A Reanalysis of 'The Town with No Poverty: The Health Effects of a Canadian Guaranteed Annual Income Field Experiment'

David A. Green *

March 5, 2021

Abstract

In 'The Town with No Poverty: The Health Effects of a Canadian Guaranteed Annual Income Field Experiment' (published in the CPP in 2011), Professor Evelyn Forget examines community level health and education impacts of the 1970's Manitoba Guaranteed Annual Income experiment (MINCOME). In this paper, I reassess the hospital use data from Forget(2011), arguing that once one takes account of pre-trends in the differences between the treatment and control samples, the data point to an increase in hospital use during MINCOME. This is the opposite of the conclusion reached in Forget(2011). Combined with results from Forget(2011) on birthweight of new borns, I argue that the MINCOME data does not support conclusions that the Guaranteed Annual Income either improved health or reduced health care costs.

*Vancouver School of Economics, UBC, and Institute for Fiscal Studies, david.green@ubc.ca. I am grateful to Evelyn Forget for an extended email discussion of this material. My hope is that I have captured that discussion faithfully. I also thank Lindsay Tedds for useful comments.

1 Introduction

In 'The Town with No Poverty: The Health Effects of a Canadian Guaranteed Annual Income Field Experiment' (Forget (2011), published in the CPP in 2011), Professor Evelyn Forget examines community level health and education impacts of the 1970's Manitoba Guaranteed Annual Income experiment (MINCOME). The analysis takes advantage of a unique element of MINCOME: that the town of Dauphin, Manitoba was a saturation site with the Guaranteed Annual Income (GAI) being available to all residents of the town and its immediate rural surroundings. The results in this paper were made possible by impressive data work by Professor Forget in gathering individual level administrative data related to both high school completion and health system use by all the residents of Dauphin. This makes possible an examination of impacts on the community as a whole, incorporating potential spillover effects beyond the direct effects on those who actually received the GAI during the experiment.

The empirical approach used in Forget (2011) (hereafter Forget11) is a difference in difference type exercise in which outcomes for residents of Dauphin are compared with those for a matched sample before, during and after the MINCOME experiment. The key result from the empirical work on health outcomes is that hospitalization rates for Dauphin residents fells by 8.5 percent relative to the comparison group. This is a very substantial reduction which leads Professor Forget to the conclusion: "While we recognize that one must be careful in generalizing potential health system savings, particularly because we use hospitals differently today than we did in 1978, the potential saving in hospital costs associated with a GAI seems worthy of consideration." (p. 300)

Following from this conclusion, this article has had a substantial impact on thinking about the effects and viability of a GAI. As of February, 2021, it has 322 citations in Google Scholar. The papers citing it tend to discuss it as either providing evidence that a GAI can cause improved health outcomes (e.g., Gibson et al. (2020)) or that it can lead to reduced health care system expenditures, freeing up money to pay for the GAI (e.g.,Widerquist (2017)) or both. Because it is one of the few papers providing evidence at the community level, it carries considerable weight in the debate over a GAI.

The data work in the paper is certainly impressive and deserves attention, but in this comment I argue that a reanalysis of the data in Forget11 alters key conclusions.[1] In particular, I argue that there are two elements of the data that, on closer examination, raise

[1]The reanalysis of this data was part of a two year project examining the effectiveness of a basic income as a policy tool. For the final report from that project and over 40 other research papers related to basic income see https://bcbasicincomepanel.ca/ .

concerns. The first is that there is evidence of a trend downward in hospital use in Dauphin relative to the comparison group before the start of MINCOME that continues after the experiment ends. This feature of the data was not considered in the original analysis and once it is, it points to the conclusion that hospital use rose (not fell) in Dauphin during the MINCOME period relative to the comparison group. The second key pattern in the data is that to the extent that hospital utilization fell substantially in Dauphin relative to the comparison group, it occurred in the years after MINCOME ended. This calls into question Professor Forget's explanation for the large estimated hospital effect, that it could have arisen because even people other than those in direct receipt of the GAI would have felt lower stress because of the insurance that potential access to the GAI provided. It is hard to see how this insurance effect could have been largest after access to the GAI had ended.

At first glance, the finding in the reanalysis that hospital use went up in Dauphin may appear to be at odds with the literature that finds a positive correlation between income and health outcomes. However, it is important to point out that health system utilization and health outcomes are different entities. Access to a GAI could allow people to take time off work to attend to long standing health concerns that require hospitalizations or even to get care for immediate injuries that would require taking time off work. This could suggest that the effects of a GAI on health are positive but that the mechanism does not imply cost reductions that could help fund the GAI - at least, not in the short to medium term. On health outcomes, the main outcome in Forget11 that is clearly beyond choice in terms of timing of when to go to hospital is newborn birthweight. There, Forget11 finds no significant effect from MINCOME. Taken together, I believe the data in the paper points to the conclusion that the MINCOME GAI did not reduce hospital costs (and may have increased them) and that it provides no conclusive evidence on positive or negative health effects. This is not to deny that a GAI could have positive effects, only that MINCOME provides no evidence that is does so.

