

Havnes, Tarjei; Mogstad, Magne

Working Paper

Is universal child care leveling the playing field?

CESifo Working Paper, No. 4014

Provided in Cooperation with:

Ifo Institute – Leibniz Institute for Economic Research at the University of Munich

Suggested Citation: Havnes, Tarjei; Mogstad, Magne (2012) : Is universal child care leveling the playing field?, CESifo Working Paper, No. 4014, Center for Economic Studies and ifo Institute (CESifo), Munich

This Version is available at:

<https://hdl.handle.net/10419/68210>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

Is Universal Child Care Leveling the Playing Field?

Tarjei Havnes
Magne Mogstad

CESIFO WORKING PAPER NO. 4014

CATEGORY 5: ECONOMICS OF EDUCATION

NOVEMBER 2012

An electronic version of the paper may be downloaded

- *from the SSRN website:* www.SSRN.com
- *from the RePEc website:* www.RePEc.org
- *from the CESifo website:* www.CESifo-group.org/wp

Is Universal Child Care Leveling the Playing Field?

Abstract

Many developed countries currently consider a move towards a universal child care program. The challenge in assessing the case for universal child care programs is that the evidence base is scarce. We analyze the staged expansion of subsidized, universally accessible child care in Norway. Our approach differs from existing literature which estimate mean impacts of universal child care programs; in contrast, we use non-linear difference-in-differences methods to estimate quantile treatment effects, showing how the child care expansion affected the earnings distribution of exposed children as adults. We complement these estimates with local linear regressions of the child care effects by family income. Our findings suggest that the effects of child care vary systematically across the earnings distribution, that the mean impact misses a lot, and that children of low income parents seem to be the primary beneficiaries of subsidized child care. These findings are important when considering the case for universal child care programs, since the benefits of providing subsidized child care to middle and upper-class children are unlikely to exceed the costs. To help understand the differential effects on earnings, we examine how the child care expansion affected the educational trajectories and cognitive test scores of the exposed children.

JEL-Code: J130, H400, I280, D310.

Keywords: universal child care, child development, non-linear difference-in-differences, heterogeneity, quantile treatment effects.

Tarjei Havnes
Department of Economics
University of Oslo
Oslo / Norway
tarjei.havnes@econ.uio.no

Magne Mogstad
Department of Economics
University College London
London / United Kingdom
magne.mogstad@gmail.com

November, 2012

We appreciate the useful comments and suggestions from Erling Barth, Manudeep Bhuller, Nabanita Datta Gupta, Hilary Hoynes, Kalle Moene, Halvor Mehlum, Mari Rege, and Kjetil Telle, participants at the “Conference on the Nordic Model” at the University of Oslo, as well as participants at numerous seminars and conferences. The project received financial support from the Norwegian Research Council (Grant Number 212305). The project is also part of the research activities at the ESOP center at the Department of Economics, University of Oslo. ESOP is supported by The Research Council of Norway.

1 Introduction

What is the case for universal child care programs? This question is important for both policy and scientific research. In the United States, most child care is provided by the private market, and publicly subsidized child care (or preschool) is predominantly targeted at low income families or single parents. An alternative model is supplied by the universal public programs found in Canada and the Scandinavian countries where all children are eligible for subsidized child care, regardless of family background. Many developed countries currently consider a move towards a universal child care program.¹ Critics of such a move argue that society should focus its investment on children from low income families, where returns are likely to be greatest; extending subsidized child care to middle and upper-class children may require a considerable increase in taxes, at the cost of economic efficiency. Proponents of universal programs counter that even if returns are greater for the poor, publicly subsidized child care may still have benefits for middle or upper-class children that exceed its costs.

The challenge in assessing the case for universal child care programs is that the direct evidence of the impact of these policies is scarce. Perhaps as a consequence, the case for universal programs is often made drawing on evidence from targeted programs, suggesting that early childhood interventions can generate learning gains in the short-run and improve the long-run prospects of children from poor families.² As emphasized by Baker (2011), using the evidence from targeted programs for the promotion of universal ones is problematic for a number of reasons. First, the widely cited targeted programs offer levels of education and care that are typically not found in programs for all children. Therefore, it is not clear that universal child care can deliver similar benefits to disadvantaged children. Second, even if universal programs could offer similar levels of care and education, it may be that the benefits for middle or upper-class children do not exceed the costs.

In this paper, we investigate the distributional effects on children's long-run outcomes of a reform from late 1975 in Norway, which led to a large-scale expansion of subsidized child care. All children 3–6 years old were eligible regardless of parents' employment and marital status, and available child care slots were in general allocated according to length of time on the waiting list.³ Our focus is on quantifying the effects of subsidized child care on the long-run outcomes of middle and upper-class children as compared to children of lower-class parents. An advantage of our long-run perspective is that we get round the

¹For example, the European Union's Presidency formulated in 2002 as a policy goal "to provide child care by 2010 to at least 90 % of children between 3 years old and the mandatory school age and at least 33 % of children under 3 years of age" (EU, 2002, p. 13).

²The Perry Preschool and Abecedarian programs are commonly cited examples of how high-quality preschool services can improve the lives of disadvantaged children. See Barnett (1995) and Karoly et al. (2005) for surveys of the literature. We refer to Baker (2011) for examples of how policy documents draw on the evidence base for targeted programs to promote universal child care programs.

³Throughout this paper, subsidized child care refers to publicly and private child care institutions, eligible for subsidies from the government because they satisfy federal quality requirements.

issues of whether short-run impacts of child care persist, and perhaps are amplified, over time.⁴

The reform we study led to a staged expansion of subsidized child care, across Norway's more than 400 municipalities. To control for unobserved differences between children born in different years as well as between children from treatment and comparison municipalities, we use a difference-in-differences (DiD) identification that exploits the temporal and spatial variation in child care coverage induced by this staged expansion. Roughly speaking, we compare the adult outcomes for 3 to 6 year olds before and after the reform, from municipalities where child care expanded a lot and municipalities with little or no increase in child care coverage. As described in detail below, subsidized child care both before the reform and during the expansion was severely rationed, with informal care arrangements (such as friends, relatives, and unlicensed care givers) servicing the excess demand. In our analysis, we will focus on the years immediately after the reform, when child care coverage increased from 10 % to 28 %, which likely reflects the abrupt slackening of constraints on the supply side caused by the reform, rather than a spike in local demand. We have found no other reforms or changes taking place in this period which could confound our estimated child care effects. Nevertheless, to increase our confidence in the empirical strategy, we perform a number of specification checks.

We begin by using non-linear DiD methods to estimate how the child care reform affected the outcome distribution of exposed children as adults. The estimated quantile treatment effects (QTE) allow us to assess the impact of the universal child care program on the lower, middle and upper part of the outcome distribution.⁵ We complement the QTE estimates with local linear regression estimates of the program effects by family income. While the QTE estimates are useful for assessing the effect on the outcome distribution, the local linear regression estimates give information important for targeting.

We find that mean impacts in the population of children as a whole miss a lot, concealing major heterogeneity in the effects of the child care reform. This is an important observation since previous empirical studies of universal child care have focused on mean impacts. While the reform had a small and insignificant mean impact on earnings, QTE-estimates are positive in the lower part of the earnings distribution and negative in the uppermost part. The local linear regression estimates support these findings: Most of the gains in earnings associated with the universal child care program relate to children of low income parents, whereas upper-class children actually experience a loss in earnings because of

⁴For example, as pointed out by Baker, Gruber, and Milligan (2008), their findings of a negative short-run impact of child care on children's non-cognitive development, could represent an initial cost of socialization, with little or no long-run consequences. Moreover, if investments in human capital have dynamic complementarities, then even a small learning gain in the short-run may improve the long-run prospects of children considerably (Heckman, Stixrud, and Urzua (2006)).

