Incentive effects of social assistance: A regression discontinuity approach

Thomas Lemieux[*], Kevin Milligan

Department of Economics, University of British Columbia, 997-1873 East Mall, Vancouver, BC, Canada V6T 1Z1

Abstract

Before 1989, childless social assistance recipients in Quebec under age 30 received much lower benefits than recipients over age 30. We use this sharp discontinuity in policy to estimate the effects of social assistance on various labour market outcomes using a regression discontinuity approach. We find strong evidence that more generous social assistance benefits reduce employment. The 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-differences estimators may perform poorly with inappropriately chosen control groups.
© 2007 Elsevier B.V. All rights reserved.

JEL classification: H52; J21

Keywords: Regression discontinuity; Welfare; Social assistance; Labour supply

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 programmes in the United States since 1967 was a concern about disincentives to work embedded in the programmes (Moffitt, 2003). 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 European 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 this paper, we study the effects of an interesting 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 effect of social assistance on the duration of social assistance

---
[*]Corresponding author. Tel.: +1 604 822 2092; fax: +1 604 822 5915.

E-mail addresses: tlemieux@interchange.ubc.ca (T. Lemieux), kevin.milligan@ubc.ca (K. Milligan).

spells using a difference-in-differences approach. The break in the policy at age 30 also provides, however, a natural setting for analysing the impact of welfare payments using a regression discontinuity approach. The key advantage of this approach is that it provides estimates that are "as credible as those from a randomized experiments" (Lee, 2007) under relatively weak conditions.[1] The regression discontinuity approach has been used to look at the effect of students aid offers on college enrolments (van der Klaauw, 2002), parents' valuation of elementary schools (Black, 1999), and electoral advantage to incumbency (Lee, 2007), to name a few examples. In this paper, we use this research design to estimate the effect of welfare payments on labour market outcomes. A unique feature of our analysis, described in detail later in the paper, is our formation of a first difference estimator in the regression discontinuity design.

The research strategies chosen over the years to study the effects of welfare have been closely intertwined with the changing policy environment.[2] In the 1970s and 1980s, most research consisted of the econometric modelling 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' programmes 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 programmes or experiments that altered the parameters and structure of welfare programmes. 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 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 into some mix of time limits, benefit rates reduction, 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 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. In 1989, those under age 30 received $185 per month compared to $507, or 175% more, for those ages 30 and over. 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 distinctive 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 the presence of young children.

---

[1]See also Hahn et al. (2001).
[2]The literature is reviewed by Moffitt (2002).

2. Social assistance in Quebec and Canada

Social assistance (as welfare is called in Canada) programmes were funded from 1967 to 1996 through the Canada Assistance Plan, which offered a 100% matching grant from the federal government for provincial spending.[3] In contrast to the federally funded welfare programmes in the United States during that period, the design of the programmes was left almost entirely to the sub-national jurisdictions, subject to weak conditions on eligibility.[4] A distinguishing feature from the case of the United States is the eligibility of singles and non-parents.

Research on incentives and 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 summarized in Ford et al. (2003) generally show strong effects of the earnings supplement on labour supply. 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 et al. (2004) study the effect of social assistance benefits on the duration 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. They find that the large increase in benefits that followed the 1989 reform increased the average duration of social assistance benefits among those under 30 by 4–8 months (depending on subgroups). Our work differs from theirs in a number of ways. First, we study static participation 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 per month in the 1980s), after which benefits were reduced dollar for dollar with income.

The unique feature of social assistance for our purposes was the differential benefit rate by age. Those under age 30 received $185 per month in 1989 (current dollars) compared to $507, or 175% more, for those age 30 and over.[5] 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

---

[3]Following 1996, a bloc grant called the Canadian Health and Social Transfer replaced the Canada Assistance Plan.

[4]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 et al. (1999).

[5]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.
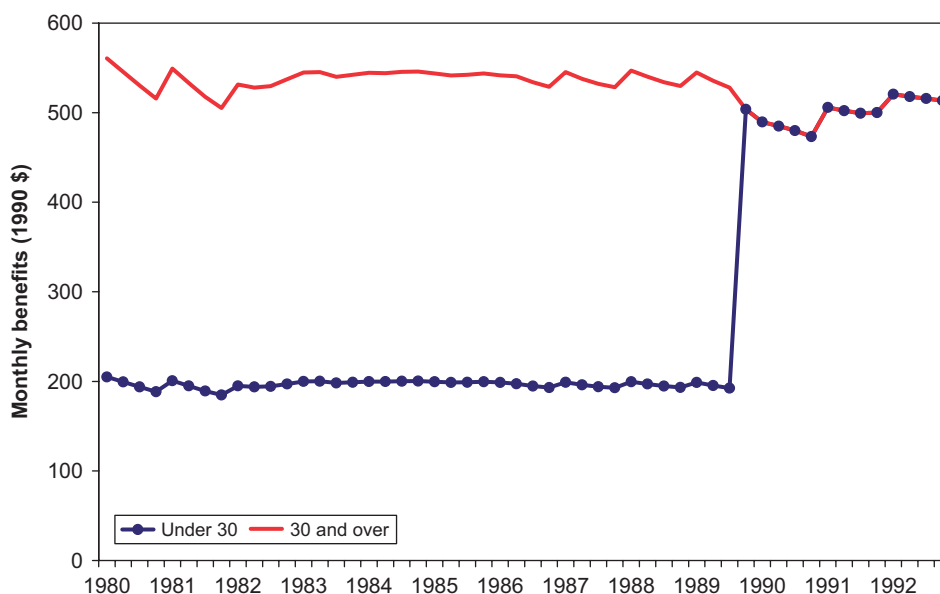
Fig. 1. Social assistance benefits, single individual.

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.[6] We graph the benefit rates in constant 1990 dollars for a single employable person without children in Fig. 1, for someone over and under age 30.[7]

It is straightforward to show that higher social assistance benefits unambiguously reduce the probability of working in a static labour supply framework. Leaving aside the (small) earnings exemption, individuals simply have to choose between (1) not working and collecting social assistance, and (2) working the optimal number of hours they would choose in absence of social assistance. All individuals whose utility at the higher (age 30 and more) benefits now exceeds the utility of work will drop out of the labour market, while those who are still better off working will keep their hours of work unchanged. An additional prediction of the model is thus that all the adjustment takes place at the extensive (participation) margin as opposed to the intensive (hours conditional on participating) margin.

