# Can Moves to Opportunity be Constructed? Evidence from the Low-Income Housing Tax Credit<sup>\*</sup>

Henry Downes<sup>†</sup> John Soriano<sup>‡</sup> George Zuo<sup>§</sup>

May 29, 2025

#### Abstract

Many low-income households face frictions in moving to opportunity-rich neighborhoods. We estimate tenant-level mobility impacts of the Low-Income Housing Tax Credit (LIHTC), the largest and fastest-growing federal program financing the construction of affordable rental housing for low-income households. We construct a novel panel of address histories for applicants to LIHTC-funded housing by linking proprietary application data from a large private developer to administrative data. Within one year of application, applicants who successfully place into a LIHTC-funded unit reside in lower-quality neighborhoods compared to similar but unsuccessful applicants, and these effects tend to increase over time. In heterogeneity analyses, we find that aggregate effects are driven by moves to especially high-poverty areas where federal incentives promote LIHTC development.

JEL Classification: H20, H31, I38, R23

<sup>\*</sup>This version: May 29, 2025. This research was supported by funding from the Smith Richardson Foundation's Domestic Policy Program. We also acknowledge research support from the University of Notre Dame's Wilson Sheehan Lab for Economic Opportunites (LEO). We thank Rob Collinson, Bill Evans, Judy Hellerstein, Dan Hungerman, Michael Hollar, Melissa Kearney, Paul Odland, David Phillips, Jason Ward, and seminar participants at the University of Maryland, the University of Notre Dame, the 2023 Association for Public Policy Analysis and Management meetings, and the 2024 Southern Economic Association meetings for helpful comments and discussions. We express special gratitude for the partnership of a private LIHTC developer who provided data for this study.

<sup>&</sup>lt;sup>†</sup>University of Notre Dame, 3085 Jenkins Nanovic Hall, Notre Dame, IN 46556. Email: hdownes@nd.edu

<sup>&</sup>lt;sup>‡</sup>University of Dallas, 223 Satish & Yasmin Gupta College of Business, 1845 E Northgate Drive, Irving, TX 75062. Email: jsoriano@udallas.edu

<sup>&</sup>lt;sup>§</sup>RAND Corporation, 1200 S Hayes St, Arlington, VA 22202. Email: gzuo@rand.org

# 1 Introduction

In the three decades since the Moving to Opportunity (MTO) housing mobility demonstration, a growing body of research has established that neighborhood quality is an important determinant of long-run outcomes, especially for low-income families (Chyn and Katz, 2021). Informed by these findings, policymakers have sought new solutions to help low-income families surmount the cost barriers that price them out of high-opportunity neighborhoods. Most existing efforts follow in the footsteps of the MTO experiment, pairing voucher-based rental assistance with individualized supports and services. Such interventions have met with some success (Bergman et al., 2024) but are costly and difficult to scale, and so remain largely at the fringes of housing policy, driven by grassroots local efforts and small-scale regional demonstrations (Mumphery et al., 2023). Overall, residential segregation by income and race has continued to grow since the MTO experiment (Reardon et al., 2018), and the available evidence suggests that demand-side mobility programs alone are unlikely to reverse these trends.

In this paper, we draw attention to a different mode of rental assistance that has been overlooked in the housing mobility discourse: affordable housing subsidized through the Low-Income Housing Tax Credit (LIHTC). The LIHTC provides roughly \$13 billion in annual tax credits to developers who agree to set aside a portion of their housing units for low-income renters at below-market prices. The LIHTC is both the largest and fastest-growing federal affordable housing program (Figure 1) and has funded more than 20% of all multi-family housing development in the U.S. over the past two decades (Diamond and McQuade, 2019).

Unlike voucher mobility programs that aim to increase demand for better neighborhoods, the LIHTC influences housing mobility through the supply and location of affordable housing. The two approaches are complementary: voucher holders enjoy the flexibility of choosing where to live, but often require extensive supports to successfully move to better neighborhoods (Bergman et al., 2024). By contrast, renters take the supply of affordable housing as given, so the mobility benefits of the LIHTC are entirely determined by where developers choose to build. This implies that federal and state governments – which wield considerable influence over where LIHTC development occurs – can theoretically create pathways to high-opportunity neighborhoods at a far greater scale than through high-touch, individualized supports for voucher holders.

Despite the LIHTC's size and importance, most existing causal studies on the LIHTC focus exclusively on the program's effects on neighborhoods (Diamond and McQuade, 2019; Davis et al., 2019; Baum-Snow and Marion, 2009; Freedman and Owens, 2011; McGuire and Seegert, 2023) and relatively little is known about the program's impacts on individual renters. This is due in large part to data constraints that stem from the LIHTC's decentralized nature: although the U.S. Department of Housing and Urban Development (HUD) ensures that LIHTC-funded properties are compliant with program affordability rules, applicant, waitlist, and tenant data are maintained as proprietary information by the thousands of developers across the country that use the tax credits. In this study, we leverage unique access to applicant and tenant data from one of the nation's largest LIHTC developers to quantify the program's effects on mobility at the *individual level*, focusing on changes in neighborhood quality as measured by poverty rates, social and economic outcomes for children, racial composition, and labor market vitality. We link personal identifiers from applicant records to individual-level migration histories from Infutor Data Solutions, a comprehensive database containing over 375 million individual records connected to nearly 1 billion address and name histories.<sup>1</sup> These data allow us to construct a novel panel that captures outcomes related to each individual's place of residence before and after their application to a LIHTC-subsidized unit.

Crucially, our sample contains information on both successful and unsuccessful applicants, which allows us to provide the first credibly causal estimates of the individual-level effect of moving to LIHTC on neighborhood quality. Specifically, we employ an event study design that compares outcomes of successful applicants and unsuccessful applicants. We leverage the panel nature of our data to eliminate sources of bias related to both observable and unobservable applicant characteristics by estimating treatment effects after absorbing person-level fixed effects. We further control for a variety of baseline application characteristics interacted with event time fixed effects. Our identification strategy rests on the assumption that, conditional on these controls, neighborhood trajectories of residents and non-residents would have evolved similarly in the absence of treatment. Event study visualizations are consistent with this assumption, revealing no discernible pre-trends.

We find that placement into LIHTC-subsidized units affects some, but not all, neighborhood out-

<sup>&</sup>lt;sup>1</sup>This dataset has been employed in numerous peer-reviewed economic studies, including Bernstein et al. (2022); Diamond et al. (2019); Pennington (2021); Asquith et al. (2023); Mast (2023); Phillips (2020), and Collinson et al. (2022).

comes. Within two years of application, successful applicants live in Census tracts characterized by higher poverty rates (1 percentage point (pp) increase, or 5.5% above the base period mean), higher predicted incarceration rates (0.175 pp increase, or 7% above baseline), and lower intergenerational income mobility (0.35 pp decrease, or 4% below baseline) compared to similar but unsuccessful applicants to the same property. These effects emerge within 6-12 months of application and tend to increase over time. The effects on poverty rates and income mobility suggest that moving to LIHTC-funded units tends to make individuals worse off with respect to neighborhood level opportunity. However, we do not detect statistically significant effects associated with moving to LIHTC-funded units for other outcomes, including measures of labor market strength, educational outcomes, and racial composition. These mixed results underscore the fact "opportunity" is a complex and multi-faceted concept that cannot be summarized by one or two measures.

To shed light on the heterogeneity underlying our average treatment effects, we first examine whether LIHTC moves improved or worsened neighborhood quality at the extensive margin. We find that successful applicants are more likely than unsuccessful applicants to move to both higher and lower opportunity neighborhoods after application. These results imply that the LIHTC creates a mix of winners and losers, but that the relative gains that some placed applicants experience are more than offset by moves to worse – and, sometimes, far worse – neighborhoods by other placed applicants. Moreover, the time path of treatment effects suggests that placement in LIHTCsubsidized units tends to "lock-in" residents in low opportunity neighborhoods over time. By contrast, counterfactual unsuccessful applicants appear to gradually make opportunity-enhancing moves in the months following their application.

We close the empirical analysis by testing for heterogeneous effects across multiple dimensions of property and applicant characteristics. Approximately 40% of all LIHTC properties are sited in Qualified Census Tracts (QCTs), especially economically disadvantaged areas where LIHTC developers are entitled to receive bonus tax credit allocations. Critics of the LIHTC have often expressed concerns that the QCT incentive structure drives low-income renters to neighborhoods characterized by low economic opportunity and few amenities. We re-estimate the event study design for subgroups of applicants who apply to properties in QCT versus non-QCT tracts. The results strongly suggest that QCTs play a central role in driving LIHTC residents to lower opportunity tracts. For the primary outcomes of interest, we find that virtually all of the variation in average outcomes between successful and unsuccessful applicants is concentrated in QCT tracts, and we estimate relatively precise null effects in non-QCT tracts. Turning next to applicant characteristics, we find that treatment effects are primarily driven by smaller households (those seeking 1 bedroom units) and applicants aged 45-64. We speculate that there are two primary mechanisms driving this heterogeneity. First, fierce competition for smaller affordable units – which are relatively scarce in major U.S. housing markets – may cause individuals who are seeking such units to trade off neighborhood quality for subsidized rents. Second, since neighborhood quality tends to most strongly impact young children (Chetty et al., 2016), households with older or adult children may be more willing to choose a "bundle" that offers lower neighborhood quality but more affordable rents.

