Volume Title: Health at Older Ages: The Causes and Consequences of Declining Disability among the Elderly

Volume Author/Editor: David M. Cutler and David A. Wise, editors

Volume Publisher: University of Chicago Press

Volume ISBN: 0-226-13231-5

Volume URL: http://www.nber.org/books/cutl08-1

Conference Date: October 8-11, 2004

Publication Date: January 2009

Chapter Title: Disability Risk and the Value of Disability Insurance

Chapter Authors: Amitabh Chandra, Andrew A. Samwick

Chapter URL: http://www.nber.org/chapters/c11118

Chapter pages in book: (295 - 336)

10

Disability Risk and the Value of Disability Insurance

Amitabh Chandra and Andrew A. Samwick

10.1 Introduction

As successive generations of Americans have access to healthier lifestyles and more advanced medical technologies, we can expect the prevalence of work-limiting disabilities to recede. A decline in disability will have important consequences for the nature of employment at older ages and the optimal design of social insurance programs. In this chapter, we take initial steps toward understanding these consequences by measuring the disability decline in the working age population over the past two decades and assessing its implications for welfare and saving. We focus on consumers' valuation of disability insurance—either as income or as an assistive technology—to protect against the risk of permanent disablement. Because the probability of disablement is small but the loss conditional on the event is large, consumers will find it difficult to self-insure substantially against the risk of disablement through precautionary saving.

To understand changes in the probability of a work-limiting disability, we

Amitabh Chandra is an Assistant Professor of Public Policy at the Kennedy School of Government at Harvard University, and a Faculty Research Fellow at the National Bureau of Economic Research. Andrew A. Samwick is a Professor of Economics and Director of the Nelson A. Rockefeller Center at Dartmouth College, and a Research Associate at the National Bureau of Economic Research.

We are grateful to the National Institute on Aging (grants P30 AG12810 and R01 AG19805) and the Mary Woodard Lasker Charitable Trust and Michael E. DeBakey Foundation for financial support. We are grateful to Andrew Houtenville for generously providing us with access to CPS data that are not publicly available. We thank David Cutler, Doug Staiger, David Wise, and participants at the NBER conference, "The Decline in Disability," for helpful comments. Any errors are our own. Address correspondence to the authors at Amitabh_Chandra@Harvard.edu and Andrew.Samwick@Dartmouth.edu

examine data from twenty-five years of the Current Population Survey (CPS). We begin by documenting the prevalence of disability in the population as a whole, as well as in subpopulations defined by age, gender, education, race, marital status, census region of residence, and metropolitan status of residence. We show that the prevalence of disability has declined substantially for men over the age of fifty-five, for whom the unadjusted declines have ranged from 15 to 25 percent of their levels twenty years ago (corresponding to an absolute decline of about 4 percentage points). Disability rates have been increasing for women, so that by 2004, disability prevalence was roughly equal for men and women. In the cross-section, the largest disparities occur across educational groups: by age sixty-two, about 17 percent of those without a college education have a work-limiting disability, compared to about 5 percent of those with a college education or higher.

We then consider the implications of differences in disability risk for welfare and saving in a stochastic life cycle model of consumption. We model disability as involuntary retirement, focusing on the economic implications of unexpectedly lower income. In this chapter, we do not consider the impact of a decline in health status on the quality of life that can be purchased with that income. We show that a typical consumer would be willing to pay about 5 percent of lifetime expected consumption to remove the average risk of disability found in the CPS, and perhaps another 4 percent to remove the highest risks we observe in our data. Our simulations also show that the share of preretirement wealth attributable to the average disability risk (for those who do not become disabled) is about 4 percent. For no demographic group that we identify in our empirical work do the simulations suggest that disability risk would account for more than 20 percent of preretirement wealth accumulation. Compared to anticipated drops in income at retirement or annual fluctuations in income that generate similar reductions in expected utility, we note that the risk of disability generates less of a saving response. The reason is that saving is a far less effective hedge against low-probability, high-impact events like disablement.

The remainder of the chapter is organized as follows. In section 10.2, we summarize the data on trends in work-limiting disability over the period from 1980–2004 in the working age population as a whole, as well as for large demographic subgroups. In section 10.3, we present graphical analyses based on logistic regressions that decompose the raw data into age profiles for disability prevalence while controlling for other demographics and year-specific shocks to disability rates. To investigate the implications of these patterns for welfare and saving, we develop a stochastic lifecycle model of consumption decisions in section 10.4. In our model, consumers face three reasons for saving: an anticipated income drop at retirement, persistent uncertainty in their annual incomes, and an annual risk of disability prior to retirement. We model disability as permanent, involuntary retirement. In section 10.5, we show that disability risk has a relatively large impact on welfare and a smaller impact on saving compared to the other motives for saving. Section 10.6 concludes and discusses directions for future work. In the appendix to this chapter, we provide a detailed account of trends in work-limiting disability for different demographic groups.

10.2 Data Description

We use data from the March Current Population Survey (CPS) from 1980–2004 for our analysis. The CPS is a monthly survey of the noninstitutionalized population that is conducted by the Bureau of Census for the Bureau of Labor Statistics. In March of each year, the standard CPS survey is supplemented with additional questions on demographic characteristics, income, program participation, employment, and health insurance. Additionally, there are questions that allow researchers to identify persons with disabilities that limit work. For example, in the last ten years of our sample, the survey asks:

> 59A. (Do you/Does anyone in this household) have a health problem or disability which prevents (you/them) from working or which limits the kind or amount of work (you/they) can do?
> 59B. If yes to 59A., who is that?

All of the measures of work-limiting disability that we use start with an affirmative answer to this question, because we require a definition of disability that changes little over the period from 1980 to 2004. Note that a measure of disability based solely on this question does not restrict respondents to be out of the labor force in order to have a work-limiting disability. In fact, it is possible for a respondent to give an affirmative response to the above question, yet also claim to be working. Furthermore, this definition allows a respondent to be working full-time but in a job other than what they may have chosen in the absence of disability (the question asks for "the *kind* or amount of work"). Therefore, our estimates of work-limiting disability should not be viewed as providing comprehensive estimates of the phenomenon of disability as defined by Americans with Disabilities Act (ADA) legislation (a definition that would include impairments that may not affect an individual's ability to work).[1] Data from the National Health Interview Survey (NHIS) demonstrate that a large fraction of those with impairments do not report having a work-limiting disability. Many of these respondents are potentially covered by ADA legislation but will not contribute to our measure of work-limiting disability.

1. An impairment is a disability under the ADA if it limits a major life activity. Such an activity is not limited to work-related spheres. For example, a college professor who loses the ability to drive a car, sit in a chair, or engage in recreational activities as a result of a back injury would generally be classified as being disabled for the purpose of ADA legislation.

While this restriction causes us to undercount a portion of the disabled population, Burkhauser et al. (2003) demonstrate that some respondents who claim to be work-limited in a single cross-section of the CPS are not permanently work-limited. The authors establish this claim by linking respondents in the CPS in consecutive years and examining the fraction that reported a work limitation in both years. They note that the prevalence of disability as measured by the more stringent two-year restriction results in lower estimates of disability than those obtained by using a single cross-section.

Burkhauser et al. (2003) also demonstrate that even though different surveys such as the CPS, NHIS, and the Survey of Income and Program Participation (SIPP) define disability differently and arrive at different estimates of the *level* of disability in a given year, each of these surveys generates very similar time *trends* in disability. This important finding suggests that there have not been substantial changes since 1980 in the prevalence of impairments (or richer measures of work limitations) that might bias our results. For example, the SIPP solicits information on respondents' Activities of Daily Living (ADLs) and Instrumental Activities of Daily Living (IADLs). These broader questions cover a richer range of limitations concerning mobility, paying bills, and doing light housework and hence raise estimates of measured disability. Therefore, the measurement of disability using SIPP data will yield a higher level of disability prevalence, albeit one with the same trend as that computed using data from the CPS or NHIS.

For the purpose of calibrating our model, we need a measure of disability that corresponds to permanent, involuntary retirement. Our strategy is to start with the work-limitation measure and impose additional conditions to ensure that we are measuring withdrawal from the labor force.[2] An alternative approach to measure permanent disability might use program participation in the Disability Insurance (DI) or Supplemental Security Income (SSI) programs. We are reluctant to pursue this approach; as carefully noted by Autor and Duggan (2003), the DI program has been greatly affected by congressionally mandated changes in the stringency with which DI applicants are screened prior to being classified as bona fide candidates

2. Mashaw and Reno (1996) note that there are over twenty definitions of disability in the literature, each being used for a specific purpose. The appropriateness of each definition should be determined by the context in which it is used. Note that our focus on measuring the probability of involuntary retirement is different from quantifying the prevalence of a physical or mental condition. The latter would be of interest if we were trying to understand improvements in (absolute) health status over time. For example, Lakdawalla, Bhattacharya, and Goldman (2004) use the NHIS to examine trends in the fraction of people with personal care and routine needs limitations. Their definition of disability deliberately abstracts from the work decision. As the nature of work becomes less physical, the prevalence of a work-limiting disability will mechanically decline. Lakdawalla, Bhattacharya, and Goldman are careful to choose a definition of disability that is robust to this transformation. In contrast to their work, if changes in the nature of work reduce the probability of involuntary retirement through lower disability rates, then it is a channel that should be measured.

for the program. Similarly, the decision to apply for DI is influenced by the attractiveness of its economic alternatives (such as unemployment, retirement, or other social programs). Changes in the returns to these alternatives could generate large fluctuations in the measured prevalence of disability. The incentives associated with these alternatives also affect our work-related definition of disability, albeit to a lesser degree, because the question that we use does not directly query program participation.

We therefore impose additional criteria to the work-limiting definition of disability to ensure that it closely corresponds to the notion of involuntary retirement. We explore four definitions, where each subsequent definition is more stringent than the previous one:

1. Respondent has a work-limiting disability.
2. Respondent has a work-limiting disability and is not presently working.
3. Respondent has a work-limiting disability, is not presently working, and did not work last year.
4. Respondent has a work-limiting disability, is not presently working, did not work last year, and is covered by Medicare.

Definition one is correlated with the prevalence of involuntary retirement but does not precisely measure it—other research shows that large numbers of respondents who claim to have a work-limiting disability are indeed working.[3] As such, definition one should be viewed as providing an upper bound on the phenomenon of involuntary retirement, overstating the impact of disability on both welfare and saving.[4] Definitions two and three are undertaken in the spirit of previously discussed work by Burkhauser et al. (2003), who recommend using a longer time frame to measure permanent disability. Finally, because SSDI beneficiaries who have been on the program for more than two years are eligible for Medicare benefits, moving from definition three to four provides us with a lower bound on the prevalence of a work-limiting disability.[5] In our empirical

3. Bound and Burkauser (1999) demonstrate that 65 percent of men who responded positively to the work-limiting disability question on the PSID (in two consecutive years) were working; only 35 percent were not working at all. For women, 52 percent of those who had a work-limiting disability in two consecutive years were working, whereas 48 percent were not working at all.

4. The Bound and Burkhauser study also notes that 38 percent of disabled men (as identified by two years of work-limiting disability in the PSID) received government transfers, whereas 26 percent of disabled women received such income. These estimates illustrate the wide range of estimates that one obtains by using alternative measures of disability. In future research, we will utilize the HRS dataset, which has been matched to the longitudinal Social Security histories of its respondents. With these data, we will be able to more precisely distinguish between disability and involuntary retirement.

5. Medicare is a secondary payer for disabled individuals who are also covered through employer-provided health insurance. If respondents forget to note this secondary coverage while responding to the CPS questionnaire we would understate the measured prevalence of disability from definition four. This is another reason for why definition four may constitute an absolute lower bound on the prevalence of a work-limiting disability.

work we focus on definition three and use definitions one and four to inform us of the maximum and minimum disability probabilities. Definitions two and three yield similar disability prevalence rates, but definition three has the advantage of being more conservative. Figure 10.1 shows the (regression-adjusted) age-disability profile for each of these definitions of disability.

We utilize all observations in the CPS for respondents who were aged twenty-two to sixty-four at the time of the survey. We choose twenty-two as our lower bound in order to minimize sampling college students while still being able to provide meaningful estimates of work-limiting disability in younger populations. We end our analysis at sixty-four, since almost all respondents claim to be retired after that age. We extract a sample of 2,166,178 respondents in this age range over the 1980–2004 time period. For a given year, our sample sizes range from 85,133 in 1980 to over 110,000 since 2002.

Fig. 10.1 Regression adjusted disability prevalence by age

Notes: Figures report predicted probability of work-limiting disability using alternative definitions from a logistic regression that controls for an unrestricted set of age and year indicator variables, gender, race, and ethnicity (four categories), education (two categories), marital status (three categories), marital status and gender interactions, census region (nine categories), metropolitan status, census division, and metropolitan status interactions, the number of children under the age of eighteen in the family, and the size of the household. The above figure adjusts for all these covariates, except for age. The regression used CPS data from the 1980–2004 files of Annual Demographic Survey (n = 2,166,178).

In table 10.1, we present trends in the prevalence of self-reported disability from 1980–2004.[6] All tables use our third definition of disability, where the respondent has a work-limiting disability, is not presently working, and did not work last year. We present separate estimates by gender and two broad education categories. We also report a regression estimate of the annual, linear trend in disability. Over these two decades, the prevalence of disability increased by 0.5 percentage points for the population aged twenty-two to sixty-four (with an annual increase of 0.036 percentage points per year).[7] For men, work-limiting disability increased by 0.7 percentage points, while it increased by 0.4 percentage points for women. Despite starting at different levels in the early 1980s, the gender gap in disability closed by 2004, with women having slightly higher disability rates than men in the most recent period. Disability rates increased for men and women in both education groups. The increase for less-educated women and the increase for less-educated men were particularly pronounced.

