CARE OR CASH? THE EFFECT OF CHILD CARE SUBSIDIES ON STUDENT PERFORMANCE

Sandra E. Black, Paul J. Devereux, Katrine V. Løken, and Kjell G. Salvanes*

Abstract—Given the wide use of child care subsidies across countries, it is surprising how little we know about the effect of these subsidies on children's longer-run outcomes. Using a sharp discontinuity in the price of child care in Norway, we are able to isolate the effects of child care subsidies on both parental and student outcomes. We find very small and statistically insignificant effects of child care subsidies on child care utilization and parental labor force participation. Despite this, we find significant positive effect of the subsidies on children's academic performance in junior high school, suggesting that the positive shock to disposable income provided by the subsidies may be helping to improve children's scholastic aptitude.

I. Introduction

MANY countries have implemented child care subsidies in an effort to help families. In the United States, the government created the Child Care and Development Fund in 1996, which provides public funds for child care assistance to low-income families. Despite the importance of the issue, little is known about the effect of child care subsidies on parent and child outcomes. Research in this area has been limited because of the difficulty identifying the causal effect of child care price on later outcomes; for example, higher child care prices may be associated with better child care or wealthier parents, in which case one cannot isolate the effect of price alone on later child outcomes. This paper uses recent data and a novel source of identifying variation—sharp discontinuities in the price of child care by income in Norway—to identify the effect of child care subsidies on parental behavior and the later academic achievement of children.

A number of papers have examined the effect of child care subsidies on female labor force participation, with the findings ranging from no effect to significant negative effects (see Blau, 2000, for a summary). More recently, Herbst and Tekin (2010b) examined the effect of child care subsidies in the United States on children's academic performance.[1] They use a unique identification strategy, applying distance to the nearest social service agency that administers the subsidy application process as an instrument for subsidy receipt. They find small negative effects of subsidy receipt the year before kindergarten on kindergarten performance, although these negative effects have generally disappeared by third grade.[2] Our work complements this existing literature, using a different (and arguably more exogenous) source of variation on a different population.

We find a significant positive effect of child care subsidies at age 5 on children's junior high school academic performance. Being eligible for lower child care prices at age 5 increases the grade point average and the grade on an oral exam by around 0.10 to 0.30 of a standard deviation. Given that take-up of child care is about 55% to 60% for the sample around the discontinuity, this suggests an effect of about 20% to 50% of a standard deviation for those who receive the child care subsidy.

Given this finding, we next investigate the mechanisms through which it is working. A child care price subsidy may have a number of effects on the family. First, it may increase the attendance at formal child care relative to less expensive and often lower-quality informal child care. A lower child care price could also reduce parental care (instead of informal care) and potentially increase parental labor supply. Alternatively, a subsidy could serve as a pure income transfer if demand for day care is inelastic. For any given gross income, families paying a lower price will have more disposable income than families paying the higher price.

While we find large effects on student performance, we find no effect of these substantial child care subsidies on the utilization of formal child care. This is consistent with a situation of excess demand for day care; it is not the price that is important but the availability of a spot.[3] Also, as we describe later, parents are not informed about the income cutoffs that determine the child care price unless their application for a child care place has been successful. As a result, the child care subsidy in Norway appears to have acted as a positive shock to disposable income in the family and, through this mechanism, improved child outcomes. We estimate the effect on disposable income at age 5 to be around 8% of yearly gross income for the families situated around the discontinuity. Given that we find significant effects on later academic performance, this suggests that early investments that increase disposable income may have long-lasting effects.[4] Interestingly, and consistent with a disposable income explanation, we also find effects of the subsidy on the academic performance of older siblings.

Received for publication May 7, 2012. Revision accepted for publication September 6, 2013. Editor: Amitabh Chandra.

* Black: University of Texas, Austin, IZA, and NBER; Devereux: University College Dublin, IZA, and CEPR; Løken: University of Bergen; Salvanes: Norwegian School of Economics, IZA, and CEE.

P.D. thanks the Irish Research Council for financial support. K.L. and K.S. thank the Research Council of Norway for financial support. We thank Jens Ludwig, Michael Baker, and seminar participants from the University of Michigan; the University of Texas, Austin; UCSB; Vanderbilt; Harris School of Public Policy; Norwegian School of Economics; AEA conference in Denver; SOLE conference attendees; and Baylor University for helpful comments.

A supplemental appendix is available online at http://www.mitpressjournals.org/doi/suppl/10.1162/REST_a_00439.

[1] Also see Tekin (2005), Tekin (2007), and Blau and Tekin (2007).

[2] There is also a substantial literature looking at the effects of programs providing universal child care. Herbst and Tekin (2010a) and Magnuson, Ruhm, and Waldfogel (2007) find negative effects of universal child care programs on children's performance, while Berlinski, Galiani, and Manacorda (2009), Berlinski, Galiani, and Gertler (2009), Fitzpatrick (2010), and Havnes and Mogstad (2010, 2011) find positive effects.

[3] Survey results strongly suggest that this was the case in Norway in the 1990s (Blix & Gulbrandsen, 2002).

[4] This relates to a large recent literature that argues that early investment in human capital matters (see Carneiro & Heckman, 2003, and Currie, 2009, for overviews).

The Review of Economics and Statistics, December 2014, 96(5): 824–837

doi:10.1162/REST_a_00439

Given this mechanism, our research also contributes to a growing literature on the effect of family income on child outcomes. The results in this literature are mixed. Using a variety of identification strategies, Oreopoulos, Page, and Stevens (2008), Dahl and Lochner (2012), and Milligan and Stabile (2007) find positive effects of family income on child outcomes, especially for poor families. This is supported by work by Duncan et al. (1998) and Levy and Duncan (2000), who apply family fixed-effects methods. However, Shea (2000) and Løken (2010), using instrumental variables, and Blau (1999) and Dooley and Stewart (2004), using fixed effects, find no or very small effects. Differences could be due to different identification strategies, data sources, countries, and institutional settings.[5]

This paper advances our understanding along two dimensions. First, we are able to convincingly separate income effects from labor force participation; most of the identification strategies used in the existing literature are likely reflecting both family income changes and labor market participation (and, for young children, child care) responses. In our paper, given that we find no effect on labor force participation or child care utilization in the short run, we are able to isolate what appears to be an income effect.[6] Second, given the recent literature suggesting the importance of investments early in a child's life (see Carneiro & Heckman, 2003, and Currie, 2009, for overviews), we are able to analyze the effects of shocks to income, through child care subsidies, when children are age 5, which is likely to be a critical period for human capital investment.

The paper unfolds as follows. Section II gives the institutional background, while section III presents the empirical strategy. Section IV describes the data, and sections V and VI present results and robustness tests. Finally, section VII concludes.

II. Institutional Background

Although the history of day care in Norway goes back a hundred years and the first law regulating day care was in 1953, there was almost no formal child care for children below age 7 (the school starting age, which was changed to 6 in 1997) in Norway until the mid-1970s.[7] However, by the 1990s, the period we study, day care center coverage had risen to 60% among 3- to 5-year-olds and continued to increase throughout the period of study.[8]

There are two types of child care centers in Norway: public (municipality level) and private. In the early 1990s, approximately 60% of the day care centers were public. The private centers were typically owned by nonprofit organizations like churches and cooperatives. However, both types of day care centers are very similar in the way they operate. Around 40% of public day care costs are directly subsidized by the central government; up to one-third is from the municipality, and the rest is paid as fees by the parents. Most of the municipalities also subsidize private day care centers, but the subsidy may be lower than one-third of the cost. Given the stringent national standards for child care, there is likely little variation in quality across private and public centers.[9] For both private and public centers, the municipality pays the difference between a full fee and a reduced fee (the discontinuity we study); the day care centers are just subsidized more in cases with a reduced fee.[10] Means-tested day care subsidies are decided at the municipality level.

There was a tremendous expansion in female labor force participation from the mid-1970s in Norway, creating excess demand for day care and leading to rationing of access to day care centers. The allocation rules determining who got access are not transparent; however, it is clear that children with special needs had priority, along with the children of single mothers (constituting 7% to 8% of children born) (OECD, 2009). Parents submitted a ranking of their preferred day care facilities to a central office in the municipality. This municipality-level institution alone allocated children based on a variety of criteria; however, tenure in line was the most important. This rule was applied to both privately owned and public day care centers (since both the state and the municipality provided subsidies).

Once the child is offered a place in child care, the family is then informed that they can apply for the subsidy if they have family income below a certain level specified in the letter. While the subsidy application is now online, eligible

[5] Dahl and Lochner (2012) argue that fixed effects (FE) estimators do not control for endogenous transitory shocks not directly related to family income and suffer from greater attenuation bias than OLS and instrumental variables (IV), because family income is measured in differences rather than levels. Løken, Mogstad, and Wiswall (2012) argue that differing estimates might be due to the use of linear FE and IV estimators. Theory suggests an increasing concave relationship between family income and child outcomes (Becker & Tomes, 1979), and different instruments might then capture different parts of the income distribution and therefore produce different effects.

