

Inputs in the Production of Early Childhood Human Capital: Evidence from Head Start[†]

By CHRISTOPHER R. WALTERS*

This paper uses data from a randomized evaluation of Head Start to answer two questions: (i) How much do short-run causal effects vary across Head Start centers? and (ii) Do observed inputs explain this variation? I find that the cross-center standard deviation of cognitive effects is 0.18 test score standard deviations, which is larger than typical estimates of variation in teacher or school effectiveness. Centers offering full-day service and home visiting are more effective, while centers that draw more children from center-based preschool have smaller effects. Other key inputs, including the High/Scope curriculum, teacher education, and class size are not correlated with Head Start effectiveness. (JEL H75, I21, I28, J13, J24)

Studies of small-scale “model” early childhood education programs show that preschool attendance can boost outcomes in the short and long run. In the High/Scope Perry Preschool Project, a randomized trial that took place in the early 1960s, 123 disadvantaged children were randomly assigned to either an intensive preschool program or a control group without access to the program. Subsequent analyses showed that participation in the Perry program increased average IQ at age 5 by nearly a full standard deviation, and had lasting impacts on educational attainment, criminal behavior, drug use, employment, and earnings (Anderson 2008; Berruta-Clement et al. 1984; Heckman et al. 2010a; Heckman, Pinto, and Savelyev 2013; Schweinhart and Weikart 1997; Schweinhart et al. 2005).¹ Heckman et al. (2010b) estimate the annual social rate of return to the Perry Project at between 7 and 10 percent. The North Carolina Abecedarian Project, another small-scale intervention, had similarly dramatic effects (Campbell and Ramey 1994, 1995). The striking success of these programs has led some analysts to argue that the returns to educational intervention peak early in life (Heckman 2011). These findings have

*Department of Economics, University of California-Berkeley, 530 Evans Hall #3880, Berkeley, California 94720, (e-mail: crwalters@econ.berkeley.edu). I am grateful to Joshua Angrist, Aviva Aron-Dine, David Autor, David Card, David Chan, Hilary Hoynes, Guido Imbens, Patrick Kline, Enrico Moretti, Christopher Palmer, Parag Pathak, Jesse Rothstein, Tyler Williams, two anonymous referees, and seminar participants at Massachusetts Institute of Technology, UC Berkeley, and the National Bureau of Economic Research (NBER) Education Program for useful comments and suggestions. This work was supported by Institute for Education Sciences award number R305A120269 and an National Academy of Education/Spencer Dissertation Fellowship.

[†]Go to <http://dx.doi.org/10.1257/app.20140184> to visit the article page for additional materials and author disclosure statement(s) or to comment in the online discussion forum.

¹Anderson (2008) argues that the Perry Project produced significant long-term benefits only for girls.

also motivated recent calls for expansion of publicly provided preschool (Obama 2013).

In contrast, evidence on the effects of large-scale early childhood programs is more mixed. Early quasi-experimental studies of Head Start, the largest early childhood program in the United States, showed positive effects on cognitive skills, child mortality, and long-term outcomes (Currie and Thomas 1995; Ludwig and Miller 2007; Garces, Thomas, and Currie 2002; Deming 2009).² More recently, results from the Head Start Impact Study (HSIS), the first randomized evaluation of Head Start, showed smaller, less persistent gains. The HSIS experiment involved random assignment of more than 4,000 children to Head Start or a control group at over 300 childcare centers throughout the United States. The HSIS treatment group outscored the control group by roughly 0.1 standard deviations on measures of cognitive skill during preschool, but these gains did not persist into kindergarten (Puma et al. 2010, 2012). Moreover, the HSIS experiment showed little evidence of effects for a wide range of noncognitive and health outcomes (Puma et al. 2010).³

Inputs and practices vary widely across Head Start centers, however, and little is known about variation in effectiveness within Head Start. This paper uses HSIS data to quantify and explain variation in causal effects across Head Start childcare centers, with an eye toward reconciling the effects of model programs and those of Head Start. Specifically, I assess the role that inputs, practices, and child characteristics play in generating differences in effectiveness across Head Start centers. Some centers use inputs more similar to successful model programs than others. For example, one-third of Head Start centers use the High/Scope curriculum, the centerpiece of the Perry Preschool experiment. Head Start centers also differ with respect to teacher characteristics, class size, instructional time, frequency of home visits, and instructor experience, all of which have been cited as central to the success of model programs (Schweinhart 2007; Chetty et al. 2011). In addition, the characteristics of Head Start applicants and the availability of alternative preschool options vary across centers. The aim of this paper is to assess the contribution of these key inputs and characteristics to cross-center differences in Head Start effects.

My analysis proceeds in two steps. First, to ask whether there is meaningful variation to be explained by program characteristics, I quantify heterogeneity in causal effects across Head Start centers. This investigation is complicated by non-compliance with random assignment in the HSIS experiment. Instrumental variables (IV) is the standard procedure for dealing with noncompliance, but IV has poor properties in small samples, and center-specific samples in the HSIS are small (Nelson and Startz 1990). To deal with this problem, I use a random coefficients version of the Heckman (1979) sample selection model to directly estimate the cross-site distribution of treatment effects, circumventing the need to work with

²Other studies finding positive effects of larger scale programs include analyses of the Chicago Child-Parent centers and some state pre-kindergarten programs (Reynolds and Temple 1998; Gormley and Gayer 2005; Wong et al. 2008). Cascio and Schanzenbach (2013) find small effects of programs in Georgia and Oklahoma for poor children, and no effects for richer children. Fitzpatrick (2008) finds small effects for Georgia's program, though some subgroups benefit.

³In other analyses of the HSIS data, Gelber and Isen (2013) show that Head Start participation increased parental involvement with children after the program ended, while Bitler, Hoynes, and Domina (2014) show larger quantile treatment effects at lower quantiles of the distribution of Peabody Picture and Vocabulary Test (PPVT) scores.

poorly behaved center-specific instrumental variables estimates. The random coefficients estimates reveal substantial heterogeneity in short-run Head Start effectiveness: the cross-center standard deviation of short-run cognitive effects is 0.18 test score standard deviations, larger than typical estimates of variation in teacher and school effectiveness (Deming 2014; Chetty, Friedman, and Rockoff 2014a; Kane, Rockoff, and Staiger 2008).

In a second step, I ask whether this variation can be explained by differences in observed program and child characteristics. My results show that some inputs play a role. Head Start centers offering full-day programs boost cognitive skills more than other centers, while centers offering frequent home visits are especially effective at raising noncognitive skills. High/Scope Head Start centers are no more effective than other centers, however, and short-run effects are uncorrelated with teacher education, class size, and center director experience. Short-run cognitive effects are larger for children with less educated mothers, but Head Start effectiveness is weakly related to other measures of family background and baseline skills. To investigate the role of alternative preschool options, I estimate the relationship between Head Start effectiveness and the share of children drawn from other preschools rather than home-based care. This analysis suggests that cognitive gains are smaller for centers that draw more children from center-based preschool. Together, observed inputs, practices, and child characteristics explain about one-third of the variation in Head Start effectiveness.

An important caveat to these findings is that inputs are not randomly assigned to Head Start centers. While the experimental variation used here eliminates selection bias in comparisons of students offered and not offered Head Start, centers with different observed characteristics may differ systematically on unobserved dimensions. As a result, relationships between inputs and effectiveness may not reflect causal impacts of changing inputs in isolation. Nevertheless, these relationships are important for two reasons. First, observed predictors of program effectiveness can help policymakers to identify high- and low-performing programs. The ability to target high or low performers is useful for policies that aim to expand effective programs or improve ineffective ones. Second, my estimates of the relationships between inputs and impacts show that some key inputs used by model programs are not sufficient to create effective preschools. For example, Schweinhart (2007) argues that the High/Scope curriculum was central to the success of the Perry Preschool Project. I find that High/Scope is not related to program effectiveness in Head Start. This shows that the High/Scope curriculum alone does not guarantee a successful preschool program.

In addition to the literature on preschool effects, this paper contributes to several other strands of research. A recent series of studies relates variation in effectiveness across education programs, including charter schools, kindergarten classrooms, and teachers to observed program characteristics (Kane, Rockoff, and Staiger 2008; Chetty et al. 2011; Hoxby and Murarka 2009; Angrist, Pathak, and Walters 2013; Dobbie and Fryer 2013). I apply a similar approach to study the relationship between inputs and Head Start effects. Hotz, Imbens, and Mortimer (2005); Raudenbush, Reardon, and Nomi (2012); and Allcott (2014) analyze variation in effects across sites in multisite randomized controlled trials, while Chandra et al. (2013) and

Syverson (2011) use empirical Bayes and random coefficients methods to measure variation in productivity across hospitals and other firms. The analysis here includes elements of each of these approaches.

The rest of the paper is organized as follows. The next section provides background on Head Start and describes the HSIS data. Section II summarizes the average impact of Head Start on summary indices of cognitive and noncognitive skills. Section III outlines the random coefficients model used to investigate effect heterogeneity, and reports the results of this investigation. Section IV analyzes the link between Head Start effectiveness and observed inputs, practices, and child characteristics. Section V concludes.

I. Data and Background

A. *Head Start and the Head Start Impact Study*

Head Start, the largest early childhood program in the United States, enrolls roughly one million 3- and 4-year-old children at a cost of about \$8 billion annually. The program awards grants to public, private, nonprofit, and for-profit organizations that provide childcare services to children below the federal poverty line, though up to 35 percent of children attending a Head Start childcare center can come from households between 100 and 135 percent of this income threshold. Grantees are required to match at least 20 percent of federal Head Start funding. Head Start is based on a “whole child” model of school readiness that emphasizes noncognitive social and emotional development in addition to cognitive skills. The grant-based nature of the program allows for a wide variety of childcare settings and practices, though all grantee agencies must meet a set of program-wide performance standards (US Department of Health and Human Services 2011, 2012).

