

CENTER FOR LABOR ECONOMICS
UNIVERSITY OF CALIFORNIA, BERKELEY
WORKING PAPER NO. 70

Incentive effects of social assistance: A regression discontinuity approach^{*}

Thomas Lemieux

University of British Columbia and UC Berkeley

Kevin Milligan

University of British Columbia and NBER

May 2004

Abstract

We examine the incentive effects of transfer programs using a unique policy episode. Prior to 1989, social assistance recipients without children in Quebec who were under age 30 received benefits 60 percent lower than recipients older than 30. We use this sharp discontinuity in policy to estimate the effects of social assistance on various labour market outcomes and on living arrangements using a regression discontinuity approach. We find strong evidence that more generous social assistance benefits reduce employment, and more suggestive evidence that they affect marital status and living arrangements. The regression discontinuity estimates exhibit little sensitivity to the degree of flexibility in the specification, and perform very well when we control for unobserved heterogeneity using a first difference specification. Finally, we show that commonly used difference-in-difference estimators may perform poorly when control groups are inappropriately chosen.

^{*} We are grateful to a Statistics Canada post-doctoral fellowship for funding this project. We thank UC Berkeley (Lemieux) and the NBER (Milligan) for hosting the authors while this paper was written. We thank Miles Corak and Garnett Picot for facilitating access to the data. Mike Haan provided skilled assistance with the Census data files. Finally, we thank David Card, Jane Friesen and seminar participants at MIT, McGill, and Berkeley for helpful comments.

1. Introduction

Links are often drawn between labour market behaviour and the generosity and structure of the transfers paid to those not working. For example, the impetus for many of the changes to welfare programs in the United States since 1967 was a concern about disincentives to work embedded in the programs.¹ In Europe, the ‘eurosklerosis’ problem of persistent high unemployment compares unfavourably to the experience in the United States. Blanchard (2004) contends that the ongoing reform of unemployment insurance systems and the introduction of in-work tax credits have improved, but not yet resolved the problems affecting European labour markets. Thus, the strength of the incentive effects of transfer policies continues to be vital to the design of policy and to the understanding of labour market behaviour.

In addition to labour market implications, transfer payments can have broader behavioural influences, such as changing family structure or living arrangements.² Living arrangements are crucial to the understanding of the effects of transfers because economic welfare is usually assessed at the household level. If living arrangements depend on transfer payments, then policy may not lead to the desired distributional consequences.

In our paper, we study the effects of an unusual policy in the province of Quebec that paid much lower social assistance benefits to individuals without children who had not yet attained the age of 30. Fortin et al. (2004) used this policy experiment to estimate the incentive effects of social assistance using a difference-in-differences approach. The break in the policy at age 30 also provides, however, the opportunity to implement a regression discontinuity analysis of the impact of welfare payments on labour market behaviour and living arrangements. This research design holds out the possibility of more

¹ See Moffitt (2003) for a history of welfare programs in the United States. The 1967 reform adjusted tax-back rates because of a concern for labour market incentives.

² This literature on family structure is reviewed in Moffitt (1998). Bitler et al. (2003) provide a detailed literature review of research on living arrangements.

credible inferences about the incentive effects of welfare policies, for reasons we make clear below.

A very large body of research has studied the labour market incentive effects of transfer programs. Moffitt (2002) provides a recent survey of the empirical evidence in the United States, which followed the exhaustive survey of the earlier literature in Moffitt (1992). He concludes that the range of estimates suggests that the counterfactual elimination of welfare would increase hours worked by 10 to 50 percent. Several recent papers have examined the effects of welfare on living arrangements.³ However, because single non-parents can receive welfare in Canada, the most relevant research for our work is the study of the living arrangement of youth in the United States and Canada found in Card and Lemieux (2000). They find that one response of young people to economic distress is to continue (or go back to) living with their parents.

The research strategies chosen over the years to study the effects of welfare have been closely intertwined with the changing policy environment. In the 1970s and 1980s, most research consisted of the econometric modeling of social experiments, such as negative income tax schemes, along with non-experimental econometric evaluations of the incentive effects of welfare. Through the 1980s and early 1990s, the '1115 Waiver' programs generated a second wave of studies, as reviewed in Harvey et al. (2000). With a waiver, states could opt out of certain provisions of the Social Security Act in order to implement demonstration programs or experiments that altered the parameters and structure of welfare programs. The study of these reforms commonly took the form of experimental evaluations, often with treatment and control groups. Finally, the Personal Responsibility and Work Opportunity Reconciliation Act (PRWORA) of 1996 generated a further wave of research attempting to evaluate the effect of reforms in the new decentralized policy environment. Much of the more recent work therefore follows a

³ Hu (2001), Bitler, Gelbach, and Hoynes (2003), and Paxson and Waldfogel (2003) look at the living arrangements of the children of welfare recipients. London (2000), in contrast, examines the arrangements of welfare-receiving mothers.

non-experimental methodology, comparing policy outcomes across states that made different policy choices in the PRWORA era.

Blank (2002) discusses three challenges confronting researchers studying the reforms of the 1990s. First, the economic environment improved dramatically contemporaneous with the reforms. Evaluating a welfare reform in the context of an improving macroeconomy makes it difficult to isolate the effect of the reform from the shifts in labour demand. Second, the dimensionality of the changes makes it difficult to understand the effect of changing one policy, *ceteris paribus*. Reforms were bundled together with some mix of time limits, benefit reduction rates, training, and sanctions, among other policies. Finally, the expansions of the Earned Income Tax Credit also improved the labour market conditions for welfare-at-risk families.

The age-based policy we exploit is able to overcome some of the challenges in the existing literature. The source of the advantage is that we do not study a reform *per se*, but a discontinuity present in a static policy. This means that there is no bundle of other reforms that may contaminate the evaluation of the low benefit policy. Moreover, we do not need to make assumptions about the comparability of the treated group to a control group located in a labour market that is temporally or geographically distinct. This helps us to avoid worries about a changing broad economic environment. Finally, the variation provided by the policy is large – an increase of 145 percent for those reaching age 30. Variation of this magnitude helps to estimate behavioural effects with better precision.

A further advantage was provided by a reform that ended the low benefit policy in August of 1989. By comparing behaviour before and after the change, and in Quebec versus other provinces of Canada, we can also evaluate the policy using a difference-in-differences empirical framework commonly used in the welfare reform literature. This allows us to assess the strengths and weaknesses of the commonly used empirical framework.

One innovative feature of our analysis is that we focus on the effects of social assistance benefits on the labour market behaviour of men without children. We think that for this group, the decision to work or to collect social assistance can be reasonably modelled using a standard labour supply approach. By contrast, employment decisions of single mothers, who are the traditional focus of the U.S. welfare, are complicated by several factors like endogenous fertility decisions and the fixed costs of working in presence of young children.

After providing some institutional details about welfare in Quebec, we describe our data and develop our empirical strategy based on the regression discontinuity approach. We then present a descriptive analysis of that data to provide preliminary evidence on the effects of the policy. Next, we present our regression discontinuity estimates, exploring the sensitivity of the estimates to several robustness and falsification tests. Finally, we compare the regression discontinuity estimates to difference-in-differences results and conclude.

2. Social assistance in Quebec and Canada

Social assistance (as welfare is called in Canada) programs were funded from 1967 to 1996 through the Canada Assistance Plan, which offered a 100 percent matching grant from the federal government for provincial spending.⁴ In contrast to the federally funded welfare programs in the United States during that period, the design of the programs was left almost entirely to the sub-national regions, subject to weak conditions on eligibility.⁵ A distinguishing feature from the case of the United States is the eligibility of singles and non-parents.

⁴ Following 1996, a block grant called the Canadian Health and Social Transfer replaced the Canada Assistance Plan.

⁵ Provinces had to cover all ‘persons in need.’ They could not set eligibility based on province of residence and could only consider the budgetary needs of the person or family, effectively ruling out work requirements. They also agreed to submit statistics to the federal government and set up an appeals process. This discussion is drawn from Baker, Payne, and Smart (1999).

Research on social assistance in Canada has been quite limited. Dooley (1999) describes the trends in social assistance receipt across demographic groups and time. Dooley et al. (2000) find no relationship between female headship and social assistance benefit levels, which is not surprising because benefits are still paid if one does not have children or is married. A large-scale social experiment, the Self-Sufficiency Project, was conducted in the 1990s and paid an earnings supplement to social assistance recipients who found work. The results of the Self-Sufficiency Project are summarized in Ford et al. (2003). Finally, Barrett and Cragg (1998) and Green and Warburton (2004) both use administrative data to study dynamics of social assistance participation in British Columbia.

More closely related to our work, Fortin, Lacroix, and Drolet (2004) study the effect of social assistance benefits on the durations of spells using administrative data. As identifying variation, they use the end of the “under age 30” social assistance rate in Quebec in 1989, comparing recipients over and under age 30 before and after the reform. Our work differs from theirs in a number of ways. First, we study static participation and living arrangements rather than dynamics. Second, using survey data rather than administrative data allows us to look at a broader range of variables and to use residents of other provinces as an additional control group. Finally, we focus our research design closely on the discontinuity of benefits at age 30, rather than making broader comparisons of those under and over age 30. If important unobservable characteristics are correlated with age, then studying behaviour at the discontinuity can improve inferences.

2.1 *Benefits in Quebec*

Social assistance payments in Quebec during the first part of the period we study were governed by the 1969 *Loi sur l'aide sociale* (Social Aid Act). Benefits were paid “. . . on the basis of the deficit that exists between the needs of, and the income available to, a family or individual . . .” Benefits were set periodically by regulation and were kept roughly constant in real terms. The number of children and adults in the family

determined the size of the benefits in a non-linear way, consistent with economies of scale within a family. The regulations also provided for a small income exemption or “disregard” (\$65 in the 1980s), after which benefits were reduced dollar for dollar with income.

The unique feature of social assistance for our purposes is the differential benefit rate by age. Those over age 30 received \$507 per month in 1989 (current dollars) compared to \$185 for those under age 30; a difference of 63 percent.⁶ Only cash benefits differed by age, so items such as subsidized dental care or medical expenses were the same for those over and under age 30. Recipients had to complete a form every month, allowing officials an opportunity to determine if age 30 had been attained. A new *Act Respecting Income Security* received Royal Assent in December 1988 and took effect on August 1st, 1989. The new Act contained a number of changes, including the end of the differential rate at age 30.⁷ We graph the benefit rates in constant 1990 dollars for a single employable person without children in Figure 1, for someone over and under age 30.⁸

The structure of benefits before 1989 is represented diagrammatically in Figure 2 in a static labour supply framework. A thirty year old faces the budget constraint ABCDE, which traces out the social assistance benefit (AB), the earnings exemption (BC), the 100 percent tax rate on earnings (CD), and finally earnings after social assistance has been exhausted (DE). A 29 year old would face a budget constraint described by AFGHE, because of the lower benefit level. A 29 year old with the preferences over consumption and leisure indicated by the indifference curves would choose to work and consume at point X. However, the same preferences for a thirty year old would result in a decision to

⁶ Under section 18 of the Act, discrimination on the basis of “race, colour, sex, religion, language, national extraction, social origin, morals, or political conviction” is not allowed. Age is not mentioned.