[6] Unfortunately we do not have good data on hours of work. Since we are able to rule out effects in participation and changes in use of formal child care, it is unlikely that hours of work change due to the subsidy.

[7] At that time, a new law was passed that aimed at a large expansion in child care as a response to the increasing labor force participation of women. The reform included subsidies to the municipalities that created incentives for municipalities to expand the sector through own establishments or providing subsidies to private nonprofit organizations (see Havnes & Mogstad, 2011). Although this reform increased the coverage, it was still only 32% in 1980 among 3- to 5-year-olds and 7% among 1- to 2-year-olds.

[8] There was also growth in the day care center coverage for 1- to 2-year-olds but at a much lower level—between 15% and 30%.

[9] The Day Care Act (Barnehageloven) gives nationwide standards along several dimensions for day care centers. There are national requirements concerning the education of the staff. For instance, the laws require that the manager and the pedagogical leader both have a college education (three years), very similar to the education requirements for teachers. There are also strict requirements when it comes to playgrounds, playground facilities, and total area within the center. The curriculum is centrally determined, with a strong focus on learning through social relationships both with other children and with adults in the day care centers (OECD, 1999; Framework Plan ("Rammeplanen")).

[10] Although there are centrally described guidelines for staffing requirements, playground requirements, and so on, there is still some room for discretion on the part of the municipalities. For instance, it is the municipality that assesses the quality of the day care facilities. As a result, there may be differences in the quality of day care centers across municipalities. A recent survey found that the share of formally qualified teachers in day care centers varied both across and within municipalities (Gulbrandsen & Winsvold, 2009). However, they did not report any differences across private and public centers.

TABLE 1.—USE OF DAY CARE AND MOTHER'S LABOR SUPPLY, 3- TO 5-YEAR-OLD CHILDREN, 1992 AND 1998

Year	1992	1998
Nannies (%)	13	8
Nannies and day care (%)	64	77
Mothers work full time (%)	32	38
Mother works part time (%)	35	41
Mother work total (%)	66	79

Source: Report from the research institute of NOVA by Gulbrandsen and Winsvold (2009).

families in the 1990s, went to the municipality office to fill out the application and document their income. If their income was below the cutoff, they received the subsidy. The relevant income measure is household income, which includes the income of the mother and, if applicable, her spouse or cohabiter.

The alternative to a formal day care center was the informal sector.[11] This could be play parks, groups run by nannies, or grandparents, other relatives, or friends. None of these informal arrangements received any subsidy from the municipality. They were also not subject to the same regulation by the municipality.[12]

Mainly due to the availability of data, we focus on child care subsidies at age 5.[13] However, given that we focus on age 5, the institutional setting suggests that the price subsidy at age 5 will work as a disposable income effect. Most children in our sample started child care before age 5; based on our own calculations, we find that 86% of those who attended child care at age 5 also attended formal child care at age 4. In addition, given the situation of excess demand, child care decisions were likely determined prior to the granting of the subsidy.

Table 1 shows information from a survey on the use of registered nannies and formal day care centers in the 1990s, in addition to labor supply of mothers.[14] We see that the labor supply of mothers with 3- to 5-year-old children matches very well the total use of formal care—either registered nannies or day care.

III. Empirical Strategy

The day care system in Norway is run at the municipality level (there are 435 municipalities), and the cost is heavily subsidized for all. Parents pay about 30% of the actual costs on average. Some municipalities have a single price that is the same for all income groups, while others have multiple prices that depend on family income. In these municipalities, pricing takes the form of a step function, with jumps in the price occurring at one or more levels of family income. These jumps suggest discontinuities in the relationship between family income and the price of child care. Assuming that other factors related to family income that affect child outcomes do not systematically change at the discontinuity points, we can identify child care subsidy effects by comparing later outcomes of children whose family income was just less than a cutoff to those of children whose family income was just above a cutoff. In this paper, we focus on the first income cutoff in each municipality because in municipalities with more than one income cutoff, the price differences at higher cutoffs are typically small.

We use a regression discontinuity (RD) approach to estimate the effect of eligibility for lower child care prices. We have a sharp design since eligibility for cheaper child care jumps from 0 to 1 at the discontinuity. However, the take-up rates of child care and the subsidy are below 100%, and we have to take this into account when interpreting the estimates. For family i, in municipality m, at time t, the eligibility for lower child care price ($E_{i,m,t}$) is a deterministic function of family income the year before ($FI_{i,m,t-1}$); if income was below a particular cutoff ($c_{m,t}$), the family received the extra subsidy and thereby paid a lower price. We can then estimate the effect of being eligible for a lower child care price on child outcomes (y) by comparing families with incomes just below and above $c_{m,t}$.

Because the level of the cutoff varies by municipality and year, in our analysis, we normalize family income by dividing it by the relevant cutoff income level in the municipality and subtracting 1:

$$I_{i,m,t-1} = (FI_{i,m,t-1}/c_{m,t}) - 1.$$

By construction, normalized family income ($I_{i,m,t-1}$) equals 0 at the cutoff and takes on positive (negative) values above (below) the cutoff.[15]

For identification, we need to assume that income and other characteristics about the family vary continuously through the cutoff point; we verify this by comparing characteristics on either side of the cutoff. We then estimate the effect of the child care subsidy by taking the difference of the boundary points of two regression functions of y on I, one for eligible families and one for ineligible families.

We estimate different versions of the following equation:

$$\begin{aligned} y_{i,m,t} = \beta_0 + \beta_1 E_{i,m,t} + \beta_2 f_l\big(I_{i,m,t-1}\big) \\ + \beta_3 f_r\big(I_{i,m,t-1}\big) + \beta_4 x_i + \lambda_{mt} + \varepsilon_{i,m,t}, \end{aligned} \tag{1}$$

[11] It was not until 2008 that Norway, through a change in the law, required municipalities to have full formal child care coverage. A law from 1998 (the so-called cash-for-care reform) gave parents the right to the state subsidy if they opted out of day care and stayed home with the child instead (see Schøne, 2004).

[12] This was true for registered nannies (who paid income taxes) as well.

[13] We have more observations for these cohorts as our data on income cutoffs and prices start in 1991, and the last cohort with observations on educational outcomes is 1992, giving us, for example, only three cohorts of 1-year-olds (1990–1992), while we have seven cohorts of 5-year-olds (1986–1992). As we rely on a regression discontinuity (RD) design for identification, we need a large sample size to get enough observations around the discontinuity. So while we would ideally like to study the effect of total child care use during childhood, we can study child care use only at age 5.

[14] See the report from the research institute of NOVA by Gulbrandsen and Winsvold (2009).

[15] We have also tried normalizing income by subtracting the cutoff level of income in the municipality, and this leads to similar results.

where i denotes individual, m denotes municipality, and t denotes year. Here $E_{i,m,t} = 1\{I_{i,m,t-1} < 0\}$, an indicator for whether family income is below the cutoff; x is a vector of individual and family control variables, and λ is a vector of cohort by municipality fixed effects. We want to estimate β_1, the effect of being eligible for lower child care prices on children's outcomes. While the presence or absence of controls should not have much impact on estimates in a regression discontinuity design, we control for pre–child care parental characteristics in order to increase the precision of our estimates.[16]

Because all our outcomes will typically vary with family income and eligibility for cheaper child care depends on income, we have to control for family income on each side of the discontinuity in a flexible way. The functions f_l and f_r allow the effect of income to differ between the left and right sides of the cutoff. We estimate equation (1) for several bandwidths and estimate specifications using linear, quadratic, cubic, and quartic functions of income. When we use linear functions of income, our method is equivalent to local linear regression (LLR), as in Fan (1992), Hahn, Todd, and Van der Klaauw (2001), and Porter (2003), using a rectangular kernel.[17] With this approach, control variables can naturally be added to the specification.

We present our results both graphically and in tables. The figures illustrate the local linear specification with rectangular kernel and bandwidth of .1. We also show the 95% confidence intervals and scatter plots with average outcomes for twenty income bins. Note that this is only to illustrate the pattern in the data; the estimation uses all the observations to estimate the discontinuity. In the tables we also present results with bandwidths ranging from .02 to .5 and income polynomials varying from 0 (no income controls) to 4 (a quartic in income) and we allow these functions to vary on each side of the discontinuity.[18]

IV. Data

We use administrative data covering the entire population of Norway. The analysis includes birth cohorts from 1986 to 1992 and links individuals to their parents through unique identifiers. We have information on parental characteristics such as parental age, educational attainment when the child was born, income, marital status, and citizenship. For children, we have grade point average and exam grades from junior high school. In addition, we match parents to their tax records, where we are able to observe whether parents take deductions for child care expenses; this allows us to identify whether a child attends formal child care. Finally, we have collected data from municipalities in Norway on child care prices and family income cutoffs in the 1990s.