⁵Bitler, Gelbach, and Hoynes (2006) use data from a randomized controlled trial to show the usefulness of QTE in situations where theory predicts systematic heterogeneity in the impact of the welfare reforms. See Athey and Imbens (2006) for a discussion of non-linear difference-in-differences methods.

the program. In line with the differential effects by family income, we find that the child care reform reduced the intergenerational income elasticity by 2.5 percentage points. We show that the estimated heterogeneity in child care effects is consistent with a simple model where parents make a tradeoff between family consumption today and investment in children.⁶

To clarify the nature of the relationship between the child care reform and the earnings of exposed children as adults, we examine the mediating role of educational attainment and cognitive test scores. We find a positive mean impact of the universal child care program on educational attainment. Most of the growth in educational attainment stems from a substantial increase in the likelihood of completing high school. The increase in educational attainment is largely driven by children of low income parents, while there is no change in schooling levels of upper-class children. We find no reform effect on children's scores on cognitive tests. All our QTE estimates on test scores are close to zero, and they are sufficiently precise to rule out any economically relevant effect. The absence of a change in cognitive test scores points to the importance of non-cognitive skill development to understand the long-run effects of universal child care.⁷

Taken together, our results have important implications for a small but rapidly growing literature on universal or large-scale child care programs.⁸ The evidence is mixed. Baker, Gruber, and Milligan (2008) and Lefebvre, Merrigan, and Verstraete (2008) analyze the introduction of subsidized, universally accessible child care in Quebec. Baker, Gruber, and Milligan (2008) report negative mean impacts on indices of behaviour, while Lefebvre, Merrigan, and Verstraete (2008) report negative mean impacts on cognitive test scores. These negative effects echo the results of subsidized child care in the United States reported in Herbst and Tekin (2008). Datta Gupta and Simonsen (2010) exploit variation across municipalities in Denmark in access to center based care. They find that, compared to home care, being enrolled in center-based care at age three does not lead to significant differences in child outcomes at age seven. Havnes and Mogstad (2011) examine the child care reform from late 1975 in Norway, reporting positive mean impacts on educational attainment and labor market attachment. Further evidence from Norway is provided by Black, Devereux, Loken, and Salvanes (2012), exploiting discontinuities in child care subsidies depending on family income. They find significant positive effects of the subsidies on the school performance of children from low income families.

Our study complements this literature by providing the first evidence on the distributional effects of subsidized child care. Our findings emphasize that the impact of subsidized

⁶Our findings are also consistent with parents overestimating the quality of care that their children receive in child care or pre-school, as found in previous studies (see e.g. Cryer and Burchinal, 1997).

⁷Evidence from targeted programs shows that even though cognitive test score advantages for children with preschool experiences tend to fade out as they go through school, the early interventions have large payoffs in educational attainment and the labor market because of improvements in non-cognitive skills (see e.g. Heckman and Masterov (2007)).

⁸For recent reviews, see Almond and Currie (2011), Baker (2011), and Ruhm and Waldfogel (2012)

child care may be both more varied and more extensive than what previous studies suggest. Our results highlight that the effects of child care vary systematically across the outcome distribution, and that children of low income parents seem to be the primary beneficiaries. This is especially important when considering the case for universal child care programs, since the benefits of providing subsidized child care to middle and upper-class children are unlikely to exceed the costs.

The paper proceeds as follows. Section 2 reviews key facts regarding the child care system in Norway, describes the child care reform and discusses its expected impact on child development. Section 3 describes the empirical strategy and Section 4 presents the data. Section 5 presents our empirical findings and provides sensitivity analyses. Section 6 concludes.

2 Background

2.1 The child care reform

In the post-WWII years in Norway, the gradual entry on the labor market of particularly married women with children, caused growing demand for out-of-home child care. In a survey from 1968, when child care coverage was less than 5 %, about 35 % of mothers with 3 to 6 year olds stated demand for subsidized child care (NOU, 1972). In the same survey, only 34 % of the latter group of respondents stated that they were in fact using out-of-home child care on a regular basis. Out of these, just 14 % were in subsidized child care, while more than 85 % were using informal arrangements. The severe rationing of subsidized child care acted as a background for political progress towards public funding of child care.

In June 1975, the Kindergarten Act was passed by the Norwegian parliament with broad bipartisan political support. It assigned the responsibility for child care to local municipalities, but included federal provisions on the price, educational content, group size, staff skill composition, and physical environment. By increasing the level of federal subsidies for both running costs in general and investment costs for newly established institutions, the government aimed at quadrupling the number of child care places to reach a total of 100,000 by 1981. In the years following the reform, the child care expansion was progressively rolled out at a strong pace, with federal funding more than doubling from USD 34.9 million in 1975 to 85.8 million in 1976, before reaching 107.3 million in 1977.⁹ This implied an increase in the federal coverage of running costs from about 10 % in 1973 to 17.6 % in 1976, and further to 30 % in 1977. From 1976, newly established child care places received additional federal funds for a period of five years. Municipalities with relatively low child care coverage rates were awarded the highest subsidies.

⁹Throughout this paper, all monetary figures are reported in 2006 levels after adjusting for inflation.

Altogether, the reform constituted a substantial positive shock to the supply of subsidized child care, which had been severely constrained by limited public funds. In succeeding years, the previously slow expansion in subsidized child care accelerated rapidly. From a total coverage rate of less than 10 % for 3 to 6 year olds in 1975, coverage had shot up above 28 % by 1979. Over the period, a total of almost 38,000 child care places were established, more than a doubling from the 1975-level. By contrast, there was almost no child care coverage for 1 and 2 year olds during this period. Figure 1 draws child care coverage rates in Norway from 1960 to 1996 for 3 to 6 year olds. As is apparent from the figure, there has been strong growth in child care coverage rates since 1975, particularly in the early years. In our analysis, we will focus on the early expansion, which likely reflects the abrupt slackening of constraints on the supply side caused by the reform, rather than a spike in the local demand.

Both private and public child care institutions that satisfied the federal provisions received subsidies for running and establishment costs from the federal government. Subsidies were determined on the basis of the number and age of children, and the amount of time they spend in subsidized child care. In general, the child care institutions were open during normal working hours. All children were eligible, and open slots were in general allocated according to length of time on the waiting list and age. Only under special circumstances could a child gain priority on the waiting list.

The quality of subsidized child care in Norway appears to be comparable to universal child care programs in other countries but substantially lower than most targeted programs (see Datta Gupta and Simonsen, 2010). The average yearly expenditure for a slot in subsidized child care was approximately USD 6,600. This is slightly higher than the expenditures for the universal Canadian child care program and somewhat lower than center based child care in Denmark. These expenditure levels are mirrored in the number of children per staff. In the Norwegian program, the average staff-child ratio was about 1:8; The corresponding ratio is 1:12 in Canada and 1:7 in Denmark.

2.2 *The expected effects*

The child care reform can be interpreted as a subsidy to parents for choosing out-of-home care of a particular quality, generating at least one and possibly two convex kinks in the family's budget constraints. Depending on counterfactual quality of care, the effects of the subsidized child care may vary in magnitude and even sign. To see this, consider a simple model of parental investment where families are unitary and can affect their child's quality Q by investing in the quality of care q and by purchasing market goods k , according to a child production function $h(q, k)$ that is increasing and concave in both arguments. Assume that both quality of care and market goods can be purchased at prices p_q and p_k , respectively. With a total time endowment L and a potential wage w , the family budget




Figure 1. Child care coverage rates for children 3 to 6 years old, 1960–1996.
Notes: Data from 1972 are based on administrative registers. Data for 1960–1972 are taken from NOU 1972:39 (1960-1972).




Figure 2. Child care coverage rates for children 3 to 6 years old by treatment (solid) and comparison group (dashed).
Notes: Treatment (comparison) municipalities are above (below) the median in child care coverage from 1976 to 1979.

is then

$$wL = c + p_q q + p_k k \quad (1)$$

where the numeraire good c is leisure or consumption not affecting child quality.

Before the introduction of subsidized care, parents minimize the cost of producing a given level of child quality by equating the marginal cost of care quality q and child goods k . The effective family budget can then be summarized as

$$wL = c + p_Q(Q)$$

where $p_Q(Q)$ is the minimum total cost in consumption units of producing a given child quality Q , which is increasing and convex. Given the budget wL , we can then map out the consumption possibility frontier of the parents, drawn in panel (a) of Figure 3.

As usual, parents optimize utility by equating the marginal benefit of child quality to the marginal cost of provision in terms of foregone consumption. In panel (a) of Figure 3, this is illustrated in the tangency between an indifference curve and the consumption possibility frontier. The upper tangency is for a family with a high value of consumption relative to child quality, while the lower tangency is for a family with low value of consumption relative to child quality.

Subsidized care makes child care of a particular quality, say \bar{q} , available at a price \bar{p}_q lower than the market price.¹⁰ A family that opts into subsidized care locks in the quality of care, higher or lower than what they would choose in the market. The family may compensate for lower quality care by purchasing more child goods k . Families may also choose not to use subsidized care, going instead to the market analyzed above.

