Early Childcare and Cognitive Development: Evidence from an Assignment Lottery

Nina Drange, *Statistics Norway*

Tarjei Havnes, *University of Oslo*

Young children are thought to be vulnerable to separation from their primary caregiver. This raises concern about whether early childcare enrollment may harm child development. We use childcare assignment lotteries to estimate the effect of enrollment at age 1–2 on cognitive development in Norway. Estimates show significant gains in language and mathematics at age 6–7 and a substantial drop in scores below publicly set thresholds for low performance. Across subsamples, we find a pattern of stronger effects on underperforming groups. We find little support for childcare quality or family income as drivers of our results.

I. Introduction

Childcare enrollment of toddlers has increased in many countries over the past decade. In 2010, the enrollment of children below age 2 stood at 43% in

This is a substantially revised version of the paper circulated as "Child Care Before Age Two and the Development of Language and Numeracy: Evidence from a Lottery." Thanks to the Oslo municipality for generously providing data, institutional detail, and feedback on the project, in particular Eli Aspelund, Thomas Bang, and Ragnhild Walberg at the Oslo City Council Department of Children and Education. The project received financial support from the Norwegian Research Council (grants 212305 and 236947). The project is also part of the research activities at the Center for the Study of Equality, Social Organization, and Performance (ESOP) at the Department of Economics, University of Oslo. ESOP is supported by the Research Council

[*Journal of Labor Economics*, 2019, vol. 37, no. 2]

0734-306X/2019/3702-0008$10.00
Submitted October 26, 2016; Accepted December 22, 2017; Electronically published February 12, 2019

the United States and at 33% in Organization for Economic Cooperation and Development (OECD) countries, up from 20% in 2003. In several countries, enrollment is much higher, with rates above 50% in countries like Denmark, Korea, the Netherlands, and Norway.[1] At the same time, there is concern among both researchers and policy makers that early separation from the primary caregiver, typically the mother, may cause stress and anxiety in the child, with potentially adverse effects on children's development (Bowlby 1969; Mercer 2006).[2] Yet evidence on how childcare affects the development of toddlers is largely missing. We are aware of only three other studies that estimate how childcare affects the development of toddlers while controlling for selection. Felfe and Lalive (2014) estimate marginal treatment effects of childcare attendance before age 3 using county-level variation in childcare coverage rates in West Germany. They find positive impacts of childcare on the youngest children, boys, and children from families of low socioeconomic status (SES). A recent study from Italy finds negative effects of early childcare for girls (Fort, Ichino, and Zanella 2016). For the United States, Herbst (2013) find negative effects of the summer dip in childcare participation on early cognitive skills measured at 9 and 24 months.[3] The relative lack of plausible evidence is worrying for policy makers, because programs are often heavily subsidized, but it is also worrying for parents, who need to decide whether and when to enroll their children in childcare.

In this paper, we provide first evidence of the impact of childcare enrollment for children age 1–2 years old (henceforth, "toddlers") on their cognitive performance in language and mathematics at age 7. Determining the impact in this age group is of key importance, both because enrollment rates are increasing at a strong rate in many countries and because children are thought to be particularly vulnerable during this period. For identification, we exploit random assignment to childcare used by the city government in the Norwegian capital Oslo in order to allocate offers of childcare places when childcare institutions are oversubscribed, similar to the strategy used by Abdulkadiroglu et al. (2011). This should provide variation in childcare enrollment that is as good as random. Indeed, we document that the alloca-

of Norway (grant 179552). Contact the corresponding author, Nina Drange, at nina.drange@ssb.no. Information concerning access to the data used in this paper is available as supplemental material online.

[1] Source: OECD Family database.

[2] Bowlby (1969) defines the attachment phase as the period when the child is from 6–8 months to 24 months old. The age period we study in this paper is largely around 12–24 months.

[3] A related strand of literature may indirectly reflect the effects of childcare attendance in looking at the effect of parental leave policies on child outcomes (e.g., Liu and Skans 2010; Dustmann and Schønberg 2012; Carneiro, Løken, and Salvanes 2015). The alternative to parental care in most of these studies is, however, likely to be informal and not formal sources of care.

tion mechanism generates balance in observable characteristics, supporting our empirical approach.

Armed with unique data on all applicants, offers, and enrollment as well as performance tests in language and mathematics at age 7, we consider the impact of lottery offers on cohorts born from 2004 to 2006. There was substantial oversubscription of toddlers to childcare institutions in Oslo among these cohorts. Our first results show that children who randomly receive an offer of public childcare perform better on both the language test (0.16 standard deviations) and the mathematics test (0.11 standard deviations). Because the distribution of the test score we consider is skewed, we assess the economic significance of our estimates by comparing them to gaps in test scores between well-defined subgroups of the population. For instance, the estimated impact on language performance of receiving an offer corresponds to about the gap between children from high- and low-income families or to about half the gap between children from high- and low-educated parents. In mathematics, the effect corresponds to about a quarter of the gap between children of low- and high-educated parents or about one-third of the gap between children from high- and low-income families. Although these effects are remarkably large, it should be noted that the tests we consider are designed to capture variation at the lower end of the distribution and may understate skills at the top.

Because test scores have no natural cardinal scale (Cunha and Heckman 2008), the previous estimates may be hard to interpret. As an alternative outcome, we therefore consider scoring below thresholds for low performance set by the national government. These thresholds are set in order to flag children for follow-ups and to be evaluated for additional resources. The estimates confirm that childcare offers in the assignment lottery improves outcomes of children, with substantial drops in the probability of performing below the thresholds. This could suggest that childcare may particularly improve the development of low-performing children.

To investigate further the potential for childcare to enhance social mobility, we estimate the impact of receiving a lottery offer on the performance of children in subsamples depending on their gender and family background. We find a striking pattern in the point estimates, with stronger effects in the group that tends to underperform on the test (as observed among children without lottery offers). For instance, children of low-educated parents are estimated to improve their performance in language by about 24% of a standard deviation, while the performance of children from high-educated parents is estimated to increase by a comparably modest and statistically insignificant 8% of a standard deviation. This suggests that if both receive a childcare offer, then the gap between the two groups is halved. Similarly, children from low-income families are estimated to improve their performance in both language and mathematics by about 26% of a standard deviation, compared with modest or no effects for children from high-income families.

The heterogeneity across subsamples motivates taking a closer look at the distributional effects of the lottery offer. To this end, we rerun the estimation in our main sample using as the dependent variable indicators for performing above 19 quantiles of the test score distribution. In line with the above, we find that the lottery offer generates substantial improvement in performance at the lower end of the distribution, while effects from the middle up are modest and statistically insignificant. Again, one should keep in mind that the tests we consider are not well designed to capture variation at the top of the distribution. We cannot rule out, therefore, that lack of effects at the top is an artifact of our outcome measure. These patterns are in line, however, with a story where educated and high-income families provide a more stimulating environment, at home or in informal care, compared with more disadvantaged families (as suggested, e.g., by the seminal study of Hart and Risley [1995]). The childcare center, if it provides a more homogeneous environment to children, may then give bigger gains to children from disadvantaged families. This is also in line with the evidence in Havnes and Mogstad (2015).

Having established the positive effects on children from receiving a lottery offer, we investigate the potential mechanisms that could be driving the effects. We consider the prime candidate to be childcare starting age. While children that receive an offer first attend childcare at about 15 months of age on average, children who randomly do not receive an offer first attend at about 19 months of age on average. Among children delayed, about two-thirds start 1 year later or more. When we estimate the impact of childcare starting age using the lottery offer as an instrumental variable (IV), we find that starting childcare 1 month earlier causes a statistically significant improvement in test scores of 0.05 and 0.03 standard deviations in language and mathematics, respectively. In support of childcare starting age as a likely mechanism, we show that differences in test performance between children who get and do not get a lottery offer are smaller when we focus on children who started childcare around the same age.

As discussed, lottery offers are a strong predictor of childcare enrollment in the year of application in our sample. It is still true, however, that many children who do not receive an offer also are not delayed in starting childcare and that almost all children do enroll in childcare before school start. This raises the issue of whether lottery offers may affect children's skill development by changing the quality of the childcare institution in which the child enrolls. To investigate the role of the alternative mode of care for children who are in childcare, we therefore study indicators of the quality of care based on structural quality indicators, staff characteristics, and peer characteristics. Estimates show that children who receive lottery offers attended centers with somewhat different characteristics, as expected. At the same time, there are few indications that children with offers attended higher-quality institutions. On the contrary, children with lottery offers seem to attend institutions with slightly lower structural quality and somewhat worse-performing peers. This suggests that differences in childcare quality are unlikely to be driving the observed impacts on test performance.

Finally, we consider the role of the alternative mode of care for children who are not in childcare. We start by considering survey data on stated demand and actual use for the population of parents with toddlers. While about 70% of parents state demand for childcare, either full time or part time, only 33% actually have their children enrolled in childcare. In comparison, while 56% of parents say that they care for their children themselves, only 17% actually prefer to do so. This suggests that parental care is the dominant alternative for Norwegian toddlers in general. To consider parental care as the counterfactual in our particular sample, we next estimate the impact of receiving a lottery offer on parents' labor supply. The estimates show a modest increase in our measure of full-time equivalent labor participation among mothers but little impact on the employment margin or on the labor supply of the father.[4] This suggests, on the one hand, that an increase in family income cannot be driving our results and, on the other, that access to informal sources of care may be more prevalent in our Oslo sample of early applicants to childcare than in the general Norwegian population.

Our results on how childcare affects the development of toddlers complement the growing recent literature on how childcare institutions affect the development of preschool age children.[5] The literature is divided in two distinct branches, one focused on targeted programs and another focused on universal programs available to the general population. While studies of targeted programs often find positive effects,[6] the literature on universal pro-

[4] These results are roughly in line with previous findings for preschool-age children in Norway (Havnes and Mogstad 2011a). Evidence from other countries is mixed. In a survey of the early literature, Blau and Currie (2006) report elasticities of maternal employment with respect to the price of childcare ranging from 0 to −1. More recently, Baker, Gruber, and Milligan (2008) find a positive effect on maternal labor supply following the introduction of heavily subsidized universally available childcare in Quebec. Meanwhile, Lundin, Mork, and Ockert (2008) find no such effect when studying a childcare reform that capped childcare prices in Sweden. See also Schlosser (2005), Cascio (2009), Havnes and Mogstad (2011b), Lefebvre and Merrigan (2008a), and Berlinski and Galiani (2007). For a review of the literature, see Blau and Currie (2006).

[5] For recent reviews, see Almond and Currie (2010), Ruhm and Waldfogel (2012), or Baker (2011).

[6] The Perry Preschool and Abecedarian programs are examples of targeted randomized programs (for surveys of the literature, see Barnett 1995; Karoly, Kilburn, and Cannon 2005), while the US Head Start program provides an example of a targeted nonrandomized program (for a review of the findings, see, e.g., McKey et al. 1985; Currie 2001). While the picture is somewhat mixed, the most robust evidence on Head Start tends to show positive effects on long-run outcomes, such as high school dropout rate, college attendance, and crime (Currie and Thomas 1995; Garces, Thomas, and Currie 2002; Ludwig and Miller 2007; Deming 2009).

grams is smaller and findings are mixed.[7] Perhaps as a consequence, the discussion on childcare policies has been based largely on the targeted literature and descriptive evidence, even when the policies discussed are universal. In contrast, we study the impact of a universally available program among applicants that are explicitly not prioritized in childcare.

