

Disability Insurance Benefits and Labor Supply

Jonathan Gruber
MIT and NBER

November, 1996
Revised July, 1999

Disability Insurance (DI) is a public program that provides income support to persons unable to continue work due to disability. The difficulty of defining disability, however, has raised the possibility that this program may be subsidizing the early retirement of workers who are not sufficiently disabled. A critical input for assessing the optimal size of the DI program is therefore the elasticity of labor force participation with respect to benefits generosity. Unfortunately, this parameter has been difficult to estimate in the context of the U.S. DI program, since all workers face an identical benefits schedule. I surmount this problem by studying the experience of Canada, which operates two distinct DI programs, for Quebec and the rest of Canada. The latter program raised its benefits by 36% in January, 1987, while benefits were constant in Quebec, providing exogenous variation in benefits generosity across similar workers. I study this relative benefits increase using both simple "difference-in-difference" estimators and more parameterized estimators that exploit the differential impact of this policy change across workers. I find that there was a sizeable labor supply response to the policy change; my central estimates imply an elasticity of labor force non-participation with respect to DI benefits of 0.28 to 0.36.

* I am grateful to Courtney Coile, Kevin Frisch and particularly Sue Dynarski for excellent research assistance, to Doug Bernheim, John Bound, Peter Diamond, Louis Kaplow, Don Parsons, Sherwin Rosen, an anonymous referee, and seminar participants at Harvard University, Brown University, and the NBER for helpful comments, to Marilyn Knock and Ging Wong for comments and invaluable assistance with data collection, to Bernard Dussault and Pierre Plamondon for endless patience in explaining the institutional features of the Canadian DI system, and to Human Resources Development Canada and the National Institute of Aging for financial support.

One of the largest social insurance programs throughout the developed world is Disability Insurance (DI). In the U.S., the DI program has over 6 million beneficiaries and benefit payments of almost \$46 billion (Social Security Administration, 1998). In theory, DI provides benefits for workers who are physically incapable of finding suitable work. Disability would seem to be an ideal targeting device, allowing program administrators to divert resources towards those truly in need of income support.

In practice, however, it is difficult to determine whether workers are truly disabled. A number of studies have revealed substantial error of both the Type I and Type II variety in the disability determination process.¹ In addition, DI benefits in the U.S. are fairly generous; on average, disability insurance replaces 42 percent of a worker's previous earnings, and these benefits are non-taxable, raising the after-tax replacement rate even further. The difficulty of appropriately identifying disability and the generous levels of benefits available have led many observers to claim that DI is largely distorting work decisions, and in essence subsidizing the early retirement of the older workers for whom appropriately defining career-ending disability is most difficult.

At the same time, other analysts have claimed that the vast majority of the DI recipient population is truly disabled and unable to pursue gainful employment, suggesting that any distortion to labor supply decisions is minimal. This argument implies that the welfare gains from redistributing resources to the low income disabled would outweigh any costs through reductions in labor supply. A critical input for evaluating this claim, and for modeling the appropriate level of DI benefits, is therefore an empirical estimate of the elasticity of response of labor supply to benefit levels.

There is a substantial U.S. based literature on the effects of DI benefits on labor supply. Evaluating this behavioral response in the context of the U.S. case has proved to be difficult, however. This is because the DI program in the U.S. provides benefits which differ across workers primarily through their past earnings

¹These studies are reviewed in Parsons (1991a).

histories. But one's earnings history will most likely be highly correlated with one's tastes for work at older ages, and it is difficult to disentangle the behavioral effects of DI from these taste differences. What is required to distinguish the effects of DI is differences in benefit levels across workers which arise independently of their underlying tastes for work at older ages.

Such differences have arisen in the context of the Canadian DI system. DI in Canada operates in much the same way as it does in the U.S., with the key difference being that the program is administered under two different plans: the Quebec Pension Plan, or QPP, which covers the province of Quebec only, and the Canada Pension Plan, or CPP, which covers the rest of Canada. These two systems are identical in most respects. Since the early 1970s, however, benefits have risen more rapidly under the QPP; by the end of 1986, benefits under the QPP were substantially more generous than were benefits under the CPP, particularly for those disabled workers who had low earnings before their disability. Then, in January 1987, the CPP raised its benefit levels to equalize the generosity of the two systems. This resulted in a rise in benefits under the CPP of almost \$2000 (Canadian) per year; relative to Quebec, there was a 36% rise in the replacement rate for the typical disabled worker. This dramatic shift in differential benefits generosity is precisely the type of change that can be used to evaluate the labor supply response to DI benefits. That is, this policy change provides an opportunity that is not available in the U.S.: the chance to study the effect of changing DI benefits differentially for some workers (those not in Quebec) and not for others (those in Quebec).

I use this policy change to estimate the elasticity of labor supply for older persons with respect to DI benefits. My data for this exercise come from the Canadian Survey of Consumer Finances (SCF), an annual cross sectional survey which collects information on demographic characteristics and work behavior. I match this data information on the benefits available under the CPP and QPP over time. And I compute two types of estimates of the policy change. The first is a standard "difference-in-difference" estimate which

focuses on the labor supply effect of the large relative change in benefits in the rest of Canada relative to Quebec. The second is a more parameterized estimate that exploits the underlying variation in the impact of this policy change across workers within the CPP and QPP plans.

For both estimators, I find that there is a large effect of benefits on the labor supply of older workers. My central estimates imply that the elasticity of labor force non-participation with respect to benefit levels is 0.28 to 0.36. This finding is robust to a variety of specification checks.

The paper proceeds as follows. In Part I, I review the key facts on the DI program in Canada, compare the system to that in the U.S., and review the empirical literature on the behavioral effects of DI. In Part II, I describe the data source, and I discuss my empirical strategy in Part III. Part IV presents my results for labor supply estimation. Part V concludes.

