# The Low-Income Housing Tax Credit and Economic Opportunity*

Henry Downes † Stephanie Karol ‡ John Soriano § George Zuo ¶

March 23, 2026

[Most recent version]

**Abstract**

We study how the Low-Income Housing Tax Credit (LIHTC), the nation's largest source of affordable rental housing, affects renters' neighborhood mobility and earnings. We construct a novel panel of applicant-level outcomes by linking proprietary application records from a major LIHTC developer to consumer reference data and federal tax records. On average, individuals who move to LIHTC properties do not experience significant changes in neighborhood quality relative to comparable unsuccessful applicants. However, null average effects mask considerable heterogeneity. In particular, residents of LIHTC properties located in Qualified Census Tracts – especially high-poverty tracts where federal rules provide additional tax credit subsidies – experience roughly 10% increases in neighborhood poverty relative to non-residents. LIHTC residence also results in slower wage growth among tenants: three years after application, residents' annual wages are roughly $3,000 lower than those of comparable non-residents.

**JEL Classification:** H20, H31, I38, R23

---

*This version: March 23, 2026. 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 Opportunities (LEO).

This research was conducted while Karol was an employee at the U.S. Department of the Treasury. The findings, interpretations, and conclusions expressed in this paper are entirely those of the authors and do not necessarily reflect the views or the official positions of the U.S. Department of the Treasury. Any taxpayer data used in this research was kept in a secured IRS data repository, and all results have been reviewed to ensure that no confidential information is disclosed.

We thank Rob Collinson, Bill Evans, Judy Hellerstein, Dan Hungerman, Michael Hollar, Melissa Kearney, Paul Odland, David Phillips, Jason Ward, Irakli Shalikashvili, and seminar participants at the University of Maryland, the University of Notre Dame, The University of Alabama, the 2023 and 2025 Association for Public Policy Analysis and Management meetings, the 2024 and 2025 Southern Economic Association meetings, and the 2025 North American Meeting of the Urban Economics Association for helpful comments and discussions. We express special gratitude for the partnership of a private LIHTC developer who provided data for this study.

†University of Georgia, B428B Amos Hall, Athens, GA 30606. Email: henry.downes@uga.edu

‡Office of Tax Analysis. U.S. Department of the Treasury, 1500 Pennsylvania Ave NW, Washington, DC 20220. Email: Stephanie.Karol@treasury.gov

§University of Dallas, 223 Satish & Yasmin Gupta College of Business, 1845 E Northgate Drive, Irving, TX 75062. Email: jsoriano@udallas.edu

¶RAND Corporation, 1200 S Hayes St, Arlington, VA 22202. Email: gzuo@rand.org

# 1 Introduction

The Moving to Opportunity (MTO) housing mobility demonstration catalyzed a wave of research showing that neighborhood environments profoundly shape outcomes for low-income families (Chyn and Katz, 2021). Now more than three decades later, policies modeled on MTO – vouchers paired with supports or incentives to move to higher-quality neighborhoods – remain fragmented and largely comprised of local pilots and small-scale regional demonstrations (Galvez and Oppenheimer, 2020; Mumphery et al., 2023). Yet in the absence of these targeted supports, voucher holders seldom leave high-poverty neighborhoods (Sard et al., 2018; Galvez and Oppenheimer, 2020; Collinson and Ganong, 2018). The persistence of post-MTO residential segregation by income and race (Reardon et al., 2018) further underscores the difficulty of achieving large-scale opportunity moves through demand-side programs alone.

In this paper, we shift attention to an understudied but important determinant of housing mobility and tenant outcomes: the Low-Income Housing Tax Credit (LIHTC). A supply-side policy, the LIHTC provides more than $14 billion annually in tax credits to developers who build and operate rental housing for low-income households. It is the nation's largest and fastest-growing source of federal affordable housing (Figure 1), financing more than 20% of all multi-family housing development over the past two decades (Diamond and McQuade, 2019) and supporting as many as 90% of newly-built affordable rental units (Axel-Lute, 2023).

We estimate how moving to LIHTC-subsidized housing affects renters' neighborhood quality and labor market outcomes. The LIHTC differs from tenant-based vouchers in two key ways with important implications for geographic and economic mobility. First, vouchers subsidize rental costs wherever families choose to live, whereas the LIHTC subsidizes a fixed set of affordable units. Accordingly, the LIHTC's mobility impacts depend crucially on where developers are incentivized to build. Because federal and state governments exercise substantial influence over LIHTC siting decisions (Ellen and Horn, 2017), they can, in principle, shape access to high-opportunity neighborhoods at scale by incentivizing development in those locations. Second, whereas voucher benefits phase out as earnings increase, LIHTC rents are generally fixed and do not mechanically decline as tenant earnings rise. The LIHTC therefore provides an opportunity to examine how housing subsidies affect labor market outcomes for low-income renters through other channels, including

income effects, neighborhood effects, and housing stability. In addition, because more than 40% of LIHTC tenants receive other forms of rental assistance (Emmanuel and Aurand, 2024), our estimates also reflect the interaction between the LIHTC and other housing programs that jointly shape outcomes for many assisted households in practice.

Most existing causal studies on the LIHTC focus on place-based effects, including impacts on local housing supply (Baum-Snow and Marion, 2009; Eriksen and Rosenthal, 2010), crime (Freedman and Owens, 2011), neighborhood composition (Diamond and McQuade, 2019; Davis et al., 2024; McGuire and Seegert, 2023), and aggregate housing markets (Soltas, 2023; Cook et al., 2023). We know far less about the LIHTC's impacts on individual renters, in large part due to limitations of existing data. As a tax credit initiative, the LIHTC is formally administered by the IRS, but detailed records on applicants, waitlists, and tenants are held as proprietary data by the thousands of developers and property management companies across the country that lease out LIHTC-subsidized units. Therefore, while the U.S. Department of Housing and Urban Development (HUD) monitors project-level compliance with federal affordability rules, the LIHTC's decentralized structure has so far inhibited the kind of individual-level analyses that are possible for other housing policies, including public housing and voucher assistance.

We address this limitation by obtaining proprietary applicant and tenant data from one of the nation's largest LIHTC developers, which we link to consumer reference data – high-frequency individual-level migration histories from Infutor[1] – as well as federal tax return data. For each applicant, we observe the timing of application, stated unit preferences, screening denials, waitlist activity, and the ultimate outcome of the application, along with a rich set of property-level characteristics. These linked data allow us to construct individual-level panels that track the outcomes of both successful applicants who become LIHTC residents and unsuccessful applicants who do not.

After excluding applicants who were denied housing due to ineligibility, we organize the data as applicant-by-event-time panels centered on the application date and estimate effects using complementary event study and synthetic difference-in-differences (SDID) designs. For neighborhood outcomes, our event study compares applicants to the same LIHTC property after adjusting for

[1]This dataset has been used 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).

time and seasonality trends. For labor market outcomes, we leverage the increased size of the linked tax data sample to implement a matched-pair event study that compares residents to similar non-residents from the same origin zip code. We use SDID in both settings because applicants who become LIHTC residents may differ from other applicants along unobserved dimensions that shape their untreated outcome trajectories even after we condition on detailed observables. In settings with such latent factor structure, Arkhangelsky et al. (2021) show that SDID can reduce bias relative to standard two-way fixed effects (TWFE) event studies by choosing unit and time weights that better reproduce the treated group's pre-treatment outcome path. This concern is especially salient for labor market outcomes, where eventual LIHTC residents and non-residents exhibit divergent earnings trajectories even before application. In augmented specifications, we include additional controls that flexibly interact application characteristics – including bedroom preference and applicant demographics – with event time indicators.

We first examine effects of LIHTC tenancy on neighborhood quality. We proxy for neighborhood quality at the Census-tract level using three primary outcomes: poverty rates, intergenerational mobility rates, and median rents.[2] Although LIHTC developments are disproportionately located in low-opportunity tracts, individuals who move into LIHTC properties do not experience significant changes in neighborhood conditions relative to comparable unsuccessful applicants, on average. Our estimates are sufficiently precise to rule out even modest aggregate effects; for example, we bound changes in neighborhood poverty to within 0 to +4 percent of the baseline mean. We show that the main results are stable across specifications and are robust to a variety of alternative sample restrictions.

Importantly, we find that average effects on mobility mask policy-relevant heterogeneity. We assess whether treatment effects on neighborhood quality differ depending on whether a property is located in a Qualified Census Tract (QCT). QCTs are especially low-income neighborhoods that are targeted for additional federal LIHTC subsidies.[3] We find that applicants who move to LIHTC properties located in QCTs experience average increases in neighborhood poverty of about 2 percentage points (pp), roughly a 10% increase relative to the baseline mean. These

[2] In supplementary analyses, we examine effects on a broader set of measures capturing other dimensions of neighborhood quality.

[3] A QCT is defined as a tract with either (1) a poverty rate of at least 25 percent or (2) median household income below 60 percent of area median gross income. Since 1990, federal rules have granted LIHTC projects located in QCTs a 30 percent increase in eligible credits.

patterns suggest that current federal siting incentives steer households toward lower opportunity neighborhoods.

We also explore heterogeneity along other margins related to project financing and location. Our results suggest that local supply-side factors play an important role in shaping mobility outcomes for LIHTC tenants. In particular, moves to LIHTC properties in higher-cost suburban tracts are associated with relative increases in neighborhood poverty. More broadly, we find that moves to properties that are more exposed to competitive pressures (as measured by an increased reliance on competitive 9% LIHTC credits) tend to be associated with increases in neighborhood poverty. These patterns are consistent with interactions between state-level credit allocation processes that reward lower-cost development and local land use constraints that limit the construction of affordable housing in high-amenity areas.

We next examine labor market outcomes using federal tax records. We focus on two primary measures, both available with annual frequency: 1) total W-2 wages and 2) an approximation to the income definition used to determine an applicant's eligibility to occupy a rent-restricted unit financed using LIHTC. This income definition includes W-2 wages, and adds non-wage income sources such as safety-net transfers and income from businesses or property.

We measure sizable negative effects of LIHTC residency on earnings. In both matched-pair event studies and SDID specifications, LIHTC residents experience substantially slower post-application earnings growth than comparable non-residents. By the third year after application, residents earn roughly $3,000 less in annual wages, and we find similarly negative effects on broader income measures. These magnitudes are consistent with prior evidence on the earnings impacts of public housing residency (Susin, 2005) and voucher utilization (Jacob and Ludwig, 2012). However, LIHTC subsidies do not necessarily vary with the tenant's own income as in those programs, which theoretically limits the role of substitution effects.

To further investigate the mechanisms that underlie the effects on earnings, we estimate the implicit housing subsidy received by LIHTC tenants by comparing applicant-level quoted rents to tract-by-bedroom market rents from the ACS. We estimate that LIHTC tenants receive an average subsidy of about $1,560 per year, or about 50% of the corresponding decline in earnings.[4] As

[4]We interpret this magnitude with some caution, because quoted rents, utility allowances, and non-residents' own housing subsidies may be measured with error.

such, income effects from the direct rent subsidy alone is unlikely to explain the full magnitude of observed earnings impacts. We therefore hypothesize that other channels – including interactions with other rental assistance programs that impact net wage rates and possible "lock-in" effects from increased housing stability – may also play an important role.

This paper provides the first causal evidence on how the LIHTC shapes residents' neighborhood environments and economic outcomes, using a novel dataset that overcomes long-standing data constraints. Given its scale and prominence, understanding how the LIHTC shapes economic mobility is central to understanding how housing policy affects access to opportunity in the United States. More broadly, our work sheds light on how supply-side housing policies impact residential mobility and access to opportunity. Our findings indicate that several key features of the LIHTC – including QCT credit boosts and other siting incentives – may work at cross-purposes with demand-side programs that seek to expand access to high-opportunity neighborhoods. These features reflect the LIHTC's supply-oriented design, in which families must accept housing where it is built rather than choose among locations in a more unconstrained manner. Our analysis thus highlights both the promise and the limitations of a system in which federal and state incentive structures shape where development occurs. Our results also highlight how the structure of housing subsidies shapes the full bundle of housing, neighborhood access, and labor market incentives available to low-income households. While the LIHTC provides measurable housing subsidies to low-income renters, a complete accounting of its welfare impacts must weigh these benefits against the harder-to-measure effects on tenant outcomes that we document.

# 2 The Low Income Housing Tax Credit (LIHTC)

## 2.1 Overview of the LIHTC

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 tax incentive has become a bedrock of the U.S. housing landscape, with over 3.3 million affordable units placed in service between 1987-2023. Today, the LIHTC costs an estimated $13.5 billion annually (Keightley, 2023). Figure 1 shows that the LIHTC is both the largest and fastest-growing federal source of affordable housing units in the U.S. These tax credits have helped to finance more than 20% of all multi-family housing

development over the past two decades (Diamond and McQuade, 2019).

Unlike federal housing programs that provide direct subsidies to renters, the LIHTC seeks to increase the supply of affordable units by offering developers tax credits to build rent-restricted housing for low-income tenants.[5] There are two distinct types of LIHTC tax credits: the 70% credit – which is awarded by state housing credit agencies through a competitive process—and the 30% credit, which is not competitive.[6] Here, "70%" and "30%" refer to the present value of the tax credits a successful LIHTC project can receive; i.e., projects financed with the 70% credit receive tax credits with a present value of 70% of the "qualified basis" of the building (i.e., the share of the building devoted to affordable housing). Upon application for LIHTC financing, prospective developers choose from a menu of options for rent-restricted units set aside for low-income 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.[7] Rents for these units cannot exceed 30% of the corresponding AMI threshold.

Because LIHTC affordability requirements bind at the property level, rents for LIHTC-subsidized units are not tied to individual tenants' incomes in the way that direct rental subsidies (e.g., vouchers) are. Under the LIHTC, tenant income is assessed at move-in to establish initial eligibility under the applicable income limits and is typically re-certified in subsequent years to verify continuing compliance. If a tenant's income increases over time to exceed 140% of their unit's applicable income limit, the tenant does not necessarily lose their rent restriction and is not evicted. Instead, properties can remain in compliance by renting the next available comparable unit to a qualified low-income tenant. LIHTC affordability requirements bind for 15 years and properties can gradually phase out of compliance over the course of the subsequent 15 years.

[5] 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 remains in compliance with the LIHTC rules.

[6] 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.

[7] There is a third criterion, added in 2018, which 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.

## 2.2 Conceptual Impacts of the LIHTC on Mobility

The LIHTC was designed to expand the supply of rental housing affordable to low-income households rather than to specifically promote moves to higher-quality neighborhoods. Nevertheless, the program may influence the mobility of low-income renters through several channels. Given the scale of LIHTC development in the United States, understanding these potential mobility effects is of clear policy importance.

First, tenants of LIHTC-funded housing typically pay below-market rents, which may allow low-income households to access high-amenity neighborhoods that would otherwise be unaffordable. Because LIHTC rent limits are tied to area-level median income rather than neighborhood-level housing costs, the program effectively flattens the housing price gradient between higher- and lower-opportunity neighborhoods. Drawing on a sample of LIHTC developments in the Chicago MSA, Cook et al. (2023) shows that predicted LIHTC rents in the lowest-opportunity neighborhoods are, on average, similar to those in the highest-opportunity neighborhoods. As a result, tenants in high-opportunity neighborhoods receive substantially larger implicit subsidies – up to 41% relative to market rents—while tenants in low-opportunity neighborhoods receive considerably smaller discounts (around 18%). In this way, LIHTC subsidies may reduce barriers to moving to high opportunity neighborhoods for low-income renters.

LIHTC may also affect mobility through its interaction with other forms of rental assistance. Although housing vouchers remain the primary policy tool for promoting mobility among low-income renters, existing research suggests that demand-side subsidies alone are often insufficient to help households access high opportunity neighborhoods (Lens et al., 2011; Collinson and Ganong, 2018). Voucher programs are highly oversubscribed, and households wait an average of 2.5 years – and as long as eight years in some markets – before receiving assistance (Acosta and Gartland, 2021). After obtaining a voucher, a household must typically secure housing within 60–90 days. Given these time constraints, and the fact that landlord refusal of vouchers remains common (Cunningham et al., 2018; Sard et al., 2018), only 60% of voucher recipients successfully utilize their voucher to lease a unit (Ellen et al., 2024), with even lower success rates in high quality neighborhoods. LIHTC properties, by contrast, are required to accept vouchers as a form of rental payment. To the extent that the limited mobility of voucher recipients to higher opportunity neighborhoods is the result

of supply-side constraints and not individual preferences, the LIHTC may complement the goals of voucher programs – particularly in high-amenity neighborhoods and in markets where landlords are permitted to discriminate on the basis of voucher status. In practice, a substantial share of LIHTC tenants receive additional rental assistance: nationally, about 40% of LIHTC tenants receive some form of other assistance, including 28.4% who receive federal assistance such as Housing Choice Vouchers or project-based housing subsidies (Emmanuel and Aurand, 2024).[8]

Because LIHTC subsidies are tied to specific properties rather than households, the program's effects on the mobility of low-income renters depend critically on where subsidized units are built. State housing agencies, in turn, play a central role in shaping the realized distribution of LIHTC development through the implementation of Qualified Allocation Plans (QAPs). QAPs provide opportunities for state agencies to encode their siting and policy preferences through the design of scoring systems, set-asides, and tiebreakers that determine how tax credits are competitively allocated to proposed developments. Many states have incorporated opportunity-based criteria into their QAPs; for example, by awarding points for constructing units in areas with high-performing schools, strong job markets, and access to health care, recreation, and transportation (Freddie Mac, 2023). Empirical work shows that strengthening these QAP incentives can measurably shift funded developments into higher-opportunity areas (Ellen et al., 2015; Ellen and Horn, 2017).

Other salient QAP features, however, may instead funnel development, and thus low-income renters, toward lower-opportunity neighborhoods. While state agencies do not simply award competitive 70% credits to the lowest-cost project, cost reasonableness and financial feasibility are typically central evaluation criteria. Since many of the factors that allow developers to achieve low costs – for example, low land prices, public-private partnerships that allow developers to access discounted or donated land, or more permissive zoning – are also likely to also steer affordable development toward low-amenity neighborhoods, cost-containment pressures may be at cross-purposes with incentives to build in high-opportunity areas. Moreover, unlike project-based rental assistance programs, LIHTC developments are generally not subject to HUD siting guidelines, including those intended to prevent the concentration of assisted housing in high-poverty neighborhoods.

Finally, federal incentives also play a part in determining the spatial distribution of LIHTC-

[8] These shares can be substantially higher in some markets. For example, the Ohio Housing Finance Agency reports that roughly 70 percent of LIHTC tenants in the state receive some form of rental assistance (Keithley, 2022).

funded units. In particular, federal rules provide a 30% tax credit boost to projects built within qualified low-income census tracts ("QCTs"). A tract is designated as a QCT if either 1) at least half of resident households have incomes below 60% of Area Median Gross Income, or 2) the tract has a poverty rate of at least 25%. Previous work shows that QCTs have historically received twice as many subsidized units as non-QCTs (Baum-Snow and Marion, 2009). Because QCTs are, by definition, the most economically disadvantaged tracts in each metropolitan area, critics argue that the incentive may concentrate the supply of affordable units in low-opportunity neighborhoods (Kawitzky et al., 2013; Galster, 2019).

Given these various misaligned incentives, it is ambiguous *ex ante* whether the LIHTC promotes or hinders moves to opportunity on average. In this paper we seek to shed new light on that question, as well as understand how these impacts may vary across policy settings, project types, and renter characteristics.

### 2.3 Conceptual Impacts of the LIHTC on Labor Market Outcomes

We conceptualize LIHTC residence as a bundled treatment that affects labor market outcomes through 1) income and substitution effects via an effective housing price change, 2) changes in housing stability, and 3) exposure to a new neighborhood.