Using the budget and production function, the total cost of child quality is now

$$\bar{p}_Q(Q) = \bar{p}_q \bar{q} + p_k(Q; \bar{q})$$

where $p_k(Q; \bar{q})$ is total expenditure on market goods to produce child quality Q , given the quality of subsidized care. Opting into subsidized care changes the consumption possibility frontier, as compared to the market. First, note that at $Q = 0$, $\bar{p}_Q(0) = \bar{p}_q \bar{q} > p_Q(0) = 0$ such that maximum family consumption is unambiguously reduced by choosing subsidized care. Second, families opt into the subsidized program only if $\bar{p}_Q(Q) \leq p_Q(Q)$ for some Q , implying that the frontier with subsidies is outside the market frontier. Finally, maximum child quality depends on the ability of parents to compensate for a lower quality of care. That is, it depends on the complementarity between quality of care and market goods

¹⁰To receive subsidies from the federal government, child care institutions were obliged to meet provisions on the maximum price to be paid by the parents, educational content and activities, group size, staff skill composition and physical environment. As a consequence, subsidized child care was quite homogenous with respect to observable characteristics such as group size, staff-child ratio, and employee training (Havnes and Mogstad, 2011). Our theoretical framework therefore assumes that subsidized child care is of homogenous quality.

in the production of child quality. If market goods are a poor substitute for quality of care, then families with a strong preference for child outcomes may not be able to fully compensate for the lower quality offered in subsidized care. The new frontier may then cross the former a second time.

Two alternative scenarios are illustrated in panels (b) and (c) of Figure 3. In panel (b), complementarity between care quality and market goods is strong, such that families with a strong preference for child quality are unable to fully compensate for the lower care quality if they opt into subsidized care. In panel (c), complementarity between care quality and market goods is weaker, and families are able to fully compensate by purchasing market goods if subsidized care is of lower quality than what they would purchase in the market.


Panels (d) and (e) of Figure 3 draw the predicted effects on child outcomes of introducing subsidized child care. For very low or very high levels of child quality, parents do not opt into subsidized care, and effects are zero. In the intermediate region, effects are positive in the lower parts of the distribution, then decrease in size, and may turn negative at the top. Families may accept some negative effects on children because of the associated gains in family consumption. There are no unambiguous predictions about the mean impacts of introducing subsidized child care: The sign and magnitude depend on the size of the different effects weighted by the relative number of children along the family's budget constraint without subsidized care. The model does, however, predict that the introduction of subsidized care equalizes the outcome distribution of children. If child care quality is a normal good, we also expect that the benefits of subsidized care is negatively correlated with family income. This means that subsidized child care should reduce the intergenerational income elasticity.

3 Empirical strategies


To estimate the effect of the expansion of subsidized child care on children's outcomes we follow Havnes and Mogstad (2011) in applying a DiD approach, exploiting that the supply shocks to subsidized child care were larger in some areas than others. Below, we will first describe the identification strategy, before discussing the econometric models we use to estimate distributional impacts and mean impacts.

3.1 Identification strategy


Our identification strategy is the following: We compare the adult outcome of interest for 3 to 6 year olds before and after the reform, from municipalities where child care expanded a lot (i.e. the treatment group) and municipalities with little or no increase in child care coverage (i.e. the comparison group).




(a) Pre-reform




(b) Post-reform – strong complementarity



(c) Post-reform – weak complementarity



(d) Effect of reform (i)



(e) Effect of reform (ii)

Figure 3. Predicted effects of introducing subsidized care.

Notes: To generate the graphs, we have parametrized the model in Section 2.2 assuming that the child production function is of the CES-form, while the parental utility function is CRRA.

The child care expansion started in 1976, affecting the *post-reform cohorts* born 1973–1976 with full force. The *pre-reform cohorts* consist of children born in the period 1967–1969. We consider the period 1976–1979 as the child care *expansion period*. Starting in 1976 gives the municipalities some time to plan and react to the policy change. Also, 1976–1979 was the period with the largest growth in child care coverage. In the robustness analysis, we make sure that our results are robust to changes in the exact choice of expansion period.

To define the treatment and comparison group, we order municipalities according to the percentage point increase in child care coverage rates from 1976 to 1979. We then separate the sample at the median, the upper half constituting the *treatment municipalities* and the lower half the *comparison municipalities*. Figure 2 shows child care coverage before and after the 1975 reform in treatment and comparison municipalities (weighted by population size). The graphs move almost in parallel before the reform, while child care coverage in treatment municipalities kinks heavily after the reform. This illustrates that our study compares municipalities that differ distinctly in terms of changes in child care coverage within a narrow time frame. In the robustness analysis, we take several steps to ensure that our results are robust to the exact child care coverage cut-off, defining treatment and comparison municipalities.

Our identification strategy raises two key selection issues. First, there could be selection of municipalities into expanding child care depending on the gains in child development from subsidized child care. For instance, it could be that child care expanded most in municipalities where child development were most responsive to subsidized child care. This means that we need to be cautious in extrapolating the DiD estimates to effects for the population of children at large.

Second, if there is selection of municipalities into expanding child care depending on the underlying time trends in child outcomes in the absence of the child care expansion, then the DiD approach will provide biased estimates of the effects on children who reside in the treatment area in the post-expansion period. For instance, it could be that the municipalities that expanded child care the most did so to counteract a negative trend in child outcomes. To address the concern for selection bias in the effects of subsidized child care, Section 4 looks closely into the determinants of the child care expansion across municipalities. We also run a battery of specification checks, which are discussed after the main results.

3.2 *Econometric models*

We will use non-linear DiD methods to estimate how the child care reform affected the outcome distribution of exposed children as adults. The estimated quantile treatment effects (QTE) allow us to assess the impact of the universal child care program on the lower, middle and upper part of the outcome distribution.

To fix ideas, suppose that subsidized child care was randomly assigned to municipalities. Then the QTE could have been identified directly from the earnings distributions of the post-reform cohorts from treatment and comparison municipalities: At any quantile $q \in (0, 1)$, the QTE is simply the difference in the q -th earnings quantile in the treatment versus the comparison group. Figure 4a illustrates by showing how the QTEs are given by the horizontal differences in the cumulative distribution functions of the outcome between the treatment and the comparison group.

With observational data, identification of QTE requires assumptions about the counterfactual earnings distribution in the absence of subsidized child care. In the spirit of the standard DiD estimator, non-linear DiD methods use the observed change in the distribution of the comparison group, from before to after treatment assignment, as an estimate of the change that would have occurred in the treatment group over this period in the absence of treatment. Let $F_t(y)$ ($G_t(y)$) be the adult earnings distribution of children from treatment (comparison) municipalities born in period t , where $t = 0$ denotes pre-reform cohorts and $t = 1$ denotes post-reform cohorts. The counterfactual earnings distribution of the post-reform cohorts from treatment municipalities is assumed to be

$$\mathcal{F}_1(y)_{RIF} = F_0(y) + (G_1(y) - G_0(y)).$$

The estimated DiD-effect in probability points is then given by the difference between the actual and the counterfactual earnings distribution at any earnings level y , i.e. $(1 - F_1(y)) - (1 - \mathcal{F}_1(y)_{RIF})$.

Figure 4b illustrates the procedure, which compares the vertical difference in the earnings distribution of pre- and post-reform cohorts from treatment and comparison municipalities. In Figure 4b, the estimated impact at $y = y'$, measured in population shares, is $-[(a - b) - (c - d)]$. Following Firpo, Fortin, and Lemieux (2009), we can divide this difference by a kernel estimate of the joint density of earnings at y' to arrive at the associated QTE. This procedure may be viewed as a recentered influence function (RIF) regression (Firpo, Fortin, and Lemieux, 2009), and we refer to it as RIF-DiD. In Section 5, we show that the QTE estimates are robust to the use of the alternative non-linear DiD methods discussed by Athey and Imbens (2006).