Table 10.1 also separates trends in disability prevalence by race/ethnicity and gender and by marital status and gender. The increase in disability for women, noted above, is similar for white and black women, but it is contrasted with a decline for Hispanic women. Disability prevalence increased more for black men than white men, but there was a decline for Hispanic men.

The last panel of table 10.1 illustrates the importance of focusing on the prevalence of disability within narrowly defined demographic groups. As noted earlier, there was an increase in the overall prevalence of disability for males. The trends described in this panel note that there was no change in disability for married males. The increase for men is the consequence of substantial increases in disability for never married and widowed, separated, or divorced men. Similarly, there has not been a change in disability rates for married women, but there have been larger increases for never married women as well as those who are widowed, separated, or divorced.

6. Our results are similar to other studies that have utilized a single year of the CPS to measure disability. For example, Burkhauser and Daly (1996) note that in 1990, the prevalence of disability among CPS respondents aged twenty-five to sixty-one was 8.1 percent for men, and 7.8 percent for women. Our sample includes respondents aged twenty-two to sixty-four, and the corresponding numbers for men and women are 8.1 and 7.5 percent, respectively. Using a similar question on work-related disability in a single cross-section from the SIPP for the same year provided estimates of 11.4 and 11.3, respectively. As noted by Bound and Burkhauser (1999), the higher rates in the SIPP are the likely consequence of explicitly referring to mental health conditions in contributing to work-limiting disability.

7. Note that an average annual increase of 0.036 percentage points over twenty-five years would result in a 0.9 percentage point increase in disability. Table 10.1 demonstrates that the actual change in disability over this period was 0.5 points. The regression line overpredicts the increase in disability for this time interval, because it imposes a linear time-trend on the data; the true trend in disability is best described by a nonlinear time trend. For most of the demographic groups examined in this chapter, we found it difficult to reject the linear specification. Therefore, we used this specification throughout the chapter.

Table 10.1 Trends in the prevalence of work-limiting disability

	By education level and gender								
	No college			College or more			Total		
Survey year	Males	Females	Total	Males	Females	Total	Males	Females	Total
1980–1984	5.9	6.0	5.9	1.2	1.5	1.3	4.7	5.2	5.0
1985–1989	5.8	5.7	5.7	1.1	1.2	1.2	4.5	4.8	4.7
1990–1994	6.0	5.8	5.9	1.1	1.1	1.1	4.7	4.8	4.7
1995–1999	6.5	6.7	6.6	1.3	1.5	1.4	5.1	5.4	5.3
2000–2004	7.0	7.1	7.0	1.4	1.7	1.6	5.4	5.6	5.5
Annual trend	0.060	0.068	0.065	0.017	0.020	0.019	0.040	0.032	0.036
	By race and ethnicity								
	Whites			Hispanic			Blacks		
	Males	Females	Total	Males	Females	Total	Males	Females	Total
1980–1984	4.2	4.6	4.4	4.6	5.8	5.2	9.0	9.3	9.2
1985–1989	3.9	4.3	4.1	5.1	5.0	5.0	8.8	7.9	8.3
1990–1994	4.1	4.2	4.2	5.2	4.9	5.0	8.8	8.1	8.4
1995–1999	4.4	4.8	4.6	4.9	5.6	5.2	10.3	9.2	9.7
2000–2004	4.9	5.1	5.0	4.3	4.9	4.6	10.3	9.3	9.8
Annual trend	0.041	0.032	0.037	–0.029	–0.021	–0.025	0.087	0.027	0.054
	By marital status and gender								
	Married			Widowed, separated, or divorced			Never married		
	Males	Females	Total	Males	Females	Total	Males	Females	Total
1980–1984	4.0	3.8	3.9	8.7	10.4	9.9	5.3	5.1	5.2
1985–1989	3.6	3.5	3.6	8.1	9.5	9.0	5.6	4.3	5.1
1990–1994	3.6	3.4	3.5	8.6	9.5	9.2	5.6	4.5	5.2
1995–1999	3.7	3.9	3.8	9.6	10.6	10.2	6.4	5.4	6.0
2000–2004	3.7	3.8	3.7	10.6	10.9	10.8	6.9	6.2	6.6
Annual trend	–0.009	0.007	*–0.001*	0.116	0.045	0.069	0.088	0.089	0.083

Notes: Tables are constructed from CPS data from the 1980–2004 files of Annual Demographic Survey (n = 2,166,178). Respondents are aged twenty-two to sixty-four at the time of the survey. Whites refer to non-Hispanic whites. Annual trends are calculated as the coefficient on a linear-regression of disability for the relevant demographic group on a linear time-trend. Italicized trends are not different from zero at the 5 percent significance level.

In interpreting these estimates, it is important to recognize that the reported trends are not necessarily causal—there is nothing in our analysis that suggests that being unmarried raises the probability of having a work-limiting disability. Rather, the fact that a respondent is unmarried at age fifty-five is better interpreted as being a marker for a particular type of person. Indeed, causality could easily flow in the other direction; it may be the case that because a person is less healthy, he or she is unmarried. With this caveat in mind, we note that the trends documented in table 10.1 (and the next three tables discussed in the following paragraphs) are possibly being driven by compositional shifts in the categorization of workers. This interpretation is best understood by reexamining the results of the first panel in table 10.1. Here, we noted an increase in the disability prevalence of respondents who did not have a college degree. As the fraction of the population with a college degree increases, the group without a college degree is increasingly representative of individuals who are drawn from the lower tails of the skill and health distributions. Such persons are more likely to have attended worse schools, lived in worse neighborhoods, or been affected by credit constraints that restricted investments in health and human capital. As a result, they also may be more likely to be disabled. This framework suggests that even if there were no changes in the prevalence of disability over time conditional on educational attainment, a growth in the college-educated population will manifest itself as increasing disability prevalence for the less-educated.[8]

Age-specific tabulations in table 10.2 demonstrate that the declines in disability have been greatest for the oldest men in our sample. Men are approaching retirement age healthier today than at any point in recent history—the prevalence of work-limiting disability has fallen substantially over the past two decades. In contrast to the declines observed for older populations, there were increases in disability for young and middle-aged men. The results for women show a more modest decline in disability rates in the oldest age group and a comparable increase in disability rates for women aged twenty-two to thirty-nine. In contrast to the results for men, there has been no change in work-limiting disability rates for women aged forty to fifty-nine.[9]

When these age patterns are analyzed by educational attainment in the second panel of table 10.2, we find that among respondents aged twenty-two to thirty-nine, there have been large increases in disability for those without a college degree and no change for the college-educated. For respondents aged forty to fifty-nine, changes in disability prevalence are

8. The same compositional shifts are consistent with the increasing disability rates of those with college educations—the subpopulation of college-educated now includes less healthy people who formerly would have lacked a college education.

9. It is instructive to compare our results to those obtained by Lakdawalla, Bhattacharya, and Goldman (2004), who use alternative measures of disability reported in the National Health Interview Survey (NHIS). Their measure, unlike ours, measures disability prevalence on an absolute scale (by examining prevalence of ADL and IADL limitations). Using these data for the 1984–1996 period, they find no decline in disability rates for those aged sixty to sixty-nine but note a 0.15 percentage point increase for those respondents aged eighteen to twenty-nine. Using the same data for the 1997–2000 period, they note a 0.19 percentage point reduction in disability for those aged sixty to sixty-nine, and a 0.2 percentage point decline for those aged eighteen to twenty-nine.

Table 10.2 **Age-specific trends in the prevalence of work-limiting disability, by gender and education**

	By gender								
	Ages 22–39			Ages 40–59			Ages 60–64		
Survey year	Males	Females	Total	Males	Females	Total	Males	Females	Total
1980–1984	1.7	1.8	1.8	6.2	7.3	6.7	18.2	17.0	17.6
1985–1989	2.0	1.7	1.9	5.8	6.6	6.2	16.4	16.1	16.2
1990–1994	2.3	2.0	2.2	6.1	6.4	6.2	15.3	14.9	15.1
1995–1999	2.5	2.4	2.4	6.5	7.1	6.8	16.1	16.5	16.3
2000–2004	2.4	2.4	2.4	6.8	7.1	7.0	13.9	14.8	14.4
Annual trend	0.040	0.037	0.039	0.041	*0.008*	0.024	–0.176	–0.085	–0.128
	By education								
	Ages 22–39			Ages 40–59			Ages 60–64		
	No college	College	Total	No college	College	Total	No college	College	Total
1980–1984	2.2	0.5	1.8	7.8	1.8	6.7	18.8	8.5	17.6
1985–1989	2.4	0.5	1.9	7.5	1.6	6.2	17.9	5.9	16.2
1990–1994	2.7	0.4	2.2	7.8	1.4	6.2	16.9	5.9	15.1
1995–1999	3.1	0.5	2.4	8.6	1.9	6.8	18.9	5.7	16.3
2000–2004	3.2	0.5	2.4	9.0	2.1	7.0	17.2	5.4	14.4
Annual trend	0.057	*0.000*	0.039	0.071	0.022	0.024	–0.050	–0.109	–0.128

Notes: Tables are constructed from CPS data from the 1980–2004 files of Annual Demographic Survey (n = 2,166,178). Respondents are aged twenty-two to sixty-four at the time of the survey. Whites refer to non-Hispanic whites. Italicized trends are not different from zero at the 5 percent significance level.

driven by an annual trend of 0.07 percentage points for those without college degrees and an annual trend of 0.02 for those with a college degree. For the oldest group in our analysis, those aged sixty to sixty-four, the prevalence of disability fell for persons regardless of educational attainment, but it fell more for the college-educated.

In table 10.3, we disaggregate the trends in disability by age and race in the first panel and by age and marital status in the second. The decline in disability prevalence that was noted for Hispanics is shown to have occurred across all age groups. For blacks, disability rates increased for those below forty and decreased for those above sixty, with the middle age group showing no significant trend. The estimated disability rate of 29.2 percent for blacks aged sixty to sixty-four in 1995–1999 is the single highest estimate in our analysis. By marital status, married people have the lowest rates of disability and the most favorable trends over time. All marital status groups had significant increases in disability rates in the lowest age group.

The most important trend in the labor market relevant to disability over this period is the changing nature of work. As the share of jobs in sectors

Table 10.3 **Age-specific trends in the prevelence of work-limiting disability, by race and marital status**

	By race											
	Ages 22–39				Ages 40–59				Ages 60–64			
Survey year	Whites	Hispanics	Blacks	Other	Whites	Hispanics	Blacks	Other	Whites	Hispanics	Blacks	Other
1980–1984	1.4	2.1	4.0	1.4	5.8	8.2	14.0	5.3	16.3	20.5	28.3	15.5
1985–1989	1.5	2.3	3.9	1.3	5.3	7.6	12.4	5.2	14.8	18.9	27.5	13.5
1990–1994	1.8	2.3	4.1	1.9	5.3	7.9	12.1	4.9	13.5	18.8	26.6	12.5
1995–1999	2.0	2.4	4.9	1.8	5.7	8.0	13.5	6.0	14.3	20.6	29.2	15.9
2000–2004	2.2	1.9	4.6	1.7	6.1	7.1	13.1	5.8	12.7	17.9	25.9	11.6
Annual trend	0.043	–0.014	0.046	0.02	0.045	–0.046	*–0.005*	0.045	–0.154	–0.086	–0.090	*–0.091*

	By marital status											
	Ages 22–39			Ages 40–59			Ages 60–64			Total		
	Married	Widowed, separated or divorced	Never married	Married	Widowed, separated or divorced	Never married	Married	Widowed, separated or divorced	Never married	Married	Widowed, separated or divorced	Never Married
1980–1984	0.9	3.2	3.2	5.0	12.1	15.3	14.9	25.2	22.7	3.9	9.9	5.2
1985–1989	0.9	3.1	3.4	4.5	10.7	14.1	13.5	23.6	22.8	3.6	9.0	5.1
1990–1994	1.1	3.6	3.5	4.4	10.7	12.9	12.1	22.7	23.7	3.5	9.2	5.2
1995–1999	1.3	4.0	3.8	4.5	11.4	14.0	12.9	24.8	25.8	3.8	10.2	6.0
2000–2004	1.2	4.2	3.6	4.5	11.6	14.6	10.6	22.9	24.4	3.7	10.8	6.6
Annual trend	0.021	0.066	0.027	–0.021	*0.009*	*0.003*	–0.183	–0.070	*0.090*	*–0.001*	0.069	0.083

Notes: Tables are constructed from CPS data from the 1980–2004 files of Annual Demographic Survey (n = 2,166,178). Respondents are aged twenty-two to sixty-four at the time of the survey. Whites refer to non-Hispanic whites. Annual trends are calculated as the coefficient on a linear-regression of disability for the relevant demographic group on a linear time-trend.

of the economy with high work-related injuries falls (for example, jobs in manufacturing and mining), it is perhaps unsurprising that the prevalence of work-limiting disability declines for older men. As the nature of work changes and fewer jobs require physical strength and dexterity, it is possible that fewer individuals are work-limited. Kutscher and Personick (1986) document a decline in manufacturing from 25.1 to 18.5 percent of total employment between 1959 and 1984. Mining fell from 0.9 to 0.6 percent of total employment over the same period of time. At the end of 2004, the Bureau of Labor Statistics (2005) reports that manufacturing represents just 10.8 percent of total employment and that mining represents just 0.4 percent of total employment. Of the manufacturing workers, only 70 percent are classified as production workers.