The strength of income effects depends on the real value of the LIHTC subsidy passed-through to tenants, measured as the gap between the LIHTC rent and the individual's counterfactual rent. If below-market rate LIHTC rents generate meaningful savings, households may reduce labor via income effects. The real value of LIHTC subsidies is difficult to measure in practice, but available estimates suggest that they may not be particularly large; for example, Soltas (2023) estimates that LIHTC rents are only 12% lower than market rate rents for comparable units.

Substitution effects in this context are likely to be more complicated. In Section 2.1, we describe how LIHTC subsidies are not as closely tied to renters' current incomes as rental subsidies that are provided through vouchers and project-based housing programs. Based on these structural differences alone, the LIHTC may not generate the same substitution effects that arise under other rental assistance programs (Jacob and Ludwig, 2012). Even so, to the extent that LIHTC rental arrangements frequently involve interactions with housing vouchers and project-based rental assistance, we would expect that the same substitution effects associated with those programs may

be observed among LIHTC tenants.

Beyond direct impacts on residents' budgets, the rent restrictions associated with LIHTC units may also affect labor market outcomes by increasing housing stability. On the one hand, existing literature finds evidence of meaningful negative impacts of eviction and housing instability on employment and earnings (Collinson et al., 2022). The LIHTC may thus raise employment and earnings by reducing forced moves, commuting disruptions, and the time and stress costs of housing insecurity. On the other hand, because LIHTC assistance is not portable but attaches to a specific unit, the LIHTC may raise the opportunity cost of relocation, inhibiting tenants from pursuing new opportunities that require moving. Such "lock-in" effects may depress earnings growth even when employment remains stable.

Finally, the neighborhoods in which LIHTC properties are located may influence tenants' labor market outcomes. Neighborhood exposure may affect labor outcomes through job accessibility, commuting costs, and local information or referral networks, and these effects may operate on a hyper-local level (Billings et al., 2026). Recent evidence from the Moving to Opportunity (MTO) experiment suggests that households induced to move to higher-opportunity neighborhoods may experience positive neighborhood effects on adult labor-market outcomes (Aliprantis and Richter, 2020). Whether moves to LIHTC-funded units tend to improve or reduce labor market opportunities through neighborhood channels therefore depends on whether such moves lead tenants to higher- or lower-opportunity areas – an empirical question we examine in this paper.

## 3 Data

Due to the highly decentralized nature of the LIHTC, no national dataset – including major surveys such as the ACS and CPS – identifies individual residents of LIHTC-funded housing (Collinson et al., 2016). Several recent studies examine tenant-level outcomes by linking the locations of LIHTC properties to administrative records (Derby, 2021; Cook et al., 2023). However, these approaches cannot identify individuals who applied to LIHTC units but were not admitted. In the absence of such a natural comparison group, it is challenging to characterize tenants' counterfactual outcomes. Since government agencies do not collect applicant data, such information can only be

obtained from proprietary developer or property management records.[9] In this section, we describe how we link application and tenant records from one of the nation's largest LIHTC developers to administrative tax data and residential address histories. The resulting dataset is the first to track outcomes for both applicants who are admitted to LIHTC-funded units and those who are not in the years immediately before and after their application.

### 3.1 Applicant and Tenant Data

We draw on applicant and tenant data from a large, private developer of affordable housing. With more than 30,000 residents and 300 LIHTC properties in over a dozen states, the developer ranks among the nation's largest developers of new LIHTC units (HUD, 2023). Importantly, the developer centrally manages all properties, ensuring that applicant and tenant data are collected and maintained under uniform standards across sites. Figure 2 shows that the developer has a presence in the Midwest, Appalachia, and the Southeast. The portfolio includes developments in several major cities – including Atlanta, Charlotte, Cincinnati, Cleveland, Columbus, and Indianapolis – but also includes large concentrations in suburban and exurban counties. The wide range of neighborhood types in our sample thus reflects the breadth of the LIHTC's presence across varied U.S. housing markets.

Since 2012, the developer has tracked all interactions with prospects, applicants, and tenants at each property using an internal client management database. Figure 3 depicts the end-to-end flow for the application and placement processes. In general, applications are reviewed on a first come, first served basis. Property staff primarily rely on unit preferences (e.g., number of bedrooms, accessibility features) and income to match applicants to available units. If no suitable unit is available, the applicant is placed on a waitlist according to their desired unit type. If a suitable unit is available, the applicant undergoes an eligibility screening, which involves criminal background and credit checks as well as third-party income verification. If the applicant is determined to be eligible, they can accept the offered unit and begin planning their move-in. At any point after an

[9] 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."

application is submitted but before a unit is accepted, an applicant may cancel their application.

We observe the following application information for all applicants: full name, date of birth, Social Security number, current address, information about the property they are applying to, unit preferences, the application date, and the terminal status of their application. Though coverage is not complete, we also often observe the monthly rent amount quoted to the applicant at the time of application ("quoted rent").

We also obtain property-level characteristics, including each property's name, main address, date placed-in-service,[10] funding sources, tax credit details, and applicable affordability rules. As some properties are "scattered sites" with units spread across multiple locations, we also observe all unit-level addresses associated with each property, allowing us to identify the full set of locations to which placed tenants may move.

### 3.2 Consumer Reference Data

To assess the impacts of moving to LIHTC-subsidized housing on neighborhood-level outcomes, we link applicant and tenant records from the client management database to migration histories available from Infutor Data Solutions' consumer reference records. Infutor aggregates information from numerous private and public record sources, including USPS change of address forms, paid subscriptions, utility records, and phone books. 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, which allows us to sequence an individual's move history over time. Because Infutor provides name aliases, we can further link records for individuals associated with multiple name variants, including those whose surname changes over time. Previous work finds that Infutor is highly representative of the U.S. adult population, especially since the 1990s.[11] Phillips (2020) shows that Infutor move dates can detect

[10] 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.

[11] 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. 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.

highly-localized migration shocks, including out-migration from Hurricane Katrina and closures of public housing projects in Chicago. Infutor data has been used in previous work 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).

In Section 5.1, we describe how we construct a panel of neighborhood-level outcomes by using personally-identifiable information to link applicants and tenants to the Infutor database.

### 3.3 Administrative Tax Data

To measure the impact of LIHTC subsidies on individual-level labor market outcomes, we draw on tax data from applicants' income tax returns (Form 1040), as well as select information returns that provide third-party reports of individuals' wages (Form W-2), unemployment compensation (Form 1099-G), and interest income (Form 1099-INT). Linking to Forms 1040, W-2, 1099-G, and 1099-INT allows us to measure income according to several different definitions. HUD uses a specific definition of income to determine LIHTC eligibility,[12] that does not correspond precisely to the definition of income used to determine income tax liability.[13] However, the Form 1040 includes a variety of line items that can be used to approximate HUD's definition of income. We construct our approximation of HUD income by adding back the non-taxable portions of interest income, dividends, IRA distributions, pensions, annuities, and Social Security benefits to total income.[14] When an applicant does not file a Form 1040 in a particular year, we impute the approximate HUD income as the sum of that individual's wages from Form W-2, unemployment compensation from Form 1099-G, and interest income from Form 1099-INT. By aggregating information from the Form W-2, we can construct an alternative definition of income: the total wages earned by the applicant and their spouse, if any.

In Section 6.1, we describe how we construct a panel of labor market outcomes by linking applicants to their tax records using Social Security Numbers.

[12] https://www.hud.gov/sites/documents/calculatingattachment.pdf

[13] For example, HUD's definition of income includes TANF and SNAP benefits, which are not taxable and so do not appear on the Form 1040.

[14] Total income appears on Form 1040, line 9. The other line items appear on Form 1040, lines 1 through 6. On each line $l$, the full value of the line item appears on line $l$a, whereas the taxable portion appears on line $l$b. This taxable portion is included in total income. To add in the non-taxable portion, we subtract $l$b from $l$a and add that amount to total income.

### 3.4 Publicly Available Data

**HUD LIHTC Property Database** HUD maintains a publicly available database of all LIHTC-funded projects placed in service since the tax credit'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, 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.[15]

**Opportunity Atlas** These data provide observed and predicted socioeconomic mobility indicators at the Census tract-level (Chetty et al., 2014) that proxy for various dimensions of neighborhood quality in the mobility analyses. From a very broad set of possible outcomes, we restrict our attention in the main analyses to two "opportunity" measures: (1) the poverty rate, a contemporaneous measure that is directly linked to LIHTC development patterns through the QCT incentive; and (2) the predicted probability of reaching the top national income quintile conditional on being born to parents in the bottom quartile of national income, a more explicit measure of intergenerational economic opportunity. In supplementary analyses we consider a wider range of neighborhood-level outcomes, including: district-level 3$^{rd}$ grade math test scores, the single parent share, the share of non-white residents, job counts within five miles, the average annualized job growth rate, and the predicted probability of adult incarceration conditional on being born to parents in the bottom quartile of national income.

**American Community Survey (ACS)** We observe median market-rate rents at the Census tract-bedroom level from ACS 5-year estimates (2015-2019). Market rents capture households' willingness to pay for a variety of local amenities and disamenities and so provide a useful summary measure of neighborhood quality in the main mobility analyses. We also use data on market rents to back out estimates of the implicit subsidy that tenants receive once they are admitted to

[15] One important shortcoming of these data are that they only include one address for each property, which is problematic for multi-site properties. These addresses also commonly reference property management offices instead of actual units.

LIHTC-funded units.

### 3.5 Summary Statistics and Sample Representativeness

More than 2% of all U.S. households live in LIHTC-funded units (Soltas, 2023). LIHTC features prominently in most housing markets and so reflects the wide variance in neighborhood and housing conditions across the country. Given that our available data from a single developer covers a relatively small share of all LIHTC-funded properties in the U.S., it is important to characterize the ways in which our sample may or may not generalize to other properties outside the sample.

We compare features of properties in our sample to a national sample of all LIHTC-funded properties in Table 1. In both samples, we restrict to the set of properties that do not specifically target elderly residents.

Properties in our sample tend to be located in areas with significantly lower population density, higher White shares of the population, and lower market rents relative to the typical LIHTC property. However, in terms of opportunity-related neighborhood measures – the poverty rate and predicted income mobility – in-sample properties more closely resemble other LIHTC properties. In general, we find that LIHTC properties tend to be located in relatively disadvantaged neighborhoods, with the average LIHTC tract falling near the 75th percentile nationally in terms of poverty rate.[16] A non-trivial share of properties both in our sample (16%) and in the national sample (21%) are located in Qualified Census Tracts (QCTs).

We also compare samples across a range of property-level characteristics. For these variables, we further limit the sample to those properties that can be successfully linked to non-missing data in the HUD PIS database.[17] Properties in our sample tend to be smaller than other LIHTC properties, as reflected by lower average unit counts. Common to both samples, however, is the very high share of total units that are subsidized: 98% in our sample and 89% in the national sample. These figures are well above HUD compliance minimums and consistent with a competitive credit allocation process that forces developers into greater low-income commitments than is minimally required. Properties in our sample are 18 percentage points (pp) more likely to be funded by competitive 70% credits

[16] 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.

[17] We fuzzy match property names in our sample to those listed in the HUD PIS database after blocking on zipcode and project stage. We manually inspect fuzzy matches and only retain high quality matches.

than properties in the national sample, which suggests that our sample is especially well-positioned to detect the impacts of siting incentives embedded in state QAPs.[18] As in the national sample, properties in our sample are often characterized by a mix of public funding sources (i.e., "layered funding"), including common interactions with Project-Based Rental Assistance Contracts (both state and federal), Rural Development 515 loans, and HOME grants. To the extent that such "layered" projects are subject to various affordability rules and siting incentives, we expect that the effects of becoming a tenant in a LIHTC-funded unit will reflect the combined features of both LIHTC and complementary programs.

Taken together, the properties in our data are located in areas that are less dense and racially diverse compared to national LIHTC averages, but appear otherwise representative in terms of key measures of economic opportunity. Properties in both samples are characterized by high shares of subsidized units and layered funding from complementary programs. These similarities suggest that our findings could credibly generalize to other settings, and especially to those where LIHTC development is concentrated outside of central cities.

# 4 Empirical Strategy

In this section, we begin by outlining a basic conceptual framework for modeling search and placement into LIHTC-subsidized units. We then present our empirical strategy for estimating the causal effects of moving to a LIHTC unit on measures of neighborhood quality and labor market outcomes.

## 4.1 Conceptual Framework

We first assume that low-income households are characterized by heterogeneous human capital ($H_i$) and preferences over housing attributes ($\phi_i$). Households initially reside in market rate housing and are induced to apply to at least one LIHTC-subsidized unit due to an increase in current rents or some other budgetary shock. At the same time that they search for a subsidized unit, households may apply to other market rate units. The probability of a vacancy opening for a LIHTC-subsidized unit or other market rate unit in any period is given by $p_i$, where $i$ indexes households. Since the

[18] It is possible for projects to be financed by both 70% and 30% credits. As a result, the sum of the means of these variables exceeds 1.

demand for and supply of subsidized and market rate units varies according to applicant and unit characteristics, $p_i$ is a function of $H_i$ and $\phi_i$.[19]

In each period, households consider only the first option $j$ that becomes available. When a vacant LIHTC-subsidized or other market rate unit becomes available, they choose according to the decision rule:

$$\max\left(\frac{\tilde{u}_{it}(j)}{1-\delta}, V_{it}\right)$$

$\tilde{u}_{it}(j)$ represents the per-period flow utility of a prospective new unit $j$, which is also a function of $H_i$ and $\phi_i$ because applicants with different underlying characteristics and preferences are likely to face different offer sets. $\delta$ is the discount factor, so the first term in the maximand represents the present value of lifetime utility gained from the new unit. For parsimony we assume that once a household moves to a new unit (LIHTC or otherwise) they remain there indefinitely. The second term represents the value of continuation: rejecting the open unit, remaining at the current address, and searching again next period. We define the present value of continuation as:

$$V_{it} = (u_i - c_i) + \delta\mathbb{E}\left[(1-p_i)(u_i - c_i) + p_i \cdot \max\left(\frac{\tilde{u}_{i,t+1}}{1-\delta}, V_{i,t+1}\right)\right]$$

where $u_i$ is per period flow utility from remaining at the current address, and $c_i$ is the per period search cost. Both are time-invariant, known to the household, and are functions of $H_i$ and $\phi_i$. When a household opts to continue searching, they receive $(u_i - c_i)$ in the current period – i.e., the utility of remaining at the current address, net of search costs. In the next period, the household would again receive $(u_i - c_i)$ if no vacancy opens, which occurs with probability $1 - p_i$. If a vacancy opens (with probability $p_i$), the household's choice would depend on their specific draw of $\tilde{u}_{i,t+1}(j)$ from the distribution of $\tilde{u}_{i,t+1}$. Importantly, the household does not know *ex ante* which draw of $\tilde{u}_{i,t+1}(j)$ they will get next period, so the continuation value is a function of the expected value of $\tilde{u}_{i,t+1}$, which is also a function of $H_i$ and $\phi_i$.

This basic model illustrates several key features of our setting. First, the value of continuation is increasing in $u_i$ and decreasing in $c_i$, which implies that households with better initial housing situations and lower search costs face a higher acceptance threshold and are thus less likely to accept a given LIHTC offer. Second, the value of continuation is increasing in the expected value

[19] We assume that the probability of a vacancy opening for a given individual $i$ is time-invariant.

of a new unit ($\mathbb{E}(\tilde{u}_{i,t+1})$) but the acceptance probability for a new vacant unit is decreasing in the realized value of that unit ($\tilde{u}_{it}(j)$). In other words, households who expect to have many desirable options in the future will be more likely to pass on vacancies in general, but a good draw will make any household more likely to accept a given vacant unit. Under an additional assumption that subsidized units tend to be more desirable than unsubsidized units, it follows that individuals who have the opportunity to rent a vacant LIHTC unit will be more likely to accept the new unit and will remain at the application address for a shorter duration, all else equal.

The model also highlights that heterogeneous households are unlikely to have comparable unconditional probabilities of placing into LIHTC-subsidized units. For example, households with higher human capital may have earnings that render them ineligible for some income-restricted LIHTC units, thereby affecting both the expected value of future offers $\mathbb{E}(\tilde{u}_{i,t+1})$ and the vacancy arrival rate $p_i$. Likewise, strong neighborhood preferences alter the set of units to which households apply, shaping the distribution of potential offers. More generally, any heterogeneity in $H_i$ and $\phi_i$ that induces variation in $p_i$, $u_i$, $c_i$, or $\mathbb{E}(\tilde{u}_{i,t+1})$ generates differential selection into LIHTC tenancy. Because these same factors are likely correlated with subsequent outcomes, failing to account for them will confound estimates of the causal effect of moving into a LIHTC unit.

Conditional on $H_i$ and $\phi_i$, the underlying (residualized) distribution of $\tilde{u}_{it}$ does not vary across applicants. Importantly, however, realized treatment status may still differ across otherwise comparable applicants due to idiosyncratic variation in the timing of offer arrival and acceptance. For example, consider two comparable applicants who apply for the same subsidized unit, just days apart. Since applications are typically reviewed in the order in which they are received, the applicant who was first to submit may have the first decision on the vacant unit. If they choose to accept it, the other applicant will be forced to continue searching. In the meantime, they would remain at their application address, and are strictly more likely to accept a market rate unit in some future period and cancel their LIHTC application. In this way, idiosyncratic factors related to the specific timing of application and churn at a given property can lead to good-as-random variation in treatment status, conditional on $H_i$ and $\phi_i$. In the next section, we discuss how our empirical design seeks to isolate variation in treatment that is orthogonal to $H_i$ and $\phi_i$.

### 4.2 Identification

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 admitted 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 household $i$'s residence at time $t$. The first term, $\mathbb{E}[Y_{i,t=2}(2)|D = 1]$, is interpretable as the mean outcome for households in the treated group after treatment, and is directly estimable from the data. In our context, this refers to the average neighborhood quality or labor market outcome 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.

A key feature of our design is that we are able to construct a panel of outcomes for both placed and non-placed applicants by linking households from the applicant and tenant sample to their own records in the Infutor and IRS databases, respectively. As we discuss in Section 4.1, unconditional comparisons of applicants who do versus do not place into LIHTC units are likely to be contaminated by various forms of selection bias, including important differences in the quality and quantity of outside options and the reservation utility of remaining in place. However, the panel structure of our data allows us to defend a weaker identification assumption; namely, we do not require that outcomes for successful and unsuccessful applicants are comparable in level terms, but only that the relative difference in mean outcomes across treatment groups would have remained constant in the absence of treatment ("parallel trends"). 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 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:

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

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 parallel trends assumption is most plausible when identifying variation comes from comparisons of applicants who are observably similar. Conditioning on human capital ($H_i$) and housing preferences ($\phi_i$) helps ensure that treated and untreated applicants face comparable offer distributions and continuation values, such that residual variation in treatment status arises from idiosyncratic timing in offer arrival and acceptance.

### 4.3 Estimation

We apply several restrictions to the applicant and tenant sample prior to estimating effects on mobility and labor market outcomes. Specifically, we exclude:

- applicants who are missing application date or have an invalid application date;
- applicants associated with multiple applications (whether at the same property or across properties)[20]
- all applicants to properties designated specifically for elderly residents;

[20]This restriction prevents contamination by individuals who are in the control group for one property but in the treatment group for another. In addition, multiple applications may signal unobservable differences that are possibly correlated with treatment.

- applicants whose terminal application status is "denied".[21]

We provide further details on sample construction for each set of outcomes in Sections 5.1 and 6.1, respectively.

After linking to outcomes, we implement our empirical tests in an applicant-by-event time panel centered on the LIHTC application date[22] and estimate treatment effects using two complementary approaches. First, we implement an OLS event study design that compares outcomes for placed and non-placed applicants after conditioning on a rich set of covariates. Second, we implement a "synthetic" event study design that re-weights the comparison group to closely match treated applicants' pre-application outcome paths. By focusing on comparisons across applicants who are most similar in observable pre-treatment dynamics, this procedure aims to isolate plausibly exogenous variation in treatment that covariate adjustment may not absorb.