3. Data 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 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.[8] We also focus on respondents without children, as parents of children were not subject to the lower social assistance benefits.[9] The bonus that would

---

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

[7]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).

[8]Recent data from the Institut de la Statistique du Québec (2004) indicates that 63% of all social assistance claimants are high school dropouts. Our own tabulations based on the 1986–1989 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% of the population.

[9]We classified people as "childless" or "without children" when they either do not have children, or have children but do not live with them.

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.[10] We discuss these sample selection issues in more detail later.

Finally, we look 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.[11] 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%, respectively. For these two reasons, the 'at risk' group is much smaller for women than men. Finally, we are more concerned about possible fertility responses in the case of women than men.

3.1. 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 20% 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% of respondents. As we are interested in obtaining large samples of individuals in narrowly defined cells, we obtained access to the full 20% master sample maintained by Statistics Canada. With this sample, we can form cells of sufficient size for meaningful analysis. For example, Table A1 shows that we have over 10,000 observations for each year of age in Quebec in the 1986 Census. Since between 26% and 32% 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 Table A1 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. 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 calendar year previous to the year of the Census. In particular, the Census collects separate income items for earnings, 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 programmes, and small provincial tax credits claimed on the tax return.[12]

Fortunately, social assistance benefits are by far the largest component of the "other transfers" variable. In Lemieux and Milligan (2004), we document using another source that 85% of income included in "other transfers" in the Census actually comes from social assistance income. So for all practical purposes, this means that we can use "other transfers" and "social assistance benefits" interchangeably in what follows. These numbers should nonetheless be interpreted with caution since existing validation studies suggest that social assistance benefits are significantly under-reported in survey data.[13]

---

[10]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 (2005) finds a response to the program, but the response is muted among low education and low income women.

[11]Among 30 years old high school dropouts in Quebec in 1986, 75.7% of women had (and lived with) children, compared to 53.4% 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.

[12]The 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).

[13]For example, Kapsalis (2001) shows that social assistance is underreported by a factor of about a third.

3.2. The labour force survey

We use the LFS to document the long-term trends in the labour market behaviour of our target population, comparing them across age groups and provinces. The LFS is available monthly back to 1976, but has too small sample sizes to be used convincingly in the regression discontinuity analysis. So, we use it instead to provide context to the discussion of labour markets in Canada and Quebec.

Fig. 2 graphs the employment rate for males using the LFS. 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 cyclicality 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 in 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 the 25–29 year old group is substantially larger than the employment rate of the 30–34 year old group 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 age 30 prior to August 1989 led to a substantial labour supply response.

Other factors could nonetheless account for the abnormally large employment rate of 25–29 year olds in Quebec in the late 1980s. Perhaps the strong economic recovery of the second half of the 1980s disproportionally benefited younger workers in Quebec. It is also not clear why the employment rates of 25–29 and 30–34 year olds 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 Fig. 2) compare to the regression discontinuity results.
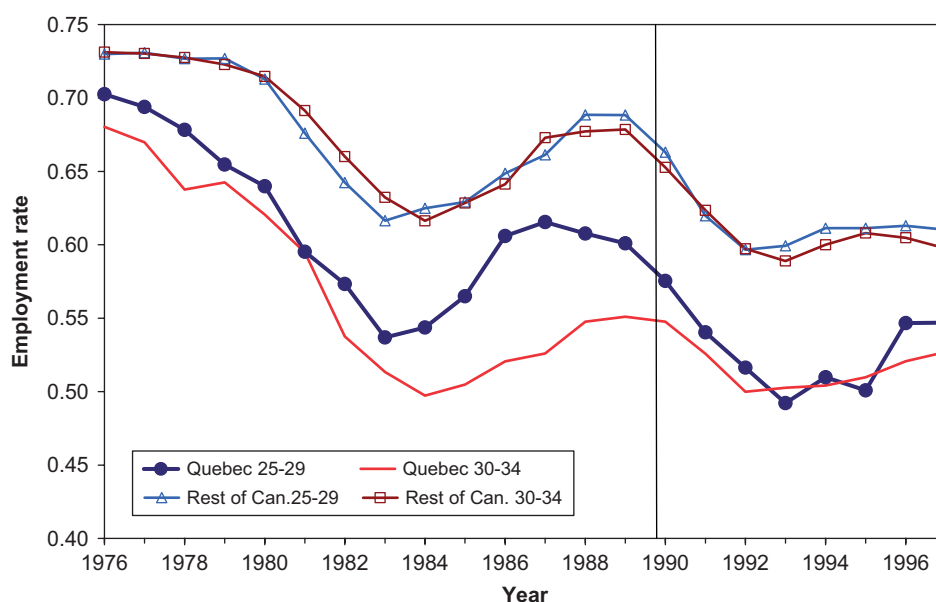

Fig. 2. Employment rates in the Labour Force Survey.

4. Empirical approach

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

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

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

$$TREAT_{ia} = \begin{bmatrix} 0 & \text{if } a < 30, \\ 1 & \text{if } a \geqslant 30. \end{bmatrix}$$

The evaluation problem consists of estimating the effect β_1 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_1 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? Most of our variables of interest like income and employment 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 that $\delta(\cdot)$ is 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 be age 30 could be systematically different from those aged 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 the authorities.

A potentially more serious problem is that we only select individuals without dependent children for most of our analysis, since only those individuals are subject to differential social assistance benefits. As shown in Table A1, 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. 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. We present evidence supporting this view in Section 6.3.

In practice, the estimated treatment effect depends on how the smooth function $\delta(\cdot)$ is itself estimated. As in any non-parametric estimation problem, there is a difficult trade-off between precision and bias. We balance this trade-off between precision and bias by estimating a variety of polynomial specifications for the regression function $\delta(\cdot)$. In Section 5, 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 present (in Section 6.1) estimates of the linear spline model for an increasingly small window around age 30 as a further robustness check.

---

[14]See Hahn et al. (2001) and Lee (2007) 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.