The paper proceeds as follows. We first discuss the institutional background in Section II. Section III presents and discusses our empirical approach before Section IV describes our data. Section V presents our main results, including discussions on heterogeneous impacts and mechanisms, while Section VI concludes.

II. Institutional Background

In this section, we provide brief institutional background about the care of young children in Norway, with a focus on the childcare sector.

A. Parental Leave

In 2005, Norwegian parents were entitled to 43 weeks of parental leave with full wage compensation (alternatively 53 weeks with 80% compensation).[8] This was expanded to 44 (54) weeks in 2006. Parents are further entitled to 1 year each of unpaid leave in immediate continuation of regular parental leave. In practice, most parents can thus stay at home with their newborn for about a year.

B. Structure and Content of Childcare in Norway

To help interpret our results, we must understand the type of care we are studying. Childcare in Norway is strictly regulated, with provisions on staff qualifications, number of children per teacher, size of play area, and educational orientation. Institutions are run by an educated preschool teacher responsible for day-to-day management and educational content. The preschool teacher education is a 4-year college degree, including supervised practice in a formal childcare institution. The head teacher is responsible for planning, observation,

[7] Several studies from Canada show a negative impact on a variety of child outcomes (Baker, Gruber, and Milligan 2008; Lefebvre and Merrigan 2008b; DeCicca and Smith 2013), while Cascio (2009) and Gupta and Simonsen (2010) find essentially no impact from childcare programs in the United States and Denmark, respectively. In contrast, positive impacts on a number of outcomes are found from childcare programs in several countries, including the United States (Fitzpatrick 2008), Uruguay (Berlinski, Galiani, and Manacorda 2008), Norway (Havnes and Mogstad 2011b, 2015), Germany (Dustmann, Raute, and Schønberg 2013; Felfe and Lalive 2014), and Spain (Felfe, Nollenberger, and Rodriguez-Planas 2015).

[8] This entitlement is conditional on maternal employment during at least six of the 10 months before the birth. About 85% of new mothers satisfy this requirement (NOU 2012:15). Remaining parents are entitled to unpaid parental leave with employment protection and receive a one-time payment of about 35,000 NOK (4,375 USD; 1 USD ≈ 8 NOK).

Table 1
Institutional Characteristics

	Mean	SD	*N*
Teacher/children	.086	.046	499
Number of adults	9.82	8.65	662
Adult/children	.34	.186	499
Minority share	.158	.207	499
Number of children 0–2	13.64	12.71	662

Source.—Childcare register, Statistics Norway.

collaboration, and evaluation of all activities. The head teacher is also responsible for communication and collaboration with parents and local authorities, including health centers and child welfare services when necessary. Childcare regulations specify that there should be at least one educated preschool teacher per 10 children below the age of 3. Each teacher typically works with two assistants. There is no educational requirements for the assistants.

In Oslo, about 60% of childcare institutions are public, while the remaining are operated privately. Both public and private institutions require municipal approval and supervision to be entitled to federal subsidies that cover around 80% of costs. Since 2003, parental copayment is capped at around 2,500 NOK per month for a full-time slot. This copayment should be the same for all centers that receive government subsidies (i.e., virtually all centers), and hence private and public centers have similar fees. For low-income families there are further subsidies, and these are similar regardless of whether the center is private or public. Childcare institutions are typically open from around 7:30 a.m. to 5 p.m.

In terms of educational content, a social pedagogy tradition has dominated childcare practices in Norway since its inception in the 1970s. According to this tradition, children should develop social, language, and physical skills mainly through play and informal learning.[9] The informal learning is typically carried out in the context of day-to-day social interaction between children and staff, in addition to specific activities for different age groups.

In table 1, we report some institutional characteristics of the institutions in our sample. We see that an average institution in our sample services 14 children aged 0–2, with about three adults per 10 children, including one teacher. The minority share among all children in the institution is about 16%, reflecting the high share of children with a minority background residing in Oslo. The enrollment of children with an immigrant background is, however, quite low for children below 3 years of age (Drange and Telle 2015).

[9] The social pedagogy tradition in early education has been especially influential in the Nordic countries and central Europe. In contrast, the so-called preprimary pedagogic approach to early education has dominated many English- and French-speaking countries, favoring formal learning processes to meet explicit standards for what children should know and be able to do before they start school.

C. Childcare Centers in Norway from an International Perspective

The provision of childcare in Norway bears resemblance with the other Nordic countries with relatively high public subsidies.[10] However, the enrollment of children below the age of 3 in Norway was 44% in 2004, substantially lower than, for example, Denmark, which had an enrollment of 83%. The Norwegian numbers are more comparable to the United States, where enrollment stood at 38% for this age group at the time (OECD 2006).

III. The Childcare Assignment Lottery

Our interest is in how early childcare enrollment affects child development. Estimating this relationship is complicated by the fact that parents and children sort into early enrollment. To circumvent this problem, we take advantage of an assignment lottery used by the Oslo city administration to distribute offers to applicants when institutions were oversubscribed.

A. Institutional Framework

Each year, the vast majority of available childcare slots in both public and private institutions are allocated in a centralized allocation round. The application deadline is around March 1 of each year, for enrollment in mid-August. Parents may apply for placement in up to seven childcare centers in their application and may list both public and private institutions.

Allocation takes place inside the city district of residence, but available slots may be allocated to children from other city districts after the main allocation round. Children may be awarded priority placement if they have, for instance, a sibling in the same childcare institution or are disabled. In our sample, 24% of children get priority placement. Children that have their first birthday after September 1 are not included in the main allocation round but may receive offers after this round is over. In our analysis, we exclude both of the former groups to focus on the main group of children that are included in the main allocation round without being assigned priority.

Based on the applications received, the municipality generates lists of nonpriority applicants to each institution. Lists for private institutions are transmitted to the institutions, which handle their own admissions based on these lists along with full details of the individual child and application. In line with Abdulkadiroglu et al. (2011), we therefore exclude from our analysis children who have a private institution ranked first on their application.

The mechanism for assignment to public institutions resembles a serial dictatorship: the order of children on the full list of applicants to each public institution is randomized in the computer before they are presented to the

[10] For children below 3 years old, parental contribution in the Nordic countries varied from 9% to 15% compared with an OECD average of 25%–30% in the mid-2000. In 2003, the state subsidy to a childcare slot for a child below 3 years old was 9,773 EUR annually in Norway (OECD 2006).

city official. Available slots are then allocated according to the random rank on the application list, and offers are sent to parents. Parents may accept or reject the offer. If they reject, the offer is conferred to the highest-ranked child on the application list who did not already receive an offer at this or some other institution. Once a child receives an offer for a childcare place, the child is taken off the lists of other public institutions to which it applied. The child may, however, choose to retain their position on the list of the institution that they ranked first in their application. Notice that this implies that the individual ranking plays no role in determining whether a child receives an offer. Conditional on receiving an offer, however, the child will have a higher likelihood of offers in highly ranked institutions.

The main allocation round ends each year around June 1. After the main allocation round, available slots may be offered to any applicant, regardless of whether they ranked the institution on their application. This process is largely at the discretion of the city officials or even childcare managers and is therefore susceptible to manipulation. We therefore use only offers dated before June 1 each year in our analysis.

B. Strategic Application Behavior

A concern with this mechanism is that it may induce strategic application behavior, that is, that applicants may not rank institutions according to their actual preferences. The allocation mechanism may spur strategic behavior along two alternative lines. The first is applying to institutions with a higher probability of offer. Listing institutions that are expected to have low oversubscription could increase the chances of receiving an offer. In our estimation, controlling for the identity of each institution on the application list should account for this kind of strategic behavior.

The second is ranking first institutions with a higher probability of offer. Notice that you may get additional draws only on the first rank after receiving a lower-ranked offer. Therefore, listing institutions that are expected to have low oversubscription first will increase the likelihood of receiving an offer from this particular institution. This may be attractive if there are large perceived gains to getting into some particular institutions on the list of acceptable institutions. For instance, say there are three institutions that you are willing to apply to. Say further that you prefer institution A over B and institution B over C and that the expected probability of getting an offer is higher in B than in A. Now assume that you much prefer B to C but are almost indifferent between A and B. In this case, you may prefer to rank B over A in order to increase your chance of avoiding C if the probability of getting an offer from B is sufficiently high compared with the probability of getting an offer from A.

Importantly, this kind of strategic behavior should not affect the internal validity of our estimates since it does not affect the probability of getting an offer, our IV. It may, however, affect the external validity of our estimates if the effect of avoiding the least preferred institutions is correlated with the likelihood of strategic behavior. Say, for instance, that only high-SES families act strategically while only low-SES children benefit from better institutions. If the strategic behavior helps high-SES children into better institutions, low-SES children will be more likely to get offers from other institutions. In this case, the estimated effect of the offer is lower than it would have been without strategic behavior. While we want to remain aware of this possibility, we do not believe that this is an important issue, both because the oversubscription rates are high for most institutions and because these oversubscription rates are not directly observable to parents.

C. Empirical Strategy

In our estimation, we follow Abdulkadiroglu et al. (2011) in including indicators for lottery-specific risk sets, denoted D below, to account for the fact that children apply to different institutions with different numbers of applicants and available slots. The extent of oversubscription determines the probability of receiving a lottery offer from a particular institution. If the extent of oversubscription is correlated with, for instance, the quality of the childcare institution while applying to good institutions is, in turn, correlated with unobservable traits that determine cognitive performance, then a comparison based on lottery offers may give biased estimates of the impact on cognitive performance of early childcare enrollment. Such endogeneity could, for instance, result from differences in strategic application behavior between families.

In addition, by ranking a higher number of institutions one may increase the likelihood of receiving an offer. If the number of institutions on the application list is correlated with the outcome, this could also be a source of bias in our estimates. In particular, ranking several institutions would probably be a signal of how strong the preference for childcare is, which should be closely related to the labor market attachment of the mother. In table A1, we report background characteristics for children with parents who ranked 1, 2–4, and 5–7 centers. As expected, families that applied to more institutions are indeed different from families that applied to less. Overall, the observed pattern is consistent with families that have a stronger attachment to the labor market, maximizing their chances to obtain a slot by applying to more institutions. Parents are better educated, mothers have higher earnings, and the child is much less likely to be of immigrant background. To guard against such bias, we control flexibly for the number and identity of institutions to which an applicant applied, in line with Abdulkadiroglu et al. (2011).[11]

[11] Specifically, the risk set D includes a full set of dummy variables for each institution by year, so that for each institution and year there is a dummy equal to 1 if child i applied to that institution in that year and 0 otherwise. In addition, the risk set includes dummy variables for the number of applications by year.

For lottery offers to be relevant, we need to have oversubscription of toddlers to childcare institutions. This is determined by the number of nonpriority applicants per remaining available slot after priority placements. Table 2 shows descriptive statistics for the number of available slots, the number of applicants, and oversubscription to public childcare institutions in Oslo in the period we consider. In the upper panel, we show oversubscription measured as the number of applicants to each childcare slot. Note that this implies that each child may be counted up to seven times, depending on how many centers have been ranked in the application. Oversubscription is both strong and widespread: the mean number of applicants to each childcare slot is about 25. In the lower panel of the table, we restrict attention to the first choice. This means that each child is now counted only once. For the first choice, on average 15 children apply for 4.4 places, with an average of 4.6 applicants to each place. The large oversubscription is mirrored in the fact that only 29% of the children in our sample receive an offer in the assignment lottery and in the strong effect that a lottery offer has on the childcare starting age documented in our first-stage estimates below.

