Tax Incentives and the Supply of Low-Income Housing

Evan Soltas*

May 11, 2024

(Click for Most Recent Version)

Abstract

This paper studies the impacts and incidence of subsidies for low-income housing development, which are often portrayed as transfers to the developers of inframarginal projects. I estimate a dynamic model of developer behavior using new data on competitions for Low-Income Housing Tax Credits and three sources of policy variation: quasi-random assignment of subsidies, shocks to subsidy generosity, and nonlinearities in scoring rules for subsidy applications. I find that subsidies add few net units to the housing stock and instead pull investment forward in time. Households benefit from modest rent discounts on subsidized units, but developers capture around half of the subsidy in profits, and another quarter is dissipated in their fixed costs of competing for subsidies. Due to lower developer incidence and entry costs, a voucher program could likely generate similar household benefits at less fiscal cost.

*MIT Department of Economics and Federal Reserve Bank of Boston. Email: esoltas@mit.edu. This project would not have been possible without the assistance of state housing finance agency staff. I am particularly grateful to Jim Poterba, David Autor, and Jon Gruber for their advice and support. I also wish to thank Rebecca Diamond, Denise DiPasquale, Ingrid Gould Ellen, Colin Higgins, Mike Hollar, Allan Hsiao, Dev Patel, Charlie Rafkin, Anna Russo, Tobias Salz, Ben Sprung-Keyser, Bill Wheaton, Paul Willen, and participants in the MIT PF Lunch. Umang Bansal, Cristine Bren, and Jayson Roxas provided excellent research assistance. I appreciatively acknowledge support from the Jerry A. Hausman Dissertation Fellowship, the Lincoln Institute of Land Policy's C. Lowell Harriss Dissertation Fellowship, the Lynde and Harry Bradley Foundation, the George and Obie Shultz Fund, and the National Science Foundation Graduate Research Fellowship under Grant No. 1122374. The views expressed in this paper are solely those of the author and not necessarily those of the Federal Reserve Bank of Boston nor the Federal Reserve System.

1 Introduction

Governments subsidize housing for lower-income people in three main ways: public housing, "tenant-based" vouchers, and "project-based" assistance. Unlike vouchers, project-based subsidies work through developers and the supply side of housing markets, supporting the construction, renovation, and operation of income-restricted housing at regulated rents. In the United States, project-based assistance is a major component of federal housing policy, despite longstanding and unresolved questions about its impacts and incidence.

In 2022, around seven million U.S. households received some form of low-income housing assistance, more than half of whom received it through project-based programs. The largest project-based subsidy, the Low-Income Housing Tax Credit (LIHTC), has funded one in five of all new U.S. multifamily units since 1987 through grants. About two percent of all U.S. households lived in LIHTC units in 2022, more than received vouchers in the same year and more than lived in public housing at its historical peak. Structured as a corporate income tax expenditure, the LIHTC reduces annual federal receipts by about $15 billion. Its rise amounts to a momentous but little-discussed reorientation of U.S. housing policy toward project-based assistance.

The key concern about project-based subsidies is that they may be transfers to the developers of inframarginal projects, delivering neither net new units nor benefits to households. Glaeser and Gyourko (2008), for instance, argue the LIHTC "essentially functions as a transfer program" to developers, and Quigley (2011) writes project-based subsidies "must be justified on some other basis" than as transfers to lower-income households. Advocates for the LIHTC dispute this characterization, arguing it is effective in both expanding the housing stock and in redistributing to lower-income households.[1] Indeed, housing-policy experts currently hold views about the LIHTC ranging from favoring its expansion to its abolition. Yet due to a host of empirical challenges, debates over the fundamental design of housing assistance have remained almost frozen in time for fifty years, when the U.S. turned away from public housing and toward vouchers.[2]

This paper revisits the impacts and incidence of project-based assistance using newly-collected data on the applications of developers competing for the LIHTC. I estimate a dynamic model of developer behavior using their responses to three sources of policy variation, each of which is linked to a model primitive in the developer's problem. These analyses show that developers respond to subsidy awards mostly by retiming investment, rather than by producing net new units, and that their applications for subsidies are highly responsive to changes in subsidy generosity

[1]Other arguments for project-based subsidies are that they may enable spatially-targeted investment, alleviate housing discrimination, and be easily combined with supportive services for populations with complex housing needs.

[2]Kazis (2022) writes that "a rough equilibrium has held for years...it has been decades since Congress or HUD has seriously reconsidered the basic forms of rental assistance." The latest review of research on U.S. housing policy (Olsen and Zabel, 2015) writes that there is "no [recent] high-quality evidence on cost-effectiveness" of project-based assistance and describes this topic as the "highest priority for housing policy research."

and scoring rules. Using the estimated model, I find the LIHTC does little to expand the housing stock on net. Households benefit from subsidies via modest rent discounts on subsidized units, but economists' concerns appear well-founded: Developers capture nearly half of the welfare gains. Developers' fixed costs to compete for subsidies are also high, further limiting the benefits to households. In counterfactuals that compare the LIHTC to a stylized voucher program, I conclude similar household benefits could likely be achieved at less fiscal cost through vouchers.

New data, introduced in Section 2, enable my analysis. Through public-information requests, I collected administrative records on 453 rounds of LIHTC competitions from 40 states, covering 22,241 applications from 2005 through 2019 with requested subsidies of $200 billion in total. I link these records to parcel-level tax assessments and neighborhood-level outcomes. Finally, I code states' grant rules from regulatory documents, which I use to compute ex-ante probabilities with which developers could expect to win by simulating the mechanism. These rich data enable careful study of developers, the key actors in project-based subsidies.

Section 3 presents the model of housing markets. In the model, each developer is associated with a land parcel and may enter it into a grant competition. The developer's problem is intrinsically dynamic. Instead of entering the competition, a developer can develop without a subsidy now or wait, possibly to apply or to develop later—and they may also reapply in the future if they lose. Developer behavior is shaped by three primitive objects: their value of winning the subsidy, their outside option, and their cost of applying for the subsidy. Households rent the entire housing stock, so they benefit from subsidies through rent savings in below-market units as well as general-equilibrium effects on market-rent units. While reductions in market rents will occur only if subsidies expand the housing stock, rent discounts allow subsidies to benefit households even if subsidized units are purely inframarginal. Developers take win probabilities as given, and equilibrium follows from profit maximization, market clearing, and rational expectations.

I start the empirical analysis by measuring the size of LIHTC rent discounts and to whom they accrue. Average monthly LIHTC rents are about 12 percent below my estimates of potential market rents for new units in the same neighborhoods, though the rent savings vary greatly across neighborhoods. To my knowledge, these are the first national estimates of LIHTC rent discounts, and they amount to less than a third of the LIHTC's fiscal cost.[3] At the same time, I also find that LIHTC tenants are poorer than the likely tenants of counterfactual developments. I estimate, for instance, that the LIHTC roughly triples the share of tenant households with incomes less than $20,000 who ultimately live there. Taken together, these facts imply the LIHTC rent discount is a progressive transfer, albeit modest in size. To understand how the LIHTC affects developers, I turn

[3]My estimates appear consistent with expert knowledge and two local estimates. Relying on their policy experience, Khadduri and Wilkins (2008) write that LIHTC rents are "indistinguishable from market rents or only slightly below market." An analysis of LIHTC properties in Tallahassee, Florida finds modest rent savings (Burge, 2011). The larger savings found by Cook et al. (2023) for the Chicago metropolitan area appear consistent with my national estimates.

to the three quasi-experimental analyses.

The first empirical strategy (Section 4) uses quasi-random assignment of subsidies to estimate the causal effects of winning on development activity. Developers apply under uncertainty about their rivals, making subsidy assignment as good as random conditional on an estimable win probability. I find small causal effects of winning on development: 75 percent of applicant parcels would be developed within ten years if they were to lose. Instead, the LIHTC pulls development forward in time, and it replaces some market-rate units with (slightly) subsidized ones. These results are confirmed by neighborhood-level event studies that compare areas where winners and losers applied with similar ex-ante win probabilities. From the model's point of view, such behavior suggests most applicants have profitable outside options.

The second empirical strategy (Section 5) estimates application supply responses to changes in subsidy generosity. Through an event study and regression discontinuity design (RDD), I exploit annual policy variation by Census tract in the LIHTC's generosity. This variation is induced by a cutoff-based rule whose inputs are measured with sampling error. Across both approaches, I find that a ten-percent reduction in the net-of-tax cost of low-income housing development leads to an increase in applications of at least three percent. The model-based implication is that entry costs are high, as many developers apply only when the tax credits cover almost the entire construction cost of low-income housing, even when these units command near-market rents.

The third empirical strategy (Section 6) infers the value to developers of winning the LIHTC from their responses to application scoring rules. In applications, developers may "bid" lower rents than the maximum allowed under federal regulations, so as to raise their win probabilities. I show that scoring rules create strong and nonlinear incentives to reduce rents. Developers respond to these incentives, bidding rents down and bunching at kink points. Such behavior implies winning is profitable, as developers would otherwise not trade rental income for a higher win probability on the margin. To quantify profit incidence, I simulate unilateral deviations to different rents and obtain developer valuations through bid inversion. The average developer behaves as if they are indifferent between a 0.9-percentage-point increase in win probability and a \$1,000 increase in present-value rental income per unit. These findings imply a developer incidence share of about 45 percent, matching the structural estimates that incorporate dynamic considerations.

I estimate the model in Section 7 to provide a unified explanation of developer behavior from the three quasi-experimental analyses.[4] Although the model is jointly estimated, I show how each analysis is linked to a primitive: outside options for the causal effects of subsidy awards, entry costs for responses to subsidy generosity, and win values for the trade-off between rental income

[4] Holmes and Sieg (2015) describe this approach as a way to unite the virtues of the quasi-experimental and structural approaches: The former credibly estimates behavioral responses, while the latter maps this description of behavior into the primitives of an equilibrium model.

and win probability. I use the estimated model to evaluate the LIHTC's impacts and incidence. Both of these objects of interest are necessarily defined relative to a counterfactual world without the LIHTC—one that is illuminated by, but is of course not implemented in, the quasi-experiments. Estimation draws on methods from dynamic discrete choice (Rust, 2000; Sweeting, 2013) and indirect inference (Gourieroux et al., 1993).

Section 8 reports the structural results. Estimates of model primitives reconcile the three empirical analyses with a simple explanation: Developers self-select into application primarily on the basis of their low entry costs, rather than their low outside options. That is, developers apply because of cost advantages in pursuing the subsidy, not because their parcels lack alternative profitable uses. Such results imply the LIHTC is ineffective in expanding the overall housing stock, instead merely retiming investment. For every ten LIHTC units, I find about eight units displace private housing that would have otherwise been built within ten years, and two units are net additions to the housing stock. Applying these displacement estimates to the LIHTC in aggregate, I conclude the LIHTC has expanded the U.S. housing stock by about 500,000 units, or by 0.4 percent.[5] Due to displacement, the fiscal cost of the LIHTC is about $1 million per net new unit on average.

Overall, the model results show that project-based subsidies provide some benefit to households. I calculate that households reap about 31 percent of the welfare gains, rejecting the hypothesis of complete incidence on developers. These gains arise mostly because LIHTC units set rents below market (23 percent of incidence) and in lesser part because the additional housing supply reduces market rents (8 percent of incidence). Yet most of the subsidy eludes households for two reasons: Developers capture a significant share (44 percent of incidence), and the remainder is competed away in the fixed costs that developers pay to compete for subsidies (25 percent of incidence). In driving profits to zero on the margin, marginal applicants pay heavy entry costs but do not eliminate the profits of inframarginal developers with entry-cost advantages.

In counterfactuals, I compare the LIHTC's incidence to that of a stylized voucher program.[6] I find vouchers could provide the same welfare benefit to households as the LIHTC at a 25-percent fiscal savings. Conversely, a balanced-budget reform that shifts from developer subsidies to vouchers could raise household welfare. These findings reflect two considerations. First, these subsidies have opposite-signed pecuniary externalities on the unsubsidized market. To offset effects on unsubsidized households, vouchers require more spending to achieve the same household benefit, all else equal. The second difference between the two policies is that vouchers reduce developer incidence and eliminate entry costs. On net, the balance of these forces favors vouchers, but the disadvantage of the LIHTC is relatively modest. Considerations beyond my analysis, many of which

[5]Impacts are larger as a share of the low-rent housing stock: 3 percent of units with monthly rents less than $1,000. The benefits to lower-income households would be further amplified by reallocation effects.

[6]I model a general subsidy for market-rate housing. This might instead be seen as a change in property taxes, depreciation schedules, or the tax treatment of rental income, and it abstracts away from voucher-specific frictions.

are project-specific in nature, suggests economists' predisposition against project-based assistance appears sensible, although some LIHTC projects may well be better than vouchers.

This paper contributes to several literatures in public finance and urban economics. The developer's choice to build subsidized housing imposes a fundamental participation constraint on the social problem of optimal housing policy (Soltas, forthcoming). Despite housing policy's shift from direct provision via public housing to market support via subsidies (Collinson et al., 2015), little research has sought to understand the supply side facing the government in this domain. This paper pursues this agenda for the largest U.S. project-based housing policy, bringing together a new model, new data, and new sources of quasi-experimental variation.

The incidence and effectiveness of project-based assistance are classic issues in housing policy. Early research on these topics was highly critical of public housing and project-based programs, and its criticisms were influential in the rise of vouchers.[7] Despite the LIHTC's ascendance, these issues have largely not been revisited. Recent research on the LIHTC has studied its displacement or "crowding-out" effects on the neighborhood-level housing stock (Sinai and Waldfogel, 2005; Baum-Snow and Marion, 2009; Eriksen and Rosenthal, 2010).[8] Yet both the economic explanation and incidence implications of such neighborhood-level displacement are fundamentally unclear. If displacement occurs because, at the parcel level, development is inframarginal to subsidy, then the subsidy's incidence on developer profits is likely substantial—whereas it is likely small if displacement reflects the spatial-equilibrium consequences of an elastic housing supply. This paper shows that displacement arises because of parcel-level inframarginality, not spatial-equilibrium effects, thus addressing both questions of impacts and incidence.[9]

Urban economists have recently adopted empirical techniques from industrial organization and mechanism design (Murphy, 2018; Diamond et al., 2019; Waldinger, 2021; Calder-Wang, 2022; Almagro and Domínguez-Iino, 2022; Hsiao, 2022; Cook et al., 2023). My approach to combining a dynamic model and quasi-experimental evidence may be well-suited for other housing-supply topics, as policy variation abounds but dynamic considerations loom large (Gyourko, 2009). Three recent related papers, all of which take static empirical approaches, find significant developer responses to housing tax and regulatory policies (Anagol et al., 2021; Levy, 2021, 2023). The dynamic framework allows me to capture key features of developer behavior in their pursuit of project-based subsidies, such as reapplication and private development after unsuccessful applications.

[7] This research primarily estimates the extent of productive and allocative inefficiency from housing policies. See Olsen (2003), Weicher (2012), and Olsen and Zabel (2015) for discussions. Economists have often argued for vouchers over project-based subsidies (Aaron, 1972; Rosenthal, 2014), although views are diverse (Favilukis et al., 2023).

[8] Scholars have also studied effects on local amenities (Freedman and Owens, 2011; Davis et al., 2019; Diamond and McQuade, 2019) and on resident outcomes (Ellen et al., 2016; Derby, 2021; Sportiche, 2022; Cook et al., 2023).

[9] It thus joins an empirical literature on the retiming responses to subsidies in markets for capital goods (House and Shapiro, 2008; Mian and Sufi, 2012; Berger et al., 2020). In modeling the project pipeline from application to completion, my framework could also be applied to study other competitive grants (e.g., Jacob and Lefgren, 2011).