The data used here come from the HSIS, a randomized evaluation of the Head Start program. The 1998 Head Start Reauthorization Act included a congressional mandate to determine the programs effects. As a result, the US Department of Health and Human Services (DHHS) conducted a nationally representative randomized controlled trial (Puma et al. 2010, 2012). The HSIS data includes information on 84 regional Head Start programs, 353 Head Start centers, and 4,442 children, each of whom applied to a sample Head Start center in Fall 2002. Sixty percent of applicants were randomly assigned the opportunity to attend Head Start (“treatment”), while the remaining applicants were denied this opportunity (“control”). Randomization took place at the Head Start center level; the HSIS data includes weights reflecting the probability of assignment for each child, which are used to adjust for these differences below.⁴

The HSIS sample includes 2 age groups, with 55 percent of students entering at age 3 and 45 percent entering at age 4. Three-year-old applicants could attend Head

⁴Some small centers were aggregated together to conduct the random assignment. Other centers conducted multiple rounds of random assignment with differing admission probabilities, and the HSIS weights do not account for these differences. The discussion in DHHS (2010) suggests that any such differences are likely to be small, however.

TABLE 1—OUTCOMES INCLUDED IN SUMMARY INDICES

Cognitive skills (1)	Noncognitive skills (2)
Peabody Picture and Vocabulary Test III (PPVT)	Takes care of personal things
Color names	Asks for assistance with tasks
Test de Vocabulario en Imágenes Peabody (TVIP adapted)	Makes friends easily
Woodcock-Johnson III Oral Comprehension	Enjoys learning
Preschool Comprehensive Test of Phonological and Print Processing (CTOPPP)	Has temper tantrums
Spanish CTOPPP	Cannot concentrate/pay attention for long
Woodcock-Johnson III Word Attack	Is very restless/fidgets a lot
McCarthy Draw-a-design	Likes to try new things
Letter naming	Shows imagination in work and play
Woodcock-Johnson III Letter-Word Identification	Hits and fights with others
Batería R. Woodcock-Muñoz Identificación de Letras y Palabras	Accepts friends' ideas in playing
Woodcock-Johnson III Spelling	
Batería R. Woodcock-Muñoz Dictado	
Woodcock-Johnson III Applied Problems	
Woodcock-Johnson III Quantitative Concepts	
Counting Bears	
Batería R. Woodcock-Muñoz Problemas Aplicados	

Notes: This table lists the cognitive and noncognitive outcomes used in the analysis. Summary indices are averages of standardized outcomes in each category.

Start for up to two years before entering kindergarten, and three-year-olds assigned to the control group could reapply to Head Start centers as four-year-olds the next year. Four-year-old applicants could attend for a maximum of one year. The data used here follow the treatment and control groups through first grade. Puma et al. (2010) provide a complete description of the HSIS experimental design and data collection procedures. The online Appendix details the procedure used to construct my sample from the HSIS data.

B. Outcomes

The HSIS data include a large number of outcomes, collected for up to four years after random assignment. I organize these outcomes into summary indices of cognitive and noncognitive skills. Table 1 lists the outcomes included in each group. Cognitive outcomes include scores on the Peabody Picture and Vocabulary Test (PPVT) and several Woodcock Johnson III (WJIII) measures of cognitive ability. Noncognitive outcomes, derived from parental surveys, include measures of social skills (making friends, hitting, and fighting) and attention span (concentration,

restlessness). I exclude noncognitive measures for which almost all respondents (90 percent or more) gave the same answer.⁵

Following Kling, Liebman, and Katz (2007) and Deming (2009), I construct indices to summarize the impact of Head Start attendance across the outcomes listed in each column of Table 1. Specifically, I define the summary index

$$(1) \quad Y_i \equiv \frac{1}{L} \sum_{\ell=1}^L \left(\frac{y_{i\ell} - \mu_{\ell}}{\sigma_{\ell}} \right),$$

where $y_{i\ell}$ is outcome ℓ for student i , and μ_{ℓ} and σ_{ℓ} are the control group mean and standard deviation of this outcome. I define outcomes so that positive signs mean better performance, and standardize them separately by year and age cohort.

C. Applicant Characteristics

Head Start applicants typically come from families with low socioeconomic status. This can be seen in the first column of Table 2, which presents mean demographic characteristics for the HSIS control group. The demographic variables come from a baseline survey of parents conducted in the Fall of 2002; parents of 3,577 HSIS applicants (81 percent) responded to this survey. The Head Start population is disadvantaged on observable dimensions: roughly two-thirds of children in the sample are nonwhite, and about half live in two-parent households. Thirty-nine percent of mothers in the sample did not complete high school, and 17 percent are teenagers. The average household income in the sample is \$1,507 per month.⁶

To check experimental balance, column 2 of Table 2 shows coefficients from regressions of baseline characteristics on assignment to Head Start, weighting by the HSIS baseline child weights to adjust for differences in the probability of assignment across centers. The treatment/control differences in means are statistically insignificant for all baseline variables except special needs status, and the joint p -value from a test of the hypothesis that assignment to Head Start is unrelated to all characteristics is 0.31. This suggests that random assignment was successful.⁷

The last two rows of Table 2 show the effects of assignment to Head Start on applicants' preschool choices. Applicants assigned to Head Start were 66 percentage points more likely to participate in the program than applicants from the control group in the first year after random assignment. Sixteen percent of students from the control group attended Head Start, most likely by applying to other nearby Head

⁵The HSIS data also includes measures of noncognitive skills reported by teachers. I do not use these measures since they are unavailable for many children before kindergarten, and my analysis focuses on outcomes during preschool.

⁶The parent survey includes two questions about household income. One question asks for exact monthly income. For parents who do not answer this question, a follow-up question asks where income falls in a set of possible categories. For parents who answer the second question, I impute income as the midpoint of the reported range.

⁷Even with successful random assignment, nonrandom attrition has the potential to bias the experimental results. Online Appendix Table A1 shows attrition rates for the HSIS sample by year and outcome group, as well as treatment/control differences conditional on the controls included in Table 4. In preschool, outcomes are observed for 82 to 84 percent of children; the follow-up rate falls slightly in elementary school. Cognitive outcomes in preschool are observed slightly more frequently for children in the treatment group (3 to 5 percentage points). This modest differential attrition seems unlikely to drive the results reported below.

TABLE 2—CHARACTERISTICS OF HEAD START APPLICANTS

Variable	Control mean (1)	Offer differential (2)
Male	0.490	0.011 (0.023)
Black	0.259	0.009 (0.011)
Hispanic	0.411	0.000 (0.014)
Home language is Spanish	0.332	−0.013 (0.014)
Special needs	0.112	0.020* (0.011)
Mother is married	0.478	−0.016 (0.020)
Both parents live at home	0.531	−0.016 (0.020)
Teen mother	0.165	−0.023 (0.016)
Mother is high school dropout	0.389	−0.022 (0.016)
Mother attended college	0.281	0.020 (0.018)
Monthly household income	1,507.124	−25.060 (61.350)
Baseline cognitive skills	−0.003	0.014 (0.023)
Baseline noncognitive skills	0.001	0.033 (0.022)
Three-year-old cohort	0.534	−0.001 (0.013)
Attended Head Start in first year	0.160	0.663*** (0.023)
Attended any preschool in first year	0.460	0.442*** (0.025)
Joint <i>p</i> -value for baseline characteristics	—	0.313
Observations (total)		4,442
Observations (completed survey)		3,577

Notes: Column 1 shows means of baseline characteristics for Head Start applicants assigned to the control group. Column 2 shows coefficients from regressions of each characteristic on assignment to Head Start. The means and regressions are weighted using the HSIS baseline child weights. The *p*-value is from a test of the hypothesis that coefficients for all baseline characteristics are zero. Standard errors are clustered at the Head Start center level.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

Start centers outside the experimental sample. Eighteen percent of children assigned to Head Start did not participate in the program. Together, these facts show that noncompliance with experimental assignments is an important feature of the HSIS data, which motivates the instrumental variables approach taken below. The last row of Table 2 shows that a Head Start offer increases the probability of attending any center-based preschool program by 44 percentage points. This implies that

two-thirds (0.442/0.663) of children induced to attend Head Start by the experimental offer would not have attended preschool otherwise, while the remaining one-third would have attended another preschool center if denied the opportunity to attend Head Start.

D. Center Characteristics

In addition to background information on applicants, the HSIS data includes detailed information on Head Start centers and their practices. I focus on inputs and practices that have been cited as central to the success of small-scale model programs. Schweinhart (2007) offers one view of the inputs that drove the success of the Perry Preschool Project:

“The external validity or generalizability of the study findings extends to those programs that are reasonably similar to the High/Scope Perry Preschool Program. A reasonably similar program is a preschool education program run by teachers with bachelors degrees and certification in education, each serving up to eight children living in low-income families. The program runs two school years for children who are three and four years of age with daily classes of 2.5 hours or more, uses the High/Scope model or a similar participatory education approach, and has teachers visiting families at least every two weeks or scheduling regular parent events.”

This account of the Perry programs effects emphasizes six key inputs: teacher education, teacher certification, class size, instruction time, the High/Scope curriculum, and home visiting. High/Scope is a participatory curriculum that emphasizes hands-on choices and experiences rather than adult-driven instruction (Epstein 2007). Schweinhart (2007) places particular weight on the High/Scope curriculum, arguing that results from the Perry Project and the follow-up High/Scope Preschool Curriculum Comparison Study “[suggest] that the curriculum had a lot to do with the findings.”

No Head Start center replicates the Perry model, which used high levels of all six inputs and spent roughly 30 percent more than the average Head Start program on a per pupil, per year basis.⁸ There is substantial variation in each of the six key Perry inputs within Head Start, however. This can be seen in Table 3, which summarizes characteristics of centers in the HSIS sample. Thirty percent of Head Start centers use the High/Scope curriculum. Thirty-five percent of Head Start teachers have bachelors degrees, and 11 percent hold teaching licenses, but the fractions with these credentials range from 0 to 100 percent across centers. The average Head Start center has 6.8 children for every staff member; the cross-center standard deviation of class size is 1.7 children. Sixty-three percent of Head Start centers provide full-day service and 20 percent offer more than three home visits per year. Table 3 also reports information on years of experience for Head Start center directors; Chetty et al. (2011) cite

⁸ Heckman et al. (2010b) report that the Perry program cost about \$17,759 per child over 2 years (2006 dollars), or \$8,880 per year. Per child expenditure in Head Start was \$7,600 in 2011, which is \$6,800 deflated to 2006 dollars using the Consumer Price Index series available at <http://www.bls.gov> (DHHS 2011).