Family income is created by adding mother's earnings to those of any spouse or cohabiter. Earnings are measured as total pension-qualifying earnings reported in the tax registry, starting from 1967. The earnings measure includes labor earnings and all taxable welfare benefits, including sick benefits, unemployment benefits, and parental leave payments. This is the same income measure that municipalities use to determine whether families are eligible for cheaper child care.[19] Our measure of disposable income is defined as family income minus the child care price faced by the family.

Our measure of child care attendance is created from the information on tax deductions for child care expenses from parents' tax records; these are available from 1993.[20] The child care tax deduction was introduced in 1948; parents are allowed to deduct up to NOK 25,000 (USD $4,310) from taxes in one calendar year for the first child for formal child care that takes place outside the home.[21] As a result, our definition of child care excludes home care and informal care by grandparents and nannies. There is a significant amount of variation in tax deductions across families due to both different costs across municipalities and different prices across income groups within municipalities. Our measure of child care is an indicator for whether or not the child is attending formal child care. The online data appendix and appendix table 1 contain details of exactly how child care use is inferred from the tax data.

Finally, we have collected data at the municipality level on the price system and actual prices of child care in the 1990s. If the municipality had variable prices across the income distribution, we asked explicitly for the income cutoffs that the municipality used to determine eligibility for cheaper child care. We received information from 69% of the municipalities, including the ten largest. This gives us information on the price system for about 85% of the total sample. Figure 1 provides a map of Norwegian municipalities showing that variable, flat, and unknown price municipalities are scattered across Norway.

[16] Control variables are parental age, parental citizenship, parental education when child is born, marital status when child is born, student and welfare recipient status of mother when child is age 4, and family income prior to age 4 (measured as the average income when the child was aged 1 to 3).

[17] Lee and Lemieux (2009) recommend doing local linear regression using one kernel and focus on estimating the model with different bandwidths. We have also tried local linear regressions with different kernels, without any significant changes to the main results.

[18] We calculate heteroskedasticity-robust standard errors (White, 1980) following suggestions by Lee and Lemieux (2009).

[19] Almost all municipalities use the previous year's tax sheet to determine subsidy eligibility. As a robustness check, we have run the regressions for municipalities where we are certain that they used last year's tax sheet and find similar results (results available on request). For identification, it is also an advantage to use the previous year's income because parents cannot manipulate the previous year's income based on today's cutoff.

[20] This means that we do not have information on child care attendance for our first two cohorts born in 1986 and 1987.

[21] The extra deduction for the second child is NOK 5,000, for a total of NOK 30,000. The annual price of child care is almost always below NOK 25,000 per child, at least for the families we study around the discontinuity.

FIGURE 1.—MAP OF NORWAY WITH INFORMATION ON PRICE SYSTEMS ACROSS MUNICIPALITIES

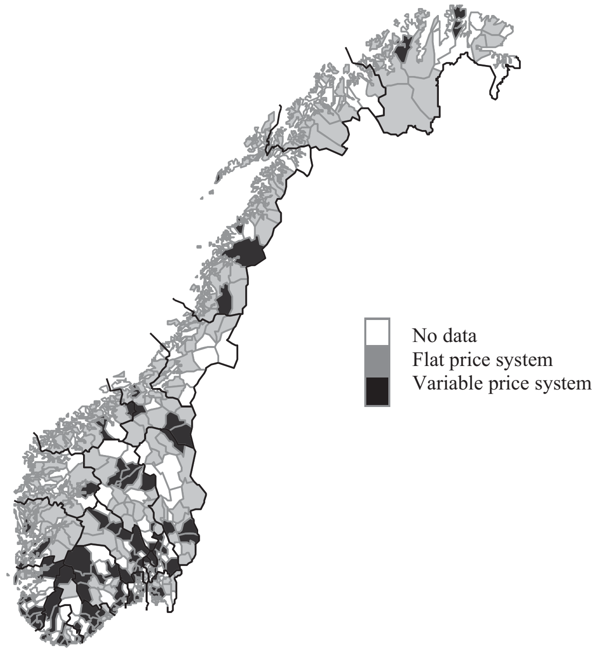

There are some missing observations in the data. We exclude the observations where parental background characteristics are missing. This reduces our sample from 448,198 observations to 367,836. We have tested that the main results on child outcomes are not sensitive to excluding these observations.

Our analysis is conducted on families located around the first price discontinuity in each municipality. We include families with income no more than 50% below or above the discontinuity—that is, with normalized income between -0.5 and $+0.5$.[22] While our maximum bandwidth is .5, we report estimates for a range of smaller bandwidths, with the smallest being .02.[23]

For children, we have information on performance on junior high school national exams and their grade point average (GPA) in tenth grade.[24] The grade point average is an average of the tenth grader's performance in all twelve graduating subjects.[25] The exam data are the grades from written and oral exams that are administered in the final year of junior high school at the national level and are externally graded. The written exam could be in math, Norwegian, or English, and the subject is randomly assigned at the school level.[26] The students are informed about which exams they will take a couple of days before the exam date. The oral exam can be in any of the twelve subjects taught in the last year of middle school, and the students are randomly allocated to subjects. These grades are important for high school admission. The grades range from 1 to 6.[27]

Table 2, column 1 gives descriptive statistics for the total sample of children born between 1986 and 1992. Column 2 provides descriptive statistics for our analysis sample (those within .5 of the cutoff), and column 3 describes those who are within our preferred bandwidth of .1. We see that the samples are very similar in terms of child characteristics such as age, gender, number of siblings, and birth order. However, in the .5 and .1 samples, parents tend to have fewer years of education and are more likely to be of non-Norwegian citizenship, highlighting the fact that these samples are composed of individuals at the bottom end of the income distribution. About 80% of the total sample is married or officially cohabiting in the year of their child's birth; this number is 72% for both the .5 and .1 samples. When comparing school performance of the total sample with our analysis samples, we see that children in the analysis samples tend to perform worse, with a mean GPA of 4 and 3.7, respectively.[28] This is not surprising as we know that children from low-income families tend to perform worse in school.

V. Results

A. Testing for Income Manipulation

Based on communications with the municipalities, we have verified that the cutoffs for the child care subsidy are not also used to determine eligibility for other social programs. Therefore, we are not worried that our discontinuity design could be picking up the effects of other welfare programs in addition to child care. The major remaining potential threat to the design is the possibility that families manipulate their income so as to locate strategically around

[22] For example, if the price discontinuity is NOK 100,000, we include families with income between NOK 50,000 and NOK 150,000. The results are robust to including more families; however, as the discontinuity is at low levels of income, we cannot move much further to the left of the discontinuity as zero income is binding from below.

[23] We have experimented with excluding families where mothers are students when the child is age 4. This is because students might have different rules for child care and might not be affected by the subsidies. Excluding students (about 5% of sample) does not change the results. We have also experimented with excluding mothers who receive social security benefits, and this also does not change any of the main results.

[24] In Norway junior high school children are ages 13 to 16.

[25] These consist of written and oral Norwegian, written and oral English, mathematics, nature and science, social science, religion, home economics, physical education, music, and handcrafts.

[26] Because we control for cohort by municipality fixed effects, differences in grading standards across cohorts or municipalities will not lead to bias. In this exercise, we control for municipality by cohort fixed effects.

[27] There are advantages and disadvantages with different measures. The exam grades are more variable because they are a one-time measure of skills. However, they are more comparable across cohorts and schools because they are graded externally. The grade point average is generally a better measure of long-term skills as it covers all subjects and averages over all grades; however, it is also more subjective because it depends on teacher assessments. However, our identification strategy compares similar families just below and above the income cutoff who will, on average, have the same schooling environment, so all three measures of academic performance should be valid.

[28] A more detailed distribution of grades for the grade point average and the written and oral exams is shown in appendix table 2.

TABLE 2.—DESCRIPTIVE STATISTICS

Samples	Total Sample	0.5 Sample	0.1 Sample
Age in 2006	18.9	19.1	19.1
	(2.0)	(2.0)	(2.0)
	[367,836]	[10,770]	[1,946]
Female	.49	.49	.48
	(.50)	(.50)	(.50)
	[367,836]	[10,770]	[1,946]
Number of siblings	1.8	1.9	2.0
	(1.1)	(1.4)	(1.4)
	[367,836]	[10,770]	[1,946]
Birth order	1.9	1.9	1.8
	(1.0)	(1.1)	(1.1)
	[367,836]	[10,770]	[1,946]
Mother's education at birth of child	11.9	10.2	10.1
	(3.4)	(4.3)	(4.3)
	[367,836]	[10,770]	[1,946]
Father's education at birth of child	11.8	9.8	9.8
	(3.5)	(4.4)	(4.4)
	[367,836]	[10,770]	[1,946]
Mother's age at birth of child	27.9	26.1	26.1
	(5.0)	(5.4)	(5.5)
	[367,836]	[10,770]	[1,946]
Father's age at birth of child	30.7	29.3	29.3
	(5.6)	(6.4)	(6.4)
	[367,836]	[10,770]	[1,946]
Mother non-Norwegian citizen at birth of child	.05	.17	.17
	(.21)	(.38)	(.38)
	[367,836]	[10,770]	[1,946]
Father non-Norwegian citizen at birth of child	.05	.18	.19
	(.21)	(.39)	(.39)
	[367,836]	[10,770]	[1,946]
Married/cohabiting at birth of child	.8	.72	.72
	(.40)	(.45)	(.45)
	[367,836]	[10,770]	[1,946]
Grade point average (scale: 1–6)	4.0	3.7	3.7
	(.82)	(.85)	(.86)
	[359,339]	[10,238]	[1,835]
Grade written exam (scale: 1–6)	3.5	3.1	3.1
	(1.1)	(1.1)	(1.1)
	[344,271]	[9,572]	[1,708]
Grade oral exam (scale: 1–6)	4.3	3.9	4.0
	(1.2)	(1.2)	(1.2)
	[318,783]	[8,823]	[1,578]

the income cutoffs (Lee & Lemieux, 2009).[29] There are several reasons that we think this is not a factor in this case.

