security system that allowed workers to achieve more adequate benefits on an individual self-help basis. Moynihan broke with the program's traditional rationale, which had emphasized from the mid–1930s the priority of providing benefit "adequacy," in the traditional system. This was a goal that some felt had only been achieved in 1972, some 37 years after the program began, when benefits were raised substantially and indexed for inflation.

A fallacy of the Moynihan position, to my mind, is that the very workers who will opt for the early benefit are likely to be ones who cannot afford to build up supplementary individual account accumulations or will not want to do so. Furthermore, making larger numbers of workers dependent on SSI would likely undermine the program in the long run (since it could make means-testing other aspects of the program a shorter step).

Looking to an individual account system within Social Security to compensate for an inadequate benefit structure also neglects the fact that there are numerous opportunities for individual savings outside of the Social Security program that lower income workers do not adequately utilize. While there might be some greater attraction for saving in an individual account system that is part of the Social Security system, it is far from certain how effective this solution might be, particularly for low-income workers, in a period in which the Social Security Administration manifests serious shortcomings in service delivery and government programs in general are on the defensive.

In summary, the point here is not to disparage Senator Moynihan's efforts, but to underscore the political difficulty of raising the early retirement age that is currently in place. Since this benchmark interacts with other important components of the U.S. retirement system, a comprehensive package of balanced reforms designed to better adapt the Social Security program to current socioeconomic conditions should deal adequately with all these components in order to make a later retirement age more attractive and feasible for a significant number of older workers.

The 1983 Social Security reforms and the 1986 tax reforms under President Reagan remain models for principled changes to these large systems. These reforms involved bipartisan, balanced packages, based on a great deal of preparation, and deft management of the political process over an extended period of time. While neither reform package fully realized the policy goals of their major sponsors, and the achievements of both degenerated subsequently, these efforts accomplished a great deal at the time and could have been platforms for further reform if there had been continuity in the political will for more reform. Assuming that this type of broad support and bipartisan cooperation are the fundamental political conditions necessary for any substantial future reforms of the Social Security and tax systems, the question is whether it is realistic to expect such conditions to emerge anytime soon. Currently, the partisanship in Congress and the polarized nature of political discourse makes it seem unlikely that such a consensus would emerge. Thus, it seems unrealistic to expect that in the near term a major Social Security bill might be enacted that would substantially change the incentives for older workers and employers.[3]

In truth, over the last few decades government institutions may be seen as having weakened in terms of the ability to respond constructively to societal needs. Professional expert leadership is often absent in govern-

ment and the policymaking process has given way frequently to ideological formulations and crass political calculations. Major progressive reforms are difficult to achieve and are less frequently undertaken, especially when the gains are very long term, but the political election cycle runs on a short-term schedule.

IV. Why the Prospects for Marginal Legislative Changes are More Likely

At this point in time, healthcare issues may well be the key to where the reconsideration of the public policy framework on entitlement programs develops. The costs of healthcare are rising and its accessibility is diminishing for all Americans. The costs of Medicare are rising and the Social Security cash benefit will inevitably be diminished over time as the rapidly growing Medicare Part B and Part D premiums are deducted from monthly Social Security payments. Private employers are curtailing healthcare benefits, particularly for retired workers, but employed workers are increasingly affected too. Healthcare costs are increasingly a problem in employing older workers. Pension provision has largely shifted from defined benefit plans to providing defined contribution plans, and there are increasingly generous provisions for allowing employees to take their accumulations long before retirement, often when they change jobs at relatively early ages.

All the presidential candidates presently are endorsing major healthcare reform and it seems likely that as we draw closer to the 2008 election the political imperatives for change will intensify. The Medicare Prescription Drug Improvement and Modernization Act enacted in November 2003 under George W. Bush presents many problems as it becomes more fully implemented, and undoubtedly there will be an effort to reconsider this law as part of healthcare reform. The Medicare Part D drug benefit will almost certainly need to be changed, as will many other aspects of the Medicare program.

Any legislation designed to reform how healthcare is provided in the United States will probably not thoroughly consider how such reforms might lead to workers being encouraged to stay in the labor force longer, and how this might affect firms employing older workers. Nonetheless, enacting healthcare reform clearly could provide the opportunity to change the incentives that would encourage longer working lives for many Americans. For example, to relieve the burden on employers of hiring older workers, Medicare could be made the primary (rather than secondary) payer of benefits for workers above age 65.

It is also possible, however, that additional incremental changes will tilt the U.S. retirement regime even further toward taking early retirement. Changes in the Medicare system could allow a buy-in before 65 years of age to make health insurance available to those who are not yet eligible for Medicare. It was a popular reform to eliminate the Social Security earnings test, and the earnings test could also be repealed for retirees under 65 years of age. Such a change might actually increase the number of early retirees who continue to work.

In general, it is easier to extend benefits than to curtail or eliminate these entitlements, which means that politicians seeking to do something in the field of Social Security and Medicare could well make the policy regime even more favorable to early retirement. Elderly Americans vote in larger numbers than do other age groups and are very important politically. As a result, incremental changes that favor older workers could be attractive, particularly if these reforms seem to involve manageable budgetary costs.

V. Real Events are Likely to be More Influential Than Changes in Laws, and Could Lead to Constructive Legislative Change

It seems unlikely that Social Security reform will emerge as a major legislative issue until the financing problems become even more acute than at present. Thus, the current public policy regime is unlikely to be greatly changed in the near term. The "social engineering" approaches in the 1983 Social Security Act and the 1986 Tax Reform Act are unlikely to be achieved in the near future because of a highly partisan political environment. This means that real events in the economy are more likely than any changes in the public policy framework to influence employees to retire, or encourage employers to employ older workers. Adverse changes

in prevailing economic conditions could stimulate changes in private sector behavior. If the economy falters and wealth prospects are diminished, so that individuals feel more insecure, there could be an impetus for working longer. On the other hand, if the economy continues to prosper, the stock market continues to do well, and the housing market remains relatively strong, workers may well accumulate wealth that leads to early retirement.

On the employer side, if major labor shortages emerge, perhaps because immigration is curtailed and outsourcing is restricted, firms may need to adjust to a diminished labor supply by taking steps to employ older workers that previously they might not have seriously considered hiring. Older workers are likely to want more flexibility and even part-time jobs. Here, a precedent that may be instructive is the way employers in many areas have adapted their practices to encourage women to work. In many cases, women have required more flexible work schedules and conditions of employment, including part-time opportunities at certain times in their work lives. Employment conditions can be changed if there is a desire to do so, and older workers would likely respond to encouragement from employers who provide greater incentives adapted to their particular needs.

The changing structure of the U.S. economy, which now is more oriented around industries providing services and knowledge, often involves less physically demanding labor, and this should enable older workers to stay in the workforce longer; in some instances this already seems to be occurring. It is also the case, however, that the newer 24/7 service economy often produces more stresses and strains and there can be greater worker "burn-out." In addition, some employers such as large law and accounting firms increasingly enable, and often force, their workers to retire at early ages in order to allow younger lawyers and accountants to rise within the firms. These are highly competitive fields where the firms do better by keeping their workforce young.

Another factor favoring younger workers is generational. The computer and digital revolution requires the constant learning of new skills. Older workers often have difficulty with this retooling process. It is possible that younger generations that are more computer literate, technically educated and well-trained, will find it easier to adapt and keep up with the rapidly changing technology that is at the heart of much of the contemporary economy.