⁷ The new Law introduced different rates for those participating in training programs. Since fewer than 10 percent of recipients participated in these programs (Fortin, Lacroix, and Drolet 2004), we focus on the benefits applicable to those who are available for work but do not participate in the training programs. Benefits fell slightly in real terms after the reform for everyone, but no other changes differentially affected those over and under age 30.

⁸ We constructed these series using the benefit rates and indexation methods described in the legislation (as reported in the Revised Statutes of Quebec and the corresponding regulations).

work and consume at point C. The higher benefit levels therefore yield an unambiguous prediction of lower labour supply.

3. Data description and Descriptive Statistics

Most of our analysis relies on data from the 1986 and 1991 Censuses. We also complement our Census numbers with some time-series data from the Labour Force Survey (LFS). For both data sets, however, the selection criteria share common features. We focus our analysis on individuals without a high school diploma (high school dropouts) who are most ‘at risk’ for being on social assistance.⁹ We also focus on respondents without children, as parents of children were not subject to the lower social assistance benefits.¹⁰ The bonus that would be received for bearing a child for those under 30 would be large, but we uncovered no evidence of a fertility response to the policy in the data.¹¹ We discuss these sample selection issues in more detail later.

Finally, the present paper looks at males only. The analysis for females is complicated by a series of factors. First, around age 30, a substantially larger fraction of women than men have children and are not, therefore, subject to the differential benefits.¹² Second, female high school dropouts are much less likely to be employed than men. The employment rate of thirty year old women and male high school dropouts in Quebec in 1986 are 39.5 and 70.4 percent, respectively. For these two reasons, the ‘at risk’ group is

⁹ Recent data from the Institut de la Statistique du Quebec (2004) indicates that 63 percent of all social assistance claimants are high school dropouts. Our own tabulations based on the 1986-89 Survey of Consumer Finance indicates that among childless men age 26 to 35 (the key group affected around the age discontinuity in the program), high school dropouts received 59.7 of social assistance payments, even though they only represented 23.5 percent of the population.

¹⁰ We classified people as “childless” or “without children” when they either do not have children, or have children but do not live with them.

¹¹ The analysis of fertility in the context of Quebec in this era is also complicated by the Allowance for Newborn Children which paid bonuses of up to \$8,000 for a new child. Milligan (2003) finds little evidence of a fertility response among low education and low income women.

¹² Among 30 years old high school dropouts in Quebec in 1986, 75.7 percent of women had (and lived with) children, compared to 53.4 percent for men. Two reasons explain this difference. First, women are much more likely than men to be single parents. Second, women have their children at a younger age than men.

much smaller for women than men. Finally, we are more concerned about possible fertility responses in the case of women than men.

4.1 Census Master Files

The bulk of our analysis is based on the master files of the Canadian Census. Statistics Canada conducts the Canadian Census quinquennially in years ending with a ‘1’ or a ‘6’, in contrast to the decennial nature of the Census in the United States. The coverage of the Census is universal. A detailed questionnaire (long form) is assigned to approximately twenty per cent of households, consisting of questions on labour market characteristics and participation, education, income, and the demographics of respondents. Some of the labour market participation questions are asked with reference to the week previous to Census day, while others refer to the previous calendar year. Because we can observe single years of age in the Census, we can implement our regression discontinuity empirical strategy with these data.

Statistics Canada typically releases a public use microdata file of between 2 and 3 per cent of respondents. As we are interested in obtaining large samples of individuals in narrowly defined cells, we obtained access to the full twenty per cent master sample maintained by Statistics Canada. With this sample, we can form cells of sufficient size for meaningful analysis. For example, Appendix Table 1 shows that we have over 10,000 observations for each year of age in Quebec in the 1986 Census. Since between 26 and 32 percent of these men have not completed high school (column 2), we get samples of around 3,000 high school dropouts for each age group (column 3). The last set of columns in Appendix Table 1 shows that the samples are further reduced when we only keep men without children. We still have, however, over 1,500 observations for each age group around the discontinuity at age 30.

The Census allows us to create a host of interesting variables for analysis. For the reference week prior to census day, we observe whether the respondent was employed, and the hours worked. We can also observe whether the relationship between the

respondent and the head of the household. When the head is a parent or a parent-in-law to the respondent, we code him as living with his parents. For marital status, we code the respondent as married if he is legally married or in a common-law relationship.

Other variables like income by source are measured over the previous calendar year. In particular, the Census collects separate income items for earnings, unemployment insurance (UI) benefits, old age security, CPP/QPP, family allowances, and child tax credits. Unfortunately, the Census does not collect a separate income item for social assistance benefits. These benefits are included in a remaining “other transfers” variable that also includes workers compensation payments, some payments under training programs, and small provincial tax credits claimed on the tax return.¹³

Fortunately, social assistance benefits are by far the largest component of the “other transfers” variable. This is illustrated in Appendix Table 2 that compares the 1986 Census to a pooled sample of the 1985, 1986, 1987, and 1988 Survey of Consumer Finances (SCF). The SCF is a much smaller survey which is, otherwise, quite similar to the Census (Boudarbat et al., 2003). Unlike the Census, however, the SCF collects a separate income item for social assistance payments. Appendix Table 2 shows that social assistance accounts for over 85 percent of “other transfers” (“SA+other” in the table) for the age and demographic groups relevant to our study. As a result, one cannot reject the null hypothesis that all the difference in “other transfers” between men just under and just over thirty is due to differences in social assistance benefits. For the remainder of the paper, we will thus assume that all of the discontinuity in “other transfers” at age thirty is due to the discontinuity in social assistance benefits at age thirty. For all practical purpose, this means that we can use “other transfers” and “social assistance benefits” interchangeably in what follows.

4.2 *The Labour Force Survey*

¹³ Few other items included in the “other transfer” category are either negligible or do not apply for the age group under consideration (e.g. veterans’ pensions).

The Labour Force Survey (LFS) is a monthly national household survey with questions about the labour market behaviour and demographic characteristics of household members, comparable to the monthly Current Population Survey in the United States. The sample size is approximately 100,000 individuals per month, with households staying in the sample for overlapping six month rotations. The population coverage of the LFS excludes residents of the territories, persons living on Indian Reserves, full-time members of the military, and inmates of penal institutions. A comparable set of surveys is available from 1976 to the present. Small provinces are oversampled, necessitating the use of survey weights to calculate representative statistics.

The primary disadvantage of the LFS for our purposes is sample size. The number of Quebecers in the appropriate age range who are high school dropouts is small in any month's sample – typically about 100 males and 50 females are between the ages of 25 and 29. In addition, we do not observe single years of age. Instead, age is reported in 5-year age groups. For these two reasons, the regression discontinuity approach cannot be successfully implemented with the LFS.

We instead exploit the frequency and long availability of the LFS to document the long-term trends in the labour market behaviour of our target population, comparing them across age groups and provinces. Figure 3 graphs the employment rate for males. We use a three-year moving average to smooth the employment rate series that otherwise show erratic movements because of small sample sizes. The top two lines trace the rate for 25-29 year olds and 30-34 year olds in provinces other than Quebec (“rest of Canada” hereafter). The two lines follow the rough contours of the business cycle, rising in the 1980s and falling with the recession of the early 1990s. Two observations are relevant. First, the cyclical nature of the employment rates makes obvious the need to have a control group in order to separate business cycle effects from policy effects. Second, the lines for the two age groups track each other quite closely. This suggests that labour market conditions for these two age groups are comparable.

The second set of lines shows the employment rate by age groups for residents of Quebec. The lines both lie approximately 10 percentage points below those for the rest of Canada, suggesting that any search for policy effects ought to consider differing labour market conditions across regions of the country. The age groups do not track each other as closely in Quebec as was the case for the rest of Canada. In particular, the employment rate of 25-29 years old is substantially larger than the employment rate of 30-34 years old prior to 1990. From 1990 on, however, the employment rates of the two age groups are much more comparable. This is consistent with the view that low social assistance benefits for men under 30 prior to August 1989 lead to a substantial labour supply response.

Other factors could nonetheless account for the abnormally large employment rate of 25-29 years old in Quebec in the late 1980s. Perhaps the strong economic recovery of the second half of the 1980s disproportionately benefited younger workers in Quebec. It is also not clear why the employment rates of 25-29 and 30-34 years old were quite similar in the early 1980s, despite the fact that Social assistance benefits for those under 30 were already much lower back then. For all these reasons, we now turn to a regression discontinuity approach. We later return to a more detailed discussion of how standard difference-in-differences estimates (like those implicit in Figure 3) compare to the regression discontinuity results.

4. Empirical Approach

Our main empirical approach exploits the discontinuity in social assistance benefits at age 30. Consider the regression model:

$$(1) \quad Y_{ia} = \beta_0 + \beta_1 TREAT_{ia} + \delta(a) + \varepsilon_{ia},$$

where Y_{ia} is an outcome variable for individual i of age a . The effect of age on the outcome variable is captured by the function $\delta(a)$, while $TREAT_{ia}$ is a treatment dummy that captures higher social assistance benefits at age 30. It is defined as:

$$TREAT_{ia} = \begin{cases} 0 & \text{if } a < 30 \\ 1 & \text{if } a \geq 30 \end{cases}$$

The evaluation problem consists of estimating the effect β_l of the treatment (higher social assistance benefits) on the outcome variable. The key identification assumption that underlies the regression discontinuity (RD) strategy is that $\delta(\cdot)$ is a smooth (continuous) function.¹⁴ Under this assumption, the treatment effect β_l is obtained by estimating the discontinuity in the empirical regression function at the point where the treatment variable switches from 0 to 1 (age 30 in our case). We have a “sharp” RD design since the treatment variable is a deterministic function of the regression variable (age).

The assumption that $\delta(\cdot)$ is a continuous function means that differential benefits are the only source of discontinuity in outcomes around age 30. How reasonable is this assumption? As is well known, most of our variables of interest like income, employment, and family arrangements exhibit well-known age profiles. For instance, log earnings are a concave function of age, which is consistent with a standard model of investment in human capital (Mincer 1974, Murphy and Welch, 1990). So while it is important to let $\delta(\cdot)$ be flexible enough to accommodate non-linearities in the age profiles, there is no reason (in human capital or related theories of behaviour over the life-cycle) to expect an abrupt change at age 30.

