JOB MARKET PAPER

The socio-economic consequences of housing assistance

Winnie van Dijk*

November 2018

Most recent version at jmp.winnievandijk.com

Abstract

This paper analyzes the effect of Europe's largest public housing program on socio-economic outcomes for low-income households. Using lotteries for housing units in the Netherlands and data linking national registers to application choices, I show that the average move into public housing negatively affects labor market outcomes and proxies for neighborhood quality, and increases public assistance receipt. However, consistent with a model of labor supply responses to conditional in-kind transfers, average impacts miss substantial heterogeneity both across neighborhoods and, within neighborhood, across recipients. Moves into high-income neighborhoods generate positive effects, which are driven by 'upward' moves made by individuals previously living in low- or middle-income neighborhoods. Lateral and 'downward' moves have the opposite effect. To evaluate whether these results generalize to non-recipients, I develop a model of application behavior that utilizes panel data on application choices and exploits variation induced by the housing allocation mechanism. Using the model, I recover the distribution of heterogeneity that drives selection into and returns from lotteries, and estimate that selection on gains is limited. This suggests that targeting public housing in high-income neighborhoods based on observable characteristics can increase economic self-sufficiency.

*Kenneth C. Griffin Department of Economics, University of Chicago, 1126 East 59th Street Chicago, IL 60637 (e-mail: wlvandijk@uchicago.edu).
I gratefully acknowledge financial support from the University of Chicago Social Sciences Division, the University of Chicago Kenneth C. Griffin Department of Economics, and the Becker Friedman Institute. I would like to thank Ali Hortaçsu, Magne Mogstad, and Azeem Shaikh for serving on my thesis committee and providing valuable guidance and feedback. I am grateful for discussions with many people, including Eliot Abrams, Mohammad Akbarpour, Dan Black, Stephane Bonhomme, Eric Budish, Caterina Calsamiglia, Jose Ignacio Cuesta, Angela Denis Pagliero, Manasi Deshpande, Michael Dinerstein, Deniz Dutz, Pieter Gautier, Michael Greenstone, James Heckman, Cody Hochstenbach, John Eric Humphries, Sonia Jaffe, Juanna Joensen, Paymon Khorrami, Scott Kominers, Steven Kromhout, Thomas Lemieux, Jacob Leshno, Steve Levitt, Claudia Macaluso, Yusuke Narita, Derek Neal, Edgar Olsen, Hessel Oosterbeek, Hyunmin Park, Canice Prendergast, Bastian Ravesteijn, Ana Reynoso, Beth Shinn, Mark Shroder, Daniel Tannenbaum, Pietro Tebaldi, Alex Teytelboym, Neil Thakral, Bas van der Klaauw, Alessandra Voena, Chris Walters, Anthony Lee Zhang, and participants of the Applied Micro Lunch and IO lunch at the University of Chicago. Ella Deeken provided excellent research assistance. I wish to thank Ivo Gorissen, Petra Slingerland-Roos, and Cees Steijn, and Michael Vermaesen at Statistics Netherlands for their advice and technical support; Jeroen van der Veer at Amsterdamse Federatie van Woningcorporaties and Alexandra de Jong at WoningNet for their help in collecting data on public housing in Amsterdam; Vera Berkers, Guus Crommelin, Kees Dignum, Loes Olgers, and Sanne van der Lelij at the Housing Department of the Amsterdam municipality for helpful feedback; and Ruben Henderiks, Jolanda Hoogeboom, Thamar Hoekema, Roelf-Jan van Til, and Jacobien Vogelzang for discussions on housing allocation. Any errors are my own.

1 Introduction

Housing assistance policies targeted at low-income households are used in most high-income countries. Examples include public housing programs, rent vouchers, and subsidized construction of below-market-rate housing. In Europe, large-scale subsidized housing programs with generous eligibility requirements are commonplace.[1] In the United States, the federal government spends around $45 billion on housing assistance annually.[2] In addition, local governments in many cities have recently responded to rising rent prices by announcing plans to increase the supply of subsidized housing for low-income households.[3] These efforts are accompanied by a reinvigorated policy debate about the potential of housing assistance as a policy instrument for stimulating upward mobility and decreasing income inequality. Yet, robust empirical evidence on the efficacy of housing assistance programs in improving recipients' economic outcomes is scarce. The few existing studies generally find zero or negative impacts on labor market outcomes. However, these studies are restricted to programs that did not substantially change housing consumption or neighborhood quality, or programs serving narrow populations with limited external validity.

This paper analyzes the effect of Europe's largest public housing program on households' socio-economic outcomes, including labor market outcomes, neighborhood and housing quality, and public assistance receipt. To the best of my knowledge, this paper is the first study of a large-scale public housing program that uses well-defined exogenous variation to address self-selection into take-up.[4] I exploit three advantages of the Dutch setting. First, unit-specific lotteries provide exogenous variation to address self-selection into take-up. Second, substantial heterogeneity in public housing units allows me to study how impacts vary with housing and neighborhood characteristics. Finally, the allocation of public housing is operationalized by means of an online platform that potential recipients use to submit applications to unit-specific lotteries.

[1] E.g., 34% of Dutch households, 26% of Austrian households, and 19% of French households live in subsidized housing, often referred to as 'social housing.' I use the term 'public housing' to refer to subsidies that are tied to specific housing units (as contrasted with rent vouchers which are tied to recipients and are portable across units). See del Pero et al. (2016) for a survey of housing assistance policies in OECD countries.

[2] See U.S. House of Representatives Committee on Ways and Means (2016).

[3] E.g., in 2016, Mayor Bill de Blasio announced 'Housing New York 2.0,' a plan that aims to build or preserve 300,000 affordable units by 2026 (City of New York, 2016). Similarly, in 2018, Mayor Sadiq Khan of London announced the 'London Housing Strategy,' which includes a commitment to invest £4.82bn ($6.15bn) in the construction of 116,000 affordable housing units by 2022 (Mayor of London, 2018).

[4] Addressing selection on unobservables is particularly important when analyzing the effect of receiving housing assistance, because decisions about residential moves often coincide with major life events (Desmond, 2012; Humphries et al., 2018), and because researchers are often not able to observe important confounding factors, for example private information about job prospects or access to financial support from family members.

The platform generates detailed panel data on choice behavior, which allows me to jointly analyze ex post impacts of the policy and ex ante search effort and take-up behavior.

The analysis in this paper is organized into two parts. In the first part, I develop an empirical framework for analyzing responses to public housing receipt, relying on randomization generated by lotteries. This framework addresses the issue that lottery offers are a function of the take-up status of other applicants by replacing the frequently-used assumption of random assignment of offers with the weaker assumption that the order in which applicants are given offers is random.[5] Under this assumption, I show that the IV estimator applied to data from a given lottery is nearly unbiased for the effect of treatment on those who receive assistance ('compliers') when the lottery has a large number of applicants. By averaging across lotteries, it is therefore possible to estimate the average effect on compliers across lotteries. This parameter corresponds to the expected impact of a policy that marginally expands the program by adding an additional 'representative' lottery. An attractive feature of the framework is that it naturally extends to the analysis of the impact of housing receipt in a specific neighborhood simply by restricting attention to lotteries in that neighborhood. Since the economic interpretation carries over to neighborhood-specific impacts, this allows for a direct comparison of impacts from marginal expansion in one neighborhood versus another.

The lottery-based estimates have high internal validity, but do not address self-selection into lotteries, limiting the policy relevance to marginal reforms. In the second part of the paper, I develop a dynamic discrete choice model of application behavior and leverage the panel data on choice histories to identify unobserved factors driving selection into and returns from lotteries. An additional source of variation is provided by random fluctuations in available choices over time. By combining estimates of ex post impacts with unobservables recovered from ex ante application behavior, I correct for self-selection of applicants into lotteries.

The results from the analysis in the first part of the paper show that the average effect of treatment on those who receive public housing is negative and substantial in size for employment and individual and household earnings, while it is positive for public assistance receipt. These findings are in line with existing work on U.S. housing vouchers and public housing (Jacob and Ludwig, 2012). In addition, average neighborhood income and average house value decrease. However, the average impact across lotteries masks substantial heterogeneity, both across neighborhoods and, within neighborhood, across recipients.

[5]The issue has previously been discussed in de Chaisemartin and Behaghel (2017) and is discussed in more detail in Section 3.

In particular, I find that individual and household income increase for recipients who move to high-income neighborhoods, but decrease for recipients who move into low-income neighborhoods. Conditioning on baseline neighborhood income shows that the positive impacts in high-income neighborhoods are driven by impacts on applicants for whom the move represents an upward move in terms of neighborhood income. In contrast, lateral and downward moves have negative effects for public housing located in all neighborhoods.

Building on labor-supply models used in the public finance literature on in-kind transfers, I lay out a simple model that rationalizes these results. Specifically, a model in which an agent maximizes utility over housing, consumption, and leisure predicts that receiving public housing could increase or decrease labor supply depending on the type and location of the housing, even in the absence of social spillovers, local labor market frictions, credit constraints, or spatial variation in the price and availability of amenities. In the absence of such mechanisms, positive labor market responses can be observed after an increase in the share of the budget spent on housing, if housing and leisure are not near-perfect complements, and the three goods are normal.[6] This model helps motivate why we may expect heterogeneous labor-market responses to public housing as well as why labor market responses have been found to be negative in past work.

In the second part of of the paper I use the dynamic discrete choice model to estimate average treatment effects by the type of house being offered and by sub-population. Specifically, I leverage panel data on individuals' lottery choices combined with variation in the available lotteries to evaluate whether individuals' select into lotteries based on their potential labor market gains. While I find that latent heterogeneity plays an important role in application behavior and earnings, there is little evidence of selection on gains based on unobservables. Average treatment effects are very similar to the lottery-specific estimates from the first part of the paper. This implies that, on average, housing in high-income neighborhoods would improve the labor market outcomes of poorer individuals. Moreover, the model suggests that individuals unlikely to apply to lotteries in high-income neighborhoods would be likely to accept housing in high-income neighborhoods, if offered.

[6] In a similar vein, theoretical work on labor supply responses to unconditional in-kind transfers predicts that the sign of the effect depends on the extent to which the transfer constrains consumption, and whether it is a complement or a substitute to leisure. Moffitt (2002) and Currie and Gahvari (2008) survey this literature. While the result does not require location-based mechanisms, and is therefore implied under very weak requirements on the recipients decision-making process, the presence of such mechanisms is likely to strengthen the prediction of heterogeneous responses to housing assistance depending on both the neighborhood in which housing is offered, and the characteristics of the recipient.

Policy implications. My results suggest two take-aways for policy design. First, the negative sign for the impact of housing assistance across all lotteries suggests that marginally expanding the Amsterdam program by adding more lotteries would tend to produce negative labor market impacts for new recipients. However, marginally expanding the program in high-income neighborhoods increases earnings and employment among those who apply. Beyond Amsterdam, these results suggest that it is possible to design in-kind transfers in such a way that they increase earnings and stimulate labor force participation the short run, something that cash transfers are unlikely achieve. This finding is out of step with the general consensus on how public housing affects labor supply, as well as with most empirical evidence for other types of in-kind transfers (e.g., food stamps, subsidized health care, or job training programs – see Currie and Gahvari, 2008, and references therein, for an overview).

Second, the heterogeneity in responses across recipients suggests that policies targeting high-quality housing to high-return individuals could result in large earnings and employment gains. This fact takes on additional importance because housing assistance programs are almost never entitlements; most programs operate under fixed budgets which allow only a fraction of the income-eligible population to be served.[7] The results in this paper also suggest specific observable characteristics on which to target. In particular, individuals living in low- and middle-income neighborhoods benefit from being offered housing in high-income neighborhoods.

Contribution to the literature. The findings in this paper contribute to several strands of the economics literature. First and primarily, they contribute to a body of research that studies the effects of housing assistance on socio-economic outcomes for low-income households. The existing experimental evidence in this area is based on a small number of experiments, and existing studies are limited in terms of the study population, in the type of housing assistance provided, or both. Several studies evaluate programs that provide low-quality housing (e.g., U.S. public housing studied in Jacob et al., 2015), programs that lead to individuals relocating far from economic activity in city centers (e.g., the housing program studied in Barnhardt et al., 2016), or assistance that did not demonstrably change housing or neighborhood quality (e.g., Jacob and Ludwig, 2012, who find no evidence of changes in neighborhood characteristics as a result of voucher receipt). A different set of studies evaluates assistance provided to narrow and highly pre-selected populations. For example, the population for the Moving to Opportunity

[7] In the U.S., for example, only one in four income-eligible households receives a form of housing assistance (JCHS, 2015), with a complex web of eligibility rules and rationing systems determining who receives assistance.

experiment consisted of families living in public housing, the vast majority of which were single mothers (Kling et al., 2004). Similarly, the Welfare to Work Voucher experiment focused on those already on welfare (Mills et al., 2006), and the Family Options study focused on homeless individuals (Gubits et al., 2016). Therefore, while these experiments each have high internal validity, they potentially have limited external validity. These studies generally find zero or negative effects of housing assistance on labor market outcomes.[8] [9]

This paper considers a new quasi-experimental setting: lotteries conducted as part of a centralized housing allocation system in the Amsterdam metropolitan area. The lotteries allocate housing across a wide spectrum of housing and neighborhood quality, and they are conducted under generous eligibility requirements. Combined with rich information on applicants' background characteristics and choice behavior, these features strengthen external validity of the findings and allow for detailed study of effect heterogeneity. My results confirm the idea that heterogeneity across types of assistance, as well as heterogeneity in responses to the same type of treatment across recipients, are indeed very important in assessing and understanding the efficacy of housing assistance programs.

Second, this paper is related to the literature on in-kind transfers and labor supply. Theoretical work on labor supply responses to in-kind transfers predicts that the sign of the effect depends on the extent to which the transfer constrains consumption, and whether it is a complement or a substitute to leisure (Murray, 1980; Leonesio, 1988; Gahvari, 1994; Shroder, 2002). The model in Section 2.2 builds on this literature, which is surveyed in Moffitt (2002) and Currie and Gahvari (2008). The latter note that, across all types of in-kind transfers, short-run effects on labor supply are often negative. My results suggest that it is possible to design in-kind transfers in such a way that they stimulate labor supply in the short run, something that theory predicts cash transfers are unlikely to achieve.

Third, in a broad sense, this paper is related to a literature that uses lotteries or cutoffs embedded in rationing mechanisms for publicly-provided goods to evaluate the impact of public programs.[10] It contributes to this literature by developing a framework that addresses endogeneity of the offer process,

[8] While the evidence for adult outcomes is inconclusive or negative, there is empirical evidence of positive long-run effects on the labor-market outcomes of children of recipients of public housing (Currie and Yelowitz, 2000) or vouchers (Chetty et al., 2016; Chyn, 2016).

[9] See Collinson et al. (2015) and references therein for an overview of empirical research on U.S. housing assistance policies.

[10] The primary example in the literature is allocation of seats in education settings, e.g., Hoxby and Murarka (2009); Angrist et al. (2010); Abdulkadiroglu et al. (2011); Deming (2011a); Dobbie and Fryer Jr. (2011); Hastings et al. (2012); Dobbie and Fryer Jr. (2013); Deming (2011b); Deming et al. (2014); Dobbie and Fryer Jr. (2015); Davis and Heller (2017); Cohodes (2016); Abdulkadiroglu et al. (2016a); Kirkeboen et al. (2016); Abdulkadiroglu et al. (2016b); Angrist et al. (Forthcoming); Walters (2017). See also Chabrier et al. (2016) and references therein.

which is typically assumed to be exogenous – an assumption that is violated when offers are made until a pre-determined number of slots is filled (as is typically the case with publicly-provided goods and services). The ideas discussed in the development of this framework are potentially relevant for a class of important allocation mechanisms, including mechanisms for allocating seats in education settings, allocation of donor organs, slots in job training programs, or slots in medical trials.

Fourth, this paper is methodologically related to studies that use external 'measurement systems' or observed choice behavior to directly model latent heterogeneity that drives correlations between choice and outcomes.[11] Within this collection of work, a closely related paper is Walters (2017), which, like this paper, uses a model of application behavior and a control function approach (Heckman, 1979) to strengthen the external validity of lottery-based estimates of treatment on the treated.

Organization. Section 2 of this paper describes Amsterdam's institutional environment and provides details on the housing lotteries and their implementation using an online platform. This section then discusses the expected impact of moving into public housing and the expected patterns of selection into lotteries. Section 3 lays out the research design and its implementation. Section 4 presents main results and documents how these results vary across neighborhoods and applicants. Section 5 develops a choice model which is used to construct average treatment effects and policy counterfactuals, and reports estimation results. Section 6 concludes.

2 The housing program and its expected impacts

This section provides details on the institutional environment of the public housing program and the details of the allocation system that are relevant for the research design. I then turn to a discussion of theory-based predictions on how the impact of take-up on consumption and labor-supply decisions will vary, depending on both the housing unit offered and the recipient's characteristics.

2.1 Public housing in Amsterdam

The Amsterdam housing market. The Netherlands has the largest public housing program in the European Union, providing housing to about a third of the population. In Amsterdam, 46% of the

[11] E.g., Carneiro et al. (2003); Heckman et al. (2006); Cunha and Heckman (2008); Heckman et al. (2016).

housing market consists of public housing. The type of housing ranges from highrise apartment buildings located far from city centers to apartments in high-cost urban areas.[12] Moreover, eligibility requirements for the program are generous: income eligibility cutoffs are set at approximately household size-adjusted median income. As a result, the program serves both households for whom the main source of income is welfare receipt, and households with labor income in the lower half of the income distribution. In terms of its target population, the program is therefore comparable to beneficiaries of affordable housing initiatives in U.S. cities, which aim to serve both the very poor and low- to middle-income families. It is also comparable to the large-scale public housing programs in Sweden, Denmark, Austria, France, and the United Kingdom (OECD, 2016).

Turnover of public housing in the Amsterdam metropolitan area is very low, and the wait list system that is used to allocate the vast majority of public housing units has resulted in expected wait times (conditional on receipt) upwards of 12 years. The local government has responded to this situation by administering housing lotteries for a small fraction of newly vacant units, under the rationale that eligible households that have not been on the wait list for several years should also have access to housing assistance. Applicants to these lotteries have an equal chance of receiving an offer of housing, regardless of how much time has passed since their enrollment in the wait list. This paper relies exclusively on such lotteries – the details of which are formalized in section 3.1 – to motivate the research design for measuring causal effects of housing receipt.