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 at Census Day (typically the first week of June). Since we only know whether people are at age 29 or age 30 at Census Day, we cannot directly compare people who "just turned 30" to people "just about to turn 30". Because age is a discrete variable in our data, we need the regression approach to extrapolate what the discontinuity would be at the precise point where people turn 30.[15]

Because of this data limitation, all the information available in the microdata 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 Eq. (1):

$$Y_a = \beta_0 + \beta_1 TREAT_a + \delta(a) + \varepsilon_a. \tag{2}$$

Regression estimates of Eq. (1) based on microdata are identical to weighted estimates of Eq. (2) if 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

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

Under the null hypothesis that model (2) is the true model, 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 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 of the previous year. This individual was thus "exposed" to higher social assistance benefits for only one of the twelve months during the previous year. This suggests assigning $\frac{1}{12}$ 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 individuals 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]:

$$TREAT'_a = \begin{cases} 0 & \text{if } a \leqslant 29, \\ 0.170 & \text{if } a = 30, \\ 0.913 & \text{if } a = 31, \\ 1 & \text{if } a \geqslant 32. \end{cases} \tag{3}$$

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

---

[15]See Lee and Card (2007) for more discussion of the regression discontinuity approach in the case where the treatment-determining covariate (age in our case) is discrete. If we observed the exact age of individuals in the data, we would not need to estimate parametric regression models and could directly estimate the discontinuity at age 30 using, for example, simple step functions.

[16]The values of the treatment variable $TREAT'_a$ for ages 30 and 31 are obtained by integrating over the uniform distribution of birth dates. It can be shown that for age 30 we get $TREAT'_a = 0.5(7/12)^2 = 0.170$. For age 30 we get $TREAT'_a = 1 - 0.5(5/12)^2 = 0.913$.

One concern is that some of the advantages of the RD design are lost because we do not have a sharp discontinuity for the outcome variables measured over the previous year. Furthermore, outcomes measured over the previous year may be under-reported because of recall biases. Fortunately, it is possible to test for the impact of these shortcomings 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.

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 5), 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 in the week of the Census is

$$ERC_a = \beta_0 + \beta_1 TREAT_a + \delta(a) + \varepsilon_a, \tag{4a}$$

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

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

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 during the week of the Census (Eq. (4a)) should drop sharply between ages 29 and 30, as $TREAT_a$ jumps from 0 to 1. By contrast, most of the drop should occur between ages 30 and 31 in the model for the employment rate in the previous year (Eq. (4b)) since, according to Eq. (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 year olds (during 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 $\beta_1 = \beta'_1$, this FD-RD estimator is obtained by estimating the regression model

$$ERC_a - ERL_a = (\beta_0 - \beta'_0) + \beta_1(TREAT_a - TREAT'_a) + \theta(a) + (\varepsilon_a - \varepsilon'_a) \tag{5}$$

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 Eq. (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 aged 29 for some unmodelled reason. The FD-RD estimator goes one step further by comparing the employment of the *same* individuals at ages 29 and 30.

5. 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 aged 25–39. The reason is the systematic difference of the age profile in most of the variables between ages 20 and 24 compared to ages 25 and 29. This suggests that data for ages 20–24 are of little use for helping to fit the model around the discontinuity point. In any case, we show below (Section 6.1) that our results are very robust to the choice of the age window.

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 squared residuals is equal to the goodness-of-fit statistic GOF (up to a normalization).

We first present some suggestive graphical evidence before showing the regression results. Figs. 3 and 4 plot the raw employment rates (along with the 95% level confidence bounds) for the employment rates at the time of the Census and last year, respectively. 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 of age 30 on Census Day. In the case of employment in the previous year, we place the discontinuity point at age 30 and $\frac{5}{12}$ths for similar reasons.

Both Figs. 3 and 4 present strong evidence that employment drops abruptly once individuals become eligible for the higher social assistance benefits. As expected, the decline in employment measured during the week of the Census happens between ages 29 and 30, while the decline in employment measured over the previous year (Fig. 4) mostly happens between ages 30 and 31.

Note also that the employment rate tends to trend down as a function of age, especially after age 30. While this may be surprising at first glance, we show in Section 6.4 that this trend is entirely driven by the fact 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. More importantly, we also show in Section 6.3 that this selection bias is a smooth function of age that does not bias the RD estimates.

Turning to the regression results, 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 Eq. (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 during the week of the Census in the more flexible models like the cubic and the quadratic spline.
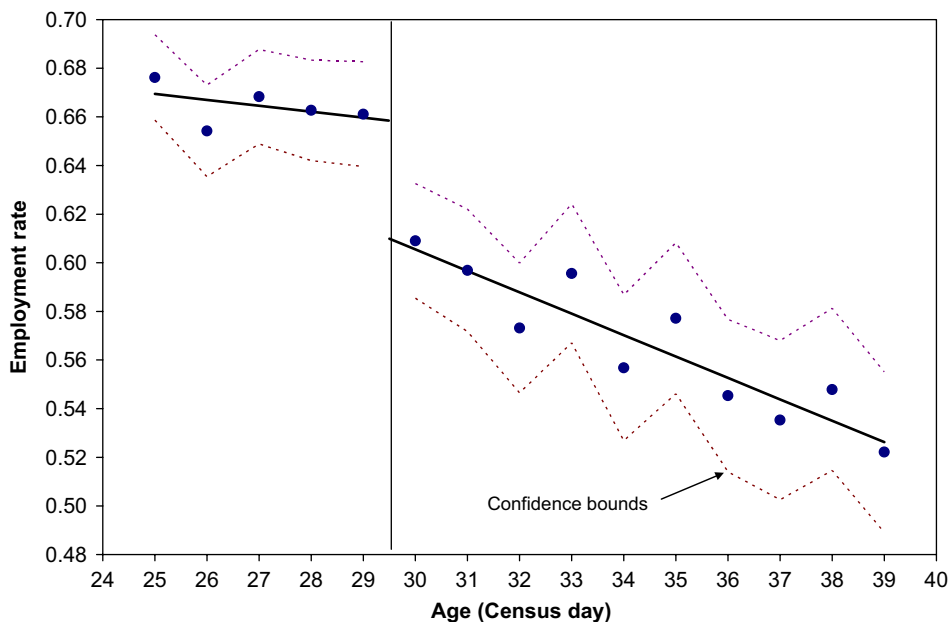

