

Making Lotteries Legible: Designing Natural Experiments

Daniel Goldstein¹ and Elyzabeth Gaumer¹

¹Center for Research on Housing Opportunity, Mobility, and Equity (HOME),
Department of Housing Preservation and Development, City of New York,
100 Gold Street, New York, NY 10038*

Abstract

Experimental studies are often considered the gold standard for evaluating the efficacy of public programs, but exploiting an existing lottery can be challenging, requiring researchers to understand all components of the process and its implementation. Moreover, researchers must assess whether the lottery provides an opportunity for clearly identifying treatment effects, for example, by considering the share of applicants who accept the offer of services and whether they continue in the program. In this paper we present details on the lottery used to allocate affordable housing in New York City and the lease-up process that follows. We discuss the implications for conducting experimental and quasi-experimental studies that leverage lotteries to identify comparison groups.

Key Words: Randomized control trial, natural experiment, quasi-experimental, lottery

1. Introduction

Randomized experiments are often celebrated as the gold standard for evaluating the effectiveness of an intervention (Hariton and Locascio 2018). Careful randomization of a given population to treatment and control conditions eliminates confounding and reduces selection bias in assessing the causal impact of an intervention (Senn 2013). Lotteries that allocate limited resources represent a rare opportunity to evaluate the effects of receiving those resources using randomized experiments, which can be uniquely effective in estimating the effects of programs in the real world. These sorts of lotteries are used extensively across a variety of areas in the United States and throughout the world, including education (Unterman 2018), healthcare (Finkelstein et al. 2012), job training (Card et al. 2011) and youth employment programs (Gelber et al. 2014), and housing (Gaumer 2021).

In this paper, we describe the housing lottery used by the City of New York to allocate income-restricted affordable housing, where demand far exceeds supply and where eligible households represent a small fraction of the initial applicant pool (lottery turn-out), which is typical for other cities across the U.S. that allocate scarce affordable housing via lottery.

* The views expressed here are those of the authors and do not necessarily reflect those of the Department of Housing Preservation and Development or the City of New York.

2. The Housing Lottery Eligibility and Screening

Since the 1980s, New York City has allocated newly constructed affordable housing through a lottery.² The lottery process remained the same from at least the early 2000s through July of 2020, with more than 100,000 units awarded.

Interested households applied to a particular development using a standardized application that was required to be submitted by the published deadline. Each application was assigned a random number shortly after the deadline, which determined the queue for being offered housing later on in the process. Applicant households underwent a two-stage eligibility screening process. First, households were assessed based on the self-reported information provided on the housing application. Second, households were invited for a face-to-face screening interview where they submitted required documentation, such as tax returns or pay stubs, that proved eligibility.³ Rejected applicants could appeal the decision; if accepted, the household re-entered the queue without penalty (i.e., they retained their original placement in the queue). Each developer or its agent conducted the screening and allocation of housing, but all followed a prescribed process that was overseen by the housing agency.⁴ As part of this oversight process, no applicant was offered housing prior to review and approval by the agency.

Generally, only a small fraction of households that passed the first screening were invited for the face-to-face interview, as this was a burdensome process for both applicant and developer. The exact number of households invited for the second stage screening interview varied according to developer resources, timing, and share of applicants that were expected to pass the second screening; however, the allocation of housing was limited to eligible households (those that had successfully passed both screenings) and the offer of housing was sequenced according to the order determined by the lottery number assigned at the beginning of the process.⁵

Each lottery prioritized a share of available units for certain types of applicants. Seven percent of units were set aside for those with a disability—five percent for mobility and two percent for visual and/or hearing. Half of the units were prioritized for eligible applicants who resided in the same community⁶ where the development was located and five percent were prioritized for municipal employees. Priority for all units was given to

² Through 2011, the housing lottery was based on paper applications; from 2012 on, the lottery was done electronically, with applications submitted via a web portal. The underlying process remained essentially unchanged. For additional details on the lottery process, please visit: <https://www1.nyc.gov/assets/hpd/downloads/pdfs/services/marketing-handbook.pdf>.

³ Eligibility requirements and corresponding documentation varied from one site to the next, but were uniformly applied to all lottery applicants at a given development.

⁴ Individual developments/lotteries were overseen by either the NYC Department of Housing Preservation or the Housing Development Corporation, both of which follow the same process.