All of these variables make it difficult to predict whether it will be easier for older workers to stay in the labor force longer, and for employers to better use the skills and experience that older workers can provide.

In the final analysis, changes in real economic and social conditions may lead to changes in the public policy framework as the private sector seeks legislative changes that can help it to adapt to changing circumstances. Over time the laws could be changed to enable workers to work longer and employers to employ older workers, and the economy would benefit from this increased labor supply. In other words, this could be a subject for which change is stimulated from the grassroots level, rather than by the policy establishment.

But the political will for enacting these changes does not seem present now because neither U.S. employers nor aging workers and the organizations that represent them seem particularly motivated to seek such changes. Employers and the institutions that represent their business interests seem to give priority to easing immigration requirements for high-tech workers and other prime-age workers from abroad. In fact, aging workers and the organizations that represent them often see early retirement opportunities as important to preserve and enhance. It will take a massive educational effort to ready the aging U.S. population for the economic realities they will face going forward.

Notes

1. The public policy framework is only one aspect of a complex situation, and what workers and employers actually do depends on a variety of economic and social factors. The public policy framework can be important beyond immediate pecuniary aspects in setting expectations for workers and employers. Non-pecuniary considerations appear to matter a great deal in influencing retirement behavior. There is considerable evidence that workers often make retirement decisions that are financially disadvantageous to their self-interests. Another factor is the information used in making retirement decisions is often inaccurate or incomplete. The government's role, particularly the Social Security Administration and Medicare, in providing useful information, and at times analysis and advice, is a subject that should be carefully explored.

2. The larger question of whether people should work longer at the expense of greater leisure is not addressed. For any individual, the choice between work and leisure is a personal decision. For the society as a whole, it is a value judgment. This comment simply assumes that greater productivity is a public good from an economic standpoint.

3. In a related issue, the disability system is inconsistent in its goals, unfair in its results, and uneven in its administration. Considering its $110 billion annual program cost, and $8 billion administrative cost, it is in imperative need of substantive and procedural reform. But without the fundamental conditions for major reform being present, little can be expected to happen legislatively.

The Seven Deadly Sins in Aging Policy and Research: A Cautionary List for Policymakers and Prognosticators

C. Eugene Steuerle

Pride, envy, gluttony, lust, anger, greed, and sloth—theologians tell us that we become better people by examining these sources of failure.

But my concern here is not with the classic seven deadly sins, but what I feel are the contemporary seven deadly sins being committed in current policy and research on aging. Reflecting on them likewise provides some warning signs for us acting as policymakers, researchers, or prognosticators.

I am not, of course, going to accuse any particular person of committing the sins I am about to discuss, since I am well aware of the Biblical injunction that only one who is without sin in these matters is allowed to cast stones. More to the point, these shortcomings, some of omission and some of commission, are social sins: these overlay the macroeconomic debate on aging even when some of us researchers and policymakers claim personally at a micro-level to have avoided them. Finally, I am sure that some of you have different religious training, and will decide that some of what I label "sins" are actually virtues.

Deadly Sin # 1: Giving Too Little Attention to the Labor Side of the Aging Debate

The first deadly sin is paying too little attention to the labor (as opposed to capital) side of the aging debate. By listing this first, I am obviously preaching to the choir assembled here. I congratulate Cathy Minehan, Bob Triest, and their Boston Fed colleagues for their leadership in taking on this most important, yet usually neglected, issue in the aging policy debate.

Look closely at projections such as those performed by the Social Security actuaries. Ignore any of the black box aspects of what these forecasters do, and it is actually quite easy to approximate their long-term projections by looking at nothing more than the projected change in the ratio of workers to beneficiaries. In a pay-as-you go system, a decrease in that ratio from 3-to-1 to 2-to-1 (more precisely, from the ratio of 3.3 workers per beneficiary in 2008 to 2.1 by 2040) means approximately that per worker revenues need to grow by one-half, or per-beneficiary spending fall by one-third. Those last fractions are fairly close to the projections by the actuaries of what would be needed in revenue increases or spending cuts to restore pay-as-you-go balance by about 2040.

But that labor force issue is not how we explain—I should say, obfuscate—the issue when talking to the American public. Instead, we discuss Social Security's trust fund balances. We talk about spending down the trust funds, even though those funds never contained more than one-tenth of Social Security's long-term liabilities. We talk about introducing individual or personal accounts—like 401(k) plans paid for directly with our Social Security taxes—and pretend that these can magically grow into large future retirement benefits with no sacrifice now. We start examining ways to reshuffle funds so that we can borrow at 0 to 3 percent real interest rates today, invest in a stock market paying 7 percent real annual returns on average, and then reap magic money through arbitrage. All of these ideas are discussed as if there are no risks involved, and no one is on the other side of each financial transaction.

In effect, through these various diversionary discussions, we convert what is primarily a labor market problem into a financial market problem, with financial solutions obtained by wielding our actuarial and economic weapons, making our present value and trust fund calculations, and pretending large gains will be painlessly wrung from better manipulation of these financial accounts.

I've been an observer of Washington policymaking for more than three decades now, and there has not been a single year in which Congress did not adopt some new incentive for saving or investment, often through policies affecting private pensions or retirement accounts or different types of saving plans. Yet national saving rates have declined, personal savings rates have even fallen below zero in recent years, and most people still go into retirement with only modest private assets. I am not going to go through the reasons why these programs have largely failed, but the history of these failures should warn us about how far we can go in thinking that we can solve this looming labor market problem merely through attempts to ramp up saving, however worthy a goal in itself.

Why have we collectively, as researchers, policymakers, and a nation, failed to pay much attention to the labor market piece of the aging issue? This failure centers in part around the way we define the problem and how we form expectations based on historic precedent. One aspect of the definitional failure, which I will return to shortly, is the simplistic and misleading definition of an "aged" person or society. Right now I want to focus on another simplistic (and related) assumption: that in a growing economy we will always want to retire for longer and longer periods and, therefore, retirement policy is largely a matter of adjusting to that inexorable force.

Perhaps you think this is not a simplistic assumption, but rather that it has well-established theoretical and historical underpinnings. In economic theory, leisure is used to provide closure to some mathematical models that economists use, but, let's be honest: in point of fact, "leisure" is almost a meaningless concept. People, at least those who are not economists, don't just do "nothing." Granted, as societies get richer, people do demand a lot more of many "good" things, and often in increasing proportions. Among those intangible items we seek are freedom from financial pressures and the dictates of superiors in our jobs; we don't like outside forces to command our use of time to do things we really don't like doing. Yet in today's service-oriented and Internet-driven economy, pleasurable activities—including intellectual challenges and enjoyable social interactions—increasingly can be sought on the job, not just through dropping out of the labor force. Fewer of us are stuck at the lathe, unable to talk to friends during the day, or unchallenged mentally on the job. The "new" economy means that we need to revise many of our ideas about what constitutes "work" and what constitutes "leisure."