We further consider this hypothesis by examining the trends in disability prevalence by census region in table 10.4.[10] Across all ages, work-limiting disability grew the most in New England and East South Central states. For all regions, there were significant increases in disability rates for respondents aged twenty-two to thirty-nine. Among respondents aged forty to fifty-nine, the only regions not to see a significant increase in disability rates are those in the East North Central and West North Central. For those aged sixty to sixty-four, all regions saw declines in disability rates, though the trends in the New England, Middle Atlantic, and East South Central regions were not significant. For this age group, the largest declines in disability are observed in the Mountain region (comprised of Arizona, Colorado, Idaho, Montana, Nevada, New Mexico, Utah, and Wyoming); this decline is consistent with an interpretation wherein new cohorts of sixty to sixty-four-year-olds are less likely to have been exposed to work in the mining sector. Similarly, for the changing nature of work hypothesis to be true, we should see increases in disability prevalence in regions of the country where manufacturing employment grew in the past twenty years. The East South Central region of the country is one such area (where large Japanese car manufacturers located in the mid-1980s). In this region, we observe an increasing trend in the level of disability (of 0.08 percentage points per year), whose magnitude is larger than that of any other region.

While suggestive, this characterization of the data is not without exceptions. First, note that table 10.1 demonstrates that there were increases in the levels of work-limiting disability for men with a college degree—a group that was never at risk for working in a coal mine or on the assembly line of a manufacturing plant. Disability prevalence has also grown for college-educated women, another group that is at low risk for job-related

10. The nine regions are defined as follows—Mountain: AZ, CO, ID, MT, NV, NM, UT, WY. New England: CT, ME, MA, NH, RI, VT. South Atlantic: DE, DC, FL, GA, MD, NC, SC, VA, WV. West North Central: IA, KS, MN, MO, NE, ND, SD. West South Central: AR, LA, OK, TX. Pacific: AK, CA, HI, OR, WA. East North Central: IL, IN, MI, OH, WI. East South Central: AL, KY, MS, TN. Middle Atlantic: NJ, NY, PA.

Table 10.4 **Age-specific trends in the prevelence of work-limiting disability, by census region**

Survey year	N. Eng	Mid Atl	E.N.Cent	W.N.Cent	S.Atl	E.S.Cent	W.S.Cent	Mount.	Pacific
				Ages 22–39					
1980–1984	1.7	2.1	1.7	1.3	2.0	2.3	1.5	1.1	1.7
1985–1989	1.2	2.2	2.0	1.2	1.8	2.7	2.0	1.5	1.9
1990–1994	1.7	2.1	2.2	1.5	2.3	3.8	2.2	1.8	2.0
1995–1999	2.4	2.8	2.5	1.6	2.6	3.8	2.3	1.7	2.1
2000–2004	2.9	3.0	2.3	1.7	2.4	3.7	2.2	1.6	2.2
Annual trend	0.075	0.049	0.036	0.028	0.033	0.084	0.036	0.024	0.026
				Ages 40–59					
1980–1984	5.3	6.6	6.6	4.9	7.4	10.2	6.9	5.3	6.5
1985–1989	4.9	6.1	6.1	4.5	7.1	9.4	6.4	5.3	5.7
1990–1994	5.0	5.7	6.0	4.3	6.6	11.2	6.9	5.2	5.8
1995–1999	6.8	6.8	6.6	4.6	7.5	10.1	7.1	5.3	6.0
2000–2004	6.3	7.0	6.5	4.8	7.6	10.9	7.4	5.8	6.5
Annual trend	0.078	0.031	*0.008*	*0.003*	0.021	0.037	0.037	0.025	0.014
				Ages 60–64					
1980–1984	14.5	15.9	16.5	14.9	19.9	21.3	20.1	17.5	17.6
1985–1989	13.4	14.0	15.9	12.7	17.5	23.8	18.9	15.5	15.6
1990–1994	13.5	13.5	14.6	11.8	16.6	19.4	17.6	14.1	14.5
1995–1999	12.7	15.0	15.4	12.3	18.2	21.6	19.4	14.8	15.6
2000–2004	13.9	14.4	13.3	11.0	14.6	20.9	17.0	11.6	13.4
Annual trend	*–0.033*	*–0.055*	–0.131	–0.163	–0.208	*–0.061*	–0.116	0.243	–0.168
				All ages					
1980–1984	4.2	5.2	4.8	3.8	5.7	6.9	4.9	3.8	4.6
1985–1989	3.6	4.8	4.7	3.3	5.2	7.1	4.8	3.9	4.2
1990–1994	3.9	4.5	4.7	3.4	5.1	8.0	5.1	4.0	4.3
1995–1999	5.1	5.5	5.2	3.7	5.9	7.9	5.5	4.1	4.5
2000–2004	5.4	5.9	5.2	3.9	5.9	8.6	5.7	4.3	4.9
Annual trend	0.082	0.041	0.027	0.016	0.027	0.084	0.046	0.026	0.023

Notes: Tables are constructed from CPS data from the 1980–2004 files of Annual Demographic Survey (n = 2,166,178). Respondents are aged twenty-two to sixty-four at the time of the survey. Whites refer to non-Hispanic whites. Annual trends are calculated as the coefficient on a linear-regression of disability for the relevant demographic group on a linear time-trend. Italicized trends are not different from zero at the 5 percent significance level.

injuries. Second, the growth in the prevalence of work limitations for women also poses a problem for an explanation of the decline based on the changing nature of work. At first glance, it might be plausible to explain the increase in female disability as stemming from growing female labor force participation (which grew from 51.6 percent in 1980 to 60.2 percent in 2003). While women of all ages increased their participation between 1980 and 2003, the largest absolute increases in participation occurred for women aged forty to fifty-four (for whom labor force participation grew from 61 percent to almost 75 percent). However, disability trends for this group stayed absolutely constant over the 1980s and 1990s (table 10.2, first panel). Finally, we note that the increase in disability in New England states, a group of states that has not witnessed an increase in manufacturing jobs, is comparable to the increase in East Central States. Therefore, no simple explanation is able to reconcile all the facts.

The preceding analysis compares the reported disability rates of respondents of different cohorts at the same age. In results not reported in this chapter, we examined changes in the age-disability profile by birth cohort. Because each cohort is observed at different points in its life cycle, we can compare the prevalence of disability at a certain age for a given cohort to that for another cohort at the same age. By examining the age-disability trend across cohorts, we can assess the extent to which different cohorts have the same age-disability profile. In general, we note only small differences across cohorts in disability through the age of fifty-four. These results are supportive of a common age profile across cohorts—disability rises monotonically with age, and the levels appear to be markedly similar at a given age across cohorts. It does not appear to be the case that a certain group of cohorts is systematically more or less disabled than another cohort, an empirical finding that reduces the potential role of cohort effects in the data.

We conclude our analysis of the trends in disability prevalence by reporting the extent to which there has been convergence (or divergence) in disability rates across the different demographic groups studied above. These results, which provide a succinct summary of the extent to which the levels of disability across different groups are becoming more or less homogeneous over time, are reported in table 10.5. The table is best explained by an example—row 1 of the table indicates that the variance in disability rates across the three age categories that we have focused on is estimated to have fallen by 25 percent in ten years (multiplying the change in variance and its standard error by 2.5 will produce the decline in variance over the twenty-five-year study period). Over the same period, the variance in disability between men and women increased by 8 percent (an estimate not statistically different from zero). Alternatively, in row 8, we note that the variance in disability across age and gender categories (that is, three age categories multiply by two gender categories = six age and gender cate-

Table 10.5 Convergence in disability prevalence by demographic group

Demographic group	Percent change in variance every 10 years	Standard error
1. Age	–25	(4.0)
2. Gender	8	(32.0)
3. Education	46	(3.0)
4. Race	22	(11.0)
5. Marital status	85	(11.0)
6. Census division	17	(15.0)
7. Metropolitan residence	177	(73.0)
8. Age and gender	–25	(4.0)
9. Age and education	–1	(4.0)
10. Age and race	–69	(14.0)
11. Age and marital status	42	(12.0)
12. Age and census division	–25	(4.0)
13. Age and metropolitan status	–32	(4.0)
14. Gender and education	45	(3.0)
15. Gender and race	31	(10.0)
16. Gender and marital status	80	(11.0)
17. Gender and census division	25	(15.0)
18. Gender and metropolitan status	–25	(29.0)

Notes: Table reports the percent change in the variance of disability over time. For each demographic group the table reports the coefficient from a regression of ln (variance of disability) on a linear time-trend. See text for details. There are three categories for age, two for education, four for race, three for marital status, two for metropolitan residence, and nine for census division.

gories) fell by 25 percent. This result implies that disability rates between older and younger men, and older and younger women, are becoming more similar. These estimates were calculated by regressing the natural logarithm of the variance in disability, calculated for each year of the data, across these demographic cells on a linear trend.[11]

The table indicates that there has been dramatic convergence in disability rates by age, but divergence across education, marital status, and metropolitan residence groups. There is no statistically significant change in variance across gender, racial groups, or by census division. In other words, it is not the case that the disability rates for blacks, Hispanics, and whites are becoming more similar, or that disability rates across the nine census regions are more uniform. In rows 8 through 18, we examine convergence in disability rates across more finely defined groupings. We are

11. In this regression, the coefficient on the linear trend measures the percent change in variance *over time*. This interpretation is important to remember, as it implies that the magnitude of the results in table 10.5 say nothing about the level of variance at a point in time. We estimate this regression by using all twenty-five years of the CPS data and computing the variance of disability in each year. We regress the log of this variance on a linear time-trend. We report the coefficient on this trend (multiplied by a factor of 10 to assist in interpretation).

primarily interested in understanding whether the convergence results noted above for age and gender groups persist in conjunction with the convergence or divergence trends for other variables. Rows 8 through 13 demonstrate that the convergence in disability rates across age groups persists even when we examine convergence in disability across age and gender groups, age and race groups, age and census division groups, and age and metropolitan status groups. In contrast, there has been no convergence across age and education groups, and there has been divergence across age and marital status categories. The latter result is unsurprising in light of the fact that there was substantial *divergence* across (univariate) educational and marital status groups, which was partially offset by the convergence across age groups. Rows 14 through 18 of table 10.5 note that there has been divergence in disability rates for groups defined by gender and education, gender and race, and gender and marital status. In other words, the disability rates of black men, white men, black women, and white women are diverging.

10.3 Regression Analysis

The results described in the previous section suffer from an important limitation in that the demographic composition of the population has been changing over time. For example, the age structure of the population has been changing—a large cohort of Baby Boomers born in the years after World War II is moving through their life cycle and affecting the age distribution of the population at each point between 1980 and 2004. Such an effect would contaminate the identification of a declining trend in disability over time, since a secular decline in disability would be partially offset by the increasing age (and consequent disability) of the Baby Boomers. In addition, as documented in section 10.2, there have been changes in education, marital status, and the composition of employment. To account for these factors and thus better understand the relationship between age and the prevalence of a work-limiting disability, we follow a regression-based approach where we estimate the following logistic regression model for the probability of being disabled for individual i, at age j in year t (Δ_{ijt}):

$$\text{(1)} \qquad \Pr(\Delta_{ijt} = 1) = F\left(\beta_0 + \mathbf{X}_i' \Theta + \sum_j \gamma_j Age_j + \sum_t \delta_t Year_t\right).$$

In this model, we have included an unrestricted set of age and year dummies, while controlling for factors such as gender, education, race, region of residence, metropolitan residence, marital status, and family size and composition. This specification allows for complete nonparametric identification of the age and year effects. In later specifications, we allow for the age effects to vary by different demographic controls by interacting each of the age indicator variables with different values of the control variables.

We begin by reporting the results of our standard specification (where the age effects are constrained to be the same across all demographic groups) in figure 10.1. The adjusted age-disability profiles (one for each of the four definitions of work-limiting disability), demonstrate a convex and monotonically increasing relationship in age. The age-disability profile that is the focus of this chapter is illustrated with an uninterrupted line. This age-disability profile is what would be observed if secular (i.e., year-specific) changes in disability, as well as in the demographic mix of the CPS sampling frame, were accounted for. The profile starts at approximately 2 percent at the age of twenty-two and increases to 15 percent at the age of sixty-four. The profile increases roughly linearly with age until the age of forty-two, at which point the relationship becomes increasingly convex. Because of the enormous sample sizes available at our disposal, the profile is very precisely estimated (the standard errors are always less than 0.5 for the estimated profiles, even at older ages). For this reason, we suppress the confidence intervals in each of the graphs.

While this model is an improvement over examining unadjusted data, it suffers from two potential criticisms. First, the effect of age on disability may vary over time. This is a testable hypothesis, and we explore its empirical content in figure 10.2. Here, we modify the logistic regression estimated in equation (1) by grouping adjacent years and ages so that we can also include year and age interactions. While these interactions are statistically significant (and have a marginal F-statistic over thirty), the age-disability profile that allows for flexibility over time is seen to be remarkably stable over time. There is a small decline of 1 percentage point in the prevalence of disability of respondents aged sixty to sixty-four over the sample period, but for ages twenty-two to fifty-nine, the profiles are stable. As such, the economic significance of the age-disability profile varying over time is negligible. These results provide persuasive evidence that the age-disability profile has remained stable despite onetime events such as the ADA and more than two complete business cycles. Furthermore, they also demonstrate that the trends in disability noted in tables 10.1 through 10.4 are the consequence of changes in the composition of the labor force. When we control for the demographic characteristics of respondents, the increase in disability for younger populations is entirely eliminated.