The first reason we think it is unlikely that individuals have strategically manipulated their income is that the cutoff is based on the previous year's income but is unknown the year before the child care subsidy is allocated; as a result, families cannot perfectly predict where the cutoff will be and adjust their income accordingly. However, a concern is that in some cases, the cutoff may be predictable. For example, of the 72 municipalities, 30 did not change cutoffs in our sample period, so the cutoff is easy to predict. The remaining 42 municipalities do change cutoffs, and we have verified that none do so in an easily predictable fashion. For example, none of them increases the cutoff by the same percentage each year. Later (in table 6), we show that our estimates are robust to restricting the sample to only municipalities with large changes in cutoffs over time.

A second reason to think there is no income manipulation is the smooth distribution of income around the cutoff. If there were income manipulation, it should show itself in a spike in the income distribution just below the cutoff. In appendix figure 1, we show the density of normalized family income for the analysis sample. We see that there is no evidence of income clustering below the cutoff.

Another concern would be that although there is no evidence of a discontinuity in income around the cutoff, there could be unobserved differences between people on opposite sides of the discontinuity. This implication is inherently untestable, but we can examine observed family characteristics. We compare presubsidy characteristics for families on opposite sides of the discontinuity to verify that observable characteristics do not change at the discontinuity. These results are presented in appendix table 3. We show esti-

[29] An alternative possibility is that even without manipulation, families below the threshold might just happen by chance to have better unobserved characteristics. The balancing tests we describe below suggest that this is unlikely to be an issue in our case.

TABLE 3.—EFFECT OF BEING BELOW THE INCOME CUTOFF ON THE GRADE POINT AVERAGE

Bandwidth/ Polynomial of Order	0.02	0.05	.10	.25	.5
0	.209	.091	.042	.043	.019
	(.156)	(.082)	(.054)	(.031)	(.022)
	[.994]	[.254]	[.330]	[.728]	[.898]
1	.207	.499*	.266*	.114	.066
	(.329)	(.167)	(.109)	(.062)	(.042)
	[.994]	[.787]	[.590]	[.708]	[.924]
2	−.025	.306	.390*	.211*	.070
	(.436)	(.235)	(.162)	(.094)	(.064)
	[.982]	[.782]	[.680]	[.863]	[.898]
3	.151	.400	.439*	.328*	.159
	(.565)	(.307)	(.214)	(.126)	(.085)
	[.554]	[.665]	[.523]	[.910]	[.950]
4	−.226	.533	.454	.513*	.289*
	(.762)	(.369)	(.258)	(.158)	(.107)
	[.553]	[.369]	[.504]	[.973]	[.984]
Optimal order of polynomials	0	1	1	2	4
N	354	912	1,835	4,727	10,211

* $p < 0.05$. Standard errors in parentheses. *P*-values from the goodness-of-fit test in brackets. The goodness-of-fit test is obtained by jointly testing the significance of a set of bin dummies included as additional regressors in the model. The bin width used to construct the bin dummies is .01. The optimal order of the polynomial is chosen using Akaike's criterion (penalized cross-validation). This table reports the coefficients from an RD regression with income polynomials on each side of the discontinuity. Control variables are cohort × municipality fixed effects, parents' age, education, citizenship, marital status at birth of child, pre–child care family income, and mother's welfare and student status when the child was age 4.

mates for three bandwidths (.1, .25, and .5), and for both the linear and quadratic models.[30] This table shows the results for balancing tests on parents' educational attainment, age at birth, citizenship and marital status in the birth year of the child, mother's welfare status when child is age 4, and average family income when child is age 0 to 3. The estimates are all statistically insignificant at the 5% level apart from three exceptions: one specification in which mother's education is significantly higher to the left of the discontinuity, one specification in which parents are less likely to be married, and one specification in which the mother is more likely to have been on welfare. Overall, there is little compelling evidence here for differences in observables around the discontinuity. Appendix figure 2 shows this graphically for several of these outcomes using the linear model with a bandwidth of .1 along with the associated 95% confidence intervals.[31]

In the last three rows of appendix table 3, we explore this issue further by considering the predicted student performance based on observable student and parent characteristics and comparing this measure for students on either side of the discontinuity. To do so, we regress each of our child outcome variables (GPA, written exam score, and oral exam score) on the parental background variables and generate a predicted GPA, predicted written exam, and predicted oral exam for each child. These new variables represent what we would predict the child outcomes to be based solely on their family background. We carry out the balancing tests with these as dependent variables. Reassuringly, the estimates are very small in all specifications and precisely enough estimated to rule out large effects. This provides support for the assumption that children's outcomes are most likely being affected only by the child care subsidy itself, and not other differences.

Finally, by using the regression discontinuity approach, we are implicitly assuming that the assignment of subsidy receipt is essentially random, conditional on observables. As a result, it should be the case that the probability of subsidy receipt prior to age 5 is equal for both the treatment and control groups. Appendix table 4 presents estimates of the probability of being below the cutoff when the child was aged 1 to 4. Importantly, we see no effect of current subsidy receipt on the probability of being below the cutoff at ages prior to the subsidy (one coefficient out of 24 is statistically significant at the 5% level). This further supports our assumption that there is no strategic income manipulation.

B. Children's Outcomes

Tables 3 to 5 present the effect of child care subsidies on children's GPA and exam grades in junior high school. When analyzing these outcomes, we standardize them to have mean 0 and standard deviation of 1 so that the coefficients are comparable across regressions. We show estimates for five bandwidths (.02, .05, .1, .25, and .5) and five different specifications (zero to four polynomials). We also indicate the optimal order of the polynomial as determined by the Akaike criterion (penalized cross-validation). Not surprisingly, the standard errors are very high for the .02 and .05 bandwidths, so we do not emphasize these results. Instead, we focus on estimates from the .1, .25, and .5 bandwidths as they tend to provide reasonably precise estimates.[32]

[30] Because of space constraints, we exclude estimates for 0, 3, and 4 polynomials and for the very small bandwidths. These estimates also show no evidence of better family characteristics on the left-hand side of the discontinuity.

[31] The values of the *y*-axis are created by always including ±1 standard deviation around the mean outcome in order to make the graphs comparable.

[32] When we use the Imbens and Kalyanaraman (2009) method to determine the optimal bandwidth, the suggested bandwidth is between .2 and .3 depending on the outcomes. However, in general, lower bandwidths imply lower bias, so we also put weight on estimates using the bandwidth of .1.

Figure 2.—Effect of Child Care Subsidy on Children's Junior High Academic Performance

