The Welfare Implications of Increasing DI Benefit Generosity

John Bound
University of Michigan, NBER

Julie Berry Cullen
University of Michigan, NBER

Austin Nichols
University of Michigan

Lucie Schmidt
University of Michigan

Prepared for the Third Annual Conference of the Retirement Research Consortium
"Making Hard Choices About Retirement"
May 17-18, 2001
Washington, DC

The research reported herein was preformed pursuant to a grant from the U.S. Social Security Administration (SSA) to the Michigan Retirement Research Center (MRRC). The opinions and conclusions are solely those of the authors and should not be construed as representing the opinions or policy of the SSA or any agency of the Federal Government or of the MRRC. The authors would like to thank Charlie Brown, Michael Chernew, Joel Slemrod, and participants in the University of Michigan Public Finance and RWJ Health Policy seminar series and the NBER Summer Institute for helpful comments and suggestions. Bound and Cullen would like to thank MECA for financial support, and Cullen also gratefully acknowledges support through the Robert Wood Johnson Foundation Scholars Program.

Abstract

The empirical literature on DI has primarily focused on the impact of program parameters on caseload growth or reduced labor force attachment. The focus on the efficiency costs of DI provides a misleading view of the social desirability of the program itself and of the adequacy of benefit levels. In order to provide a more comprehensive view, we develop a framework that allows us to simulate the benefits as well as the costs associated with a marginal increase in benefit generosity using a representative cross-sectional sample of the population. Using the 1991 March CPS, we estimate the total cost of providing an additional $1 of income to current DI recipients to be $1.42. While the load factor due to moral hazard is fairly high, we demonstrate that it is moderate enough that representative workers should be willing to "buy" additional insurance through reduced take-home pay at this price. The reform looks less attractive, however, once the financial benefits and costs are distributed across individuals in the sample. First, the average implicit price of an additional dollar of insurance is actually much higher than $1.42 for all but the least educated category of workers due to the redistributive nature of the program. We predict that the reform leads to a net welfare loss for these more highly educated groups, regardless of the level of risk aversion. Second, despite an average implicit price of less than $1, the expected utility gain also turns negative for high school dropouts under high levels of risk aversion. This counterintuitive finding arises since the utility calculation weights low income individuals more heavily as risk aversion increases and individuals with income below the floor provided to current DI recipients help to finance the benefit increase.

1. Introduction

Disability status is increasingly used as a means for targeting resources in the United States. In 1999, 5.6 million disabled individuals and 1.6 million spouses and dependent children received Disability Insurance (DI) benefits for a total outlay of $4.2 billion, representing over one-eighth of the Social Security budget for benefit payments. Given its importance and continuing growth, an evaluation of whether the DI program provides adequate insurance against the income losses associated with the onset of severe limitations is overdue. In this paper, we address this issue by evaluating the welfare implications of marginal changes in the level of benefit generosity.

The empirical literature on DI has primarily focused on the impact of program parameters on the costs of the program, either in terms of caseload growth or reduced labor force attachment. The focus on the efficiency costs of DI provides a misleading view of the social desirability of the program itself and of the adequacy of benefit levels. The effectiveness of the program depends on how the costs relate to the social benefits. What is striking is that there has not been any explicit valuation of the benefits associated with providing DI.

A literature on the economic disadvantage of the disabled, particularly the "doubly disabled" who are black, women, or have low levels of education does exist (e.g., Haveman and Wolfe 1990; Burkhauser, Haveman, and Wolfe 1993). This research also shows that publicly provided transfers are an increasingly important source of income for people with disabilities. However, these descriptive facts do not adequately capture the value of disability insurance for potential beneficiaries. This value depends on answers to counterfactual questions about what the income and work effort (leisure) of disabled individuals would be in the absence of DI or under different DI regimes.

In order to provide a more comprehensive view, our analysis accounts for the benefits as well as the costs associated with a marginal increase in benefit generosity. How an individual worker fares under this expansion depends on the expected change in lifetime utility. This expectation in turn depends on the impact on the path of family income and work effort both in the case where the worker becomes disabled and applies for DI and in the case where the worker never applies, and on the likelihood of each of these outcomes. Our approach is based on the intuition that if the population is in steady state, a representative cross-sectional sample of the population can be thought of as capturing the distribution of potential life-cycle paths for a representative individual or cohort. We can therefore use observed earnings patterns to simulate the impact of increasing DI benefits on family income for current recipients, potential new applicants, and workers in the sample. Combining data from the 1991 March Current Population Survey (CPS) with plausible assumptions about behavioral responses, we calculate the expected financial benefits accruing to and financial costs borne by each individual in the sample as a results of a one percent increase in DI benefits.

We then conduct the welfare analysis in two stages. In the first stage, we compare the increase in transfers to current recipients to the total costs associated with providing this increase. The ratio of total costs to the direct benefits received yields an estimate of the average implicit price of providing an additional dollar of income to current beneficiaries in the presence of moral hazard. This price differs from one only to the extent that there are behavioral responses to the benefit increase that generate additional tax and benefit costs, including changes in work effort and program participation among current beneficiary and potential applicant families. Using this price, we explore whether representative workers should find purchasing additional insurance through reduced-take home pay attractive. Whether different types of workers gain or lose from this reform provides evidence on whether they are under or over-insured against career-ending disabilities, given the existing level of generosity and degree of moral hazard in the system. By holding the degree of actuarial unfairness of marginal insurance constant across worker types, we are able to isolate the insurance-value of DI.

In our second approach, we recognize that like other social insurance programs, the value of DI derives from its dual role as a program that redistributes resources across individuals as well as insures individuals against adverse events. Thus, for example, for individuals at all levels of education, DI acts as privately provided insurance against the reduction in earnings capacity associated with the onset of disability. At the same time, DI redistributes from those with higher to those with lower levels of education both because benefits are relatively more generous for lower earners (replacement rates are higher) and because the less well educated are substantially more likely to become DI beneficiaries.[1] In order to determine the lifetime incidence accounting for the underlying distribution of costs and benefits across individuals, we value the change in income for each member of our sample. The fact that we are considering a marginal change in program generosity allows us to ignore changes in utility for individuals who are induced to apply to the program since these individuals are assumed to be indifferent between applying and not applying. The sum of individual valuations describes whether or not individuals gain on average from the benefit increase, where the average puts more weight on changes in income for individuals at the low end of the income distribution.

To summarize our findings, we estimate the total cost of providing an additional $1 of income to current DI recipients to be $1.42. While the load factor due to moral hazard is fairly high, we demonstrate that it is moderate enough that representative workers should be willing to "buy" additional insurance through reduced take-home pay at this price. The reform looks less attractive, however, once the financial benefits and costs are distributed across individuals in the sample. First, the average implicit price of an additional dollar of insurance is actually much higher than $1.42 for all but the least educated category of workers due to the redistributive nature of the program. We predict that the reform leads to a net welfare loss for these more highly educated groups, regardless of the level of risk aversion. Second, despite an average implicit price of less than $1, the expected utility gain also turns negative for high school dropouts under high levels of risk aversion. This counterintuitive finding arises since the utility calculation weights low income individuals more heavily as risk aversion increases and individuals with income below the floor provided to current DI recipients help to finance the benefit increase.

The remainder of the paper proceeds as follows. The next section provides background on the DI program. Section 3 develops the conceptual framework we use to evaluate the welfare impact of increasing DI benefit generosity. Section 4 describes the implementation using the 1991 March CPS and additional details on the assumptions that

[1] Much of the theoretical work relies on static one-period models, with variation across individuals in the degree to which they suffer a disability (modeled as the disutility of work). Productivity differences across

we make. Section 5 presents and discusses the results, and Section 6 concludes.

2. Background

DI is currently the most important disability program in terms of its scope and the magnitude of expenditures involved. The program provides benefits to disabled workers (and their spouses and children) in amounts related to the disabled worker's former earnings in covered employment. DI benefit payments are calculated in essentially the same way as are retirement benefits. Monthly benefits rise as a function of past earnings, but less than proportionately. Funding is provided through a Social Security payroll tax, a portion of which is allocated to a separate DI trust fund.

To be eligible for DI benefits, a worker must have worked a sufficient number of quarters in Social Security covered employment.[2] In addition, he or she must not be gainfully employed and must pass a medical screening. The determination of whether an applicant meets the medical requirements for disability involves a sequential process, under which about half of all applicants are turned down. Generally, the worker must be found incapable of gainful employment, taking into account age, education and work experience. The person must be disabled for a five-month waiting period before he or she can receive benefits. Given continuing restrictions on work, the vast majority of recipients refrain from working at all. After two years on the program, beneficiaries become eligible for Medicare. Successful applicants can expect to continue to receive cash and health benefits for the remainder of their lives. At the age of 65, beneficiaries are officially transferred from the disability to the retirement program, but their retirement benefits are equivalent to the benefits they received under the disability program.

The fact that DI is targeted to eligible workers distinguishes DI from the second most prominent disability program—Supplemental Security Income (SSI). While DI is essentially an early retirement program, SSI is a means-tested program providing benefits to needy disabled and blind as well as needy aged individuals, regardless of work history. The two programs serve different populations. The typical DI beneficiary tends to be

---

individuals are assumed away. In this context, the equity/insurance distinction disappears since everyone is identical.

[2] Roughly speaking, an individual has to have worked in covered employment for 5 of the 10 years prior to

older and is less likely to be female, a high school graduate, or white than is a typical member of the working aged population. The typical SSI recipient tends to be somewhat younger than the DI population and is even less well educated and more likely to be female. We focus our attention on the DI program because it functions more directly as an insurance program.

The case for mandatory public insurance for career-ending disabilities is relatively strong. First, given that these types of disabilities represent very low probability events, it is inefficient for individuals to self-insure through own savings. Second, in the absence of government intervention, adverse selection would likely prohibit a market for comprehensive long-term disability insurance from forming. Firms offering private disability insurance would attract those individuals with the highest risk of becoming disabled. At any given price, this form of consumer self-selection may prevent insurance companies from breaking even. The government can overcome problems of adverse selection by requiring participation, so that low-risk individuals are forced to pool with high risk individuals. While supplementary private insurance packages are currently available to some classes of workers, this market presumably owes its existence partly to the primary insurance offered by DI. In fact, the early private insurance market established before the public system was introduced collapsed during the Great Depression and was eventually resurrected with more limited coverage (cite).

Though public insurance can solve adverse selection problems, DI is not necessarily any more immune to moral hazard than private programs. The striking historical correlation between trends in DI participation and reduced labor supply among elderly males has led to a large literature measuring the disincentive effects of DI. While early generation estimates attributed nearly all the decline in reduced labor force participation to DI, more recent studies assign closer to one-third of the fall to DI (cites). Though DI is associated with large behavioral distortions, there is little in the literature that would permit one to say whether these sizeable effects outweigh or are outweighed

---

the point in time the worker becomes disabled.

by the benefits of the program.[3]

3. Conceptual Framework

Determining whether typical workers gain

The goal of the first phase of our analysis is to establish whether representative workers should be willing to trade reduced take home pay for increased DI benefits, given the degree of actuarial unfairness associated with the public insurance program. The program is actuarially unfair on average because of moral hazard. That is, if higher benefits lead to behavioral responses that increase the cost beyond the amount that is actually transferred to current recipients, it will cost more than $1 to increase the income of the disabled by $1. However, workers may still want to purchase additional insurance at this price if the marginal utility of income is high enough in the disabled states of the world.

The experiment that we consider is a one percent increase in DI benefit payments. In order to calculate the implicit price of additional insurance, we first calculate the static costs due to the increase in transfers to DI beneficiaries holding behavior constant. Then, we add the dynamic tax and transfer costs associated with the behavioral responses of these and other individuals to the benefit increase. From this point forward, we refer to the static costs as "direct" costs and the dynamic costs as "indirect costs."

Consider the difficulties that arise when determining even the increase in DI payments to current beneficiaries. Individuals who enter the program when they are 35 continue to receive benefits from then on. We would, therefore, need to project benefits forward to determine the full cost.[4] We greatly simplify our task by making a steady state assumption. As long as the overall population and the fraction receiving DI are not growing, the age distributions of both the overall population and of DI recipients are not changing, and there is no productivity growth, we can measure the costs from a random

[3] A notable exception is Gruber's (1996) analysis of a discrete change in benefit generosity in Canada that concludes with a back-of-the-envelope calculation of the net utility benefits to current beneficiaries, marginal applicants, and workers.

[4] Longitudinal Social Security records that track individual earnings and benefits over time might appear to solve the informational problem. However, these records do not provide information on family income amounts and sources over time that is integral to our calculations.

sample of the population. Each individual in the cross-section can be conceptualized as depicting the circumstances at one age within a possible life cycle path for an individual. Then, payments to 45-year olds on the program today can be thought of as picking up the payments to 35-year olds ten years hence. While these assumptions are not particularly realistic in the context of the age structure in the United States, they make the analysis easiest conceptually. We test the sensitivity of our results to relaxing these assumptions, and as we discuss later, find that the steady state assumption is not terribly misleading.

To measure the direct cost of increasing DI benefit payments, we identify current recipients in our sample and increase family benefits by one percent. The resulting change in public transfers to these families, however, will be moderated by mechanical interactions with other transfer programs. We use benefit algorithms to reduce cash SSI benefits and in-kind food stamp benefit amounts received. For those low-income recipients who are currently receiving SSI benefits, the net change in transfers will actually be zero since SSI benefits are reduced one-for-one with DI benefits. Finally, the resulting net increase in income will lead to higher income taxes paid by current beneficiaries. We use TAXSIM to estimate the change in federal and state income taxes paid.[5] The total direct costs of the reform are, therefore, the increase in DI benefits paid less any reduction in other transfers and any increases in taxes paid summed across families. This sum represents the net additional public outlay to families with current beneficiaries, as well as the net increase in income for these families.

We use this same steady state framework to calculate two sets of indirect costs. First, we allow current beneficiaries and their families to respond to the lump sum increase in income. Own labor supply has little room to respond given the restriction on gainful work imposed by DI, although family labor supply responses could potentially be sizeable (Cullen and Gruber 2000). Any reduction in labor supply represents a program cost since it will result in reduced income taxes paid by the family. Relying on the broad thrust of the literature on the long run income elasticity of labor supply (see Killingsworth 1983), we assume that husbands of current beneficiaries continue to earn the same as

[5] TAXSIM is the National Bureau of Economic Research's FORTRAN program for calculating liabilities under the U.S. federal and state income tax laws from individual data. We use version 4.0 available at URL (http://www.nber.org/taxsim/). For more information about TAXSIM, see Feenberg and Coutts (1993).

before, but apply an income elasticity of negative one for wives of current beneficiaries. After calculating the reduction in spousal earnings, we again use TAXSIM to estimate the associated tax cost.

The second set of indirect costs that we consider have been emphasized in the prior literature on DI—those associated with new applications. The magnitude of these costs will depend on how many more individuals apply for DI and how many of these new applicants pass the medical screening. Those new applicants who are successful will receive benefits, but since we are analyzing a marginal change and these individuals are assumed to be just indifferent between applying or not, we can ignore the value of these benefits and account for them solely on the cost side. Further, changes in the income and labor supply of successful and unsuccessful applicants and their families may lead to increased participation in other transfer programs as well as lower income tax payments. Workers not only have to pay for the DI benefit payments to new successful applicants, but must also pay for the higher public transfers and reduced tax revenue associated with successful and unsuccessful applications.