The paper proceeds in two main sections. In the first, I describe the MINCOME experiment as well as the data, methodology, and results from Forget11. I focus on elements that are essential to the re-analysis, referring readers to the original article for more details on all other elements. In the second, I present a re-analysis of the results. Following the approach in Forget11, I do this both graphically and through a regression analysis of hospitalization usage rates.

2 Elements of Forget (2011)

2.1 The MINCOME Experiment

The MINCOME GAI experiment was a joint initiative of the Canadian federal government and the Manitoba provincial government. More details on the experiment can be found in Forget (2011) and Forget (2018). For our purposes, there are several key features of the experiment that require emphasis. First, MINCOME had a regular, randomized treatment component but also, unique among GAI experiments, it included a saturation site. Every family in the western Manitoba town of Dauphin and its adjacent rural municipality was eligible to take part in MINCOME. Second, the payments under MINCOME amounted to 60% of the LICO low income line, varying by family size, and had an associated tax back rate of 50% (i.e., for every dollar a family earned, their MINCOME benefit was reduced by 50 cents until they reached the break even level beyond which no benefits were paid). Forget (2011) states that about a third of families in Dauphin had incomes below the break even income level (i.e., were eligible for some benefit payment) but that for many of them the 'stipends would have been quite small.'(p. 291). Third, in a difference-in-difference analysis, the precise timing of the 'treatment' period (the period during which MINCOME benefits were paid out, in this case) is crucial. A baseline survey was conducted in late 1974 to gather background information then participants were enrolled between November, 1974 and October, 1975, with payments made for three years. That means that there is a small set of payments made in 1974, with the final payments made in 1979 and that 1974 should be considered a non-treatment year while 1975 through 1979 are all treatment years.[2] Fourth, the main related policy change near the time of the experiment was the introduction of nationalized health care. On May 28, 1967, the Manitoba legislature passed Bill 68, formally joining the new federally directed medicare system, however actual implementation was post-dated to July 1, 1969. Thus, 1970 was the first full year of medicare in Manitoba.

2.2 Data

The health data used in Forget11 is drawn from the Manitoba Population Health Research Data Repository and contains data on 'almost every physician and hospital contact in the

[2]Surprisingly, getting precise information on the timing of MINCOME is not straightforward. For example Hum et al. (1979), the first technical report put out for MINCOME states that the benefits stopped flowing in December of 1977, but this appears not to be true. For an re-analysis of MINCOME data on labour supply, including a discussion of related data issues, see Riddell and Riddell (2020). Chris Riddell kindly pulled up the MINCOME data he is using to confirm that few of the experiment participants received benefits in 1974 and that payments continued well into 1979.

province' (p. 292). Because the data comes from the universal health insurance system, it starts in the first full year of that system, 1970. The main outcome variable examined in the paper is the rate of hospitalization, measured as the number of hospital separations examined primarily as a rate relative to the community population. In a sub-analysis, this is broken down into separations by cause of visit, specifically looking at accidents and injuries, and non-congenital mental health. There is also data on physician visits and low birth weight for newborns. Using this data to examine the impact of MINCOME is an intelligent approach to assessing the impacts of the experiment and bringing the data to bear on these questions is one of the substantial contributions of Forget11.

2.3 Methodology

The paper adopts a generalized difference-in-differences approach to examining the impact of MINCOME on health outcomes.[3] The logic of this approach is straightforward. One could consider estimating the impact by comparing the outcomes for residents of Dauphin during the MINCOME experiment with their outcomes before it was implemented. The obvious difficulty with this is that there could be a trend in hospital usage independent of the MINCOME experiment (for example, as the new universal health care system developed). But, in that case, one could not tell whether any difference in usage during the MINCOME period was due to the effects of the experiment or simply reflected the general trend. The response in the difference-in-differences approach is to calculate usage rates for a comparison group who are intended to capture what would have happened to the residents of Dauphin during the MINCOME period if the experiment had not happened. Calculating the difference in their usage before versus during the experiment period and subtracting it from the same difference for the residents of Dauphin (the 'difference in differences') is then argued to remove the general trend and reveal the actual effect of the experiment.