Allocation mechanism and operationalization using online platform. Housing units become available when an occupant decides to move out and when units are newly constructed. After a unit becomes available, it is listed on an online platform, along with pictures of the unit and details including size, number of rooms, address, floor, eligibility restrictions (if any), and monthly rental price. Appendix A.1 contains a screenshot of the online platform's user interface. Potential tenants can submit up to two electronic applications per application period, and can only do so after signing in to the platform with a personal account. After the application period ends, the unit is sequentially offered to applicants in a random order until someone accepts. Applicants who do not receive an offer and those who reject the unit face no restrictions to participation in future application periods.

[12]This dispersion over a large and heterogeneous set of neighborhoods is the result of requirements set by local governments to maintain a minimum share of public housing in every neighborhood.

Data generated by the online platform and linkage to administrative data. The analysis in this paper combines confidential data sets originating from two institutional sources: application data provided by the associations that maintain and distribute public housing in the Amsterdam metropolitan region, and administrative data from Statistics Netherlands, the national statistical agency. Statistics Netherlands relies on other government agencies for data collection. For example, address histories are derived from municipal registers, data on individuals' annual earnings is derived from tax returns, and data on public assistance receipt is derived from the agencies responsible for distribution of such assistance. Data linkage experts at Statistics Netherlands performed the linkage between the newly-collected application data and their previously-held data collection. For individuals, the linkage was implemented based on date of birth, gender, and address. To link housing units to administrative data on house and neighborhood characteristics, a linkage was performed based on address. This process resulted in a high quality linked data set where the match rate for individuals was about 96%. For units, it was about 99%. Further details on data sources, sample selection, and linkage are provided in the data appendix (Appendix C.1). For the analysis sample, I follow Jacob and Ludwig (2012), and include only working-age, able-bodied adults.

Conceptualizing treatment and counterfactual. In order to interpret the empirical results in this paper and connect them to predictions from a simple economic model, it is important to understand what changes when an applicant receives an offer of housing, as well as the counterfactual situation, since this is what the treatment effect will be measured against.

The components of housing assistance allocated in a given lottery include the right to rent a specific housing unit at a price below the market rate, plus an income-dependent subsidy from the national government. The monthly rental price of the unit is partly pegged to market rent prices, up to a rent ceiling.[13] For this reason, subsidized units in more areas with higher market rents also have higher controlled rents. The subsidy is available to all renters, including those in the private market, and the majority of applicants receive it before being allocated public housing. Therefore, I consider the right to rent the unit at a rate below market to be the dominant component of treatment.

Figure 1 illustrates the lottery process for a given lottery. The probability of receiving an offer in any lottery is very low. For applicants who don't receive an offer in a given lottery and who choose to re-apply, the possibility of receiving an offer in a future lottery is therefore a near-negligible determinant of their

[13]The rent price is a function of the unit's tax-assessed value, which in turn is estimated using recent transactions of similar units in the same area.

Figure 1: The lottery process.

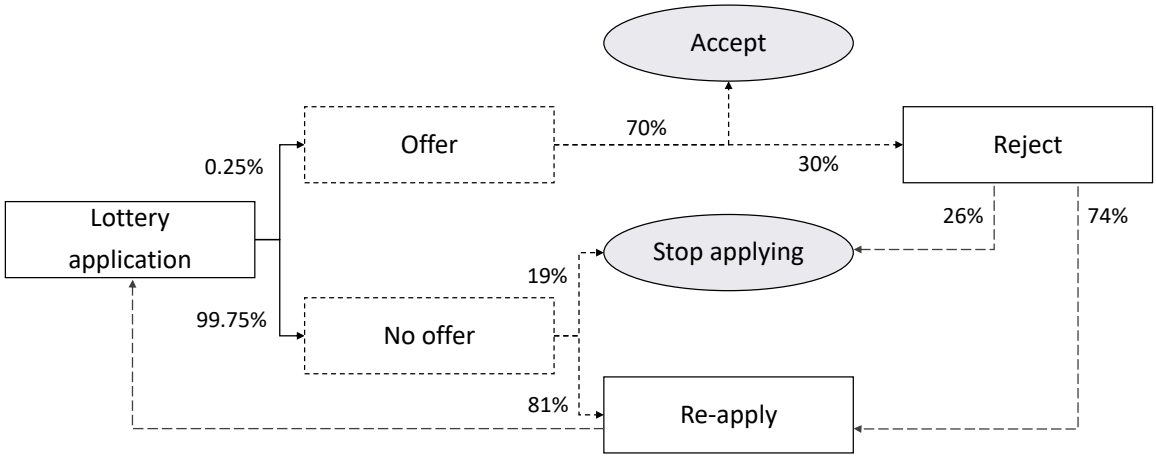

Notes: Figure illustrates the process of applying to housing lotteries. Probabilities are calculated by averaging sample frequencies across lotteries in the baseline analysis sample.

counterfactual outcome. For this reason, we may be inclined to interpret the treatment effects reported in this paper as describing the effect of assistance receipt versus no receipt throughout the period of analysis. However, because individuals who do not receive lottery offers may also obtain public housing through the simultaneously-operated wait list, the results in this paper may still be thought to be subject to some 'substitution bias' through that channel. Table 6 in Appendix A.2 shows that substitution through the wait list is observed for a larger but still relatively small fraction of individuals, mostly due to the fact that the wait list is highly over-subscribed, and most lottery applicants have not reached the point where they can expect to successfully apply to housing units allocated through the wait list. Moreover, there is no evidence of differential substitution rates across lotteries in different neighborhoods, which would complicate a comparison of effects across neighborhoods. I therefore propose that the treatment effects reported in this paper are interpreted as by and large determined by the impact of receiving housing assistance versus not receiving housing assistance, at least for the duration of the time over which outcomes are measured.

2.2 Expected impacts

Expected impacts of housing assistance receipt. As noted above, a distinguishing feature of the Dutch housing program is the fact that the neighborhoods in which public units are located are very heterogeneous. Due to generous eligibility requirements, applicants are also very heterogeneous,

for example in terms of current living situation and in baseline economic outcomes. Combined, these two facts imply that housing lotteries can induce large changes in housing consumption. Moreover, the program is likely to induce both 'upward' and 'downward' transitions. It follows that studying the average transition may obscure important features of the treatment.

Figure 2 illustrates this point. The top panel shows average neighborhood income over time for applicants who did and did not receive an offer. This panel shows that individuals tend to move into poorer neighborhood upon receiving public housing. The second panel adds two additional comparisons by splitting up applicants applying to housing in low-income neighborhoods and applicants applying to housing in high-income neighborhoods. This panel clearly shows that treatment is heterogeneous across lotteries: lotteries for housing in high-income neighborhoods tend to induce 'upward' moves, while lotteries for housing in low-income neighborhoods induce 'downward' moves.

Economic theory predicts that recipients will have different consumption and labor-supply responses to public housing receipt, depending on the type of transition it represents. From the perspective of the recipient, receiving an offer of public housing translates into changes to his economic incentives in two ways. First, by affecting the shape of the budget constraint, public housing receipt can cause changes in consumption and labor supply decisions.[14] Second, public housing receipt induces a residential move, which may come with costs or benefits such as moving costs, changes in commuting costs, changes in access to job opportunities, or social spillovers.

Consider a stylized utility-maximization model where agents solve:

$$\max_{H,C,L} U(H,C,L) \quad \text{s.t.} \quad wT = wL + p_h H + p_c C,$$

where L is leisure, H is housing, C is consumption, and T is the agent's time endowment. Similarly, w is the wage rate, p_h is the price of a unit of housing, and p_c is the price of a unit of other consumption.

[14]This channel is similar to the channel through which in-kind transfers are usually modelled to influence labor supply, but with the distinguishing feature that housing is typically a substantial share of household expenditure. For that reason, housing assistance is likely to alter consumption of housing and other goods post-receipt, in contrast to other forms of in-kind transfers (for example, food stamps).

Figure 2: Neighborhood transitions.

(a) On average, across lotteries.

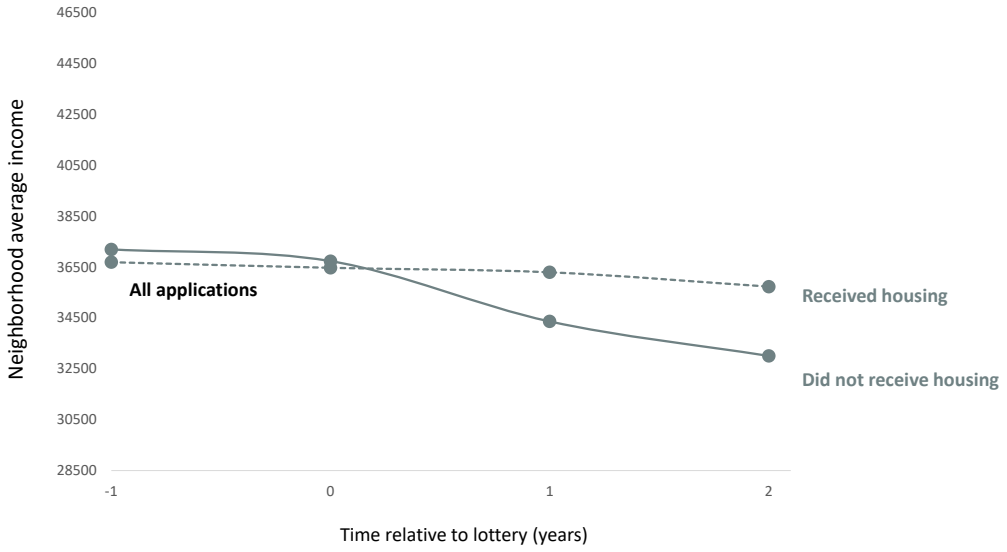

(b) Lotteries in high-income vs. low-income neighborhoods.

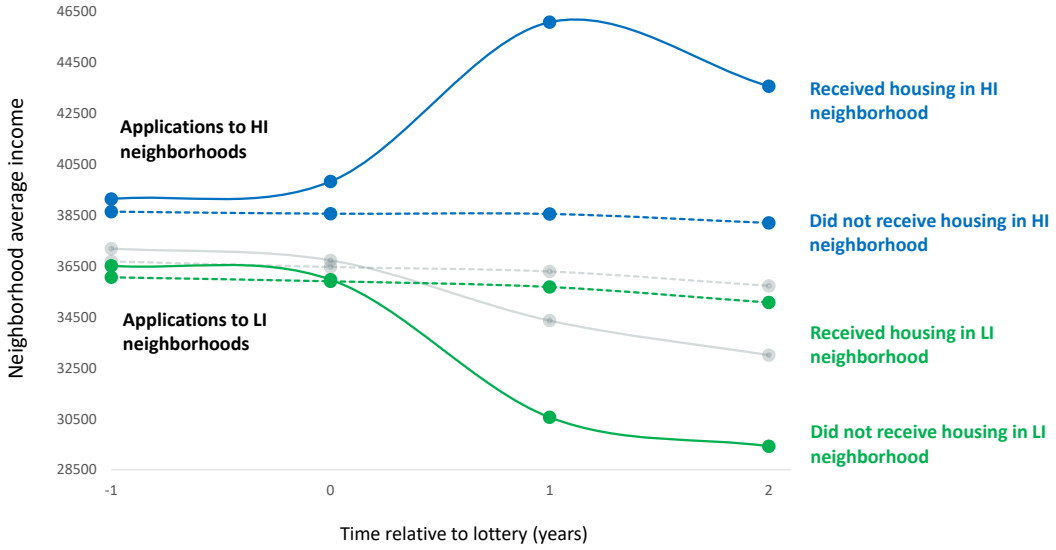

Notes: Year 0 is the year in which the lottery was administered. Neighborhood average income is in 2015 USD. Neighborhood definitions follow administrative boundaries as determined by Statistics Netherlands ('buurten'). Neighborhoods are comparable in population size to U.S. Census tract definitions. High-income neighborhoods are neighborhoods with average income in the top quartile. Low-income neighborhoods are neighborhoods with average income in the bottom quartile. Dashed lines show the trends for applicants who do not receive offers, while solid lines show the trends for those who receive housing offers through the lottery.

Being offered a public housing unit gives agents the option to consume public housing H_s at subsidized

price $\tilde{p}_h$, where $\tilde{p}_h < p_h$.

$$\max_{C,L} U(H_s, C, L) \quad \text{s.t.} \quad wT = wL + \tilde{p}_h H_s + p_c C.$$

While all housing units are subsidized, the quality of the housing (i.e. the amount of H_s) varies substantially from lottery to lottery. As a result, the budget share spent on housing could increase ($\tilde{p}_h H_s > p_h H$) or decrease ($\tilde{p}_h H_s < p_h H$) depending on the housing offered. In the case where the budget share goes up, the share spent on other goods ($wL + p_c C$) must go down by ($\tilde{p}_h H_s - p_h H$). This point is illustrated in Figure 3, where the top figure shows a hypothetical agent's budget constraint and indifference curve for housing consumption and all other consumption. The figure shows an agent's initial choice of housing and other goods (point A), as well as two new alternative consumption bundles involving public housing. Point B shows the agent receiving an offer for a low-quality house where the total budget share spent on housing decreases and, as a result, consumption of other goods increases. Point C shows the agent receiving an offer for a high-quality house where total budget share spent on housing increases and, as a result, non-housing consumption must decrease. In both cases the agent is better off, but the two treatments have opposite implications for how consumption of other goods must change.[15]

The bottom two panels in Figure 3 show how non-housing consumption and leisure could hypothetically change if the agent moved from A to B (left plot) or from A to C (right plot). If the agent moves into house B, the budget share spent on housing decreases and thus remaining budget for non-housing consumption and leisure increases and the budget constraint shifts out. In contrast, if the agent moves into house C, their budget constraint for non-housing consumption and leisure shifts in. As illustrated in Figure 3, agents decrease (increase) both leisure and non-housing consumption when the budget constraint shifts in (out), but this is not the only possible response. While the combined consumption of C and L must decrease in response to a shift in of the budget constraint, the size of the reduction in C or L will depend on (1) if C and L are normal goods and (2) if C and L are substitutes or complements to H. In particular, leisure will fall when the budget contstraint shifts in, except in the case where either leisure and housing are highly complementary goods or leisure is an extremely inferior good.

Next, consider the case where public housing may additionally provide better labor market opportunities by moving the agent into a different neighborhood. Assume that $\tilde{p}_h H_s > p_h H$, but that by accepting

[15]Note that the agent cannot simply choose a new optimal bundle of housing and other goods after the housing subsidy is in place, as subsidies are tied to specific units.

Figure 3: The effect of public housing receipt on non-housing consumption and leisure.

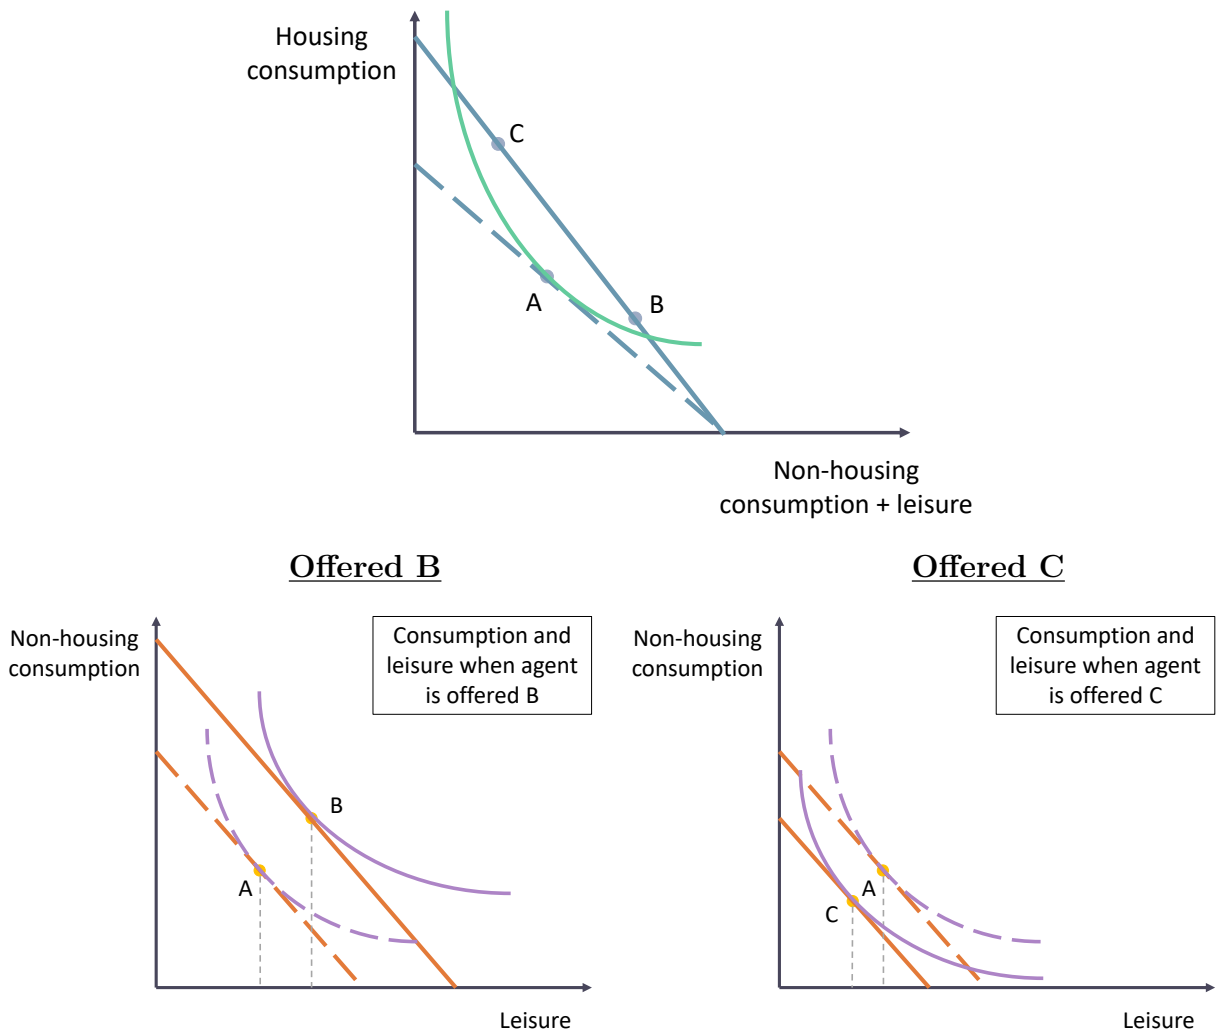

Notes: Top figure shows the utility curve and budget constraint for a hypothetical recipient comparing housing consumption and all other consumption where prior to receiving housing the agent chooses to consume at point A. The hypothetical agent is then offered one of two houses B and C, both of which are preferred to A. The bottom left figure shows how the constrained optimization between non-housing consumption and leisure would change if the agent accepted house B, while the bottom right figure shows how this optimization problem changes after accepting house C.

public housing, wage changes to $w_p > w$. Now the impact on the budget available for leisure and consumption is ambiguous as the overall budget increases by $(w_p - w)T$ while the share spent on housing increases by $\tilde{p}_h H_s - p_h H$. In the case where the overall budget left for consumption and housing decreases ($(w_p - w)T < \tilde{p}_h H_s - p_h H$), the same results apply as before, but with the additional result that leisure has become more expensive, causing the recipient to further reduce leisure.

