The Human Capital Effects of Housing Assistance*

Crossan Cooper†
Yale University

Job Market Paper

November 12, 2026

[Updated regularly. Please click here for the latest version]

Abstract

This paper estimates the impact of means-tested housing subsidies on children's human capital accumulation and early adult outcomes. I link applications to public housing and housing voucher wait lists in New Haven, Connecticut, to statewide education, labor market, and criminal justice records, and exploit variation in when households reach the top of randomly-ordered wait lists as an instrument for childhood exposure to assistance. Receiving housing assistance significantly reduces school-reported homelessness and housing instability at younger ages and raises measures of early childhood development among boys. Housing assistance also improves longer-run outcomes. Each additional year of assistance raises high school graduation and college enrollment by 2 to 4 percentage points, and children that receive housing assistance are less likely to be convicted of a felony in early adulthood. To distinguish among mechanisms driving these results, I estimate short-run effects on school continuity, school quality, and adult labor supply. Finally, translating the attainment effects into predicted impacts on lifetime earnings and accounting for reductions in incarceration and sheltered homelessness, I calculate that providing housing assistance to a household with children produces roughly $2 in benefits per government dollar spent. Because of larger impacts on educational attainment, housing vouchers are more cost-effective than public housing units.

*I thank Joe Altonji, Winnie van Dijk, John Eric Humphries, and Seth Zimmerman for providing valuable guidance and feedback. I also thank Cody Cook, Jason Abaluck, and Cormac O'Dea for their helpful comments. This project was made possible by a data sharing agreement between DataLinkCT and the local public housing agency. I thank the staff of the PHA for their advice, technical support, and patience. I also thank Scott Gaul, Katie Breslin, Ashley Dhaim, Rachel Leventhal-Weiner, Rachel Moskowitz, the rest of the DataLinkCT team, and state employees at DataLinkCT Participating Agencies for their help in linking administrative records. Finally, I thank the Yale Center for Research Computing for assistance in computation run on the Milgram cluster. This work would not be possible without data provided by the DataLinkCT participating agencies. The findings do not necessarily reflect the opinions of the State of Connecticut or the organizations and agencies contributing data. Any errors are my own.

†Cooper: Department of Economics, Yale University (email: crossan.cooper@yale.edu)

1 Introduction

Almost 4 million children in the United States live in households receiving means-tested federal housing subsidies. These transfers are large and long-lasting: the average household receives more than $13,600 annually and nearly two-thirds of non-elderly recipients retain assistance for at least a decade (Husock and Meyer 2025). Although these transfers can span much of childhood, whether they improve children's long-run outcomes – and do so cost-effectively – remains unclear.

Housing assistance may raise children's human capital by preventing homelessness and residential instability, enabling moves to higher-opportunity neighborhoods or higher-quality schools, or increasing households' net-of-housing resources.[1] In practice, the impacts of housing assistance are limited by program design and by frictions in the housing market. Long wait lists for these non-entitlement transfers delay the timing of assistance (Collinson et al. 2015; Waldinger 2021), while search frictions and jurisdictional boundaries limit voucher recipients' access to high-opportunity neighborhoods (Bergman et al. 2024; Garboden 2024). Poor unit conditions in public housing may also harm children, as evidenced by longstanding concerns among policymakers that public housing is "bad for kids" (Currie and Yelowitz 2000; Collinson and Gibbs 2025). Finally, the subsidy schedule disincentivizes parental labor supply and earnings (Jacob and Ludwig 2012; van Dijk 2019), with potentially offsetting effects on the time and money parents invest in children.[2]

In this paper, I combine new data with a natural experiment to provide new evidence on the effects of housing vouchers and public housing on children, focusing on their human capital accumulation, educational attainment, and early adult criminal justice outcomes. First, I construct a new administrative panel linking more than a decade of housing assistance wait list applications in New Haven, Connecticut, to statewide education, labor market, and, separately, criminal justice records. Second, I develop a research design based on a series of lotteries that randomly ordered applicants on wait lists for housing vouchers and public housing units. I use lottery-induced variation in the date a household is selected from these wait lists as an instrument for whether a child receives assistance and, for longer-run outcomes, the number of childhood years spent with housing subsidies.

My primary analysis relies on application records for more than 5,500 children living in

[1] The neighborhood channel may be particularly important, given evidence that childhood exposure to neighborhoods causally impacts long-run outcomes (Chetty et al. 2016; Chyn 2018; Katz 2026).

[2] Recipients of housing assistance typically pay 30% of income, including cash transfers, in rent.

New Haven county linked to administrative data on annual school-reported homelessness and attendance, kindergarten readiness, standardized achievement tests in elementary and middle school, high school graduation, and college enrollment. These linked records allow me to evaluate the effects of housing assistance on residential stability, human capital accumulation, and educational attainment. Because lotteries occur over time and children differ in age at the time of application, I use different subsamples across outcomes. I additionally link adult household members to quarterly earnings records to characterize changes in household resources, and children to felony and misdemeanor conviction records to examine early adult criminal justice outcomes.[3]

I estimate the causal impacts of housing vouchers and public housing on children using nine wait list position lotteries held between 2003 and 2018. Because randomly-assigned positions were often overwritten, I instrument for receiving assistance with the timing of selection for an eligibility interview, which I refer to as an "offer".[4] In separate estimates for children from ages six to seventeen, an offer raises annual assistance receipt by more than 45 percentage points at any age, with first-stage F-statistics close to 1,000. Because the timing of offers depends on an applicant's randomly-assigned position and on unit or voucher turnover, I check for balance with respect to household characteristics that may also affect child outcomes. After including fixed effects that absorb differences in offer timing across lotteries for households with differing unit eligibility, I find no remaining significant differences in pre-application demographic and labor market characteristics between children in offered and not-offered households at any age.

Using this research design, I first show that housing assistance has long effects on the home environment by reducing children's exposure to school-reported homelessness and housing instability. Pooling annual observations from ages six through eleven, receiving assistance by a given age reduces the probability of school-reported homelessness or housing instability by 4.3 percentage points ($p = 0.003$), relative to a complier control mean of 6.7%. The reduction is largest at ages six to seven, when assistance lowers this probability by 6.9 percentage points ($p < 0.001$), relative to a complier control mean of 9.0%. Estimated reductions remain negative at older ages, although precision declines with age. Pooling ages six through eleven, assistance reduces non-shelter housing instability by 2.8 percentage points ($p = 0.027$), driven almost entirely by a 2.9 percentage

[3]Because of my Data Use Agreement (DUA) with DataLinkCT, I am unable to directly combine the criminal justice records, an external data linkage, with the linked education and labor market data provided through DataLinkCT.

[4]Roughly 60% of children in the sample are in households that eventually receive an offer before they reach the age of 18. Binary treatment specifications use an offer indicator, while exposure specifications use years since selection for the orientation interview.

point reduction in doubled-up homelessness ($p = 0.012$). The estimated reduction in sheltered homelessness is 1.6 percentage points ($p = 0.037$), while effects on living in a hotel or motel – a rare outcome – are small and statistically insignificant. These findings complement Collinson and Gibbs 2025, who find that both housing vouchers and public housing reduce sheltered homelessness and doubling up among families with young children in Miami-Dade county.

Second, using records for more than 2,000 children that enter kindergarten after these wait list lotteries, I find evidence of improvements in boys' early developmental outcomes. Boys whose families receive assistance by the kindergarten assessment year score 0.21 standard deviations higher on an index averaging teachers' assessments across six categories ($p = 0.113$). The largest category-level point estimates for boys are in creativity, physical and motor development, and personal and social development. The overall estimate for girls is small and imprecise. In the joint program specification, the unit-minus-voucher difference in the developmental skills average is positive but imprecisely estimated ($p = 0.105$). My results are similar to Collinson and Gibbs 2025, who find that positive effects of housing vouchers on early literacy and kindergarten readiness are driven primarily by boys. I also find no consistent effects on elementary or middle school achievement for either gender, consistent with evidence of test score "fade out"from studies of other early childhood interventions (Currie and Thomas 1995; Deming 2009).

Third, I examine longer-run outcomes using education and criminal justice records. I focus primarily on educational attainment and criminal justice involvement between ages twenty and twenty-two. Among compliers, the estimated effects of housing assistance are positive but imprecise: assistance increases high school graduation by age 20 by 14.9 percentage points ($p = 0.145$) and college enrollment by age 20 by 19.1 percentage points ($p = 0.063$). To exploit the wide variation in childhood exposure to assistance, I estimate a dose-response design and find that, under a constant marginal effects assumption (Chyn 2018; Chaudhry and Eng 2024), each additional year of assistance increases high school graduation by age 22 by 2.9 percentage points ($p = 0.099$) and college enrollment by age 22 by 4.2 percentage points ($p = 0.020$). The effect on college enrollment is concentrated in the four-year sector. The attainment gains are especially large for voucher recipients, and I find suggestive evidence of larger college enrollment effects among girls. Finally, housing assistance reduces the probability of a felony conviction by age 22 by 7.4 ($p = 0.097$) to 9.8 ($p = 0.045$) percentage points – depending on whether the control group is limited to applicants from New Haven county or across Connecticut – with larger negative point estimates among boys. Together, these results indicate that housing assistance can generate lasting gains in educational attainment and reduce involvement with the criminal

justice system early in adulthood.

To better understand these longer-run effects, I examine mechanisms related to school attachment, school quality, and household resources. Housing assistance strengthens children's attachment to school early in childhood, raising attendance and reducing chronic absenteeism from ages six to eleven. These improvements do not appear to operate through access to higher-performing schools, and I find small and imprecise effects on school test score value added or average proficiency in elementary and middle school. If anything, assisted children attend schools with more economically disadvantaged peers at older ages. For example, assistance increases the share of students eligible for free or reduced-price lunch at the child's school by approximately 6 to 7 percentage points at ages fourteen to seventeen. In sum, the attainment gains appear more consistent with increased housing stability and school continuity than with moves to observably better schools.

Finally, to benchmark the cost-effectiveness of government spending on housing assistance, I map the estimated effects on educational attainment into predicted impacts on lifetime earnings using estimates from Zimmerman 2014 and Hendren and Sprung-Keyser 2020, and I use other external estimates to calibrate the reductions in government costs from sheltered homelessness and crime (McCollister et al. 2010). Under an assumption that children's projected after-tax earnings enter willingness to pay, I find that housing assistance returns between $2.05 and $2.26 per dollar of net government spending, depending on whether assistance is valued at $0.83 per dollar or equivalently to cash (Reeder 1985). When vouchers and units are valued by households equivalently to cash, I estimate that each dollar spent on housing vouchers returns around $2.35 in benefits, while each dollar spent on public housing returns around $1.49. In this setting, differences in the MVPF arise from larger effects on college enrollment among voucher recipients.

Related literature This paper contributes to two strands of research. First, it adds to a large literature that studies how means-tested housing assistance affects children. Researchers often exploit variation in childhood exposure to public housing or housing vouchers and find positive effects on outcomes ranging from test scores to early-adult earnings and health outcomes (Currie and Yelowitz 2000; Schwartz et al. 2020; Pollakowski et al. 2022; Chaudhry and Eng 2024; Han and Schwartz 2024; Beauregard 2025; Ribeiro and Mariante 2025). A smaller set of studies uses lottery-ordered wait lists to estimate the effects of housing assistance on adults (Jacob and Ludwig 2012; van Dijk 2019; Carr and Koppa 2020) and children (Jacob et al. 2013, 2015), finding mixed results. The closest paper to this one is Collinson and Gibbs 2025, which studies a combined housing voucher

and public housing lottery in Miami-Dade County and finds improvements in early childhood skills, concentrated among families that move to higher-opportunity neighborhoods or higher-quality public housing developments. I add to this literature by combining data on longer-run schooling and criminal justice outcomes, allowing me to examine changes in housing stability, school attachment, educational attainment, criminal justice involvement, and household resources within the same setting.

Second, this paper contributes to a broad literature on in-kind transfers and children's long-run outcomes, especially their human capital development. Most transfers provided to low-income families in the United States are in kind rather than in cash (Currie and Gahvari 2008), and a large empirical literature finds that childhood access to the social safety net improves later-life health, educational attainment, and earnings (Currie and Thomas 1995; Kline and Walters 2016; Hoynes et al. 2016; Brown et al. 2020; Bailey et al. 2021; Gray-Lobe et al. 2023; Bailey et al. 2024; Currie and Van Parys 2026). My results suggest that protection against adverse shocks – such as homelessness or housing instability – may be one channel through which in-kind transfers improve long-run outcomes, even in the absence of significant improvements in intermediate outcomes like elementary and middle school test scores.

The remainder of the paper proceeds as follows: Section 2 describes the institutional context. Section 3 describes the data sources and samples and provides descriptive statistics. Section 4 introduces the research design. Section 5 presents the results on short-run outcomes. Section 6 presents the results on longer-run outcomes. Section 7 discusses the cost effectiveness of housing assistance. Section 8 concludes.

2 Background

2.1 Housing assistance programs in New Haven

The Department of Housing and Urban Development (HUD) provides long-term rental assistance for low-income households in two categories: rental vouchers, known as *Housing Choice Vouchers* (HCVs), and subsidies tied to specific units. HCVs subsidize leases of qualifying units in the private rental market. Unit-based assistance includes traditional *Low Income Public Housing* (LIPH) and *Project Based Vouchers* (PBVs), both of which subsidize leases in specific developments. These programs are funded by HUD but administered by more than 3,300 local Public Housing Agencies (PHAs). In 2023, HUD provided rental subsidies to around 4.5 million households through these programs at an annual

cost of more than $54 billion.[5]

These two forms of assistance share similar eligibility and subsidy rules but differ in how they provide housing subsidies.[6] For both units and vouchers, eligibility depends primarily on household income, and recipients contribute 30% of household income toward rent.[7] For tenant-based vouchers, the PHA pays a participating landlord the difference between the tenant's contribution and the unit's contract rent, up to a locally determined payment standard, typically set at an estimate of the 40th percentile rents in a metropolitan area. In order to successfully use a voucher, recipients must locate a qualifying unit and a landlord willing to accept the voucher.[8] In the LIPH and PBV programs, the subsidy is attached to an apartment either owned by the PHA or operated through a long-term agreement between the PHA and the property owner. Recipients of unit-based assistance face lower search frictions in the rental housing market but have less choice over the unit and neighborhood in which they live (Ellen et al. 2024).

Housing assistance in New Haven This study examines housing assistance administered by the PHA serving New Haven, Connecticut. New Haven is a relatively low-income city surrounded by a more affluent metropolitan area: in 2023, median household income in New Haven was roughly $51,000, compared with almost $81,000 in the surrounding area, and around one-in-four residents lived below the poverty line. The PHA provided assistance to nearly 13,500 residents in 2023, around 10% of the population of New Haven.

Housing assistance in New Haven is a large transfer relative to recipients' resources. In 2023, the average PHA household received an annual subsidy of $14,292 and paid $5,520 toward rent. Pooling years from 2012 to 2023, the average PHA voucher subsidy was in the 88th percentile of all PHAs, while the average public housing subsidy was in the 99th percentile. The full distribution of average monthly rents and subsidies across PHAs is reported in Appendix Figures A.1 and A.2. Average PHA household income in 2023 – including cash transfers and other non-wage income – was only $19,234. That household

[5]National and local program statistics in this subsection are calculated using HUD's Picture of Subsidized Households dataset.

[6]Households are eligible for HCVs when their annual income is below 50% of the median income in their metropolitan area (AMI). Households are eligible for PBVs and LIPH when their annual income is below 80% of AMI.

[7]Annual household income may be adjusted for dependents, elderly or disabled household status, and certain eligible expenses. Households receive a deduction of $480 per eligible dependent and a single $400 deduction if they qualify as an elderly or disabled household. Qualifying elderly or disabled households may also deduct medical expenses in excess of 3% of annual income.

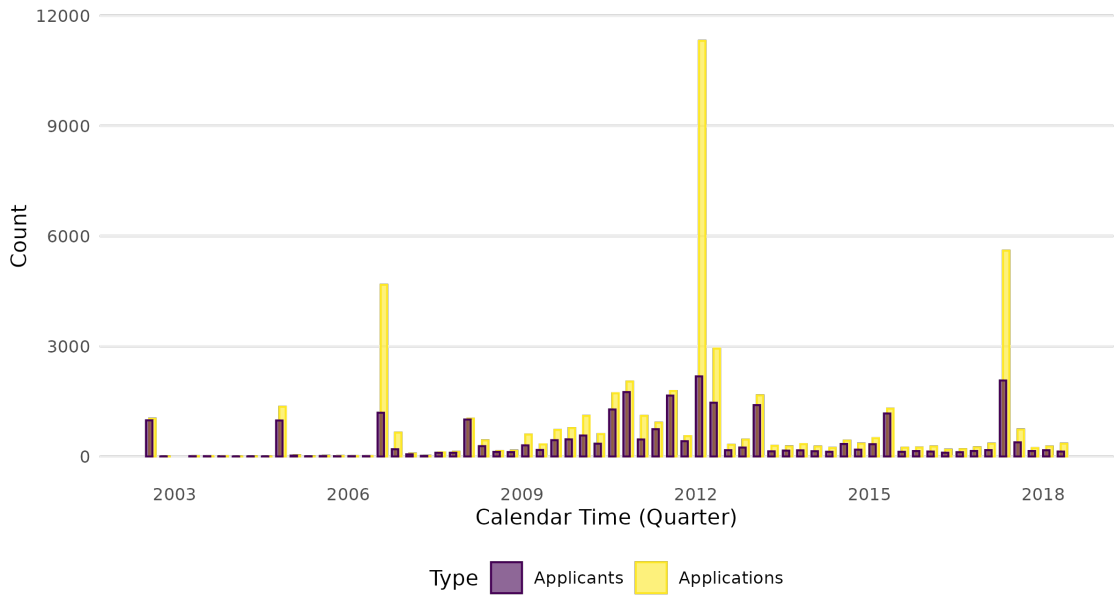
[8]A unit will be declared ineligible if it either (i) fails a health and safety inspection or (ii) is determined to not have a "reasonable" contract rent (i.e., one in line with market rents for comparable units).

incomes are well below even the voucher eligibility cutoff for one-person households of $40,150 is consistent with nationwide evidence that housing assistance recipients are negatively-selected on income relative to the eligible population (Rafkin et al. 2025).

2.2 PHA wait lists

Unlike many social safety net programs in the United States, housing assistance is not an entitlement: because of limited federal funding, income-eligible households are not guaranteed a voucher or unit, and local PHAs use wait lists to ration assistance. Wait list policies vary across PHAs, from indefinitely open, first-in-first-out queues to lists that accept applications only during infrequent opening periods (Collinson et al. 2015; Carr and Koppa 2020; Waldinger 2021). During the period that I study, the PHA typically opened wait lists only when it either expected that the existing queue could be processed within a few months (Bromage 2005) or needed to fill units in recently rehabilitated developments. In keeping with HUD guidelines, the PHA announced the beginning and end of each application period through newspapers, flyers, and online communications. With limited exceptions, applicants could join a list only while it remained open. Appendix B describes these opening periods in more detail.

Figure 1: PHA wait list applications over time



Notes: The y-axis is the count of applications (applicants) to one of the 24 primary wait lists maintained by the PHA for non-elderly households. The x-axis is the calendar time in quarters from January 2003 to December 2018. The data source is the universe of PHA wait list applications.

The PHA uses two methods to order applicants during the study period. During some openings, applicants were ordered by the exact date and time of application (i.e., first-in-first-out). During others, the PHA assigned each applicant a random number that determined whether they were placed on the wait list and, conditional on placement, their initial wait list position. Applicants assigned numbers below a predetermined cutoff were added to the list in random number order behind any existing applicants. The PHA does not retain records of unplaced applicants.

The PHA manages a combination of development-specific and general wait lists in order to allocate vouchers and unit-based assistance. From 2003 to 2018, the PHA received more than 75,000 applications from more than 25,000 households across 81 wait lists. As shown in Figure 1, the number of applications noticeably spikes during, for example, a set of 2007 public housing wait list openings. Across the 24 primary wait lists for non-elderly households, the average number of applications per wait list is 2,160, equivalent to 34 per quarter.

3 Data and matching

The data in this paper combine wait list application and income certification records from the PHA with administrative data accessed through the Connecticut Office of Policy and Management's DataLinkCT initiative and public records from the Connecticut Judicial Branch. I describe these data sources below.

Housing assistance data I obtain data on applications to housing assistance wait lists through a partnership with the PHA in New Haven, Connecticut. For each applicant-by-wait list combination, I observe the exact date and time of application, the name of the wait list, the household's wait list position after application, and the household's most recent reported address. Each application record contains information on every household member's gender, age, and race.[9] For applicants that receive housing assistance from the PHA, I additionally observe annual income certification records that are used to verify household eligibility and calculate rent payments.[10] Certification records also contain information on point-in-time household address and formal and informal earnings. Finally, I observe dates of participation in various community and economic development (CED) programs run by the PHA.

[9]Each application contains the date of birth for all applicant household members. My Data Use Agreement (DUA) with DataLinkCT does not allow me to observe individual ages in the matched sample.

[10]These records form the basis for the Form 50058 submissions made to HUD.

Education data I obtain data on school attendance and test scores from the Connecticut State Department of Education (SDE), which maintains records on public K-12 educational programs and postsecondary enrollment for Connecticut public high school graduates. DataLinkCT matches the children of applicants to PHA wait lists to SDE records on the universe of New Haven county public school children from the 2010-2011 to 2023-2024 school years using name, date of birth, and the last 4 digits of the Social Security number (SSN).[11] Importantly, a student is observed in the matched file in each year they attend *any* Connecticut public school if they *ever* attend a public school within New Haven county. SDE records include information on kindergarten achievement tests, end-of-year test scores for students in grades 3 through 8, annual measures of students' school attendance and housing insecurity, and the date of high school graduation. Kindergarten test scores, unlike the remaining education records, are available as early as the 2008-2009 school year. Test scores in grades 3 through 8 are only available beginning in the 2014-2015 school year.

In addition to the K-12 education records, I observe information on enrollment in postsecondary education for the universe of Connecticut public high school graduates that ever attend a New Haven county school. These students are matched to the National Student Clearinghouse (NSC) if they have postsecondary enrollment dates from May 2011 to August 2024. The NSC file contains information on enrollment dates, enrollment status, and institution level. For students who graduate before August 2024, I observe the degree type, major, and graduation date. The data cover both two-year and four-year degree programs. My analysis of higher education outcomes in Section 6 is restricted to children that ever match to SDE records. I describe the SDE data files, including the NSC records, in more detail in Appendix C.1.

Labor market data I also obtain data on quarterly earnings and employment from 2001-2021 from the Connecticut Department of Labor (DOL), which maintains records of earnings reported by employers in Connecticut. DataLinkCT matches all household members in PHA wait list applications to DOL records on the universe of Connecticut residents. I keep only unique matches and inflate earnings to be in 2023 dollars. Employer identifiers are not available, but each observation, which is at the employer-by-individual-by-quarter level, does record the employment sector (i.e., 6-digit NAICS code). I describe the DOL data file in more detail in Appendix C.2. I use these records to measure pre-application household employment and earnings, which enter some specifications as controls, and to evaluate changes in adult earnings and employment among applicant

[11] The matching is conducted by DOL staff through the DataLinkCT initiative.

households.