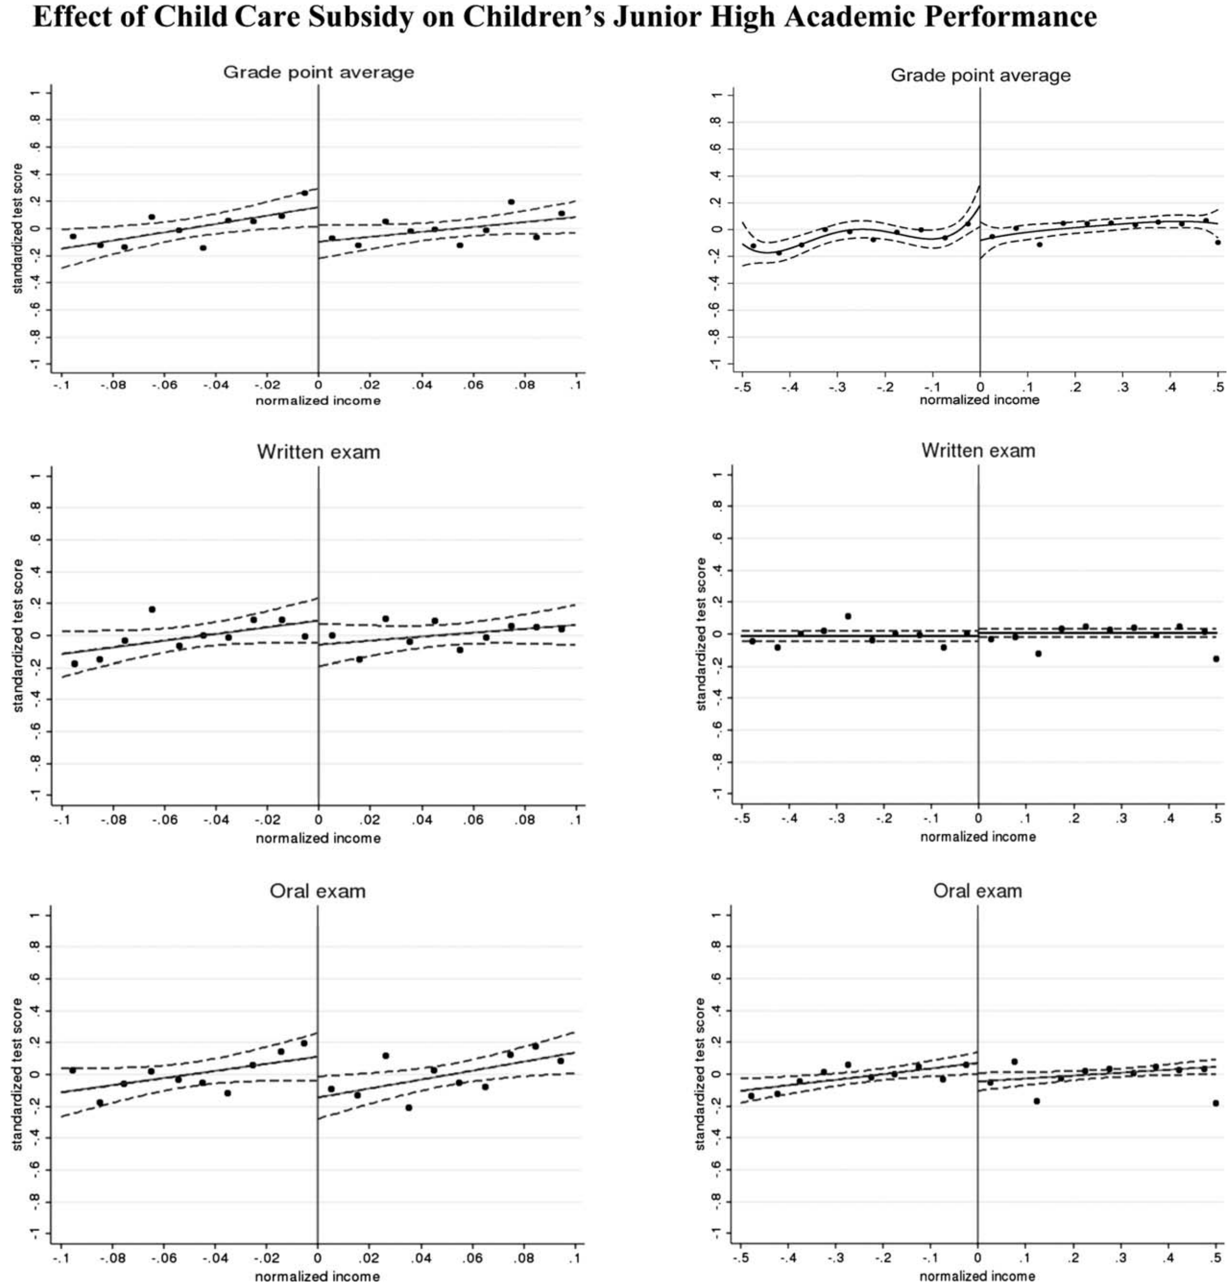

(Left) The solid line is the local linear regression with a rectangular kernel and bandwidth .1. (Right) The solid line is the optimal polynomial regression from tables 3 to 5 with a window of .5 (quartic for GPA, no polynomial for written exam, and linear for oral exam). The dashed lines are 95% confidence intervals. The scatter plot is the average standardized outcome for twenty income bins.

In table 3, we see a statistically significant positive effect of subsidy receipt on children's GPA that is generally between .1 and .5 of a standard deviation. The range of estimate sizes across bandwidths and specification implies that we cannot be certain about the exact magnitude. Because not everyone is in child care and hence affected by the child care subsidy, this is an intention-to-treat effect. Given that take-up of child care for our analysis sample is about 55% to 60%, this means that the effect for the treated is even larger. If we conservatively take an estimate of .1 to .3 for the intention to treat, this implies an effect on the treated of about .17 to .50 of a standard deviation.

Table 4 has estimates for the written exam. Here there is not much evidence for a positive effect as all estimates are statistically insignificant except for the largest bandwidth of .5. Even for this bandwidth, the significant effect is small at .05, implying an effect on the treated of about .1 of a standard deviation. Results for the oral exam are in table 5. Once again, the estimates are generally positive but statistically significant in only a few specifications. The statistically significant estimates using the optimal number of

Table 4.—Effect of Being below the Income Cutoff on the Written Exam Score

Bandwidth/ Polynomial of order	0.02	0.05	.10	.25	.5
0	.078	.006	.013	.041	.053**
	(.149)	(.085)	(.054)	(.032)	(.022)
	[.617]	[.627]	[.637]	[.910]	[.683]
1	−.105	.224	.165	.067	.023
	(.298)	(.170)	(.110)	(.063)	(.043)
	[.484]	[.244]	[.452]	[.860]	[.672]
2	−.412	−.052	.140	.084	.037
	(.411)	(.233)	(.162)	(.095)	(.065)
	[.438]	[.480]	[.416]	[.750]	[.604]
3	.238	.018	.232	.129	.064
	(.501)	(.306)	(.215)	(.127)	(.087)
	[.916]	[.593]	[.458]	[.752]	[.635]
4	.058	.181	.087	.237	.202*
	(.707)	(.348)	(.258)	(.157)	(.108)
	[.970]	[.388]	[.601]	[.910]	[.647]
Optimal order of polynomials	0	1	1	1	0
N	334	852	1,708	4,401	9,546

* $p < 0.05$. Standard errors in parentheses. *P*-values from the goodness-of-fit test in brackets. The goodness-of-fit test is obtained by jointly testing the significance of a set of bin dummies included as additional regressors in the model. The bin width used to construct the bin dummies is .01. The optimal order of the polynomial is chosen using Akaike's criterion (penalized cross-validation). This table reports the coefficients from an RD regression with income polynomials on each side of the discontinuity. Control variables are cohort × municipality fixed effects, parents' age, education, citizenship, marital status at birth of child, pre–child care family income, and mother's welfare, and student status when the child was age 4.

Table 5.—Effect of Being below the Income Cutoff on the Oral Exam Score

Bandwidth/ Polynomial of Order	0.02	0.05	.10	.25	.5
0	.132	.117	.029	.089*	.063*
	(.162)	(.091)	(.058)	(.034)	(.024)
	[.929]	[.225]	[.257]	[.364]	[.080]
1	.021	.377*	.194	.091	.113*
	(.314)	(.175)	(.117)	(.067)	(.046)
	[.231]	[.356]	[.594]	[.424]	[.279]
2	.108	.329	.259	.073	.062
	(.449)	(.245)	(.171)	(.102)	(.069)
	[.217]	[.333]	[.653]	[.380]	[.304]
3	.612	.431	.327	.199	.117
	(.586)	(.314)	(.225)	(.135)	(.092)
	[.735]	[.593]	[.521]	[.480]	[.345]
4	−.144	.666	.462	.421*	.161
	(.758)	(.358)	(.269)	(.165)	(.116)
	[.736]	[.619]	[.601]	[.619]	[.343]
Optimal order of polynomials	3	3	1	0	1
N	312	793	1,578	4,401	9,546

* $p < 0.05$. Standard errors in parentheses. *P*-values from the goodness-of-fit test in brackets. The goodness-of-fit test is obtained by jointly testing the significance of a set of bin dummies included as additional regressors in the model. The bin width used to construct the bin dummies is .01. The optimal order of the polynomial is chosen using Akaike's criterion (penalized cross-validation). This table reports the coefficients from an RD regression with income polynomials on each side of the discontinuity. Control variables are cohort × municipality fixed effects, parents' age, education, citizenship, marital status at birth of child, pre–child care family income, and mother's welfare, and student status when the child was age 4.

polynomials for the .25. and .5 bandwidths suggest a magnitude size of about .1 of a standard deviation.

Figure 2 presents the results for the standardized scores graphically. For each outcome, we show the graph for the linear model with a bandwidth of .1 and a graph for the optimal number of polynomials (quartic for GPA, no income control for written, and linear for the oral exam) using the largest bandwidth of .5. For the bandwidth of .1, the linear model appears to fit quite well. This is unsurprising as the optimal polynomial at this bandwidth is linear for both GPA and the oral exam. With the larger bandwidth of .5, there is more curvature in the GPA data, and the optimal polynomial is quartic. The right-hand panel of figure 2 suggests that the quartic fits the data well. Interestingly, even with this larger bandwidth, the optimal polynomials for the written and oral exams are 1 and 0, respectively. The linearity of the data for these outcomes is apparent in figure 2.

