The Effects of Free Public Housing on Children*

Adriana Camacho[1], Valentina Duque[2], Michael Gilraine[3], and Fabio Sanchez[1]

[1]University of Los Andes

[2]University of Sydney

[3]New York University

July 27, 2021

Abstract

This paper analyzes the effect of Colombia's ambitious "Free Housing" program on children's educational outcomes. The program was generous, giving *free* housing units to beneficiaries. We evaluate the program by leveraging housing unit lotteries and linking lottery applicants to their children's education data. We find that receiving free public housing increases high school graduation rates by six percentage points, a *fifteen* percent increase relative to the control mean. Contrary to other well-known housing programs, the "Free Housing" program both provided adequate housing and was purposefully located in desirable areas. We show that much of the gain in graduation rates can be attributed to the better schools attended by lottery winners.

Keywords: Public Housing; Education.
JEL codes: I38, II4, I24.

*Jorge Caputo and Sofia Collante provided excellent research assistance. All remaining errors are our own. Contact emails: Camacho, adcamach@uniandes.edu.co; Duque, valentina.duque@sydney.edu.au; Gilraine, mike.gilraine@nyu.edu; Sanchez, fasanche@uniandes.edu.co.

1 Introduction

Children growing up in high-poverty areas fare worse than those who live in lower-poverty neighborhoods on a wide range of economic, health, and educational outcomes, with quasi-experimental and experimental evidence indicating this relationship is causal (Aaronson, 1998; Currie and Yelowitz, 2000; Chetty et al., 2016; Nakamura et al., Forthcoming; Chetty and Hendren, 2018; Chyn, 2018; Deutscher, 2020; Laliberté, 2021). Motivated by such findings, policymakers often provide housing assistance to improve the economic opportunity of children from low-income households. The impact of household assistance, however, is likely to vary with housing, neighborhood, and school characteristics (van Dijk, 2019).

Evidence on the efficacy of housing assistance has generally been mixed. For instance, evidence from the United States – which is usually delivered in the form of housing vouchers – has found both positive and null effects on child outcomes (see Collinson et al. (2015) for a review). In the developing world, researchers have found minimal benefits from public housing receipt (Barnhardt et al., 2017; Picarelli, 2019; Franklin, 2019). Key to discovering the source of heterogeneity is understanding how the impacts of public housing vary with the characteristics of the local neighborhoods and schools. For example, changes to school quality among recipients could explain why some experiments in the United States have found null effects for adult economic outcomes (Katz et al., 2001; Ludwig et al., 2013) but positive effects among children (Chetty et al., 2016; Chyn, 2018).

This paper leverages public housing lotteries in Colombia to estimate the impact of public housing assistance on child outcomes. It does so in the context of Colombia's "Free Housing Program," which provides high-quality housing to over 100,000 highly-disadvantaged families. The program is unique in two dimensions: (i) the public housing units were built in desirable areas near city centers, (ii) the housing unit was given to recipients for *free*. As the housing was oversubscribed, thirty percent of units were randomly assigned via lottery. We leverage these lotteries to show the causal impact of public housing receipt on children's schooling

outcomes.

We find that receiving free public housing increases high school graduation rates by six percentage points, a staggering *fifteen* percent increase relative to the control mean of forty-two percent. We also find that public housing receipt increases the probability of taking and raises test scores for the ICFES, Colombia's high school exit exam which is used for university admissions. Given that our education data only cover up to 2019 and the lotteries were conducted in 2012-14, these effects are for children who resided in public housing for 4-6 years before graduating.

To investigate mechanisms underlying our results, we next use data on the universe of public school students before the lotteries occurred to estimate the quality of schools using value-added methods. We then link the children of lottery winners and losers to the schools that they attended before and after lottery receipt so that we can calculate the proportion of the increase in graduation rates coming from changes to school quality (results coming soon).

This paper connects with a growing literature that examines the impact of public housing on child outcomes. In the United States, the literature has focused on the impact of housing vouchers on child outcomes, finding mixed results. Jacob et al. (2015) take advantage of a randomized housing voucher lottery in Chicago and find little impact of housing assistance on a wide variety of child outcomes. Similarly, Jacob (2004) does not find any effect of housing assistance in the form of vouchers for students affected by high-rise public housing demolitions in Chicago. In contrast, Schwartz et al. (2020) find that housing vouchers in New York City raise students' test score performance. Similarly, Chyn (2018) finds that children affected by public housing demolitions who were given vouchers to move to less disadvantaged neighborhoods were more likely to complete high school, be employed, obtain higher earnings, and commit fewer violent crimes. Chetty et al. (2016) come to similar conclusions in their analysis of the Moving to Opportunity experiment, finding that young children (below age 13) who moved to better neighborhoods had higher

levels of college attendance and earnings.[1] Similar to our setting, children affected by housing vouchers were affected in multiple dimensions, often including increased wealth as well as improvements in housing, neighborhood, and school quality.