Criminal justice data Finally, I use court records from the Connecticut Judicial Branch to measure arrests that result in misdemeanor or felony convictions. I identify convictions from recorded guilty findings and measure the timing of these outcomes using the corresponding arrest date. I link court records to individuals in PHA application files using name and date of birth, separately from the DataLinkCT linkage. I require exact matching on date of birth and fuzzy matching on name. Conviction records are available from January 2000 to August 2024. Appendix C.3 describes these records and the matching process in more detail.

3.1 Sample construction

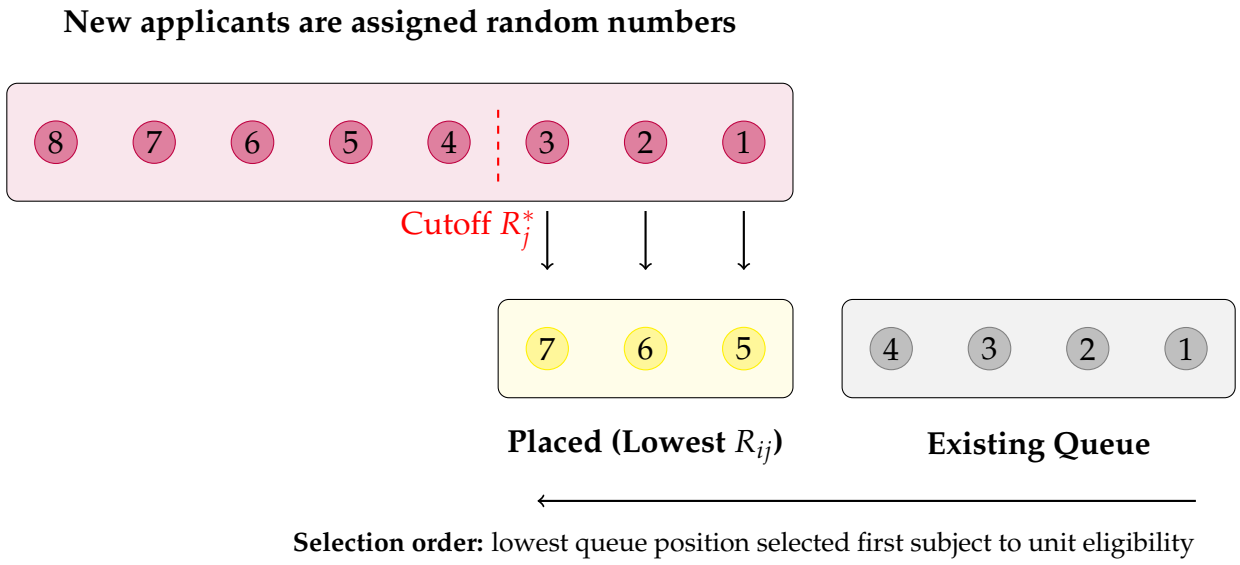
To construct the analysis sample, I begin with the universe of successful applications to the PHA's various housing assistance wait lists as of March 13, 2025, which is the sample provided by the PHA for matching through DataLinkCT.[12] This dataset contains almost 329,000 applications to 100 different wait lists by more than 80,000 households. I make a number of restrictions to arrive at my final analysis sample.

First, because the PHA indefinitely opened many of its housing assistance wait lists beginning in June 2019, I limit the sample to applications from 2003 to 2018, prior to the indefinite opening policy and after the PHA begins recording new program admissions in 2002. This restriction reduces the sample to around 75,000 applications to 81 different wait lists by more than 25,000 households. Second, I remove applications from outside of New Haven county. In Appendix Figure A.3, I report the total number of unique applicants and the share of applicants from outside New Haven County in each year from 1998 to 2024. Only 10.3% of applicants from 2003 to 2018 are removed under this restriction. I then remove specialty wait lists – such as the "LIPH Community Re-Entry" wait list for applicants exiting the correctional system – wait lists for developments that are restricted to elderly or disabled applicants, and wait lists for internal transfers and portability vouchers.[13] I refer to this limited set of applications as the *descriptive sample*. This sample contains more than 44,000 applications to 24 wait lists by more than 16,000 households. Roughly 46% of these applications are submitted by households containing children.

[12]Successful applications are applications that result in placement onto a wait list. During wait list position lotteries, many applicants are not placed and their applications are discarded. In non-lottery periods, all applications are successful.

[13]Portability is a feature of the Housing Choice Voucher program that allows recipients to move - while retaining their voucher - to the jurisdiction of a new PHA after 12 months in the program.

Figure 2: Wait list position lotteries



Notes: Figure 2 depicts a stylized version of the wait list assignment process used during position lotteries. Numbers in purple are generated using a random number generator. Queue position depends on both random number and the size of the existing queue. Only applicants with a random number that falls below the predetermined lottery cutoff R_j^* are placed (1, 2, and 3 in this illustrative example). My dataset contains the *placed* (yellow) and *existing* (grey) wait list participants **but not** applicants to the left of the random number cutoff (4-8 in this illustrative example) for each wait list lottery j. I observe the entire household of all placed and existing applicants for each wait list opening.

Openings and lotteries I next restrict the sample to applications that occur during clearly-defined wait list opening periods. To identify wait list openings, I cross-reference monthly application counts with local news articles and the PHA's annual reports. I describe the process of identifying wait list openings and lotteries – opening periods during which applicants are placed onto a wait list in random order – in Appendix B. Figure 2 presents a stylized example of a wait list lottery, in which applicants are assigned random numbers and a subset are placed onto an existing wait list in random order. Using a subsample of applicants with wait list position information, I am able to identify lotteries by testing for a relationship between an applicant's assigned position and their implied position given the exact date and time they apply to a wait list. I treat an opening as a lottery if I both fail to reject the null that the correlation between these two positions is zero and estimate a correlation of less than 0.25 in magnitude. Point estimates and 95% confidence intervals from this regression test are presented in Appendix Figure B.2.

Because my focus is on the effects of housing vouchers and public housing relative to a counterfactual of no housing assistance, I additionally exclude households living in public housing or using a voucher at baseline. The final lottery sample contains 7,399 children

from 4,194 households. Wait list lotteries occur between March 2003 and January 2018 across 9 opening periods and 14 wait lists.[14] I summarize the number and characteristics of children in each lottery in Appendix Table A.1.

Matching I work with DataLinkCT and the PHA to match wait list applications to employment records from DOL and education records from SDE using name, date of birth, and the last 4 digits of the Social Security number. I refer to the *analysis sample* as the set of lottery sample applications for which I match a child to any SDE record. The matched SDE analysis sample contains 5,590 children from 3,480 households (the child-level match rate is 75.6%). The matched SDE and DOL linked sample contains 5,505 children from 3,417 households. My sample includes all outcomes for students that enroll in any New Haven county school in any school year, so I observe educational outcomes for students who exit New Haven county but remain in Connecticut.[15] I fail to observe educational outcomes for children who never attend a public school in New Haven county (e.g., the child of an applicant who moves to Bridgeport after the lottery). My main results in Section 5 use different subsets of the matched SDE sample, depending on child age at the time the outcome is measured.[16] Finally, I confirm that children in households that are selected from a wait list are not more likely to match to the administrative records. In Appendix Table A.2, I show that match rates are nearly identical ($p = 0.305$) for children in offered and not offered households after adjusting for differences in child characteristics. Unadjusted match rates are actually lower for offered children.

3.2 Summary statistics

Table 1 reports summary statistics across the descriptive, lottery, and matched samples. The children of applicants appear similar along observable dimensions across samples. In the sample of lottery applicants that can be matched to DOL earnings records, average household W-2 earnings are just above $16,500, though many families likely supplement these earnings with cash transfer programs. The typical child is young at the time of first application, averaging five years of age across samples. Appendix Figure A.5 plots the full distribution of child ages at the time of first application in the descriptive and lottery samples. Across each subsample, around half of children are female, and 74 percent of children in the matched lottery sample live in households with non-zero recorded W-2

[14] Two opening periods contain multiple development-specific wait lists.

[15] I include a map of Connecticut counties in Appendix Figure A.4.

[16] My Data Use Agreement (DUA) with DataLinkCT does not allow me to observe child age in the matched sample. I impute age from information on child grade derived from test score and school attendance records under the assumption that each child is 5 at the time they first enroll in kindergarten.

earnings in the year prior to application. While only 30 percent of lottery children are in households that receive assistance, these transfers are long-lived conditional on receipt. Limiting to lottery applicants that receive assistance by 2015, I find that 62 percent of voucher and 69 percent of unit participants are still receiving assistance nine calendar years after entry. The assisted share in each post-entry year is reported in Appendix Figure A.6.

Table 1: Applicant child summary statistics

	Descriptive	**Lottery**		
		All	**SDE**	**SDE & DOL**
Panel A: Child and household characteristics				
Black	0.55	0.57	0.58	0.58
Female	0.49	0.49	0.49	0.49
Age	5.62	5.68	5.12†	5.11†
Household size	3.89	3.79	3.79	3.81
Household head is citizen	0.50	0.51	0.54	0.54
Household head is disabled	0.09	0.06	0.07	0.07
Household head age	31.09	30.44	–	–
Household earnings (2023 dollars)	–	–	–	$16,524.87
Household employment	–	–	–	0.74
Panel B: Assignment outcomes				
Offered by age 18	0.62	0.61	0.61	0.61
Assisted by age 18	0.30	0.26	0.30	0.30
Wait time (yrs)	3.63	4.13	4.16	4.17
Panel C: Sample counts				
# Children	13,815	7,399	5,590	5,505
# Applications	8,164	4,328	3,587	3,524
# Wait Lists	24	14	14	14
# Households	7,565	4,194	3,480	3,417

Notes: Panel A of Table 1 presents averages for the children of applicants to PHA wait lists. Age is measured for each child at the time of application. Earnings is total household earnings reported by DOL in the year prior to the application year (measured in 2023 dollars). Employed is an indicator equal to 1 if household earnings in the year prior to application are non-zero. Panel B of Table 1 reports the share of children in each sample in households that receive interview offers or housing assistance by age 18 (or the last recorded age). Average wait time is the average time between the child's earliest application and their earliest date of program participation, reported conditional on receiving assistance. The data source is PHA wait list application records from 2003-2018 and DOL earnings records from 2001-2021. The SDE column includes all households with at least one child aged 0-17 at the time of application who matches to the SDE records. I impute age from the earliest observed child grade in these columns, which are marked with a †. The SDE & DOL column is limited to households with at least one adult who matches to the DOL records. DOL and SDE matches require exact matching.

4 Empirical strategy

4.1 Instrumental variables specification

I use variation in when households reach the top of lottery-ordered waiting lists to estimate the effects of housing assistance on children. Let Y_{ijg} denote an outcome for child i, whose household applied during wait list opening j, measured at age g. I define the indicator $Z_{ijg} = 1$ if the household has been selected from the wait list for an eligibility interview before the outcome is measured at age g. I refer to this selection as an offer. Offer status is absorbing: once a household has been selected from list j, Z_{ijg} remains equal to one at all future ages. I next define $D_{ijg} = 1$ if the household has received housing assistance at any point between opening j and the measurement of the outcome at age or grade g. Like the offer instrument, treatment is absorbing, and a household that enters and subsequently exits housing assistance remains treated.

For each age group or grade, I estimate the following first and second-stage equations by two-stage least squares:

$$D_{ijg} = \pi_b + \delta_b Z_{ijg} + X'_{ijg}\Psi_b + \lambda_{j,s(i)} + \nu_{ijg}, \tag{1}$$

$$Y_{ijg} = \alpha_b + \beta_b D_{ijg} + X'_{ijg}\Gamma_b + \lambda_{j,s(i)} + \varepsilon_{ijg}, \tag{2}$$

where $b = b(g)$ denotes the age bin containing age g. The fixed effects $\lambda_{j,s(i)}$ are defined by interacting wait list opening j with the household-size-by-disability-status cell $s(i)$. These interacted fixed effects account for differences across openings and the administrative rules that match households to vouchers or available units based on household size and disability status. The vector X_{ijg} contains predetermined child and household characteristics, including race, gender, and age at application, as well as the outcome year. Standard errors are clustered at the household level to allow for dependence across siblings and repeated observations of the same child at different ages.

I also estimate the reduced form effect of an offer by replacing D_{ijg} in Equation (2) with Z_{ijg}. The reduced form estimate captures the effect of being selected from the wait list by age g, while δ_b measures the first-stage effect of selection on assistance. In an exactly identified specification, the 2SLS coefficient β_b scales the reduced form effect by the first stage.

For annual outcomes constructed from school records, I often stack child-year observations within the two-year age bins 6–7, 8–9, 10–11, 12–13, 14–15, and 16–17 and estimate Equations (1)–(2) separately in each bin. For other outcomes, the age bin equals the child

age. The coefficient β_b allows the effect of entering assistance to vary with the age at which the outcome is measured. For outcomes measured at a terminal age, such as college enrollment by age twenty or twenty-two, I retain one observation per child-lottery and estimate a separate specification for each outcome and terminal age. In these specifications, I use $D_{ij,18}$, an indicator for any assistance receipt through age 18, and $Z_{ij,18}$, an indicator for any offer through age 18, as the endogenous treatment and binary instrument.

4.2 Identification and interpretation

A causal interpretation of β_b requires four conditions. First, conditional on the $\lambda_{j,s(i)}$ fixed effects and X_{ijg} controls, offer timing must be independent of children's potential outcomes. While initial position is randomly assigned within each lottery, the Z_{ijg} instrument also depends on voucher or unit turnover and household eligibility for particular units. By assumption, then, unit availability and wait list processing must not select households based on unobserved determinants of children's outcomes after conditioning on unit eligibility.[17]

Second, wait list selection must have a nonzero effect on assistance receipt. I assess this relevance condition by evaluating the first stage directly for each sample. Third, for β_b to have a treatment effect interpretation, selection from the wait list cannot cause households to select out of assistance. This monotonicity assumption holds by design. Fourth, wait list selection can affect children's outcomes only through receiving housing assistance. The exclusion restriction in this setting rules out direct effects of the orientation interview or access to other services associated with the offer.

Under these assumptions, β_b identifies the local average treatment effect of receiving housing assistance by age $b(g)$ for households induced to receive assistance by wait list selection. Because D_{ijg} records having any assistance by age g, each β_b captures the effect of gaining access to assistance, including the timing and duration of the resulting assistance spell.

[17] A useful benchmark is $Z_{ijg} = \mathbf{1}\{R_{f(i)j} \leq K_{j,s(f(i)),t(g)}\}$, where $f(i)$ denotes child i's household, $R_{f(i)j}$ is the household's randomly assigned position in opening j, $s(f(i))$ is its baseline household-size-by-disability-status cell, and $K_{j,s,t}$ is the largest position processed from that cell by date t. If $K_{j,s,t}$ is determined by voucher or unit availability and administrative capacity alone, then, conditional on j, s, and t, offers depend only on random position. In practice, the PHA may extend the cutoff when earlier applicants do not respond, fail eligibility verification, or decline assistance. Realized offers are then not exactly randomized. I assume the resulting dependence is negligible under a large wait list approximation (De Chaisemartin and Behaghel 2020).

4.3 Cumulative exposure specifications

For some long-run outcomes, I complement the design in Equations (1)–(2) by replacing the binary endogenous treatment D_{ijg} with the number of years A_{ij} a child has received assistance before age eighteen and instrument it with the number of years since the household's first offer E_{ij}. In practice, I estimate the following system of equations by two-stage least squares:

$$A_{ij} = \pi + \rho E_{ij} + X'_{ij}\Psi + \lambda_{j,s(i)} + \nu_{ij}, \tag{3}$$

$$Y_{ij} = \alpha + \beta A_{ij} + X'_{ij}\Gamma + \lambda_{j,s(i)} + \varepsilon_{ij}. \tag{4}$$

This design allows me to exploit the total variation in exposure to assistance generated by offer timing. Interpreting β as the effect of one additional year with housing assistance requires stronger assumptions than in the binary treatment specification. In particular, offer timing must weakly increase the number of years with assistance, and the effect of each year with assistance must be constant. This "dosage" model assumes that the impacts of assistance do not depend on the child's age at receipt. Other studies of housing assistance programs find support for this assumption in similar contexts (Chyn 2018; Chaudhry and Eng 2024; Beauregard 2025; Ribeiro and Mariante 2025).

4.4 Balance and first stage

I begin by validating the primary research design, checking that the realized timing of offers is unrelated to households' economic circumstances or labor market trajectories before application. In practice, I test whether household offer status at different ages is related to pre-application household employment and earnings or child demographics.

For each age bin b, I use the same sample and offer instrument Z_{ijg} as in Equations (1)–(2). For each predetermined characteristic X_i^k, I estimate

$$X_i^k = \alpha_b^k + \beta_b^k Z_{ijg} + A'_{ijg}\rho_b^k + \lambda_{j,s(i)}^k + \varepsilon_{ijg}^k, \qquad g \in b, \tag{5}$$

where A_{ijg} contains application age, calendar year, and years since lottery fixed effects, and $\lambda_{j,s(i)}^k$ are the lottery-by-size-by-disability fixed effects, which control for characteristics used by the PHA to match applicants to available units. I cluster standard errors at the household level.

I check for balance based on child gender and race, household employment and earnings in the year before application, and changes in household employment and earnings over

the two years preceding application. I report the coefficient on offer status for each characteristic and the p-value from a joint test that all β_b^k coefficients are zero in Table 2. While pre-application changes in earnings are marginally significant at older ages, I fail to reject the joint null that all coefficients are equal to zero at any age. The results are similar if I consider balance in the timing of waiting list selection after a lottery.

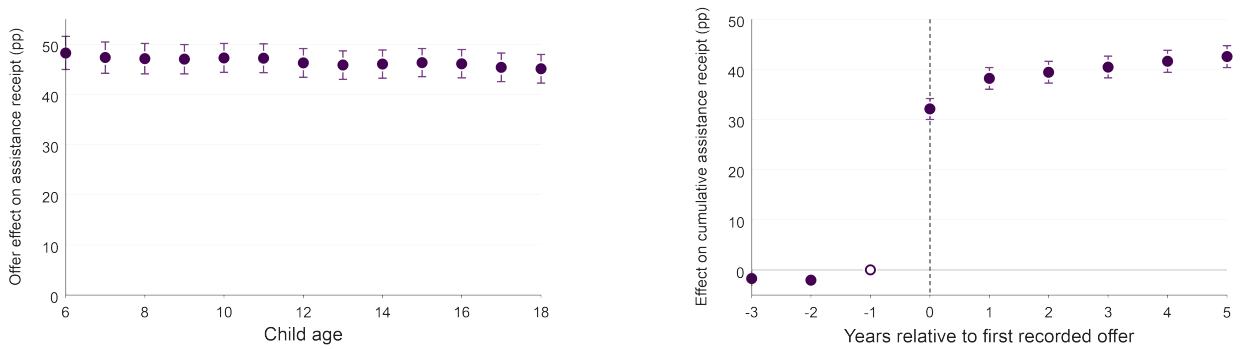
Table 2: Offer timing balance by child age

	$Z_{ijg} = \mathbf{1}\{\text{Selected from list by age } g\}$					
	Ages $6-7$	**Ages** $8-9$	**Ages** $10-11$	**Ages** $12-13$	**Ages** $14-15$	**Ages** $16-17$
Female	-0.007	-0.021	-0.017	-0.013	-0.014	-0.008
	(0.019)	(0.018)	(0.017)	(0.017)	(0.016)	(0.016)
Black	0.009	0.006	-0.003	-0.006	-0.004	0.002
	(0.020)	(0.020)	(0.019)	(0.020)	(0.020)	(0.020)
Any employment, prior year	-0.023	-0.015	-0.015	-0.018	-0.019	-0.014
	(0.020)	(0.019)	(0.019)	(0.019)	(0.019)	(0.020)
Employment change, years -2 to -1	-0.006	-0.008	-0.008	-0.014	-0.009	-0.010
	(0.019)	(0.019)	(0.019)	(0.019)	(0.019)	(0.020)
Earnings, prior year ($1,000)	0.460	0.090	0.037	-0.166	-0.406	-0.309
	(0.782)	(0.773)	(0.784)	(0.844)	(0.858)	(0.878)
Earnings change, years -2 to -1 ($1,000)	0.648	0.698	0.934*	0.941*	0.819	0.895*
	(0.472)	(0.469)	(0.478)	(0.489)	(0.501)	(0.510)
Joint test p-value	0.584	0.583	0.364	0.305	0.461	0.492
Child-year observations	6,988	8,087	9,059	9,829	10,360	10,769
Children	3,663	4,167	4,639	4,991	5,239	5,431
Households	2,643	2,879	3,086	3,221	3,318	3,384
Share offered	0.476	0.521	0.554	0.579	0.598	0.606

Notes: Table 2 reports conditional balance tests for the binary offer instrument Z_{ijg}. The sample excludes households receiving assistance before the lottery. Each entry is the coefficient from a separate regression of the listed predetermined characteristic on an indicator for selection from a wait list by the age at which the observation is measured. All specifications include age at lottery and age at outcome year, outcome year, and lottery-by-household-size-by-disability-status fixed effects. Standard errors are clustered at the household level and appear in parentheses. Earnings variables measure W-2 earnings and are reported in thousands of 2023 dollars at the household level. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Relevance Reaching the top of a lottery-ordered wait list has a large, positive effect on housing assistance take-up. Figure 3 reports the static and dynamic relationships between offer timing and the receipt of assistance. Because the offer represents selection from the top of the wait list, incomplete take-up is driven by a combination of non-response, ineligibility, and failure to lease up: Waldinger 2021 documents significant non-response to notifications on housing assistance wait lists in Cambridge, Massachusetts. The static estimates show that receiving an offer raises the probability of receiving assistance by a given age by more than 45 percentage points. In the dynamic event study specification (Sun and Abraham 2021), the probability of receiving assistance after an offer rises slightly over time to over 40 percentage points at year five. Table 3 reports corresponding first-

Figure 3: Wait list selection first stage



(a) Static first stage

(b) Dynamic first stage

Notes: Figure 3 plots the first stage relationship between reaching the top of a wait list and subsequent receipt of housing assistance. Figure 3a measures assistance at each age from six to eighteen. Figure 3b measures the dynamic effect of being selected from a wait list on being assisted, using the method in Sun and Abraham 2021. Whiskers report 95% confidence intervals, where standard errors are clustered at the household level. The data source is the lottery sample of PHA applications from 2003-2018.

stage estimates of Equation (1) for each two-year age bin. Being offered by age bin b raises the probability of receiving assistance by roughly 45 percentage points in the matched sample, with first stage F-statistics above 900 at all ages.