There are, nonetheless, at least two reasons why the assumption $\delta(\cdot)$ is a continuous at age 30 may be violated. First, while the true age of an individual is predetermined, it is conceivable that some people could find ways to “cheat” on their age by, for example, falsifying their birth certificates. If such manipulations were possible, people claiming to

¹⁴ See Hahn, Todd, and van der Klaauw (2001) and Lee (2003) for a more formal discussion of the conditions under which the RD design is as valid as if it were based on a randomized experiment.

be age 30 could be systematically different from those age 29. In particular, people age 29 with a higher propensity to receive social assistance (because of low earnings capacity, etc.) could systematically claim they are 30, thus generating a spurious correlation between age and the error term. This problem is unlikely to occur here since the true age of an individual can be easily verified by social assistance authorities.¹⁵

A potentially more serious problem is that we only select individuals with dependent children for most of our analysis, since only those individuals are subject to differential social assistance benefits. As shown in Appendix Table 1, the fraction of men with children increases steeply as a function of age. To the extent that these fertility and living arrangements decisions (live with your children or not) are endogenous, this generates a problem of non-random selection in our main analysis sample. For instance, we explain in Section 6 that the steep decline in employment rates as a function of age most likely reflects the fact that men without children are an increasingly “negatively selected” group of individuals. As long as these selection biases are a smooth function of age, however, they will be captured by the function $\delta(\cdot)$ and the RD approach will remain valid.

The RD approach may not be valid, however, if the decision to have children and live with them was itself influenced by social assistance benefits. For instance, an unemployed man living with his wife and children could decide to leave home once he turns 30 because he can now get much higher social assistance benefits as a “single”. Appendix Table 1 and Appendix Figure 1 show, however, that there is no evidence of a discontinuity at age 30 in the fraction of men with children in Quebec in 1986. In fact, the increase in this fraction between age 29 and 30 is essentially identical to what is observed in situations where there is no discontinuity in social assistance benefits at age 30 (Quebec in 1991, Rest of Canada in 1986 or 1991). We also present some additional results below where we estimate our models for all men instead of conditioning on men without children. Using all men “solves” the selection problem but leads (presumably) to

¹⁵ Note that it was relatively easy to falsify one’s age in Quebec in the 1980s since baptismal certificates issued by local parishes were used as official birth certificates (and proof of identity). By the time individuals were in their late 20s, however, their “official” birth date had long been recorded by tax, social insurance, citizenship (passport) and other government authorities. It is thus highly unlikely that more than a handful of individuals managed to get higher social assistance benefits by cheating on their age.

a smaller estimated treatment effect since we now add a group of individuals known to be unaffected by the differential benefits (men with children) to the main analysis sample of men without children.

In practice, the estimated treatment effect depends on how the smooth function $\delta(\cdot)$ is itself estimated. One possible route is to estimate $\delta(\cdot)$ using non-parametric methods, with the usual trade-offs in the choice of the bandwidth. When a very small bandwidth is used, the estimate of β_l ends up being the difference in the mean value of the outcome variable just to the right and just to the left of the discontinuity point. But unless very large amounts of data are available, such estimates may be very imprecise. With a larger bandwidth, however, a bias can be introduced if people further away from the exact discontinuity point are systematically different from those at the discontinuity point.

We balance this trade-off between precision and bias by estimating a variety of polynomial specifications for the regression function $\delta(\cdot)$. In Section 6, we present estimates of the treatment effect using five different specifications for the regression function. The specifications include standard linear, quadratic, and cubic functions, as well as linear and quadratic splines (separate regressions on both sides of the discontinuity).

We also need to adapt our RD approach to some of the data limitations discussed in the previous section. One problem is that we only know the age in years of the individual at census day (typically the first week of June). This means that the best we can do is to compare all individuals age 29 on census day to all individuals age 30 at census day. In other words, we cannot compare people who “just turned 30” to people “just about to turn 30”.

Because of this data limitation, all the information available in the micro data can be summarized in the age-specific means of the variables (sufficient statistics). The empirical model we work with is the age-cell version of equation (1):

$$(2) \quad Y_a = \beta_0 + \beta_1 TREAT_a + \delta(a) + \varepsilon_a.$$

Regression estimates of equation (1) based on micro data are identical to weighted estimates of equation (2) when the weight used is the number of observations by age group.

Another advantage of working with age cells is that is straightforward to test how well the model “fits” the data. Since the outcome variable Y_a is a cell mean, its sampling variance V_a can be readily computed. Under the assumption that model (2) is correct, the only source of error in the model should be the sampling error. This assumption can be tested using the goodness-of-fit statistic

$$GOF = \sum_a (\hat{\varepsilon}_a^2 / V_a).$$

Under the null hypothesis that model (2) is well specified, GOF should follow a chi-square distribution with $N-k$ degrees of freedom.

Up to now, we have implicitly assumed that the outcome variable Y was measured at the time of the Census. As discussed in the previous section, some of the outcome variables like current employment and hours of work, marital status, and family arrangements are indeed measured at the time of the census. However, other variables like transfer income, earnings, and weeks worked pertain to the previous year. As a consequence, the regression discontinuity is not “sharp” for these outcome variables.

To see this, consider the receipt of social assistance transfers in the previous year. Take the case of an individual age 30 at census day who turned 30 on the first of December in the previous year. This individual was thus “exposed” to higher social assistance benefits for only one of the twelve months during the previous year. We deal with this problem by assigning $1/12$ to the treatment variable for this specific individual.

If we knew the exact birth date of individuals, we could use the fraction of the previous year during which the individual was age 30 as the treatment variable. The treatment variable $TREAT_{ia}$ would be equal to zero for all individuals age 29 or less at census day, one for all individuals age 32 or more at census day, and a number ranging from zero to one for those age 30 or 31 at census day (depending on their exact birth date).

Since we only know the age in years at census day, we need to average $TREAT_{ia}$ over all individual of a certain age. We do so by assuming that Census day is June 1st and that birth dates are uniformly distributed over the year. Under those assumptions, it is easy to show that the average treatment $TREAT'_a$ takes the following values for the different age groups:¹⁶

$$(3) \quad TREAT'_a = \begin{cases} 0 & \text{if } a \leq 29 \\ 0.170 & \text{if } a = 30 \\ 0.913 & \text{if } a = 31 \\ 1 & \text{if } a \geq 32 \end{cases}$$

By contrast, in the models for outcomes at the time of the Census, $TREAT_a$ is simply 0 for all individuals age 29 or less at census day, and 1 for individuals age 30 or more.

One concern is that some of the advantages of the RD design are lost because we do not have a sharp discontinuity for the outcomes variables measured over the previous year. Fortunately, it is possible to test for the impact of this shortcoming when looking at employment. In the Census, we know both the employment status in the reference week, and the number of weeks worked in the previous year. For a given age group, we can construct an employment rate in the Census reference week, ERC_a , and an employment rate based on the fraction of weeks worked in the previous year, ERL_a .

¹⁶ The values of the treatment variable $TREAT'_a$ for age 30 and 31 are obtained by integrating over the uniform distribution of birth dates. It can be shown that for age 30 we get $TREAT'_a = .5(7/12)^2 = 0.170$. For age 31 we get $TREAT'_a = 1 - .5(5/12)^2 = 0.913$.

We can thus compare the “sharp” RD results based on the analysis of the outcome variable ERC_a , to the “fuzzy” RD estimates based on the variable ERL_a . We find that both specifications give very similar results (Section 6), which suggests that the RD approach yields valid estimated treatment effects despite the “fuzziness” introduced in outcome variables measured over the previous year. More specifically, the model for the employment rate on census week is

$$(4a) \quad ERC_a = \beta_0 + \beta_1 TREAT_a + \delta(a) + \varepsilon_a,$$

while the model of the employment rate in the previous year is

$$(4b) \quad ERL_a = \beta'_0 + \beta'_1 TREAT'_a + \delta'(a) + \varepsilon'_a.$$

We can then compare the alternative estimates of the treatment effect β_1 and β'_1 . The two estimates should be the same if the models are well specified. If the labour supply impact of social assistance benefits is large, the employment rate at census week (equation 4a) should drop sharply between age 29 and 30, as $TREAT_a$ jumps from 0 to 1. By contrast, most of the drop should occur between age 30 and 31 in the model for the employment rate in the previous year (equation 4b) since, according to equation (3), $TREAT'_a$ increases from 0.170 at age 30 to 0.913 at age 31.

This suggests another estimator of the treatment effect based on the difference between the two employment rates, which is in fact the change in the employment rate between the previous year and the Census reference week. If individuals truly reduce their labour supply once social assistance benefits become more generous, the employment rate of 30 years old (on census week) should be unusually low compared to their employment rate in the previous year (when they were mostly 29).

This alternative estimator is essentially a first-difference (FD) estimator that exploits the longitudinal nature of the information about employment in the census. Under the

assumption that $\beta_l = \beta'_l$, this FD-RD estimator is obtained by estimating the regression model

$$(5) \text{ERC}_a - \text{ERL}_a = (\beta_0 - \beta'_0) + \beta_1(\text{TREAT}_a - \text{TREAT}'_a) + \theta(a) + (\varepsilon_a - \varepsilon'_a),$$

by (weighted) OLS. Note that $\theta(a)$, the difference between $\delta(a)$ and $\delta'(a)$, is once again a smooth function of age that can be captured by the same functions as before. As in a standard FD model, one advantage of this model is that individual-specific fixed effects are eliminated by taking differences in the error term in equation (5).

The RD estimator is based on the assumption that people close to the discontinuity are similar. While the assumption is highly plausible in our case, it usually remains untestable at some basic level. Perhaps people just above 30 are different from those age 29 for some unmodelled reason. The FD-RD estimator goes one step further by comparing the employment of the *same* individuals at age 29 and 30.

Taken together, the quasi panel nature of the Census (for employment behaviour) and the discontinuity in social assistance benefits at age 30 provide a variety of estimation strategies that can be used to validate our basic RD research design. In Section 6, we present these alternative estimators and argue that the RD estimates of the impact of social assistance benefits on employment are indeed very robust across estimation methods. This gives considerable confidence in the RD estimates for other outcomes of interest.

5. Cross-sectional age profiles

Before turning to the RD estimates, we first graph a host of outcomes against age, looking separately at Quebec and the rest of Canada over the 1986 and 1991 censuses. In principle, all we need in our RD design are the data from Quebec in 1986. It is

nonetheless useful to see whether the raw data confirm the basic prediction that age profiles are discontinuous around age 30 in Quebec, but not in the other cases.