⁵ Assignment of units to households thus did not vary based on the efficiency of the interview process. The fact that some developers assessed households further down the queue than others may have resulted in a differential ratio of households-offered-units to households-eligible-for-but-not-offered-units across individual lotteries.

⁶ The community preference is based on residence in the same Community District, an administrative boundary that is closely approximated by Census Bureau Public Use Microdata Areas (PUMAs). There are 59 Community Districts in New York City and 55 PUMAs.

current New York City residents.⁷ All remaining eligible applicants who were not offered housing through a set-aside or a community or municipal employee “preference” were considered for the remaining units.

Eligible applicant households were sorted along two dimensions: according to unit-types for which they qualified based on household size and income criteria⁸ and according to the set-aside and preference(s) for which they qualified. Eligible households were offered housing based on the random number assigned, from the lowest to highest number, first among those that were eligible for set-asides and then among those eligible for one or more preference until such priorities were filled. The remaining units were then allocated within unit-types, again based on the random number assigned, until all units had been filled. Eligible households that were not offered housing were “next in line” for housing, at initial lease-up as well as later on if any household moved out of the select affordable housing unit.

3. Lottery and Lease-Up: An Illustrative Example


Figures 1A and 1B show a schematic of the lottery and lease-up process using a simple example of a 10-unit building with only one priority type of up to 50 percent of available units (applicants who are eligible for this priority are indicated by a black circle). In this example, the lottery produced 20 eligible applicants (8 of whom were eligible for the priority and 12 of whom were not). For simplicity, it is assumed that all 20 applicants passed the first and second screening and were therefore eligible for housing and either offered a unit or classified as eligible but not offered a unit. As discussed in the following section, these groups could be deemed “treatment” and “control,” or “TX” and “CT” in figure 1B. In practice, some applicants would be deemed ineligible and their lottery numbers would be removed from the queue. For example, if Number 3 were deemed ineligible, the sorted number list in Stage 2 would be sequenced as 1, 2, 4, and so on.

The lease-up or allocation of units occurs after the lottery (Figure 1B). In Stage 3, applicants are sorted in numeric order along two dimensions: by the unit-type for which they qualify (A, B, or C) and priority status (black or white circles). In Stage 4, priority applicants are offered housing first, based on the queue, up to the maximum number of units specified for a given set-aside or preference category in that building. In this example, there is a 50 percent priority, or 5 units, so the first 5 priority applicants are offered units in sequence (here, applicants 3, 8, 11, 12, and 16 in that order). If there were not enough applicants to fill the priority, all remaining units would be filled in Stage 6. Here, if there were only four applicants, then only 4 units would be filled with a priority; or, if all eight priority applicants were only eligible for Type A units, then only 3 units would be filled with the priority. If, on the other hand, priority applicant 8 turned down the offer of a Type C unit, then priority applicant 17 would be offered a Type A unit. Non-priority applicant 4 would not be offered a Type A unit, but instead would be next in line.

⁷ Although non-New York City residents may be eligible for housing, the large number of city residents that apply to the lottery means that it is rare for anyone outside of New York City to be awarded housing through this process.

⁸ Individual units at a given site have specific eligibility criteria, including minimum and maximum household sizes and incomes ranges. This varies by site, but is published as part of the public notice soliciting applications to the lottery. Applicants with secondary rental assistance are a special case that is not addressed in this paper.

1. A lottery is held for a building with 10 affordable units. The building has 3 unit types: A, B, and C.



2. The lottery results in 20 applicants (here, all are eligible): 8 priority (in black); 12 non-priority (in white). Each applicant is assigned a random number.

Applicants are sorted in random number order to form the queue (i.e., 1-20).




Figure 1A: Toy example of an affordable housing lottery (Stages 1 and 2)

In Stage 5, the queue is re-sorted so that all remaining applicants are ordered according to the random number assigned by the lottery. Any applicants who were eligible for a priority group but were not offered a unit are considered for the remaining units along with all applicants who did not qualify for a priority. In this example, applicant 17 was eligible for a priority and was originally 2nd in line for a Type A unit in Stage 3; however, the priority was filled before 17's number was reached; therefore, they continue in the process but are now lower down in the queue for Type A units. In Stage 6, all remaining units are offered to applicants that qualified for that unit-type based on the queue defined in Stage 5. Here, applicants 2 and 4 are offered Type A units, applicants 7 and 10 are offered Type B units, and applicant 1 is offered the last remaining Type C unit. If applicant 1 declines the offer of housing, then applicant 6 who is next in line would be offered that unit.