The impact of public housing from developing countries has been more negative, with researchers finding small to null effects. Barnhardt et al. (2017) examine the impact of a randomized housing lottery in India that provided winning slum dwellers the opportunity to move into improved housing and find no improvements in family income or human capital. Similarly, Franklin (2019) exploits a lottery for government housing in Ethiopia and finds no impact on earnings. Picarelli (2019) also finds little impact of a housing relocation program in South Africa in terms of household earnings. All these papers, however, study government provided housing that was located far from city centers to save costs.

Our contribution is to twofold. First, we provide new evidence of the benefits of public housing on children's educational outcomes. Our results are of particular interest given the generosity of the public housing assistance program we study. In particular, the units were high-quality, located in desirable areas, and provided for *free*. Second, our unique data allow us to investigate possible mechanisms underlying the improvements in child outcomes, namely from changes in school quality. In contrast to the decidedly mixed results in the literature, we find staggering improvements in the educational outcomes of recipients in our context. Our results highlight the importance of the location and quality of public housing to help raise the economic opportunities of recipients.

The rest of the paper is organized as follows: The next section describes the "Free Housing" program. Section 3 then sets out our empirical methodology and introduces our data. These are followed by our results in Section 4 with mechanisms underlying these results being discussed in Section 5. Section 6 concludes.

[1] In contrast, earlier Moving to Opportunity analyses found little impact of improved neighborhood quality (Katz et al., 2001; Kling et al., 2007; Ludwig et al., 2013). Chetty et al. (2016) suggest that this was driven by negative treatment effects among older (above age 13) children. One potential explanation for these differential effects is that moving to a very different environment as an adolescent could disrupt social networks and have other adverse effects on child development.

2 Background

On April 23, 2012 President Manuel Santos introduced Law 1537, establishing the *Programa Vivienda Gratuita* or "Free Housing" program which provided a *free* residence for the disadvantaged. The law was in line with the government of Colombia's long-standing support home ownership[2] and received broad political support with Congress quickly passing the legislation. The program was ambitious in scope, aiming to build and deliver 100,000 homes to the disadvantaged for *free* within two years.[3]

To build the necessary housing units, the federal government allocated 4 trillion pesos (roughly 2.2 billion USD using 2012 exchange rates). Given the amount of money allocated and the number of housing units required, a limit for construction costs of 40 million pesos (roughly $22,000 USD in 2012) per unit was set.[4] The government then opened up a call for mayors and governors to identify properties for the new housing units, setting an application deadline of July 3, 2012. The properties had to meet certain criteria set out by the government, such as: nearby availability of public services, have the necessary zoning and construction permits, be on 'urban' land, and not be located in 'risky' areas. These criteria were set to avoid endemic problems in Colombia's previous public housing programs whereby subsidized housing was located in peripheral land that lacked public services or in regions with high flood risks.[5]

A total of 650 properties were put forward for consideration of which 298 were deemed suitable for development. Private builders then submitted bids with a point

[2]See Gilbert (2014) for a detailed description of public housing programs in Colombia and the political context of the program's introduction. We rely on Departamento Nacional de Planeación (2014) for the technical details of the program.

[3]After the construction of the initial 100,000 homes, a second phase of the program commenced in 2016 aiming to build an additional 30,000 housing units (25,000 completed to date) for 1.9 trillion pesos, with a particular focus on small municipalities. Our analysis focuses on the first phase of the program.

[4]Even though construction costs are higher in bigger cities, this limit did not vary across the country. Given this, smaller municipalities generally constructed larger housing units in terms of square footage.

[5]For example, one-fifth of Colombia's subsidized housing in 2011 was found to be on land highly-susceptible to flooding (Gilbert, 2014).

system determining the winning bid, with bids being evaluated on: services provided, development layout, and the size and quality of the homes. Over one hundred companies obtained contracts, although over half of the housing units were built by ten companies which included the three largest construction companies in the country.

Project Locations and Quality: Figure 1a displays the locations of the development projects across the country built by the end of 2014, with the size of the pin indicating the number of housing units in the project. In the end, 225 developments were built across 191 municipalities between 2012-14, which created a total of 66,242 housing units.[6] The location of these developments mirror Colombia's spatial population distribution where most of the population lives between the Cordillera Occidental and Cordillera Oriental mountain ranges (i.e., the triangle formed by the major cities of Bogota, Cali and Medellín) or on the Caribbean coast near Cartagena.

The housing units usually involved two-bedroom units in apartment blocks or single-story row houses. In smaller towns, these developments would number a few hundred housing units, while in larger cities the developments would be full-sized neighborhoods or apartment complexes with 3,000 to 15,000 people. These developments were also prioritized for social infrastructure through an agreement with various ministries. For example, the Ministry of Technology provided internet connection points, the Department of Sport built sport fields, the Ministry of the Interior installed security cameras, and the Ministry of Culture provided 12 books for each housing unit. The only stipulation for recipients was that they could not sell or rent the house for a period of ten years after receiving the deed.