Table 3: Wait list selection first stage by child age

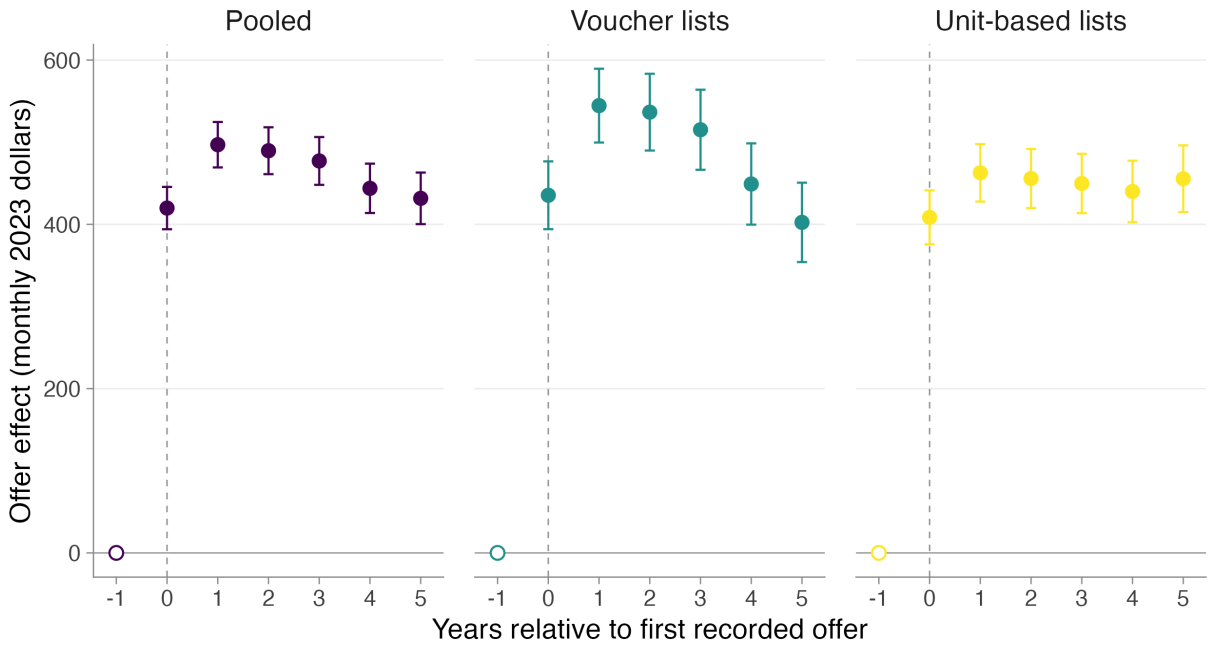
	Outcome: assistance received by age b, D_{ijb}					
	Ages 6–7	Ages 8–9	Ages 10–11	Ages 12–13	Ages 14–15	Ages 16–17
Offer indicator (Z_{ijb})	0.476***	0.468***	0.469***	0.456***	0.457***	0.452***
	(0.016)	(0.015)	(0.014)	(0.014)	(0.014)	(0.014)
Outcome mean	0.24	0.26	0.28	0.29	0.29	0.30
Demographic controls	Y	Y	Y	Y	Y	Y
Lottery × size × disability fixed effects	Y	Y	Y	Y	Y	Y
Age and timing fixed effects	Y	Y	Y	Y	Y	Y
Adjusted R^2	0.40	0.39	0.39	0.38	0.37	0.37
First stage F-statistic	912.4	983.3	1,060.5	998.0	1,022.1	976.9
# Observations	6,988	8,087	9,059	9,829	10,360	10,769
# Applications	2,748	2,985	3,193	3,328	3,425	3,491
# Children	3,663	4,167	4,639	4,991	5,239	5,431
# Households	2,643	2,879	3,086	3,221	3,318	3,384

Notes: Table 3 presents first stage estimates using the instrument Z_{ijb}, which is an indicator for whether the household had received an offer by that age bin. The sample is the matched sample of children linked to SDE records with any DOL household link. Specifications include race and sex controls and fixed effects for age at application, age, outcome year, years since lottery, and lottery × family size × disability status. Standard errors are clustered at the household level. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

5 Short-run outcomes

5.1 Effects on housing outcomes

Figure 4: Offer effects on monthly housing subsidies



Notes: Figure 4 reports reduced-form effects of a housing assistance offer on rental subsidies, inclusive of program exits. The y-axis is the reduced-form effect, measured in monthly 2023 dollars. The x-axis is calendar time (in years) relative to the offer year. Whiskers report household-clustered 95% confidence intervals.

Housing subsidies and rents I begin by adapting the research design in Section 4 to examine how offers of housing assistance affect the receipt of these in-kind transfers. I first estimate effects on the monthly rental subsidy, then examine how the transfer changes total housing consumption and households' out-of-pocket spending on rent.

For each lottery j and offer cohort g, I compare households offered assistance in year g with applicants from the same lottery who are either not yet or never offered. I stack these comparisons and estimate

$$Y_{ijgt}^{k} = \alpha_{ijg}^{k} + \lambda_{g,j,s(i),t}^{k} + \sum_{r \neq -1} \beta_{r}^{k} D_{ij}(g) \mathbf{1}\{t - g = r\} + \varepsilon_{ijgt}^{k}, \tag{6}$$

where i indexes households, j indexes lottery openings, g indexes offer cohorts, and t

indexes calendar years. I let the k superscript denote different outcomes. The indicator $D_{ij}(g)$ equals one for households offered assistance through lottery j in year g and zero otherwise. The fixed effects α^k_{ijg} absorb permanent household differences across offer cohorts, while the $\lambda^k_{j,s(i),g,t}$ fixed effects absorb household characteristics that may affect the timing of offers.

For the subsidy outcome, I assign zero in years when a household is unassisted and use either the observed or calibrated monthly subsidy while assisted. For HCV participants, I observe the subsidy directly through the PHA's payments to landlords. For unit-based assistance, I construct a rental subsidy index equal to the occupied unit's bedroom-specific fair market rent (FMR) less the household's recorded tenant payment (TTP).[18] Figure 4 reports the resulting offer effects. Pooling across programs, offers increase monthly subsidies by \$420 in the offer year, \$497 one year later, and \$432 five years later, in 2023 dollars.

While these subsidy effects capture the direct value of the in-kind transfers, they fail to identify if the transfers lead to changes in housing consumption and out-of-pocket payments. I examine these margins next using the same offer event study specification. The additional challenge for evaluating effects on these outcomes is that the administrative records contain no rent information at the time of application and no rent information for unassisted households. To address this, I construct an external benchmark for the rent households would pay without assistance using the publicly-available microdata from the five-year American Community Survey (ACS) samples ending between 2009 and 2018. I first calculate average gross rent among New Haven County renter households with children, matching applicants on household size at the time of application. I next restrict the ACS sample to households with annual income between \$10,000 and \$30,000 in 2023 dollars to approximate applicants' resources (see Table 1). For each offer cohort g, I use the ACS window ending in $g-1$, which limits the rent analysis to offer cohorts from 2010 through 2019.

Formally, I let B_{ig} denote this imputed pre-offer rent benchmark for household i in offer cohort g. For total housing consumption and out-of-pocket payments, I construct

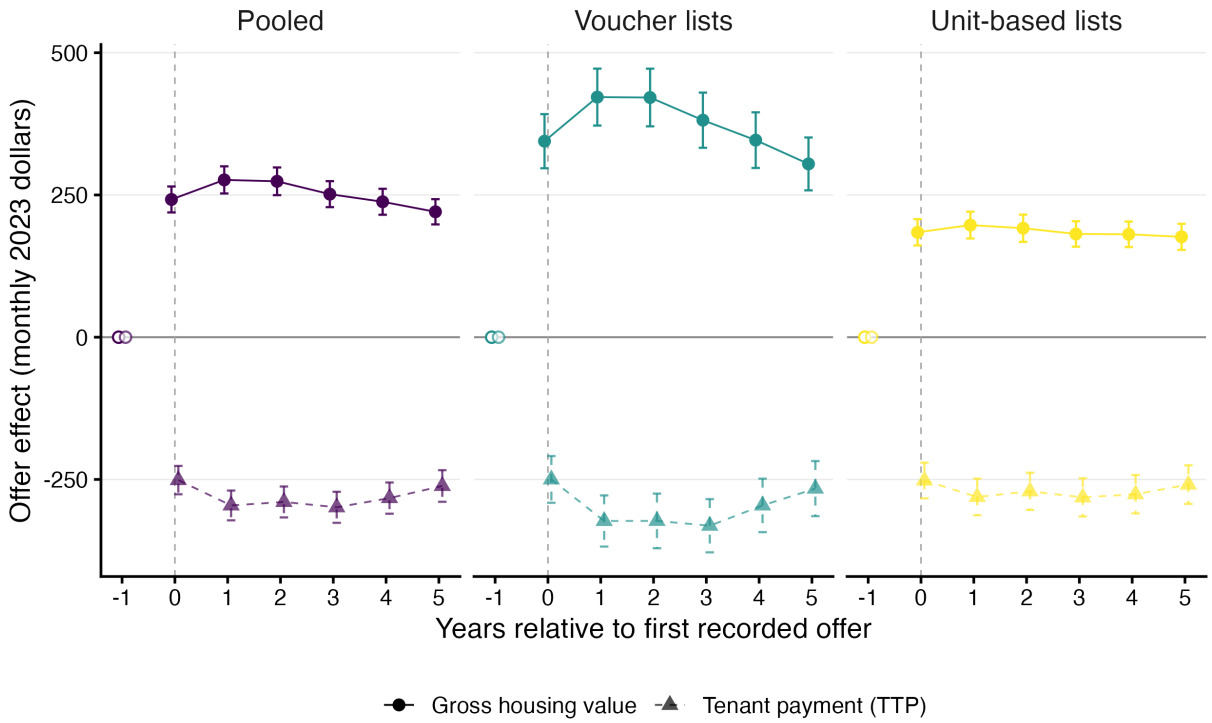
$$Y^k_{ijgt} = \begin{cases} R^k_{it}, & \text{if assisted,} \\ B_{ig}, & \text{if unassisted,} \end{cases}$$

[18] Because FMR approximates the 40$^{\text{th}}$ percentile of rents in the metropolitan area, this benchmark may overstate the rental value of lower-quality developments and understate that of higher-quality developments.

where R^k is either recorded voucher gross rent or the occupied unit's bedroom-specific FMR (for housing consumption) or recorded TTP (for out-of-pocket payments).

I use these constructed outcomes in Equation (6). Holding the unassisted outside option fixed in real dollars assigns zero changes to unassisted households, including households that exit assistance. In other words, the β_r^k estimates are relative to the ACS benchmark. Interpreting these estimates as changes in actual housing consumption or rent payments requires this benchmark to accurately capture the untreated path of rental market outcomes.

Figure 5: Offer effects on housing consumption and tenant payments



Notes: Each panel in Figure 5 reports reduced-form offer effects on total gross rent and tenant payments. Solid lines and circles show gross rent. Lighter dashed lines and triangles show recorded tenant payment (TTP). The y-axis is the monthly change in 2023 dollars relative to an imputed pre-offer rent constructed from ACS data. For assisted households, total housing consumption is either voucher gross rent or the bedroom-specific FMR. Tenant payment is recorded TTP. Whiskers report household-clustered 95% confidence intervals.

Homelessness and housing insecurity The preceding results describe how housing assistance changes housing consumption and out-of-pocket rent payments. Assistance may also protect households against the loss of stable housing. To examine this margin, I use school records collected under the McKinney-Vento Act, which identify children reported

Figure 6: Effects on school-reported housing instability



Notes: Figure 6 plots the dynamic reduced-form effects of a housing assistance offer on school-reported homelessness or housing instability, constructed using McKinney-Vento records. The figure reports estimates for children linked to education records in a year strictly before the lottery (purple) or just before the offer year (yellow). Separate estimates are shown for voucher and unit lists. Both samples require observed homelessness, attendance, and chronic absenteeism measures in the year before the first recorded offer. The year before the offer is the reference period. The y-axis is in percentage points. Standard errors are clustered at the household level. Whiskers represent 95% confidence intervals.

as experiencing homelessness or housing insecurity in a given school year.

I begin by estimating an offer event study comparing children offered assistance in year g with not-yet and never offered children from the same lottery:

$$H_{ijgt} = \alpha_{ig} + \lambda_{j,s(i),g,t} + \sum_{r \neq -1} \beta_r D_{ij}(g) \mathbf{1}\{t - g = r\} + \varepsilon_{ijgt}, \tag{7}$$

where H_{ijgt} is the school-reported homelessness or housing insecurity indicator, and $D_{ij}(g)$ indicates an offer from lottery j in year g. The individual-by-cohort fixed effects α_{ig} absorb permanent differences across children within each cohort comparison.

Figure 6 plots estimates of β_r, the reduced-form offer effect at each event year relative to the year before the offer. The estimates suggest larger reductions among applicants to voucher wait lists than among applicants to unit-based assistance lists.

To make use of the full sample of applicants, including children never observed in school prior to the offer year, I report estimates of Equations (1)-(2), using two-year age bins, in Table 4. These static estimates suggest that the reduction in homelessness is concentrated among younger children. At ages six to seven, assistance reduces school-reported

homelessness by 6.88 percentage points ($p < 0.001$), relative to a complier control mean of 8.99 percent. The estimate is 3.42 percentage points at ages eight to nine ($p = 0.084$) and becomes smaller and less precise at older ages. In Appendix Table A.3, I report estimates from pooling annual school observations from ages six through eleven and twelve through seventeen. For younger children, I find that receiving housing assistance by a given age reduces school-reported homelessness or housing instability by 4.34 percentage points ($p = 0.003$), relative to a complier control mean of 6.70 percent.

Table 4: Effects of housing assistance on school-reported homelessness by child age

Child ages	OLS	RF	2SLS	CM	CCM	Observations	# Children	First stage F
6–7	-3.73***	-3.07***	-6.88***	5.16	8.99	4,702	2,620	660.5
	(0.89)	(0.81)	(1.81)	–	(1.72)	–	–	–
8–9	-3.15***	-1.45*	-3.42*	4.05	5.77	5,527	3,016	740.7
	(0.95)	(0.84)	(1.98)	–	(1.86)	–	–	–
10–11	-2.28**	-1.00	-2.40	3.43	5.05	6,120	3,303	741.0
	(0.90)	(0.79)	(1.89)	–	(1.75)	–	–	–
12–13	-2.18***	-1.22*	-3.13*	3.63	4.74	6,205	3,374	580.6
	(0.78)	(0.71)	(1.84)	–	(1.72)	–	–	–
14–15	-1.48*	-0.79	-2.13	3.46	5.04	5,938	3,209	512.5
	(0.87)	(0.69)	(1.85)	–	(1.68)	–	–	–
16–17	-1.42	-0.57	-1.59	3.55	4.57	4,029	2,707	333.4
	(1.00)	(0.97)	(2.71)	–	(2.47)	–	–	–

Notes: Each row reports an estimate of β_b after stacking annual child-year observations within the indicated two-year age bin. Coefficient estimates and means are in percentage points. OLS reports the coefficient on receiving assistance by the corresponding age. RF denotes the reduced-form effect of having been offered assistance by that year. The binary 2SLS specification instruments assistance receipt with the offer indicator. CM is the mean outcome among observations not yet offered assistance. CCM is the estimated outcome mean among untreated compliers. Standard errors are shown in parentheses below coefficient estimates and CCMs. Specifications include fixed effects for age-at-application, current age, calendar year, years since application, and lottery-by-family-size-by-disability allocation cells, in addition to controls for child race and gender. Standard errors are clustered at the household level. $^{*}p < 0.10$, $^{**}p < 0.05$, $^{***}p < 0.01$, respectively.

The estimated reductions are concentrated in the forms of school-reported instability most common in the sample. Pooling observations from ages six through eleven, assistance reduces non-shelter housing instability by 2.77 percentage points ($p = 0.027$) and doubled-up homelessness by 2.86 percentage points ($p = 0.012$). The estimated effect on shelter homelessness is also negative, at 1.56 percentage points ($p = 0.037$). I find no detectable effect on living in a hotel or motel (0.08 percentage points, $p = 0.865$), a rare outcome in the SDE records. In other words, the aggregate decline in school-reported housing instability largely reflects reductions in less visible forms of housing instability – particularly doubled-up homelessness – that are recorded under the McKinney-Vento Act but

generally do not appear in other administrative datasets used to study homelessness (i.e., HMIS records). This pattern is consistent with studies finding large effects of housing-related interventions on broad measures of housing instability, not only homelessness shelter utilization (Collinson and Gibbs 2025; Liao et al. 2025).

5.2 Effects on early childhood development

Table 5: Effects on kindergarten readiness, by gender

Outcome	Pooled	Boys	Girls
KEI index	0.066	0.206	0.028
	(0.094)	(0.130)	(0.136)
Developmental skills average	0.110	0.288**	0.033
	(0.100)	(0.141)	(0.142)
Academic skills average	0.022	0.124	0.024
	(0.100)	(0.138)	(0.146)
Developmental categories			
Creativity/aesthetic	0.117	0.306*	0.003
	(0.111)	(0.160)	(0.157)
Physical/motor development	0.047	0.231	−0.039
	(0.110)	(0.159)	(0.156)
Personal/social development	0.166	0.328**	0.136
	(0.111)	(0.158)	(0.160)
Academic categories			
Numeracy	0.054	0.083	0.095
	(0.112)	(0.156)	(0.163)
Literacy	0.044	0.173	0.007
	(0.108)	(0.151)	(0.155)
Language	−0.032	0.115	−0.032
	(0.110)	(0.155)	(0.158)
Children	2,117	1,078	1,038
Households	1,734	972	952

Notes: Table 5 reports binary 2SLS estimates for the effects of housing assistance on early childhood development. The endogenous treatment is receipt of housing assistance by the kindergarten assessment year, instrumented with an offer by that year. The KEI index averages standardized language, literacy, numeracy, creativity/aesthetic, physical/motor, and personal/social scores. Each score is standardized within a given school year. KEI scores are available beginning in the 2008-2009 school year. Standard errors are clustered at the household level. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$, respectively.

Housing assistance may also affect early childhood development. To examine whether

these effects translate into greater school readiness, I use Connecticut's Kindergarten Entrance Inventory (KEI), a statewide observational assessment administered each fall since 2010 to measure incoming kindergartners' readiness. Teachers assess children in six categories: language, literacy, numeracy, creativity, physical and motor development, and personal and social development. Scores range from 1 (needs substantial instructional support) to 3 (consistently demonstrates the skill). These scores are used to evaluate the effects of universal pre-kindergarten programs in Humphries et al. 2024.

I estimate effects on each child's raw score KEI_{it}^{k} and standardized score z_{it}^{k} in category k, as well as an overall index combining the six standardized scores. To distinguish different dimensions of school readiness, I also report separate indices of academic skills (language, literacy, and numeracy) and developmental skills (creativity, physical and motor development, and personal and social development).

Averaged across all children, the estimated effect on the overall KEI index is small and imprecise, at 0.066 standard deviations ($p = 0.482$). This pooled estimate, however, hides a more favorable pattern for boys, for whom receiving housing assistance before kindergarten increases the average of the developmental skill measures – creativity, physical and motor development, and personal and social development – by 0.29 standard deviations ($p = 0.042$). A latent factor measurement model that separates these scores into noisy measures of development and academic skill produces a similar estimate: assistance increases boys' latent developmental skill by 0.35 standard deviations ($p = 0.040$). In contrast, the estimated effect on boys' latent academic skill is small and statistically insignificant.[19] The corresponding latent developmental skill estimate for girls is 0.01 standard deviations ($p = 0.888$), though the estimated difference between boys and girls is imprecise ($p = 0.148$). I report the full set of estimates for the standardized category outcomes in Table 5 and the measurement model estimates for the developmental factor in Appendix Table A.4.

5.3 Effects on elementary and middle school test scores

Despite the significant increases in measures of early childhood development among boys, I find no consistent evidence that housing assistance improves measured academic achievement in elementary and middle school. In the binary specification, estimates across both math and ELA that pool boys and girls are positive in grades 3 through 6, ranging from 0.045 to 0.178 standard deviations, but none is statistically significant. Estimates in grades 7 and 8 range from -0.208 to 0.033 standard deviations and are also

[19] Additional details on the latent factor model are available in Appendix D.1.

statistically insignificant. The confidence intervals remain wide enough to accommodate sizable increases or decreases in test scores across grades and subjects. Among boys, assistance increases third-grade ELA scores by an estimated 0.493 standard deviations ($p = 0.032$) and fifth-grade ELA scores by 0.346 standard deviations ($p = 0.060$). I report the full set of estimates in Appendix Table A.5.

A challenge for interpretation in this setting is that the ELA point estimates are very large in comparison with estimated effects of increased school spending. For example, Jackson and Mackevicius 2024 find that increasing annual school spending by $1,000 per pupil for four years raises test scores by only 0.032 standard deviations. Some point estimates are also large relative to difference-in-differences estimates of the effects of housing assistance in New York City. Schwartz et al. 2020 find that housing voucher receipt raises ELA and math scores by around 0.05 standard deviations, while Han and Schwartz 2024 estimate gains of 0.09 standard deviations in ELA and 0.11 standard deviations in math in the years following entry into public housing. Notably, the confidence intervals on my effect estimates are always wide enough to include gains of the magnitudes documented in these studies.

The absence of consistently positive achievement effects in elementary and middle school places these New Haven estimates closer to the findings from the Chicago voucher lottery studied by Jacob et al. 2015 than the significant, persistent positive elementary school effects documented in the Miami-Dade lottery studied by Collinson and Gibbs 2025. At the same time, the positive effects on early childhood development among boys suggest that standardized achievement tests administered in grades 3 through 8 may not fully capture the developmental impacts of increased housing stability. This interpretation is consistent with evidence that changes in non-cognitive skills explain a sizable share of the long-run benefits of the Perry Preschool intervention (Heckman et al. 2013) and with evidence that teachers affect educational attainment through improvements in attendance and behavior (Jackson 2018).

5.4 Effects on other outcomes

The preceding results show that housing assistance reduces housing instability and improves early developmental skills among boys without producing significant improvements in test scores in elementary and middle school. This pattern suggests that assistance may affect children's home environments and their attachment to school without substantially changing other inputs to the process of human capital accumulation. I next

examine three potential channels and intermediate outcomes: school attachment, household resources, and school quality.

School attachment I find evidence that housing assistance improves children's attachment to school, particularly during elementary school. Pooling child-year observations from ages 6–11, receiving assistance by a given age increases attendance by 1.14 percentage points ($p = 0.027$) and reduces chronic absenteeism by 6.46 percentage points ($p = 0.028$). Receipt of assistance also reduces unplanned school switches by 2.70 percentage points ($p = 0.014$). The effect on unplanned school district switches is also negative but smaller and statistically insignificant, at -1.37 percentage points ($p = 0.153$). Together, these results are consistent with a school continuity channel for housing assistance at early ages. This may have important implications for future outcomes: recent studies find that increased school attendance has positive longer-run effects (Liu et al. 2021; Goldman and Gracie 2024).