Now you may assert that regardless of abstract economic theory, empirical and historical work supports the notion that leisure, so to speak, is and always has been a superior "good." I don't think so. Until roughly

a century ago, almost no one retired at all from the labor force. Veblen's 1899 book, *The Theory of the Leisure Class*, offers one anecdotal piece of evidence for this assertion, when he contrasts the small leisure class with the much broader working class. Back then, most people were farmers and did some share of the family's chores almost until the day they died. I guess starting from zero, time off from hard work is a superior good for a while, but that does not offer much evidence for its retaining its same value over time, especially when the structure of work is changing rapidly.

Put another way, we don't always seek more leisure. We seek more freedom and more enjoyable and stimulating activities. At various margins those good things might just as well be found in the workplace as off the job. Even if we did want more time off the job, it's not clear why we would want exactly 100 percent time off from work and only for the last third or so of our adult lives.

The industrial revolution, of course, changed work and family structures largely by separating paid employment from family, social, and volunteer life. But the traditional industrial economy has been on the wane for quite some time. Today, in the information age, growth in output more and more centers on services that can be provided in the profit-making sector, the non-profit sector, or the household sector alike. A similar story can be told of growth rates within occupations. Think of occupations in healthcare and the knowledge economy. This change in the economy implies that many aspects of life are becoming less compartmentalized over time. Devices like a period of retirement—a consecutive rather than simultaneous approach to providing the benefits of greater freedom—are less necessary to enable people, while participating in the workforce, to fulfill many different sides of their personal development.

Now let's look beyond theory to actual labor data. Over the last half of the twentieth century, we find that the percentage of adults who were employed in the United States increased rather than declined in just about every non-recession year. Hours of work also increased over this same period. Does this data indicate that we always demand more and more leisure?

Well, some might retort that sociological, not economic, factors were at play during this period. American women entered the labor force

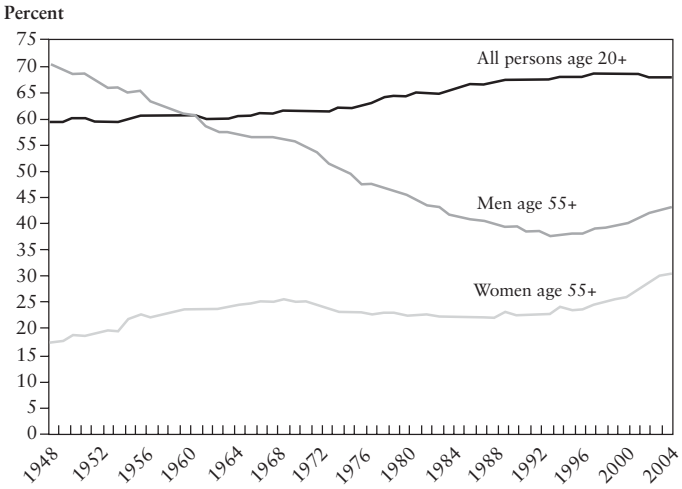

Figure 7.1
U.S Labor Force Participation Rates for Men and Women Aged 55 Years and Older Compared with the Entire Adult Population, 1948–2004
Source: U.S. Bureau of Labor Statistics, 2005.

faster than American men dropped out of it. The Pill meant having fewer kids, while new types of durable goods, like dishwashers and microwave ovens, and institutional changes, like equal rights and reduced discrimination, were all factors that enabled women, often much better educated than their predecessors, to enter the workforce in increasingly larger numbers.

In point of fact, these sociological factors were important, but mostly by empowering women to provide the labor supply that we as a society demanded. That is, there still was a labor demand curve reflecting our nation's demand for goods and services relative to what could be produced.

It is common these days to use time series data of age/sex specific actual labor supply to project future labor supply. A few assumptions are added—for instance, as they age, younger women will have a larger labor force participation rate when they are older than do older women today—but these additions don't have huge effects on the ultimate estimates of aggregate labor supply. Basically, new cohorts of the same age

and sex are largely assumed to behave the same way as do current cohorts (with some modest upticks). For instance, if 70 percent of men aged 62 currently don't work, then it is assumed that 50 years from now roughly the same percentage of men that age will also not work. Since there will be a much larger percentage of people in retirement age, as traditionally defined, under these kinds of projections the percent of adults working will decline quite significantly (for instance, from three workers per beneficiary to two per beneficiary).

Yet, in addition to misrepresenting the demand for freedom as a demand for leisure, as discussed above, the complication with this approach is that it essentially fails to assume much, if anything, about even the existence of a labor demand function. For instance, the data over the last half of the twentieth century is roughly consistent with a relatively inelastic labor demand curve boosted by an increase in labor supplied in the formal versus informal labor market, as women responded to the sociological factors already mentioned. When labor demand is introduced, all of a sudden different age/sex groups can be seen to be responding to that demand function according to the sets of relative opportunities facing them, thus helping explain the relative shifts in male and female participation rates. Labor demand also provides a partial explanation for the conversion from defined benefit to defined contribution plans, as employees tend to work longer when defined contribution plans are in place. Traditional defined benefit plans often pushed them out of the workforce at ages as early as 50, 55, and almost always by 60 or 62, as these private plans usually provide negative economic accruals of pension benefits for working longer.

I am not asserting that labor demand is inelastic, only that our projections do not do a good job of dealing with the influence of labor demand on labor supply.

As a prognosticator, I will stick my neck out here. I predict that in the first half of the twenty-first century, older Americans aged 55 years or more are going to occupy the position that women occupied in the last half of the twentieth century—the largest pool of underutilized talent and human capital in the economy. And that in response to a potential shortage of workers, their labor force participation will increase much more than implied by almost all projections made today.

Deadly Sin # 2: Making Policy Without Setting Real Targets

The policy debate over many aging issues usually proceeds without any real targets based upon established principles, such as progressivity, efficiency, or the equal treatment of individuals. For instance, take many of the recent debates that have centered either on preserving the current Social Security system or taking a portion of existing Social Security taxes and directing these toward individual personal accounts that would operate somewhat like individual retirement accounts (IRAs). What do those objectives even mean? "The current system" is hardly a principled target, and neither are "personal accounts." These concepts may be the means to some ends, but the ends need to be specified.

For instance, consider progressivity or "vertical equity," as opposed to equal justice for those in equal circumstances, or "horizontal equity." Presumably one goal of Social Security is to protect people from poverty or to redistribute income in some manner. The establishment of a progressive benefit schedule clearly indicates a desire to instill a fair degree of progressivity, or income redistribution to those less well-off, into the program. Fascinatingly, however, until the past few years almost no one even tried to measure what progressivity the Social Security system actually achieved. Or how this goal has changed over time. Like so many policy debates, the focus seemed to be more on symbol than on substance.

When the 2001 Social Security commission established by President Bush was in the midst of its debate, groups on both left and right tried to lay claim to concerns over progressivity by playing the race card. One side claimed that blacks and other shorter-lived groups lost out in Social Security because of forced annuitization. Opponents of this view asserted that increasing the retirement age would disadvantage these same groups. Yet neither side asked for or used any empirical work to back up their a priori assertions. You'd think those policymakers and economists making such proposals would want to know those facts so they could modify their proposals to achieve some realistic target for progressivity. Ideology trumped empirics.