Our work makes two primary contributions. First, we shed new light on the impacts that LIHTC subsidies have on low-income renters. Prior research recognizes the LIHTC as an effective place-based policy that can increase local housing supply (Baum-Snow and Marion, 2009; Eriksen and Rosenthal, 2010), reduce crime (Freedman and Owens, 2011), and revitalize low-income neighborhoods (Diamond and McQuade, 2019) – benefits that housing vouchers do not offer. Other recent work focuses on modeling developers' response to LIHTC incentives and aggregate impacts on local housing markets (Soltas, 2023; Cook et al., 2023). However, the program's impacts on individual tenants remain underexplored, in large part because data on applicants to LIHTC-funded units is not publicly available. Drawing on our novel data, we provide the first credibly causal estimates of the individual-level effect moving to LIHTC on neighborhood quality. Second, we contribute new evidence on the role that programs that increase the supply of affordable units play in shaping access to opportunity. Our findings suggest that as currently designed, the LIHTC may not be as effective on average as previously-studied voucher programs in inducing moves to better neighborhoods. However, we observe that LIHTC subsidies have highly heterogeneous effects and create a mix of winners and losers. Since LIHTC is fundamentally a project-based rental assistance program and applicants take the supply of subsidized units as given, it therefore matters a great deal where developers are incentivized and awarded credits to build. Accordingly, we find that program incentives play an outsized role in mediating the impacts of moving to LIHTC on neighborhood opportunity.

# 2 The Low Income Housing Tax Credit (LIHTC)

### 2.1 How do the subsidies work?

The LIHTC was established as part of the Tax Reform Act of 1986 to expand the supply of affordable housing via developer subsidies. Since its passage, the program has become a bedrock of the U.S. housing landscape with over 3.5 million units placed in service between 1987-2021. Today, the program costs an estimated \$13.5 billion annually (Keightley, 2023). Figure 1 shows that the LIHTC is both the largest and fastest-growing housing affordability program operated by HUD. Program funding has been used in the construction of more than 20% of all multi-family housing development over the past two decades (Diamond and McQuade, 2019).

The LIHTC differs from other HUD programs in that it provides developers with incentives (in the form of tax credits) to build rent-restricted housing that is set aside for low-income tenants.<sup>2</sup> The program subsidizes up to 70 percent of a project's development costs for low-income units.<sup>3</sup> Prospective developers choose from a menu of options for rent-restricted units set aside for lowincome residents, including: 1)  $\geq$ 20% of units occupied by tenants earning below 50% of the Area Median Income (AMI), or 2)  $\geq$ 40% of units occupied by tenants earning below 60% of AMI.<sup>4</sup> Rents for these units cannot exceed 30% of the corresponding AMI threshold. The low-income requirements bind for 15 years, and properties can gradually phase out of the LIHTC program over the course of the subsequent 15 years.

#### 2.2 How does the LIHTC affect mobility to high-opportunity places?

Currently, housing vouchers remain the predominant approach for promoting housing mobility. The landmark Moving to Opportunity (MTO) experiment demonstrated that housing vouchers tied to moving to a low-poverty neighborhood increased the future incomes and social outcomes of young children (Chetty et al., 2016). Three decades later, voucher-based housing mobility programs

 $<sup>^{2}</sup>$ Developers sell tax credits to investors, who in turn receive a dollar-for-dollar credit against their federal tax liability for a period of 10 years if the property continues to comply with program requirements.

<sup>&</sup>lt;sup>3</sup>The 30 percent subsidy, covers either the acquisition cost of existing buildings or the construction of new buildings when paired with other subsidies. The 70 percent subsidy supports only new construction and does not allow additional federal subsidies to be used.

 $<sup>^{4}</sup>$ There is a third criteria added in 2018 that allows developers to set aside 40% of units such that the average AMI threshold among these units does not exceed 60%. The units contributing to this average must range between 20% and 80% AMI.

still remain limited in size and prevalence (Mumphery et al., 2023). The broader housing voucher system also faces significant constraints. Just one quarter of housing voucher applicants are selected by lottery, only to be placed on waitlists averaging two and a half years in length (and as long as eight years in some areas) before ultimately receiving their vouchers (Acosta and Gartland, 2021). Widely-documented discrimination by landlords against voucher holders has also negatively impacted housing mobility among low-income households (Sard et al., 2018).

In theory, the LIHTC could provide an attractive supply-side complement to housing vouchers by aligning financial incentives of private developers with the ever-growing need for affordable housing. The subsidy may therefore provide similar benefits as a housing mobility voucher in helping low-income renters afford to live in opportunity-rich neighborhoods. In addition, owners of LIHTC-funded properties are prohibited from discriminating against voucher holders, which creates important complementarities between the LIHTC and voucher programs (Emmanuel and Aurand, 2024). In practice, however, the LIHTC's role in promoting economic mobility crucially depends on state-specific and federal incentives that determine the spatial allocation of LIHTC development. Each state determines how to allocate its share of tax credits through its own competitive scoring process and discloses its scoring criteria through an annual "Qualified Allocation Plan" (QAP). Project location plays a central role in scoring, and many states have incorporated opportunity-based criteria into their QAPs, including access to education, economic vitality, health care access, and access to transportation (Freddie Mac, 2023). While most state QAPs award points to developers who propose projects in high-amenity neighborhoods, these incentives can be countervailed by preferences for lower-cost developments, which tend to be most feasible in low-amenity neighborhoods. In addition, any project built within qualified low-income census tracts ("QCTs") is eligible for a federally-mandated 30% tax credit boost, further incentivizing development in high-poverty neighborhoods. Given this complicated landscape with potentially misaligned federal and state incentives, it is not clear ex ante whether the LIHTC promotes or hinders moves to opportunity, on average, or whether the program's mobility impacts vary across policy settings. These fundamental empirical questions motivate our main analyses.

# 3 Data

Due to the highly decentralized nature of the LIHTC program, there is no national individual-level dataset that identifies residents of LIHTC-funded housing.<sup>5</sup> Moreover, existing large microdatasets that capture outcomes of interest (e.g., ACS, CPS) do not capture whether an individual is living in a LIHTC-funded unit (Collinson et al., 2016). Some very recent work has made progress in this area by linking known locations of LIHTC-funded properties to administrative records (Derby, 2021; Cook et al., 2023). However, such linkages are unable to identify *unsuccessful* applicants to LIHTC-funded units, who form a natural comparison group at the individual-level for placed tenants. Since government agencies do not collect administrative data on applicants to LIHTC-funded properties, it is only possible to identify applicants by accessing proprietary data maintained by developers.

Our key empirical innovation in this paper is to link proprietary application, waitlist, and tenant data from a large private developer of LIHTC-funded housing to individual-level outcomes. These novel data allow us to surmount the obstacles that have impeded previous work by identifying plausible counterfactual outcomes for LIHTC tenants. In this section, we provide background information on the private developer who supplied the applicant and tenant data, including institutional details relating to the application and placement processes. We also discuss salient features of the applicant and tenant data, provide summary statistics, and detail the various other datasets that we link to the applicant and tenant sample.

#### 3.1 Applicant and Tenant Data

We obtain applicant and tenant data from a large private developer of affordable housing based in the Midwest. With more than 30,000 residents and 300 LIHTC-funded properties spanning over a dozen states, this developer ranks as one of the top 20 largest developers of newly constructed lowincome LIHTC units in the U.S. (HUD, 2023). Crucially, the developer is fully vertically integrated – the development, design, construction, and management of each property is all conducted in-house.

<sup>&</sup>lt;sup>5</sup>In fact, HUD only began collecting project-level demographic and economic information on tenants in LIHTC units in 2008. The annual "tenant reports" produced from these aggregated data are limited by incomplete coverage and the lack of uniform reporting requirements. Most of the data collection is carried out by the respective state housing finance agencies (HFAs), and per HUD documentation: "although income and rent information was collected across HFAs using fairly uniform standards and definitions, the demographic information was not standardized and, for some HFAs, not collected at all."

As a result, applicant and tenant data for all properties is subject to uniform internal standards of reporting and stewardship, a key advantage in our setting.

Figure 2 depicts the geographic distribution of units associated with all LIHTC-funded properties managed by the developer. Extending from its base in the Midwest, the developer also has a significant presence in parts of Appalachia and the Southeast. While the portfolio includes developments in several large cities (e.g., Columbus, Cincinnati, and Cleveland, OH; Atlanta, GA; Charleston, SC; Norfolk, VA), the majority of LIHTC-funded units are located in suburban or exurban counties. This distribution highlights an important feature of LIHTC more broadly: while much of the literature on other affordable housing policies tends to focus on dynamics in urban areas, LIHTC is embedded in the housing landscape across a wide variety of neighborhood types and locations, often including suburban and rural areas. In this context, migration to LIHTC-subsidized units may look quite different than migration to public housing or project-based Section 8 units. In particular, migration to LIHTC units likely occurs over longer distances on average, and thus may imply even greater changes in neighborhood characteristics such as school district quality and access to jobs.

Since 2012, the developer has tracked all interactions with prospective applicants, applicants, and tenants at each property using Yardi, an internal client management database. Figure 3 depicts the end-to-end flow for the application and placement processes. Prospective applicants ("prospects") can initiate the process by contacting the property manager directly or by expressing interest through one of several online listing platforms (e.g., Zillow, Apartments.com). At this stage, prospects supply basic contact information and wait for the property staff to follow up with details on how to apply.<sup>6</sup> The application collects additional information including family size, unit preferences (e.g., # bedrooms, special needs), current address, and self-reported income. Once the application is submitted and the application fee is paid,<sup>7</sup> the application is forwarded for processing. In general, applications are reviewed on a first come, first served basis. Property staff use observed characteristics including unit preferences and income to match applicants to suitable units. Income is especially important, since units within each property have varying income restrictions in order to ensure compliance with HUD affordability rules. If no suitable unit is currently available, the

<sup>&</sup>lt;sup>6</sup>Alternatively, an interested individual can simply submit a full application without first expressing interest.