TABLE 3—CHARACTERISTICS OF HEAD START CENTERS

Variable	Head Start centers				Other centers
	Mean (1)	SD (2)	Min. (3)	Max. (4)	Mean (5)
Fraction of teachers with bachelor's degree	0.35	0.40	0.00	1.00	0.41
Fraction of staff with teaching license	0.11	0.23	0.00	1.00	0.30
Student/staff ratio	6.79	1.71	2.33	13.50	8.76
Full day service	0.63	0.48	0.00	1.00	0.67
More than three home visits per year	0.20	0.40	0.00	1.00	0.13
High/Scope curriculum	0.30	0.46	0.00	1.00	0.28
Center director experience (years)	18.17	10.12	0.00	52.00	13.96
Number of randomized applicants	12.90	10.46	2.00	79.00	—
Fraction of applicants assigned to Head Start	0.59	0.06	0.25	0.83	—
Observations (centers)	302				319

Notes: This table summarizes characteristics of Head Start centers in the HSIS data. Means and standard deviations are student-weighted for variables other than number of applicants and fraction assigned to Head Start. The HSIS sample excludes centers where the center director did not answer the HSIS survey, and centers where the fraction of students assigned to Head Start was zero or one. Column 5 shows mean characteristics for other preschools attended by children not offered a seat.

teacher experience as a strong predictor of classroom effectiveness in the Tennessee STAR class size experiment. The average center director has 18 years of experience working in center-based preschools, and the standard deviation of director experience across centers is 10 years. In Section IV, I explore whether this variation in inputs can explain differences in effectiveness across Head Start centers.

II. Pooled Estimates

Before investigating heterogeneity in causal effects, I summarize the average impact of Head Start using pooled equations of the form

$$(2) \quad Y_i = \alpha + \beta D_i + \mathbf{X}_i' \boldsymbol{\lambda} + \epsilon_i,$$

where Y_i is a summary index of outcomes for student i ; D_i is a dummy for Head Start attendance; and \mathbf{X}_i is a vector of the baseline controls from Table 2, included to increase precision. The attendance dummy is instrumented with an indicator for assignment to Head Start, Z_i , with first-stage equation

$$(3) \quad D_i = \kappa + \pi Z_i + \mathbf{X}_i' \boldsymbol{\delta} + \eta_i.$$

I estimate these equations by weighted two-stage least squares using the HSIS baseline child weights to account for differences in the probability of assignment across centers. These weights multiply the inverse probability of a child's experimental assignment by the probability that a child's center was sampled from the national population (Puma et al. 2010). Estimates using other weighting schemes, or including center fixed effects in equations (2) and (3), were very similar to those reported below. The coefficient β can be interpreted as a weighted average of center-specific local average treatment effects (LATEs), defined as effects of Head Start attendance

TABLE 4—EFFECTS OF HEAD START ON COGNITIVE AND NONCOGNITIVE SKILLS BY COHORT AND YEAR

Time period	Cohort	Cognitive skills		Noncognitive skills	
		First stage (1)	IV estimate (2)	First stage (3)	IV estimate (4)
Spring 2003	3-year-olds	0.679*** (0.031) 2,070	0.171*** (0.040) 2,070	0.679*** (0.031) 2,062	0.053 (0.035) 2,062
	4-year-olds	0.684*** (0.034) 1,638	0.088** (0.037) 1,638	0.685*** (0.032) 1,631	-0.041 (0.032) 1,631
Spring 2004	3-year-olds	0.362*** (0.031) 2,046	0.152* (0.079) 2,046	0.358*** (0.031) 2,032	0.083 (0.071) 2,032
	4-year-olds	0.693*** (0.033) 1,535	-0.080* (0.045) 1,535	0.693*** (0.032) 1,555	-0.035 (0.041) 1,555
Spring 2005	3-year-olds	0.375*** (0.033) 1,927	-0.014 (0.090) 1,927	0.379*** (0.033) 1,996	0.045 (0.088) 1,996
	4-year-olds	0.668*** (0.034) 1,527	0.003 (0.062) 1,527	0.668*** (0.034) 1,576	-0.064 (0.044) 1,576
Spring 2006	3-year-olds	0.367*** (0.032) 1,876	0.058 (0.104) 1,876	0.372*** (0.032) 1,957	0.030 (0.076) 1,957

Notes: This table reports estimates of the effect of Head Start attendance on summary indices of cognitive and noncognitive skills. Estimates come from instrumental variables models using assignment to Head Start as an instrument for Head Start attendance. All models use the HSIS baseline child weights and control for the baseline covariates listed in Table 2. Missing covariates are set to zero, and dummies for missing values are included. Standard errors are clustered at the Head Start center level.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

on students induced to attend by the experimental offer (Angrist and Imbens 1995). Standard errors for these and all subsequent models allow for clustering by center of random assignment.

Estimates of equations (2) and (3) reveal that Head Start attendance boosts outcomes during preschool, but these effects fade out quickly once children leave the program. Table 4 reports estimates of effects for cognitive and noncognitive skills, separately by grade and assignment cohort. Column 1 shows that in the first year after random assignment, applicants assigned to treatment were 68 percentage points more likely to attend Head Start than applicants in the control group. The corresponding second-stage estimates for cognitive skills, reported in column 2, show that Head Start attendance increased cognitive skills by 0.17 standard deviations for 3-year-olds and 0.09 standard deviations for 4-year-olds. These estimates are statistically significant at the 5-percent level. In contrast, estimates for noncognitive

skills, reported in column 4, show no evidence of an effect: the point estimate for three-year-olds is positive, the estimate for four-year-olds is negative, and neither is statistically significant.

In Spring 2004, members of the three-year-old cohort were still enrolled in Head Start. The cognitive point estimate for this time period is comparable to the Spring 2003 estimate (0.15 standard deviations), but is less precise ($s.e. = 0.08$). The decline in precision between 2003 and 2004 is driven by a decline in compliance for the 3-year-old cohort: many children in the control group reapplied to Head Start and were admitted at age 4, reducing the first stage from 0.68 to 0.36.⁹ Similarly, the noncognitive estimate for three-year-olds in Spring 2004 is positive but imprecise.

The remaining rows of Table 4 show that the effects of Head Start attendance dissipate once children exit the program. The cognitive estimate for the three-year-old cohort in Spring 2005 is close to zero, and the estimate for four-year-olds in Spring 2004 is negative and marginally significant. Estimates for both cohorts are small and statistically insignificant in later periods. Noncognitive estimates are not statistically distinguishable from zero in any time period for either cohort. Together, these results show little evidence of cognitive or noncognitive effects of Head Start after children leave preschool.

III. Variation in Head Start Effects

A. Variation in Instrumental Variables Estimates

I next turn to the primary contribution of this paper: quantifying and explaining variation in short-run effects across Head Start centers. As a first look at cross-center heterogeneity, Figure 1 plots center-specific reduced form coefficients against first stages. These coefficients come from regressions of cognitive skills and Head Start attendance in Spring 2003 on the Head Start offer indicator, pooling the three- and four-year-old cohorts. In the absence of treatment effect heterogeneity, reduced forms should be proportional to first stages with the same constant of proportionality for every center, so a single line through the origin should fit all points in Figure 1 up to sampling error. The red line shows a weighted least squares regression through the origin, with weights proportional to sample size times the variance of the Head Start offer. The χ^2 statistic from a test that all points lie on this line is equal to the overidentification test statistic from a two-stage least squares model using all center-by-offer interactions as instruments for Head Start attendance. The χ^2 statistic is equal to 421.4 and has 318 degrees of freedom, so the null hypothesis of no cross-center effect heterogeneity is rejected ($p < 0.01$).

The evidence in Figure 1 suggests that effects vary across Head Start centers. The magnitude of this variation is also of interest. Empirical Bayes (EB) methods are the conventional approach to quantifying cross-site variation in treatment effects (Morris 1983). The EB approach involves specifying a prior distribution for the cross-site distribution of parameters and then estimating the hyperparameters of the

⁹Head Start participation in Spring 2004 is measured from the parental survey since an administrative measure of participation is only available in Spring 2003. See the online Appendix.




FIGURE 1. CENTER-SPECIFIC REDUCED FORMS AND FIRST STAGES

Notes: This figure plots center-specific reduced form differences in cognitive skills in Spring 2003 against first-stage differences in Head Start attendance rates. The line comes from a weighted least squares regression through the origin, with weights proportional to $NP(Z)[1 - P(Z)]$, where N is sample size and $P(Z)$ is the fraction of applicants offered Head Start. The slope is 0.14 ($SE = 0.03$). The chi-squared statistic from a test that all points lie on the line is 421.4 (degrees of freedom = 318, $p = 0.00$).

prior. In cases where site-specific estimates are unbiased and have a known sampling variance, the EB estimator takes an especially simple form: The variance of treatment effects can be consistently estimated by subtracting the average squared standard error from the sample variance of site-specific estimates (Jacob and Lefgren 2008). With enough data at each site and many sites, this estimator nonparametrically identifies the cross-site variance of effects. An efficient “shrinkage” estimator of the effect at a particular site can then be constructed as a weighted average of the estimate for that site and the overall average effect.

This approach is inappropriate for the HSIS data. Figure 1 reveals substantial variation in compliance with random assignment across centers; to account for this variation, it is necessary to study instrumental variables estimates rather than intent-to-treat effects of assignment to Head Start. Instrumental variables estimates have no finite moments and are not centered at the true parameter in finite samples (Nelson and Starz 1990). In addition, conventional asymptotic standard errors provide a poor approximation to their behavior in small samples (Mariano 1977). Center-specific samples in the HSIS are often small, so the finite-sample behavior of IV is relevant for center-specific IV estimates. This can be seen in Figure 2, which shows a histogram of the distribution of sample sizes across HSIS centers. More than half of the centers have fewer than 10 applicants, and few have more than 25.