Overall, these housing projects represented a substantial improvement in terms of both physical structure and location compared to recipient's prior residences. As an example, Figure 2 displays the pre-lottery housing for an applicant compared to the government provided housing units that the applicant eventually received. The

[6] A further 70 developments that contained roughly 35,000 housing units were completed in 2015. Given the sample restrictions we make (see Section 3.2), we only include pre-2015 development in our sample and so we focus on the 225 developments constructed before 2015.

photos make clear the poor housing conditions that the household faced before the lottery and the substantial improvement the housing units from the "Free Housing" program represented.

In general, the housing projects were well-built across the country. The government put in place several quality controls, including that the units had to pass inspections before builders were paid for their work. In addition, given that the properties had to meet several criteria in terms of proximity to public services, most of the projects were located in desirable areas with many amenities: 75% of projects are located near main avenues, 76% are located near a park, and 80% are near a school. For example, a major free housing project in Bogota, *La Hoya*, is located directly next to a station on the *TransMilenio* – the city's key public transportation system – and is only 20 minutes away from Bolivar Square in central Bogota via public transit. In contrast, the majority of recipients previously lived in "*comunas*," which are often located in the hilly suburban regions of major Colombian cities. Houses in these *comunas* are poorly-built and the neighborhoods themselves feature high crime rates and are located far from city centers (i.e., roughly equivalent to the notorious *favelas* in Brazil).

Program Eligibility: Three groups of individuals were eligible for the program: (i) victims of natural disasters, (ii) internally displaced persons (usually due to armed conflict), and (iii) the 'extreme poor.' The vast majority of recipients came from the latter two groups: 47% were the 'extreme poor' and 45% were internally displaced persons. These groups were then further subdivided into five or six priority tiers based on need. The three eligibility groups were not mutually exclusive as individuals could belong to the 'extreme poor' and either be victims of natural disaster or internally displaced.[7]

Identification of beneficiaries and their priority tier was conducted across several government agencies which identified 250,000 potential beneficiaries. The Ministry of Housing then constructed project-specific lists of beneficiaries as only current

[7]The natural disaster and internally displaced groups were, however, mutually exclusive.

residents of the municipality were eligible for a given project. Based on this list, the Ministry of Housing then opened a calls for applications of potential beneficiaries when each project neared completion and instructed an organization called the Family Compensation Funds to contact each household on the list to apply. Potential beneficiaries were usually notified of their eligibility via a phone call, although they could also have been informed about the program via radio, television, newspaper, billboards, and informational campaigns in their local communities. Applications for each project could also be made by households not on the potential beneficiary list, with auditors then determining their eligibility for the program. Given the use of federally-determined beneficiary lists, the selection process was free of political interference, making many local politicians hoping to use the program to curry favor with voters unhappy (Gilbert, 2014). That said, fraud in the program inevitably occurred with some beneficiaries who received houses being subsequently found ineligible and evicted.[8]

Assignment of Beneficiaries to Houses: As each project was nearing completion, the project's housing units were assigned to one of the three beneficiary groups. Decisions for the exact distribution of units across the three groups were made jointly by the Ministry of Housing and the mayor of the municipality. The decision-makers generally aimed to match the distribution of units to the distribution of beneficiaries in that municipality, although favored internally displaced persons due to the government's focus on reparations for victims of long-standing conflicts.

Once the supply of units for each beneficiary group was set, the assignment of units among each group was conducted according to priority tier until all units had been assigned. If there were more applicants than units, a lottery would be held to determine the recipients. Approximately three-quarters of recipients were directly assigned to housing, while one-quarter of recipients were assigned via lottery.

[8]For example, 13 of the 91 beneficiaries of the first public housing project to open in La Pradera, Valle were later found to be ineligible and were evicted.

We clarify the assignment mechanism with an illustrative example of a housing project with 100 units designated for the 'extreme poor' (other units in the project would be divided among the other two groups). Suppose that 200 'extreme poor' apply for housing, with the applicants evenly divided among five priority tiers. Then, all eighty individuals belonging to the first two priority tiers receive housing, while the eighty individuals in the last two priority tiers do not. Among the third priority tier, however, there are forty applicants for the twenty remaining housing units. Housing for these individuals was then assigned via lottery.

The lotteries were run by the Department of Social Prosperity. To ensure fairness, the draws were publicized via radio and local press with potential beneficiaries invited to attend the draw. The draw was then conducted at a suitable site (e.g., soccer stadium), with chairs and water provided for attendees. By law, the draw had to be attended by several public officials (or their designees): i) the Governor of the department,[9] ii) the Mayor of the municipality, iii) the Director of Social Prosperity, iv) the Executive Director of the National Housing Fund, and v) the Municipal Representative (the Colombian equivalent of an ombudsman).

After the lottery to determine recipients, another draw was conducted to assign recipients to housing units. To do so, the project's housing units are placed in a physical urn and recipients are invited up one at a time to draw their housing unit. If a recipient does not physically attend the lottery, one of the public officials draws their housing unit at the end of the draw for them. Once assigned to a unit, the recipient is able to inspect the unit and then signs the deed in the presence of a notary. The average time between unit assignment and delivery of the house averages 26.2 days.

[9]There are 32 departments in Colombia, so the Governor of one is roughly equivalent to a state Governor in the United States.