2 Setting and Data

2.1 The Low-Income Housing Tax Credit

Established in the Tax Reform Act of 1986 and defined in Section 42 of the Internal Revenue Code, the LIHTC is an investment tax credit issued to developers against federal corporate income tax liabilities.[10] The credits reduce tax liabilities over ten years by a specified share of a project's "qualified basis." The basis typically includes construction costs associated with low-income units, as well as a developer fee, but it excludes land and predevelopment expenses.

Context. Figure 1 depicts the evolution of three leading housing programs as shares of the U.S. housing stock: public housing, rent vouchers, and the LIHTC. The era of public housing prevailed from the Great Depression until the 1970s, when it was ended by a construction moratorium and a watershed report, the National Housing Policy Review (1973), which argued that public housing was not cost-effective. A rapid transition to tenant-based subsidies followed, with sporadic use of project-based assistance (Orlebeke, 2000). By the 1990s, debates over the proper form of housing assistance were viewed as definitively resolved in favor of vouchers (Winnick, 1995). The rise of the LIHTC has quietly overturned this consensus. Since 2000, the LIHTC's growth has coincided with the stagnation of vouchers, reorienting policy toward project-based subsidies.

Applications. Each year, state housing finance agencies may issue credits up to a per-capita maximum. In 2022, the maximum is \$2.60 in ten-year credits per state resident, with an additional allowance for small states, amounting to a national annual tax expenditure of about \$10 billion (Joint Committee on Taxation, 2023). Agencies award credits through competitions in which they receive proposals and select some for funding.

Each agency awards funding according to its Qualified Allocation Plan (QAP), a public document of selection criteria. In 2019, all but two states used numerical rubrics that prevent discretion at the level of individual applications. An application i's tax-credit assignment depends upon its QAP score q_i and set-asides z_i, along with those of other applications $(\boldsymbol{Q}_{-i}, \boldsymbol{Z}_{-i})$. The primary purpose of set-asides is to balance the distribution of subsidies over geographies and demographic constituencies (e.g., seniors). The assignment $W_i \in \{0, 1\}$ is thus characterized by

$$W_i = W(q_i, z_i; \boldsymbol{Q}_{-i}, \boldsymbol{Z}_{-i}), \tag{1}$$

where I refer to $W(\cdot)$ as the "grant rule." I programmed the grant rules for each round in my data

[10] Throughout this paper, I focus on the "9%" (also known as "70%") LIHTC, which is awarded competitively. Another program, the "4%" (a.k.a. "30%") LIHTC, funds developments at a lower credit rate. In general, this program funds rehabilitation projects, rather than new construction, usually with tax-exempt bonds and without a competitive process. Collinson et al. (2015) review the LIHTC alongside other low-income housing programs.

Figure 1: U.S. Housing Policy in Historical Perspective

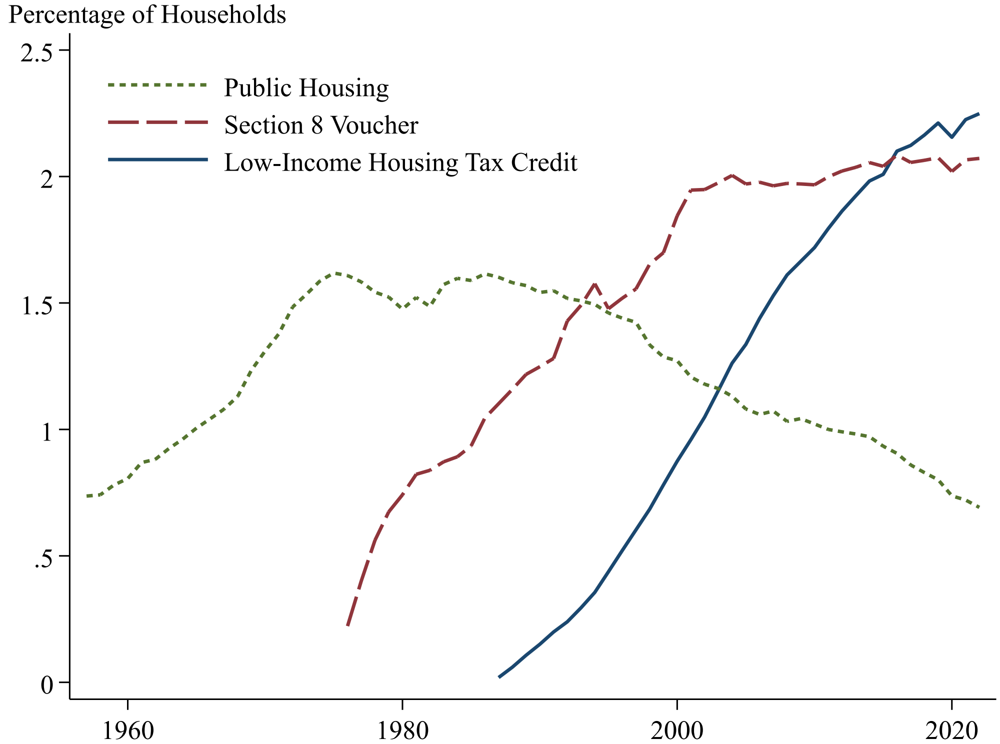

Notes: This figure displays the assisted shares of households for three housing policies: public housing (built under Section 9 of the U.S. Housing Act of 1937), tenant-based rental assistance (Section 8 vouchers and certificates), and project-based developments financed by Low-Income Housing Tax Credits (9% and 4% LIHTC). Annual assisted unit counts are drawn from Olsen (2003) and Vale and Freemark (2012). I extend the data to the present using the HUD Picture of Subsidized Households and LIHTC Databases. I also adjust for incomplete reporting in the final three years of the LIHTC time series. Household counts are from Census Table HH-1.

following my reading of the QAP. For instance, Georgia's 2019 QAP reserved 10 percent of credits for nonprofits, set aside another $15 million for rural projects that were already awarded another project-based subsidy (HOME), and then split the remaining subsidy between rural and non-rural regions, funding projects in score order within these regions.

Scoring rules $q_i = q(x_i)$ determine scores from application characteristics x_i, which I take to contain z_i and other variables. These rules are complex and differentiated across states (Shelburne, 2021), as federal regulations allow them to be "appropriate to local conditions." Common criteria include site selection, building amenities, and developer characteristics. Developers are thought to " 'chase points' when they make decisions" (Ellen and Horn, 2018). A key federal stipulation on QAPs is that they must incorporate a preference for "projects serving the lowest income tenants." As U.S. housing policy links income and rent levels, states meet this requirement with rules favoring applications that set rents below federally-determined levels.[11] By observing official scores, I avoid

[11]The federal maximum on monthly rent works out to 1.25 percent of the local Area Median Income (AMI), although property-level "income averaging" introduced at the end of my sample allows some units to charge higher rents.

score computations in most of the analysis. Typically, the scoring rules are additively separable in characteristics x_i. This feature makes it easy to compute a counterfactual score $q_i' = q(x_i')$ for a deviation $x_i' - x_i$ from an actual application x_i and score q_i.

Most states hold annual or semiannual application rounds. Rounds are open to a variety of entrants: for-profit developers, non-profit developers, and public housing agencies. Application appears costly, and developers must show their "readiness to proceed" if funded. Satisfying this criterion usually entails "site control"—that is, ownership or a purchase option contract contingent on winning. Applications also often include construction plans, zoning approvals, pro-forma income statements, financing contracts, and letters of support from local politicians.[12]

Awards. Upon winning, developers transfer the credit to taxable equity limited-partner investors who finance construction (Desai et al., 2010). Units are income- and rent-restricted for at least 30 years. If an applicant loses, they may reapply in subsequent cycles. There is also a landscape of other subsidies they may pursue, including a less-generous credit that is not awarded competitively in most states. Losing applicants are also free to not develop or to develop for alternative uses.

Winning applicants receive nonrefundable tax credits with a face value of 70 percent of their project's basis in present value, applied to tax liabilities over ten years. In some areas, winners qualify for a 30-percent "boost" to their basis. Thus, when boosted, each basis dollar translates to \$0.91 in credits. Areas can qualify for the boost in two ways: as a Qualified Census Tract (QCT) or as a Difficult Development Area (DDA). Federal regulations assign the boost based on Census data, and the regulations are complicated. A tract can be designated a QCT according to its poverty rate or its income distribution relative to the median income of the metropolitan area. Areas are designated DDAs based on a ratio of rents to median incomes. Until 2016, this designation was set by metropolitan area, after which it shifted to combinations of zipcodes and counties.[13]

2.2 Data

The next three subsections introduce the data. For further information, see Appendix B.

Applications. I built a new database of LIHTC applications covering 40 U.S. states from 2005 to 2019. The data are from the administrative records of state housing finance agencies, which I compiled, digitized, and standardized.[14] The database contains records of 22,241 LIHTC applications, both winning and losing, with a total funding request of approximately \$200 billion in 2022

[12]The detail in applications fundamentally shapes this paper's approach. On the one hand, enough is known about losers to provide a counterfactual to winners. On the other hand, self-selection into application is a major concern.

[13]Since 2009, states are authorized to assign basis boosts discretionarily beyond federal requirements. However, QAPs usually regulate the assignment of "discretionary" boosts. I intend to incorporate these rules in a future draft.

[14]To minimize the burden of my requests on agencies, I requested records produced in their review processes. Most agencies were highly cooperative. I won or settled appeals against several agencies after they denied my initial requests.

constant dollars, adjusted for inflation using the Consumer Price Index.

For each of the 453 application rounds in the data (see Appendix Figure A1), I have every application considered. Due to agencies' varying record-keeping practices, other rounds cannot be fully reconstructed and are thus excluded from the data. For each application, I have the proposed name and address of the development, the primary demographic group and income levels to be served, the unit count (below-market and market-rate), the value of tax credits requested, the identity of the applicant (their name, contact details, and nonprofit status), and whether the funds are for new construction or rehabilitation. The data further include scores and all set-asides, enabling me to simulate tax-credit assignments. I identify reapplications by linking across rounds.

Parcel Outcomes. By combining several data sources, I observe whether and what, if anything, was developed at the level of the application parcel. The sources are CoreLogic and the National Housing Preservation Database (NHPD). I manually link the applications to both datasets.

The CoreLogic data collate tax assessments from localities and therefore include variables that are consistently available in such records. In particular, the data contain the assessed values of land and improvements, year built (the end year of construction), land use (e.g., multifamily residential), floor space (in square feet), and ownership information.

The NHPD is the most extensive database of U.S. subsidized housing, covering nearly all federal programs. I use the NHPD to measure if a property is subsidized, and by which program if not the LIHTC. If an application is never funded in my data (including reapplications), is not matched to an NHPD record, but appears in CoreLogic, I assume it is not subsidized.

Neighborhood Data. I use data by Census tract and block group as outcomes and as covariates. The outcome data come from the U.S. Postal Service (USPS) and a mail-marketing firm (Data Axle, formerly Infogroup). I align both to consistent Census 2000 geographies. I use covariates from the 2000 Census, not from later years, to avoid confounding from the developments themselves.

The USPS data are counts of residential addresses by tract and quarter since 2006. The Data Axle files contain annual address-level microdata, also since 2006. I collapse the files to the block group and tract, so as to complement the USPS counts with details on demography and land use.

2.3 Win Probabilities

A key input in my analysis is an application's ex-ante win probability, taking rival applications as a random variable whose distribution but not realization is known when a developer applies. These win probabilities are used to balance winners and losers on unobservables in the quasi-experimental analysis (Section 4) and as developers' beliefs in the structural estimation (Section 8).

I estimate the win probabilities of applicants through a simulation procedure that I take from research on auctions (Hortaçsu and McAdams, 2010) and market design (Abdulkadiroğlu et al., 2017). Following the auctions literature, I assume developers make application choices independently within round, consistent with an informational assumption of independent private values.

The estimated win probability is a valid balancing score in the sense of Rosenbaum and Rubin (1983) under two assumptions. First, the grant rule $W(\cdot)$ must be correctly specified. Second, developers must know no more than their own application and the distribution of potential rivals. The propensity-score theorem then gives that any residual variation in tax-credit assignment is independent of potential outcomes.

Variation in assignment conditional on win probability, however, is not guaranteed in non-stochastic mechanisms. A necessary condition for the existence of such variation is that the grant or scoring rule is multidimensional, so that applicants with the same propensity score can differ in their characteristics and thus assignments. Set-asides create such variation in my context.

To obtain win probabilities, I first estimate the distribution $\hat{\Psi}_{it} = \hat{\Psi}(\boldsymbol{Q}_{-it}, \boldsymbol{Z}_{-it})$ of potential rivals by resampling from applications within the same round. That is, for a bootstrap replication $b = 1, \ldots, B$, I draw $N-1$ applications uniformly with replacement from each applicant's distribution of actual rivals. This yields simulated rivals $(\boldsymbol{Q}^b_{-it}, \boldsymbol{Z}^b_{-it})$ for an application i in replication b.[15] Next, I run the mechanism on the applicant and each simulated draw of rivals, which assumes developers know their own score and set-asides.[16] Running the mechanism, I obtain B simulated assignments $\hat{W}^b_{it} = W\left(q_{it}, z_{it}; \boldsymbol{Q}^b_{-it}, \boldsymbol{Z}^b_{-it}\right)$, and I calculate the estimated win probabilities by $\hat{p}_{it} = \frac{1}{B}\sum_{b=1}^{B} \hat{W}^b_{it}$.

2.4 Summary Statistics

Application Characteristics. Table 1 summarizes the data. Columns 1 and 2 show means of characteristics for winners and losers respectively. A typical application is for a development of roughly 60 units, is entirely income-restricted, and charges a monthly rent of approximately \$900 in 2019. About 45 percent of developments are intended for elderly or other non-family groups (typically disabled, homeless, or supportive housing). The competitions are dominated by "repeat players" (i.e., developers specialized in low-income housing), and 28 percent of applications involve nonprofit organizations. Relative to the national average, neighborhoods with applications are poor and densely-populated but do not have elevated rental vacancy rates.

Differences between winners and losers on observed characteristics are statistically significant but of small magnitude. Without mechanism-based win probabilities, a researcher might rely on simple comparisons of outcomes between winners and losers. However, Table 1 also suggests

[15] I implement this procedure in a way that is more computationally efficient, so as to avoid BN draws of the sample. Before each bootstrap replication, I split realized applications into K "folds." I retain the realized applications in one fold and resample among the $K-1$ others. In calculating $\hat{p}_{it}$, I only those replications in which application i is in the retained fold. This procedure only requires BK draws of the sample.

[16] Using the "self-scores" that developers submit in some states before their applications are reviewed, I find that developers are highly informed about their own scores (see Appendix B).