The identifying assumption underlying the RIF-DiD estimator is that *the change in population shares* from before to after treatment around a given *level of earnings* would be the same in the treatment group as in the comparison group, in the absence of treatment. In comparison, the standard DiD-estimator assumes a common trend in mean earnings in the absence of treatment. The counterfactual mean earnings of the treatment group in the post-reform period is then $\mu_{\mathcal{F}_1}^{DD} = \mu_{F_0} + (\mu_{G_1} - \mu_{G_0})$, where μ_F denotes the mean of the distribution F . The mean impact estimate is then $\mu_{F_1} - \mu_{\mathcal{F}_1}^{DD}$. In contrast to the standard DiD-estimator, the identifying assumption of RIF-DiD is invariant to monotonic

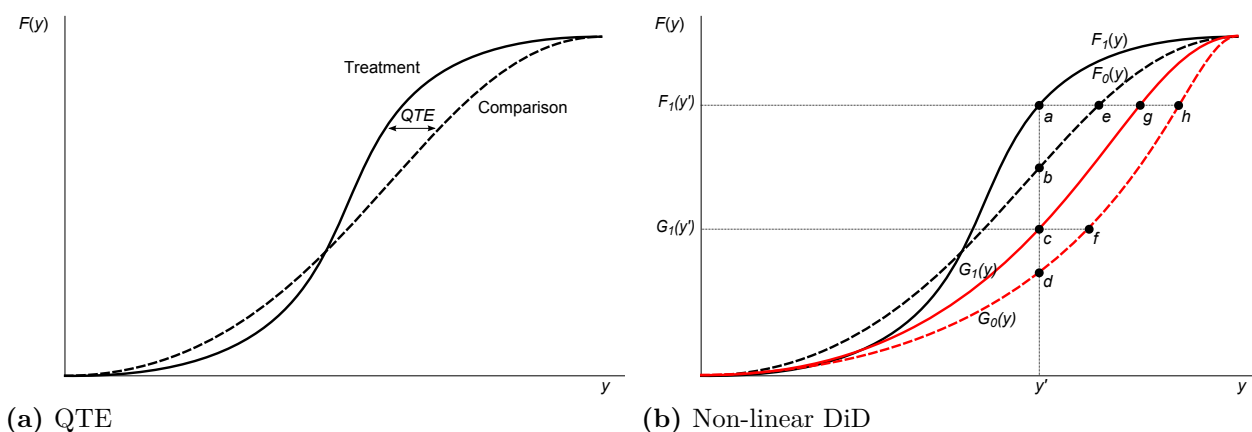


Figure 4. Illustration of the empirical strategy: Quantile treatment effects and non-linear DiD-estimators.

Notes: $F_t(y)$ ($G_t(y)$) are simulated outcome distributions for treated (comparison) before ($t = 0$) and after ($t = 1$) treatment. Indicated points illustrate alternative non-linear DiD estimators at $y = y'$: RIF-DiD: $-[(a - b) - (c - d)]$; CiC: $(a - e) - (c - f)$; QDiD: $(a - e) - (g - h)$.

transformations of the outcome (e.g. considering the log of earnings instead of earnings levels).

As in Baker, Gruber, and Milligan (2008) and Havnes and Mogstad (2011), we will focus on intention-to-treat effects of the child care reform, since the DiD strategy estimates the reduced form impact on all children from post-reform cohorts who reside in treatment municipalities. An advantage of the intention-to-treat effects is that they capture the full reform impact of the introduction of subsidized child care, including any changes in parental behavior and spillover effects on children who were not attending subsidized care. The intention-to-treat effects may be scaled by the increase in child care coverage following the reform in the treatment group relative to the comparison group. The scaled estimates give the impact of the introduction of subsidized child care – per child care place – on children born to post-reform cohorts in treatment municipalities.

4 Data

Data and sample restrictions. Our analysis employs several register data sets from Norway that we can link through unique identifiers for each individual.¹¹ We begin with a rich longitudinal database that covers every resident from 1967 to 2009. For each year, it contains individual demographic information (including gender, date of birth, and municipality of residency) and socio-economic data (including education and earnings). The data set includes unique family identifiers that allow us match parents to their children. We link this data with the national military records with information on cognitive test

¹¹The coverage and reliability of Norwegian register data is considered to be exceptional, as illustrated by the fact that they received the highest rating in a data quality assessment conducted by Atkinson et al. (1995).

scores. We also use an administrative register that covers all child care institutions eligible for public subsidies from from 1972 to 2009. In each year, the child care institutions report the number of children in child care by age. Merging this data with the demographic files containing information about the total number of children according to age and residency, we construct a time series of annual child care coverage (by age of child) in each municipality.

We focus on children born 1967–1969 (pre-reform cohorts) and children born 1973–1976 (post-reform cohorts). We restrict the sample to children of married mothers (prior to the reform), which makes up about 92 % of the above sample. The reason for this choice is that our data does not allow us to distinguish between cohabitants and single parents in these years. Family background characteristics (including family income and parental education) are measured prior to the reform. Rather than dropping a small number observations where information on family background characteristics is missing, we include separate categories for missing values. The final sample used in the estimations consists of 341,170 children.

The primary outcome variable of interest is the child’s average earnings over the period 2006–2009. Our earnings measure includes wages and income from self-employment. We also consider the child’s educational attainment and ability test scores. The ability score is only available for males because they are collected from military records, and military service is compulsory for men only.¹² Before entering the military, their cognitive ability is assessed; this occurs for the great majority between their eighteenth and twentieth birthdays. The ability test score at these ages is particularly interesting as it is about the time of entry to the labor market or to higher education. The test score is a composite score from three timed tests—arithmetic, word similarities, and figures (see Thrane 1977, and Sundet et al. 2004, 2005). The composite test score is an unweighted mean of the three subtests. The raw scores are standardized into a a nine point scale that has a discrete approximation to a normal distribution, a mean of 100, and a standard deviation of 15.

Descriptive statistics. Table 1 shows mean earnings and earnings levels at different percentiles. As is evident from the table, there are systematic changes in the differences in the earnings of the treatment and comparison group, between the pre-reform and post-reform cohorts. In our DiD approach, this pattern is suggestive of significant effects of child care on children’s earnings. In particular, Table 1 indicates positive effects in the lower and central parts, and negative effects in the upper parts of the earnings distribution. In contrast, there is no evidence of changes over time in the characteristics of children and

¹²We have test scores for about 84 % of all Norwegian men born in the years we consider. Eide, Øyen, Skjærven, Nilsen, Bjerkedal, and Tell (2005) examine patterns of missing scores for the men in the 1967–1987 cohorts. Of those, 1.2 % died within 1 year and 0.9 % died between 1 year of age and registering with the military at about age 18. About 1 % of men had emigrated before age 18, and 1.4 % were exempted because they were permanently disabled. An additional 6.2 % of scores are missing for a variety of reasons, most notably foreign citizenship.

Table 1. Distribution of earnings.

	– Level –	– Differences –		– DiD –
	Treated	Treated	– Comparison	Treated
	<i>Pre-reform</i>	<i>Pre-reform</i>	<i>Post-reform</i>	<i>Post-reform - Pre-reform</i>
5th percentile	4,965	596	1,381	785
10th percentile	85,068	-654	6,666	7,320
25th percentile	243,990	3,817	6,888	3,071
50th percentile	358,769	4,082	7,004	2,922
75th percentile	476,324	9,336	9,787	451
90th percentile	652,850	22,057	17,388	-4,669
95th percentile	798,376	31,454	18,198	-13,256
Mean	383,897	9,942	10,176	234

Notes: Pre-reform cohorts are born 1967–1969 and post-reform cohorts are born 1973–1976. Treatment (comparison) municipalities are above (below) the median in child care coverage growth from 1976 to 1979. The earnings variable is defined as the child’s average earnings over the period 2006–2009 measured in NOK ($NOK/USD \approx 6$).

parents from the treatment municipalities as compared to the comparison municipalities (see Table A1 in the Appendix). This finding is reassuring since substantial changes over time in the differences in the observable characteristics of the two groups may suggest unobserved compositional changes, which would call our empirical strategy into question.

Because we use a DiD approach, it is not necessary that the child care expansion is unrelated to municipality characteristics. It is useful, however, to understand the determinants of the expansion across municipalities. The treatment and comparison municipalities are quite well spread out over Norway, covering urban and rural municipalities.¹³ In our baseline specification, five of the ten largest cities – by the number of children in our sample – are defined as treatment municipalities (Oslo, Bergen, Stavanger, Bærum and Fredrikstad), while the others are defined as comparison municipalities (Trondheim, Kristiansand, Tromsø, Skien and Drammen). Furthermore, there appears to be no substantial differences in terms of local government expenditure per capita, in total or on primary school in particular.¹⁴ This is most likely because of strict federal provisions for minimum standards of different local public services, and considerable ear-marked grants-in-aid from the central government. The same holds for local government income, consisting largely of grants-in-aid from the central government, local income taxes, and user fees. This comes as no surprise, as the federal government determines the tax rate and the tax base of the

¹³See Figure A1 reported in the Appendix.