3 Empirical Strategy and Data

We describe our empirical strategy which leverages the public housing lotteries to estimate the intent-to-treat impact of public housing by comparing outcomes of winners and losers. The data sources used for this project are also detailed.

3.1 Empirical Strategy

As public housing for a subset of applicants was assigned by lottery, we can intuitively compare outcomes between those who won the lottery and those that did not to provide an unbiased estimate of the effects of being offered public housing. As we have many lotteries in our data, we include lottery fixed effects to ensure that only winners and losers within the same lottery are compared. Fortunately, each lottery occurring at a housing project is given a unique identifier and so project-by-lottery-identifier groupings uniquely identify lotteries in our data. These combinations roughly correspond to housing project-by-eligibility-group-by-priority-tier fixed effects but help us identify which lottery an applicant participated in if they belong to two eligibility groups.[10] Hereafter, we call these project-by-lottery-identifier groupings 'lottery fixed effects.'

Our analysis is somewhat complicated by some municipalities having multiple projects, implying that applicants could apply to multiple projects. Since each project's lottery is independent, the probability that an applicant wins will rise with the number of applications. Fortunately, our data contain the date of application and so we only use the lottery outcome from each applicant's first application (Ketel et al., 2016).[11] Formally, we estimate the impact of receiving public housing on

[10]Fortunately, the lottery identifiers solve the multiple group issue by clearly stating which lottery they took part in as while applicants could belong to multiple groups, they would be assigned to the group with the highest probability of winning. The project-by-lottery-identifier grouping may also not exactly correspond to grouping by project-by-eligibility-group-by-priority-tier as some projects had several lotteries, most often caused by some units becoming available after some recipients were later deemed ineligible and were evicted.

[11]Alternatively, one could define lottery risk sets as the group of non-degenerate lotteries to which an applicant applied (Abdulkadiroğlu et al., 2011). Unfortunately, while our data include date of application and date of housing receipt, they do not contain date of lottery. Therefore, there are a few cases where we are unsure if the applicant has applied to multiple lotteries simultaneously or applies to the subsequent lottery after losing the first making it difficult to define the risk sets.

child outcomes using the following regression:

$$y_i = \alpha + \beta D_i + \delta X_i + LC_i + \epsilon_i \,, \tag{1}$$

where y_i is the outcome of child i, D_i is a dummy variable equal to one if the child's family won the *first* lottery they applied for and X_i is a vector of controls which includes applicant's age at first lottery along with pre-lottery characteristics (e.g., gender, family wealth, etc.). We also include lottery fixed effects for the *first* lottery that child i's family applied for, LC_i (i.e., project-by-lottery-identifier fixed effects), which ensures that, conditional on lottery fixed effects, the probability of receiving housing is identical among individuals. Our parameter of interest is β which is the impact of winning the lottery on child outcome y. Compliance with the admission lottery is ninety-three percent, so the effect of winning the lottery can effectively be interpreted as the impact of receiving public housing.

3.2 Data

We start with data on the universe of public housing applications. These applications contain information on each application made by an individual along with whether they were determined to be eligible,[12] the groups that they were eligible for, their priority tier, how public housing assignment was determined (i.e., by lottery or directly admitted), the lottery identifier (if applicable), the lottery outcome (if applicable), and the date of housing receipt (for those directly admitted or lottery winners).

The application data contain information on each member of the household, including children. For each household member, we have unique identifiers, including: national ID number, full name, and date of birth. We use these identifiers to merge individuals to our administrative data sources, which we now detail. We focus on the 71,974 households whose public housing receipt was determined by lottery.

[12] If the individual was determined to be ineligible, the data stated the reason for ineligibility.

SISBEN III: The SISBEN or the "Census of the Poor" is a census of Colombia's poor which aims to capture the wealth of individuals for means tested social programs, such as free health care and conditional cash transfers. The data are collected door-to-door by surveyors and include rich demographic and socioeconomic information of all household members including sex, age, exact date of birth, education, marital status, occupation, income, household size, dwelling characteristics, and indicators of household wealth (e.g., has a fridge). We use the third wave of the SISBEN or "SISBEN III" which was conducted in 2009-10, a few years before the first housing lottery. The SISBEN III covers roughly 28.5 million people, corresponding to about 62 percent of the population.

We match our public housing application data to the SISBEN III using national ID, name, and date of birth. We were able to match 100 percent of our lottery sample to the SISBEN III, allowing us to examine baseline characteristics of the lottery participants (see Table 1) and control for several pre-lottery covariates in our empirical models. We also use the fourth wave of the SISBEN or "SISBEN IV" which was conducted in 2019-20 to compare the post-lottery housing characteristics of lottery winners and losers. Unfortunately, the data collection for the SISBEN IV was interrupted by COVID so we are only able to match about 45 percent of our sample.