A similar complaint can be lodged against the debate over whether to carve some personal accounts, almost like mandated 401(k)s, out of existing Social Security contributions. The goal of these accounts is

presumably to try to increase the supply of saving. But if that is the goal, then one has to be careful about the deficits the government would likely run to finance these accounts. In addition, even if the government did not finance its saving subsidies out of it own deficits or dissaving, one has to worry about the extent to which individual saving in such mandated Social Security accounts would be offset by the ability of people to put less money in other accounts or to borrow from other retirement accounts like 401(k)s, hence reducing any net saving. These behavioral issues, again, were largely ignored in most recent public debates on Social Security, including those surrounding the 2001 Commission set up by President Bush, as well as his efforts after the 2004 election to generate interest in Social Security reform. To a large extent, these reforms have been debated without agreeing on precise targets as to what the system should do, other than eventually be in some sort of financial balance.

Deadly Sin #3: Limiting the Debate So as to Be Politically Correct

The third deadly sin committed by so many in our economics profession is to limit the debate about aging to issues that are politically correct, at least within each political party or ideology. Topics such as progressivity and increasing retirement savings at least have been on the table, even if these issues were often engaged in a symbolic rather than in a real way. But many other important issues and principles are not even on the table: these considerations have largely been ignored because these are not the politically correct fights of the day.

For instance, whatever happened to equal justice? The current Social Security system sanctions a broad and blanket discrimination against abandoned parents and other single heads of household, one that would be clearly illegal for private pensions and private property. Largely women, these individuals often work more hours, pay more taxes, raise more children, and get fewer benefits than other individuals who just happened to marry somebody with higher earnings. All other sorts of related glitches occur in the current system because of its outdated structure of family benefits designed around an Ozzie-and-Harriet-type household with only one working adult, who is male. Furthermore, consider the inequities surrounding the ability of a worker to generate benefits for multiple spouses without paying a dime for any of them, the extra benefits offered largely to men who have children late in life, the discrimination in favor of someone working 30 years and earning $40,000 a year over someone working 40 years and earning $30,000 a year, and the denial of any spousal rights to a women who is married to the father of her children for less than 10 years.

I have approached some well-known expert authors of Social Security books and proposals over the years, and all admit that the current system of family benefits, as well as some other aspects of Social Security, are both unfair and inefficient. These shortcomings are slam dunk issues for any reform based on principles of equal justice or horizontal equity, regardless of how one comes out on some tougher issues, such as how large the system should be and the extent to which greater progressivity reduces economic efficiency. But these authors invariably respond that these issues are difficult politically because any remedies would create some losers along with winners, and so they decide to dodge considering them. I find this avoidance, particularly by experts, disturbing. It would be one thing if these problems were confined to a limited period of time. However, in theory we build these systems to be permanent—so the discrimination will last eternally, or at least decades until the next major reform, and will grow in absolute terms along the way. I believe we have obligations as analysts to present a comprehensive and objective picture of the current failings in the Social Security system, and save the political compromises for the politicians to work out.

Deadly Sin #4: Misusing "Aging" as a Term

The fourth deadly sin relates to how we misuse the word "aging," and some potential analytical errors that follow from this misnomer. To begin with, a crucial distinction must be made between the concept of people simply living longer, and the aging of an entire population due to lower birth rates. In the first instance, we are largely measuring improvements in health status at given ages. For instance, a 65-year-old person alive today has more well-being than a 65-year-old living in 1940, at least as measured by having more years of life expectancy at that age and, as it turns out, better reported health status. The phenomenon of more people

living longer does not by itself "age" a population; that is, there is no necessary increase in the proportion of the population closer to death or in worse health because of an increase in longevity. Put another way, if "old" means being in the last ten (or fifteen) years of life, or bearing the disabilities that often accompany "old" age, then the fact that we are living longer does not mean that we are spending more years as "aged" or that the population is "aging."

Here's an exaggerated example to make this point. If the population were suddenly to live on average to 100 years of age, rather than to 70 years, would one assert that age 62 is "old" in both scenarios?

On the other hand, the decline in the birth rate clearly does decrease the percent of the population in their early years of life, and increase the percent who are in their last years. For example, the *proportion* of the population within five years of death will eventually rise. In this scenario, social needs might indeed be relatively greater for items like assisted living vis-à-vis transportation or housing.

How "aging" is defined has strong implications for our research and the conclusions drawn from these studies. Many current analyses assume that age from birth represents the same variable in measuring a person's status in 1940 as it does in 2007. To the contrary, I suggest that it is more appropriate in later ages to measure people of equivalent age over time by their remaining life expectancy, or at least their relative age (such as being in the last 10 percent of their lives, as measured by life expectancy).

In one simple analysis, I compared the labor force participation rate of men aged 65 years with the participation rate of men having close to 17 years of remaining life expectancy for the 1940–2001 period. While the former showed a more steady drop in participation over most years of that period, the latter showed almost no drop in participation for about 25 years until the early 1960s, when two major events occurred—in 1962 men were granted an early retirement age of 62, and Medicare was made available in 1965. These different curves imply very different weights for the influence of policy relative to an ever-increasing demand for leisure over time.

Our use of words associated with age and aging also have powerful signaling effects. If 62-year-old individuals are told they are "old," they may act accordingly. I once made an appeal in a column to reform Social

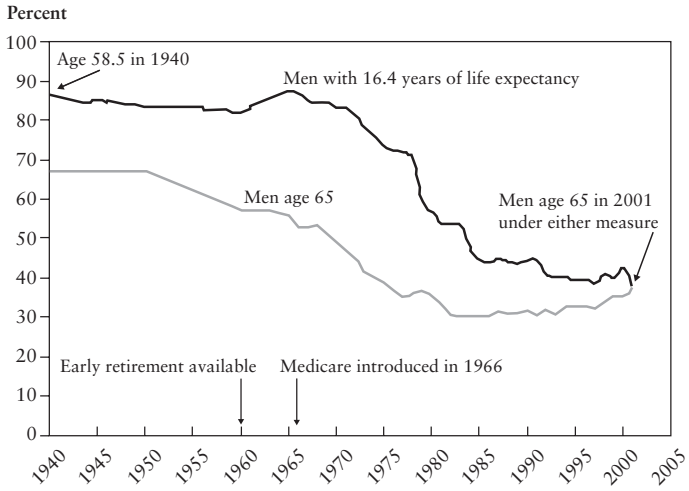

Figure 7.2
U.S. Labor Force Participation Rates for Men, 1940–2001
Source: C. Eugene Steuerle, Adam Carasso, and Meghan Bishop. 2002. The Urban Institute. Based on data from the U.S. Bureau of Labor Statistics, the U.S. Census of Population, and the U.S. Social Security Administration.

Security's Old Age and Survivors Insurance by re-labeling it "The Middle-Age, Old Age, and Survivors Insurance System" when it provided benefits beyond some expected number of years, such as 15. I was not being facetious. The right label would probably go a long way toward achieving reform, since "old" is still correlated in many people's minds with disability and incapacity. Providing early retirement benefits for the middle-aged is a very different animal than providing needed retirement help for the truly elderly portion of the population.

Deadly Sin #5: Ignoring the Balance Sheet

Committing deadly sin #5, ignoring the budget balance sheet, is as easy a trap for policymakers and researchers as the Sirens' call is for sailors. Sometimes, when the amounts are small and the returns to marginal efforts are large, the budgetary issue is less important. However, when programs like Social Security and Medicare become huge relative to the

rest of the federal budget, their movement has dramatic effects on other programs, just as the movement of an elephant in a room would affect everyone else in it.