Before modeling the behavioral responses, we first must determine how many individuals will be induced to apply by the increase in DI benefit generosity, and how many of those applications will be successful. Imagine that the 1% increase in benefits induces an x% increase in awards that is neutral with respect to age groups. In the new steady state, the age distribution of those on DI would be the same as before, but the overall size of the DI population would be x% larger. We rely on evidence from prior literature to predict how many new applicants and beneficiaries will result from the policy reform. Although estimates of the elasticity of applications with respect to benefit generosity vary considerably (Bound and Burkhauser 1999), plausible bounds exist. The estimates we find the most credible are based on time series evidence from the 1960s and 1970s, a period of time when replacement ratios rose appreciably (Halpern 1979; Lando, Coate and Kraus 1979). These estimates suggest elasticities in the neighborhood of 0.5. Thus, our baseline assumption is that a 1% increase in benefits leads to a 0.5% increase in applications.

To link the application elasticity to the size of the current pool of recipients, we need evidence on the ratio of the number of those who ever applied in the population to the number of current beneficiaries. Historically, roughly half of those who apply for DI are awarded benefits (Bound and Burkhauser 1999), which would imply a ratio of ever applicants to beneficiaries of two. The ratio will be greater than this estimate to the extent that mortality is higher among successful than unsuccessful applicants, and will be smaller to the extent that many individuals apply several times. Using matched data from the 1990-93 Survey of Income and Program Participation (SIPP) panels and Social Security Administration (SSA) administrative disability determination records, we are able to estimate the ratio directly for individuals included in the SIPP panels whose applications were acted upon between 1977 and 1997. This evidence suggests that the two effects mentioned above largely cancel. Therefore, we assume that the steady state number of applicants is approximately twice as large as the recipiency base. This implies that the elasticity of beneficiaries with respect to benefits is the same as the application elasticity.[6] We describe in the next section how we identify these new unsuccessful and successful applicants in our sample and assign DI family benefits to them.

We rely on the same matched SIPP and SSA data to guide the assumptions that we make about the behavioral responses of marginal applicants and their families. Our regression framework combines information on individuals who are observed in the SIPP either at some point before after applying, during the application process, or at some point after to create a pseudo panel. This methodology, which is described in detail in Appendix A, allows us to trace out the impact of both successful and unsuccessful applications. Relative to a control group of similar individuals with work limitations who never apply, we attribute changes in the level and incidence of earnings and transfers from one year before to one year after application to the application itself. Our approach will, if anything, tend to overstate the causal impact of application since some of the change is likely due to the onset of disability.

The time patterns for earnings and transfer receipts are shown in figures A1 a-f through A5 a-f for awarded and denied applicants separately by gender. The panel on the

[6] It seems likely that those induced to apply by an increase in benefits would tend to be less severely impaired than the typical applicant. In this case, they would be less likely to pass the screening. Since unsuccessful applicants cost less in terms of taxes and transfers than do successful applicants, our

left in each set of figures corresponds to successful applicants ("awarded"), and the right-hand side panel shows the results for unsuccessful applicants ("denied"). The y-axis in all cases shows the average baseline value of the variable of interest and the x-axis indicates the number of months relative to the time of application, with a negative sign indicating months before application. Each of the five sets of graphs is based on one of following monthly income measures: own earnings conditional on positive earnings, own labor force participation, spouse's earnings (unconditional), spouse's labor force participation, and other transfer income (excluding AFDC, Food Stamps, and SSI).

As expected, we find that nearly all successful applicants leave the labor force permanently (Figure A2). For this group, there is no evidence of a significant change in average spousal earnings[7] or in average transfers (Figures A3 and A5). All successful applicants will by definition now be eligible for SSI if family income is low enough, since the screening process is the same for both disability programs. We use the 1990 algorithm to predict SSI benefit eligibility and amounts and assume full take-up. We use Food Stamps and AFDC algorithms to predict eligibility and benefit amounts after the reform. For the Food Stamps program, we use the matched SIPP-SSA data to estimate average take-up among eligible awarded applicants more than one year after application (53.5%), and apply this estimate to randomly assign eligibles to the program. The sample is too small to generate reliable estimates of take-up for those eligible for AFDC, so we assume full take-up. Therefore, to measure the public costs associated with successful applications, we remove own earnings, estimate changes in SSI, AFDC, and Food Stamps benefit receipt, and then calculate any changes in state and federal income taxes paid. For now, we ignore the fact that these individuals become eligible for Medicare after two years on the program, and address the role of health insurance in a separate section below.

Though labor force participation rates of unsuccessful applicants drop sharply prior to the month of application, participation rebounds somewhat over time (Figure A2). For denied males, 31.8% as many are working by one year after application as were working one year before. The same estimate for females is 49.9%. These rates reflect

---

assumption should lead to an overestimate of the costs of a benefit increase.

monthly participation and could be consistent with some fraction returning to work full time and some fraction returning to work part time or intermittently. In the absence of evidence on the appropriate combination, we assume the above fractions of denied applicants who are initially working in our sample return to work and that the remaining individuals leave the labor force. The figure showing the pattern in earnings over time for denied applicants (Figure A1) suggests that denied applicants who return to work earn the same as before. There is no evidence of a sizeable impact of application on either spousal labor supply or other transfers (Figures A3 and A5), so we ignore these. Finally, these individuals will not be eligible for SSI, but we handle AFDC and Food Stamps in the same way as for successful applicants, except that we use the estimated post-application Food Stamps take-up rate specific to denied applicants (70.9%).

Though we account for the most obvious behavioral responses, we ignore several classes—some of which we expect to have a negligible impact on our calculation and others that are potentially important but about which we can do little. We ignore the deadweight burden associated with raising taxes to pay for the increased benefits. Empirical evidence (Gruber 1994; Anderson and Meyer 1995) suggests that payroll taxes that support programs that are restricted to workers are essentially benefit taxes, so that the deadweight burden of the tax increase will be negligible. However, as we demonstrate, DI is far from a pure benefit tax since it involves substantial redistribution across worker types. There is not an easy solution to account for deadweight loss in this case, such as by scaling the tax costs upward using factors estimated in the literature, since some individuals receive subsidies and others are taxed. Ignoring this cost is equivalent to assuming the net burden is zero.

We also ignore a variety of other responses to changes in DI benefit generosity, such as changes in savings or marriage rates or participation in private disability insurance programs. We suspect that none of these effects is likely to be large. The occurrence of a long-term disability during a person's working years is a relatively rare event, but is likely to represent a reasonably permanent state. Because of this, saving to

---

[7] While this appears to be inconsistent with our assumption that wives of current beneficiaries respond to the lump sum DI payment, application to DI involves competing changes in the value of leisure and in income that apparently generate a net zero response for spouses.

self-insure against that eventuality is not very effective, and is therefore unlikely to be heavily used and unlikely to respond to changes in benefit levels. Moreover, applicants for DI typically have little savings and are typically not covered by any type of private disability program, leaving little room for crowd-out along these margins.[8] While marriage does significantly mitigate the costs of becoming disabled, and as such, does help insure individuals against the associated income losses, we do not expect marriage rates to respond to marginal changes in DI benefit levels.

In order to determine whether a representative individual would be willing to purchase additional insurance given the degree of actuarial unfairness, we first estimate the total cost of transferring an additional \$1 to current DI recipients. The average implicit price of insurance is equal to the ratio of total direct and indirect costs to the net increase in transfers to current DI recipients. Again, this price differs from one only to the extent that behavioral distortions are important, reflecting the impact that moral hazard has on the cost of insurance.[9]

In order to determine willingness to "buy" more insurance at these prices, we need to determine the relative value of income to an individual across able-bodied and disabled

$$U(C,H) = v(C) - f(H) = \frac{C^{1-\theta}}{1-\theta} - f(H), \qquad (1)$$

states of the world. We assume a utility function that exhibits constant relative risk aversion, and is separable in consumption and leisure:

where C is consumption, H represents hours worked, and θ is the coefficient of relative risk aversion. Given this utility function, the marginal utility of income is $C^{-\theta}$. The relative value of \$1 across the two states of the world can be expressed as $\left(\frac{C_d}{C_{nd}}\right)^{-\theta}$, where C_d represents consumption when disabled and C_{nd} represents consumption when

[8] Only 8.5% of DI recipients in the March 1991 CPS report receiving private disability insurance. Those who receive private benefits tend to have more years of education than those who do not (12.5 vs. 10.3). Median savings among those who apply for DI benefits in the matched SIPP-SSA sample are only \$600.

[9] The theoretical literature on disability insurance has been explicit about viewing disability insurance as a form of mandatory income insurance where benefit generosity is limited by moral hazard (e.g. Diamond and Sheshinski 1995).

able-bodied. We rely on outside studies to approximate the relative consumption levels. Then, we compare willingness to pay calculated at varying levels of risk aversion to the implicit price of additional insurance.

Determining whether workers benefit on average

Given our estimates of the direct and indirect costs, how these costs are distributed across workers determines how the average worker fares. While DI transfers resources from younger to older individuals and from able-bodied to disabled individuals in the cross-section, from our steady-state perspective DI transfers resources from earlier to later years and across different states of the world. If the cross-sectional distribution captures the full distribution of potential life-cycle paths, then the net benefits to the random sample measure the average life-cycle incidence.

The approach that we use is to first simulate the change in utility following the reform for each individual in our sample, and then to aggregate across all individuals. To clarify our approach, it is helpful to decompose the sum of utilities in the following way:

$$W = \sum_{ijt} v(C_{ijt}) - f(H_{ijt}) = \sum_{i}\sum_{jt} v(C_{ijt}) - f(H_{ijt})$$
$$dW = \sum_{i}\sum_{jt} \frac{dv}{dC_{ijt}} dC_{ijt} - \frac{df}{dH_{ijt}} dH_{ijt} \quad , \qquad (2)$$

where i indicates individual, j indicates possible realizations (e.g. disabled or not), and t indicates the time period or age. Each member of our sample can be conceptualized as representing a point in time on one possible life path for an individual. Imagine all of the observations representing a given individual's possibilities, and imagine that individual at some base period early in her life. Also, focus on the financial component of utility, which we justify is reasonable for our analysis below. Our calculation of $\frac{dv}{dC_{ijt}} dC_{ijt}$ for these observations can be thought of as that person's report of how she values the change in consumption in time period t and state j caused by the reform. The sign of the sum across j and t describes whether she benefits from the reform or not. Subsequently aggregating across different persons weights their net valuations by the relative marginal values of income at the base period.

While conceptually we can separate individuals and life paths, it is impossible to do so in our cross-section. The role of the risk aversion parameter is complicated in this setting. From the perspective of a given individual, the utility gain from reducing the potential income loss from the onset of disability depends on the level of risk aversion in the traditional manner. A higher level of risk aversion increases the weight on low income periods, making insurance more desirable. But, in the welfare calculation, a higher level of risk aversion also increases the weight that low income individuals receive. The social net benefit of the reform will vary with the coefficient of relative risk aversion through both of these channels, so that redistribution and insurance effects are entangled.

We derive the individual-, realization-, and period-specific components of equation (2) by simulating the point in time change in utility for each person in our CPS sample. We begin by distributing the costs of the increase in DI benefit generosity according to current funding patterns. That is, increases in DI benefit payments are shared across workers in proportion to earnings below the Social Security maximum. Net tax costs due to lost tax revenue and other transfers are distributed according to workers' shares of federal and state income taxes paid. We can then simulate the impact of the reform on family income for current beneficiaries, potential new applicants, and workers. In this case, workers' views of the desirability of increasing DI generosity includes the effect of the progressivity of DI replacement rates and of the tax code on their relative shares of these new costs.[10] The degree to which the increased insurance is actuarially unfair is folded into the analysis since the predicted shifts in family income absorb the tax increases due to direct and indirect financial costs.

The next step is to translate the changes in income and leisure into changes in utility. Assuming that utility takes the form described above, we are able to focus exclusively on the financial component. The only changes in leisure considered in our analysis are among families of current beneficiaries and marginal applicants. For the

[10] While the above method of distributing costs is the most immediately policy-relevant approach, a second method would instead distribute the net costs across workers in proportion to expected net benefits. Expected net benefits would be based on a worker's age-adjusted position in the earnings distribution and would require information on the pre-application position of DI applicants. This would have the advantage of more effectively isolating the insurance value of increased benefit generosity because the degree to which the policy change is actuarially unfair (e.g. expected financial payments exceed expected benefit receipts) reflects only the relative magnitude of the costs arising from moral hazard.

families of current beneficiaries, the increase in transfers (net their share of costs) are valued ignoring any changes in spousal earnings and leisure. Though these families choose to "spend" some of the increase in income on increased leisure, the marginal utility gain can be captured through the income effect. While the direct benefits accrue to those already on DI through the increase in income, indirect benefits accrue to the individuals induced to apply for DI benefits. However, since these individuals are at the margin, revealed preference suggests that they will gain relatively little from the increase. Thus, we assume that the change in utility for marginal applicants is zero.[11] Other families experience only a change in income according to their share of costs.

The table below shows the changes in family income and utility for three different possible point-in-time family types: families with a current beneficiary (b), with a marginal applicant (a), and with only able-bodied individuals (w). Think of these family types as representing specific realizations, so that representative workers in the base period have some chance of finding themselves in any one of these families at some point.

Changes in Income and Utility by Family Type

Family Type	Change in Family Income	Change in Utility
Current Beneficiary	$\Delta Y_{ibt} = \Delta D_{ibt} + \Delta T_{ibt} + \Delta E_{ibt}$	$v(f(Y_{ib0} + \Delta D_{ibt} + \Delta T_{ibt})) - v(f(Y_{ib0}))$
Marginal Applicant	$\Delta Y_{iat} = p^s_{iat}(\Delta D^s_{iat} + \Delta T^s_{iat} + \Delta E^s_{iat}) + p^u_{iat}(\Delta T^s_{iat} + \Delta E^s_{iat})$	0
Able-Bodied	$\Delta Y_{iwt} = \Delta T_{iwt}$	$v(f(Y_{iw0} + \Delta T_{iwt})) - v(f(Y_{iw0}))$

In the table, D represents DI benefits, T represents taxes and transfer payments, E represents earnings, Y_0 stands for family income before the regime change, and p^s and p^u are the probability of successful and unsuccessful DI applications, respectively. Families

[11] For the benefit to marginal applicants to be equal to zero requires additional assumptions on the distribution of preferences across individuals. Consider the demand curve for DI benefits. As the cost of applying goes down, more and more individuals apply for benefits, with the demand for benefits curve tracing out the distribution of individuals' willingness to pay. As long as the aggregate demand curve for benefits is smooth and continuous in the level of benefit generosity, the marginal applicant will be indifferent between applying and not applying. To the extent this does not hold, we ignore a non-zero gain to this group and thus underestimate the benefits associated with increasing DI benefit generosity.

of current beneficiaries witness a one percent increase in DI benefits (ΔD_{ibt}), lose some interrelated transfers and pay a share of the costs (ΔT_{ibt}). Also included in the change in income for these families is any reduction of spousal labor supply in response to the net increase in transfers (ΔE_{ibt}). The change in period utility for individuals in these families is determined by the change in individual consumption, modeled as a function of current family income ($f(Y)$), ignoring the spousal labor supply response. Families of successful and unsuccessful applicants both experience changes in non-disability transfer and tax payments and in earnings, but only successful applicants receive DI benefits. At the point of the decision to apply, these families are assumed to be indifferent between applying or not applying and to have no change in expected utility. And, families with all able-bodied individuals simply pay their share of the costs of the reform, valued according to the associated reduction in individuals' consumption.

The final step is to specify how changes in family income are converted to changes in individual consumption. One alternative is to simply proxy for consumption using per capita family income. However, there is evidence that families are capable of buffering themselves against short-term fluctuations in income. Therefore, we also use an alternative proxy to account explicitly for divergence between income and consumption due to saving. We use data from the Consumer Expenditure Survey (CEX) to calculate the relationship between concurrent family income and consumption, and then smooth family income accordingly.[12] In addition, in both the cases where family income is smoothed and unsmoothed, we calculate individual-level measures using a per capita equivalency in addition to a strict per capita measure. Family income is divided by a factor that accounts for both differences in consumption between children and adults and economies of scale.[13]

The estimated changes in consumption are then fed through the utility function to calculate dU_{ijt} for each CPS sample member. The sum of these net benefits determines the social value of increasing net benefits. The same calculation for workers with given

[12] The methodology is described in detail in Appendix B.