A second criticism of the model underlying figure 10.1 is that it is possible that trends in work-limiting disability are a consequence of cohort effects. Such effects would represent cohort-specific changes in the disability profile that are not related to age or period effects. For example, if there is a disease (say, the influenza epidemic of 1918) or medical breakthrough that affects some cohorts but not others, such effects will be cohort-specific. As is well known in the economics literature (Heckman and Robb 1985), it is not possible to separate age, period, and cohort effects simultaneously without further assumptions. For example, in Deaton and Paxson

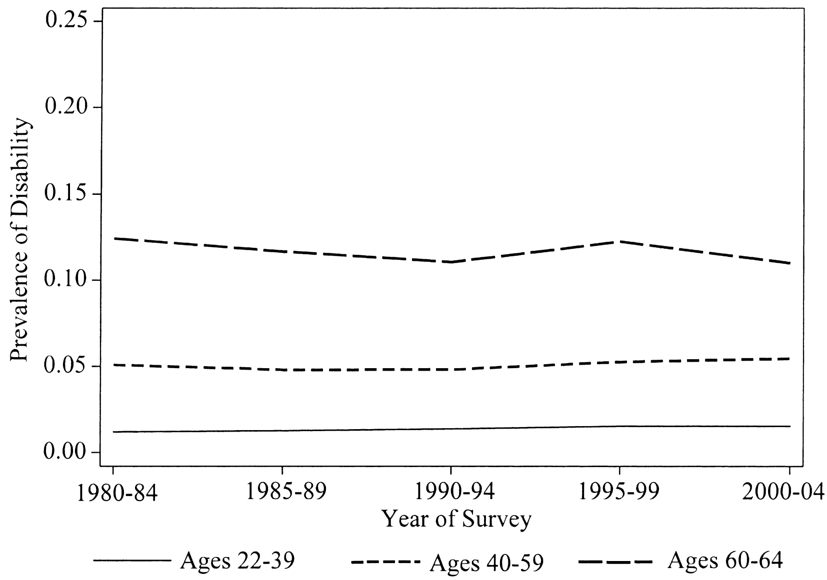

Fig. 10.2 Regression adjusted disability prevalence by age over time

Notes: Figures report predicted probability of work-limiting disability from a logistic regression that controls for an unrestricted set of age and year indicator variables, gender, race, and ethnicity (four categories), education (two categories), marital status (three categories), marital status and gender interactions, census region (nine categories), metropolitan status, census division, and metropolitan status interactions, the number of children under the age of eighteen in the family, and the size of the household. The above figure adjusts for these covariates, except the ones used in the figure. The regression used CPS data from the 1980–2004 files of Annual Demographic Survey (n = 2,166,178).

(1994), the key identifying assumption is that the year effects sum to zero and are orthogonal to a linear time trend. The Deaton-Paxson assumption is plausible for their study of consumption, but is probably inappropriate for the study of disability.

In our analysis, we assume that cohort effects are zero. Our logic is as follows: there is a powerful biological case for including age effects in our models—disability increases with age because of the onset of illness and muscular-skeletal deterioration that is a consequence of the aging process. In addition, the deleterious effects of disability are often cumulative—although not all disabilities are irreversible, many of the more severe ones probably are. Furthermore, we argue that there is also a case for including year effects in our model. Economy-wide changes in the nature of work (for example, from manufacturing and mining jobs to service jobs) will affect the prevalence of disability. Additionally, legislative interventions such as the ADA will manifest themselves through year effects. Therefore, while there are strong a priori reasons for including age and year effects, the case for cohort effects is less clear on a prima facie basis.

Having established that the age-disability profile is stable over time, we explore trends in prevalence of work-limiting disability by gender in figure 10.3. In this figure (and subsequent ones) we allow the age profile to vary by demographic characteristics (here, gender and time.) The graphs adjust for time effects, education attainment, race, marital status, census region of residence, metropolitan residence, and family size and composition. The graphs clearly illustrate the decline in disability for men over the age of fifty-five in the past twenty years. In contrast; the disability profile for women demonstrates the increase in work-related disability for women aged thirty-five to fifty. In the unadjusted data reported in table 10.1, men and women had similar disability rates in 2004, with women reporting substantially larger increases in disability prevalence between 1980 and 2004. In contrast, figure 10.3 demonstrates that those trends are the consequence of compositional changes; after year effects and other demographic characteristics are controlled for, the change in disability profiles for women is not evident. It is interesting to note that these profiles are similar for men and women until the age of fifty. After that, the age-disability profile for men accelerates upwards.

We study differences in the age-disability profile by educational attainment in figure 10.4 and by racial group in figure 10.5. Less-educated workers have a higher level of disability at all points in their life cycle. Additionally, the disability hazard is greater for these workers—at any point in the age distribution, less-educated workers have a larger probability of becoming disabled. Therefore, a small difference in the initial level of

Fig. 10.3 Regression adjusted disability prevalence by gender over time
Note: See note to fig. 10.2.

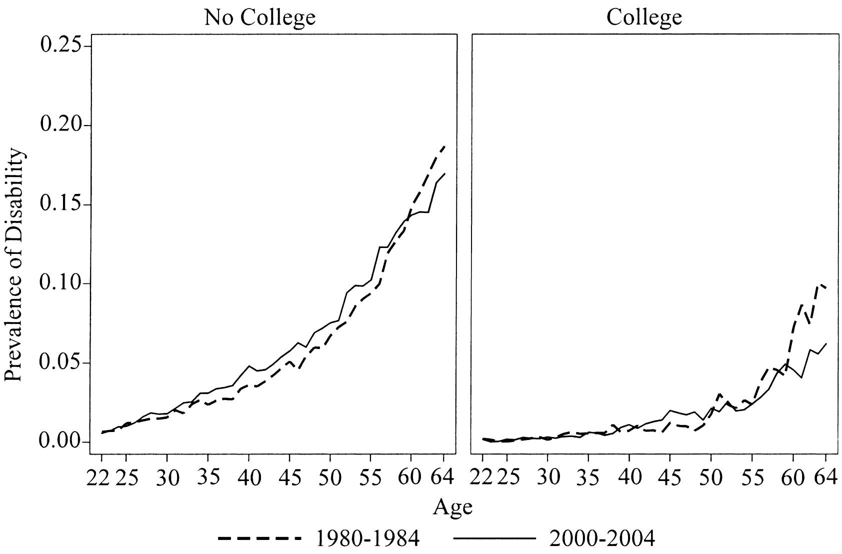

Fig. 10.4 Regression adjusted disability prevalence by education over time
Note: See note to fig. 10.2.

disablement at the age of twenty-two is converted to a 10-percentage point difference at the age of sixty-four. For both educational levels, the decline in disability is most pronounced for those over the age of sixty. Figure 10.5 illustrates disability prevalence profiles by racial group. Despite starting at very similar levels at the age of twenty-two, the hazard of reporting a work-limiting disability is much greater for blacks and Hispanics relative to whites (the differences between whites and nonblacks/non-Hispanics, i.e., those grouped together as "other," are not statistically significant). On the eve of retirement, the probability of having a disability for whites is approximately 12 percent in 2004. Yet it is 15 percent for Hispanics and 20 percent for blacks. In the appendix figures, we describe age-disability profiles by marital status, over time (figure 10A.1), marital status and gender (figure 10A.2), education and gender (figure 10A.3), and metropolitan residence (figure 10A.4).

To summarize, the prevalence of a work-limiting disability at a typical age of retirement, across all demographic groups and for the average year in our sample, is 15 percent (fig. 10.1). The lowest prevalence of 6 percent at retirement is observed for college-educated persons in the 2000–2004 period. The highest observed prevalences are 26 percent for never married men (see appendix fig. 10.A2), and 20 percent for black males. These estimates span the range of disablement probabilities and will therefore constitute a central input to our model of intertemporal consumption and saving.

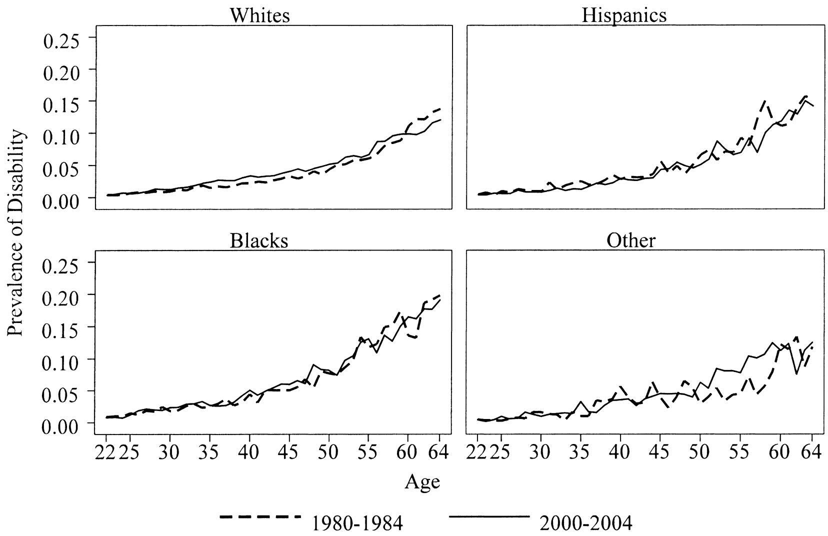

Fig. 10.5 Regression adjusted disability prevalence by race over time
Note: See note to fig. 10.2.

10.4 Optimal Consumption in a Stochastic Life Cycle Model with Disability

Our empirical work with the CPS documents cross-sectional and time series differences in disability prevalence. We now seek to quantify the importance of these differences by evaluating their impact on expected utility (welfare) and saving for representative consumers.

10.4.1 Intuition

Our main analytical tool is a model of household consumption decisions based on a life cycle framework, in which a primary motivation for saving is to transfer resources (income) from the working years, when income is predictably high, to the retirement years, when income is predictably low. In this model, the household makes a consumption decision in each period (taken to be a year) based on resources currently available (assets) and expectations of income to be received from continued work and retirement in the future. The household will be more inclined to consume resources today when it has more assets available or when it has a higher predicted income in the future.

We add to this basic framework two sources of uncertainty in future income, to make it a stochastic life cycle model. The first is the annual variability in income that is attendant to risks of unemployment and the trajectory along which people advance in their careers. Annual variability in income will lower consumers' well-being if they are risk averse (i.e., if they would be willing to pay a premium for actuarially fair insurance to remove the uncertainty that they face). With a suitable choice of preferences, consumers will also choose to increase their saving in response to increased uncertainty as a way to buffer their consumption against income uncertainty. It is now standard in the economics literature to incorporate this sort of precautionary saving against annual income variability into analyses of consumption and saving.

Unlike the existing literature on consumption and saving, we add a second source of uncertainty to the model in the form of an annual probability that working-age consumers will experience a disabling event that requires them to permanently withdraw from the labor force. The consumer must then rely on accumulated saving and future payments from disability insurance programs to support future consumption. As with annual variability in income, the risk of becoming disabled makes consumers worse off and may elicit higher savings to help mitigate the reduction in future consumption that results from disability.

The remainder of this section presents the more technical specification of the consumption model and a derivation of its key results. The latter are answers to two key questions: by how much are consumers worse off when their risk of disability increases, and by how much do they adjust their saving in response to that higher risk? As a starting point, we develop a stylized model of intertemporal consumption that solves for optimal profiles of consumption while disabled (C_s^d) and nondisabled (C_s^n). The consumer's value function is defined as:

$$V_t(A_t, Y_t, \Delta_t) \equiv \max_{\{C_s^d, C_s^n\}} E_t \sum_{s=t}^{T} \beta^{s-t}[\Delta_s u(C_s^d) + (1 - \Delta_s)\, u\,(C_s^u)]$$

such that:

$$(2) \qquad (a)\ u(C) = \frac{C^{1-\gamma}}{1-\gamma}$$

$$(b)\ A_{s+1} = [1 + r(1 - \tau)]\,[A_s + Y_s\,(1 - \tau) - C_s]$$

$$(c)\ A_s \geq 0\ \forall s$$

$$(d)\ Pr(\Delta_s = 1) = \begin{cases} 1 - \prod_{q=t}^{s}(1 - \delta_q), \text{ for } \Delta_t = 0 \\ 1, \text{ for } \Delta_t = 1. \end{cases}$$

The value function in each period, $V_t(A_t, Y_t, \Delta_t)$, has three state variables: the level of assets (A), the level of income (Y), and an indicator for whether the consumer is disabled (Δ). The value function is equal to the expected discounted utility of consumption in each period from the current period t to the final period T, discounted by a factor of β each period. The rate of time preference is equal to $1/\beta - 1$ and is similar to an interest rate in governing the utility tradeoff across periods. The within-period utility function is assumed to be additively separable in consumption and all other factors that affect utility, including health and leisure, so that these factors can be omitted from the optimization problem. A more detailed model would allow for health and leisure in a given period to affect the willingness of the consumer to spend versus save resources for a later period. As equation (2) makes explicit, our analysis is also predicated on the assumption that agents do not exhibit state-dependent utility. In other words, our calculations are based on the potentially strong supposition that disability shocks do not alter utility from a given level of consumption.[12]

The utility of consumption in each period is assumed to take the Constant Relative Risk Aversion (CRRA) form: $u(C) = C^{1-\gamma}/(1-\gamma)$, where γ is the coefficient of relative risk aversion. With a utility function such as CRRA that has a convex marginal utility function (i.e., $u'''(C) > 0$), there is a precautionary motive for saving, and greater uncertainty in the income process will induce greater saving.[13]

Assets in the next period are equal to the excess of current assets and income over consumption, augmented at the after-tax interest rate, as in equation (2b). For simplicity, the income tax system is assumed to be linear at a rate, τ, and the portfolio choice is subsumed in the form of a constant interest rate, r. Assets are constrained to be nonnegative in each period. The uncertainty in this model comes from two sources. The first is the risk of becoming disabled. The probability that a nondisabled person becomes disabled in period s is δ_s, thereby leading to the expression in equation (2d) for the probability that a person is disabled in that period. Disability is assumed to be an absorbing state. The second source of uncertainty is the variance in current income while working.