<sup>&</sup>lt;sup>7</sup>Application fees at these properties are typically low (less than \$25). They are used to cover administrative costs as well as costs associated with background checks used for eligiblity screening.

applicant is assigned to a wait list for units of that type. If a suitable unit is available, the applicant progresses to the eligibility screening, which involves criminal background and credit checks as well as a third-party verification of income. If an applicant is determined to be eligible, they can accept the offered unit and begin planning their move-in. At any point after an application is submitted (but before a unit is accepted), an applicant may cancel their application, which would remove them from any waitlist.

For all applicants, we observe the applicant's full name and date of birth, the address they applied from, information about the property they are applying to, unit preferences, the application date, and the outcome of their application.<sup>8</sup> For placed tenants, we observe all of the applicant characteristics as well as race,<sup>9</sup> verified income, unit characteristics, and move date.

We observe a variety of property-level characteristics from the developer, including each property's name, date placed-in-service,<sup>10</sup> funding sources, tax credit details, and affordability rules. We also observe all unit-level addresses associated with each property in our sample. Since some properties are "scattered sites" with units spread across multiple locations, this allows us to identify the exact location to which placed residents move. Finally, we observe several unit-level characteristics, including unit size (sq ft), number of bedrooms, rent and utility allowance, and whether the unit is designated as subsidized or market rate. For all subsidized units, we observe the income restriction rule associated with the unit.

### 3.2 Supplementary Data

We supplement the applicant and tenant data from the LIHTC developer with data from various sources:

• LIHTC Placed-In-Service Database (HUD, 2023). HUD maintains a publicly available database of all LIHTC-funded projects placed-in-service since the program's founding in 1986. These data include developer identifiers, property names and addresses, placed-in-service and credit allocation dates, whether the property targets specific populations (e.g., families,

<sup>&</sup>lt;sup>8</sup>In some cases we observe additional application details, including self-reported income. Unfortunately, we are not able to use the income data at this stage due to a high proportion of missing values.

<sup>&</sup>lt;sup>9</sup>Properties are not permitted to record the race of non-residents.

<sup>&</sup>lt;sup>10</sup>The LIHTC "placed-in-service" date corresponds to the date that a property is certified as ready for occupancy. For new construction, this date generally corresponds to the inspector's certification. LIHTC-funded properties can begin earning tax credits for investors only after the units are placed-in-service and occupied.

elderly, disabled, homeless), unit counts (total and subsidized), funding sources, and several variables related to the siting of properties. We use these data to compare the properties in our sample to other LIHTC properties. We also fuzzy match the sample of properties to the HUD database to supplement the set of property-level variables available from the developer.

- **Opportunity Atlas**. The Opportunity Atlas is constructed collaboratively by researchers at the Census Bureau and Opportunity Insights, a research and policy group based at Harvard University. The Opportunity Atlas builds upon the literature spawned by Chetty et al. (2014) in providing Census tract-level data for a wide array of observed and predicted socioeconomic mobility indicators. We use these data to construct tract-level measures of "opportunity." We follow previous work in selecting a diverse set of measures intended to proxy for various dimensions of neighborhood opportunity. We include two measures that relate to generational mobility, a key focus of the place-based opportunity literature: (1) the predicted probability of incarceration as an adult conditional on having parents in the bottom quartile of the national income distribution as a child, and (2) the predicted probability of reaching the top quintile of the national income distribution as an adult conditional on having parents in the bottom quartile of the national income distribution as a child. A drawback of these mobility measures is that they represent neighborhood opportunity lagged by several decades, since they are constructed based on the childhood locations of current adults. Therefore, we supplement the generational mobility measures using several contemporaneous measures of opportunity, including: the poverty rate, standardized 3<sup>rd</sup> grade math scores (at the district-level), the single parent share, jobs within five miles, and the average annualized job growth rate. We also include the non-white share of the population at the tract-level to analyze effects on the racial composition of neighborhoods.
- Infutor Consumer Reference Database. We draw on personally-identifiable information available from the Yardi data to link LIHTC applicants and tenants to migration histories compiled by Infutor Data Solutions. Infutor contains over 375 million individual-level records linked with nearly 1 billion address and name histories. Each address record in Infutor is associated with an "effective date", interpretable as a move-in date, allowing us to sequence an individual's move history over time. Because Infutor provides name histories, we can further

link records for individuals associated with multiple aliases and for those whose surname changes over time. Previous work finds that Infutor is highly representative of the U.S. adult population,<sup>11</sup> especially since the 1990s.<sup>12</sup> Infutor data has been previously used by economists to study immigration (Bernstein et al., 2022), rent control (Diamond et al., 2019), housing markets (Pennington, 2021; Asquith et al., 2023; Mast, 2023), and housing stability (Phillips, 2020; Collinson et al., 2022).

### 3.3 Property Summary Statistics and Sample Representativeness

The LIHTC features prominently in most housing markets in the U.S.,<sup>13</sup> and so reflects the heterogeneity in neighborhood and housing conditions that characterize those various markets. Within this complex landscape, we acknowledge that there are limitations to what can be learned by analyzing data from a single developer. Therefore, in this section, we compare features of properties in our sample to a national sample of all LIHTC-funded properties to shed light on the degree to which our results may generalize more broadly.

Table 1 provides summary statistics at the property level. Several features are notable from panel A (Property Characteristics). First, in both our sample and the national sample, over 90% of total units are set-aside for low-income families and have applicable income restrictions. This greatly exceeds the compliance standards imposed by HUD, but is consistent with a competitive credit allocation process that forces developers to commit to building more subsidized units than is minimally required. Second, the properties in our sample are characterized by low vacancy rates: point-in-time measurement from 2023 shows that only 4.2% of all subsidized units are vacant.<sup>14</sup> High occupancy rates reflect the fact that HUD compliance requires all subsidized units to be occupied within one year of the placed-in-service data, absent extenuating circumstances. High

<sup>&</sup>lt;sup>11</sup>Infutor aggregates information from numerous private and public record sources, including USPS change of address forms, county assessor records, magazine subscriptions, and phone books. Due to the nature of the source data, Infutor is generally unable to detect individuals who do not have a "paper trail", such as minors, undocumented immigrants, or homeless populations. Given that our population consists of adults who are able to afford LIHTC-subsidized housing, we believe that Infutor is well-suited to measure migration-related outcomes in our context.

<sup>&</sup>lt;sup>12</sup>Using 2017 ACS 5-year estimates, Asquith et al. (2023) finds a median of 0.88 Infutor observations per Census individual aged 25+ across all Census tracts. Bernstein et al. (2022) find that variation in county populations estimated from Infutor can account for 99% of the variation observed in county populations in the 2000 Census. Phillips (2020) shows that Infutor move dates can detect highly-localized migration shocks, including out-migration from Hurricane Katrina and closures of public housing projects in Chicago.

<sup>&</sup>lt;sup>13</sup>More than 2% of all U.S. households lived in LIHTC-funded units in 2022 (Soltas, 2023).

<sup>&</sup>lt;sup>14</sup>For comparison, the vacancy rate for all LIHTC two-bedroom units from 2010-2020 was approximately 4%.

occupancy rates also imply that units are oversubscribed. We will leverage this oversubscription of units and the first come, first-served nature of application review in our empirical design.

In order to characterize the types of neighborhoods in which LIHTC properties are located, we geocode each property-level address, attribute Census tracts, and merge in 2010 tract-level observables from the Opportunity Atlas data. The results are presented in panel B. Consistent with the geographic distribution of our provider's development portfolio, we find that properties in our sample tend to be located in areas that have lower population density, higher White shares of the population, and higher homeownership rates relative to other LIHTC-funded properties. However, in terms of the three measures that we interpret as proxies for neighborhood-level opportunity – share with a bachelor's degree or higher, poverty rate, and single parent share – neighborhoods with properties in our sample are comparable to the average neighborhood with LIHTC properties. The average neighborhood with LIHTC-funded housing is near the national 75<sup>th</sup> percentile in terms of the poverty rate and single parent share, which suggests that LIHTC development is concentrated in areas of relatively low opportunity.<sup>15</sup>

We also fuzzy match sample properties to the HUD LIHTC Placed-in-Service Database to analyze additional property-level characteristics (panel C). We successfully match 270/342 properties (79%) to the HUD data. From this sub-sample, we are particularly interested in the proportion of properties that are built in Qualified Census Tracts, or "QCTs". QCTs are tracts that have (1) > 50% of households with incomes below 60% of the Area Median Gross Income and/or (2) a poverty rate of 25% of more. Since 1990, federal statute has stipulated that properties in QCTs are to be afforded a 30% tax credit boost, which incentivizes additional development.<sup>16</sup> Many state QAPs also encode special preferences for development in QCTs (Ellen et al., 2015). However, as QCTs are defined to be the most economically disadvantaged neighborhoods in an area, they may play a role in further entrenching LIHTC families in low opportunity neighborhoods. We find that while properties in our sample are less likely to be sited in QCTs than other LIHTC properties, a non-trivial proportion – approximately 15% – are located in a QCT.

Overall, the results in Table 1 indicate that properties in our sample tend to be in neighborhoods

<sup>&</sup>lt;sup>15</sup>This is consistent with Ellen et al. (2018), who find that LIHTC units are located in neighborhoods with higher poverty rates, lower quality schools, and more polluted environments relative to other rental units.