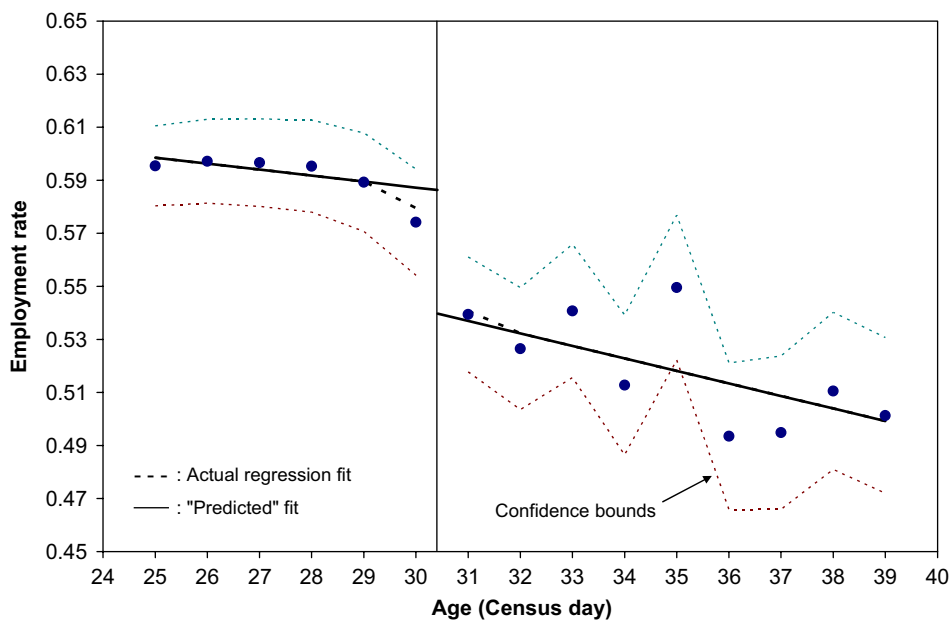Fig. 3. Employment rate in Census week, Quebec 1986.

Fig. 4. Employment rate in previous year, Quebec 1986.

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***	−0.041***	−0.029**	−1.45**
	(0.012)	(0.012)	(0.011)	(0.54)
Quadratic	−0.048***	−0.051***	−0.031**	−1.75**
	(0.013)	(0.012)	(0.012)	(0.61)
Cubic	−0.043**	−0.048***	−0.030**	−1.47*
	(0.018)	(0.014)	(0.013)	(0.70)
Linear spline	−0.047***	−0.049***	−0.032**	−1.72***
	(0.013)	(0.011)	(0.013)	(0.55)
Quadratic spline	−0.038	−0.056**	−0.035*	−1.66
	(0.024)	(0.018)	(0.016)	(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	0.55

Note:
***Indicate statistical significance at the 1% level.
**For the 5% level.
*For the 10% level.

One nice feature of the results is that the two employment rate measures yield remarkably similar estimates that range from −0.038 to −0.056. 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.
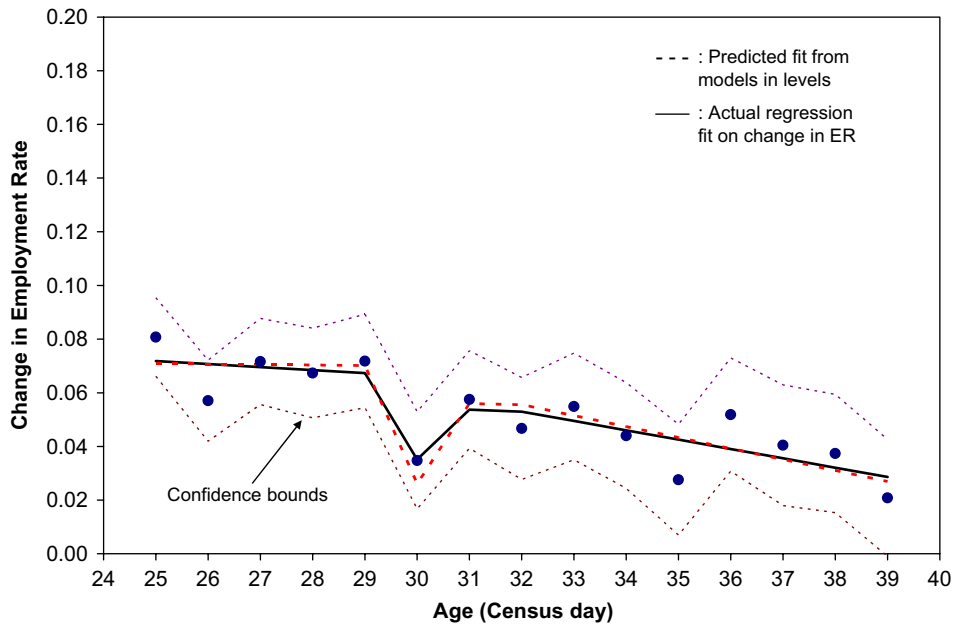
Fig. 5. Change in employment rate, Quebec 1986.

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 Figs. 3 and 4 for the linear spline models. In Fig. 4, 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). As shown in Table 1, the estimated effects are almost identical for employment last year and in the reference week (−0.047 and −0.049, respectively).

Another interesting feature of Figs. 3 and 4 is that the last observation before the discontinuity point is right on the regression line. Since people know prior to turning 30 that benefits will greatly increase once they turn 30, some of the labour supply adjustment may have been expected to take place before age 30.[17] If these adjustments were important, we would expect the regression model to overpredict the employment rate at age 29 (for Census week employment) or 30 (for employment last year). The results thus suggest a very quick response to the higher benefits shortly after individuals turn 30.

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 Eq. (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 (in the −0.03 to −0.04 range) 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 aged 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. Fig. 5 confirms this prediction that the employment rate difference is abnormally low for individuals aged 30 at the time of the Census (95% level confidence bounds are also shown in the figure). 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 Figs. 3 and 4).