In the equations below, we state the processes that describe income uncertainty and the evolution of current income:

Before retirement or disability:

$$(3) \qquad (a)\ \ln(Y_s) = \ln(P_s) + u_s$$

$$(b)\ u_{s+1} = \rho u_s + \varepsilon_{s+1}$$

$$(c)\ \varepsilon_s \sim i.i.d.\ N(0, \sigma^2)$$

12. Empirical support for our characterization is provided in Gertler and Gruber (2002). In their analysis of Indonesian households on the effect of changes in the disability status of a head of household on changes in per-capita (nonmedical) consumption, they found no evidence for the most plausible manifestations of state-dependent utility. We return to this point in the concluding section of our chapter.

13. The use of the CRRA utility function is standard in both the empirical and theoretical literature on precautionary saving. Constant Relative Risk Aversion utility means that a consumer remains equally willing to engage in gambles over a constant proportion of current wealth as wealth increases. An alternative, and perhaps more realistic assumption, might be that the consumer will accept larger proportional risks as wealth increases. See Kimball (1990) for a discussion and derivation of the key results for precautionary saving.

At retirement or disability:

$$(d)\ Y_{s+1} = gY_s$$

After retirement or disability:

$$(e)\ Y_{s+1} = Y_s.$$

Prior to retirement or disability, the natural log of current income (Y_s) is equal to the natural log of permanent income (P_s) plus a shock to income (u_s) that follows an autoregressive (AR)(1) process. The innovations to that AR(1) process are assumed to be independently and identically draws from a normal distribution with mean zero and variance σ^2.[14] In the eventuality of either retirement or disablement, income is reduced to a replacement rate (g) of its most recent value. After retirement or disability, income is unchanged at this new level and is no longer uncertain. In the model below, we model life cycle labor supply as the consumer starting work at age twenty-two, retiring at age sixty-two, and living with certainty until age eighty-two (implying that voluntary retirement is taken in the fortieth year out of sixty in the assumed lifetime.)

In this framework, we have included a stylized version of the current Social Security Disability Insurance program. When a worker becomes disabled, he or she gets the same replacement rate that he or she would get at retirement, though calculated on income through the year of disability. In both cases, the real value of benefits stays constant over time. More detailed formulations of both the retirement benefit and the disability benefit are possible, though not without substantially increasing the complexity of the model with an additional state variable (e.g., the average index monthly earnings of the worker to date).

The solution method for stochastic optimization problems with multiple state and control variables is discussed in detail in Carroll (2001). The solution begins in the last period of life, T, when the problem is trivial because the household simply consumes all of its assets and after-tax income, yielding optimal values for the control variables, C_T^d and C_T^n, as a function of the state variables A_T and Y_T. These solutions generate the value function, $V_T(A_T, Y_T, \Delta_T)$, and the partial derivative, $V_T^A(A_T, Y_T, \Delta_T)$, which represents the marginal value of an additional dollar in assets at the beginning of period T.[15] Moving back to the period $T-1$ problem, we can rewrite the objective function as:

$$(4)\quad V_{T-1}(A_{T-1}, Y_{T-1}, \Delta_{T-1}) \equiv \max_{\{C_{T-1}\}} u(C_{T-1}) + \beta E_{T-1}[V_T(A_T, Y_T, \Delta_T)].$$

14. In our simulations, we normalize the mean of the lognormal shock to be one in all periods.

15. The partial derivative, $V_T^Y(A_T, Y_T, \Delta_T)$, is not needed, since the value of Y_T is not influenced by the choice variable in period $T-1$.

The problem in period $T - 1$ is a special case, since there is no income uncertainty or further risk of disability. More generally, given the function $V_{t+1}(A_{t+1}, Y_{t+1}, \Delta_{t+1})$ and the associated partial derivatives, the problem at period t is:

$$(5a) \qquad V_t(A_t, Y_t, 1) \equiv \max_{\{C_t\}} u(C_t) + \beta E_t[V_{t+1}(A_{t+1}, Y_{t+1}, 1)]$$

for disabled consumers or:

$$(5b) \quad V_t(A_t, Y_t, 0) \equiv \max_{\{C_t\}} u(C_t) + \beta E_t[\delta_{t+1} V_{t+1}(A_{t+1}, Y_{t+1}, 1) + (1 - \delta_{t+1}) V_{t+1}(A_{t+1}, Y_{t+1}, 0)]$$

for nondisabled consumers. These one-period problems have first-order conditions given by:

$$(6a) \qquad u'(C_t) - \beta[1 + r(1 - \tau)]\, E_t[V^A_{t+1}(A_{t+1}, Y_{t+1}, 1)] = 0$$

and

$$(6b) \quad u'(C_t) - \beta[1 + r(1 - \tau)] \cdot E_t[\delta_{t+1} V^A_{t+1}(A_{t+1}, Y_{t+1}, 1) + (1 - \delta_{t+1}) V^A_{t+1}(A_{t+1}, Y_{t+1}, 0)] = 0.$$

The first term in each first-order condition is the marginal utility of an additional dollar of consumption in period t. The second term is the expected discounted value of saving that dollar to be used in period $t + 1$. The dollar grows by the after-tax interest rate and has a marginal value of V^A at that time, where in the case of a nondisabled consumer V^A is evaluated at both possibilities for period $t + 1$—disabled or nondisabled—and weighted appropriately by the probability of disability or its complement. The expected marginal utility of a dollar of assets at time $t + 1$ is discounted back to period t utility at a rate β. The difference between the marginal utility of consumption and the expected marginal utility of assets in the next period is zero at the optimal level of consumption.[16]

Once the optimal consumption rules have been obtained, the models can be simulated forward by specifying initial values of the state variables, drawing random shocks to income and disability status, and applying the consumption rules to generate distributions of asset balances in each period. In our simulations, we construct average consumption and asset profiles based on 5,000 independent random draws.

The baseline model consists of assumptions about the income process

16. The solution method is complicated by the liquidity constraint. The constraint that A_{t+1} cannot be negative implies that the maximum amount of consumption in the prior period is $C_t = A_t + Y_t(1 - \tau)$.

and the preference parameters. We assume a starting value of income at age twenty-two of \$20,000 and allow permanent income to grow similarly to the profile for noncollege graduates specified in Hubbard Skinner, and Zeldes (1995).[17] We also adopt their parameters of $\rho = 0.95$ and $\sigma = 0.15$ for the AR(1) income process. We specify a replacement rate of 40 percent at retirement or disability, corresponding to a typical replacement rate from the Social Security system. The constant, pretax interest rate, r, is assumed to be 5 percent. The tax rate in the linear tax system, τ, is taken to be 20 percent, applied to both labor and investment income. The after-tax interest rate is therefore 4 percent per year. In alternative models, we consider replacement rates of 20 and 60 percent (with an income standard deviation of 15 percent) and income standard deviations of 10 and 20 percent (with a replacement rate of 40 percent).

There are two main preference parameters. The coefficient of relative risk aversion, γ, is assumed to be 3. In a CRRA model, this results in an intertemporal elasticity of substitution of 1/3. The discount rate, β, takes one of two values: For simulations of a patient consumer, β is assumed to be 1/1.04 percent. For simulations of an impatient consumer, β is assumed to be 1/1.08 percent. In the absence of income uncertainty, an impatient consumer would seek to borrow against future income to finance current consumption, whereas a patient consumer would not. A patient consumer begins saving for retirement early in the life cycle, while an impatient consumer typically engages in buffer stock saving for several periods before saving for retirement. The difference in behavior results from the comparison of the rate of time preference to the interest rate—the patient consumer has $1/\beta = 1 + (1 - \tau)r$, whereas the impatient consumer has $1/\beta > 1 + (1 - \tau)r$. The two values chosen are close to the median estimates of the discount rate in Samwick (1998).

The key parameter that we vary in our simulations is the age-disability profile. For all simulations, we adopt the empirical age profile from figure 10.1 that reflects the regression-adjusted age profile in our CPS data over the period from 1980 to 2004. We smooth out the initial disability prevalence at age twenty-two linearly over all ages, so that the profile starts out at a zero probability of disablement and then rises to the 12 percent prevalence at age sixty-two that we observe in the data. To consider variation around this baseline, we scale the entire profile up or down proportionally to achieve alternative disability rates at age sixty-two, including the values of 0, 6, 12, 18, and 24 percent, with the latter corresponding roughly to the maximum preretirement disability rate that we observe for any of our groups empirically.

17. We approximate this profile by having real income grow at annual rates of 2.5, 1.7, 0.5, and –1.3 percent over the four decades of the working life. Total income growth by the peak (in the thirtieth year of the working life) is about 55 percent, with a subsequent drop of about 12 percent. The choice of the starting income level is immaterial here, as the entire optimization problem scales linearly with income.

Figure 10.6 shows the impact of disability risk on consumption. Each panel graphs the average level of consumption by age for consumers who do not become disabled under each of four different rates of preretirement disability. The top panel is for the impatient consumer, and the bottom panel is for the patient consumer. In each panel, the solid curve indicates the profile that average consumption would take in the absence of a disability risk. Even in this baseline case, it is upward sloping during most of the working life due to the need for precautionary saving against annual income fluctuations in the early years. After the peak, it declines toward the

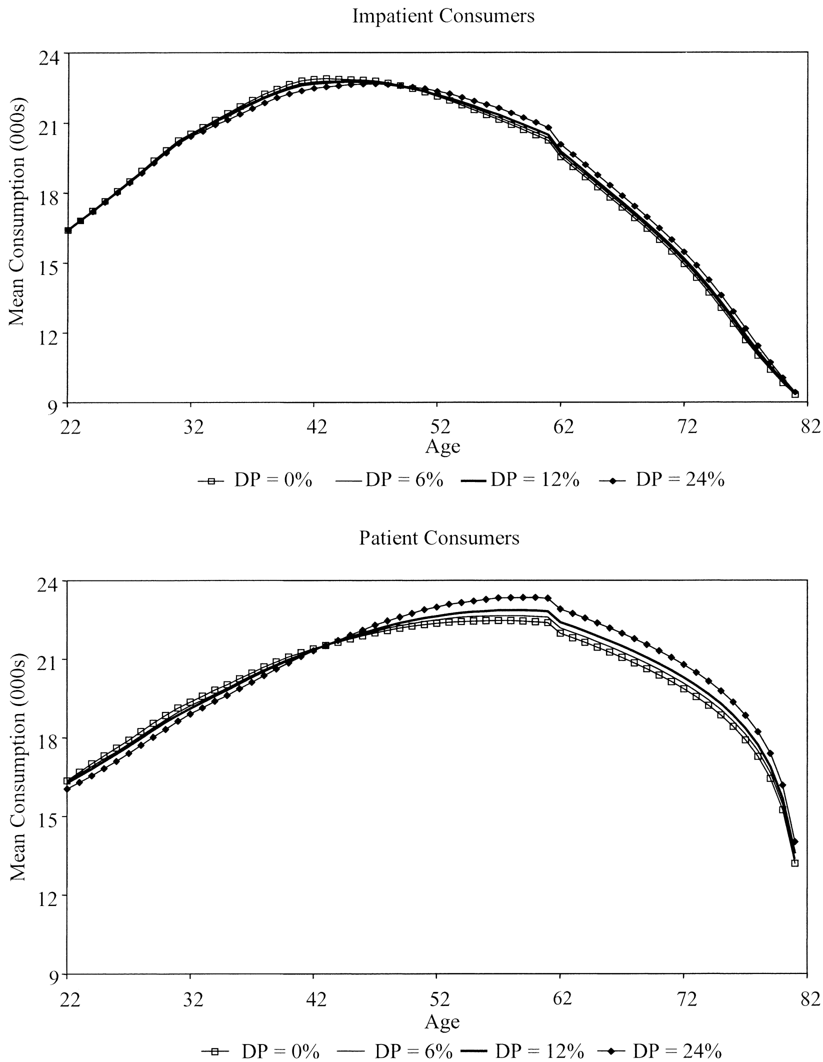

Fig. 10.6 Age-consumption profiles for consumers who do not become disabled

end of the life cycle. There is no important change at retirement in this model, since the voluntary retirement is completely anticipated. The baseline profile is flatter and peaks later for the patient consumer than for the impatient consumer.