Figure 4 presents employment rates (reference week) by age, from 20 to 39. Like all other figures discussed in this section, Figure 4 is based on the sample of male high school dropouts with no dependent children. Separate lines are drawn for Quebec and the rest of Canada, and for the 1986 and 1991 Censuses. A vertical line marks the split between those under and over age 30. The top two lines are for the rest of Canada. After rising 9.6 percentage point to 70.5 percent at age 23 for 1986, the employment rates are generally flat with a slight downward trend. The smaller sample size in Quebec adds more sampling variation to the Quebec lines, but a striking change in the relative position of the 1986 and 1991 lines is evident at age 30. The drop at age 30 in 1986 is 5.2 percentage points. After age 30, both Quebec lines trend downward.

A very similar pattern can be seen for hours worked in the reference week in Figure 5. Between ages 23 and 29, hours worked in Quebec in 1986 is constant at around 26 hours per week. At age 30, there is a dramatic drop to 24 hours per week, a decrease of 7.2 per cent. Together, these two figures suggest that most of the variation in labour market participation for this sample of males is on the extensive margin.

Figures 6 and 7 explore two measures of living arrangements. The four lines tracing out the proportion of respondents living with their parents in Figure 7 are virtually on top of each other across all ages. The rate falls from around 70 percent at age 20 to around 20 percent at age 39. Clearly, the sharp discontinuity in social assistance benefits at age 30 in Quebec in 1986 appears to have little impact on this dimension of living arrangements.

The proportion of respondents who are legally married or in a common-law partnership is graphed in Figure 7. After age 30, the four lines are close to each other and constant just under a rate of 30 percent. However, before age 30 the line for Quebec in 1986 shows an increasing gap, reaching 4.3 percentage points relative to Quebec in 1991 at age 29 before falling to near zero at age 30. Breaking the data into separate analyses (not shown

here) for legally married and common-law partnerships reveals that much of the pattern is driven by legal marriages. This may indicate that single males were more willing to enter into a marriage when the social assistance rate was low than when it was high. Furthermore, the absolute drop in being married at age 30 may indicate that these social assistance-induced marriages did not persist once the male had the possibility of a higher government transfer payment at age 30. Moreover, the age 29-30 jump in the proportion of respondents who report being separated or divorced was 2.9 percent in 1986, compared to only 1.6 in 1991, providing further suggestive evidence that social assistance had some impact on marital choices.

The next set of figures displays results from variables based on income data from the calendar year prior to Census day. As discussed earlier, only those who are age 32 or higher on Census day spent the entire previous calendar year over age 30. For this reason, we draw an extra line in the figures between ages 31 and 32. The ages between the lines correspond to ages at which some time was spent at age 29 and some at age 30 in the previous year.

The first income graph in Figure 8 shows the dollar value of “other government transfers” (in 1990 Canadian dollars). As discussed earlier, this variable mostly captures social assistance benefits. Before age 30, the 1986 and the 1991 lines for Quebec follow each other very closely. By age 32, a large gap between them opens. By contrast, social assistance receipts only grow slowly as a function of age in the rest of Canada in either 1986 or 1991.

Figure 9 shows the level of earned and self-employment income. In both Quebec and the rest of Canada, the age-earnings profile grows steeper between 1986 and 1991. This shift is more prominent in Quebec where men age 32 to 39 earn less in 1986 than in 1991, while men age 25 to 29 earn more in 1986 than 1991. A natural explanation for this pattern of results is that younger men (age 25-29) in Quebec worked more (Figures 4 and 5) and earned more in response to the very low social assistance benefits that prevailed in

1986. Note, however, that there is not a sharp decline in earnings between age 29 and 32 that mirrors the sharp increase in social assistance receipts documented in Figure 8.

6. Regression Discontinuity Estimates

We now formally exploit the discontinuity in social assistance benefits by estimating the RD models discussed in Section 4. After several experiments, we decided to limit our analysis to men age 25 to 39. The reason for this choice is that the age profile in most of the variables in Figures 4 to 9 is systematically different between age 20 and 24 than between age 25 and 29. This suggests that data for age 20 to 24 are of little use for helping to fit the model around the discontinuity point.

Note also that all the regression models are estimated by (weighted) OLS using the inverse of the sampling variances (V_a) as weights. The resulting estimates are very similar to those obtained using the number of observations in each age cell as weights. The advantage of using the inverse of the sampling variances instead is that the sum of square residuals is equal to the goodness-of-fit statistic GOF (up to a normalization).

6.1 *Employment Effects*

Table 1 shows the estimated treatment effects for the labour supply variables in Quebec in 1986. Column 1 shows the RD estimates for the employment rate in the previous year (1985). This model corresponds to equation (4b) in Section 4. The employment impacts are precisely estimated in the first four models, but less precisely estimated when the richest model, the quadratic spline, is used.

The results are even stronger in the model for employment at Census week reported in column 2. In this model, the employment effect remains precisely estimated even when the quadratic spline is used (the most flexible model). Remember that we have a sharp discontinuity in this latter model, while the discontinuity is not sharp in the model based

on the employment rate in the previous year. This may explain why the effect of social assistance is more precisely estimated for employment at census week in the more flexible models like the cubic and the quadratic spline.

One nice feature of the results is that the two employment rate measures yield remarkably similar estimates. This suggests that the RD approach is appropriate for the models of previous year outcomes despite some of the data shortcomings discussed in Section 4. Note also that the goodness-of-fit tests suggest that even the simpler models (linear or linear spline) fit the data very well.

To get a better sense of how the models fit the data, we compare the predicted regression models to the actual data for the two employment measures in Figures 10 and 11 for the linear spline models. In the case of the employment rate in the Census reference week, we place the discontinuity point at age 29.5. Since people coded as “age 30” on census day are 30.5 years old, on average, we need to move the discontinuity point by half a year to get people who are exactly age 30 on census day. In the case of employment in the previous year, we place the discontinuity point at age 30 and $5/12^{\text{th}}$ for similar reasons. In all the “previous year” models, we both show the linear regression lines (solid lines predicted by the linear splines) and the actual fit obtained using the *TREAT* variable (dotted lines).

Both Figures 10 and 11 present strong evidence that employment drops abruptly once individuals become eligible for the higher social assistance benefits. As expected, the decline in employment measured at census week happens between age 29 and 30, while the decline in employment measured over the previous year (Figure 11) mostly happens between age 30 and 31. Interestingly, the estimated employment effect of the higher social assistance benefits is almost identical for the two measures of the employment rate in the linear spline models illustrated in Figures 10 and 11. Table 1 shows that the estimated effects are -0.047 and -0.049 for employment last year and in the reference week, respectively.

As discussed earlier, an even more stringent test of the disincentive effects of social assistance on labour supply is based on the difference between the two employment measures. The FD-RD estimates of equation (5) are reported in column 3 of Table 1. The estimated employment effects are very robust across specifications and tend to be a bit smaller than the standard RD estimates reported in columns 1 and 2 of Table 1.

Remember that the key group used to identify the FD-RD estimates are individuals age 30 at the time of the Census. Since these individuals were mostly 29 in the previous year, we should see their census week employment drop relative to their previous year employment as they become exposed to the higher benefits after turning 30. By contrast, all other age groups (except for a few of the 31 year olds) are exposed to the same social assistance benefits at census week and in the previous years. Figure 12 confirms this prediction that the employment rate difference is abnormally low for individuals age 30 at the time of the Census. The figure also shows that the regression fit based on the difference model (solid line) is quite similar to the fit implied by the two models for employment levels (dotted line defined as the difference between the regression lines in Figures 10 and 11).

The last column of Table 1 shows that the effect of higher social assistance benefits on hours of work at census week (including zeros) is similar to the estimated effect on the employment rate. The estimated effect on hours in the linear spline model (-.1.72) represents about 7.1 percent of average hours of work (24.39). This is very similar to the 7.9 percent effect on employment probability obtained for the most comparable employment rate model (linear spline model for employment at census week). The results suggest that all of the impact of social assistance benefits on labour supply happens at the extensive margin (participation) as opposed to intensive margin (hours of work conditional on employment), which is consistent the model presented in Figure 2.

We run a series of “falsification experiments” in Table 2 to present further evidence on the robustness of our findings. Since there is no discontinuity in social assistance benefits in Quebec in 1991 or in the rest of Canada in either 1986 or 1991, RD estimates

for these alternative samples should not show significant employment effects. Table 2 indeed indicates a sharp contrast between Quebec in 1986 where employment effects are always significant at the 95 percent level (except in one case where it is significant at the 90 percent level), and other regions or years where employment effects are generally insignificant. The contrast is particularly striking for the FD-RD estimates (second panel of Table 2) where only one of the fifteen estimates for other regions or years is significant (at the 90 percent level). Interestingly, estimates based on employment at Census week (first panel) are more erratic. They are even statistically significant in a number of cases.

One interesting methodological finding is thus that the availability of panel data (quasi panel in the Census) does indeed make the RD estimates more reliable. This is illustrated in the last column of Table 2 that shows the difference-in-differences in the RD estimates.¹⁷ These difference-in-differences estimates range from -.055 to -.098 for the regular RD estimates, but only from -.048 to -.053 for the FD-RD estimates. Putting all these results together, we conclude that higher social assistance benefits reduce the employment rate by at least three percentage points, and perhaps as much as five percentage points. Furthermore, the similarity in the results for the different employment specifications suggest that the RD approach “works” despite the fuzziness introduced in the models based on the reporting of outcomes over the previous calendar year.

6.2 *Other outcome variables*

Table 3 shows the estimated effects for a variety of other outcomes variables. The first column reports the estimates for total social assistance income based on the “other transfers” variable in the Census. The results indicate a precisely estimated effect in the range of \$450 to \$500 per year for the different specifications. The second column shows that the effect on total transfer dollars comes from both a higher take-up rate of social assistance (column 2), and higher social assistance receipts conditional on receiving

¹⁷ The rationale for such an estimator is that there may be “spurious” region and year specific discontinuity at age 30 that can be removed by taking difference in differences.

positive transfers. Both of these effects are precisely estimated and robust across specifications.

More importantly, the magnitude of the estimated effects is consistent with other results presented in the paper. For example, we find that the higher social assistance benefits increase the take-up rate of social assistance (column 2) by about 4 percentage points. This is right in the range of employment effects (3 to 5 percentage points) documented in Table 1. Once again, this is consistent with the labour supply model of Figure 2 that predicts all workers who quit employment in response to higher social assistance benefits end up receiving social assistance payments.