Table 1: Covariate Balance Between Winning and Losing Applications

	Means		Differences (SE)	
	Winners (1)	Losers (2)	Uncontrolled (3)	Controlled (4)
Panel A: Project Characteristics				
Unit Count	61.62	64.94	-3.32*** (0.55)	0.20 (0.66)
% Low-Income Units	97.5	98.1	-0.6*** (0.1)	-0.0 (0.1)
Monthly Rent	$885.77	$894.37	-8.60** (4.16)	4.44 (5.40)
% New Construction	61.9	64.6	-2.8*** (0.8)	-2.6*** (0.9)
Family	54.1	56.7	-2.6*** (0.9)	1.1 (1.0)
Elderly	31.4	33.5	-2.1*** (0.8)	-1.3 (1.0)
Other	14.5	9.8	4.7*** (0.6)	0.2 (0.7)
Win Probability	69.9	23.6	46.3*** (0.5)	·
PDV Tax Credits Per Unit	$144,900	$155,182	-10,282*** (1,056)	-2,093 (1,314)
Panel B: Applicant Characteristics				
Application Count in State	7.11	8.97	-1.87*** (0.21)	-0.47* (0.24)
Win Rate in State	45.6	37.5	8.2*** (0.5)	2.3*** (0.6)
Nonprofit	28.0	28.1	-0.1 (0.7)	-0.4 (0.8)
Panel C: Location Characteristics				
Med. Income Per Capita	$16,932	$16,826	106 (111)	-101 (141)
% Poor	20.7	20.1	0.6** (0.2)	1.1*** (0.3)
% Less than HS	27.1	27.1	0.0 (0.2)	0.4 (0.3)
% HS Graduate	30.2	30.9	-0.6*** (0.1)	-0.6*** (0.2)
% Some College	25.3	25.4	-0.1 (0.1)	-0.1 (0.2)
% College Graduate	11.5	11.0	0.5*** (0.1)	0.2 (0.2)
% More than College	5.9	5.7	0.3*** (0.1)	0.1 (0.1)
% Non-Hispanic White	59.9	60.5	-0.6 (0.5)	-1.0* (0.6)
% Non-Hispanic Black	19.9	21.5	-1.6*** (0.4)	0.5 (0.5)
% Hispanic	14.3	13.3	1.0*** (0.3)	-0.5 (0.4)
% Asian	2.2	1.8	0.5*** (0.1)	0.4*** (0.1)
Pop. Density (per sq. mi.)	3,620	3,447	173** (82)	307*** (107)
% Rentals Vacant	8.0	8.3	-0.3*** (0.1)	-0.1 (0.1)
% In QCT or DDA	42.2	44.2	-2.1*** (0.7)	1.0 (0.8)
Observations	9,022	11,968		
P-val. of Balance Test			0.000	0.000

Notes: This table reports means and differences in means between winning and losing applications. Columns 1 and 2 respectively report means of characteristics for winning and losing applications. Column 3 reports the unconditional win-minus-lose difference in means, and Column 4 reports the difference controlling nonparametrically for the win propensity score. Standard errors are clustered by tract. Observations indicate counts of winning and losing applications, but some covariates are unavailable for some applications. $^{*} = p < 0.10$, $^{**} = p < 0.05$, $^{***} = p < 0.01$.

Figure 2: Calibration and Distribution of Simulated Win Probabilities

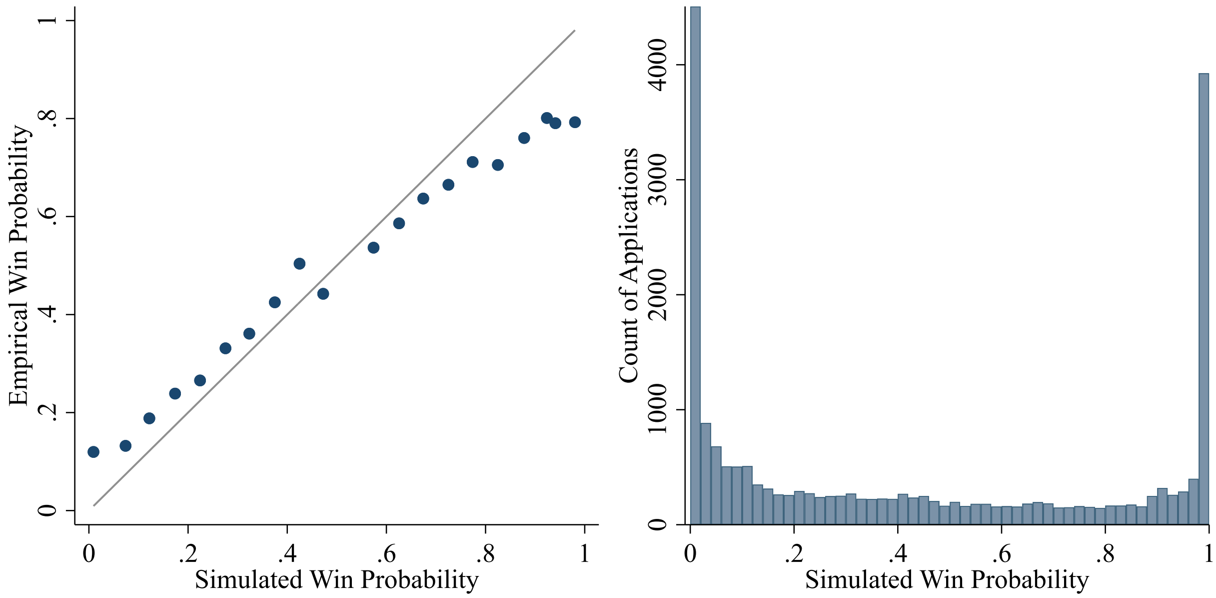

Notes: This figure plots, in the left panel, a binned scatterplot of empirical versus simulated probabilities of winning a tax credit. In the right panel, the figure plots a histogram of the simulated win probabilities. The data is split into twenty equal-interval bins in the left panel and fifty equal-interval bins in the right panel.

that such a design would be dangerous. Despite observable similarity, winners and losers differ sharply in their ex-ante win probabilities: on average, by 46 percentage points. By contrast, if one naively predicted win probabilities from a regression on all other application characteristics in this table, the analogous difference would be 7 percentage points. The gap reflects that the information contained in scores is not closely related to other observed characteristics.

Table 1 also shows that I can easily reject balance on observables between winners and losers, even conditional on the win probabilities. This likely reflects my misspecification of QAP rules, not agencies' secret discretion.[17] I provide two checks regarding these imbalances, since I use winners-versus-losers comparisons in Section 4. First, in the tract event study, I control for variables where there is clear imbalance (Appendix Figure A2). This approach, however, does not address residual imbalance on unobservables. I therefore also instrument for the actual assignment using my version of the grant rule to simulate the assignment, while controlling for the win probability (Appendix Figure A3). Both tests suggest the remaining imbalance is innocuous.

Figure 2 shows the calibration and distribution of the simulated win probabilities. In the left panel, I present a binned scatterplot of the empirical win probabilities against the simulated values, along with a 45-degree line for reference. The simulation would be perfectly calibrated if the scatter points fell exactly along the reference line. The win probabilities are well calibrated but not

[17]QAPs occasionally provide opaque descriptions of aspects of their grant rules $W(\cdot)$, and I have simplified a few very complex provisions.

Figure 3: Rent Discounts in Tax Credit Units

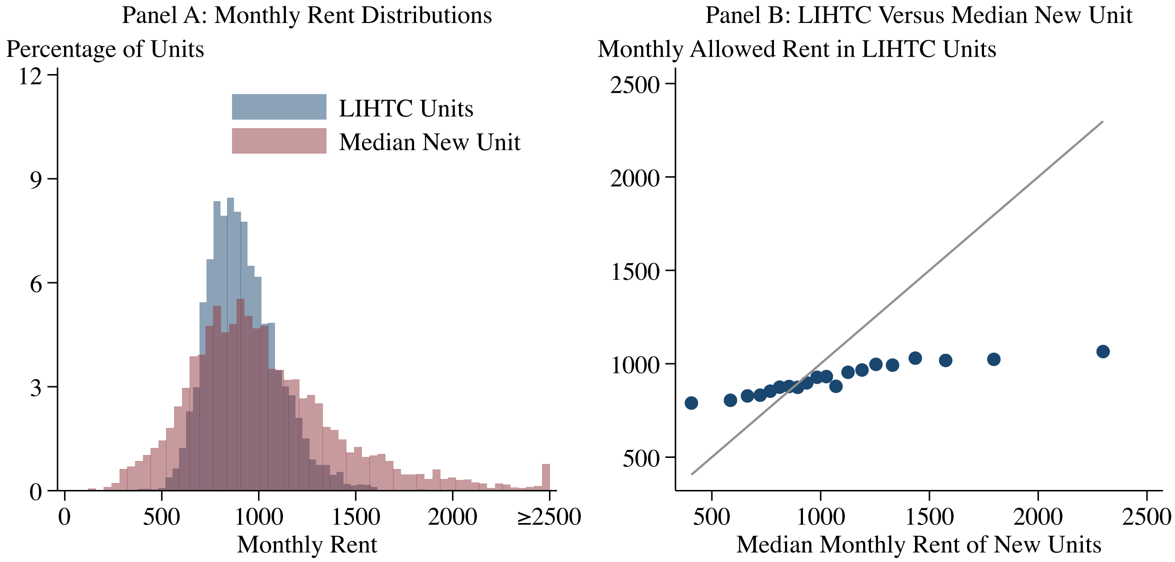

Notes: This figure compares rents in LIHTC units to the median rents of new units in the same Census block group. See the text for explanations of the panels.

perfectly so. As above, the issue appears to be difficulties in coding the QAP rules.

Rent Discounts. Figure 3 describes the LIHTC in terms of the rent discounts it provides to tenants. Panel A show that on average, the monthly rents of subsidized units are close to the rent that I estimate could be charged by the median new unit in the same Census block group and with the same number of bedrooms.[18] All values are as of the year 2019.

Panel B shows a binned scatterplot of regulated LIHTC rents versus new-unit market rents. Rent regulations bind strongly in high-rent neighborhoods, generating large discounts. However, few LIHTC developments are built in such areas. I estimate that, on average, a new market-rate unit in a neighborhood with a LIHTC application would rent for $990 per month, which is about 40 percent below the national-average rent for new units. Accounting for place—and, in particular, that low-income housing is built mostly in places with low willingness to pay for new housing—is crucial to estimate rent savings. Indeed, I find regulated rents do not bind for about half of applications, so considerable subsidy goes to developments charging local market rents.

Assuming that LIHTC units charge the minimum of the market rent and their regulated rent, the average monthly rent discount is 12 percent. In present-value terms, the LIHTC thus buys little in rent savings. At a 5-percent annual discount rate, $1 in the present value of tax credits reduces

[18] "New" means built between 2010 and 2019. New-unit rents are observed in only 17 percent of Census tracts. I therefore impute new-unit rents using a hedonic regression that incorporates local rents on older units and a sample-selection correction that accounts for selective development with respect to potential rents. Other processing steps project these rents to the block-group level and align regulated and market rents on bedrooms. See Appendix B.

Figure 4: Tenant Incomes in LIHTC Units and Non-LIHTC Counterfactual Units

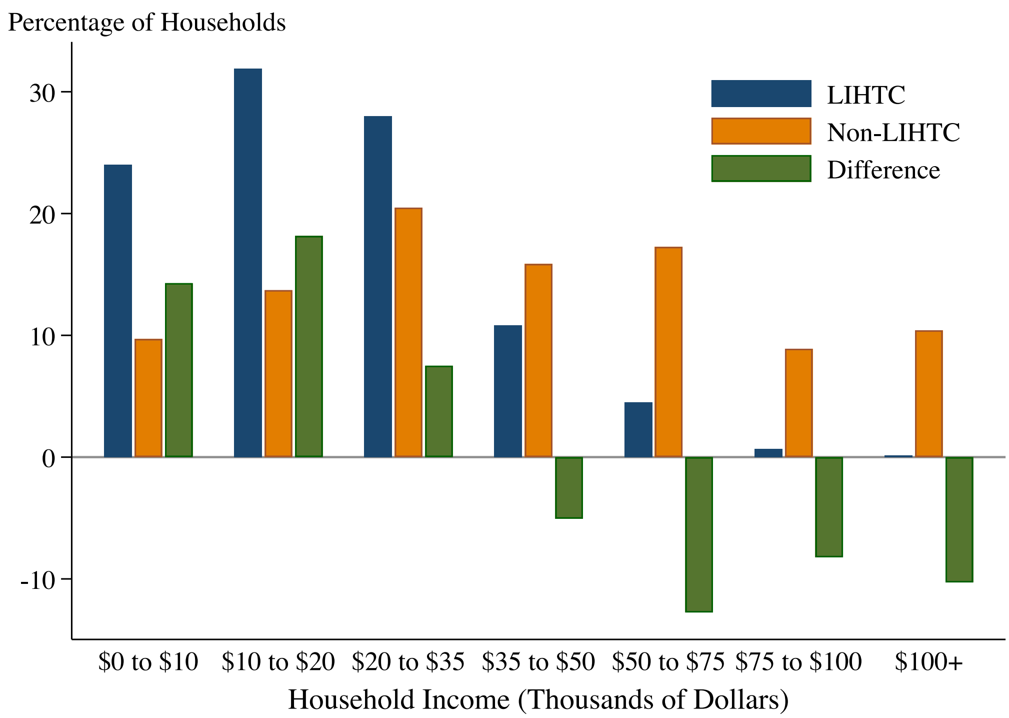

Notes: This figure compares household incomes in LIHTC units with the predicted distribution in counterfactual non-LIHTC units. The counterfactual assumes the building charges the estimated local market rent for new units.

the present value of rents by 27 cents.

Several caveats are relevant for these estimates. First, benchmarking LIHTC units to the median new rental unit nearby assumes comparability on other dimensions, ignoring differences in quality as well as any services provided by low-income units. Second, interpreting rent discounts as a measure of tenant welfare assumes tenants treat the benefit as equivalent to cash. Third, my estimates are for rent discounts when the building is new. As the building depreciates, its counterfactual market rents fall, but regulated rents do not, reducing the rent savings. Finally, some LIHTC units surely appear the new-unit comparison group, attenuating the estimated rent savings.

Tenant Composition. I also examine whether LIHTC tenants differ in income from the likely tenants of counterfactual market-rate developments. This analysis uses property-level HUD data on the household incomes of LIHTC tenants, along with ACS data on tract-level joint distributions of rent and income to predict household income from a counterfactual rent (see Appendix B).

Figure 4 shows that LIHTC tenants have lower incomes than counterfactual non-LIHTC tenants. The share of households earning less than $20,000 roughly triples, and few LIHTC households earn more than $50,000, as compared to roughly a third of counterfactual residents.[19]

[19]The fact in the raw data that drives this result is that the median LIHTC household earns much less than the median

These results raise the question of whether housing-cost effects reasonably summarize the household welfare effects of the LIHTC. If housing markets facing low-income households are competitive, this approach is sensible: Households then face no barriers other than price, so the reallocation of new housing across households is not a source of welfare gains.[20] However, the assumptions underpinning this view are hardly beyond question. For instance, Bergman et al. (forthcoming) and Christensen and Timmins (2023) find barriers for voucher recipients and minorities in accessing high-quality neighborhoods.

3 A Dynamic Model of Housing Markets

This section introduces a dynamic equilibrium model of the markets for subsidized and unsubsidized housing. Its primary aim is to jointly explain the application and building behavior of developers. The model's household side is kept simple, with the minimum needed for equilibrium effects of subsidies in the unsubsidized market. I then provide welfare and incidence measures. Appendix C discusses the model's nonparametric identification.

3.1 Setup

Choices. A developer $i \in I$ has the exclusive right to develop a land parcel. They make two profit-maximizing choices in each time period t. The first, $A_{it} \in \{0, 1\}$, is whether to submit their parcel to a grant competition. If they win, the developer must build low-income housing. If they lose the competition, they may reapply in the future. The second choice, made subsequently, is whether to build if they do not win or if they do not apply for the grant: $B_{it} \in \{0, 1\}$. Developers can defer the building decision indefinitely. Whether subsidized or not, development is an absorbing state. Figure 5 depicts the structure of choices in the model.

State Variables. The developer's vector of state variables is denoted by $\boldsymbol{s}_{it} = \{d_{it}, h_{it}, x_{it}, \xi_i, r_t^m\}$. The variable d_{it} indicates whether the parcel is developed by t. The "history" variable h_{it} indicates whether the developer has previously applied for the grant by t. The developer also has characteristics x_{it} and scalar unobservable characteristic ξ_i, which has a distribution G_ξ and is fully persistent through time. The market rent for housing is r_t^m.

Each period, the developer also draws temporary unobservable shocks $\boldsymbol{\varepsilon}_{it} = (\varepsilon_{it}^A, \varepsilon_{it}^B)$ from a distribution F_ε. The shocks are assumed to be independently and identically distributed (i.i.d.) and additively separable from payoffs. Developers' draws are private information, so that each faces

household in the same tract (see Appendix Figure A23). In Section 4, I present some direct evidence consistent with reallocation of housing to voucher households.