Part I: Background

The Canadian DI Program

The Canadian DI program dates from January 1, 1966, when it was introduced along with work related retirement pensions under the QPP and CPP. Eligibility is conditioned on working and contributing to the program in 2 of the previous 3 years, or 5 of the previous 10 years. Eligibility is also conditioned on an inability to pursue gainful employment due to a physical disability.² This is determined by a medical examiner; individuals who are denied claims have the right to appeal their decisions at least twice to higher levels of adjudication. About 40% of claims were denied at the initial determination stage under the CPP in 1989, the last year of my sample (and the earliest year for which data are available); the denial rate for the QPP at this time was 33%. While the CPP has a higher initial denial rate, it has a lower denial rate during the appeals

²Under the CPP, gainful employment means any job. Under the QPP gainful employment means "usual job" since 1993; it was any job before then. Since 1984, for those over age 60 in the QPP, gainful employment means one's last job, but this paper focuses on those below age 60 only.

process, so that after successful appeals are factored in the overall denial rate is quite similar across the two plans (32% under CPP vs. 30% under QPP).³ There is a three to four month waiting period from the onset of disability before benefit receipt begins.⁴ The DI program currently has approximately 340,000 beneficiaries, with benefit payments of over \$3 billion.⁵

Under both the CPP and QPP, benefits consist of three parts. The first is a (lump sum) flat rate portion available to all eligible workers. The second is an earnings related portion. This portion is calculated by first inflating the workers earnings history (back to 1966) to current dollars using a wage index, dropping the lowest 15% of months of real earnings, and taking 18.75% of the average of the remaining series.⁶ The final portion is a child allowance, which is a fixed amount per month per child under the age of 18. Averaging across both the CPP and QPP, benefit levels replaced approximately 26% of the average earnings of 50-59 year old workers in Canada in 1986.⁷

While the computation of the earnings related portion has been identical across the CPP and QPP since the programs' inception, there have been differences in the other two parts of the benefits computation. The flat rate portion was identical in the two provinces until 1972, at which point it began to rise more rapidly in Quebec. This time series pattern is illustrated in Figure 1, which graphs the flat rate over time. There is

³Based on unpublished administrative data from the CPP and QPP. It is difficult to infer relative differences in screening stringency across the programs from these figures, since the underlying pool of applicants at any point in time may differ in their health; see Gruber and Kubik (1997) for a further discussion of the interpretation of denial rate data.

⁴Technically, benefits flow on the first day of the third full month after the month of disability - so that if the injury occurs on the first of the month, the waiting period is four months.

⁵Canadian data from unpublished tabulations by CPP and QPP.

⁶Months of previous receipt of disability insurance are also excluded, as are months where the worker had primary child-bearing responsibility. Since I focus only on older men, I ignore the second of these in the benefits calculation.

⁷Based on author's computation, using the potential benefits calculation methodology described below.

a growing gap between the two provinces over time, which by 1987 was over \$150 per month. Then, in January, 1987, the CPP raised its flat rate portion to be identical to that of the QPP, a rise of over 150%.⁸ On average, this represented a rise of 36% in the replacement rate of the CPP relative to the QPP. The two series have moved in tandem ever since. There have also been differences in the computation of the child benefit over time; this benefit became more generous in the CPP, rising steadily from \$57 per child per month in 1981 to \$155 in 1993, while it remained low (\$29) until 1993 under the QPP. This counteracted some of the time series gap in flat rate portions for those disabled workers with children, but had little effect on the huge relative change in benefits in January, 1987.

It is important to note that the increase in benefits under the CPP was not the only policy change of 1987; there were two other changes that are potentially relevant for this analysis. The first was a reduction in the required earnings history to qualify for CPP Disability benefits. Before 1987, eligibility was conditioned on having contributed in the lesser of 10 years or 1/3 of one's career; in 1987, the requirements were eased to those described above. While making a number of younger workers eligible for DI, however, this had little practical effect on the older population on which my study will focus, since these workers generally had enough experience to be eligible under either system.

The second policy change is potentially more problematic: the introduction of the early retirement option (at age 60) for retirement benefits under CPP.⁹ This means that even in the absence of a change in DI benefits there may be reduced labor force participation among those aged 60-64. This motivates my focus on workers below age 60 for this analysis. It seems unlikely *ex ante* that this change had important effects on workers below age 60, since Baker and Benjamin (1996) find little effect on workers in the age 60-64

⁸Note that this change applied to both new applicants and existing beneficiaries, so that there was no incentive from this new law to delay applications for disability insurance.

⁹Individuals who chose to retire before 65 see their benefits reduced by 0.5% per month for each month before 65 that they claim, for a total reduction in benefits of 30% for those claiming at age 60.

group who were directly affected by the policy change. Nevertheless, in a life-cycle labor supply model it is certainly possible that changes in the opportunity set after age 60 can have impacts on decisions made before that point. I therefore provide direct evidence below that this early retirement change is not driving my results for the 45-59 year old sample, by exploiting the fact that Quebec changed its age of early retirement several years earlier.

Of course, I cannot rule out the hypothesis that this benefits increase was itself motivated by underlying (relative) changes in the (non-Quebec) economy that affected the relative job prospects of older workers.¹⁰ After presenting my basic results, I therefore also discuss a number of tests which suggest that this is not the case, justifying the use of this policy change as an instrument for DI benefits.

Comparison to the U.S. Program

The DI programs in the U.S. and Canada are quite similar, with only two major differences. The first is in the structure of benefits. Benefits in the U.S. consist primarily of an earnings related portion, without any lump sum component. On the other hand, the schedule translating past earnings to benefits is much more progressive than in Canada, so that the two countries have a similar redistributive structure to their benefits schedules. Benefits are much higher in the U.S. on average, however, with a replacement rates of 42% for the average worker (U.S. Congress Committee on Ways and Means, 1990). Moreover, income from DI is not taxable for most households, whereas it is fully taxable in Canada. As a result, after tax replacement rates are much higher in the U.S.