[13] We use the adjustment factor suggested by the National Research Council: $(N_a + .8{*}N_c) \times .7$, where N_a is the number of adults and N_c is the number of children.

characteristics determines the value of the policy change to subsets of individuals. We present results at both low and high levels of risk aversion in order to test the sensitivity of our findings.

The Role of Health Insurance

Up until this point we have ignored the role that medical insurance plays in both welfare analyses. For the first analysis, if potential applicants become newly eligible for public health insurance programs, the total costs of the reform may increase. However, determining this net fiscal externality is very difficult. Even individuals who are insured through private plans may be receiving subsidies from other lower cost individuals covered by the same plan. Similarly, individuals who are uninsured may impose costs on others through the system.

Consider two extreme cases for successful applicants, who are able to move onto the Medicare program after receiving DI for two years. In the first case, assume that the net subsidy through health insurance or uncompensated health care prior to application is as great as the subsidy provided by Medicare. Then, the net increase in the financial burden on others will not change, though clearly the distribution of the burden will differ. In the second case, assume that the applicants bear the full cost of their health care prior to the reform, perhaps through lower wages, so that Medicare expenditures on these individuals after they move onto DI are new public costs. The truth is certainly somewhere in between these two extremes. Assuming the latter case holds would provide an upper bound on the share of Medicare costs in the total costs of a marginal increase in DI benefit generosity. We consider both this case and the case where only applicants who were not previously covered by publicly provided health insurance generate net new health care costs.

We use information on Medicare expenditures to approximate these costs for new successful applicants. Data published in the Social Security Bulletin (1991) imply that the average ratio of Medicare benefits to DI family benefits is 0.475.[14] We will therefore

[14] The Social Security Bulletin (1991) reports that total annual Medicare spending for DI beneficiaries was $11,239 million in 1990. In that same year, monthly payments to DI worker beneficiaries, their spouses, and children were $1,768 million, 40 million, and 162 million, respectively.

scale the value of DI benefits to successful applicants by 1.475 to incorporate Medicare costs.

The role that health insurance plays in the second welfare analysis is much more complicated. Valuing medical insurance for marginal applicants, albeit implicitly since we continue to assume they are just indifferent between applying and not applying, requires valuing medical insurance for the working population as well in order to determine baseline utility. The insurance value is not the dollar value of benefits but the insurance value. Imagine that insurance status has little effect on what kind of medical procedures one has, but affects only the out-of-pocket costs. Then, the cost of being uninsured is the possibility of having to pay large medical bills and can be modeled as a loss of income. With data from the CEX on the distribution of the income share of out-of-pocket medical costs by insurance status and age (see Appendix C), we randomly reduce income by these amounts before proceeding.

The specific implementation in the CPS of the methodologies outlined above is described in the next section.

4. Implementation

The random sample of the population that serves as the basis for our analysis comes from the March 1991 Current Population Survey (with income data for calendar year 1990). We take both the size of the disability population and the implementation of the program (e.g. screening stringency) as given when evaluating the desirability of increasing benefit generosity. Since 1990 predates the recent surge in program growth, [15] we conduct the analysis with the DI population inflated to current levels in an extension.[16] Our sample consists of 158,318 individuals residing in 73,638 tax-filing units.[17]

An issue immediately arises with our assumption that the population is in steady

[15] See Stapleton et al (1998) for a discussion of the various factors behind this recent growth.
[16] As expected, inflating the size of the DI population to current levels increases the implicit price of insurance, but only from \$1.42 to \$1.43. It has no qualitative effect on the welfare implications of the policy reform.
[17] Our definition of tax-filing units essentially matches the CPS definitions of subfamilies, except that adult unmarried children residing with their parents are considered to be their own tax-filing units.

state. The hump-shaped age distribution we observe in the CPS is not consistent with a feasible steady state. That is, the fraction of the population that falls within each age category has to be monotonically declining, or at least non-increasing. As a sensitivity test, we re-weight our sample assuming constant fertility and mortality rates to approximate a feasible steady state distribution.[18]

Identifying Recipients and Potential Applicants

Among the families in the CPS, we identify three categories that will be treated differently in the calculations: those with current DI beneficiaries, those with potential applicants, and those families with neither. In order to identify DI beneficiaries and potential applicants, a further distinction is made between "working aged" (aged 21-61) and "elderly" (aged 62 and older) individuals. While it is clear that both beneficiaries and applicants should be found among the working aged population, understanding why both are among the elderly requires explanation. Recall that our conceptual framework views an elderly person as the aged version of a working aged individual. Therefore, following the reform, the status of a 75-year old individual may change if the individual either received DI benefits or was induced to apply for DI at an earlier age. In order to simulate the full impact of the reform, we need to identify which elderly individuals would have been DI beneficiaries and applicants in their working years. We refer to such individuals as elderly DI beneficiaries or marginal applicants, though technically they are elderly individuals who belonged in these categories at younger ages.

Since the CPS does not explicitly identify which individuals are DI recipients, we first need to impute DI beneficiary status. This is fairly straightforward for the working aged population. The CPS includes self-reports of whether or not an individual has a condition that limits his/her ability to work. In addition, it includes information on individual receipt of Social Security income. We assume that individuals aged 21 to 61 who report having such a limitation and who receive Social Security benefits are receiving those benefits on the basis of their disability. This imputation appears to be fairly accurate, capturing over 96% of the number of recipients reported in the Social

[18] Reweighting the sample to approximate a feasible steady state distribution lowers the implicit price of an additional dollar of insurance slightly, from $1.42 to $1.37, but has no qualitative effect on the welfare implications of the policy reform.

Security Bulletin administrative data for 1990.[19] Table 1 summarizes the characteristics of DI beneficiaries relative to working aged individuals for both males and females. For both men and women, DI beneficiaries are over 50% more likely to be black than other working aged individuals. They also tend to be older and have lower levels of education. These types of demographic differences drive the redistribution within the DI program.

The process of identifying current beneficiaries is more difficult for individuals over the age of 61. We would like to identify those who were receiving DI benefits during their working years, before the disability and retirement programs merge. However, for those over 61, it is impossible to distinguish Social Security Disability Income from Social Security retirement income. In addition, since self-reported disability rates increase with age, they cannot be used effectively to determine which individuals were DI beneficiaries during their working years. Rather than identify a subset of elderly individuals as DI beneficiaries with certainty, we impute the probability that each elderly individual would have been a DI beneficiary by assuming that elderly DI beneficiaries are similar to the older working aged beneficiaries already identified.

In order to calculate this probability, we first relate the probability that individuals aged 51-61 are worker beneficiaries to their individual characteristics, using a logit specification run separately by gender. Results from these regressions can be found in the first two columns in Table 2. We then use the coefficients from this regression to predict the probability that individuals aged 62 and over are DI recipients, calculated by setting their age to 61. Then, any individual who does not report Social Security income is assumed to have a zero probability of DI receipt.[20] Assuming similar mortality rates across disabled and non-disabled individuals, the probabilities are rescaled proportionately so that the fraction of the population imputed to be DI recipients at any age over 61 is the same as the fraction at age 61 (7.6% for males and 6.5% for females).[21]

[19] The Social Security Bulletin documents 2,479,270 worker beneficiaries between the ages of 21 and 61 in 1990. Weighting the count of imputed beneficiaries by the CPS person weight yields an estimate of 2,387,258 worker beneficiaries.

[20] If a family has more than one elderly individual, we assume that only the individual with the highest predicted probability has a non-zero probability of being a DI recipient. We assign the probability that either is a DI recipient to this individual.

[21] In fact, mortality rates among DI beneficiaries tend to be higher than mortality rates among the newly retired. Thus, our assumption is likely to exaggerate the fraction of the elderly who would have been DI

Of the 22,966 individuals in our sample over the age of 61, 14,842 are assigned a non-zero probability of being a DI beneficiary. Summary statistics for the elderly population as a whole and for imputed beneficiaries are found in Table 3.

Once we have identified current DI beneficiaries, the next step is to identify those who would be marginal applicants under the new regime. To this end, we have had to rely on extreme assumptions about who would be induced to apply. Among working aged individuals, we assume that the pool of potential applicants consists of all persons who self-report having disability-related work limitations but who are not identified as current DI beneficiaries.[22] The last column for both men and women in Table 1 presents summary statistics for these individuals. Among elderly individuals, we assume that the pool of potential applicants is the same as the pool of individuals imputed to have a non-zero probability of having been DI beneficiaries.

Following the regime change, each potential applicant may follow one of three paths. They may choose not to apply for DI, they may apply and be accepted to the program, or they may apply but be rejected. Each of these scenarios involves different indirect costs associated with changes in benefit outlays, labor supply, transfers, and tax payments. Rather than attribute a specific path to each individual within the potential applicant pool, we assign individuals a probability of each of these three outcomes.

In order to predict the likelihood of each of these three scenarios for working aged limited individuals, we first simulate the likelihood that each non-beneficiary would be induced to apply for DI benefits given estimates of the elasticity of application and awards with respect to benefit generosity from the literature. The simplest way to calculate the likelihood that any given limited individual applies would be to assume that each potential applicant is equally likely to apply. Our assumption that one-half of applicants are successful implies that the size of the applicant pool is twice the size of the pool of current beneficiaries. With the additional assumption that the elasticity of

---

recipients. Calculations using existing estimates of the mortality rate among DI beneficiaries suggest the magnitude of this effect is small.

[22] We exclude limited individuals who report receipt of Supplemental Security Income since these individuals would have had to have passed the same screening used for DI recipients and must have been determined to be ineligible for DI due to inadequate work history. While true in general, this ignores the possibility that some SSI recipients are in the process of applying for DI.

applications with respect to benefits is .5, the probability that a given limited individual applies in response to a 1% increase in benefits is $.01 \times .5 \times (2 \times D)/L$, where D is the number of current beneficiaries and L is the number of limited non-beneficiaries.

However, it is unlikely that each limited person is equally likely to apply. To account for individual differences in this proclivity for working aged individuals, we adjust the above probability using information on the likelihood that limited individuals with differing observable characteristics are on DI. The results of logit regressions predicting DI receipt among limited individuals separately by gender are reported columns 2 and 3 in Table 2. The predicted probability that any given individual applies is as shown above, weighted by a term that is equal to that individual's predicted likelihood of being on DI relative to the average probability among the limited non-recipient population. The resulting estimates of the probability of applying rely on the assumption that marginal applicants are observably similar to current recipients.

For elderly individuals, the probability of applying is derived directly from the imputed likelihood that the individual is a current DI beneficiary. We scale the probability that an elderly individual is a current DI beneficiary by $.01 \times .5 \times 2$ to take into account the magnitude of the benefit increase, the application elasticity, and the likelihood of acceptance.

With estimates of the likelihood that an individual is induced to apply on hand, we derive the likelihood of each outcome for a potential applicant. If we define the probability that an applicant applies as a_i, then, given our assumption that half are accepted, the probabilities of applying and being accepted and of applying and being rejected are both equal to $.5 \times a_i$. The likelihood of not applying is simply $(1-a_i)$.

Our final category of individuals includes those who are not limited.[23] These individuals are affected only indirectly by the reform through tax increases.

Calculating the Direct and Indirect Costs

[23] We assume that those who do not report a limitation in the baseline are not in the pool of potential applicants. This ignores the possibility that self-reported disability might respond to changes in the parameters of DI. Previous work has suggested that the fraction of the working aged population identifying themselves as unable to work is quite sensitive to the availability and generosity of disability benefits, but that the fraction identifying themselves as in some way limited in their capacity for work is much less sensitive to such factors (Waidmann et al 1995).

The increase in DI benefit generosity involves a variety of costs, which can be grouped into two main categories. Direct (static) costs result from the transfer of additional income to current beneficiaries, and involve no behavioral response. Indirect (dynamic) tax and transfer costs are generated by the behavioral responses of individuals to the regime change.

The first component of direct costs is calculated by scaling family benefits for current working aged and elderly DI beneficiaries to reflect the 1% increase in benefit generosity. We first impute family benefits for both worker and elderly beneficiaries. Since all of the working aged beneficiaries in our sample receive DI benefits on their own account (that is, due to their own disability and work history), we impute family benefits for our working aged beneficiaries based on the reported level of worker benefits and the rules followed by the Social Security system. The baseline level of worker benefits associated with the working aged population is estimated by our methodology to total $13.8 billion, as compared with the value of $17.3 billion determined from the administrative data reported by the Social Security Bulletin. To fully account for the costs associated with the benefit increase despite apparent underreporting of DI benefits in the CPS, we scale all reported SS benefits by a factor of 1.25.[24] The 1% increase, therefore, generates direct benefits costs among working aged individuals of $191.1 million.

For the elderly, all of the individuals assigned some positive probability of being a DI beneficiary report receiving Social Security income. However, in the scenario under which these individuals are treated as having been worker beneficiaries, an increase in DI benefit generosity will only affect family income if these individuals are receiving benefits due to their own work history. If the individual currently receives a Social Security benefit amount determined by a retired or deceased spouse's benefit eligibility, the reform will have no impact on benefit payments.

To identify those receiving Social Security on their own account, we first group elderly beneficiaries into two subsets: those in families with only one elderly Social

[24] The undercount of Social Security benefits in the CPS appears to affect all age levels. An equivalent scaling factor derived for the population aged 62 and over is equal to 1.21. Due to the similarity of the implied scaling factors, we apply the working aged scaling factor to all Social Security benefits for

Security recipient and those in families with more than one. Among those in families with only one (70% of elderly beneficiaries), we assume that female elderly beneficiaries receive benefits on their own account only if they were never married, or if they are currently married to a man reporting no Social Security income. We assume that male beneficiaries in families with only one recipient are all receiving benefits on their own account. For the 30 percent of elderly beneficiaries in families with two Social Security recipients, we look at the ratio of their own benefits to the highest benefit level within the family. We assume that they are receiving benefits on their own account if this ratio is greater than 0.5, since spouses are eligible for the maximum of the benefit amount determined based on their own work history and half of their spouse's benefit payment. Then, we impute family DI benefits for those who receive benefits on their own account based on their reported Social Security benefits.[25] The direct benefit costs for elderly beneficiaries are equal to 1% of current family DI benefits multiplied by the probability the individual had been a working aged recipient. Total costs of increased benefits paid to elderly beneficiaries are estimated to be $129.6 million.

Additional costs are generated mechanically by the interaction of the DI program with other transfer programs. DI beneficiaries who were also receiving SSI benefits find their SSI benefit amounts reduced dollar for dollar with the increase in DI benefits. This leads to a cost savings of $28.9 million. In addition, DI beneficiaries currently receiving food stamps will have the value of these in-kind benefits reduced at a rate of 30% with additional DI income. This interaction leads to a cost savings of $0.9 million. Finally, additional income to DI beneficiaries means that these families will owe additional federal and state taxes. We determine the net tax savings by first using actual family non-DI income and DI benefits to calculate baseline federal and state income taxes owed, and then recalculate taxes owed adding the net increase in transfers. We rely on the National Bureau of Economic Research's TAXSIM program for this and all following income tax

consistency.