Expected patterns of selection into lotteries. The discussion above traced out theoretical predictions on labor-supply and consumption responses to an offer of public housing. Yet, both of the cases considered in the discussion above are cases where the agent receiving the offer prefers accepting public housing over his current bundle of consumption. Potential applicants in the Amsterdam housing lotteries will not in general prefer all of the public housing units over their current consumption bundle. Moreover, the rules of the allocation system specify that they can participate in at most two lotteries per week (see Section 2.1). For these reasons, potential applicants repeatedly self-select into lotteries among the available options. Both of these factors make it unlikely that selection into lotteries will be 'as good as random.' This complicates use of the model laid out above for the interpretation of lottery-based estimates of the effect of receiving public housing.

Indeed, Figure 2 clearly shows that there is strong selection on level into lotteries for housing in different neighborhoods, with applicants to lotteries in low-income neighborhoods having lower baseline neighborhood income than applicants to lotteries in high-income neighborhoods. This suggests that there may be further observed or unobserved differences between these applicant pools. In the context of the simple model above, these sorting patterns could arise from agents with stronger preferences for leisure or non-housing consumption choosing to apply for housing in lower-income neighborhoods (which have lower rent prices). Similarly, agents with lower wage rates (and thus smaller overall budgets) will apply for housing in lower-income neighborhoods if the three goods are complements, as low earners would have to commit too large of a share of their budget to housing if they moved into a subsidized unit in a high-income neighborhood.

The lottery mechanism also creates strategic incentives for those in more immediate need of housing to apply to lotteries in which they are more likely to receive an offer. As shown in Figure 4, lotteries for housing in low-income neighborhoods receive substantially fewer applications than lotteries in high income neighborhoods,regardless of the size of the housing unit. Thus, individuals currently living in low-income neighborhoods may strategically apply to lotteries in low-income neighborhoods if they are in higher need of housing.

Figure 4: Strategic incentives in the application decision.

Notes: This figure shows the average number of applications by neighborhood income for the housing unit being allocated. The left-most cluster of bars shows the average number of applicants across all types of apartments, by neighborhood-income bin. The three additional groups of bars show the same quantities, conditional on the number of rooms in the housing unit up for lottery.

Summary of model predictions. The discussion above generated three theoretical predictions. First, the same person may respond to being offered housing that he prefers over his current situation by increasing or decreasing his labor supply, depending on the type of housing he is offered. In particular, if housing receipt lowers the share of the budget spent on rent, it will have a different effect on leisure decisions than if it increases the share of the budget spent on rent. Under mild requirements on the recipient's utility function, we would expect a positive labor-supply response when the share of the budget spent on rent increases and a negative labor-supply response when it decreases. To the extent that 'upward' ('downward') moves correspond to increases (decreases) in budget shares spent on rent, we would therefore expect these moves to have a positive (negative) impact on labor supply. Second, different people may have opposite labor supply responses to being offered the same house depending on their wage rate, preferences over leisure and non-housing consumption, and whether agents have different complementarities between housing consumption and leisure. Finally, different people may select into different lotteries depending on their preferences over neighborhoods, their wage rates, and their preferences over housing and non-housing consumption. People may also select into different housing based on their current need for stable housing arrangements as different lotteries will have different probabilities of receiving an offer.

3 Research design and implementation

In this section, I develop an empirical framework for using the lotteries embedded in Amsterdam's centralized assignment system to identify and estimate the effect of moving into public housing on compliers

(i.e., those who would take up assistance if offered). This framework accounts for non-compliance with offers of assistance, and for the issue that lottery offers are a function of the take-up status of other applicants.[16]

3.1 Empirical framework

Consider a setting with lotteries indexed by j, and applicants to lottery j indexed by i. For a given lottery j, let N_j be the number of applicants, let Z_{ij} be an indicator for applicant i receiving an offer in lottery j, let $D_{ij}(z)$, for $z \in \{0, 1\}$, be an indicator for treatment of applicant i given offer receipt,[17] and let $Y_{ij}(d)$, for $d \in \{0, 1\}$, denote potential outcomes for applicant i given treatment status. Let D_{ij} be an indicator for the observed treatment status of applicant i to lottery j, i.e., $D_{ij} = D_{ij}(1)Z_{ij} + D_{ij}(0)(1 - Z_{ij})$. Similarly, let Y_{ij} be the observed outcome of interest for applicant i in lottery j, i.e., $Y_{ij} = Y_{ij}(1)D_{ij} + Y_{ij}(0)(1 - D_{ij})$. In this type of setting, researchers often take a two-stage least-squares approach to evaluating causal treatment effects, where the first- and second-stage equations take the form:

$$D_{ij} = \gamma + \delta Z_{ij} + L_j\rho + \nu_{ij}, \tag{1}$$

$$Y_{ij} = \alpha + \beta D_{ij} + L_j\pi + \varepsilon_{ij}. \tag{2}$$

Here, L_j is a fixed effect for lottery j.

The identifying assumption frequently adopted in this context is the assumption of *random assignment of offers*. Letting $\mathbf{Z_j} = (Z_{ij})_{i\in\{1,..,N_j\}}$, this assumption can be formally stated as

$$(D_{ij}(0), D_{ij}(1), Y_{ij}(0), Y_{ij}(1))_{i\in\{1,..,N_j\}} \perp\!\!\!\perp \mathbf{Z_j}\ \forall j.$$

However, because the offer process in lottery j ends once the unit has been accepted, offer assignment status for the ith applicant in lottery j is a function of both his own priority in the random order in which offers are made, and the complier status of those who were randomly assigned a higher priority. To operationalize the offer process, suppose that applicants in lottery j are assigned a random number that

[16] This issue has previously been discussed in de Chaisemartin and Behaghel (2017), who propose a different solution that is not feasible in my setting.

[17] Implicit in the notation is the assumption that take-up responses $D_{ij}(z)$, $z \in \{0, 1\}$, do not depend on the number of rejections that may have preceded an offer to applicant i in lottery j. That is, an individual receiving an offer for a given unit does not consider his assigned priority when deciding whether to accept or reject the offer.

determines the order in which the housing unit will be offered. Let R_{ij} denote the position of applicant i in the order. I will refer to R_{ij} as i's *priority*. For example, $R_{ij} = 1$ means applicant i will be the first to be offered the unit, $R_{ij} = 2$ means applicant i will be offered the unit if the applicant who was assigned first priority declines, and so on. Let $\mathbf{R_j}$ collect the R_{ij} in a vector.

From the discussion above it follows that Z_{ij} equals one for all applicants assigned a priority higher than or equal to the complier assigned the highest priority, and zero otherwise:[18]

$$\mathbf{Z_j} = \left(\mathbb{1}\{R_{ij} \leq \min_{k \in \{1,..,N_j\}: D_{kj}(1)=1} \{R_{kj}\}\} \right)_{i \in \{1,..,N_j\}}.$$

It is clear that random assignment of offers is not satisfied in this setting, because $\mathbf{Z_j}$ is a function of $\mathbf{D_j(1)}$, the random vector that stacks complier status for applicants to lottery j. This violation of *random assignment of offers* will bias the usual estimators of the intent-to-treat parameter and the probability of take-up given an offer, even in samples with large applicant pools, or settings where treatment effects are assumed constant across recipients but with selection on levels. This issue has been pointed out and discussed in detail in de Chaisemartin and Behaghel (2017). The authors propose an approach to causal inference that is not feasible in the current setting, because it requires each lottery to allocate at least two seats, whereas the lotteries in the current setting only allocate a single unit. Hence, in the next section, I propose a different approach that replaces the assumption of *random assignment of offers* with the assumption of *random assignment of priority*.

Assumption 3.1 (Random assignment of priority). $(D_{ij}(0), D_{ij}(1), Y_{ij}(0), Y_{ij}(1))_{i \in \{1..N_j\}} \perp\!\!\!\perp \mathbf{R_j}\ \forall j$.

By construction, $\mathbf{R_j}$ is a randomly drawn element from the set of permutations of $(1, .., N_j)$. Assumption 3.1 is therefore satisfied. Table 1 reports empirical tests of the assumption. Under the null hypothesis of random assignment, demographics and past application behavior cannot be used to predict whether a given applicant will be assigned in a given priority position. To test this hypothesis across all lotteries, I consider whether covariates can predict priority after controlling for lottery fixed effects. The results from this exercise, reported in Table 1 (columns 1-3), show no evidence of a statistically significant relationship between observed covariates and applicants being assigned first, second, or third priority.

In addition to random assignment of priority I rely on a relevance condition that is standard in instrumental-variable frameworks. Specifically, receiving an offer must impact individuals' probability of

[18] This definition implicitly assumes that lottery j has at least one complier, an assumption that holds in my empirical setting. In the remaining discussion, I implicitly assume this is true for every lottery, i.e. $\sum_{j=1}^{N_j} D_{ij}(1) > 0\ \forall j$.

moving into the unit.

Assumption 3.2 (Relevance). $\mathbb{P}\left(D_{ij}(1) \neq D_{ij}(0)\right) > 0 \ \forall i, j$.

Table 1 provides results of empirical tests of Assumption 3.2, showing that it is convincingly satisfied with a statistically significant first-stage of a 0.593, which implies 60% of the first offers are accepted.

As in other instrumental-variable frameworks, interpretation of results requires a monotonicity assumption.

Assumption 3.3 (One-sided non-compliance). $D_{ij}(0) = 0$.

Because applicants cannot move into housing for which they have not received an offer, my setting permits only one-sided non-compliance. Therefore, Assumption 3.3 is satisfied. This rules out the existence of 'always-takers' (those who would take up treatment even if it is not offered to them). Yet, it still leaves substantial room for non-compliance, as individuals who are offered housing may choose to decline. Indeed, the data clearly shows imperfect take-up of lottery offers; Table 1 shows that compliance by those who are assigned first priority is about 60%.

Finally, let C_j denote the number of compliers in lottery j, i.e., $C_j = \sum_{i=1}^{N_j} D_{ij}(1)$, and let $\rho_j := \frac{C_j}{N_j}$ denote the fraction of compliers in lottery j as ρ_j. I assume that complier shares are always positive.

Assumption 3.4 (Positive complier share). $\rho_j > 0 \ \forall j$.

This assumption holds in my empirical setting as all lotteries result in the allocation of the house to an applicant. In general, this assumption is likely to hold in most empirical settings where the number of applicants is large.

3.2 Causal inference under random assignment of priority

In this section I show that the causal effect of treatment on compliers for a given lottery is identified up to a bias term that vanishes with the size of the applicant pool. I then argue that estimation can be implemented using a simple approach that is common in the literature.

Consider a single lottery j. Let O_j denote the number of offers for that lottery before someone accepts the housing unit, i.e., $O_j = \sum_{i=1}^{N_j} Z_{ij}$. Assume that potential outcomes and complier status for applicants to the lottery are sampled i.i.d. according to a distribution F_j, and that the fraction of compliers in the

pool of applicants in expectation is always ρ_j, regardless of sample size. The parameter of interest is the average effect of treatment on compliers, which I denote by Δ_j:

$$\Delta_j := \int [Y_j(1) - Y_j(0)]\, dF_{j \mid D_j(1)=1}.$$

Finally, define

$$\hat{\Delta}_{j,N_j} := \sum_{i=1}^{N_j} Y_{ij} Z_{ij} - \tfrac{O_j}{N_j - O_j} \sum_{i=1}^{N_j} Y_{ij} \left(1 - Z_{ij}\right). \tag{3}$$

This is the within-lottery Wald estimator.[19] The following proposition states that the Wald estimator is asymptotically unbiased for Δ_j as the size of the applicant pool grows large.[20] This essentially restores our ability to use lottery-specific contrasts between treated and untreated individuals to identify causal effects, at least up to a 'small' bias term.[21]

Proposition 3.1. *Under Assumptions (3.1)–(3.4),*

$$\lim_{N_j \to \infty} \mathbb{E}\left[\hat{\Delta}_{j,N_j}\right] = \Delta_j.$$

Proof. See Appendix section B.2.

This result is not immediately obvious. Intuition is helped by two observations. First, for finite applicant pools, offer assignment is determined by first drawing a permutation of applicant labels and then offering treatment in the order determined by the permutation until someone accepts. As the size of the applicant pool grows, while keeping the fraction of compliers constant, this process becomes similar to offering treatment to individuals drawn at random and with replacement from the population, until someone accepts. For large applicant pools, the number of offers therefore follows a geometric distribution with its parameter equal to the fraction of compliers in the pool. Second, drawing with replacement until a complier is drawn in expectation generates the same number of compliers (one) as a two-step process where the number of offers is first determined by a realization of the geometric distribution, and

[19] I derive this expression in Appendix B.1.

[20] The expectation in Proposition 3.1 is shorthand notation for the expectation over the distribution of $\mathbf{R_j} \mid N_j$. I.e., $\mathbb{E}\left[\hat{\Delta}_{j,N_j}\right] = \int \hat{\Delta}_{j,N_j} dG_{j,N_j}$, where G_{j,N_j} denotes the CDF of the permutation distribution of $\mathbf{R_j} \mid N_j$.

[21] It may be possible to derive tight bounds on the identified set and develop bias reduction techniques, based on the expression for the bias term developed in the proof to Proposition 3.1. This falls outside the scope of the current paper.

then offering treatment to that number of applicants at random. The number of non-zero terms in the first sum appearing in equation (3) is determined according to the geometric sampling process, while the second term weights outcomes in the control group so as to represent a randomly drawn group of size O_j. Thus, the estimator mimics a difference in sample means, where the means are weighted to have the same effective size, and are generated according to two sampling processes that each draw a single complier in expectation.[22]

I now discuss how this result can be used as a building block for characterizing the estimand for a frequently-used approach to pooling lotteries. Let $\underline{N} := \min_j N_j$, and let $\hat{\beta}_{\underline{N}}$ be the two-stage least-squares estimator that corresponds to the model in (1)–(2), for a fixed number of lotteries J.

Corollary 3.2. *Under Assumptions (3.1)–(3.4),*

$$\lim_{\underline{N}\to\infty} \mathbb{E}\left[\hat{\beta}_{\underline{N}}\right] = \tfrac{1}{J}\sum_{j=1}^{J} \Delta_j.$$

Proof. See Appendix section B.3.

Corollary 3.2 characterizes the two-stage least-squares estimand from model (1)–(2) as approximating the equally-weighted average of lottery-specific effects of treatment on the treated. This result is useful for three reasons. First, it shows that despite the endogeneity of the offer process, causal inference is still possible.[23] Second, it states that the estimand has a clear economic interpretation, corresponding to a feasible program change that is of direct interest to policy makers.

Third, the estimand can be constructed for a sub-sample of lotteries and interpreted as the effect on a randomly drawn complier in the average lottery from that sub-sample. For some sub-samples, such as lotteries from a particular neighborhood, this directly corresponds to a policy, such as a marginal increase in the number of lotteries in a given neighborhood. It also provides us with neighborhood-specific parameters that are directly comparable, as the estimands have equal weights, rather than weights that are difficult to interpret and vary across sub-samples of lotteries.

[22]Note that, following the same logic, we could have constructed an unbiased estimator that only uses three randomly drawn observations for the 'control group term' in equation (3). However, the proposed estimator makes more efficient use of the available data.

[23]But note that, while the Wald estimator becomes unbiased as the applicant pool grows large, it is not true that the numerator provides an unbiased estimate of the intent-to-treat parameter even in the limit, nor does the denominator provide an unbiased estimator of the first-stage probability of offer take-up.

Table 1: Testing for random assignment of priority and relevance.