As the disability risk increases, the consumption profile starts lower early in life (to allow for more precautionary saving) and remains higher later in life. All of the consumption profiles (in both panels) have the same present value, because the income draws are the same and these are the consumers who do not become disabled before retirement. The distortions in the timing of consumption over the life cycle—induced by the risk of becoming disabled—contribute to the welfare loss discussed in the following paragraphs. The distortions are larger for the patient consumer than for the impatient consumer, because the former reacts more strongly and sooner to the higher risks of preretirement disability.

Figure 10.7 illustrates the effect of a disabling event on the accumulation of assets. Each panel graphs the average asset values by age for preretirement disability risks of 12 and 24 percent. For consumers who do not become disabled before retirement, the asset profiles rise steadily to a peak just prior to retirement, after which they are spent down gradually to zero. Consistent with the consumption profiles, the patient consumer consumes less and accumulates more assets prior to retirement. The graphs also show the asset profiles of consumers who become disabled at age fifty-two, after thirty years of work but ten years prior to voluntary retirement. These profiles track those of those who do not get disabled through that age[18] and are fairly quickly spent down to zero after disablement. It takes about fifteen years for the impatient consumer to exhaust the assets, and about thirty years for the patient consumer to do so. After the assets are exhausted, the consumer simply consumes the income provided by the disability insurance program.

Figure 10.8 shows the range of age-asset profiles corresponding to preretirement disability rates up to 24 percent. In each panel, the solid (bottom) profile corresponds to the baseline case with no disability risk. Preretirement wealth peaks at $87,000 for the impatient consumer and $147,700 for the patient consumer, compared to starting income of $20,000 and peak average income of about $32,000 (at age fifty-two). Successively higher profiles show the effect of increasing the disability risk. With the sample average risk of 12 percent, the peak preretirement wealth rises to $90,100 (3.6 percent) and $154,000 (4.3 percent) for the impatient and patient consumers, respectively.[19]

18. The graphs should overlap exactly prior to disablement. The disparity is due to the small sample variation in income draws for the subset of random draws that first become disabled at exactly age fifty-two.

19. The numbers reported for assets are for the mean profile. Median asset amounts are 15 to 20 percent lower on the eve of voluntary retirement. The increase from zero to 12 percent disability risk raises median assets from $70,200 to $73,900 for the impatient consumer and from $126,300 to $133,300 for the patient consumer.

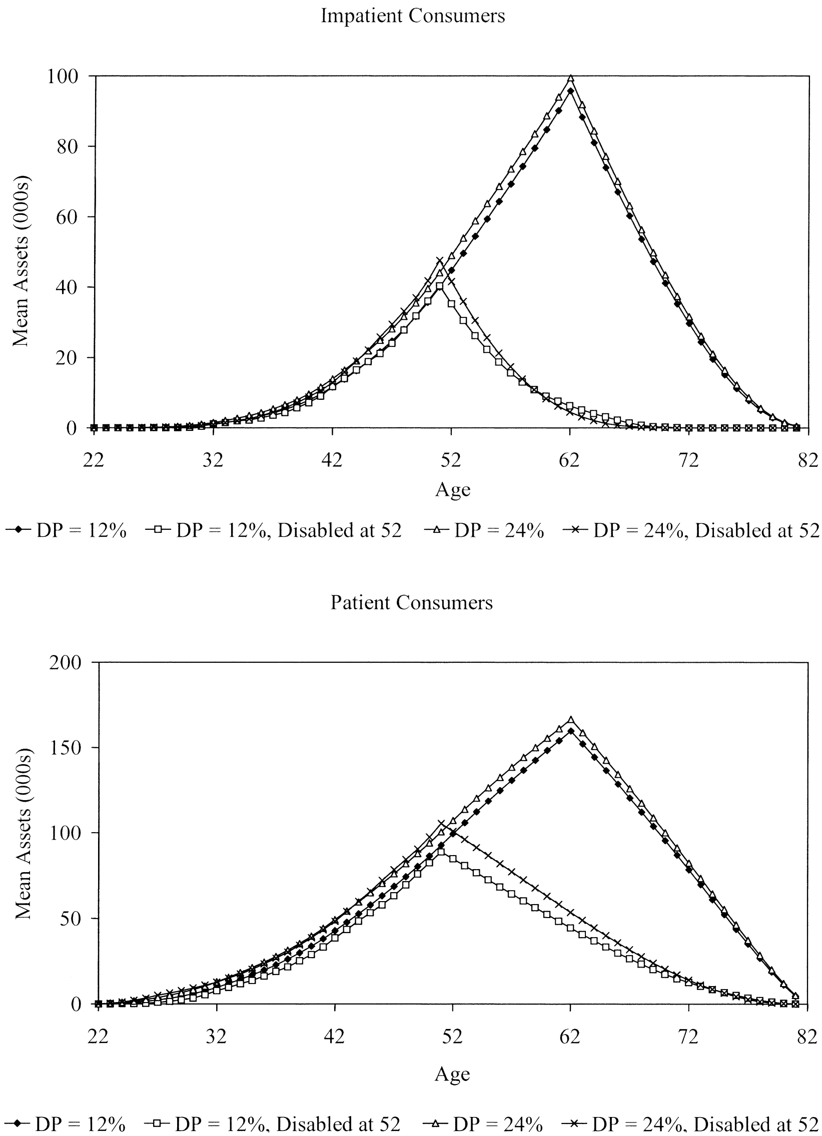

Fig. 10.7 Age-asset profiles by age of disability

10.5 Implications for Expected Utility

Table 10.6 shows the impact of disability risk on the expected lifetime utility and preretirement asset accumulation of the consumer across a range of values for the replacement rate and income uncertainty. The first panel shows the results for our baseline assumptions of a 40 percent replacement rate at retirement or disability and a 15 percent standard devia-

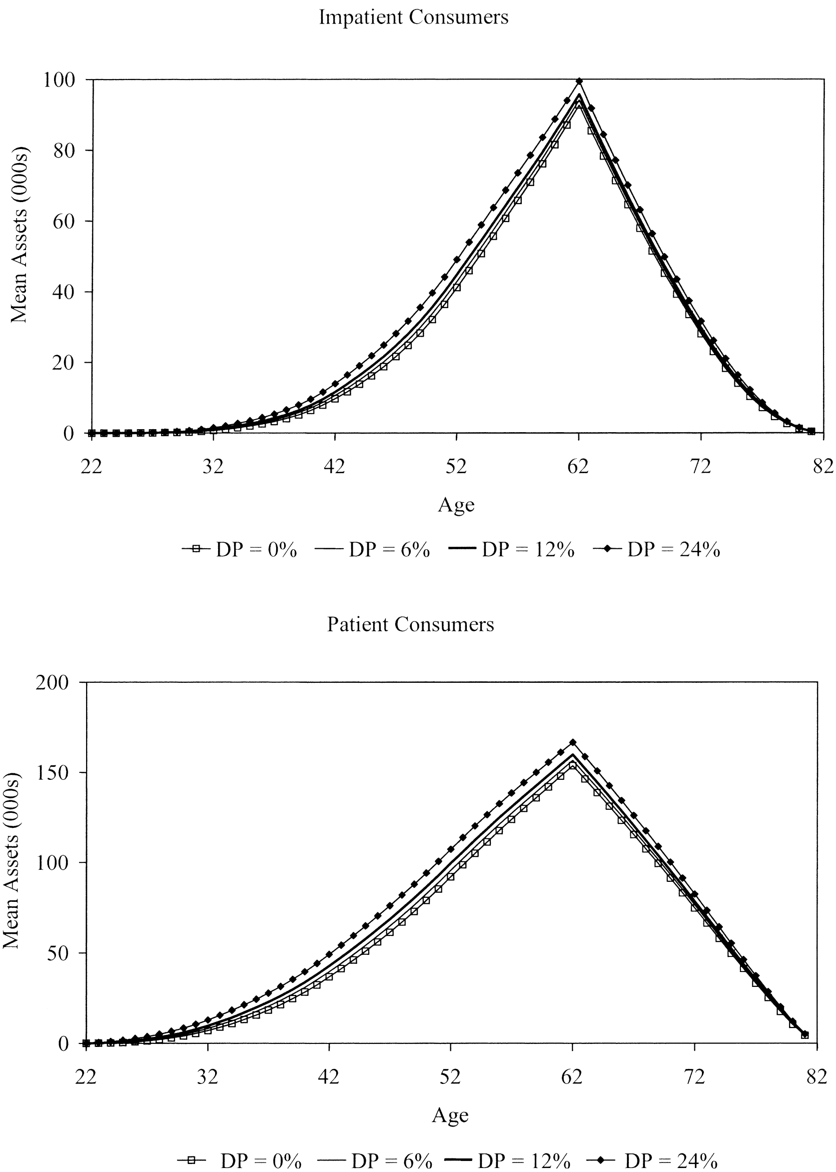

Fig. 10.8 Age-asset profiles for consumers who do not become disabled

tion of the shocks to annual income. The next two panels vary the replacement rate to 20 and 60 percent of preretirement income, and the last two panels vary the standard deviation of the annual income shock to 20 and 10 percent.

Within each panel, the first column of numbers identifies the probability

Table 10.6 **Impact of disability risk on expected utility and saving**

Preretirement disability rate (%)	Loss in present value of consumption (%)	Equivalent variation (%)		Asset reduction in baseline (%)	
		Impatient	Patient	Impatient	Patient
	Replacement rate = 40%, Standard deviation = 15%				
6	1.00	2.34	2.70	1.69	2.01
12	1.92	4.36	5.19	3.47	4.09
18	2.92	6.41	7.58	5.71	6.45
24	3.94	8.05	9.34	7.40	8.27
	Replacement rate = 20%, Standard deviation = 15%				
6	1.32	8.41	8.45	5.59	6.33
12	2.60	14.44	13.81	10.66	11.37
18	3.93	19.41	18.27	15.41	15.43
24	5.27	22.29	21.18	19.21	18.40
	Replacement rate = 60%, Standard deviation = 15%				
6	0.69	0.76	1.07	0.19	0.47
12	1.27	1.50	2.10	0.41	1.06
18	1.96	2.41	3.10	1.35	2.10
24	2.68	3.12	3.98	1.32	2.45
	Replacement rate = 40%, Standard deviation = 20%				
6	1.00	2.02	2.40	1.12	1.19
12	1.89	3.92	4.71	2.33	2.48
18	2.90	5.84	6.86	4.06	4.09
24	3.94	7.20	8.50	4.98	4.95
	Replacement rate = 40%, Standard deviation = 10%				
6	0.99	2.22	2.95	2.23	2.65
12	1.93	4.29	5.54	4.57	5.38
18	2.93	6.43	8.04	7.29	8.33
24	3.93	8.26	10.16	9.78	10.94

Notes: Preretirement disability rate is the prevalence of disability in the last period of the working life (here modeled as forty working periods and twenty retirement periods). Loss in PV of consumption is the percent reduction in the present value of expected consumption due to the specified disability risk. Equivalent variation is the percent reduction in consumption relative to the zero-disability case that would reduce expected utility by the same amount as facing the specified disability risk. Asset reduction in the baseline is the percent reduction in asset accumulation in the last period of the working life due to the absence of the specified disability risk. Patient and impatient consumers have rates of time preference equal to 4 and 8 percent, respectively. See the text for a description of the other parameter assumptions.

of becoming disabled prior to the age of anticipated retirement. The numbers range from 6 to 24 percent, with a value of 12 percent corresponding most closely to the average disability prevalence measured in the CPS for those aged sixty to sixty-four. There are several demographic groups with disability prevalence lower than that average, and the highest value of 24 approximates the highest prevalence estimated in the data (fig. 10.A2 demonstrates that is the prevalence of disability for never married men at age sixty-two). The 6 percent number represents the prevalence of a work-limiting disability at age sixty-two for college-educated workers in 2004 (fig. 10.4). The estimate of 18 percent is observed for black men at age sixty-two (fig. 10.5). The impact of shifting disability prevalence across time or group—conditional on our choice of baseline parameters—is therefore captured by the simulations shown in the table.

With a CRRA utility function, multiplying consumption in each period (and disability state) by a constant, k, multiplies expected utility by a factor of $k^{1-\gamma}$. We can therefore compare consumer welfare across two parameterizations by solving for the value of k such that if consumption in the baseline case is multiplied by k, the expected utility would equal the expected utility of the optimal consumption profile under the alternative set of parameters. If k is less (greater) than one, then the consumer is worse (better) off under the alternative parameterization. Since we consider the impact of increases in the disability rate, which necessarily make the consumer worse off, we refer to $l - k$ as an equivalent variation, because it measures the amount of money (as a share of lifetime consumption in the baseline case) that a consumer would pay to avoid facing the higher disability risk in the alternative parameterization. The middle two columns show the equivalent variations separately for impatient and patient consumers.

Because consumers have a precautionary motive for saving, increasing disability risk will generate an increase in asset accumulation. As a means of calibrating the equivalent variation, the last two columns of the table show the percentage by which asset accumulation is lower in the baseline case with no disability risk compared to the asset accumulation in the alternative case (for those who do not become disabled) with the specified disability risk. Alternatively, this is the percentage of wealth in the alternative case that is attributable to the nonzero risk of disability.