The results in column 3 indicate that people on social assistance experience, on average, a \$1,200 increase in social assistance receipts when they become eligible for the higher benefits at age 30. This is considerably smaller than the roughly \$3,300 annual increase in benefits that an individual on social assistance for a full year should experience after turning 30.¹⁸ As is well known, however, people move in and out of social assistance and typically spend less than a full year on social assistance. There is unfortunately no information on the number of months an individual spent on social assistance in the Census. Fortin et al. (2004) show, however, that the median spell of social assistance for men age 25 to 29 lasted between 6 and 9 months in the 1980s and early 1990s. Furthermore, existing validation studies suggest that social assistance (welfare) receipts are underreported by a factor of about a third in standard government surveys.¹⁹ Correcting the \$3,300 figure for underreporting and the fact that people do not typically spend their whole year on social assistance yields an expected effect much closer to \$1,200.²⁰

The social assistance results for the linear spline models are illustrated in Figures 13 to 15. In all three cases, there is clear visual evidence of a discontinuity around age 30.

¹⁸ The difference in monthly benefits in 1985 is about \$280 in Figure 1.

¹⁹ See Kapsalis (2001).

²⁰ Some back-of-the-envelope calculations based on the numbers reported in Fortin et al. (2004) suggest that social assistance claimants spend, on average, about 7 months on social assistance during a calendar year. \$3,300 multiplied by 7/12 and 2/3 is equal to \$1,283, which is very close to the estimated effect.

Note that both the total social assistance money transfers (Figure 13) and the take-up rate of social assistance (Figure 14) are trending up as a function of age. The explanation for this upward trend is, once again, that men without children are negatively selected in terms of their labour market prospects, and that the magnitude of the bias increases as a function of age.²¹ By contrast, the total dollar value of social assistance benefits conditional on receiving some benefits (Figure 15) is roughly a constant function of age except for the discontinuity at age 30. This is consistent with the administrative regulations of social assistance that do not link benefits to age, except for the differential benefits for individual under the age of 30.

The last set of columns of Table 3 show the impact of social assistance benefits on a few other outcomes. Column 4 shows that there is a negative but not statistically significant effect of higher social assistance benefits on the amount of unemployment insurance (UI) benefits. This suggests, at best, weak substitution effects between social assistance and UI.

Column 5 shows that higher social assistance benefits generally a negative impact on annual earnings (including zeros). This is consistent with expectations since earlier results show significant impacts on employment. However, the effect is imprecisely estimated and not statistically significant (except in the quadratic model where it is significant at the 90 percent level). In fact, the standard errors are too large to make it possible to distinguish among some reasonable null hypotheses. One first hypothesis is that workers affected by the higher benefits are representative of all workers. The means in Table 1 and 3 indicate that people earn, on average, about \$11 an hour.²² Given the estimated effect on weekly hours in Table 1 (1.72 for the linear spline model), the expected effect on annual earnings is $\$11 \times 1.72 \times 52 = \980 , which is basically the same as the estimated coefficient in Table 3 (\$975). By contrast, if workers affected at the margin

²¹ This interpretation of the evidence is based on the fact that social assistance does not trend up as a function of age when all high school dropouts, as opposed to only those without children, are used in the analysis.

²² This is obtained by dividing average earnings (\$13,924) by average weekly hours (24.39) times 52 weeks.

are earning the minimum wage (\$4), the expected effect would be \$360. Unfortunately, \$980 is not statistically different from \$360 because of the large standard errors.

Another interesting hypothesis is that workers who drop out of the labour force are the ones who previously earned more than the lower “under age 30” benefits, but now earn less than the higher “over age 30” benefits and decide to drop out of the labour market. The implication of this (Ashenfelter, 1983) model of program participation is that the decrease in earnings should be smaller than the increase in social assistance payments of about \$500 (column 1 of Table 3). Once again, this hypothesis cannot be rejected because of the imprecision of the earnings estimates.

The last column of Table 3 shows that, consistent with Figure 7, higher benefits have a negative effect on the probability of being married. One possible explanation is that people have incentives to live together and share expenses when forced to live on the low social assistance benefits, but then can afford to live alone once they become eligible for the higher benefits at age 30.

Finally, Table 4 shows the “falsification test” for total social assistance payments in Quebec in 1991 and in the Rest of Canada in 1986 and 1991. As in the case of the employment models reported in Table 2, there is generally no significant discontinuity in social assistance payments except in Quebec in 1986.

6.3 Broadening the target group

As discussed in Section 4, one concern with our main estimation results is that we may be creating a sample selectivity problem by only looking at men without children. We are not too concerned with this problem since Appendix Figure 1 shows no evidence of a discontinuity at age 30 in the fraction of men who have children. As an additional piece of evidence, we show the employment rate (at census week) for *all* high school dropouts in Appendix Figure 2. The figure shows once again a sharp drop in employment between age 29 and 30 in Quebec in 1986, but no comparable change in the other regions or years.

Running the linear spline model yields an estimated employment effect of -0.030 (standard error of 0.008). As expected, this is smaller than the corresponding effect for childless men only (-0.049) since we are now also including men with dependent children who are not affected by the differential benefits. In fact, since about half of the men around age 30 have dependent children, the estimate for the broader sample should be about half of the estimate for the larger sample, which is consistent with our findings.²³

Appendix Figure 3 shows the same graphs for all men, irrespective of their level of education and of the presence of children. Note that each data point in the figure is now based on very large samples of more than 10,000 individual per year of age in Quebec, and around 30,000 individuals in the rest of Canada. Twenty percent of the Canadian population at each age is used to compute these employment rates. Perhaps surprisingly, there is still clear evidence of an abrupt decline in the employment rate at age 30 in Quebec in 1986. The point estimate in the linear spline model is about one percentage point (-0.012, standard error of 0.06). This is roughly a third of the estimate for high school dropouts only (Appendix Figure 2). Once again, this is consistent with our expectations since the fraction of high school dropouts receiving some social assistance is about three times as large as the corresponding fraction for all individuals (footnote 8).

The robustness of our findings to the choice of sample and estimation method gives us considerable confidence in our findings that more generous social assistance benefits have an adverse impact on employment. While the impact is relatively modest for the whole population, we find quite substantial impacts for the group most affected by the differential benefits (high school dropouts with no dependent children).

7. Comparing RD and Difference-in-Differences results

²³ Accounting for the standard errors, we cannot reject the null hypothesis that the estimate for the broader sample, -0.030, is a half of the estimate for the narrower sample, -0.049.

In this section, we compare the results from the regression discontinuity estimator to results from more traditional difference-in-difference estimators. In the PRWORA era, many researchers have pursued difference-in-difference strategies to measure the effect of welfare reform across states and years.²⁴ Our goal is to assess the effectiveness of these commonly used methods in estimating treatment effects.

The cancellation of the low benefit policy in 1989 makes a pre- and post-1989 comparison natural. In addition, we can use the other provinces in Canada as a control for any common economic shocks hitting the entire country. We present statistical tests featuring comparisons with the two additional control groups, using both cell means and regressions to draw inferences.

8.1 Difference-in-difference estimators

We start with the group of Quebecers at age 29 in the 1986 Census. The first comparison we draw looks at Quebecers age 30 in 1986. This discontinuity cell mean estimator will be unbiased if there is no trend in the age profile for the dependent variable. Reflecting on the figures discussed in Section 6, this assumption seems reasonable for several of the labour supply variables, but less so for the living arrangements measures.

To account for age effects, three different control group strategies may be pursued. First, comparisons can be made to the age 29 versus age 30 gap in the rest of Canada in 1986. If the age profile of the dependent variable is the same in Quebec and the rest of Canada, this estimator will be unbiased. Second, we can use Quebec in 1991 as a control group. This control group will produce unbiased if the age profile for the dependent variable is unchanged through time. As 1991 saw the onset of a sharp recession, this assumption will not hold if younger labour market participants were differentially affected by the recession. Finally, we can compare Quebecers at age 29 in 1986 to those age 29 in 1991, then compare the result to the same difference in the rest of Canada. To be unbiased, this

²⁴ See the discussion in Blank (2002), or the extensive literature cited in Bitler et al. (2003).

estimator requires the assumption that the 1991 recession had the same impact on the behaviour of residents of Quebec as the residents of other provinces.

The ultimate cell mean estimator combines all the control groups into one triple-difference estimator. The 1986 difference in the age 29-age 30 gap between Quebec and the rest of Canada is compared to the same difference in the 1991 Census. This estimator is unbiased so long as the age profile of the dependent variable does not shift differentially between 1986 and 1991 in Quebec and the rest of Canada.

The final estimator we present places the triple differences in a regression framework. For the sake of consistency with the RD analysis, we define an indicator variable $TREAT$ which takes the value 0 for observations at age 29 in Quebec in 1986 (low benefits), and 1 otherwise (high benefits). We then create an indicator variable for age (α_a), an indicator for residing in the province of Quebec (π_p), and an indicator for observations from the 1986 Census (τ_t). We include these indicators, along with their second order interactions in the regression equation. Finally, we include a vector of observable characteristics X_{apt} for each observation. The extra control variables included are dummies for completed years of education, dummies for mother tongue (French, English, and other), a dummy for living in an urban region, and a dummy for being born outside Canada.

The β_l coefficient is the regression analogue to the cell mean triple difference estimator. The dependent variables (Y_{apt}) are regressed on all of the explanatory variables by ordinary least squares using the following equation, with a indexing age, p indexing province, and t indexing time.

$$Y_{apt} = \beta_0 + \beta_1 TREAT_{apt} + \beta_2 \alpha_a + \beta_3 \pi_p + \beta_4 \tau_t + \beta_5 \alpha_a \times \pi_p + \beta_6 \alpha_a \times \tau_t + \beta_7 \pi_p \times \tau_t + \beta_8 X_{apt} + \varepsilon_{apt}.$$

8.2 Results from the Census

Table 5 reports the results for employment in the reference week and for having positive social assistance transfers in the previous year (take-up rate). At the top, we show the mean of each dependent variable for the treatment group, 29 year old Quebecers in 1986. The next five rows display cell means for the treatment group, compared to different control groups. Finally, the last row reports the regression coefficients for the triple-difference regression.

The first column has the results for being employed in the reference week. The discontinuity cell means estimator compares the mean of 29 year olds to 30 year olds. The estimated effect is a strongly significant 5.2 percentage points, which is 7.9 per cent of the mean. The next row shows the effect of a difference in differences estimator using the age 29 – age 30 difference in the rest of Canada as a control for any age-related trend in employment. The difference between employment for those ages in the rest of Canada is a small and insignificant 0.7 percentage points, leading to little change in the estimated effect relative to the discontinuity measure.