**OLS TWFE Event Study:** To measure the effect of moving to LIHTC on neighborhood-level outcomes, we estimate the following two-way fixed effects (TWFE) regression using OLS:

$$Y_{it} = \alpha + \mu_i + \delta_t + \theta_{it(p)} + \sum_{t=-24}^{-2} u_t D_i + \sum_{t=0}^{24} u_t D_i + \mathbf{X}_{it}'\Pi + \upsilon_{it} \tag{1}$$

where $i$ indexes individual applicants and $t$ indexes months in event time.

$Y_{it}$ is a continuous neighborhood-level outcome variable. The variables $\mu_i$ and $\delta_t$ represent applicant and event time fixed effects, respectively, so the $u_t$ coefficients are identified by variation in outcomes within an applicant's address history over time. We also include $\theta_{it(p)}$, a fixed effect for the property $p$ that individual $i$ applied to interacted with event time, which further restricts identifying variation to comparisons of applicants within the same LIHTC property. The vector $\mathbf{X}_{it}'$ 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

[21] A denial signifies that an applicant was deemed ineligible for a unit, typically due to income verification, but also possibly due to issues with a background check or eligibility according to the property's specified target population. Including such applicants in the control group could mechanically lead to bias because all treatment group individuals will have been deemed eligible.

[22] For mobility analyses, we observe outcomes at a monthly frequency, so event time is measured relative to the month of application; for analyses using tax data, we observe outcomes at an annual frequency, so event time is measured relative to the year of application.

year of application interacted with event time. In augmented specifications, $\mathbf{X}'_{it}$ includes additional controls for observable characteristics that proxy for $\phi_i$, including bedroom preference and age at time of application, each separately interacted with event time fixed effects.[23] Because treatment assignment and vacancy timing operate at the property level, applicants competing for units within the same property may face correlated shocks. Accordingly, we cluster standard errors at the property level (Abadie et al., 2023).

To estimate the effects of moving into LIHTC housing on labor market outcomes, we implement a modified OLS two-way fixed effects (TWFE) event study design that leverages the larger sample available through higher match rates in the tax data. Specifically, we use a two-step procedure that seeks to match each successful applicant to a comparable unsuccessful applicant. We form pairs between treated and control units using a logit model that predicts the probability of treatment using baseline characteristics:

$$Treat_i = f(\beta_x X_i) \tag{2}$$

where $f(\cdot)$ represents the logistic function. All covariates included in $X_i$ are measured as of the year prior to the application year. These covariates include the applicant's wages, approximate HUD income, the count of unique W-2 forms received per filer in the tax unit, indicators for a set of age bins,[24] the year and month of application, stated preference for number of bedrooms in the LIHTC unit, and a property-level fixed effect.[25] We depict predicted propensities for both treatment and control groups in Figure 4.

After estimating predicted treatment probabilities, we construct matches between treatment and control units. We require that each matched pair apply to a LIHTC unit from the same origin zip code, which we argue eliminates unobservable differences between treated and control units that are a function of $H_i$ and $\phi_i$ and may be correlated with treatment status. To ensure that the matched units are sufficiently similar, we further require that the absolute value of the difference in predicted treatment probability between treatment and control units not exceed the 90th percentile

[23] 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.

[24] There are four bins: [19,25), [25,35), [35,55), [56,84].

[25] Note that the predicted probability of treatment is undefined when some combination of covariates perfectly predicts treatment status. This removes approximately half of applicants from our sample.

observed in the data.[26] We note that while each member of the treatment group is assigned to exactly one control unit, it is possible for a control unit to be assigned to multiple treatment units. As a result, the effective sample size is larger than the number of unique applicants due to repeated observations of certain control units.

We then estimate the treatment effect of LIHTC residency on income and wages using the following variant of Equation (1):

$$Y_{it} = \alpha + \mu_i + \delta_t + \lambda_{it(g)} + \sum_{t=-5}^{-2} u_t D_i + \sum_{t=0}^{3} u_t D_i + \upsilon_{it} \tag{3}$$

In addition to previously defined variables, $\lambda_{i(g)}$ represents a fixed effect for the pair group $g$ to which applicant $i$ belongs. As such, identifying variation in this specification comes only from within-pair variation. We cluster standard errors at the property level.

**Synthetic Event Study** We complement the OLS TWFE event study design with a synthetic difference-in-differences design (SDID) that assigns weights to non-residents such that their weighted pre-application outcome path closely matches that of residents, and places greater emphasis on the pre-periods most informative about post-treatment trends (Arkhangelsky et al., 2021). SDID strengthens credibility when the donor pool can replicate the treated group's pre-application dynamics with diffuse weights, but bias may arise if pre-treatment fit is poor (counterfactual error) or if the number of pre-periods is limited (overfitting). For inference, we estimate bootstrapped standard errors.

# 5 Effects on Mobility

## 5.1 Sample Construction

We link the common sample defined in Section 4.3 to the Infutor database using personal identifiers. We retain only high quality matches, defined as those for which:

- an applicant's last name, year of birth, and application address exact match[27] to an Infutor

[26] This value is equal to 0.37.

[27] To avoid complications associated with matching address strings across sources, we geocode addresses from Infutor and from the applicant data and define an "exact match" on address according to 3-digit latitude and longitude coordinates.

record; or

- an applicant's first name, last name, and application address exact match to an Infutor record.

Given the requirements of our matching algorithm, we drop applicants with missing data on required application identifiers including name, date of birth, and current address from the mobility sample. We additionally drop applicants whose application address is a PO box, since these addresses cannot be geocoded and do not necessarily correspond to the neighborhood in which an applicant resides. Overall, we successfully match about 60% of our applicant sample to Infutor address records, consistent with match rates from other studies that use Infutor to link to external, individual-level data (Collinson et al., 2024).

For 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 address start dates, following Phillips (2020).[28] We attribute each geocoded address to 2010-definition Census tracts and merge in tract-level attributes.[29]

We apply several final sample restrictions on the linked applicant-Infutor mobility sample. We exclude applicants with multiple addresses in their Infutor history that share the same effective date, as these records cannot be unambiguously sequenced. 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 disambiguate these cases using other available data.[30] Finally, we construct a balanced panel by restricting to applicants for whom we observe residential Census tracts for at least two years before and after the application date.[31]

[28] 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.

[29] We use 2010 Census tract boundaries. 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.

[30] 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, the same effective date may be imputed 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.

[31] Note that this last restriction drops applicants with a PO box in their Infutor address history in the two years immediately before and after application.

Our final analysis sample contains approximately 10,800 applicants, of which 36.5% are eventual residents who constitute the treatment group; i.e., "successful applicants" who move to a unit in a LIHTC-funded property. Individuals who ultimately canceled their application (38.5%) and individuals who remain on a waitlist at last observation (24%) make up the "control" group of unsuccessful applicants.[32]

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, 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:

$$y_i = \beta_0 + \beta_1 Resident_i + \theta_{it(p)} + \mathbf{X}'_{it} + \varepsilon_i \tag{4}$$

where the dependent variable $y_i$ is an observable characteristic at baseline, $Resident_i$ is a treatment indicator equal to 1 if the applicant is a placed resident, $\theta_{it(p)}$ is a property-level fixed effect, and $\mathbf{X}'_{it}$ includes additional fixed effects for calendar year interacted with calendar month of application. Identifying variation for $\beta_1$ therefore comes from comparisons of applicants who apply to the same property, after accounting for common seasonality trends.[33] Column (4) reports the $\beta_1$ coefficient for each baseline characteristic $X_i$, which is interpretable as the "regression-adjusted" difference in means (resident minus non-resident) between groups.

The results indicate that there is considerable selection into treatment based on applicant characteristics. Relative to non-residents, eventual residents tend to be older and more likely to apply from an address that is in the same zip code or tract as the LIHTC property address. In all cases, the difference in raw means attenuates substantially after including property-level and calendar month-year fixed effects. This pattern suggests that most selection into treatment operates at the application stage – specifically through decisions about where and when to apply – consistent with the conceptual framework outlined in Section 4.1. Notably, more than half of successful applicants apply from the same zip code where the LIHTC-funded property is located,

[32] We also observe a small number of applicants whose application is "approved" but who have not yet moved into a unit (0.6%); we assign these individuals to the control group, but our results are not sensitive to dropping them from the sample.

[33] Note that when comparing property tract characteristics, we do not include a property-level fixed effect as that would absorb all identifying variation.

and more than a quarter apply from the same tract. This is consistent with existing evidence that most LIHTC movers are local (Derby, 2021).[34]

We also find that residents have significantly higher quoted rents than non-residents, and this difference persists after controlling for property and calendar time fixed effects. Because stated unit size preferences (as measured by number of bedrooms) are similar across groups, this pattern is unlikely to reflect systematic differences in desired unit types. Instead, the difference may reflect greater competition for more heavily subsidized units within a given property, or differences in baseline income or rental assistance receipt across groups. In our main event-study specifications, we include applicant-level fixed effects to account for potential imbalances between treatment and control groups arising from time-invariant differences of this nature.

Moving beyond application characteristics, we find that eventual residents and non-residents apply from similar neighborhoods across a range of tract-level measures. After regression adjustment, differences are small in magnitude and typically statistically insignificant, particularly for our three primary neighborhood quality outcomes: poverty rate, income mobility, and median rent. Although similarity in outcome levels at the time of application is not required for identification in our event study framework, these results provide reassurance that the baseline fixed effects mitigate several key threats to the credibility of our treatment-control comparisons. Similarly, we do not detect any significant differences between residents and non-residents in terms of the characteristics of the LIHTC property tract associated with the application.

### 5.2 Descriptive Results

Before turning to main outcomes, we first analyze move activity to assess the reliability of the Infutor data for tracking month-level move timing and to understand how migration outcomes evolve before and after treatment for residents and non-residents. Figure 5 plots monthly move rates (i.e., moves per 100 applicants) separately for residents (solid line) and for non-residents (dashed line), with $t$ = -1 signifying the month immediately before application. Move activity is virtually identical across groups prior to the month of application, suggesting that non-residents

[34] 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.

provide a compelling counterfactual for residents in terms of baseline migration propensity in the absence of treatment. At $t = 0$, migration activity sharply increases among residents and further increases in $t = 1$. This result suggests that despite known measurement issues in the Infutor data, it nonetheless appears that Infutor's effective dates are reliable proxies for move timing at monthly frequencies. Move activity quickly decays but remains elevated in the treatment group for up to 7 months after the month of application, which likely reflects the fact that individuals may successfully obtain units from wait lists over time. Non-residents also exhibit an increase in move activity around the time of application. However, in the months immediately following application, their move rate is less than half that of residents. This pattern suggests that, conditional on not securing a unit, non-residents are substantially more likely to remain at their application address, consistent with predictions from our conceptual framework in Section 4.1. Over longer horizons, however, many non-residents ultimately relocate and experience lower overall housing stability than residents during the post-application period.

We next examine descriptive evidence on outcomes by plotting group-specific mean trends over time for each of the main neighborhood quality measures in Figure 6.[35] Prior to application, outcome levels differ modestly between residents and non-residents, consistent with the raw differences reported in Table 2. In the post-application period, these differences remain largely stable. There is some indication that non-residents gradually move to lower-poverty tracts over time, while residents do not. However, the magnitude of these changes is small relative to baseline levels. Overall, the raw trends suggest that individuals who move into LIHTC units experience neighborhood quality similar to that of counterfactual unsuccessful applicants, on average.

### 5.3 Event Study Results

We separately estimate Equation 1 for each of our three primary neighborhood quality outcomes and present the results in Figure 7.[36] For each outcome, we report both OLS two-way fixed effects estimates (panels (a), (c), and (e)) and SDID estimates (panels (b), (d), and (f)), plotting point estimates and 95% confidence intervals.[37] Coefficients are interpretable as regression-adjusted

[35] We also present time series by treatment group for supplementary tract-level outcomes in Figure A1.

[36] Event study results for supplementary tract-level outcomes appear in Figures A2 and A3.

[37] We note that the available sample size for estimating effects on median rents is reduced by about 50%. This is because we attribute median rents at the tract-bedroom level, and tract-bedroom cells are often suppressed in less populous tracts in the ACS 5-year estimates.

differences in mean outcomes between eventual residents and unsuccessful applicants at event time $t$.

The key identifying assumption is that outcomes for residents and non-residents would have evolved similarly absent treatment. As an indirect test of this assumption, we assess trends in pre-application outcomes across groups. We do not find evidence of differential pre-trends across groups in the pre-application period, which supports a causal interpretation of the research design.

Comparing specifications, the SDID estimates are more precise than the OLS TWFE estimates in the pre-period. This difference arises mechanically from the SDID procedure, which uses pre-period data to construct optimal unit and time weights that improve pre-treatment fit.

One potential concern with SDID specifications is that identification could rely disproportionately on a small number of heavily weighted observations. To assess this possibility, we examine the distribution of applicant-level SDID weights across bootstrap repetitions. We find that the sum of squared applicant-level weights is close to the total number of unique applicants in the sample, indicating that weights are approximately uniform and that no small subset of applicants drives the results.[38]

Turning to the post-application period, we find little evidence that moving into a LIHTC-funded unit leads to economically meaningful changes in neighborhood quality, on average. Point estimates across specifications and outcomes are generally comparable, small, and statistically indistinguishable from zero. However, our post-period estimates are sufficiently precise to rule out even relatively modest effects on neighborhood quality. For example, one year on from application, we can bound changes in the tract-level poverty rate from a lower bound of 0pp to an upper bound of +0.7pp, which corresponds to about a 4% increase over the base period mean. Estimated effect sizes for income mobility and median rents are similarly small. To the extent that post-application estimates differ from zero, they do not present a consistent pattern of neighborhood change. For example, estimates for tract-level poverty rates suggest that residents move to slightly lower-quality neighborhoods, whereas estimates for median rents suggest moves toward slightly higher-quality neighborhoods.

[38] We also examine the distribution of time-period weights in the SDID procedure. We find that identifying variation in the pre-period is concentrated in the event month immediately preceding application ($t-1$). When we re-estimate the model excluding $t-1$, weights shift primarily to $t-2$, with little change in substantive results. This pattern reflects the high degree of within-applicant persistence in neighborhood outcomes prior to application; i.e., outcomes change only when individuals move, and move activity is relatively low throughout the pre-period.

We next collapse the panel into single pre- and post-periods and summarize ATT effects for each outcome in Table 5. In addition to SDID estimates, we report OLS TWFE specifications both with baseline controls and with an augmented set of application controls that includes applicants' age at time of application and bedroom preference, each interacted with event month. Although difference-in-difference estimates are more precise than the event study estimates, we do not detect statistically significant effects on neighborhood quality. For example, SDID estimates imply that we can rule out increases in poverty rates exceeding +0.54pp (3% of the baseline mean), increases in income mobility rates exceeding +0.18pp (2% of the baseline mean), and increases in median rents exceeding $21 per month (2.5% of the baseline mean). Within each outcome, results are stable across specifications, suggesting that our findings are not sensitive to particular modeling choices or estimator selection.

We apply varying sets of additional sample restrictions to probe the robustness of the main results:

- We exclude applicants to "unmonitored" properties. Over time, LIHTC properties may age out of formal affordability compliance monitoring (see Section 2.1). In principle, properties that are no longer subject to monitoring could operate as market-rate developments.[39]

- We exclude applicants to acquired or rehabilitated properties, as well as to properties that become inactive during the sample period. Restricting the sample to newly constructed developments ensures that properties maintain stable management and affordability characteristics throughout the measurement period. In some cases, properties initially managed by our partner developer later become "inactive" in the client management database when ownership or management is transferred to another entity. Excluding such properties avoids potential inconsistencies in treatment exposure and record continuity.

- We exclude properties that cannot be linked to the HUD Placed-in-Service (PIS) database. Unmatched properties likely reflect incomplete coverage in the HUD data. However, since subsequent analyses examine heterogeneity according to variables observed only in the HUD

[39] In practice, such properties comprise a very small share of our sample. Moreover, the developer that provided the applicant and tenant data for this study has indicated that it continues to adhere to affordability restrictions beyond the minimum required compliance period.

data, restricting the sample to matched properties ensures that those results are not confounded by changes in sample composition.

- We exclude applicants with application dates in 2020 or later to assess whether our results are sensitive to the substantial supply- and demand-side shocks that affected rental markets during the COVID-19 pandemic.

We re-estimate ATTs for each outcome after applying each additional sample restriction separately. Figure 8 reports SDID difference-in-differences estimates and 95% confidence intervals for each restricted sample alongside the baseline estimates from the main analysis sample. Overall, the results are highly stable across samples. The similarity of estimates across alternative restrictions suggests that the main results are not sensitive to sample composition or data linkage decisions, reinforcing the conclusion that average neighborhood quality is largely unchanged by LIHTC placement.

Our analysis thus far has focused on average effects of LIHTC tenancy on neighborhood-level outcomes. However, average effects may mask substantial heterogeneity in the distribution of moves to LIHTC-subsidized units. To better characterize this variation, we examine the full distribution of individual-level changes in neighborhood outcomes for residents in the main analysis sample. Figure 9 presents binned scatterplots of origin (application) tract characteristics against destination (LIHTC property) tract characteristics.[40] In each plot, the x-axis corresponds to the origin tract and the y-axis to the destination tract; distance from the 45-degree line reflects the magnitude of the change associated with the move.

For tract-level poverty rates (panel (a)), there is substantial dispersion around the 45-degree line, indicating considerable heterogeneity in move-level changes. The distribution is also somewhat asymmetric: a large mass of moves are to slightly higher-poverty tracts (above the 45-degree line), though a non-trivial share of moves are to substantially lower-poverty tracts. This pattern suggests that LIHTC placement generates both upward and downward neighborhood transitions, even when average effects are close to zero. For income mobility (panel (b)) and median rents (panel (c)), the distributions are more symmetric and exhibit less dispersion around the 45-degree line, which implies that most residents tend to end up in tracts that are fairly similar to their application tracts

[40] We use Stata's *binscatter* program to non-parametrically group individual-level observations into bins of equal size.

along these dimensions.

### 5.4 Heterogeneity Analyses

Motivated by the substantial move-level variation that underlies average treatment effects, we conclude the mobility analysis by examining heterogeneous effects across a range of property-level characteristics.[41]

We first assess whether treatment effects on neighborhood quality differ depending on whether a property is located in a Qualified Census Tract (QCT). Prior work shows that the QCT basis boost strongly influences the spatial distribution of LIHTC developments; historically, QCTs have received roughly twice as many subsidized units, on average, as non-QCT tracts (Baum-Snow and Marion, 2009). Because QCTs are, by definition, the most economically disadvantaged neighborhoods within a given metropolitan area, critics argue that the incentive may further concentrate low-income renters in low-opportunity areas (e.g., Galster (2019); Kawitzky et al. (2013)).

We re-estimate SDID event studies separately for applicants to properties located in QCT and non-QCT tracts and present the results in Figure 10. Panel (a) shows that individuals who move to LIHTC properties in QCTs experience increases in tract-level poverty rates of approximately 2pp in the years following application – roughly a 10% increase relative to the baseline mean for this sub-sample. In contrast, moving to LIHTC properties in non-QCT tracts is associated with null effects, or at most modest declines, in neighborhood poverty (panel (b)). These results provide strong evidence that QCTs play a role in driving applicants to high-poverty tracts. By contrast, QCT status appears to play a more limited role when neighborhood quality is measured in terms of income mobility rates or median rents.