Increasing the preretirement disability risk from zero to the sample average value of 12 percent generates equivalent variations of 4.36 and 5.19 percent for the impatient and patient consumers, respectively. Even though the present value of consumption falls by only 1.92 percent with this disability risk, the consumers would forego about 5 percent of their lifetime consumption to avoid that risk. Recall that disability risk is already partially insured through a stylized DI program in this model—income does not go to zero upon disablement. The 5 percent equivalent variation captures the amount that consumers would be willing to pay to remove the risk that they will ever have to receive payments from that program. The reduction in assets accumulated upon retirement (for those who do not become disabled) is about 1.85 percent in the baseline compared to the alternative.

For every 6-percentage point increase in the risk of disability, the present value of preretirement consumption falls by about 1 percent, the equivalent variation increases by about 2 percentage points, and the gap in preretirement asset accumulation for those who do not become disabled increases by about 2 percentage points. Considering the alternative definitions of disability in figure 10.1, the equivalent variation in the strictest definition, corresponding to the additional requirement of being on Medicare and a disability risk at age sixty-two of 6 percent, is as low as 2.5 percent, with 1 percentage point representing lost consumption. For the least restrictive definition, based solely on the answer to the work-limiting disability question and having a disability risk at age sixty-two of 18 percent, the equivalent variation is about 7 percent, with 3 percentage points representing lost consumption. Alternatively, the different simulations can correspond to the extremes of our data measured using our baseline definition of disability. College-educated workers with only a 6 percent disability risk have equivalent variations of 2.5 percent, while men who never marry have equivalent variations in excess of about 8.5 percent.

The next two panels show the analogous calculations when the replacement rate at retirement or disability is changed by 20 percentage points, to 20 and 60 percent, respectively. Focusing on a preretirement disability risk of 12 percent, cutting the replacement rate in half increases the amount of lost consumption by about 35 percent (from 1.92 to 2.60 percentage points), but the equivalent variation more than doubles to about 14 percent of baseline consumption. The asset reduction in the baseline case also increases to about 11 percent. With the higher replacement rate of 60 percent, the equivalent variation falls to about 1.4 percent and the asset reduction similarly falls to about 0.7 percent. In all cases, the asset reduction in the baseline due to the absence of disability risk is a bit below the equivalent variation.

The bottom two panels show the calculations when the standard deviation of the annual shock to preretirement income is changed by 5 percentage points, to 20 and 10 percent, respectively. The present value of the consumption losses are the same when uncertainty increases (apart from sampling variance), as the shocks are constrained to have a mean of 1. The equivalent variations change very little as income uncertainty changes—the 12 percent disability risk generates an equivalent variation of about 4 to 5 percent, again with about 1 percentage point of difference between impatient and patient consumers. The asset reductions in the baseline case without disability risk are lowest when there is high income uncertainty—the added precautionary saving against annual income fluctuations diminishes the relative importance of precautionary saving against the disability risk.

Table 10.7 shows the impact on expected utility and preretirement asset accumulation of lowering the replacement rate (from 60 to 40 to 20 percent) and increasing the standard deviation of the annual income shock (from 10 to 15 to 20 percent), while holding the preretirement disability risk fixed at 0 or 12 percent. In both sets of comparisons, the changes should serve to decrease expected utility and increase asset accumulation. The results in this table can be used to gauge the magnitudes in table 10.6.

Table 10.7 **Impact of replacement rate and income uncertainty on expected utility and saving**

Replacement rate or Standard deviation (%)	Loss in present value of consumption (%)	Equivalent variation (%) Impatient	Equivalent variation (%) Patient	Asset reduction in baseline (%) Impatient	Asset reduction in baseline (%) Patient
	Disability risk = 0%, Standard deviation = 15%				
RR = 40	2.75	1.50	2.00	44.69	28.42
RR = 20	5.50	2.94	3.69	62.77	44.95
	Disability risk = 12%, Standard deviation = 15%				
RR = 40	3.38	4.36	5.09	46.39	30.62
RR = 20	6.77	15.69	15.21	66.60	50.69
	Disability risk = 0%, Replacement rate = 40%				
SD = 15	0.00	9.81	9.98	30.39	26.16
SD = 20	0.00	22.18	20.92	51.68	44.89
	Disability risk = 12%, Replacement rate = 40%				
SD = 15	–0.02	9.87	9.64	29.59	25.16
SD = 20	–0.04	21.88	20.22	50.54	43.20

Notes: Disability Risk is the prevalence of disability in the last period of the working life (here modeled as 40 working periods and 20 retirement periods). Comparisons in each panel are made relative to the highest replacement rate (60 percent) and the lowest income uncertainty (10 percent). Loss in PV of consumption is the percent reduction in the present value of expected consumption due to the specified change in replacement rate or income uncertainty. Equivalent variation is the percent reduction in consumption relative to the lowest-saving case that would reduce expected utility by the same amount as facing the lower replacement rate or higher income uncertainty. Asset reduction in the baseline is the percent reduction in asset accumulation in the last period of the working life due to the absence of the specified reduction in the replacement rate or increase in income uncertainty. Patient and impatient consumers have rates of time preference equal to 4 and 8 percent, respectively. See the text for a description of the other parameter assumptions.

With no risk of disability, reductions in the replacement rate affect only the level of income received after retirement. Each reduction of 20 percentage points in the replacement rate lowers the present value of consumption by 2.75 percent. However, the equivalent variations show that neither the impatient nor the patient consumer would be willing to forego that amount of consumption—deducted proportionally in each year—to avoid these reductions. The reason is that the consumers can optimally (rather than merely proportionally) adjust their consumption to offset the lower retirement replacement rate by saving more. The last two columns show that there is a substantial savings response to the reductions in the replacement rate, of 45 and 63 percent for the two increments by the impatient consumer and 28 and 45 percent by the patient consumer. (The proportion is lower for the patient consumer, who is doing more life cycle saving in the baseline when the replacement rate is 60 percent.)

The loss in the present value of consumption is 3.38 percent for each 20 percentage point reduction in the replacement rate when the disability risk is increased to the sample average of 12 percent. The equivalent variations increase substantially to 4.4 to 5.1 percent for the drop to a 40 percent replacement rate, and 15.2 to 15.7 percent for the drop to a 20 percent replacement rate. Compared to the zero disability risk case, however, the asset reductions in the baseline increase only 2 to 4 percentage points for the impatient consumer (to 46.4 and 66.6 percent) and 2 to 6 percentage points for the patient consumer (to 30.6 and 50.7 percent).

For changes in the standard deviation of the income shock, the results are comparable for disability risks of zero and 12 percent. In both cases, and for both the impatient and the patient consumer, raising the standard deviation from 10 to 15 percent generates an equivalent variation of 9.6 to 10 percent and asset reductions in the baseline of 25 to 30 percent. An increase in the standard deviation from 10 to 20 percent generates an equivalent variation of about 20 to 22 percent and asset reductions in the baseline of 43 to 52 percent.

The distinction between the impacts of disability risk, income uncertainty, and the replacement rate on saving and expected utility is summarized in figure 10.9. The bottom profile remains the average assets by age for consumers who face a retirement replacement rate of 40 percent, a standard deviation of income shocks of 15 percent, and zero risk of disability. The next lowest profile reflects the impact (on those who do not become disabled) of facing a preretirement disability risk of 36 percent, or three times the sample average risk. In calculations analogous to those shown in table 10.6, the consumer would be willing to forego about 12 percent of consumption in each period to avoid that risk, and about 12 percent of the wealth accumulated at the preretirement peak is attributable to that disability risk.

The top profile in the graphs shows the impact of keeping the disability risk at zero but increasing the standard deviation of the annual income shocks from 15 to 20 percent. The equivalent variation for this shift is about 13 percent, just slightly higher than for the increase in the disability risk to 36 percent. However, about 28 percent of the wealth accumulated at the preretirement peak is attributable to the higher income uncertainty, which is more than twice as much as for the profile that increased the disability risk to 36 percent.

The last profile in the graph, with the shortest dashes, shows the impact of keeping the disability risk at zero and the standard deviation of the income shock at 15 percent, but lowering the replacement rate from 40 to 20 percent. Note that because the disability risk is zero, this reduction affects only the income received after retirement. The peak asset accumulation for this profile is comparable to the peak when income uncertainty is raised to 20 percent. However, the equivalent variation relative to the baseline case is less than 2 percent of lifetime consumption.

What explains these differences? Increases in the disability risk have a large effect on expected utility but a comparatively small effect on asset ac-

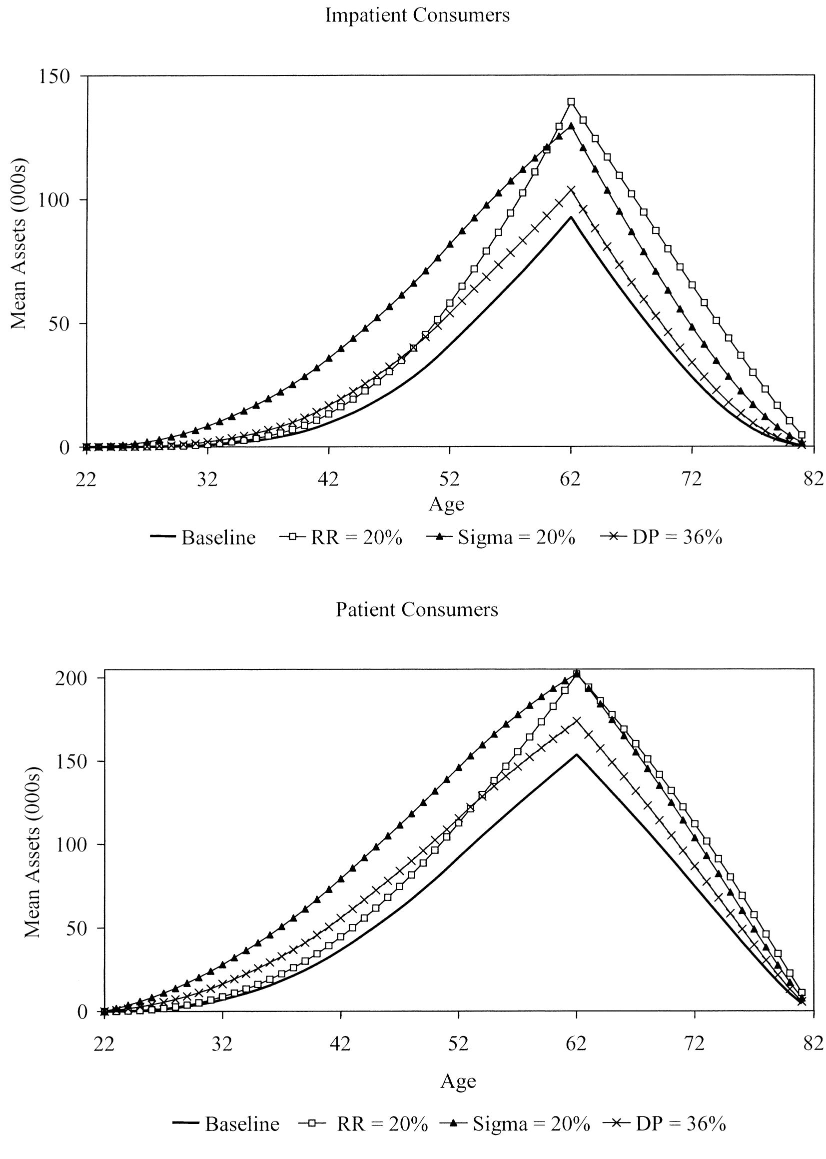

Fig. 10.9 Age-asset profiles for alternative parameterizations for consumers who do not become disabled

cumulation. The reason is that saving is less useful in protecting against a low-probability, high-impact risk than the events in the two alternatives. Saving is a perfect response against the certain drop in income at retirement, and because such an ideal response exists, the equivalent variation of reductions to the retirement replacement rate (when there is no disability risk) is quite small. Additional saving is an effective though imperfect hedge against year-to-year income fluctuations. For a change in parameters of a given equivalent variation, the consumer does less saving against annual income risk than against a planned drop in income, but more saving than against a risk of disability. Preretirement saving is least attractive as a hedge against a disabling risk.

The age-asset profiles also show that when disability or income risk increases, the impact on saving is immediate. The consequences of disability early in life are critical, but the impact on assets fades over time as the number of years over which the disability could occur is reduced and the income level that would be replaced increases. In contrast, when the replacement rate at planned retirement is lowered, the consumer has forty years to overcome the loss. With an upward sloping income profile over the first thirty years of the working life, the lower replacement rate can be accommodated by small reductions in consumption at each age, which accelerate as the date of retirement approaches and income growth slows.

10.6 Conclusion

Using the Current Population Survey over the period from 1980 to 2004, we document the decline in disability over the past two decades, along with cross-sectional differences in the prevalence of disability by gender, education, and other demographic groups. Once we account for compositional changes in demographic characteristics, as well as trends in disability that affect all groups equally, we find that the age-disability profile is fairly stable across time and demographic groups. The exception to this finding is a slight decline in disability prevalence for Americans aged sixty to sixty-four. These findings suggest that demographic changes and year-specific changes in the disability rate (that affect all groups equally) generate the observed changes in the level of disability over time. In our stochastic life cycle model of consumption, we estimate that a typical consumer would be willing to pay a premium equal to about 5 percent of lifetime consumption to avoid the average risk of disability found in our data, even in the presence of a stylized disability insurance program that provides the same replacement rate upon disability as at retirement.