¹⁴Table A2, reported in the Appendix, displays characteristics of the municipalities in 1976, in the treatment and comparison area. In addition, when examining the pre-reform trends in time-varying municipality characteristics (such as municipal expenditures, primary school expenditures, tax income, average education, labor force participation, family patterns), we find a good coherence between the treatment and the comparison municipalities. As expected, we find a good coherence between the time trends of the treatment and the comparison municipalities.

income tax. Also, the federal government used equalization transfers to redistribute income from rich to poor municipalities, such that local differences in revenues are largely offset. Interestingly, there are no noticeable differences in the share of female voters between the municipalities of the treatment and comparison area, nor is there significant disparity in the socialist shares of voters. This conforms well to the fact that there was broad bipartisan support for child care expansion in Norway in the 1970s. Further, we do not find any substantial differences in population size or the population shares of neither 0 to 6 year olds, nor females of fecund age, 19–35 or 36–55 years old.

There are, however, some notable differences between treatment and comparison municipalities. Most importantly, the ratio of child care coverage to employment rate of mothers of 3–6 year olds pre-reform, is substantially lower in treatment than in comparison municipalities. In treatment municipalities, there is on average more than four employed mothers for each child care place, while the same ratio is less than three-to-one in comparison municipalities. This is not surprising, since federal subsidy rates were higher for municipalities with low child care coverage prior to the reform, but also because local political pressure for expansion of subsidized care is likely to be stronger in areas where child care was severely rationed. Two of the variables indicating a rural municipality also have a small positive relationship with the child care expansion (average distance to zone center and ear marks per capita). This might be due to the discreteness of child care expansion; Establishing a typical child care institution increases the child care coverage rate more in smaller than in larger municipalities. In Norway, there was a very slow process of urbanization until the mid 1980s, which implies that rural status is likely to be more or less constant during the period we consider, and should, therefore, be controlled for in the DiD approach.¹⁵

5 Empirical results

5.1 *Distributional impacts*

Figure 5 graphs the QTE estimates from the RiF-DiD estimator. The figure suggests substantial heterogeneity in the child care effects. The point estimates are zero or positive for all percentiles until the 82nd percentile. The effect peaks at the 11th percentile, suggesting an impact of about 9,000 NOK from the universal child care program (or 50,000 NOK per child care place). Between the 15th and the 60th percentile estimates remain positive, at about half the size, but then fade out and turn negative in the upper part of

¹⁵We have also regressed the change in the municipality’s child care coverage between 1976 and 1979 on the characteristics of the municipalities. Consistent with the descriptive statistics, there is little evidence of systematic relationships between the child care expansion and the characteristics of the municipalities. A notable exception is the ratio of child care to maternal employment rate prior to the reform. In addition, we have examined the pre-reform trends in time-varying municipality characteristics.

the distribution.¹⁶

Taken together, the QTE estimates suggest that the introduction of subsidized care equalized the outcome distribution of children, in line with our simple model of parental investments. In order to quantify the equalizing effect of the child care reform, we compute the Gini coefficient in the counterfactual ($\mathcal{F}_1(y)_{RIF}$) and the actual ($F_1(y)$) earnings distribution of the post-reform cohorts from the treatment municipalities. We find that the child care reform reduced the Gini coefficient from .306 to .296. Put into perspective: this one point reduction in the Gini coefficient corresponds to introducing a one percent proportional tax on earnings and then redistributing the derived tax as equal sized amounts to the individuals (Aaberge, 1997). This indicates that universal child care has a small but non-negligible equalizing effect.

Another prediction from the simple model is that the benefits of subsidized care should be negatively correlated with family income (insofar child care quality is a normal good). To assess this prediction, Panel (a) of Figure 6 graphs the results from local linear regressions of childrens' earnings on (log) family income. We run the local linear regression separately for pre- and post-reform cohorts in the treatment and the comparison municipalities. We can then map out the DiD effect of the child care expansion by family income as the difference in earnings between the pre- and post-reform cohorts in the treatment municipalities as compared to the comparison municipalities. In Panel (b), we graph these DiD estimates by family income alongside the density of family income in treatment municipalities before the reform. We see that most of the gains in earnings associated with the universal child care program relate to children of low income parents, whereas upper-class children experience a loss in earnings because of the program.

The differential effects by parental income suggest that the child care reform reduced intergenerational income persistence. To quantify the effect on intergenerational income persistence, we extend the standard DiD-model with a full set of interactions between family income ($Faminc$) and dummy variables for post-reform cohort (P), treatment municipalities (T), and their interaction:

$$y_{it} = \gamma_t + \alpha_1 T_i + \alpha_2 T_i \cdot P_t + (\beta_0 + \beta_1 P_t + \beta_2 T_i + \beta_3 T_i \cdot P_t) Faminc_{it} + \epsilon_{it}$$

where y_{it} represents average earnings in 2006–2009 of child i born in year t , γ_t denotes the birth cohort fixed effect, and ϵ is the error term. Our estimates suggest that the child care reform reduced the intergenerational income elasticity by 2.5 percentage points.¹⁷ This

¹⁶We have estimated QTE at quantiles 96–99, but do not include them in the figures below because of little overlap in the uppermost part of the earnings distributions of the pre and post-reform cohorts.

¹⁷We estimate that $\hat{\beta}_3 = -.0239$ with $SE = 0.010$, while $(\hat{\beta}_0, \hat{\beta}_2) = (0.248, 0.0165)$ with $SE = (0.005, 0.008)$. If we define the intergenerational income elasticity as $(\partial y / \partial Faminc) \times (Faminc / y)$, then $\hat{\beta}_0 + \hat{\beta}_2$ translates into an intergenerational income elasticity of .275 among pre-reform cohorts in the treatment group, while $\hat{\beta}_3$ suggests a reduction of .0248 percentage points in this elasticity. Alternative specifications, using the log of family income and/or the log of child's earnings, give very similar estimates.




Figure 5. QTE estimates on earnings.

Notes: This figure graphs QTE estimates from the RIF-DID estimator, including a 90 % confidence interval based on a non-parametric bootstrap with 500 replications. Pre-reform cohorts are born 1967–1969 and post-reform cohorts are born 1973–1976. Treatment (comparison) municipalities are above (below) the median in child care coverage growth from 1976 to 1979. The earnings variable is defined as the child’s average earnings over the period 2006–2009 measured in NOK ($NOK/USD \approx 6$).

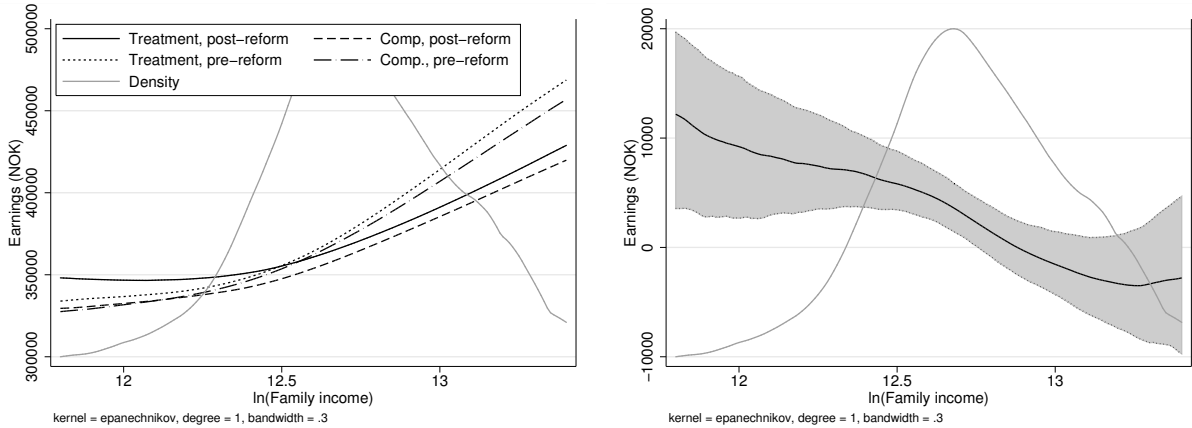
indicates that universal child care levels the playing field by reducing intergenerational income persistence.

5.2 Mean impacts

Table 2 shows the standard DiD estimates for the population of children as a whole and in subgroups by child and parental characteristics. The first row of Panel (a) shows that the overall mean impact misses a lot. While the child care program had a small and insignificant mean impact on earnings, there are sizable and significant effects in both the lower and the upper part of the earnings distribution.

In the other rows of Panel (a), we estimate subgroup mean impacts by gender and family income. As expected, the subgroup mean impacts by family income mirror the results produced by the local linear regressions. While children from high-income families suffer a mean loss of about 8,000 NOK (or 46,000 NOK per child care place), children from low income families experience a gain of over 9,000 NOK (or 53,000 NOK per child care place). In contrast, there is no evidence of differential effects by child gender.¹⁸

¹⁸Because child gender is as good as random, the subgroup estimates by gender are not confounded by differences in family income.