Forget11 calls all the residents of Dauphin the 'subjects' and the comparison group, the 'comparators'. To form the comparator group, each resident of Dauphin is matched (using propensity matching) with 3 residents of other parts of the province (excluding Winnipeg, First Nations reserves, and the north since their circumstances are judged to be too different from Dauphin) based on a list of characteristics, including age, sex, number of people in the family, whether the residence was rural, and whether it was a lone parent family. The paper includes a very complete exercise examining whether there are differences between Dauphin and the comparators' communities in features such as income, employment, and religion using

[3] Forget11 uses terminology from the biostatistics literature while I will use terminology from econometrics in this paper.

the 1971 Census. This exercise is somewhat reassuring in that the comparators' communities are quite similar in terms of the levels of these characteristics. But similarities in the levels of variables are less important for establishing whether the estimator works than similarities in the trends in hospital usage before the experiment (what the literature calls no difference in pre-trends) because what we use the comparator group for in the difference-in-differences approach is not to establish the level of hospital usage in the absence of the experiment but how usage would have changed in Dauphin if the experiment had not occurred. I will return to that point in the reanalysis section.

It is worth highlighting that the subjects group consist of individuals who lived in Dauphin throughout the 1974 through 1978 period and does not include people who either moved in or out. This is a wise design choice that makes sure that differences between the subjects and controls can't arise because of, say, choices to move to Dauphin in order to qualify for MINCOME.

2.4 Results

Forget11 presents the examination of the hospitalization impact of receipt of MINCOME benefits in two ways. The first is graphical, plotting the number of hospitalization separations per 1,000 residents of either Dauphin (the 'subjects') or comparator communities (the matched 'comparisons'). This is shown in Figure 2 in Forget11. In Figure 1, here, I re-plot those lines using data over the period 1970 through 1985 from Forget (2013).[4][5]

[4]Figure 2 in Forget11 shows the same data for the period 1970 through 1980, inclusive. In an email discussion, Professor Forget expressed some concern about the post 1980 data because the Dauphin sample, constrained as it was to include only people who were resident in Dauphin for the entire 1974-78 period, could become less and less representative of the town over time. However, the regression analysis appears to use data for the longer 1971 through 1985 period. This is not specified in Forget (2011) but is stated as the sample period in the header to Table 1 in Forget (2013), which contains identical estimates to those in Table 2 in Forget (2011). For consistency, I present both graphical and regression results using the entire 1970-1985 period. The fact that regression results are nearly identical to those in Forget11 when using the entire period supports the conclusion that the estimation reported in the article used the entire period. The key conclusions drawn here are not altered if I truncate the sample in 1980.

[5]The numbers shown in this figure and used in the regression analysis in this paper are based on reading from Figure 1 in Forget (2013), which appears to be identical to Figure 2 in Forget (2011). The figure from Forget (2013) has horizontal lines at each 10 digits, making it easy to get quite close to the true values. Professor Forget kindly dug up files from over a decade ago with values for the rates for the years 1970 through 1974. Those values do not exactly match what is in Figure 1 from Forget (2013). I decided to work entirely with my readings from Figure 1 from Forget (2013) since that appears to be what was used in the figures and tables in Forget11. The main results and conclusions presented here are not altered if I use the pre-1974 data supplied by Professor Forget.