As discussed earlier, our estimates are derived under the assumption that agents do not exhibit state-dependent utility. Under state-dependent utility, the utility of consumption is no longer independent of health. If invalid, this assumption would cause us to either overstate or understate an agent's willingness to pay for insurance against disability. If we assume the utility of consumption if disabled is lower than if nondisabled, which is equivalent to assuming that health and income are complements in consumption, individuals would be willing to pay less for insurance against disability than the situation that we have considered. Alternatively, it could be that the marginal utility of income is higher when disabled since the disabled need more resources to purchase healthcare and assistive technologies. If so, individuals would be willing to pay more for disability insurance. In the absence of evidence that supports state-dependent utility, we have delegated further exploration of this potentially important issue to future work.

Compared to other motives for saving, like an anticipated drop in income at retirement or annual fluctuations in income, disability risk generates little additional saving for a given welfare loss. This is because precautionary saving is less useful as a hedge against low-probability, high-impact events like disability. As a result, it is unlikely that the precautionary saving that occurs specifically due to the empirically observed probabilities of disability is large enough to be of macro- or microeconomic importance. We estimate that no more than 20 percent of assets accumulated before voluntary retirement are attributable to disability risks observed in our data. Because the probability of disablement is too small, and the average size of the loss (conditional on becoming disabled) is large, disability risk is not effectively insured through precautionary saving. Therefore, the value of disability insurance, whether in the form of income or assistive technology, is likely to be high. The magnitude of our finding is consistent with the related work of Fuchs-Schundeln (2003), who estimates that precautionary savings constitute between 5 and 12 of aggregate savings in a calibrated model that measures the welfare effects of uninsured labor risks on consumption.

Our chapter is silent on the specifics of the optimal disability-insurance system. While the value of such an insurance program may be high, and the typical worker is willing to pay for the program, it is important to ensure that its design is actuarially fair by income group. As noted by Bound et al. (2004), the current DI program permits undesirable transfers from low-income able-bodied workers and non-DI-eligible disabled persons to comparatively better off DI beneficiaries. These transfers render the program less fair to low-income persons, a group who would have been predicted to be the greatest beneficiaries of this insurance program.

Throughout our analysis, we have defined disability as a health problem or condition that limits the kind or amount of work that a person can do. This definition is focused on the link between disability and work, but work-limiting disabilities as defined in the CPS may include more conditions than those that lead to involuntary retirement. To mitigate this possibility, our main estimates in the paper follow the methodology of Burkhauser, Houtenville, and Wittenburg (2003) and require that the respondent not be working at present nor have worked in the prior year. Compared to our baseline estimate of a 5 percent equivalent variation, our lower and upper bounds are 2.5 and 7 percent, respectively. Once again, these estimates are in the range obtained by Fuchs-Schundeln (2003), who finds an equivalent variation of 2 to 5 percent of consumption as the welfare consequence of incomplete insurance against idiosyncratic labor shocks.

Relative to the definition of disability chosen in our analysis, other definitions based on health status, such as the inability to perform certain Activities of Daily Living, would suggest an alternative modeling framework in which the impact of disability would lower the utility from a given amount of consumption in addition to its impact on the ability to work. Developing such a model, in which the savings response could differ markedly from that presented here, is an important topic for further research.

Appendix

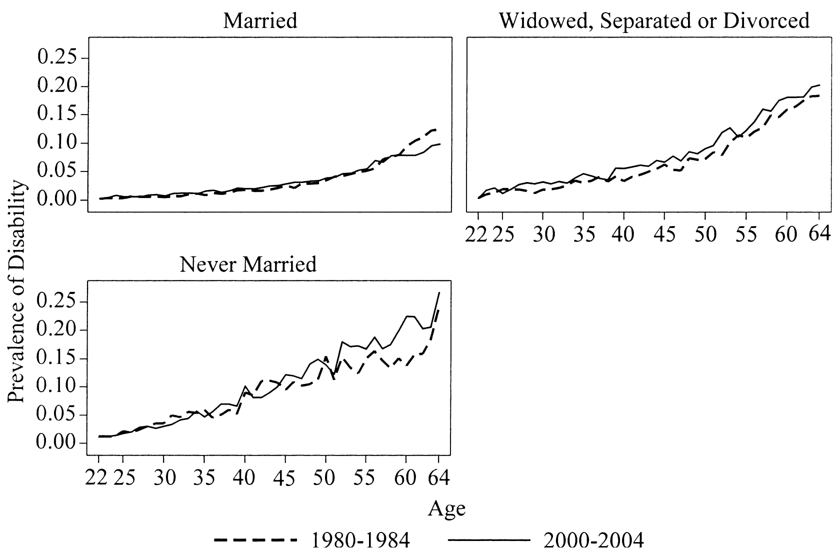

Fig. 10.A1 Regression adjusted disability prevalence by marital status over time

Notes: Figures report predicted probability of work-limiting disability from a logistic regression that controls for an unrestricted set of age and year indicator variables, gender, race and ethnicity (four categories), education (two categories), marital status (three categories), marital status and gender interactions, census region (nine categories), metropolitan status, census division and metropolitan status interactions, the number of children under the age of eighteen in the family, and the size of the household. The above figure adjusts for these covariates, except the ones used in the figure. The regression used CPS data from the 1980–2004 files of Annual Demographic Survey (n = 2,166,178).

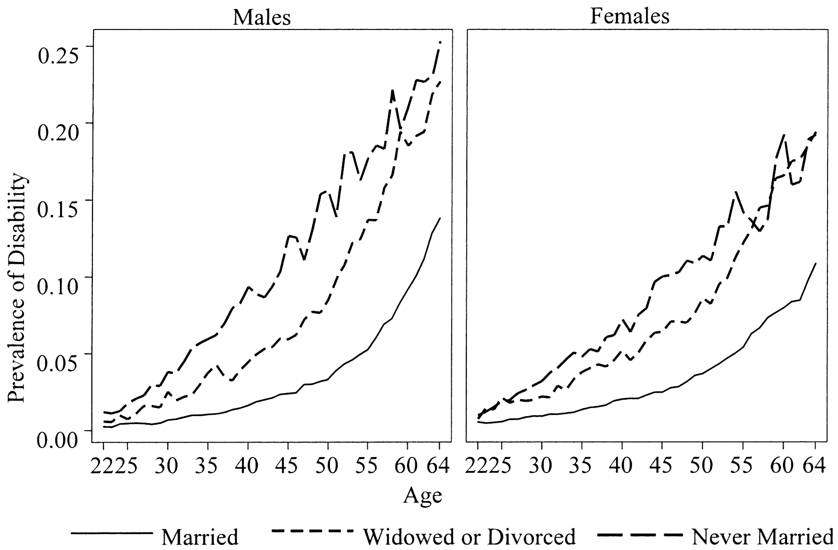

Fig. 10.A2 Regression adjusted disability prevalence by marital status and gender
Notes: See note to figure 10.A1.

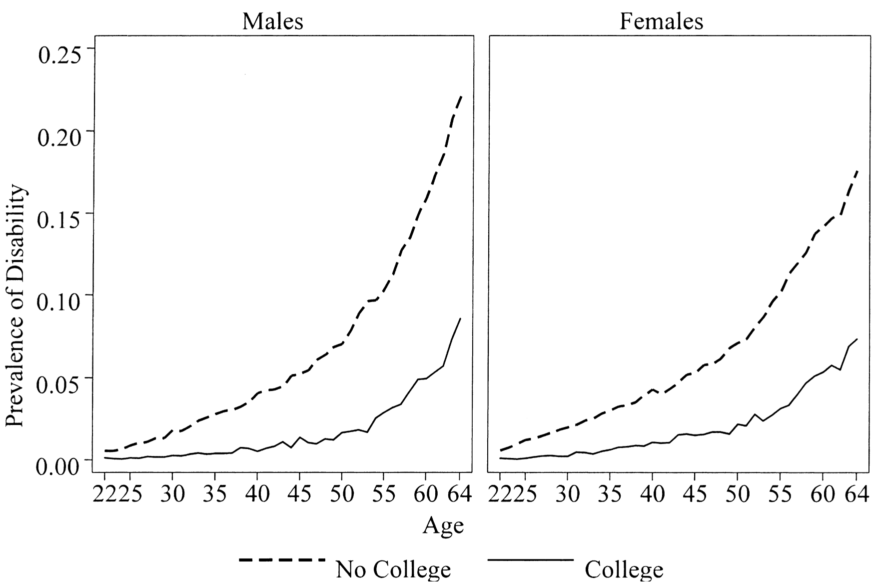

Fig. 10.A3 Regression adjusted disability prevalence by education and gender
Notes: See note to figure 10.A1.

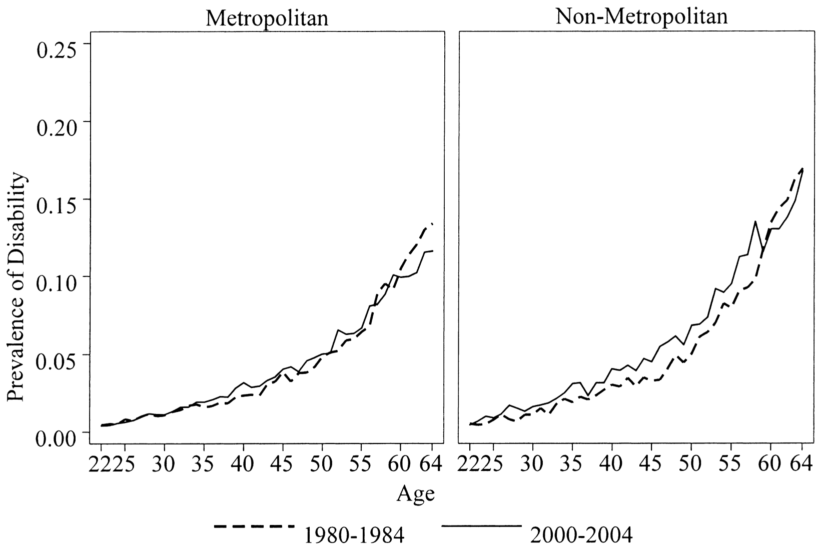

Fig. 10.A4 Regression adjusted disability prevalence by metropolitan status
Notes: See note to figure 10.A1.

References

Autor, David H., and Mark G. Duggan. 2003. The rise in disability rolls and the decline in unemployment. *Quarterly Journal of Economics* 118 (1): 157–206.

Bound, John, and Richard V. Burkhauser. 1999. Economic analysis of transfer programs targeted on people with disabilities. In *Handbook of labor economics, Vol. 3*, ed. Orley Ashenfelter and David Card, 3417–3528. Amsterdam: Elsevier.

Bound, John, Julie Berry Cullen, Austin Nichols, and Lucie Schmidt. 2004. The welfare implications of increasing disability insurance benefit generosity. *Journal of Public Economics* 88 (12): 2487–2514.

Bureau of Labor Statistics. 2005. The Employment Situation: January 2005. Accessed February 4, http://www.bls.gov/ces/

Burkhauser, Richard V., and Mary C. Daly. 1996. Employment and economic well-being following the onset of a disability: The role for public policy. In *Disability, work, and cash benefits*, ed. Jerry Manshaw, Virginia Reno, Richard Burkhauser, and Monroe Berkovitz, 59–102. Kalamazoo, MI: W. E. Upjohn Institute for Employment Research.

Burkhauser, Richard V., Andrew J. Houtenville, and David C. Wittenburg. 2003. A users guide to current statistics on the employment of people with disabilities. In *The decline in employment of people with disabilities*, ed. David C. Stapleton and Richard V. Burkhauser, 23–85. Kalamazoo, MI: W. E. Upjohn Institute for Employment Research.

Carroll, Christopher D. 2001. Lecture notes on solution methods for microeconomic dynamic stochastic optimization problems. Johns Hopkins University, Department of Economics. Unpublished Manuscript, April.

Deaton, Angus, and Christina H. Paxson. 1994. Intertemporal choice and inequality. *Journal of Political Economy* 102 (3): 437–67.
Fuchs-Schundeln, Nicola. 2003. Welfare loss and precautionary savings due to uninsured idiosyncratic labor risk. Harvard University, Department of Economics. Unpublished manuscript.
Gertler, Paul, and Jonathan Gruber. 2002. Insuring consumption against illness. *American Economic Review* 91 (1): 51–70.
Heckman, James J., and Richard Robb. 1985. Alternative methods for evaluating the impact of interventions. In *Longitudinal analysis of labor market data*, eds. James Heckman and Burton Singer, 156–245. New York: Cambridge University Press.
Hubbard, Glenn R., Jonathan S. Skinner, and Stephen P. Zeldes. 1995. Precautionary saving and social insurance. *Journal of Political Economy* 103 (2): 360–99.
Kimball, Miles S. 1990. Precautionary Saving in the Small and the Large. *Econometrica* 58 (1): 53–73.
Kutscher, Ronald E., and Valerie A. Personick. 1986. Deindustrialization and the Shift to Services. *Monthly Labor Review* 109 (6): 3–13.
Lakdawalla, Darius, Jayanta Bhattacharya, and Dana. P. Goldman. 2004. Are the young becoming more disabled? *Health Affairs* 23 (1): 168–76.
Mashaw, Jerry L., and Virginia P. Reno. 1996. Balancing Security and Opportunity: the Challenge of Disability Income Policy. Report of the Disability Policy Panel. Washington, DC: National Academy of Social Insurance.
Samwick, Andrew A. 1998. Discount rate heterogeneity and social security reform. *Journal of Development Economics* 57 (1): 117–46.