[25] For those elderly who were awarded DI benefits, the DI benefit formula is used to calculate their benefit level instead of the retirement formula. However, without detailed work histories for these individuals, we are unable to use these formulas to calculate a true level of benefits under each program. Instead, we assume both DI and retirement benefits to be equal to reported Social Security benefits for those determined to be receiving benefits on their own account.

calculations. These changes in tax liabilities lead to savings of $4.1 billion in federal taxes and $0.5 billion in state taxes. The sum of these components is the net new transfer to current DI recipients assuming no behavioral response, and totals $286.4 million.

We account for two main types of indirect costs generated from the behavioral responses to the regime change. First, we account for spousal labor supply responses to the increase in transfers to families of current beneficiaries. As mentioned earlier, we assume that husbands of current DI beneficiaries do not reduce their labor supply, but that wives of beneficiaries reduce their earnings by the same percentage as the increase in other family income. This labor supply response reduces the amount of taxes owed by the family, leading to a loss of $1.2 million in federal taxes and $0.3 in state income taxes.

The second type of indirect cost is generated from the response of those individuals induced to apply by the policy change. We must first account for the new benefits paid to individuals who are induced to apply for benefits and are successful. Ideally we would like to measure the Social Security benefits to which a given individual would be entitled based on their work history. However, since the CPS does not include earnings histories, we instead predict the benefits that working aged potential applicants would receive if successful by returning to our sample of current working aged DI beneficiaries. We regress their actual log benefits on their individual characteristics by gender (regression results can be found in columns 5 and 6 of Table 3), then use the resulting coefficients to predict individual benefits for our marginal applicants. For elderly potential applicants, we simply use family Social Security income benefits as calculated before from reported benefits. This ignores changes in benefit payments that would result from altered work histories. The total simulated cost of DI benefits paid to marginal applicants is $89.7 million, which is about one-fourth as large as the total direct benefit cost.[26]

Among these marginal applicants, we allow reductions in own labor supply and

[26] As a plausibility test, we compare the indirect benefit cost for working aged individuals to an alternative crude calculation. Indirect benefit costs should be approximately equal to total benefits paid to current working aged beneficiaries multiplied by .05 (the 1% benefit increase times the application elasticity). The ratio between the amount resulting from the more involved calculation detailed in the text and this amount is 0.94.

account for changes in public transfers and taxes based on the corresponding changes in family income. We assume that successful applicants stop working and never work again. For unsuccessful applicants, estimates from the SIPP-SSA data (see Appendix A) suggest that 31.8% of denied male applicants and 49.9% of denied female applicants will return to work. We assume that those who return to work do so at their previous earnings. Those marginal applicants who do not return to work generate a decrease in both federal and state taxes paid.

Finally, we consider changes in three transfer programs—AFDC, food stamps, and SSI. Under the scenario where applicants are rejected, we remove the individual's wage and salary and self-employment earnings. We use the new income to calculate AFDC benefit levels, assuming 100% take-up rates.[27] We then calculate food stamp benefit levels, assuming a 70.9% take-up rate estimated from the SIPP for denied applicants. The change in transfer payments is calculated as the difference between baseline amounts and the imputed post-reform amounts.

The scenario where marginal applicants are accepted is similar, except that we add predicted family DI benefits to family income. We again calculate the new level of AFDC and food stamp transfers to which the family is entitled, now using a 53.5% take-up rate for food stamps estimated from the SIPP for accepted applicants. In this case, since applicants have successfully gone through the disability determination process, they may also be eligible for SSI benefits if their DI benefits are low enough. We use the SSI algorithm to calculate these additional benefits. The net expected transfer costs due to the behavioral responses of the marginal applicants includes these changes in AFDC, food stamps, and SSI benefits paid.

The costs associated with the behavioral responses of the potential marginal applicants depend on how likely successful and unsuccessful applications are. As a result, the *expected* cost of the increase in benefit generosity due to marginal applicants is based on the predicted likelihood of being in a given scenario. We weight the costs for the respective scenarios by the probability that the marginal applicant applies and is rejected ($.5\times a_i$) and is accepted ($.5\times a_i$). The probabilities for the rejected applicants are

[27] We attribute AFDC benefits only to eligible single parent families and ignore the AFDC-UP program,

further adjusted to reflect the likelihood that they return to work.

These expected indirect costs of the reform are summarized along with the total costs in Table 4. The federal tax costs are $22.8 million, and the state tax costs are $5.1 million. We actually predict savings on AFDC benefit payments, since few marginal applicants receive AFDC in the base case, few become eligible as unsuccessful applicants, and most become ineligible as successful applicants. Total savings are 0.5 million, with the shares to the federal and state governments determined by the 1990 AFDC cost sharing formulas. The federal food stamp cost is $1.2 million, driven partially by less restrictive income limits for disabled individuals. Finally, SSI benefits to new successful applicants add $0.3 million to the costs of the policy change. The costs generated by behavioral responses to the regime change total $120.0 million.

Distributing Costs to Determine the Impact on Family Income

Once we have estimated the costs associated with the benefit increase, we distribute these to families in a way that attempts to mimic the current (as of 1990) financing of the relevant programs. Since the costs are based on the post-regime-change income, for probabilistic cases we distribute costs based on expected income. The costs associated with the increase in DI benefits paid are distributed to families in proportion to their share of earnings below the Social Security cap ($51,300 per individual in 1990). The indirect tax and transfer costs due to the labor supply responses of the marginal applicants are distributed to families based on their shares of the federal and state income tax payments to capture both their share of the tax bases and tax code progressivity.

5. Results and Discussion

Representative Workers' Willingness-to-Pay

Our calculations of total direct and indirect costs yield an estimate of the average degree of actuarial unfairness associated with the "purchase" of an additional dollar of DI benefits through reduced income for non-recipients. The ratio of total costs to the direct costs, which capture the net transfer to current DI recipients, is equal to 1.42. For every $1 of income transferred to current recipients, an additional 42 cents are incurred in

---

which represented only 5.3% of total AFDC caseloads in 1990.

indirect costs. Given the focus in the empirical literature on moral hazard, it is surprising that our estimates suggest that these distortions do not lead to a greater load factor.[28]

We use our estimate of the average implicit price to calculate whether individuals should be willing to pay for increased benefits. These calculations are presented in Table 5. For example, consider a typical working aged man with positive earnings, whose share of household earnings is 2/3 and average replacement rate is 0.5.[29] If we assume that when he becomes disabled, he loses his own earnings and they are replaced by DI benefits, the income drop suffered by the household when he moves to DI is 1/3. As discussed in the prior section, given a separable constant relative risk aversion utility function, the relative value of a dollar in the two states can be expressed by the ratio of income in the two states raised to the power of the negative relative risk aversion parameter. The relative value of $1 in the disabled state relative to the able-bodied state is $1.50, $2.25, and $5.07 for levels of risk aversion equal to 1, 2, and 4 respectively. The average working aged man would be willing to "buy" more insurance at a price of $1.42.

Consider the same calculation for the typical working aged woman who has positive wage and salary income. She contributes 1/3 of household income when working so that household income would fall by 1/6 if the replacement rate is also 1/2. The relative value of $1 across a state where income is only five-sixths as high as another is equal to $1.2, $1.4, and $2.1 for levels of relative risk aversion of 1, 2, and 4. Because the fall in household income is less severe, additional insurance only looks attractive at the individual level when risk aversion is at least moderately high.

Next, we consider a typical applicant. The typical applicant not only tends to earn less than the typical worker, but also tends to face higher replacement rates. Calculations using the SIPP imply that a typical male beneficiary will have household income of about 5/6 of the level experienced a year prior to the application for benefits. This matches the case for typical working aged women outlined above. For female applicants, household income appears to fall by about 12%. The relative value of $1 across the two states

[28] Adding the upper bound estimate of new Medicare costs due to successful applications would increase the estimated implicit price to $1.57.

[29] These numbers are based on the CPS and from published data on Social Security benefits and average

would be $1.14, $1.29, and $1.66 for our three levels of risk aversion. Again, the relatively small fall in income means that only relatively higher levels of risk aversion are associated with additional insurance being individually attractive.

Now, we consider the fact that different individuals must pay different prices for increased insurance given the redistribution within the program. Rather than relating overall direct costs to total costs, we consider how these costs are distributed across individuals with higher and lower levels of education. Table 6 shows the average share of costs borne by families according to years of education of the family head. The majority of costs are distributed in proportion to earnings under the payroll tax cap. The distribution of costs is highly progressive in terms of dollar amounts, as the least-educated families pay $1.44 in total costs (annually) and the most educated pay $5.57 on average through reduced take-home pay. However, the distribution of costs is much less progressive from the perspective of income shares, as shown in the fourth column in Table 6. Those with low levels of education, who are likely to have the highest marginal value of income, face income drops that are small in terms of dollars but large as a proportion of income.

Comparing these average costs to the expected benefits by level of education provides a crude sense of variation in actuarial unfairness. The expected increase in DI benefit receipts is much higher for less educated families due to a likelihood of receipt that is nearly eight times as great as for the most educated families. The benefit increase provides an additional dollar to these families at a price of only 26 cents. All other families pay more than one dollar, with the college-educated paying over six dollars. Applying these "prices" to the cases above suggests that only certain subgroups of workers would find increasing DI benefit generosity financially attractive.

Life-Cycle Incidence for Average Workers

In this section, we convert the changes in family income to changes in individual utility to evaluate the net welfare impact of the increase in DI benefits. Before converting the dollar values to utils, we first summarize the changes in family income in Table 7. The first three columns of Table 7 present average pre-reform income for individuals

---

earnings.

classified by family type. We identify three types of families (at any given point in time): families with no DI beneficiary, families with a working aged DI beneficiary, and families with an elderly DI beneficiary. The six rows are based on three concepts of income: family income, per capita income, and adjusted per capita income. Each of these three measures are presented unsmoothed and smoothed to account for the relationship between concurrent family income and consumption.

Baseline unadjusted family after-tax income is larger for individuals in families with no DI beneficiary ($30,773) than for those in families with either a worker beneficiary ($23,043) or an elderly beneficiary ($21,969). However, the gap narrows when family income is smoothed. The difference between families with and without beneficiaries becomes even smaller when income is expressed in per capita terms. In fact, the average income for individuals in families with an elderly DI beneficiary is larger than that for families with no limited individuals for three of the four per capita measures. The next three columns of Table 7 show the post-reform change in the various measures of income for individuals in the three types of families. As expected, individuals in families with no beneficiary face income losses while individuals in families with either a working aged or elderly beneficiary experience increases in income. The initial level of income and the change in income determine the measured impact on the financial component of utility for each individual.

Table 8 shows changes in social welfare for different levels of risk aversion and different measures of individual income[30]. The first column indicates the value of θ, the coefficient of relative risk aversion. The second column presents the average of the monetary component of baseline utility. The next three columns show the average change in utility for individuals in three different types of families – those without a DI beneficiary, those with a working age beneficiary, and those with an elderly beneficiary. The final column shows the average change in utility over all individuals. The sum of the changes over all individuals in a particular group represents the net welfare impact to that group, and these numbers are presented in square brackets immediately below the

[30] Tables 8 and 9 present results for per capita family income adjusted for economies of scale and the number of adults and children as described above. Results for pure per capita income do not differ qualitatively.

averages.

Recall that social welfare is calculated using a utility function that exhibits constant relative risk aversion of the form shown in equation (1). If individuals are risk neutral (θ equal to zero), utility is simply measured by individual income. Since our reform throws away 42 cents for every dollar transferred to beneficiaries, there must be a net loss in welfare under risk neutrality. What becomes clear from the results presented in Table 8 is that the policy reform leads to a net welfare loss for larger values of θ as well.

The pattern of net welfare impacts (as normalized by the relevant average pre-reform utility) across levels of risk aversion in Table 8 is surprising. The net welfare loss at first shrinks and then becomes larger as the level of risk aversion increases. It appears that a representative individual from behind the "veil of ignorance" would be less willing to pay for higher disability benefits when very risk averse. This is counterintuitive, since at higher levels of risk aversion the insurance motive should become more important and we would expect the policy reform to become more attractive.

For the welfare of a given representative individual to fall as the level of risk aversion rises, the individual must not want to transfer resources from their able-bodied to their disabled (and on DI) states of the world. This could be the case if there is a strong upward life-cycle earnings path, since DI transfers from younger to older years. An alternative explanation would be prolonged negative earnings shocks that are not associated with disability, since DI benefits are a function of average earnings and could be higher than earnings during that phase. Evidence for these effects would be replacement rates that exceed one. Using the Health and Retirement Survey data for those who were on DI as of 1991, we calculated the ratio of average earnings between the year the individual turned 25 and the year they applied for DI to average individual benefits since application.[31] While this ratio exceeds one for only two percent of men in the sample, one in five women have replacement rates above one. This is evidence that at least some individuals are over-insured by the DI program.

Another explanation for the counterintuitive program is that the DI program

[31] Financial variables were adjusted for inflation using the CPI for all goods.

combines insurance with substantial redistribution. The program redistributes from low, moderate, and high income populations toward a moderate income population. At the bottom end of the distribution, non-beneficiaries have lower baseline income than beneficiaries since they are not guaranteed an income floor like the one provided by DI benefits. As the level of risk aversion increases, this population is weighted disproportionately in the welfare calculation.

Table 9 shows how the net welfare impact varies for more and less well-educated individuals at different levels of risk aversion. The entries are comparable to the numbers in square brackets in the last column of Table 8. They represent the aggregate change in utility, normalized by average baseline utility, associated with the reform. These results again stress the redistributive component of the DI program. For all individuals with education of a high school diploma or greater, the policy change causes a loss in net welfare. For individuals with less than a high school diploma, there is a net welfare gain for all but the highest levels of risk aversion. Given the favorable implicit price of insurance found for this group in the last section, this finding merits explanation. While the price is favorable on average for individuals with low education due to low moral hazard costs and redistribution from the more highly educated, the gain turns negative at high levels of risk aversion due to redistribution from low to modest income individuals within the low education sample.

Incorporating Health Insurance

We incorporate health insurance first on the cost side. Marginal successful applicants who become eligible for Medicare after two years on the program may generate additional costs. Under the assumption that each receives the average level of Medicare benefits, marginal applicants would generate an additional $42.6 million in indirect costs. This would raise the implicit price of an additional dollar of insurance by 10%, from $1.42 to $1.57. Under the more moderate assumption that only those who were not previously covered by public insurance generate marginal costs, the price would still increase to $1.53.

To conduct the welfare analysis, we distribute these costs across the population according to their share of earnings subject to the Medicare tax. Though the marginal applicants do not enter into our welfare calculation because of our indifference assumption, they implicitly place some value on Medicare coverage. To determine the appropriate baseline utility for all of the other families in the sample, we treat insurance status as merely exposing individuals to different distributions of out-of-pocket medical expenditures. For families by age of the heads and type of insurance coverage, we randomly reduce after-tax income by a draw from the appropriate distribution as estimated by the CEX.

The method we use to calculate the empirical distributions of out-of-pocket health care expenditures is described in Appendix C. As the figures in Appendix C show, health insurance status in the cross-section is largely a proxy for heterogeneity in health. Out-of-pocket expenditures are consistently highest for the group covered by Medicare, which for younger households consists of disabled individuals. Expenditures are actually lowest for the uninsured. Therefore, our adjustment will tend to make DI families appear relatively worse off in the baseline and would be expected to make the reform more attractive. Table 10 presents the sign of the net welfare impact when incorporating health insurance, and confirms that this is the case. Overall, and for the two lowest educational categories, the policy reform now becomes welfare-enhancing at high levels of risk aversion.