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 of −1.72 in the linear spline model represents about 7.1% of the average hours of work (24.39). This

---

[17]For example, an individual who loses his job and collects unemployment insurance at age 29 may not bother looking for a new job in expectation of the higher social assistance benefits available at age 30.

is very similar to the 7.9% effect (effect of -0.049 multiplied by an average employment rate of 0.618) on employment probability obtained for the most comparable employment rate model (linear spline model for employment at Census week). The results suggest that, consistent with a simple labour supply model, 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).

Putting all these results together, we conclude that higher social assistance benefits reduce the employment rate in our sample 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. Robustness checks and extensions

In order to check the robustness of our results, we offer in this section a selection of robustness checks and extensions to our core analysis.

6.1. Narrowing the window

In Table 1, we show that the results are very robust to the choice of functional form for the regression equations. An alternative approach for checking the robustness of the results is to estimate a more parametric model like the linear spline for an increasingly narrow age window around the discontinuity point. This arguably captures better the spirit of the regression discontinuity approach by relying only on observations that are increasingly close to the discontinuity point.

Table 2 shows RD estimates with varying window width for the four outcome variables reported in Table 1. The specification used for all models is the linear spline. The first row reproduces the results reported in Table 1 for all childless dropout men in the age range 25–39. The following set of rows show the results for increasingly narrow age windows. For instance, row 2 reports the results for a "±5 years" window (five years above and five year below the discontinuity point).[18] Note also that the smallest window we can use with the linear spline model is "±2 years" since we need at least two observations on each side to identify separate regression lines. These models perfectly fit the data since there is no degree of freedom left.

Broadly speaking, the results are very robust the age window used in the estimation. For example, the more robust results based on the differenced model (FD-RD) remain in the -0.032 to -0.038 range for the different specifications. Estimates for other outcomes vary a little more but always remain negative (and statistically significant in most cases).

6.2. Falsification tests

We then run a series of "falsification experiments" in Table 3 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 3 indeed indicates a sharp contrast between Quebec in 1986 where employment effects are always significant at the 95% level (except in one case where it is significant at the 90% 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 15 estimates for other regions or years is significant (at the 90% level). Estimates based on employment at Census week (first panel) are more erratic. They are even statistically significant in a number of cases.

---

[18]In the case of outcomes measured at the Census date (employment rate and hours of work), the discontinuity point is between ages 29 and 30. The "±5 years" window this includes men ages 25–29 for the "below" group, and men ages 30–34 for the "above" group. In the case of employment last year, the "±5 years" window includes men ages 26–30 and 31–35 since the discontinuity occurs between ages 30 and 31. In the case of difference in employment rates, we use men ages 25–35 in order to have a five-year window both above (ages 31–35) and below (25–29) men age 30 (the critical group used for the identification).

Table 2
Linear spline regression discontinuity estimates with different age windows in Quebec, 1986

Window width	Empl. rate last year	Empl. rate at Census	Difference in empl. rate	Weekly hours
All ages 25–39	−0.047***	−0.049***	−0.032**	−1.72***
	(0.013)	(0.011)	(0.013)	(0.55)
±5 years	−0.056***	−0.046**	−0.037**	−1.49**
	(0.014)	(0.014)	(0.015)	(0.66)
±4 years	−0.042**	−0.057**	−0.038**	−2.09**
	(0.013)	(0.015)	(0.010)	(0.62)
±3 years	−0.050*	−0.039**	−0.034*	−1.37*
	(0.014)	(0.006)	(0.012)	(0.34)
±2 years	−0.033 (−)	−0.045 (−)	−0.044 (−)	−1.60 (−)

Note:
***Indicate statistical significance at the 1% level.
**For the 5% level.
*For the 10% level.

Table 3
Falsification test: comparing labour supply response in Quebec and rest of Canada in 1986 and 1991

Specification for age	Quebec, 1986	Rest of Canada, 1986	Quebec, 1991	Rest of Canada, 1991
Regression discontinuity estimates: employment rate on Census week				
Linear	−0.041***	−0.013**	0.041*	0.005
	(0.012)	(0.006)	(0.022)	(0.011)
Quadratic	−0.051***	−0.013*	0.012	−0.017***
	(0.012)	(0.007)	(0.023)	(0.006)
Cubic	−0.048***	−0.009	0.037**	−0.016**
	(0.014)	(0.007)	(0.015)	(0.007)
Linear spline	−0.049***	−0.014*	0.010	−0.010
	(0.011)	(0.006)	(0.017)	(0.007)
Quadratic spline	−0.056**	−0.007	0.042*	−0.007
	(0.018)	(0.010)	(0.022)	(0.007)
Regression discontinuity estimates: difference in employment rate				
Linear	−0.029**	−0.009	0.022*	−0.007
	(0.011)	(0.007)	(0.011)	(0.006)
Quadratic	−0.031**	−0.006	0.022	−0.005
	(0.012)	(0.007)	(0.013)	(0.006)
Cubic	−0.030**	−0.004	0.020	−0.002
	(0.013)	(0.006)	(0.014)	(0.006)
Linear spline	−0.032**	−0.004	0.021	−0.003
	(0.013)	(0.008)	(0.014)	(0.006)
Quadratic spline	−0.035*	0.001	0.012	−0.005
	(0.016)	(0.009)	(0.016)	(0.008)

Note:
***Indicate statistical significance at the 1% level.
**For the 5% level.
*For the 10% level.

6.3. Testing for the absence of "manipulation" effects

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. In particular, the RD approach may not be valid 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". The implication of this kind of "manipulation" is that the fraction of men with children is discontinuous around the discontinuity point (Lee, 2007; McCrary, 2007).

Table A1 and Fig. A1 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 ages 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).

McCrary (2007) suggests a formal test for the presence of manipulation in a RD design. Following Lee (2007), the idea is that the density of the treatment-determining variable (age in our case) should be continuous at the discontinuity point for the RD design to be valid. Given the discrete nature of our age data, it is very easy to implement McCrary's test by running local linear regression of the fraction of men in each age group below and above the discontinuity point.[19] The test consists of checking whether the fraction of men predicted to be at age 30 is the same for the two local linear regressions. When looking either at the fraction or log fraction of men, we find no evidence of a discontinuity at age 30. In other words, we cannot reject the null that the density of age is the same just below and just above the discontinuity point. The *p*-values are 0.57 and 0.44 for the fraction and log fraction, respectively.