Let me give some examples of how incomplete balance sheet accounting distorts policy predictions.

Consider the issue of taking earlier or later retirement. This decision affects not just benefits but the revenues that fund entitlement programs and the rest of government. Some calculations mainly attend to the effect on benefits, or benefits relative to past cash wages, as in calculations of replacement rates. The Social Security actuaries do better when they look to the trust funds, since those calculations at least incorporate Social Security revenues. But then they stop there.

In point of fact, when a person retires and drops out of the workforce, there is a decline in output and income roughly equal to that person's marginal output or income from labor. As an example, take a worker making $50,000 a year who retires and collects benefits. If there is no new worker to replace this person, national income falls by $50,000, but almost all of the loss must be borne by others. That is, not only must someone come up with the extra $24,000 to pay Social Security and Medicare benefits to this retiree, but existing programs (including Social Security and Medicare) must get by with $16,000 less in income tax, Social Security tax, and other revenues that were used in part to fund other individuals' benefits from government programs. The Social Security Administration does not estimate most of these revenue effects when analyzing reform proposals, nor does it report on the worker's loss of other income from the wages that are foregone.

Another example derives from looking at the federal budget as a whole. Consider first that growth in benefits is so rapid that an average-income 50-year-old couple today is scheduled to get about one million dollars' worth of Social Security and Medicare benefits across the 26 or so years that one or the other is expected to be drawing benefits. That is, this couple would need an interest-earning bank account of $1 million (in today's dollars) by age 65 to cover the cost of its future Social Security and Medicare benefits. Younger couples get even more because benefits are scheduled to increase continually over time.

But we don't need to wait to see what this growth toward million-dollar-plus packages of benefits, plus the decline of the percent of adults working, does to the government balance sheet right now. Between 2006 and 2010, for instance, the cost of paying Social Security, Medicare, and Medicaid benefits is scheduled to grow by $326 billion. The growth in revenues implied by current law is only expected to be $494 billion over the same period. In other words, these three programs will absorb two-thirds of all revenue growth. (Social Security, Medicare, and Medicaid already constitute about 40 percent of total spending.) If defense expenditures fall moderately relative to GDP, and we add in interest on the federal debt, my calculations show that by sometime before 2020 there are no revenues left for any other function of government other than providing defense, Medicaid, Medicare, and Social Security. Basically, middle-aged people today are scheduled, by the time they retire, to get almost everything government provides—for themselves. Their children and grandchildren are not scheduled to get much of anything, until they retire. Clearly, this situation is not morally, economically, or politically sustainable, but currently we are doing nothing to address the balance sheet implications of the policies now in place.

Deadly Sin # 6: Assuming Away Arbitrary Aspects of the Status Quo

Many economists love to argue that the status quo represents some logical equilibrium of market forces striving toward balance. The complication was that, at least until recently, they would discount the influence of arbitrary accidents, herd instincts, and other commonly recognized irrational and unpredictable human behaviors when making policy recommendations.

In the case of aging research, this status quo approach often assumes away the importance of downright bad policy design. Former Reagan press secretary Jim Brady had a term for how these types of situations come about, the BOGSAT method of decisionmaking: Bunch of Guys Sitting Around a Table.

It wouldn't be so bad if we just acknowledged that the BOGSAT method is how much existing Social Security and health entitlement pol-

icy was made, and then tried to adjust it pragmatically from then on. But once policy gets enacted, all sort of interest groups will assert that the process for arriving at that result was arrived at in some totally rational way. They will usually succeed in getting some economists to fall in line as defenders of the faith, to bolster their argument that some crazy line of arbitrary policymaking represents some ideal economic equilibrium.

Let me give you one concrete example—defined benefit pension plans. If one thinks about the history of defined benefit (DB) plan design, the main rationale seems to be that early on, when pension schemes were young and seldom fully funded, some employers, or employers in bargaining with unions, decided to reward their long-term, retiring workers. They latched onto this DB design as a way to try to replace in retirement a significant portion of those workers' last years of salary. Gradually the design was extended to the next generations, and government started insisting that such plans be funded.

This design took off and for a time dominated private pension policy. But certain aspects of it were and always have been crazy. Calculate, if you will, the value of the DB benefit in defined contribution (DC) terms—that is, what a person gets if he or she stays on the job for one more year. I initially did this type of calculation when I was first at the Treasury Department in the 1970s, a period when rising inflation kept lowering the economic value of the DB plan for almost all younger employees—but especially those who left a firm before an age that benefits are paid out. (The measure of the last or highest years of salary is not adjusted for inflation, so someone who leaves a firm when 50 years old will suffer a benefit erosion due to 15 years of inflation if the benefit, based on final salary, is provided at age 65.)

This DB plan design creates other problems as well, encouraging some employers to get rid of older workers and encouraging some older workers to quit. Because staying on the job for awhile compounds benefits in an exponential fashion (for reasons I will not explain here), employers offering DB plans have an extraordinary incentive to get rid of senior employees in the fast-accruing years immediately prior to retirement. After all, the $30,000-a-year 50-year-old worker might be accruing $9,000 in pension benefits, whereas the 25-year-old worker paid the same salary might be accruing next to nothing. It might make sense, for

instance, to move a plant from Michigan to Tennessee when this noncash pension compensation gets very high for an aging workforce that has been in place for years.

After their plan reaches peak economic value, on the other hand, employees often move into a situation where they accrue negative economic benefits in their pension plan, so they retire early. Think of school teachers working for state governments, who get a maximum pension benefit of 50 percent of pay after a certain number of years on the job. In this case, if a fully vested teacher paid $50,000 in her final years of employment retires, she gets $25,000 of pension benefit annually. If she continues to teach another year, she gets $50,000 of pay (assuming this salary is capped) but totally loses the pension benefit for that year.

So what is the sin here? Some economists like to argue that DB plans were appropriate for an older economic era where people did not work beyond the age of 55 or 60, stayed with the same employer through most of their lives, and did not have such long life expectancies. Perhaps DB plans were more appropriate in the past than today, but were many of these features of DB plans, such as lack of inflation adjustment, ever really appropriate? The historical data simply do not support the view that in some halcyon earlier period, most American workers were employed by the same firm for most of their lives, and by the time they were 60 years old retired after 35-40 years with a gold watch and a golden pension. Workers who fit this description were the exception, not the rule. Many of the other workers suffered discrimination, albeit legal discrimination, in their compensation packages when they achieved much lower rates of accrual in their pension plans for doing the same job. We need to admit it when crazy sets of incentives like the ones that have evolved from DB plans have real world implications, and accordingly adjust our plan designs, policies, and laws defining what is truly discriminatory.

Deadly Sin #7: Hubris about Knowing the Future

Now let me turn to the most deadly sin of all. Dante believed that Pride was the root of all sins and in *Purgatorio* he suggested that all souls must be first "purged" of that sin. The root of our sins in aging policy and research—the sin from which most of the others flow—is our hubris

about being able to predict the future accurately. We then overconfidently enact long-term entitlement policies that hamstring our children from following their own vision for how their society should evolve.