The second difference is the stringency of the screening process for DI. While the basic structure is the same (with an initial claiming stage and an appeals process), the denial rate at the initial stage in the

¹⁰No such motivation is mentioned by either the law itself or by narratives describing the political economy of the DI program (Human Resources Development Canada, 1995).

U.S. is 69%; the ultimate denial rate (factoring in appeals) is 51%, as opposed to roughly 30% under the Canadian system (U.S. Congress, 1998). Also, the waiting period for benefits receipt (5 months) is somewhat longer than in Canada. Despite more stringent screening (and perhaps because of the more generous benefits), the incidence of DI receipt is somewhat higher in the U.S.; 4.8% of 45-59 year old men are on this program, as compared to 3.9% of 45-59 year old men in the CPP provinces.¹¹ It is unclear, of course, whether this difference represents underlying differences in screening stringency, application propensity, or the health of the population.

DI and the Behavior of Older Workers

The literature on the effects of DI on the labor market in the U.S. is motivated by a striking time series fact: the almost exactly parallel increase in the DI rolls and decline in the labor force participation of older men in the 1960s and 1970s. DI enrollment grew from 455,000 in 1960 to 2.9 million by 1980 (U.S. Department of Health and Human Services, 1993). Over this same period, the non-participation rate among 45-54 year old men rose by 105%, and the non-participation rate for 55-64 year old men rose by 111% (Bound, 1989). But drawing causal inferences from this time series data is problematic, as there were a number of other changes in the labor market and non-labor market opportunities of older males during this era.¹²

A sizeable literature has attempted to use cross-sectional variation to identify the role that DI plays

¹¹Data for 1993. Data for Canada from Human Resources Development Canada (1996); data for U.S. from U.S. Department of Health and Human Services (1994). The cost of the DI program is, as a result, roughly 10% higher as a share of GNP in the U.S. than in Canada.

¹²For example, there was rapid growth in retirement incomes in this era, both due to increased Social Security benefit levels, and increased coverage of the labor force by pensions (Lumsdaine and Wise, 1990). There was also a rapid rise in the labor force participation of wives, which could either increase (through the income effect) or decrease (through complementary leisure effects) non-participation.

in the labor force participation decisions of older men. These studies generally proceed by modeling labor force participation or DI reciprocity as a function of potential DI benefit levels. The first study to do so was Parsons (1980), who estimated an elasticity of labor force non-participation with respect to DI benefit levels of 0.49 to 0.93. His upper bound estimate implied that increases in DI benefits (as well as in benefits from other welfare programs for older workers) over the 1960s and 1970s could explain the entire time series trend in non-participation. Other estimates have supported the contention that DI has a significant disincentive effect, although the estimated magnitudes have generally been much smaller than that of Parsons; see Leonard (1986) and Bound (1989) for reviews of this evidence, which estimates elasticities of non-participation in the range of 0.1 to 0.2.

Bound (1989) argues, however, that this type of strategy is likely to yield misleading inferences for the effect of DI generosity on labor force participation. Since DI benefits are a redistributive function of past earnings which is common to all workers, variation in potential benefits comes primarily from differences in earnings histories across workers. This leads to a fundamental identification problem in modeling the effect of potential DI benefits on work decisions: a finding that workers with higher potential DI replacement rates are more likely to leave their jobs may simply reflect the fact that low earning workers have less of a desire to continue working.¹³ What is clearly needed to identify the behavioral impact of DI benefits is variation in program generosity which is independent of underlying tastes for work. This variation is provided by the large relative benefits increase under the CPP in 1987.

I am aware of only one article which has analyzed the behavioral incentives of the Canadian DI system. Maki (1993) pursues two different strategies in analyzing the effects of benefits on labor force attachment. First, he uses a panel of aggregate province-level data for the 1975-1983 period, and he finds

¹³Studies such as Haveman and Wolfe (1984) attempt to correct for this omitted variables bias, but Bound (1989) argues that the problem has not been convincingly resolved because of the strict distributional assumptions necessary to achieve their solution.

a strong negative correlation between benefits (normalized by average wages) and participation. But this effect disappears when he includes province and year fixed effects in the regression, which may be necessary to control for underlying trends in labor supply and fixed differences in tastes for work across areas. Second, he uses a cross-section of micro-data for 1985 to estimate a structural model of the effect of DI, along the lines of most of the U.S. literature. With this approach, his estimates are very sensitive to the exact specification of his model. But this technique is once again subject to Bound's (1989) criticism, since the variation here mostly comes from differences in individual characteristics that may otherwise be correlated with tastes for work.

Part II: Data

The Canadian Survey of Consumer Finances (SCF) is an annual supplement to the nationally representative monthly Labor Force Survey (LFS), conducted each April. Comparable to the March Current Population Survey in the U.S., the SCF contains data on labor force attachment, demographics, and income. There are survey data collected for individuals from April 1982 onwards, with the exception of April 1984. Family level data were also collected every other year from 1976 to 1980.¹⁴ I use the surveys from April 1985-86 as the "before" period, and those from April 1987-89 as the "after" period.¹⁵ I do not use earlier surveys in the base case analysis because there is no April 1984 survey; I do use the 1982 and 1983 data in a specification check below. I do not use later surveys because there was a major change in the classification of the education variable in April 1990, rendering it difficult to follow precise education groups from before

¹⁴There are actually some family surveys for some earlier years, but differences in the definition of the education variable render them useless for my purposes.

¹⁵The policy change of interest was enacted in July, 1986, and became effective in January, 1987; since my before period ends in April, 1986, I avoid any anticipatory labor force leaving behavior between the enactment and effective dates.