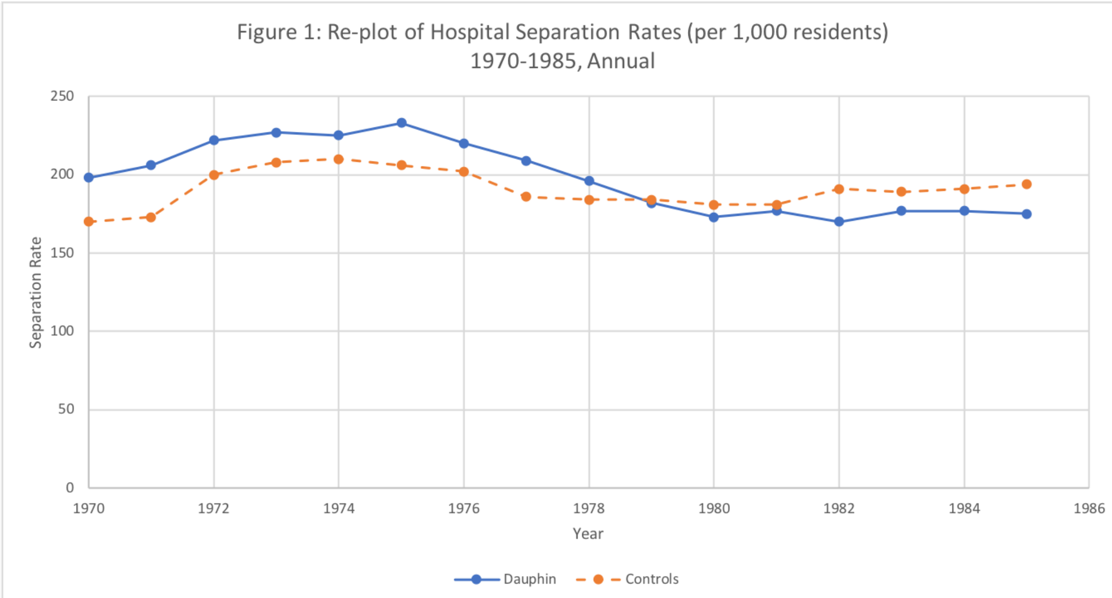

As Forget11 discusses, Figure 1 shows that residents of Dauphin had higher hospitalization rates after the start of universal healthcare compared to the matched comparison group. In assessing this difference, Forget11 states, 'None of the health or census variables that we examined could explain the persistent gap between subjects and controls before 1974, but we note that there was a fairly new hospital, which may have led to some supply-induced demand.'(p. 295) Hospitalization rates were 8.5 percent higher than the comparator group in the years up to 1974 but declined relative to the control group to a point of equality in 1979. This is the basis of the key conclusion in Forget11 that 'hospitalization rates among Dauphin subjects fell by 8.5 percent relative to the comparison group' because of the MINCOME experiment.

Forget11 also includes a regression version of a difference-in-differences analysis that is referred to in biostatistics as a 'segmented time-series model'. This confirms a declining trend in the hospitalization rate in Dauphin relative to the controls after 1974. The form of that analysis is non-standard, at least relative to the econometrics literature, in ways I will describe in the re-analysis section but simply working with a more standard econometric specification does not change the results. What does turn out to matter is how one treats a potential pre-trend and how one defines the end of the treatment period.

The breakdowns of the regression analysis by type of admission show the same patterns for both the accidents and injuries category and the mental health category: Dauphin has higher hospitalization rates than the controls for both categories at the start of the period but that difference has disappeared by the end of the MINCOME experiment. In contrast, no effect is found for low birth weight, at-risk birth weight, or small-for-gestational-age rates for newborns. Forget11 argues this non-effect is likely due to a lack of food insecurity issues in rural communities at the time.

3 Re-Analysis

In this section, I present a re-analysis of both the graphical and regression results in Forget11, arguing that the re-analysis alters key conclusions drawn in the paper.

3.1 Graphical Approach

Recall that the difference-in-differences approach involves a comparison of how differences between the Dauphin residents and the matching controls change over time - from before to during and, potentially, after the experiment. Figure 1 presents the data involved in that comparison but is difficult to read in terms of what we really care about - whether hospitalizations rates in Dauphin evolved differently from those among the controls - because our eye is caught by long term trends common to both groups that are not a reflection of the experiment effects. In Figure 2, I present the difference between the two lines in Figure 1. This effectively removes trends over time that are common to both groups.