Let me explain this hubris in the context of today's federal budget. We have had budget problems before, but these were always easily contained in the following years simply by showing a little less exuberance in enacting new legislation. Today, however, so many entitlement benefits have been promised for so far into the future that government has lost almost complete control over the future direction of policy. Decisions made in the past continually deter today's policymakers from better allocating resources to meet the current needs of society and the new desires of voters in a democratic society. Those policy decisions that attempt to preordain the future are mainly in the areas of retirement and healthcare, and in both cases the spending growth largely applies to the older adult population (children's healthcare, for instance, is relatively inexpensive).

Fundamentally, the federal government is oriented ever more each year to serving us when we are older. Middle-aged people, for instance, can expect to get larger and larger shares of government revenues for themselves, but smaller and smaller shares for their children or grandchildren. Among other considerations, this approach treats future generations of adults like permanent minors, whose future options must be controlled in advance. If they ever want to do something new with the taxes they will be paying, we budgetarily box them into first having to retract on promises already made. We or our predecessors have essentially put into the law requirements for how all future revenues of the government should be spent for all eternity.

Pride and Prejudice

In sum, aging policy and research needs to pay attention to its seven deadly sins if it is to advance policy in the interest of all U.S. citizens, young and old, living and still not born.

Atoning for these sins of omission and commission will not be easy, but with the aging of the baby boom generation—the first cohorts, born in 1946, start collecting Social Security in 2008 and Medicare in 2011—we will be increasingly hard-pressed every year we delay fixing those programs. Reform is going to mean swallowing our pride, and giving up our natural prejudices toward old ways of doing things. These ingrained habits are considerable, and include: 1) accounting for a labor market problem mainly in financial terms; 2) fighting over retirement policy without setting targets based on firmly articulated principles; 3) excluding principled reforms because these are not part of today's politically correct debate; 4) poorly defining what it means to be aged; 5) ignoring both sides of the budgetary balance sheet; 6) defending parts of the status quo as natural when these policies derive from little more than past arbitrary decision-making, and 7) maintaining the hubris that we should restrict our children's ability to determine how government should best meet the needs of their future society.

Contributors

Katharine G. Abraham is a professor in the Joint Program in Survey Methodology and an adjunct professor in the economics department at the University of Maryland. Abraham's current research agenda is concentrating on the development of the recently funded American Time Use Survey Data Access System, and continuing her ongoing work on time use, the quality of new jobs, and the transition to retirement. She chairs the Committee on Economic Statistics for the American Economic Association and is a research associate at the National Bureau of Economic Research. From 1993 to 2001 she was Commissioner for the Bureau of Labor Statistics. Abraham was previously a visiting professor and research scholar at Princeton University, a research associate at the Brookings Institution, and an associate professor at the Massachusetts Institute of Technology's Sloan School of Management. She has a B.S. in economics from Iowa State University and a Ph.D. in economics from Harvard University.

David H. Autor is a professor of economics at the Massachusetts Institute of Technology. He specializes in labor economics, and is a research associate at the National Bureau of Economic Research. Autor is currently working on two research programs, one on the growth of labor market intermediation, and the other on technological change, job skill demands, and earnings inequality. He received a National Science Foundation CAREER award for his research on labor market intermediation. Autor has a B.A. from Tufts University and a Ph.D. in public policy from the Kennedy School of Government at Harvard University. Before going

to graduate school, Autor spent three years directing efforts in San Francisco and South Africa to teach computer skills to economically disadvantaged children and adults.

Susanto Basu joined the economics faculty at Boston College in 2005 as a full professor; previously he had been a professor in the economics department at the University of Michigan since 1992. He is also a visiting scholar at the Federal Reserve Bank of Boston, a research associate at the National Bureau of Economic Research, and a member of the executive committee of the Conference on Research in Income and Wealth. His scholarship centers on macroeconomics, with special emphasis on technological change and productivity growth. Along with Miles S. Kimball and John G. Fernald, Basu recently published "Are Technology Improvements Contractionary?" in the *American Economic Review* (December 2006). Basu holds an A.B. and a Ph.D. from Harvard University.

Richard Berner is a managing director and chief U.S. economist at Morgan Stanley, where he is responsible for directing the analysis and forecasting of the U.S. economy and financial markets. He is a member of the Economic Advisory Panel of the Federal Reserve Bank of New York, a member of the Panel of Economic Advisors of the U.S. Congressional Budget Office, and a member of the Board of Directors of the National Bureau of Economic Research. Before joining Morgan Stanley in 1999, Berner was executive vice president and chief economist at Mellon Bank Corporation, where he also served on the senior management committee. He has held previous positions at Salomon Brothers, Morgan Guaranty Trust Company, and Wharton Econometrics. Berner received a bachelor's degree from Harvard College and a Ph.D. from the University of Pennsylvania.

Jared Bernstein is senior fellow and director of the Living Standards Program at the Economic Policy Institute in Washington, DC. His areas of research include income inequality and mobility, trends in employment and earnings, low-wage labor markets and poverty, international comparisons, and the analysis of federal and state economic policies. Between 1995 and 1996, Bernstein was the deputy chief economist at the U.S. Department of Labor. He wrote the book *All Together Now: Common Sense for a Fair Economy* (2006) and has co-authored eight editions of *The State of Working America*. Bernstein has published extensively in popular venues, including the *New York Times*, the *Washington Post*, and the *American Prospect*. He has a Ph.D. in social welfare from Columbia University.

Katharine Bradbury is a senior economist and policy advisor at the Federal Reserve Bank of Boston. Prior to joining the Bank in 1981, Bradbury worked as a research associate at the Brookings Institution in Washington, DC and at the Institute for Research on Poverty at the University of Wisconsin. Earlier, she served as a Peace Corps volunteer in the Mariana Islands. Bradbury's research focuses on income inequality and income mobility, labor force participation and other labor economics topics, state aid to local governments and other issues of state and local public finance, and the New England regional economy. In addition to her duties at the Boston Fed, Bradbury is an editorial board member for *Massachusetts Benchmarks* and serves on the Local Aid Subcommittee of the Massachusetts Municipal Finance Task Force. Bradbury earned her B.A. from Carleton College and her Ph.D. from the Massachusetts Institute of Technology.

Erik Brynjolfsson teaches at the Massachusetts Institute of Technology (MIT), where he is the Director of the MIT Center for Digital Business and the George and Sandra Schussel Professor at the MIT Sloan School of Management. Brynjolfsson was among the first researchers to measure the productivity contributions of information technologies, and his research has been recognized with six best paper awards and five patents. He is the author or co-editor of several books, including *Understanding the Digital Economy*, *Intangible Assets*, and *Strategies for eBusiness Success*. He is a research associate at the National Bureau of Economic Research and serves on *Time* magazine's Board of Economists and on the Academic Advisory Council of the Federal Reserve Bank of Boston. Brynjolfsson has an A.B. and an S.M. from Harvard University in applied mathematics and decision sciences, and a Ph.D. from MIT in managerial economics.

Gary Burtless holds the John C. and Nancy D. Whitehead Chair in Economic Studies at the Brookings Institution in Washington, DC. His recent research has focused on the sources of growing wage and income inequality in the United States, the influence of international trade on income inequality, the job market prospects of public aid recipients, social insurance reforms in developing countries and in formerly socialist economies, and the implications of privatizing the U.S. social security system. Burtless co-authored *Globaphobia: Confronting Fears about Open Trade* (1998), has edited and contributed to other books, and has written numerous scholarly and popular articles on the economic effects of social security, public welfare, unemployment insurance, and taxes. Before joining Brookings in 1981, he served as an economist at the U.S. Department of Labor (1979–1981) and at the U.S. Department of Health, Education, and Welfare (1977–1979). His undergraduate degree is from Yale College and he earned a Ph.D. in economics from the Massachusetts Institute of Technology.