<sup>&</sup>lt;sup>16</sup>The impact of the QCT rule on the local supply of LIHTC units has been exploited by Baum-Snow and Marion (2009) and subsequent literature.

that are whiter and less dense compared to the national LIHTC average, but are fairly representative in terms of key measures of economic opportunity. Moreover, the properties in our sample exhibit many of the same features, including high shares of subsidized units and excess demand for units, that characterize LIHTC-funded properties across the U.S. The results therefore suggest that our findings are likely to generalize to other settings, and especially to those where LIHTC development is concentrated outside central cities.

### 4 Empirical Strategy

Our objective is to estimate the causal effect of moving to a LIHTC-subsidized unit on measures of neighborhood quality and opportunity. To fix ideas, consider a simple setup with 2 time periods (t=1 and t=2) and 2 groups: a treatment group (D=1) and an untreated group (D=0). Treatment is defined as being placed into a subsidized LIHTC unit and occurs in period 2. In a potential outcomes framework, the target parameter is the average treatment effect on the treated, which is given by the estimand:

$$ATT = \mathbb{E}[Y_{i,t=2}(2)|D=1] - \mathbb{E}[Y_{i,t=2}(\infty)|D=1]$$

where  $Y_{i,t}$  is a neighborhood level measure of opportunity that corresponds to individual *i*'s residence at time *t*. The first term,  $\mathbb{E}[Y_{i,t=2}(2)|D=1]$ , is interpretable as the mean outcome for individuals in the treated group after treatment, and is directly estimable from the data. In our context, this refers to the average neighborhood quality that individuals experience after moving to LIHTC-funded units. The second term,  $\mathbb{E}[Y_{i,t=2}(\infty)|D=1]$ , however, is not directly estimable, since we do not observe the counterfactual outcomes for treated individuals in the absence of treatment.

One possible simplifying assumption would be to impute  $\mathbb{E}[Y_{i,t=2}(\infty)|D=1] = \mathbb{E}[Y_{i,t=1}(2)|D=1]$ ; i.e., assume that in the absence of treatment, individuals would have remained in their sending neighborhoods, or at least in measurably similar neighborhoods. A fundamental problem with this approach is that there may be unobserved selection into the decision to apply to a LIHTC-funded unit. By definition, applicants have a demonstrated interest in moving from their current location – whether because of an economic shock, life event, or some other factor – which signals that their

outside option in the absence of LIHTC may not be to remain in place. For example, some applicants may have a preference to move to a different neighborhood with higher economic opportunity, but the rent discounts associated with a LIHTC-subsidized unit may induce them to remain in similar or lower opportunity neighborhoods.<sup>17</sup> In such cases,  $\mathbb{E}[Y_{i,t=2}(\infty)|D = 1] > \mathbb{E}[Y_{i,t=1}(2)|D =$ 1]. Alternatively, some applicants might remain in their sending neighborhood (or a comparable neighborhood) with the LIHTC subsidy, but would have had to migrate to a lower opportunity neighborhood absent the subsidy. For these applicants,  $\mathbb{E}[Y_{i,t=2}(\infty)|D = 1] < \mathbb{E}[Y_{i,t=1}(2)|D = 1]$ . Ultimately, whether one of these potential sources of selection bias would dominate, or if they would cancel out on average, is an empirical question.

Alternatively, one may consider comparing LIHTC applicants to other low-income individuals who did not apply to LIHTC. i.e., impute  $\mathbb{E}[Y_{i,t=2}(\infty)|D = 1] = \mathbb{E}[Y_{i,t=2}(\infty)|D = 0]$ . Such comparisons, however, may also be biased by unobservable characteristics such as personal ambition, the availability of social networks and supports, prior history with housing searches, and skills or ability – all of which could simultaneously influence an individual's economic well-being and their decision to seek LIHTC housing in the first place. Indeed, the available evidence suggests that there is a considerable amount of selection into various types of LIHTC properties (Cook et al., 2023).

We make progress on this question by linking applicant data from the Yardi sample to Infutor to construct a panel that contains address histories for both successful and unsuccessful applicants. We then compare the evolution of outcomes for successful applicants versus those of unsuccessful applicants in an event study design. Formally, we estimate the following regression:

$$Y_{i,t} = \alpha_i + \delta_t + \theta_p + \sum_{t=-24}^{-2} u_t D_i + \sum_{t=0}^{24} u_t D_i + \mathbf{X}'_{i,t} \Pi + v_{i,t}$$
(1)

where *i* indexes individual applicants and *t* indexes months in event time; i.e., relative to the month of application (t=0). In our preferred specifications,  $Y_{i,t}$  is a continuous variable, which yields estimates of effects on the intensive margin. In supplementary results, we test effects on the extensive margin by transforming  $Y_{i,t}$  to be dummy indicator that compares the outcome

 $<sup>^{17}</sup>$ Using data from the Chicago metropolitan area, Cook et al. (2023) find that LIHTC-subsidized rents are 18-41% lower than rents for comparable market rate units.

for individual *i* at event time *t* to that individual's observed outcome at time of application. The variables  $\alpha_i$  and  $\delta_t$  represent applicant and event time fixed effects, respectively, so the  $u_t$  coefficients are identified by variation in outcomes within an individual's address history over time. The vector  $\mathbf{X}'_{i,t}$  contains additional controls that are meant to improve precision and strengthen the case for the (conditional) exogeneity of treatment. Specifically, to account for possibly endogenous factors related to application timing, we include fixed effects for combinations of the calendar month and year of application interacted with event time. We also include a baseline measure of the outcome (corresponding to the application tract) interacted with event time fixed effects to account for the fact that the characteristics of an applicant's sending neighborhood may determine their available outside options and may thus be correlated with selection into treatment. Finally, motivated by the differences between treatment and control individuals that we observe in Table 2, we control for application, all interacted with event time.<sup>18</sup> We cluster standard errors at the individual level, which reflects the unit of randomization into treatment (Abadie et al., 2023).

The key identifying assumption for the event study design is that, in the absence of treatment, outcomes for successful and unsuccessful applicants would have evolved on similar paths over time. Formally, returning to a 2 group  $\times$  2 period setup:

$$\mathbb{E}[Y_{i,t=2}(\infty)|D=1] - \mathbb{E}[Y_{i,t=1}(\infty)|D=1] = \mathbb{E}[Y_{i,t=2}(\infty)|D=0] - \mathbb{E}[Y_{i,t=1}(\infty)|D=0]$$

Assuming that there is no anticipation of treatment by treated individuals, this can be rewritten as:

$$\mathbb{E}[Y_{i,t=2}(\infty)|D=1] - \mathbb{E}[Y_{i,t=1}(2)|D=1] = \mathbb{E}[Y_{i,t=2}(\infty)|D=0] - \mathbb{E}[Y_{i,t=1}(\infty)|D=0]$$

Rearranging terms:

$$\mathbb{E}[Y_{i,t=2}(\infty)|D=1] = \mathbb{E}[Y_{i,t=1}(2)|D=1] + \mathbb{E}[Y_{i,t=2}(\infty)|D=0] - \mathbb{E}[Y_{i,t=1}(\infty)|D=0]$$

<sup>&</sup>lt;sup>18</sup>Unit preference and age variables occasionally contain missing values. When data is missing, we impute the variable to = 0 and include dummy variables in the regression that indicate the presence of missing data.

The term on the left-hand side is the counterfactual component that we cannot directly estimate from the data. However, it is now a function of terms that are all directly estimable. Plugging back into the expression for the ATT:

 $ATT = \mathbb{E}[Y_{i,t=2}(2)|D=1] - (\mathbb{E}[Y_{i,t=1}(2)|D=1] + \mathbb{E}[Y_{i,t=2}(\infty)|D=0] - \mathbb{E}[Y_{i,t=1}(\infty)|D=0])$ 

The panel structure of our data and our ability to observe both successful and unsuccessful applicants, together with the parallel trends assumption, therefore allow us to recover a causal estimate of the ATT that is purged of possible sources of selection bias. While it is not possible to directly test the parallel trends assumption, a common indirect test of this assumption is to compare trends in observed outcomes prior to treatment. We supply evidence on pre-treatment trends and the time path of treatment effects using regression output from our event study models in Section 5.

### 5 Results

### 5.1 Sample Construction

To construct our main analysis sample, we begin with an extract from the provider's Yardi system that contains individual-level records covering all applicants to their LIHTC-funded properties since 2012. We drop observations for which full name, date of birth, or application at address are missing, since we are unable to link individuals without personally-identifiable information to Infutor. We additionally drop observations that are missing an application date, and any individuals who are associated with more than one application, whether at the same property or across multiple properties. The latter sample restriction aids the analysis in three ways. First, this helps ensure that our treatment effect estimates are not contaminated by individuals who are in the control group for one property but in the treatment group for another property. Second, individuals associated with multiple applications may be unobservably different compared to one-time applicants, and such differences are plausibly correlated with treatment. Finally, in some cases, individuals who are already in LIHTC-funded units can re-apply to move to a different unit within the same property. Moves that result from such re-applications do not provide variation of interest in the context of our study.