	Dependent variables				*Explanatory variables*	
	1st priority (1)	2nd priority (2)	3rd priority (3)	Take-up (4)	Mean (5)	St. Dev. (6)
1st priority				0.5929***	0.0015	0.039
				(0.0184)		
2nd priority				0.1900***	0.0015	0.039
				(0.0147)		
3rd priority				0.1032***	0.0015	0.039
				(0.0115)		
I. *Demographics*						
Male	0.0001	-0.0001	-0.0000		0.495	0.5
	(0.0001)	(0.0000)	(0.0001)			
Age	-0.0000	0.0000	0.0000		29.522	7.755
	(0.0000)	(0.0000)	(0.0000)			
Foreign born	0.0000	0.0002	0.0002		0.351	0.477
	(0.0002)	(0.0001)	(0.0002)			
Second gen. immigrant	-0.0001	0.0000	0.0001		0.387	0.487
	(0.0002)	(0.0001)	(0.0001)			
HH size	0.0000	-0.0000	-0.0000		2.71	19.34
	(0.0000)	(0.0000)	(0.0000)			
II. *Housing and neighborhood characteristics at time of application*						
NH average income	0.0000	0.0000	0.0000		37,136	8,859
	(0.0000)	(0.0000)	(0.0000)			
NH fraction non-western	0.0005	0.0004	-0.0005		0.236	0.164
	(0.0004)	(0.0005)	(0.0004)			
NH fraction low-income	-0.0002	-0.0006	0.0012		0.208	0.045
	(0.0015)	(0.0015)	(0.0015)			
III. *Labor market and residential mobility history*						
Personal labor income	0.0000	-0.0000	-0.0000		16,252	14,826
	(0.0000)	(0.0000)	(0.0000)			
Personal total income	-0.0000	0.0000	-0.0000		20,375	13,045
	(0.0000)	(0.0000)	(0.0000)			
HH labor income per adult	0.0000	0.0000	-0.0000		16,751	15,242
	(0.0000)	(0.0000)	(0.0000)			
HH total income per adult	-0.0000	-0.0000	0.0000		21,003	13,479
	(0.0000)	(0.0000)	(0.0000)			
Employed	0.0000	-0.0002	-0.0000		0.493	0.500
	(0.0002)	(0.0002)	(0.0002)			
Residential move, $t-1$	-0.0001	0.0001	-0.0001		0.270	0.444
	(0.0001)	(0.0001)	(0.0001)			
Residential move, $t-2$	0.0000	0.0001	-0.0000		0.262	0.440
	(0.0001)	(0.0001)	(0.0001)			
F-stat for joint test	0.5260	0.7911	0.6576	1,039		
p-value	0.9354	0.6975	0.8380	0.0000		
Lottery FE	✓	✓	✓	✓		
N	452,014	452,014	452,014	470,172	470,172	470,172

Notes: The first three columns show the coefficients from a regression of demographics, prior neighborhood characteristics, and prior labor market and residential mobility outcomes on the probability of being assigned first, second, or third priority in a given lottery. This provides evidence that ran is randomly assigned. The fourth column regresses take-up of treatment (i.e. moving into a lotteried unit) on indicators for being given first, second, or third priority and controls. The statistically significant coefficients and large F-statistic provide evidence that assigned rank is relevant. The F-statistics are reported for the model after projecting out lottery fixed effects. The F-statistic for the take-up regression (model (4)) is computed for the regression with '1st priority' as the dependent variable. As an additional check for random assignment of priority, running the regression of take-up on priority with the covariates used in the balance regression as controls shows that the coefficients are very stable, which they should be if the pre-determined outcomes are indeed uncorrelated with the instrument. The coefficients for the regression with controls are 0.5929 ($p<.01$), 0.1888 ($p<.01$), and 0.1026 ($p<.01$). The F-stat for the full projected model is 59.17 for the regression using '1st priority.' Columns (5) and (6) report the mean and standard deviation of the various demographic and prior outcomes in the sample. Robust standard errors clustered at the individual level. * = $p<.1$; ** = $p<.05$; *** = $p<.01$.

4 Main results

This section reports estimates on the causal effects public housing receipt. The first section focuses on the main results, while the following section discusses results for several sub-populations of interest. I consider short-run impacts on household income, personal labor income, employment, cumulative subsequent moves, tax-assessed house value, and average neighborhood income.

4.1 Mean impacts

I find that the average causal impact of housing assistance is negative with a statistically significant reduction in household income of over $2,000 per adult. Recipients of housing also end up living in lower-valued housing in lower-income neighborhoods and are less likely to move in the future. These results are consistent with estimates of the impact of housing voucher receipt in the literature (Jacob and Ludwig (2012)). Eligibility is not periodically re-evaluated after assistance is awarded, suggesting that the reduction in income is not driven by a desire to maintain eligibility for the current home.

Table 2 summarizes the mean impacts across all lotteries 2 years after receipt. The table is broken down into four sets of results: the first shows the effects on labor market outcomes, the second on housing quality and expenditure, the third on neighborhood characteristics, and the fourth on fiscal costs such as other public assistance receipt. The column labeled 'Base' reports the average outcome for applicants who did not receive an offer.[24] The column labelled 'OLS' reports simple OLS estimates, while the column labeled 'IV' reports two-stage least squares results using the ever-offer instrument discussed in Section 3.1. Both the OLS and IV estimates include lottery fixed effects. The final column labeled 'N' reports the number of observations.

On average, public housing receipt is associated with a notable decrease in labor market outcomes. As shown in Table 2, OLS estimates among applicants show that those who move into public housing experience a -2,250 dollar reduction in household labor income per adult, a -1,800 reduction in individual labor income, and a four percentage point reduction in employment. Column (2) of the table shows the estimated effect of treatment on the treated using the 'ever offer' instrument (Z_{ij}). Similar to the OLS estimates, household labor income falls by 2,300 dollars per adult and personal labor income falls by 1500 dollars. The estimated impact on employment is slightly smaller, but statistically insignificant and

[24] All monetary results are reported in 2015 US dollars throughout this paper.

statistically indistinguishable from the OLS estimate.

Table 2: Effects of public housing receipt: two-year impacts.

	Base	OLS (1)	IV (Z_{ij}) (2)	N
I. *Income and employment*				
HH labor income	17,711	-2,279***	-2,375***	439,381
		(647.3)	(871.8)	
HH total income	22,834	-1,371**	-1,323*	439,381
		(547.2)	(751.7)	
HH disposable income	20,129	-1,028**	-1,042*	439,381
		(440.6)	(615.0)	
Individual labor income	17,042	-1,795***	-1,513*	467,027
		(635.7)	(615.0)	
Individual total income	22,040	-810.1	-382.9	467,027
		(542.2)	(729.1)	
Employment	0.58	-0.0440***	-0.0341	439,877
		(0.0169)	(0.0227)	
II. *Housing quality and rent expenditure*				
House value (per m^2)	2,635	-105.5***	-154.5***	463,373
		(30.50)	(43.32)	
Monthly rent expenditure (net of subsidy)	445.8	10.70	-5.998	464,181
		(13.48)	(19.11)	
Cum. two-year residential move rate	0.51	-0.2754***	-0.2987***	467,027
		(0.0168)	(0.0239)	
III. *Neighborhood characteristics*				
NH average income	36,470	-2,913***	-3,359***	467,027
		(237.3)	(361.7)	
NH fraction non-western	0.24	0.0505***	0.0533***	467,027
		(0.004)	(0.0084)	
Social distance	0.50	0.0404***	0.0349***	467,027
		(0.0067)	(0.0085)	
Distance to city center (m)	11,752	-3,555***	-3,975***	467,027
		(483.2)	(683.2)	
IV. *Fiscal cost*				
Public assistance receipt	0.23	0.0540***	0.0367**	439,877
		(0.0140)	(0.182)	
Fiscal externality	-2,456	-1,207***	-1,160***	467,027
		(294.9)	(401.0)	
Lottery FE		✓	✓	

Notes: Individual income is income from paid or self-employment before taxes and transfers, in 2015 US dollars. HH (household) income is average labor income for adults in the household, excluding adult children. Employment is an indicator for paid employment or self-employment being the individual's main source of income. Social distance is the fraction of neighborhood residents that are not of the applicant's ethnic group. Fiscal externality is individual taxes paid net of non-housing transfers. The 'Base' column reports average outcomes for those not treated. The 'OLS' column reports results from a simple linear regression while the 'IV' column reports two-stage least squares results using the ever-offer instrument developed in the previous section. The 'N' column reports the total number of observations. Regressions include controls for baseline outcomes and lottery fixed effects. * = $p<.1$; ** = $p<.05$; *** = $p<.01$. 'Base' is the sample mean for those who do not receive an offer, at the time of application.

Part II of Table 2 reports results on housing quality and rent expenditure. OLS estimates show that, on average, those who receive housing through lotteries move into lower-value homes. IV estimates similarly find a reduction in house value of -160 dollars per square meter, a reduction of approximately 6%. Overall, this evidence suggests that, on average, compliers who receive housing through lotteries move into lower-quality housing suggesting that rent expenditures decrease.

Part III of Table 2 reports results on neighborhood characteristics. On average, receiving housing through the lottery is associated with a substantial reduction in neighborhood quality and an increase in the fraction of non-western residents in the neighborhood. IV estimates find that average neighborhood income falls by 3,400 dollars and that the fraction of non-western residents increases by 5 percentage points (more than a 25% increase). Consistent with the model in Section 2.2, individuals move into lower-income neighborhoods, which may decrease their proximity to jobs and economic activity.

Part IV of Table 2 reports results on how accepting public housing affects other public assistance receipt. On average, for those who receive housing through a lottery, public assistance receipt increases by approximately 4 percentage points using the ever-offer instrument, with similar but somewhat larger estimates from the OLS estimates. Part IV also reports effects on the 'fiscal externality' of an individual receiving public housing which I define as the amount paid in taxes net of non-housing government transfers.[25] The fiscal externality is on average around -1,200 dollars, showing that the net transfers from the government increase, implying that, on average, reliance on the government increases above and beyond the housing subsidy.

4.2 Heterogeneity in responses by neighborhood

On average, housing recipients end up living in poorer neighborhoods and earning less, yet this masks substantial heterogeneity in the impacts. Specifically, receipt of housing in high-income neighborhoods causes large and statistically significant labor market gains.

Figure 5 displays the mean impact of lotteries on household labor income, employment, and fiscal externality, with each row representing a different outcome. Within each row, the first bar reports the IV estimates for the average effect for individuals who receive housing through the lottery, as reported in Table 2. The second through fourth bars report the IV estimates, but conditional on the lottery being in a low-income neighborhood (second bar), middle-income neighborhood (third bar) or high-income neighborhood (fourth bar). All monetary amounts are in 2015 US dollars and statistical significance is indicated by stars above the bars as explained in the Figure notes.

While the overall IV effects are negative for labor-market outcomes, the results conditional on the average neighborhood income of the public housing units show a clear pattern. Public housing units

[25]Note, here I exclude housing transfers as these mechanically make this estimate more negative, though these would need to be included for a cost-benefit analysis. Taking the cost of public housing into account will lead to a negative net benefit to the government across all neighborhoods.

offered in low-income neighborhoods have much stronger negative effects on income and labor supply, while housing units offered in high-income neighborhoods have large and statistically significant *positive* effects.[26] Similarly, we find that those moving to low-income neighborhoods have larger negative fiscal externalities, while those moving to high-income neighborhoods have positive fiscal externalities, though this estimate is not statistically significant.

These heterogenous effects are consistent with the predictionts from the price theory model in Section 2.2. In the model, the labor-market response will likely differ depending on whether the share of the budget spent on housing increases or decreases. Consistent with the model's prediction, the bottom plot in Figure 5 shows that the average monthly rent expenditures fall for those who receive housing in low-income neighborhoods (who earn less), while it increases for those receiving housing in high-income neighborhoods (who earn more).

For those moving to high-income neighborhoods, the increase in household neighborhood income is notably larger than the annual estimated increase in rent. To explain this finding in the context of the simple model in Section 2.2, we need additional conditions such as housing and non-housing consumption being strong complements or high-income neighborhoods lowering the fixed-costs of working or raising labor market returns. As shown in the second panel of figure 5, individuals receiving housing in high-income neighborhoods increase employment by almost 15 percentage points, which is also consistent with high-income neighborhoods lowering the fixed costs or raising the returns to work.

The results in Figure 5 are only for those who chose to apply to lotteries in a given neighborhood and then accept a housing offer, yet these results are still directly policy relevant. In particular, these estimates correspond to a policy which marginally expanded the stock of housing up for auction in low-, medium-, or high-income neighborhoods.

4.3 Heterogeneity in responses by applicants' neighborhood of origin

In addition to variation in lottery impacts across different neighborhoods, the model laid out in Section 2.2 predicts that impacts may vary if we condition on the characteristics of the recipients. Conditioning on neighborhood of origin shows that this is indeed the case.

Figure 6 shows the IV estimates on household labor income conditioning on both the average income

[26]Tables with results for additional outcomes of interest for recipients in high and low-income neighborhoods are in Appendix Sections A.3 and A.4

in the neighborhood of the lottery and the average income of the origin neighborhood of the applicant. Each sub-figure reports results based on the average neighborhood income in which the lottery is located, with (a) showing high-income destinations, (b) showing middle-income destinations, and (c) showing low-income destinations. Within each sub-figure, the first bar shows the average treatment on the treated, while the last three bars show the treatment on the treated conditional on the average income of the the applicant's origin neighborhood.

The figure shows that the impact for individuals starting in low or middle income neighborhoods and receiving housing through the lottery in a high-income neighborhoods have very large positive effects with estimated increases household and personal labor income. In contrast, for households already in high-income neighborhoods, receiving public housing in a high-income neighborhood has a small, negative, and statistically insignificant impact on household labor income. This suggests that the positive effects of receiving housing in high-income neighborhoods previously shown in Figure 5 are driven by 'upward transitions' from those originating in low- or middle-income neighborhoods.

Importantly, the results in this subsection are conditional on individuals who apply to and accept the offer for a given lottery. While the returns are estimated to be large for applicants from low-income neighborhoods who apply to lotteries in high-income neighborhoods, we do not know if these returns would generalize to individuals from low-income neighborhoods who are less likely to apply. Section 5 aims to generalize these results.

4.4 Policy relevance of findings

The estimand characterized in Corollary 3.2 describes an easily-interpretable and policy-relevant treatment effect, while allowing for rich patterns of heterogeneity: the assumptions required for causal inference do not restrict variation in impacts across potential recipients, nor do they restrict heterogeneity in treatment across lotteries. This is both an advantage and a disadvantage: on the one hand, this parameter is highly relevant for policy decisions, as it describes expected impacts that would result from a small expansion of the program. Moreover, by conditioning on lotteries offered in a given neighborhood, it can be straightforwardly adapted to study the impact of marginally expanding the number of lotteries in that neighborhood. A policy maker whose objective is to maximize impact by choosing neighborhoods in which to offer housing, regardless of distributional concerns, requires only this information to make his decisions (possibly adjusting for differences in the cost of providing housing in different neighborhoods).

Figure 5: Treatment on the treated, by neighborhood income of public housing unit.

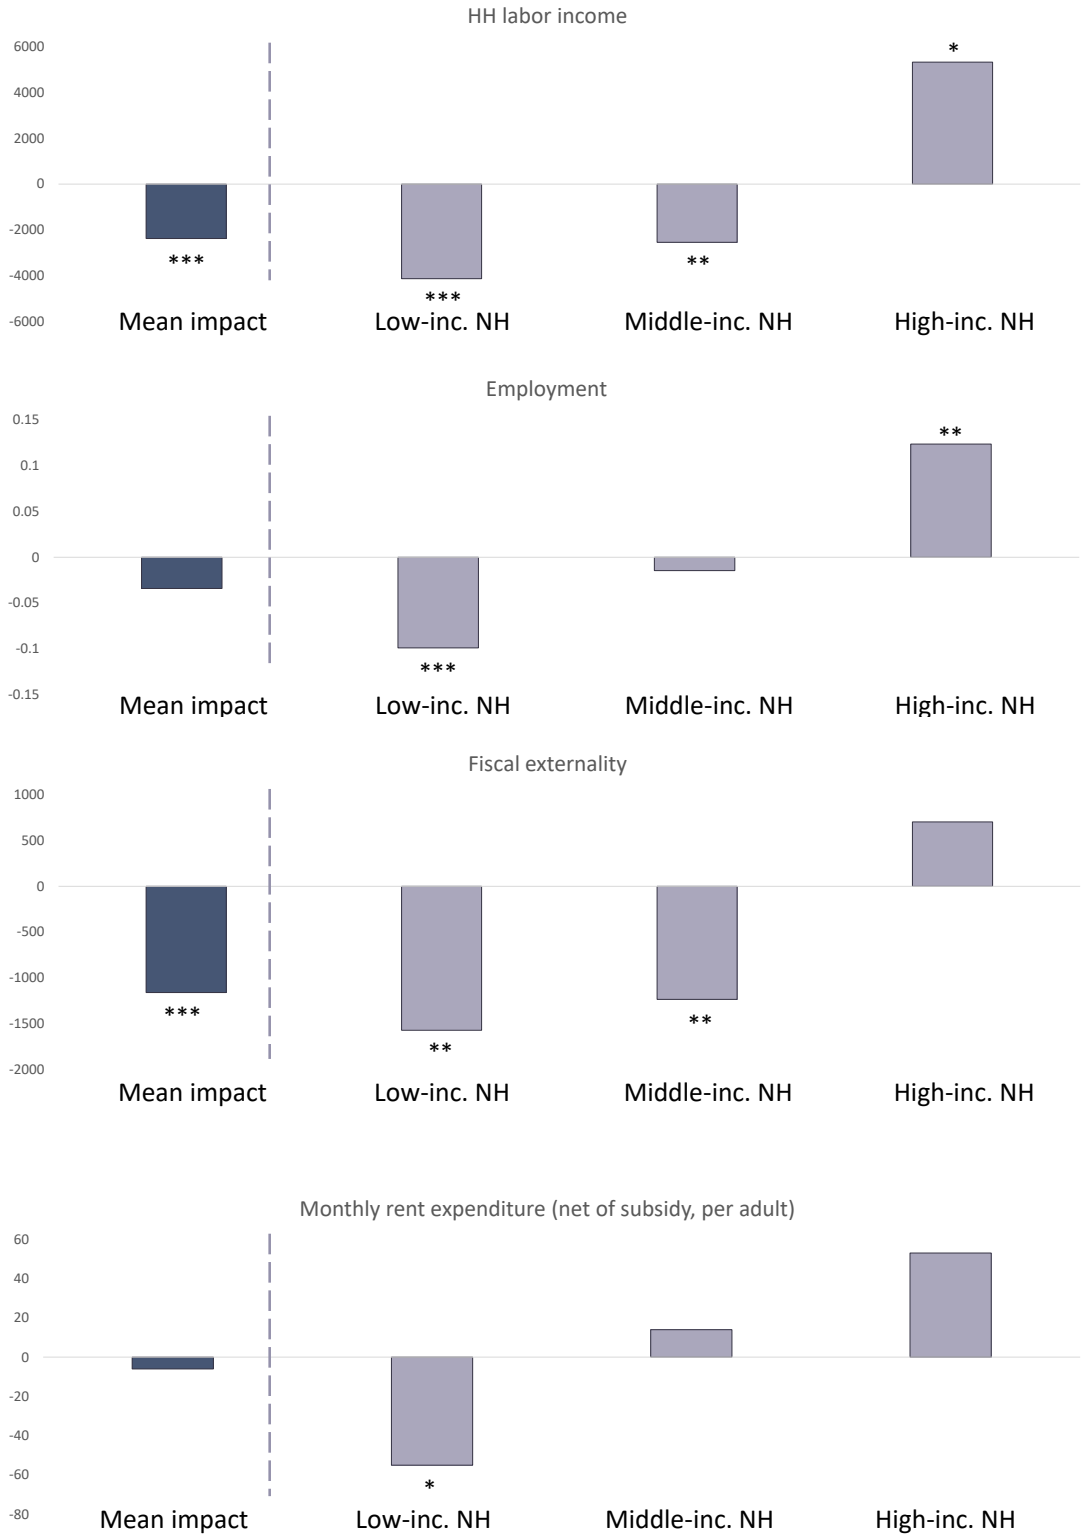

Notes: Table reports the IV estimates of the effect of receiving housing through the lottery. Each row in the figure shows treatment on the treated for a single outcome. The left-most bars show impacts for all lotteries, while the bars to the right of the dashed line show those impacts conditional on destination neighborhood type (high, medium, or low income). HH labor income is average household labor income per adult applicant. Fiscal externality is defined as the total amount paid in taxes net non-housing transfers. * = $p<.1$; ** = $p<.05$; *** = $p<.01$.