Table 6: Effects of housing assistance on school continuity at ages 6–11

Outcome	OLS	RF	2SLS	CM	CCM	Observations	# Children	First stage F
Attendance rate	-0.13	0.49**	1.14**	92.36	90.67	18,043	4,110	946.4
	(0.26)	(0.22)	(0.52)	–	(0.46)	–	–	–
Chronic absenteeism	-0.31	-2.80**	-6.46**	25.24	34.06	18,043	4,110	946.4
	(1.45)	(1.26)	(2.93)	–	(2.69)	–	–	–
Switched schools within year	-1.76***	-1.17**	-2.70**	5.05	6.04	18,043	4,110	946.4
	(0.52)	(0.48)	(1.10)	–	(1.04)	–	–	–
Switched districts within year	-1.84***	-0.60	-1.37	3.40	3.63	18,043	4,110	946.4
	(0.44)	(0.42)	(0.96)	–	(0.91)	–	–	–

Notes: Table 6 reports OLS, reduced form, and 2SLS estimates of the effects of housing assistance on school continuity using Equations (1)-(2). The sample pools child-year observations from ages 6 through 11. Coefficient estimates and means are in percentage points. OLS reports the coefficient on receiving assistance by the corresponding child age. RF denotes the reduced form effect of having received an offer by that year. The 2SLS specification instruments assistance receipt with the offer indicator. CM denotes the mean among observations not yet offered assistance, and CCM denotes the estimated complier control mean. Standard errors are shown in parentheses below coefficient estimates and CCMs and are clustered at the household level. Specifications include controls for child characteristics and fixed effects for age at application, age, calendar year, years since application, and lottery-by-family-size-by-disability-status. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

School quality Housing assistance may also affect children by changing the quality of the schools they attend. For elementary and middle schools, I examine how housing vouchers and public housing impact school quality, measured using school-by-year value added in ELA and math. To construct school-level value added, I estimate a lagged achievement model following Hanushek 1971. I also examine differences in school proficiency rates and student characteristics. Appendix C.1.4 describes the value added model

and presents the distribution of the school-by-year value added estimates.[20]

I find no evidence of effects on school value added. But older students that receive assistance do appear to attend schools with different peers. Receiving assistance increases the share of students eligible for free or reduced-price lunch at the child's school by 5.92 percentage points at ages 14–15 ($p = 0.008$) and by 7.21 percentage points at ages 16–17 ($p = 0.012$). At least in part, this difference reflects the 20 to 30 percentage point higher FRPL shares at public high schools in New Haven compared to surrounding communities. Consistent with this interpretation, I find no corresponding effects at younger ages. That children do not move to schools with observably different peers or test school value added is consistent with the limited effects on school quality in Jacob et al. 2015. This does distinguish my setting from studies of mobility interventions like the Moving To Opportunity experiment that shift the types of schools and neighborhoods experienced by families. I report the full set of school quality estimates in Appendix Table A.7.

Household resources Finally, housing assistance may affect children by changing the financial resources and time available within the household. Lower rent frees resources for other consumption, while the income-based subsidy formula – rent is typically 30% of income – likely alters adult incentives to supply labor through both income and substitution effects. To illustrate these two margins, let m_j denote consumption net of housing costs and h_j denote hours worked without ($j = 0$) or with ($j = 1$) housing assistance. Holding the hourly wage w and non-wage income n fixed, and abstracting from saving and other tax and benefit adjustments, the household's budget constraints are

$$m_0 = wh_0 + n - r_0, \qquad m_1 = wh_1 + n - 0.3(wh_1 + n),$$

where r_0 is unassisted rent and assisted rent is set to 30% of income. The change in resources available for non-housing consumption is

$$m_1 - m_0 = \underbrace{r_0 - 0.3(wh_0 + n)}_{\text{rent relief at baseline income}} + \underbrace{0.7w(h_1 - h_0)}_{\text{earnings response net of rent adjustment}}.$$

Under the assumption that households spend more than 30% of income on rent at baseline – which is implied by the ACS estimates – housing assistance must increase net-of-housing resources if labor supply increases in response to the transfer. In turn, these resources will only fall if the reduction in labor supply, scaled by the implicit subsidy tax rate, exceeds the magnitude of rent relief.

[20] Because the administrative records do not identify teachers, the value added measures combine school effects with the effects of teacher sorting across schools or student sorting across classrooms.

Existing studies of housing assistance lotteries find average reductions in adult employment or earnings (Jacob and Ludwig 2012; van Dijk 2019). I examine these responses using Connecticut UI earnings records for adult household members identified at the selected lottery application. Appendix Table A.6 reports 2SLS estimates by child age. The estimated effects of receiving assistance on employment and earnings are positive in every age bin. Employment effects are statistically insignificant throughout, and earnings effects are statistically insignificant through ages 14–15. The estimated earnings increase reaches $12,432 at ages 16–17 (SE = $6,236). These estimates provide some suggestive evidence of earnings gains at younger ages, but are consistent with reductions or increases in employment or earnings given the wide confidence intervals.

6 Long-run outcomes

The preceding section reports how housing assistance affects children's housing stability, early childhood development, and other short-run outcomes. Next, I examine how housing assistance affects educational attainment and criminal justice involvement as children enter adulthood.

6.1 Effects on educational attainment

In contrast to the test score results discussed above, housing assistance has large positive effects on educational attainment. I report the full set of estimates in Table 7 using two complementary specifications. The binary specification instruments receiving any assistance by age 18 with an indicator for receiving an offer by age 18. The exposure specification instead instruments years of assistance through age 18 with years since an offer. Under the constant marginal effects assumption discussed in Section 4, the latter coefficient can be interpreted as the effect of one additional year of childhood assistance (Chyn 2018; Chaudhry and Eng 2024; Beauregard 2025). Because this interpretation requires stronger assumptions than those needed for a LATE interpretation of the binary offer design, I report both estimates and treat the exposure results as complementary.

Effects on high school graduation are positive but imprecisely estimated. An additional year of assistance increases on-time graduation by 1.42 percentage points ($p = 0.180$) and graduation by age 20 by 1.98 percentage points ($p = 0.140$). Among children old enough to be observed through age 22, the estimate rises to 2.88 percentage points ($p = 0.099$), relative to a control mean of 63.1 percent. The corresponding binary estimates are also

positive, but none is statistically significant at the 5 percent level. These results are suggestive of increases in high-school completion, especially at later ages, which is consistent with estimates of the effects of eviction on child educational attainment (Collinson et al. 2025).

The evidence for attainment effects is stronger along the postsecondary enrollment margin. An additional year of assistance increases college enrollment by age nineteen by 1.94 percentage points ($p = 0.079$) and college enrollment by age twenty-two by 4.19 percentage points ($p = 0.020$). The effect at later ages is concentrated in the four-year sector. For example, an additional year of assistance by age twenty-two raises four-year enrollment by 3.92 percentage points ($p = 0.017$), compared with a smaller and statistically insignificant 1.29 percentage point effect on enrollment in the two-year sector ($p = 0.412$). I find some evidence that this translates into gains in college completion, though this analysis is limited by child age. I show that each year with assisted leads to a 1.99 percentage point increase in graduation in either sector by age twenty-two ($p = 0.054$) and a 2.92 percentage point increase in graduation by age twenty-four ($p = 0.085$).

This pattern – larger effects on attainment than on standardized tests – is similar to findings in other papers that study the effects of in-kind transfers in early childhood on long-run outcomes (Deming 2009; Gray-Lobe et al. 2023). For example, Gray-Lobe et al. 2023 use lotteries to show that enrolling in free pre-kindergarten programs in Boston increases high school graduation by 6.0 percentage points and college enrollment by 5.4 percentage points despite no detectable effect on state achievement tests. Within the literature on housing assistance, my results differ from the null effects in Jacob et al. 2015, though the research designs and treatment definitions in the two studies are not directly comparable. Jacob et al. 2015 estimate the effect of ever using a voucher, whereas my preferred dose-response specification scales the effect by the number of years that children spend with assistance. My results are broadly consistent with evidence that longer childhood exposure to higher-quality housing environments can improve adult outcomes (Chyn 2018; Chaudhry and Eng 2024; Beauregard 2025).

6.1.1 Heterogeneity in attainment effects

Gender Table 8 reports the dose-response estimates separately by child gender. The point estimates for any college enrollment and college completion are especially large for girls relative to boys, providing suggestive evidence that gains in postsecondary attainment may be concentrated among girls. While these results are similar to the pattern fo short run gains documented in the Moving To Opportunity (MTO) experiment (Kling et al.

Table 7: Housing assistance effects on high school and college outcomes

				Exposure		Binary	
Outcome	CM	Children	Households	2SLS	First stage F	2SLS	First stage F
Panel A: High school graduation							
HS graduation by age 18	53.6	3,057	2,096	1.42	222.8	16.34*	268.2
	–	–	–	(1.06)	–	(9.00)	–
HS graduation by age 20	62.4	2,315	1,633	1.98	168.1	14.91	196.5
	–	–	–	(1.34)	–	(10.23)	–
HS graduation by age 22	63.1	1,727	1,252	2.88*	114.5	10.74	146.2
	–	–	–	(1.75)	–	(12.11)	–
Panel B: College enrollment and completion							
College enrollment by age 19	34.0	2,685	1,853	1.94*	196.7	15.38*	227.5
	–	–	–	(1.10)	–	(9.18)	–
College enrollment by age 20	35.1	2,315	1,633	3.09**	168.1	19.07*	196.5
	–	–	–	(1.32)	–	(10.24)	–
College enrollment by age 22	38.2	1,727	1,252	4.19**	114.5	14.53	146.2
	–	–	–	(1.80)	–	(12.42)	–
Four-year college enrollment by age 22	22.1	1,727	1,252	3.92**	114.5	14.73	146.2
	–	–	–	(1.65)	–	(10.87)	–
Two-year college enrollment by age 22	23.2	1,727	1,252	1.29	114.5	3.40	146.2
	–	–	–	(1.58)	–	(10.65)	–
College degree by age 22	6.6	1,727	1,252	1.99*	114.5	10.42	146.2
	–	–	–	(1.03)	–	(6.56)	–
College degree by age 24	9.1	1,173	897	2.92*	75.9	19.44**	102.4
	–	–	–	(1.70)	–	(9.86)	–

Notes: Table 7 reports 2SLS estimates of the effect of housing assistance on high school and college outcomes using offer timing as an instrument. All coefficient estimates and means are in percentage points. Exposure columns report the effect of one additional year of assistance through age 18, instrumenting assisted years with offered years. Binary columns report the effect of any assistance through age 18, instrumenting assistance receipt with an offer. CM is the outcome mean among children whose households were not offered assistance by age 18. Standard errors, shown in parentheses, are clustered at the household level. Specifications include controls for child and household characteristics and fixed effects for age at application, outcome year, and lottery-by-family-size-by-disability-status. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

2007), they differ from the insignificant gender differences in long-run outcomes found from the MTO intervention (Chetty et al. 2016).

Table 8: Attainment effects by child gender

Outcome	Male	Female	Female − male	p-value
High school graduation by age 22	3.47	2.31	-1.17	0.758
	(2.32)	(2.94)	(3.78)	–
Any college enrollment by age 22	1.42	8.65**	7.23*	0.090
	(2.21)	(3.57)	(4.26)	–
Four-year college enrollment by age 22	2.76	5.73*	2.97	0.413
	(1.78)	(3.20)	(3.63)	–
College degree by age 24	0.61	6.13	5.51	0.184
	(1.49)	(3.84)	(4.15)	–

Notes: Table 8 reports 2SLS estimates, in percentage points, of the effect of one additional year of housing assistance through age 18. Assisted years are instrumented with offered years. The final two columns report the female-minus-male difference and its two-sided p-value from a pooled interaction specification. Standard errors, shown in parentheses, are clustered at the household level. The age twenty-two samples contain 870 male and 857 female children. The age twenty-four samples contain 589 male and 584 female children. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Program In Table 9, I report the effects of housing assistance separately for tenant-based vouchers and unit-based assistance. The estimates come from the following model:

$$Y_{ij} = \beta_V A_{ij} V_j + \beta_U A_{ij} U_j + X_{ij}' \Gamma_{p(j)} + \lambda_{j,s(i)} + \varepsilon_{ij}, \tag{8}$$

where V_j and U_j are indicators for applying through either a voucher or unit-based lottery, respectively, and A_{ij} denotes years of housing assistance between lottery j and age eighteen. I jointly instrument $A_{ij}V_j$ and $A_{ij}U_j$ with $E_{ij}V_j$ and $E_{ij}U_j$, where E_{ij} denotes offer exposure from lottery j through age eighteen. The relationship between household covariates X_{ij} and the outcome is allowed to vary by lottery type $p(j)$, and all specifications include lottery-by-household-size-by-disability-status fixed effects. Thus, β_V and β_U capture the effects of additional assistance exposure induced by offers within voucher and unit-based lotteries, respectively. The voucher estimates in Table 9 are positive for all four outcomes, while the unit-based estimates are close to zero or negative. For high school graduation and any college enrollment, the differences are statistically significant. Taken at face value, these exposure measures imply enormous returns to housing vouchers, with little effect from unit-based assistance.

Table 9: Attainment effects by housing assistance program

Outcome	Voucher	Unit-based	Unit − voucher	p-value
High school graduation by age 22	4.34**	-0.44	-4.79**	0.017
	(1.79)	(1.91)	(2.00)	–
Any college enrollment by age 22	5.40***	1.36	-4.03*	0.050
	(1.82)	(1.96)	(2.06)	–
Four-year college enrollment by age 22	4.16**	1.71	-2.45	0.189
	(1.64)	(1.78)	(1.87)	–
College degree by age 24	2.81*	0.63	-2.19	0.196
	(1.57)	(1.64)	(1.69)	–

Notes: Table 9 reports 2SLS estimates, in percentage points, of the effect of one additional year of housing assistance through age 18, separately for applicants to voucher and unit-based lotteries. The endogenous variables are total assisted years interacted with indicators for voucher and unit-based lotteries. The instruments are the corresponding interactions of offered years with lottery type. Covariate coefficients are allowed to vary by lottery type. The final two columns report the unit-lottery-minus-voucher-lottery difference and its two-sided p-value. Standard errors, shown in parentheses, are clustered at the household level. The age twenty-two sample contains 1,727 children in 1,252 households. The age twenty-four sample contains 1,173 children in 897 households. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

6.2 Effects on crime

I next examine whether childhood housing assistance affects criminal justice involvement in early adulthood. I separately link children in the wait list sample to Connecticut criminal court records and construct an indicator for having any post-lottery docket with an explicit guilty finding on a felony charge by ages 19 through 22. Outcomes are dated using the age at arrest and capture criminal convictions.

I estimate these effects using the full sample of applicants from New Haven County, regardless of SDE or DOL match.[21] Receiving assistance by age 18 reduces the probability of a felony conviction by age 20 by 3.88 percentage points ($p = 0.190$) and by age 22 by 7.45 percentage points ($p = 0.097$). The final estimate is large relative to the estimated 8.86 percent conviction rate among untreated compliers. The exposure estimates follow a similar but imprecise pattern, with the age twenty-two estimate suggesting an annual reduction in felony conviction risk of 1.07 percentage points ($p = 0.175$).

The conclusion is strengthened when I expand the geographic sample from New Haven County applicants to all applicants from Connecticut. In this larger sample, receiving assistance by age 18 reduces felony convictions by age 22 by 9.76 percentage points ($p = 0.045$), while an additional assisted year reduces convictions by 1.34 percentage points

[21] My DUA with DataLinkCT prevents me from combining the criminal justice sample with the linked education and employment records, limiting my ability to assess differential out-migration within the criminal justice sample.

($p = 0.097$). Earlier estimates are also negative but less precise. Overall, the results provide evidence that childhood housing assistance reduces serious criminal justice involvement in early adulthood, complementing prior evidence that housing and neighborhood interventions can affect later criminal involvement (Kling et al. 2007; Chyn 2018). I report the full set of estimates in Table 10.

Table 10: Housing assistance effects on felony convictions

					Exposure		Binary	
						First-stage		First-stage
Outcome	CM	CCM	Children	Households	2SLS	*F*	2SLS	*F*
Panel A: New Haven County sample								
Felony conviction by age 19	1.63	1.95	3,474	2,189	0.16	192.1	−0.54	243.5
					(0.28)		(1.93)	
Felony conviction by age 20	3.21	4.83	3,083	1,960	−0.46	177.7	−3.88	216.2
					(0.41)		(2.96)	
Felony conviction by age 21	3.91	5.52	2,683	1,727	−0.61	143.9	−4.40	182.8
					(0.58)		(3.72)	
Felony conviction by age 22	5.08	8.86	2,309	1,513	−1.07	122.8	−7.45*	156.4
					(0.79)		(4.49)	
Panel B: Connecticut sample								
Felony conviction by age 19	1.65	2.66	4,217	2,619	0.15	193.6	−1.37	238.5
					(0.28)		(2.03)	
Felony conviction by age 20	3.21	6.58	3,743	2,359	−0.57	178.9	−5.63*	214.0
					(0.43)		(3.22)	
Felony conviction by age 21	4.02	8.48	3,255	2,085	−0.84	144.6	−7.43*	182.6
					(0.59)		(4.03)	
Felony conviction by age 22	4.99	11.06	2,807	1,823	−1.34*	120.9	−9.76**	154.6
					(0.81)		(4.86)	

Notes: Entries are 2SLS estimates in percentage points. The outcome is any felony conviction with an explicit guilty finding by the given age. Exposure estimates instrument assisted years through age 18 with offered years through age 18. Binary estimates instrument any assistance through age 18 with any offer through age 18. CM is the conviction rate among children whose households were not offered assistance by age 18. CCM is the untreated conviction rate among compliers in the binary receipt specification. Panel A uses all applicants from New Haven County. Panel B uses all applicants from Connecticut. Specifications include demographic controls and fixed effects for age at application, outcome year, and lottery-by-family-size-by-disability-status. Parentheses contain household-clustered standard errors. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$ represent significance at the 10%, 5%, and 1% significance levels, respectively.

6.3 Robustness

I now check whether the main results are sensitive to the treatment definition. I first show that the main conclusions are qualitatively similar when assistance and offer exposure are instead calculated through age twenty-two (rather than age eighteen) using household assistance records. The estimated effects are 2.33 percentage points for high school graduation by age twenty-two ($p = 0.102$), 2.96 percentage points for any college enrollment by age twenty-two ($p = 0.041$), and 2.67 percentage points for four-year enrollment by age twenty-two ($p = 0.039$). Measuring treatment through age twenty-four yields a 2.38 percentage point effect on college graduation by age twenty-four ($p = 0.048$).

7 Cost-benefit analysis

My results show that housing assistance programs increase children's educational attainment, reduce their criminal justice involvement, and reduce housing instability. However, housing assistance is costly to provide. In this section, I evaluate the costs and benefits of providing one year of housing assistance to a household with children using the marginal value of public funds (MVPF) framework from Hendren and Sprung-Keyser 2020. The MVPF equals recipients' willingness-to-pay divided by net program costs:

$$\text{MVPF} = \frac{\Delta W}{\Delta E - \Delta C},$$

where ΔW represents recipients' willingness-to-pay for one year of housing assistance, ΔE measures the upfront government cost of providing that year of assistance, and ΔC measures reductions in government expenditures or increases in government revenue generated by that one-year assistance spell. I compute each term using estimates from my quasi-experimental design and external estimates from Reeder 1985, McCollister et al. 2010, Zimmerman 2014, and Hendren and Sprung-Keyser 2020. All dollar amounts in this section are in 2023 dollars unless otherwise noted.

7.1 Policy unit and direct costs

The policy unit is one household with children assisted for one year. The PHA's average annual subsidy in 2023 was \$6,496.36 per assisted person. I estimate that complier households contain 3.459 people and 1.607 children. The direct government cost is therefore

$$G^{hh} = 3.459 \times \$6{,}496.36 = \$22{,}472.32.$$

7.2 Fiscal effects and benefits

College enrollment One additional assisted year through age 18 increases four-year college enrollment by age 22 by 3.92 percentage points ($p = 0.017$). To value this effect, I follow the Hendren and Sprung-Keyser 2020 extrapolation from Zimmerman 2014. Their calibration assigns an induced FIU admission $142,757 in lifetime pre-tax earnings, $2,617 in public college costs, and $2,851 in private tuition in 2012 dollars. The admissions cutoff studied in Zimmerman 2014 increases FIU admission by 23.4 percentage points but state university attendance by 11.9 percentage points. Because my outcome is four-year college enrollment rather than admission, I multiply all three components of college benefits and costs by $0.234/0.119 = 1.966$. After CPI conversion, the resulting values per induced state-university attendee are $372,547 in earnings, $6,829 in public costs, and $7,440 in private tuition.

The dose-response estimate of the effect of housing assistance on four-year college enrollment implies

$$1.607 \times 0.03924 = 0.06308$$

additional four-year enrollees per assisted household-year. This generates $23,500.87 in projected lifetime earnings, $4,700.17 in tax revenue under a 20% effective tax rate, and $430.81 in public college costs. Projected after-tax child earnings net of private tuition equal

$$B^{child} = (1 - 0.20) \times \$23{,}500.87 - \$469.34 = \$18{,}331.36.$$

Homelessness Among children ages 6–11, one additional assisted year reduces the annual probability of any school-recorded homelessness by 0.862 percentage points ($p < 0.001$). The effect on shelter homelessness, the component with direct public costs, is a 0.235 percentage point reduction ($p = 0.034$). To price this reduction in sheltered homelessness, I use Connecticut HMIS records. I find that family shelter stays last 148.6 days on average. I then apply an estimated $5,552.50 public cost per individual stay (Stober 2026) and the 3.459 complier-equivalent household members gives

$$H = 3.459 \times 0.002354 \times \$5{,}552.50 = \$45.21$$

in expected shelter savings per household-year.

Crime In the Connecticut-wide sample, one additional year with housing assistance through age 18 reduces the probability of a felony conviction by age 22 by 1.34 percentage points ($p = 0.097$). To evaluate the welfare implications of this reduction in crime, I separate victim benefits from government savings. First, I construct a victim-cost measure

– following McCollister et al. 2010 – using the national felony conviction mix, excluding criminal justice costs, offender career costs, and homicide risk. I maintain the conservative assumption that each reduced conviction has no future effects on criminal involvement, though I also assume that each conviction reduces custody episodes one-for-one. Given these assumptions, I discount the reductions in victim harm and custody costs by ten years to the year of assistance. This implies that victim hard falls by \$8,193.62 and custody costs fall by \$2,704.10.[22] As a result, the reduction in criminal justice involvement yields

$$\begin{aligned} V &= 0.02160 \times \$8{,}193.62 = \$177.01 && \text{in victim benefits,} \\ J &= 0.02160 \times \$2{,}704.10 = \$58.42 && \text{in government savings.} \end{aligned}$$

I allow victim benefits to enter willingness to pay for the transfers. Custody savings, however, only reduce net government costs of providing housing assistance.