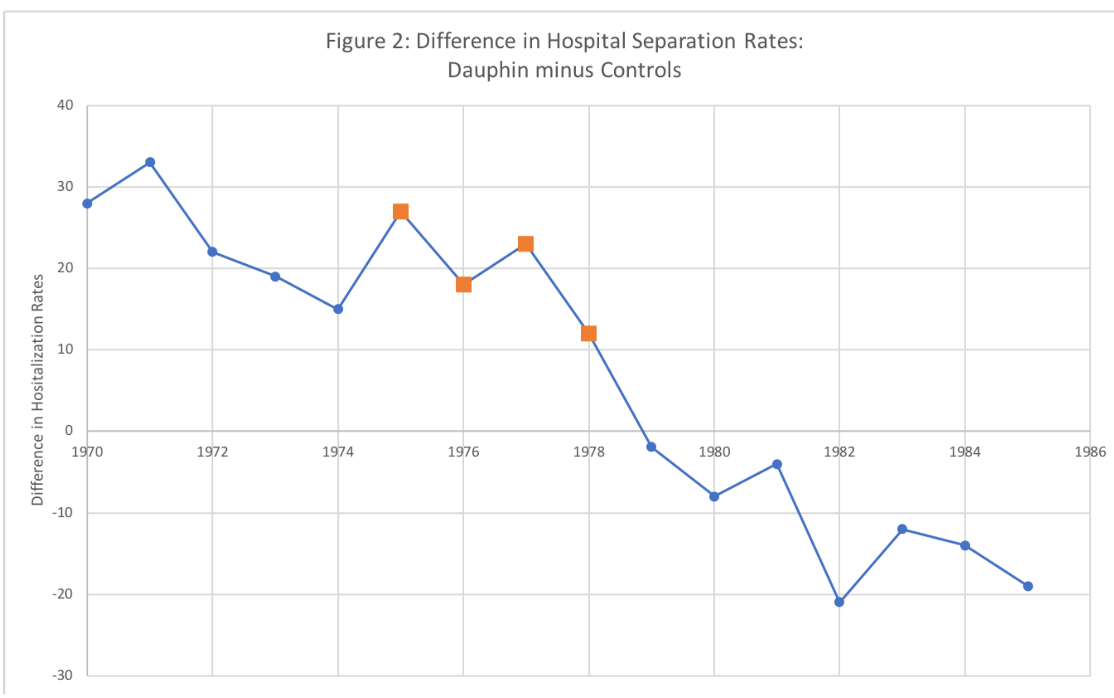

In Figure 2, the points with squares denote the MINCOME experiment years (1975 through 1978, inclusive). It is worth pointing out what we would expect to see in this figure if it were to provide clean evidence of MINCOME reducing hospital use. We use the controls as benchmark of how hospital use would have changed over time in Dauphin in the absence of MINCOME. To have confidence in that benchmark, we would want to see that changes over time before MINCOME were the same for the controls and Dauphin. Otherwise, what we claim is a MINCOME effect could just be part of the relative difference in trends between the two groups. In the figure, having identical changes over time would mean that the difference line would be a flat line right on the zero axis in the years before MINCOME. Then, if MINCOME reduced hospital use, we would see the difference line move into negative territory during the MINCOME years. Once the experiment was over, we would expect to see the line return to being flat at zero difference, though there might be lingering effects of MINCOME that would mean it would take a few years to get back to zero difference.

This zero-negative-zero pattern is not what is seen in Figure 2. Instead, several points jump out. The first is that there is an evident pre-trend according to which the Dauphin hospitalization rates were converging to those in the control group in the years before the experiment - the line is not flat in the pre-MINCOME years as we would want in a clean analysis. Indeed, extending a rough line following the pre-trend would seem to imply that the difference between Dauphin and the controls would have been eliminated by approximately 1979 or 1980 even in the absence of the experiment. I don't have an explanation for the source of this trend but its source may have to do with the reason hospital use was higher in Dauphin than among the controls in the early 1970s - something for which Forget11 is unable to find an explanation in spite of careful investigation.

Second, the downward trend in the difference continues for the 7 years after MINCOME ended, leading to Dauphin hospitalization rates falling below those among the controls. This, too, does not fit with an ideal difference-in-differences analysis, in which the line in Figure 2 would (at least eventually) become flat at zero. This is important because, as we will see, Forget11 defines the experimental treatment period in the regression analysis to include all years after 1974, i.e., to include these post-experimental years when hospitalization rates were particularly low for Dauphin. It is hard to see how this is justifiable. Forget11 notes that the 8.5% decline in hospitalizations that the paper argues happened in Dauphin because of MINCOME seems very large given that 'only about a third of Dauphin families qualified for MINCOME stipends at any point, and because of the structure of the payment scheme, many of those stipends would have been quite small.'(p.291) The argument provided for this is that the GAI acted as a form of insurance such that even people not receiving the benefits (or receiving small amounts) knew it was there for them if they had troubles. This led to a reduction in stress which, in turn, led a reduction in mental health related hospital visits and a reduction in actions that could lead to physical injury. While this is plausible, it is an explanation that should only operate while the GAI is actually available. It is hard to see how the memory of having an accessible back-stop would continue to have an effect for years afterward. In fact, in the data the largest declines in hospitalisation use in Dauphin relative to the controls happens in the first year after the experiment was over and between 1981 and 1982.

Third, the MINCOME years appear to be, if anything, a departure from the trend in a positive direction - the opposite of what we said a clean difference-in-differences analysis with a negative hospitalisation effect would show. That is, the difference in differences approach actually appears to imply that MINCOME led to an increase in hospital use. Of course, there is a huge body of work that establishes a positive correlation between income and health outcomes [6] and this result may seem to be in opposition to those results. One possibility is that a basic income could cause adverse health outcomes. Gibson et al. (2020), in their review of the literature on GAI experiments state that some adverse effects have been found, notably related to substance abuse. But this seems unlikely to explain effects at the level of a whole community. A more plausible explanation lies in the distinction between health outcomes and health care system usage. A more secure income base might allow people to take time off work to address their health concerns. This is a plausible channel for explaining the positive relationship between income and health outcomes, but one that would involve an increase rather than a decrease in health system usage.