(a) Earnings levels

(b) DiD estimates

Figure 6. Child earnings and reform effects by family income.

Notes: Panel (a) shows predictions from a local linear regression of earnings on the natural logarithm of family income (at age 3–6) by treatment status for children born in 1967–1969 (pre-reform) and 1973–1976 (post-reform), using an epanechnikov kernel with bandwidth = 0.3. Birth cohort effects were removed from earnings prior to estimation, and the residuals were rescaled to fit the level of earnings in the 1967-cohort. Panel (b) shows the density of log family income in the treatment group pre-reform, and DiD-estimates constructed from the graphs in panel (a), including a 90 % confidence interval based on a non-parametric bootstrap with 500 replications. The earnings variable is defined as the child’s average earnings over the period 2006–2009 measured in NOK ($NOK/USD \approx 6$).

5.3 Mediating relationships

To clarify the nature of the relationship between the child care reform and the earnings of exposed children as adults, we examine the mediating role of educational attainment and cognitive test scores.

Panel (b) of Table 2 shows the effect of the universal child care program on years of schooling. While there is a positive mean impact on years of schooling, the effect is largely driven by children of low income parents. There is no change in the educational attainment of upper-class children. In Panel (a) of Figure 7, we report QTE estimates from the RiF-DiD estimator for years of schooling. The estimates reveal that most of the growth in educational attainment stems from a substantial increase in the likelihood of completing high school and attending college. This finding aligns well with the differential effect on years of schooling by family income, as there is little room for improvement in high school completion and college attendance among upper-class children.¹⁹

Panel (b) of Figure 7 displays the RIF-DID estimates on cognitive tests scores. Recall that the test scores are normalized to mean 100 and standard deviation of 15. We see that the QTE estimates on test scores are close to zero, and they are sufficiently precise to rule out any economically relevant effect. The combination of substantial effects on

¹⁹For the pre-reform cohorts from treatment municipalities, the high school completion rate is 65.3 % for children of low-income parents and 86.2 % for children of high-income parents.

Table 2. Estimated mean impacts on earnings and education.

	A. Earnings		B. Years of schooling		N
	Coef (SE)	Mean	Coef (SE)	Mean	
Overall	333 (1,596)	361,860	0.074 (0.017)	12.83	498,947
Child gender					
Boys	-978 (2,504)	440,020	0.084 (0.024)	12.62	253,677
Girls	631 (1,616)	281,020	0.066 (0.025)	13.06	245,270
Family income					
High	-8,095 (4,575)	410,678	0.018 (0.041)	13.82	99,761
Mid	-1,785 (3,403)	355,887	0.081 (0.037)	12.69	99,793
Low	9,287 (3,331)	325,523	0.241 (0.039)	12.20	99,801


Notes: This table reports mean impacts estimates from the standard DiD estimator. Pre-reform cohorts are born 1967–1969 and post-reform cohorts are born 1973–1976. Treatment (comparison) municipalities are above (below) the median in child care coverage growth from 1976 to 1979. The earnings variable is defined as the child’s average earnings over the period 2006–2009 measured in NOK ($NOK/USD \approx 6$). Standard errors are in parentheses.

earnings and educational attainment and no impact on cognitive test scores points to the importance of non-cognitive skill development. This conforms with evidence from targeted programs showing that early interventions have large returns in the labor market because of improvements in non-cognitive skills, even though cognitive test score advantages for children with preschool experiences tend to fade out quickly as they go through school (see e.g. Heckman and Masterov, 2007). Unfortunately, we do not have information on non-cognitive skills, and therefore cannot assess this conjecture.

5.4 Sensitivity analysis

This section reports results from a battery of specification tests, demonstrating that the effect of child care on the distribution of children’s adult earnings are similar across specifications.

Identifying assumption. The identifying assumption underlying the RIF-DiD estimator is that the change in population shares from before to after treatment around a given level of earnings would be the same in the treatment group as in the comparison group, in the absence of treatment. A concern is that the QTE estimates may reflect differential time trends between the treatment and comparison municipalities in the population shares, rather than true policy impacts. For instance, the QTE estimates would be biased if the time trend in earnings differs by, say, parent’s education, while there are differential changes over time in parental education between treatment and comparison municipalities.



(a) Years of schooling

(b) Cognitive test scores (males only)

Figure 7. Estimated impact on years of schooling and cognitive test scores.


Notes: This figure graphs QTE estimates from the RIF-DID estimator, including a 90 % confidence interval based on a non-parametric bootstrap with 500 replications. Pre-reform cohorts are born 1967–1969 and post-reform cohorts are born 1973–1976. Treatment (comparison) municipalities are above (below) the median in child care coverage growth from 1976 to 1979. The cognitive test scores are measured at age 18–19, and are only available for males. The raw scores are standardized into a nine point scale that has a discrete approximation to a normal distribution, a mean of 100, and a standard deviation of 15.

To address such concerns for omitted variables bias, we first include a set of controls for child and parental characteristics, and then also add municipality-specific fixed effects. Panels (b) and (c) of Figure A2 reveal that QTE estimates change little across these specifications.²⁰

Although the DiD approach controls for the direct effects of pre-determined factors of the municipalities, like differences in rationing of subsidized child care prior to the reform, bias can arise from the determinants of the child care expansion being systematically related to underlying trends in child outcome. To address this concern, we follow Duflo (2001) in interacting cohort fixed effects with observable pre-reform municipality characteristics. This allows for different underlying trends in earnings growth depending on the pre-reform characteristics of the municipality. Panel (d) of Figure A2 shows that these QTE estimates conform well to the results from the baseline specification.

Alternative treatment definition. In our main analysis, we define treatment by ordering municipalities according to the increase in child care coverage rate in the period 1976–1979, and then separating them at the median. To evaluate the sensitivity of our estimates to the exact definition of the treatment, Figure A3 reports QTE-estimates for alternative definitions. In Panel (b), treated municipalities are over the 67th percentile and comparison

²⁰We include linear and quadratic terms for mother’s and father’s age, age at first birth, income and education at age 2; the number of older siblings; as well as dummy variables for gender and immigrant status. Rather than dropping a small number observations where information on family background characteristics is missing, we include separate categories for missing values.



(a) CIC-estimates

(b) QDiD-estimates

Figure 8. QTE estimates on earnings from alternative non-linear DiD estimators.

Notes: This figure graphs QTE estimates, including a 90 % confidence interval based on a non-parametric bootstrap with 500 replications. Pre-reform cohorts are born 1967–1969 and post-reform cohorts are born 1973–1976. Treatment (comparison) municipalities are above (below) the median in child care coverage growth from 1976 to 1979. The earnings variable is defined as the child’s average earnings over the period 2006–2009 measured in NOK ($NOK/USD \approx 6$). Panel (a) reports estimates using the CiC-estimator, while panel (b) reports estimates using the QDiD-estimator.

municipalities are below the 33rd percentile in child care coverage growth. In Panels (c) and (d), we again divide the sample at the median, but change the expansion period to 1977–1979 and 1976–1978, respectively. Panel (a) compares the estimates, revealing a strong similarity in the QTE estimates across the different definitions.

Alternative non-linear DiD estimators. To evaluate the sensitivity of the QTE estimates to the assumption about the counterfactual earnings distribution, we also implement two alternative non-linear DiD estimators. Similar to the RIF-DiD procedure, quantile DiD (QDiD) and Changes-in-Changes (CiC) use the observed changes in the earnings distribution of children in the comparison group to estimate the counterfactual earnings distribution of treated children post-reform, in the absence of treatment. However, the identifying assumptions differ. Below, we briefly describe the QDiD and CiC estimators; a detailed discussion can be found in Athey and Imbens (2006).²¹

The CiC-estimator is a three-step procedure. For every value y on the support of $F_1(y)$, first find the quantile of y in G_1 , then find the earnings level at that quantile in G_0 , and finally find the quantile of that earnings level in F_0 . Specifically, the CiC-estimator assumes that the counterfactual earnings distribution is

$$\mathcal{F}_1(y)_{CIC} = F_0(G_0^{-1}(G_1(y)))$$

²¹We use RIF-DiD as our baseline estimator both because of the considerable gains in computation time, and because of the relative ease with which we can include covariates (see Firpo, 2007).