Combining the direct cost and the three downstream fiscal channels gives

$$\begin{aligned} \Delta E - \Delta C &= G^{hh} + P - T - H - J \\ &= \$22{,}472.32 + \$430.81 - \$4{,}700.17 - \$45.21 - \$58.42 \\ &= \$18{,}099.33. \end{aligned}$$

7.3 Willingness to pay and MVPF

I parameterize willingness to pay as

$$\Delta W_\lambda(\theta) = \theta G^{hh} + \lambda B^{child} + V,$$

where θ is the cash-equivalent value of one dollar of housing assistance to the recipient household. I report estimates under $\theta = 1$, the dollar-for-dollar benchmark, and $\theta = 0.83$, the estimated willingness-to-pay in Reeder 1985. The parameter $\lambda \in \{0, 1\}$ determines whether children's future after-tax earnings (net of tuition) enter willingness to pay. If parents fully internalize children's future earnings and can purchase their optimal housing bundle, an envelope theorem argument favors $\lambda = 0$ (Humphries et al. 2024). However, if housing market constraints or incomplete financial markets prevent families from choosing optimally absent housing assistance, $\lambda = 1$ may be appropriate. Because house-

[22] The conviction mix comes from the Bureau of Justice Statistics' Felony Sentences in State Courts, 2006. The custody calculation uses national sentencing shares, a 3% discount rate, and Connecticut's 2024 marginal custody cost estimate, converted to 2023 dollars.

holds in my sample are low-income, I prefer the assumption $\lambda = 1$. In either case, victim benefits enter willingness to pay for the transfer program.

Table 11 reports the resulting estimates. Under the dollar-for-dollar transfer benchmark, the MVPF is 1.25 when child earnings are excluded from willingness to pay and 2.26 when they are included. Under $\theta = 0.83$, the corresponding estimates are 1.04 and 2.05.

Table 11: MVPF of one household-year of housing assistance

Transfer valuation θ	Transfer and victim WTP $\lambda = 0$	Transfer, victim, and child WTP $\lambda = 1$
0.83	1.04	2.05
	[0.87, 1.29]	[1.02, 3.58]
1.00	1.25	2.26
	[1.05, 1.55]	[1.20, 3.84]

Notes: Brackets report 95% parametric bootstrap confidence intervals from one million independent normal draws of the four-year college, Connecticut-wide felony conviction, and pooled ages 6–11 shelter effects.

In Section 6, I showed that housing vouchers have larger effects on postsecondary attainment than public housing. Incorporating these differences, as well as program-specific estimates on crime and sheltered homelessness, into program-specific MVPF estimates, I find larger returns to public investment in voucher programs. I summarize the results in Table 12.

Table 12: Voucher- and unit-specific MVPFs

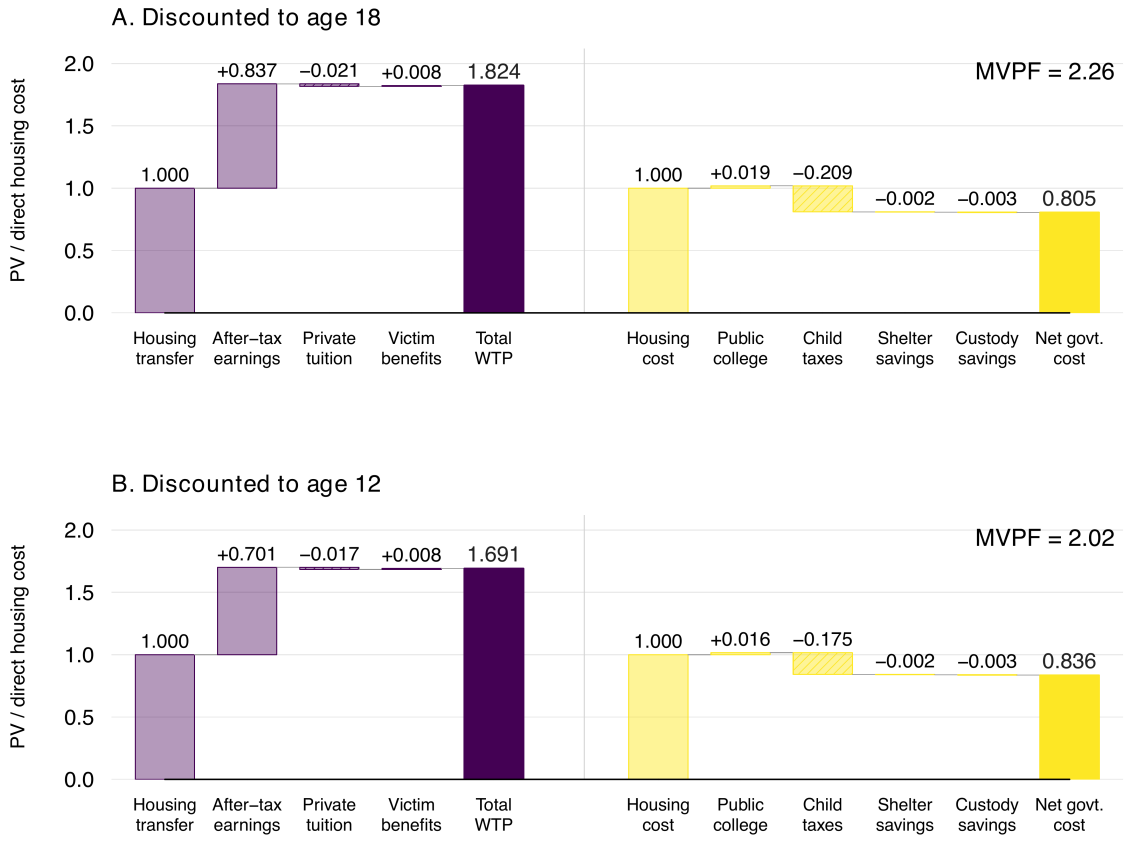
	$\theta = 0.83$		$\theta = 1.00$	
Program	$\lambda = 0$	$\lambda = 1$	$\lambda = 0$	$\lambda = 1$
Voucher	1.05	2.14	1.26	2.35
	[0.88, 1.30]	[1.09, 3.70]	[1.06, 1.57]	[1.26, 3.96]
Unit	0.92	1.31	1.10	1.49
	[0.77, 1.13]	[0.43, 2.58]	[0.93, 1.35]	[0.59, 2.81]

Notes: Brackets report 95% parametric bootstrap confidence intervals from one million independent normal draws of the four-year college, Connecticut-wide felony conviction, and pooled ages 6–11 shelter effects.

Robustness I additionally assess the sensitivity of these MVPF estimates to the particular earnings calibration, the persistence of college returns, and the timing of benefits and costs in Appendix D.2. First, I replace the benchmark FIU earnings estimates with

projections using ACS earnings profiles, allowing the earnings level to vary with the lottery sample's gender composition and parental income. Under $\theta = \lambda = 1$, the pooled MVPF rises from 2.26 to 2.54 using sample-weighted ACS earnings and falls to 1.75 if child earnings are reweighted to reflect the lower parent earnings in this sample. Second, halving the projected earnings gain after age 25 lowers the pooled MVPF to 1.60. Finally, discounting education benefits and costs to child age 12, rather than age 18, lowers the benchmark estimate to 2.02. I report the full decomposition of the MVPF components in Figure 7 for the pooled sample.

Figure 7: Components of willingness to pay and net government cost



Notes: Each panel in Figure 7 decomposes the MVPF into the specific components that contribute to private willingness to pay (purple) and net government costs (yellow). The labels are shown per dollar of transfer, in 2023 dollars. Bolded columns denote the overall value. Panel A assumes that a year of housing assistance is provided at 18. Panel B assumes that a year of housing assistance is provided at 12.

Comparison with existing estimates My results place the returns to housing assistance in the middle of existing MVPF estimates in the public and urban economics literature.

In Appendix Figure A.7, I compare my MVPF estimates with those from other social insurance programs. The returns to housing assistance that I estimate are low relative to recent MVPF estimates from some programs, like UPK, but large relative to historical estimates of the return to public spending on either housing vouchers or public housing. The most conservative specification in Table 11, which uses $\theta = 0.83$ and excludes children's future earnings from willingness-to-pay, yields an MVPF of 1.04. This is above the baseline housing voucher estimate in Hendren and Sprung-Keyser 2020, which uses results from the Chicago voucher lottery studied by Jacob et al. 2015. In that setting, the implied MVPF is 0.65, rising only to 0.68 when child test score effects are used to forecast future earnings. The lower Chicago estimate reflects the combination of modest estimated child effects and negative fiscal externalities from reduced adult earnings and increased transfer receipt. My direct transfer estimates are larger because I find no evidence of adult labor supply reductions, and because reduced shelter use and incarceration and children's projected future tax payments partly offset the upfront subsidy cost.

The MVPF estimates that include child earnings in willingness-to-pay, 2.05 under the Reeder 1985 valuation and 2.26 under dollar-for-dollar valuation of the transfers, are larger than recent estimates for housing assistance programs based on their insurance value alone. Grasley 2026 estimates that federal rental assistance has an MVPF of 1.51 once rental assistance is modeled as insurance against uninsurable income risk. I abstract from this channel.

The estimates are also close to Chaudhry and Eng 2024, who study childhood exposure to New York City public housing. They estimate that each dollar of government spending on public housing generates about $1.40 in benefits for children overall, and about $2.30 for children from the most disadvantaged families. My results should also be interpreted in light of evidence that the effects of housing assistance depend on how these transfers change consumption. Collinson and Gibbs 2025 study housing assistance lotteries in Miami-Dade county and find that vouchers improve early literacy, kindergarten readiness, behavioral development, and even elementary school test scores, while public housing has no detectable average effect on children. They also show that these differences are predicted by changes in housing and neighborhood quality: vouchers move families to higher-quality housing and safer neighborhoods, while public housing moves families, on average, to lower-quality neighborhoods.

Finally, it is worth considering whether housing assistance is more cost-effective than cash. In principle, a cash transfer allows households to optimally choose consumption, and recent evidence from Krause et al. 2025 shows that large temporary cash transfers

increase short-run consumption – including spending on rent – and increase neighborhood mobility. However, the same study finds little persistent improvement in household finances from the transfer. In other words, a low or moderate MVPF from housing vouchers and public housing does not necessarily imply that housing assistance is undesirable if policymakers are paternalistic and cash transfers fail to deliver spillovers to other socially-valuable outcomes (Currie and Gahvari 2008).

8 Conclusion

This paper presents new quasi-experimental evidence on the impacts of housing assistance on children from low-income households. To estimate the effects of these transfers, I construct a new administrative panel dataset spanning nearly two decades that links applications to housing assistance wait lists in New Haven, Connecticut, to records on housing assistance participation, housing instability, children's educational and criminal justice outcomes, and adult labor market outcomes. I first show that housing assistance significantly increases housing consumption while reducing school-reported homelessness and housing instability. I then show that receiving assistance does not improve end-of-grade test scores in elementary and middle school but does raise early childhood development among boys. Next, I show that housing assistance raises rates of high school graduation and college enrollment, with the results driven by children in households that receive tenant-based vouchers. I also find evidence that housing assistance reduces the rate of early-adult felony convictions. Finally, I calculate the benefit per net dollar spent on federal housing assistance. I find that housing assistance yields roughly $2 in benefits for every net dollar spent. Because of larger effects on educational attainment, tenant-based vouchers have larger a larger estimated MVPF.

While the setting I study is similar along observable dimensions to many other medium-sized American cities, my results apply to a particular set of applicants and study period. As highlighted by evidence from voucher design, cash transfer, and mobility experiments (Collinson and Ganong 2018; Bergman et al. 2024; Phillips and Sullivan 2025), a promising area for future work is to compare the relative return to short- or long-term housing assistance that stabilizes families and housing policies that change children's exposure to higher-opportunity neighborhoods. Additional studies of housing assistance programs, using information on wait list applicants from different settings and with more detailed measurements of non-cognitive skills, other long-run outcomes, or information on parental investments in children, are especially welcome.

References

Bailey, Martha J., Hilary Hoynes, Maya Rossin-Slater, and Reed Walker, "Is the Social Safety Net a Long-Term Investment? Large-Scale Evidence from the Food Stamps Program," *The Review of Economic Studies*, 2024.

_ , Shuqiao Sun, and Brenden Timpe, "Prep School for Poor Kids: The Long-Run Impacts of Head Start on Human Capital and Economic Self-Sufficiency," *American Economic Review*, December 2021.

Beauregard, Pierre-Loup, "It's About Time: Social Housing, Parental Labour Supply, and Long-term Child Outcomes," *Working Paper*, 2025.

Bergman, Peter, Raj Chetty, Stefanie DeLuca, Nathaniel Hendren, Lawrence Katz, and Christopher Palmer, "Creating Moves to Opportunity: Experimental Evidence on Barriers to Neighborhood Choice," *American Economic Review*, 2024, *114* (5), 1281–1337.

Bromage, Andy, "Section 8 Waiting List to Reopen," New Haven Register 2005. Published March 26, 2005. Accessed September 11, 2026.

Brown, David W., Amanda E. Kowalski, and Ithai Z. Lurie, "Long-Term Impacts of Childhood Medicaid Expansions on Outcomes in Adulthood," *The Review of Economic Studies*, 2020, *87* (2), 792–821.

Carr, Jillian B and Vijetha Koppa, "Housing vouchers, income shocks and crime: evidence from a lottery," *Journal of Economic Behavior & Organization*, 2020.

Caspi, Aviv and Charlie Rafkin, "Legal Assistance for Evictions," *Working paper*, 2026.

Chaisemartin, Clement De and Luc Behaghel, "Estimating the effect of treatments allocated by randomized waiting lists," *Econometrica*, 2020.

Chaudhry, Raheem and Amanda Eng, "From Marcy to Madison Square?: The Effects of Growing Up in Public Housing on Early Adulthood Outcomes," *Working Paper*, 2024.

Chetty, Raj, John N Friedman, and Jonah E Rockoff, "Measuring the impacts of teachers II: Teacher value-added and student outcomes in adulthood," *American Economic Review*, 2014, *104* (9), 2633–2679.

_ , Nathaniel Hendren, and Lawrence F Katz, "The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment," *American Economic Review*, 2016, *106* (4), 855–902.

_ , _ , Maggie R Jones, and Sonya R Porter, "Race and economic opportunity in the United States: An intergenerational perspective," *The Quarterly Journal of Economics*, 2020.

Chyn, Eric, "Moved to Opportunity: The Long-Run Effects of Public Housing Demolition on Children," *American Economic Review*, 2018, *108* (10), 3028–3056.

Collinson, Robert and Chloe R Gibbs, "In-Kind Transfers and Early Childhood Development: Evidence from Housing Assistance Lotteries," *Working Paper*, 2025.

_ **and Peter Ganong**, "How Do Changes in Housing Voucher Design Affect Rent and Neighborhood Quality?," *American Economic Journal: Economic Policy*, 2018.

_ **, Deniz Dutz, John Eric Humphries, Nicholas S Mader, Daniel Tannenbaum, and Winnie van Dijk**, "The Effects of Eviction on Children," *NBER Working Paper*, 2025.

_ **, Ingrid Gould Ellen, and Jens Ludwig**, "Low-Income Housing Policy," in "Economics of Means-Tested Transfer Programs in the United States," University of Chicago, 2015.

Cunha, Flavio, James J Heckman, and Susanne M Schennach, "Estimating the technology of cognitive and noncognitive skill formation," *Econometrica*, 2010.

Currie, Janet and Aaron Yelowitz, "Are public housing projects good for kids?," *Journal of public economics*, 2000.

_ **and Duncan Thomas**, "Does Head Start Make a Difference?," *American Economic Review*, 1995.

_ **and Firouz Gahvari**, "Transfers in Cash and In-Kind: Theory Meets the Data," *Journal of Economic Literature*, 2008.

_ **and Jessica Van Parys**, "Affordable Housing During Childhood Improves Long-Term Outcomes of Women and Their Children," *Journal of Human Resources*, September 2026.

Cutuli, J. J., Sandra Torres Suarez, Aaron Truchil, Tyler Yost, and Ciani Flack-Green, "Strategies to Better Identify Student Homelessness Using Data in an Urban School District," *Educational Researcher*, 2024.

Deming, David, "Early Childhood Intervention and Life-Cycle Skill Development: Evidence from Head Start," *American Economic Journal: Applied Economics*, 2009, *1* (3), 111–134.

Ellen, Ingrid Gould, Katherine O'Regan, and Sarah Strochak, "Race, Space, and Take Up: Explaining Housing Voucher Lease-Up Rates," *Journal of Housing Economics*, 2024.

Finkelstein, Amy and Matthew J Notowidigdo, "Take-Up and Targeting: Experimental Evidence from SNAP," *The Quarterly Journal of Economics*, 2019.

Garboden, Philip, "You Can't Get There from Here: Mobility Networks and the Housing Choice Voucher Program," *Journal of Planning Education and Research*, 2024.

Goldman, Benjamin and Jamie Gracie, "Every Day Counts: Absenteeism and the Returns to Education in High-Poverty Schools," *Working Paper*, 2024.

Grasley, Nick, "The Insurance Value of Federal Rental Assistance," *Working Paper*, February 2026.

Gray-Lobe, Guthrie, Parag A Pathak, and Christopher R Walters, "The long-term effects of universal preschool in Boston," *The Quarterly Journal of Economics*, 2023.

Han, Jeehee and Amy Ellen Schwartz, "Are Public Housing Projects Good for Kids after All?," *Journal of Policy Analysis and Management*, 2024.

Hanushek, Eric, "Teacher characteristics and gains in student achievement: Estimation using micro data," *AEA Papers & Proceedings*, 1971.

Heckman, James J., Jora Stixrud, and Sergio Urzua, "The Effects of Cognitive and Noncognitive Abilities on Labor Market Outcomes and Social Behavior," *Journal of Labor Economics*, 2006.

Heckman, James, Rodrigo Pinto, and Peter Savelyev, "Understanding the Mechanisms through Which an Influential Early Childhood Program Boosted Adult Outcomes," *American Economic Review*, 2013.

Hendren, Nathaniel and Ben Sprung-Keyser, "A unified welfare analysis of government policies," *The Quarterly journal of economics*, 2020.

Hoynes, Hilary, Diane Whitmore Schanzenbach, and Douglas Almond, "Long-run impacts of childhood access to the safety net," *American Economic Review*, 2016.

Humphries, John Eric, Christopher Neilson, Xiaoyang Ye, and Seth D Zimmerman, "Parents' Earnings and the Returns to Universal Pre-Kindergarten," *NBER Working Paper*, 2024.

Husock, Howard and Bruce Meyer, "Subsidized Housing and Upward Mobility," *National Affairs*, 2025.

Jackson, C. Kirabo, "What Do Test Scores Miss? The Importance of Teacher Effects on Non–Test Score Outcomes," *Journal of Political Economy*, 2018.

Jackson, C Kirabo and Claire L Mackevicius, "What impacts can we expect from school spending policy? Evidence from evaluations in the United States," *American Economic Journal: Applied Economics*, 2024.

Jacob, Brian A and Jens Ludwig, "The Effects of Housing Assistance on Labor Supply: Evidence from a Voucher Lottery," *American Economic Review*, 2012, *102* (1), 272–304.

_ , _ , and Douglas L Miller, "The Effects of Housing and Neighborhood Conditions on Child Mortality," *Journal of Health Economics*, 2013, *32* (1), 195–206.

_ , Max Kapustin, and Jens Ludwig, "The Impact of Housing Assistance on Child Outcomes: Evidence from a Randomized Housing Lottery," *The Quarterly Journal of Economics*, 2015, *130* (1), 465–506.

Katz, Lawrence F., "Neighborhood Effects and Missing Markets for Opportunity," *NBER Working Paper*, 2026.

Kline, Patrick and Christopher Walters, "Evaluating Public Programs with Close Substitutes: The Case of Head Start," *The Quarterly Journal of Economics*, 2016.

Kling, Jeffrey R, Jeffrey B Liebman, and Lawrence F Katz, "Experimental analysis of neighborhood effects," *Econometrica*, 2007.

Krause, Patrick K, Elizabeth Rhodes, Sarah Miller, Alexander W Bartik, David E Broockman, and Eva Vivalt, "The Impact of Unconditional Cash Transfers on Parenting and Children," *Working Paper*, 2025.

Liao, Hsi-Ling, Grace Ortuzar, Alison Lodermeier, and Stephen Stapleton, "Is There a "Scarlet E"? The Effects of Public Eviction Records on Low-Income Households," *Working Paper*, 2025.

Liu, Jing, Monica Lee, and Seth Gershenson, "The short-and long-run impacts of secondary school absences," *Journal of Public Economics*, 2021.

McCollister, Kathryn E., Michael T. French, and Hai Fang, "The Cost of Crime to Society: New Crime-Specific Estimates for Policy and Program Evaluation," *Drug and Alcohol Dependence*, 2010, *108* (1–2), 98–109.

Phillips, David C and James X Sullivan, "Do homelessness prevention programs prevent homelessness? Evidence from a randomized controlled trial," *Review of Economics and Statistics*, 2025.

Pollakowski, Henry O, Daniel H Weinberg, Fredrik Andersson, John C Haltiwanger, Giordano Palloni, and Mark J Kutzbach, "Childhood Housing and Adult Outcomes: a Between-Siblings Analysis of Housing Vouchers and Public Housing," *American Economic Journal: Economic Policy*, 2022.

Rafkin, Charlie, Adam Solomon, and Evan Soltas, "Self-Targeting in US Transfer Programs," *Working Paper*, 2025.

Reeder, William J, "The Benefits and Costs of the Section 8 Existing Housing Program," *Journal of Public Economics*, 1985.

Ribeiro, Bernardo and Gabriel Leite Mariante, "Subsidized housing and intergenerational mobility: evidence from Brazil," *Working Paper*, 2025.

Ruby, Sasha, "Targeting Transfers Through Priority Rules: Evidence from the Housing Choice Voucher Program," *Working Paper*, 2025.

Schwartz, Amy Ellen, Keren Mertens Horn, Ingrid Gould Ellen, and Sarah A Cordes, "Do Housing Vouchers Improve Academic Performance? Evidence from New York City," *Journal of Policy Analysis and Management*, 2020.

Stober, Preston, "New Haven Homeless Shelters Emphasizing Need for State Assistance," WTNH News 8 2026. Published April 2026. Accessed July 3, 2026.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, December 2021.

van Dijk, Winnie, "The socio-economic consequences of housing assistance," *Working Paper*, 2019.

Waldinger, Daniel, "Targeting In-Kind Transfers through Market Design: A Revealed Preference Analysis of Public Housing Allocation," *American Economic Review*, 2021, *111* (8), 2660–2696.

Zimmerman, Seth, "The returns to college admission for academically marginal students," *Journal of Labor Economics*, 2014.