The observed effects of moving to a QCT may correspond to substantial impacts on the welfare of households residing in LIHTC units. Using data from the Moving to Opportunity (MTO) experiment, Ludwig et al. (2012) find that a 13pp reduction in neighborhood poverty increased adult subjective well-being by an amount comparable to the difference associated with a $13,000 increase in annual income (in 2009 dollars). Although such effects may not scale linearly and the MTO intervention differs substantially from the LIHTC context, these estimates suggest that a

[41] We also analyze heterogeneity according to application characteristics, including bedroom preference and age at time of application, in Appendix Figures A4 and A5, respectively.

2pp increase in neighborhood poverty may nonetheless have meaningful effects for adult tenants. Similarly, Chetty et al. (2016) find that children who moved to lower-poverty neighborhoods before age 13 earned $3,477 more annually as early-career adults. While the poverty differences in MTO were considerably larger than those observed here, the direction of these estimates implies that moves into higher-poverty QCT neighborhoods could have also have modest negative long-run effects for children.

To assess the role of supply-side competitive forces in shaping neighborhood impacts, we also examine heterogeneity by LIHTC credit type (Figure 11) and by urbanicity (Figure 12).

Approximately half of the properties in our main analysis sample are financed exclusively with 70% credits, while the remainder receive a combination of 70% and 30% credits. Because 70% credits are allocated through competitive state-level Qualified Allocation Plans (QAPs), whereas 30% credits are typically awarded non-competitively with tax-exempt bond financing, projects funded exclusively with 70% credits are more directly exposed to the incentive structures embedded in QAP scoring systems. Although QAP criteria are multidimensional and vary across states and over time, splitting the sample by credit structure allows us to assess whether more competitively allocated projects are systematically associated with opportunity-improving moves (e.g., if QAPs reward amenity-rich or higher-opportunity locations) or opportunity-reducing moves (e.g., if cost containment pressures or local political constraints steer development toward lower-cost neighborhoods).

Panel (a) of Figure 11 shows that individuals who move to properties funded exclusively with 70% credits experience modest increases in tract-level poverty rates (approximately +0.75 percentage points). While we do not directly observe the mechanisms underlying this heterogeneity, the direction of the estimates is consistent with competitive allocation processes operating alongside cost containment considerations or local land-use constraints that may limit development in higher-opportunity neighborhoods. By contrast, effects on tract-level income mobility rates and median rents are more mixed and generally small in magnitude.

The properties in our sample are primarily located in suburban and exurban/rural markets. These environments differ meaningfully in housing market structure, land costs, and local political constraints. Suburban jurisdictions typically feature higher baseline housing costs and more restrictive land use regimes, which may limit the set of feasible development sites for LIHTC projects.

In such settings, local opposition or zoning constraints may funnel subsidized developments toward lower-cost tracts within suburbs, potentially concentrating projects in lower-opportunity neighborhoods. By contrast, in exurban and rural areas, housing cost gradients tend to be flatter and land availability less constrained. In these markets, LIHTC projects may more closely resemble marginal market-rate developments, implying weaker sorting into particularly advantaged or disadvantaged tracts.

We observe clear heterogeneity in impacts on neighborhood poverty when separately analyzing properties in suburban (Figure 12, panel (a)) versus exurban / rural (Figure 12, panel (b)) markets. Individuals who move to suburban LIHTC properties experience significant increases in tract-level poverty rates (approximately +1.75pp). We do not detect meaningful changes in neighborhood poverty among applicants to properties located in exurban or rural areas. These patterns provide supporting evidence that neighborhood impacts are more pronounced in higher-cost, more constrained suburban markets.

Finally, we examine heterogeneity according to whether LIHTC properties are paired with project-based rental assistance (PBRA) contracts.[42] Properties with PBRA typically serve lower-income households, as rents are tied to tenant income and supplemented through federal rental subsidies. In contrast, LIHTC properties without project-based assistance rely primarily on the tax credit rent structure and may serve a somewhat higher-income segment of the eligible population. Because the presence of PBRA affects both the financial structure of the development and the likely composition of tenants, it may also influence where projects are sited and the types of neighborhoods to which applicants relocate. In the absence of applicant-level rental assistance data, splitting the sample by property-level PBRA status provides a useful proxy for differences in subsidy depth and tenant income targeting, allowing us to assess whether deeper rental assistance is associated with systematically different neighborhood exposures.

Figure 7 depicts a clear divergence in neighborhood poverty exposure by PBRA status. Among applicants to properties with project-based rental assistance (panel (a)), moving to LIHTC is associated with declines in tract-level poverty rates that increase over time (approximately -1pp after 2 years). Because most successful applicants move shortly after the month of application,

[42]The variable that indicates PBRA status is frequently missing in the HUD PIS database, so our sample size for this analysis is significantly reduced.

the post-period trend likely reflects continued moves among non-residents into higher-poverty neighborhoods over time. By contrast, applicants to properties without PBRA contracts (panel (b)) experience increases in neighborhood poverty that widens somewhat over time (approximately +0.5pp after 2 years). These results suggest that deeper rental subsidies may shape neighborhood access through two related channels. First, by reducing rent burdens for the lowest-income households, PBRA may enable tenants to access neighborhoods they would otherwise be unable to afford. Second, non-PBRA LIHTC properties may attract applicants who already reside in relatively higher-opportunity tracts and therefore exhibit less upward neighborhood mobility. As with other dimensions of heterogeneity, we do not detect corresponding differences in tract-level income mobility rates or median rents.

# 6 Effects on Labor Market Outcomes

## 6.1 Sample Construction

Given the differences in coverage, frequency, and richness between the Infutor data and tax data, we briefly describe differences in sample construction. To construct the sample for labor market outcomes, we begin by matching individual-level records from the LIHTC applicant data to IRS tax records available through the U.S. Department of the Treasury. We match individuals to their tax records by exact matching on Social Security Number.[43]

As tax data are available at the annual level instead of monthly, we have fewer observations per individual in the linked applicant-tax data sample than in the linked applicant-Infutor sample. To maximize statistical power, we restrict attention to applicants to LIHTC properties during the period with the highest application volume, 2012–2021. Given data availability at the time of writing, this choice limits the post-period to three years following the application year. We include a longer pre-period—beginning five years prior to the application year—to allow for a more comprehensive assessment of pre-treatment outcome trajectories.

We compile a balanced panel of outcomes for each individual, which ranges from five years prior to application to three years after application. In the instances where an applicant does not

[43] 71.7% of individuals in the LIHTC applicant sample have valid SSNs (must be non-missing and in an acceptable SSN format).

appear as a primary or secondary filer on a Form 1040, and does not receive a Form W-2, 1099-G, or 1099-INT, we impute that their wages and approximate HUD income are equal to zero.[44] This imputation allows us to make internally valid comparisons that are not confounded by changes in sample composition over time that may result from selection into measurement.[45]

The composition of our analysis sample differs across estimation strategies. As described in Section 4, we employ a nearest-neighbor matching procedure to eliminate unobservable differences between residents and non-residents that may be correlated with treatment status. Constructing matched pairs results in significant sample attrition; the matched-pair sample includes only 11,158 unique applicants. We note that since it is possible for a single control unit to be paired with multiple treatment units, our "effective" sample size exceeds the number of unique applicants.[46]

By contrast, our synthetic difference-in-difference strategy does not rely on constructing matched pairs. The sample size available in these specifications is therefore much larger, consisting of 37,166 unique applicants.[47] This "SDID" sample is a superset of the matched-pairs analysis sample.

We present summary statistics for the matched-pairs sample of applicants in Table 3 and for the SDID sample in Table 4. Columns (1) and (2) present unconditional means for non-residents and residents, respectively, and column (3) reports the raw difference in means. In the matched-pairs sample, we also test for differences at time of application using Equation 4, as in Section 5.1, and report the regression-adjusted differences in means in column (4). We measure each outcome in the year prior to application. In the SDID sample, we observe statistically significant and large differences between residents and non-residents along a set of application characteristics, including age and number of dependents. LIHTC residents are also significantly less likely to have filed a Form 1040 in the year prior to application, compared to unsuccessful LIHTC applicants. These results suggest that there is considerable selection into LIHTC tenancy, an issue that motivates our use of both the SDID estimation approach and the nearest-neighbor matching design. The

[44] When we observe an individual's W-2(s) but not 1040, we impute income to equal the sum of W-2 wages, unemployment compensation from Form 1099-G, and interest income from Form 1099-INT.

[45] While we observe W-2s for any individual who is listed on at least one employer's payroll, we only observe 1040s for individuals who file taxes in a given year, or who are claimed as a dependent of or co-file with a tax filer in the same household. While filing rates are generally high among those required to file ($> 90\%$), not all Americans are required to file. The most important exception in this context is the minimal income filing threshold, which exempts very low-income households from the tax filing requirement. In 2024, the threshold for household heads without children under age 65 was \$21,900. Previous work has shown that filing rates for households under the minimal income threshold can be substantially lower (Hauck and Wallossek, 2024).

[46] 72.5% of unique applicants in the matched-pair sample are treated.

[47] 58% of unique applicants in this sample are treated.

pair-matching procedure, combined with our fixed effects structure, successfully eliminates most observable differences between groups.

### 6.2 Descriptive Results

We plot unconditional mean outcomes for residents (treatment) and non-residents (control) by event time in Figure 14.[48] In the SDID sample, applicants who become residents earn less than applicants who do not become residents. In the year prior to application, control applicants out-earn treated applicants by roughly $5,000.

Although the raw data for the SDID sample reveal differential pre-trends, these imbalances are no longer present after implementing our matching procedure (Figure 15). Among this matched-pairs sample, pre-period outcomes are very similar across the two treatment groups, both in levels and in changes.

In each sample, we observe that mean income and wages increase for *both* residents and non-residents starting in the year of application ($t = 0$).[49] However, non-residents experience much greater growth in outcomes relative to LIHTC residents in the post-application period. Focusing on the matched-pair sample, average wages are nearly identical between the two groups in the year prior to application, but non-residents earn about $2,500 more on average during the year of application and about $3,200 more on average two years after application.

Why do income and wages rise after application for both groups? One possibility is that applicants experience a contemporaneous shock –such as a rent increase – that prompts both a search for more affordable housing and adjustments in labor supply. A joint housing and job search of this kind could lead to increases in earnings, particularly for successful applicants who receive only modest LIHTC rent discounts. In future work, we aim to further clarify the mechanisms underlying post-application income and wage growth.

[48] Both 1040 income and W-2 wages are measured in 2019 USD.

[49] Given the annual frequency of the tax data, outcomes measured in the year of application ($t = 0$) may capture dynamics both before and after the application date. We therefore interpret the application year as a partially-treated time unit.

### 6.3 Event Study Results

The descriptive results in the previous section indicate that applicants who place into LIHTC-subsidized units experience lower income and wage growth relative to comparable unsuccessful applicants, and that this divergence reflects slower earnings growth among residents rather than declines in earnings levels. In this section, we formally test for differences in outcomes using the matched-pair event studies outlined in Section 4. We supplement the matched-pair design with results from SDID event studies that benefit from a larger sample and employ a set of data-driven weights to achieve balanced pre-trends.

We first report results from the matched-pair event study specification. Figure 16, panels (a) and (b), presents the estimated effects of LIHTC residency on income and wages, respectively. In each case, we plot point estimates and 95% confidence intervals across event time, with the base year (the year prior to application) omitted. We do not find evidence of differential pre-application trends across treatment groups for either outcome. Starting in the year of application ($t = 0$), however, we detect negative and statistically significant effects of moving to LIHTC on both income and wages. During the application year, residents' wages decline by approximately \$1,900 relative to those of comparable non-residents, an 11.2% difference compared to the base period mean. By the second year after application, the difference in wages grows to roughly \$3,000 (18.1% relative to the base period mean). The time path and magnitude of treatment effects are very similar for both income and wages. Since HUD income is comprised of both wages and income from various other sources, the loss of labor market earnings by LIHTC tenants does not appear to be compensated by gains from other sources.

Figure 17 reports the SDID estimates. The patterns align closely with the matched-pair results, though the synthetic approach yields more precise estimates. In the year of application, residents' wages and incomes decline by roughly \$2,000, relative to those of non-residents. By the third year after application, the gap between residents' and non-residents' wage earnings increases further to approximately \$3,000. As in the matched-pair sample, the estimates of treatment effects on HUD income and wages are very similar in magnitude, suggesting that the drop in wage earnings is not accompanied by an increase in other earnings, such as self-employment income. The SDID estimates correspond to an overall ATT of LIHTC residency on wage income of –\$2,552 in the three

years after application. For our measure of HUD income, the estimated ATT is –$3,109.

### 6.4 Interpretation

The magnitudes of our estimated treatment effects are comparable to those found in other studies on the effects of housing subsidies on income and wages. For example, Susin (2005) finds that moving to public housing is associated with a 19% decline in earnings, while Jacob and Ludwig (2012) estimate that voucher receipt reduces quarterly earnings by 10%. However, as discussed in Section 2.1, LIHTC subsidies differ from public housing and voucher programs in that they do not mechanically reduce benefits as earnings rise and therefore do not directly distort net wage rates. In the absence of such substitution effects, we would expect the earnings impacts of LIHTC subsidies to be smaller than those associated with vouchers or public housing.

On the one hand, we cannot rule out that our larger estimates may be attributable in part to our selected sample of individuals, who either file tax returns and/or have labor market earnings for seven consecutive years around the time of application. Such individuals are likely to be more attached to the labor market and may thus have relatively high levels of – and variation in – earnings compared to individuals who are out of the labor force for multiple years. On the other hand, it is possible that our estimates do not necessarily represent pure income effects due to possible complementarities between LIHTC and housing voucher programs. As shown in Table 1, more than 20% of the properties in our sample have project-based rental assistance contracts, and it is likely that many residents also receive vouchers or other forms of tenant-based rental assistance. In practice, access to LIHTC-funded units may generate substitution effects for these individuals.

Given that the mobility analysis in Section 5.3 yields null average effects on neighborhood quality, it is unlikely that neighborhood changes account for the magnitude of the income and wage impacts we observe. However, the direction of the earnings effects is consistent with increased housing stability among LIHTC tenants, which may generate "lock-in" effects that dampen post-period earnings growth.

To further contextualize the labor market impacts, we estimate the size of the implicit housing subsidy received by LIHTC tenants by combining internal applicant-level quoted rents with ACS microdata on market rents. Because residents do not experience increases in average neighborhood rents after moving to LIHTC relative to non-residents (see Section 5.3), the implicit subsidy

represents the primary source of housing cost savings associated with LIHTC tenancy.

We define the implicit subsidy as:

$$s_{i,u,j} = rent^{mkt}_{u,j} - rent^{LIHTC}_{i}, \tag{5}$$

where $i$ indexes applicants, $u$ indexes unit type (bedroom count), and $j$ indexes Census tract. We impute market rents based on the applicant's stated bedroom preference and the tract of the LIHTC property, merged with tract-by-bedroom median rents from the 2015-2019 ACS 5-year estimates. We convert quoted rents – observed at the applicant level at the time of application – to 2017 USD for comparability.[50]

The average quoted rent for residents in our main sample is $625 per month, compared to an average imputed market rent of $755. This implies an average implicit subsidy of $130 per month, or approximately $1,560 annually. (We plot the full distribution of resident-level implicit subsidies in Figure A7). This annual subsidy corresponds to roughly 50% of the annual ATT on income and 61% of the annual ATT on wages. While meaningful, this magnitude suggests that income effects alone are unlikely to account for the full labor market impacts we estimate. Instead, we speculate that substitution effects from other rental assistance programs are likely to play an important role, along with possible "lock-in" effects of housing stability.

To benchmark our estimates, Cook et al. (2023) calculates implicit LIHTC subsidies in the Chicago metropolitan area as the difference between predicted market rents and regulated LIHTC rents. Because average rents in Chicago exceed those in our sample, we rescale Cook et al. (2023)'s estimates to match the rent distribution observed in our markets.[51] After adjustment, Cook's estimates imply an annual subsidy of approximately $1,500 – nearly identical to our estimate based on quoted rents and ACS data.

There are several important limitations of this exercise. First, quoted rents may differ from

[50] We use quoted rents and not unit-level LIHTC rent restrictions to define the value of the implicit subsidy because we observe that rent restrictions often do not bind in our sample. We plot the full distribution of differences between predicted market rents and rent restrictions in Figure A6. On average, LIHTC rent restrictions exceed predicted market rents by roughly $110 per month. This is consistent with Cook et al. (2023) who find that the average rent ceiling for LIHTC units in Chicago between 2016-2018 was $954, but the average rent that tenants report to the ACS in the same period was only $742. These findings suggest that below-market LIHTC rents are often generated by competition for 9% credits and not necessarily by binding affordability restrictions.

[51] This approach implicitly assumes that LIHTC subsidies scale proportionally with market rents. While strong, this assumption is broadly consistent with Cook et al. (2023)'s finding that subsidy magnitudes increase with neighborhood opportunity.

rents actually paid by residents after move-in. In particular, quoted rents may not reflect other forms of rental assistance; in such cases, quoted rents should be interpreted as an upper bound. Moreover, we do not observe how rents may evolve in the years after move in. In addition, ACS tract-by-bedroom estimates of median rents do not adjust for unit quality and therefore may over- or under-state the true market value of LIHTC units. ACS rents are also interpretable as "gross rents" – the total amount a tenant pays, including base rent plus additional operating costs like utilities. Although we interpret quoted rents as gross rents, our subsidy estimates will be biased upward if quoted rents do not consistently incorporate utility allowances. Finally, because we do not observe rent payments for non-residents, our approach may overstate relative housing cost savings if non-residents also receive rental assistance or reside in below-market units.

# 7 Conclusion

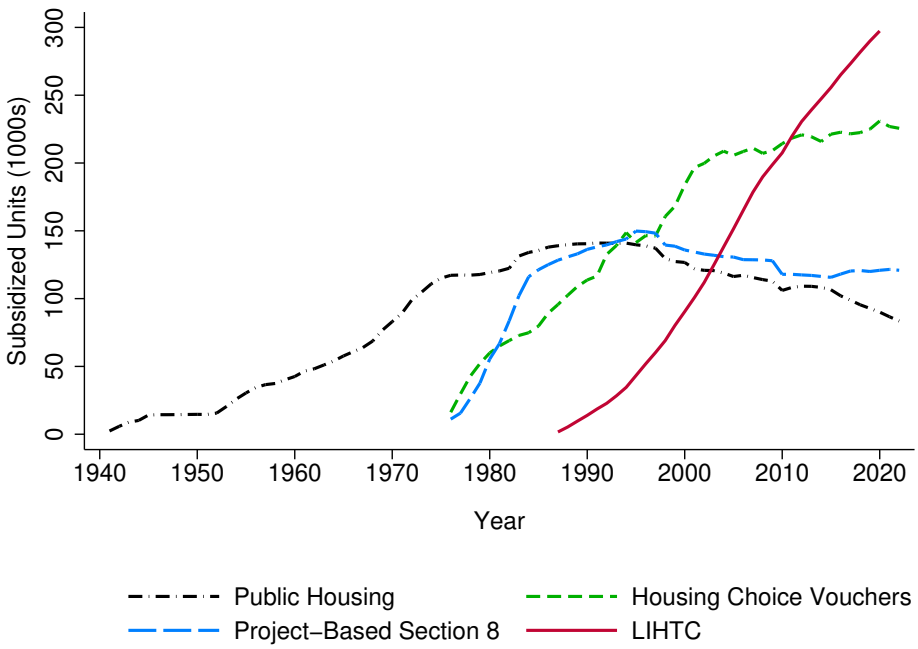
This paper provides causal evidence on how residence in LIHTC-subsidized housing shapes neighborhood trajectories and labor market outcomes. Although LIHTC developments are disproportionately located in lower-opportunity tracts, individuals who move into LIHTC properties do not experience significant changes in neighborhood quality relative to comparable applicants, on average. However, these average effects mask important policy-relevant heterogeneity. Individuals who move to LIHTC properties located in Qualified Census Tracts – economically disadvantaged neighborhoods where federal rules provide additional incentives for LIHTC development – experience significant increases in neighborhood poverty. Our results also suggest that local constraints and state-level policy incentives play an important role in shaping mobility outcomes for LIHTC tenants. In particular, moves to LIHTC properties in higher-cost suburban tracts are associated with increases in neighborhood poverty. This pattern is consistent with an interaction between competitive state credit allocation processes that reward lower-cost development and local land use constraints that limit the construction of affordable housing in high-amenity areas.