Figure 6: Treatment on the treated for household income, by neighborhood income of applicant at time of application.

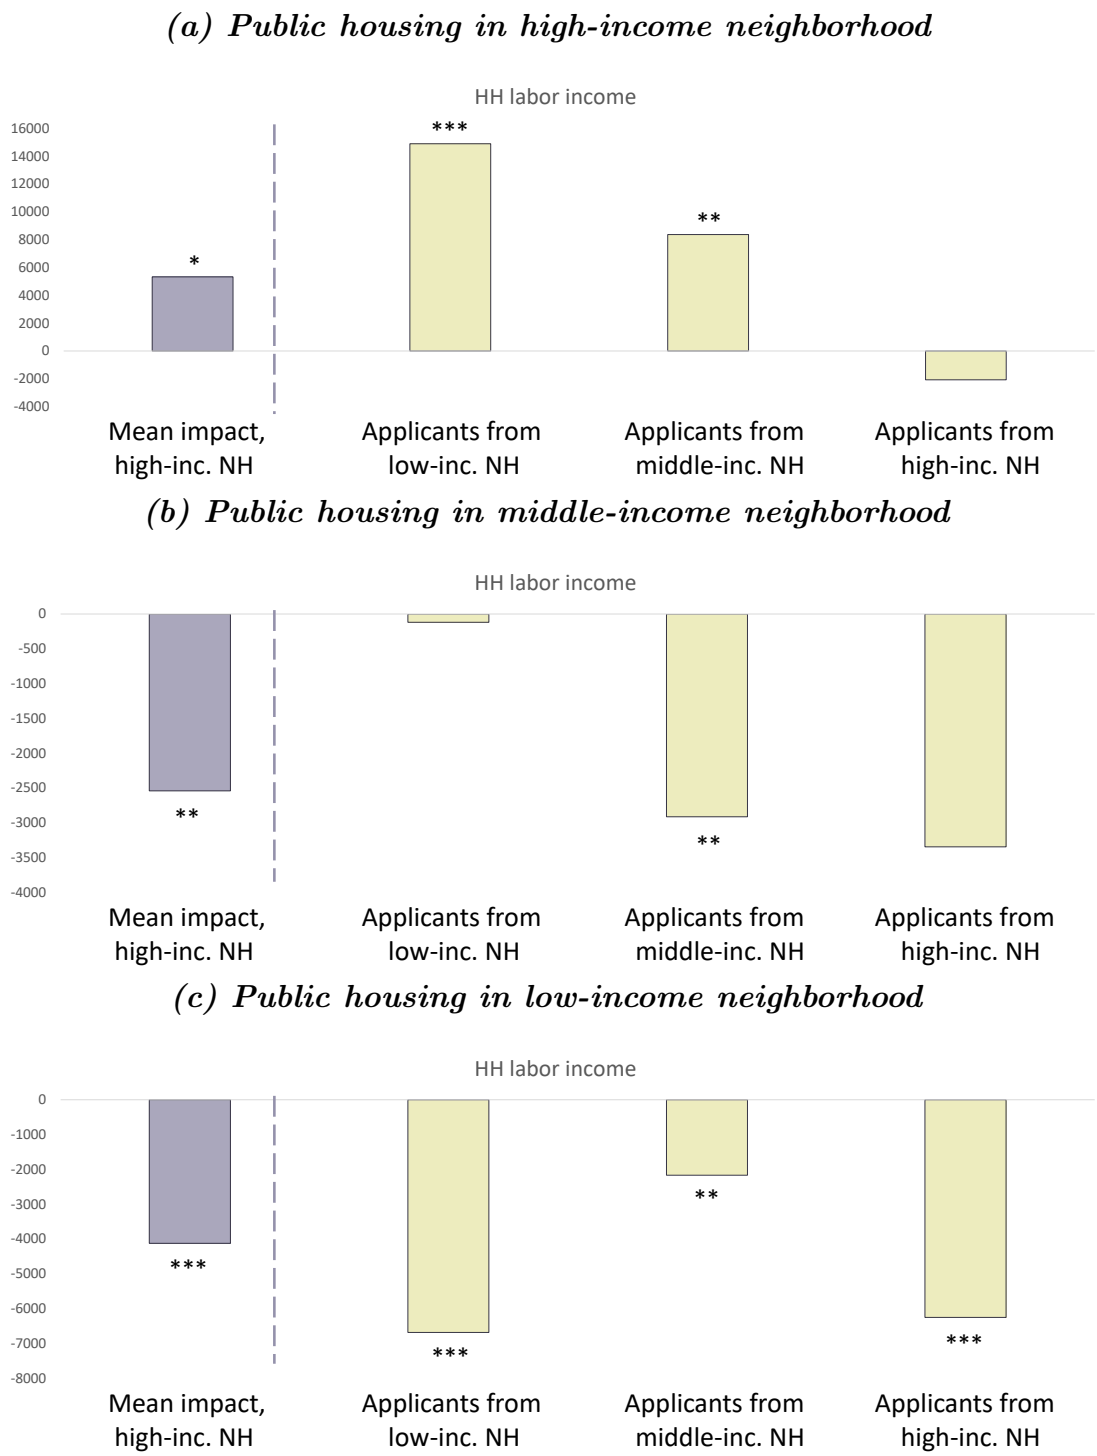

Notes: Table reports the IV results of the treatment on the treated conditional on origin- and destination neighborhood for household labor income. Each row in the figure shows treatment on the treated for (size-adjusted) household income, with the top row showing results for lotteries in high-income neighborhoods, the second row showing results for lotteries in middle-income neighborhoods, and the third row showing results for lotteries in low-income neighborhoods. The left-most bars show impacts for all lotteries in a given neighborhood type (high, medium or low income), while the bars to the right of the dashed line show those impacts conditional on applicants' neighborhood of origin.

Figure 7: Comparing IV and ATE based on observables.

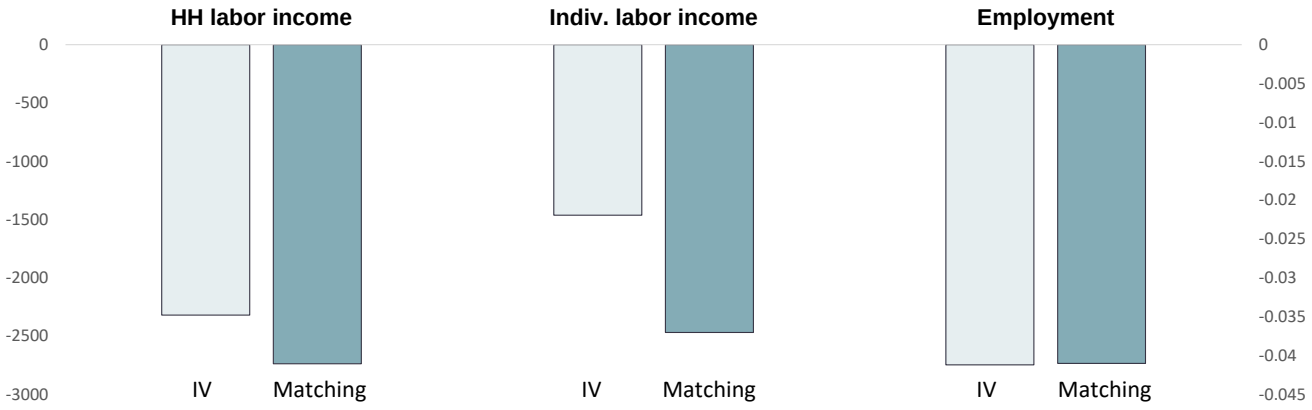

Notes: Figure compares IV estimates of the average treatment on the treated across all lotteries and estimated average treatment effects assuming no selection on unobservables. Specifically, the estimated average treatment effects are generated by estimating by a flexible regression of outcomes on treatment status, observable characteristics, and their interactions. Using this regression, I estimate ATE (observables) = $\int\int \hat{Y}_{i,j}(1) - \hat{Y}_{i,j}(0) dF(\boldsymbol{W}) dG(\boldsymbol{X})$ where $F(\boldsymbol{W})$ is the full distribution of housing characteristics and $G(\boldsymbol{X})$ is the population distribution of observables among all participants. This provides an estimate of the average treatment effect among all applicants that is sometimes called 'regression adjustment' in the literature and is equivalent to parametric matching on observables.

On the other hand, we may be interested to learn more about the drivers of the patterns of heterogeneity documented in the previous section. In particular, we may be interested to learn whether these patterns arise from differences in treatment, versus differences in the individuals receiving treatment. This is particularly important if policy makers wish to target policies based on observable characteristics of individuals.

If application decisions are entirely driven by observable characteristics of applicants, then a reweighting exercise can provide treatment effect estimates that describe impacts of housing in a given neighborhood for a specified composition of applicants.[27] Figure 7 compares the IV estimates with reweighting estimates of the average treatment effect that assumes no sorting on unobservables.[28] Each set of bars compares the IV estimate and the average treatment effect estimated on observables for a given outcome. Overall, the figure shows that the IV estimates and the reweighting estimates are qualitatively similar. However, this adjustment does not account for applicants sorting into lotteries based on unobservable factors. To account for this source of heterogeneity, the next section introduces a model of application behavior and uses that model to correct for selection on unobservables.

[27]Assuming the composition is made up of types who apply to housing in that neighborhood with positive probability.
[28]See Table 4 in the next section for comparisons conditional on destination or origin neighborhood.

5 Understanding heterogeneous responses

This section first develops a model of participation, application, and acceptance decisions to characterize ex-ante search behavior and selection into lotteries that can be utilized to extract latent heterogeneity from the choice history of applicants. I then use the estimated model to correct for selection into lotteries, and describe the resulting average treatment effect estimates.

5.1 Choice model

The application and acceptance behavior of potential recipients is modelled as a three-stage decision problem that repeats itself in each application period. At the start of every period, applicants decide whether or not to participate. Conditional on participation, individuals apply to up to two lotteries. Finally, in the event that they receive a lottery offer, they decide whether to accept or reject the housing being offered to them. The process is illustrated in Figure 8. To define notation, let $i \in \{1, ..., N\}$ index applicants, $t \in \{1, ..., T\}$ index application periods, and $j \in \mathcal{J}_t = \{1, ..., J_t\}$ index the housing units up for lottery in period t. Let W_{jt} be a vector of characteristics of house j in period t and let $\mathbf{W_t}$ stack the W_{jt} and X_{it} be a vector of characteristics of individual i in period t. Let M_{ijt} be a vector of instruments for participation and application decisions that affect the probability individual i applies to a given house j in period t, but that are uncorrelated with i's potential outcomes, and let $\mathbf{M_{it}}$ stack the M_{ijt}. Finally let θ_i be persistent unobserved heterogeneity of individual i which has distribution G_θ in the population of applicants.

Figure 8: Sequence of events in the choice model.

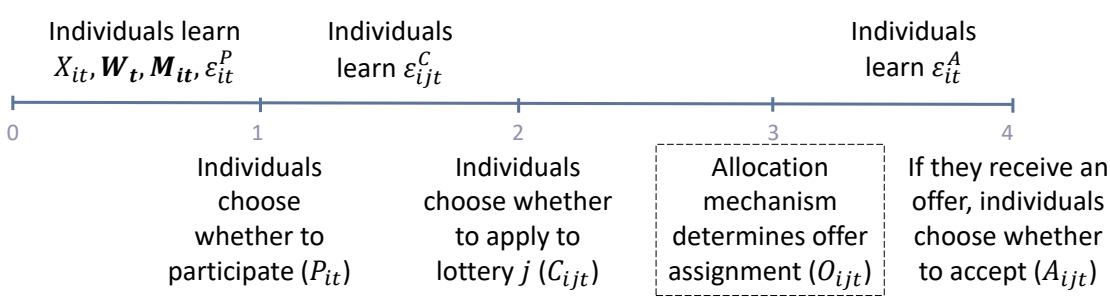

Notes: Figure represents the within-period decisions and timing of information as individuals make participation, application, and acceptance decisions.

Participation decision. Each period t, an individual must decide if it is worth looking into and applying to lotteries. I model this decision as a index crossing model where individuals choose to participate if

$$f^P(\mathbf{W_t}, X_{it}, \mathbf{M_{it}}, \theta_i) + \epsilon_{it}^P > 0.$$

ε_{it}^P may be interpreted as an 'awareness shock' – an idiosyncratic factor determining whether or not the individual searches for housing on the platform in that period.

When modeling this as a logistic regression, the likelihood contribution of individual i in period t is given by:

$$L_{it}^P(\boldsymbol{W}_t, X_{it}, \mathbf{M_{it}}, \theta_i; \boldsymbol{\gamma}) = \left(\frac{\exp(f^P(\mathbf{W_t}, X_{it}, \mathbf{M_{it}}, \theta_i; \boldsymbol{\gamma}))}{1 + \exp(f^P(\mathbf{W_t}, X_{it}, \mathbf{M_{it}}, \theta_i; \boldsymbol{\gamma}))}\right)^{P_{it}} \left(\frac{1}{1 + \exp(f^P(\mathbf{W_t}, X_{it}, \mathbf{M_{it}}, \theta_i; \boldsymbol{\gamma}))}\right)^{(1-P_{it})},$$

with P_{it} being an indicator for participation.

Application decision. If an individual participates in a given period, he is allowed to submit applications for up to two housing units and their associated lotteries. Let the expected benefit of applying to lottery j in period t be given by:

$$b_{ijt} = \pi_{jt} f^C(W_{jt}, X_{it}, M_{ijt}, \theta_i) + \epsilon_{ijt}^C$$

where π_{jt} is the probability of receiving an offer for house j at time t if the individual applies. Let $B_{it} = \{b_{i1t}, .., b_{iJ_t t}\}$ be the set of expected benefits.

I assume that ϵ_{ijt}^C are i.i.d. across individuals and choices. This assumption simplifies the portfolio choice problem as individuals will choose to apply to the two choices with the highest expected benefit b_{ijt}. Suppose an individual chooses to apply to j' and j''. Since we do not observe how they would rank j' and j'', we know that:

$$b_{ij't} > b_{ijt} \ \forall j \in \mathcal{J}_t \setminus \{j', j''\}$$

and

$$b_{ij''t} > b_{ijt} \ \forall j \in \mathcal{J}_t \setminus \{j', j''\}$$

Note that the two inequalities above provide the same information as data generated in a situation where an individual faces two multinomial choice problems, one in which they choose j' from $\mathcal{J}_t \setminus j''$ and another in which they choose j'' from $\mathcal{J}_t \setminus j'$ and, under the assumptions made above, the likelihood contributions will be equivalent. If the error terms are assumed to be type-1 extreme value (Gumbel) distributed, then the contribution of individual i to the period t likelihood is equivalent to two multinomial logistic regressions each with $J_t - 1$ choices. This simplifies the estimation, yet two challenges persist. First, the choice sets shaping these two multinomial logistic regressions will differ depending on which two lotteries the individual applies to. Second, the housing units that are available differ across application periods. To address this, I use a hedonic model of house value, where each house is characterized fully by its characteristics W_{jt}, there is no house-specific intercept and individual characteristics $\{X_{it}, \theta_i\}$ only enter the application likelihood by allowing the value of observable housing characteristics to vary by individuals' characteristics. Multinomial logits with these restrictions are commonly referred to 'conditional logits". Conditional logits have the property that the set of choices does not need to be the same over time or across individuals.

Thus, individual i's likelihood contribution in period t is:

$$L_{it}^{C}(\boldsymbol{W}_t, X_{it}, \mathbf{M_{it}}, \theta_i; \boldsymbol{\beta}) = \prod_{j \in \{\mathcal{J} \setminus j''\}} 1\{j = j'\} \frac{\exp(\pi_{jt} f^C(W_{jt}, X_{it}, M_{ijt}, \theta_i; \boldsymbol{\beta}))}{\sum_{k \in \{\mathcal{J} \setminus j''\}} \exp(\pi_{kt} f^C(W_{kt}, X_{it}, M_{ikt}, \theta_i; \boldsymbol{\beta}))} \times$$
$$\prod_{j \in \{\mathcal{J} \setminus j'\}} 1\{j = j''\} \frac{\exp(\pi_{jt} f^C(W_{jt}, X_{it}, M_{ijt}, \theta_i; \boldsymbol{\beta}))}{\sum_{k \in \{\mathcal{J} \setminus j'\}} \exp(\pi_{kt} f^C(W_{kt}, X_{it}, M_{ikt}, \theta_i; \boldsymbol{\beta}))}$$

where this likelihood is set to 1 if the individual does not participate in the lottery.

Acceptance decision. Suppose individual i wins lottery j in period t. I model their acceptance decision as an index crossing model where the person accepts if:

$$f^A(W_{jt}, X_{it}, \mathbf{M_{ijt}}, \theta_i) + \epsilon_{it}^A > 0.$$

This assumption produces the likelihood contribution for the acceptance decision:

$$L_{it}^{A}(\mathbf{W_t}, X_{it}, \mathbf{M_{it}}, \theta_i; \boldsymbol{\alpha}) = \left(\frac{\exp(f^A(\mathbf{W_t}, X_{it}, \mathbf{M_{it}}, \theta_i; \boldsymbol{\alpha}))}{1 + \exp(f^P(\mathbf{W_t}, X_{it}, \mathbf{M_{it}}, \theta_i; \boldsymbol{\alpha}))} \right)^{A_{it}} \left(\frac{1}{1 + \exp(f^P(\mathbf{W_t}, X_{it}, \mathbf{M_{it}}, \theta_i; \boldsymbol{\alpha}))} \right)^{(1 - A_{it})}$$

where this likelihood is set to 1 if an individual does not win a lottery.