1990 to after; following educational groups is a key feature of my approach to measuring potential DI benefits. Another advantage of using this set of years is that it avoids the contamination of the estimates by the recessions of the early 1980s and early 1990s, which might affect older workers' propensity to apply to the DI program.¹⁶

I focus on men aged 45-59 for this analysis. My focus on men follows the previous literature on DI. Since I only have cross-sectional data on a worker's labor force attachment, I do not know whether that worker has the requisite earnings history to be eligible for the DI program. This problem should be minimal for men, who generally have sufficient earnings histories to qualify, but may be more of a problem for women.

My choice of age group is dictated by two considerations. First, I wanted to use workers old enough so that DI was a relevant option in their choice set. For this age group the incidence of DI benefits for men in the CPP is 3.9%; this is 4 times as high as the incidence rate among those age 40-44. Second, as was noted earlier, the increase in DI benefits under the CPP was not the only important policy change in 1987; there was also a reduction in the age of eligibility for CPP retirement benefits to 60, which I hope to avoid by focusing on those below age 60.

Part III: Empirical Methodology

Difference-in-Difference Estimation

The most straightforward means of analyzing this policy change is through the "difference-in-difference" framework (Card, 1992; Gruber, 1994). This involves a simple comparison of the change in behavior outside of Quebec, where benefits increased, with the change in behavior inside Quebec, where

¹⁶See Lewin-VHI (1996) for evidence on the cyclical responsiveness of DI applications.

benefits did not.¹⁷ This comparison can be implemented in a straightforward manner by estimating logistic regressions of the form:¹⁸

$$(1) \quad NE_i = f(a + \beta_1 CPP + \beta_2 AFTER + \beta_3 CPP * AFTER + \beta_4 X_i + e_i)$$

where NE_i is dummy for non-employment of person i
 CPP is an indicator for whether the individual lives in CPP province
 AFTER is an indicator for whether the year is after the policy change
 X_i is a set of covariates for person i (age, married, education, number of children)

In this regression framework, I control for location by including a dummy for whether an individual lives in a CPP province or in Quebec. And I control for time by including a dummy for whether this observation is from before or after the policy change. The coefficient of interest (β_3) therefore measures the effect of being covered by the CPP, relative to being covered by the QPP, after the benefits increase, relative to before.

The dependent variable is a dummy for whether the 45-59 year old man was not working during the week of the SCF survey. Thus, the coefficient β_3 measures the effect of the policy change on non-participation, defined as non-work. I also include controls for education, age, marital status, and number of children to control for any observable differences between workers that might confound the analysis. Education is measured by four dummy variables for less than 9 years of education, 9-10 years of education, 11-13 years of education, and some post-secondary education. Age is measured by a set of dummies for single years of age from 45-59. There are separate dummies for each number of co-residing children under

¹⁷Note that I assume that there is not migration across the Quebec border in response to DI benefits differences. Under CPP or QPP rules, if a worker moves from a CPP region to Quebec and immediately files for benefits, he receives the benefits he was entitled to under the CPP (similarly QPP benefit rules apply for moves from inside to outside Quebec). If, however, this worker moved and then worked in Quebec before applying, he would be eligible under the QPP rules. So workers would have to anticipate a future application need for there to be a migration incentive.

¹⁸I use the logistic function to follow previous literature in this area. The results are similar if either probit models or linear probability models are used instead.

age 18 (up to a maximum of 8 children).

This approach is attractive because it allows me to cleanly identify the effects of the benefit change. But it has two limitations. First, it does not allow me to directly measure the elasticity of response to the change in DI benefits, since I have measured only the numerator of this elasticity (the change in labor supply) and not the denominator (the change in potential benefits). Second, this is a very rough categorization of the data that does not fully take advantage of this policy change, since there is further variation available in potential benefits *within provinces at a point in time*. Since only the flat rate portion was increased by the CPP, the percentage point increase in the replacement rate is much larger for those with a low lifetime level of earnings, as the flat-rate portion is a larger share of their DI benefits. I can use this fact to further identify the effect of the benefit change, by exploiting the differential impact of the benefits change across workers of different lifetime earnings levels.

Parameterized Models

To address both of these points, I must measure the change in potential benefits for each person in the SCF sample. In theory, calculating potential DI benefits requires longitudinal information on workers' earnings since 1966, which is not available in the SCF (an annual snapshot of earnings). Thus, I instead calculate "synthetic earnings histories" for groups of workers in order to impute their potential DI benefits. This is done in several steps. I begin by creating a database using each of the individual SCF's for April 1982-1989, and using data on the male heads of families from the family SCF for April 1976, 1978, and 1980. In each of these data sets, I then divide workers into cohort cells according to their age, location (four regions: Quebec, Ontario, the Atlantic Provinces, and the remainder of Canada), and their educational attainment (the

four groups described above). I then tabulate the median earnings in each cohort cell for each year.¹⁹ By stringing together the median earnings in each cohort cell through time, I can form a proxy for the earnings history of a worker in that cohort cell.

These surveys contain annual earnings data for the years 1981-1988, with the exception of 1983 when no survey was carried out, and biannual data from 1975-1979; for the missing years, earnings is imputed as an average of the surrounding years. To backcast from 1975 to 1966, before cross-sectional survey data is available, I first estimate cross-sectional age-earnings profiles by education group in the 1975 survey. I then apply these estimates to "un-age" the workers in the 1975 survey back to 1966, and deflate these pre-1975 profiles by average wage growth by region, using data from Gruber and Hanratty (1995).

With these synthetic earnings histories in hand, it is then straightforward to compute potential DI benefits using the legislative rules in place in CPP and QPP in a given year. The key regressor, the replacement rate, is this potential benefit over the synthetic earnings for the cell in the year before the survey. This measure does not vary individual-by-individual, but rather only cell-by-cell, where the cells are defined by each education/region/year group.²⁰

I then estimate regression models of the form:

$$(2) \quad NE_i = f(a + \beta_1 RR_i + \beta_2 X_i + \beta_3 t_t + \beta_4 ED_i * d_j + \beta_5 ED_i * t_t + e_i)$$

¹⁹That is, for 45-59 year old in 1989, I use 44-58 year old in 1988, 43-57 year old in 1987, and so on back through time. I have also computed benefits using the mean; the results are quite similar.