Universe of Students in Colombia's Public Schools: The second administrative data source we use is the core database of the Ministry of Education, which provides information on school progression for all students in public schools.[13] In particular, the R-166 allows us to observe the first year that a child entered the school system (e.g., first grade) up to high school graduation (or dropout) for everyone who was ever enrolled in the public school system. The data provide key educational outcomes that capture a child's progression through the school system (although it does not contain information on test scores), as well as the specific

[13]While Colombia has a vibrant private school sector with approximately a twenty percent market share, over ninety-three percent of children from families that are eligible for public housing attend public schools.

school that a child attends each year. In addition, the data indicates whether a student has received a high school diploma. We use data up to 2019, the last year available.[14]

End-of-High School Exam (ICFES): The Icfes is the national high school exit exam administered by the *Instituto Colombiano para el Fomento de la Educación Superior* (Icfes). The exam is mandatory for all high school seniors regardless of whether they intend to apply to college, although the ICFES score is also used for admission purposes for those who apply to college. The ICFES includes separate tests on math, Spanish, social studies, sciences, and an elective subject. We aggregate the subject-specific scores into a continuous variable that captures the average score across all individual subjects and standardize these scores to have mean zero and standard deviation one each year. The data is available up to 2019. We use both test-taking and ICFES scores as outcomes in our analysis.

Sample Restrictions: We make two sample restrictions. First, we restrict our data to children who are at least 18 by the end of our data in 2019 to ensure that the child had the opportunity to finish high school. In Colombia, high school ends after eleventh grade when students are usually 17. Restricting our data to 18 and older therefore ensures that these children have reached the age to graduate, allowing for one grade repetition. Second, we restrict our sample to children aged 14 or below at the time of the first lottery application.[15] The restriction is made so that the child has not already dropped out of school at the time of the lottery since the legal dropout age in Colombia is 15. An added benefit of this restriction is that it ensures that children have been in public housing a sufficient time period for effects to appear.

Our analysis sample consists of 10,779 children, of whom 2,868 won the lottery and 7,911 lost the lottery. For the most part these children are from different families, although our data does include 1100 siblings. Figure 1a shows the location of

[14]Colombia's academic calendar mirrors the calendar year.

[15]Combined with the restriction that children must reach the age of 18 by 2019 ensures that children are aged 12 to 14 at the time of the first lottery application.

our sample, with the size of the pin indicating the size of our sample in a given project. Compared to the location of the projects (see Figure 1a), our sample is overrepresented in cities along the Caribbean coast.

Descriptive Statistics: Column (1) of Table 1 shows summary statistics for our sample of children. The average age of a child at first lottery is 13.4 years and about half come from families where the parents are married. Households tend to have an average of 5.8 members. While the program was not targeted to rural households, we do see that about twenty percent of the sample resided in rural areas prior to the lottery. Regarding assets, we see that only three percent of the families own a vehicle, less than half have a fridge, and 11 percent have a washing machine. Ninety-five percent of the sample have access to electricity and only 80 percent have access to water and sewage. (Comparison to eligible and all public housing recipients coming soon).

4 Results

We first discuss the validity of our empirical design based on lotteries and then present the first-stage and reduced-form results. Given the high levels of compliance to the lottery, our results are reported as intent-to-treat estimates. Throughout, standard errors are adjusted for two-way clustering at the municipality and family levels to account for the fact that children face common municipality-level shocks and our data sometimes feature multiple entries per family (Cameron et al., 2012).

4.1 Validity

The validity of the empirical design laid out in Section 3 relies on the fact that the lotteries were indeed random (conditional on lottery fixed effects). Given the publicity surrounding these lotteries and the fact they were well-attended by both public officials and potential recipients (see Section 2), we suspect there is limited scope for cheating. Regardless, we verify that these lotteries appear to be random

by checking for covariate balance.

Table 1 checks for covariate balance among the lottery winners and losers. The second and third columns show treatment and control means of demographics of the child, the household head, as well as characteristics of the house that they lived in and the durable goods the household owned before the lottery. Differences are shown in column (4), with the p-value from a formal test of equality between the lottery winners and losers reported in column (5). Reassuringly, the table shows that there are few statistically significant differences between lottery winners and losers. Only one characteristics ("house has natural gas") is significant at the five percent level, which is to be expected by chance given that we are testing balance for nineteen covariates.

First-Stage: While we expect almost all lottery winners to accept the free public housing given the generosity of the program, lottery losers may still receive public housing since they can apply to another housing project in the same municipality. Table 2 shows the 'first-stage' results of winning the public lottery on receiving public housing both in terms of ever receiving public housing and the number of years the child was in public housing (up to age 18). We find that winning the lottery raises the probability of receiving public housing by ninety-three percent and increases the number of years the child resides in public housing by roughly 4.5 years. Given the high rate of compliance, we report intent-to-treat or 'reduced-form' estimates hereafter as they very similar to treatment-on-the-treated estimates which can be obtained by dividing our estimates by 0.93.

4.2 Results

Table 3 reports our main estimates of the impact of free housing on children's educational outcomes, with column (1) reporting results of equation (1) without any controls (aside from lottery fixed effects) while column (2) includes detailed pre-lottery demographic controls. As expected, the inclusion of controls has little effect on our results and so we treat results from column (2) as our preferred estimates. We

find that the children of lottery winners have substantially improved educational outcomes by age 18 compares to lottery losers. The point estimates reveal that winning the lottery increases high school graduation rates by six percentage points, a staggering *fifteen* percent increase relative to the control mean of forty-two percent. Similarly, public housing receipt increases years of education by 0.43 (compare to the control mean of 9.0).