The next control group is Quebecers from the 1991 Census. The age gap for employment in 1991 was 2.7 percentage points in favour of those age 30, compared to 5.2 percentage points in the other direction in 1986. Looking at the graph for employment in Figure 4 for the rest of Canada, it appears that the 1991 recession had a differential impact on younger males relative to older ones. To the extent this macro shock influences the age 29 – age 30 employment gap in Quebec, the resulting difference-in-differences estimator may be biased.

Restricting attention to 29 year olds only provides the next difference-in-differences comparison. Here, the estimated impact of the policy is an insignificant 1.0 percentage points. Why is this estimate so different than the others? To be unbiased, the estimator requires that the 1991 recession have the same impact on employment in Quebec and in the rest of Canada. Hoynes (2000) and Black et al. (2003) provide evidence that local economic conditions influence welfare expenditures. To the extent that conditions differ

in Quebec and the rest of Canada, using the rest of Canada as a control may be a poor choice.

The final cell mean result for employment in Table 6 is the triple-difference estimator. The estimated effect of the policy was a decrease in employment of 7.9 percentage points, or 12.0 percent of the mean. Directly beneath is the corresponding regression estimator, which shows a similar negative response of 7.4 percentage points. Regression estimators analogous to the other cell mean estimators also showed little change in the estimated impact of the policy. Since the extra control variables are unlikely to be strongly different in the age 29 and age 30 groups, this is unsurprising.

The next set of results in Table 6 uses the presence of transfer income as the dependent variable. The discontinuity cell mean estimator suggests an effect of 5.8 percentage points. However, Figure 8 makes clear that transfer receipt trends up with age, so the comparison of 29 to 32 year olds may be upward biased. The next two estimators control for the upward age trend using the rest of Canada and Quebecers in 1991 as control groups, respectively. The estimated impact here is lower, at 3.9 percentage points and 4.0 percentage points. These estimates are economically large, representing 21 per cent of the mean for this variable.

The analysis of difference-in-differences results in this section has shown that additional control groups do not necessarily improve on the regression discontinuity estimator. In particular, without a control group placed in the same labour market as the treatment group, the difference-in-difference estimates can diverge greatly from the regression discontinuity estimates.

8.3 *Comparing to the LFS*

One difficulty with using Census data for performing a difference in differences analysis is the long time lag (five years) between each Census. Ideally, a difference in differences analysis should compare outcomes just before and just after the change in policy. To do

so, we turn back to the LFS data used in Figure 3. Table 6 presents cell means and cell differences using a window of 36 months on either side of the policy change in August 1989.²⁵ We have to use these relatively wide windows because of small sample sizes. The first row shows the mean of the dependent variable for residents of Quebec in the 25-29 age group (low benefits group) in the 36 months before the reform. The employment rate for this group is 61.6 per cent. We also look at another outcome variable available in the LFS, the fraction of people still living with their parents. The second column of the table shows that 42.4 per cent of men in this group live with their parents.

The next three rows of numbers show the results from differencing estimators, comparing the cell mean in the treatment group to those in control groups. Below each difference is the standard error. The top difference estimator compares the two age groups in Quebec in the 36 months before and after the reform. This estimator will produce an unbiased estimate if there is no time trend difference in the behaviour of the two age groups over this period. The estimated mean decrease in the employment rate is 5.6 percentage points, which is a 9.1 per cent decrease on the mean value.

The next differencing estimator uses the rest of Canada as a control group for the difference between the age groups in Quebec, just looking at the period before the reform. This estimator will produce an unbiased estimate if there is no fixed difference in the age gap in Quebec versus the rest of Canada. The estimated impact of the policy on employment is slightly larger than with the previous estimator, and the estimate for living with parents is much larger. The 9.5 percentage point decrease in living with parents corresponds to a 22 per cent decrease relative to the mean.

Finally, the last row combines the two control groups using a triple difference estimator. The difference between age groups, Quebec and the rest of Canada, and before and after the reform is calculated, yielding a decrease of 5.0 per cent for employment and 6.2 per

²⁵ We exclude July and August 1989 from the analysis as the policy was actively changing at that point and we want to avoid any adjustment process. The 36 month band was chosen to gather sufficient sample sizes for precise inferences. Tighter bands around the reform led to similar, but less precise estimates.

cent for living with parents. Relative to the mean, these are economically significant magnitudes.

Which estimator performs best depends on the nature of the unobservable characteristics, as the estimators identify the effect of policy only by assuming the absence of fixed differences across the control groups. For example, the difference estimator in the second row assumes that, in the absence of the policy, there would be no difference in the behaviour of respondents in the two age groups. If the onset of a recession had a differential impact on one age group relative to another, this estimator would be biased.

While the employment results are quantitatively similar to the RD estimates reported in the previous section, we find a large and significant effect on living arrangements in the LFS which is at odds with the RD estimates. As mentioned earlier, there is clearly no abrupt decline between age 29 and 30 in the fraction of men living with their parents in Quebec in 1986. This suggests that the difference in differences estimates from the LFS are driven by some spurious province-age specific shocks that are not appropriately controlled for.

9 Conclusions

Using a unique policy episode involving lower social assistance payments to those under 30 in Quebec, we studied the effects of a transfer program on several measures of labour market behaviour, income and living arrangements. Our main finding is that more generous social assistance benefits substantially reduce the employment probability of less-educated men without dependent children. The employment rate for this group of men drops by three to five percentage points in response to the higher benefits. Perhaps more surprisingly, we also find that higher benefits also reduce the employment rate of *all* men by about one percentage point. From a broader perspective, this suggests that work disincentives embodied in to social programs may explain some (but certainly not all) of the difference in employment rates across OECD countries.

We also find that, as expected, the take-up of social assistance increases when benefits rise. We do not find significant impacts, however, in the case of most other outcome variables. One exception is marital status. Higher social assistance benefits appear to reduce the probability of being married. We also find some (imprecise) evidence that higher social assistance payments substituted for decreased earnings.

These findings are limited for several reasons. In particular, all our effects are identified for men at age 29-30. This might not generalize to other ages. Furthermore, our results for single employable males may not be relevant for similar programs in countries like the United States where this group is not eligible for welfare benefits.

We also have several interesting methodology findings. Most importantly, we find that the RD approach gives sensible results that are not very sensitive to whether we control very flexibly or just with a linear specification. We also find that exploiting the longitudinal nature of the Census (FD-RD method) improve inferences. Finally, we conclude that difference-in-differences estimators are more sensitive to specification issues than the RD estimates. While the difference-in-differences approach works well when we use a control group in the same labour market, it does not work very well when we use other regions to control for common economic trends.

References

Ashenfelter, Orley (1983), "Determining Participation in Income-Tested Social Programs," *Journal of the American Statistical Association*, Vol. 78, No. 383, pp. 517-525.

Baker, Michael, A. Abigail Payne, and Michael Smart (1999), "An empirical study of matching grants: the 'cap on CAP'," *Journal of Public Economics*, Vol. 72, pp. 269-288.

Barrett, Garry F. and Michael I. Cragg (1998), "An untold story: The characteristics of welfare use in British Columbia," *Canadian Journal of Economics*, Vol. 31, No. 1, pp. 165-188.

Bitler, Marianne P., Jonah B. Gelbach, and Hilary W. Hoynes (2003), "Welfare reform and children's living arrangements," Manuscript.

Black, Dan A., Terra G., McKinnish, and Seth G. Sanders (2003), "Does the availability of high-wage jobs for low-skilled men affect welfare expenditures? Evidence from shocks to the steel and coal industries," *Journal of Public Economics*, Vol. 87, pp. 1921-1942.

Blanchard, Olivier (2004), "The economic future of Europe," NBER Working Paper No. 10310, February.

Blank, Rebecca M. (2002), "Evaluating welfare reform in the United States," *Journal of Economic Literature*, Vol. 40, No 4, pp. 1105-1166.

Boudarbat, Brahim, Thomas Lemieux and W. Craig Riddell, "Recent Trends in Wage Inequality and the Wage Structure in Canada," TARGET working paper no. 6, September 2003

Card, David and Thomas Lemieux (2000), "Adapting to circumstances: the evolution of work, school, and living arrangements among North American youth," in *Youth Employment and Joblessness in Advanced Countries*, David Blanchflower and Richard Freeman (eds.). Chicago: University of Chicago Press.

Dooley, Martin D. (1999), "The evolution of welfare participation among Canadian lone mothers, 1973-1991," *Canadian Journal of Economics*, Vol. 32, No. 3, pp. 589-612.

Dooley, Martin D., Stephane Gascon, Pierre Lefebvre, and Philip Merrigan (2000), "Lone female headship and welfare policy in Canada," *The Journal of Human Resources*, Vol. 35, No. 3, pp. 587-602.

Ford, Reuben, David Gyarmati, Kelly Foley, and Doug Tattrie (2003), "Can work incentives pay for themselves? Final report on the Self-Sufficiency Project for welfare applicants." Social Research and Demonstration Corporation, Ottawa.

Fortin, Bernard, Guy Lacroix, and Simon Drolet (2004), "Welfare benefits and the duration of welfare spells: evidence from a natural experiment in Canada," *Journal of Public Economics*, Vol. 88, pp. 1495-1520.

Green, David A. and William P. Warburton (2004), "Tightening a welfare system: the effects of benefit denial on future welfare receipt," *Journal of Public Economics*, Vol. 88, pp. 1471-1493.

Hahn, Jinyong, Wilbert van der Klaauw, and Petra Todd (2001), "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design," *Econometrica*, 69(1), January, 201-209.

Harvey, Carol, Michael J. Camasso, and Radha Jagannathan (2000), "Evaluating welfare reform waivers under Section 1115," *Journal of Economic Perspectives*, Vol. 14, No. 4, pp. 165-188.

Hoynes, Hilary W. (2000), "Local labor markets and welfare spells: Do demand conditions matter?" *Review of Economics and Statistics*, Vol. 82, No. 3, pp. 351-368.

Hu, Wei-yin (2001), "Welfare and family stability: Do benefits affect when children leave the nest?" *Journal of Human Resources*, Vol. 36, No. 2, pp. 274-303.

Institut de la Statistique du Québec (2004), « Adultes actifs au programme d'assistance-emploi selon la scolarité, Québec, » April.

Kapsalis, Constantine (2001) "An Assessment of EI and SA reporting in SLID", Statistics Canada Analytical Research Branch Working Paper no. 166, August

Lee, David (2003) "Randomized Experiments from Non-random Selection in U.S. House Elections," UC Berkeley manuscript, September

London, Rebecca A. (2000), "The interaction between single mothers' living arrangements and welfare participation," *Journal of Policy Analysis and Management*, Vol. 19, No. 1, pp. 93-117.