5.2 Estimation

The model is estimated via maximum likelihood. The total likelihood can be written as:

$$\int \prod_{i \in N} \prod_{t \in T} L_{it}^{A}(\boldsymbol{W}_t, X_{it}, \mathbf{M_{it}}, \theta_i; \boldsymbol{\alpha}) L_{it}^{C}(\boldsymbol{W}_t, X_{it}, \mathbf{M_{it}}, \theta_i; \boldsymbol{\beta}) L_{it}^{P}(\boldsymbol{W}_t, X_{it}, \mathbf{M_{it}}, \theta_i; \boldsymbol{\gamma}) dF_\theta$$

where the interdependence between choices within period and choices over time is captured through θ and π_{jt} can be estimated non-parametrically from the data. I assume that θ can take on one of M unique values, where a discrete distribution such as this can approximate any distribution of unobserved heterogeneity given a sufficient number of values M.

I estimate the model using the Expectation-Maximization (EM) algorithm, which uses an iterative procedure to integrate out θ as described in Arcidiacono and Jones (2003) and Train (2008). In the EM algorithm, the likelihoods for the participation, application, and acceptance decisions are log separable conditional on the estimated posterior probabilities of θ_i which further simplifies estimation.

5.3 Identification

The key identifying assumption is that weekly participation decisions, application decisions, and acceptance decisions are conditionally independent when conditioning on the sets of observable W_{jt} and X_{it} as well as the unobserved latent heterogeneity θ_i. These assumptions are similar to those used in Arcidiacono and Jones (2003) and Arcidiacono and Miller (2011). As discussed in Kasahara and Shimotsu (2009) access to many choices per agent, combined with the fact that factors have different loadings in the participation, application, and acceptance decisions, allows for the factor shares and loadings to be non-parametrically identified.

While not strictly necessary for identification, the institutional setting also offers exogenous variation in the probability a given applicant applies to a given lottery. As discussed in the next subsection, choice sets vary randomly across application period providing exogenous variation over time in the value of applying to a given unit relative to other units. This source of variation (combined with the random nature of lottery-based allocation) provides additional variation in application and acceptance decisions. Given sufficient variation in these instruments, identification can alternatively be established without conditional independence assumptions as in Heckman and Navarro (2007).

5.4 Variation in differential distance as an instrument for application decisions

Each week approximately 12 houses are offered through the lottery mechanism. The set of offered housing varies in its characteristics and location, and the set of units that are offered each week is plausibly exogenous. This natural variation in the choice set means that some weeks people will choose to apply to housing in better or worse neighborhoods based on what is offered. Individuals are also more likely to apply to housing that is closer to their current residence – a result that may be driven by information about neighborhoods closer to current residence, or a lower cost of visiting the location of the house being offered. As a result, an individual may be more likely to apply to a house in a specific neighborhood with a given set of characteristics based on its relative distance from other units offered that week. This suggests that a natural source of exogenous variation in the probability of applying to a given house is its relative distance to current residence compared to the other offerings that week. For individual i and lottery j in week t, I define differential distance $M_{i,j,t}$ as the distance from current residence to house j minus the leave-j-out average distance from current residence, i.e. $M_{i,j,t} = d_{i,j,t} - \bar{d}_{i,-j,t}$.

Table 3 evaluates the exclusion and relevance of differential distance as a source of exogenous variation in application choice. The first column shows that differential distance $M_{i,j,t}$ is not related to demographics (panel I), the housing and neighborhood characteristics of the applicant (panel II), or baseline labor market outcomes (panel III). A concern that is sometimes raised about instruments based on differential distance is that they may violate the exclusion restriction, because individuals with a strong preference for housing in high-income neighborhoods may already be living near such neighborhoods. The exogeneity of the composition of the choice set that is realized each week arguably mitigates this concern. The second column of Table 3 shows a reduced-form first stage of if individual i chooses to apply to lottery j in week t (i.e. $C_{i,j,t} = 1$) regressed on differential distance and the standard set of covariates and lottery fixed effects. Differential distance is strongly statistically significant and shows that an applicant is 2 percentage points less likely to apply to a lottery for every kilometer increase in differential distance $M_{i,j,t}$.

5.5 Results from the choice model

This section reports results from the choice model. First I discuss the role of latent heterogeneity (θ) in the participation, application, and acceptance decisions. Second, I compare the average treatment effects

Table 3: Testing for exogeneity and relevance of M_{ijt}.

	Dependent variables	
	M_{ijt} (1)	C_{ijt} (2)
M_{ijt}		-0.0117***
		(0.0000)
I. *Demographics*		
Male	0.0000	
	(0.0042)	
Age	0.0000	
	(0.0003)	
Foreign born	-0.0000	
	(0.0056)	
Second gen. immigrant	0.0000	
	(0.0055)	
HH size	0.0000	
	(0.0002)	
II. *Housing and neighborhood characteristics at time of application*		
NH average income	0.0000	
	(0.0000)	
NH fraction non-western	0.0000	
	(0.0160)	
NH fraction low-income	-0.0000	
	(0.0590)	
III. *Labor market and residential mobility history*		
Personal labor income, $t-1$	0.0000	
	(0.0000)	
Personal total income, $t-1$	0.0000	
	(0.0000)	
HH labor income per adult, $t-1$	0.0000	
	(0.0000)	
HH total income per adult, $t-1$	-0.0000	
	(0.0000)	
Employed, $t-1$	-0.0000	
	(0.0059)	
Residential move, $t-1$	0.0000	
	(0.0047)	
F-stat for joint test	0.0000	248,200
p-value	1	0.0000
Application period FE	✓	✓
N	6,456,905	6,456,905

Notes: This table provides evidence on the exogeneity and relevance of the differential distance instrument. M_{ijt} is the relative distance of lottery j form individual i given the other lotteries in week t. C_{ijt} is an indicator of if individual i chose to apply to lottery j in week t. The first column shows residual distance of lotteries regressed on demographics, neighborhood characteristics and labor market outcomes. The second column shows the application decision regressed on residual distance, controlling for the full set of covariates. All regressions include application-period fixed effects. As an additional check on the exclusion restriction, running the regression of C_{ijt} on M_{ijt} with the covariates used in the balance regression as controls shows that the coefficient is stable, which they should be if the pre-determined outcomes are indeed uncorrelated with the instrument. The coefficient for the regression with controls is -0.0117 ($p<.01$). The F-stat for the full projected model is 16,150. * = $p<.1$; ** = $p<.05$; *** = $p<.01$.

estimated using the choice model to the IV and OLS estimates. Third, I use the choice model to estimate several relevant treatment effects.

Figure 9 documents the role of the latent heterogeneity in choice behaviors. The model estimates three types, with the first type accounting for 25%, the second type accounting for 57%, and the third type accounting for 19%. Choice behavior is found to vary substantially based on latent type. Type 2 individuals participate only in 7% of application periods, while types 1 and 3 participate more than 30% of the time. Among the frequently-participating types, type 1 is less likely to apply to housing in high-income neighborhoods, while type 3 is more likely to apply to those lotteries. Overall, behaviors differ substantially based on the identified latent heterogeneity suggesting it may play an important role in explaining the benefits of housing.

Figure 9: Characterizing the latent types.

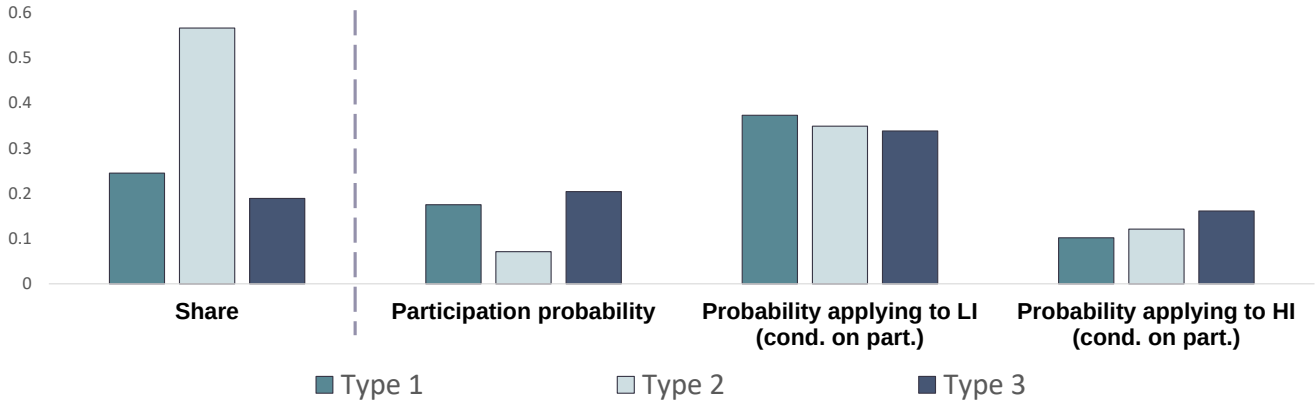

Notes: This figure shows selected results based on the estimated choice model. The left-most cluster of bars shows the estimated shares of each latent type within the population of participants. The other clusters show types' participation probability and probabilities of applying to high- and low-income neighborhoods, conditional on participating. A more detailed description of the latent types is available in Table 9 in Appendix A.5.

Next, I use the choice model to directly estimate the average treatment effect. Table 4 compares the OLS and TOT (ever-offer) estimates of the treatment on the treated from Section 4. The three column groups break down results by 'all applicants' applicants from a 'low income origin' neighborhood and applicants from a 'high income origin' neighborhood. The three subsections break down the results for all lotteries, lotteries in high-income neighborhoods, and lotteries in low-income neighborhoods. Overall, I find that the ATE estimates produced from the model are very similar to the Treatment-on-the-treated estimates produced by the ever-offer IV regression. These results provide additional evidence that largely

positive IV results for high income destination neighborhoods (and largely negative IV results for low income destination neighborhoods) are largely due to differences in returns and not due to selection.

Table 5 calculates various treatment effects using the choice model. Specifically, the table reports the ATE, ATE for those from low-income neighborhoods, and ATE for those from high-income neighborhoods as in Table 4. Then, the table also calculates the treatment on the treated (TT), treatment on the untreated (TUT), the average treatment effect of offering a house to a random participant, and the ATE conditional on latent type θ. I find that for lotteries in high-income neighborhoods, the TT tends to be smaller than the TUT which suggests that housing in high-income neighborhoods would increase labor market outcomes of those in low-income neighborhoods and that the IV results were not driven by selection.

6 Conclusion

This paper analyzes the effect of Europe's largest public housing program on households' socio-economic outcomes, including labor market outcomes, neighborhood and housing quality, and public assistance receipt. To the best of my knowledge, this is the first paper to study such a large scale program while addressing confounding due to self-selection. Moreover, I am able to use random variation in the quality and location of housing units to analyze how quality of housing assistance may mediate the efficacy of the program.

The first part of the paper exploits lotteries embedded in the Amsterdam centralized assignment mechanism to identify and estimate the effect of moving into public housing on those who receive assistance. Consistent with the past literature, I find that, on average, receiving public housing is associated with lower earnings, reduced labor supply, and increased public assistance. Yet I find that there is substantial heterogeneity in returns depending on the neighborhood in which housing is offered, with large positive labor market effects for those who receive housing in high-income neighborhoods, and large negative effects for those who receive housing in low-income neighborhoods. Moreover, I find that positive effects to lotteries in high-income neighborhoods are driven by 'upward' moves by individuals previously living in low- or middle-income neighborhoods.

The results from the first part of the paper only describe treatment-on-the treated parameters that are conditional on applying to and accepting an offer from a lottery. Thus, they are not sufficient to

determine whether the heterogeneous labor market responses across neighborhoods would generalize to the larger population of applicants to any of the lotteries. The second part of the paper uses a choice model to characterize individuals' application decisions and to recover the distribution of heterogeneity driving selection into and returns from lotteries. Using the model, I am then able to estimate average treatment effects by the type of house being offered and by sub-population. Specifically, I leverage panel data on individuals' lottery choices combined with variation in the available lotteries to evaluate whether individuals' select into lotteries based on their potential labor market gains. While I find that latent heterogeneity plays an important role in application behavior and earnings, there is little evidence of selection on gains. Average treatment effects are very similar to the lottery-specific estimates for those who apply from the first part of the paper. This implies that, on average, housing in high-income neighborhoods would improve the labor market outcomes of poorer individuals. Moreover, the model suggests that individuals unlikely to apply to lotteries in high-income neighborhoods would be likely to accept housing in high-income neighborhoods, if offered.

Across methodologies, my results suggest that policies providing access to higher quality neighborhoods would generate positive labor-market outcomes. For Amsterdam, a policy that targets high-quality housing to high-return individuals could result in earnings gains and potentially reduce public assistance receipt. More broadly, my findings suggest that policies promoting access to public housing could increase labor force participation for the poor, a finding that contrasts with the general consensus on how public housing affects labor supply.

Table 4: Comparing treatment effects.

		All applicants			Low income orig.			High income orig.		
	Base	OLS	IV	ATE	OLS	IV	ATE	OLS	IV	ATE
I. *All destination neighborhoods*										
HH labor income	17,671	-2,252***	-2,320***	-2,656	-2,829**	-2,150	-1,453	-2,831**	-4,006***	-1,695
		(651.3)	(878.1)		(1,382)	(1,478)		(1,345)	(1,446)	
Individual labor income	17,000	-1,771***	-1,463*	-2,337	-1,852	-1,181	-1,354	-2,126	-3,522**	-980
		(641.9)	(850.0)		(1,365)	(1,455)		(1,380)	(1,476)	
Employment	0.58	-0.0440***	-0.0412	-0.042	-0.028	-0.018	0.0284	-0.084**	-0.087**	0.0514
		(0.0171)	(0.0227)		(0.0336)	(0.0369)		(0.0361)	(0.0395)	
II. *Low income destination NHs*										
HH labor income	17,489	-4,437***	-4,127***	-3,874	-7,465***	-6,687***	-4,269	-4,423**	-6,256***	-867.5
		(1,145)	(1,431)		(2,227)	(2,334)		(2,169)	(2,377)	
Individual labor income	16,847	-3,588***	-3,223**	-2,858	-5,548**	-4,867**	-3,321	-2,788	-5,014*	1,332
		(1,116)	(1,462)		(2,337)	(2,424)		(2,452)	(2,671)	
Employment	0.581	-0.091***	-0.099***	-0.0743	-0.094*	-0.055	-0.0660	-0.157***	-0.135**	-0.0733
		(0.0287)	(0.0371)		(0.0520)	(0.0552)		(0.0590)	(0.0656)	
III. *High income destination NHs*										
HH labor income	18,827	4,770**	5,318*	3,212	10,840***	14,880***	7,765	-2,334	-2,074	-3,383
		(2,139)	(2,842)		(3,810)	(5,426)		(3,476)	(3,595)	
Individual labor income	18,249	4,431**	4,363*	2,854	8,745***	12,980***	6,904	-1,749	-2,086	-3,656
		(2,049)	(2,637)		(3,201)	(4,980)		(3,339)	(3,381)	
Employment	0.6135	0.1104**	0.1234**	0.1424	0.332***	0.364***	0.2874	-0.060	-0.021	-0.0124
		(0.046)	(0.059)		(0.0540)	(0.0839)		(0.0889)	(0.0951)	
Lottery FE		✓	✓		✓	✓		✓	✓	

Notes: Figure reports the OLS estimates, treatment on the treated estimates (based on the 'ever offer' instrument), and the average treatment effect (based on the decision model) for various labor market outcomes. The 'All applicants' columns include the full population of applicants. The 'Low income orig.' columns include those originating in low-income neighborhoods. The 'High income orig.' columns include those originating in high-income neighborhoods. The first sub-table considers lotteries in all destination neighborhoods while the second sub-table shows results only for lotteries in high-income neighborhoods and the third sub-table shows results only for lotteries in low-income neighborhoods. IV and OLS specifications control for the lagged outcome and lottery fixed effects. * = $p<.1$; ** = $p<.05$; *** = $p<.01$. 'Base' is the sample mean for those who do not receive an offer, at the time of application.

Table 5: Treatment effects from the choice model.

	ATE	*ATE (low orig)*	*ATE (high orig)*	*TT*	*TUT*	*ATE (offer)*	*ATE ($\theta = 1$)*	*ATE ($\theta = 2$)*	*ATE ($\theta = 3$)*	*ATE (observables)*
I. *All destination neighborhoods*										
HH labor income	-2,656	-1,453	-1,695	-2,639	-2,656	-1,426	-9769	-3,615	-3,885	-2,736
Indiv. labor income	-2,337	-1,354	-980	-2,375	-2,337	-1,255	-705.4	-3,362	-3,042	-2,469
Employment	-0.0420	0.0284	0.0514	-0.041	-0.041	-0.022	0.000	-0.070	-0.056	-.041
II. *Low income destination NHs*										
HH labor income	-3,874	-4,269	-867.5	-3,953	-3,874	-2,017	-2,631	-5,889	-3,132	-4,391
Indiv. labor income	-2,858	-3,321	1,332	-1,719	-5,227	-1,638	-1,453	-2,858	-3,126	-3,521
Employment	-0.0743	-0.0660	-0.0733	-0.080	-0.074	-0.039	-0.024	-0.137	-0.064	-0.074
III. *High income destination NHs*										
HH labor income	3,212	7,765	-3,383	3,172	3,212	1,671	11,930	5,501	-8,359	4,236
Indiv. labor income	2,854	6,904	-3,656	2,721	2,854	1,459	1,185	4,123	-7,130	3,704
Employment	0.1424	0.2874	-0.0124	0.109	0.142	0.080	0.278	0.107	0.037	0.142

Notes: All numbers are based on estimates from the choice model. 'ATE' is the estimated average treatment effect weighting all participants and lotteries equally. 'ATE (low orig)' is similar to 'ATE', but only considers applicants originating from low-income neighborhoods, while 'ATE (high orig)' only considers applicants originating in high-income neighborhoods. 'TT' is Treatment on the Treated and is the estimated average treatment effects conditional on appling, while 'TUT' is the estimated average treatment effect conditional on not applying. 'ATE (offer)' is the average treatment effect of offering the unit, but not requiring the individual to accept the offer. 'ATE ($\theta = j$)' corresponds to the ATE conditional on the most likely posterior unobserved typ of the individual being j. 'ATE (observed)' is the estimated average treatment effect not accounting for the latent heterogeneity θ.