The validity of the lottery offer as an instrument for childcare starting age relies on the quality of the assignment lottery. While the city administration assures us that the lottery was randomized by a computer algorithm, as described earlier, there is always the possibility that the randomization failed or that there was manipulation between the actual randomization and the sending out of offers. Above, we also noted the possibility of strategic application behavior; that could be a further threat to the randomization.

To verify that the randomization was successful, the first two columns of table 3 report means and standard deviations of background characteristics for children in our estimation sample separately by whether the child received an offer. Table 3 shows that the two groups look well balanced. We also test this formally in the context of our econometric model by regressing the offer dummy on all characteristics, controlling for the risk set D. The final column of table 3 reports t-statistics of the individual coefficients from this regression, which are usually very low. In a joint test of

Table 2
Applications, Places, and Oversubscription in Public Childcare Institutions in the Centralized Admission Process in Oslo, 2005–7

	Mean	SD	Min	Max
Oversubscription:				
Number of places	3.99	3.73	0	38
Number of applicants	70.89	46.28	1	265
Applicants/places	25.13	24.38	.5	156
Oversubscription by first-choice institution:				
Number of places	4.37	3.47	0	21
Number of applicants	14.81	9.95	1	82
Applicants/places	4.60	4.47	.2	30

Table 3
Balance in Background Characteristics between Children With and Without a Lottery Offer

	Offer	No Offer	t-Value
Girl	.507 (.50)	.498 (.50)	.05
Age	14.74 (2.18)	14.75 (2.19)	−1.34
Immigrant	.113 (.32)	.103 (.30)	−.79
Mother:			
Years of education	14.59 (3.06)	14.71 (3.10)	.17
Earnings	293,247 (162,959)	302,394 (162,135)	−.61
Age	33.03 (4.38)	33.42 (4.372)	−1.26
Age first birth	29.64 (4.47)	29.87 (4.47)	.26
Father:			
Years of education	14.26 (3.50)	14.55 (3.41)	−1.08
Earnings	420,629 (416,577)	416,434 (347,658)	1.28
Age	35.02 (6.55)	35.51 (6.15)	−.30
Age first birth	31.00 (5.90)	31.45 (5.61)	.17
	N = 852	N = 2,036	F = .95

NOTE.—The table reports means and standard deviations of covariates by whether the child received an offer in the assignment lottery. The final column reports t-statistics of the individual coefficients from a regression of the offer dummy on all characteristics, controlling for the risk sets, and the F-statistic from a joint test of whether coefficients on all covariates are equal to zero. Age refers to the age in months of the child in August of the year of application. Earnings are pensionable income from work and self-employment. Detailed descriptions of the background characteristics are provided in Sec. IV.

whether coefficients on all covariates are equal to zero, we get an F-value of 1.22, confirming that the two groups are indeed well balanced. This suggests that the randomization is successful.

In line with Abdulkadiroglu et al. (2011), we exclude from our analysis children who have a private institution ranked first on their application. Including children with private institutions on the first rank would risk introducing bias into the estimates. To see this, recall that private institutions handle their own admissions based on the application lists, including potentially full details on the individual child and application. Offers made by private institutions are therefore likely to be correlated with other determinants of the outcome and, hence, endogenous. If children who receive offers from private institutions withdraw from the public application process, then they may be less likely to receive lottery offers in the main allocation round. Including them in the estimation sample may then cause selection into the control group. This problem is clearly most relevant for children with a private institution on the first rank. In contrast, children with a private institution on lower ranks are likely to remain in the lottery to retain the possibility of getting an offer from a higher-ranked institution. This is supported by the data, where the probability of getting a lottery offer is much lower among children with a private institution on the first rank, at about 18%. Meanwhile, children with a private institution on lower ranks are quite similar to children without private institutions on their application lists in their probability of receiving a lottery offer (28% vs. 31%).

IV. Data and Estimation Sample

A. Data

Our data are based on several different administrative registers from the Oslo city government and Statistics Norway. First, we have access to the municipal database used in the centralized application system for childcare in Oslo. This provides information on applications for and enrollments in virtually all childcare institutions in Oslo for the years 2005–10, including both public and private childcare institutions. Applicants who list several institutions in their applications are registered as separate coincident applications. The database also provides information about offers of slots in public childcare centers. Applications, enrollment, and offers are recorded with date of receipt, date of first attendance, and date the offer was made, respectively.

Second, we have access to a database with information about performance on tests made available by the school authority in the Oslo municipality. This provides information about enrollment in primary school and scores on performance tests in the Norwegian language and mathematics, conducted in April of first grade. The tests are designed nationally and are intended to help identify underperforming children, enabling schools to allocate resources to these children. The language test maps the ability to write letters, to recognize written letters, to identify spoken letters, to combine sounds, to write words, to read words, and to read sentences. The mathematics test maps the ability to count, to compare numbers, to rank numbers, to recognize sequences of numbers, to count forward and backward from a given number, to split a number into two other numbers (i.e., 4 = 1 + ...), to solve textual assignments, and to add two numbers. We provide further detail on these tests in the appendix.

Each test is scored on a relatively fine scale, where students may score from 0 to 105 in language and from 0 to 50 in mathematics. Because tests are designed to identify children with problems, test score distributions are skewed,[12] with about 10% and 15% of children in our sample getting the top score in language and mathematics, respectively. This is important to keep in mind when interpreting our results. To verify that this truncation does not affect our estimates, we have estimated Tobit models with the maximum attainable score as the upper truncation point. It is reassuring that estimates are virtually unchanged.

In our analysis, we consider two outcomes from each test. First, we normalize the scores to have mean 0 and standard deviation 1. In addition to separate estimations for the standardized test scores in language and mathematics, we also consider the unweighted average of the two tests as a summary measure of cognitive skills. Second, we use dummy variables for per-

[12] Figure A1 draws the distribution of test scores in our sample.

formance below a nationally determined threshold. Thresholds are set for individual parts of each test from a trial of the test on a panel of children conducted prior to actual testing. The thresholds are intended to identify the bottom 15%–20% of children nationally. From these we define the dummy variable *Below threshold* equal to 1 if the child has one or more test parts with scores below the threshold and 0 otherwise, separately for language and mathematics. There are two advantages with this outcome over the test scores. First, while the economic significance of test scores may be hard to assess, the thresholds are constructed in order to identify children with potential development problems. Second, these dummy variables are robust to problems associated with outliers.

Third, we can link both databases to rich Norwegian administrative registers available from Statistics Norway with individual information on demographics (e.g., sex, age, immigrant status, marital status, number of children), SES (e.g., years of education, income, employment status), and residence. Income and employment data are collected from tax records and other administrative registers. The household information is from the Central Population Register, which is updated annually by the local population registries and verified by the Norwegian Tax Authority. We also have access to national registry data on municipal childcare coverage reported by the childcare institutions themselves. The reliability of Norwegian register data is considered to be very good, as is documented by the fact that they received the highest rating in a data quality assessment prepared for the OECD by Atkinson, Rainwater, and Smeeding (1995). Importantly, all data sources contain personal identifiers that allow us to link individuals across all registers.

B. Estimation Sample

We start with the universe of children born from 2004 to 2006,[13] for whom parents apply for a childcare slot in Oslo in the calendar year the children turn 1 year old. Because our identification comes from offers of public childcare slots, we focus attention on children with a public institution on the first rank, while we allow both private and public institutions on slots 2–7. As discussed above, we also exclude children who had priority in childcare or who turn 1 after September 1 in the application year, since our identification does not influence these children. We finally exclude a handful of children with missing values on our dependent variables and a handful of children registered as starting in childcare before 10 months old. Rather than excluding children with missing values on control variables, we construct dummy variables for missing and include these in our regressions. Our final estimation sample consists of 2,888 children. To explore how representative this sample is, table A3 re-

[13] Due to a restrictive storage policy in the municipality, data on children born in January and February 2004 were deleted from the application database before we received access to it. We are therefore not able to include these children in our sample.

Table 4
Performance in Language and Mathematics Tests at Age 7 for Children With and Without a Lottery Offer

	Offer	No Offer
Average score	72.97 (5.72)	72.00 (8.34)
Language	100.31 (7.14)	99.02 (11.74)
Below limit	.10 (.30)	.12 (.33)
Mathematics	45.63 (5.66)	44.97 (6.59)
Below limit	.04 (.20)	.06 (.24)
N	852	2,036

Note.—"Offer" indicates children who received an offer in the assignment lottery, and "No Offer" indicates children who did not receive an offer in the assignment lottery (see Sec. III).

ports means of observable covariates for (1) all children in the relevant cohorts residing in Oslo, (2) all families applying for a childcare slot the year their child turns 1, and (3) the final estimation sample. We see that, in line with what we would expect, parents who apply for a childcare slot the year their child turns 1 do on average have somewhat higher education and earnings than the average parents with a 1-year-old. The share with immigrant background is also substantially lower in the sample of parents who apply.[14] However, comparing families who apply with families in the final estimation sample, we see only small differences in average background characteristics.

V. Empirical Results

A. Reduced Form: Effects of Offers from the Assignment Lottery

We now turn to our main analysis of how early enrollment in childcare affected the cognitive performance of children at age 7. We start with a reduced-form analysis, where we compare outcomes of children who received a lottery offer to children who did not receive a lottery offer. Table 4 shows means and standard deviations of our main outcome variables separately for children who received an offer in the assignment lottery and children who did not receive an offer. The mean test scores show that children who receive a lottery offer perform about 10% of a standard deviation better than children who do not receive a lottery offer. Meanwhile, just over 12% of children are below the threshold for low performance in language, while about 6% are below the threshold in mathematics. In both subjects, children with a lottery offer are about 2 percentage points less likely to score below the threshold for low performance. This is the first evidence that early childcare enrollment has a positive impact on children's cognitive development.

Next, we consider this reduced-form model formally by estimating the impact of receiving an offer on test performance, controlling for risk sets as in equation (3). Specifically, we run the regression

[14] It is known from other studies that immigrant parents in Oslo are less likely than native parents to enroll their toddler in childcare (Drange and Telle 2015).

Table 5
Reduced-Form Estimates of the Impact of a Lottery Offer on Performance in Language and Mathematics

	No Controls		With Controls		
	b	SE	*b*	SE	Mean
Average score	.130	.044	.138	.042	.00
Language	.158	.048	.165	.047	.00
Below limit	−.032	.016	−.033	.016	.12
Mathematics	.103	.051	.111	.048	.00
Below limit	−.020	.011	−.021	.011	.06
N	2,888		2,888		

NOTE.—Effects are reported as percentage of the standard deviation. Standard errors are robust to heteroskedasticity and are clustered at the first-choice institution level. All regressions include a risk set with a full set of dummy variables for each institution by year and the number of childcare institutions listed. We also include cohort fixed effects. "No Controls" reports estimates without covariates, whereas "With Controls" reports estimates including the controls listed in table 3.

$$y_{it} = \tilde{\gamma} OFFER_{it} + D'_{it}\tilde{\alpha} + X'_{it}\tilde{\beta} + \tilde{\epsilon}_{it}, \tag{1}$$

where i denotes the individual, t denotes the cohort, and D_{it} is the risk set of individual i (see n. 11); $OFFER_{it}$ is a dummy equal to 1 if the child received an offer of a public childcare place generated in the assignment lottery, and X_{it} are a set of sociodemographic characteristics of the child and parents, measured the year before the child was born.[15] We report estimates both including and excluding these covariates. While including covariates should not change our estimates when the explanatory variable of interest is as good as random, it may be helpful to improve precision in our estimates. The residual $\tilde{\epsilon}$ is clustered at the level of the first-choice institution.[16]

Results reported in table 5 clarify our observations from above and indicate that receiving a lottery offer improved the average performance of children by about 13% of a standard deviation overall. This effect was driven both by an improvement in language of about 16% of a standard deviation and by an improvement in mathematics of about 11% of a standard deviation. When we consider the impact on the probability of scoring below the limit for low performance, we find a decrease of 3 percentage points on the language test and 2 percentage points on the mathematics test. Both of these effects are large compared with the mean in the control group, corresponding to a drop in the probability of about 30%. As expected from the above balancing analysis, estimates barely move when we include covariates.