Table 5 illustrates the poor finite-sample behavior of center-specific IV estimates for cognitive skills in Spring 2003. The IV estimate for center j , $\hat{\beta}_j$, is the ratio of the center-specific reduced form and first stage. The sample standard deviation of these estimates is large (1.44 test score standard deviations), and estimates for some

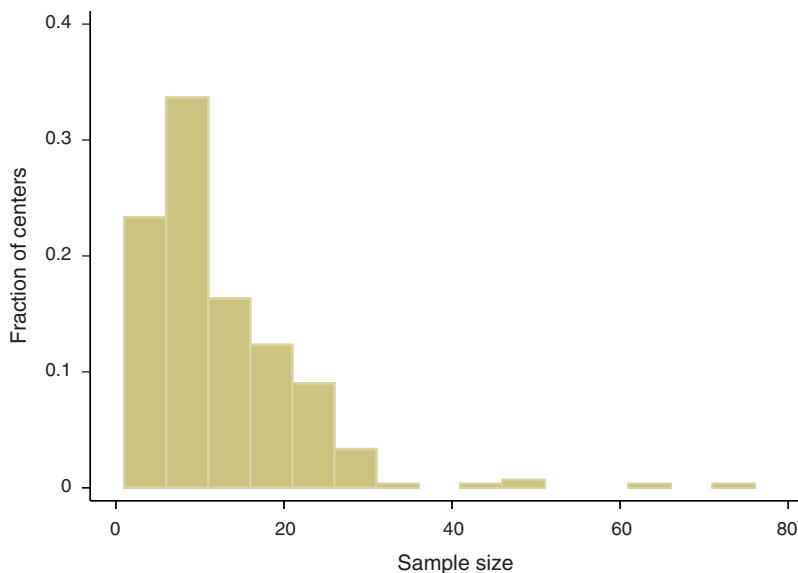


FIGURE 2. HISTOGRAM OF SAMPLE SIZE ACROSS HEAD START CENTERS

Notes: This figure shows the histogram of center-specific sample sizes in the HSIS experiment. The data are grouped into bins of 5 children (0 to 5, 6 to 10, etc.).

TABLE 5—FINITE-SAMPLE BEHAVIOR OF CENTER-SPECIFIC INSTRUMENTAL VARIABLES ESTIMATES

	Mean (1)	SD (2)	Min. (3)	Max. (4)
IV estimate	0.238	1.437	-4.541	14.804
IV asymptotic standard error	1.304	6.299	0.047	91.122
Implied cross-center variance of effects	Unweighted: -39.18 (1,195.02) Weighted: -35.98 (34.98)			

Notes: This table summarizes the distribution of center-specific instrumental variables estimates for cognitive skills in Spring 2003. The estimate for each center comes from a separate IV regression of cognitive skills on Head Start attendance instrumented by Head Start assignment, pooling the three- and four-year-old cohorts and using the HSIS child weights. The sample excludes centers with less than three applicants and centers with first stages equal to exactly zero. Two other centers with small samples and first stages very close to zero are also dropped. The sample includes 286 centers. The implied cross-center variance of effects is the sample variance of the IV estimates minus the average squared standard error. The weighted variance calculation weights observations by the reciprocal of the IV standard error. Standard errors of variance estimates are in parentheses.

centers are implausible (as large as 14.8 standard deviations). The wide dispersion in center-specific estimates is evident in Figure 3, which shows a histogram of $\hat{\beta}_j$, excluding estimates in excess of 2 in absolute value to keep the scale reasonable.

Moreover, the asymptotic standard errors associated with these estimates yield nonsensical results. The average standard error is 1.3 standard deviations. An estimate of the variance of β_j is given by

$$(4) \quad \hat{\sigma}_{\beta}^2 = \frac{1}{J} \sum_j \left((\hat{\beta}_j - \bar{\beta})^2 - SE(\hat{\beta}_j)^2 \right).$$




FIGURE 3. HISTOGRAM OF CENTER-SPECIFIC IV ESTIMATES

Notes: This figure plots the histogram of center-specific IV estimates for cognitive skills in Spring 2003. Estimates greater than two in absolute value are excluded.

As a result of extremely large standard errors for some centers, this estimate is negative and large (-39.2 standard deviations), and the associated standard error shows that it is almost completely uninformative. Since the IV asymptotic standard errors may be most inaccurate for the smallest centers, Table 5 also shows a variance estimate that weights centers by sample size. This estimate is negative and similar in magnitude to the unweighted estimate (-36 standard deviations); weighting improves precision, but the standard error of the weighted variance estimate is still extremely large (35 standard deviations). These negative variance estimates, and the associated sampling uncertainty, make it clear that the $\hat{\beta}_j$ and their asymptotic standard errors are not informative about the extent of effect heterogeneity across centers. I next describe a framework that consistently quantifies variation in Head Start effects despite small within-center sample sizes.

B. Random Coefficients Framework

My approach to quantifying effect variation uses a sample selection model to describe potential outcomes and Head Start participation conditional on center-specific parameters. I treat the parameters at each center as draws from a prior distribution of random coefficients, and derive an integrated likelihood function for the sample that depends only on the hyperparameters of this distribution. I then estimate the hyperparameters by maximum likelihood. This approach circumvents the need to compute $\hat{\beta}_j$ for every Head Start center.

Let $Y_{ij}(1)$ and $Y_{ij}(0)$ denote potential outcomes in and out of Head Start for student i applying to Head Start center j . Potential outcomes can be written as

$$(5) \quad Y_{ij}(d) = \alpha_{dj} + \epsilon_{idj}, \quad d \in \{0, 1\},$$

where $E[\epsilon_{idj}] = 0$. The Head Start participation decision is described by

$$(6) \quad D_{ij} = 1\{\lambda_j + \pi_j Z_{ij} > \eta_{ij}\}.$$

The vector of parameters at center j is therefore

$$(7) \quad \theta_j \equiv (\alpha_{1j}, \alpha_{0j}, \lambda_j, \log \pi_j)'$$

The average effect of Head Start attendance at center j is $\alpha_{1j} - \alpha_{0j}$. Note that the parameter vector is defined in terms of $\log \pi_j$, which guarantees that a Head Start offer weakly increases the probability of Head Start participation for any value of θ_j .

I assume the following parametric structure for the within-center distribution of potential outcomes:

$$(8) \quad (\epsilon_{i1j}, \epsilon_{i0j}, \eta_{ij})' | Z_{ij} \sim N(0, \Sigma).$$

Conditional on the center-specific parameters θ_j , assumption (8) yields a two-sided version of the Heckman (1979) sample selection (Heckit) model. The likelihood of the observed outcomes for student i is given by

$$(9) \quad \mathcal{L}_{ij}(Y_{ij}, D_{ij} | Z_{ij}; \theta_j) = \left[\Phi \left(\frac{\sigma_1(\lambda_j + \pi_j Z_{ij}) - \rho_1(Y_{ij} - \alpha_{1j})}{\sigma_1 \sqrt{1 - \rho_1^2}} \right) \frac{1}{\sigma_1} \phi \left(\frac{Y_{ij} - \alpha_{1j}}{\sigma_1} \right) \right]^{D_{ij}} \\ \times \left[\left(1 - \Phi \left(\frac{\sigma_0(\lambda_j + \pi_j Z_{ij}) - \rho_0(Y_{ij} - \alpha_{0j})}{\sigma_0 \sqrt{1 - \rho_0^2}} \right) \right) \frac{1}{\sigma_0} \phi \left(\frac{Y_{ij} - \alpha_{0j}}{\sigma_0} \right) \right]^{1-D_{ij}},$$

where σ_d is the standard deviation of ϵ_{idj} and ρ_d is its correlation with η_{ij} .¹⁰

Next, I assume that the cross-center distribution of parameters follows a normal distribution

$$(10) \quad \theta_j | Z_j \sim N(\theta_0, \mathbf{V}_0),$$

¹⁰There are two standard concerns with the Heckit model. First, without excluded instruments, the model is identified only by functional form restrictions (Heckman 1990). This is not a problem in the present context because the Head Start offer is a strong instrument. Second, even with an excluded instrument, the functional form assumptions may be incorrect. As a check on the plausibility of assumption (8), online Appendix Table A2 compares estimates from a version of the Heckit model with no center heterogeneity to results from instrumental variables estimation. The maximum likelihood estimates of the first- and second-stage parameters closely match the IV estimates, suggesting that the Heckit model is not badly misspecified.

where \mathbf{Z}_j is the vector of experimental offers for children at center j . The variance matrix \mathbf{V}_0 captures heterogeneity in outcome distributions and experimental compliance across Head Start centers. To estimate $\boldsymbol{\theta}_0$ and \mathbf{V}_0 , I integrate the site-specific parameters out of the likelihood function. The integrated likelihood for center j is

$$(11) \quad \mathcal{L}_j^I(\mathbf{Y}_j, \mathbf{D}_j | \mathbf{Z}_j; \boldsymbol{\theta}_0, \mathbf{V}_0) = \int \prod_i \mathcal{L}_{ij}(Y_{ij}, D_{ij} | Z_{ij}; \boldsymbol{\theta}) \phi_m(\boldsymbol{\theta}; \boldsymbol{\theta}_0, \mathbf{V}_0) d\boldsymbol{\theta},$$

where $\phi_m(\mathbf{x}; \boldsymbol{\mu}, \mathbf{V})$ is the multivariate normal density function. The integral in equation (11) does not have a closed form, so I approximate it by simulation, using 1,000 draws of $\boldsymbol{\theta}_j$ for each Head Start center. An empirical Bayes (EB) estimator of $\boldsymbol{\theta}_0$ and \mathbf{V}_0 maximizes the sum of logarithms of simulated likelihoods across Head Start centers.