[20]Indeed, the literature on household valuations of in-kind transfers usually interprets this distortion in housing consumption as a source of welfare loss. Estimates reviewed in Olsen (2003), for instance, suggest tenants value housing benefits at about 80 cents on the dollar.

Figure 5: Decision Tree for the Developer

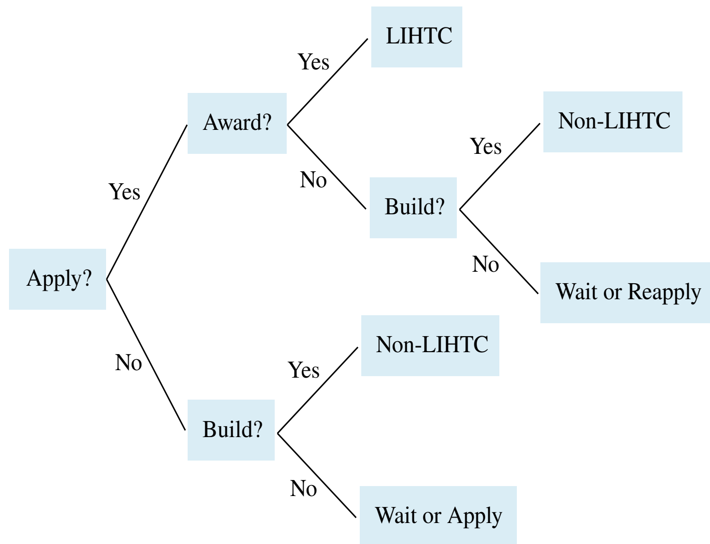

uncertainty as to which rivals will apply. They therefore perceive a probability p_{it} of winning the grant if they apply. Developers are atomistic, taking their individual win probabilities and rents as given but determining both collectively, following Hopenhayn (1992).

3.2 Application and Building Decisions

I express the developer's problem through two Bellman equations. The first, which I call the *application* value function, concerns the developer's application choice. The second, which I call the *building* value function, concerns the developer's building choice. Developers move between these value functions by deciding not to apply and by deciding not to build.

Application. The developer's application value function in the state $\boldsymbol{s}_{it}$ is given by:

$$V^A(\boldsymbol{s}_{it}, \boldsymbol{\varepsilon}_{it}) = \max_a \left\{\Pi^A(a, \boldsymbol{s}_{it}) + \varepsilon^A_{it}(a)\right\}, \tag{2}$$

where their application choice is a and $\Pi^A(a, \boldsymbol{s}_{it})$ denotes their flow payoff upon taking the action a from their state. Application choices are therefore $A_{it} = A(\boldsymbol{s}_{it}, \boldsymbol{\varepsilon}_{it}) = \arg\max_a \left\{\Pi^A(a, \boldsymbol{s}_{it}) + \varepsilon^A_{it}(a)\right\}$.

The developer's expected payoffs are

$$\Pi^A(a, \boldsymbol{s}_{it}) = \begin{cases} p_{it}\pi_1(\boldsymbol{s}_{it}) + (1 - p_{it})\beta \mathrm{E}[V^A(\boldsymbol{s}_{it+1}, \boldsymbol{\varepsilon}_{it+1}) \mid \boldsymbol{s}_{it}] - \kappa(\boldsymbol{s}_{it}) & \text{if } a = 1 \\ \mathrm{E}[V^B(\boldsymbol{s}_{it}, \boldsymbol{\varepsilon}_{it}) \mid \boldsymbol{s}_{it}] & \text{if } a = 0. \end{cases} \tag{3}$$

In choosing to apply, the developer compares the expected value of applying to a reservation value, their building value function. If they apply, they receive an average of the win payoff $\pi_1(\boldsymbol{s}_{it})$

and the discounted continuation value $\beta \mathrm{E}[V^A(\boldsymbol{s}_{it+1}, \boldsymbol{\varepsilon}_{it+1}) \mid \boldsymbol{s}_{it}]$, weighted by the win probability, less their entry cost $\kappa(\boldsymbol{s}_{it})$. They pay entry costs whether they win or lose. Upon winning, the developer receives no further payoffs. If they apply, the state is updated such that $h_{it} = 1$; the update is to $h_{it+1} = d_{it+1} = 1$ if they win.

If the developer does not apply, their payoff is the building value function $\mathrm{E}[V^B(\boldsymbol{s}_{it}, \boldsymbol{\varepsilon}_{it}) \mid \boldsymbol{s}_{it}]$, taking the ex-ante expectation over ε^B_{it}. By not applying, the continuation value $\mathrm{E}[V^A(\boldsymbol{s}_{it+1}, \boldsymbol{\varepsilon}_{it+1}) \mid a = 0, \boldsymbol{s}_{it}]$ is also set to zero for all states $\boldsymbol{s}_{it}$. A decision not to apply thus moves the developer over to the building value function, but it does not rule out applications in future periods. Not applying preserves the developer's history and development state variables: $h_{it+1} = h_{it}$ and $d_{it+1} = d_{it}$.

Building. The building value function is

$$V^B(\boldsymbol{s}_{it}, \boldsymbol{\varepsilon}_{it}) = \max_b \left\{ \Pi^B(b, \boldsymbol{s}_{it}) + \varepsilon^B_{it}(b) \right\}, \tag{4}$$

where b is the building choice. These are given by $B_{it} = B(\boldsymbol{s}_{it}, \boldsymbol{\varepsilon}_{it}) = \arg\max_b \left\{ \Pi^B(b, \boldsymbol{s}_{it}) + \varepsilon^B_{it}(b) \right\}$. The expected payoffs are

$$\Pi^B(b, \boldsymbol{s}_{it}) = \begin{cases} \pi_0(\boldsymbol{s}_{it}) & \text{if } b = 1 \\ \beta \mathrm{E}[V^A(\boldsymbol{s}_{it+1}, \boldsymbol{\varepsilon}_{it+1}) \mid \boldsymbol{s}_{it}] & \text{if } b = 0. \end{cases} \tag{5}$$

The developer's building choice thus involves a comparison of their outside option and their application value function. In particular, if the developer chooses to build, their payoff is $\pi_0(\boldsymbol{s}_{it})$. They receive no further payoffs, and their development state variable is updated to $d_{it+1} = 1$. If they do not build, their payoff is the (ex-ante) application value function $\beta \mathrm{E}[V^A(\boldsymbol{s}_{it+1}, \boldsymbol{\varepsilon}_{it+1}) \mid \boldsymbol{s}_{it}]$. They also move back to the application value function and preserve history and development states.

3.3 Grant Mechanism

Applications are scored according to $q_{it} = q(x_{it}, \xi_i, \varepsilon_{it})$. The dependence on the unobservables induces both selection in awards and self-selection at application, as potential applicants anticipate their win probabilities. If a developer applies, their grant assignment is given by Equation 1.

To compute their win probability, a developer evaluates the conditional expectation of their assignment over the distribution of potential rival applicants Ψ_{it}:

$$p_{it} = \int W(q_{it}, x_{it}; \boldsymbol{Q}_{-it}, \boldsymbol{X}_{-it})\, d\Psi_{it}(\boldsymbol{Q}_{-it}, \boldsymbol{X}_{-it}). \tag{6}$$

The dependence on both the score q_{it} and characteristics x_{it} means the mechanism need not have a single-index representation in which the score fully determines winners and losers. The potential-rival distribution is given by the product of the distribution φ of potential-applicant characteristics and the conditional probability of application for such a potential applicant: $\Psi_{it} = \prod_{j \in I \setminus i} \Pr(A_{jt} \mid \boldsymbol{s}_{jt}) \varphi(\boldsymbol{s}_{jt})$, where j indexes potential applicants. As the shocks $\boldsymbol{\varepsilon}_{it}$ are private information, developers cannot determine their rivals' application decisions before making their own, and so $\Pr(A_{jt} \mid \boldsymbol{s}_{jt})$ takes values betwen zero and one.

3.4 Housing Demand

A representative household rents the entire housing stock each period. They view subsidized and unsubsidized housing as perfect substitutes, and they have constant elasticity of substitution (CES) preferences over housing H_t and a non-housing numeraire good C_t:

$$\max_{H_t} u(H_t, C_t) = [H_t^{\frac{\rho-1}{\rho}} + (\phi C_t)^{\frac{\rho-1}{\rho}}]^{\frac{\rho}{\rho-1}} \quad \text{s.t.} \quad r_t H_t + C_t = Y_t, \tag{7}$$

where ρ is the elasticity of substitution and ϕ is a preference weight on the numeraire. The household has an exogenous nominal income of Y_t. These preferences give rise to a price index

$$P(r_t) = [r_t^{1-\rho} + \phi^{\rho-1}]^{1/(1-\rho)},$$

where $r_t = [\lambda_t(1-\delta) + (1-\lambda)]r_t^m$ is the average rent. There is an equilibrium share λ_t of subsidized units, renting at a discount of δ, set by the government, relative to the market rent.

I measure household welfare by the present value of its indirect utility, $V^H = \sum_t \beta^t (Y_t/P(r_t))$. For the counterfactuals, I assume that changes in rents enter developer primitives as present values: $\Delta\pi_0(\boldsymbol{s}_{it}) = \sum_t \beta^t \Delta r_t^m H(\boldsymbol{s}_{it})$ and $\Delta\pi_1(\boldsymbol{s}_{it}) = (1-\delta) \sum_t \beta^t \Delta r_t^m H(\boldsymbol{s}_{it})$, where $H(\boldsymbol{s}_{it})$ is the developer's potential number of units. This adjustment would result from treating the outside option as a difference of construction cost and the present value of rent: $\pi_0(\boldsymbol{s}_{it}) = -C(\boldsymbol{s}_{it}) + \sum_t \beta^t r_t^m H(\boldsymbol{s}_{it})$.

3.5 Equilibrium and Welfare

Equilibrium. Given primitives $\{\pi_0, \pi_1, \kappa, F_\varepsilon, G_\xi, \varphi\}$, parameters $\{\beta, \phi, \rho, Y_t\}$, and an initial state, an equilibrium is defined by a sequence of endogenous developer quantities $\{A_{it}, B_{it}\}$, household quantities $\{H_t, C_t\}$, and probabilities and prices $\{p_{it}, r_t^m\}$ such that

1. *Developers and households optimize*: Application and development choices A_{it} and B_{it} maximize Equations 2 and 4, and housing consumption H_t maximizes Equation 7.

2. *Housing market clears*: Market rents $r_t^m = r^m$ balance supply and demand in expectation.

3. *Rational expectations*: Win probabilities p_{it} satisfy Equation 6.

Welfare. Social welfare is the sum of developer and household welfare:

$$\mathcal{W} = \mathrm{E}[V^A(\boldsymbol{s}_{it}, \boldsymbol{\varepsilon}_{it})] + V^H.$$

Let Δ's indicate differences with respect to some policy change, and let the present value of entry costs be $K = \mathrm{E}[\sum_t \beta^t A_{it} \kappa(\boldsymbol{s}_{it})]$. In the incidence analysis, I refer to developer and household shares, defined respectively as $\Delta \mathrm{E}[V^A(\boldsymbol{s}_{it}, \boldsymbol{\varepsilon}_{it})]/(\Delta \mathcal{W} + \Delta K)$ and $\Delta V^H/(\Delta \mathcal{W} + \Delta K)$. I also refer to the share lost to entry costs, $\Delta K/(\Delta \mathcal{W} + \Delta K)$. The denominator of the incidence shares is the sum of these money-metric changes in welfare and application costs, $\Delta \mathcal{W} + \Delta K$.[21]

4 Causal Effects of Tax Credit Awards

This section estimates the causal effects of winning the LIHTC on development outcomes. It is the first of three empirical analyses that I aim to match in estimating the model in Section 3. By estimating what winning developers would have done had they counterfactually lost, such an analysis is informative about their outside options to the subsidy.

4.1 Setup

The probability that, from a state $\boldsymbol{s}_{it}$, a developer builds without the LIHTC is

$$\log \frac{\Pr(B_{it} = 1 \mid \boldsymbol{s}_{it})}{1 - \Pr(B_{it} = 1 \mid \boldsymbol{s}_{it})} = \frac{1}{\sigma_b} \left[\pi_0(\boldsymbol{s}_{it}) - \beta \mathrm{E}[V^A(\boldsymbol{s}_{it+1}, \boldsymbol{\varepsilon}_{it+1}) \mid \boldsymbol{s}_{it}] \right], \tag{8}$$

where I suppress the probability's condition of having applied and lost, $A_{it} = 1$ and $W_{it} = 0$.[22] All else equal, a developer with a strong outside option (large $\pi_0(\boldsymbol{s}_{it})$) is likelier to develop if they lose. However, the build probability does not directly reveal outside options, in that it also depends on the value of waiting to apply or build later ($\beta \mathrm{E}[V^A(\boldsymbol{s}_{it+1}, \boldsymbol{\varepsilon}_{it+1}) \mid \boldsymbol{s}_{it}]$). The identification would be one-to-one in a static model, where there is no value of waiting.

By conditioning on the unobservable part (ξ_i) of the state $\boldsymbol{s}_{it}$, Equation 8 also warns about naive "winners-versus-losers" estimators of grant impacts that omit a win-probability control. For instance, if applications viewed by the government as strong also tend to have strong outside options, then naive estimates will overstate true causal effects and understate winners' outside options.

I therefore estimate winners-versus-losers comparisons that compare applications with similar win probabilities. Depending on context explained below, I estimate various event-study specifica-

[21] This sum differs from the fiscal cost. In the model, developers may value a subsidy dollar at something other than a dollar, due to the additional costs (e.g., construction) and benefits (e.g., other subsidies) that come with the LIHTC.

[22] This equation also assumes that ε^B_{it} is distributed Type I extreme value with dispersion σ_b, as in Section 8.

tions, but all comparisons fundamentally take form

$$Y_i = \beta(\hat{p}_i)W_i + f(\hat{p}_i) + u_i, \tag{9}$$

where Y_i is a development outcome for application i, W_i is an indicator for i's tax-credit assignment and $\hat{p}_i$ is its win probability. To keep $f(\cdot)$ flexible, I use a cubic basis spline with four knots spaced evenly through the distribution of win probabilities. The $\beta(\hat{p}_i)$ notation allows for heterogeneous effects by win probability.

Let $\hat{\beta}$ denote the estimator in Equation 9 imposing constant win effects. Under conditions discussed in Section 2.3, the win probability fully controls for the government's information set with respect to application i. It thus balances winners and losers on ξ_i, with $E[u_i W_i] = 0$. By consequence, $\hat{\beta}$ is a valid estimator of a weighted-average causal effect of awards on winners.

4.2 Parcel Effects

Figure 6 shows the results of a parcel-level event study of development outcomes.[23] The left half of Panel A shows effects of winning and losing on the annual probability of development, relative to earlier and later applicants. The right half of Panel A shows win-versus-lose estimates consistent with Equation 9. In both, the outcome is an indicator for whether the parcel is recorded as having construction or rehabilitation in that year.

Taken together, the two halves of Panel A tell a clear story. When developers win, construction follows almost immediately. When developers lose, construction typically still happens, albeit at a lag of several years and not substituting one-for-one with the LIHTC. I estimate that, at a horizon of ten years after the LIHTC round, about 75 percent of winning parcels would have been counterfactually developed had they lost. I compute this "displacement rate" by summing the lose and win coefficients from Panel A and dividing the former by the latter.[24]

What gets built? I answer this in Panel B. Not surprisingly, winners build LIHTC units. There is also a small effect on unsubsidized development on parcels I associate with winning applications. When developers lose, some reapply, succeed in their reapplication, and therefore end up with LIHTC developments. Around half of development among losers is due to reapplication. However,

[23] Some applications concern multiple parcels. When construction years differ across parcels, I code the outcome for the largest parcel by floor space. Weighting results by floor space yields essentially identical results (see Appendix Figure A4). There is some limited evidence for heterogeneous effects by win probability, with greater displacement among applications that are likelier to win (see Appendix Figure A5).