The second panel of Table 3 reports results from the high school exit exam (ICFES). We find that winning the lottery increases high school exit exam taking by five percentage points, which are similar to our results for high school graduation (to be expected given that the exit exam is required for graduation). In terms of performance on the exam, we find that lottery winners score 0.04 standard deviations higher than lottery losers. While this increase in test scores is not statistically significant, we note that these improvements are occurring even though public housing receipt also increases exam-taking (coming soon: results by ICFES subject).

5 Mechanisms

To investigate mechanisms underlying our results, we next use data on the universe of public school students before the lotteries occurred to estimate the quality of schools using value-added methods. We then link the children of lottery winners and losers to the schools that they attended before and after lottery receipt so that we can calculate the proportion of the increase in graduation rates coming from changes to school quality (results coming soon).

6 Conclusions

This study investigates the effects of Colombia's "Free Housing" program on children's educational achievement. To do so, we leverage housing lotteries and find that receiving free public housing increases high school graduation rates by six percentage points, a *fifteen* percent increase relative to the control mean. Large im-

provements in grade repetition, years of education, and exit exam taking and scores are also seen.

The program that we study is highly-generous, providing housing units for free and locating them in desirable areas. This generosity drives the large effects we find on children's educational achievement. In contrast, much of the literature – especially in developing countries – evaluates public housing projects in less than ideal locations (e.g., far from city centers) which greatly reduce their desirability and ability to generate economic opportunity.

Our results indicate that public housing can improve the economic opportunity of children. Policymakers need to take care, however, to ensure that the public housing provided to these children is high-quality and located in desirable neighborhoods near excellent public schools. Unpacking the exact features of the housing environment alongside their costs can help governments better-locate these housing projects to best improve the economic opportunities of the disadvantaged.

References

Aaronson, Daniel (1998), "Using sibling data to estimate the impact of neighborhoods on children's educational outcomes." *Journal of Human Resources*, 915–946.

Abdulkadiroğlu, Atila, Joshua D. Angrist, Susan M. Dynarski, Thomas J. Kane, and Parag A. Pathak (2011), "Accountability and flexibility in public schools: Evidence from Boston's charters and pilots." *Quarterly Journal of Economics*, 126, 699–748.

Barnhardt, Sharon, Erica Field, and Rohini Pande (2017), "Moving to opportunity or isolation? Network effects of a randomized housing lottery in urban India." *American Economic Journal: Applied Economics*, 9, 1–32.

Cameron, A. Colin, Jonah B. Gelbach, and Douglas L. Miller (2012), "Robust inference with multiway clustering." *Journal of Business & Economic Statistics*.

Chetty, Raj and Nathaniel Hendren (2018), "The impacts of neighborhoods on intergenerational mobility I: Childhood exposure effects." *Quarterly Journal of Economics*, 133, 1107–1162.

Chetty, Raj, Nathaniel Hendren, and Lawrence F. Katz (2016), "The effects of exposure to better neighborhoods on children: New evidence from the Moving to Opportunity experiment." *American Economic Review*, 106, 855–902.

Chyn, Eric (2018), "Moved to opportunity: The long-run effects of public housing demolition on children." *American Economic Review*, 108, 3028–56.

Collinson, Robert, Ingrid Gould Ellen, and Jens Ludwig (2015), *Low-Income Housing Policy*, 59–126. University of Chicago Press, URL http://www.nber.org/chapters/c13485.

Currie, Janet and Aaron Yelowitz (2000), "Are public housing projects good for kids?" *Journal of Public Economics*, 75, 99–124.

Departamento Nacional de Planeación (2014), "Evaluación de operaciones y l inea de base de los programas de subsidio familiar de vivienda y subsidio familiar de vivienda en especie: Volumen 1: Informe final resumen ejecutivo." Technical report.

Deutscher, Nathan (2020), "Place, peers, and the teenage years: Long-run neighborhood effects in Australia." *American Economic Journal: Applied Economics*, 12, 220–49.

Franklin, Simon (2019), "The demand for government housing: Evidence from lotteries for 200,000 homes in Ethiopia.", URL https://www.dropbox.com/s/kbrq60h9l93v7px/Franklin_JMP.pdf?dl=0. Unpublished.

Gilbert, Alan G. (2014), "Free housing for the poor: An effective way to address poverty?" *Habitat International*, 41, 253–261.

Jacob, Brian A. (2004), "Public housing, housing vouchers, and student achievement: Evidence from public housing demolitions in Chicago." *American Economic Review*, 94, 233–258.

Jacob, Brian A., Max Kapustin, and Jens Ludwig (2015), "The impact of housing assistance on child outcomes: Evidence from a randomized housing lottery." *Quarterly Journal of Economics*, 130, 465–506.