We link the remaining Yardi observations to the Infutor database by executing a match on name, month and year of birth, and the zipcode associated with the application address.<sup>19</sup> Overall, we successfully match about 45% of our applicant sample to Infutor address records, which is consistent with results from other studies that have linked external individual-level data to Infutor (Collinson et al., 2024). For all matched records, we construct a panel of address histories by sequencing Infutor addresses according to the "effective date", which corresponds to the date at which Infutor first observes an individual as residing at a given address. We interpret effective dates as proxies for start dates at each address, following Phillips (2020).<sup>20</sup> We attribute Census tracts by first geocoding each address in the panel and then using the output latitude and longitude coordinates to link addresses to 2010 definition Census tracts.<sup>21</sup>

We apply several final restrictions at the applicant level to generate our main analysis sample. First, for each applicant, we validate that their application address from Yardi exact matches to at least one record in their linked Infutor address history, and drop any applicants whose application address cannot be validated. This restriction serves two purposes. First, we are concerned that individuals with application addresses that appear in Yardi but not in Infutor may have listed an address for the purposes of the application that does not correspond to their actual place of residence (e.g., a relative or partner's address). Alternatively, a non-validated address may be evidence of measurement error in Infutor; i.e., a "missed move." In such cases, we cannot make confident inferences about the sequencing or timing of moves within an individual's Infutor address history, which is crucial in our setting. For similar reasons, we drop any applicants who ever have multiple addresses in their Infutor history that are associated with the same effective date. Our understanding, based on conversations with Infutor data specialists, is that this issue is related to the manner in which Infutor ingests data from its various sources, and it is not possible to

<sup>&</sup>lt;sup>19</sup>We allow for fuzzy matches on name, but require exact matches on month and year of birth and on zipcode.

<sup>&</sup>lt;sup>20</sup>As Phillips (2020) shows in validation exercises from Hurricane Katrina migration and public housing closures in Chicago, Infutor's effective date appears to be a reasonable proxy for the timing of moves. For any non-terminal address in an individual's history, we can use the effective date as the "begin date" for any address and the effective date of the subsequent address as the "end date." However, for terminal addresses (i.e., an individual's most recent address) we cannot directly define an "end date." In such cases, we assume indefinite continuity at the most recent address, unless Infutor explicitly observes that the individual is deceased.

<sup>&</sup>lt;sup>21</sup>Tract-level data are from the TIGER shape files provided by the U.S. Census Bureau. We employ *geoinpoly*, a user-written Stata command (Picard, 2015), to execute the spatial join of latitude and longitude to shape files containing tract-level information.

disambiguate these cases using other available data.<sup>22</sup> Finally, to ensure that we have a balanced panel, we restrict to applicants for whom we observe address histories for at least two years before and after the application date.

Our final analysis sample contains 4,556 applicants who are associated with 351 unique properties. Of the 4,556 applicants, 41.2% (1,878) are identified as residents; i.e., successful applicants who moved in to a unit in a LIHTC-funded property. The remaining group of non-resident applicants have a variety of application statuses: 29% ultimately canceled their application, 15.7% were denied due to ineligibility, and 13.9% remain on a waitlist. In the analyses that follow, we classify residents in the treatment group and non-residents of any application status in the untreated group. In future work, we hope to increase our sample size by reducing attrition from matching and address validation, which in turn will allow us to be more selective about the composition of the untreated group.

### 5.2 Applicant Summary Statistics by Treatment Group

We present summary statistics for the main sample of applicants in Table 2. Columns (1) and (2) present unconditional means for non-residents and residents, respectively, and column (3) reports the raw difference in means. We also test for differences at time of application using the following regression equation, estimated by OLS:

$$X_i = \beta_0 + \beta_1 Resident_i + \theta_p + \varepsilon_i \tag{2}$$

where the dependent variable  $X_i$  is an observable characteristic at baseline,  $Resident_i$  is a treatment indicator equal to 1 if the applicant is a placed resident, and  $\theta_p$  is a property-level fixed effect. Column (4) reports the  $\beta_1$  coefficient for each baseline characteristic  $X_i$ , which is interpretable as the "regression-adjusted" difference in means from columns (1) and (2).  $\beta_1 > 0$  implies that the resident group mean is greater than the non-resident group mean when using only within-property variation.

In panel A, we characterize applicants based on features of their application data. Residents

 $<sup>^{22}</sup>$ The effective date corresponds to the date that Infutor first observes an individual at an address. However, if Infutor ingests more than one novel address for an individual from the same source at the same time, they attribute the same effective date for each address. In some but not all cases, records with identical effective dates may be resolved over time as additional data is ingested and linked to an individual's history.

tend to be slightly older than residents, although much of this difference dissipates after controlling for property fixed effects, which suggests that age differences are mostly driven by selection at the property level. Even within properties, however, residents have significantly higher quoted rent than non-residents. Since bedroom preferences are comparable across groups, we interpret this difference as indicative of increased competition for units attached to larger subsidies, or "set asides". LIHTC rents are not based on a tenant's income, but on affordability rules stipulated by HUD. Property owners have some flexibility to set rents at the unit level, as long as the overall composition of rents remains compliant with affordability rules at the property level. Although we do not observe the specific unit associated with each application, these results suggest that more subsidized units receive more applications than less subsidized units within the same property. We also characterize applicants based on the the geographic proximity of their application address and the address of the property to which they are applying. Nearly half of successful applicants apply from the same zipcode and more than a quarter apply from the same tract. This is consistent with other results in the literature that show that most moves to LIHTC-funded housing are local moves (Derby, 2021).<sup>23</sup> By comparison, non-residents are much less likely to apply from the local area, which may be a signal that unsuccessful applicants are unsuccessful in part because they are less interested in units outside their local area and so are less likely to follow through at each step of the application process, and/or may be less willing to wait longer on wait lists for units to open up. To the extent that this imbalance reflects unobservable differences between treated and untreated applicants, we view this as an area of concern and hope to address it in future work.

In panel B, we compare characteristics of the application tract using outcome data from Opportunity Atlas. Overall, the results indicate that successful and unsuccessful applicants reside in similar neighborhoods at the time of application, as regression-adjusted differences tend to be statistically insignificant or small in magnitude. Although our main empirical analysis relies on the comparability of outcomes between treated and untreated applicants in changes and not in levels, this is nonetheless an encouraging result because it suggests that the outside option of remaining in place offers similar opportunity levels for both groups. Interestingly, although we do

 $<sup>^{23}</sup>$ Using administrative data from tax records, Derby (2021) finds that people who move to LIHTC-funded units often move from nearby addresses. As we discuss in Section 4, however, we cannot infer from this fact that moves to LIHTC are unlikely to impact neighborhood quality on average. Such a conclusion could only be drawn by making comparisons to a relevant counterfactual, as we do in the main analysis that follows.

not directly observe race in our data, we find some evidence that there may be racial selection at the property level. Specifically, unsuccessful applicants tend to come from significantly less white tracts on average, but this imbalance does not survive property fixed effects. To the extent that tract-level racial composition is a proxy for applicants' race, this suggests that non-white applicants are disproportionately applying to more competitive properties. As of this writing, it is not clear whether such selection effects reflect increased demand for properties in desirable neighborhoods,<sup>24</sup> increased demand for properties with more generous subsidies, or some other unobserved preference. We hope to probe these dynamics further in future work.

In panel C, we report breakdowns of property-level characteristics across treatment groups. Most differences are small or statistically insignificant. A notable exception is that, compared to residents, non-residents are much more likely to have applied to a property that specifically targets families. Such properties make up about one-third of our the properties in our sample. It is plausible that these properties face excess demand and so have longer waitlists, leading more individuals to cancel their application before being placed in a unit compared to less targeted properties.

#### 5.3 Descriptive Results

Before turning to analyses of the primary outcomes of interest, we first analyze move activity over time for applicants in the main analysis sample. This exercise is useful for two reasons. First, it allows us to assess whether the Infutor data provides reliable measurement of move timing at a monthly frequency, a necessary feature for our event study design. In addition, the analysis sheds light on how migration outcomes evolve before and after treatment between the treated and untreated groups. Given that effects on neighborhood level outcomes can only be identified from variation resulting from moves, it is not possible to make sense of the changes in outcomes that we study later without first understanding how treatment impacts move activity.

We plot the results in Figure 4. The outcome is the move rate (per 100 applicants). For each treatment group and event month t, we define the move rate as the number of moves made by individuals in the specified group during that month, divided by the total number of individuals in the group. We plot move rates separately for residents (solid line) and for non-residents (dashed

 $<sup>^{24}</sup>$ It is hard to reconcile this explanation with the evidence from Cook et al. (2023), who find that: "many moderate-need and white (non-Hispanic) households will *only* (italics theirs) apply to affordable housing built in high-opportunity neighborhoods, reducing the odds that other applicants are allocated."

line). The vertical line corresponds to event month t = -1, the month immediately preceding the month of application. Move activity is virtually identical across groups prior to the month of application, which suggests that non-residents provide a valid counterfactual for residents in terms of the propensity to migrate in the absence of treatment. At t = 0, however, migration activity sharply increases among residents, and increases further in t = 1, the month after application. This is an encouraging result; despite known measurement issues in the Infutor data, it nonetheless appears that Infutor's effective dates are reasonable proxies for the timing of moves to new addresses at the month-level. Move activity quickly decays but remains elevated in the treatment group for up to 9 months after the month of application, which likely reflects the fact that some applicants are placed into units only after spending time on a waitlist. Non-residents also experience an uptick in move activity around the time of application, but the profile is different. First, the share of non-residents who move in the months immediately after application is less than half the share of residents who move in the same timeframe. It is reassuring but not surprising that, conditional on failing to secure placement into a unit, non-residents are more likely to remain at their application address. However, increased levels of move activity decay more slowly for non-residents, which suggests that many of these individuals eventually do move from their application address, and in general experience less housing stability than placed applicants in the post period.

Moving to the main outcomes of interest, we plot mean outcomes for each treatment group by event time in Figure 5. This analysis allows us to assess *absolute* changes in outcomes in the pre and post period for each group. In the next section, we formalize these comparisons by estimating event study regressions following Equation 1, which yield credibly causal estimates of *relative* changes in outcomes between groups.