Bruce Fallick is a senior economist in the macroeconomic analysis section, Division of Research and Statistics, at the Board of Governors of the Federal Reserve System. His research interests concentrate on labor market economics and macroeconomics. Prior to joining the Board in 1993, he was an assistant professor at the University of California, Los Angeles from 1988 to 1994. Fallick has a B.A. in economics from the University of Pittsburgh, and a Ph.D. in economics from the University of Pennsylvania.

Christopher L. Foote is a senior economist and policy advisor in the Research Department at the Federal Reserve Bank of Boston, where he currently serves as Acting Director of the Center for Behavioral Economics and Decisionmaking. He graduated from the College of William and Mary in 1987, and then worked for two years as a newspaper reporter in Harrisonburg, Virginia. He received a Ph.D. in economics from the University of Michigan in 1996. From 1996 to 2002, Foote taught at Harvard University's department of economics, where he also served as director of undergraduate studies. In July 2002, he accepted a position as senior staff economist with the Council of Economic Advisers, becoming chief economist at the CEA in February 2003. From May 2003 to September 2003, he served as an economic adviser to the Coalition Provisional Authority in Baghdad, Iraq, returning briefly to Iraq in January and February of 2004. He joined the Boston Fed in October 2003.

Richard J. Goettle is the president, co-founder, and principal of Cambridge Planning and Analytics, Inc., developer and marketer of DATADISK Information Services. Goettle also serves as a senior economist with Dale W. Jorgenson Associates and previously was with the National Center for Analysis of Energy Systems at Brookhaven National Laboratory. He has written extensively on the general equilibrium consequences of U.S. energy, environmental, and tax policies. Goettle holds a B.A. in mathematics and computer science from Miami University, a M.B.A. from Northwestern University, and a Ph.D. in economics from the University of Cincinnati.

Robert E. Hall is the Robert and Carole McNeil Hoover Senior Fellow and a professor of economics at Stanford University. He is an applied economist with interests in technology, competition, employment issues, and economic policy in the aggregate economy and in particular markets. His current research focuses on activity, employment levels, and stock-market valuations in market economies, and on the economics of high technology. Hall serves as the director of the National Bureau of Economic Research's program on economic fluctuations and growth in the U.S. economy. He is a member of the National Academy of Sciences, a fellow of the American Academy of Arts and Sciences, and a fellow of the Econometric Society. Before joining Stanford's department of economics, he taught at the Massachusetts Institute of Technology and at the University of California, Berkeley. He received a B.A. in economics from the University of California, Berkeley, and a Ph.D. in economics from the Massachusetts Institute of Technology.

Mun S. Ho is a senior economist at Dale W. Jorgenson Associates, where he has worked since 1993. His research interests focus on environmental economics, public economics, technological change, and economic growth. Since 2006 Ho has been a visiting scholar at the Institute of

Quantitative Social Science at Harvard University, where he is working on productivity measurement and environmental policy analysis. Since 1999 he also has been a visiting scholar at Resources for the Future in Washington, DC. From 1992 to 2005 he was a fellow in the program on Technology and Economic Policy at Harvard University's Kennedy School of Government. He was an assistant professor of economics at SUNY-Buffalo from 1989 to 1992. Ho earned his A.B. in mathematics at Northwestern University and his Ph.D. in economics at Harvard University.

Robert Hutchens is a professor in the Department of Labor Economics at Cornell University's School of Industrial and Labor Relations. His teaching focuses on labor economics and on the economics of government tax and transfer programs. His current research focuses on long-term implicit contracts and on employer policy toward older workers. As part of this research focus, he is involved in a survey of firms that examines employer policies toward phased retirement. Hutchens has served as a policy fellow at the Brookings Institution, as an associate editor at the *Industrial and Labor Relations Review*, as chairman of the Department of Labor Economics, and as a research fellow at the Institute for the Study of Labor (IZA). He has held visiting positions at the University of British Columbia, at the U.S. Department of Health and Human Services, and at the International Labor Organization. He received his Ph.D. in economics from the University of Wisconsin, Madison.

Dale W. Jorgenson is the Samuel W. Morris University Professor at Harvard University. After teaching at the University of California, Berkeley, he joined the Harvard faculty in 1969 and was appointed the Frederic Eaton Abbe Professor of Economics in 1980. He served as chairman of the department of economics from 1994 to 1997. Jorgenson has conducted groundbreaking research on information technology and economic growth, energy and the environment, tax policy and investment behavior, and applied econometrics. Jorgenson received the prestigious John Bates Clark Medal of the American Economic Association in 1971. He was elected to Fellowship in the American Association for the Advancement of Science (1982), the American Statistical Association (1965), and the Econometric Society (1964). He received a B.A. in economics from Reed College and a Ph.D. in economics from Harvard.

Chinhui Juhn has been teaching at the University of Houston since 1991. She was named the Henry Graham Professor of Economics in 2005, and serves as co-director of the graduate program. Juhn is a research fellow at the National Bureau of Economic Research and a research fellow at the Institute for the Study of Labor (IZA) in Bonn, Germany. A specialist in labor economics, she has published papers on a wide range of topics, including wage inequality and how gender and race drive employment trends. Her current research focuses on how working spouses impact cyclical fluctuations in the labor force and on the impact of economic reforms and trade liberalization in Mexico. Juhn was a visiting scholar at ERMES and the University of Paris II in fall 2001, and a visiting scholar at the Australian National University in summer 2000. She was an economist at the Federal Reserve Bank of New York in 1998–1999. Juhn has B.A. in economics from Yale University and a Ph.D. in economics from the University of Chicago.

Lisa M. Lynch is the William L. Clayton Professor of International Economic Affairs and former Academic Dean at the Fletcher School at Tufts University. She is a research associate at the National Bureau of Economic Research, the Economic Policy Institute, and the Institute for the Study of Labor (IZA) in Bonn, Germany. Lynch is currently chair of the Board of Directors of the Federal Reserve Bank of Boston, chair of the American Economic Association Committee on the Status of Women in the Economics Profession, and a member of the executive board of the Labor and Employment Relations Association. From 1995 to 1997 she was the chief economist at the U.S. Department of Labor, and she has been a faculty member at the Massachusetts Institute of Technology, The Ohio State University, and the University of Bristol. She is an associate editor of the *Journal of Population Economics* and was a co-editor of the *Journal of Labor Economics* and an associate editor of *Labour Economics*. Lynch received her undergraduate degree from Wellesley College and her M.Sc. and Ph.D. in economics from the London School of Economics.

Joyce Manchester is the Unit Chief for the Long-Term Modeling Group in the Health and Human Services Division of the Congressional Budget Office. Previously, she was the Director of the Division of Economic Research in the Office of Policy at the Social Security Administration (SSA). Before joining the SSA, she was the staff economist for the Social Security Advisory Board, where she was the principal staff person for the 1999 Technical Panel on Assumptions and Methods. Prior to her work at the Advisory Board, Manchester was an analyst in the Macroeconomic Analysis Division of the Congressional Budget Office, working on social security, consumption taxes, and baby boomers in retirement. She also spent six months as a visiting scholar at the World Bank, looking at tax evasion within social security programs around the world. After receiving a B.A. from Wesleyan University and a Ph.D. in economics from Harvard University, she was an assistant professor at Dartmouth College.