In Figure 4b, the QTE estimate at y' is therefore $(a - e) - (c - f)$. In comparison, the QDiD-estimator takes the counterfactual earnings distribution to be

$$\mathcal{F}_1^{-1}(F_0(y))_{QDiD} = y + [G_1^{-1}(F_0(y)) - G_0^{-1}(F_0(y))]$$

In Figure 4b, the QDiD-estimate at y' is therefore $(a - e) - (g - h)$. While the CiC-estimator assumes common growth in earnings between the treatment and comparison group in the absence of the reform, the QDiD-estimator assumes a common trend in absolute earnings.

Figure 8 shows the sensitivity of the QTE estimates to the choice of non-linear DiD estimator. It is reassuring to find that the QTE estimates are quite similar to those we found with the RiF-DiD estimator.

6 Conclusion

Many developed countries currently consider a move towards a universal child care program. The challenge in assessing the case for universal child care programs is that the evidence base is scarce. We analyzed the staged expansion of subsidized, universally accessible child care in Norway. Our approach differs from existing literature which estimates mean impacts of universal child care programs; in contrast, we use non-linear difference-in-differences methods to estimate how the child care expansion affected the earnings distribution of exposed children as adults. We complement these estimates with local linear regressions of the child care effects by family income. Our findings suggest that the effects of child care vary systematically across the earnings distribution, that mean impact misses a lot, and that children of low income parents seem to be the primary beneficiaries of subsidized child care. To help understand the differential effects on earnings, we examined how the child care expansion affected the educational trajectories and cognitive test scores of the exposed children.

Taken together, our results have important implications for a small but rapidly growing literature on universal or large-scale child care programs. Our study points out that the effects of subsidized child care may be both more varied and extensive than what previous studies suggest. Our findings highlight that the effects of child care vary systematically across the outcome distribution, and that children of low income parents seem to be the primary beneficiaries. This is especially important when considering the case for universal child care programs, since the benefits of providing subsidized child care to middle and upper-class children are unlikely to exceed the costs.

References

- AABERGE, R. (1997): “Interpretation of changes in rank-dependent measures of inequality,” *Economics Letters*, 55(2), 215–219.
- ALMOND, D., AND J. CURRIE (2011): *Human Capital Development before Age Five* vol. 4 of *Handbook of Labor Economics*, chap. 15, pp. 1315–1486. Elsevier.
- ATHEY, S., AND G. W. IMBENS (2006): “Identification and Inference in Nonlinear Difference-in-Differences Models,” *Econometrica*, 74(2), 431–497.
- BAKER, M. (2011): “Innis Lecture: Universal early childhood interventions: what is the evidence base?,” *Canadian Journal of Economics*, 44(4), 1069–1105.
- BAKER, M., J. GRUBER, AND K. MILLIGAN (2008): “Universal Child Care, Maternal Labor Supply, and Family Well-Being,” *Journal of Political Economy*, 116(4), 709–745.
- BITLER, M. P., J. B. GELBACH, AND H. W. HOYNES (2006): “What Mean Impacts Miss: Distributional Effects of Welfare Reform Experiments,” *American Economic Review*, 96(4), 988–1012.
- BLACK, S. E., P. J. DEVEREUX, K. V. LOKEN, AND K. G. SALVANES (2012): “Care or Cash? The Effect of Child Care Subsidies on Student Performance,” IZA Discussion Papers 6541, Institute for the Study of Labor (IZA).
- CRYER, D., AND M. BURCHINAL (1997): “Parents as child care consumers,” *Early Childhood Research Quarterly*, 12(1), 35 – 58.
- DATTA GUPTA, N., AND M. SIMONSEN (2010): “Non-cognitive child outcomes and universal high quality child care,” *Journal of Public Economics*, 94(1-2), 30–43.
- DUFLO, E. (2001): “Schooling and Labor Market Consequences of School Construction in Indonesia: Evidence from an Unusual Policy Experiment,” *American Economic Review*, 91(4), 795–813.
- EIDE, M., N. ØYEN, R. SKJÆRVEN, S. NILSEN, T. BJERKEDAL, AND G. TELL (2005): “Size at Birth and Gestational Age as Predictors of Adult Height and Weight,” *Epidemiology*, 16, 175.
- EU (2002): “Presidency Conclusions,” Barcelona European Council 15 and 16 March 2002, Barcelona.
- FIRPO, S. (2007): “Efficient Semiparametric Estimation of Quantile Treatment Effects,” *Econometrica*, 75(1), 259–276.

- FIRPO, S., N. M. FORTIN, AND T. LEMIEUX (2009): “Unconditional Quantile Regressions,” *Econometrica*, 77, 953–973.
- HAVNES, T., AND M. MOGSTAD (2011): “No Child Left Behind: Subsidized Child Care and Children’s Long-Run Outcomes,” *American Economic Journal: Economic Policy*, 3(2), 97–129.
- HECKMAN, J. J., AND D. V. MASTEROV (2007): “The Productivity Argument for Investing in Young Children,” NBER Working Papers 13016, National Bureau of Economic Research, Inc.
- HECKMAN, J. J., J. STIXRUD, AND S. URZUA (2006): “The Effects of Cognitive and Noncognitive Abilities on Labor Market Outcomes and Social Behavior,” *Journal of Labor Economics*, 24(3), 411–482.
- HERBST, C. M., AND E. TEKIN (2008): “Child Care Subsidies and Child Development,” IZA Discussion Papers 3836, Institute for the Study of Labor (IZA).
- LEFEBVRE, P., P. MERRIGAN, AND M. VERSTRAETE (2008): “Childcare Policy and Cognitive Outcomes of Children: Results from a Large Scale Quasi-Experiment on Universal Childcare in Canada,” Cahiers de recherche 0823, CIRPEE.
- RUHM, C. J., AND J. WALDFOGEL (2012): “Long-Term Effects of Early Childhood Care and Education,” *Forthcoming in Nordic Economic Policy Review*.

A Additional material (NOT FOR PUBLICATION)

A.1 Background and descriptive statistics

Table A1. Descriptive statistics of the treatment and control group.

	– Level –		– Differences –	
	Treated (pre-reform)		Treated – Comparison	
	<i>Mean</i>	<i>SD</i>	<i>Pre-reform</i>	<i>Post-reform</i>
Male	0.51	0.50	-0.002	0.002
No. of older siblings	2.13	1.06	-0.082	-0.112
Family income	353,402	161,834	18,938	19,562
Immigrant	0.06	0.23	0.011	0.016
Mother's income	24,916	89,322	3,009	6,723
– years of schooling	9.54	2.58	0.240	0.234
– age at first birth	23.33	3.84	0.566	0.647
Father's income	328,486	134,382	15,928	12,839
– years of schooling	10.23	3.14	0.333	0.332
– age at first birth	26.46	4.82	0.488	0.546
Children – treatment			77,906	84,041
Children – comparison			87,825	91,398

Notes: Pre-reform cohorts are born 1967–1969 and post-reform cohorts are born 1973–1976. Treatment (comparison) municipalities are above (below) the median in child care coverage growth from 1976 to 1979.