[15] Child characteristics include gender, month of birth, and birth order. Parental characteristics include dummy variables for full-time work, receipt of social assistance, high school completion, college degree, missing parental education, and missing parent identifier.

[16] Note that we do not want to cluster on the childcare center attended since this is likely to be endogenous to the outcome of the assignment lottery.

To understand the economic significance of our estimates, note first that there is a substantial negative impact on the probability of scoring below the limit for low performance, where the school should consider specific measures to support the development of the child. This suggests that the offer of early childcare promotes skill development among the most needy kids in particular. These estimates are also reassuring, since they are not affected by the particular shape of the test score distribution or problems associated with outliers.

Understanding the economic significance of changes in test scores is in general somewhat difficult because they do not have a meaningful cardinal scale (Cunha and Heckman 2008). This may be particularly true in our case, where the distribution of the test score is skewed and quite different from the often bell-shaped test scores considered in the literature. Although comparisons of estimates across different test score outcomes is always risky, this means that it could be particularly misleading to compare our estimates directly to those found in other studies.

To interpret our estimates, we therefore need to map them into a metric that is more easily interpretable in other contexts. Ideally, we would use an estimate of the effect or association of the test score and a long-term outcome like income or education. Unfortunately, our outcome measure is not available for older cohorts. Instead, we compare our estimates to the gaps in test scores that we observe between well-defined subgroups of the population. Specifically, the effect on the language test corresponds to the gap in language performance between children from high- and low-income families or to about half the gap between children from high- and low-educated parents (see table 6). In mathematics, our reduced-form estimate corresponds to about one-quarter of the gap between children of low- and high-educated parents or about one-third of the gap between children from high- and low-income families. This confirms

Table 6
Reduced-Form Estimates of the Impact of a Lottery Offer on Performance in Language and Mathematics in Different Subgroups

	Language			Mathematics			
	b	SE	Mean	*b*	SE	Mean	*N*
Boys	.209	.098	−.119	.085	.088	.017	1,442
Girls	.128	.076	.056	.158	.088	−.061	1,446
Parents' education low	.244	.116	−.182	.113	.107	−.201	1,551
Parents' education high	.081	.067	.137	.044	.083	.179	1,337
Mom age low	.274	.098	−.135	.162	.091	−.081	1,596
Mom age high	.042	.065	.093	.043	.089	.050	1,292
Family income low	.264	.123	−.127	.256	.106	−.166	1,444
Family income high	.085	.068	.062	.010	.078	.120	1,444

Note.—Effects are reported as percentage of the standard deviation. Estimates are from eq. (1), controlling for the risk set by including a full set of dummy variables for each institution by year and the number of childcare institutions listed in the application (see Sec. III). Standard errors are clustered at the first-choice institution level and are robust to heteroskedasticity. Parents' education is the maximum education of the mother and father. Family income is the sum of pensionable income of the mother and father in the year before birth. The socioeconomic groups in rows 3–8 are constructed by splitting the sample at the median.

that the estimated effects of a public childcare offer are substantial on children's language and numeracy development.[17]

B. Heterogeneous Effects

One important argument for why governments want to subsidize childcare is that they may help counter differences in school readiness between children from different socioeconomic backgrounds. It is therefore natural to next consider whether the starting age in childcare has a different impact on toddlers from different socioeconomic groups.

Table 6 reports estimates from our reduced-form including covariates, estimated separately for children by gender and by different socioeconomic groups.[18] Although estimates are not sufficiently precise to statistically rule out equality, the pattern in the point estimates is striking. Throughout, effects are two to three times stronger in the underperforming group. For instance, children of low-educated parents are estimated to improve their performance in language by about 25% of a standard deviation, almost the same as the overall gap in performance compared with children from high-educated parents. The performance of these latter children is estimated to increase by a comparably modest and statistically insignificant 8% of a standard deviation. This suggests that if both receive a childcare offer, then the gap between the two groups is halved. Similarly, children from low-income families are estimated to improve their performance in both language and mathematics by about 25% of a standard deviation, compared with modest or no effects for children from high-income families.

One hypothesis for why some groups respond more than others might be that the impact of the treatment is stronger, that is, the delay in childcare start is longer. However, heterogeneity on the first stage turns out to be relatively modest, and the IV estimates across groups mirror closely the pattern we see in the reduced-form estimates (see table A4).

The heterogeneity observed in the impacts on the mean test score motivates looking more closely at how impacts vary over the distribution of the test score. To this end, we have estimated our reduced form where the dependent variables are dummy variables for scoring above cutoffs defined from

[17] It should be noted, however, that the test scores we consider are skewed and may reflect poorly the ability distribution at the top. The observed gaps are therefore likely smaller than what we would observe with more symmetrically distributed test scores. Note also that we have verified that this apparent truncation of the ability distribution does not affect our estimates by estimating Tobit models with upper truncation points at the maximum attainable score. The estimated impacts on the test scores are virtually identical in this and our baseline model.

[18] The socioeconomic groups are constructed by splitting the sample at the median for family income (the sum of pensionable income of the mother and the father in the year before birth). Parents where both the mother and the father have completed a degree beyond high school are included in the high education group. Mother's age is defined as high if the mother was above 30 when she gave birth to her first child.

quantiles of the test score distribution among control children. In figure 1, we draw estimates for 19 quantiles with associated 90% confidence interval (for the estimates and standard errors, see table A5). Specifically, the first point in the graph suggests that getting an offer increases the probability of scoring above the fifth percentile by 3 percentage points in language and 2 percentage points in mathematics, both statistically significant. Overall, figure 1 supports the pattern suggested by the subsample estimates of positive impacts for children who perform at the lower end of the distribution, while effects at the top are less precise and closer to zero. Remember, however, that the tests we consider are not designed to capture well the variation at the top of the distribution. We should therefore be cautious in reading too much into the estimates at the top of the distribution.

C. Mechanisms

Having established that receiving a random offer causes a substantial improvement in performance on the tests, the question is what is driving this ef-

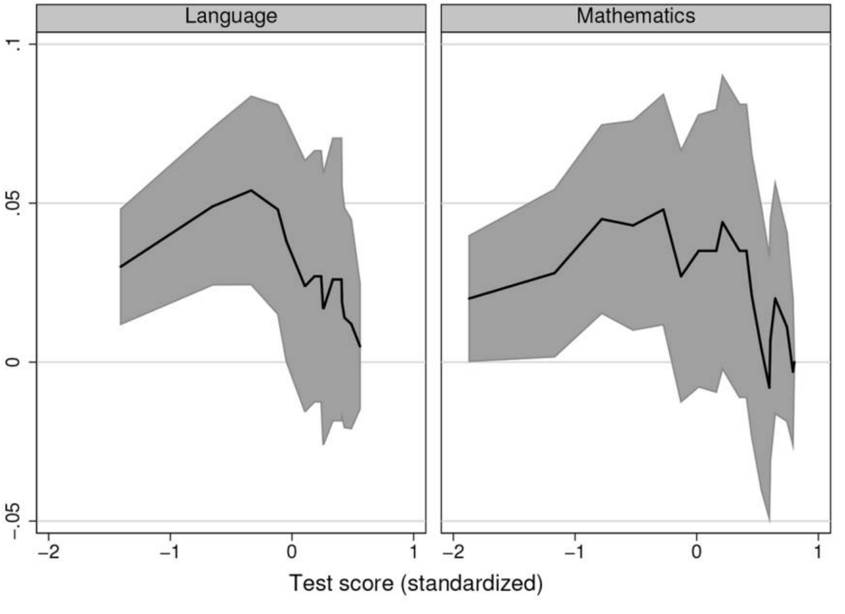

FIG. 1.—Reduced-form estimates of the impact of a lottery offer on performance in language and mathematics at 19 quantiles and associated 90% confidence interval. The dependent variable at each point is a dummy for performing above the level of the test score indicated by the X-axis. Estimates are from equation (1), estimated at each of 19 quantiles, controlling for the risk set by including a full set of dummy variables for each institution by year and the number of childcare institutions listed in the application (see Sec. III). Standard errors are clustered at the first-choice institution level and are robust to heteroskedasticity.

fect. As usual in natural experiments, the estimated reduced-form effect may immediately be attributed to more than one causal channel. We therefore think it is useful to separate clearly between, on the one hand, the reduced-form estimates, which are unbiased as long as the instrument is as good as randomly assigned, and, on the other, the analysis of mechanisms, which will always rely more on theory, interpretation, and suggestive evidence. In our case, we believe that receiving an offer in the assignment lottery can impact the child's development broadly through three alternative channels:

1. It can allow the child to start childcare at an earlier age.
2. It can allow the child to attend a higher-quality institution (in general or for the individual child).
3. It can lead to changes in the family environment, in particular family income and parental labor supply.

Below, we consider each of these channels in turn.

1. Childcare Starting Age

To investigate starting age as a potential mechanism explaining our estimates, we need to understand more in detail how receiving a lottery offer affects childcare starting age. To this end, figure 2*A* shows the cumulative distribution of children having started childcare at different ages separately for children who receive and do not receive an offer. While 91% of children who received an offer had started childcare by the time they turned 18 months old, this was the case for only 65% of the comparison group. On average, children with public offers started childcare at around 15 months of age, while children without public offers started closer to 19 months of age.

Figure 2*B* shows the corresponding cumulative density function among children with and without an offer of months of delay before childcare start, compared with their age in August in the year of application. Children who start normally should be delayed less than 1 or 2 months. Among children who receive offers in the assignment lottery, less than 5% are delayed, compared with almost 40% of children without offers. Among those children who were delayed because of a lack of a lottery offer, about two out of three were delayed by more than 12 months.

The strong impact of getting an offer on the starting age in childcare motivates using the offer as an IV for starting age in childcare.[19] To this end, we follow closely the approach of Abdulkadiroglu et al. (2011). Specifically, we wish to estimate the following IV model, where equation (2) is the second stage and equation (3) is the first stage.

[19] Notice that we condition on the age of the child in August, so that our specification is econometrically equivalent to a specification in terms of delay rather than starting age.