[24] To consider the importance of pulling development forward in time, I first adjust the displacement rate for discounting at five percent per year. I obtain a rate of 71 percent, as compared to 75 percent without discounting. I also calculate the number of additional "building–years" within the 10-year horizon (linear discounting). The result is 2.1 building–years, implying a displacement rate in these terms of 67 percent. Incorporating time preference thus modestly reduces estimates of displacement.

Figure 6: Parcel Effects of Tax Credit Awards

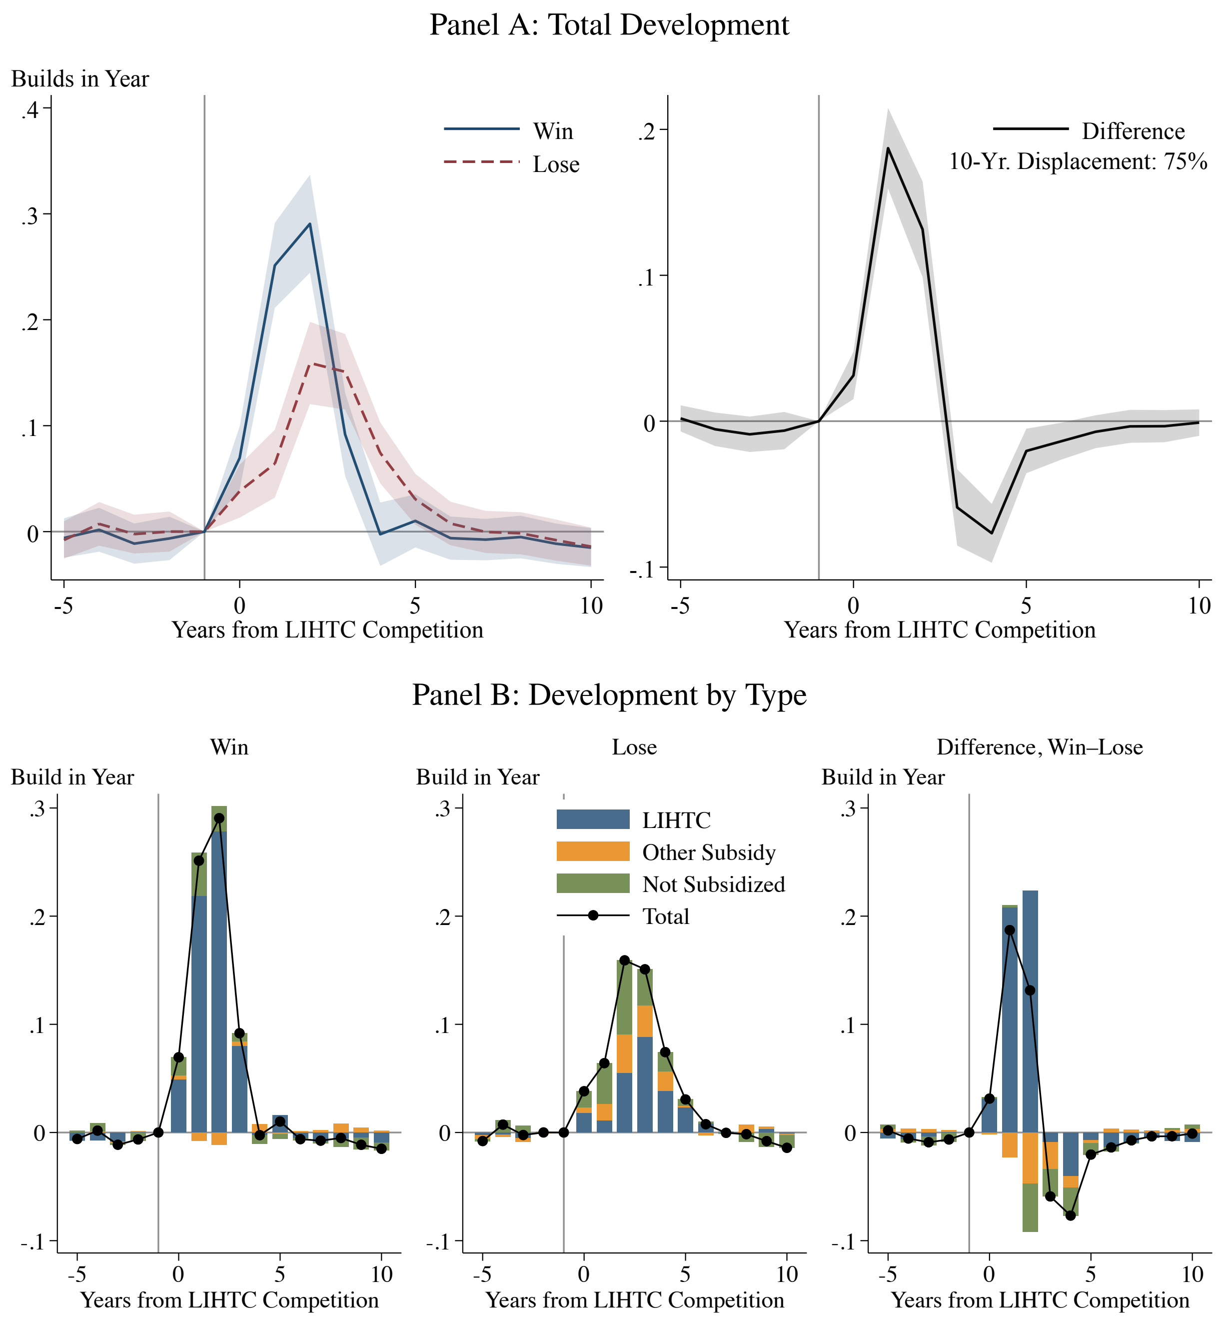

Notes: Panel A plots win and lose effects (on left) and the win–lose difference (on right), of tax credit awards on the property construction is recorded as completed in a given year. Panel B decomposes these development impacts into developments funded by the LIHTC, those funded by another subsidy, and those not subsidized. In both, the specification is an event-study analog of Equation 9 on applicant parcels: $B_{it} = \alpha_i + \alpha_t + \sum_k [\beta_k \text{Win}_{it} + \gamma_k \text{Lose}_{it} + f_k(\hat{p}_{it})] + e_{it}$, for k years from the competition. Standard errors in Panel A are clustered by the consolidated parcel.

many losers also find other ways to develop the property. About 21 percent of developments among losers secure a project-based subsidy other than the LIHTC, while the remaining 29 percent appears not to be subsidized. Overall, reapplication, substitution across subsidies, and substitution to unsubsidized development are all relevant contributors to the high displacement rate.

The rightmost figure of Panel B provides the win–lose difference by type of development. Winning the LIHTC pulls forward development in time and changes its nature from one subsidy to another or from unsubsidized to subsidized. It largely does not induce development where it otherwise would not happen within a few years. These results generally suggest applicants have strong outside options to the subsidy. On the other hand, the key role of reapplication makes a clear case for the model, as such behavior otherwise would confound attempts to estimate the value of the non-LIHTC outside option.

4.3 Neighborhood Effects

Figure 7 shows the results of event-study comparisons of neighborhoods with winning and losing applications to similar non-applicant tracts. All specifications control for win probabilities in applicant tracts and for pre-award neighborhood observables.[25]

Panel A estimates quarterly impacts on a tract's total occupied housing stock. It confirms the parcel-level findings. Following a LIHTC win, there is a tract-level development "boom," relative to ex-ante similar non-applicant tracts, which occurs sharply with the timing of the LIHTC round. The increase in its occupied housing stock corresponds to 150 additional households, which is about three times the average number of proposed units.[26]

On its own, the comparison of winners to ex-ante similar non-applicants would suggest that awards have large causal effects on development. Yet the comparison group of tracts with losing applications suggests it is not so: The losers also "boom." Indeed, at standard levels of statistical significance, I cannot reject zero net impact of winning the LIHTC on a tract's occupied housing stock (see Appendix Figure A6). The null is sufficiently precise as to rule out effects larger than 30 households, or half of the average size of a single application.

Panel B estimates annual impacts on block groups' household counts, splitting up each tract's block groups into those with LIHTC applications and those without. Nearly all of the tract-level development "boom" occurs within the block group of the application, whereas little occurs in adjacent block groups.[27] These findings help to establish that we are observing the same economic

[25]The neighborhood controls are those listed in Panel C of Table 1. I take the logs of population density and household income, and I also include the tract's cumulative population growth rate from 1990 through 2000. In a few tracts with multiple applications in the same year, I compute the probability any application wins in the tract, assuming independence across applications.

[26]Some of the excess increase in the housing stock appears to reflect follow-on applications. It may also be attributable in part to "crowded-in" developments that are unaffiliated with the applicant.

[27]Block groups are small geographic units. In the 2000 U.S. Census, there were 211,267 block groups, or 3.2 block groups per tract, each containing about 500 households on average.

Figure 7: Neighborhood Effects of Tax Credit Awards

Panel A: Tract (USPS)

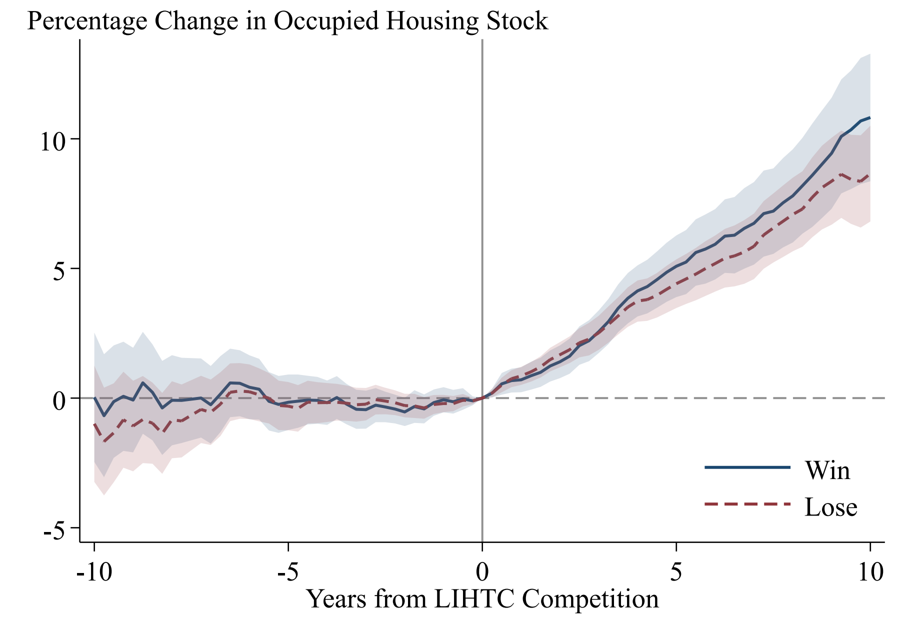

Panel B: Block Group (Data Axle)

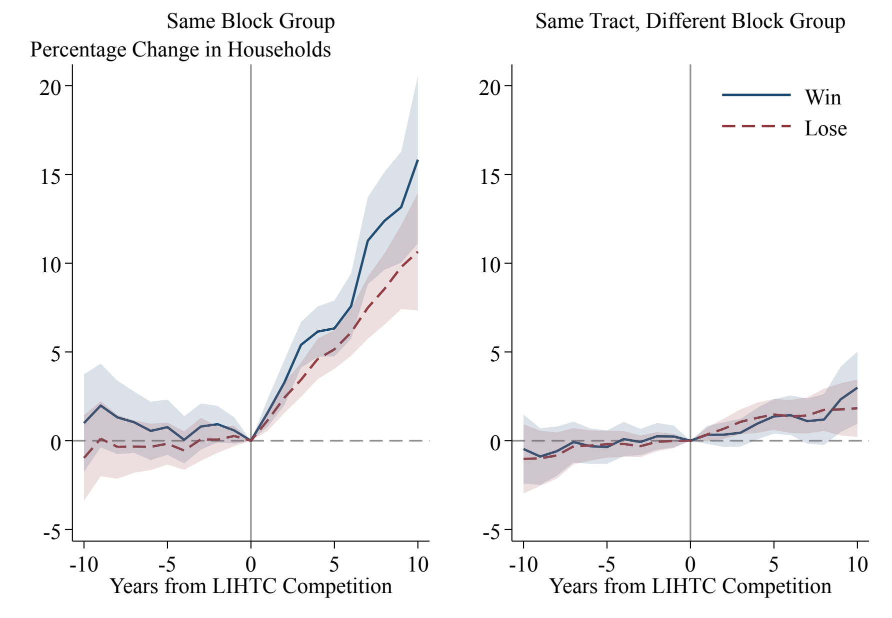

Notes: Panel A plots win and lose effects on the occupied housing stock (change in percentage points). Panel B plots win and lose effects effects on the change (in percentage points) in the household count in the same Census block group as the LIHTC application, or in the same tract but other block groups. In both, the specification is an event-study analog of Equation 9 that includes non-applicant tracts: $\Delta \log \mathrm{E}[Y_{i,t+k}] = \alpha_k \mathrm{Win}_{it} + \beta_k \mathrm{Lose}_{it} + \boldsymbol{X}_{it}\boldsymbol{\gamma}_k$, for k quarters or years from the award. The specification is estimated via Poisson regression, and it includes controls X_{it} for the win probability, pre-award tract characteristics, and the baseline level of the outcome. Winning and losing are both defined as a neighborhood's first event in sample. In both panels, the color bands depict pointwise 95-percent confidence intervals, and standard errors are clustered at the tract level.

phenomenon, of parcel-level displacement, in the parcel- and tract-level results.

Additional impacts on local demographics and land use suggest that LIHTC tenants are poorer than the likely residents of counterfactual non-LIHTC developments, though the buildings themselves are observably similar, as Section 2 anticipates. In HUD administrative data, I find the share of voucher-recipient households rises more in winning tracts than in losing tracts (see Appendix Figure A8). Around one in four tenants in LIHTC developments also receives vouchers, and so these results suggest the LIHTC can expand access to neighborhoods beyond those the voucher holders could otherwise reach.[28] The changes in local income distributions in winning and losing tracts are more suggestive, perhaps because the income data are noisy (see Appendix Figure A9). On land use, I find both winning and losing tracts see strong growth in the stocks of multi-family housing and rental housing (Appendix Figure A7). Appendix B contains robustness checks.

5 Application Responses to Subsidy Generosity

This section estimates application responses to variation in the LIHTC's value. From the perspective of the model, developer responses to subsidy generosity are informative about the entry cost $\kappa(\boldsymbol{s}_{it})$. Holding fixed outside options and win values, developers are likelier to apply when entry costs are lower, and they are more responsive to changes in win values when entry costs are less dispersed. To see this, consider the developer's application probability:

$$\log \frac{\Pr(A_{it} \mid \boldsymbol{s}_{it})}{1 - \Pr(A_{it} \mid \boldsymbol{s}_{it})} = \frac{1}{\sigma_a} \left[p_{it}\pi_1(\boldsymbol{s}_{it}) + (1 - p_{it})\beta \mathrm{E}[V^A(\boldsymbol{s}_{it+1}, \boldsymbol{\varepsilon}_{it+1}) \mid \boldsymbol{s}_{it}] - \kappa(\boldsymbol{s}_{it}) - \mathrm{E}[V^B(\boldsymbol{s}_{it}, \boldsymbol{\varepsilon}_{it}) \mid \boldsymbol{s}_{it}] \right],$$

where I impose the assumption (as I will in Section 7) that ε_{it}^A is distributed Type I extreme value with dispersion σ_a. This equation implies that, when the value of winning π_1 rises, developers become more likely to apply. As before, however, one cannot directly "read off" $\kappa(\boldsymbol{s}_{it})$ from these behavioral responses. In the dynamic setting, this policy variation also affects the value functions V^A and V^B, whereas in a static setting, it would directly identify $\kappa(\boldsymbol{s}_{it})$.