C. Random Coefficients Estimates

Table 6 reports key parameter estimates from the normal random coefficients model for Spring 2003, pooling the three- and four-year-old cohorts.¹¹ The full set of parameter estimates is reported in online Appendix Table A3. I focus on Spring 2003 because effects for this period are largest and most precisely estimated; in addition, the evidence in Chetty et al. (2011) suggests that immediate impacts of early childhood programs may predict long-run effects better than impacts in later time periods. Results for Spring 2005 are reported in online Appendix Tables A3 and A4.

The estimated parameter distributions reveal substantial heterogeneity in parameters across Head Start centers. Consistent with the first-stage estimates in Table 4, the mean compliance probability is 0.74. Compliance rates vary substantially across sites: The cross-site standard deviation of the compliance probability is 0.22. This implies that about 20 percent of centers have compliance probabilities below 0.5.

Table 6 also shows estimates of the cross-center distribution of causal effects. The estimate of the average effect for cognitive skills is 0.11 standard deviations, while the mean noncognitive effect is 0.02. The cross-center standard deviation of Head Start effects, given by $\sqrt{\text{Var}(\alpha_{1j} - \alpha_{0j})}$, is estimated to be 0.18 standard deviations for cognitive skills. This implies substantial treatment effect variation across Head Start centers. For comparison, estimates of the standard deviations of school and teacher effectiveness are typically around 0.1 test score standard deviations (Chetty, Friedman, and Rockoff 2014a; Deming 2014; Kane, Rockoff, and Staiger 2008). My estimates therefore suggest that variation in short-run Head Start effectiveness is larger than variation in value-added across teachers or schools. The standard

¹¹ Within a center, three- and four-year-old applicants sometimes faced different probabilities of assignment to Head Start. I reweight likelihood contributions to account for these differences. Specifically, the likelihood contribution of child i is $\mathcal{L}_{ij}^{w_i}$, where \mathcal{L}_{ij} is the expression for the likelihood given in equation (11) and w_i is a weight proportional to child i 's base HSIS weight, normalized to sum to the total sample size.

TABLE 6—RANDOM COEFFICIENTS ESTIMATES FOR SPRING 2003

Parameter	Description	Cognitive skills		Noncognitive skills	
		Estimate (1)	Standard error (2)	Estimate (3)	Standard error (4)
$E[\Phi(\lambda_j + \pi_j) - \Phi(\lambda_j)]$	Mean compliance probability	0.743***	0.022	0.744***	0.021
$[\text{Var}(\Phi(\lambda_j + \pi_j) - \Phi(\lambda_j))]^{1/2}$	Standard deviation of compliance probability	0.220***	0.011	0.203***	0.011
$E[\alpha_{1j}]$	Mean treated outcome	0.105***	0.026	0.024	0.017
$E[\alpha_{0j}]$	Mean nontreated outcome	-0.009	0.029	0.000	0.016
$E[\alpha_{1j} - \alpha_{0j}]$	Mean Head Start effect	0.114***	0.035	0.024	0.021
$[\text{Var}(\alpha_{1j} - \alpha_{0j})]^{1/2}$	Standard deviation of Head Start effects	0.184***	0.016	0.068***	0.007

Notes: This table lists maximum simulated likelihood estimates of parameters of the cross-center distribution of Head Start effects in Spring 2003. The sample pools the three- and four-year-old cohorts, and observations are weighted using the HSIS baseline child weights. The MSL procedure uses 1,000 simulations for each Head Start center. Standard errors are robust to misspecification and are clustered at the Head Start center level.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

deviation of effects for noncognitive skills is smaller (0.068 standard deviations). Figure 4 summarizes the estimated random coefficient distributions, comparing them to histograms of center-specific, first-stage, and IV estimates.¹² The estimated parameter distributions show much less dispersion than the distributions of center-specific estimates; nonetheless, these distributions display substantively important heterogeneity.

The random coefficients estimates suggest that some Head Start centers have negative effects: 27 percent of centers ($\Phi(-0.11/0.18)$) are estimated to have cognitive effects below 0. To some extent, this is an artifact of the assumed distribution for θ_j , which has full support on the real line. There is no reason to expect Head Start effects to be positive for all centers or children, however. Head Start does not charge tuition, and some parents who would otherwise spend money or time on higher quality childcare may be willing to forego quality in exchange for this subsidy.¹³

As a check on the robustness of the random coefficient results to changes in functional form assumptions, I estimated an alternative version of the model assuming that θ_j is drawn from a finite set of possible types rather than a normal distribution.

¹²The first stage for center j is the difference in attendance probabilities between offered and non-offered applicants, given by

$$FS_j = \Phi(\lambda_j + \pi_j) - \Phi(\lambda_j).$$

Since $(\lambda_j, \log \pi_j)$ is assumed to be multivariate normal, the expression inside the first CDF is the sum of a normal random variable and a correlated log normal, which is not normally distributed. This functional form implies that the first stage is between 0 and 1 for all centers, a key assumption of instrumental variables models.

¹³Cohodes and Goodman (2014) find evidence of this phenomenon in the higher education sphere: the Massachusetts' Adams Scholarship program induces some students to substitute from expensive private institutions to less expensive public ones, reducing degree attainment in the process.

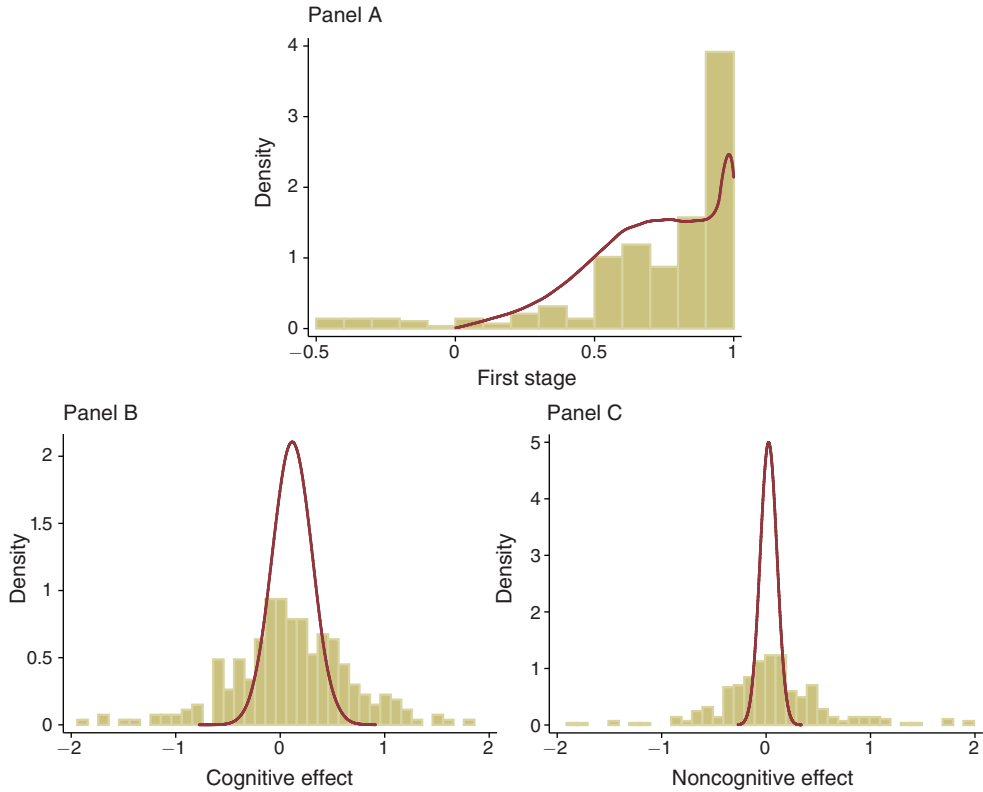


FIGURE 4. ESTIMATES OF CROSS-CENTER PARAMETER DISTRIBUTIONS

Notes: This figure plots maximum simulated likelihood estimates of the cross-center distributions of parameters in Spring 2003. The bars are histograms of center-specific first-stage and IV estimates. The curves are kernel density estimates produced using 200,000 draws from the distributions listed in online Appendix Table A3. The densities are estimated with a triangle kernel. The bandwidth is 0.05 for panel A and 0.1 for panels B and C.

The finite-type estimates are reported in online Appendix Table A5. These estimates also suggest substantial effect heterogeneity across Head Start centers. The implied cross-center standard deviations of effects for 3- and 5-type models are 0.12 and 0.22 standard deviations, roughly similar to the normal estimate of 0.18. This result implies that the key conclusions of the random coefficients analysis are not sensitive to the assumed functional form for the distribution of θ_j .

To provide further context for these estimates, I next compute the implied earnings effect of an improvement in Head Start quality, using the relationships between test score effects and lifetime earnings reported by Chetty, Friedman, and Rockoff (2014b). Chetty, Friedman, and Rockoff (2014b) show that a 1 standard deviation increase in teacher value-added in a single grade translates into a 1.3 percent increase in lifetime earnings. If the mapping between the short-run effect of Head Start on test scores and its effect on earnings is the same as this mapping for teachers, my results imply that a Head Start center at the eighty-fourth percentile of program quality (1 standard deviation above average) will boost lifetime earnings by 1.8 percent more than an average Head Start center. Assuming that children in the HSIS data will earn roughly the same amount as their parents

relative to the national median (a conservative assumption since earnings revert to the mean), and using the same assumptions on lifetime earnings trajectories used by Chetty, Friedman, and Rockoff (2014b), this translates into an earnings effect of about \$3,400 per child in 2010 dollars.¹⁴ This calculation shows that the magnitude of cross-center variation in Head Start effectiveness is large enough to matter for later outcomes, and is also large relative to the per child cost of the program (roughly \$7,600; DHHS 2011).

IV. Explaining Head Start Effects

A. Definitions of Inputs

The estimates reported above show that some Head Start programs are substantially more effective than others. In the remainder of the paper, I ask whether this variation in effectiveness can be explained by observed inputs. I assess the contributions of three sets of variables: Head Start center characteristics, child characteristics, and counterfactual preschool choices.