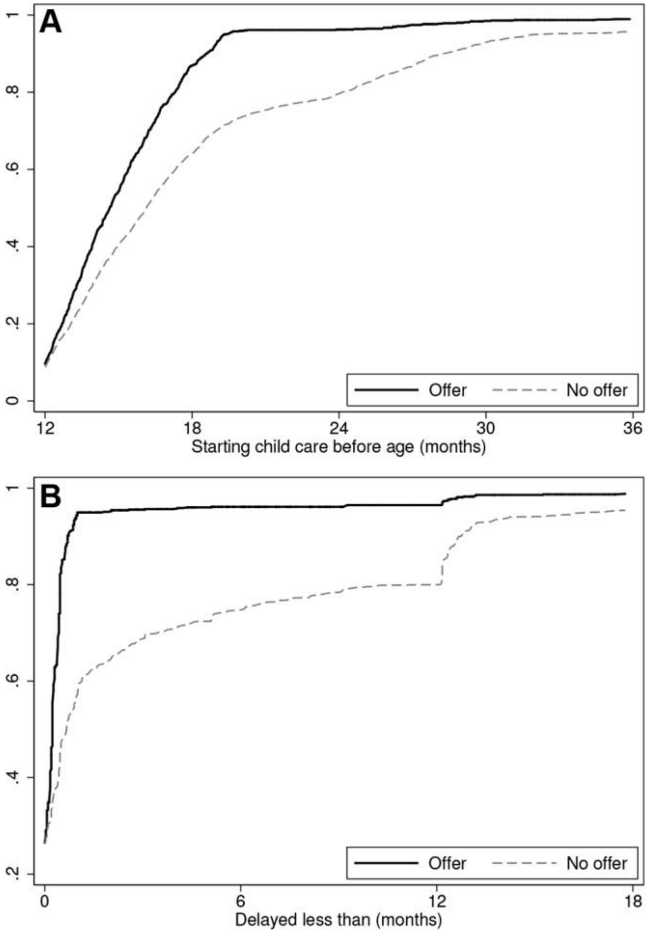

FIG. 2.—Cumulative distribution of age at childcare start and delay in childcare start for children with and without a lottery offer. *A*, Starting age. *B*, Delay. Childcare starting age is the age in months on the day when the child first attends childcare. Delay in childcare start is calculated as the childcare starting age less the age of the child in August of the year of application.

$$y_{it} = \gamma AGE_{it} + D'_{it}\alpha + X'_{it}\beta + \epsilon_{it}, \tag{2}$$

$$AGE_{it} = \pi OFFER_{it} + D'_{it}\eta + X'_{it}\psi + \omega_{it}, \tag{3}$$

where AGE_{it} is the age of child i in months when he or she first attends any childcare institution, public or private, and all other variables are defined as before. Again, we always condition on the risk set D (see n. 11) and estimate the model with and without the set of control variables X_{it} (see n. 15). In both the first stage and the second stage, residuals are clustered at the level of the first-choice institution, as before.

Note that estimating equation (2) directly will likely yield biased estimates. This is because enrollment in childcare is likely to be determined in part by parental preferences and by child innate characteristics and because starting age is likely to be correlated with unobserved determinants of cognitive performance. For instance, we might expect more able parents to be more closely tied to the labor market and therefore enroll their children in childcare earlier. If so, then we may expect that children who are enrolled early would perform better in any case. On the other hand, we might expect more child-centered parents to enroll their children in childcare later. If so, being enrolled early could be a marker for a poor home environment, which would suggest that these children should perform worse. This implies that direct estimation of equation (2) will give biased estimates of the impact of childcare starting age on cognitive performance.

Panel A of table 7 reports estimates from our full IV model, where the receipt of a lottery offer is used to instrument for the age at first attendance in a childcare institution. The two last rows of the table report estimates from our first-stage equation and show that the lottery offer decreased starting age by about 3.2 months on average. The *F*-statistic on the instrument is above 100, which implies that we need not worry about problems associated with weak instruments.

Turning now to the IV estimates in table 7, the estimates without controls suggest that starting childcare 1 month later causes a drop in school performance of about 4% of a standard deviation. As before, this is driven by a

Table 7
Instrumental Variable Estimates of the Impact of Childcare Starting Age on Performance in Language and Mathematics

	No Controls		With Controls		
	b	SE	*b*	SE	Mean
A. First stage:					
Offer	−3.235	.315	−3.137	.310	18.62
F-value (instrument)	105.5		102.4		
B. Second stage:					
Average score	−.040	.012	−.044	.012	.000
Language	−.049	.013	−.053	.013	.000
Below limit	.010	.004	.010	.004	.121
Mathematics	−.032	.013	−.035	.013	.000
Below limit	.006	.003	.007	.003	.062
N	2,888		2,888		

Note.—Effects are reported as percentage of the standard deviation. Estimates are from eqq. (2) and (3). Standard errors are clustered at the first-choice institution level and are robust to heteroskedasticity. All regressions control for the risk set by including a full set of dummy variables for each institution by year and the number of childcare institutions listed in the application (see Sec. III). Control variables are listed in Sec. III. Mean refers to the mean of the dependent variable among children in the control group.

drop in both the language score and the mathematics score, and effects are somewhat stronger on the language score. Quantitatively, the estimates suggest an increase in performance of about 5% and 3.5% of a standard deviation in language and mathematics, respectively. All of these estimates are significant at the 5% level. When we consider the impact on the probability of scoring below the limit for low performance, we find an increase of around 10% on both tests compared with the mean in the control group. As before, estimates barely move when we include covariates.

Finally, to evaluate the plausibility of starting age as a mechanism, we consider children's performance as a function of the delay in childcare start separately among children who received a lottery offer and among those who did not receive an offer. In figure 2*B*, we saw the cumulative distribution of children having started childcare with different delays compared with the start of term in August. Figure 3 reports the mean test score among children who were delayed less than 0–18 months, approximated by a local linear regression. Specifically, the rightmost point on the solid curve gives the expected performance of a child in the treatment group who is delayed by less than 18 months. The corresponding point on the dotted line gives the expected performance of a child in the control group who is delayed by less than 18 months. Since almost no children are delayed by more than this, the gap between the two groups corresponds roughly to the overall difference in performance among treated and control children.

As we move left, the treatment and comparison groups are becoming more homogeneous in terms of the age at which they start childcare. If starting age is an important mechanism, we expect the performance of children to become more similar as starting ages move closer together. This is largely confirmed in both panels: children who start childcare early, regardless of whether they receive a lottery offer, perform similarly on the tests, while children who start later and who did not receive an offer tend to perform worse. This is striking in the case of the mathematics test, where the difference between the groups is entirely driven by children in the control group who start late. For the language test, the gap also widens with starting age, but here the gap between early starters suggests that the treatment also generates other relevant differences.

2. *Childcare Quality*

Above we considered childcare starting age to be the primary candidate for how an offer of a place in public childcare affects children's language and numeracy skills. An alternative candidate is the characteristics or quality of the childcare institution attended. If children with offers not only started earlier but also attended better-quality childcare institutions, then this could be driving the improvement in performance we observed in our main estimates.

To understand the potential for childcare quality as a mediator for the estimated effects, we investigate whether the characteristics of the childcare

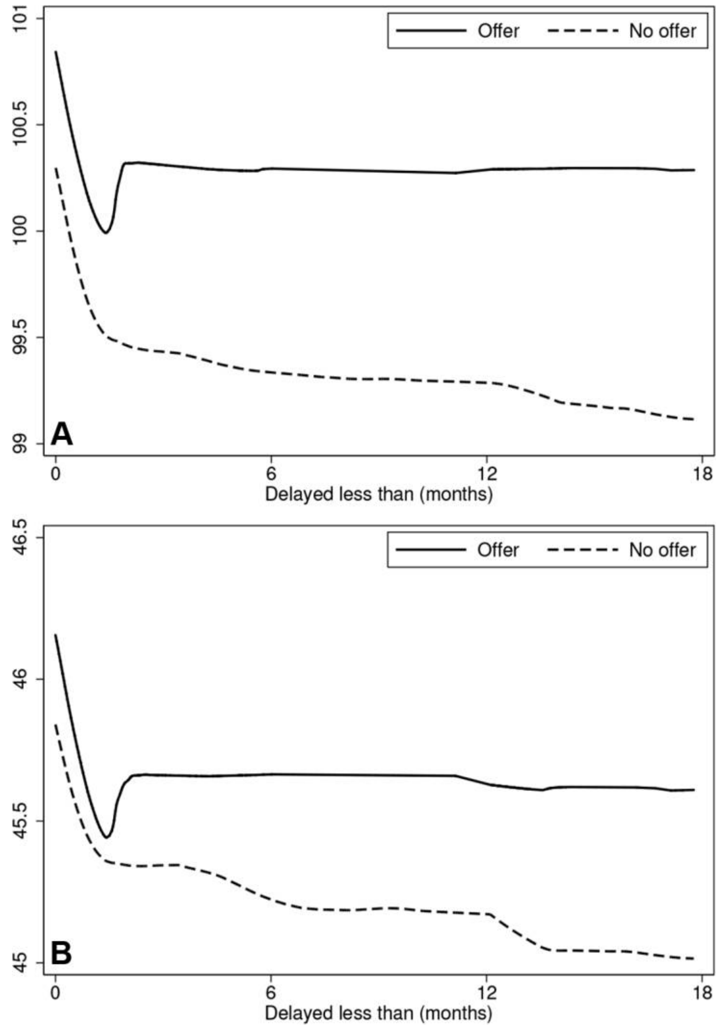

FIG. 3.—Performance in language and mathematics by childcare starting age, for children with and without a lottery offer. *A*, Language. *B*, Mathematics. Figures show local linear regression estimates of mean language and mathematics performance against starting age in childcare, using an Epanechnikov kernel with the bandwidth set to 1 month.

institution that the child first attends depend on whether the child received a lottery offer. Table 8 lists a wide set of characteristics of the childcare institutions that children first attend. Structural characteristics are mean test scores of all children in our data who first attended the institution, the teacher-to-child ratio, and travel distance from home. Staff characteristics are mean characteristics of staff employed at the institution, while peer char-

Table 8
Impact of a Lottery Offer on Characteristics of the First Childcare Institution Attended

	b	SE	Mean	*N*
A. Structural characteristics:				
Language	−.053	.011	.100	2,888
Mathematics	−.034	.013	.087	2,888
Teacher/child	−.007	.002	.082	2,888
Distance (km)	−.515	.257	3.268	2,581
Distance (min)	−1.221	.440	6.300	2,581
Private center	−.408	.025	.442	2,888
B. Staff characteristics:				
Income	12,638	2,546	248,791	1,998
College graduates	.004	.008	.356	1,998
Immigrants	−.029	.012	.248	1,998
Males	−.004	.007	.108	1,998
Age	1.754	.405	36.808	1,998
C. Peer characteristics:				
Family income	−13,715	11,787	755,752	2,742
College graduates	−.004	.011	.494	2,742
Immigrants	.011	.006	.079	2,742
Males	.004	.011	.497	2,742
Young	−.054	.012	.496	2,742

NOTE.—The table reports the impact of a lottery offer on characteristics of the first institution a child attends. The first two columns report coefficients and standard errors from a regression of the offer dummy on various institutional characteristics (see Sec. III). Mean reflects the average characteristics of the childcare center that the child first attends. Variations in sample size are due to missing dependent variables.

acteristics are mean background characteristics of children in the institution in the year of application.[20]

In table 8, we report reduced-form estimates of the impact of a lottery offer on each characteristic. As one should expect, children who got offers attended centers with somewhat different characteristics. At the same time, there are few indications that children with offers attended higher-quality institutions. On the contrary, children with lottery offers seem to attend institutions with slightly lower structural quality and somewhat worse performing peers.