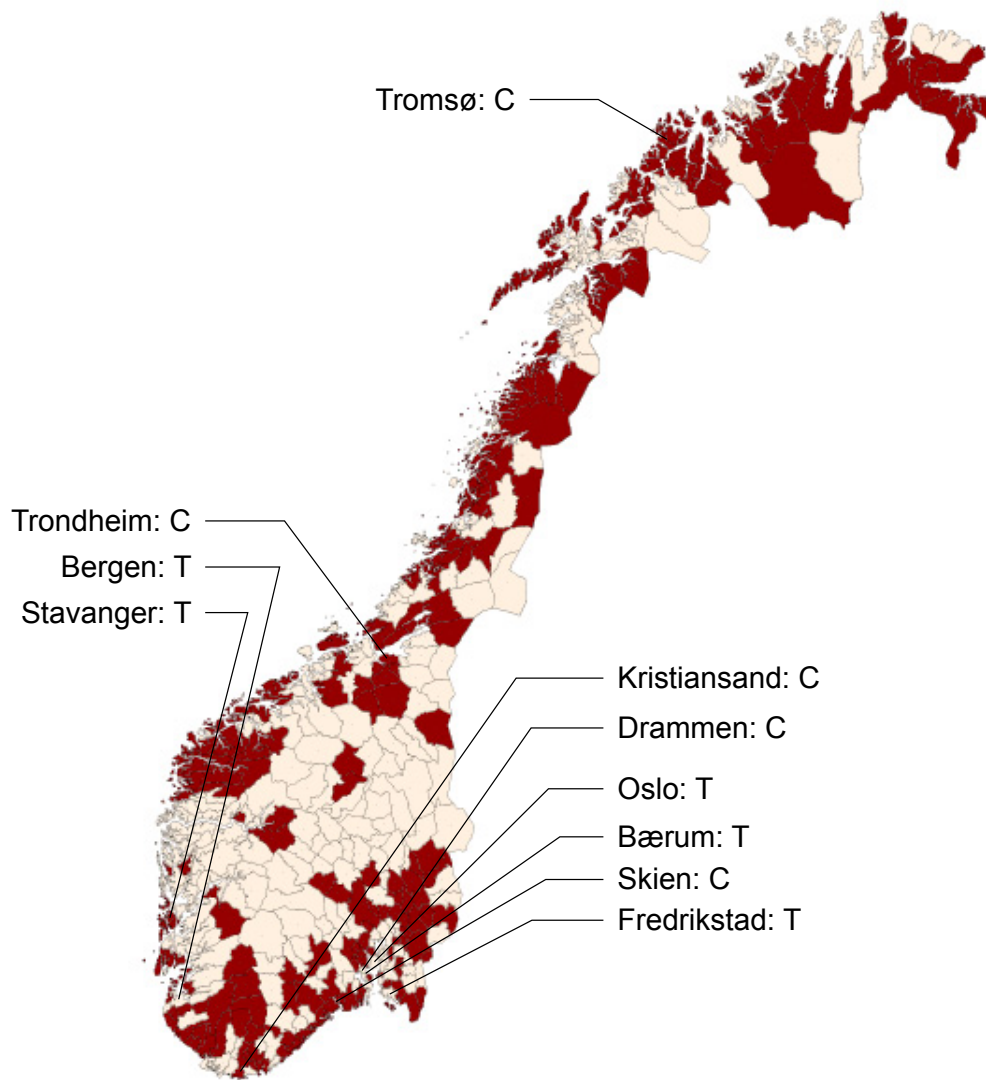


Figure A1. Treatment (light) and comparison (red) municipalities, cutoff at median growth in child care from 1976 to 1979.

Notes: Treatment (comparison) municipalities are above (below) the median in child care coverage growth from 1976 to 1979.

Table A2. Descriptive statistics for treatment and comparison municipalities in 1976.

	Treatment		Comparison	
Child care/maternal employment rate	0.2471	[0.4596]	0.3542	[0.6213]
Child care coverage rate	0.0534	[0.0899]	0.0695	[0.0968]
Years of education, males	9.2256	[0.5514]	9.2174	[0.4699]
–, females	8.8198	[0.3820]	8.7672	[0.3313]
Earnings, males	0.3917	[0.0762]	0.4081	[0.0734]
–, females	0.108	[0.0349]	0.1121	[0.0321]
Labor force part., males	0.8324	[0.0591]	0.8367	[0.0665]
–, females	0.2919	[0.0844]	0.2997	[0.0813]
–, mothers of 3–6 year olds	0.1903	[0.0753]	0.1953	[0.0710]
Expenditure (2006-NOK/capita)				
Total	6237.73	[1891.83]	5913.18	[1492.08]
Primary school	1571.57	[698.10]	1452.10	[590.40]
Revenue (2006-NOK/capita)				
Ear marks, total	3702.01	[1416.68]	3552.19	[1131.72]
–, primary school	570.57	[216.52]	568.43	[234.78]
Fees, total	810.23	[585.20]	682.63	[414.70]
–, primary school	5.59	[8.97]	6.31	[11.70]
Taxes	2463.50	[682.50]	2464.09	[745.88]
Geography				
In densely populated areas	0.4049	[0.2915]	0.4827	[0.2979]
Ave. distance to zone center	0.8876	[0.7789]	0.7732	[0.6788]
– to neighboring center	3.7768	[2.6130]	3.4297	[2.8039]
Population				
Total	9846	[36400]	9476	[13267]
Married	0.4664	[0.0277]	0.4618	[0.0346]
Divorced	0.0144	[0.0081]	0.0153	[0.0080]
Immigrant	0.0098	[0.0096]	0.0095	[0.0086]
0 to 6 years old	0.1077	[0.0177]	0.1141	[0.0177]
7 to 10 years old	0.0673	[0.0099]	0.0708	[0.0097]
11 to 18 years old	0.1293	[0.0127]	0.1314	[0.0133]
Females: 19 to 35 years old	0.1021	[0.0187]	0.1082	[0.0170]
–: 36 to 55 years old	0.1027	[0.0101]	0.1019	[0.0095]

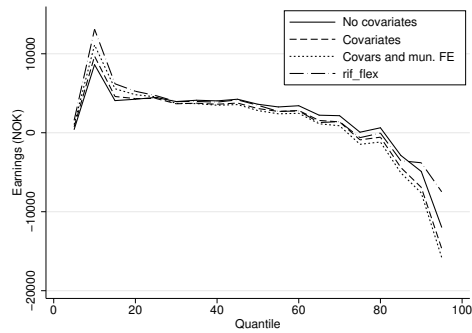
Males: 19 to 35 years old	0.1175	[0.0152]	0.1227	[0.0137]
–: 36 to 55 years old	0.1096	[0.0091]	0.1077	[0.0092]

Politics


Registered voters	6480	[26654]	5863	[8488]
–, female	0.4896	[0.0169]	0.4928	[0.0167]
Election participation	0.7243	[0.0587]	0.7093	[0.0563]
–, females	0.7094	[0.0666]	0.6962	[0.0632]
Female elected representatives	0.1521	[0.0807]	0.1394	[0.0635]
Socialist vote share	0.3864	[0.1654]	0.4031	[0.1651]
Socialist mayor	0.314	[0.4652]	0.3671	[0.4832]
Female mayor	0.0097	[0.0981]	0.0145	[0.1198]

Notes: Columns 1–4 report means and standard deviations across treatment and comparison municipalities. Earnings denotes pensionable income in NOK 100,000. Socialist parties are defined as RV, SV and DNA. Densely populated areas are contiguous zones of at least 200 people where the distance between houses is generally less than 50 meters (400 meters if separated by e.g. parks, rivers or industrial zones). Average distance to zone center is the mean predicted travel distance in km from a citizen’s home to the most populous census area in a contiguous zone of more than 2,000 people within the municipality. Distance to neighboring center is the mean travel distance from the center of a census area to the closest center in another census area within the same economic zone. Treatment (comparison) municipalities are above (below) the median in child care coverage growth from 1976 to 1979. Standard deviations are in brackets.


A.2 Robustness analysis




(a) Comparison



(b) Covariates




(c) Municipality FE




(d) Flexible trends

Figure A2. QTE estimates on earnings with alternative sets of controls.


Notes: This figure graphs QTE estimates from the RIF-DID estimator, including a 90 % confidence interval based on a non-parametric bootstrap with 500 replications. The earnings variable is defined as the child's average earnings over the period 2006–2009 measured in NOK ($NOK/USD \approx 6$). Panel (a) compares results across the alternative specifications. Panel (b) includes the following set of covariates: linear and quadratic terms for mother's and father's age, age at first birth, income and education at age 2; the number of older siblings; and dummy variables for gender and immigrant status. Panel (c) includes these covariates and adds municipality-specific fixed effects. Panel (d) also allows for differential inter-cohort time trends across municipalities by adding interaction terms between cohort dummies and the municipal level of female education, female labor supply, primary school expenditures, and the child care coverage rate in 1976. Pre-reform cohorts are born 1967–1969, and post-reform cohorts are born 1973–1976. Treatment (comparison) municipalities are above (below) the median in child care coverage growth from 1976 to 1979.




(a) Comparison



(b) 1976-1979, 33rd vs 67th



(c) 1977-1979, 50th vs 50th



(d) 1976-1978, 50th vs 50th

Figure A3. QTE estimates on earnings with alternative treatment definitions.

Notes: This figure graphs QTE estimates from the RIF-DID estimator, including a 90 % confidence interval based on a non-parametric bootstrap with 500 replications. The earnings variable is defined as the child's average earnings over the period 2006-2009 measured in NOK ($NOK/USD \approx 6$). Panel (a) compares results across the alternative treatment definitions. In Panel (b), the expansion period is 1976-1979, and the treatment (comparison) group are above the 67th (below the 33rd) percentile in child care growth over the expansion period. In Panel (c), the expansion period is 1976-1978, and the treatment (comparison) group are above (below) the median in child care growth over the expansion period. In Panel (d), the expansion period is 1977-1979, and the treatment (comparison) group are above (below) the median in child care growth over the expansion period.