To estimate application supply responses, I use the "basis boost" introduced in Section 2, which raises the rate at which construction expenses reduce tax liabilities. I take two empirical approaches to this variation: an event study around entry and exit from the boost, and an RDD at the boost threshold. In structural estimation (Section 8), I draw on both approaches for targeted moments.

[28]This analysis weighs against one critique of Figure 4, which is that LIHTC units could create building-level sorting without neighborhood-level changes. LIHTC advocates have previously argued for its complementarity with vouchers. For example, one writes "it's crucial that the LIHTC property exists to enable those lucky enough to have a voucher find somewhere to use it" (https://www.novoco.com/notes-from-novogradac/hud-lihtc-tenant-report-highlights-47-lihtc-residents-earn-or-below-30-ami).

5.1 Event Study

Approach. Areas are assigned to the boost according to the tract-level poverty rate and median income (see Section 2). Changes in boost thus reflect, in principle, a combination of sampling variation and genuine changes in local economic conditions. As changes in conditions may independently affect the desirability of local LIHTC investment, the credibility of the event-study approach depends upon whether sampling variation or true economic variation predominates.

Simulations suggest the identifying variation in the boost is, fortunately, almost entirely due to sampling. To arrive at this conclusion, I reassign the boost using the actual assignment rule but drawing two sets of new values for its assignment variables. I draw the values using Census estimates of tract-level standard errors (see Appendix B for details). In 2019 data, for instance, around 3.9 percent of tracts are differently assigned between two simulation runs. By comparison, 8.0 percent of tracts changed their actual designation from 2018 to 2019, as I show in the left panel of Figure 8. Sampling variation is so large because tract-level estimates from national surveys are noisy. Furthermore, my event-study approach exploits variation in the timing of switches, compounding the importance of sampling variation.[29]

I estimate impacts of the boost by the following generalized event-study specification:

$$y_{it} = \alpha_i + \alpha_t + \sum_s \beta_s \Delta\text{Boost}_{i,t+s} + \boldsymbol{X}_{it}\boldsymbol{\gamma}_t + \varepsilon_{it}, \tag{10}$$

where tracts are indexed by i and time t is in years. The terms α_i and α_t are thus tract and year fixed effects, and $\boldsymbol{X}_{it}$ is a vector of controls with potentially time-varying coefficients $\boldsymbol{\gamma}_t$. Standard errors are clustered by tract. The specification is estimated via a Poisson regression.

I note several implementation details. The event variable, ΔBoost, can take three values: 0 (no year-to-year change in boost), 1 (gains boost), or -1 (loses boost). This specification thus assumes symmetric effects of entry and exit from boost, although I relax this assumption by allowing for differential effects. I include leads of the event of up to eight years ahead and for lags up to six years after. I bin any events occurring beyond these endpoints and thereby impose constant effects in these ranges. I exclude always-boosted tracts from the sample, so that never-treated tracts form the only pure control group in Equation 10. Due to a timing convention that LIHTC funds notionally for year t are sometimes committed in the fourth quarter of $t-1$ (and are therefore recorded in my data as applications in $t-1$), the base event-time period in this event study is two years prior to the year a tract gains or loses its boost.

[29] In related work, Freedman and Owens (2011) use panel variation in the boosted share of a county's Census tracts as an instrument in a two-way fixed-effects model to study the effects of LIHTC on crime. In other policy contexts, Feiveson (2015), Suárez-Serrato and Wingender (2016), and Chodorow-Reich et al. (2019) all exploit measurement error and nonlinear transformations of data embedded in policy rules as sources of variation.

Figure 8: Application Supply Responses to Changes in Eligibility for Basis Boost

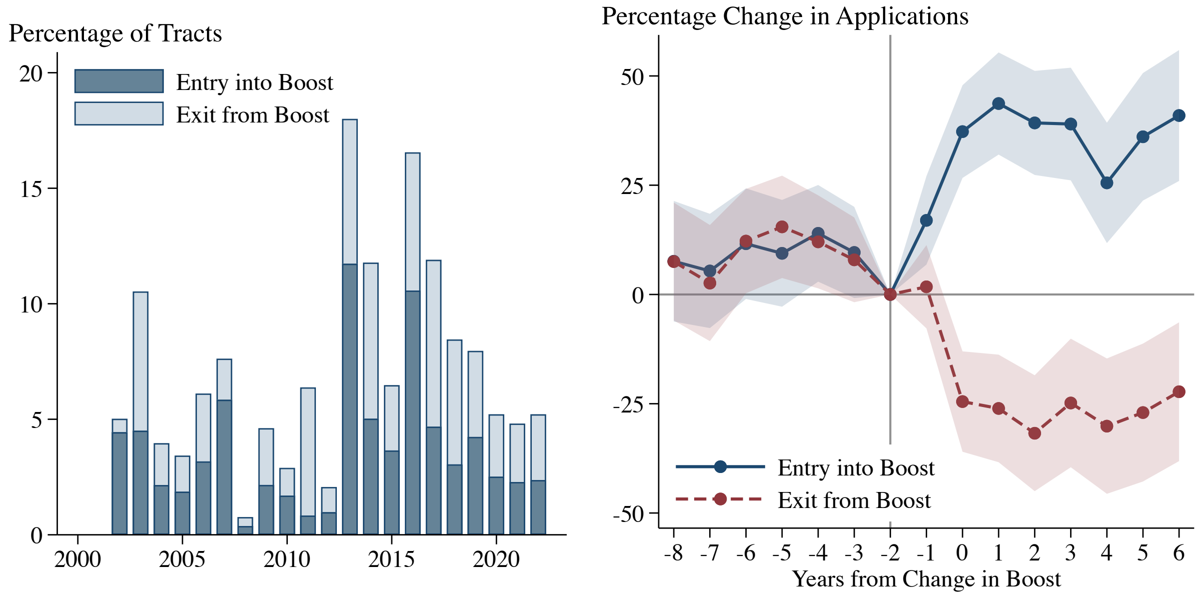

Notes: This figure plots, in the left panel, the share of Census tracts which enter or exit boost in years from 2001 to 2022. The right panel plots event-study coefficients from Equation 10 which allow for heterogeneous effects of entry and exit from boost. I include state–year fixed effects as the controls. The bands show pointwise 95-percent confidence intervals, with standard errors clustered by tract.

Results. The right panel of Figure 8 shows estimates of Equation 10, splitting effects into entry- and exit-specific estimates. The count of applications from a tract changes by approximately 30 percent in the years immediately after a change in boost designation, with symmetric effects of entry and exit. Pre-period coefficients appear steady before the change in boost. These results imply the application elasticity with respect to the net-of-tax price is roughly 0.25.

I have also examined whether contemporaneous local shocks confound the event study (see Appendix Figure A13). As a test, I predict application volume from the running variables among the non-boosted tracts and then use these predicted volumes as the event-study outcome. I find that year-to-year variation in the running variables is essentially unrelated to application volumes in this sub-sample, consistent with noise. The expected change in applications around boost entry and exit is thus essentially zero. Appendix B presents additional robustness checks.

5.2 Fuzzy Regression Discontinuity Design

Approach. The RDD directly exploits the cutoff rules in the boost designation. Some tracts are boosted despite being on the "unboosted" side of the cutoff and vice versa, due to other rules, making the discontinuity slightly fuzzy (see Appendix B). Relative to the event study, the main strengths of this strategy is that I can investigate changes in application characteristics, as well as its perhaps-finer control for local economic conditions.

I define the running variable as a tract's distance from the cutoff with respect to its rank in the running-variable distribution for its metropolitan area. Thus, a running variable value of -0.1 implies a tract is 10 percentiles away in its metro-specific distribution from being designed for the boost. This approach consolidates cutoffs in the assignment rule. The median metro area contains 64 tracts, so comparisons within a bandwidth of 0.1 of the running variable represent comparisons of roughly the six tracts nearest to either side of the cutoff in a typical metro area.

Both Baum-Snow and Marion (2009) and Davis et al. (2019) use the same discontinuity to identify effects of LIHTC projects on local areas. I make several adjustments to the definition of the running variable that strengthen the sharpness of the first stage (see Appendix B). These adjustments incorporate steps in the boost assignment rule omitted in prior work.

I model the effect of being boosted as

$$Y_{it} = \beta \text{Boost}_{it} + f_2(c_{it}, \delta) + u_{it},$$

where $f_2(c_{it}, \delta)$ is a locally-linear specification in the running variable c_{it} estimated with a bandwidth δ. The coefficient β captures the effect on LIHTC outcomes of a tract just barely qualifying for the boost as a Qualified Census Tract (QCT). The "first stage" in the fuzzy RD is

$$\text{Boost}_{it} = \gamma \text{QCT}_{it} + f_1(c_{it}, \delta) + v_{it},$$

where $f_1(c_{it}, \delta)$ is also locally linear in the running variable and QCT_{it} indicates QCT status.

Results. Figure 9 presents plots of the treatment and main outcome variable, the annual count of applications per 100,000 households in the tract, around the cutoff. Table 2 reports accompanying estimates for a broader array of application-supply outcomes. Overall, I find an increase in application volumes of 90 percent at the threshold, exceeding the event-study impacts.

The left panel of Figure 9 shows treatment assignment around the threshold. The probability a tract is a QCT jumps from approximately 20 percent to 80 percent at the threshold, as shown by the hollow blue dots. Some tracts are boosted for other reasons, and so the solid blue dots show that approximately 30 percent of tracts are boosted just below the QCT threshold. About 80 percent of tracts are boosted just above the QCT threshold.

The right panel of Figure 9 visualizes the application response to the boost. There are an additional 1.6 applications per 100,000 households in barely-boosted tracts, relative to a base level of 3.4 applications in barely-unboosted tracts, yielding an IV-estimated increase of 90 percent. The general upward slope in the outcome variable in Figure 9 reflects that tracts are poorer as the running variable rises. Table 2 shows similar increases in other measures of application supply.

Appendix A includes supporting analyses. Appendix Figure A18 plots the discontinuity for all

Figure 9: Application Volume Around the Qualified Census Tract Threshold

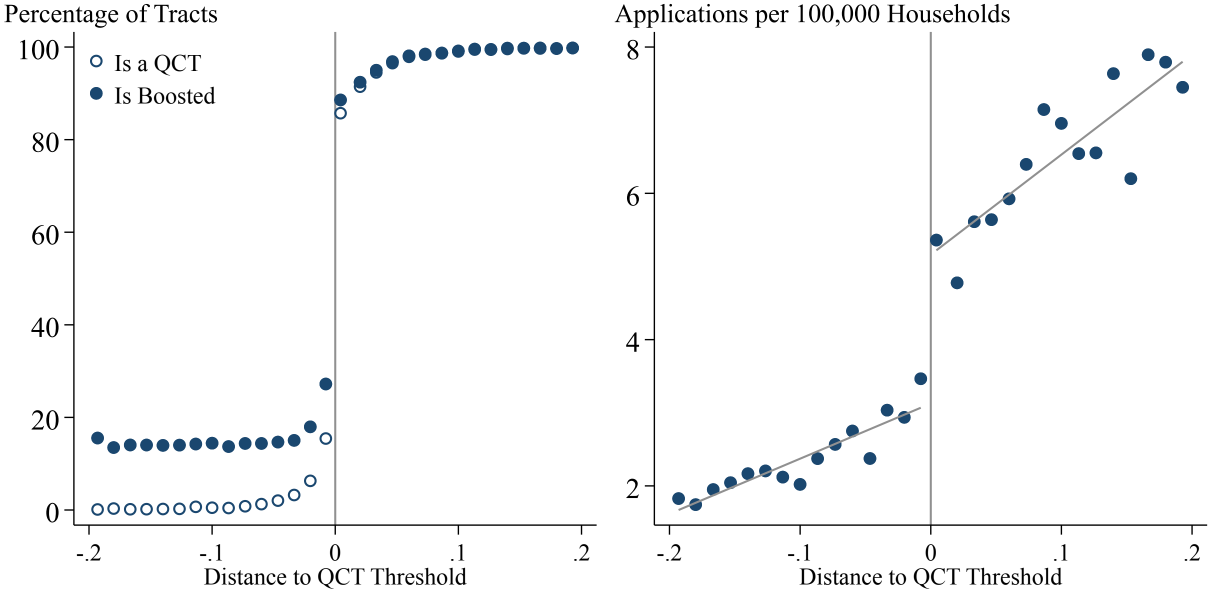

Notes: In the left panel, this figure shows that the probabilities that a tract is designated a Qualified Census Tract (QCT) and is boosted both rise discontinuously in its distance to the QCT threshold. In the right panel, this figure shows that the count of LIHTC applications per 100,000 households living in the Census tract (in the year of application) also jumps at the QCT threshold. In both panels, the running variable is a distance defined with respect to tract rank within its metropolitan area. Both panels split the data into 15 equal-interval bins on either side of the QCT threshold.

outcomes in Table 2. The running variables in the assignment rule evolve smoothly through the cutoff, along with four other tract characteristics (see Appendix Figure A19). Tracts just above and below the cutoff appear similar on observable characteristics.[30] I also find no evidence of heterogeneity by six application characteristics (Appendix Figure A20) and modest, if imprecise, evidence of heterogeneity across tracts (Appendix Table A1). Taken together, these analyses suggest increases in the subsidy rate would not attract more, but not "better," applications.

Implications. Both sets of results in this section suggest entry costs are high on average. Many applicants apply only when the net-of-tax price of low-income housing is almost zero, as it is in boosted tracts. There are two possible explanations for that result. One is that low-income housing is not valuable to developers, but this is belied by the "bidding" analysis in Section 6. The remaining possibility is that, although winning is valuable, entry costs negate most of the win value.

[30]Due to transformations involved in the running variable, there is a discrete change in the mass of tracts at the cutoff, but institutional considerations rule out potential concerns of precise control.

Table 2: Application Responses to Subsidy Generosity: QCT Threshold Estimates

	Is Boosted (1)	Applications (2)	Wins (3)	Proposed Units (4)	Funded Units (5)
Panel A: RD Estimates					
QCT Threshold	0.541***	0.155***	0.069**	8.37**	3.46**
	(0.014)	(0.052)	(0.027)	(3.26)	(1.68)
Panel B: Fuzzy RD Estimates					
Is Boosted		0.308***	0.123***	16.34***	5.50*
		(0.104)	(0.046)	(6.06)	(2.82)
Bandwidth	0.030	0.044	0.056	0.044	0.054
Tracts in Bandwidth	7,619	9,392	10,633	9,392	10,442
Untreated Mean at Threshold	0.340	0.341	0.134	17.47	6.72
Estimate / Mean		0.904	0.914	0.935	0.819

Notes: This table reports estimates of LIHTC application responses to the Qualified Census Tract (QCT) designation, which raises subsidy generosity discontinuously at the QCT threshold. Outcomes in Columns 2–4 are in terms of levels per 10,000 Census tract households per year. I define the running variable as a tract's distance in ranks from the QCT designation within the distribution of Census tracts in its metropolitan area. The untreated mean at the threshold is estimated by local linear regression to the left of the QCT threshold using the corresponding bandwidth in Panel B. The bottom row divides the fuzzy-RD estimate by the untreated mean at the threshold. $^{*} = p < 0.10$, $^{**} = p < 0.05$, $^{***} = p < 0.01$.

6 Bidding for Subsidies

Developers face a trade-off between their win probability and the rent they may charge if they win. This section characterizes their behavior in the face of this trade-off and estimates preferences over these objects. Using the preferences, I obtain a semi-structural estimate of the LIHTC's incidence on developers. This provides a useful check on the structural estimation in Section 8, where I will target this section's reduced-form results directly, rather than its incidence estimate.