We also find that moving to a LIHTC-funded unit is associated with sizable negative effects on earnings. Within three years of application, LIHTC residents earn roughly $3,000 less in W2 wages per year than comparable non-residents – an 18% decline – with similar relative reductions across broader measures of income. These treatment effects reflect slower wage growth among LIHTC

tenants rather than absolute declines in earnings levels. We estimate that the implicit LIHTC rent subsidy averages approximately $1,560 per year, suggesting that income effects alone are unlikely to fully explain the magnitude of the observed earnings impacts. We therefore hypothesize that other channels – including interactions with other rental assistance programs and potential "lock-in" effects associated with increased housing stability – may play an important role.

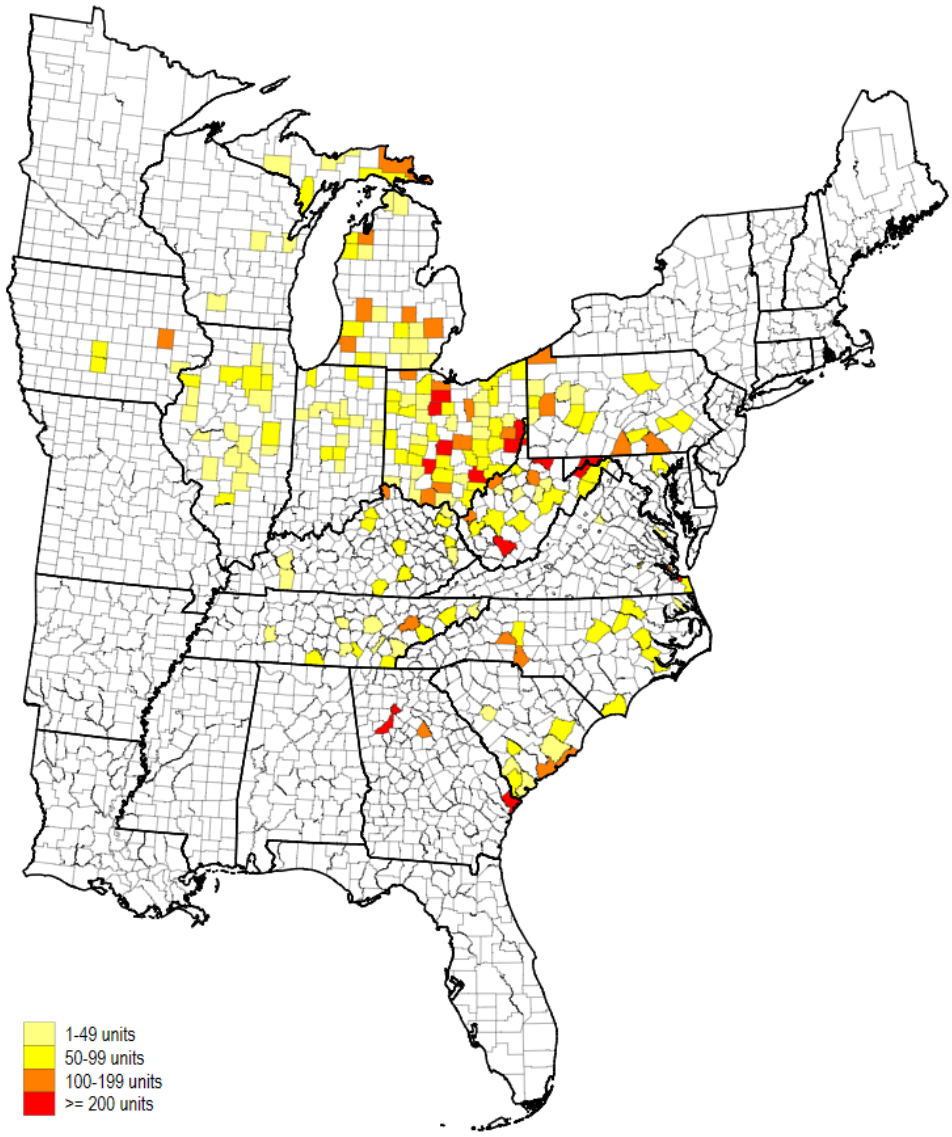
The LIHTC is currently the largest federal affordable housing program and costs more than $14 billion annually in forgone tax revenue. The program's primary objective is to increase the supply of affordable rental housing, and we find evidence that LIHTC tenants benefit from meaningful rent discounts relative to comparable market-rate units. However, by providing location-based below-market rents, the LIHTC also alters tenants' labor supply incentives and may steer some households toward lower-quality neighborhoods. Consequently, although individual renters may experience welfare gains by accepting the "bundle" offered by a LIHTC-subsidized unit, these gains must be weighed against the broader costs of the program. In addition to its direct cost, the LIHTC may generate longer-run inefficiencies through the spatial misallocation of workers. Our findings indicate that these outcomes are not inherent to policies that expand the supply of affordable housing, but instead arise from specific federal and state policy incentives that interact with local constraints on low-income development.

**Figure 1.** Subsidized Units per Year, by HUD Program



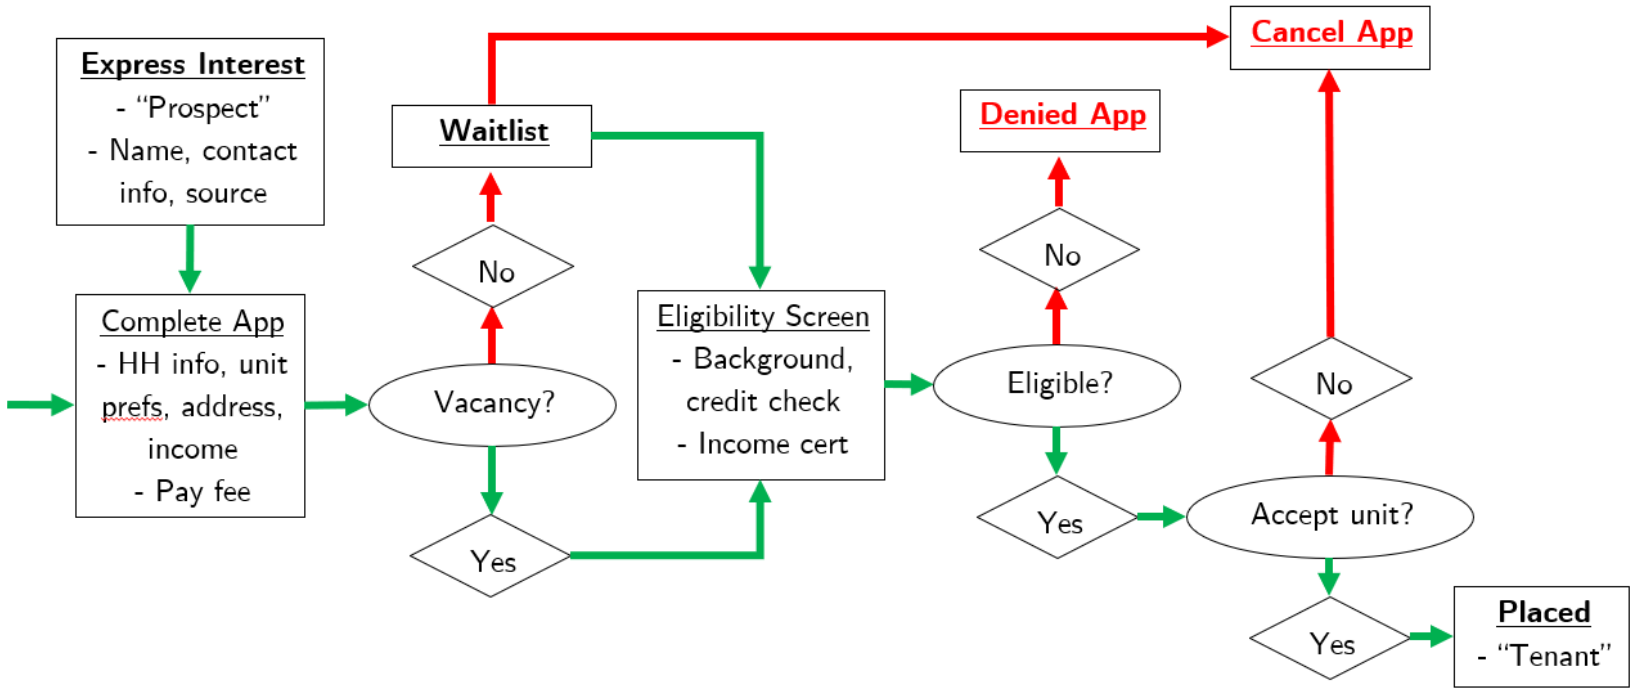

*Notes:* LIHTC series data are from HUD's database of all LIHTC units placed in service 1987-2021, and reflect all low-income units in LIHTC-subsidized properties. Data on public housing, housing choice voucher, and project-based Section 8 are from Olsen (2003) (pre-1998), HUD Annual Performance Reports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

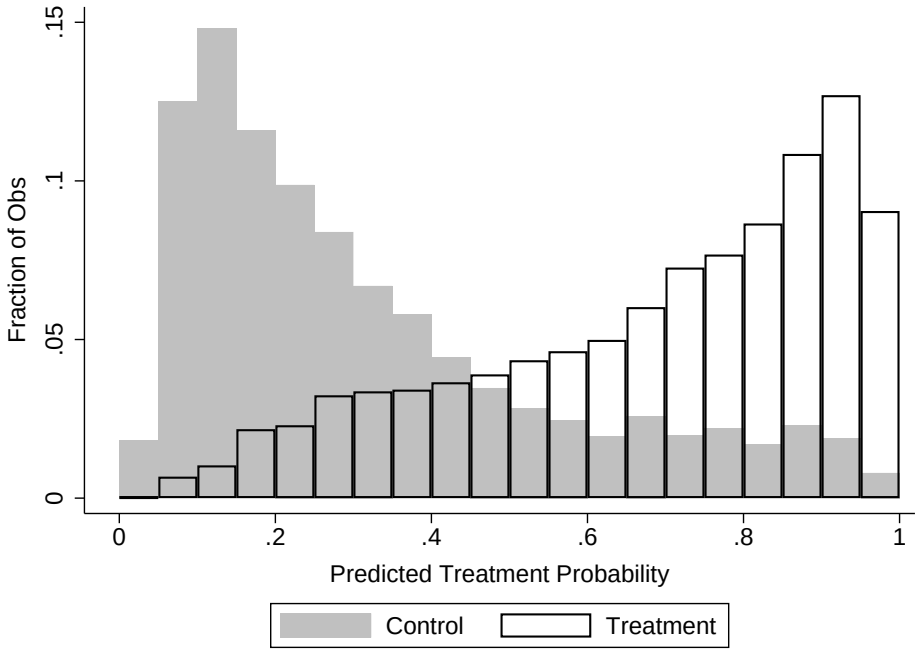



*Notes:* This chart describes the LIHTC application process for our sample. Prospective applicants ("prospects") initiate the process by contacting the property manager directly or through an online listing platform (e.g., Zillow, Apartments.com). At this stage, prospects submit basic contact information and wait for property staff to follow up with details on how to apply.[52] 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,[53] the application is forwarded for processing.

[52] Alternatively, an interested individual can simply submit a full application without first expressing interest.

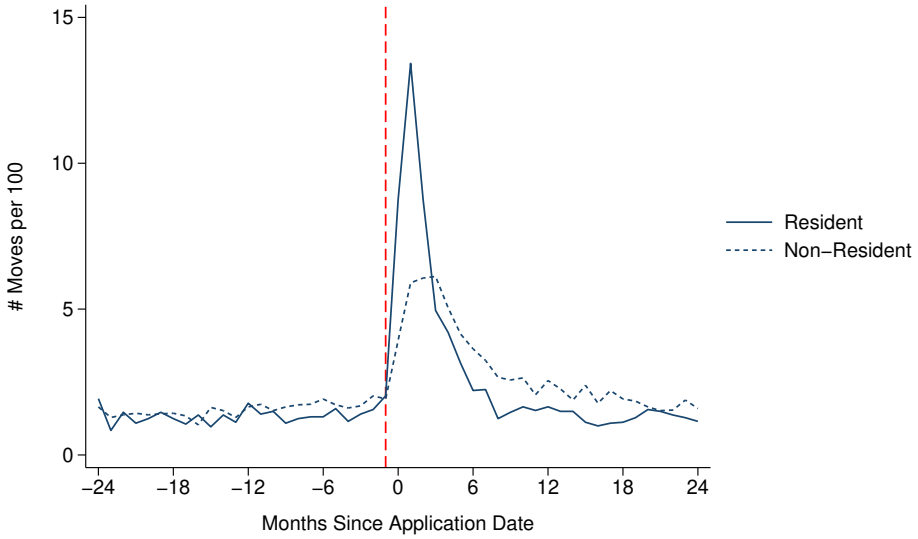
[53] Application fees at these properties are typically low (less than $25). They are used to cover administrative costs as well as costs associated with background checks used for eligibility screening.

**Figure 4.** Tax Data Sample:
Predicted Treatment Distribution by Group



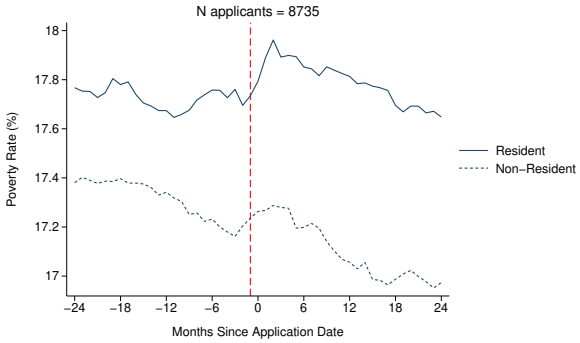

*Notes:* The sample contains applicants in the matched-pairs analysis sample as described in Section 6.1. The "Control" group includes applicants whose applications were canceled. The "Treatment" group refers to applicants who become LIHTC residents. We predict treatment status following Equation 2.

**Figure 5.** Measuring Move Activity in Infutor

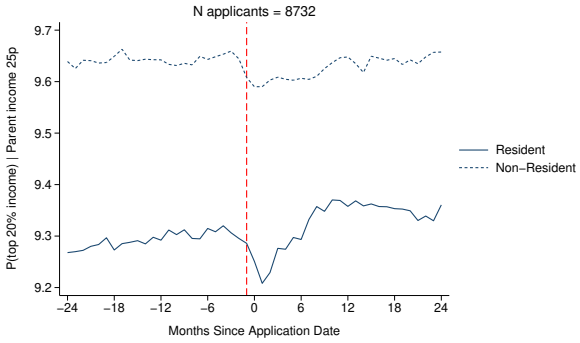



*Notes:* The sample contains applicants in the main analysis sample as described in Section 5.1. The month of application is defined to be event month $t = 0$.

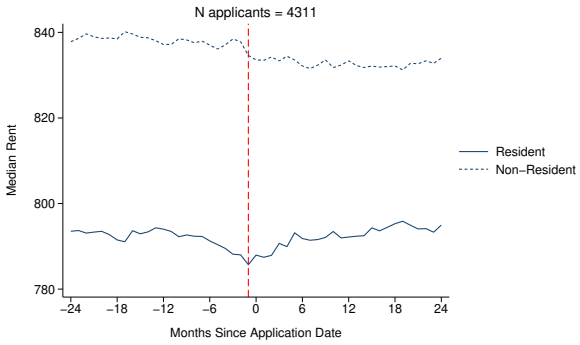
**Figure 6.** Mobility Outcomes:
Unconditional Means by Event Time and Treatment Group




**(a)** Poverty rate



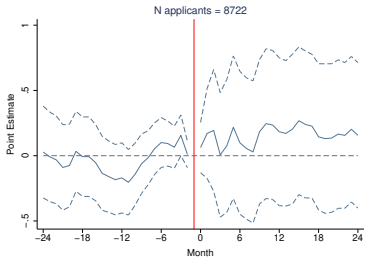

**(b)** P(top 20% income) | parent income 25p




**(c)** Median rents

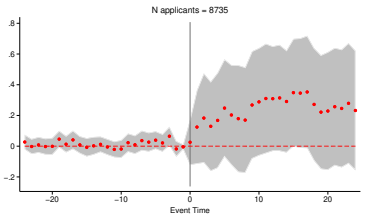
*Notes:* The sample contains applicants in the main analysis sample as described in Section 5.1. For each specified outcome, we further restrict to a balanced panel of applicants with non-missing outcome data in the 24 months before and after the month of application.

**Figure 7.** Mobility Outcomes:
Event Study Results



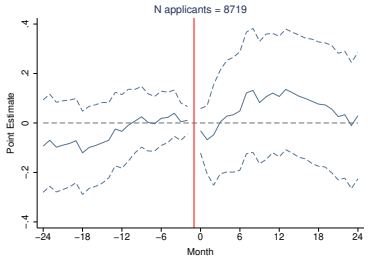

**(a)** Poverty rate

OLS TWFE



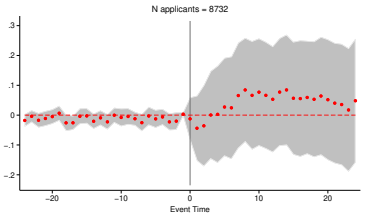

**(b)** Poverty rate

SDID



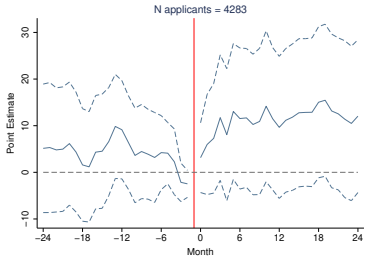

**(c)** P(top 20% income) | parent income 25p

OLS TWFE



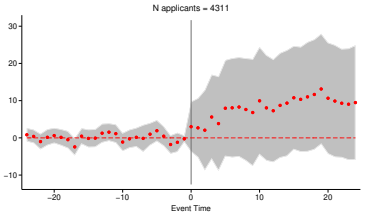

**(d)** P(top 20% income) | parent income 25p

SDID




**(e)** Median rents

OLS TWFE

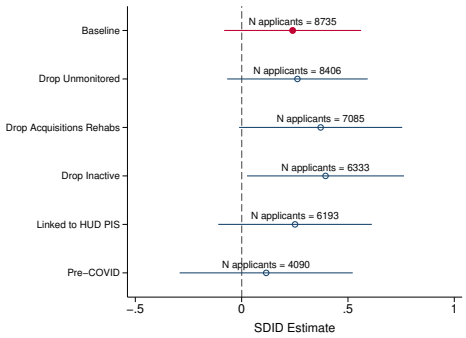



**(f)** Median rents

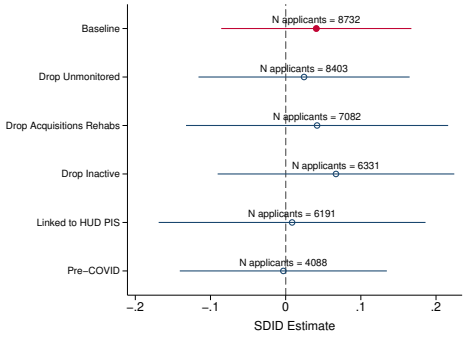
SDID

*Notes:* The sample contains applicants in the main analysis sample as described in Section 5.1. For each specified outcome, we further restrict to a balanced panel of applicants with non-missing outcome data in the 24 months before and after the month of 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 the continuous specification of Equation 1. We cluster standard errors at the property level in OLS TWFE specifications and estimate bootstrapped standard errors (100 reps) in SDID specifications.