A Appendix Tables and Figures

A.1 Appendix Tables

Table A.1: Child characteristics across lotteries

Lottery	Children	Offered by $t+5$	Assisted by $t+5$	Age	Female	Black	Adult Employed
HCV 2003	730	0.94	0.32	5.66	0.48	0.46	0.73
HCV 2005	602	0.73	0.20	5.93	0.50	0.56	0.77
LIPH 2007	1,271	0.74	0.24	4.74	0.50	0.59	0.74
LIPH 2008	575	0.48	0.20	4.38	0.46	0.56	0.77
HCV 2011	1,702	0.40	0.19	7.55	0.49	0.47	0.68
LIPH 2012	549	0.31	0.19	3.85	0.50	0.78	0.71
LIPH 2013	671	0.40	0.19	4.57	0.50	0.66	0.69
HCV 2015	132	0.36	0.23	6.83	0.48	0.60	0.90
HCV 2018†	1,167	0.31	0.07	5.86	0.49	0.64	0.83

Notes: Table A.1 reports characteristics of children at the level of the lottery. The offered and assisted columns report the share of children in households selected from the wait list (offered) or receiving assistance (assisted) by five years after the lottery year, the final year available for all lotteries in the sample. Age is the child's age at the time of the lottery. Female and Black are means of child demographic indicators. Adult employed is the share of children in families with any adult observed employed in the year before application, calculated among children in the linked SDE & DOL sample. The data source is the restricted lottery sample, which excludes child-lottery observations whose household received housing assistance before the lottery year. The HCV 2018 lottery is denoted with † because it is a "project based voucher" lottery.

Table A.2: Offer timing and linkage to administrative records

	(1) Raw	(2) Regression adjusted
SDE linked	-3.77^{***}	1.35
	(1.24)	(1.30)
SDE and DOL linked	-4.12^{***}	1.36
	(1.26)	(1.33)
Lottery fixed effects	No	Yes
Age fixed effects	No	Yes
Lottery × family size × disability FE	No	Yes
Demographic controls	No	Yes
SDE linkage rate (%)	75.55	75.55
SDE and DOL linkage rate (%)	74.40	74.40
Offered by year five (%)	52.32	52.32
Children	7,399	7,399
Households	4,194	4,194

Notes: This table reports OLS estimates of the impact of receiving an offer by five years after a lottery on administrative data linkage. Coefficients are reported in percentage points. Column (1) is unadjusted. Column (2) includes age fixed effects, fully interacted lottery × family size × household head disability fixed effects, and controls for race, sex, and citizenship. Standard errors, shown in parentheses, are clustered by household. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$, respectively.

Table A.3: Effects of housing assistance on school-reported homelessness

Child ages	OLS	RF	2SLS	CM	CCM	Observations	# Children	First stage F
6–11	-3.07***	-1.86***	-4.34***	4.22	6.70	16,349	3,794	885.8
	(0.68)	(0.63)	(1.47)	–	(1.38)	–	–	–
12–17	-1.63**	-0.90	-2.38	3.55	4.82	16,172	3,942	582.8
	(0.69)	(0.57)	(1.52)	–	(1.34)	–	–	–

Notes: This table reports the effects of housing assistance on school-reported homelessness, pooling across age bins. I estimate a single β_b coefficient within each bin. Coefficients and means are in percentage points. Receipt and offer are cumulative binary indicators through the outcome year. Each specification includes fixed effects for age at lottery, age, calendar year, years since lottery, and lottery × family size × disability. Demographic controls follow the main specification. Household-clustered standard errors appear in parentheses, including for the complier control mean (CCM). $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$, respectively.

Table A.4: Effects of housing assistance on latent developmental factor by gender

	Estimate	Bootstrap SE	p-value
Pooled	0.14	0.12	0.224
Boys	0.35**	0.17	0.040
Girls	0.01	0.17	0.888
Difference (Boys − girls)	0.34	0.24	0.148
N		2,117	

Notes: Appendix Table A.4 reports binary 2SLS effects in standard deviations of the latent developmental readiness factor. Creativity, physical/motor development, and personal/social development load on this factor. The endogenous treatment is assistance receipt by the kindergarten assessment year, instrumented with an offer by that year. Standard errors come from 500 bootstrap replications that resample households and re-estimate the measurement and causal models. $^{*}p < 0.10$, $^{**}p < 0.05$, and $^{***}p < 0.01$, respectively.

Table A.5: Effects of housing assistance on test scores, by grade and gender

	Pooled		Boys		Girls		Girls − boys
Grade	2SLS	Children	2SLS	Children	2SLS	Children	Difference
Panel A: Math							
Grade 3	0.149	1,369	0.300	666	0.056	703	-0.244
	(0.133)		(0.211)		(0.184)		(0.279)
Grade 4	0.066	1,676	0.130	858	0.018	817	-0.113
	(0.118)		(0.188)		(0.158)		(0.246)
Grade 5	0.045	1,797	0.166	912	-0.047	884	-0.213
	(0.107)		(0.166)		(0.147)		(0.219)
Grade 6	0.163	1,923	0.118	970	0.248	952	0.130
	(0.115)		(0.173)		(0.161)		(0.234)
Grade 7	-0.001	1,892	0.019	950	0.026	942	0.007
	(0.121)		(0.195)		(0.159)		(0.242)
Grade 8	0.033	1,790	0.010	858	0.043	932	0.034
	(0.135)		(0.205)		(0.170)		(0.256)
Panel B: English language arts							
Grade 3	0.178	1,373	0.493**	667	0.012	706	-0.481*
	(0.136)		(0.230)		(0.178)		(0.286)
Grade 4	0.102	1,685	0.293	865	-0.064	819	-0.358
	(0.124)		(0.197)		(0.170)		(0.261)
Grade 5	0.089	1,797	0.346*	914	-0.078	882	-0.424*
	(0.117)		(0.184)		(0.157)		(0.240)
Grade 6	0.049	1,936	-0.022	975	0.164	960	0.187
	(0.119)		(0.178)		(0.165)		(0.239)
Grade 7	-0.044	1,921	0.044	967	-0.046	954	-0.090
	(0.123)		(0.199)		(0.162)		(0.249)
Grade 8	-0.208	1,830	-0.264	876	-0.173	954	0.091
	(0.144)		(0.229)		(0.183)		(0.284)

Notes: Table A.5 reports 2SLS estimates of the effect of receiving any housing assistance by the assessment year on standardized math and English language arts scores. Standard errors, shown in parentheses, are clustered at the household level. First-stage F statistics range from 325.3 to 518.9 in the pooled models, from 152.1 to 257.7 for boys, and from 184.5 to 292.0 for girls. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

Table A.6: Effects of housing assistance on adult employment and earnings, by child age

	Ages 6–7	Ages 8–9	Ages 10–11	Ages 12–13	Ages 14–15	Ages 16–17
Panel A: Effect of receiving assistance by the corresponding child age						
Household adult earnings	2,322	2,488	2,973	3,181	7,036	12,432**
	(2,389)	(2,552)	(3,169)	(4,039)	(4,932)	(6,236)
IHS household adult earnings	0.175	0.653	1.012*	0.771	0.891	1.221
	(0.435)	(0.473)	(0.569)	(0.685)	(0.822)	(1.029)
Any adult employed	0.23	5.72	8.41	7.37	8.30	8.52
	(3.94)	(4.28)	(5.12)	(6.15)	(7.31)	(9.11)
First stage *F*-statistic	906.2	874.1	714.7	565.2	408.3	255.6
Panel B: Effect of one additional year of assistance by the corresponding child age						
Household adult earnings	417	434	746	400	1,065	2,073**
	(514)	(458)	(520)	(558)	(717)	(964)
IHS household adult earnings	0.021	0.067	0.181**	0.157*	0.120	0.156
	(0.100)	(0.086)	(0.091)	(0.094)	(0.118)	(0.151)
Any adult employed	0.06	0.57	1.52*	1.62*	1.11	1.13
	(0.91)	(0.78)	(0.82)	(0.84)	(1.04)	(1.34)
First stage *F*-statistic	623.3	695.8	585.3	499.2	327.4	210.6
Observations	6,850	7,511	7,672	7,507	6,540	5,474
Children	3,672	3,999	4,095	3,992	3,585	2,971
Households	2,556	2,642	2,598	2,505	2,244	1,844

Notes: Table A.6 reports 2SLS estimates of the effects of housing assistance on employment and earnings among adult household members identified at the child's selected lottery application. Each column pools observations within the indicated age bin. Panel A instruments an indicator for receiving assistance by the corresponding age with an indicator for having received an offer by that year. Panel B instruments cumulative years of assistance with cumulative years of offer exposure. Earnings are measured in 2023 dollars. Employment coefficients and standard errors are reported in percentage points. Standard errors, shown in parentheses, are clustered at the household level. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

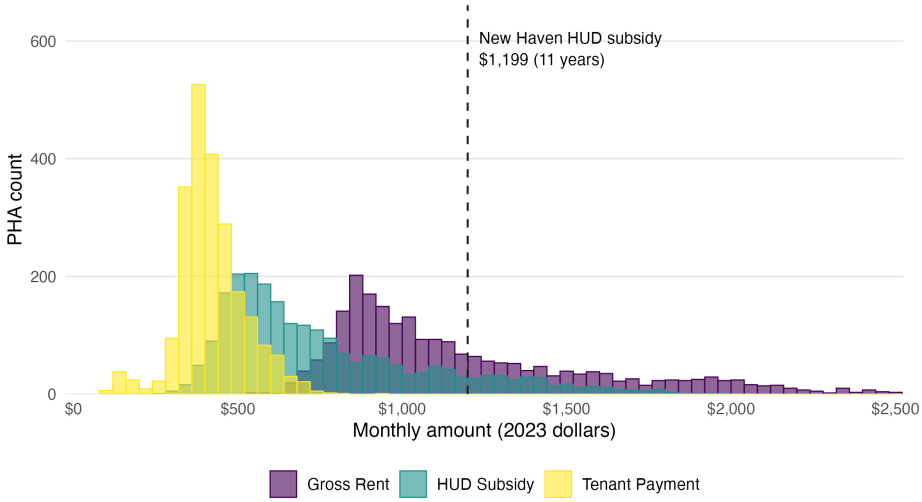
Table A.7: Effects of housing assistance on school quality and composition, by child age

	Ages 6–7	Ages 8–9	Ages 10–11	Ages 12–13	Ages 14–15	Ages 16–17
Panel A: Effect of receiving assistance by the corresponding child age						
ELA value added	-0.007	-0.007	-0.001	0.019	—	—
	(0.019)	(0.016)	(0.013)	(0.016)	—	—
Math value added	-0.007	-0.022	-0.005	0.005	—	—
	(0.022)	(0.018)	(0.013)	(0.016)	—	—
School proficiency rate	0.06	0.09	-1.60	-2.37	—	—
	(1.92)	(1.74)	(1.61)	(1.85)	—	—
School FRPL share	-0.55	-0.74	1.95	3.71*	5.92***	7.21**
	(1.54)	(1.54)	(1.57)	(1.91)	(2.21)	(2.85)
School homeless share	-0.28	-0.14	0.04	0.30	0.42*	0.30
	(0.23)	(0.25)	(0.23)	(0.20)	(0.23)	(0.31)
First stage F: value added sample	366.8	481.4	569.5	315.9	—	—
First stage F: proficiency sample	467.1	588.7	612.0	436.1	—	—
First stage F: composition sample	749.9	754.1	737.8	584.7	443.3	269.7
Panel B: Effect of one additional year of assistance by the corresponding child age						
ELA value added	0.001	0.000	0.004**	0.002	—	—
	(0.004)	(0.002)	(0.002)	(0.002)	—	—
Math value added	-0.003	-0.002	0.003	0.003	—	—
	(0.005)	(0.003)	(0.002)	(0.002)	—	—
School proficiency rate	0.44	0.15	0.26	0.06	0.16	-0.46
	(0.39)	(0.28)	(0.23)	(0.23)	(0.41)	(0.57)
School FRPL share	-0.28	-0.27	0.04	0.38	0.55**	0.99***
	(0.31)	(0.26)	(0.23)	(0.25)	(0.27)	(0.37)
School homeless share	-0.11**	-0.09**	-0.05	-0.02	0.00	0.03
	(0.05)	(0.04)	(0.03)	(0.03)	(0.03)	(0.04)
First stage F: value added sample	262.8	378.0	435.1	320.7	—	—
First stage F: proficiency sample	368.4	469.9	516.4	404.4	—	—
First stage F: composition sample	542.5	584.4	615.5	531.1	394.1	228.7
Value added observations	1,977	2,916	3,344	2,560	—	—
Proficiency observations	2,910	3,864	4,222	3,483	—	—
Composition observations	5,031	5,760	6,110	5,853	5,031	3,149

Notes: Table A.7 reports 2SLS estimates of the effects of housing assistance on the leave-one-out characteristics of children's schools. Each column pools observations within the indicated age bin. Panel A instruments an indicator for receiving assistance by the corresponding age with an indicator for having received an offer by that year. Panel B instruments cumulative years of assistance with cumulative years of offer exposure. Standard errors, shown in parentheses, are clustered at the household level. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively.

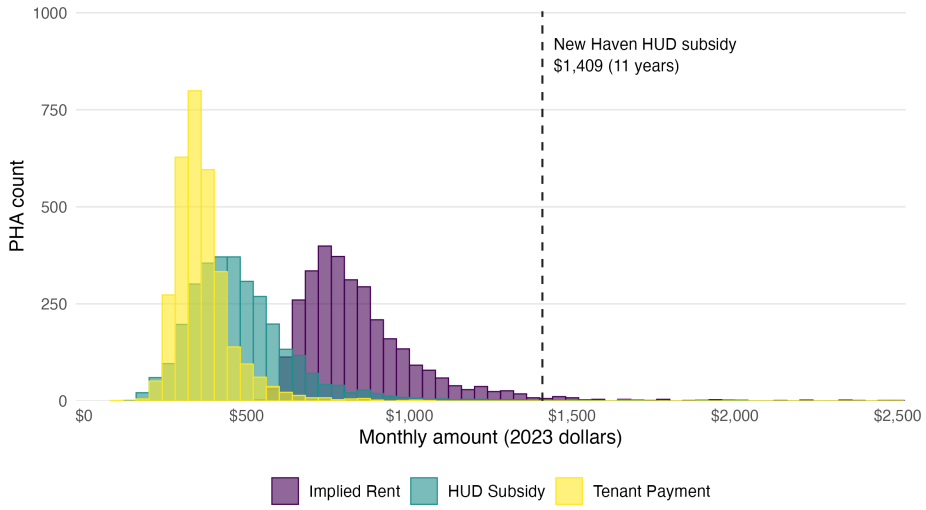
A.2 Appendix Figures

Figure A.1: Tenant payments, gross rents, and **voucher** subsidies across PHAs



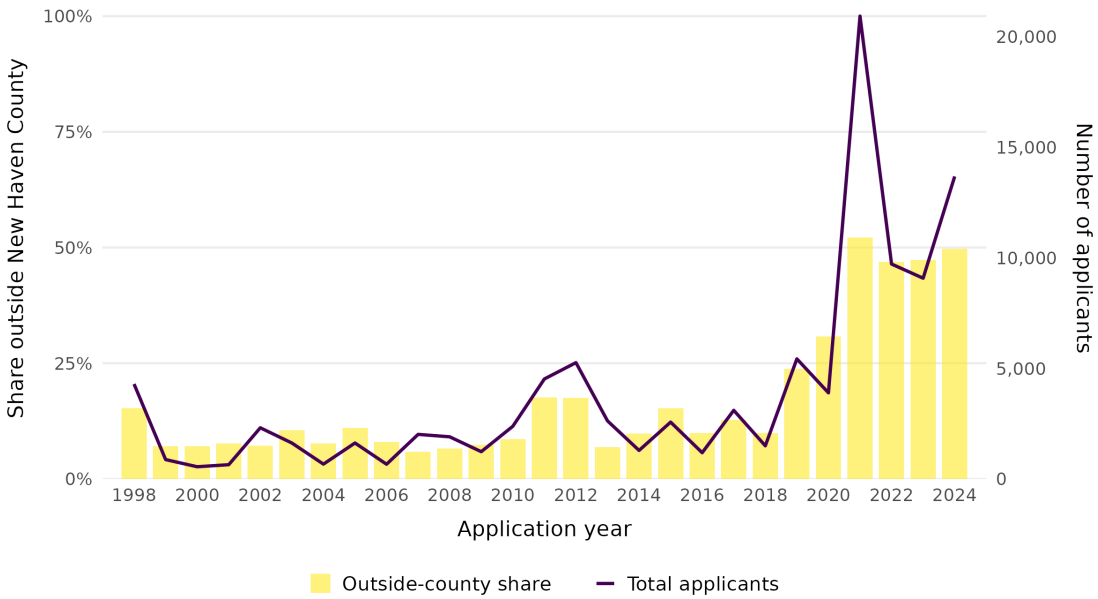
Notes: Figure A.1 plots the distribution of average monthly tenant payments, gross contract rents, and PHA expenditures (HUD subsidies) in the federal Housing Choice Voucher program from 2009 to 2023. All dollar values are in 2023 dollars. The results are shown in $40 bins. HUD records contain 2,283 PHAs with non-missing housing voucher data and average gross rent at or below $2,500. I exclude data from 2020. The New Haven PHA has an average subsidy of $1,199.

Figure A.2: Tenant payments, implied rents, and **public housing** subsidies across PHAs



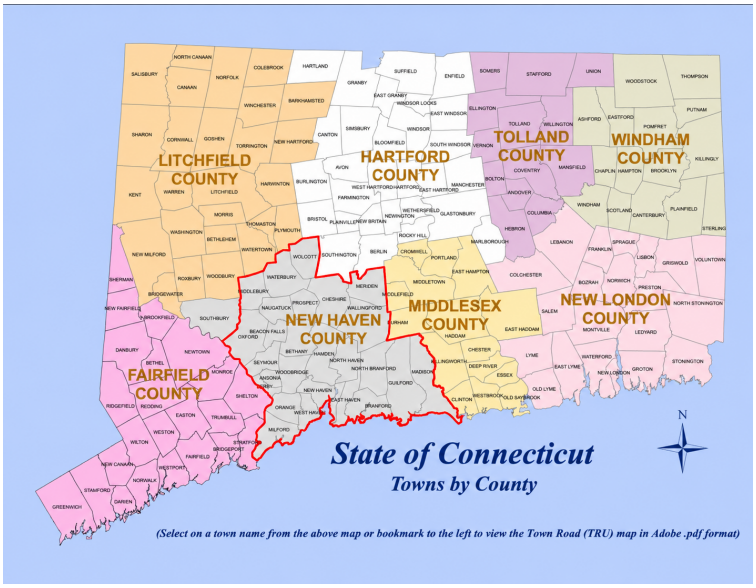
Notes: Figure A.2 plots the distribution of average monthly tenant payments, implied rents, and PHA expenditures (HUD subsidies) in the federal Low Income Public Housing program from 2009 to 2023. The results are shown in $40 bins. All dollar values are in 2023 dollars. HUD records contain 3,090 PHAs with non-missing public housing data and average implied gross rent (the sum of HUD subsidy and tenant payment) at or below $2,500. I exclude data from 2020. The New Haven PHA has an average subsidy of $1,409.

Figure A.3: Time series of PHA applicant counts and outside county share



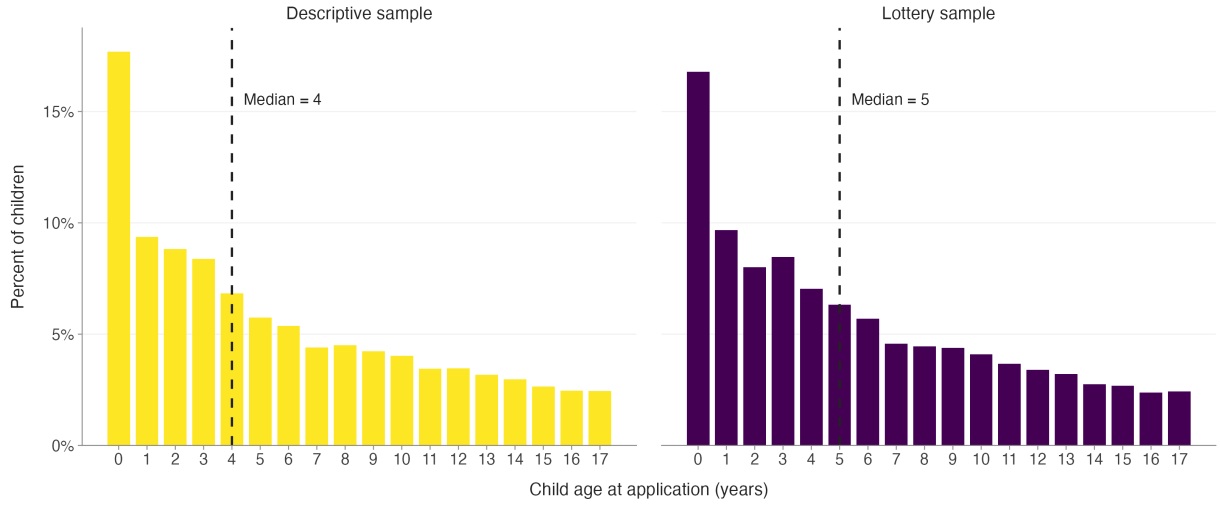
Notes: The left y-axis in Figure A.3 is the share of applicants from outside of New Haven county, measured using the last recorded ZIP per applicant. Assisted applicants are assumed to come from New Haven County. The right y-axis is the number of applicants. Each series is recorded annually from 1998 to 2024. The data source is the universe of PHA applications.

Figure A.4: Map of Connecticut counties



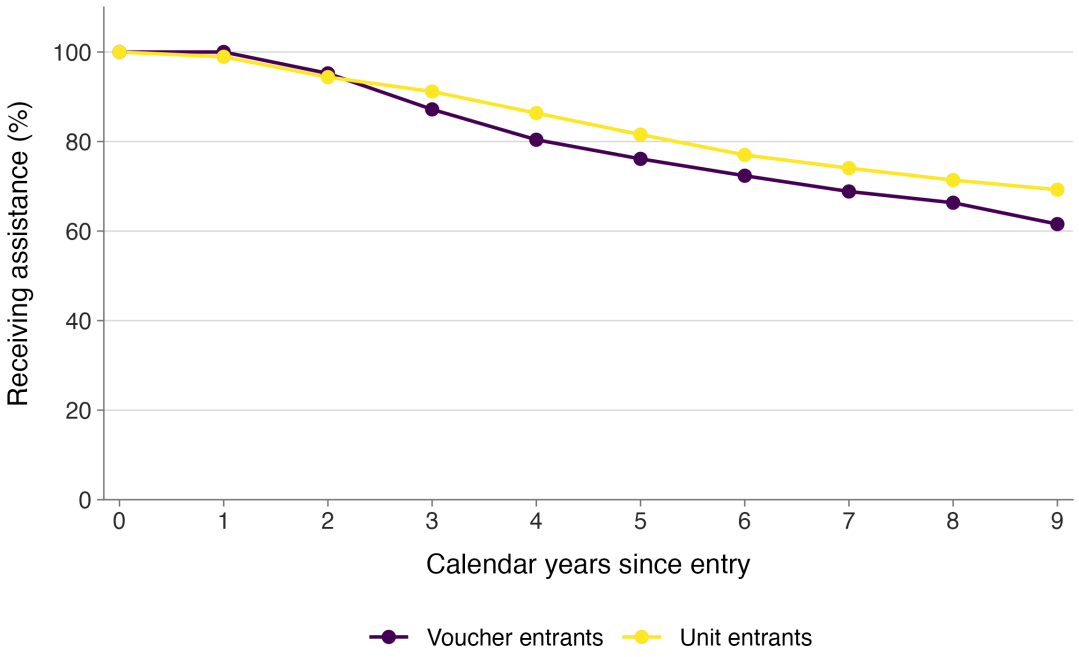
Notes: Figure A.4 displays a map of Connecticut counties. New Haven county is shaded in grey and outlined in red. The matched SDE attendance sample contains all Connecticut students that ever enroll in a public school in New Haven county between the 2010-2011 and 2023-2024 school years.