6.1 Trade-Off Between Win Probability and Rental Income

Incentives in QAPs for lower rents, as noted in Section 2, allow developers to adjust their win probability by setting rents higher or lower. Three considerations make rent a useful dimension of applications to study. First, rent is particularly flexible as a choice variable, while other dimensions may be inflexible or require adjustments on other dimensions. Second, rent is a key component of application scores. In the median state in 2019, I calculate that rent preferences amounted to one quarter of the maximum point score. Third, the cost to developers of rent restrictions is measured

Figure 10: Trade-Off Between Win Probability and Rental Income

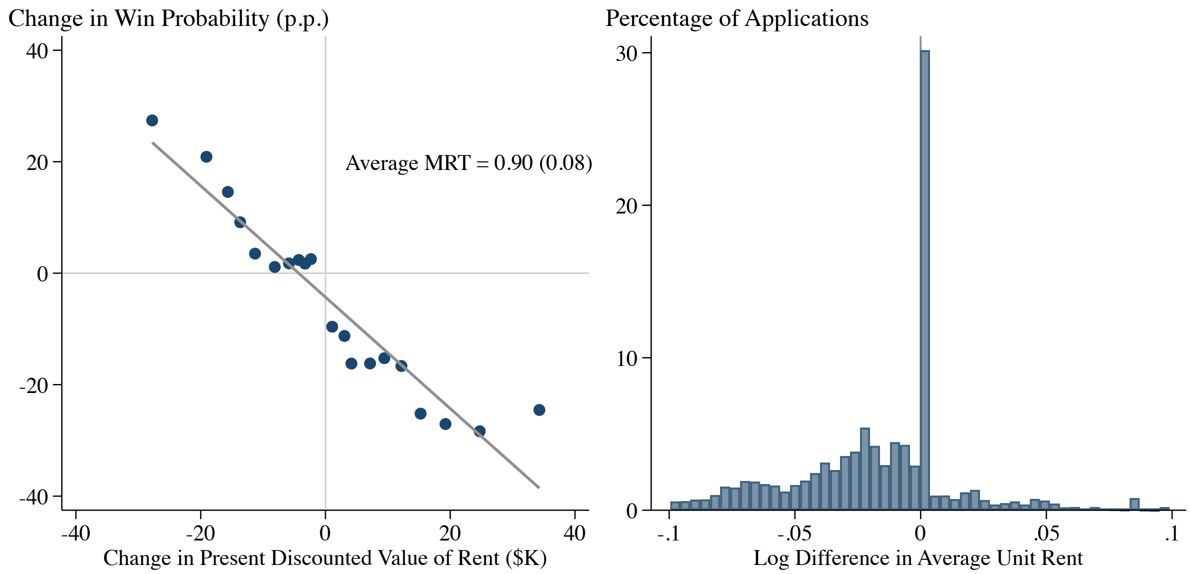

Notes: The left panel is a binned scatterplot of two dimensions of deviations from the applications that developers actually submitted. The horizontal axis is the present-value rent difference in thousands of dollars per unit. The vertical axis is the change in win probability in percentage points. The right panel is a histogram of the difference between developers' proposed average unit rent and the score-maximizing rent level.

in dollars, facilitating conclusions about their valuations of winning the LIHTC.[31]

To analyze bidding decisions, I compute win probabilities if developers were to deviate unilaterally by setting rents modestly higher or lower than those they actually propose. I implement this approach by re-scoring 6,793 applications from 13 states (30 percent of my sample). For each, I searched for "up" and "down" deviations that were compliant with the QAP, not strictly dominated by another choice, and which solely involved a change in rents. I also determined the highest rent that would still receive full points in the rent category. With the new scores in hand, I simulate new win probabilities for applications as in Section 2. For additional details, see Appendix B.

The left panel of Figure 10 shows the rent-versus-win-probability frontier facing developers, visualized as a binned scatterplot of both up and down deviations from actual applications. On the horizontal axis, I plot the change in applications' present values of rental income per unit. On the vertical axis, I plot the change in applications' win probabilities. The (negative of the) slope of the line through these points can be interpreted as an average marginal rate of transformation (MRT) between rents and probabilities. For a $1,000 reduction in the present value of rent per unit over a project's first thirty years of occupancy, developers can raise their probability of winning the

[31] An objection to rent as a trade-off variable is that, when rent regulations do not bind, rent concessions are free to developers on the margin. Using the market rents estimated in Section 2 to adjust for bindingness, I estimate higher marginal rates of substitution and lower developer incidence (see Appendix Table A2).

LIHTC by about 0.9 percentage points on average.[32]

The right panel of Figure 10 shows developers respond to rent incentives. It presents a histogram of applications' proposed rents, computed as an application-level average, relative to the highest rent level that would qualify for full points. Providing rent reductions beyond this level has no direct QAP-score benefit, and thus QAPs generate strongly "kinked" incentives with respect to rent around this level. About 30 percent of applications bunch precisely at the score-maximizing kink point.[33] Using variation in the location of the kink across states and over time also uncovers substantial bunching under weaker parametric assumptions (Appendix Figure A21). Finally, using application-level variation in local marginal rates of transformation, I show that developers with stronger local incentives are more likely to set lower rents (Appendix Figure A22).

These results immediately suggest developer incidence of the LIHTC. Developers routinely bid rents down to the minimum so as to maximize their win probabilities. If equilibrium profits from the LIHTC were minimal, developers would be unwilling to trade lower rents for higher win probabilities at the margin. Intuitively, if the benefit of a higher win probability exceeds the cost of less rent if won, the prize must be worth winning.

6.2 Estimating Developer Preferences and Incidence

Approach. Consider a developer facing a menu of possible rents and win probabilities for their application $\{r_{it}^j, p_{it}(r_{it}^j)\}$, holding fixed all other dimensions. Denote the rent difference between two alternatives by Δr_{it} and the win-probability difference by Δp_{it}. In general, rent levels could be related to other developer costs, which I represent as $\Delta e_{it} = \Delta\pi_1(\boldsymbol{s}_{it}) - \Delta r_{it}$.

From Equation 2, the difference between these alternatives in the application value function is

$$\Delta V^A(\boldsymbol{s}_{it}) \approx \Delta p_{it}\left[\pi_1(\boldsymbol{s}_{it}) - (1-p_{it})\beta \mathrm{E}[V^A(\boldsymbol{s}_{it+1}, \boldsymbol{\varepsilon}_{it+1}) \mid \boldsymbol{s}_{it}]\right] + p_{it}\left[\Delta r_{it} + \Delta e_{it}\right], \qquad (11)$$

where $\pi_1(\boldsymbol{s}_{it})$ and p_{it} are respectively the win value and the win probability associated with an arbitrary base level for the rent.

This equation formalizes the developer's trade-off between win probability and rental income. In particular, indifference between two alternatives such that $\Delta p_{it} \neq 0$ reveals the win value:

$$\pi_1(\boldsymbol{s}_{it}) = (1-p_{it})\beta \mathrm{E}[V^A(\boldsymbol{s}_{it+1}, \boldsymbol{\varepsilon}_{it+1}) \mid \boldsymbol{s}_{it}] - \frac{p_{it}(\Delta r_{it} + \Delta e_{it})}{\Delta p_{it}}. \qquad (12)$$

[32] Federal regulations do not require commitments to rent restrictions beyond 30 years, though some states encourage such extended commitments. I do not analyze these, as the present value of further years of rent regulation is highly sensitive to assumptions on developer discount rates.

[33] Many developers restrict rent even beyond what the LIHTC directly encourages. My inspection of applications suggests that developers may do so to qualify for other point categories in some QAPs (e.g., SRO housing), to comply with external rent restrictions (e.g., in public housing rehabilitation), or to pursue additional sources of funding.

The indifference condition in Equation 12 is a bid inversion, akin to those in research on auctions. Ignoring the value-function term, for example, a developer indifferent between a \$1,000 per-unit rise to the present value of rent and a one-percentage-point increase in win probability, starting at a baseline probability of 50 percent, must value winning at \$50,000 per unit in present value.[34]

Here I consider a simplified version of Equation 11, and in particular, the following auxiliary model of developer choice:

$$\Delta u_{ij} = \beta_1 \Delta p_i(r_{ij}) + \beta_2 \Delta \log r_{ij} + \Delta e_{ij}, \tag{13}$$

where j denotes an alternative (that is, the actual choice, the "up" deviation, or the "down" deviation) and Δe_{ij} is an error term. I use the latent utility index Δu to clearly distinguish from Equation 11, which I estimate as a conditional logit and as a fixed-effects linear probability model (LPM).

The parameters to be estimated are β_1, the developer's preference for a higher win probability, and β_2, their preference for higher rents. Without further assumptions about the structural error Δe_{ij}, these parameters are not identified. That is, applications likely to win may be unattractive to developers on dimensions other than rents (e.g., high construction costs). To obtain identification, I therefore assume that within the comparisons I have constructed, differences in unobservables are uncorrelated with differences in rents and win probabilities: $\mathrm{E}[\Delta e_{ij} \Delta p_{ij}] = \mathrm{E}[\Delta e_{ij} \Delta r_{ij}] = 0$. The justification for this assumption is that I control the comparison choices and consider a tightly restricted subset of alternatives.[35]

Results. Table 3 presents estimates of Equation 13. Both the LPM (Column 1) and the conditional logit (Column 3) find an average marginal rate of substitution between win probability and rents of approximately 0.9.[36] On average, developers would be indifferent between a 0.9-p.p. rise in their win probability and a \$1,000 rise in present-value rent. Encouragingly, this marginal rate of substitution matches the marginal rate of transformation in Figure 10. The final two rows of Table 3 use the approach in Equation 12 to estimate the mean per-unit value of winning the grant and for the incidence of the grant on developer profits. I find developers value winning at approximately \$50,000 per unit in present value. By comparison, the average grant amount per unit is \$120,000, implying that developers capture around 45 percent of the grant.

I explore a key concern with the use of simulated win probabilities to estimate preferences in

[34] While the static value of winning is $\pi_1(\boldsymbol{s}_{it})$, the dynamic equivalent is $\pi_1(\boldsymbol{s}_{it}) - (1 - p_{it})\beta \mathrm{E}[V^A(\boldsymbol{s}_{it+1}, \boldsymbol{\varepsilon}_{it+1}) \mid \boldsymbol{s}_{it}]$. Thus, the estimates in Table 3 apply to this dynamic object.

[35] This identification assumption follows Pakes (2010), who uses a similar comparison-based sample to difference out a structural error term. It rules out the possibility of changes in other costs that flow from serving a different tenant population. For instance, extremely-low-income tenants can require supportive services in addition to housing; this would lead to some overstatement of developer incidence.

[36] Across columns, the coefficients fluctuate in level due to differences in econometric approach. Incidence and the other objects of interest are determined by the ratio of coefficients, which is stable across approaches.

Table 3: Estimating Valuations from Bidding Behavior

	(1) OLS	(2) IV	(3) Cond. Logit.	(4) + Ctrl. Funct.
Win Probability	0.416***	1.708***	1.278***	6.423***
	(0.021)	(0.042)	(0.071)	(0.195)
Log Average Rent	0.620***	2.900***	1.826***	10.249***
	(0.049)	(0.098)	(0.147)	(0.395)
Applications	6,785	6,785	6,779	6,779
Marg. Rate of Substitution	0.923	1.051	0.884	0.987
	(0.040)	(0.021)	(0.039)	(0.019)
Mean Win Value Per Unit	\$54,957	\$48,292	\$57,385	\$51,385
	(2,441)	(1,022)	(2,608)	(1,040)
Developer Incidence Share	0.456	0.401	0.476	0.427
	(0.021)	(0.010)	(0.023)	(0.010)

Notes: This table reports estimates of coefficients in Equation 13. In Column 4, I take a control-function approach to instrument for the win probability in the conditional logit. Marginal rates of substitution are calculated as a ratio of coefficients; see the text for detail on the other calculations. Standard errors are clustered by application. $^{*} = p < 0.10$, $^{**} = p < 0.05$, $^{***} = p < 0.01$.

Columns 2 and 4 of Table 3. If the simulated probabilities provide noisy estimates of developers' beliefs, as seems likely, then choices will be more responsive to win probabilities than Columns 1 and 3 suggest. This bias will cause me to understate the marginal rate of substitution and, in turn, to overstate developer incidence. As a solution, I instrument for the win probability using the application's rank in the round-specific distribution of QAP scores, which is free from simulation-based measurement error. This approach does in fact yield slightly higher estimates of marginal rates of substitution and slightly lower developer valuations and incidence shares.

7 Structural Estimation

This section estimates the model introduced in Section 3. First, I introduce the parameterization of model primitives, along with other significant choices. I then explain the estimation procedure, which combines parametric policy iteration (PPI, see Rust, 2000; Sweeting, 2013) with simulated minimum distance (SMD, see Gourieroux et al., 1993). Finally, I list the empirical moments used in estimation, and I report the coefficient estimates and model fit.

7.1 Setup

I parameterize the win value $\pi_1(\boldsymbol{s}_{it})$, the outside option $\pi_0(\boldsymbol{s}_{it})$, the entry cost $\kappa(\boldsymbol{s}_{it})$, the distribution G of persistent unobservables ξ_i, and the distribution F of temporary unobservables $\boldsymbol{\varepsilon}_{it}$. I collect the parameters in the vector $\boldsymbol{\theta}$. The matrix $\boldsymbol{x}_i$ contains observable application characteristics. The parameterizations are as follows:

- *Outside Option, Win Value, and Entry Cost:* The primitives $\pi_0(\boldsymbol{s}_{it})$, $\pi_1(\boldsymbol{s}_{it})$, and $\kappa(\boldsymbol{s}_{it})$ are assumed to be linear in the history h_{it}, the unobservable ξ_i, and the characteristics of applications and tracts. For instance, outside options are $\pi_0(\boldsymbol{s}_{it}) = \pi_0^h h_i + \pi_0^\xi \xi_i + \boldsymbol{\pi}_0^x \boldsymbol{x}_i$, with structural parameters $(\pi_0^h, \pi_0^\xi, \boldsymbol{\pi}_0^x)$. The characteristics are the unit count, the per-unit value of tax credits, the tract poverty rate, and the tract (log) population density.

- *Application Characteristics:* I assume that potential applicants' number of units and credits per unit are distributed independent log-normally. To avoid drawing from the distributions of scores and set-asides—an extremely high-dimensional object—I directly simulate applications' chains of win probabilities for their initial application and any reapplications. I use a bivariate beta distribution and assume the win probabilities follow a Markov process.

- *Unobservables:* I assume the persistent unobservable heterogeneity ξ_i is distributed normally, setting the mean to zero and variance to unity, as it is rescaled in the primitive objects. For the temporary unobservables, developers draw i.i.d. Type I extreme value shocks $\boldsymbol{\varepsilon}_{it} = (\varepsilon_{it}^A, \varepsilon_{it}^B)$ each period. The shocks are additively separable from mean payoffs. The dispersion parameters are respectively σ_a and σ_b.

The unit of simulation is a tract–year. Each tract draws a new potential applicant each year, which then makes a sequence of application and building decisions until reaching a terminal state. In estimating the model, I define parameters in units of dollars of the potential application's qualified basis, normalizing for differences in scale. For instance, a value $\kappa(\boldsymbol{s}_{it}) = 0.1$ would mean that the entry cost is a tenth of the basis, or about 14 percent of tax-credit value.

The model's household side is calibrated. In particular, I take the elasticity of substitution $\rho = 0.691$ between housing and the non-housing good from Albouy et al. (2016), and I set ϕ to achieve a housing consumption share of 0.15.[37] From Section 2, I set the rent discount $\delta = 0.12$. For further details regarding the structural setup and counterfactuals, see Appendix B.

7.2 Procedure

Parametric Policy Iteration (PPI). The first part of the estimation procedure forms choice probabilities for simulated developers. The challenge this part solves is that, in dynamic models, optimal

[37] Estimated from U.S. Bureau of Economic Analysis data on the housing share of personal consumption expenditures.

choices and thus choice probabilities can be costly to compute. My implementation of PPI, which closely follows Sweeting (2013), is summarized below and is presented fully in Appendix C.