4. Lottery Use for Research and Evaluation

The majority of New York City affordable housing lotteries screened and identified many applicant households that were eligible for affordable housing but not offered units because demand exceeded supply. For programmatic purposes, these eligible households were often considered “back-ups” to the households offered units and generally were placed on a waiting list. In the research context, however, these eligible households represent a control group that is comparable to the households that were offered a unit. Ultimately, the lottery randomly assigns, or does not assign, housing to a group of households that have all been classified as eligible for housing. This creates the conditions for a naturally occurring randomized experiment to evaluate the impact of newly constructed affordable rental housing on this group of eligible households.

3. Applicants are sorted by unit-type and random number; priority comes first, followed by non-priority.

A	B	C
3	12	8
17	19	11
2	7	16
4	10	20
5	18	1
9		6
15		13
		14

4. Priority applicants are offered units in order of their random number up to maximum preference (here, 50% or 5 units).

A	B	C
		8
3	12	11
		16

5. Remaining applicants are re-sorted by random number and unit-type (regardless of priority).

A	B	C
2	7	1
4	10	6
5	18	13
9	19	14
15		20
17		

6. Remaining units are offered in order of their random number.

A	B	C
2	7	
4	10	1

7. Final housing allocation and identification of "treatment" and "control."

A	B	C
		8
3	12	11
2	7	16
4	10	1
5	18	6
9	19	13
15		14
17		20

↑ TX
 ↓ CT

Figure 1B: Toy example of lease-up / allocation (Stages 3-7)

Every applicant who was offered an affordable housing unit in a given lottery, in this context the intervention, could be deemed a “treatment” household, and every applicant who passed the first and second screening for an affordable unit in the same lottery but was not offered one could be classified as a “control” household. Further, each lottery independently identifies a group of households that are eligible for housing, thus the assignment of a household to treatment or control based on the lottery would be independent of the size of the lottery—turn out, the number of applicants that passed the screening process, or the length of time it took to complete the lottery and lease-up.

In practice, an individual household could be eligible for more than one unit-type if the household size and income requirements overlap for two or more units. This could result in a differential probability of being awarded a unit among eligible households that could be addressed through various analytic strategies or exploited for research purposes. Regardless, depending on the estimand(s) of interest and the study design, researchers should understand both the design and implantation of the lottery that they are using.

There are, however, some individual lotteries where the demand among *eligible* households does not exceed the supply of available units. These lotteries that do not identify a sufficiently large control group may also be used by researchers employing other methods beyond experimental comparisons. Because income qualification uses a sharp cut-off at an exact dollar amount, a regression discontinuity design (RDD) may be employed to compare households offered a unit with households that did not qualify only because their incomes were too high (or too low) (Imbens and Lemieux 2008; Jacob et al. 2012). Other quasi-experimental designs may also be used, such as propensity score matching between households offered housing and those not deemed eligible, or even those that did not apply (Caliendo and Kopeinig 2008). And, to the extent that the aim is pool across multiple lotteries, other forms of matching may be more or less appropriate.

The lottery process for affordable housing in New York City generally results in nearly every household that is offered an affordable unit accepting the offer and the overwhelming majority of households remaining in their units for a fairly long period of time, implying high “take-up” and “compliance” from an experimental perspective. This, along with the high ratio of households seeking affordable housing to units available to them, generally results in few control households ever moving into a unit in the specific affordable housing development whose lottery they applied to (“interference”) or, in fact, moving into any affordable housing development at all, “interference” and “contamination.” The standard analytic approach for a randomized trial is intent-to-treat, as it is the most common way to retain the original randomization and ensure that the groups are comparable (Gupta 2011). For the affordable housing lottery, this would entail comparing all households offered a unit—regardless of whether they accepted the offer—with all households eligible for but not offered a unit. But to the extent that they are interested in the effects of affordable housing on people living in a unit of affordable housing, researchers may use two-stage least-squares to generate instrumental variables (IV) estimates of the “dose-response” to affordable housing (Ludwig et al. 2012).