Specifically, in panel A the first two rows give estimates of the effect of a lottery offer on the average performance of children who attended the center before starting school. The estimates suggest that children with lottery offers attended institutions where children performed about 5 and 3 percentage points worse on the language and mathematics test in first grade, respectively. The third row shows that the teacher-to-child ratio was also lower. At the same time, institutions were about 500 meters closer to the child's home, or about a 1.25-minute drive. Panel B of table 8 shows that the staff characteristics were largely the same in the first institutions of both treated and control

[20] To match staff with childcare centers, we rely on matched employer-employee registers that cover about 80% of Norwegian employees.

children, although treated children were somewhat less likely to have staff with immigrant background. Finally, panel C shows that the peer characteristics were also largely similar in terms of both gender and family background, although children with lottery offers attended institutions with less young children (below 3 years old) and more children with an immigrant background. Overall, table 8 suggests that there are few substantial differences in childcare quality between children with and without lottery offers. The main exception is a much higher probability of attending a private institution. This effect is automatic in our setting, where the treatment is defined as receiving an offer in a public childcare institution. Notice that public and private institutions both require municipal approval and supervision to be entitled to federal subsidies that cover around 80% of costs. Both are also subject to the same regulations in terms of quality and price. Ownership of private institutions can be nonprofit or for profit. Below we discuss further whether quality differences between public and private centers is a likely channel for the increase in test scores among children with lottery offers.

As we have seen, perhaps the most distinct difference from table 8 was the substantial increase in the likelihood of attending a private institution. A closer look at our data reveals that while public and private institutions differ in some dimensions, these differences go in different directions with regard to the expected quality of the institution (see table A2). On the one hand, private institutions have slightly higher teacher-to-child ratios, cater to children from higher SES backgrounds, and achieve better scores on early language and mathematics tests, suggesting higher quality. At the same time, private institutions have younger and less well paid staff and are situated further from the childrens' home on average, suggesting lower quality. Overall, while public and private institutions do differ along some dimensions, we do not believe that the evidence supports that the higher likelihood of ending up in a private institution could be an important contributor to the treatment effect.

A concern with these estimates is that they may not capture the unobservable quality differences between childcare institutions. While estimates on the performances of children who have attended the institution should capture the overall quality, observable and unobservable, it cannot capture the individual quality that an institution may offer a particular child. In particular, we might be concerned that children who receive an offer in the assignment lottery are able to attend a childcare institution that matches better with their or their parents' needs.

To get at the match quality, we exploit the fact that parents of children who do not receive an offer in the assignment lottery should differ in the importance that they assign to the quality of the institution. Parents who assign less weight to quality should be more willing to accept alternative childcare placings. If we assume that the arrival rate of quality and nonquality offers is the same for both types of parents, it follows that children in the control group who start childcare at a similar age to children in the treatment group should attend childcare of lower quality than children in the control group who delay starting childcare. This implies that if the quality of the childcare institution is the primary driver of our reduced-form estimates, then we should expect that children in the control group who do not delay childcare attendance compared with treatment children suffer the most from not getting an offer in the assignment lottery. Among children who do delay childcare, there should be little or no difference depending on the assignment offer.

To investigate this, we consider, separately for children in the control group and the treatment group, the mean performance of children by the delay in childcare starting age (compared with starting in August, at the start of the childcare year). We turn again to figure 3, where the mean test scores among children who were delayed less than 0–18 months is reported. If match quality is driving our results, then the gap between the performance of control and treatment children should be shrinking as the delay increases. In figure 3, we see that the contrary is true: the performance gap is increasing as the delay increases. This suggests that match quality is not the prime driver of our reduced-form estimates. On the other hand, this pattern is, of course, well in line with expectations if starting age is the prime driver. In this case, the performance gap should be small for children who do not delay and larger as the delay grows, just as observed in figure 3.

3. *Family Income and Alternative Mode of Care*

The effect of childcare enrollment is related to the alternative mode of care had the children not been enrolled in childcare. In our case, our instrument pushes childcare enrollment forward by about 3 months on average, or about 1 year for a fourth of the sample. After this period, children who do not receive a lottery offer are also on average enrolled in childcare. For both, enrollment in childcare is largely an absorbing state, with the vast majority of children who enroll staying enrolled until the school starting age (which is 6 years in Norway). To understand the drivers behind our estimates, we therefore need to consider these two modes of care. That is, what type of care are control children in before they start regular childcare, and what are the characteristics of the formal care that they attend once they do start regular childcare?

Above we analyzed the differences in the formal care that the treatment and control children attend. We now take a look at the mode of care attended. Typically, one considers three alternatives: parental care, formal childcare, and other, more informal sources of care (Blau and Currie 2006). Unfortunately, we do not have direct information about alternative modes of care in our sample. As an alternative, we consider survey data on stated demand and actual use of childcare for a nationally representative sample of parents with 1-year-olds in 2003, presented in table 9. While about 70% of parents

Table 9
Survey Evidence on the Demand for and Use of Childcare among 1-Year-Olds in Norway, 2003

	Stated Demand	Actual Use
Parents	.17	.56
Relatives	.04	.04
Unlicensed caregivers	.08	.13
Childcare	.42	.22
Combined/other	.28	.05

SOURCE.—Pedersen 2003.

in the survey state demand for childcare, either full time or part time in combination with other forms of care, only 27% actually have their children enrolled in childcare. In comparison, while 56% of parents say that they care for their children themselves, only 17% actually prefer to do so. According to this evidence, then, parental care is the dominant alternative for Norwegian toddlers in general. While it is not clear that this is representative for our sample, anecdotal evidence from Oslo points in the same direction. Typically, one or both parents may, if they do not get a place in childcare for their child, take unpaid leave (with employment protection), or they may patch together childcare by working reduced hours and relying on grandparents or other sources of informal care.

We next consider how a lottery offer affects the earnings and labor force participation of the parents. On the one hand, this addresses the extent to which family income may explain part of the estimated effects. On the other hand, the effect on labor force participation may give us an idea about the importance of parental care as the counterfactual mode of care in our particular sample. To measure labor force participation, we rely on information about annual earnings, including wages and income from self-employment. Specifically, we construct dummy variables for employment based on the basic amounts in the Norwegian Social Insurance Scheme (used to define labor market status, determining eligibility for unemployment benefits as well as disability and old age pension). In 2006, one basic amount was about 80,000 NOK, or about 10,000 USD. Following Havnes and Mogstad (2011a), parents are defined as employed if they earn more than two basic amounts and full-time equivalent if they earn more than four basic amounts.

Because the childcare year starts in August, the impact on parental labor supply may materialize both in the fall of the application year and in the spring of the following year. We have therefore estimated the impact on outcomes in both years. Estimates from the reduced form are reported in table 10. While estimates are not sufficiently precise to draw strong conclusions, they suggest that receiving a lottery offer may modestly increase labor supply of mothers around the margin of full-time employment but has little impact on employment of fathers. Also, effects on fathers earnings go in the expected direction, pointing to a positive impact of a lottery offer on the intensive margin of

Table 10
Estimates of the Impact of a Lottery Offer on Parental Labor Supply in the Application Year and in the Following Year

	Application Year			Following Year		
	b	SE	Mean	*b*	SE	Mean
Mother:						
Earnings	1,828	10,141	333,887	8,941	13,944	383,061
Employment	−.009	.020	.848	−.011	.017	.880
Full-time equivalent	.024	.026	.583	.044	.024	.692
Father:						
Earnings	45,218	43,340	606,905	74,575	51,882	643,043
Employment	−.014	.013	.946	−.004	.014	.952
Full-time equivalent	−.010	.020	.852	−.015	.017	.878

Note.—Effects on earnings are reported in NOK (1 USD ≈ 8 NOK). Standard errors are clustered at the first-choice institution level and are robust to heteroskedasticity. All regressions control for the risk set by including a full set of dummy variables for each institution by year and the number of childcare institutions listed in the application (see Sec. III). Control variables are listed in Sec. III.

fathers labor supply. All in all, however, effects are moderate and, we believe, unlikely to explain our main estimates. The modest estimates may also be interpreted to imply that parents in our sample have access to informal sources of care in case they do not get a place in childcare, at least part time, in line with the anecdotal evidence discussed above.

VI. Concluding Remarks

Childcare enrollment of young children is substantial and growing, and childcare is often heavily subsidized by the government. At the same time, there is concern among both researchers and policy makers that separation from the primary caregiver, typically the mother, may cause stress and anxiety in the child, with potentially adverse effects on children's development (Bowlby 1969; Mercer 2006). Yet convincing evidence on how childcare affects the development of toddlers remains scarce.

In this paper, we present evidence on the impact of early childcare enrollment on the cognitive performance of children at age 7. Results indicate that early childcare enrollment has a beneficial effect for children's performance both on a language test and on a mathematics test. This is true both when we consider the mean test score and when we consider the probability of scoring below thresholds for low performance set by the national government.

To investigate further the potential for childcare in alleviating differences between children, we estimate the impact of receiving a lottery offer on the performance of children in subsamples depending on their gender and family background. The estimates suggest that childcare may particularly improve the development of low-performing children. On the one hand, the estimated impacts are larger in subsamples of children that tend to perform worse on the test. Our estimates suggest, for instance, that the gap in language performance between children from low- and high-educated families is halved when both attend childcare. On the other hand, the estimated impacts are clustered in the lower end of the test score distribution, while our estimates become modest and insignificant in the middle and particularly in the upper end of the distribution. This is true both on the language test and on the mathematics test. It should be noted, however, that our outcome measure is not designed to capture well variation in the upper part of the distribution. One should therefore be careful in applying our estimates on the upper end of the distribution in other contexts. These patterns are in line with a story where educated and high-income families provide a more stimulating environment, at home or in informal care, compared with more disadvantaged families (as suggested, e.g., by the seminal study of Hart and Risley [1995]). The childcare center, if it provides a more homogeneous environment to children, may then give bigger gains to children from disadvantaged families. This is also in line with the evidence in Havnes and Mogstad (2015).

Our results on how childcare affects the development of toddlers extends the growing recent literature on how childcare institutions affect the development of preschool-age children. While results are mixed, several studies have shown positive effects, in particular for children from disadvantaged families. Our study shows that positive effects of childcare are not unique to preschool children but can be extended to toddlers below 18 months of age. Importantly, our estimates lend no support to the concerns about childcare at early ages having a detrimental impact.