6.4. Broadening the target group

We also present some additional results below for our models estimated on all men instead of conditioning on men without children. Using all men "solves" the selection problem but leads (presumably) to 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.

Fig. A2 shows the employment rate (at Census week) for *all* high school dropouts. The figure shows again a sharp drop in employment between ages 29 and 30 in Quebec in 1986, but no comparable change in the other regions or years (Quebec in 1991 or the rest of Canada in 1986 and 1991). 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). 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.[20]

Interestingly, there is no longer a declining trend in the employment rate when all dropouts are considered (Fig. A2). This confirms our earlier conjecture that the negative trend in the employment rate of high school dropouts without children (Figs. 3 and 4) is driven by an increasingly negative selection (as age increases) in this group of individuals.

In Lemieux and Milligan (2004), we also graphed the same lines for all men, irrespective of their level of education and of the presence of children. Perhaps surprisingly, we still found clear evidence of an abrupt decline in the employment rate at age 30 in Quebec in 1986, with a point estimate in the linear spline model of about one percentage point (-0.012, standard error of 0.006). This is roughly a third of the estimate for high school dropouts only (Fig. A2). 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).

---

[19]As suggested by McCrary (2007), we estimate the local linear regressions using a triangular kernel. We also use a window width of 5, i.e. use the five age groups to the left and the five age groups to the right to estimate the regression. Using a triangular kernel means that we simply estimate weighted regressions where the weight attached to an age group linearly declines from 1 at the discontinuity point to 0 five years away from the discontinuity point.

[20]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.

Table 4
Regression discontinuity estimates of the effect of higher social assistance benefits on other outcomes in Quebec, 1986

Specification for age	Transfers ($1,000)	Fraction with transfers >0	Transfers cond. on transfers >0	UI ($1,000)	Earnings ($1,000)
Mean of the dependent variable					
	1.065	0.212	4.885	1.126	13.924
Regression discontinuity estimates					
Linear	0.477***	0.041***	1.248***	−0.106	−0.921
	(0.086)	(0.012)	(0.232)	(0.065)	(0.595)
Quadratic	0.477***	0.041***	1.135***	−0.120	−1.059*
	(0.089)	(0.012)	(0.233)	(0.072)	(0.575)
Cubic	0.496***	0.042**	1.273***	−0.085	−0.461
	(0.137)	(0.018)	(0.320)	(0.098)	(0.792)
Linear spline	0.481***	0.041***	1.165***	−0.121	−0.975
	(0.088)	(0.012)	(0.218)	(0.068)	(0.587)
Quadratic spline	0.445**	0.033	1.169**	−0.074	0.202
	(0.166)	(0.022)	(0.423)	(0.130)	(0.913)

Note:
***Indicate statistical significance at the 1% level.
**For the 5% level.
*For the 10% level.

6.5. Other outcome variables

Table 4 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–$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 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 (3–5 percentage points) documented in Table 1. The results are consistent with a simple labour supply model that predicts that 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.[21] 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 aged 25–29 lasted between 6 and 9 months in the 1980s and the early 1990s. As mentioned earlier, existing validation studies also suggest that social assistance (welfare) receipts are under-reported by a factor of about a third in standard government surveys. Correcting the $3,300 figure for under-reporting and the fact that people do not typically spend their whole year on social assistance yields an expected effect much closer to $1200.[22]

Fig. 6 shows graphically the social assistance (total money transfers) results for the linear spline models. The confidence bounds (95% level) are also reported in the figure. There is clear visual evidence of a discontinuity

---

[21]The difference in monthly benefits in 1985 is about $280 in Fig. 1.

[22]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 $\frac{7}{12}$ and $\frac{2}{3}$ is equal to $1,283, which is very close to the estimated effect.
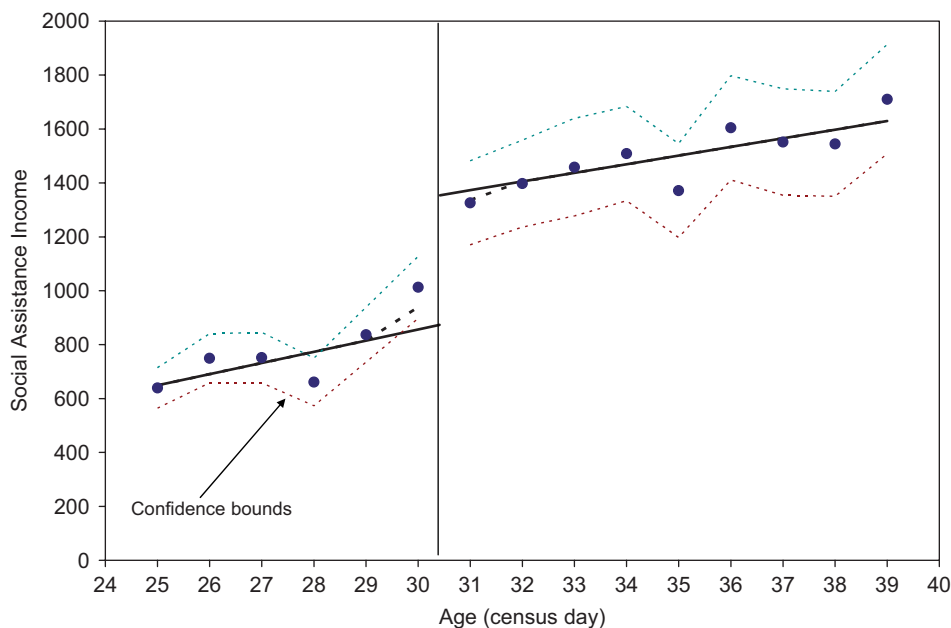
Fig. 6. Social assistance income, Quebec 1986.

around age 30. Note that social assistance money transfers are trending up as a function of age because of the negative selection problem discussed earlier. By contrast, the total dollar value of social assistance benefits conditional on receiving some benefits (not reported here) 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 in Table 3 shows 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 UI benefits. This suggests, at best, weak substitution effects between social assistance and UI.

Column 5 shows that higher social assistance benefits generally have 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% level). In fact, the standard errors are too large to make it possible to distinguish among some reasonable null hypotheses. A first hypothesis is that workers affected by the higher benefits are representative of all workers. Given the estimated effect on weekly hours in Table 1 (1.72 for the linear spline model) and the average hourly wage of \$11, 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).[23] By contrast, if workers affected at the margin 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 programme 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.