3.2 Regression Analysis

The examination of Figure 2 generates some interesting observations but, of course, are necessarily imprecise. One would like to know, for example, whether the apparent pre-trend is statistically significantly different from a flat trend and whether the higher usage rates relative to trend during the MINCOME experiment are statistically significant. In this subsection, I examine Forget11's regression analysis that is intended to allow precise answers to these statistical questions. For those who are uninterested in the technical details, let me list the key conclusions from my regression analysis. First, I am able to replicate Forget11's results, showing a MINCOME effect in the form of a negative trend in the difference between Dauphin and the controls using her specification. Second, this result begins to change once one moves to a more standard specification in which MINCOME is allowed to affect the level of hospital use as well as its trend and its effect is restricted to the actual MINCOME years rather than all the years from 1974 through 1985. Third, once one allows for a long term trend

[6] See, for example, Schmidt et al. (2020), who review the literature on the relation between child poverty and a variety of outcomes, arguing for an epigenetic mechanism.

in the difference between Dauphin and the controls in hospital use, a statistically significant negative trend emerges. During the MINCOME years, the relative trend in hospital use between Dauphin and the controls is no different from the long term, negative trend but there is a jump up in the level of hospital use in Dauphin. That is, what your eye tells you in looking at Figure 2 is true and cannot be explained as simple sampling variation. Thus, the conclusions reached from looking at Figure 2 stand: MINCOME did not lower hospital usage rates, it raised them.

Now, I turn to the details behind these conclusions. Forget11 implements a segmented time series model in order to allow for tests of significance within the context of the difference-in-differences approach. The model is implemented using individual level data and a negative binomial estimator. This is an estimator that is appropriate for count data (in this case, counting the number of hospital separations) and is implemented in Forget11 in such a way that it effectively shows impacts on the rate of hospital usage at the community level. I will work with the community level rates recorded in Figure 1. Working at this aggregate level focuses attention on the actual level of the identifying variation - across time by community.[7] The matching procedure implies that no further controls for individual characteristics are necessary.

Column 1 of Table 1 contains the estimated coefficients reported in the first column of Table 2 in Forget11. Recall that we need to control for general time trends in order to make sure that the changes observed in Dauphin over time are not just reflections of ongoing trends. One could accomplish this by including a simple linear trend but Forget11, appropriately, includes a more flexible specification that allows the trend to change in level and slope starting in 1974 and again from 1979 onward. The Constant and the following five coefficients in the first column represent this flexible 'spline' time trend. The spline is identified by the comparison sample, representing a guess of what would have happened to the time trend of hospitalization use in Dauphin in the absence of MINCOME. It is not particularly interesting in its own right.

The 7th coefficient and standard error pair in the first column of the table corresponds to a dummy variable indicating whether the hospitalization rate is from Dauphin and shows that hospitalization rates were statistically significantly higher in Dauphin than among the comparison group. The next entry (Dauphin*Mtrend) is an interaction of the Dauphin dummy variable with a trend the begins in Forget11's claim of the first year of MINCOME (1974) and continues for the remainder of the sample period. Forget11 takes the negative

[7]This has the advantage of automatically providing standard errors at the correct level of clustering. I use robust standard errors in order to account for any potential heteroskedasticity arising from the use of means for the dependent variable.

Table 1: Regression Re-Analysis

Variable	1 Forget11	2 Spline	3 Year Effects	4 Replications 1974 Not Treatment	5 Incl. Treatment Level Effect	6 Only MINCOME years in Treatment	7 With Pre-trend
Constant	-2.50**	-1.83**	.	.	.	.	.
	(0.023)	(0.023)	.	.	.	.	.
Trend	0.029**	0.031**	.	.	.	.	.
	(0.0042)	0.0038)	.	.	.	.	.
D74+	0.044	0.054	.	.	.	.	.
	(0.025)	(0.03)	.	.	.	.	.
Trend74+	-0.045**	-0.044**	.	.	.	.	.
	(0.0052)	(0.005)	.	.	.	.	.
D79+	-0.041	-0.071*	.	.	.	.	.
	(0.023)	(0.021)	.	.	.	.	.
Trend79+	0.021**	0.021**	.	.	.	.	.
	(0.0036)	(0.0038)	.	.	.	.	.
Dauphin	0.13**	0.13**	0.13**	0.13**	0.12**	0.012	0.17**
	(0.015)	(0.017)	(0.015)	(0.015)	(0.020)	(0.031)	(0.014)
Dauphin*	-0.011**	-0.010**	-0.010**	-0.011**	-0.012**	-0.0074*	0.0015
Mtrend	(0.0013)	(0.0010)	(0.0010)	(0.0011)	(0.0013)	(0.0027)	(0.0029)
Dauphin*	.	.	.	.	0.025	0.12**	0.053**
MINCOME	.	.	.	.	(0.026)	(0.034)	(0.016)
Dauphin*	.	.	.	.	.	.	-0.0089**
Trend	.	.	.	.	.	.	(0.00073)
Year Effects	No	No	Yes	Yes	Yes	Yes	Yes
Obs	.	16	16	16	16	16	16