The analysis of center characteristics focuses on the seven variables listed in Table 3: The High/Scope curriculum, teacher education and certification, class size, instructional time, home visiting, and center director experience. These variables are often cited as key contributors to the success of model preschool programs (Schweinhart 2007; Chetty et al. 2011). Child characteristics include mother's education, family income, and baseline cognitive and noncognitive skills. These variables seem likely to be closely linked with human capital. The Perry Preschool Project enrolled a population of very disadvantaged children (Schweinhart et al. 2005). The analysis here asks whether differences in child characteristics partly explain the difference in effectiveness between Head Start and model programs.

I also investigate the role of differences in private preschool attendance rates across centers. Children in the HSIS sample can participate in three types of child-care: Head Start, other center-based preschool, or home care (no preschool). As shown in Table 2, the effect of a Head Start offer on the probability of Head Start attendance is larger than its effect on preschool attendance. This implies that some applicants would attend other preschools in the absence of Head Start. If private preschool affects cognitive skills relative to no preschool, differences in private preschool participation rates may drive cross-center variation in Head Start effects even if Head Start programs are of uniform quality.

¹⁴Chetty, Friedman, and Rockoff (2014b) report that the standard deviation of teacher quality is 0.13 test score standard deviations. They argue that a 1 standard deviation move upward in this teacher quality distribution for 1 year raises students earnings by 1.3 percent. The implied earnings gain per standard deviation of test scores is therefore $(1.3/0.13) = 10$ percent. I estimate that the standard deviation of Head Start quality is 0.18 test score standard deviations, so a 1 standard deviation increase in Head Start quality boosts earnings by $0.18 \cdot 10 = 1.8$ percent. Chetty, Friedman, and Rockoff (2014b) estimate that the mean present value of lifetime earnings is roughly \$522,000 at age 12 in 2010 dollars, which is \$434,000 discounted back to age 5 at a 3 percent rate. The average HSIS family earned \$18,085 per year, or 44 percent of the US median in 2002 (see <http://www.census.gov/prod/2003pubs/p60-221.pdf>). The average present discounted value of earnings at age 5 for children in the HSIS sample can therefore be conservatively estimated as $0.44 \cdot \$434,000 = \$190,960$. The earnings impact of a 1 standard deviation increase in Head Start quality can then be approximated as $\$190,960 \cdot 0.018 = \$3,437.28$.

To investigate this issue, I estimate the share of students drawn into Head Start from other preschools at center j using the regression

$$(12) \quad C_{ij} = \tau_j^C + \rho_j^C Z_{ij} + u_{ij}^C,$$

where C_{ij} is an indicator for attending non-Head Start center-based preschool. The coefficient ρ_j^C measures the reduction in other center-based preschool attendance caused by a Head Start offer. Similarly, the share of students drawn from no preschool is estimated using the regression

$$(13) \quad N_{ij} = \tau_j^N + \rho_j^N Z_{ij} + u_{ij}^N,$$

where N_{ij} is an indicator for attending no preschool. Under the assumption that a Head Start offer does not affect the choice of private versus no preschool,¹⁵ the share of Head Start compliers drawn from other preschool centers is given by

$$(14) \quad S_j^C = \frac{(-\rho_j^C)}{(-\rho_j^N) + (-\rho_j^C)}.$$

I estimate equations (12) and (13) by weighted least squares using the HSES child weights, setting positive coefficients to zero to keep S_j^C between zero and one. Figure 5 shows a histogram of S_j^C . This figure reveals that the share of compliers who would attend other preschools in the absence of Head Start varies across centers. At about 10 percent of centers, all compliers attend other preschools if denied the opportunity to attend Head Start. About 20 percent of centers appear to draw children only from home care. The remaining 70 percent draw children from a mix of private preschool and no preschool.

I investigate the relationship between inputs and Head Start effects using two approaches. First, I estimate interacted two-stage least squares models, with second- and first-stage equations of the form

$$(15) \quad Y_{ij} = \alpha + \mathbf{P}'_{ij}\phi + \beta D_{ij} + D_{ij} \cdot \mathbf{P}'_{ij}\psi + \mathbf{X}'_{ij}\gamma + \epsilon_{ij},$$

$$(16) \quad D_{ij} = \kappa + \mathbf{P}'_{ij}\nu + \pi Z_{ij} + Z_{ij} \cdot \mathbf{P}'_{ij}\tau + \mathbf{X}'_{ij}\delta + \eta_{ij},$$

where \mathbf{P}_{ij} is a vector of child i 's characteristics and the characteristics of her center of random assignment. The first-stage equations for the interactions of D_{ij} and \mathbf{P}_{ij} are analogous to equation (16). This approach compares IV estimates for groups of centers and children with different values of \mathbf{P}_{ij} . Since samples at groups of centers using different inputs are larger than samples at individual centers, this IV analysis

¹⁵This assumption can be motivated by a revealed preference argument: The availability of private preschool is unaffected by a Head Start offer, so preferences for private versus no preschool should not be affected by the offer. A shift between private and no preschool in response to a Head Start offer would violate the exclusion restriction required for the offer to be a valid instrument for Head Start attendance.




FIGURE 5. DISTRIBUTION OF CENTER-BASED PRESCHOOL COMPLIER SHARE

Notes: This figure shows the histogram of center-specific shares of compliers attending non-Head Start center-based preschool. The data are grouped into bins of width 0.05.

is not subject to the finite-sample issues discussed in Section II. The vector ψ captures the relationship between the effect of Head Start attendance and observed inputs. I estimate two sets of interaction models: bivariate models that include inputs in \mathbf{P}_{ij} one at a time, and multivariate models that include all inputs simultaneously. Equations (15) and (16) are estimated using binary measures of each input; for continuous variables, these indicators equal one for centers above the sample median. An analysis using continuous measures yielded similar but less precise results.

Second, I extend the selection model to incorporate dependence between inputs and causal effects. The potential outcome and selection equations are

$$(17) \quad Y_{ij}(d) = \alpha_{dj} + \mathbf{P}'_{ij}\psi_d + \epsilon_{ijd}, \quad d \in \{0, 1\},$$

$$(18) \quad D_{ij} = 1\{\lambda_j + \mathbf{P}'_{ij}\boldsymbol{\nu} + \exp(\log \pi_j + \mathbf{P}'_{ij}\boldsymbol{\tau}) \cdot Z_{ij} > \eta_{ij}\},$$

where $(\epsilon_{i1j}, \epsilon_{i0j}, \eta_{ij})$ and $(\alpha_{1j}, \alpha_{0j}, \lambda_j, \log \pi_j)$ are assumed to be normally distributed as before. The vector $(\psi_1 - \psi_0)$ measures the relationship between inputs and Head Start effects. This approach relies in part on parametric assumptions, so it is likely to be less robust than two-stage least squares. The advantage of the random coefficients approach is that it generates an estimate of \mathbf{V}_0 , the residual variation in center-specific parameters remaining after accounting for observed inputs. It can therefore be used to measure the share of effect heterogeneity explained by \mathbf{P}_{ij} .

B. Relationships between Inputs and Head Start Effects

Table 7 reports the results of the analysis of inputs. Panel A shows estimates of relationships between Head Start effectiveness and center characteristics. The estimates reveal that centers offering full-day service and frequent home visiting are more effective. On average, cognitive effects of full-day Head Start centers are 0.14 standard deviations larger than effects of centers that do not offer this service. Corresponding estimates for the multivariate interaction and maximum likelihood models are somewhat smaller but still statistically significant. This implies that the relative effectiveness of full day centers is not explained by other inputs. Centers that offer frequent home visits per year are especially effective at raising noncognitive skills: The bivariate model shows that centers offering more than 3 home visits per year boost noncognitive skills by 0.11 standard deviations more than centers providing 3 or less visits, and this estimate is statistically significant. The multivariate and maximum likelihood estimates show that frequent home visiting is also associated with larger effects on cognitive skills.

The remaining estimates in panel A of Table 7 show that other center characteristics are mostly unrelated to Head Start effectiveness, though these estimates vary in precision. High/Scope centers do not boost scores more than non-High/Scope centers; the interaction terms associated with High/Scope are close to zero in all models. Moreover, this difference is precisely estimated. The hypothesis that High/Scope centers are 0.15 standard deviations more effective than other centers is rejected at the 5 percent confidence level for both cognitive and noncognitive skills. This result weighs against the view that the High/Scope curriculum alone generated the success of the Perry Preschool Project.

Estimates of relationships between Head Start effectiveness and teacher education and licensing are statistically insignificant in most models. This result is consistent with studies of teacher value-added, which typically find weak relationships between teacher effectiveness and credentials (Kane, Rockoff, and Staiger 2008). The estimate for teacher education is reasonably precise. In the bivariate model, the interaction coefficient on an indicator for any staff with a bachelor's degree is 0.026, with a standard error of 0.063. The upper bound of the 95 percent confidence interval associated with this estimate is 0.15. The mean share with a bachelor's degree among centers with any bachelor's degrees is 0.6. This implies that I can reject relatively small differences in effects between centers that differ substantially in mean teacher education. The results for licensing are less clear. Licensing estimates are positive in all models; the cognitive bivariate estimate is marginally significant, and the 95 percent confidence interval in the multivariate model includes effects as large as 0.22 standard deviations. These estimates suggest that there may be a relationship between teacher licensing and Head Start effectiveness, but the research design used here does not have the power to detect it.