In Lemieux and Milligan (2004), we also report results from "falsification tests" 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.

---

[23]The average hourly wage is obtained by dividing average earnings (\$13,924) by average weekly hours (24.39) times 52 weeks.

Table 5
Cell means and regression results from the Census

	Sample/specification	Obs.	Employed	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)

Note:
[***]Indicate statistical significance at the 1% level.
[**]For the 5% level, and [*]for the 10% level.

7. Comparing RD and difference-in-differences results

In this section, we briefly compare the results from the regression discontinuity estimator to results from more traditional difference-in-difference estimators. A more detailed comparison is provided in Lemieux and Milligan (2004). In the PRWORA era, many researchers have pursued difference-in-difference strategies to measure the effect of welfare reform across states and years. 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. We focus our analysis on Census data but also refer to additional evidence based on the LFS.

7.1. 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). We start with the group of Quebecers at age 29 in the 1986 Census. Means for this base group are shown in the first row of Table 5. The first "discontinuity" comparison we draw looks at Quebecers age 30 in 1986. This simple cell mean estimator (or simple difference estimator) will be unbiased if there is no trend in the age profile for the dependent variable. The estimated effect for employment is substantial (−0.052) a strongly significant (we discuss the effect on social assistance take-up below).

To control for this age effect, we use three difference-in-differences strategies. 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. The difference-in-differences estimate reported on the third row of Table 5 (−0.044) is similar to the simple discontinuity estimator, suggesting a small age effect.

Second, we can use Quebec in 1991 as a control group. This control group will produce unbiased estimates 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.[24] The difference-in-differences estimate reported on the fourth row of Table 5 is now substantially larger (−0.079) than the simple discontinuity estimator. The estimator is larger than the others because the age gap for employment in 1991 was 2.7 percentage points in favour of those aged 30, compared to 5.2 percentage points in the other direction in 1986. This suggests 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.

Third, we can compare Quebecers at age 29 in 1986 to those age 29 in 1991, and then compare the result to the same difference in the rest of Canada. To be unbiased, this 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. Here, the difference-in-differences estimate is positive (0.01) but not significant. 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 three difference-in-differences strategies suggest an alternative estimator that combines all the control groups into a 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 as long as the age profile of the dependent variable does not shift differentially between 1986 and 1991 in Quebec and the rest of Canada. This triple-difference strategy yields an estimated effect of −0.079 on employment. We also report at the bottom of Table 5 an alternative version of the triple-difference estimator where the micro data are used to control for additional covariates: 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.[25] This regression version of the triple-difference model yields a very similar estimate (−0.074) to the model based on cell means only.

We then report the results with the presence of transfer income used as the dependent variable in the second column of Table 5. The discontinuity cell mean estimator suggests an effect of 0.058. However, Fig. 6 makes clear that transfer receipt trends up with age, so the comparison of 29–32 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 (0.039 and 0.040). These estimates are nonetheless economically large, representing 21% of the mean for this variable. As in the case of the employment rate, however, the third difference-in-differences estimate has the wrong sign and is not statistically significant. This is the estimator where Quebecers at age 29 in 1986 are compared to those age 29 in 1991, and then compared to the result to the same difference in the rest of Canada, without any control in the same labour market.

In results detailed in Lemieux and Milligan (2004), we use the LFS to augment our Census results. Because the Census requires a five year lag between data points, we are unable with that data to compare outcomes just before and just after the change in policy. With the LFS, we can form a window of 36 months on either side of the policy change as a base for estimation. Using the LFS, we find results similar to the RD estimates in Sections 4 and 5: the triple difference estimator in the LFS yields an estimate for employment of −0.050 (standard error 0.016).

In summary, 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. Conclusions

Using a unique policy episode involving lower social assistance payments to those under 30 in Quebec, we studied the effects of a transfer programme on labour market outcomes. Our main finding is that more

---

[24]Factors other than aggregate shocks (like the 1991 recession) may also make 1991 a questionable control group. For example, because of secular changes in the level of education, high school dropouts in 1991 may not be comparable in terms of their abilities to those in 1986. Unlike the difference-in-differences approach, the RD approach is robust to this type of compositional change.

[25]The treatment variable in this regression model is a dummy for men age 30 in Quebec in 1986. This effect is estimated controlling for the covariates mentioned in the text and for age effects, province effects, year effects, and bivariate interactions between these three dummies. See Lemieux and Milligan (2004) for more details.

generous social assistance benefits reduces 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. While these effects are precisely estimated, some care must be taken in interpreting their economic magnitude. After all, a three to five percentage point response to a 175% increase in benefits is consistent with relatively modest behavioural effects.

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 social programmes 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, as well as 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 ages 29–30 only. More generally, all we can identify here is a "treatment effect on the treated" that might not generalize to other groups. Furthermore, our results for single employable males may not be relevant for similar programmes in countries like the United States where this group is not eligible for welfare benefits.

We also have several interesting methodological 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) improves inferences. Finally, we conclude that 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.

Acknowledgements

Appendix A

Cell size and sample composition for men in the 1986 census are shown in Table A1. Figs. A1 and A2 show the fraction of HS dropouts with children and employment rate for all HS dropouts, respectively.

Table A1
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

Table A1 (*continued*)