5. Conclusion

This paper has developed two basic insights into the desirability of increasing DI benefits. First, indirect costs appear smaller than the emphasis in the prior empirical literature would lead one to believe. We estimate a ratio of total costs to direct costs of 1.42, suggesting that the implicit degree of actuarial unfairness due to moral hazard is not very substantial. Second, while the moral hazard costs appear to be relatively low, the implicit price of additional insurance can be quite large for some subgroups of the population due to redistribution. The implicit price of an additional dollar varies from less than $1 for individuals with less than a high school degree to more than $6 for individuals with a college degree. While the importance of redistribution across broad education categories has not been previously quantified, this pattern in implicit prices is not surprising given differences in the incidence of disability. What is less acknowledged is that an increase in benefits leads to transfers from low income individuals to individuals with more moderate income, such that these transfers lead to net welfare losses at high levels of risk aversion. These results highlight that DI treats the certifiably disabled as more "deserving" than able-bodied individuals who may be financially worse off. From an individual insurance perspective, a marginal increase in DI benefits is less attractive because of this potential redistribution to bad states of the world from worse states of the world.

Appendix A: The Impact of Applying to DI

In order to identify the impact of an application for Social Security Disability Insurance (DI) benefits, we estimate a series of regressions aimed at tracking changes in key variables around the time of initial application. These variables include the applicant's earnings and labor force participation, the applicant's spouse's earnings and labor force participation, and transfer income other than Food Stamps, AFDC, or SSI.[1] The data and methods used in these regressions are discussed in next two sections below. In the third section, we describe the strategy we use to estimate Food Stamps take-up rates.

Data

The Surveys of Income and Program Participation (SIPP) are a series of United States Census panel surveys of representative populations of the United States. New panels were fielded in 1990, 1991, 1992, and 1993.[2] For the 1990 SIPP panel, the Census interviewed a new rotation group each month for four months starting in February, 1990. Each of the four rotation groups is a random sample of the United States population. These groups were interviewed eight times at four-month intervals. Each interview contains monthly information for the preceding four months. Hence, the 1990 SIPP panel provides monthly data for up to 32 months on each individual covering a 35-month period from October 1989 through August 1992. The 1991 SIPP panel was fielded using the same panel design, and in 1992 and 1993 a similar design including an additional ninth wave of interviews was used to provide a total of 36 months of data.

We match the 1990-1993 SIPP panels to the Social Security disability determination records. The records are from Disability Determination Service (DDS) and Administrative Law Judge (ALJ) stages of the determination process for those who applied for DI benefits. We have information for individuals whose applications were acted upon between 1986 and 1994 for the 1990 and 1991 panels and between 1977 and 1997 for the 1992 and 1993 panels.[3] We use only those administrative records that begin with an initial consideration or reconsideration by the DDS and contain valid date of birth, filing date, and decision date information. The matching procedure produced a total of 5,594 SIPP respondents who are identified as having applied for DI, with the bulk of the applications occurring during the late 1980s and 1990s.

In what follows, we will consider only the first application date observed in the administrative records for each respondent, and we will consider only the final decision observed for that application. Thus, a respondent who is initially denied benefits but is awarded benefits at a later stage of the same application process will be referred to as having been awarded benefits. A respondent who is denied benefits, then later re-applies and is awarded benefits as a result of the later application, will be referred to as having been denied benefits.

[1] We model changes in transfers from these three programs using the benefit algorithms for 1990. We assume full take-up for SSI and AFDC and apply our estimates of Food Stamps program take-up post-application.

[2] New SIPP panels were fielded in earlier and later years but not all of these data have been matched to Social Security Administration administrative records.

[3] These data are essentially the same as those used by Lahiri, Vaughan and Wixon (1995) and Hu, Lahiri, Vaughan and Wixon (1997) in their work studying applications to the DI and SSI programs.

In our analysis, we merge our four SIPP panels but do not do so along a calendar time dimension. Instead, we focus on an event—initial application month and year for DI benefits and array our respondent panel data by individual from the months prior to application $(t-k)$ through the months following applications $(t+k)$ where t is the month of application. For those who applied for benefits prior to, or in the early waves of, the SIPP, we have information on their household income following application. For those who applied for benefits in a middle wave, we have information on their household income in the months just prior to application and just following application. For those who applied for benefits in the later waves of the SIPP or just afterwards, we have information on their household income in the months prior to application. Using this approach, we are able to obtain snapshots of respondents' average household income in the months and years prior to and following their application for DI benefits that extend beyond the maximum of 36 months that any one respondent is followed in a given SIPP panel.

We include a control group of SIPP respondents who report having experienced a work limitation, but whom we do not identify as having applied for federal disability benefits (either DI or SSI) at any time. To ensure that we compare applicants to non-applicants that have similar demographic characteristics, we assign weights to non-applicants that correspond to the proportion of the DI applicant population within the relevant age, race, sex, martial status, and educational attainment cell. There are 288 separate cells in the underlying matrix, and the same number of applicants and non-applicants are assigned to each cell when DI applicants are weighted using the given population weights and non-applicants are weighted using the newly defined weights.

Methods

We set up a regression model that allows us to estimate how a variety of income measures and indicators are affected by application. We include fixed effects for each individual, i, and for each calendar month, t, in the sample period. A set of indicator variables for months relative to application month is also included. These variables (month zero, for example, representing the month of application, or month twelve, to take another example, representing the month one year subsequent to application) are equal to one if observation of income data is taken from the month corresponding to the indicator variable, and zero otherwise. Thus the equations have the form:

$$y_{it} = \alpha_i + \gamma_t + \sum_j \beta_j m_{ij} + \varepsilon_{ij} \ , \qquad (1^A)$$

where y_{it} represents monthly income (in January 1990 dollars) or an indicator for positive income for individual i in time period t, depending on the case. The α's and the γ's represent applicant and calendar month fixed effects and the m_{ij}'s represent the indicator variables for each month relative to the application month. There is a separate indicator variable for each month representing data from 35 months before to 39 months after application. Observations from 39 or more months prior to application or 40 or more months subsequent to application are excluded from the regression sample. Therefore, the coefficients on application month indicator variables (the β's) can be interpreted as the average change from the level observed 36 to 38 months prior to application.

Graphs in the figures that follow are constructed using estimated coefficients from regressions with the format given by equation (1^A). For purposes of presentation, we graph the β's relative to average baseline incomes (technically, relative to the average α for applicants). Thus, the graphs show average change in income from a particular source relative to baseline. Since a control group is included in the regressions (and application month dummies will be identically zero for all individuals in the control group), the calendar month fixed effects remove time trends resulting from age or economy-wide trends.

For regressions in which applicant's earnings is the dependent variable, we exclude any month in which an applicant reported zero earnings. The graphs in Figures A1 can therefore be thought of as tracking average earnings conditional on working. These graphs exhibit no clear trend, whereas the labor force participation graphs in Figures A2 do follow a clear pattern. We estimated other regressions, not reported here, of applicants' earnings that do not exclude reported zero earnings, and therefore track average earnings not conditioned on work. The regression results for unconditioned earnings mimic the patterns of labor force participation, with a clear drop at the time of application and a partial recovery for those applicants who are eventually denied benefits.

Take-up Rates for Food Stamps

To estimate Food Stamps take-up rates, we restrict our sample to those data reported for the calendar year 1990, and estimate Food Stamp eligibility based on reported income month by month.[4] We then calculate the fraction of respondents we determine to be eligible who are receiving a positive benefit, and report this fraction as the take-up rate. We estimate the fraction separately for individuals we observe more than a year prior to application (i.e. those who applied in 1991 or later) and those we observe more than a year after application (i.e. those who applied in 1989 or earlier). The Food Stamp take-up rates are reported separately by decision type, either awarded or denied, since these populations are likely to be different even before application, but certainly can be expected to exhibit different patterns afterward.

Table A1: Food Stamps Take-up More Than One Year Before and After an Application for DI Benefits

	Take-up rate	Std. deviation	Observations
Awarded Applicants:			
Before, weighted	.634	.490	29
After, weighted	.535	.511	22
Denied Applicants:			
Before, weighted	.512	.504	63
After, weighted	.709	.459	50

For individuals who are induced to apply for DI and are successful, we apply a take-up rate of 53.5%. For those who are denied, we apply a rate of 70.9%. For both groups, we use actual take-up observed in the CPS for our pre-application baseline.

[4] We calculated AFDC take-up rates from the SIPP in a similar manner, but chose not to use those rates due to the extremely small sample size involved. Though we use the take-up rates observed in the CPS to establish baseline participation for marginal applicants, we assume 100 percent take-up of AFDC among eligibles after application.

Figures A1 a-f: Applicant's Earnings (Conditional on Reporting Positive Earnings)