Our first observation is that there are often differences in level terms between the group-level means. However, these differences are often very small relative to the group means, and overall present a complicated picture with respect to neighborhood opportunity. Although some measures suggest that residents tend to live in lower opportunity neighborhoods (higher poverty rates, lower predicted income mobility, lower job growth rates), others suggest that they live in higher opportunity neighborhoods (lower single parent share, lower predicted incarceration, higher math scores). Ultimately, our empirical strategy allows for such level differences across groups as long as the parallel trends assumption is satisfied. Comparing changes in outcomes over time, there is clear evidence that the average poverty rate, which evolves similarly across groups prior to the application month, diverges across groups after the application month. This divergence is largely attributable to moves by non-residents to tracts that have lower poverty rates relative to their pre-application baseline. This is an example of a case where making inferences about the effect of moving to LIHTC housing without a valid counterfactual group would lead to significantly attenuated estimates. There is also evidence, though not as stark, from the predicted income mobility, predicted incarceration, and math score measures that residents do worse in terms of neighborhood opportunity relative to non-residents after placement in a LIHTC unit. The time series for the job concentration (jobs within 5 miles) measure complicates this narrative, however, and suggests that placed applicants end up in tracts that have somewhat stronger job markets.

The results in Figure 5 are purely descriptive and so should be interpreted with caution. Moreover, as we showed in Section 5.2, property level selection appears to play an important role in this context, and accounting for property level fixed effects can even flip the sign of a raw difference in means. For a causal analysis of treatment effects that formally addresses property level selection and other sources of possible bias, we turn next to the event study results.

#### 5.4 Event Study Results

We separately estimate Equation 1 for each outcome, where outcomes are continuous measures of neighborhood opportunity levels, and report the results in Figure 6. In each case, point estimates (solid lines) are interpretable as the regression-adjusted mean difference in treatment group outcomes (residents minus non-residents) at event time t for the specified outcome. Dashed lines represent the upper and lower bounds of the 95% confidence interval.

We find that within about one year of application, individuals who move to LIHTC-funded units reside in tracts with higher poverty rates (panel a), higher predicted incarceration rates (panel c), and lower income mobility (panel f) relative to comparable non-residents. There is no evidence of pre-trends for these outcomes, but we detect statistically significant treatment effects that tend to increase over time. Two years after application, residents live in tracts with 1 pp higher poverty rates (a 5.5% increase over the base period mean), 0.175 pp higher predicted incarceration rates (a 7% increase over the base period mean) and 0.35 pp lower income mobility (a 4% decrease relative

to the base period mean). In each case, we observe a clear trend break in outcomes that is coincident with the month of application. Though not statistically significant, we also find suggestive evidence that individuals who move to LIHTC-funded units experience relative declines in local 3<sup>rd</sup> grade math scores (panel d) in the months following application.

Back-of-the-envelope calculations suggest that these effect sizes correspond to non-trivial impacts on long-run outcomes for both adult tenants of LIHTC-subsidized units and their children. Using data from the MTO experiment, Ludwig et al. (2012) find that a 13 pp decrease in neighborhood poverty increases subjective well-being by an amount equal to the gap in subjective well-being between people whose annual incomes differ by \$13,000 (2009 USD). Although this estimate is unlikely to scale linearly across treatment doses, it nonetheless suggests that a 1 pp *increase* in neighborhood poverty may have measureable negative effects on adults' well-being. Moreover, Chetty et al. (2016) find that children who move to lower poverty neighborhoods before age 13 earn \$3,477 more, on average, in annual income as adults (aged in mid-20s). Although the difference in neighborhood poverty rates was much more pronounced for movers in that study than in our context, we expect that there would likely be detectable negative impacts on adult earnings for children who move to LIHTC-funded properties in higher poverty neighborhoods, though these effects may only be on the order of a few hundred dollars per year.

The effects on poverty rates, predicted incarceration rates, and income mobility suggest that moving to LIHTC-funded units tends to make individuals worse off with respect to neighborhood level opportunity. However, "opportunity" is a complex and multi-faceted concept that cannot be summarized by one or two measures. For other outcomes, we do not detect statistically significant effects associated with moving to LIHTC-funded units. Unfortunately, given the current limitations of our sample size, these null effects are often not estimated with enough precision to rule out economically meaningful treatment effects in either direction. For example, for the local job concentration measure (jobs within 5 miles), we cannot rule out positive or negative effects up to 20% of the base period mean two years after application.

Our event study results in Figure 6 identify the relative difference between treatment and control groups over time on the intensive margin. However, we cannot directly infer from these results whether the average effects are driven by changes in neighborhood quality in the treatment group, in the control group, or in both groups. For example: can we attribute the positive effect on neighborhood poverty rates to: (1) moves by residents to LIHTC-funded units in poorer neighborhoods, (2) moves by non-residents to less poor neighborhoods in the post period, or some combination of the two? It is challenging to interpret the aggregate effects, and thus determine the policy implications of the results, without some understanding of the relative importance of these two margins.

To shed light on this question, we first analyze differences in outcomes between an individual's application tract and the tract of the LIHTC property associated with their application for all residents in our main sample. We plot the full distribution of individual-level changes separately for each outcome in Figure 7.<sup>25</sup> The x-axis corresponds to the value of the specified outcome in the application tract, and the v-axis corresponds to the value of the specified outcome in the LIHTC property tract. Intuitively, the further away a data point is from the 45 degree line, the greater is the implied change in the outcome that results from the move to a LIHTC-funded unit. For most measures, we observe considerable heterogeneity in individual-level outcomes, as represented by dispersion about the 45 degree line. In some cases (e.g., single parent rate) there is a wide range in observed changes from sending to receiving tract, but this variation is masked in the aggregate because the prevalence of opportunity-improving moves is roughly symmetrical to the prevalence of opportunity-reducing moves. This underlying heterogeneity suggests that moving to LIHTC creates a mix of winners and losers with respect to neighborhood opportunity, even in cases where there is little to no average change. In addition, we find evidence that the relative increases in the poverty rate from our event study analysis is driven, at least in part, by moves by treated individuals to LIHTC units in poorer tracts: in panel (a), the mass of data points lie above the 45 degree line, which indicates that placed individuals tend to end up in higher poverty tracts than they applied from. The results for the income mobility measure (panel f) are more mixed; most residents end up in tracts that are fairly similar to their application tracts, but there is a subset of individuals who start off in very high mobility tracts and do much worse after moving to LIHTC.

With the insights from this distributional analysis, we return to our event study design but now re-specify the outcomes of interest to capture effects on the *extensive* margin. Specifically, we construct a set of dummy indicators that compare the outcome for individual i at event time t to

 $<sup>^{25}</sup>$ To make the graphs legible, we use Stata's *binscatter* program to non-parametrically group individual-level data into bins of equal observations.

that individual's observed outcomes at time of application. We present results in Figure 8, focusing on the three opportunity measures for which we detect statistically significant treatment effects on the intensive margin. For each measure, we estimate separate regressions using indicators for (1) an increase relative to the application tract and (2) a decrease relative to the application tract, respectively, as the outcomes of interest.

Consistent with our results from Figure 7, we find evidence that average effects mask considerable heterogeneity on the extensive margin. In all cases, treated individuals are more likely to end up in both higher opportunity neighborhoods and lower opportunity neighborhoods compared to non-residents in the post period. In part, this reflects the fact that placed residents have a generally higher propensity to move after application compared to non-residents (see Figure 4). We also note, however, that the time path of treatment effects is suggestive of more complicated dynamics. Focusing on the poverty rate outcome, for example (panels a and b), we see that magnitudes for effects on both extensive margins (relative increases and relative decreases) are typically comparable for a few months immediately after application. Given that effects on the intensive margin are increasing, and that moves by residents drive most of the variation in this timeframe, we interpret this result as evidence that some placed residents move to LIHTC properties in much poorer neighborhoods and a roughly equal share move to LIHTC properties in slightly less poor neighborhoods. This interpretation is broadly consistent with the heterogeneity we observe in Figure 7. However, over time, moves by non-residents account for an increasing share of the variation, and we begin to see the effects on opportunity-improving moves (i.e., decreases in the poverty rate relative to application tract) converge back toward zero. At the same time, effects on the intensive margin continue to increase. A leading explanation for results in this later period is therefore that non-residents increasingly make opportunity-improving moves while residents tend to remain in place. This interpretation is also supported by the time series in Figure 5. We observe similar dynamics for the predicted incarceration rate and income mobility measures, although the signs are flipped for the latter since neighborhood opportunity is decreasing in poverty rates and predicted incarceration rates but increasing in predicted income mobility.

#### 5.5 Heterogeneity Analyses

We close the analysis by testing for heterogeneous effects based on property and applicant characteristics. In these analyses, we focus on the three outcomes for which we detect statistically significant effects in the aggregate analysis.

For most of our measurement period, projects built in Qualified Census Tracts, or QCTs, were awarded a 30% increase in tax credits. Previous work has shown that these extra incentives had measureable impacts on the spatial distribution of LIHTC-funded housing, as QCTs received twice as many subsidized units, on average (Baum-Snow and Marion, 2009). As we discuss in Section 3.3, QCTs are of particular interest in our setting because they are by definition the most economically disadvantaged neighborhoods in each metro area. Critics of the LIHTC have often expressed concerns that the QCT incentive structure may therefore play a role in "locking-in" LIHTC tenants in low-opportunity neighborhoods.