Appendix

Additional Material

A1. Tables and Figures

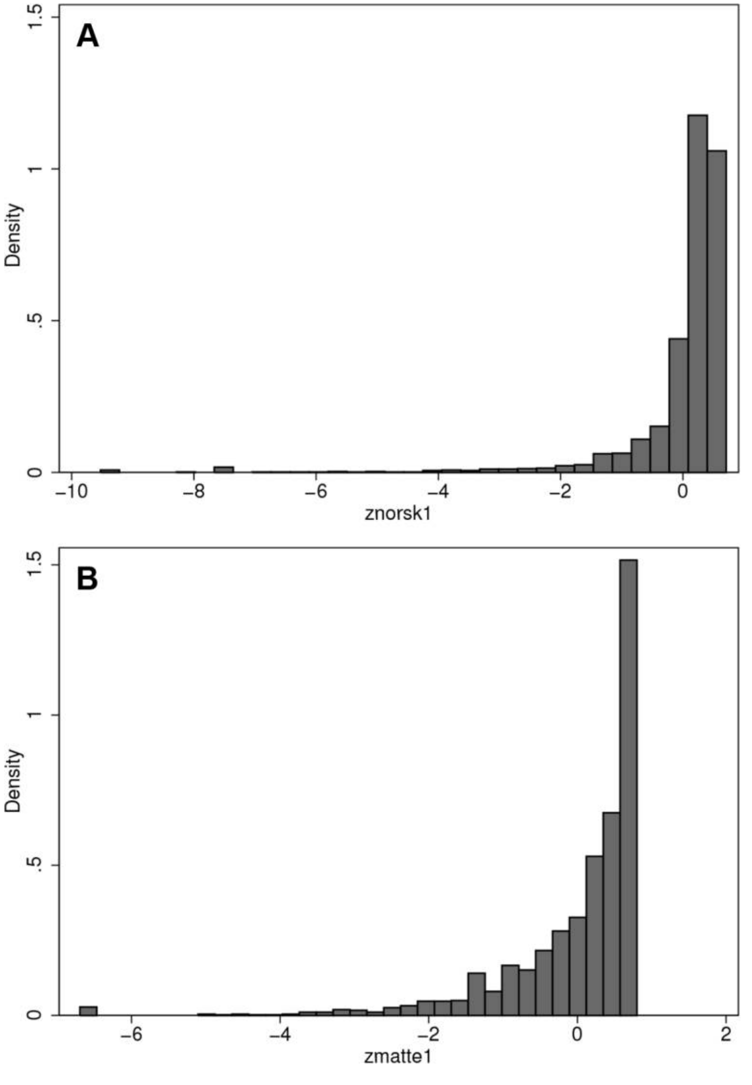

FIG. A1.—Distribution of test scores among children who did not receive an offer in the assignment lottery. *A*, Language. *B*, Mathematics.

Table A1
Background Characteristics by Number of Institutions to Which a Family Applied

	Applied to 1 Institution	Applied to 2–4 Institutions	Applied to 5–7 Institution
Girl	.53 (.50)	.47 (.50)	.51 (.50)
Age	15.56 (2.27)	14.84 (2.22)	14.67 (2.16)
Immigrant	.35 (.48)	.24 (.43)	.05 (.23)
Mother:			
Years of education	13.01 (3.51)	13.63 (3.34)	15.06 (2.88)
Earnings	211,916 (148,572)	261,565 (173,629)	315,170 (156,815)
Age	33.39 (5.07)	33.05 (5.17)	33.36 (4.10)
Age first birth	27.81 (5.04)	28.33 (4.86)	30.31 (4.20)
Father:			
Years of education	13.38 (3.67)	13.59 (3.74)	14.76 (3.29)
Earnings	404,797 (371,221)	370,251 (356,906)	430,743 (371,399)
Age	36.96 (7.96)	35.68 (7.34)	35.18 (5.82)
Age first birth	30.27 (5.94)	30.15 (6.41)	31.69 (5.44)
N	140	564	2,184

Table A2
Characteristics of Public versus Private Childcare Centers

	Mean (SD)	
	Public	Private
A. Structural characteristics:		
Language	98.14 (6.56)	99.46 (5.93)
Mathematics	44.27 (4.03)	45.09 (3.80)
Teacher/child	0.08 (0.03)	0.09 (0.06)
Distance (km)	2.84 (3.65)	3.63 (3.74)
Distance (min)	5.22 (5.48)	7.13 (6.68)
B. Staff characteristics:		
Income	260,547 (31,642)	219,405 (63,015)
College graduates	0.35 (0.12)	0.35 (0.24)
Immigrants	0.23 (0.16)	0.26 (0.28)
Males	0.09 (0.10)	0.13 (0.17)
Age	39.30 (5.31)	31.65 (7.11)
C. Peer characteristics:		
Family income	705,425 (247,601)	798,225 (256,022)
College graduates	0.43 (0.24)	0.56 (0.27)
Immigrants	0.11 (0.15)	0.04 (0.10)
Males	0.50 (0.20)	0.49 (0.24)
Young	0.46 (0.22)	0.54 (0.26)

Table A3
Background Characteristics by Sample Restrictions

	All Children	All Applying	Lottery Sample
Girl	.48	.47	.50
Age	21.64	15.96	17.69
Immigrant	.28	.07	.11
Mother:			
Years of education	13.06	14.69	14.68
Earnings	227,883	297,943	299,696
Age	32.73	33.11	33.30
Age first birth	28.66	30.04	29.80
Father:			
Years of education	13.08	14.30	14.42
Earnings	337,624	402,660	417,671
Age	34.11	33.68	35.36
Age first birth	29.64	30.27	31.32
N	26,129	13,250	2,888

NOTE.—The table reports means of covariates by different sample restrictions; all children in Oslo belonging to the relevant cohorts (first column), all children applying for a childcare slot (second column), and finally the lottery sample (third column). Detailed descriptions of the background characteristics are provided in Sec. IV.

Table A4
Instrumental Variable Estimates of the Impact of Childcare Starting Age on Test Performance in Different Subgroups

	A. Language			B. Mathematics			C. Childcare Starting Age			
	b	SE	Mean	*b*	SE	Mean	*b*	SE	Mean	*N*
Boys	−.066	.024	−.119	−.027	.020	.017	−3.167	.680	18.644	1,442
Girls	−.036	.016	.056	−.044	.019	−.060	−3.587	.695	18.590	1,446
Parents' education low	−.060	.021	−.182	−.028	.019	−.201	−4.038	.657	19.219	1,551
Parents' education high	−.027	.017	.137	−.015	.020	.179	−2.984	.552	17.944	1,337
Mom age low	−.067	.019	−.135	−.040	.017	−.081	−4.088	.681	18.879	1,596
Mom age high	−.016	.017	.093	−.016	.023	.050	−2.695	.573	18.301	1,292
Family income low	−.065	.021	−.127	−.063	.018	−.166	−4.082	.712	19.535	1,444
Family income high	−.0284	.017	.062	−.003	.019	.120	−2.990	.633	17.717	1,444

NOTE.—Panels A and B report two-stage least-squares estimates from eq. (2). Panel C reports first-stage estimates from eq. (3). Effects are reported as percentage of standard deviation. Standard errors are clustered at the first-choice institution level and are robust to heteroskedasticity. All regressions control for the risk set by including a full set of dummy variables for each institution by year and the number of childcare institutions listed in the application (see Sec. III). Control variables are listed in Sec. III.

Table A5
Reduced-Form Estimates of the Impact of a Lottery Offer on Performance in Language and Mathematics at 19 Quantiles

	Language			Mathematics		
Percentage	*b*	SE	Score	*b*	SE	Score
5	.030	.011	−1.408	.020	.012	−1.87
10	.049	.015	−.651	.028	.016	−1.168
15	.054	.018	−.335	.045	.018	−.78
20	.048	.020	−.116	.043	.020	−.523
25	.038	.023	−.046	.048	.022	−.274
30	.024	.024	.105	.027	.024	−.129
35	.024	.024	.110	.035	.026	.016
40	.027	.024	.185	.035	.027	.161
45	.027	.024	.240	.044	.028	.212
50	.017	.026	.256	.035	.028	.352
55	.017	.026	.261	.035	.028	.409
60	.026	.027	.336	.021	.027	.452
65	.026	.027	.336	.004	.027	.533
70	.026	.027	.407	−.008	.025	.597
75	.019	.022	.411	.007	.023	.606
80	.014	.021	.431	.020	.022	.645
85	.012	.020	.487	.011	.018	.742
90	.005	.012	.559	−.003	.014	.791
95	.005	.012	.559	.000	.000	.803

NOTE.—The dependent variable in each row is a dummy for performing above the percentile of the test score indicated by the row. Estimates are from eq. (1), estimated at each of 19 quantiles, controlling for the risk set by including a full set of dummy variables for each institution by year and the number of childcare institutions listed in the application (see Sec. III). Standard errors are clustered at the first-choice institution level and are robust to heteroskedasticity.

A2. Background Information about the Language Test

Every teacher who is responsible for carrying out tests in his or her class receives a teacher's instruction manual. The following text about the test is from this manual (Oslo Department of Education 2011b).

Conditions for Learning
Test Part 1

The student's attitude when it comes to reading. This part should provide information about the student's attitude and interests related to different activities related to the written language. Results on this part of the test should not been given a score but are meant as information to the teacher.

Level of Knowledge about the Alphabet
Test Part 2

To write letters. This measures the student's ability to link sound and letter as well as their ability to construct the letters in question. First, the students hear a word. Second, the teacher repeats the first sound of the word and asks the students to write the letter that goes with that sound. There is a picture supporting the word in the student book. The students' results on this test should be scored by the teacher.

Test Part 3

To recognize letters. This part measures one of the basic skills in reading. Students have 4 minutes at their disposal. Starting from a capital letter the students shall recognize the same lowercase letter among several other lowercase letters. The students' results on this test should be scored by the teacher. This test part does not have a "critical threshold" since many students may have been exposed to only one type of letter throughout the first school year. The score registrations are meant as information to the teacher about which of the capital letters and lowercase letters the students can recognize and link.

Test Part 4

To identify the initial sound. This measures the students' ability to do exactly this. The teacher reads a word and asks the students to identify the first sound of the word and write it down. There is a picture supporting the word in the student book. The students' results on this test should be scored by the teacher.

Test Part 5

To draw together sounds. Maps the students' abilities in phonological synthesis. Each part of this test contains a sequence with four pictures that illustrate different words. The teacher instructs the students by first presenting the word that illustrates each picture and thereafter the target word, sound by sound, with a break between every sound. The task of the students is to carry through the synthesis process and determine which picture that goes with the target word. The students' results on this test should be scored by the teacher.

Understanding Words

Test Part 6

To write words. Consists of a word dictation where each word is presented for the students in a sentence. This test part comprises 8 sentences in total. The students' results on this test should be scored by the teacher.

Test Part 7

To read words. To read words implies that the students should compare an illustration with four written words and subsequently identify the word that fits with the illustration. The students should identify as many words as possible (total possible words, 19) within 5 minutes. The students' results on this test should be scored by the teacher.

Understanding Sentences
Test Part 8

To read sentences. To read sentences consists of nine subparts. In each part the student reads a sentence and marks the picture that illustrates the entire content of the sentence among four alternative pictures with similar content. The length of the sentences increases from two to five words as the test proceeds. The students should link as many pictures and sentences as possible within 5 minutes. The students' results on this test should be scored by the teacher.

A3. Background Information about the Mathematics Test

Every teacher responsible for carrying out tests in his or her class receives a teacher's instruction manual. The following text about the test is from this manual (Oslo Department of Education 2011a).

This test consists of nine pages with several different tasks (see below). The points scored on each page should be added together. The critical threshold for the mathematics test is based on the aggregated sum of points.

Page 1. Maps the students' ability to count and determines whether they know the numbers and can link a number of items to a certain figure.

Page 2. Determines whether the students understand the idea "equally many," that is, that they can compare the number in two different countable sets.

Page 3. Investigates whether the students' can rank the numbers in two different countable sets and whether they understand the concept "most" (see fig. A2).

Page 4. Maps the students' knowledge about a sequence of numbers (linear).

Page 5. Maps the students' knowledge about a sequence of numbers (linear), determines whether they can count forward and backward from a given number, and determines whether they understand the sequence concepts "prior to" and "subsequent to."

Page 6. Investigates the students' knowledge of a series of numbers and the ranking of given numbers within the series.

Page 7. Tests the students' ability to split a number into two other numbers (e.g., $4 = 1 + \ldots$).