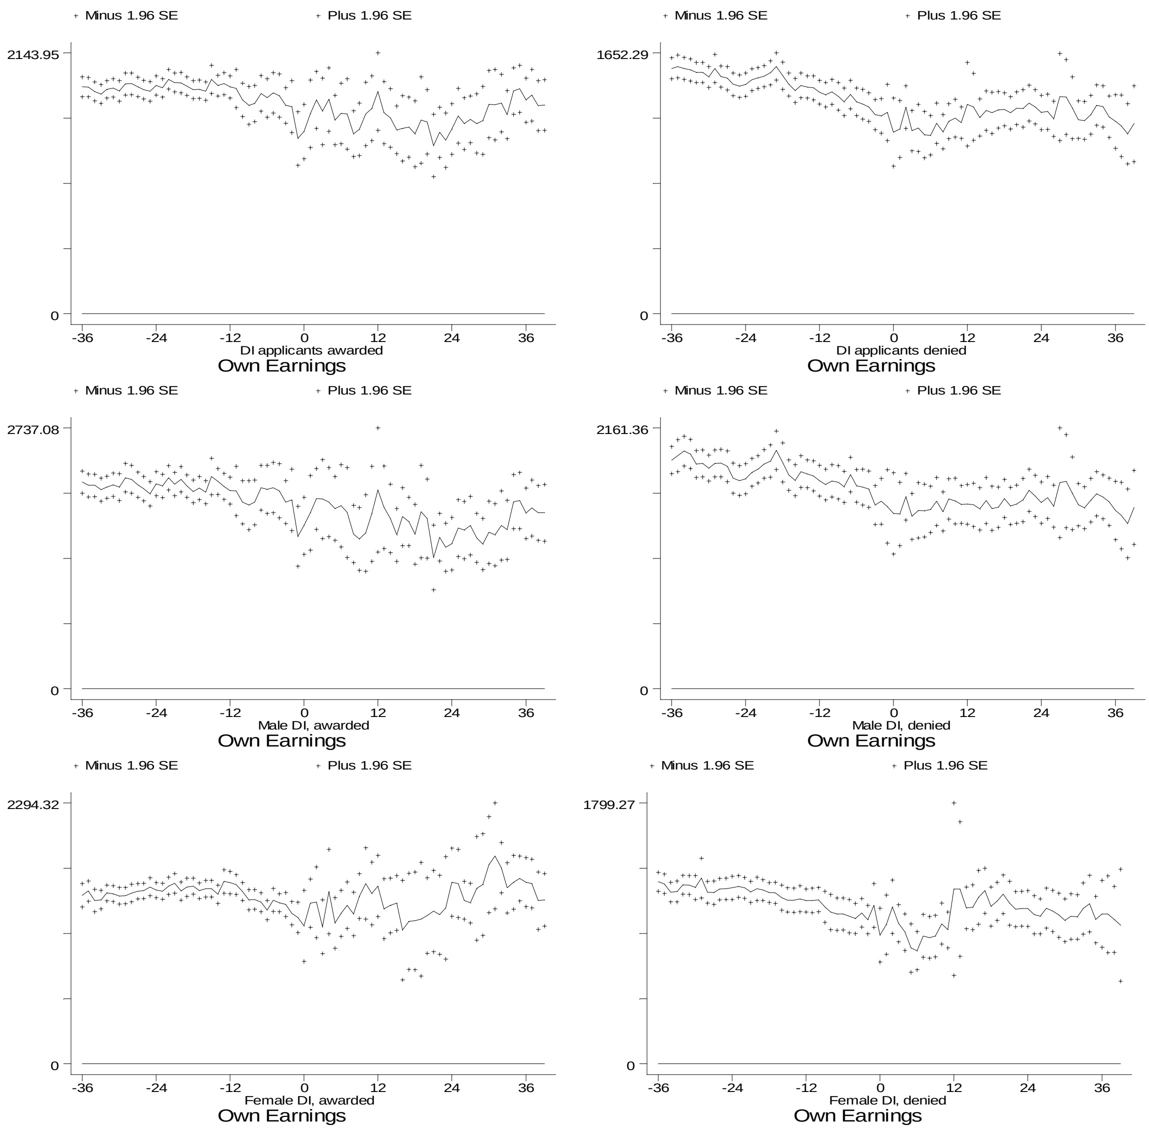

Figures A2 a-f: Applicant's Labor Force Participation

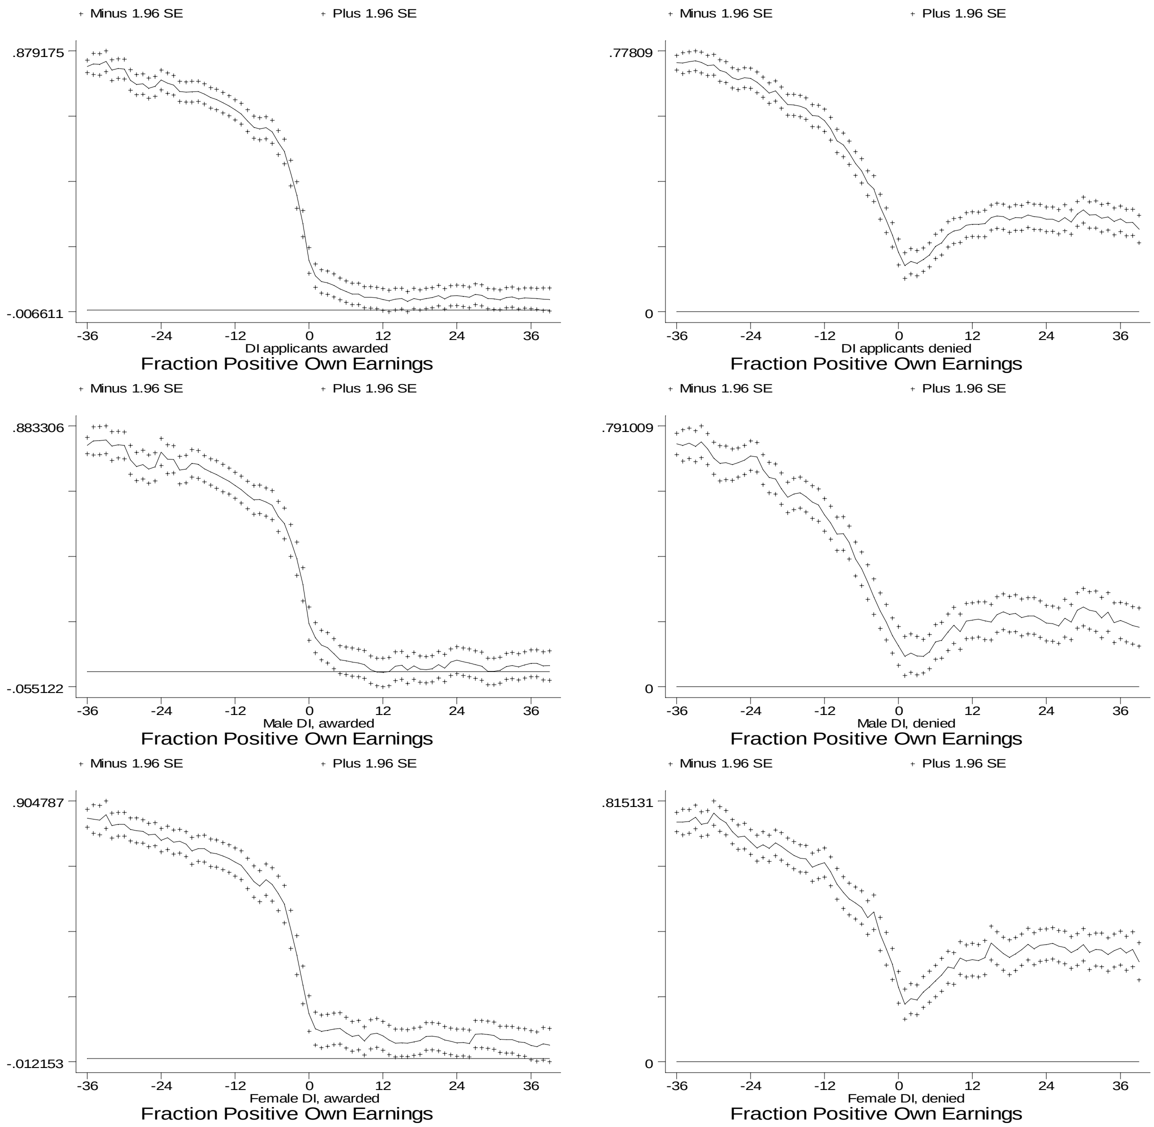

Figures A3 a-f: Applicant's Spouse's Earnings

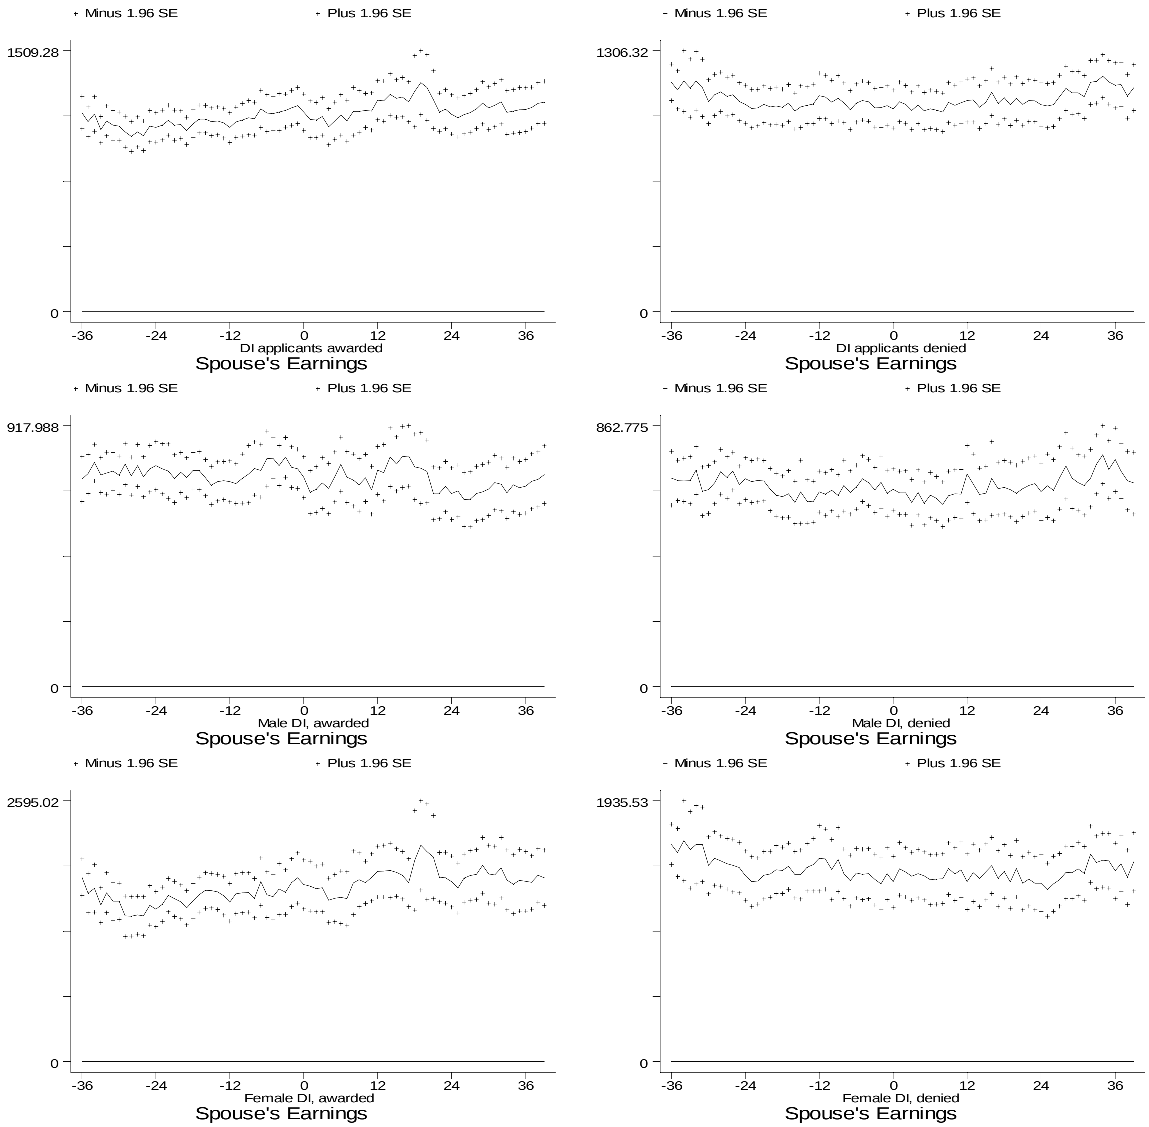

Figures A4 a-f: Applicant's Spouse's Labor Force Participation

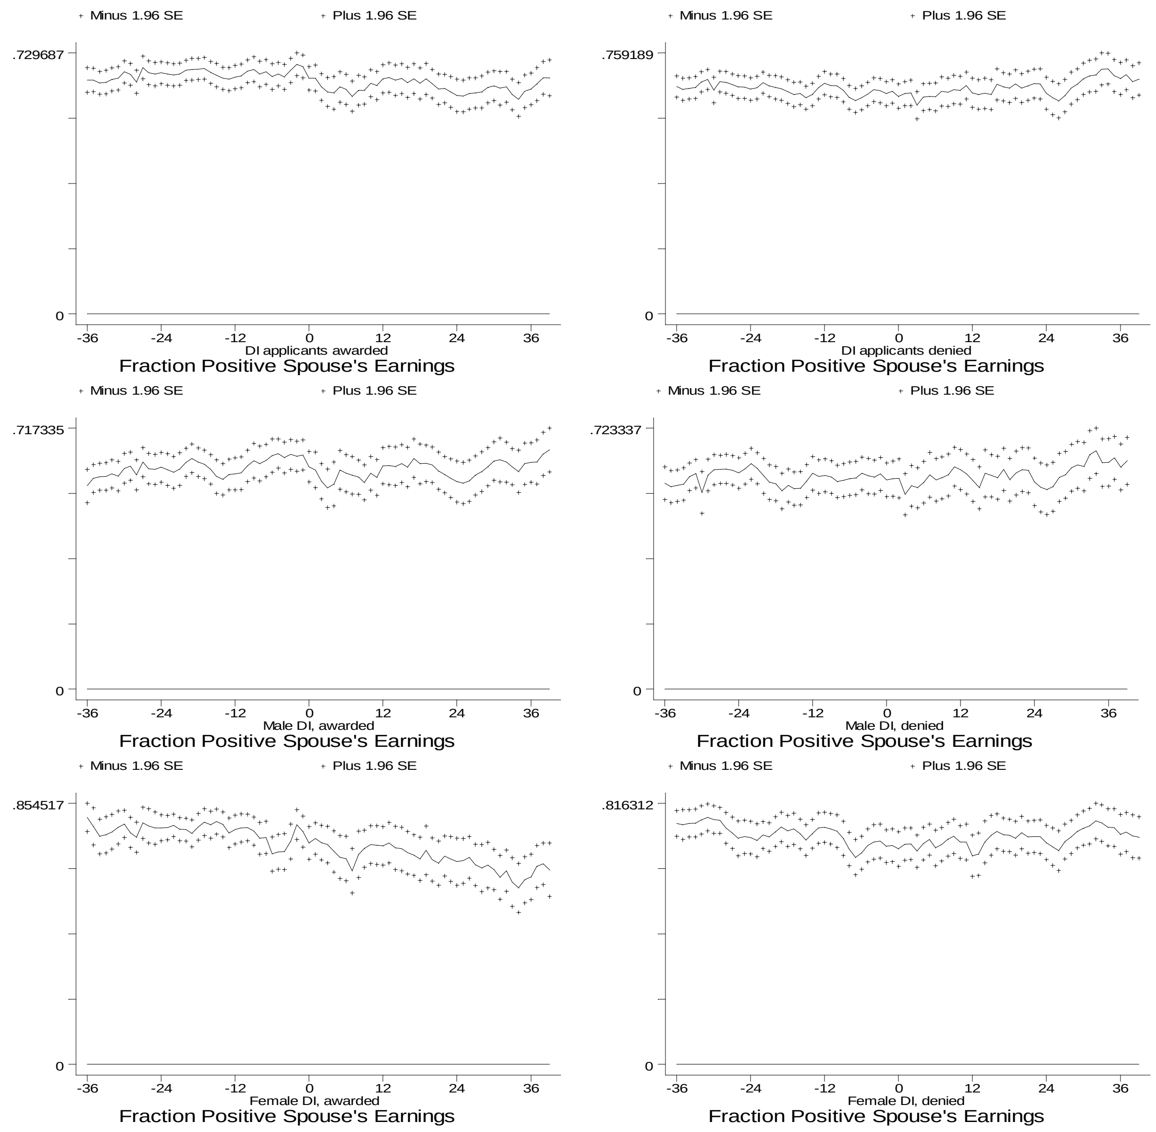

Figures A5 a-f: Applicant's Transfer Income Excluding AFDC, Food Stamps, and SSI

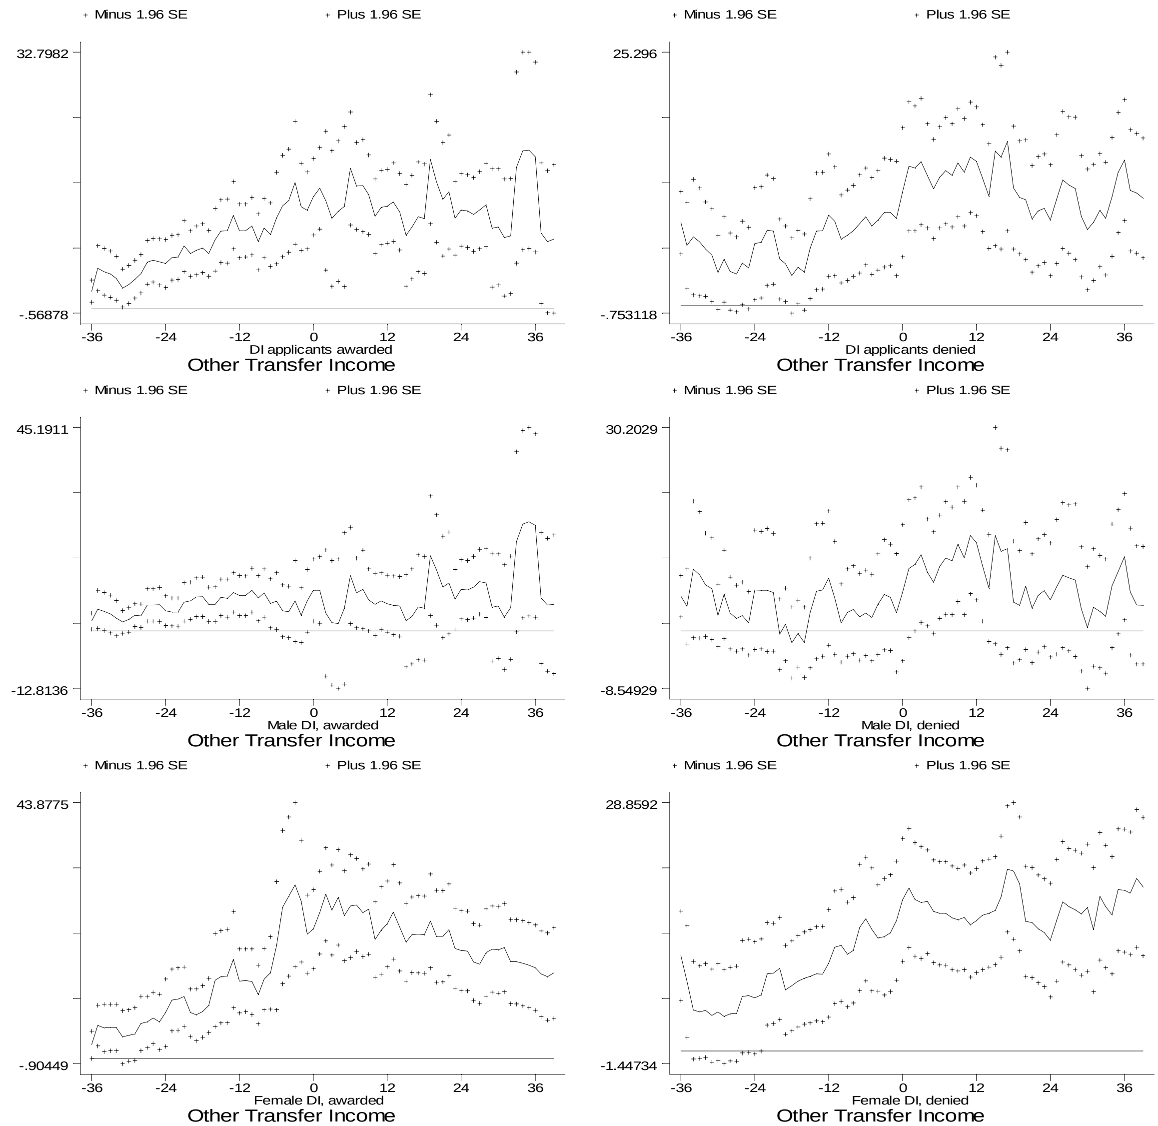

Appendix B: Predicting Household Consumption

We would like to account for the fact that variation in the distribution of income overstates variation in consumption since households can use savings and other means to smooth transitory shocks to income. Therefore, we estimate the relationship between consumption and income from the Bureau of Labor Statistics' Consumer Expenditure Survey (CEX). We then use this relationship to adjust family income in the CPS.

Data

The CEX interviews approximately 5,000 households every quarter on a rotating basis. Households are interviewed for four consecutive quarters and then dropped from the sample. Each quarter, households provide detailed reports of expenditure by category for the prior three months. Information on income for the prior year, which coincides with the timing of annual consumption, is gathered in the final interview.

We use the family-level extracts compiled by the Harris and Sabelhaus (2000). In order to provide a ten-year window around 1990, we include households interviewed from the first quarter of 1985 through the last quarter of 1995. We restrict the sample to the subset of 40,848 households that provide complete income reports and remain in the sample for a full year. Weights are provided that adjust the national CEX weights to account for attrition by age and housing tenure.

Our pre-tax income measure is calculated to match after-tax income in the CPS as closely as possible. It includes all components of earned and unearned income (including Food Stamps benefits) except for payments to employees in-kind and insurance refunds, which are not available in the CPS. For the same reason, we do not subtract alimony or child support payments made by the household. After-tax income is pre-tax income less reported federal and state income taxes paid.

Our total consumption measure is comprehensive and includes expenditures on the following categories of goods and services: food, tobacco, alcohol, clothing, personal care, medical, recreation, education, household, utilities, personal business, vehicles, transportation, and non-mortgage interest payments. We include rent paid for renters and home rental equivalent value for homeowners (and exclude home maintenance and financial costs).