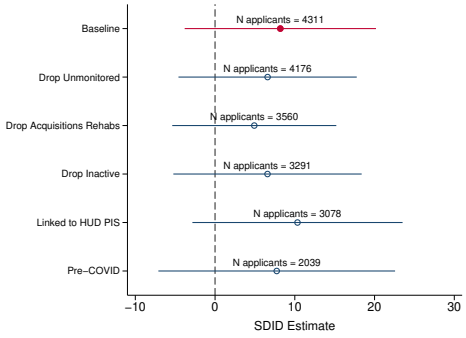
**Figure 8.** Mobility Outcomes:
Robustness of Difference-in-Differences Estimates




**(a)** Poverty rate



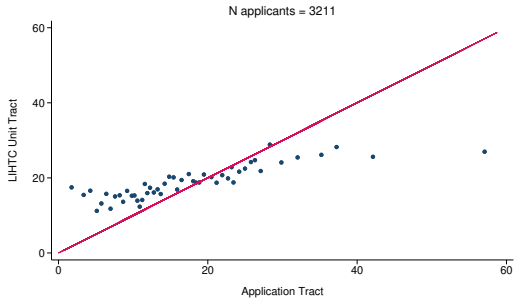

**(b)** P(top 20% income) | parent income 25p



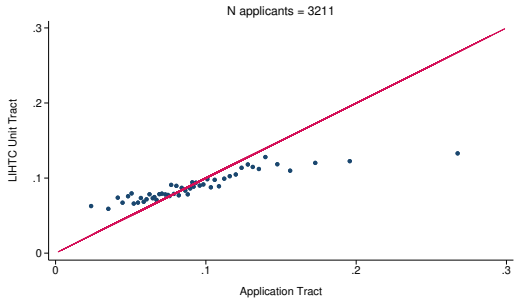

**(c)** Median rents

*Notes:* The sample contains applicants in the main analysis sample as described in Section 5.1. For each specified outcome, we further restrict to a balanced panel of applicants with non-missing outcome data in the 24 months before and after the month of application. Each specification differs from the baseline specification only in the specified manner. We depict ATT point estimates and 95% confidence intervals estimated using SDID regressions with bootstrapped SEs (100 reps).

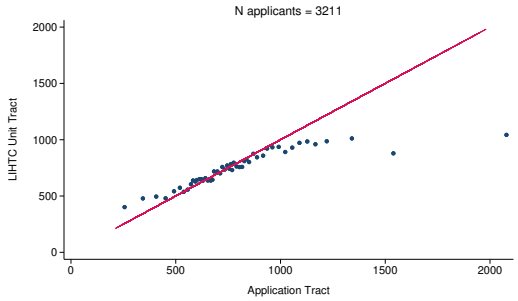
**Figure 9.** Mobility Outcomes:
Distribution of Individual-Level Moves to LIHTC Units




**(a)** Poverty rate



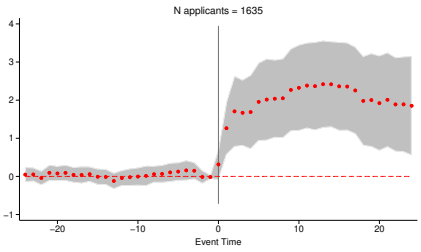

**(b)** P(top 20% income) | parent income 25p




**(c)** Median rents

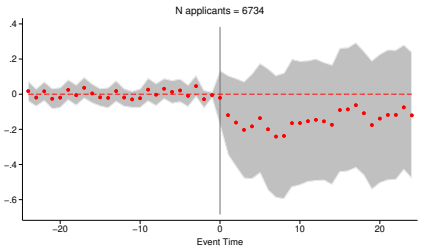
*Notes:* The sample restricts to residents in the main analysis sample as described in Section 5.1. We use Stata's *binscatter* program to non-parametrically group individual-level observations into bins of equal size.

**Figure 10.** Mobility Outcomes:
Heterogeneity by QCT Status



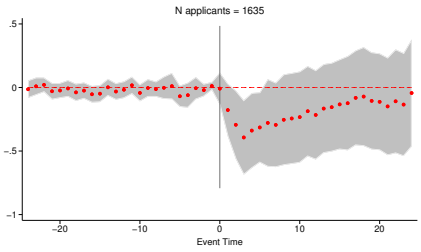

**(a)** Poverty Rate

QCT



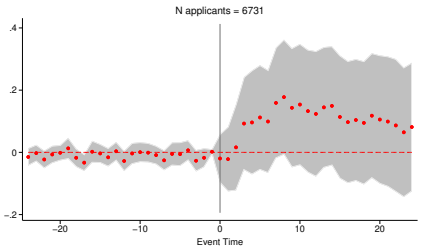

**(b)** Poverty Rate

Non QCT



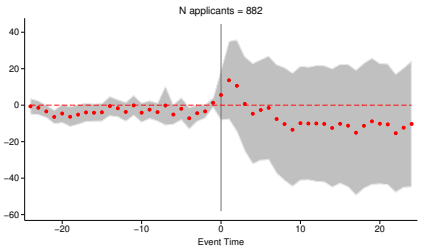

**(c)** P(top 20% inc) | parent inc 25p

QCT



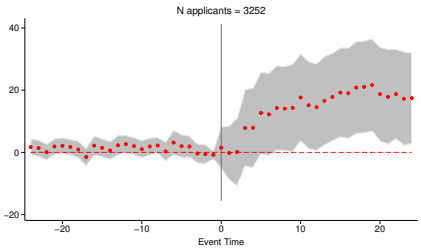

**(d)** P(top 20% inc) | parent inc 25p

Non QCT




**(e)** Median rents

QCT

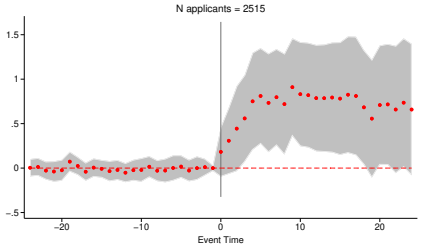



**(f)** Median rents

Non QCT

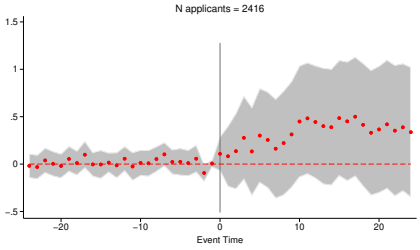
*Notes:* The sample contains applicants in the main analysis sample as described in Section 5.1. For each specified outcome, we further restrict to a balanced panel of applicants with non-missing outcome data in the 24 months before and after the month of application. Each event study plots point estimates and 95% confidence intervals from separate SDID regressions of the outcome on event month treatment indicators and the baseline set of controls, following the continuous specification of Equation 1 with bootstrapped SEs (100 reps).

**Figure 11.** Mobility Outcomes:
Heterogeneity by LIHTC Credit Type



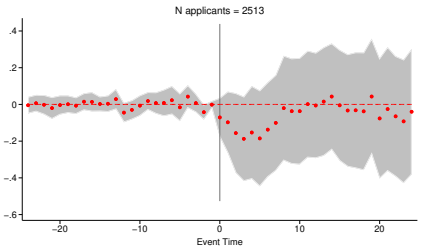

**(a)** Poverty Rate

70% only



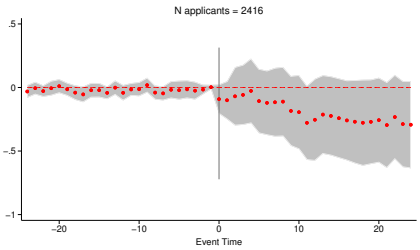

**(b)** Poverty Rate

Mix of 70% and 30%



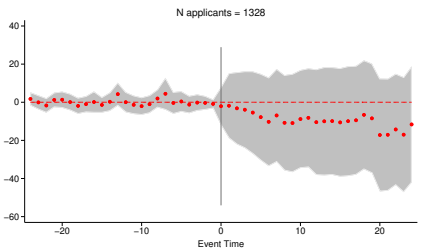

**(c)** P(top 20% inc) | parent inc 25p

70% only



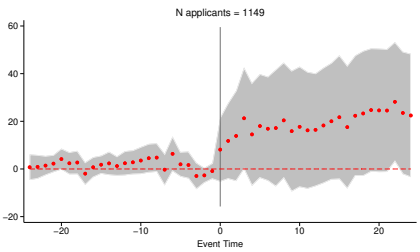

**(d)** P(top 20% inc) | parent inc 25p

Mix of 70% and 30%




**(e)** Median rents

70% only

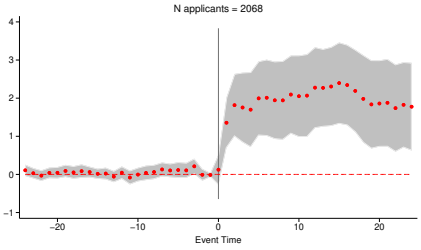



**(f)** Median rents

Mix of 70% and 30%

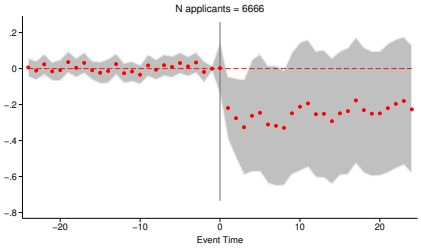
*Notes:* The sample contains applicants in the main analysis sample as described in Section 5.1. For each specified outcome, we further restrict to a balanced panel of applicants with non-missing outcome data in the 24 months before and after the month of application. Each event study plots point estimates and 95% confidence intervals from separate SDID regressions of the outcome on event month treatment indicators and the baseline set of controls, following the continuous specification of Equation 1 with bootstrapped SEs (100 reps).

**Figure 12.** Mobility Outcomes:
Heterogeneity by Urbanicity



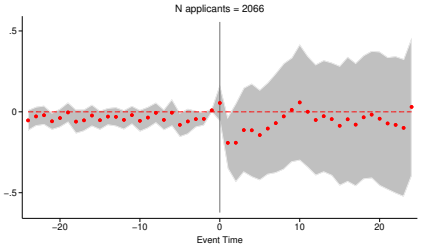

**(a)** Poverty Rate

Suburban



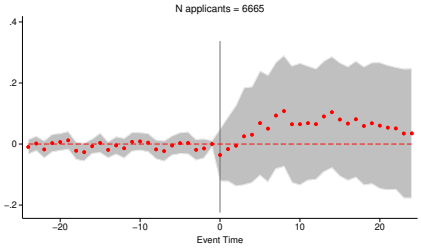

**(b)** Poverty Rate

Exurban / Rural



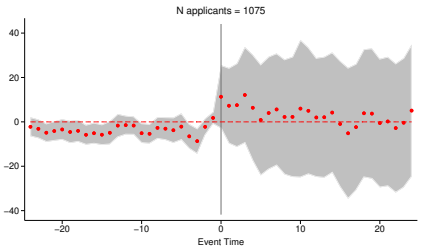

**(c)** P(top 20% inc) | parent inc 25p

Suburban



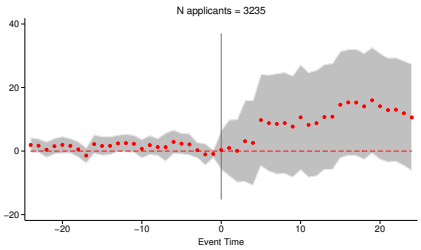

**(d)** P(top 20% inc) | parent inc 25p

Exurban / Rural




**(e)** Median rents

Suburban

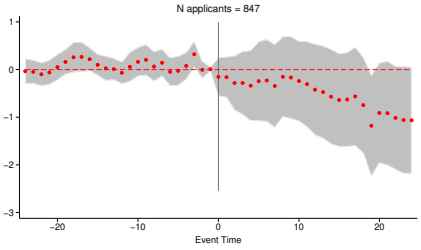



**(f)** Median rents

Exurban / Rural

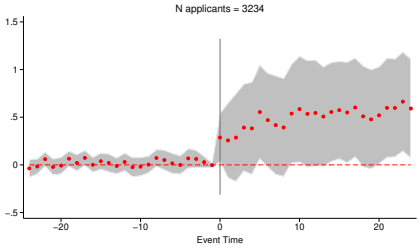
*Notes:* The sample contains applicants in the main analysis sample as described in Section 5.1. For each specified outcome, we further restrict to a balanced panel of applicants with non-missing outcome data in the 24 months before and after the month of application. Each event study plots point estimates and 95% confidence intervals from separate SDID regressions of the outcome on event month treatment indicators and the baseline set of controls, following the continuous specification of Equation 1 with bootstrapped SEs (100 reps).

**Figure 13.** Mobility Outcomes:
Heterogeneity by Project-Based Voucher Status



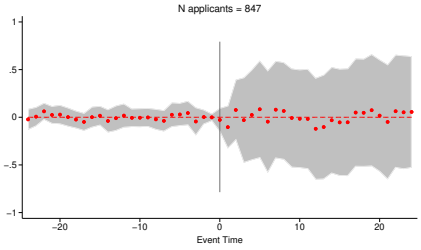

**(a)** Poverty Rate

PBRA



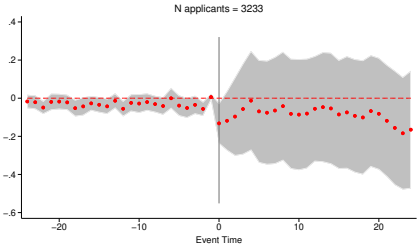

**(b)** Poverty Rate

Non PBRA



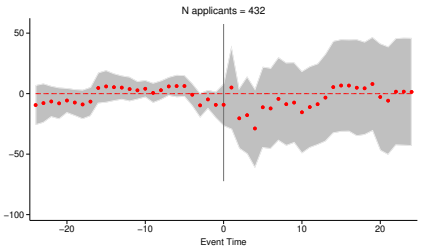

**(c)** P(top 20% inc) | parent inc 25p

PBRA



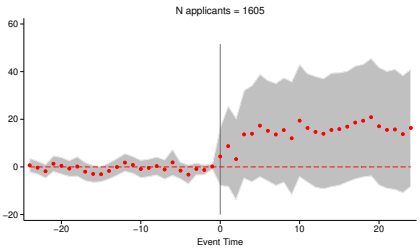

**(d)** P(top 20% inc) | parent inc 25p

Non PBRA




**(e)** Median rents

PBRA

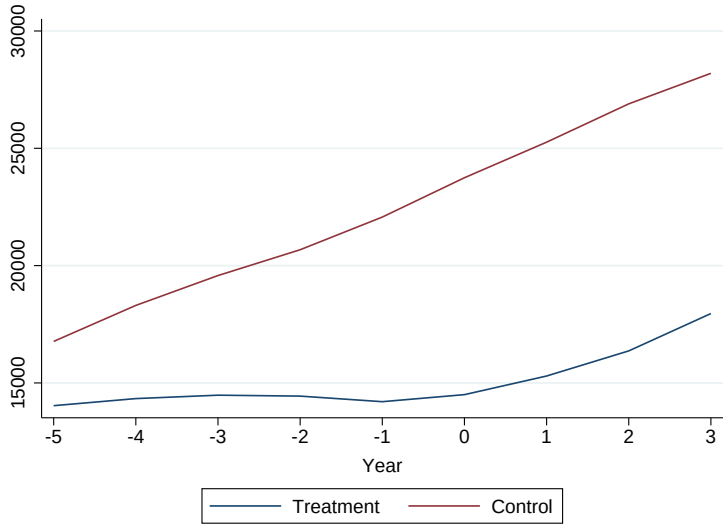



**(f)** Median rents

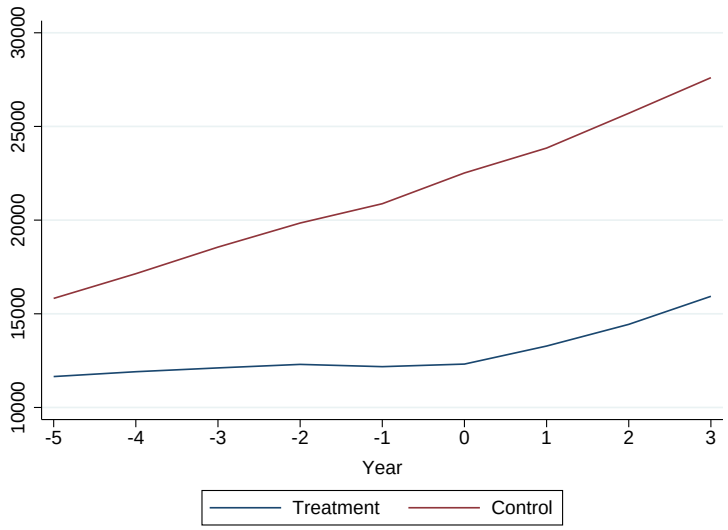
Non PBRA

*Notes:* The sample contains applicants in the main analysis sample as described in Section 5.1. For each specified outcome, we further restrict to a balanced panel of applicants with non-missing outcome data in the 24 months before and after the month of application. Each event study plots point estimates and 95% confidence intervals from separate SDID regressions of the outcome on event month treatment indicators and the baseline set of controls, following the continuous specification of Equation 1 with bootstrapped SEs (100 reps).

**Figure 14.** Labor Market Outcomes:
Unconditional Means by Event Time and Treatment Group,
SDID Sample



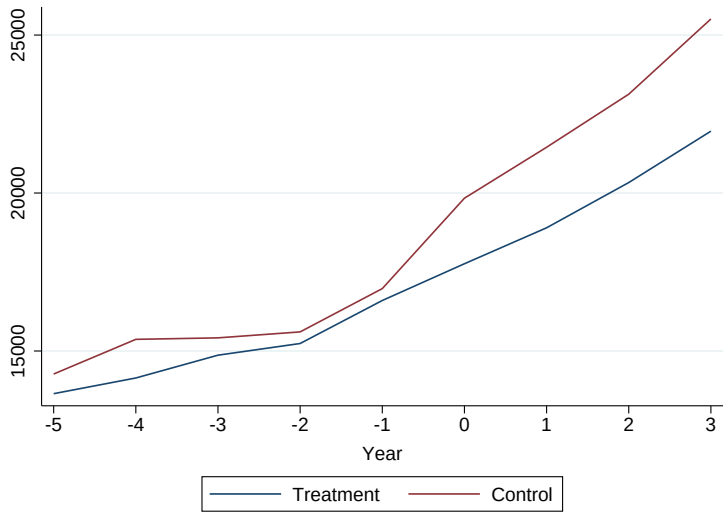

**(a)** 1040 HUD Income



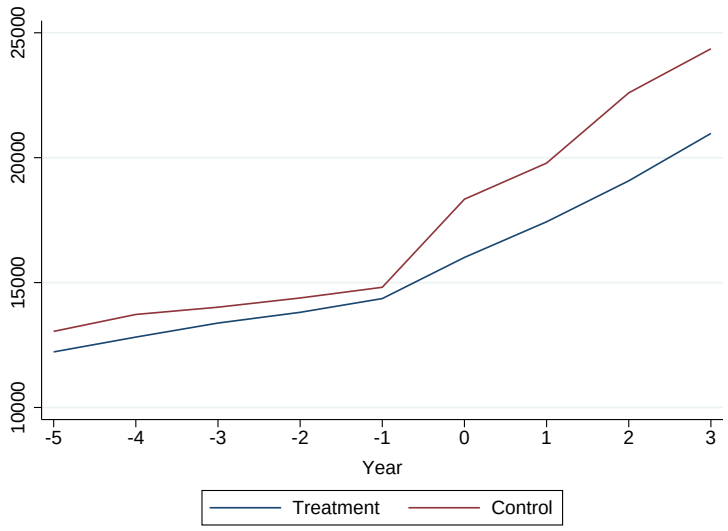

**(b)** W-2 Wages

*Notes:* The sample includes all applicants in the synthetic difference-in-differences analysis sample as described in Section 6.1. The "Control" group includes applicants whose applications were canceled. The "Treatment" group refers to applicants who become LIHTC residents. Income and wages are measured in 2019 USD.