²⁰I do not include the worker's potential child benefits in the computation of the replacement rates, for two reasons. First, this preserves the variation in potential benefits only at the cell level, which is important for my identification strategy. Second, it is not clear how to combine child benefits, which for these older workers will only be paid for the presumably small number of years until the child turns age 17, with the other benefit components, which will be paid until age 65 (at which point all disabled are shifted to the retirement income system). In practice, this is not a very important consideration, as only 1/3 of my sample has any children. Adding child benefits to the computed benefit total, based on the actual number of children, raises the level of the replacement rate somewhat, but not the relative change; and the estimated elasticities reported below are similar whether or not child benefits are accounted for in calculating replacement rates.

where RR is potential replacement rate
 ED is a set of dummies for education categories (four categories)
 d_j is set of region dummies (four regions)
 t_t is set of year dummies

This model controls for fixed effects for year, for each of the 16 education*region cells in each year, and for education*year. The first of these is included to capture secular trends in labor market opportunities in Canada, as in equation (1). The second of these is included to account for the fact that there is a potential spurious correlation between the labor supply choices of these 16 groups and their potential replacement rate; this is just a restatement of the criticism levelled by Bound (1989) against the U.S. literature. By taking out fixed effects for each group, I only use changes in each group's potential replacement rate over time, to identify the effect of DI. Finally, I am potentially concerned about identification from changes in the return to education over this period, which would affect both the replacement rate and the decision to work, so I include the set of education*time interactions.

Conditional on this set of controls, the model is identified by two sources of variation: changes over time in the CPP provinces relative to Quebec (region*time), and how those changes evolve differentially across these 16 groups (region*education*time). The first of these is the difference-in-difference variation that was used to identify model (1); the second is additional variation from the differential impact of this policy change across groups. This additional variation is potentially useful in pinning down the elasticity of labor supply; indeed, in one specification check below, it allows me to control for relative shocks to the labor markets in Quebec and the rest of Canada over this period. Moreover, the resulting coefficient β_1 is now directly interpretable as the benefit semi-elasticity of labor supply.

Part IV: Results

Means

Table 1 presents the means of the data set, divided into the CPP regions and the QPP region, before

the law change and afterwards. The final column of the table shows a first pass difference-in-difference estimate of the policy effect. There are two findings of interest from Table 1. First, as the first two rows show, the policy change was associated with a significant increase in benefits. While the replacement rate was roughly constant in Quebec, it rose substantially in the rest of Canada; the relative rise was 8.8 percentage points, or 36% of the baseline average replacement rate.

Second, there is strong evidence of a labor supply response to the benefits increase. Non-participation rises from before to after in the CPP regions, and falls in the QPP regions; the latter finding reflects the underlying improvements in the Canadian economy over this period. As a result, there is a large relative rise in non-participation in the CPP regions of 2.7 percentage points.

DD Regression Results

The next table formalizes the inferences from the table of means in a regression model, including as well the set of covariates in (1). Recall that the regression also includes a full set of dummies for age and number of children which are not reported in the table. The regression is estimated as a logistic model; the last row shows the effect of the DD interaction on the probability of being non-employed, which is the average effect across the sample on the predicted probability of non-participation.

These findings confirm the conclusion from Table 1 that there is a response to the policy change. The effect is slightly smaller than in Table 1, with a relative rise in non-employment in the CPP regions of 2.3%; it is statistically significant. This is still a quite sizeable response, indicating that the 36% benefits rise led to a rise in non-employment of 11.5% from the baseline value, for an implied (arc) elasticity of non-participation of 0.36. Thus, this straightforward DD estimate is very supportive of a strong labor supply response to the benefits increase. The control variables in the regression have their expected effects, with married and more educated workers less likely to be non-participants. The age dummies (not shown) have

the expected upwards trend, while there is no clear pattern from the dummies for number of children (also not shown).

Parameterized Model

As noted above, these DD estimates do not fully exploit the available variation in potential benefits across workers in Canada. To do so, in Table 3 I present estimates of the replacement rate model (2). For each model, I show the coefficient of interest, the implied effect of the 8.8 percentage point replacement rate rise, and the implied elasticity of non-employment.

The first row presents the basic model. There is a sizeable and significant effect of the potential replacement rate. The estimate implies that this policy change raised the non-employment rate by 1.2 percentage points, which is substantially below the DD estimate, but is more precisely estimated. The implied arc elasticity of non-participation with respect to benefits is 0.19.

One potential concern about the identification of this model, however, is that the variation in benefits does not arise solely from the policy change, as it impacts the 16 different education*region groups, but rather also from year to year changes in replacement rates within the before and after periods. Some of this year to year variation is legislative, arising from evolving system parameters over time (ie. changes in the flat rate). But some of it also arises from year to year differences in earnings across education*region cells, which induce changes in the potential replacement rate, but which might also be independently correlated with the labor supply decisions of individuals in those cells. Moreover, this year to year variation may reduce the signal to noise ratio in my key regressor, since the true variation of interest comes from the policy change only.

In order to purge the model of these year to year changes and focus solely on the before/after comparison, in the next row of Table 3 I present instrumental variables estimates of the model. The instruments are a set of interactions of education*region*AFTER, where as in equation (1) AFTER is an

indicator for being after the policy change. When instrumented in this way, the only variation in benefits that is used by the regression model is the before/after difference in benefits, on average and as it impacts differentially these 16 education*region groups. That is, this IV strategy provides the means of extending the DD estimation to account for variations in the impact of the policy by education and region.²¹ The first stage fit is excellent; the F statistic is 5500.