5. Conclusion

The internal validity of randomized experiments depends on randomization to achieve two equivalent groups. While this is relatively straight-forward in a laboratory-controlled setting or when the research team controls randomization in a field experiment, natural experiments rely on a “found” method of identifying treatment and control groups. The

existing housing lottery used by the City of New York is one clear example.

Although there is continued debate over the value of randomized experiments, they remain an essential tool for producing unbiased estimates of the average effect of an intervention—here, the offer of affordable housing. It is important to remember that, just like in any investigator-randomized study, any findings from a natural experiment remain limited to the population being studied, the affordable housing sites enrolled in the study, and the time period in which data were collected (Cartwright 2007). Like other randomized experiments, its purpose is to describe the effect of the intervention rather than the mechanisms that explain the effect—that is, the “what” and not “why.” In this sense, any study resulting from the housing lottery would be a valuable but incomplete contribution to the overall understanding of affordable housing and its impact. As a community of housing researchers, policymakers, and practitioners, we need replications of the study findings across different populations and periods that may leverage similar housing lotteries in other cities or regions. We also need future research—including non-experimental quantitative and qualitative studies—on the underlying mechanisms of affordable housing.

Acknowledgements

We would like to acknowledge the marketing staff at HPD and at the Housing Development Corporation as well as the many affordable housing developers and lease-up agents who graciously supported our observation of the lottery and housing allocation process across multiple sites and over many years. Without their generosity we would have not gained the invaluable information about the on-the-ground implementation of these processes, which have been a vital part of our own process for leveraging the NYC lotteries in our studies.

References

- Caliendo, M., and S. Kopeinig. 2008. “Some Practical Guidance for the Implementation of Propensity Score Matching.” *Journal of Economic Surveys* 22(1):31–72. doi: 10.1111/j.1467-6419.2007.00527.x.
- Card, D., P. Ibarrraran, F. Regalia, D. Rosas-Shady, and Y. Soares. 2011. “The Labor Market Impacts of Youth Training in the Dominican Republic.” *Journal of Labor Economics* 29(2):267–300. doi: 10.1086/658090.
- Cartwright, N. 2007. “Are RCTs the Gold Standard?” *BioSocieties* 2(1):11–20. doi: 10.1017/S1745855207005029.
- Finkelstein, A., S. Taubman, B. Wright, M. Bernstein, J. Gruber, J. P. Newhouse, H. Allen, and K. Baicker. 2012. “The Oregon Health Insurance Experiment: Evidence from the First Year.” 1057(3):1057–1106.
- Gaumer, E. 2021. “Residential Context and Well-Being: An Experimental Study of Affordable Housing in New York City.” University of Chicago.
- Gelber, A., A. Isen, and J. B. Kessler. 2014. *The Effects of Youth Employment: Evidence from New York City Summer Youth Employment Program Lotteries*. Working Paper 20810. Cambridge (MA).
- Gupta, S. K. 2011. “Intention-to-Treat Concept: A Review.” *Perspectives in Clinical Research* 2(3):109–12.
- Hariton, E., and J. J. Locascio. 2018. “Randomised Controlled Trials—the Gold Standard for Effectiveness Research.” *BJOG* 125(13):1716. doi:

<https://doi.org/10.1111/1471-0528.15199>.

- Imbens, G., and T. Lemieux. 2008. "Regression Discontinuity Designs: A Guide to Practice." *Journal of Econometrics* 142(2):615–35.
- Jacob, R., P. Zhu, M. Somers, and H. Bloom. 2012. *A Practical Guide to Regression Discontinuity*. New York.
- Ludwig, J., G. J. Duncan, L. A. Gennetian, L. F. Katz, R. C. Kessler, J. R. Kling, and L. Sanbonmatsu. 2012. "Neighborhood Effects on the Long-Term Well-Being of Low-Income Adults." *Science* 337:1505–10. doi: 10.1126/science.1224648.
- Senn, S. 2013. "Seven Myths of Randomisation in Clinical Trials." *Statistics in Medicine* 32(9):1439–50.
- Unterman, R. 2018. *What Are Naturally Occurring School Lotteries and How Do We Identify Them?* New York.