**Figure 15.** Labor Market Outcomes:
Unconditional Means by Event Time and Treatment Group,
Matched-Pairs Sample



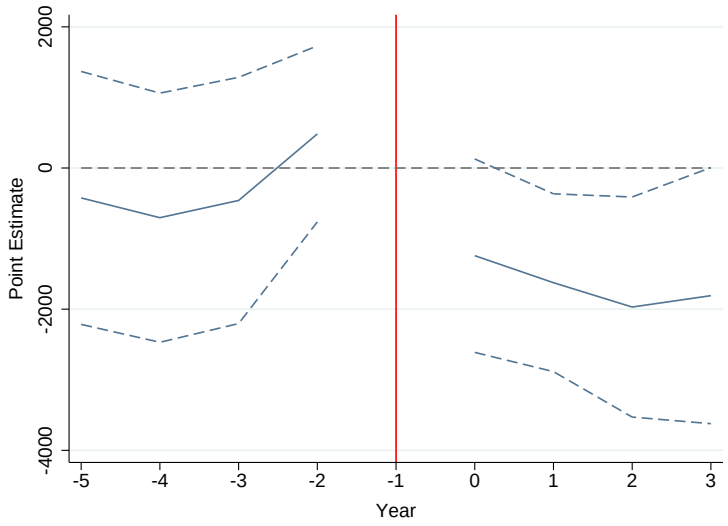

**(a)** 1040 HUD Income



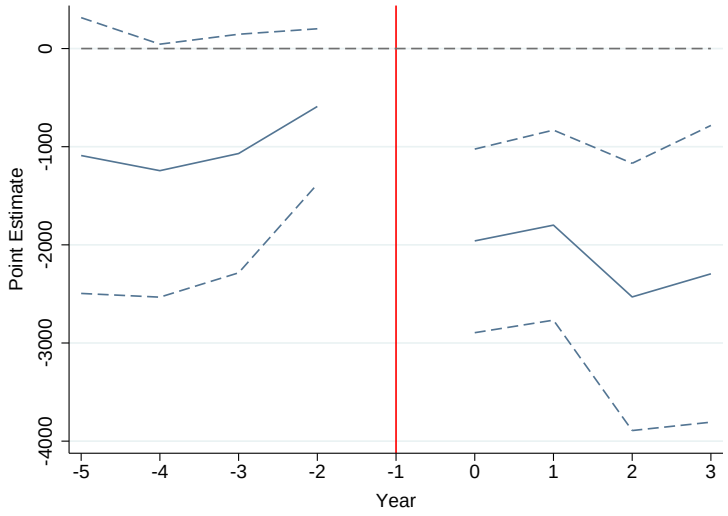

**(b)** W-2 Wages

*Notes:* The sample includes all applicants in the matched-pairs analysis sample as described in Section 6.1. The "Control" group includes applicants whose applications were canceled. The "Treatment" group refers to applicants who become LIHTC residents. Income and wages are measured in 2019 USD.

**Figure 16.** Labor Market Outcomes:
Event Study Results, OLS TWFE




**(a)** 1040 HUD Income

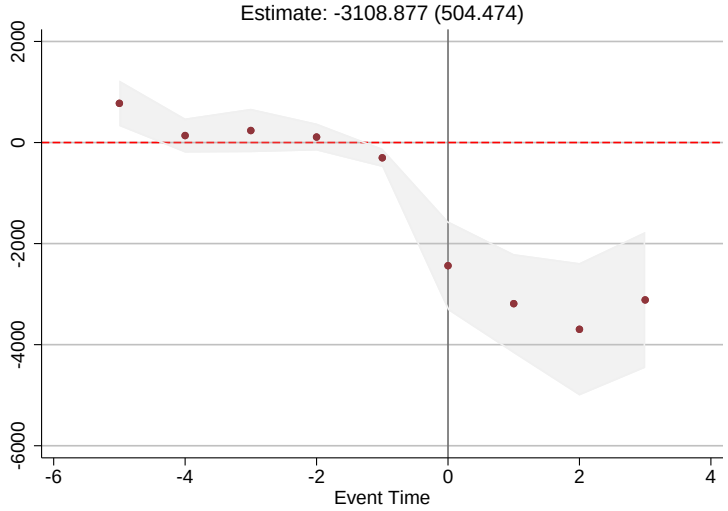



**(b)** W-2 Wages

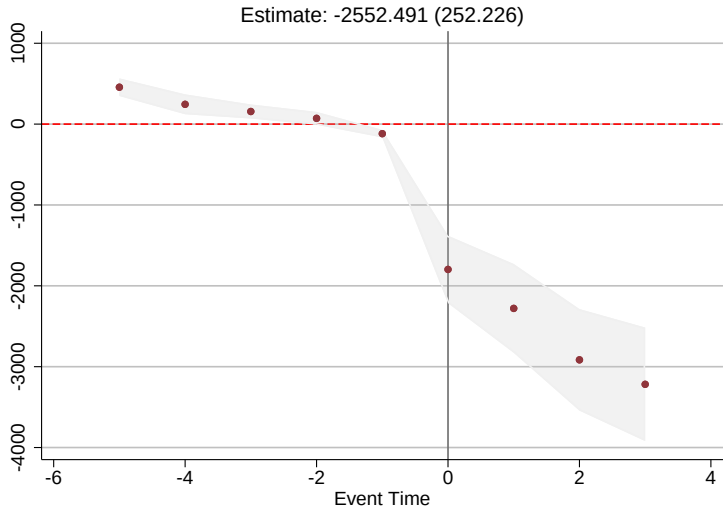
*Notes:* The sample includes all applicants in the matched-pairs sample as described in Section 6.1. Event study estimation follows Equation 3. Income and wages are measured in 2019 USD.

.

**Figure 17.** Labor Market Outcomes:
Event Study Results, SDID




**(a)** 1040 HUD Income




**(b)** W-2 Wages

*Notes:* The sample includes all applicants in the synthetic difference-in-differences sample as described in Section 6.1. Bootstrap standard errors calculated using 100 replications. Income and wages are measured in 2019 USD.

.

**Table 1.** Summary Statistics: LIHTC-Funded Properties

| | Sample Properties | Other LIHTC |
|---|---|---|
| *Tract Characteristics* | | |
| Poverty rate | 0.18 | 0.24 |
| P(top 20% income \| parent income 25p) | 0.093 | 0.11 |
| Median rent | 744.9 | 956.2 |
| Pop per sq mi | 1055.0 | 8308.3 |
| Share non-white | 0.17 | 0.51 |
| Qualified Census Tract | 0.16 | 0.21 |
| N | 357 | 41661 |
| *Property Characteristics* | | |
| # units | 40.0 | 68.3 |
| # low income units | 39.3 | 61.1 |
| New construction | 0.62 | 0.62 |
| 70% LIHTC credit | 0.83 | 0.65 |
| 30% LIHTC credit | 0.50 | 0.44 |
| Project-Based Rental Assistance | 0.21 | 0.27 |
| Rural Development 515 | 0.33 | 0.10 |
| HOME | 0.22 | 0.14 |
| Target pop = families | 0.26 | 0.36 |
| Metro, outside city center | 0.96 | 0.83 |
| Metro, in city center | 0.028 | 0.13 |
| N | 193 | 38112 |

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

**Table 2.** Mobility Sample:
Applicant Summary Statistics

| | (1) Non-Resident Mean | (2) Resident Mean | (3) Raw Diff | (4) Adj. Diff |
|---|---|---|---|---|
| *Application Characteristics* | | | | |
| BR preference | 2.15 | 2.15 | 0.007 | 0.061*** |
| Quoted rent | 599.4 | 676.0 | 76.6*** | 102.4*** |
| Age at application | 39.5 | 44.8 | 5.36*** | 2.49*** |
| Applied from same zip | 0.26 | 0.51 | 0.25*** | 0.096*** |
| Applied from same tract | 0.10 | 0.32 | 0.22*** | 0.079*** |
| *Application Tract Characteristics* | | | | |
| Poverty rate | 0.17 | 0.18 | 0.006** | 0.0065** |
| P(top 20% income \| parent income 25p) | 0.097 | 0.093 | -0.003*** | -0.003** |
| Median rent | 835.2 | 793.4 | -41.9*** | -5.24 |
| Pop per sq mi | 2059.0 | 1644.4 | -414.6*** | 15.1 |
| Share non-white | 0.25 | 0.22 | -0.037*** | -0.003 |
| *Property Tract Characteristics* | | | | |
| Poverty rate | 0.19 | 0.19 | -0.002 | 0.007 |
| P(top 20% income \| parent income 25p) | 0.088 | 0.090 | 0.002** | 0.003 |
| Median rent | 768.9 | 750.8 | -18.1*** | -8.87 |
| Pop per sq mi | 1280.9 | 1133.9 | -147.1*** | 2.18 |
| Share non-white | 0.23 | 0.20 | -0.029*** | -0.008 |
| Qualified Census Tract (QCT) | 0.20 | 0.19 | -0.010 | 0.032 |
| N | 6936 | 3885 | 10821 | 10808 |

*Notes:* The sample includes all applicants in the main analysis sample as described in Section 5.1. In the final column, we estimate regression adjusted differences with property fixed effects and cluster standard errors at the property level.
* – significant at 10%, ** – significant at 5%, *** – significant at 1%.

**Table 3.** Tax Data Sample:
Applicant Summary Statistics, Matched Pairs

| | (1) Non-Residents | (2) Residents | (3) Raw Diff | (4) Adj. Diff |
|---|---|---|---|---|
| Bedroom Preference | 1.867 | 1.815 | -0.052 | -0.052 |
| Age | 39.628 | 40.515 | 0.887*** | 0.739* |
| Form 1040 Filer | 0.733 | 0.702 | -0.030*** | -0.025** |
| Married Filing Jointly | 0.158 | 0.187 | 0.030*** | 0.025** |
| Number of Dependents | 0.981 | 1.144 | 0.163*** | 0.145*** |
| N | 3,073 | 8,085 | 11,158 | 11,158 |

*Notes:* The sample includes all applicants in the matched-pairs analysis sample as defined in Section 6.1. In the final column, we estimate regression adjusted differences with property fixed effects and cluster standard errors at the property level.
* – significant at 10%, ** – significant at 5%, *** – significant at 1%.

**Table 4.** Tax Data Sample:
Applicant Summary Statistics

| | (1)<br>Non-Residents | (2)<br>Residents | (3)<br>Raw Diff |
|---|---|---|---|
| Bedroom Preference | 1.503 | 1.506 | 0.003 |
| Age | 34.856 | 41.684 | 6.828*** |
| Form 1040 Filer | 0.826 | 0.660 | -0.166*** |
| Married Filing Jointly | 0.166 | 0.191 | 0.025*** |
| Number of Dependents | 0.825 | 1.016 | 0.191*** |
| N | 15,594 | 21,572 | 37,166 |
| Filers | 12,879 | 14,227 | 27,106 |

*Notes:* The sample includes all applicants in the SDID analysis sample as defined in Section 6.1.
* – significant at 10%, ** – significant at 5%, *** – significant at 1%.

**Table 5.** Mobility Outcomes:
Difference-in-Difference ATT Estimates

| | Poverty rate | | | P(top 20% income) \| parent inc 25p | | | Median rent | | |
|---|---|---|---|---|---|---|---|---|---|
| | OLS | | SDID | OLS | | SDID | OLS | | SDID |
| | (1) | (2) | (3) | (4) | (5) | (6) | (7) | (8) | (9) |
| LIHTC Resident | 0.192 | 0.160 | 0.240 | 0.097 | 0.114 | 0.041 | 7.243 | 7.521 | 8.184 |
| | (0.275) | (0.285) | (0.155) | (0.115) | (0.116) | (0.069) | (7.758) | (7.978) | (6.542) |
| | [0.485] | [0.575] | [0.123] | [0.401] | [0.325] | [0.552] | [0.351] | [0.347] | [0.211] |
| Baseline Controls | X | X | | X | X | | X | X | |
| App Controls | | X | | | X | | | X | |
| Base Dep. Var. Mean | 17.33 | 17.33 | 17.33 | 9.56 | 9.56 | 9.56 | 822.47 | 822.47 | 822.47 |
| N | 427378 | 427378 | 428015 | 427231 | 427231 | 427868 | 209867 | 209867 | 211239 |
| N (applicants) | 8722 | 8722 | 8735 | 8719 | 8719 | 8732 | 4283 | 4283 | 4311 |

*Notes:* The sample contains applicants in the main analysis sample as described in Section 5.1. For each specified outcome, we further restrict to a balanced panel of applicants with non-missing outcome data in the 24 months before and after the month of application. We estimate ATT point estimates and 95% confidence intervals using SDID regressions. Baseline controls include: applicant, event month, event month × property, and event month × calendar year of application × calendar month of application FEs. App controls include: age at application × event month and bedroom preference × event month FEs. SEs (bootstrapped, 100 reps) are in parentheses; p-values are in brackets.

# 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* .

Aliprantis, D. and Richter, F. G.-C. (2020), 'Evidence of neighborhood effects from moving to opportunity: Lates of neighborhood quality', *The Review of Economics and Statistics* **102**(4), 633–647.

Arkhangelsky, D., Athey, S., Hirshberg, D. A., Imbens, G. W. and Wager, S. (2021), 'Synthetic difference-in-differences', *American Economic Review* **111**(12), 4088–4118.
**URL:** *https://www.aeaweb.org/articles?id=10.1257/aer.20190159*

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.

Axel-Lute, M. (2023), 'The only tool in the box? what it means that lihtc dominates affordable housing', *Shelterforce* . Accessed: 2025-09-23.
**URL:** *https://shelterforce.org/2023/12/08/the-only-tool-in-the-box-what-it-means-that-lihtc-dominates-affordable-housing/*

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.

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* .

Billings, S. B., Hoekstra, M. and Rotger, G. P. (2026), 'Good neighborhoods, good neighbors, good jobs?', *NBER Working Paper 34872* .

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. and Ganong, P. (2018), 'How do changes in housing voucher design affect rent and neighborhood quality?', *American Economic Journal: Economic Policy* **10**(2), 62–89.

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* .

Cunningham, M. K., Galvez, M. M., Aranda, C., Santos, R., Wissoker, D., Oneto, A. D., Pitingolo, R. and Crawford, J. (2018), 'A pilot study of landlord acceptance of housing choice vouchers', *Urban Institute* .

Davis, M., Hartley, D. and Gregory, J. (2024), 'Preferences over the racial composition of neighborhoods: Estimates and implications', *Working Paper* .

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. and Horn, K. (2017), 'Points for Place: Can State Governments Shape Siting Patterns of Low-Income Housing Tax Credit Developments?', *New York University Furman Center* .

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.

Ellen, I. G., O'Regan, K. and Strochak, S. (2024), 'Race, space, and take up: Explaining housing voucher lease-up rates', *Journal of Housing Economics* **63**, 101980.

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.

Galster, G. C. (2019), 'Neighborhoods and national housing policy: Toward circumscribed, neighborhood-sensitive reforms', *Housing Policy Debate* **29**(1), 217–231.

Galvez, M. M. and Oppenheimer, S. (2020), Taking neighborhood mobility to scale through the housing choice voucher program, Technical report, Urban Institute.

Hauck, T. and Wallossek, L. (2024), 'Optional (non-) filing and effective taxation', *Journal of Public Economics* **238**, 105187.

HUD (2023), 'Low-Income Housing Tax Credit (LIHTC): Property Level Data'.
**URL:** *https://www.huduser.gov/portal/datasets/lihtc/property.html*

Jacob, B. A. and Ludwig, J. (2012), 'The effects of housing assistance on labor supply: Evidence from a voucher lottery', *American Economic Review* **102**(1), 272–304.

Kawitzky, S., Freiberg, F., Houk, D. L. and Hankins, S. (2013), 'Choice constrained, segregation maintained: Using federal tax credits to provide affordable housing', *New York, NY: Fair Housing Justice Center* .

Keightley, M. P. (2023), An Introduction to the Low-Income Housing Tax Credit, CRS Report RS22389, Congressional Research Service.

Keithley, D. (2022), 'Housing tax credit resident profile', *Ohio Housing Finance Agency* .

Lens, M. C., Ellen, I. G. and O'Regan, K. (2011), 'Do vouchers help low-income households live in safer neighborhoods? evidence on the housing choice voucher program', *Cityscape* pp. 135–159.

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.

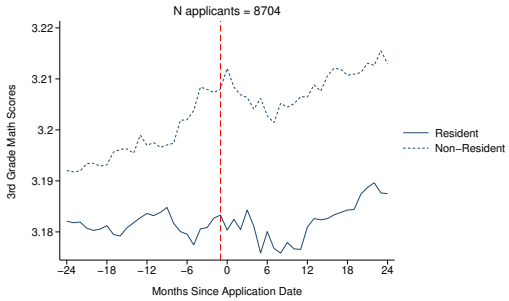
Sard, B., Rice, D., Bell, A. and Mazzara, A. (2018), 'Federal policy changes can help more families with housing vouchers live in higher-opportunity areas', *Washington, DC: Center on Budget and Policy Priorities* .

Soltas, E. (2023), Tax incentives and the supply of affordable housing, Working paper.

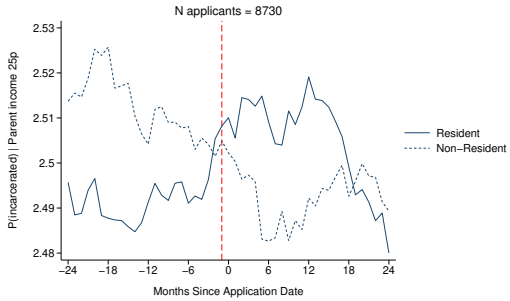
Susin, S. (2005), 'Longitudinal outcomes of subsidized housing recipients in matched survey and administrative data', *Cityscape* pp. 189–218.

# A Additional Tables and Figures

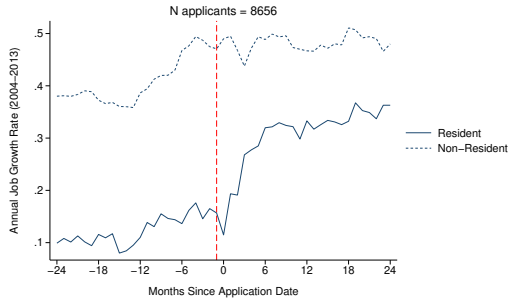
**Figure A1.** Mobility Outcomes:
Unconditional Means by Event Time and Treatment Group
Additional Outcomes



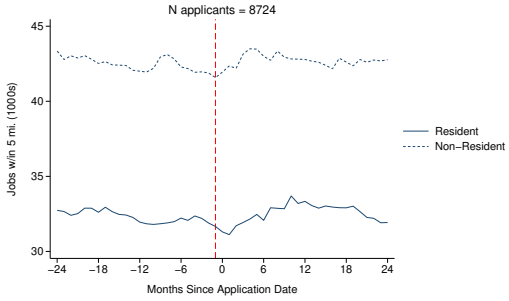

**(a)** 3rd grade math scores



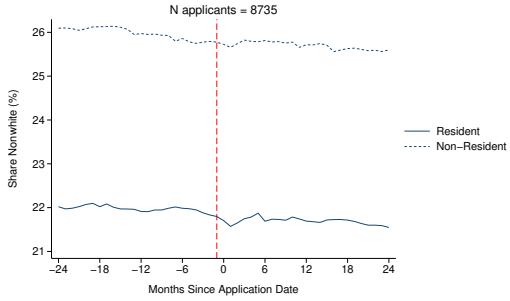

**(b)** P(incarcerated | parent inc 25p)



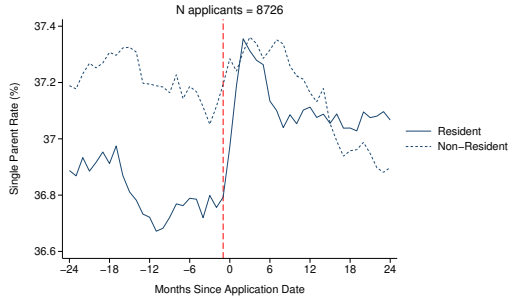

**(c)** Annual job growth rate




**(d)** Jobs w/in 5 miles



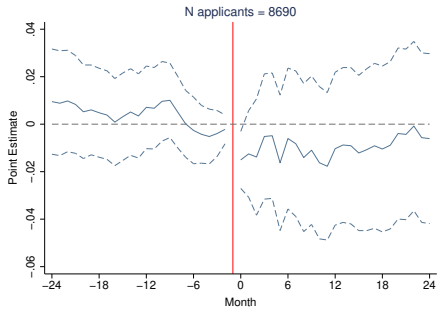

**(e)** Share non-white



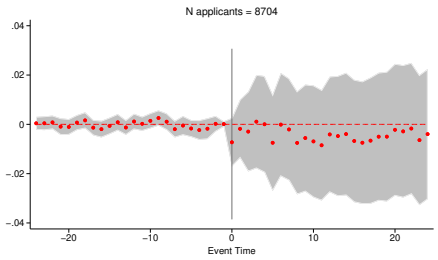

**(f)** Share single parent

*Notes:* The sample contains applicants in the main analysis sample as described in Section 5.1. For each specified outcome, we further restrict to a balanced panel of applicants with non-missing outcome data in the 24 months before and after the month of application.