In fact, this instrumental variables approach raises the estimates substantially, consistent with the notion that noise in the year-year replacement rate changes was biasing the estimate downwards. At this new point estimate, the implied effect on non-participation from the policy change, 1.8 percentage points, is close to the DD estimate. The implied arc elasticity of non-employment with respect to benefits rises to 0.28.²² This is higher than the post-Parsons literature in the U.S., but is much less than the lower bound of Parsons' estimates.²³

Addressing Alternative Hypotheses

The fundamental identification assumption embodied in the estimation thus far is that there was no

²¹In terms of the discussion above, in this model the identification comes solely from region*AFTER and region*education*AFTER.

²²I focus on non-employment as the outcome of interest, because there is a vague distinction between unemployment and non-participation in the labor force in this age group. If the results are replicated using non-participation (e.g. moving the unemployed from 1 to 0 in the dependent variable), the estimated response is about 85% as large. Since the mean of non-participation is only 63% as large as for non-employment, this implies elasticities of non-participation that are 35% larger than the elasticities of non-employment reported here.

²³Note also that my estimates are consistent with aggregate relative movements in the DI rolls over this period. From 1984 to 1989, the number of persons on the CPP program, relative to the QPP program, rose by 56,576. Unfortunately, I only have aggregate enrollment data over time for both provinces, so I cannot distinguish the share of this increase due to 45-59 year old men. But assume that this group represented the share of the increase that they represent of the 1993 CPP rolls (30%); the rise for this group was then 16,973 workers. 1.8% of the 45-59 year old male population in the CPP provinces, times a 68% average acceptance rate, is 16,340 workers, which is quite close to this administrative figure.

other change in the CPP provinces, relative to Quebec, that was correlated with the labor supply decisions of older workers. In this section, I consider the two natural alternatives to this identifying assumption. The first is that the policy was itself responding to a trend in relative labor supply across the provinces. That is, perhaps there was an underlying trend towards lower labor force participation among men in the CPP provinces, relative to Quebec, and the policy was passed in response to this trend.

I can test for this underlying trend by pursuing a falsification exercise: reestimating the model on data from four years earlier. That is, I construct a new sample of 45-59 year old men, with data from April 1982 and 1983 as the "before" period, and April 1985 and April 1986 as "after". There was no significant change in DI policy around 1984. Thus, if I estimate the DD model on this data set, and there is a significant positive effect on non-participation, then it suggests that there was a pre-existing trend. If there is no effect, however, it demonstrates that labor supply was moving in parallel in Quebec and the rest of Canada in this pre-policy change period, and that the break in the series arose only when the benefits were increased under the CPP.

The result of this falsification exercise are presented in the first row of Table 4. In fact, there is a small and insignificant positive coefficient. As the second column shows, this coefficient indicates that non-participation rose by 0.3 percentage points in the CPP (relative to the QPP) before the policy change, as opposed to the roughly 2 percentage point increase around the time of the policy change. That is, there was no relative trend before the policy change; the differential between the CPP and QPP grew only after 1987. This timing evidence supports the contention that the policy change caused the relative growth in non-participation, and not the other way around.

Moreover, this finding provides a means of confirming that the contemporaneous change in the early retirement age under the CPP is not driving my results. The effect of this change in retirement age on 45-59 year old is testable because there is a "reverse experiment": Quebec first lowered its retirement age from 65 to 60 in 1984, without changing its DI benefits. As a result, if the early retirement age change is driving the

behavior that we see for 45-59 year olds, there should be a similar change in behavior for this group in Quebec, relative to the rest of Canada, around 1984. But this is exactly the hypothesis that is tested, and rejected, by the falsification exercise; there is no relative change in labor supply across these regions around 1984. This rules out the early retirement age change as an explanation for my finding.

The second alternative is that there was some other contemporaneous change in the relative labor market prospects of older workers in Quebec and the rest of Canada, perhaps due to a relatively faster recovery from the recession of the early 1980s in Quebec. I can assess the importance of contemporary economic conditions in driving my results by making use of a within-region control group: workers aged 25-39. This younger group should be subject to the same economic shocks that affected older workers, but is unlikely to be affected in an important way by changes in DI policy, since the incidence of DI is so much lower for young workers.²⁴ Thus, by rerunning the basic models for this group, I can assess whether there are omitted variables driving the findings.

In fact, as the next two rows of Table 4 show, there is little correlated change in behavior among younger workers. The DD coefficient is positive, but it is fairly small relative to the magnitude for older workers. In the next row, I reestimate the (instrumental variables) parameterized model for this population, assigning to younger workers the benefits for 45-59 year olds in that region/education/year cell. In fact, applying this method to younger workers yields a negative and insignificant coefficient.

Thus, considering the two specification checks together, my finding is that there was a relative change in labor supply of older workers in the CPP provinces, relative to Quebec, that arose only after benefits increased, and that was present only for the older workers to which the program primarily applies (and not for younger workers). That is, the only potential factors which could be confounding my conclusions are sudden changes in the relative economic opportunities or tastes for work of older workers (relative to younger

²⁴The incidence of DI among male workers age 25-39 in is less than 0.2%.

workers), in the CPP provinces (relative to Quebec), around January, 1987.

In fact, there is one further test that can even rule out alternatives in this category: I can explicitly include a CPP*AFTER interaction in the parameterized model, and estimate a "difference-in-difference-in-difference" model (Gruber, 1994) which is identified solely from differences in the effects of this policy change across these 16 groups of workers. That is, this model controls for any changes on average in the economic circumstances or tastes for work of older workers in the CPP regions relative to Quebec, ruling out most plausible alternative explanations for the results. After controlling for average relative changes in labor supply across Quebec and the rest of Canada, this model asks whether the groups that saw the largest replacement rate increase were the groups that increased their non-participation the most.