Standard errors in parentheses. *,** significantly different from 0 at 5%, 1% level of significance
Trend is 2 times annual count with 1970 = 1; D74+ and D79+ are dummy variables for all years 1974 and after, and 1979 and after; Trend74+ is 2 times annual count with 1974=1; Trend79+ is defined analogously; Dauphin is a dummy for residents of Dauphin; Mtrend is 2 times annual count with the first year of MINCOME treatment = 1; MINCOME = 1 for treatment years, which is all years from 1974 onward in columns 1 and 2, 1975 onward in column 3, and the experiment years (1975 through 1978) in the remaining columns.

and significant coefficient on this interaction as showing that MINCOME caused a relative trend downward in hospital use in Dauphin.

Column 2 of the table is a replication of column 1 using the usage rates taken from Figure 1 in Forget (2013) for the years 1970 through 1985. The direct analogy to the negative binomial estimator used with individual data and implemented in SAS as in Forget11 is a simple regression with the log of the proportion of the community using a hospital as the dependent variable, and that is the specification I use in the remainder of the table. A comparison of columns 1 and 2 shows that this specification yields estimates that are very similar to those in Forget11.[8] Most importantly, the estimated coefficients for the Dauphin dummy variable and the Dauphin*Mtrend interaction are identical in size and statistical significance. Thus, working with the aggregate data is sufficient to capture the same data patterns as in Forget11.

In column 3 of the table, I replace the spline in time with a complete set of year fixed effects. This is a more standard in the econometrics literature and even more flexible way to capture general trends. Working with it reduces clutter in the table and any complications in interpreting trends.[9] Fortunately, it also yields identical estimates to those in the second column and in Forget11. Given that the dependent variable is the log of the utilization rate, the coefficient on Dauphin implies that usage rates were approximately 13% higher in Dauphin than the comparisons. The coefficient on the Dauphin*Mtrend interaction imply that over the 5 years after where Forget11 marks the start of the experiment, usage rates declined by 5% in Dauphin relative to the comparisons.

In column 4, I change the definition of the onset of MINCOME to be 1975, given the argument earlier that too few people received MINCOME benefits in 1974 to count it as part of the experiment. Making this change has almost no effect on the key parameters.

The specification in Forget11 is non-standard relative to the econometrics literature in two important ways. The first is that it allows the experiment to alter the trend in hospitalization usage but not its level. There is no reason to restrict the impact of the experiment in this way and doing so has the potential to generate distortions in other estimated coefficients, such as the one on the Dauphin variable. In column 5, I include the interaction of the the Dauphin dummy variable and a dummy variable for the years 1975 and after (the way Forget11 defines the treatment period). A standard difference-in-differences approach would not include the Dauphin*Mtrend variable and would interpret the coefficient no Dauphin*MINCOME as

[8]Forget11 has semi-annual observations on hospitalization usage while the data used here is at the annual level. In order to scale coefficiens on the trend variables the same, I multiply them by 2.

[9]I do not report the estimated coefficients for the year dummy variables to ease in the reading of the table.

the estimate of the treatment effect. To stay as close as possible to Forget11, I include both Dauphin*Mtrend and Dauphin*MINCOME. This yields only small changes in the coefficients on Dauphin and Dauphin*Mtrend.

The second non-standard element of the Forget11 specification is that it defines the whole period after the start of the experiment (1974 in the paper but 1975 in the columns beginning with column 4 in Table 1) as the period within which MINCOME has its effects. This implies that differences in the hospitalization rates between Dauphin and the comparisons in 1985 - 7 years after the end of the experiment - reflect the impact of the experiment. A much more standard approach is to assume that MINCOME had an effect only while it was actually in operation. As discussed earlier, this fits better with Forget11's argument that a GAI provides a backstop that reduces stress even when a family does not actually draw benefits from it. When I alter the definition of the treatment period in this way (in column 6), the results change substantially. In particular, there is now a significant and positive effect of being treated (i.e., living in Dauphin during the experiment) that amounts to a 12% increase in use relative to the controls in the same years. There is still a negative slope effect of the experiment on the trend but this decline would only amount to about a quarter of the increase in level over the course of the experiment.