C. *Robustness Checks*

To verify our findings, we run a number of robustness checks. A key concern is whether individuals can manipulate their income to be on the desirable side of the cutoff. As noted before, this is unlikely, given that we see no bunching at the favorable side of the income cutoff and there is little evidence of observable differences in characteristics on opposite sides of the discontinuity. As a further check, we divide our sample into municipalities with no changes or small changes in cutoffs over time and those with larger changes. Presumably if individuals are manipu-

TABLE 6.—ROBUSTNESS TESTS: EFFECT OF BEING BELOW THE INCOME CUTOFF ON CHILDREN'S JUNIOR HIGH ACADEMIC PERFORMANCE

	No or Small Changes in Cutoffs, 1993–1997	Big Changes in Cutoffs, 1993–1997	Municipalities with Large Price Jumps	Municipalities with Smaller Price Jumps	Municipalities with Cutoff at Low Levels of Family Income	Municipalities with Cutoff at Higher Levels of Family Income
Grade point average	.147	.361	.405*	−.034	.336*	.207
	(.167)	(.189)	(.131)	(.200)	(.149)	(.161)
	[762]	[774]	[1,288]	[547]	[909]	[926]
Written exam	.130	.152	.260*	−.047	.239	.119
	(.169)	(.189)	(.132)	(.199)	(.151)	(.161)
	[706]	[723]	[1,205]	[503]	[858]	[850]
Oral exam	.177	.177	.248	.099	.275	.123
	(.182)	(.196)	(.140)	(.215)	(.165)	(.168)
	[668]	[658]	[1,095]	[483]	[774]	[804]

* $p < 0.05$. This table reports the coefficients from an RD regression with linear trends on each side of the discontinuity and a window/bandwidth of .1. Control variables are cohort × municipality fixed effects, parents' age, education, citizenship, marital status at birth of child, pre–child care family income, and mother's welfare and student status when the child was aged 4. For the first subgroup on changes in cutoffs, small changes in cutoffs 1993–1997 are defined as no more than a 5% increase in cutoffs in the period from 1993 to 1997, while big changes are more than 5%. The median change in cutoffs from 1993 to 1997 is 5%.

lating their income, the effects should be driven by the more predictable, no- or small-change municipalities.[33] The first two columns of table 6 show that this is not the case.[34]

As a further check, it should be the case that municipalities with smaller jumps in prices at the cutoff should also experience smaller changes in children's performance. As a test, we split municipalities into those with small jumps in price (where there is little effect on disposable income) and those with larger jumps in price. Table 6, columns 3 and 4, presents the results, and it is clear that the main effects come from municipalities with the largest price cuts (also see appendix figure 3 for the pictures for GPA). We find essentially no effect for municipalities with small price jumps and large, significant effects for municipalities with larger price jumps.

Finally, when it comes to municipalities with different cutoffs across family income, we might expect the effect to be larger for municipalities where the cutoff is at lower levels of family income, since this is consistent with the literature finding larger effects for the most disadvantaged families. However, we see from table 6, columns 5 and 6, that the effects are very similar across both types of municipalities. This may be because all the cutoffs are at reasonably low levels of income. [35]

We also take advantage of the fact that some municipalities have no variation in the price of child care; municipalities with a flat price system do not give us variation across family income to identify an effect of differences in child care prices across income. However, as a placebo test, we assign the flat price municipalities the average cutoff of the variable price municipalities to check whether there are any systematic differences across child outcomes for our cohorts that are unrelated to the price discontinuity. (We should observe no effect of this "placebo" discontinuity on any outcomes.) Appendix table 5 presents these results; it is reassuring to note that there is no effect on children's grades or test scores.

Finally, in appendix table 6, we present results when we move the cutoff ± 5% from the true cutoff and estimate the effects using these placebo cutoffs; again, it is reassuring to see no effects on any of the outcomes whether we use ± 5%.

D. *Mechanisms*

Given the observed effect of the child care subsidy on children's junior high school performance, the next question becomes what factors are driving these effects. The first part of table 7 shows the effects of the child care subsidy on various intermediate outcomes, and figure 3 plots the variables for the linear model with bandwidth of .1. Importantly, from the first row in table 7, we see no evidence of any effect of the child care subsidy on child care utilization.[36]

Despite this, it is clear that the price jumps, and hence disposable income falls, significantly at the discontinuity. From the next row in table 7 we see that families below the discontinuity pay, on average, almost NOK 10,000 (USD $1,700) less for child care per year. Given that the average value of the cutoff in our analysis samples is NOK 125,000, this implies that families just below the cutoff have just over 8% more yearly disposable income when the child is in child care at age 5.

[33] We classify municipalities as large-change municipalities if the average change in the cutoff in the municipality over our period exceeds the median across municipalities. As such, this group excludes all municipalities that never change cutoffs and also excludes some municipalities that make small changes.

[34] Due to space constraints, table 6 shows only estimates using the linear model with a bandwidth of .1. The findings are broadly similar using other bandwidths and specifications.

[35] To understand more about which children are most affected by the child care subsidy, we have split the sample into subgroups based on pre–child care characteristics (mothers having ten years of education or fewer when the child was born compared to more than ten years, parents who were married or cohabiting when the child was born compared to not married or not cohabiting, and females compared to males and mothers with and without Norwegian citizenship status when the child was born). There is very little evidence of large differences by subgroup.

[36] We do not have information on child care deductions at age 5 for all cohorts, so we have fewer observations for this variable than for others. For the cohorts for which we have child care information, we have tried estimating the effect of the subsidy on academic outcomes by child care status. Unfortunately, the estimates are not informative, with high standard errors for both groups and no significant differences.