We re-estimate effects on the intensive margin separately for applicants to properties in QCT and non-QCT tracts, and present the results in Figure 9. The results strongly support the notion that QCTs play a central role in driving applicants to low-opportunity neighborhoods. For all three measures, we find that virtually all of the difference in aggregate outcomes between residents and non-residents is driven by properties located in QCT tracts. By contrast, we estimate relatively precise null effects when comparing applicants to properties in non-QCT tracts. Overall, we view these results as clear evidence that the design of developer incentives mediates the impacts of moving to LIHTC on neighborhood opportunity.

We also use our rich individual-level data to explore whether effects are heterogeneous by applicant characteristics. Although we do not directly observe the number of individuals in each applicant's household, we use applicants' unit size preferences – specifically, the preferred number of bedrooms – as a proxy for family size. We re-estimate effects on the intensive margin separately for individuals seeking 1, 2 and 3+ bedroom units, respectively, and present the results in Figure 10. In each case, we find that effects are concentrated among applicants seeking 1 bedroom units, with lesser or null effects for applicants seeking larger units. One interpretation that is consistent with these results is that applicants for smaller affordable units face more competition given the relative

scarcity of such units,<sup>26</sup> and so are more willing to trade off neighborhood quality for subsidized rents.

We also decompose the main effects by age group and present the results in Figure 11. This is an important dimension of heterogeneity because previous work has shown that the effects of neighborhood quality are highly heterogeneous by age. In particular, neighborhood quality likely has greater impacts on young children than on older children and adults (Chetty et al., 2016). The key pattern that emerges from this analyses is that effects tend to be driven by individuals aged 45-64. Taken together with the unit size analysis in Figure 10, the results suggest that the individuals who are most likely to be on the margin for making opportunity-reducing moves to subsidized LIHTC units are couples or single adults who are either childless or whose adult children have left the household, but who have not yet reached retirement age. It is possible that such households are more willing to trade off neighborhood quality for affordable rents than younger families, who may be more concerned about local factors that influence the development of children. Therefore, to the extent that treatment effects are concentrated among households that are less likely to have young children, some of the negative effects associated with moving to lower opportunity neighborhoods may be muted.

# 6 Discussion

Neighborhoods play a decisive role in shaping economic opportunity and access to valuable amenities for families. However, low-income renters are frequently priced out of living in high-opportunity neighborhoods. While the available evidence suggests that demand-side subsidies (e.g., vouchers) can be effective in promoting moves to better neighborhoods, such programs are expensive, oversubscribed, and constrained by resistance from private landlords. This project evaluates the potential of the Low-Income Housing Tax Credit (LIHTC), the nation's largest and fastest-growing housing affordability program, to promote moves to opportunity for low-income families by subsidizing the supply of affordable units. We construct a novel panel of address histories for applicants to LIHTC-funded housing by linking proprietary application data from a large LIHTC developer to

 $<sup>^{26}</sup>$ There is some empirical evidence that shortages of affordable units in major U.S. housing markets tend to be most severe for smaller units (i.e., studios and 1 bedroom units) (Scott, 2014). Anecdotally, property managers interviewed for this study also noted that smaller units tend to be easier to lease up than larger units.

administrative records. Within one year off application, applicants who successfully place into a LIHTC-funded unit reside in lower-quality neighborhoods compared to similar but unsuccessful applicants, and these effects tend to increase over time. In heterogeneity analyses, we find that aggregate effects are driven (1) by moves to especially high-poverty areas where federal incentives promote LIHTC development and (2) by smaller households that likely have few or no young children.

Our work makes two primary contributions. First, we shed new light on the impacts that LIHTC subsidies have on low-income renters. Prior research recognizes the LIHTC as an effective place-based policy that can increase local housing supply (Baum-Snow and Marion, 2009; Eriksen and Rosenthal, 2010), reduce crime (Freedman and Owens, 2011), and revitalize low-income neighborhoods (Diamond and McQuade, 2019) – benefits that housing vouchers do not offer. Other recent work focuses on modeling developers' response to LIHTC incentives and aggregate impacts on local housing markets (Soltas, 2023; Cook et al., 2023). However, the program's impacts on individual tenants remain highly underexplored, in large part because data on applicants to LIHTC-funded units is not publicly available. Drawing on our novel data, we provide the first credibly causal estimates of the individual-level effect moving to LIHTC on neighborhood quality. Second, we contribute new evidence on the role that programs that increase the supply of affordable units play in shaping access to opportunity. Our findings suggest that supply-side interventions such as the LIHTC may not be as effective, on average, as demand-side subsidies in inducing moves to better neighborhoods. However, we observe underlying heterogeneity that suggests that LIHTC subsidies create a mix of winners and losers. Since LIHTC is fundamentally a project-based rental assistance program and applicants take the supply of subsidized units as given, it therefore matters a great deal where developers are incentivized and awarded credits to build. Accordingly, we find that program incentives play an outsized role in mediating the impacts of moving to LIHTC on neighborhood opportunity.

Given the scale of LIHTC-funded housing in the U.S., these results have important implications for policies that seek to address issues related to housing affordability and economic opportunity. We find strong evidence that the design of specific incentives within the LIHTC program plays a central role in driving applicants to low-opportunity neighborhoods. Such policies may be defensible on other grounds, but a full accounting of their welfare impacts must include a consideration of these costs. Our findings also suggest that there is a trade off between the capacity of the LIHTC to increase the supply of affordable units and the potential of LIHTC to induce moves to better neighborhoods, since the supply response appears to be most concentrated in low-opportunity tracts. Finally, to the extent that program incentives "lock-in" low-income renters to low-opportunity neighborhoods, they may create spatial mismatches that have negative long run impacts on economic growth.

In future work, we will continue to study the individual-level impacts of LIHTC subsidies by linking our proprietary applicant sample to other outcomes from administrative data. In particular, we plan to analyze effects on financial health using credit data and effects on labor market outcomes using tax data.





*Notes:* LIHTC series data are from HUD's database of all LIHTC units placed in service 1987-2021, and reflect all low-income units in LIHTC-subsidized properties. Data on public housing, housing choice voucher, and project-based Section 8 are from Olsen (2003) (pre-1998), HUD Annual Performance Reports (1998-2015), and HUD's "Assisted Housing: National and Local" database (2016-2021).



Figure 2. LIHTC-Funded Properties Managed by Developer

*Notes:* The sample includes all LIHTC-funded properties managed by the developer who provided data for our analysis, regardless of target population or year placed-in-service. Unit counts reflect both LIHTC-subsidized units and market rate units.

Figure 3. Applicant Process Flow Diagram







*Notes:* The sample contains applicants in the main sample. The "Non-Resident" group includes applicants with the following statuses: Applied, Approved, Canceled, and Denied.



Figure 5. Mean Outcomes by Event Time and Treatment Group

Notes: The sample includes all applicants in the main analysis sample. For each specified outcome, we further restrict to a balanced panel of applicants with non-missing outcome data 2 years pre/post application.



Figure 6. Event Study Results: Intensive Margin

*Notes:* The sample includes all applicants in the main analysis sample. For each specified outcome, we further restrict to a balanced panel of applicants with non-missing outcome data 2 years pre/post application. Each figure plots point estimates and 95% confidence intervals from separate regressions of the outcome on event month treatment indicators and the baseline set of controls, following Equation 1. We cluster standard errors at the applicant level.



Figure 7. Characterizing Moves to Low-Income LIHTC Units

Notes: The sample includes all residents in the main analysis sample. We use Stata's binscatter program to collapse into bins that each contain 50 applicants.



Figure 8. Event Study Results: Extensive Margin

*Notes:* The sample includes all applicants in the main analysis sample. For each specified outcome, we further restrict to a balanced panel of applicants with non-missing outcome data 2 years pre/post application. Each figure plots point estimates and 95% confidence intervals from separate regressions of the outcome on event month treatment indicators and the baseline set of controls, following Equation 1. We cluster standard errors at the applicant level.



Figure 9. Event Study Results: Heterogeneity by QCT Status

*Notes:* The sample includes all applicants in the main analysis sample. For each specified outcome, we further restrict to a balanced panel of applicants with non-missing outcome data 2 years pre/post application. Each figure plots point estimates and 95% confidence intervals from separate regressions of the outcome on event month treatment indicators and the baseline set of controls, following Equation 1. We cluster standard errors at the applicant level.



Figure 10. Event Study Results: Heterogeneity by Unit Preference

*Notes:* The sample includes all applicants in the main analysis sample. For each specified outcome, we further restrict to a balanced panel of applicants with non-missing outcome data 2 years pre/post application. Each figure plots point estimates and 95% confidence intervals from separate regressions of the outcome on event month treatment indicators and the baseline set of controls, following Equation 1. We cluster standard errors at the applicant level.



Figure 11. Event Study Results: Heterogeneity by Age Group

Notes: The sample includes all applicants in the main analysis sample. For each specified outcome, we further restrict

*Notes:* The sample includes all applicants in the main analysis sample. For each specified outcome, we further restrict to a balanced panel of applicants with non-missing outcome data 2 years pre/post application. Each figure plots point estimates and 95% confidence intervals from separate regressions of the outcome on event month treatment indicators and the baseline set of controls, following Equation 1. We cluster standard errors at the applicant level.