Both income and consumption are converted to 1990 dollars using the CPI for all goods. Mean annual after-tax income in the CEX sample is $29,712 (with a standard error of $118), which is somewhat higher than the mean for our CPS sample ($25,367 ($75)). This difference can at least be partly explained by the fact that the CE definition of household (individuals who share major expenses) is more expansive that the CPS definition. Mean annual consumption in the CEX sample is $25,715 ($81).

Predicting Consumption

Our specification relates the natural logarithm of consumption to the natural logarithm of income and a constant term. We replace income to $1 if income is less than $1 and also include an indicator for these low-income households. We allow for heterogeneity in the relationship between consumption and income by age and disability status. We split households into three categories according to age of the household head:

less than 45 years, 45-61 years, and 62 years and over. For the two non-elderly groups, we split the sample further into households with and without a disabled working-aged individual.[5]

We began with a specification that allowed for a full set of interactions between the other control variables and indicators each demographic group. We could not reject the hypothesis that the slopes were the same across these five groups. While the coefficients on the low income indicator were significantly different from one another for the groups that had low income individuals, not all of our groups were represented. Therefore, the specification that we implement constrains the coefficient on log income and the coefficient on the low income indicator to be the same across groups, and includes separate constant terms for each group.

The ordinary least squares results from this specification are shown below:

Table B1: Predicting Annual Consumption

Independent Variable	Sample Share	Coefficient	Standard Error
Constant		5.123	0.056
Ln(annual after-tax income)		0.484	0.005
Indicator for income less than $1		5.352	0.079
Head under 45 years of age			
No working-aged disabled in CU	51.8%	-	-
Working-aged disabled member	2.5%	-0.110	0.019
Head between 45 and 61 years of age			
No working-aged disabled in CU	20.8%	0.073	0.006
Working-aged disabled member	2.4%	-0.058	0.018
Head aged 62+	23.3%	-0.010	0.007

The dependent variable is log consumption and the omitted demographic category is household head under 45 with no working-aged disabled members. The regression is weighted using the adjusted weight. The number of observations is 40,848 and the adjusted R-squared is 0.499.

The coefficient on log income is surprisingly low, a robust finding that is not an artifact of pooling the sample. There is certainly a great deal of measurement error in income in the CEX that would lead to a downward bias in the estimated income elasticity. As long as measurement error in the CPS is similar to that in the CEX, then this prediction is still useful for our purposes. It is important to keep in mind, however, that applying these estimates to predict consumption in the CPS likely leads to an overly narrow distribution. One other finding that is worth noting is that disabled households apparently consume less than non-disabled households with the same income.

In order to predict consumption for our CPS sample, we apply the estimated coefficients shown in Table B1. However, we first make one adjustment. Since we conduct welfare analyses both for income and for consumption ("smoothed" income), we want to ensure that any differences in the inferences are not driven by differences in mean levels. We, therefore, redefine the constant term to equate mean income and mean predicted consumption.

[5] Households with working-aged disabled members are identified as households that either report SSI receipts or have no member over 61 years of age and report receiving Social Security income.

Appendix C: Estimating Out-of-pocket Medical Expenditures

We model the role of health insurance in our welfare calculations through its impact on the distribution out-of-pocket expenditures. In this section, we describe how we estimate this distribution for different demographic groups by insurance status.

Data

We rely on the same CEX extracts that are described in Appendix B. We combine these data with additional data on health insurance status from the detailed expenditure files on hospitalization and health insurance included in the CEX available since the first quarter of 1988. From these files, we are able to identify whether any members in the household are covered by private insurance, Medicare, Medicaid, and/or another program that provides free health care (e.g. CHAMPUS or military health care). To match these quarterly data on insurance status to the annual income files, we identify a household as covered by one of these plans if any member was covered at any point during the year. Restricting the sample to households that are complete income reporters and remain in the sample for a full year, the analysis sample includes 28,718 households.

Out-of-pocket Health Care Expenditures

We use the empirical distributions of out-of-pocket health care expenditures in the CEX to randomly reduce income for households in the CPS. We estimate the distribution separately for households with heads under 45 years, between 45 and 61 years, and over 62 years of age. We further distinguish among households within these groups by four mutually exclusive classes of insurance coverage: no insurance (12.2%), private insurance only (52.6%), Medicare (22.8%), and Medicaid or any other free care program (12.4%). The Medicare category may include individuals covered by supplementary private insurance plans. The Medicaid category includes any household covered by Medicaid, including households that are also covered by Medicare or private insurance.

Health care expenditures include expenditures on medical goods and services, as well as payments for health insurance premiums and copayments. We calculate the income share of health care expenditures by taking the ratio of expenditures to after-tax income, as calculated in Appendix B. For each of our three age groups by the four insurance statuses, we then identify the percentiles of the distribution of the income share (for every two percentiles from the 2^{nd} to the 98^{th}) for each of our three demographic groups by the four insurance statuses. Each household in the CPS sample is then randomly assigned the out-of-pocket income share associated with one of the percentiles of the distribution that is appropriate to them. For the welfare calculations, we then subtract from (or add to) household income the difference between the overall average health care expenditure share and the assigned health care expenditure share. The empirical distributions from the CEX are shown in the following figures.

Figure C1: Income Share of Out-of-pocket Health Care Expenditures for Households with Head Under 45 Years of Age

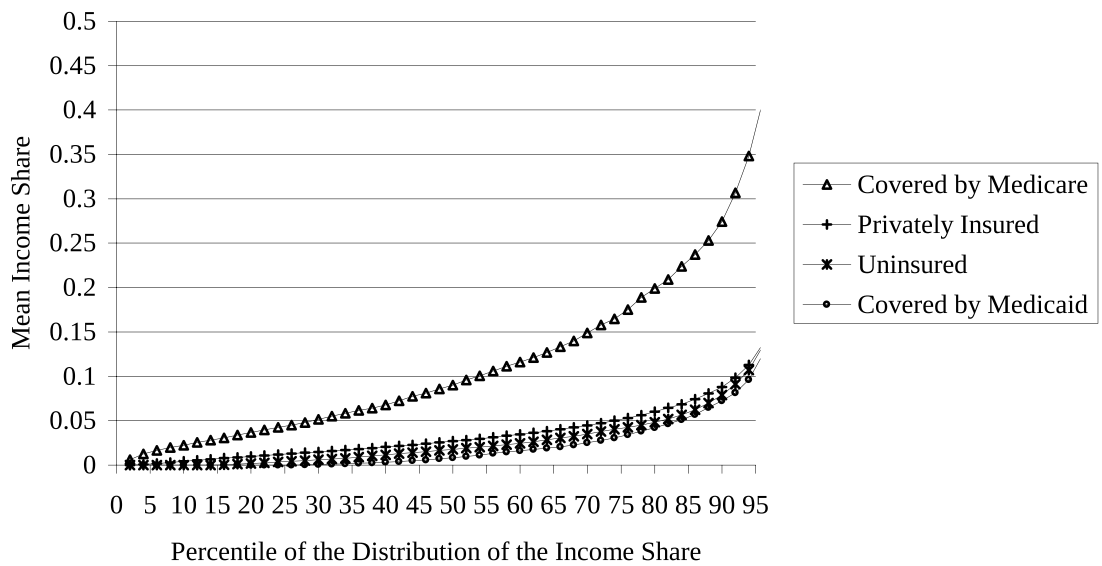

Figure C2: Income Share of Out-of-pocket Health Care Expenditures for Households with Head 45-61 Years of Age

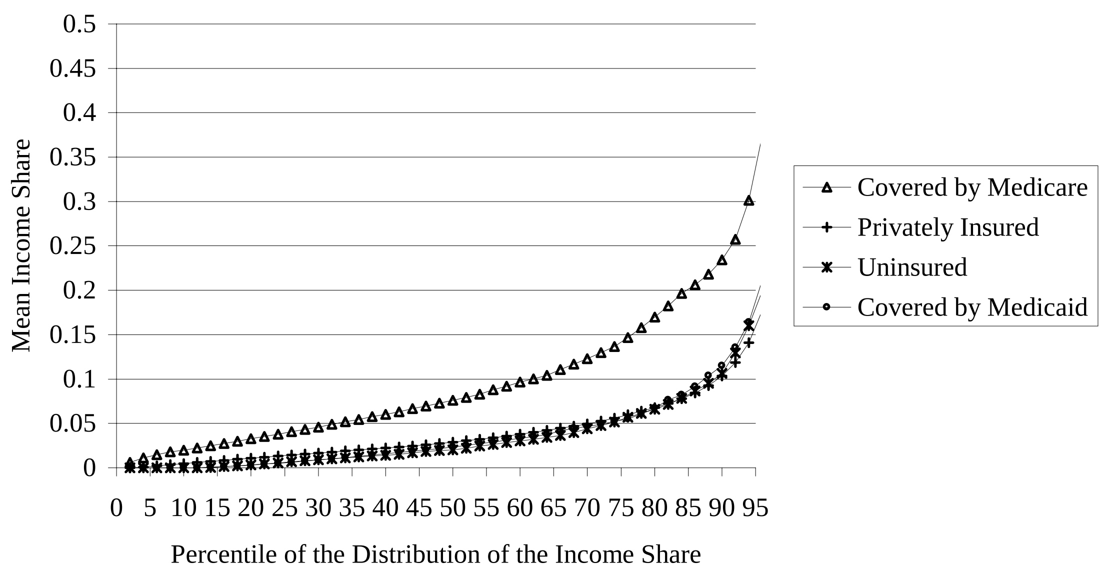

Figure C3: Income Share of Out-of-pocket Health Care Expenditures for Households with Head Over 61 Years of Age

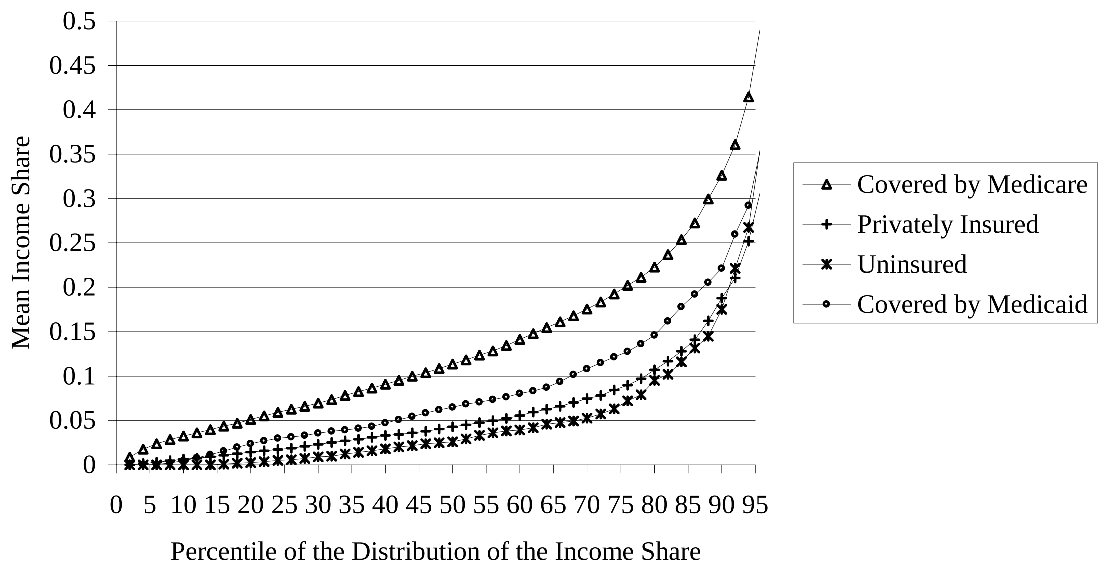

References

Anderson, Patricia M. and Bruce D. Meyer, 1995, "The Incidence of a Firm-Varying Payroll Tax: The Case of Unemployment Insurance," NBER Working Paper No. 5201 (Cambridge, MA: National Bureau of Economic Research).

Barsky, Robert B., Thomas F. Juster, Miles S. Kimball and Matthew D. Shapiro, 1997, "Preference Parameters and Behavioral Heterogeneity: An Experimental Approach in the Health and Retirement Survey," *Quarterly Journal of Economics*, 112(2): 537-79.

Bound, John, 1989, "The Health and Earnings of Disability Insurance Applicants," *American Economic Review*, LXXIX: 482-503.

Bound, John and Richard Burkhauser, 1999, "Economic Analysis of Transfer Programs Targeted on People with Disabilities" *Handbook of Labor Economics, Vol. 4*, Ashenfelter and Card, editors.

Bound, John and Timothy Waidmann, 1992, "Disability Transfers, Self-Reported Health, and the Labor Force Attachment of Older Men: Evidence from the Historical Record," *Quarterly Journal of Economics*, 107(4): 1393-1419.

Burkhauser, Richard V., Robert H. Haveman, and Barbara L. Wolfe, 1993, "How People with Disabilities Fare when Public Policies Change," *Journal of Policy Analysis and Management*, 12(2) (Spring): 251-269.

Cullen, Julie Berry and Jonathan Gruber, 2000, "Does Unemployment Insurance Crowd Out Spousal Labor Supply?," *The Journal of Labor Economics*, 18(3) (July): 546-72.

Diamond, Peter and Eytan Sheshinski, 1995, "Economic Aspects of Optimal Disability Benefits," *Journal of Public Economics*, 57(1) (May): 1-23.

Feenberg, Daniel Richard and Elizabeth Coutts, 1993, "An Introduction to the TAXSIM Model," *Journal of Policy Analysis and Management*, 12(1) (Winter): 189-194.

Feldstein, Martin, 1999, "Tax Avoidance and the Deadweight Loss of the Income Tax," *Review of Economics and Statistics* 81(4): 674-80.

Gallicchio, Sal and Barry Bye, 1980, "Consistency of Initial Disability Decisions among and within States," Staff Paper No. 39, SSA Publication No. 13-11869, Office of Research and Statistics. Washington, DC: Department of Health and Human Services, Social Security Administration.

Gruber, Jonathan, 1994, "The Incidence of Mandated Maternity Benefits," *American Economic Review*, 84(3) (June): 622-641.

Halpern, Janice H., 1979, "The Social Security Disability Insurance Program: Reasons for Its Growth and Prospects for the Future," *New England Economic Review* (May/June): 30-48.

Harris, Ed and John Sabelhaus, 2000, "Consumer Expenditure Survey Family-Level Extracts, 1980:1 -- 1998:2," Congressional Budget Office, (September).

Haveman, Robert H. and Barbara L. Wolfe, 1990, "The Economic Well-Being of the Disabled, 1962-1984," *Journal of Human Resources*, 25(1): 32-55.

Hu, Lahiri, Vaughan, and Wixon, 1997, "A Structural Model of Social Security's Disability Determination Process," SSA Office of Research, Evaluation, and Statistics, Working Paper #72.

Killingsworth, Mark, 1983, *Labor Supply* (Cambridge and New York: Cambridge University Press).

Lahiri, Vaughan, and Wixon, 1995, "Modeling SSA's Sequential Disability Determination Process Using Matched SIPP Data," *Social Security Bulletin*, 58(4) (Winter): 3-42.

Lando, Mordechai E., Malcolm B. Coate and Ruth Kraus, 1979, “Disability Benefit Applications and the Economy,” *Social Security Bulletin*, 42: 3-10.

Nagi, Saad, 1969, *Disability and Rehabilitation: Legal, Clinical and Self-Concepts of Measurement* (Columbus, Ohio: Ohio State University Press).

Parsons, Donald O., 1991, “Self-screening in targeted public transfer programs,” *Journal of Political Economy*, 99: 859-76.

Stapleton, David C., Kevin Coleman, Kimberly Dietrich, and Gina Livemore, 1998, “Empirical Analysis of DI and SSI Application and Award Growth,” in Kalman Rupp and David C. Stapleton, eds., Growth in Disability Benefits (Kalamazoo, MI: W.E. Upjohn Institute).