Cathy E. Minehan is President and Chief Executive Officer of the Federal Reserve Bank of Boston. An expert in payment systems, she is a member of the Payment System Policy Advisory Committee, which considers issues related to systemic risk in national and international payment systems. Currently, she serves on the Steering Committee of the Conference of Presidents of the Federal Reserve Banks and as a voting member of the Federal Open Market Committee. Minehan is vice chair of the National Council on Economic Education, which promotes economic literacy with students and teachers. Within New England, she serves on the boards of many civic, professional, and educational organizations, including the Boston Private Industry Council, Massachusetts General Hospital, the United Way of Massachusetts Bay, Jobs for Massachusetts, the University of Rochester, and the Carroll School of Management at Boston College. She has received numerous leadership awards and honorary degrees, including the 2006 Pinnacle Award for Lifetime Achievement presented by the Greater Boston Chamber of Commerce. Minehan began her career with the Federal Reserve Bank of New York, following her graduation from the University of Rochester, and holds an M.B.A. from New York University.

Alicia H. Munnell is the Peter F. Drucker Professor of Management Sciences at Boston College's Carroll School of Management. She also serves

as the Director of the Center for Retirement Research at Boston College. Before joining Boston College in 1997, Munnell was a member of the President's Council of Economic Advisers (1995–1997) and assistant secretary of the Treasury for economic policy (1993–1995). Previously, she spent 20 years at the Federal Reserve Bank of Boston (1973–1993), where she became senior vice president and director of research in 1984. She has published many articles, authored numerous books, and edited several volumes on tax policy, Social Security, public and private pensions, and productivity. Munnell was co-founder and first president of the National Academy of Social Insurance and is currently a member of the American Academy of Arts and Sciences, the Institute of Medicine, and the Pension Research Council at Wharton. She is a member of the board of The Century Foundation, the National Bureau of Economic Research, and the Pension Rights Center. Munnell earned a B.A. from Wellesley College, an M.A. from Boston University, and a Ph.D. in economics from Harvard University.

Jonathan Pingle is an economist at Brevan Howard Asset Management in Washington, DC. From 2003 to 2007 he was an economist at the Board of Governors of the Federal Reserve System. His research interests center on labor economics, with a particular emphasis on how demographic change affects the supply of labor. He is a research affiliate with the Institute for the Study of Labor (IZA) in Bonn, Germany. Pingle has a B.A. from the University of Connecticut, an M.P.P. from the University of Chicago, and a Ph.D. in economics from the University of North Carolina at Chapel Hill.

Stanford G. Ross is a lawyer/consultant based in Washington, DC. He has served as chairman of the Social Security Advisory Board, Commissioner of Social Security, and Public Trustee of the Social Security and Medicare Trust Funds, and has served in the U.S. Treasury Department, on the White House domestic policy staff, and as general counsel of the Department of Transportation. The author of many papers on federal taxation and income security subjects, Ross has taught at the law schools of Georgetown University, Harvard University, New York University, and the University of Virginia, and has been a visiting fellow at the Hoover

Institution, Stanford University. He has provided technical assistance on social security and tax issues to many foreign countries under the auspices of the IMF, World Bank, and U.S. Treasury. Ross has served as chairman of American Bar Association Tax Section committees and on the Advisory Committee for the federal income tax laws of the American Law Institute. He received his B.A. from Washington University (St. Louis) and his J.D. from Harvard Law School.

Steven A. Sass joined the Center for Retirement Research at Boston College in January 2002. Previously, he was an economist at the Federal Reserve Bank of Boston. Sass is the author of *The Promise of Private Pensions: The First Hundred Years* (Harvard University Press, 1997), *The Pragmatic Imagination: A History of the Wharton School 1881–1981* (University of Pennsylvania Press, 1982), and with Alicia H. Munnell, *Social Security and the Stock Market* (Upjohn, 2006). Sass was an assistant research professor at the Wharton School, a research associate of the Pension Research Council, and has taught at Rutgers University and Brandeis University. He earned his B.A. from the University of Delaware and his Ph.D. from The Johns Hopkins University.

Daniel T. Slesnick is a professor of economics and an associate dean of research in the College of Liberal Arts at the University of Texas at Austin. Slesnick's recent research has examined issues related to the consumption behavior of the elderly, price measurement in healthcare markets, and the role of age and cohort effects on observed patterns of inequality. His book titled *Consumption and Social Welfare* (Cambridge University Press, 2001) reports on a research program that uses consumption data to measure inequality, poverty, and the standard of living in the postwar United States. He received a B.S. in mathematics from the University of Washington, and a Ph.D. in economics from Harvard University.

C. Eugene Steuerle is a senior fellow at The Urban Institute in Washington, DC, and co-director of the Urban-Brookings Tax Policy Center. He serves on the National Committee on Vital and Health Statistics and on advisory panels or boards for the U.S. Congressional Budget Office, the

General Accounting Office, the Joint Committee on Taxation, the Actuarial Foundation, and the Independent Sector. Steuerle's previous positions include president of the National Tax Association (2001–2002), chair of the 1999 Technical Panel advising Social Security on its methods and assumptions, president of the National Economists Club Educational Foundation, deputy assistant secretary of the Treasury for tax analysis (1987–1989), and resident fellow at the American Enterprise Institute. Between 1984 and 1986, he served as economic coordinator and original organizer of the Treasury's tax reform effort. Steuerle has a Ph.D. in economics from the University of Wisconsin, Madison.

Robert K. Triest is a senior economist and policy advisor in the Research Department at the Federal Reserve Bank of Boston. He is also a visiting scholar at the Center for Retirement Research at Boston College. Before joining the Bank, he was an assistant professor at the Johns Hopkins University and an associate professor at the University of California at Davis. Triest's research on topics in public finance and labor economics has been published in various professional journals, including the *American Economic Review,* the *Journal of Economic Perspectives*, and the *Review of Economics and Statistics*, as well as in Boston Fed publications. He is currently working on projects related to gross flows in the labor market, and to the interaction of housing, savings, and portfolio choices. Triest earned his bachelor's degree from Vassar College and his Ph.D. in economics from the University of Wisconsin.

Peter J. Wilcoxen is an associate professor of economics and public administration at the Maxwell School of Citizenship and Public Affairs at Syracuse University, which he joined in 2003. He is also the director of the Maxwell School's Center for Environmental Policy and Administration, and is a nonresident senior fellow at the Brookings Institution in Washington, DC. Wilcoxen's principal area of research is how environmental and energy policies affect economic growth, international trade, and the performance of individual industries. His work often involves the design, construction and use of large-scale intertemporal general equilibrium models. He collaborated with Dale E. Jorgenson on the Jorgenson-

Wilcoxen model, a 35-sector econometric general equilibrium model of the U.S. economy that has been used to study a wide range of environmental, energy and tax policies. He taught at the University of Texas at Austin from 1991 until 2003. Wilcoxen received a B.A. in physics from the University of Colorado in 1982, and a Ph.D. in economics from Harvard University.

Author Index

Subject Index