FIGURE 3.—MECHANISMS

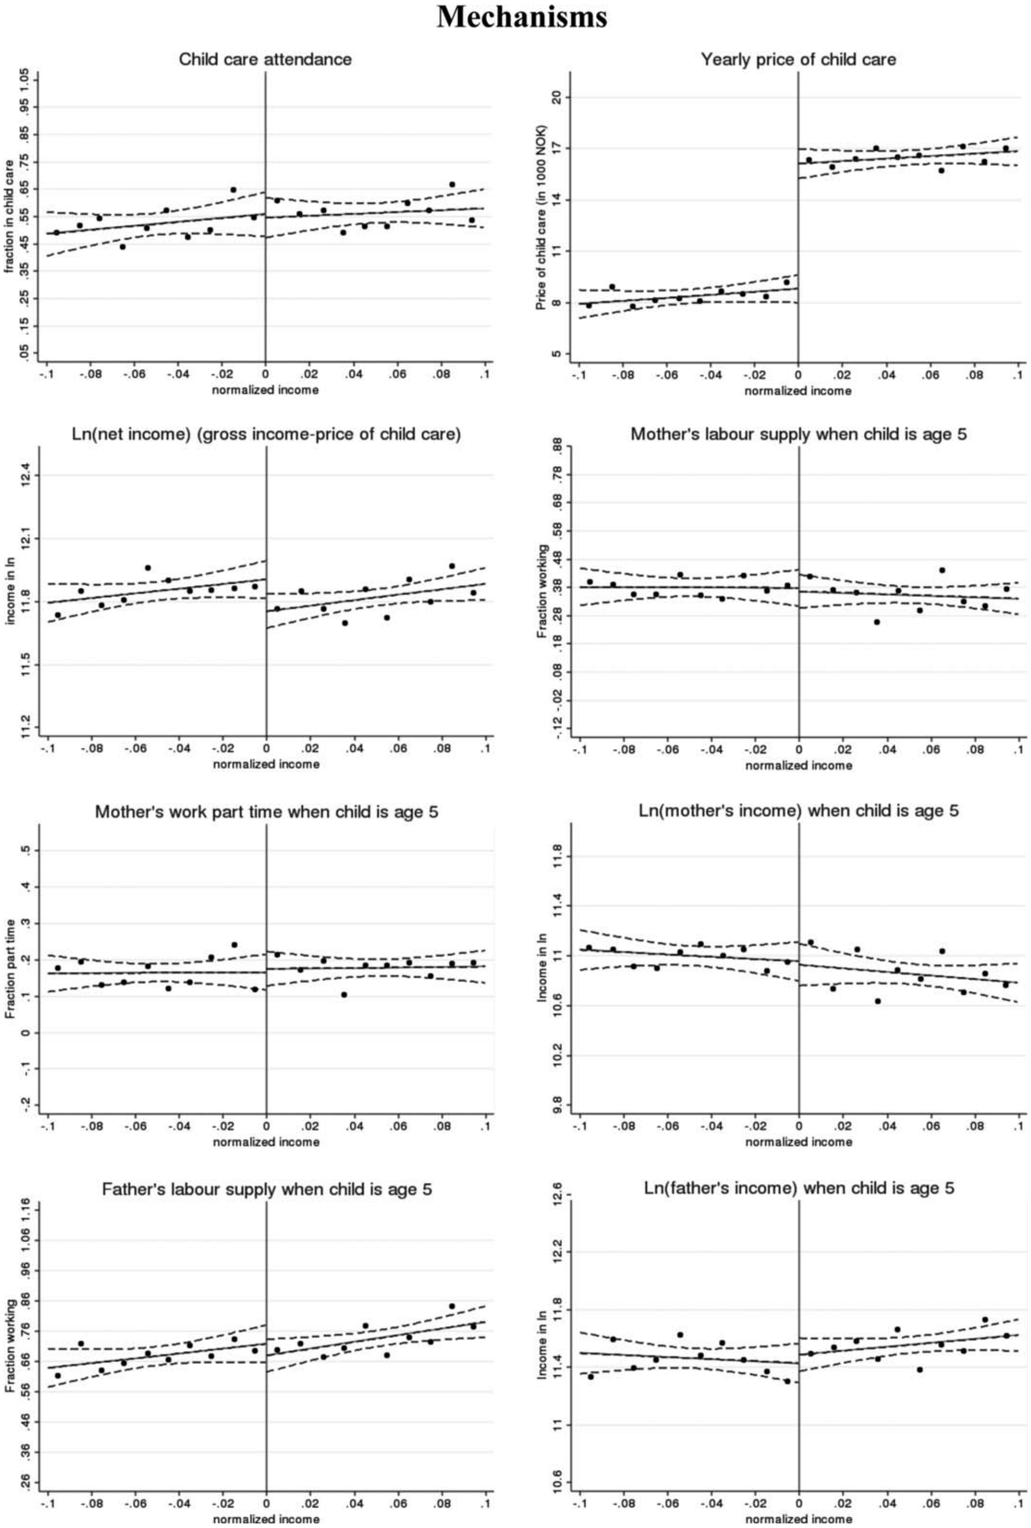

The solid line is the local linear regression with a rectangular kernel and bandwidth .1. The dashed lines are 95% confidence intervals. The scatter plot is the average standardized outcome for twenty income bins.

We next study whether the subsidy affects parental labor supply and income. We see little evidence of effects of the subsidy on mother's or father's labor supply (although the coefficient is positive and statistically significant for women in one specification). This suggests that there are no responses by the parents on the extensive margin and is consistent with no effects on child care utilization; the subsidy does not appear to affect time allocation between the

TABLE 7.—MECHANISMS: EFFECT OF BEING BELOW THE INCOME CUTOFF ON VARIOUS OUTCOMES

Polynomials	1	1	1	2	2	2
Window/Bandwidth	.1	.25	.5	.1	.25	.5
Child care attendance	−.026	−.015	−.018	−.032	−.039	−.026
	(.061)	(.036)	(.024)	(.090)	(.054)	(.036)
Price of child care (in NOK)	−9,328*	−9,596*	−9,265*	−8,840*	−9,307*	−9,426*
	(374)	(194)	(131)	(613)	(313)	(208)
Mother's labor supply	.013	.043	.059*	−.031	.033	.043
	(.048)	(.028)	(.019)	(.072)	(.042)	(.329)
Mother work part time	−.014	−.008	−.007	.005	−.017	−.006
	(.040)	(.023)	(.015)	(.062)	(.035)	(.023)
Ln(Mother's income)	−.040	.102	.087	−.201	.154	.095
	(.129)	(.075)	(.050)	(.196)	(.115)	(.075)
Father's labor supply	.026	−.010	−.006	.012	−.013	.025
	(.045)	(.026)	(.018)	(.067)	(.039)	(.027)
Ln(Father's income)	.002	.032	−.000	−.189	−.025	.076
	(.099)	(.059)	(.040)	(.152)	(.089)	(.061)
Ln(Annuity of family income)	.183*	.092*	.042*	.124	.138*	.076*
Age 6–15	(.050)	(.030)	(.021)	(.073)	(.043)	(.030)
N	1,341/1,946	3,424/4,991	7,460/10,743	1,341/1,946	3,424/4,991	7,460/10,743
	/1,946/1,946	/4,991/4,991	/10,743/10,743	/1,946/1,946	/4,991/4,991	/10,743/10,743
	/1,400/1,946	/3,677/4,991	/10,743/8,024	/1,400/1,946	/3,677/4,991	/10,743/8,024
	/1,781/1,924	/4,574/4,932	/9,790/10,609	/1,781/1,924	/4,574/4,932	/9,790/10,609

* $p < 0.05$. This table reports the coefficients from an RD regression with income polynomials on each side of the discontinuity. Control variables are cohort × municipality fixed effects, parents' age, education, citizenship, marital status at birth of child, pre–child care family income, and mother's welfare and student status when the child was age 4.

labor market and care for the children. We also look at mother's part-time work when the child is aged 5 and see no effect (almost no fathers work part time). When we study mother's and father's income at age 5, we see no significant effect on mother's or father's income.

But could one year of higher disposable income have this large an effect on student's performance many years later? Given the availability of data on parental income and labor force participation throughout the child's school years, we are able to see if there are longer-run effects on family resources. Table 7 also presents the results when we examine average family income, calculated as the average from the years when the child was between 6 and 15. Importantly, we find that the child care subsidy at age 5 appears to have a significant and substantial positive effect on average family income. The exact magnitude differs by specification but is generally in the region of 10%.[37]

In summary, we find that the child care subsidy at age 5 leads to higher disposable income at that age but also higher family incomes at ages 6 to 15. We are not sure why it has this effect. One possibility is that these low-income people are liquidity constrained, and the extra disposable income allows them to move into self-employment. For example, Blanchflower and Oswald (1998) show that windfall gains are very important factors in enabling low-income people to make a transition to self-employment. Another possibility is that the extra income allows them to invest in human capital, and they receive the payoff from this in terms of higher earnings in later years. Unfortunately, we do not have the required data to test these hypotheses.

Assuming that the child care subsidy affects school scores through its effects on family income, we can estimate the implied effect of income on scores. Given that disposable income at age 5 increases by about 8% due to the subsidy if the child is in child care and the effect on test scores is about 17% to 50% of a standard deviation, this would imply that a 1% increase in family income at age 5 would increase scores by about 2% to 6% of a standard deviation. This would be a large effect of a once-off income shock, and it is substantially larger than the short-run effect reported in Dahl and Lochner (2012), where an increase in income of about 20% from a tax credit increased test scores in math and reading test scores by about 6% of a standard deviation. However, in our case, persons below the cutoff also have about 10% higher family income when the child is aged 6 to 15. So the magnitude of our effects needs to be interpreted in the context of a permanent increase in family income rather than that of a once-off shock to income.[38]

[37] We have also examined parents' labor force participation, calculated as the total years of employment when the child was between the ages of 6 and 15. Importantly, we find that the child care subsidy at age 5 appears to have a significant positive effect on father's labor force participation.

[38] We also calculated the OLS estimates from a regression of the outcome variables on family income at age 5. To avoid using variation from the discontinuity, we included a dummy control for whether the case was below the discontinuity. When we used the level of family income, the estimates implied that an increase in family income of NOK 10,000 (approximately the average value of being just below the cutoff), increases scores by about .002 to .01 of a standard deviation. Using log family income, the estimates implied that an income change of about 10% (again, approximately the difference between being above and below the cutoff) leads to a change in grades of about .005 to .01 of a standard deviation. Thus, the OLS income effects are much lower than our point estimates using the discontinuity. This is not surprising given the existing literature. Dahl and Lochner (2012) find a similar pattern, with their IV results being larger than their OLS. There are a number of possible explanations for this difference. The most obvious is that the RD is estimating local average treatment effects, and large shocks to income mean more for low-income individuals. Another explanation may be that the one-time large income subsidy for child care attendance has a longer-run effect on the income trajectories of the parents beyond just the initial income shock. This explanation is consistent with our findings on the mechanisms.

Table 8.—Effect of Being below the Income Cutoff on (Closest) Older Sibling's Junior High Academic Performance

Polynomials	1	1	1	2	2	2
Window/Bandwidth	.1	.25	.5	.1	.25	.5
Grade point average	.266*	.114	.066	.390*	.211*	.070
	(.109)	(.062)	(.042)	(.162)	(.094)	(.064)
Written exam	.165	.067	.012	.140	.084	.037
	(.110)	(.063)	(.043)	(.162)	(.095)	(.065)
Oral exam	.194	.091	.113*	.259	.073	.062
	(.117)	(.067)	(.046)	(.171)	(.102)	(.069)
N (gpa/wri/oral)	1,835 /1,708/1,578	4,727 /4,401/4,064	10,211 /9,546/8,796	1,835 /1,708/1,578	4,727 /4,401/4,064	10,211 /9,546/8,796

* $p < 0.05$. This table reports the coefficients from an RD regression with income polynomials on each side of the discontinuity. Control variables are cohort × municipality fixed effects, parents' age, education, citizenship, marital status at birth of child, pre–child care family income, and mother's welfare and student status when the child was age 4.