Waidmann, Timothy, John Bound, and Michael Schoenbaum, 1995, "The Illusion of Failure: Trends in the Self-Reported Health of the U.S. Elderly," NBER Working Paper 5017.

Table 1: Summary Statistics: Working-Aged Population (21-61 years)

	Males			Females		
	Non-Limited Individuals	DI Beneficiaries	Potential Applicants	Non-Limited Individuals	DI Beneficiaries	Potential Applicants
Black	.105	.174	.160	.122	.225	.151
American Indian	.005	.010	.010	.006	.015	.010
Asian	.028	.010	.018	.031	.013	.019
Hispanic	.086	.061	.095	.082	.058	.068
Age	37.9	46.3	43.0	38.2	48.4	42.5
	(10.9)	(10.9)	(11.2)	(11.0)	(10.7)	(11.0)
Years of Education	13.3	10.4	11.9	13.1	10.6	12.3
	(2.9)	(3.9)	(3.1)	(2.7)	(3.8)	(2.9)
Married	.631	.465	.566	.643	.332	.514
Never Married	.257	.340	.237	.186	.223	.161
Number of Children	0.77	0.33	0.63	0.91	0.32	0.68
	(1.12)	(0.88)	(1.09)	(1.17)	(0.76)	(1.08)
Individual DI Benefits	-	7973*	7383†	-	6254*	5708†
		(3854)	(1521)		(3088)	(1003)
Family DI Benefits	-	9016†	8626†	-	7448†	6788†
		(5097)	(2622)		(4640)	(1713)
Probability Induced to Apply	-	-	.005†	-	-	.004†
			(.002)			(.002)
Number of observations	38561	816	1776	41946	618	1651

Notes: The mean and standard deviation are shown for key characteristics of the working-aged population, broken down by gender and DI status. The second column reports statistics for individuals who are imputed to be current DI worker beneficiaries. The third column reports statistics for individuals who report a limitation but are not currently receiving Social Security income, individuals who we assume are within the pool of potential applicants. The first column presents statistics for all other working-aged individuals. *The reported DI benefits have been scaled up by 25% to account for underreporting in the CPS. †These are predicted values, using the methods described in the text.

Table 2: Predicting DI Participation and Benefits

	Predicting the Likelihood Elderly Individuals are DI Beneficiaries		Predicting the Likelihood a Limited Working-Aged Individual is a DI Recipient		Predicting Individual DI Benefits for Working-Aged Applicants	
	Males	Females	Males	Females	Males	Females
	(1)	(2)	(3)	(4)	(5)	(6)
Constant	-4.62	-6.55	-3.48	-1.84	8.35	7.08
	(1.08)	(1.11)	(0.83)	(0.38)	(0.15)	(0.55)
Black	--	.617	-.246	--	-.116	--
		(.167)	(.152)		(.058)	
American Indian	.604	1.31	--	--	--	--
	(.414)	(0.47)				
Asian	-1.84	--	-.751	--	.390	--
	(0.80)		(.465)		(.112)	
Hispanic	-.853	-.553	-.701	-.475	--	--
	(.232)	(.262)	(.159)	(.192)		
Age	.062	.081	.128	.041	.006	.051
	(.018)	(.018)	(.038)	(.006)	(.002)	(.023)
Age Squared (coefficient × 100)	--	--	-.110	--	--	-.054
			(.040)			(.025)
Years of Education	-.071	-.010	-.080	-.087	-.014	-.012
	(.067)	(.070)	(.014)	(.015)	(.016)	(.025)
Years of Education Squared	-.007	-.008	--	--	.002	.003
	(.003)	(.004)			(.001)	(.001)
Married	--	--	--	-.757	--	-.077
				(.122)		(.059)
Never Married	.547	.542	.570	.267	-.311	-.117
	(.204)	(.229)	(.140)	(.151)	(.059)	(.080)
Number of Children Under 18	--	--	-.145	-.157	--	--
			(.066)	(.070)		
Number of Observations	7990	8741	2872	2871	816	618

Notes: The first two columns are based on separate logistic regressions for males and females aged 51-61 in our March 1991 CPS sample. In both cases, the dependent variable is equal to 1 if the individual is a worker beneficiary. The next two columns present results from separate logistic regressions for males and females aged 21-61 where the dependent variable is equal to 1 if the individual is a worker beneficiary. Columns 5 and 6 report the results form ordinary least squares regressions for male and female DI beneficiaries aged 21-6. The dependent variable is the logarithm of reported individual DI benefits. Each observation is weighted by the CPS person weight and the standard errors are corrected using White's generalized method. Observable characteristics that did not enter significantly were dropped from the regressions since our goal is to obtain the best prediction.

Table 3: Summary Statistics: Elderly Population (62+ years)

	Males			Females		
	Non-DI Beneficiaries	DI Beneficiary, Own Benefits	DI Beneficiary, Spouse Benefits	Non-DI Beneficiaries	DI Beneficiary, Own Benefits	DI Beneficiary, Spouse Benefits
Black	.089	.114	.013	.071	.181	.197
American Indian	.005	.008	.024	.004	.021	.013
Asian	.032	.001	0	.021	.016	.014
Hispanic	.047	.031	.009	.050	.037	.029
Age	69.8 (6.8)	71.1 (6.9)	71.8 (7.7)	69.4 (6.4)	70.4 (6.4)	72.9 (7.6)
Years of Education	12.8 (3.7)	8.6 (3.6)	9.7 (3.4)	11.5 (3.4)	9.8 (3.3)	9.4 (3.3)
Married	.893	.711	1	.810	.766	.287
Never Married	.034	.073	0	.032	.234	0
Number of Children under 18	0.04 (0.26)	0.04 (0.28)	0	0.02 (0.18)	.034 (.222)	.032 (.268)
Probability is a DI Beneficiary	-	.171[†] (.101)	.154[†] (.070)	-	.204[†] (.140)	.143[†] (.112)
Individual DI Benefits	-	8390* (3424)	0	-	6098* (3089)	0
Family DI Benefits	-	9331[†] (4418)	0	-	6227[†] (3406)	0
Probability Induced to Apply	-	.0017[†] (.0010)	.0015[†] (.0007)	-	.0020[†] (.0014)	.0014[†] (.0011)
Number of observations	3559	6023	66	4565	1649	7104

Notes: The mean and standard deviation are shown for key characteristics of the elderly population, broken down by gender and DI status. The second and third columns report statistics for individuals who are imputed to be current DI worker beneficiaries. Individuals included in the sample for the second column generate Social Security, or what we term DI benefits, on their own account. Individuals included in the third column receive Social Security benefits due to another family member and will therefore not generate additional benefits due to the policy reform. The first column presents statistics for all other elderly individuals who are assigned a zero probability of having been a working-aged DI recipient. *The reported DI benefits have been scaled up by 25% to account for underreporting in the CPS. [†]These are predicted values, using the methods described in the text.

Table 4: Summary of the Costs of the Policy Reform

	Millions of Dollars
Net transfer to current DI recipients (static costs)	**286.4**
Increased benefits to worker beneficiaries	**191.1**
Increased benefits to elderly beneficiaries	**129.6**
Decreased SSI benefits to current beneficiaries	**(28.9)**
Decreased food stamp benefits to current beneficiaries	**(0.9)**
Increase in federal taxes paid by current beneficiaries	**(4.1)**
Increase in state taxes paid by current beneficiaries	**(0.5)**
Costs associated with behavioral response (dynamic costs)	**120.0**
Benefits paid to successful new applicants	**89.7**
Expected costs due to changes in taxes paid	**29.3**
Decrease in federal taxes paid due to spousal labor supply response	1.2
Decrease in state taxes paid due to spousal labor supply response	0.3
Decrease in federal taxes paid by marginal applicants	22.8
Decrease in state taxes paid by marginal applicants	5.1
Expected transfer costs (due to marginal applicants)	**1.0**
Change in SSI payments	0.3
Change in food stamp benefits	1.2
Change in AFDC payments	(0.5)

Notes: The values in bold represent totals, and the values that are not in bold are subtotals. Parentheses () are used to indicate cost savings. All values are in 1990 millions of dollars.

Table 5: Representative Individuals' Willingness to Pay

	Working Age Male	Working Age Female	Male Applicant	Female Applicant
Willingness to pay				
θ=1	\$1.50	\$1.20	\$1.20	\$1.10
θ=2	\$2.30	\$1.40	\$1.40	\$1.30
θ=4	\$5.10	\$2.10	\$2.10	\$1.70
Share of household earnings	2/3	1/3	--	--
Average replacement rate	0.5	0.5	--	--
Drop in household income	1/3	1/6	1/6	.12

Notes: The first three rows of Table 5 present the willingness to pay for \$1 of DI benefits at different levels of risk aversion, given the separable constant relative risk aversion utility function presented in the text. The parameters used in this calculation are found in the next three rows. The share of household earnings and average replacement rate for working age males and females are based on the CPS, and from published data on Social Security benefits and average earnings. These are used to calculate the drop in household income for these groups of individuals, in the final row of the table. The SIPP was used to calculate the drop in household income for male and female applicants.

Table 6: Implicit Price of Additional Insurance by Education Category

	Direct Cost	Indirect Cost	Total Cost	Total Cost/ Income (*1000)	Pr(DI Recipient)	Benefit Increase	Implicit Price
Less than High School	1.05	0.39	1.44	0.081	.092	61.1	0.26
	(0.01)	(0.01)	(0.02)	(0.001)	(.002)	(0.6)	
High School Degree	2.31	0.91	3.22	0.115	.036	70.3	1.27
	(0.01)	(0.01)	(0.02)	(0.001)	(.001)	(0.7)	
Some College	2.90	1.18	4.08	0.121	.023	72.2	2.46
	(0.02)	(0.01)	(0.03)	(0.001)	(.001)	(1.3)	
Finished College	3.80	1.77	5.57	0.120	.012	74.2	6.26
	(0.02)	(0.01)	(0.03)	(0.001)	(.001)	(1.5)	

Notes: The unit of observation is the family. The first two columns show the average direct and indirect costs borne by families without current working-aged or elderly DI beneficiaries, separately for families broken down by the education level of the head. The third column presents the average total cost (sum of the direct and indirect costs) for families in each educational category. The fourth column reports the average of the total cost as a fraction of pre-regime change income for families (these values are multiplied by 1000 for ease of viewing). The next two columns show the probability that a family has a worker or elderly beneficiary and the average increase in family benefits. The last column is calculated by summing the direct and indirect costs and dividing by the expected benefits (the probability the family has a current beneficiary multiplied by the expected increase in benefits). Standard errors are in parentheses.

Table 7: Pre-reform and Post-reform Income

	Individuals in families with: No DI beneficiary	DI worker beneficiary	Elderly DI beneficiary	Individuals in families with: No DI beneficiary	DI worker beneficiary	Elderly DI beneficiary
	In the base period (pre-reform):			Change in income following the reform:		
Family after-tax income	30773	23043	21969	-4.48	78.3	54.5
	(21728)	(16707)	(15722)	(3.73)	(63.7)	(56.2)
Smoothed family after-tax income	27959	21921	23047	-1.87	40.3	30.4
	(10149)	(7858)	(7595)	(1.21)	(29.0)	(30.6)
Per capita after-tax income	12141	9831	12197	-1.68	36.5	30.0
	(9659)	(7087)	(8134)	(1.58)	(32.1)	(32.6)
Per capita smoothed after-tax income	12101	10464	13476	-0.75	21.4	17.4
	(7360)	(4897)	(4861)	(0.65)	(21.4)	(20.0)
Per capita equivalent after-tax income	18002	14398	17462	-2.50	53.5	43.0
	(13773)	(10106)	(11602)	(2.28)	(46.4)	(46.5)
Per capita equivalent smoothed after-tax income	17941	15273	19295	-1.11	31.2	24.9
	(10392)	(6778)	(6891)	(0.92)	(30.6)	(28.5)
Number of Observations	155633	2685	21929	155633	2685	21929
Weighted Number of Obs	240.2 million	4.3 million	4.1 million	240.2 million	4.3 million	4.1 million

Notes: The first three columns present mean income pre-reform for individuals classified by family type. The last three columns show the average change in income following the reform. The rows correspond to different measures of income. Rows 1, 3, and 5 are based on income measures that have not been smoothed using the relationship between income and consumption in the CEX, while rows 2, 4, and 6 are based on smoothed versions. Rows 1 and 2 use family income, rows 3 and 4 use per capita family income, and rows 5 and 6 use adjusted (for economies of scale and number of adults relative to children) per capita family income. All cells are weighted using the March CPS person weight.

Table 8: Welfare Analysis (Adjusted per capita income)

θ	Average Utility Pre-reform	Normalized Change in Monetary Component of Utility for Individuals in Families w/: No DI Beneficiary	DI Worker Beneficiary	DI Elderly Beneficiary	Net Welfare Impact
1) Individual income = adjusted per capita family income, not smoothed					
0	17929	-.0001 [-33485]	.0030 [12870]	.0024 [9752]	-.00004 [-10862]
1	9.39	-.00001 [-3303]	.00045 [1937]	.00031 [1257]	-4.39×10^{-7} [-109]
2	.00027	-.0002 [-44130]	.0019 [8234]	.0013 [5220]	-.00012 [-30677]
4	.00002	-.0042 [-984727]	6.11×10^{-9} [.026]	2.42×10^{-8} [.098]	-.00396 [-984727]
2) Individual income = adjusted per capita family income, smoothed					
0	17916	-.0001 [-14894]	.0017 [7513]	.0014 [5645]	-.00001 [-1736]
1	9.64	-.00001 [-1559]	.00021 [914]	.00015 [594]	-2.05×10^{-7} [-51]
2	.00008	-.00006 [-14312]	.00217 [9349]	.00119 [4856]	-4.30×10^{-7} [-107]
4	1.62×10^{-11}	-.00020 [-48362]	.00018 [760]	.00006 [229]	-.00019 [-47353]

Notes: Per capita family income is adjusted for economies of scale and the number of adults and children as described in the text. The top panel shows the welfare analysis based on unsmoothed adjusted per capita family income, while the bottom panel is based on adjusted per capita family income that is smoothed according to the relationship between after-tax income and consumption in the CEX. Each row presents results using a different parameter of relative risk aversion (θ). The second column shows average utility in the base period, the next three columns show the average change in utility before and after the reform (normalized by average base utility) by family type, and the final column shows the average overall change in utility (also normalized by average base utility). The aggregate welfare impact by each family type and overall is shown in brackets.

Table 9: Net Welfare Impact by Education Level (Adjusted Per Capita Income)

θ	Level of Educational Attainment Less than high school	High school degree	Some college	Completed college
0	5346	-711	-2048	-3184
1	563	-72	-201	-325
2	3939	-750	-1727	-3193
4	-10820	-5187	-10971	-2755

Notes: Positive values indicate that the net welfare impact of a 1% increase in DI benefit generosity is positive for the subgroup and level of risk aversion (θ) indicated, while negative values indicate a net welfare loss. Individual utility is calculated using per capita family income smoothed using the relationship between consumption and income in the CEX, which is then adjusted for economies of scale and the number of adults and children.

Table 10: Net Welfare Impact Incorporating Health Insurance, by Education Level

θ	All individuals	Level of Educational Attainment: Less than high school	High school degree	Some college	Completed college
0	--	++	--	--	--
1	--	++	--	--	--
2	--	++	--	--	--
4	++	++	++	--	--

Notes: Positive values indicate that the net welfare impact of a 1% increase in DI benefit generosity is positive for the subgroup and level of risk aversion (θ) indicated, while negative values indicate a net welfare loss. Individual utility is calculated using per capita family income smoothed using the relationship between consumption and income in the CEX, which is then adjusted for economies of scale and the number of adults and children.