Figure A.5: Child ages at application



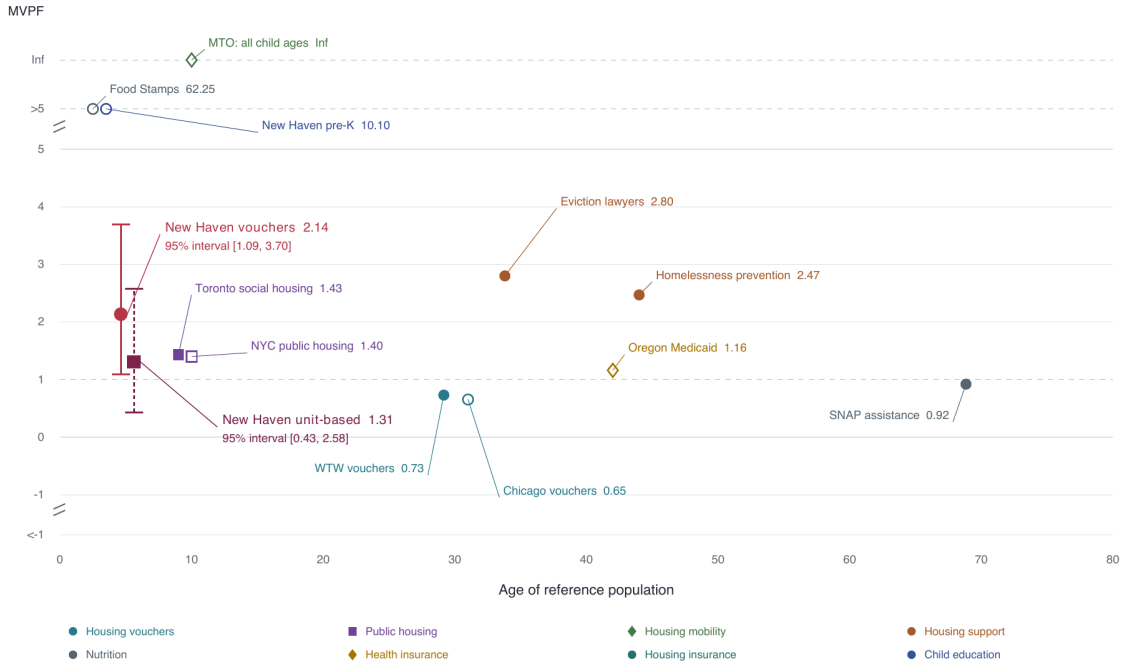
Notes: Figure A.5 displays the distribution of child age at application in the descriptive sample (left panel) and lottery sample (right panel). Bars report percentages within each sample. Dashed vertical lines indicate the sample medians.

Figure A.6: Assistance duration by program



Notes: Figure A.6 plots the percentage of households receiving any housing assistance in each calendar year relative to entry, separately for households initially entering the voucher and unit-based programs. The sample is restricted to lottery applicants that receive a voucher or move into a unit by 2015.

Figure A.7: Marginal Values of Public Funds across social insurance programs



Notes: Appendix Figure A.7 plots marginal values of public funds (MVPF) estimates for various social insurance programs against recipient ages. Chicago vouchers, Moving to Opportunity, and Oregon Medicaid use Hendren and Sprung-Keyser 2020. Welfare to Work vouchers use the equal-priority estimate in Ruby 2025. NYC public housing uses Chaudhry and Eng 2024. Toronto social housing uses Beauregard 2025. New Haven pre-K uses Humphries et al. 2024. Food Stamps uses Bailey et al. 2024. SNAP application assistance uses Finkelstein and Notowidigdo 2019. Homelessness prevention uses Phillips and Sullivan 2025. Eviction lawyers use Caspi and Rafkin 2026.

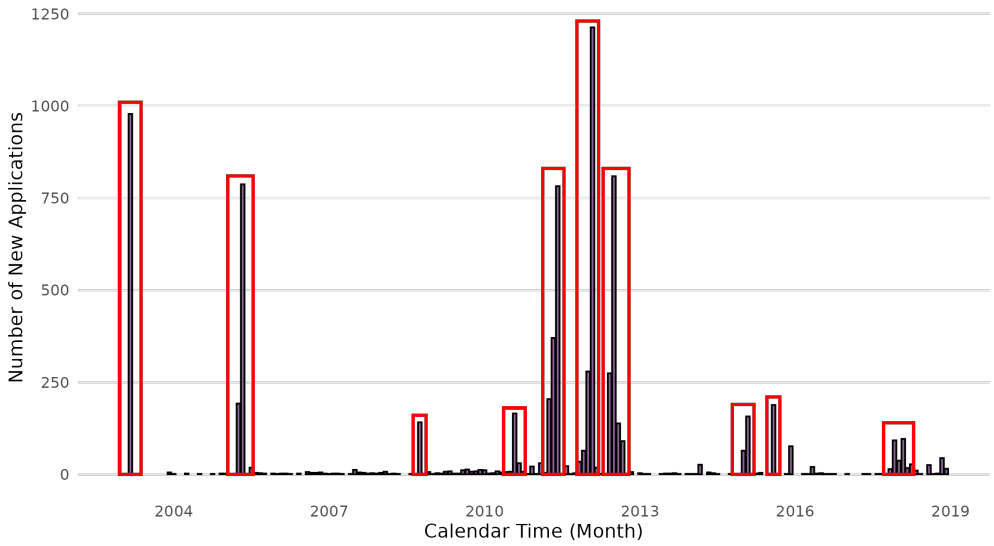
B Classifying wait list opening periods and lotteries

Wait lists for housing assistance are always open to households that are displaced by government action, include victims of domestic violence or sexual assault, are in a documented Witness Protection Program, or have direct referrals from the homelessness Coordinated Access Network. Households that do not meet these conditions may only apply for housing assistance from a particular PHA when its wait lists are open.

Applicants from outside traditional opening periods likely differ along observable and unobservable dimensions from other applicants. To address this concern, I identify wait list opening events by manually isolating structural breaks in the wait list-by-month count of new applications and cross referencing these structural breaks with local news articles announcing wait list openings. I use these wait list opening periods to define a subset of the descriptive sample that uses only clearly-defined wait list openings. I next combine these wait list openings with information on the assigned wait list position of applicants. To identify lotteries, I require that an applicant's wait list position be uncorrelated with the exact date and time of their application *within* the wait list opening period.

Opening periods In Figure B.1, I illustrate wait list opening periods on the tenant-based HCV wait list from 2003 to 2018.

Figure B.1: HCV wait list applications



Notes: The y-axis is the count of new applications to the PHA's tenant-based HCV wait list. The x-axis is the calendar time in months from January 2003 to December 2018. The red rectangles highlight wait list opening periods. Data comes from the PHA wait list application records.

I identify 10 wait list opening periods: March 2003, April to May 2005, October 2008, August 2010, April to June 2011, January to February 2012, June to September 2012, January to February 2015, August 2015, and December 2017 to February 2018. Whenever possible, I corroborate the identified openings using local news announcements or announcements of wait list opening events in PHA reports submitted to HUD annually from FY 2006 to FY 2023 (excluding FY 2007). For example, attached here is the body an article describing a lottery for 1,000 spots on the housing voucher wait list.

Lotteries To determine whether wait list opening events are lotteries, I check that the lottery ordering is not equivalent to ordering households by the exact date and time of application. Separately for each wait list opening j, I estimate – using the subsample of applicants without overwritten positions – the correlation β_j between the system-assigned position R_{ij} and the position an applicant would receive given the exact date and time of their application R_{ij}^{Date}

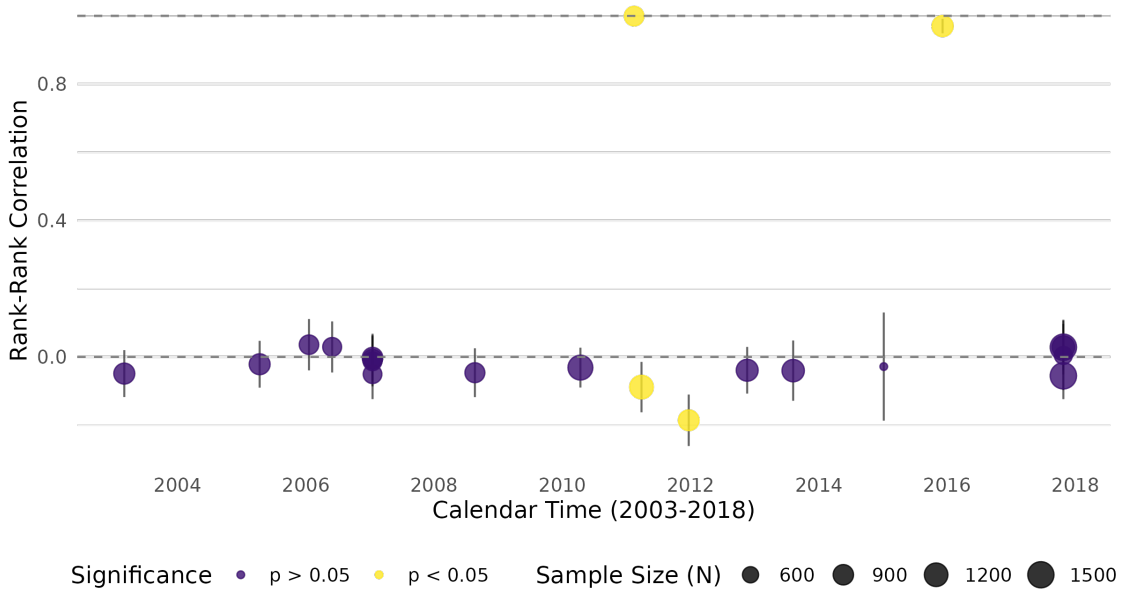
$$R_{ij} = \alpha_j + \beta_j R_{ij}^{\text{Date}} + \epsilon_{ij}. \tag{9}$$

I identify an opening period as a lottery whenever I fail to reject the null hypothesis that $\beta_j = 0$ at the 5% significance level. Some applications have lottery positions that are overwritten at the time of interview offer or when the household re-applies in wait list openings after 2018. In Figure B.2, I graphically present the coefficient estimates and standard errors from separately estimating Equation (9) for a subset of the manually-identified wait list opening periods from 2003-2018 across all wait lists.[23] The sample sizes are for the entire set of applicants regardless of overwritten position.

Absent concerns over selection, it appears from examining the $\hat{\beta}_j$ point estimates that most well-defined opening periods from 2003 to 2018 are lotteries. The regression test identifies 9 lottery opening periods across 14 wait lists and 4 non-lottery opening periods.

[23] I remove wait list openings restricted to elderly and disabled applicants. For example, the October 2008 HCV wait list opening is restricted to elderly and disabled applicants.

Figure B.2: System position - date position correlation



Notes: The y-axis is the $\hat{\beta}_j$ point estimate and 95% confidence interval from estimating Equation 9. The upper dashed line is at 1 where date position and system position are identical. The lower dashed line is at 0 where date position and system position are unrelated. The significance level is for a hypothesis test of $\beta_j = 0$. Data comes from the PHA wait list application records.

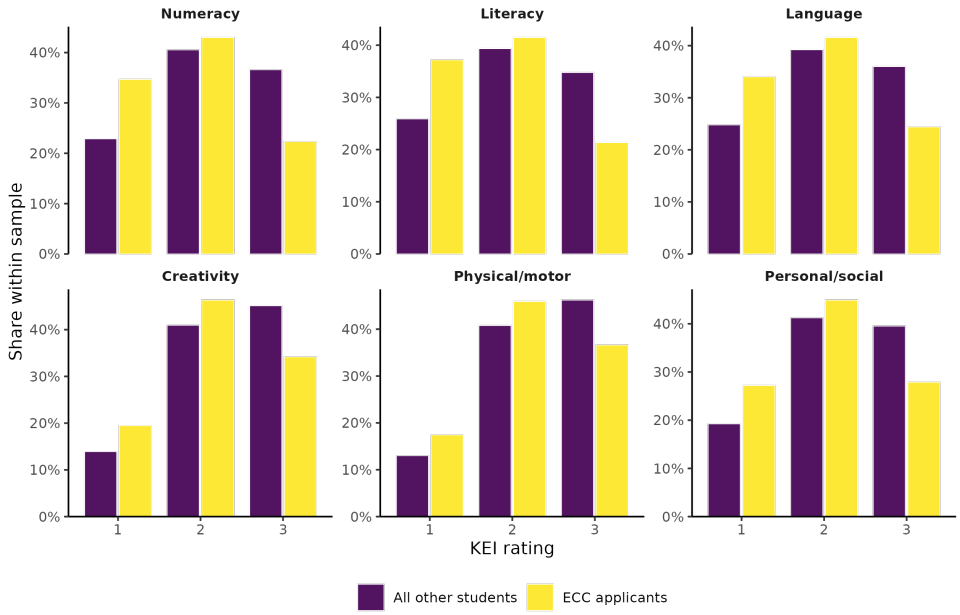
C Data appendix

C.1 Education data

State employees match children in PHA applicant households to State Department of Education (SDE) records using first name, last name, date of birth, and the last 4 digits of the Social Security number. They match 5,590 children in the lottery sample to unique SDE records (75.6%). Of these children, 5,505 are in households with a unique adult DOL match (98.4%).

C.1.1 Kindergarten test scores

Figure C.1: Distribution of kindergarten readiness ratings across categories



Notes: Appendix Figure C.1 reports the share of students at each raw KEI rating across categories. PHA applicants are in yellow, while all other students are in purple. Ratings range from one to three, with higher values indicating more consistent demonstration of the corresponding skills (and less need for instructional support). Each student contributes the earliest assessment with all six categories observed. The overall sample contains 136,502 students, while the PHA applicant sample contains 16,642 students.

This section describes my measures of student achievement for kindergarteners, which are the earliest measures of achievement observed in the SDE records. This section closely follows Humphries et al. 2024, which uses identical achievement measures to study the effects of universal pre-kindergarten (UPK) programs in New Haven.

I use the Connecticut Kindergarten Entrance Inventory (KEI) to measure student achievement in kindergarten. Kindergarten teachers conduct this assessment in the fall of each academic year. Students receive scores ranging from one to three on six readiness measures: language proficiency, numeracy, literacy, physical/motor skills, creativity, and personal/social skills. A higher score represents both a more consistent demonstration of the skills and less need for instructional support. I observe KEI scores for each school year from 2008-2009 through 2023-2024. I plot the raw scores by category (pooling across years) in Figure C.1.

Table C.1: Correlations across kindergarten readiness categories

	Numeracy	Literacy	Language	Creativity	Physical	Personal
Numeracy	1.000	—	—	—	—	—
Literacy	0.847	1.000	—	—	—	—
Language	0.734	0.744	1.000	—	—	—
Creativity	0.652	0.629	0.666	1.000	—	—
Physical	0.665	0.637	0.660	0.780	1.000	—
Personal	0.629	0.617	0.717	0.725	0.681	1.000
Observations			136,503			

Notes: Appendix Table C.1 reports the correlations among the six KEI categories in the overall SDE student population, using scores from the 2008-2009 to 2023-2024 school years. Each category is standardized within assessment year. The sample contains each student's earliest assessment.

To construct the measures of kindergarten achievement used in the paper, I take three approaches:

1. **Raw scores:** I directly use students' ordinal scores in each category.

2. **Standardized scores:** I standardize each KEI category score within each cohort. For category k, the standardized score for student i in year t is given by

$$z_{it}^k = \frac{KEI_{it}^k - \overline{KEI}_t^k}{\sigma_t^k},$$

where the mean and standard deviation are taken over the population of New Haven county kindergarten students. I aggregate these scores into a KEI achievement index

$$\bar{z}_{it} = \frac{\sum_{k=1}^{6} z_{it}^k}{6}.$$

3. **Latent factors:** Student scores in the six categories may represent noisy measures of a lower-dimensional vector of underlying student skills. I use a two-factor latent

factor model – described in Appendix D.1 – to decompose each observed score into variation attributable to these two latent skills and category-specific measurement error (Heckman et al. 2006; Cunha et al. 2010).

I report the student-level correlation across scores in Table C.1. Scores are positively correlated across categories, with the strongest correlation between literacy and numeracy (0.85) and physical and creativity (0.78).

C.1.2 End-of-year test scores

This section describes my measures of student achievement for students in grades 3 through 8, which are the most consistent measures of achievement observed in the SDE records.

I use the Connecticut Smarter Balanced Assessment (SBA) to measure student achievement in grades 3 through 8. SBA records are available beginning in the 2014-2015 school year. The SBA measures student end-of-year achievement in English Language Arts (ELA) and Mathematics. The SBA consists of a computer adaptive test (in ELA and Mathematics) and a performance task (in Mathematics). Students are assigned a vertical score (which ranges from 2,114-2,862) and a performance level (which ranges from 1-4) in each category. Students assigned a performance level of 1 or 2 do not meet or are approaching expectations. Students assigned a performance level of 3 or 4 meet or exceed expectations. In this paper, I use the vertical score s_{it} to measure the achievement of student i in year t. I standardize this score s_{it} within each grade-year for the full sample of New Haven county students so that the ELA and Mathematics scores are given by

$$z_{it}^{E} = \frac{s_{it}^{E} - \overline{s_{g(i)t}^{E}}}{\sigma_{g(i)t}^{E}} \quad \text{and} \quad z_{it}^{M} = \frac{s_{it}^{M} - \overline{s_{g(i)t}^{M}}}{\sigma_{g(i)t}^{M}}$$

where $g(i)$ denotes the grade of student i and s_{it}^{E} and s_{it}^{M} are the vertical scores in ELA and Mathematics.

C.1.3 Other outcomes

This section describes the other outcomes that I measure using the SDE records. I focus on student attendance and chronic absenteeism in addition to homelessness and housing instability. These data are available from the 2010-2011 to 2023-2024 school years.

Attendance I observe the number of school days attended by each student, the number of possible school days, and an indicator for chronic absenteeism. In the full sample of

New Haven county students, 17.4% of student-years involve chronic absenteeism (i.e., missing more than 10% of instructional days in a school year). The average attendance rate is 93.4%, and 44.6% of students are chronically absent in at least one school year.[24] Among all matched PHA applicants, the average attendance rate is 89.6% and 73.2% of students are chronically absent in at least one school year. Within the matched lottery sample, 79.5% of students are chronically absent in at least one school year. The average attendance rate in the lottery sample is 90.1%.

Homelessness The McKinney-Vento Act of 1987 requires public schools to collect information on student homelessness and housing insecurity. To satisfy the requirements of the Act, school districts appoint local homeless-education liaisons who attempt to identify students experiencing homelessness – often through enrollment questionnaires, staff referrals, and outreach – to connect them with academic and support services. There is some evidence that homelessness is underreported in McKinney-Vento records (Cutuli et al. 2024). The SDE records indicate whether students are experiencing one of four different types of housing instability: doubling up, living in a shelter, living in a hotel or motel, and experiencing unsheltered homelessness. In the lottery sample, 7.1% of children ever experience doubling up, 4.4% of children are ever recorded in a homeless shelter, and 2.2% of children are ever recorded living in a hotel or motel. Rates of unsheltered (i.e., street) homelessness are very low.

[24]The average attendance rate is taken across all student-years. As such, this number is a weighted average across students with the weights in proportion to the number of school years per student.

C.1.4 Test score value added

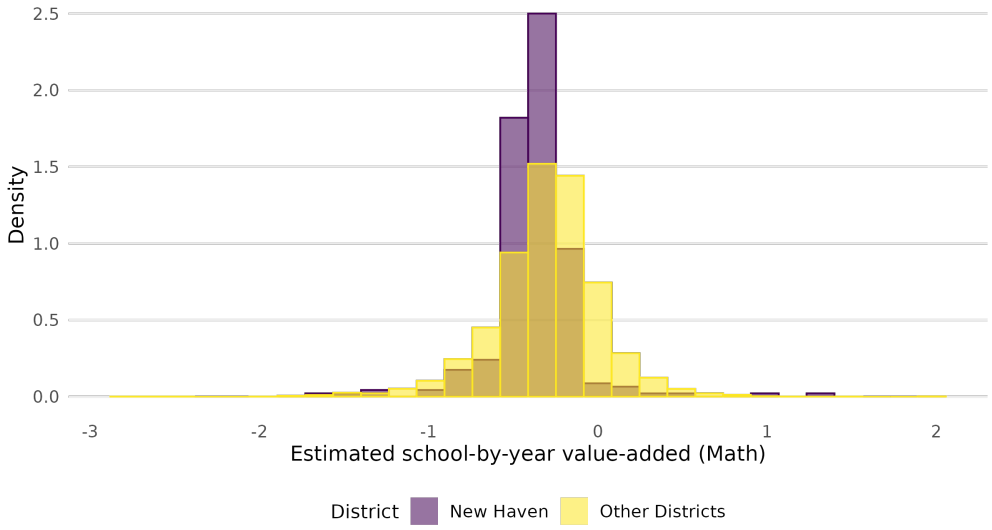
In this section, I present a simple statistical model of test score value added in the spirit of Hanushek 1971, which I use to characterize the quality of schools available to students in New Haven county in grades 3 through 8.

Let A_{igst} denote student i's end-of-grade test score in grade g at school s in year t. I standardize scores within each grade-year to obtain Z_{igst}. I estimate school value added from the lagged-score model

$$Z_{igst} = \theta Z_{i,g-1,s',t-1} + X'_{it}\beta + \gamma_{st} + \epsilon_{igst}, \tag{10}$$

where X_{it} is a vector of student-level controls, and γ_{st} is the school-by-year value added parameter.[25] Because information on teachers is not observed in the state administrative records, the $\hat{\gamma}_{st}$ estimates combine pure "school" effects and the impacts of teacher sorting across schools (Chetty et al. 2014). I construct the $\hat{\gamma}_{st}$ estimates using OLS. I plot the distribution of school-by-year fixed effects separately for math and ELA in Figure C.2 (math) and Figure C.3 (ELA) below, with separate series for New Haven schools and other schools in New Haven county.