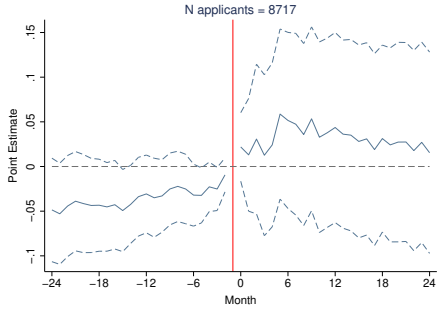
**Figure A2.** Mobility Outcomes:
Event Study Results, Continuous
Additional Outcomes



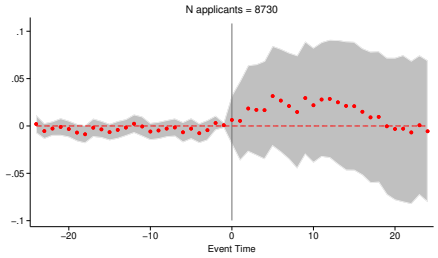

**(a)** 3rd grade math scores
OLS TWFE



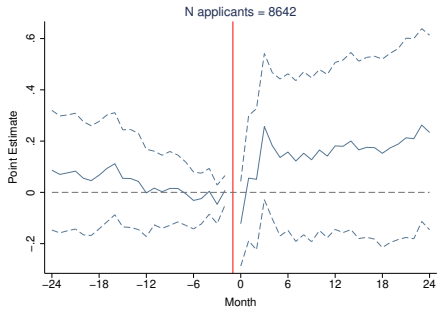

**(b)** 3rd grade math scores
SDID



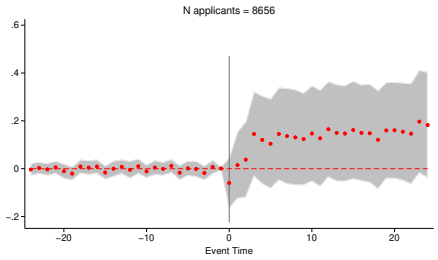

**(c)** P(incarcerated | parent inc 25p)
OLS TWFE




**(d)** P(incarcerated | parent inc 25p)
SDID



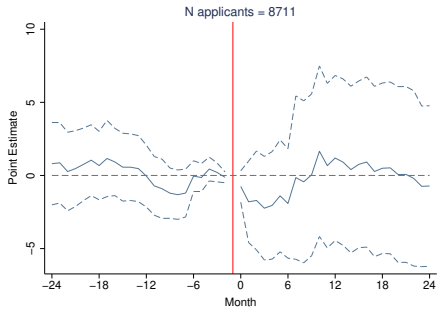

**(e)** Annual job growth rate
OLS TWFE



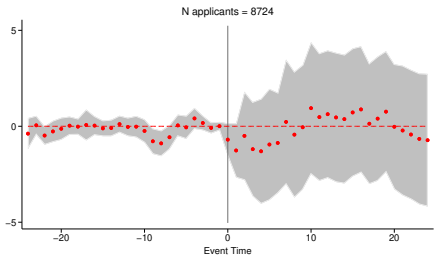

**(f)** Annual job growth rate
SDID

*Notes:* The sample contains applicants in the main analysis sample as described in Section 5.1. For each specified outcome, we further restrict to a balanced panel of applicants with non-missing outcome data in the 24 months before and after the month of 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 the continuous specification of Equation 1. We cluster standard errors at the property level in OLS TWFE specifications and estimate bootstrapped standard errors (100 reps) in SDID specifications.

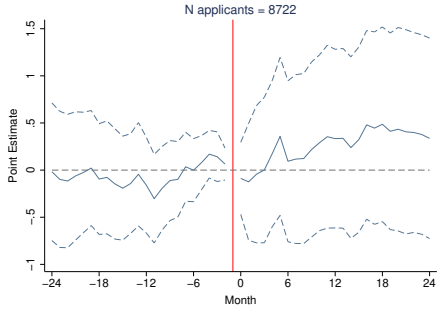
**Figure A3.** Mobility Outcomes:
Event Study Results, Continuous
Additional Outcomes



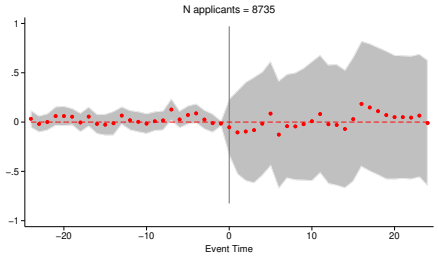

**(a)** Jobs w/in 5 miles
OLS TWFE



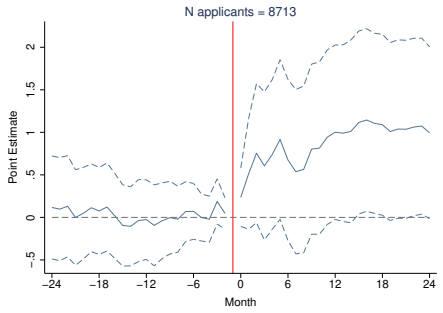

**(b)** Jobs w/in 5 miles
SDID



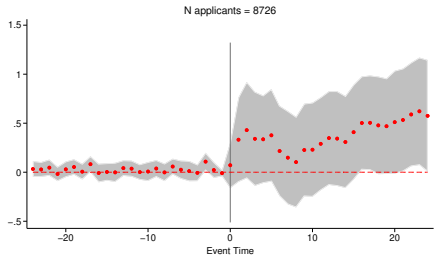

**(c)** Share non-white
OLS TWFE




**(d)** Share non-white
SDID



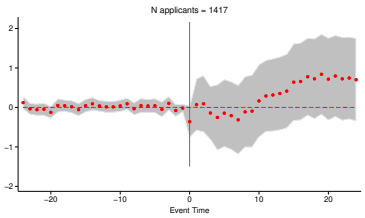

**(e)** Share single parent
OLS TWFE



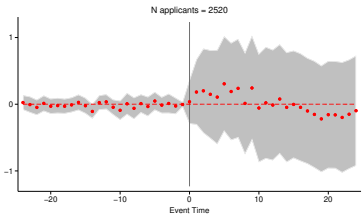

**(f)** Share single parent
SDID

*Notes:* The sample contains applicants in the main analysis sample as described in Section 5.1. For each specified outcome, we further restrict to a balanced panel of applicants with non-missing outcome data in the 24 months before and after the month of 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 the continuous specification of Equation 1. We cluster standard errors at the property level in OLS TWFE specifications and estimate bootstrapped standard errors (100 reps) in SDID specifications.

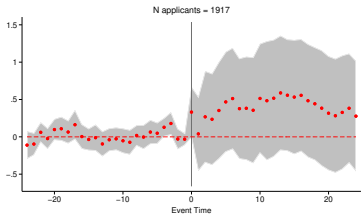
**Figure A4.** Mobility Outcomes:
Heterogeneity by BR



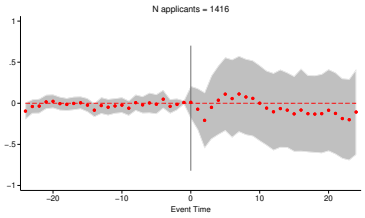

**(a)** Poverty Rate
1 BR



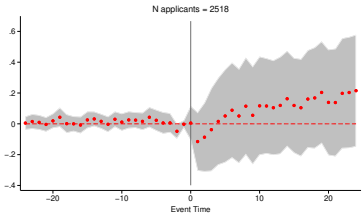

**(b)** Poverty Rate
2 BR



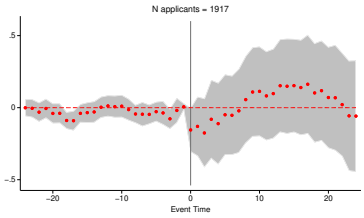

**(c)** Poverty Rate
3+ BR



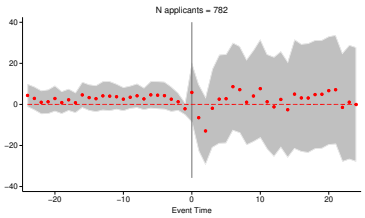

**(d)** P(top 20% inc) | parent inc 25p
1 BR



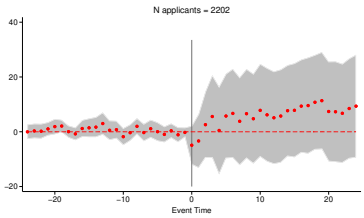

**(e)** P(top 20% inc) | parent inc 25p
2 BR



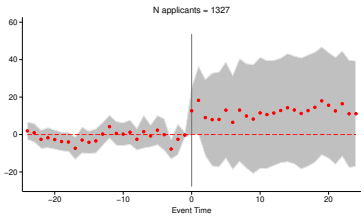

**(f)** P(top 20% inc) | parent inc 25p
3+ BR




**(g)** Median rents
1 BR



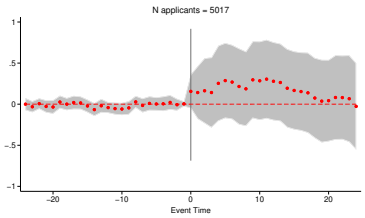

**(h)** Median rents
2 BR



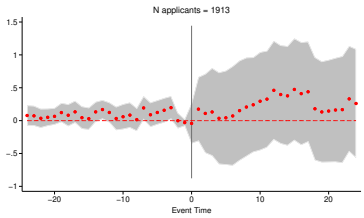

**(i)** Median rents
3+ BR

*Notes:* The sample contains applicants in the main analysis sample as described in Section 5.1. For each specified outcome, we further restrict to a balanced panel of applicants with non-missing outcome data in the 24 months before and after the month of application. Each event study plots point estimates and 95% confidence intervals from separate SDID regressions of the outcome on event month treatment indicators and the baseline set of controls, following the continuous specification of Equation 1 with bootstrapped SEs (100 reps).

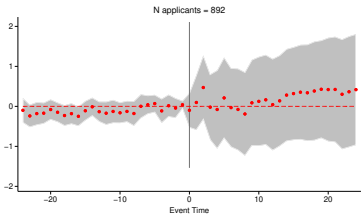
**Figure A5.** Mobility Outcomes:
Heterogeneity by Age



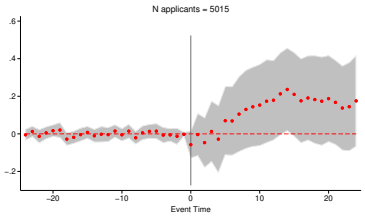

**(a)** Poverty Rate
Under 45



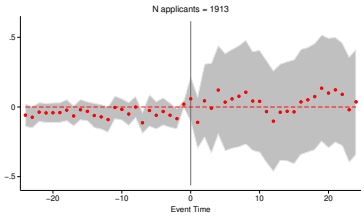

**(b)** Poverty Rate
45-64



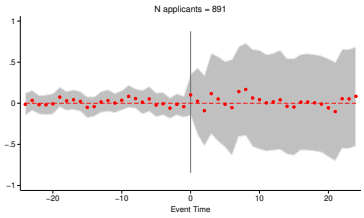

**(c)** Poverty Rate
65+



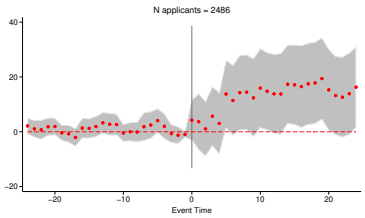

**(d)** P(top 20% inc) | parent inc 25p
Under 45



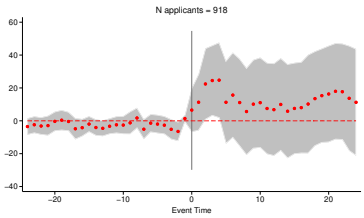

**(e)** P(top 20% inc) | parent inc 25p
45-64



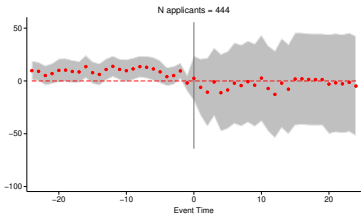

**(f)** P(top 20% inc) | parent inc 25p
65+




**(g)** Median rents
Under 45



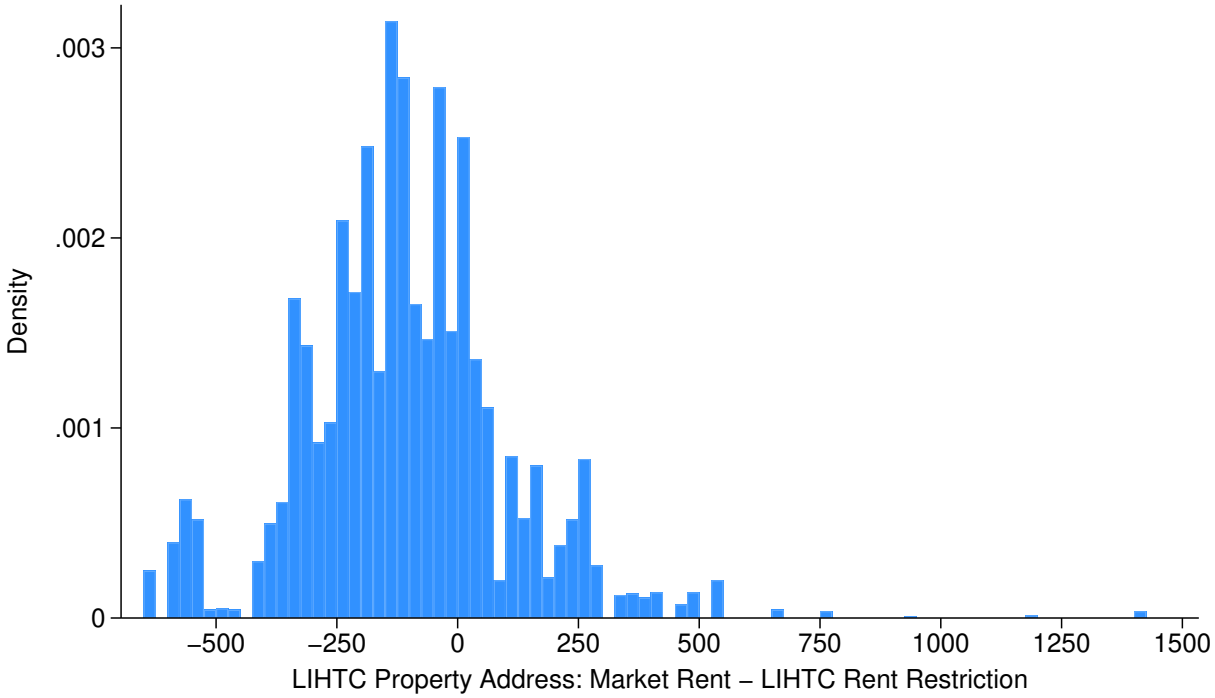

**(h)** Median rents
45-64




**(i)** Median rents
65+

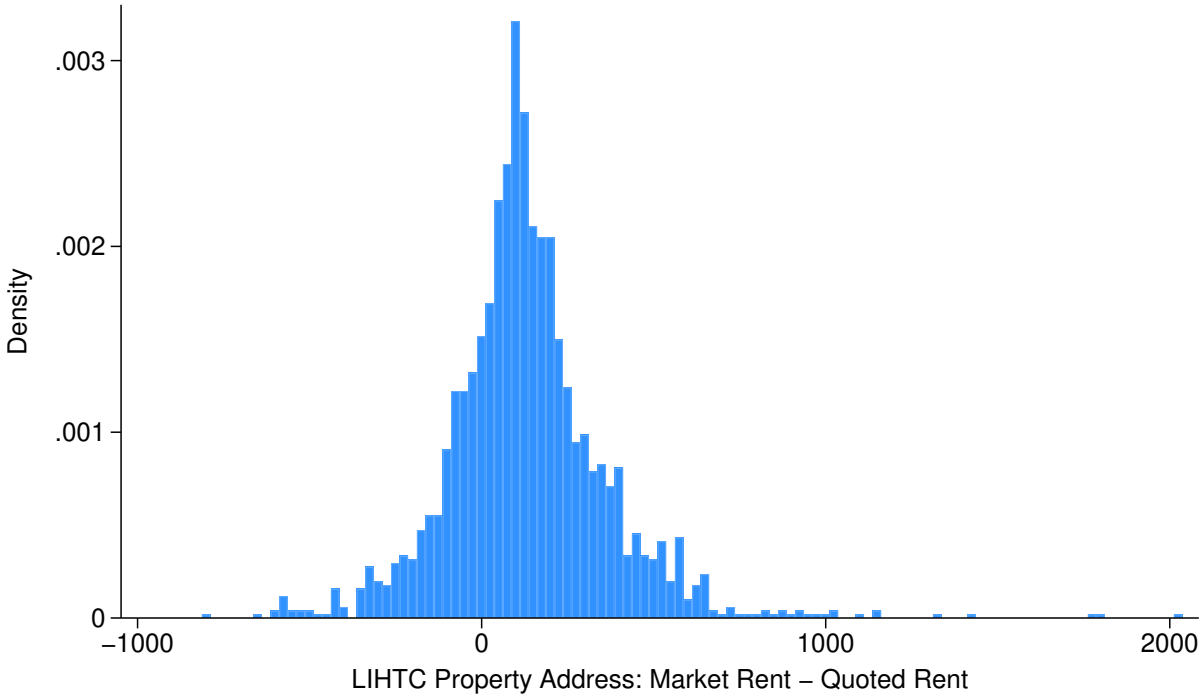
*Notes:* The sample contains applicants in the main analysis sample as described in Section 5.1. For each specified outcome, we further restrict to a balanced panel of applicants with non-missing outcome data in the 24 months before and after the month of application. Each event study plots point estimates and 95% confidence intervals from separate SDID regressions of the outcome on event month treatment indicators and the baseline set of controls, following the continuous specification of Equation 1 with bootstrapped SEs (100 reps).

**Figure A6.** Market Rents vs. LIHTC Rent Restrictions




*Notes:* The sample contains applicants in the main analysis sample as described in Section 5.1. Values pertain to per month gross rents in 2017 USD.

**Figure A7.** Market Rents vs. LIHTC Quoted Rents




*Notes:* The sample contains applicants in the main analysis sample as described in Section 5.1. Values pertain to per month gross rents in 2017 USD.