Katz, Lawrence F., Jeffrey R. Kling, and Jeffrey B. Liebman (2001), "Moving to opportunity in Boston: Early results of a randomized mobility experiment." *Quarterly Journal of Economics*, 116, 607–654.

Ketel, Nadine, Edwin Leuven, Hessel Oosterbeek, and Bas van der Klaauw (2016), "The returns to medical school: Evidence from admission lotteries." *American Economic Journal: Applied Economics*, 8, 225–54.

Kling, Jeffrey R., Jeffrey B. Liebman, and Lawrence F. Katz (2007), "Experimental analysis of neighborhood effects." *Econometrica*, 75, 83–119.

Laliberté, Jean-William (2021), "Long-term contextual effects in education: Schools and neighborhoods." *American Economic Journal: Economic Policy*, 13, 336–77.

Ludwig, Jens, Greg J. Duncan, Lisa A. Gennetian, Lawrence F. Katz, Ronald C. Kessler, Jeffrey R. Kling, and Lisa Sanbonmatsu (2013), "Long-term neighborhood effects on low-income families: Evidence from Moving to Opportunity." *American Economic Review: Papers & Proceedings*, 103, 226–31.

Map data: Google (2021), "Colombia. Google Maps." URL https://jsfiddle.net/byrupej6/. (accessed 28 March 2021).

Nakamura, Emi, Jósef Sigurdsson, and Jón Steinsson (Forthcoming), "The gift of moving: Intergenerational consequences of a mobility shock." *Review of Economic Studies*.

Picarelli, Nathalie (2019), "There is no free house." *Journal of Urban Economics*, 111, 35–52.

Schwartz, Amy Ellen, Keren Mertens Horn, Ingrid Gould Ellen, and Sarah A. Cordes (2020), "Do housing vouchers improve academic performance? Evidence from New York City." *Journal of Policy Analysis and Management*, 39, 131–158.

van Dijk, Winnie (2019), "The socio-economic consequences of housing assistance.", URL https://drive.google.com/file/d/1nxOHg4uKgVtRiyNjCjMq4ScsJo8uX97F/view. Unpublished.

FIGURE 1: Location of Projects and Estimation Sample

A Location of Housing Projects

B Location of Estimation Sample

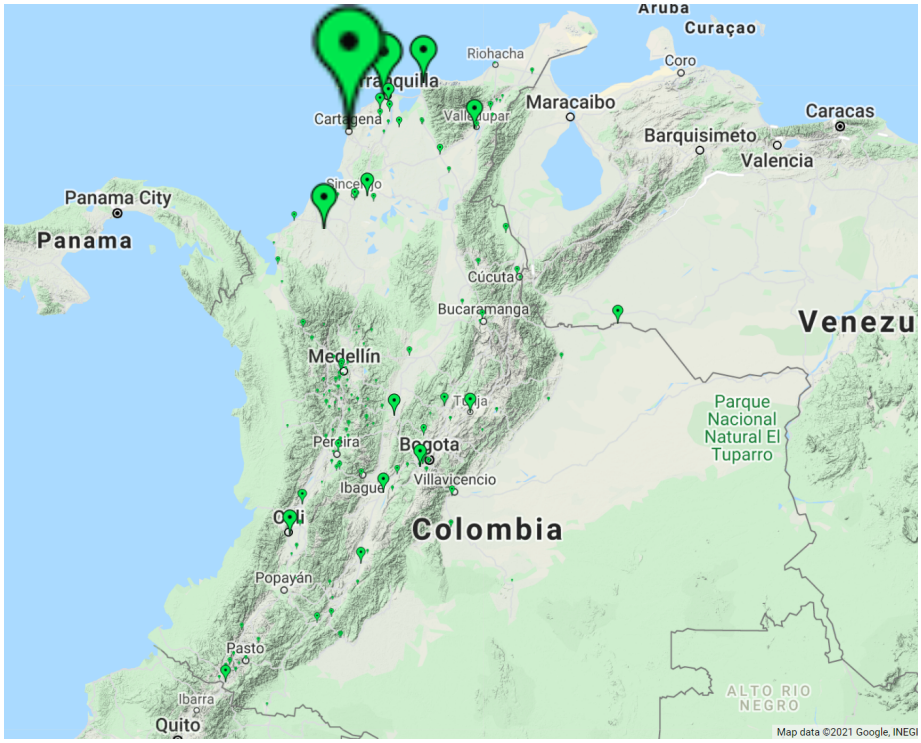

Notes: Map of Colombia (Map data: Google, 2021). Figure 1a displays the location of the 225 projects in our data while Figure 1b shows the location of the 10,779 individuals in our estimation sample. The size of the pin in each figure represents relative size of the project in terms of number of housing units, with a minimum size imposed for projects with few observations to ensure that they are visible.

FIGURE 2: Example of Applicant Housing Relative to Government Provided Housing

A Exampe of Applicant Housing in Lorica

B Government Housing Project in Lorica

Notes: Figure shows an example of pre-lottery housing for an applicant compared to the government provided housing units that the applicant eventually received from the "Free Housing" program. The photos for this example come from the city of Lorica which is located in the department of Córdoba on the Caribbean coast. Figure 2a shows the residence of an applicant for the public housing project *Urbanización La Victoria en Lorica*. Figure 2b then shows housing units in the *Urbanización La Victoria en Lorica* housing project where the applicant moved to after winning the lottery.