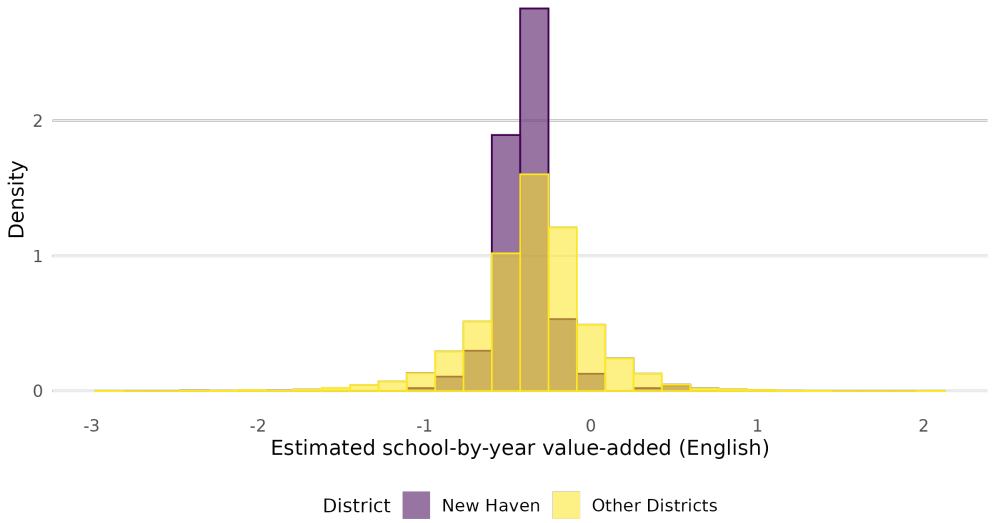
Figure C.2: Grade 3-8 math value added distribution



Notes: Appendix Figure C.2 displays the distribution of school-by-year fixed effects from estimating Equation 10 using end-of-year math test scores for students in grades 3-8. School-year fixed effects are calculated for the 2015-2016 to 2023-2024 school years.

[25] The X_{it} vector contains information on student i's contemporaneous FRPL status, homelessness status, and gifted and talented program status, as well as lagged attendance.

Figure C.3: Grade 3-8 English value added distribution



Notes: Appendix Figure C.3 displays the distribution of school-by-year fixed effects from estimating Equation 10 using end-of-year ELA test scores for students in grades 3-8. School-year fixed effects are calculated for the 2015-2016 to 2023-2024 school years.

C.2 Earnings and employment data

In this section, I describe the earnings and employment records provided by the Connecticut Department of Labor (DOL). DOL provides me with a panel of earnings at the individual-by-quarter-by-job level (employer information is not provided), containing all wages earned in Connecticut covered by unemployment insurance.[26] The file includes anonymous individual identifiers, each job's 6-digit NAICS code, and the individual's earnings in each job-quarter with nonzero recorded earnings. The observations run from the first quarter of 2001 to the final quarter of 2021.

DOL attempts to match the full universe of PHA applicants to their earnings records using a combination of name, date of birth, and the last 4 digits of the Social Security Number. DOL successfully assigns a unique matched identifier to 90.6% of all PHA applicants.

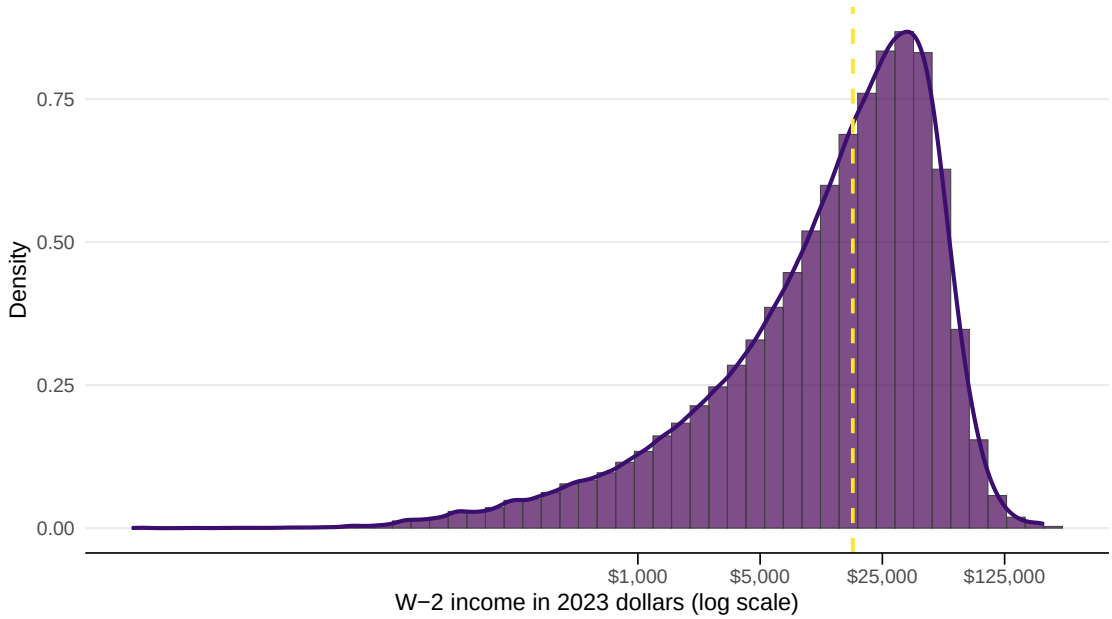
Data processing Starting from the 149,885 uniquely-matched individuals, I construct an individual-by-quarter panel from Q1 2001 - Q4 2021. The value for individual earnings in each quarter is the sum of quarter-by-job earnings for each individual. In periods when an individual is missing from the DOL earnings panel, I assume that the individual's earnings in that quarter are \$0.[27] To standardize the earnings measures over time, I deflate quarterly earnings using the January 2023 CPI, meaning that all earnings measures are in 2023 dollars.

I produce a balanced panel of 12,590,340 observations for 149,885 individuals. I top-code individual earnings in the full panel at the 99.9^{th} percentile of quarterly real earnings in the conditional distribution of earnings given employment in the sample (\$52,400.07 in 2023 dollars). Figure C.4 shows the conditional distribution of earnings in the matched PHA sample, pooled across all years from 2001-2021 and restricted to working age adults. The median annual earnings in this sample are \$16,939.88, and the modal primary sector of employment is "Temporary Help Services" (NAICS code 561320).

[26] This data covers the vast majority of payroll employment, but there are a few entities (e.g., some churches) that are exempt from unemployment insurance. There are other limitations to wage records in Connecticut that are common to many state administrative data systems: DOL records do not contain W-2 income from working for employers in other states, wage records for federal employees, or wage records for individuals employed in the informal or gig economy.

[27] An alternative explanation is that the individual's state of residence is not Connecticut or that the individual is either not yet born or deceased.

Figure C.4: Conditional earnings density among PHA applicants

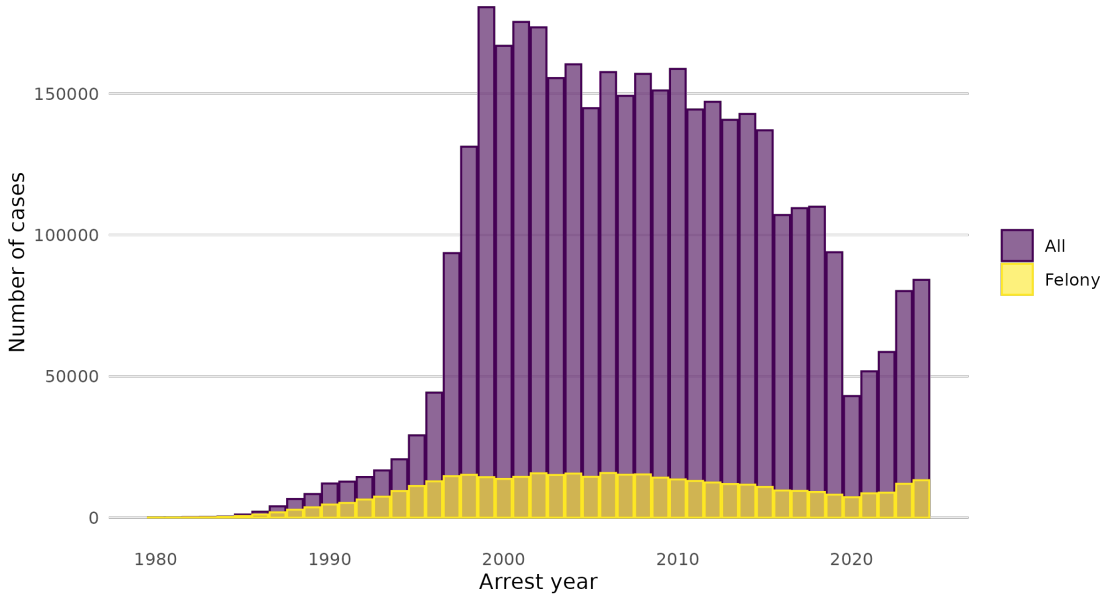


Notes: The x-axis in Appendix Figure C.4 is individual annual earnings, reported in 2023 dollars and shown on a log scale. The density is at the individual-by-year level for all years from 2001-2021. The dashed yellow line reports the median earnings conditional on employment, which is $16,939.88. The data source is the sample of PHA applicants that are uniquely matched to DOL earnings records, restricted to individuals ages 18 to 65 in a given year. The unemployment rate (i.e., the share of individual-year combinations with zero earnings) is 58.2%.

C.3 Criminal court data

In this section, I describe the criminal case records provided by the State of Connecticut Judicial Branch (CJB). I use the September 9, 2024 daily refresh of cases from CJB containing 3,778,246 criminal and motor-vehicle infractions. These records contain pending cases and conviction records for criminal and motor-vehicle matters, including felonies, misdemeanors, infractions, and violations. The data include case-level, defendant-level, charge-level, and sentencing files, which I link using docket identifiers, but exclude sealed, erased, pardoned, clean slate, and otherwise non-disclosable cases. Historical coverage is limited for cases disposed before 2000, which were often archived after sentence completion and may not appear in the current refresh. Because my sample begins in 2003, this does not affect data availability for my sample. Figure C.5 shows the distribution of convictions by arrest year.

Figure C.5: Connecticut criminal and motor vehicle infractions



Notes: The y-axis in Appendix Figure C.5 is the count of convictions involving criminal or motor vehicle infractions in Connecticut. Separate series are included for all cases and cases involving at least one felony charge. The x-axis is the calendar time in years from January 1980 to 2024. The data source is the full sample of pending cases and conviction records available in the State of Connecticut Judicial Branch's daily criminal database refresh. Sealed, erased, pardoned, archived, and otherwise non-disclosable cases are not included. Cases disposed prior to 2000 are incompletely observed due to archival rules.

Matching To link PHA applications to records from CJB, I first standardize individual names in the two data sources by performing basic string processing: I transliterate non-ASCII characters, convert names to lowercase, remove punctuation and other

non-alphabetic characters, collapse leading or trailing whitespace, and remove common suffixes. I then discard CJB records with missing defendant names and form potential matches between individuals in the PHA records and CJB defendants that exactly match on date of birth (DOB). Within each DOB group, I compute Jaro-Winkler similarities between surnames and given names across all potential individual pairs. I retain potential matches if name similarity is at least 0.90. For each CJB defendant record, I then select the highest-scoring candidate in the PHA records with Jaro-Winkler similarity of at least 0.9 and link the resulting individual identifier back to the defendant's court dockets.

D Measurement appendix

D.1 Latent factor model

This appendix describes the latent factor estimates reported in Table A.4. I summarize the six Kindergarten Entrance Inventory (KEI) categories using two correlated factors: academic readiness and developmental readiness. Estimation proceeds in two steps. First, I estimate how the six category-specific scores load on these two factors. Second, I use the estimated loadings to combine the category-specific effects of an assistance offer into effects on the factors, which I scale by the effect of an offer on assistance receipt.

D.1.1 Scores and estimation sample

The KEI records teacher assessments of numeracy, literacy, language, creativity, physical/motor development, and personal/social development. Each assessment uses three ordered categories, with higher values indicating greater readiness. I use the category scores standardized within assessment year in the source data and treat these standardized scores as continuous outcomes.

I require observed scores in all six categories and complete data for the controls, fixed effects, treatments, and instruments used in the binary and exposure specifications, including their program-specific versions. The resulting sample contains 2,117 children from 1,734 households. There are 1,078 boys, 1,038 girls, and one child without recorded gender.

D.1.2 Measurement model

Let Y_i be the vector of six standardized category scores for child i. I model their covariance using

$$Y_i = \mu + \Lambda \eta_i + \varepsilon_i, \qquad \eta_i = \begin{pmatrix} \eta_i^A \\ \eta_i^D \end{pmatrix}, \tag{11}$$

where η_i^A denotes academic readiness and η_i^D denotes developmental readiness. Numeracy, literacy, and language only load on academic readiness. Creativity, physical/motor development, and personal/social development only load on developmental readiness.

Both factors are assumed to have mean zero and unit variance in the joint measurement model, but the correlation between the two factors is estimated. The category-specific residuals ε are uncorrelated with the latent factors and with one another by assumption.

Any residual may reflect category-specific skills and variation in teacher assessment practices, as well as measurement error. Throughout, I orient the factors so that the numeracy loading on academic readiness and the creativity loading on developmental readiness are positive.

I estimate the loadings, residual variances, and factor correlation by fitting the implied covariance matrix

$$\Sigma = \Lambda \Phi \Lambda' + \Theta$$

to the pooled covariance matrix of the six scores, using Gaussian maximum-likelihood. Here, Φ is the factor correlation matrix and Θ is the residual covariance matrix. I use the same pooled loading matrix for the pooled, boys', and girls' treatment effect estimates, allowing all three estimates to be placed on a common scale.

D.1.3 Effects of housing assistance

For the binary treatment estimates, let D_i indicate any recorded assistance receipt between the lottery and the assessment year, and let Z_i indicate an offer by that year. Treatment is absorbing.

Within each estimation group g – the pooled sample, boys, or girls – I estimate an offer reduced form for each of the six categories and a linear first stage for assistance receipt:

$$Y_{ijk} = b_{kg} Z_{ij} + X_{ij}' \gamma_{jkg} + u_{ijk}, \tag{12}$$

$$D_{ij} = \pi_g Z_{ij} + X_{ij}' \gamma_{Djg} + v_{ij}. \tag{13}$$

Here, j indexes lotteries and k indexes KEI categories. The control vector X_{ij} includes fixed effects for age at the lottery, age at assessment, assessment year, and lottery-by-family-size-by-disability allocation cells.

Let $\widehat{b}_g$ collect the six estimated offer effects and let $\widehat{\Lambda}$ be the pooled loading matrix. I assume that the offer-induced changes in category-specific scores operate through the two modeled factors:

$$b_g = \Lambda \delta_g,$$

where δ_g contains the offer effects on academic and developmental readiness.

I estimate δ_g by minimum distance:

$$\widehat{\delta}_g = \left(\widehat{\Lambda}' \widehat{W}_g \widehat{\Lambda} \right)^{-1} \widehat{\Lambda}' \widehat{W}_g \widehat{b}_g, \qquad \widehat{\tau}_g = \frac{\widehat{\delta}_g}{\widehat{\pi}_g}, \tag{14}$$

where $\widehat{W}_g$ is the inverse of the household-clustered covariance matrix of the six estimated reduced-form coefficients. The two elements of $\widehat{\tau}_g$ are the estimated assistance effects on academic and developmental readiness, expressed in standard deviations of the latent factor. This procedure estimates effects directly from the category-specific reduced forms and the measurement model.

D.1.4 Inference and model assessment

To conduct inference, I use 500 household bootstrap replications. Each replication samples households with replacement, retaining all children within each selected household. I then re-estimate the pooled measurement model, the category-level reduced forms, and the receipt first stages, and recompute the factor effects.

The estimated correlation between the academic and developmental factors is 0.827. To assess the fit of the measurement model, I report standard fit statistics: CFI= 0.966, TLI= 0.935, RMSEA= 0.137, and SRMR= 0.041. RMSEA is relatively high, so I treat this exercise as a complement to my other results. A separate comparison of models with common versus gender-specific loadings does not reject equal loadings across boys and girls ($p = 0.236$).

D.2 Earnings projections and the MVPF

This appendix describes how I map the estimated effects of housing assistance on college enrollment into lifetime earnings effects. I first describe the benchmark calibration from Hendren and Sprung-Keyser 2020 that predicts lifetime earnings effects using the point-in-time earnings estimates from Zimmerman 2014. I then use cross-sectional data on earnings from the American Community Survey (ACS) to examine how sensitive my results are to differences in household and child characteristics. All dollar amounts are expressed in 2023 dollars.

D.2.1 Baseline earnings calibration

The baseline MVPF estimates use the Florida International University (FIU) calibration of Hendren and Sprung-Keyser 2020, based on Zimmerman's estimates of the returns to college admission. They assign an induced admission $142,757 in lifetime pretax earnings, $2,617 in public college costs, and $2,851 in private tuition, in 2012 dollars. Each amount is already discounted to age 18. I convert these values to 2023 prices and multiply them

by $s = 0.234/0.119 = 1.966$, the ratio of the admission effect to the state-university attendance effect in Zimmerman 2014. This produces values per induced attendee of $372,547 in earnings, $6,829 in public costs, and $7,440 in private tuition.

It is useful to distinguish the FIU estimates from the general method for valuing intermediate earnings outcomes in Hendren and Sprung-Keyser 2020. For the calculation of the returns to admission into a four-year higher education institution, Hendren and Sprung-Keyser 2020 forecast the estimated effects on point-in-time earnings from Zimmerman 2014 by constructing a life cycle earnings profile from the ACS. An alternative method would express the same earnings effects in percentage terms – rather than in dollars – when constructing the counterfactual earnings profile. At baseline, I use the dollar value estimates and the percentage effects as a robustness check.

D.2.2 Earnings profiles and projected college returns

I construct representative earnings profiles from the national 2015 ACS one-year person-level public-use microdata sample. The sample includes 1,948,411 individuals aged 18 to 65, including people with zero wages. Earnings combine wage and salary income, adjusted to survey-year dollars using the ACS income adjustment factor. I use person weights to calculate mean earnings at each age.

Let μ_{ag} denote mean earnings at age a for group g, measured in 2015 dollars. I construct a cohort entering college at age 18 in 2020, allow for annual real wage growth $\gamma = 0.005$, and express earnings in 2023 dollars:

$$w_{ag} = \mu_{ag} \frac{\text{CPI}_{2023}}{\text{CPI}_{2015}} (1+\gamma)^{2020+a-18-2015}. \tag{15}$$

In other words, I use the cross-sectional age profile of earnings in the ACS to construct the life-cycle earnings profile used in my MVPF estimates.

Following Hendren and Sprung-Keyser 2020, the projected earnings effect at age a for each group g college attendee is

$$\Delta w_{ag} = s\, r_a w_{ag}, \qquad r_a = \begin{cases} 0, & a = 18, \\ -12,294/94,368, & 19 \le a \le 25, \\ 1,593/7,241, & 26 \le a \le 65. \end{cases} \tag{16}$$

The early period between ages 19 and 25 captures foregone earnings during college. The later period between ages 26 and 65 applies the proportional earnings gain through age

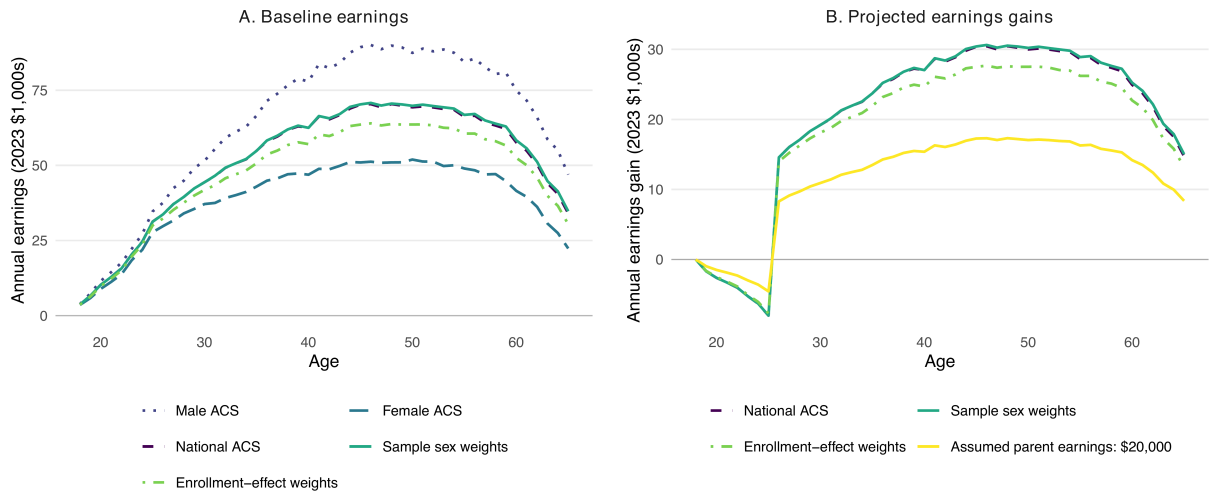
65. Before the attendance conversion, these changes are approximately −13% and 22%, respectively. At age p, the present value earnings effect is

$$L_g(p) = \sum_{a=18}^{65} \frac{\Delta w_{ag}}{(1+r)^{a-p}}, \qquad r = 0.03. \tag{17}$$

D.2.3 Alternative earnings calibrations

In Figure D.1, I report ACS earnings profiles and projected annual earnings gains using different sample weights. In addition to the national profile, I use the female and male shares of the college estimation sample, 49.62% and 50.38%, to construct a sample-weighted profile. I also use weights that are proportional to each group's sample size scaled by its estimated enrollment response in the sample. These enrollment-effect weights are 67.15% female and 32.85% male.

Figure D.1: ACS earnings profiles and projected college earnings gains



Finally, I consider an earnings path that accounts for the low adult (i.e., parent) earnings in my sample. Assuming that annual combined parental pretax earnings are $20,000, close to the average earnings reported in Table 1, I use the national relationship between parental household-income rank and children's adult individual-income rank from Chetty et al. 2020 to predict each child's individual-income rank, and convert that rank to income at age 34, rescaling the life-cycle ACS earnings profile.

Table D.1 reports two further sensitivity checks. First, I use earnings profiles constructed from Connecticut residents in the 2015 ACS one-year PUMS, weighted to the sex composition of my sample. Second, I halve the projected returns to college enrollment between

ages 26 and 65. Finally, I show how discounting the future college costs, earnings gains, and income tax revenues affects the MVPF estimates across every specification under the assumption that a year of housing assistance is provided at age 12 rather than 18.

Table D.1: Earnings calibration and MVPF sensitivity

Calibration	Earnings PV ($1,000)	MVPF: age 18 $\lambda = 0$	MVPF: age 18 $\lambda = 1$	MVPF: age 12 $\lambda = 1$
Imported HSK/FIU package	372.5	1.25	2.26	2.02
FIU observed effects + ACS forecast	438.0	1.31	2.56	2.25
National ACS	431.8	1.31	2.53	2.23
Sample sex weights	434.1	1.31	2.54	2.24
Enrollment effect weights	396.4	1.27	2.37	2.10
Connecticut ACS: sample sex weights	576.2	1.46	3.30	2.80
Half the later earnings return	203.5	1.12	1.60	1.50
Assumed parent earnings: $20,000	245.5	1.15	1.75	1.62

Notes: Earnings present values are in 2023 dollars per induced attendee, discounted to age 18. MVPFs use one assisted household-year and $\theta = 1$. At $\lambda = 0$, child earnings enter tax revenue but not willingness to pay. The age 12 column discounts earnings, public college costs, and private tuition by 1.03^{-6} and recomputes the ratio. The parental income row assumes parents are aged 35.