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
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% sample of the Canadian Census. The acronym "RoC" stands for the "rest of Canada".
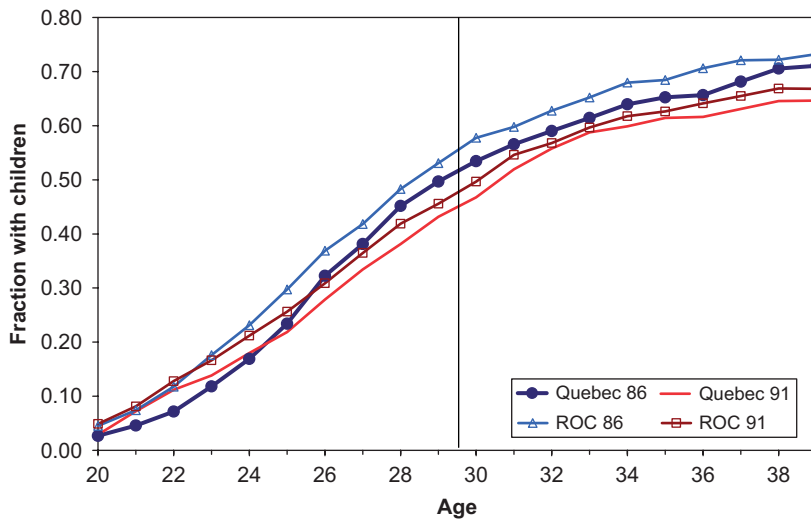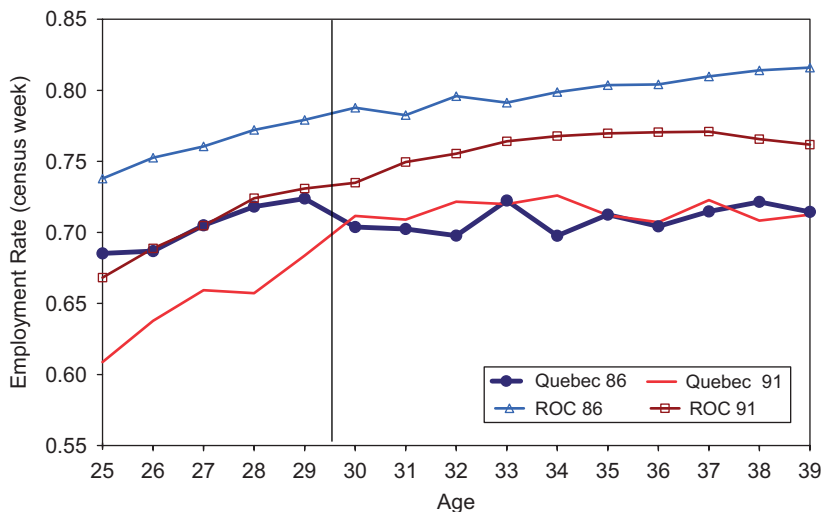

Fig. A1. Fraction of HS dropouts with children.

Fig. A2. Employment rate for all HS dropouts.

References

Ashenfelter, O., 1983. Determining participation in income-tested social programs. Journal of the American Statistical Association 78, 517–525.

Baker, M., Payne, A.A., Smart, M., 1999. An empirical study of matching grants: the 'cap on CAP'. Journal of Public Economics 72, 269–288.

Barrett, G.F., Cragg, M.I., 1998. An untold story: the characteristics of welfare use in British Columbia. Canadian Journal of Economics 31, 165–188.

Black, D.A., McKinnish, T.G., Sanders, S.G., 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 87, 1921–1942.

Black, S.E., 1999. Do better schools matter? Parental valuation of elementary education. Quarterly Journal of Economics 114, 577–599.

Blanchard, O.J., 2004. The economic future of Europe. Journal of Economic Perspectives 18, 3–26.

Blank, R.M., 2002. Evaluating welfare reform in the United States. Journal of Economic Literature 40, 1105–1166.

Dooley, M.D., 1999. The evolution of welfare participation among Canadian lone mothers, 1973–1991. Canadian Journal of Economics 32, 589–612.

Dooley, M.D., Gascon, S., Lefebvre, P., Merrigan, P., 2000. Lone female headship and welfare policy in Canada. The Journal of Human Resources 35, 587–602.

Ford, R., Gyarmati, D., Foley, K., Tattrie, D., 2003. Can work incentives pay for themselves? Final report on the self-sufficiency project for welfare applicants. Social Research and Demonstration Corporation, Ottawa.

Fortin, B., Lacroix, G., Drolet, S., 2004. Welfare benefits and the duration of welfare spells: evidence from a natural experiment in Canada. Journal of Public Economics 88, 1495–1520.

Green, D.A., Warburton, W.P., 2004. Tightening a welfare system: the effects of benefit denial on future welfare receipt. Journal of Public Economics 88, 1471–1493.

Hahn, J., van der Klaauw, W., Todd, P., 2001. Identification and estimation of treatment effects with a regression-discontinuity design. Econometrica 69, 201–209.

Harvey, C., Camasso, M.J., Jagannathan, R., 2000. Evaluating welfare reform waivers under Section 1115. Journal of Economic Perspectives 14, 165–188.

Hoynes, H.W., 2000. Local labor markets and welfare spells: do demand conditions matter? Review of Economics and Statistics 82, 351–368.

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

Kapsalis, C., 2001. An assessment of EI and SA reporting in SLID. Statistics Canada Analytical Research Branch Working Paper no. 166.

Lee, D.S., 2007, Randomized experiments from non-random selection in U.S. House elections. Journal of Econometrics, this issue, doi:10.1016/j.jeconom.2007.05.004.

Lee, D.S., Card, D., 2007, Regression discontinuity inference with specification error. Journal of Econometrics, this issue, doi:10.1016/j.jeconom.2007.05.003.

Lemieux, T., Milligan, K., 2004. Incentive effects of social assistance: a regression discontinuity approach. National Bureau of Economic Research Working Paper No. 10541.

McCrary, J., 2007. Manipulation of the running variable in the regression discontinuity design: a density test. Journal of Econometrics, this issue, doi:10.1016/j.jeconom.2007.05.005.

Milligan, K., 2005. Subsidizing the stork: new evidence on tax incentives and fertility. The Review of Economics and Statistics 87, 539–555.

Mincer, J., 1974. Schooling, Experience and Earnings. Columbia University Press, New York.

Moffitt, R.A., 2002. Welfare programs and labor supply. In: Auerbach, A., Feldstein, M. (Eds.), Handbook of Public Economics, vol. 4. North-Holland, Amsterdam, pp. 2393–2430.

Moffitt, R.A., 2003. The temporary assistance for needy families program. In: Moffitt, R.A. (Ed.), Means-tested Transfer Programs in the United States. University of Chicago Press, Chicago, pp. 291–363.

Murphy, K.M., Welch, F., 1990. Empirical age–earnings profiles. Journal of Labor Economics 8, 202–229.

van der Klaauw, W., 2002. Estimating the effect of financial aid offers on college enrollment: a regression-discontinuity approach. International Economic Review 43, 1249–1287.