An alternative is to compare the magnitudes to the effects of early preschool intervention programs. These have often been found to have large effects. For example, the Abecedarian Project (IQ increases of .75 to 1.0 standard deviations), the Perry Preschool project (IQ increases of .60 standard deviations), and the Tennessee class size experiment (increased achievement by approximately .2 standard deviations) (Duncan, Morris, & Rodrigues 2011).

E. *Siblings*

Finally, given that it seems the most likely mechanism is disposable income during childhood, an important check of our results would be to look at the effects of the subsidy on the other children in the family. If the subsidy is in fact increasing disposable income for the family, then all children should benefit, and not just the child who generates the subsidy. In table 8, we report effects for older siblings and see that there are tendencies toward positive effects, although they are more imprecisely estimated due to smaller sample sizes. When we look at younger siblings, the estimates are too imprecisely estimated to draw any conclusions.

VI. Conclusion

Given the wide use of child care subsidies across countries, it is surprising how little we know about the effect of these subsidies on children's longer-run outcomes. Using a sharp discontinuity in the price of child care in Norway, we are able to isolate the effects of child care subsidies on both parental and student outcomes. We find very small and statistically insignificant effects of child care subsidies on child care utilization and parental labor force participation. Despite this, we find significant positive effects of the subsidies on children's academic performance in junior high school, suggesting that the positive shock to disposable income provided by the subsidies may help to improve children's scholastic aptitude.

Policy recommendations based on the results in this paper point toward increasing disposable income for low-income families. Norway subsidizes child care with NOK 28 billion (USD $4.5 billion) yearly, and most of these subsidies are universal. A move toward more income-means-tested subsidies may be beneficial for children. Our findings suggest that the child care subsidy in Norway for 5-year-olds works as an in-kind transfer, providing families with more disposable income for a period of early childhood. It also leads to greater labor force participation in later years and therefore a permanent positive shock to family income.

As in Norway, child care subsidies are available to certain groups in the United States (Herbst & Tekin, 2010a). Given the low level of the income cutoffs for subsidies in both the United States and Norway, the affected population looks quite similar. Subsidy recipients tend to be significantly less educated and poorer than the population as a whole. They are substantially more likely to be a minority in the United States and more likely to be a nonnative in Norway. That said, the institutional environments are sufficiently different that one might be reluctant to generalize findings across these countries. These results are more likely to translate to other countries with an extended welfare state, such as the countries in Northern Europe.

We cannot rule out that subsides targeted at other ages or in other settings might give different parental responses. A key limitation of this paper is that although we are able to identify a significant and positive effect of the child care subsidy on children's performance, we are ultimately unable to disentangle the underlying causes of this effect. This suggests the need for future work focusing on understanding these mechanisms.

REFERENCES

Becker, Gary S., and Nigel Tomes, "An Equilibirum Theory of the Distribution of Income and Intergenerational Mobility," *Journal of Political Economy* 87 (1979), 1153–1189.

Berlinski, Samuel, Sebastian Galiani, and Paul Gertler, "The Effect of Pre-Primary Education on Primary School Performance," *Journal of Public Economics* 93 (2009), 219–234.

Berlinski, Samuel, Sebastian Galiani, and Marco Manacorda, "Giving Children a Better Start: Preschool Attendance and School-Age Profiles," *Journal of Public Economics* 92 (2009), 1416–1440.

Blanchflower, David G., and Andrew J. Oswald, "What Makes an Entrepreneur?" *Journal of Labor Economics* 16 (1998), 26–60.

Blau, David M., "The Effect of Income on Child Development," this REVIEW 81 (1999), 261–276.

——— "Child Care Subsidy Programs," NBER working paper 7806 (2000).

Blau, David, and Erdal Tekin, "The Determinants and Consequences of Child Care Subsidies for Single Mothers in the USA," *Journal of Population Economics* 20 (2007), 719–741.

Blix, Karin W. S., and Lars Gulbrandsen, "Child Care," Ministry of Education Report (2002).
Carneiro, Pedro M., and James J. Heckman, "Human Capital Policy," in J. Heckman and A. Krueger, eds., *Inequality in America: What Role for Human Capital Policy?* (Cambridge, MA: MIT Press, 2003).
Currie, Janet, "Healthy, Wealthy, and Wise: Socioeconomic Status, Poor Health in Childhood, and Human Capital Development," *Journal of Economic Literature* 47 (2009), 87–122.
Dahl, Gordon B., and Lance Lochner, "The Impact of Family Income on Child Achievement Evidence from Changes in the Earned Income Tax Credit," *American Economic Review* 102 (2012), 1927–1956.
Dooley, Martin, and Jennifer Stewart, "Family Income and Child Outcomes in Canada," *Canadian Journal of Economics/Revue canadienne d'economique* 37 (2004), 898–917.
Duncan, Greg J., Pamela A. Morris, and Chris Rodrigues, "Does Money Really Matter? Estimating Impacts of Family Income on Children's Achievement with Data from Random-Assignment Experiments," *Developmental Psychology* 47 (2011), 1263–1279.
Duncan, Greg J., W. Jean Yeung, Jeanne Brooks-Gunn, and Judith R. Smith, "How Much Does Childhood Poverty Affect the Life Chances of Children?" *American Sociological Review* 63 (1998), 406–423.
Fan, Jianqing, "Design-Adaptive Nonparametric Regression," *Journal of the American Statistical Association* 87 (1992), 998–1004.
Fitzpatrick, Maria D., "Preschoolers Enrolled and Mothers at Work? The Effects of Universal Pre-Kindergarten," *Journal of Labour Economics* 28 (2010), 51–85.
Gulbrandsen, Lars, and Winsvold, Aina, "Kvalitet og kvantitet Kvaliteten i barnehagen i sterk vekst.," *NOVA rapport 2/2009* (2009). ["Quality and Quantity: The Quality in day care is increasing".]
Hahn, Jinyong, Petra Todd, and Wilbert Van der Klaauw, "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design," *Econometrica* 69 (2001), 201–209.
Havnes, Tarjei, and Magne Mogstad, "Is Universal Child Care Levelling the Playing Field? Evidence from Non-Linear Difference-in-Differences," IZA discussion paper 4978 (2010).
——— "No Child Left Behind: Subsidized Child Care and Children's Long-Run Outcomes," *American Economic Journal: Economic Policy* 3 (2011), 97–129.
Herbst, Chris M., and Erdal Tekin, "Child Care Subsidies and Child Development," *Economics of Education Review* 29 (2010a), 618–638.
——— "The Impact of Child Care Subsidies on Child Well-Being: Evidence from Geographic Variation in the Distance to Social Service Agencies, NBER working paper 16250 (2010b).
Imbens, Guido, and Karthik Kalyanaraman, "Optimal Bandwidth Choice for the Regression Discontinuity Estimator," NBER working paper 14726 (2009).
Lee, David S., and Thomas Lemieux, "Regression Discontinuity Design in Economics," NBER working paper 14723 (2009).
Levy, Dan M., and Greg Duncan, "Using Sibling Samples to Assess the Effect of Childhood Family Income on Completed Schooling," Northwestern University, JCPR working paper 168 (2000).
Løken, Katrine V., "Family Income and Children's Education: Using the Norwegian Oil Boom as a Natural Experiment," *Labour Economics* 17 (2010), 118–129.
Løken, Katrine V., Magne Mogstad, and Matthew Wiswall, "What Linear Estimators Miss: The Effects of Family Income on Child Outcomes," *American Economic Journal: Applied Economics* 4 (2012), 1–35.
Magnuson, Katherine A., Christopher Ruhm, and Jane Waldfogel, "Does Prekindergarten Improve School Preparation and Performance?" *Economics of Education Review* 26 (2007), 33–51.
Milligan, Kevin, and Mark Stabile, "The Integration of Child Tax Credits and Welfare: Evidence from the Canadian National Child Benefit Program," *Journal of Public Economics* 91 (2007), 305–326.
OECD, "Thematic Review of Early Childhood Education and Care Policy" (Paris: OECD, 1999).
Oreopoulos, Philip, Marianne Page, and Ann H. Stevens, "The Intergenerational Effects of Worker Displacement," *Journal of Labor Economics* 26 (2008), 454–483.
Porter, Jack, "Estimation in the Regression Discontinuity Model," unpublished manuscript, Department of Economics, University of Wisconsin at Madison (2003).
Schøne, Pål, "Labour Supply Effects of a Cash-for-Care Subsidy," *Journal of Population Economics* 17 (2004), 703–727.
Shea, John, "Does Parents' Money Matter?" *Journal of Public Economics* 77 (2000), 155–184.
Tekin, Erdal, "Child Care Subsidy Receipt, Employment, and Child Care Choices of Single Mothers," *Economic Letters* 89 (2005), 1–6.
——— "Childcare Subsidies, Wages, and Employment of Single Mothers," *Journal of Human Resources* 42 (2007), 453.
White, Halbert, "A Heteroskedasticity-Consistent Covariance Matrix Estimator and a Direct Test for Heteroskedasticity," *Econometrica* 48 (1980), 817–838.