The results for student/staff ratios and director experience are more surprising. Estimates from both experimental and quasi-experimental settings suggest that smaller classes and more experienced teachers boost test scores (Krueger 1999; Angrist and Lavy 1999; Chetty et al. 2011). In contrast, the results in Table 7 suggest that Head Start centers with smaller classes and more experienced directors

TABLE 7—RELATIONSHIPS BETWEEN INPUTS AND HEAD START EFFECTS

	Cognitive skills			Noncognitive skills		
	Two-stage least squares		Maximum likelihood (3)	Two-stage least squares		Maximum likelihood (6)
	Bivariate (1)	Multivariate (2)		Bivariate (4)	Multivariate (5)	
<i>Panel A. Center characteristics</i>						
Any staff with bachelor's degree	0.026 (0.063)	0.050 (0.048)	0.001 (0.043)	0.012 (0.051)	-0.041 (0.050)	-0.022 (0.032)
Any staff that have a teaching license	0.127* (0.068)	0.090 (0.064)	0.034 (0.044)	0.087 (0.052)	0.085 (0.056)	0.030 (0.040)
Low student/staff ratio	-0.044 (0.059)	-0.061 (0.050)	-0.068 (0.049)	0.033 (0.053)	0.046 (0.051)	0.031 (0.032)
Full day service	0.138** (0.055)	0.089* (0.047)	0.083* (0.043)	-0.043 (0.053)	-0.031 (0.050)	-0.002 (0.031)
More than three home visits per year	0.024 (0.070)	0.110* (0.064)	0.092* (0.055)	0.112** (0.050)	0.088 (0.056)	0.094** (0.037)
High/Scope curriculum	-0.009 (0.066)	0.005 (0.054)	-0.023 (0.048)	0.042 (0.053)	0.085 (0.057)	0.024 (0.035)
High center director experience	0.022 (0.061)	0.055 (0.053)	0.021 (0.044)	-0.011 (0.052)	-0.007 (0.053)	-0.013 (0.034)
<i>Panel B. Child characteristics</i>						
Mother graduated high school	-0.127** (0.062)	-0.077 (0.057)	-0.024 (0.042)	0.015 (0.049)	-0.010 (0.046)	-0.034 (0.032)
High income	-0.011 (0.061)	-0.003 (0.054)	-0.079* (0.044)	0.008 (0.047)	0.041 (0.043)	0.020 (0.027)
High baseline skills	-0.085 (0.055)	-0.004 (0.051)	0.019 (0.035)	0.023 (0.047)	-0.005 (0.049)	-0.020 (0.032)
<i>Panel C. Counterfactual preschool choices</i>						
High center-based preschool complier share	-0.099* (0.054)	-0.117** (0.051)	-0.076* (0.045)	-0.011 (0.047)	-0.019 (0.049)	-0.004 (0.032)
Residual standard deviation of Head Start effects	—	—	0.150	—	—	0.053
R^2			0.337			0.393

Notes: This table reports estimates of relationships between Head Start effects and inputs in Spring 2003. Two-stage least squares models instrument Head Start attendance and its interactions with inputs using assignment to Head Start and its interactions with inputs, with the same weighting scheme and controls as in Table 4. High (low) values of inputs are values above (below) the sample median. The bivariate models in columns 1 and 4 estimate a separate interaction model for each input, while the multivariate models in columns 2–3 and 5–6 include all interactions simultaneously. Main effects of interacting variables are included as controls. Bivariate models exclude observations with missing values for the relevant input; multivariate models exclude observations with missing values for any input. Standard errors are clustered at the Head Start center level.

***Significant at the 1 percent level.

**Significant at the 5 percent level.

*Significant at the 10 percent level.

are not more effective. In fact, the point estimates associated with a below-median student/staff ratio are negative in all models. The 95 percent confidence interval rules out differences in effects as small as 0.074 standard deviations between above- and below-median centers. I can also reject reasonably small differences (around 0.1 standard deviations) between centers with more experienced and less experienced directors.

Panel B of Table 7 reports relationships between child characteristics and Head Start effects. The estimates show that Head Start has larger effects for children with less educated mothers: Children of high school graduates gain 0.13 standard deviations less than children of high school dropouts, and this estimate is statistically significant at the 5 percent level. Corresponding estimates from the multivariate

models are smaller and insignificant, however. This implies that the larger effect for children of less educated mothers is mostly explained by other observed characteristics. The IV estimates for baseline skills and family income are statistically insignificant, though the point estimates suggest slightly larger effects for lower skilled and lower income students.¹⁶ Together, the estimates in Panel B suggest that Head Start is more effective for more disadvantaged students, but this relationship is fairly weak and is therefore unlikely to explain large differences in effects between Head Start and model programs.

Panel C reveals a significant negative relationship between Head Start effectiveness and the share of experimental compliers drawn from other center-based preschools. Head Start centers above the median of S_j^C boost cognitive skills by about 0.1 standard deviations less than centers below the median. Moreover, this relationship is not explained by other observed characteristics: the estimate is highly statistically significant and of similar magnitude in the multivariate two-stage least squares model. The choice between private preschool and home care is endogenous, so effects on subgroups of children drawn from these two sources cannot be directly estimated without further assumptions. The estimates in Panel C provide suggestive evidence that children drawn from home care rather than other preschools may benefit more from Head Start attendance.

The bottom of Table 7 reports estimates of $\sqrt{\text{Var}(\alpha_{1j} - \alpha_{0j})}$, the residual standard deviation of Head Start effects after accounting for observed inputs. Residual standard deviations are 0.150 for cognitive skills and 0.053 for noncognitive skills. A comparison with Table 6 reveals that in an R^2 sense, the inputs and practices examined here explain a significant proportion of cross-center effect variation. Specifically, inputs explains 34 percent of the variation in cognitive effects and 39 percent of the variation in noncognitive effects.¹⁷ Nevertheless, a majority of the variation in Head Start effects is left unexplained, and several of the key inputs emphasized by Schweinhart (2007) are unrelated to program effectiveness. This suggests that some important drivers of successful preschool programs have yet to be identified.¹⁸

V. Conclusion

Studies of small-scale model early childhood programs show that early intervention can boost outcomes in the short and long run. Randomized evidence from the Head Start Impact Study (HSIS) suggests that the Head Start program produces smaller short-run gains. This paper uses data from the HSIS to quantify impact

¹⁶ Bitler, Hoynes, and Domina (2014) find larger effects of Head Start on PPVT scores for children with lower baseline PPVT scores. This is consistent with the negative point estimate for baseline skills in column 1 of Table 7; however, I find that baseline skills are less related to Head Start effects for the other components of the summary index used here.

¹⁷ The proportions of variation in cognitive and noncognitive effects explained by inputs are $1 - (0.150/0.184)^2$ and $1 - (0.053/0.068)^2$.

¹⁸ While short-run effects are the focus of this paper, online Appendix Tables A4 and A6 repeat the analysis of heterogeneity and inputs for Spring 2005. There is much less effect variation in Spring 2005 than in Spring 2003, and relationships with inputs are less precisely estimated in this period. The inputs that predict short-run gains do not seem to predict longer run gains, which suggests that larger short-run effects are not associated with less fadeout.

variation across Head Start centers and asks whether differences in key inputs used by model programs can explain this variation. Estimates of a random coefficients selection model reveal substantial variation in effectiveness across Head Start centers, particularly with respect to cognitive skills. Centers that offer full day service and frequent home visiting are more effective than other centers, as are centers that draw more students from home care rather than center-based preschool. Other inputs typically cited as important to the success of small-scale programs, including the High/Scope curriculum, teacher education, and class size, do not predict program effectiveness in Head Start. Children of high school dropout mothers benefit more from Head Start, but family income and baseline skills weakly predict gains. Together, observed inputs and characteristics explain about one third of the variation in short-run cognitive effects across Head Start centers.

It is important to emphasize that educational practices and applicant populations are not randomly assigned to Head Start centers, so the estimates reported here may not reflect causal impacts of changing inputs in isolation. Since Head Start centers face budget constraints, spending more on observed inputs may require cutting spending on unobserved dimensions. As a result, my estimates may be biased toward zero relative to the causal effects of improving inputs. Nonetheless, this analysis shows that some inputs predict Head Start effectiveness, while others do not. The results provide no evidence that adoption of the High/Scope curriculum or teacher education requirements would improve program effectiveness in Head Start. This finding is relevant to recent policy changes that mandate increased education levels for Head Start teachers (DHHS 2008). My results show that full-day service and home visiting are most predictive of short-run Head Start effectiveness, and that efforts to target children who would not otherwise attend preschool might boost the effects of the program. Identifying factors that explain the large residual variation in program effectiveness is an important task for future research.

REFERENCES

- Allcott, Hunt. 2014. "Site Selection Bias in Program Evaluation." https://files.nyu.edu/ha32/public/research/Allcott_SiteSelectionBias.pdf.
- Anderson, Michael L. 2008. "Multiple Inference and Gender Differences in the Effects of Early Intervention: A Reevaluation of the Abecedarian, Perry Preschool, and Early Training Projects." *Journal of the American Statistical Association* 103 (484): 1481–95.
- Angrist, Joshua D., and Guido W. Imbens. 1995. "Two-Stage Least Squares Estimation of Average Causal Effects in Models with Variable Treatment Intensity." *Journal of the American Statistical Association* 90 (430): 431–42.
- Angrist, Joshua D., and Victor Lavy. 1999. "Using Maimonides' Rule to Estimate the Effect of Class Size on Scholastic Achievement." *Quarterly Journal of Economics* 114 (2): 533–75.
- Angrist, Joshua D., Parag A. Pathak, and Christopher R. Walters. 2013. "Explaining Charter School Effectiveness." *American Economic Journal: Applied Economics* 5 (4): 1–27.
- Berruta-Clement, John R., Lawrence J. Schweinhart, W. Steven Barnett, Ann S. Epstein, and David P. Weikart. 1984. *Changed Lives: The Effects of the Perry Preschool Program on Youths Through Age 19*. Ypsilanti, MI: HighScope Press.
- Bitler, Marianne P., Hilary W. Hoynes, and Thurson Domina. 2014. "Experimental Evidence on Distributional Effects of Head Start." National Bureau of Economic Research (NBER) Working Paper 20434.
- Campbell, Frances A., and Craig T. Ramey. 1994. "Effects of Early Intervention on Intellectual and Academic Achievement: A Follow-up Study of Children from Low-Income Families." *Child Development* 65 (2): 684–98.