The results of this estimation are presented in the final row of Table 4, for the IV model (instrumented once again by region*education*AFTER). In fact, the estimated effect here is somewhat larger than in Table 3, indicating an arc elasticity of 0.36; the coefficient is marginally significant. Taken together with the findings for younger workers, this result suggests that other general changes in the CPP provinces relative to Quebec are not driving my estimates.

Part V: Conclusions

A critical parameter for the design of DI policy is the responsiveness of labor supply with respect to benefits generosity. Estimating this parameter in the U.S. context has proved difficult, but the substantial relative benefits rise under the CPP program provides a mechanism for doing so. I do so using both straightforward difference-in-difference models and more parameterized models. In both cases I find a large labor supply effect of the benefits increase: my central estimates imply an elasticity of non-employment with respect to benefits of 0.28 to 0.36.

Is this estimate large or small? There are two benchmarks against which it can be compared. The

first is the previous literature on the U.S. My estimate is closer to the post-Parsons evidence on this elasticity than it is to Parson's estimates, confirming the notion that DI benefits changes along cannot explain the dramatic time series trend among older men in the 1970s.²⁵

Second, this estimate can be compared to the estimated welfare gains from this transfer to the relatively poor off population of disabled. This policy change did not simply distort labor supply decisions; it also potentially offered some benefits to those who now qualified for more generous DI benefit levels. Disability is the kind of large random event for which individuals would ideally hold insurance, but private insurance markets for disability are incomplete. As a result, individuals may suffer a substantial reduction in their standard of living when they become disabled. This is particularly true under the CPP before this benefits increase, where replacement rates averaged only 25% of previous earnings. From the perspective of a social planner, it might be therefore be welfare improving to tax workers somewhat more highly in order to provide a more level consumption stream for those becoming disabled. Thus, while the effects on labor supply were large, it is hard to gauge their importance without reference to the gains to those persons who benefitted from the more generous benefits regime under the CPP.

Computing the welfare gains from this benefits increase in a comprehensive manner is a difficult task, requiring a number of assumptions on the form of the utility function, the social welfare function, and the extent to which other sources of support were crowded out by these benefits increases. In Gruber (1996), I undertook a rudimentary calculation of the social costs and benefits of this policy change. I found that for

²⁵More specifically, from 1960 to 1980, potential DI replacement rates rose by 53% (U.S. Congress Committee on Ways and Means, 1990). At my central elasticity estimates of 0.28 to 0.36, this increase would induce a rise in non-participation of 15 to 19%. But over this time period, as noted above, the non-participation rate of 45-54 year old men rose by over 100%, so that the increase in DI benefits can explain at most only about one-sixth of the increase in non-participation. This does not rule out a role for the DI program per se, since increased program awareness or easing disability standards may have played a stronger role in this era. See Bound and Waidmann (1992) for a more detailed interpretation of these time series trends.

sensible parameter values, so long as the benefits increase was not fully crowding out other sources of support, there were welfare gains from this policy change. While these calculations have some limitations, they raise the possibility that even the substantial distortions to labor supply documented in this paper can be offset by welfare gains when benefits start from such a very low level.

It is also important to note that this analysis has ignored dynamic considerations, so that my findings may misstate the steady state elasticity of response to benefits levels. In particular, by examining behavior for only several years after the benefits change, I may be understating the response if there is some adjustment to this new higher level of benefits. The effect on the long run stock of disabled workers may be substantially larger if there is a now a higher elasticity of labor supply with respect to health shocks which slowly accumulate among older workers in the CPP. On the other hand, my estimated elasticity may overstate the steady state elasticity if there are "announcement effects", whereby large benefits increases affect behavior more strongly than do incremental benefits differences.

Finally, it is interesting to compare the U.S. and Canadian systems in terms of their benefits generosity versus screening stringency. The U.S. has much a higher level of disability benefits than Canada, even after this policy change; Daly (1996) reports that the family income of the disabled in the U.S. is 80% as large as for the non-disabled. At the same time, the much higher denial rates in the U.S. means that a larger share of the disabled population, who may have difficulty working but are denied by the DI program, are living on very low incomes; as Bound (1989) highlighted, fewer than half of those denied by DI return to work, and those that do return do so at much lower wages than before their disability. Moreover, while my finding is roughly in the mid-range of previous estimates of the benefits elasticity, it is much larger than Gruber and Kubik's (1998) estimate of the response of labor supply to denial rates (elasticities of non-participation of 0.12 to 0.17). Given the high effective replacement rates through DI, the poor living standards of denied DI applicants, the much lower elasticity of response to denial rates, and a presumption that the marginal utility

of consumption is declining, a shift in the U.S. towards both lower denial rates and lower benefits seems likely to raise welfare. Exploring the welfare implications of this tradeoff is an important priority for future work.

Table 1: Means

	CPP, Before	CPP, After	QPP, Before	QPP, After	Diff-in- Diff
Benefits	5134	7776	6878	7852	1668 (17)
Replacement Rate	0.245	0.328	0.336	0.331	0.088 (0.003)
Not Employed last week	0.200	0.217	0.256	0.246	0.027 (0.013)
Married?	0.856	0.856	0.817	0.841	-0.024
Any Kids < 17?	0.367	0.351	0.354	0.336	0.002
Less than 9 Years of Education	0.303	0.274	0.454	0.421	0.004
9-10 Years of Education	0.202	0.199	0.179	0.178	-0.002
11-13 Years of Education	0.246	0.254	0.169	0.187	-0.010
Post-Secondary Education	0.249	0.273	0.198	0.214	0.008
Number of Obs	11349	18059	2134	3113	

Notes: Based on author's tabulations. QPP refers to Quebec; CPP refers to the remainder of Canada. Before is 1985-1986; After is 1987-1989. Standard deviations in parentheses.