Milligan, Kevin (2003), "Subsidizing the stork: new evidence on tax incentives and fertility," updated version of NBER Working Paper No. 8845.

Mincer, Jacob (1974), *Schooling, Experience and Earnings*, Columbia University Press: New York.

Moffitt, Robert (1992), "Incentive effects of the US welfare system: A review," *Journal of Economic Literature*, Vol. 30, No. 1, pp. 1-61.

Moffitt, Robert (1998), "The effect of welfare on marriage and fertility," in *Welfare, the family and reproductive behavior*, R. Moffitt (ed.). Washington: National Research Council, National Academy of Sciences Press.

Moffitt, Robert (2002), "Welfare programs and labor supply," in *Handbook of Public Economics Vol. 4*, A. Auerbach and M. Feldstein (eds.).

Moffitt, Robert (2003), "The Temporary Assistance for Needy Families Program," in *Means-tested Transfer Programs in the United States*, Robert A. Moffitt (ed). Chicago: University of Chicago Press.

Murphy, Kevin M. and Finis Welch (1990), "Empirical Age-Earnings Profiles," *Journal of Labor Economics* 8, 202-229.

Paxson, Christina H. and Jane Waldfogel (2003), "Welfare reforms, family resources, and child maltreatment," *Journal of Policy Analysis and Management*, Vol. 22, No. 1, pp. 85-113.

**Appendix Table 1:
Cell Size and Sample Composition in the 1986 Census, Men**

Age	All men in Quebec 1986		All HS Dropouts in Quebec 1986		Cell size for High School Dropouts without children			
	Cell size	Fraction HS Dropouts	Cell size	Fraction w/ children	Quebec 1986	Quebec 1991	RoC 1986	RoC 1991
20	10,945	0.274	3,004	0.027	2,923	2,471	10,203	8,553
21	11,939	0.268	3,202	0.046	3,055	2,206	9,890	7,303
22	11,908	0.275	3,272	0.072	3,037	1,933	9,362	6,299
23	11,838	0.279	3,299	0.118	2,909	1,764	8,736	5,990
24	11,701	0.284	3,318	0.169	2,757	1,733	8,118	5,664
25	12,006	0.297	3,564	0.234	2,730	1,833	7,557	5,677
26	11,841	0.307	3,634	0.323	2,461	1,778	6,594	5,719
27	11,594	0.315	3,650	0.381	2,258	1,797	5,770	5,384
28	11,812	0.311	3,672	0.452	2,013	1,758	5,186	5,294
29	11,939	0.310	3,696	0.497	1,859	1,676	4,506	4,939
30	11,593	0.305	3,537	0.534	1,647	1,637	3,912	4,784
31	11,253	0.299	3,364	0.566	1,461	1,579	3,722	4,219
32	11,156	0.289	3,219	0.590	1,319	1,506	3,241	3,959
33	10,549	0.279	2,943	0.614	1,135	1,425	2,872	3,741
34	10,806	0.271	2,930	0.640	1,056	1,367	2,551	3,406
35	10,451	0.268	2,796	0.652	972	1,236	2,499	3,275
36	10,438	0.270	2,823	0.656	970	1,195	2,384	3,036
37	10,331	0.271	2,804	0.682	893	1,150	2,232	2,761
38	10,260	0.283	2,908	0.706	856	1,013	2,371	2,504
39	10,327	0.295	3,046	0.711	881	1,012	2,382	2,477

Note: These descriptive statistics are based on the 20 percent sample of the Canadian Census. The acronym “RoC” stands for the “rest of Canada”.

**Appendix Table 2:
Comparison of 1986 Census and 1985-88 SCF**

A. 1985-88 Survey of Consumer Finance

Age group	Observations	SA+other	SA	other
26-30	313	848 (98)	764 (95)	84 (24)
31-35	212	1263 (147)	1072 (131)	191 (77)
Difference:		-415 (177)	-308 (162)	-107 (80)

B. 1986 Census

Age group	Observations	SA+other
26-30	10,238	791 (22)
31-35	5,943	1407 (45)
Difference:		-616 (44)

Note: Computed from a sample of male high school dropouts without dependent children.

Table 1
Regression Discontinuity Estimates of the Effect of Higher Social Assistance Benefits on Labour Supply in Quebec, 1986

Specification for age	Empl. rate last year	Empl. Rate at census	Difference in empl. rate	Weekly hours
Mean of the dependent variable				
	0.562	0.618	0.056	24.39
Regression discontinuity estimates				
Linear	-0.045 *** (0.012)	-0.041 *** (0.012)	-0.029 ** (0.011)	-1.45 ** (0.54)
Quadratic	-0.048 *** (0.013)	-0.051 *** (0.012)	-0.031 ** (0.012)	-1.75 ** (0.61)
Cubic	-0.043 ** (0.018)	-0.048 *** (0.014)	-0.030 ** (0.013)	-1.47 * (0.70)
Linear spline	-0.047 *** (0.013)	-0.049 *** (0.011)	-0.032 ** (0.013)	-1.72 *** (0.55)
Quadratic spline	-0.038 (0.024)	-0.056 ** (0.018)	-0.035 * (0.016)	-1.66 (0.94)
Goodness of fit statistic (p-value)				
Linear	0.48	0.52	0.91	0.48
Linear spline	0.47	0.72	0.85	0.00

Three asterisks indicate statistical significance at the 1% level; two asterisks for the 5% level, and one asterisk for the 10% level.

Table 2
Regression Discontinuity Estimates of Labour Supply
Response: Comparison of Quebec and Rest of Canada in 1986
and 1991

Specification for age	Quebec 1986	***	Rest of Canada 1986	**	Quebec 1991	*	Rest of Canada 1991	Diff-in-Diff	
Regression discontinuity estimates: Employment rate on Census week									
Linear	-0.041	***	-0.013	**	0.041	*	0.005	-0.064	***
	(0.012)		(0.006)		(0.022)		(0.011)	(0.028)	
Quadratic	-0.051	***	-0.013	*	0.012		-0.017	***	-0.067
	(0.012)		(0.007)		(0.023)		(0.006)	(0.028)	***
Cubic	-0.048	***	-0.009		0.037	**	-0.016	**	-0.092
	(0.014)		(0.007)		(0.015)		(0.007)	(0.023)	***
Linear spline	-0.049	***	-0.014	*	0.010		-0.010		-0.055
	(0.011)		(0.006)		(0.017)		(0.007)	(0.022)	**
Quadratic spline	-0.056	**	-0.007		0.042	*	-0.007		-0.098
	(0.018)		(0.010)		(0.022)		(0.007)	(0.031)	***
Regression discontinuity estimates: Difference in Employment rate									
Linear	-0.029	**	-0.009		0.022	*	-0.007		-0.049
	(0.011)		(0.007)		(0.011)		(0.006)	(0.018)	**
Quadratic	-0.031	**	-0.006		0.022		-0.005		-0.052
	(0.012)		(0.007)		(0.013)		(0.006)	(0.020)	**
Cubic	-0.030	**	-0.004		0.020		-0.002		-0.048
	(0.013)		(0.006)		(0.014)		(0.006)	(0.021)	**
Linear spline	-0.032	**	-0.004		0.021		-0.003		-0.052
	(0.013)		(0.008)		(0.014)		(0.006)	(0.022)	**
Quadratic spline	-0.035	*	0.001		0.012		-0.005		-0.053
	(0.016)		(0.009)		(0.016)		(0.008)	(0.026)	**

Three asterisks indicate statistical significance at the 1% level; two asterisks for the 5% level, and one asterisk for the 10% level.

Table 3
Regression Discontinuity Estimates of the Effect of Higher Social Assistance Benefits on other Outcomes in Quebec, 1986

Specification for age	Transfers (\$1000)	Fraction with transfers>0	Transfers cond. on transfers>0	UI (\$1000)	Earnings (\$1000)	Married					
Mean of the dependent variable											
	1.065	0.212	4.885	1.126	13.924	0.323					
Regression discontinuity estimates											
Linear	0.477 (0.086)	*** (0.012)	0.041 (0.012)	*** (0.232)	1.248 (0.232)	*** (0.065)	-0.106 (0.595)	-0.921 (0.575)	-0.027 (0.021)		
Quadratic	0.477 (0.089)	*** (0.012)	0.041 (0.012)	*** (0.233)	1.135 (0.233)	*** (0.072)	-0.120 (0.575)	-1.059 (0.575)	* (0.013)	-0.064 (0.013)	***
Cubic	0.496 (0.137)	*** (0.018)	0.042 (0.018)	** (0.320)	1.273 (0.320)	*** (0.098)	-0.085 (0.792)	-0.461 (0.792)	-0.050 (0.012)	***	
Linear spline	0.481 (0.088)	*** (0.012)	0.041 (0.012)	*** (0.218)	1.165 (0.218)	*** (0.068)	-0.121 (0.587)	-0.975 (0.587)	-0.055 (0.012)	***	
Quadratic spline	0.445 (0.166)	** (0.022)	0.033 (0.022)	1.169 (0.423)	** (0.423)	-0.074 (0.130)	0.202 (0.913)	-0.036 (0.016)	*		

Three asterisks indicate statistical significance at the 1% level; two asterisks for the 5% level, and one asterisk for the 10% level.

Table 4
Regression Discontinuity Estimates of the Effect of Higher Benefits on Social Assistance Payments (dollar amount)

Specification for age	Quebec 1986	Rest of Canada 1986	Quebec 1991	Rest of Canada 1991
Mean of the dependent variable				
	1.065	0.646	0.851	0.625
Regression discontinuity estimates				
Linear	0.477 *** (0.086)	0.071 (0.049)	0.153 * (0.082)	-0.074 * (0.039)
Quadratic	0.477 *** (0.089)	0.062 (0.049)	0.192 ** 0.083	-0.067 (0.041)
Cubic	0.496 *** (0.137)	0.047 (0.073)	0.071 (0.103)	-0.112 * (0.055)
Linear spline	0.481 *** (0.088)	0.066 (0.049)	0.180 ** (0.079)	-0.071 (0.041)
Quadratic spline	0.445 ** (0.166)	0.010 (0.094)	0.034 (0.144)	-0.154 ** (0.067)

Three asterisks indicate statistical significance at the 1% level; two asterisks for the 5% level, and one asterisk for the 10% level.