Finally, the specification developed to this point still does not account for the apparent ongoing relative decline in hospital usage in Dauphin relative to the comparisons in both the pre and post-MINCOME years seen in Figure 2. It allows for a flexible general trend in hospital usage (captured in the year effects) but requires them to be the same for Dauphin residents and the controls in all years. In column 7, I include an interaction of the overall trend with the Dauphin dummy variable. This allows for the trends in usage to be different between Dauphins residents and the comparison group in a simple, linear way across the sample period. Importantly, the coefficient on that interaction is negative and statistically significantly different from zero, showing that the trend that is evident to the eye in Figure 2 is statistically significant. Further, the Dauphin*Mtrend coefficient is now small and statistically insignificant, implying that the relative usage trend in Dauphin within the 1975-78 MINCOME period is the same as in the pre and post MINCOME periods. The negative coefficient on this variable estimated in the other specifications is just a reflection of a more general negative trend that is not attributable to MINCOME. Instead, MINCOME has a statistically significant positive effect on the level of hospital use. Its coefficient implies that hospital use maintained the same relative decline as in the pre-MINCOME years but jumped up by about 5% relative to the comparisons during those years. I view this as the preferred specification because it allows for MINCOME to induce changes in both the level and trend of hospital use and for the pre-trend that is evident in Figure 2. As it turns out, that pre-trend is real and taking account of it yields regression estimates that ratify what the eye sees in Figure 2: hospital usage jumped up by about 5% in Dauphin relative to the controls in the MINCOME years.[10]

4 Conclusion

Forget (2011) contains an influential examination of the impact of the MINCOME Guaranteed Annual Income on the usage rates of hospitals, among other outcomes. The paper concludes that the availability of a GAI caused an 8.8% reduction in hospital usage and suggests that this may point to health improvements with associated cost reductions that could help pay for a GAI. In this paper, I re-analyse the impressive data from Forget (2011). I argue that there is evidence of a decline in hospital use in Dauphin relative to comparison communities beginning before the MINCOME experiment and continuing well after it concluded. When one takes account of this long term trend and only allows MINCOME to have an effect on outcomes while benefits were being paid out, the conclusion changes to MINCOME causing a 5% *increase* in hospital use. I argue this could be plausible if the availability of a GAI allows people to take the time to attend to health issues (much as people today argue that sickness benefits would allow workers who have Covid to stay away from work).

The positive impact on hospitalizations highlights the distinction between health outcomes and health care usage. In this case, MINCOME may have had positive effects on both. However, it is worth pointing out that there is one health indicator described in Forget (2011) that can likely be viewed as a direct measure of a health outcome: newborn birthweight. Forget (2011) finds no significant impact on this outcome. This is not at all to say that a GAI cannot have positive health effects. But it does say that the MINCOME data does not reveal such effects. What MINCOME does appear to reveal is that a basic income does not yield reductions in hospital costs, at least in the short to medium term. One cannot argue, based on MINCOME data, that a basic income could partly pay for itself

[10] I carried out a variety of robustness checks that do not alter the main conclusions here. In particular, in an email conversation expressed concern both about the quality of both the post-1979 data (because the sample potentially became less representative of the population of Dauphin after 1978) and in the first few years of the data (because the nationalised health system from which the data was drawn was just starting). I note that the year effects should capture these elements but, nonetheless, implemented versions of the empirical specification while dropping either the post-1979 or pre-1972 data without any changes in the conclusions. In another robustness check, I include the interaction of the Dauphin variable with the overall trend in a specification that includes all post-1975 years in the treatment period. In this case, too, the estimated treatment level effect is positive and significant while the treatment slope effect is small and insignificant.

through a reduction in such costs.

References

Forget, E. L. (2011). The town with no poverty: The health effects of a canadian guaranteed annual income experiment. *Canadian Public Policy 37*(3), 283–305.

Forget, E. L. (2013). New questions, new data, old interventions: The health effects of a guaranteed annual income. *Preventive Medicine 57*, 925–928.

Forget, E. L. (2018). *Basic income for Canadians: The key to a healthier, happier, more secure life for all.* James Lorimer and Company.

Gibson, M., W. Hearty, and P. Craig (2020). The public health effectso of interventions similar to basic income: A scoping review. *Lancet Public Health 5*, e165–e176.

Hum, D., M. E. Laub, and B. J. Powell (1979). The objectives and design of the manitoba basic annual income experiment, technical report no. 1. Technical report, Mincome Manitoba.

Riddell, C. and W. C. Riddell (2020). Does a universal basic income reduce labour supply for all groups? evidence from canada?s negative income tax experiment. Technical report, Research paper commissioned by the Expert Panel on Basic Income, British Columbia.

Schmidt, K. L., R. Gill, A. M. Gadermann, and M. S. K. Kobor (2020). Society to cell: How child poverty gets 'under the skin? to influence health now and forever. Technical report, Research paper commissioned by the Expert Panel on Basic Income, British Columbia.

Widerquist, K. (2017). The cost of basic income: Back-of-the-envelope calculations. *Basic Income Studies 16*, 1–13.