	(1)	(2)		
	Data Provider	Other LIHTC		
Panel A. Property Characteristics	45 4	71.0		
# Units	45.4	71.2		
# Low Income Units	42.5	64.7		
Vacancy Rate - Low Income Units	0.042	•		
Applicants per Units	2.67			
Multiple Buildings	0.49			
HOME funding	0.40			
Rural Development funding	0.31	•		
Local funding	0.032			
Panel B. 2010 Census Tract Characteristics				
Pop per sq mi	1094.0	7080.8		
% Black	0.11	0.25		
% Hispanic	0.038	0.17		
% owner occupied units	0.56	0.44		
% college plus	0.18	0.21		
% in poverty	0.19	0.23		
% single parent HHs	0.38	0.46		
Ν	342	48103		
Panel C. HUD Characteristics				
New construction	0.63	0.55		
QCT	0.15	0.19		
Metro, outside city center	0.89	0.81		
Metro, in city center	0.037	0.12		
Non-metro	0.059	0.048		
Target $pop = families$	0.35	0.32		
N	270	48227		

# Table 1. Summary Statistics: LIHTC-Funded Properties

*Notes:* For HUD characteristics, the sample is restricted to properties that could be fuzzy matched to HUD's LIHTC Placed-in-Service Database.

	(1)	(2)	(3)	(4)
	Non-Residents	Residents	Raw Diff	Adj. Diff
Panel A. Application Characteristics				
BR preference	2.16	2.11	-0.045	$0.059^{**}$
Quoted rent	590.3	626.4	$36.1^{***}$	$53.8^{***}$
Age at application	45.1	50.2	$5.12^{***}$	$1.55^{***}$
Applied from same zip	0.29	0.47	0.18***	$0.14^{***}$
Applied from same tract	0.093	0.29	0.20***	0.18***
Panel B. Application Tract Characteristics				
Poverty rate	0.18	0.19	0.0039	0.014***
Annual job growth rate	0.0027	0.000	-0.0029	-0.0015
Pop per sq mi	2198.5	1986.8	-211.7	31.9
% Non-White	0.31	0.24	-0.071***	-0.0042
$P(\text{incarcerated}) \mid \text{parent inc } 25p$	0.027	0.025	-0.0028***	0.0005
$P(top 20\% income) \mid parent inc 25p$	0.094	0.092	-0.0026*	-0.0045***
Panel C. Property Characteristics				
# Low Income Units	49.6	49.5	-0.084	
Year placed in service	2011.9	2011.5	-0.40**	
Target pop. $=$ families	0.63	0.54	-0.094***	
QCT	0.25	0.25	0.000	
Income Limit = $50\%$ AMGI	0.10	0.099	-0.0054	
Income Limit = $60\%$ AMGI	0.88	0.87	-0.0059	
Metro: central city	0.045	0.050	0.0055	
Non metro	0.034	0.050	$0.016^{**}$	
HOME funding	0.44	0.48	0.031*	
N	2677	1878	4555	4524

 Table 2.
 Summary Statistics: Applicant Sample

*Notes:* The sample includes all applicants in the main analysis sample. In the final column, we estimate regression adjusted differences with property fixed effects and cluster standard errors at the applicant level. \* – significant at 10%, \*\* – significant at 5%, \*\*\* – significant at 1%.

# References

- Abadie, A., Athey, S., Imbens, G. W. and Wooldridge, J. M. (2023), 'When should you adjust standard errors for clustering?', The Quarterly Journal of Economics 138(1), 1–35.
- Acosta, S. and Gartland, E. (2021), 'Families wait years for housing vouchers due to inadequate funding', *Center on Budget and Policy Priorities*.
- Asquith, B. J., Mast, E. and Reed, D. (2023), 'Local effects of large new apartment buildings in low-income areas', The Review of Economics and Statistics **105**(2), 359–375.
- Baum-Snow, N. and Marion, J. (2009), 'The Effects of Low-Income Housing Tax Credit Developments on Neighborhoods', Journal of Public Economics 93(5-6), 654–666.
- Bergman, P., Chetty, R., DeLuca, S., Hendren, N., Katz, L. F. and Palmer, C. (2024), 'Creating Moves to Opportunity: Experimental Evidence on Barriers to Neighborhood Choice', *American Economic Review* 114(5), 1281–1337.
- Bernstein, S., Diamond, R., Jiranaphawiboon, A., McQuade, T. and Pousada, B. (2022), 'The contribution of high-skilled immigrants to innovation in the United States', *National Bureau of Economic Research*.
- Chetty, R., Hendren, N. and Katz, L. F. (2016), 'The Effects of Exposure to Better Neighborhoods on Children: New Evidence from the Moving to Opportunity Experiment', *American Economic Review* **106**(4), 855–902.
- Chetty, R., Hendren, N., Kline, P. and Saez, E. (2014), 'Where is the Land of Opportunity? The Geography of Intergenerational Mobility in the United States', *The Quarterly Journal of Economics* **129**(4), 1553–1623.
- Chyn, E. and Katz, L. F. (2021), 'Neighborhoods Matter: Assessing the Evidence for Place Effects', Journal of Economic Perspectives 35(4), 197–222.
- Collinson, R., DeFusco, A. A., Humphries, J. E., Keys, B. J., Phillips, D. C., Reina, V., Turner, P. S. and van Dijk, W. (2024), 'The Effects of Emergency Rental Assistance During the Pandemic: Evidence from Four Cities', *National Bureau of Economic Research* (No. w32463).
- Collinson, R., Ellen, I. G. and Ludwig, J. (2016), 'Low-income housing policy', Economics of Means-Tested Transfer Programs in the United States, Volume II 2, 59.
- Collinson, R., Humphries, J. E., Mader, N. S., Reed, D. K., Tannenbaum, D. I. and Van Dijk, W. (2022), 'Eviction and poverty in American cities', *National Bureau of Economic Research*.
- Cook, C., Li, P. and Binder, A. (2023), 'Where to Build Affordable Housing? The Effects of Location on Tenant Welfare and Segregation', *Working Paper*.
- Davis, M., Hartley, D. and Gregory, J. (2019), 'The long-run effects of low-income housing on neighborhood composition', *IDEAS Working Paper Series from RePEc*.
- Derby, E. (2021), 'Does Growing Up in Tax-subsidized Housing Lead to Higher Earnings and Educational Attainment?', *Working Paper*.
- Diamond, R. and McQuade, T. (2019), 'Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Low-Income Property Development', Journal of Political Economy 127(3), 1063–1117.
- Diamond, R., McQuade, T. and Qian, F. (2019), 'The effects of rent control expansion on tenants, landlords, and inequality: Evidence from San Francisco', American Economic Review 109(9), 3365–3394.
- Ellen, I. G., Horn, K., Kuai, Y., Pazuniak, R. and Williams, M. D. (2015), 'Effect of QAP Incentives on the Location of LIHTC Properties', U.S. Dept. of Housing and Urban Development, Office of Policy Development and Research
- Ellen, I. G., Horn, K. M. and Kuai, Y. (2018), 'Gateway to opportunity? Disparities in neighborhood conditions among Low-Income Housing Tax Credit residents', *Housing Policy Debate* 28(4), 572–591.

- Emmanuel, D. and Aurand, A. (2024), 'The Role of Vouchers in the Low-Income Housing Tax Credit Program', Cityscape 26(2), 209–222.
- Eriksen, M. D. and Rosenthal, S. S. (2010), 'Crowd out effects of place-based subsidized rental housing: New evidence from the LIHTC program', *Journal of Public Economics* **94**(11), 953–966.
- Freddie Mac (2023), Spotlight on Underserved Markets: Opportunity Incentives in LIHTC Qualified Allocation Plans, Technical report, Freddie Mac Multifamily, National Housing Trust, and Duty to Serve.
- Freedman, M. and Owens, E. G. (2011), 'Low-income housing development and crime', *Journal of urban economics* **70**(2), 115–131.
- HUD (2023), 'Low-Income Housing Tax Credit (LIHTC): Property Level Data'. URL: https://www.huduser.gov/portal/datasets/lihtc/property.html
- Keightley, M. P. (2023), An Introduction to the Low-Income Housing Tax Credit, CRS Report RS22389, Congressional Research Service.
- Ludwig, J., Duncan, G. J., Gennetian, L. A., Katz, L. F., Kessler, R. C., Kling, J. R. and Sanbonmatsu, L. (2012), 'Neighborhood effects on the long-term well-being of low-income adults', *Science* **337**(6101), 1505–1510.
- Mast, E. (2023), 'JUE Insight: The effect of new market-rate housing construction on the low-income housing market', *Journal of Urban Economics* **133**, 103383.
- McGuire, T. J. and Seegert, N. (2023), 'Is the Low-Income Housing Tax Credit an Effective Policy for Increasing Neighborhood Income Diversity?'.
- Mumphery, D., Johnson, C., Weismann, G., Knudsen, B. and DeFord, K. (2023), 'Housing Mobility Programs in the U.S. 2022', Poverty & Race Research Action Council.
- Olsen, E. O. (2003), Housing programs for low-income households, *in* 'Means-tested transfer programs in the United States', University of Chicago Press, pp. 365–442.
- Pennington, K. (2021), 'Does building new housing cause displacement? The supply and demand effects of construction in San Francisco', *Working Paper*.
- Phillips, D. C. (2020), 'Measuring housing stability with consumer reference data', Demography 57(4), 1323–1344.
- Picard, R. (2015), 'GEOINPOLY: Stata module to match geographic locations to shapefile polygons'.
- Reardon, S. F., Bischoff, K., Owens, A. and Townsend, J. B. (2018), 'Has income segregation really increased? Bias and bias correction in sample-based segregation estimates', *Demography* 55, 2129–2160.
- Sard, B., Rice, D., Bell, A. and Mazzara, A. (2018), 'Federal policy changes can help more families with housing vouchers live in higher-opportunity areas', Washington, DC: Center on Budget and Policy Priorities.
- Scott, M. (2014), 'Could our super-sized rental market be crowding out affordable housing opportunities?', Urban Institute, Urban Wire.
- Soltas, E. (2023), Tax incentives and the supply of affordable housing, Working paper.