Table 5: Cell Means and Regression Results from the Census

	sample / specification	obs.	Emp- loyed		obs.	Positive Transfers	
Mean of dependent variable	Quebec, age 29, 1986	1859	0.661		1859	0.190	
Cell means estimator	Discontinuity Quebec, 1986: Age 30 vs. Age 29	3506	-0.052 (0.016)	***	3178	0.058 (0.015)	***
	Diff-diff 1986: 29-30 vs Que-RoC	11924	-0.044 (0.019)	**	10925	0.039 (0.016)	**
	Diff-diff Quebec: 29-30 vs 1986-1991	6819	-0.079 (0.023)	***	6360	0.040 (0.020)	**
	Diff-diff Age 29: 1986-1991 vs Que-RoC	12908	0.010 (0.018)		12908	-0.019 (0.015)	
	Diff-diff-diff	24960	-0.079 (0.026)	***	23005	0.045 (0.023)	**
Regression estimator	Diff-diff-diff	24960	-0.074 (0.026)	***	23005	0.034 (0.023)	

Three asterisks indicate statistical significance at the 1% level; two asterisks for the 5% level, and one asterisk for the 10% level.

Table 6
Cell Means Results from the Labour Force Survey

	obs.	Emp- loyed	Lives with parents
Mean of dependent variable			
Quebec, Age 25-29, Before	3875	0.616	0.424
Cell Mean Difference Estimators			
Diff-diff	18748	-0.056 ***	-0.028 **
Quebec: Before vs. After and 25-29 vs. 30-34		(0.015)	(0.014)
Diff-diff	38061	-0.078 ***	-0.095 ***
Before: Quebec vs. RoC and 25-29 vs. 30-34		(0.013)	(0.013)
Diff-diff-diff	93996	-0.050 ***	-0.062 ***
		(0.016)	(0.016)

Three asterisks indicate statistical significance at the 1% level; two asterisks for the 5% level, and one asterisk for the 10% level. All results are weighted.

Figure 1: Social Assistance Benefits, Single Individual




Figure 2: Static Labour Supply Model




Figure 3: Employment Rates in the Labour Force Survey

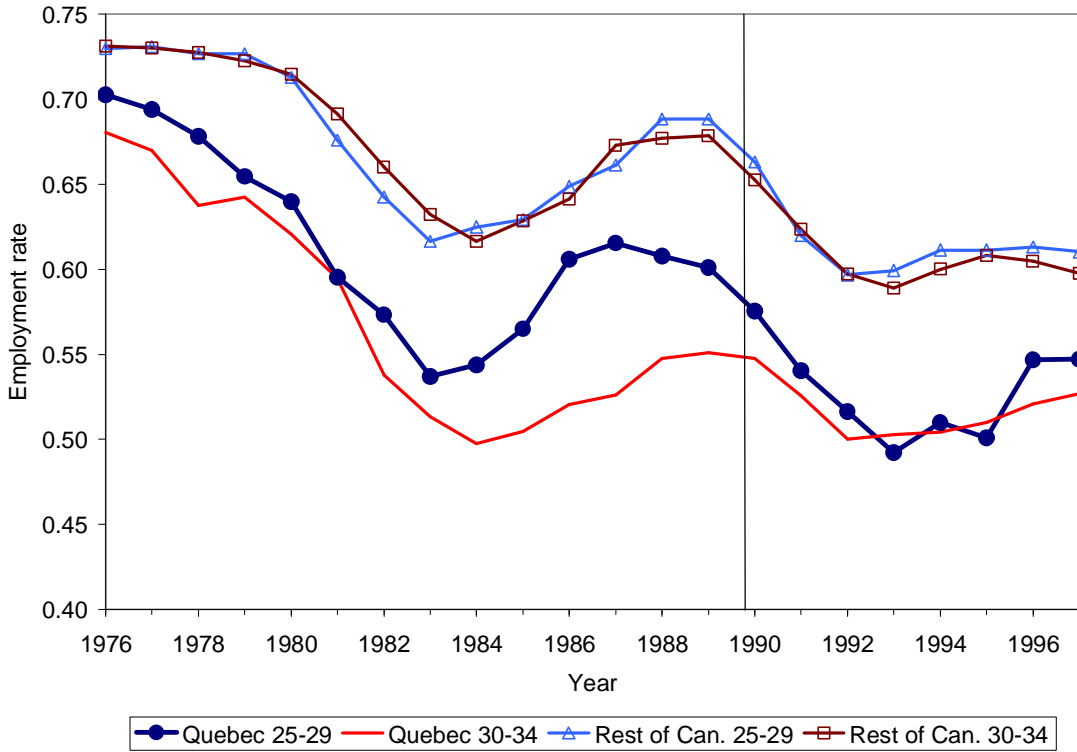


Figure 4: Employment Rate in Census Week




Figure 5: Weekly Hours of Work




Figure 6: Fraction Living with Parents




Figure 7: Fraction Married

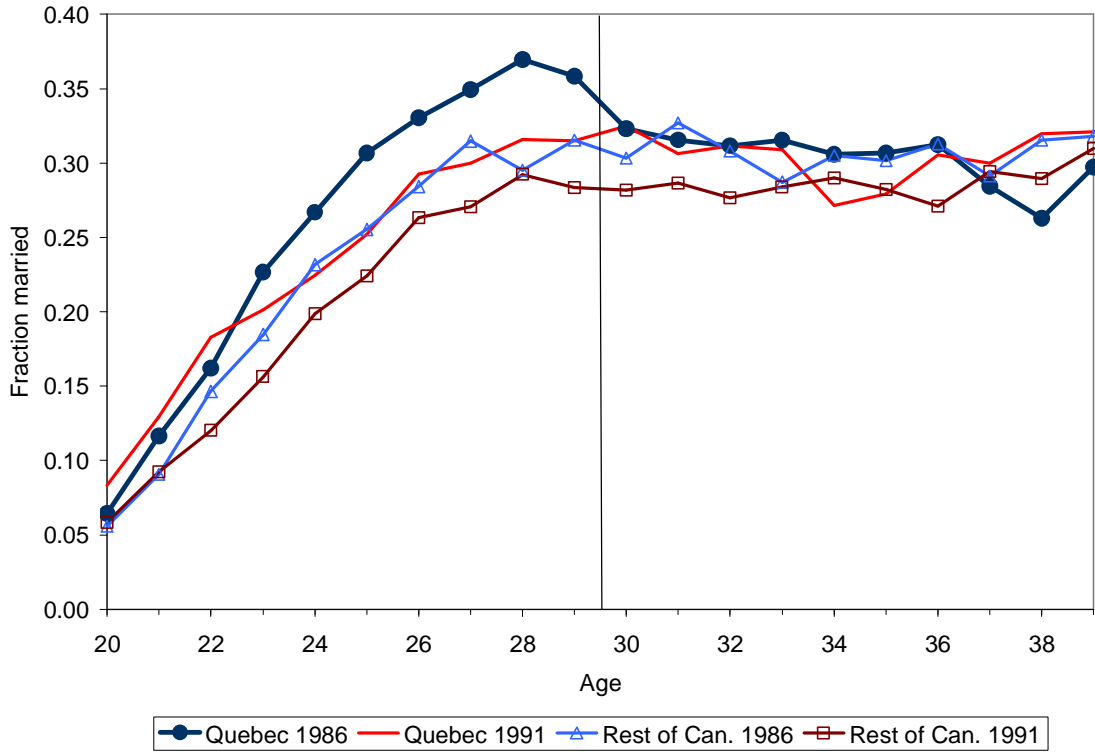


Figure 8: Transfer Income (Social Assistance)




Figure 9: Annual Earnings




Figure 10: Employment Rate in Census Week, Quebec 1986




Figure 11: Employment Rate in Previous Year, Quebec 1986




Figure 12: Change in Employment Rate, Quebec 1986

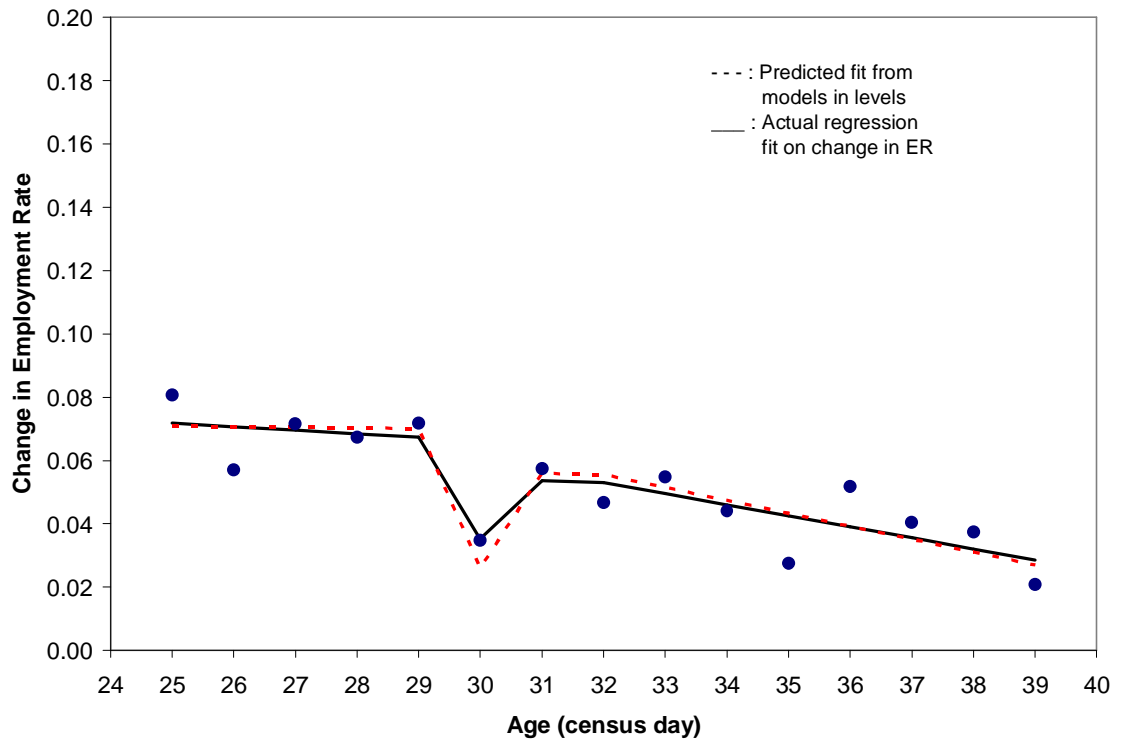


Figure 13: Social Assistance Income, Quebec 1986




Figure 14: Social Assistance Take-up Rate, Quebec 1986







Figure 15: Social Assistance Income for SA>0, Quebec 1986



Appendix Figure 1: Fraction of HS Dropouts with Children



Appendix Figure 2: Employment Rate for All HS Dropouts



Appendix Figure 3: Employment Rate for All Men