References

Abdulkadiroglu, Atila, Joshua D. Angrist, Peter Hull, Parag A. Pathak, and Christopher R. Walters, "Charters without Lotteries: Testing Takeovers in New Orleans and Boston," *American Economic Review*, 2016, *106* (7), 1878–1920.

_ , _ , Susan M. Dynarski, Thomas J. Kane, and Parag A. Pathak, "Accountability and Flexibility in Public Schools: Evidence from Boston's Charters and Pilots," *Quarterly Journal of Economics*, 2011, *126* (2), 699–748.

_ , _ , Yusuke Narita, and Parag A. Pahtak, "Research Design Meets Market Design: Using Centralized Assignment for Impact Evaluation," *Working paper*, 2016.

Angrist, Joshua D., Peter Hull, Parag A. Pathak, and Christopher R. Walters, "Leveraging Lotteries for School Value-Added: Testing and Estimation," *Quarterly Journal of Economics*, Forthcoming.

_ , Susan M. Dynarski, Thomas J. Kane, Parag A. Pathak, and Christopher R. Walters, "Inputs and impacts in charter schools: Kipp lynn," *American Economic Review*, 2010, *100* (2), 239–243.

Arcidiacono, Peter and John Bailey Jones, "Finite Mixture Distributions, Sequential Likelihood and the EM Algorithm," *Econometrica*, 2003.

_ and Robert A. Miller, "Conditional Choice Probability Estimation of Dynamic Discrete Choice Models with Unobserved Heterogeneity," *Econometrica*, 2011, *79* (6), 1823–1867.

Barnhardt, Sharon, Erica Field, and Rohini Pande, "Moving to Opportunity or Isolation? Network Effects of a Randomized Housing Lottery in Urban India," *American Economic Journal: Applied Economics*, 2016, *9* (1), 1–32.

Carneiro, Pedro, Karsten T. Hansen, and James J. Heckman, "Estimating Distributions of Treatment Effects with an Application to the Returns to Schooling and Measurement of the Effects of Uncertainty on College," *International Economic Review*, 2003, *44* (2).

CBS Statline. Centraal Bureau voor de Statistiek (Statistics Netherlands) public database. URL: http://statline.cbs.nl/.

Chabrier, Julia, Sarah Cohodes, and Philip Oreopoulos, "What Can We Learn from Charter School Lotteries?," *Journal of Economic Perspectives*, 2016, *30* (5), 57–84.

Chetty, Raj, 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.

Chyn, Eric, "Moved to Opportunity: The Long-Run Effect of Public Housing Demolition on Labor Market Outcomes of Children," *Working paper*, 2016.

City of New York, *Housing 2.0* 2016. Report by NYC Dept. of Housing Preservation & Development. URL: https://www1.nyc.gov/assets/hpd/downloads/pdf/about/hny-2.pdf.

Cohodes, Sarah, "Teaching to the Student: Charter School Effectiveness in Spite of Perverse Incentives," *Education Finance and Policy*, 2016, *11* (1), 1–42.

Collinson, Robert, Ingrid Gould Ellen, and Jens Ludwig, "Low Income Housing Policy," in "Economics of Means-Tested Transfer Programs in the United States," Vol. 2, University of Chicago Press, 2015.

Cunha, Flavio and James J. Heckman, "Formulating, Identifying and Estimating the Technology of Cognitive and Noncognitive Skill Formation," *Journal of Human Resources*, 2008, *43* (4), 738–782.

Currie, Janet and Aaron Yelowitz, "Are public housing projects good for kids?," *Journal of Public Economics*, 2000, *75* (1).

_ **and Firouz Gahvari**, "Transfers in Cash and In-Kind: Theory Meets the Data," *Journal of Economic Literature*, 2008, *46* (2), 333–383.

Davis, Matthew and Blake Heller, "'No excuses' Charter Schools and College Enrollment: New Evidence from a High School Network in Chicago," *Working paper*, 2017.

de Chaisemartin, Clement and Luc Behaghel, "Next please! Estimating the effect of treatments allocated by randomized waiting lists," *Working paper*, 2017.

del Pero, Angelica Salvi, Willem Adema, Valeria Ferraro, and Valerie Frey, "Policies to promote access to good-quality affordable housing in OECD countries," *OECD Social, Employment and Migration Working Papers 176*, 2016.

Deming, David J., "Better Schools, Less Crime?," *The Quarterly Journal of Economics*, 2011, *126* (4), 2063–2115.

_ , "Using School Choice Lotteries to Test Measures of School Effectiveness," *American Economic Review*, 2011, *104* (5), 406–411.

_ **, Justine S. Hastings, and Douglas O. Kane Thomas J. Staiger**, "School Choice, School Quality, and Postsecondary Attainment," *American Economic Review*, 2014, *104* (3), 991–1013.

Desmond, Matthew, "Eviction and the Reproduction of Urban Poverty," *American Journal of Sociology*, 2012, *118* (1), 88–133.

Dobbie, Will and Roland G. Fryer Jr., "Are High-Quality Schools Enough to Increase Achievement Among the Poor? Evidence from the Harlem Children's Zone," *American Economic Journal: Applied Economics*, 2011, *3* (3), 158–187.

_ **and** _ , "Getting Beneath the Veil of Effective Schools: Evidence from New York City," *American Economic Journal: Applied Economics*, 2013, *5* (4), 28–60.

_ **and** _ , "The Medium-Term Impacts of High-Achieving Charter Schools," *Journal of Political Economy*, 2015, *123* (5), 985–1037.

Frölich, Markus, "Nonparametric IV Estimation of Local Average Treatment Effects with Covariates," *Journal of Econometrics*, 2007, *139* (1), 35–75.

Gahvari, E., "In-kind Transfers, Cash Grants and Labor Supply," *Journal of Public Economics*, 1994, *55*, 495–504.

Gubits, Daniel, Marybeth Shinn, Michelle Wood, Stephen Bell, Samuel Dastrup, Claudia Solari, Scott Brown, Debi McInnis, Tom McCall, and Utsav Kattel, "Family Options Study: 3-Year Impacts of Housing and Services Interventions for Homeless Families," *SSRN Working Paper*, 2016.

Hastings, Justine S., Chris A. Neilson, and Seth D. Zimmerman, "The effect of school choice on intrinsic motivation and academic outcomes," *NBER Working Paper 18324*, 2012.

Heckman, James J., "Sample Selection Bias as a Specification Error," *Econometrica*, 1979, *47*, 153–161.

_ **and Salvador Navarro**, "Dynamic Discrete Choice and Dynamic Treatment Effects," *Journal of Econometrics*, 2007, *136* (2), 341–96.

_ **, John Eric Humphries, and Gregory Veramendi**, "Dynamic Treatment Effects," *Journal of Econometrics*, 2016, *191* (2), 276–292.

_ **, Jora Stixrud, and Sergio Urzua**, "The Effects of Cognitive and Noncognitive Abilities on Labor Market Outcomes and Social Behavior," *Journal of Labor Economics*, 2006, *24* (3).

Hoxby, Caroline M. and Sonali Murarka, "Charter Schools in New York City: Who Enrolls and How They Affect Their Students' Achievement," *NBER Working Paper*, 2009.

Humphries, John Eric, Nick Mader, Daniel Tannenbaum, and Winnie van Dijk, "Does Eviction Cause Poverty? Quasi-Experimental Evidence from Cook County, IL," *Working paper*, 2018.

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.

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

JCHS, *America's Rental Housing: Expanding Options for Diverse and Growing Demand* December 2015. Report by the Joint Center for Housing Studies of Harvard University. URL: http://www.jchs.harvard.edu/americas-rental-housing.

Kasahara, Hiroyuki and Katsumi Shimotsu, "Nonparametric Identification of Finite Mixture Models of Dynamic Discrete Choices," *Econometrica*, 2009, *77* (1), 135–175.

Kirkeboen, Lars, Edwin Leuven, and Magne Mogstad, "Field of Study, Earnings, and Self-Selection," *Working paper*, 2016.

Kling, Jeffrey, Jeffrey Liebman, Lawrence Katz, and Lisa Sanbonmatsu, "Moving to Opportunity and Tranquility: Neighborhood Effects on Adult Economic Self-Sufficiency and Health from a Ran domized Housing Voucher Experiment," *Princeton University Working Paper*, 2004.

Leonesio, M., "In-Kind Transfers and Work Incentives," *Journal of Labor Economics*, 1988, *4*, 515–529.

Mayor of London, *London Housing Strategy* May 2018. Report published by the Greater London Authority. URL: https://www.london.gov.uk/sites/default/files/2018_lhs_london_housing_strategy.pdf.

Mills, Gregory, Daniel Gubits, Larry Orr, David Long, Judie Feins, Bulbul Kaul, Michelle Wood, Amy Jones & Associates, Cloudburst Consulting, and the QED Group, *The Effects of Housing Vouchers on Welfare Families* 2006. Washington, DC: US Department of Housing and Urban Development, Office of Policy Development and Research.

Moffitt, Robert A., "Welfare Programs and Labor Supply," in "Handbook of Public Economics," Vol. 4, Elsevier Science B.V., 2002.

Murray, M., "A Reinterpretation of the Traditional Income-Leisure Model, with Application to In-Kind Subsidy Programs," *Journal of Public Economics*, 1980, *14* (69-81).

OECD, *OECD Affordable Housing Database* December 2016. URL: http://www.oecd.org/social/affordable-housing-database.htm.

Shroder, Mark, "Does Housing Assistance Perversely Affect Self-Sufficiency? A Review Essay," *Journal of Housing Economics*, 2002, *11*, 381–417.

Train, Kenneth E, "EM Algorithms for nonparametric estimation of mixing distributions," *Journal of Choice Modelling*, 2008, *1* (1), 40–69.

U.S. House of Representatives Committee on Ways and Means, *Green Book* 2016. Appendix A: Federal Benefits and Services for People with Low Income. URL: https://greenbook-waysandmeans.house.gov/2016-green-book/appendix-a-federal-benefits-and-services-for-people-with-low-income.

Walters, Christopher R., "The Demand for Effective Charter Schools," *Journal of Political Economy*, 2017.

A Tables and figures

A.1 Online platform

Figure 10: Online platform for housing allocation.

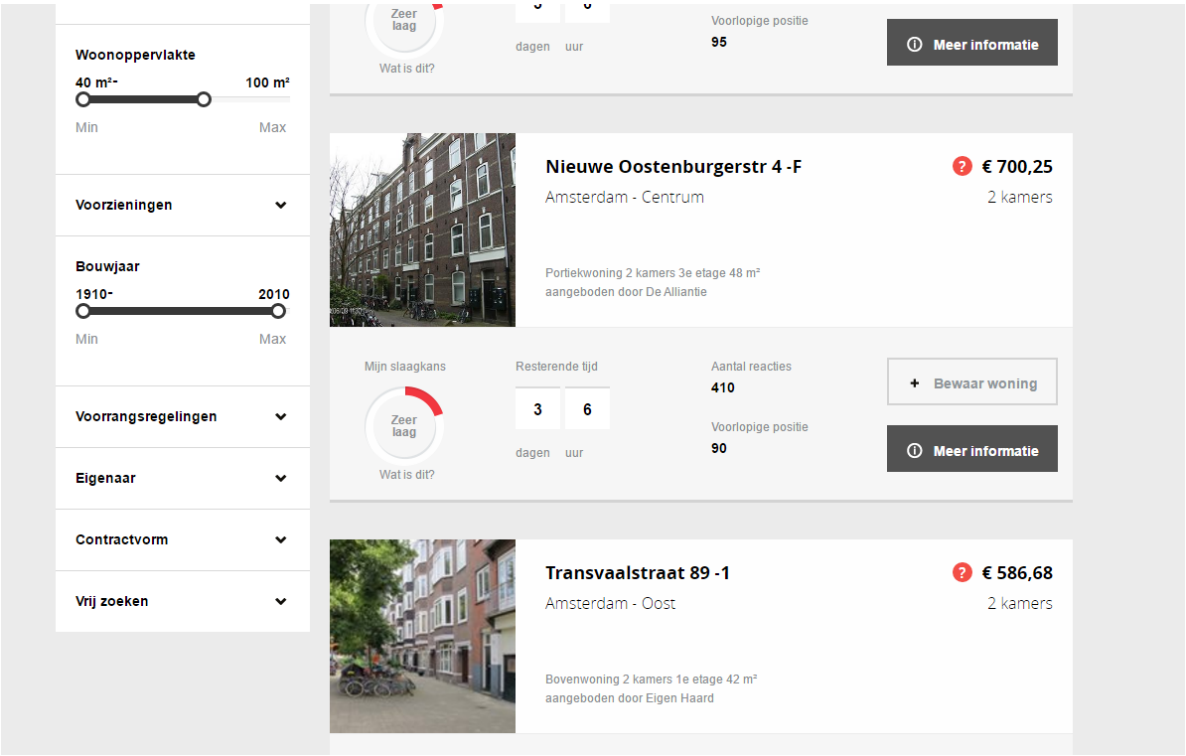

Notes: This figure shows a screenshot of the online platform's search interface. Each listing shows the address of the housing unit, the rental price, an indication of the probability of successfully applying ('Mijn slaagkans', bottom left), time left in this application period (three days and six hours), and the number of applications currently received (410). Clicking on 'Meer informatie' (more information) leads to a page with additional details on the housing unit, an example of which is shown in Figure 11.

Figure 11: Details for a listing on the online platform.

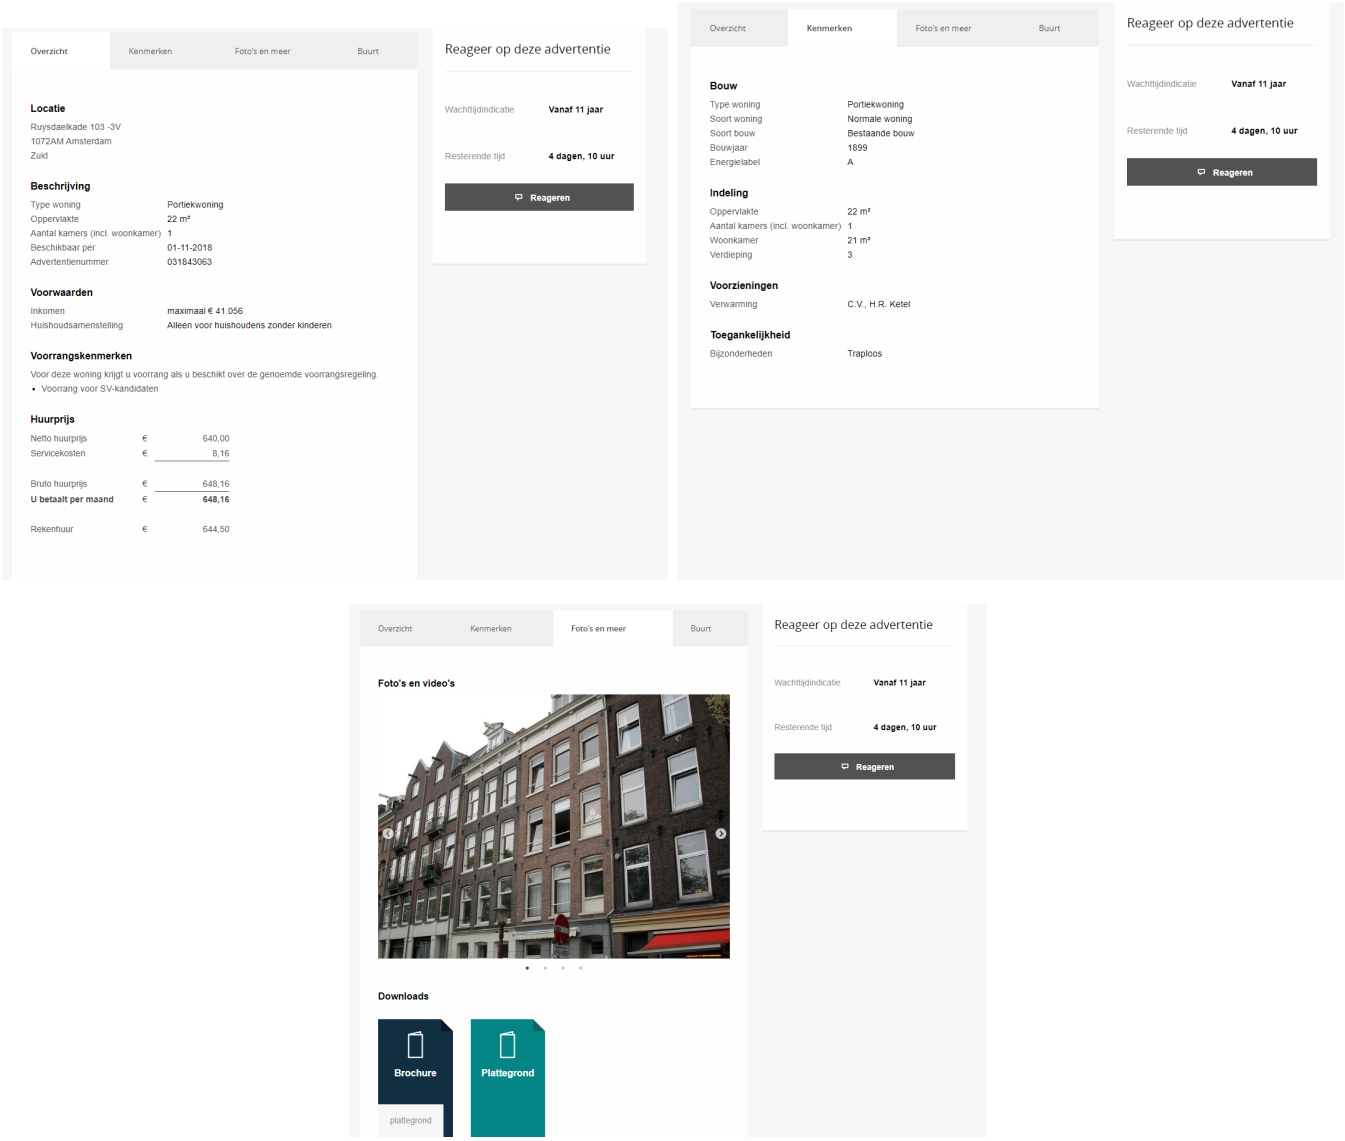

Notes: These figures show an example listing on the online platform. Details provided in the listing include address, size of the unit, number of rooms, date it will become available, monthly rent price, year of construction, floor, type of heating system, and a downloadable floor plan.