Table 2: DD Model

Married	-0.952 (0.035)
< 9 Years of Education	1.291 (0.041)
9-10 Years of Education	0.835 (0.045)
11-13 Years of Education	0.390 (0.046)
CPP Region	-0.173 (0.058)
After Policy Change	-0.005 (0.068)
CPP Region *	0.150
After Policy Change	(0.075)
Implied Probability Effect	0.023
Arc Elasticity	0.36
Number of Observations	34655

Notes: Table presents logistic estimation of equation (1) in text. Standard errors in parentheses. Regressions also include full set of dummies for age and number of children.

Table 3: Parameterized Models

Specification:	Not Employed		
	Estimate	Policy Effect	Arc Elasticity
Basic Model	0.927 (0.469)	0.012	0.19
IV Model	1.344 (0.563)	0.018	0.28
Number of Obs	34655		

Notes: Coefficients are those on replacement rate from logistic models such as (2); standard errors in parentheses. Regression includes all of the control variables listed in Table 2, as well as a full set of dummies for number of children, age, year, region, education*region, and education*year. IV model uses as instruments a set of education*region*AFTER dummies. Policy effect is impact of relative replacement rate increase in CPP in 1987; elasticity is percentage change in dependent variable (relative to average of ex-ante and ex-post CPP values) relative to percentage change in replacement rate (relative to average of ex-ante and ex-post CPP replacement rates).

Table 4: Alternative Hypotheses

Specification:	Estimate	Implied Policy Effect	Arc Elasticity
Falsification Exercise: Preexisting Trends?	0.023 (0.080) 28756	0.003	
DD for Younger Workers	0.055 (0.060) 60483	0.007	
Parameterized Model - Younger Workers - IV	-0.303 (0.605) 60483	-0.003	
DDD Model with CPP*AFTER, IV Estimate	1.710 (0.891) 34655	0.023	0.36

Notes: Standard errors in parentheses; number of observations in final row of each cell. First row shows results of a DD regression of the form of (1), with 1982 & 1983 as before, and 1985-1986 as after. Second row shows DD regressions for younger (25-39 years old) male workers; third row shows parameterized model of the form of (2) for this sample. Final row shows regression of the form of (2), but also including a CPP*AFTER interaction; this is IV model, using as instruments a set of education*region*AFTER dummies. Rows (1) and (2) include control variables listed in Table 2 and footnote to that table. Rows (3)-(5) include all of the control variables listed in Table 2, as well as a full set of dummies for number of children, age, year, region, education*region, and education*year. Policy effect is impact of relative replacement rate increase in CPP in 1987; elasticity is percentage change in dependent variable (relative to average of ex-ante and ex-post CPP values) relative to percentage change in replacement rate (relative to average of ex-ante and ex-post CPP replacement rates).

References

- Baker, Michael, and Dwayne Benjamin (1996). "Early Retirement Provisions and the Labour Force Behavior of Older Men: Evidence from Canada". Mimeo, University of Toronto
- Bound, John (1989). "The Health and Earnings of Rejected Disability Applicants," American Economic Review, 79, 482-503.
- 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, 1393-1420.
- Card, David, "The Effects of Minimum Wage Legislation: A Case Study of California, 1987-89," Industrial and Labor Relations Review, 46 (1992), 38-54.
- Daly, Mary (1996). "Characteristics and Income Patterns of Non-Elderly OASDI and SSI Beneficiaries". Memo prepared for the National Academy of Social Insurance Committee on the Privatization of Social Security.
- Gruber, Jonathan, and Maria Hanratty (1995). "The Labor Market Effects of Introducing National Health Insurance: Evidence from Canada," Journal of Business and Economics Statistics, 13, 163-174.
- Gruber, Jonathan, and Jeffrey D. Kubik (1997). "Disability Insurance Rejection Rates and the Labor Supply of Older Workers," Journal of Public Economics, 64, 1-13.
- Halpern, Janice, and Jerry A. Hausman (1986). "Choice Under Uncertainty: A Model of Applications For the Social Security Disability Insurance Program," Journal of Public Economics, 31, 131-161.
- Haveman, Robert H., and Barbara L. Wolfe (1984). "The Decline in Male Labor Force Participation: Comment," Journal of Political Economy, 92:532-541.
- Human Resources Development Canada (1995). "CPP Disability Incidence Study". Ottawa: Income Security Programs Branch, HRDC.
- Human Resources Development Canada (1996). "Statistics Related to Income Security Programs". Ottawa: HRDC.
- Leonard, Jonathan S. (1986). "Labor Supply Incentives and the Disincentives for Disabled Persons," In Disability and the Labor Market: Economic Problems, Policies, and Programs, Ed. Berkowitz, M. and M. A. Hill. Ithaca, NY: ILR Press.
- Lewin-VHI (1996). "Labor Market Conditions, Socioeconomic Factors, and the Growth of Applications and Awards for SSDI and SSI Disability Benefits". Final Report to U.S. Department of Health and Human Services.
- Lumsdaine, Robin and David Wise (1990). "Aging and Labor Force Participation: A Review of Trends and

Explanations," NBER Working Paper 3420, August.

Maki, Dennis R. (1993). "The Economic Implications of Disability Insurance in Canada," Journal of Labor Economics, 11, S148-S169.

Parsons, Donald (1980). "The Decline of Male Labor Force Participation," Journal of Political Economy, 88, 117-134.

Perry, J. Harvey (1984). Taxation in Canada: Fourth Edition. Ottawa: Candian Tax Foundation.

Perry, J. Harvey (1990). Taxation in Canada: Fifth Edition. Ottawa: Candian Tax Foundation.

Social Security Administration (1998). Annual Statistical Supplement to the Social Security Bulletin. Washington, D.C.: SSA.

U.S. Congress, House Committee on Ways and Means (1990,1993, 1998). Overview of Entitlement Programs, U.S. Department of Health and Human Services.