- Campbell, Frances A., and Craig T. Ramey.** 1995. "Cognitive and School Outcomes for High-Risk African-American Students at Middle Adolescence: Positive Effects of Early Intervention." *American Educational Research Journal* 32 (4): 743–72.
- Cascio, Elizabeth U., and Diane Whitmore Schanzenbach.** 2013. "The Impacts of Expanding Access to High-Quality Preschool Education." *Brookings Papers on Economic Activity* 43 (2): 127–92.
- Chandra, Amitabh, Amy Finkelstein, Adam Sacarny, and Chad Syverson.** 2013. "Healthcare Exceptionalism? Productivity and Allocation in the U.S. Healthcare Sector." National Bureau of Economic Research (NBER) Working Paper 19200.
- Chetty, Raj, John N. Friedman, Nathaniel Hilger, Emmanuel Saez, Diane Whitmore Schanzenbach, and Danny Yagan.** 2011. "How Does Your Kindergarten Classroom Affect Your Earnings? Evidence from Project STAR." *Quarterly Journal of Economics* 126 (4): 1593–1660.
- Chetty, Raj, John N. Friedman, and Jonah E. Rockoff.** 2014a. "Measuring the Impacts of Teachers I: Evaluating Bias in Teacher Value-Added Estimates." *American Economic Review* 104 (9): 2593–2632.
- Chetty, Raj, John N. Friedman, and Jonah E. Rockoff.** 2014b. "Measuring the Impacts of Teachers II: Teacher Value-Added and Student Outcomes in Adulthood." *American Economic Review* 104 (9): 2633–79.
- Cohodes, Sarah R., and Joshua S. Goodman.** 2014. "Merit Aid, College Quality, and College Completion: Massachusetts' Adams Scholarship as an In-Kind Subsidy." *American Economic Journal: Applied Economics* 6 (4): 251–85.
- Currie, Janet, and Duncan Thomas.** 1995. "Does Head Start Make a Difference?" *American Economic Review* 85 (3): 341–64.
- Deming, David.** 2009. "Early Childhood Intervention and Life-Cycle Skill Development: Evidence from Head Start." *American Economic Journal: Applied Economics* 1 (3): 111–34.
- Deming, David J.** 2014. "Using School Choice Lotteries to Test Measures of School Effectiveness." National Bureau of Economic Research (NBER) Working Paper 19803.
- Dobbie, Will, and Roland G. Fryer.** 2013. "Getting Beneath the Veil of Effective Schools: Evidence from New York City." *American Economic Journal: Applied Economics* 5 (4): 28–60.
- Epstein, Ann S.** 2007. *Essentials of Active Learning in Preschool: Getting to Know the High/Scope Curriculum*. Ypsilanti, MI: HighScope Press.
- Fitzpatrick, Maria Donovan.** 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 (1): Article 46.
- Garces, Eliana, Duncan Thomas, and Janet Currie.** 2002. "Longer-Term Effects of Head Start." *American Economic Review* 92 (4): 999–1012.
- Gelber, Alexander, and Adam Isen.** 2013. "Children's Schooling and Parent's Behavior: Evidence from the Head Start Impact Study." *Journal of Public Economics* 101: 25–38.
- Gibbs, Chloe, Jens Ludwig, and Douglas L. Miller.** 2011. "Does Head Start Do Any Lasting Good?" National Bureau of Economic Research (NBER) Working Paper 17452.
- Gormley, William T., Jr. and Ted Gayer.** 2005. "Promoting School Readiness in Oklahoma: An Evaluation of Tulsa's Pre-K Program." *Journal of Human Resources* 40 (3): 533–58.
- Hanushek, Eric A.** 2009. "Teacher Deselection." In *Creating a New Teaching Profession*, edited by Dan Goldhaber and Jane Hannaway. Washington, DC: Urban Institute Press.
- Heckman, James.** 1979. "Sample Selection Bias as a Specification Error." *Econometrica* 47 (1): 153–61.
- Heckman, James.** 1990. "Varieties of Selection Bias." *American Economic Review* 80 (2): 313–18.
- Heckman, James.** 2011. "The American Family in Black and White: A Post-Racial Strategy for Improving Skills to Promote Equality." *Daedalus* 140 (2): 70–89.
- Heckman, James, Seong Hyeok Moon, Rodrigo Pinto, Peter A. Savelyev, and Adam Yavitz.** 2010a. "Analyzing Social Experiments as Implemented: A Reexamination of the Evidence from the High/Scope Perry Preschool Program." *Quantitative Economics* 1 (1): 1–46.
- Heckman, James, Seong Hyeok Moon, Rodrigo Pinto, Peter A. Savelyev, and Adam Yavitz.** 2010b. "The Rate of Return to the High/Scope Perry Preschool Program." *Journal of Public Economics* 94 (1–2): 114–28.
- Heckman, James, Rodrigo Pinto, and Peter A. Savelyev.** 2013. "Understanding the Mechanisms through Which an Influential Early Childhood Program Boosted Adult Outcomes." *American Economic Review* 103 (6): 2052–86.
- Heckman, James, Sergio Urzua, and Edward Vytlacil.** 2006. "Understanding Instrumental Variables in Models with Essential Heterogeneity." *Review of Economics and Statistics* 88 (3): 389–432.
- Hotz, V. Joseph, Guido W. Imens, and Julie H. Mortimer.** 2005. "Predicting the Efficacy of Future Training Programs Using Past Experiences at Other Locations." *Journal of Econometrics* 125 (1–2): 241–70.

- Hoxby, Caroline, Sonali Murarka, and Jennie Kang.** 2009. *How New York City's Charter Schools Affect Achievement*. Cambridge, MA: New York City Schools Evaluation Project.
- Imbens, Guido W., and Joshua D. Angrist.** 1994. "Identification and Estimation of Local Average Treatment Effects." *Econometrica* 62 (2): 467–75.
- Jacob, Brian A., and Lars J. Lefgren.** 2008. "Can Principals Identify Effective Teachers? Evidence on Subjective Performance Evaluation in Education." *Journal of Labor Economics* 26 (1): 101–36.
- Kane, Thomas J., Jonah E. Rockoff, and Douglas O. Staiger.** 2008. "What Does Certification Tell Us about Teacher Effectiveness? Evidence from New York City." *Economics of Education Review* 27 (6): 615–31.
- Kling, Jeffrey R., Jeffrey B. Liebman, and Lawrence F. Katz.** 2007. "Experimental Analysis of Neighborhood Effects." *Econometrica* 75 (1): 83–119.
- Krueger, Alan B.** 1999. "Experimental Estimates of Education Production Functions." *Quarterly Journal of Economics* 114 (2): 497–532.
- 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 (1): 159–208.
- Mariano, Roberto S.** 1977. "Finite Sample Properties of Instrumental Variable Estimators of Structural Coefficients." *Econometrica* 45 (2): 487–96.
- Morris, Carl N.** 1983. "Parametric Empirical Bayes Inference: Theory and Applications." *Journal of the American Statistical Association* 78 (381): 47–55.
- Nelson, Charles R., and Startz, Richard.** 1990. "Some Further Results on the Exact Small Sample Properties of the Instrumental Variable Estimator." *Econometrica* 58 (4): 967–76.
- Obama, Barack.** 2013. "State of the Union Address." Remarks by the President. Washington, DC. February 12, 2013. <http://www.whitehouse.gov/the-press-office/2013/02/12/president-barack-obamas-state-union-address>.
- Puma, Michael, Stephen Bell, Ronna Cook, and Camilla Heid.** 2010. *Head Start Impact Study: Final Report*. U.S. Department of Health and Services. Administration for Children and Families. Washington, DC, January.
- Puma, Michael, Stephen Bell, and Camilla Heid.** 2012. *Third Grade Follow-up to the Head Start Impact Study*. U.S. Department of Health and Human Services. Washington, DC, October.
- Raudenbush, Stephen W., Sean F. Reardon, and Takako Nomi.** 2012. "Statistical Analysis for Multisite Trials Using Instrumental Variables with Random Coefficients." *Journal of Research on Educational Effectiveness* 5 (3): 303–32.
- Reynolds, Arthur J., and Judy A. Temple.** 1998. "Extended Early Childhood Intervention and School Achievement: Age Thirteen Findings from the Chicago Longitudinal Study." *Child Development* 69 (1): 231–46.
- Schweinhart, Lawrence.** 2007. "How to Take the High/Scope Perry Preschool to Scale." Paper presented at the National Invitational Conference of the Early Childhood Research Collaborative, Minneapolis, December 7–8.
- Schweinhart, Lawrence J., Jeanne Montie, Zongping Xiang, W. Steven Barnett, Clive R. Belfield, and Milagros Nores.** 2005. *Lifetime Effects: The High/Scope Perry Preschool Study Through Age 40*. Ypsilanti, MI: HighScope Press.
- Schweinhart, Lawrence J., and David P. Weikart.** 1997. *Lasting Differences: The High/Scope Preschool Curriculum Comparison Study Through Age 23*. Ypsilanti, MI: HighScope Press.
- Syverson, Chad.** 2011. "What Determines Productivity?" *Journal of Economic Literature* 49 (2): 326–65.
- U.S. Department of Health and Human Services.** 2008. "Statutory Degree and Credentialing Requirements." Administration for Children and Families. http://eclkc.ohs.acf.hhs.gov/hslc/standards/pdf/PDF_IMs/IM2011/ACF-IM-HS-11-03.pdf (accessed March 27, 2013).
- U.S. Department of Health and Human Services.** 2011. *Head Start Program Facts: Fiscal Year 2011*. Administration for Children and Families. Washington, DC. <http://eclkc.ohs.acf.hhs.gov/hslc/mr/factsheets/docs/hs-program-fact-sheet-2011-final.pdf>.
- U.S. Department of Health and Human Services.** 2012. Head Start Services. <http://www.acf.hhs.gov/programs/ohs/about/head-start> (accessed March 27, 2013).
- Walters, Christopher R.** 2015. "Inputs in the Production of Early Childhood Human Capital: Evidence from Head Start: Dataset." *American Economic Journal: Applied Economics*. <http://dx.doi.org/10.1257/app.20140184>.
- Wong, Vivian C., Thomas D. Cook, W. Steven Barnett, and Kwanghee Jung.** 2008. "An Effectiveness-Based Evaluation of Five State Pre-kindergarten Programs." *Journal of Policy Analysis and Management* 27 (1): 122–54.