A.2 Substitution patterns

Table 6: Substitution patterns.

	All lotteries	Low-inc. NH	Middle-inc. NH	High-inc. NH
Fraction receiving wait list housing 2013-2014	7%	7%	8%	7%
Fraction receiving wait list housing 2015-2016	6%	6%	6%	6%
Fraction receiving lottery housing 2013-2014	4%	4%	4%	3%
Fraction receiving lottery housing 2015-2016	4%	4%	3%	3%

Notes: Table entries are averages across lotteries for individuals in the analysis sample.

A.3 Effects for high-income neighborhoods

Table 7: Effects of public housing receipt: two-year impacts (high-income destination neighborhoods).

	Base	OLS (1)	IV (Z_{ij}) (2)	N
I. *Income and employment*				
HH labor income	18,827	4,770**	5,318*	55,623
		(2,132)	(2,829)	
HH total income	23,168	3,947**	5,026**	55,623
		(1,942)	(2,564)	
HH disposable income	20,312	3,235**	4,375**	55,623
		(1,642)	(2,214)	
Individual labor income	18,250	4,431**	4,363*	58,579
		(2,043)	(2,624)	
Individual total income	22,454	3,628*	4,179*	58,579
		(1,862)	(2,361)	
Employment	0.6135	0.1104**	0.1234**	55,677
		(0.046)	(0.059)	
II. *Housing quality and rent expenditure*				
House value (per m^2)	2,869	635.1***	734.4***	57,914
		(129.5)	(165.9)	
Monthly rent expenditure (net of subsidy)	444.6	107**	52.9	58,238
		(49.26)	(67.7)	
Cum. two-year residential move rate	0.502	-0.243***	-0.247***	58,579
		(0.0558)	(0.0736)	
III. *Neighborhood characteristics*				
NH average income	38,564	5,181***	6,583***	58,579
		(747.9)	(1,329)	
NH fraction non-western	0.212	-0.094***	-0.0998***	58,579
		(0.0129)	(0.01864)	
Social distance	0.430	0.002	-0.015	55,348
		(0.0203)	(0.0243)	
Distance to city center (m)	11,811	-4,431***	-5,407***	58,579
		(638.7)	(1,627)	
IV. *Fiscal cost*				
Public assistance receipt	0.184	-0.037	-0.030	55,677
		(0.0307)	(0.0413)	
Fiscal externality	-1,476	1,284	702.0	58,579
		(789.3)	(1090)	
Lottery FE		✓	✓	

Notes: Individual income is income from paid or self-employment before taxes and transfers, in 2015 US dollars. HH (household) income is average labor income for adults in the household, excluding adult children. Employment is an indicator for paid employment or self-employment being the individual's main source of income. Social distance is the fraction of neighborhood residents that are not of the applicant's ethnic group. Fiscal externality is individual taxes paid net of non-housing transfers. 'Base' is the sample mean for those who do not receive an offer, at the time of application. Regressions include controls for baseline outcomes. Standard errors are clustered at the individual level. * = $p<.1$; ** = $p<.05$; *** = $p<.01$.

A.4 Effects for low-income neighborhoods

Table 8: Effects of public housing receipt: two-year impacts (low-income destination neighborhoods).

	Base	OLS (1)	IV (Z_{ij}) (2)	N
I. *Income and employment*				
HH labor income	17,489	-4,437***	-4,127***	153,224
		(1,139)	(1,427)	
HH total income	22,790	-2,976***	-2,685**	153,224
		(957.3)	(1240)	
HH disposable income	20,120	-2,307***	-2,219**	153,224
		(772.4)	(1017)	
Individual labor income	16,847	-3,588***	-3,223**	163,930
		(1,166)	(1,462)	
Individual total income	22,030	-2,167**	-1,666	163,930
		(990.9)	(1,265)	
Employment	0.581	-0.091***	-0.099***	153,440
		(0.0287)	(0.0371)	
II. *Housing quality and rent expenditure*				
House value (per m^2)	2,597.8	-355.3***	-418.3***	162,960
		(36.89)	(58.94)	
Monthly rent expenditure (net of subsidy)	455.35	-34.63*	-54.64*	162,909
		(20.08)	(29.57)	
Cum. two-year residential move rate	0.52	-0.2639***	-0.3320***	163,930
		(0.0274)	(0.0392)	
III. *Neighborhood characteristics*				
NH average income	35,905	-5,822***	-6,039***	163,930
		(320.2)	(544.4)	
NH fraction non-western	0.263	0.1654***	0.1544***	163,930
		(0.0104)	(0.0133)	
Social distance	0.540	0.0647***	0.0580***	152,461
		(0.0128)	(0.0151)	
Distance to city center (m)	11,405	-4,397***	-4,235***	163,930
		(685.0)	(990.7)	
IV. *Fiscal cost*				
Public assistance receipt	0.248	0.0893***	0.0659**	153,440
		(0.0251)	(0.039)	
Fiscal externality	-2,664	-1,771***	-1,571**	163,930
		(498.4)	(645.6)	
Lottery FE		✓	✓	

Notes: Individual income is income from paid or self-employment before taxes and transfers, in 2015 US dollars. HH (household) income is average labor income for adults in the household, excluding adult children. Employment is an indicator for paid employment or self-employment being the individual's main source of income. Social distance is the fraction of neighborhood residents that are not of the applicant's ethnic group. Fiscal externality is individual taxes paid net of non-housing transfers. Regressions include controls for baseline outcomes. * = $p<.1$; ** = $p<.05$; *** = $p<.01$. 'Base' is the sample mean for those who do not receive an offer, at the time of application.

A.5 Descriptive details on latent types

Table 9: Description of latent heterogeneity.

	Type 1	**Type 2**	**Type 3**
Type shares:	0.245	0.566	0.189
Probabilities by type:			
Participation probability	0.175	0.071	0.204
Application probability (low income dest NH)	0.373	0.349	0.338
Application probability (middle income dest NH)	0.525	0.530	0.501
Application probability (high income dest NH)	0.102	0.121	0.161
Acceptance probability (low income dest NH)	0.514	0.523	0.515
Acceptance probability (middle income dest NH)	0.527	0.543	0.566
Acceptance probability (high income dest NH)	0.568	0.606	0.623

Notes: Top sub-table reports the shares of each type in the population of applicants. The bottom sub-table reports various probabilities conditional on type. Note that application probabilities are conditional on participating and acceptance probabilities are conditional on being made an offer.

B Proofs and derivations

B.1 Derivation of expression for Wald estimator

We would like to show that $\hat{\Delta}_{j,N_j}$ equals the within-lottery Wald estimator.

Observe first that the Wald estimator is

$$\frac{\frac{1}{O_j}\sum_{i=1}^{N_j} Y_{ij}Z_{ij} - \frac{1}{N_j-O_j}\sum_{i=1}^{N_j} Y_{ij}(1-Z_{ij})}{\frac{1}{O_j}\sum_{i=1}^{N_j} D_{ij}Z_{ij} - \frac{1}{N_j-O_j}\sum_{i=1}^{N_j} D_{ij}(1-Z_{ij})}.$$

The denominator in this expression equals $\frac{1}{O_j}$. This follows from the fact that D_{ij} is zero for individuals with $Z_{ij} = 0$, and from the fact that, under Assumptions 3.3 and 3.4, exactly one individual receives treatment, and therefore $\sum_{i=1}^{N_j} D_{ij}Z_{ij} = 1$.

Hence, as desired, the Wald estimator equals:

$$\sum_{i=1}^{N_j} Y_{ij}Z_{ij} - \frac{O_j}{N_j-O_j}\sum_{i=1}^{N_j} Y_{ij}\left(1-Z_{ij}\right).$$

B.2 Proof of Proposition 3.1

We would like to show:

$$\lim_{N_j\to\infty} \mathbb{E}\left[\hat{\Delta}_{j,N_j}\right] = \Delta_j.$$

Here, the expectation is taken over the joint distribution of the sampling process and $\mathbf{R_j}$. It is straightforward to derive

$$\begin{aligned}\hat{\Delta}_{j,N_j} = &\sum_{i=1}^{N_j}\left(Y_{ij}(1)-Y_{ij}(0)\right)Z_{ij}D_{ij}(1) + \sum_{i=1}^{N_j} Y_{ij}(0)Z_{ij}\left(1-D_{ij}(1)\right) + \sum_{i=1}^{N_j} Y_{ij}(0)Z_{ij}D_{ij}(1)\\ &- \frac{O_j}{N_j-O_j}\sum_{i=1}^{N_j} Y_{ij}(0)\left(1-D_{ij}(1)\right)\left(1-Z_{ij}\right) - \frac{O_j}{N_j-O_j}\sum_{i=1}^{N_j} Y_{ij}(0)D_{ij}(1)\left(1-Z_{ij}\right). \qquad (4)\end{aligned}$$

Under Assumption 3.1,

$$\mathbb{E}\left[Z_{ij}|O_j, C_j, N_j, D_{ij}(1)\right] = \begin{cases} \frac{1}{C_j} & \text{if } D_{ij}(1) = 1 \\ \frac{O_j - 1}{N_j - C_j} & \text{if } D_{ij}(1) = 0 \end{cases}.$$

Using this fact and the law of iterated expectations, equation (4) implies

$$\mathbb{E}\left[\hat{\Delta}_{j,N_j}\right] = \mathbb{E}\left[\frac{1}{C_j}\sum_{i=1}^{N_j}\left(Y_{ij}(1) - Y_{ij}(0)\right)D_{ij}(1)\right] + B_{j,N_j}, \tag{5}$$

where

$$B_{j,N_j} := \mathbb{E}\left[\left(\frac{O_j(C_j-1)}{N_j-O_j} - 1\right)\widetilde{Y}_N(0) + \left(1 - \frac{O_j(C_j-1)}{N_j-O_j}\right)\widetilde{Y}_C(0)\right],$$
$$\widetilde{Y}_N(0) = \frac{1}{N_j-C_j}\sum_{i=1}^{N_j} Y_{ij}(0)(1 - D_{ij}(1)) \quad \text{and} \quad \widetilde{Y}_C(0) = \frac{1}{C_j}\sum_{i=1}^{N_j} Y_{ij}(0)D_{ij}(1).$$

Observe that, under regularity conditions on second moments of potential outcomes,

$$\frac{1}{C_j}\sum_{i=1}^{N_j}\left(Y_{ij}(1) - Y_{ij}(0)\right)D_{ij}(1) \xrightarrow{p} \Delta_j.$$

It remains to show that $B_{j,N_j} \to 0$. Again assuming second moments for potential outcomes are bounded, $\widetilde{Y}_C(0)$ and $\widetilde{Y}_N(0)$ are $O_p(1)$. Furthermore, $\widetilde{Y}_C(0)$ and $\widetilde{Y}_N(0)$ are mean-independent of O_j, C_j, and N_j. It will therefore suffice to show that $\mathbb{E}\left[1 - \frac{O_j(C_j-1)}{N_j-O_j}\right] \to 0$. To see this, observe that

$$1 - \frac{O_j(C_j-1)}{N_j-O_j} = \frac{N_j - O_jC_j}{N_j-O_j} = \frac{1-O_j(C_j/N_j)}{1-O_j/N_j}.$$

Using Assumption 3.4 and the fact that $O_j \in \{1, 2, .., N_j - C_j\}$,

$$\frac{1-O_j\rho_j}{C_j/N_j} \leq \frac{1-O_j\rho_j}{1-O_j/N_j} \leq 1 - O_j\rho_j.$$

It can be shown that $O_j \xrightarrow{d} \text{geom}(\rho_j)$.[29] Hence $\mathbb{E}\left[O_j\right] \to \frac{1}{\rho_j}$, and the expectations of both the upper and

[29] O_j has the same distribution as a random variable $X + 1$ where X follows a negative hypergeometric distribution with parameters N_j, $N_j - C_j$, and 1. $X + 1$ has expectation $\frac{N_j+1}{C_j+1}$, and approaches a geometric random variable as $N_j \to \infty$. Intuitively, for large pools of applicants the offer process that relies on permutations comes to approximate random sampling with replacement.

the lower bound approach zero as $N_j \to \infty$, forcing $\mathbb{E}\left[1 - \frac{O_j(C_j-1)}{N_j-O_j}\right] \to 0$, as required.

These bounds also suggest which term governs the speed with which the bias term approaches zero: this will be determined mainly by how fast the expectation of O_j approaches $\frac{1}{\rho_j}$, which in turn will depend on the size of ρ_j (intuitively, if ρ_j is closer to one, the size of the applicant pool will seize to matter sooner).

B.3 Derivation for Corollary 3.2

We would like to show

$$\lim_{N\to\infty} \mathbb{E}\left[\hat{\beta}_N\right] = \frac{1}{J}\sum_{j=1}^{J} \Delta_j.$$

The 2SLS estimator is:[30]

$$\hat{\beta}_N = \frac{\sum_{j=1}^{J} O_j \left(1 - \frac{O_j}{N_j}\right)\left[\bar{Y}_{j,Z=1} - \bar{Y}_{j,Z=0}\right]}{\sum_{j=1}^{J} O_j \left(1 - \frac{O_j}{N_j}\right)\left[\bar{D}_{j,Z=1} - \bar{D}_{j,Z=0}\right]}$$

Note that $\bar{D}_{j,Z=0} = 0$ due to Assumption 3.3. Also, since offers are given until one person accepts, $\bar{D}_{j,Z=1} = \frac{1}{O_j}$. Therefore:

$$\begin{aligned}\hat{\beta}_N &= \frac{\sum_{j=1}^{J} O_j \left(1 - \frac{O_j}{N_j}\right)\left[\bar{Y}_{j,Z=1} - \bar{Y}_{j,Z=0}\right]}{\sum_{j=1}^{J} \left(1 - \frac{O_j}{N_j}\right)} \\ &= \frac{1}{J}\sum_{j=1}^{J} \hat{\Delta}_{j,N_j}\left[1 - \frac{1}{J}\sum_{j=1}^{J}\frac{O_j}{N_j}\right]^{-1} - \frac{1}{J}\sum_{j=1}^{J}\frac{O_j^2}{N_j}\left[\bar{Y}_{j,Z=1} - \bar{Y}_{j,Z=0}\right]\left[1 - \frac{1}{J}\sum_{j=1}^{J}\frac{O_j}{N_j}\right]^{-1}.\end{aligned}$$

It would be sufficient to show that $\frac{O_j^2}{N_j} = o_p(1)$ for each j. Recall that $O_j \xrightarrow{d} \text{geom}(\rho_j)$. It can be shown that the second moment of the $\text{geom}(\rho_j)$ distribution is $\frac{2-\rho_j}{\rho_j^2}$, and by assumption 3.4, this is a constant $\in (0,1)$. Thus, as $N_j \to \infty$, $\frac{O_j^2}{N_j} = o_p(1)$.

[30] See Frölich (2007) or de Chaisemartin and Behaghel (2017).

C Data collection and sample construction

C.1 Data sources

The analysis in this paper combines several confidential and public data sets originating from two institutional sources.

1. **Statistics Netherlands.** *Centraal Bureau voor de Statistiek* (CBS) is the independent national statistical agency of the Netherlands. It is responsible for the collection, processing and publishing of official statistics that are used by national and local governments, researchers, the private sector, and the general public. It makes aggregated data available for public use on its website: http://statline.cbs.nl/Statweb/ (available in English), cited in this paper as CBS Statline. Under strict agreements, Statistics Netherlands allows academic researchers to work with the confidential micro data sets that underlie its aggregate statistics.

2. **Amsterdam housing associations.** *Platform Woningcorporaties Noordvleugel Randstad* is an umbrella organization representing all public housing associations in the Amsterdam metropolitan area. The data is collected by the organization that implements the allocation of housing units on their online platform (WoningNet).

C.2 Data linkage and sample construction

This paper considers housing lotteries conducted in 2013 and 2014. The data set contains information on all allocated units and the universe of applicants to any of the lotteries. I also use information on housing lotteries conducted in 2015 and 2016 to study whether applicants who do not receive housing assistance in a given lottery still obtain it in the future, but these lotteries are not included in the main analysis. Statistics Netherlands implemented the de-identification and linking of the different data sets used. Lottery applicants were linked based on date of birth, gender, and address. Housing units were linked based on address. The rate of successful linkage to administrative data for individuals applying to lotteries in 2013-2014 is 96.4%; for housing units allocated through lotteries in that period it is 99.4%. Individuals or housing units are dropped from the analysis sample if the linkage failed.

Table 10: Variables included in each data set.

	Amsterdam	Statistics Netherlands	
		Public	**Confidential**
Household characteristics			
Residential location	○		✓
Household size	○		✓
Income	○		✓
Wealth			✓
Parental wealth			✓
Rent subsidy received			✓
Public assistance received			✓
Demographics			✓
Unit characteristics			
Surface area	✓		✓
Number of rooms	✓		✓
Year of construction	✓		✓
Location (block)	✓		✓
Type of heating system	✓		✓
Tax-assessed value	✓		✓
Rent price	✓		○
Neighborhood characteristics			
Population density		✓	✓
Avg. tax-assessed value		✓	✓
Avg. income		✓	✓
Avg. education level			✓
Avg. family size		✓	✓
Choice data			
Applications to public housing	✓		
Date of enrollment	✓		
Date of exit	✓		
Date of move into public housing	✓		
Local housing and labor markets			
Local private rental rate			✓
Local house prices		✓	✓
Local wages		✓	✓

Notes: This table summarizes the contents of the data sets used in the analysis. Checkmarks (✓) indicate the variable is available. ○ means self-reported. □ means data is limited to a subset of the sample under consideration. (*) To be augmented with data for 2010-2012. (†) start year depends on the administrative source from which the data is drawn, most variables (notably those based on tax returns) in the administrative micro data set are available from at least 1998. (‡) Computed as number of households registered to the wait list during the period 2013-2014, divided by total population (as the denominator does measures individuals rather than households, this figure is a lower bound).