TABLE 1. Covariate Balance

	Overall Mean (1)	Treatment (2)	Control (3)	Difference (Treatment-Control) (4)	Test of Equality (p-value) (5)
Child Demographics					
Age at First Lottery	13.41	13.46	13.40	-0.06	0.599
Female	0.49	0.50	0.49	0.01	0.575
Household Size	5.81	5.80	5.82	-0.02	0.158
Lived in Urban Area	0.78	0.75	0.80	-0.05	0.551
Household Head Characteristics					
Age at Birth	27.83	27.99	27.77	0.22	0.511
Married	0.53	0.51	0.53	-0.02	0.889
Employed	0.51	0.50	0.51	-0.01	0.440
High School Graduate	0.74	0.73	0.74	-0.01	0.987
Housing Characteristics					
Number of Rooms	2.76	2.74	2.77	-0.03	0.942
Number of Bathrooms	0.89	0.89	0.89	0.00	0.994
Has Kitchen	0.80	0.82	0.79	0.03	0.448
Access to Services					
Electricity	0.95	0.94	0.96	-0.02	0.821
Water/Sewage	0.80	0.79	0.81	-0.02	0.051
Natural Gas	0.39	0.34	0.41	-0.07	0.040
Trash Collection	0.76	0.71	0.78	-0.07	0.441
Household Wealth					
Has Vehicle	0.03	0.03	0.03	0.00	0.496
Has Fridge	0.43	0.42	0.43	-0.01	0.750
Has Washing Machine	0.11	0.10	0.11	-0.01	0.483
Has TV	0.73	0.72	0.78	-0.06	0.216
# of children	10,779	2,868	7,911	10,779	10,779

Notes: This table reports means and treatment-control differences in pre-lottery characteristics of children that applied for public housing in Columbia between 2012 and 2014 and whose housing assignment was determined via lottery. The pre-lottery characteristics come from the SISBEN III and were collected in 2010. The sample is restricted to children who are age 12-14 at the time of lottery and age 18-21 in 2019. The sample includes one observation per child, with children being assigned to treatment according to their first application. Column (4) reports the p-value on receiving a public housing offer for their first application in a regression controlling for lottery fixed effects, with two-way clustered standard errors the municipality and family level.

TABLE 2. Impact of Winning Housing Lottery on Public Housing Receipt

Sample:	Main Sample (1)	ICFES Sample (2)
Ever Received Public Housing	0.931***	0.929***
	(0.015)	(0.017)
Years in Public Housing (up to age 18)	4.481***	4.357***
	(0.211)	(0.298)
# of Observations	10,779	5,352

Notes: This table reports the effect of winning the public lottery on receiving public housing and so represents 'first-stage.' We report the 'first stage' results both in terms of ever receiving public housing and the number of years of public housing the child experienced up to the age of 18. All regressions include lottery fixed effects to ensure that only individuals in the same lottery are being compared. Of the 10,779 children in the main sample, 2,868 won the lottery and 7,911 lost the lottery. The sample is smaller for the 'ICFES sample' outcome as many children did not take the ICFES as they dropped out. Standard errors are two-way clustered at the municipal and family level. ***,** and * denote significance at the 1%, 5% and 10% levels, respectively.

TABLE 3. Impact of Winning Housing Lottery on Educational Outcomes

Impact of Winning First Housing Lottery on:	No Controls (1)	Demographic Controls (2)	Control Mean (3)	# of Observations (4)
Panel A. Schooling Outcomes				
Years of Education	0.499***	0.427***	9.02	10,779
	(0.080)	(0.074)		
High School Graduation	0.071***	0.055***	0.42	10,779
	(0.020)	(0.019)		
Panel B. High School Exit Exam (ICFES) Outcomes				
Took ICFES	0.071***	0.050***	0.48	10,779
	(0.015)	(0.015)		
ICFES Score	0.050	0.035	-0.34	5,352
	(0.035)	(0.033)		

Notes: This table reports intent-to-treat estimates of the effect of winning the public lottery on schooling outcomes as described by equation (1). All regressions include lottery fixed effects to ensure that only individuals in the same lottery are being compared. Column (1) includes reports results when no controls are included (aside from lottery fixed effects), while column (2) contain controls for a child's gender, age at first lottery, whether a family lived in urban/rural area, household size, along with characteristics of the household head including age at birth, marital status, employment status and education (all measured pre-lottery in 2010). The sample includes children who were 18 or older in 2019 and 14 or below on the date of applications, which restrict our data to children aged 12 to 15 at the time of the first lottery application. Of the 10,779 children in the main sample, 2,868 won the lottery and 7,911 lost the lottery. The sample is smaller for the 'ICFES score' outcome as many children did not take the ICFES as they dropped out. Standard errors are two-way clustered at the municipal and family level. ***,** and * denote significance at the 1%, 5% and 10% levels, respectively.