Page 8. Tests the students' ability to solve text assignments. The assignments have different additive structures, and they deal with both addition and subtraction (but rather low numbers).

Page 9. Maps the students' ability to add two numbers. The additive structure deal with the combination of two sums of money, to allow the students to employ their knowledge about money and coins to arrive at the correct answer.

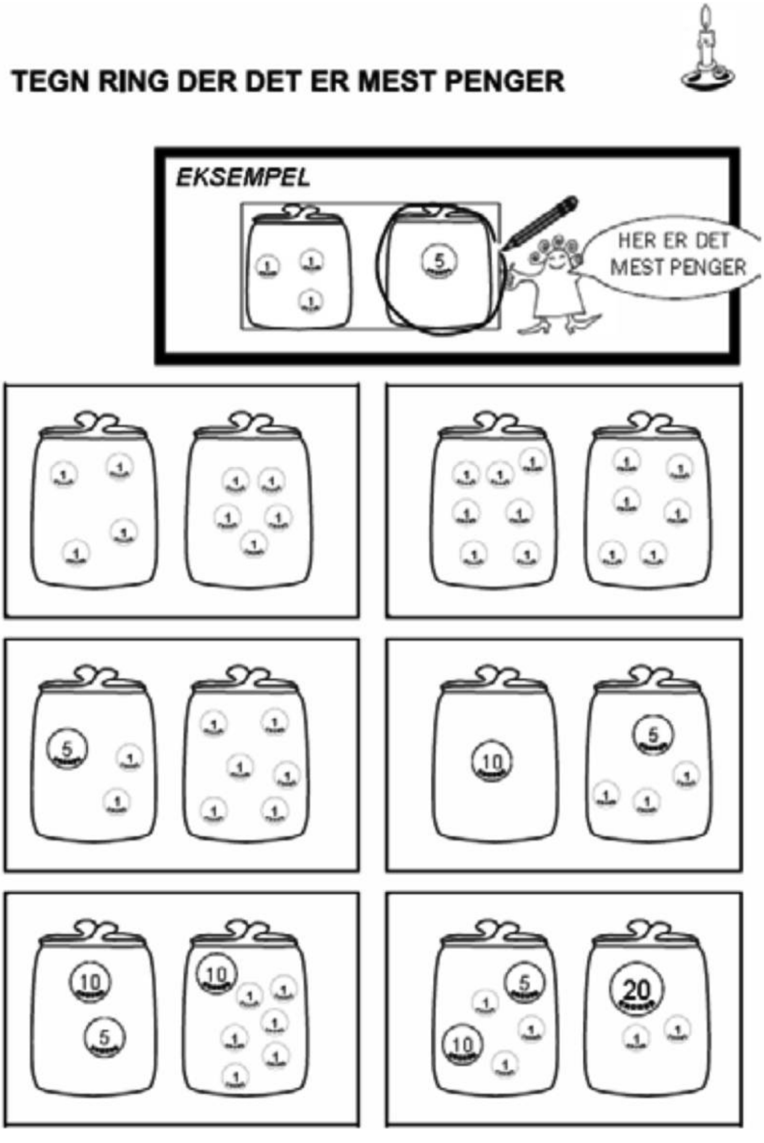

FIG. A2.—Exercise from page 3 of the math test. In the exercise, the children are asked to draw a circle around the wallet with the most money in each of the six rectangles. The top box is precompleted as an example. Direct translations: "Draw a ring where there is the most money," "Example," and "Here is the most money." Illustration by Inger Landsem.

References

Abdulkadiroglu, Atila, Joshua D. Angrist, Susan M. Dynarski, Thomas J. Kane, and Parag A. Pathak. 2011. Accountability and flexibility in public schools: Evidence from Boston's charters and pilots. *Quarterly Journal of Economics* 126, no. 2:699–748.

Almond, Douglas, and Janet Currie. 2010. Human capital development before age five. NBER Working Paper no. 15827, National Bureau of Economic Research, Cambridge, MA.

Atkinson, A. B., Lee Rainwater, and Timothy M. Smeeding. 1995. *Income distribution in OECD countries: Evidence from the Luxembourg income study*. Paris: OECD Publications and Information Center.

Baker, Michael. 2011. Innis lecture: Universal early childhood interventions: What is the evidence base? *Canadian Journal of Economics* 44, no. 4:1069–105.

Baker, Michael, Jonathan Gruber, and Kevin Milligan. 2008. Universal child care, maternal labor supply, and family well-being. *Journal of Political Economy* 116, no. 4:709–45.

Barnett, W. Steven. 1995. Long-term effects of early childhood programs on cognitive and school outcomes. *Future of Children* 5, no. 3:22–50.

Berlinski, Samuel, and Sebastian Galiani. 2007. The effect of a large expansion of pre-primary school facilities on preschool attendance and maternal employment. *Labour Economics* 14, no. 3:665–80.

Berlinski, Samuel, Sebastian Galiani, and Marco Manacorda. 2008. Giving children a better start: Preschool attendance and school-age profiles. *Journal of Public Economics* 92, no. 5/6:1416–40.

Blau, David, and Janet Currie. 2006. Pre-school, day care, and after-school care: Who's minding the kids? In *Handbook of the economics of education*, vol. 2, chap. 20, 1163–1278. Amsterdam: Elsevier.

Bowlby, John. 1969. *Attachment and loss*, vol. 1. *Attachment*. London: Hogarth.

Carneiro, Pedro, Katrine V. Løken, and Kjell G. Salvanes. 2015. A flying start? Maternity leave benefits and long-run outcomes of children. *Journal of Political Economy* 123, no. 2:365–412.

Cascio, Elizabeth U. 2009. Do investments in universal early education pay off? Long-term effects of introducing kindergartens into public schools. Working Paper no. 14951, National Bureau of Economic Research, Cambridge, MA.

Cunha, Flavio, and James J. Heckman. 2008. Formulating, identifying and estimating the technology of cognitive and noncognitive skill formation. *Journal of Human Resources* 43, no. 4:738–82.

Currie, Janet. 2001. Early childhood education programs. *Journal of Economic Perspectives* 15:213–38.

Currie, Janet, and Duncan Thomas. 1995. Does Head Start make a difference? *American Economic Review* 85, no. 3:341–64.

DeCicca, Phillip, and Justin Smith. 2013. The long-run impacts of early childhood education: Evidence from a failed policy experiment. *Economics of Education Review* 36:41–59.

Deming, David. 2009. Early childhood intervention and life-cycle skill development: Evidence from Head Start. *American Economic Journal: Applied Economics* 1, no. 3:111–34.

Drange, Nina, and Kjetil Telle. 2015. Promoting integration of immigrants: Effects of free child care on child enrollment and parental employment. *Labour Economics* 34:26–38.

Dustmann, Christian, Anna Raute, and Uta Schønberg. 2013. Does universal childcare matter? Evidence from a large expansion in pre-school education. Working paper, University College London.

Dustmann, Christian, and Uta Schønberg. 2012. Expansions in maternity leave coverage and children's long-term outcomes. *American Economic Journal: Applied Economics* 4, no. 3:190–224.

Felfe, Christina, and Rafael Lalive. 2014. Does early child care help or hurt children's development? IZA Discussion Paper no. 8484, Institute of Labor Economics, Bonn.

Felfe, Christina, Natalia Nollenberger, and Nuria Rodriguez-Planas. 2015. Can't buy mommy's love? Universal childcare and children's long-term cognitive development. *Journal of Population Economics* 28, no. 2:393–422.

Fitzpatrick, Maria D. 2008. Starting school at four: The effect of universal pre-kindergarten on children's academic achievement. *B. E. Journal of Economic Analysis and Policy* 8, no. 1:46.

Fort, Margherita, Andrea Ichino, and Giulio Zanella. 2016. Cognitive and non-cognitive costs of daycare 0–2 for girls. IZA Discussion Paper no. 9756, Institute of Labor Economics, Bonn.

Garces, Eliana, Duncan Thomas, and Janet Currie. 2002. Longer-term effects of Head Start. *American Economic Review* 92, no. 4:999–1012.

Gupta, Nabanita Datta, and Marianne Simonsen. 2010. Non-cognitive child outcomes and universal high quality child care. *Journal of Public Economics* 94, no. 1/2:30–43.

Hart, Betty, and Todd R. Risley. 1995. *Meaningful differences in the everyday experience of young American children*. Towson, MD: Paul H. Brookes.

Havnes, Tarjei, and Magne Mogstad. 2011a. Money for nothing? Universal child care and maternal employment. *Journal of Public Economics* 95, no. 11/12:1455–65.

———. 2011b. No Child Left Behind: Subsidized child care and children's long-run outcomes. *American Economic Journal: Economic Policy* vol. 3, no. 2:97–129.

———. 2015. Is universal child care leveling the playing field? *Journal of Public Economics* 127:100–114.

Herbst, Chris M. 2013. The impact of non-parental child care on child development: Evidence from the summer participation dip. *Journal of Public Economics* 105:86–105.

Karoly, Lynn A., M. Rebecca Kilburn, and Jill S. Cannon. 2005. *Early childhood interventions: Proven results, future promise*. Santa Monica, CA: RAND.

Lefebvre, Pierre, and Philip Merrigan. 2008a. Child-care policy and the labor supply of mothers with young children: A natural experiment from Canada. *Journal of Labor Economics* 26, no. 3:519–48.

———. 2008b. Family background, family income, cognitive tests scores, behavioural scales and their relationship with post-secondary education participation: Evidence from the NLSCY. Cahiers de recherche 0830, CIRPEE.

Liu, Qian, and Oskar Nordstrom Skans. 2010. The duration of paid parental leave and children's scholastic performance. *B. E. Journal of Economic Analysis and Policy* 10, no. 1:1–35.

Ludwig, Jens, and Douglas L. Miller. 2007. Does Head Start improve children's life chances? Evidence from a regression discontinuity design. *Quarterly Journal of Economics* 122, no. 1:159–208.

Lundin, Daniela, Eva Mork, and Bjorn Ockert. 2008. How far can reduced childcare prices push female labour supply? *Labour Economics* 15, no. 4: 647–59.

McKey, Ruth Hubbell, Larry Condelli, Harriet Ganson, Barbara J. Barrett, Catherine McConkey, and Margaret Plantz. 1985. The impact of Head Start on children, families and communities. Final report of the Head Start Evaluation, Synthesis and Utilization Project. Technical Report, US Department of Health and Human Services.

Mercer, Jean. 2006. *Understanding attachment: Parenting, child care, and emotional development*. Westport, CT: Praeger.

OECD (Organization for Economic Cooperation and Development). 2006. Starting strong. II: Early childhood education and care. Technical report.

Oslo Department of Education. 2011a. Kartlegging av tallforståelse og regneferdighet på 1. trinn. laererveiledning.

———. 2011b. Obligatorisk kartleggingsprøve: Kartlegging av leseferdighet 1. årstrinn del 1: Laererveiledning.

Pedersen, Silje Vatne. 2003. Barnefamiliers tilsynsordninger, yrkesdeltakelse og bruk av kontantstøtte våren 2002. Report 3, Statistics Norway.

Ruhm, Christopher J., and Jane Waldfogel. 2012. Long-term effects of early childhood care and education. *Nordic Economic Policy Review* 1:23–51.

Schlosser, Analia. 2005. Public preschool and the labor supply of Arab mothers: Evidence from a natural experiment. Technical report, Hebrew University of Jerusalem.