PPI methods form choice probabilities by iterating between two steps, "policy valuation" and "policy improvement" (Rust, 2000). In the valuation step, I compute an approximation to the expectation of the applicant value function at each state, $\mathrm{E}[V^A(\boldsymbol{s}_{it}, \boldsymbol{\varepsilon}_{it}) \mid \boldsymbol{s}_{it}]$, using a linear regression of flow payoffs $\Pi^A(a, \boldsymbol{s}_{it})$ on a polynomial basis of the state variables. This computation uses initialized values for structural parameters as well as choice probabilities.

In the improvement step, I then use the approximation of $\mathrm{E}[V^A(\boldsymbol{s}_{it}, \boldsymbol{\varepsilon}_{it}) \mid \boldsymbol{s}_{it}]$ to update the choice probabilities. In particular, application choice probabilities $\Pr(a \mid \boldsymbol{s}_{it})$ at each state depend upon $\Pi^A(a, \boldsymbol{s}_{it}) + (1 - p_{it})\beta\mathrm{E}[V^A(\boldsymbol{s}_{it+1}, \boldsymbol{\varepsilon}_{it+1}) \mid a, \boldsymbol{s}_{it}]$, the sum of the flow payoffs and the value function. Building choice probabilities follow similarly. With the updated choice probabilities in hand, I return to the valuation step. The procedure is repeated until revisions to the choice probabilities become small in absolute magnitude.

Simulated Minimum Distance (SMD). The second part of the estimation procedure chooses structural parameters to minimize a distance between moments of the actual data and simulated data generated above. I target quasi-experimental moments and key descriptive patterns in the data, so that simulated developers respond as do the actual developers to policy variation.

Formally, the SMD step finds parameter estimates $\hat{\boldsymbol{\theta}}$ that jointly minimize a distance between empirical moments $\hat{\boldsymbol{\beta}}$ and their simulation analogs $\tilde{\boldsymbol{\beta}}(\boldsymbol{\theta})$. In particular, the estimates $\hat{\boldsymbol{\theta}}$ solve

$$\arg\min_{\boldsymbol{\theta}} \, [\hat{\boldsymbol{\beta}} - \tilde{\boldsymbol{\beta}}(\boldsymbol{\theta})]' \boldsymbol{\Sigma}^{-1} [\hat{\boldsymbol{\beta}} - \tilde{\boldsymbol{\beta}}(\boldsymbol{\theta})], \tag{14}$$

where $\boldsymbol{\Sigma}^{-1}$ is a block-diagonal weight matrix. Each block is formed from the sum of the covariance matrices of the empirical and simulated moments.

7.3 Moments

I target three groups of empirical moments $\hat{\boldsymbol{\beta}}$: (1) causal effects of the LIHTC using quasi-random assignment, from Section 4; (2) application responses to changes in subsidy generosity, from Section 5; and (3) developer bidding behavior in response to incentives to reduce rents, from Section 6. These are joined by moments to match descriptive patterns. In greater detail, the moments are:

- *Winners Versus Losers.* In the simulated data, I can compare winners and losers, controlling for win probability (Equation 9). The model presumes that winners must build, so I estimate an equivalent specification among losers only: $B_{it} = \boldsymbol{X}_{it}\boldsymbol{\gamma}_0 + \gamma_1 p_{it} + u_{it}$, where B_{it} is an indicator for the development outcome and $\boldsymbol{X}_{it}$ contains application characteristics. The coefficients $\boldsymbol{\gamma}_0$ and γ_1 are targeted moments.

- *Application Supply.* First, I match application responses to the boost. The simulated response is calculated as an average derivative across all potential applicants, and I target the event-study estimate for the boost's effect on the annual probability (in percentage points) of any application from the tract. Second, I match the boost's (small) effects on the composition of the applicant pool from the RD estimates, targeting effects on win probabilities, credits requested per unit, and proposed unit counts (Appendix Figure A19). Third, I match the coefficients of two cross-sectional regressions. For all applications, I regress whether a tract has any applications in a given year on tract characteristics. For reapplications, I regress whether a losing applicant reapplies on both application and tract characteristics.
- *Bidding.* I target the fixed-effects LPM coefficients from Column 1 of Table 3. I enrich the regression specification in the table by interacting the win-probability and rent variables with the two tract characteristics, poverty and population density.
- *Distributional Moments.* I target the means and variances of application characteristics: the number of units and the credit amount requested per unit. I also target the mean and variance of the distribution of win probabilities. Among reapplicants, I match the coefficients of a regression of their current-round win probabilities on their prior-round probabilities.

7.4 Coefficient Estimates and Model Fit

I briefly discuss the parameter estimates in Appendix Table A3, reserving further discussion for the distributions of model primitives in Section 8, which are more readily interpreted.

Parameter estimates imply that potential applicants with high unit counts or in dense areas have stronger outside options on average. Applications with larger credit requests per unit generally have weaker outside options. There is state dependence: Applying improves applicants' outside options, and the cost of reapplying is lower than the cost of the initial application. The tax credit makes up a larger share of win value in low-income areas and for projects with high credits per unit. There are substantial returns to scale in entry costs, measuring size both by unit count and in credits requested per unit. Entry costs are lower in poor and low-density areas. The persistent unobservable induces a strong positive correlation in outside options and win values, implying there is advantageous selection on outside options into the LIHTC.

The model fits the targeted moments reasonably well (see Appendix Table A4). This is informative in itself: The model does not seem to struggle to explain developer behavior across three distinct empirical analyses. The largest gaps in fit occur in the cross-sectional regression of application supply on tract characteristics.

Table 4: Estimates of Model Primitives and Potential-Applicant Characteristics

	(1)	(2)	(3)	(4)	(5)	(6)
	Applicants			Non-Applicants		
	Median	P25	P75	Median	P25	P75
Panel A: Model Primitives as a Share of Basis						
Ex-Ante Value	1.41	1.13	1.76	0.56	0.25	0.94
Outside Option	1.25	0.88	1.65	-0.03	-0.71	0.63
Win Value	1.45	1.01	1.89	0.47	-0.11	1.06
Entry Cost	0.21	0.08	0.31	0.55	0.40	0.65
Panel B: Model Primitives in Thousands of Dollars Per Unit						
Ex-Ante Value	193	136	285	65	30	115
Outside Option	165	102	263	-3	-94	75
Win Value	194	129	278	55	-15	120
Entry Cost	24	8	43	65	41	94
Panel C: Potential-Applicant Characteristics						
Win Probability	0.48	0.05	0.71	0.29	0.10	0.55
Unit Count	51	34	89	43	31	59
Tax Credits Per Unit	218	137	301	184	133	259

Notes: Panels A and B report estimates of the model primitives, the ex-ante value of the application value function $\mathrm{E}[V^A(\boldsymbol{s}_{it}, \boldsymbol{\varepsilon}_{it}) \mid \boldsymbol{s}_{it}]$, the outside option $\pi_0(\boldsymbol{s}_{it})$, the win value $\pi_1(\boldsymbol{s}_{it})$, and the entry cost $\kappa(\boldsymbol{s}_{it})$. Panel C reports estimates of potential-applicant characteristics. Tax credits per unit are in thousands of dollars. Appendix Table A6 provides a similar analysis for winning and losing applicants.

8 Results and Counterfactuals

This section reports model results. After reviewing estimates of model primitives, I turn to the paper's two main questions: the LIHTC's housing-market impacts and its incidence. I also assess its cost-effectiveness in comparison to a stylized voucher program.

8.1 Model Primitives and Potential-Applicant Characteristics

Panels A and B of Table 4 reports estimates of the key model primitives. These are the (ex-ante) application value function $\mathrm{E}[V^A(\boldsymbol{s}_{it}, \boldsymbol{\varepsilon}_{it}) \mid \boldsymbol{s}_{it}]$, the outside option $\pi_0(\boldsymbol{s}_{it})$, the win value $\pi_1(\boldsymbol{s}_{it})$, and the entry cost $\kappa(\boldsymbol{s}_{it})$. To summarize the heterogeneity, I report medians and interquartile ranges. In Panel A, I report these values as shares of the potential application's eligible basis. Panel B rescales the values into dollar terms.

Estimated win values show winning the LIHTC is valuable to applicants. Taking the reciprocal, the credits are worth about the full project value for the median applicant ($1/(0.7 \cdot 1.45) = 1.01$).

Beyond the grant, the project value reflects construction costs, future rents, and other subsidies, so it could be greater or less than the grant amount in principle.

For the median applicant, winning is about 16 percent ($1.45/1.25 - 1 = 0.16$) better than the outside option. For such an applicant, the LIHTC is narrowly "buying out" their next-best-alternative land use. Moreover, outside options are positive (i.e., statically preferred to not building that period) for nearly all applicants. In dollar terms, outside options appear plausible by comparison to capitalized rents—that is, the former is less than the latter under reasonable discount rates. This discussion illuminates how the model reconciles the application responses and the award effects: by making applicants marginal to applying but not marginal to development. Applicants vary greatly in their primitives; such heterogeneity is necessary for developer incidence in equilibrium.

I now turn to the win values and outside options of non-applicants. Win values generally exceed outside options, and outside options are negative for approximately half of non-applicants, implying much land with little productive use. Model results suggest the key driver of selection into application is variation in entry costs: The median applicant has small entry costs, whereas the median non-applicant would burn the subsidy's entire value in applying. Such an applicant would not find it profitable to apply even if they would win with certainty. This conclusion is natural: Had selection into application been driven by variation in outside options, for instance, losers would be unlikely to develop privately. Overall, estimated entry costs appear large, at 10 percent of win value for the median applicant.[38] The table also examines potential applicants on observable characteristics, finding that applicants are positively selected on win probability relative to non-applicants.

8.2 Incidence and Impacts

Table 5 shows the model-based estimates of incidence and impacts. Column 1 reports overall averages. Columns 2 to 5 split the simulation according to whether a potential applicant's tract is above or below the median on two characteristics, the poverty rate or the population density.

Overall, I find that households receive about 30 percent of the welfare gains from the LIHTC. Much of the LIHTC instead goes to developer profits and entry costs. Entry costs absorb a quarter of the LIHTC. The share is more greater in the comparisons of primitives above, as these costs are paid by both winning and losing applicants, whereas households and developers only benefit when applications win. The gains to households are split between rent savings and general-equilibrium effects but mostly arise from the former.[39]

[38] These estimated entry costs are difficult to benchmark against real-world costs and are ultimately a structural object needed to explain entry behavior. Although I take entry costs as arising per application, there is likely a large fixed component at the developer corporate level.

[39] LIHTC residents are poorer than the typical resident of the same tract, so the general-equilibrium component is less targeted to lower-income households than the rent-savings component (see Appendix Figure A23).

Table 5: Model-Based Incidence and Impacts

	(1)	(2)	(3)	(4)	(5)
		Poverty Rate		Population Density	
		Above Med.	Below Med.	Above Med.	Below Med.
Panel A: Incidence					
Household Share	0.314	0.265	0.361	0.391	0.396
	(0.027)	(0.026)	(0.029)	(0.047)	(0.031)
Rent Savings	0.233	0.180	0.284	0.265	0.177
	(0.019)	(0.016)	(0.022)	(0.022)	(0.014)
Gen. Eqm. Effects	0.081	0.085	0.076	0.125	0.219
	(0.022)	(0.023)	(0.021)	(0.045)	(0.024)
Developer Share	0.437	0.480	0.401	0.397	0.382
	(0.033)	(0.032)	(0.034)	(0.058)	(0.032)
Entry Cost Share	0.250	0.255	0.239	0.212	0.222
	(0.025)	(0.030)	(0.022)	(0.045)	(0.020)
Panel B: Impacts					
Displacement Rate	0.777	0.758	0.788	0.579	0.845
	(0.018)	(0.024)	(0.023)	(0.103)	(0.028)
Cost Per Net Unit	960	870	1019	515	1371
	(100)	(98)	(126)	(97)	(388)

Notes: This table reports model-based estimates of the incidence and impacts of the LIHTC. The displacement rate is one minus the net change in the housing stock per LIHTC unit. For tract characteristics, I split at the median among simulated applicants. Costs per unit are in thousands of dollars. Bootstrap standard errors are reported in parentheses.

The model also finds that subsidized units substantially displace other construction on the same parcel and are not one-for-one net additions to the local housing stock. For every ten subsidized units, there are about two net units added to the housing stock, an estimate that is consistent with Baum-Snow and Marion (2009). Unlike the results in Section 4, however, these estimates account for reapplication, general-equilibrium effects, and state dependence in behavior.

There is spatial variation in the displacement rate and thus in costs per net unit. Displacement is weaker in high-density than in low-density areas. As a result, the LIHTC's cost per net new unit varies notably, and essentially inversely with its cost per subsidized unit. Spatial heterogeneity in incidence is also considerable. In low-poverty areas, the LIHTC provides larger transfers to households, consistent with spatial variation in rent savings. Elsewhere, the LIHTC more generally subsidizes development with less redistributive benefit.

I explore the sensitivity of the structural results to changes in the empirical moments in Appendix Table A5. Following Andrews et al. (2017), I report the changes in estimates of model primitives,

Figure 11: Model-Based Comparison of the LIHTC and Rent Vouchers

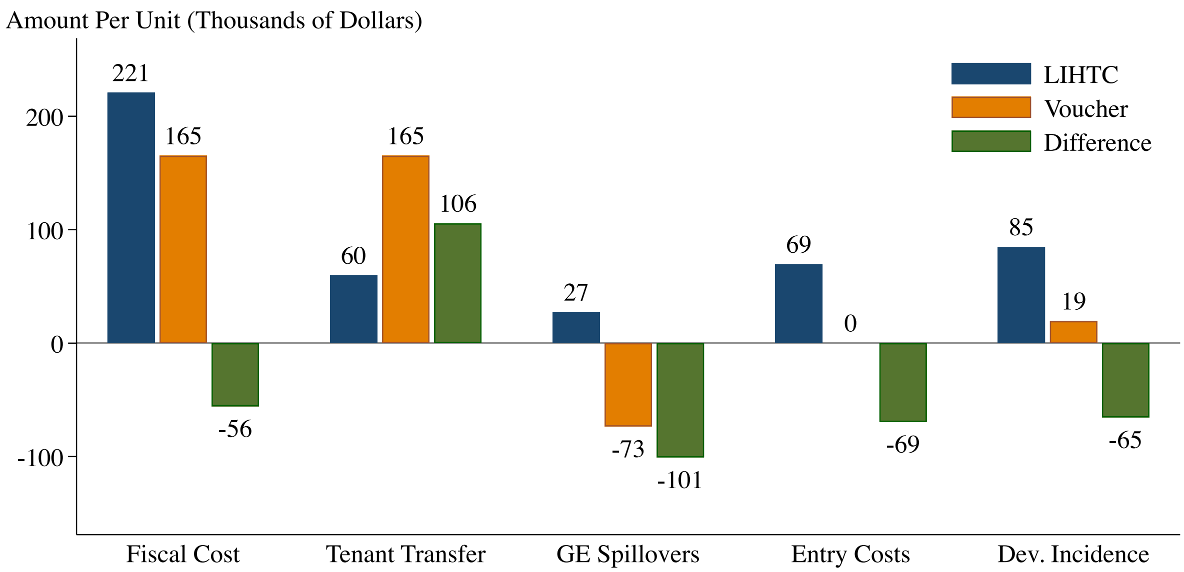

Notes: This figure contrasts the LIHTC to rent vouchers using the model. The voucher is specified to achieve the same aggregate benefit to households as the LIHTC. All values are expressed in thousands of dollars per LIHTC unit.

incidence, and impacts that result from small perturbations to moments. There are three important lessons. First, the direction of sensitivity of parameters to moments is generally intuitive. For instance, a higher application-supply elasticity would push down estimated entry costs. Second, specific equations do not identify specific structural parameters in isolation, so estimating the model jointly over the moments does matter. Third, incidence estimates are especially sensitive to reapplication behavior, a result that suggests the value of dynamic models in this context.

8.3 Vouchers as a Policy Alternative

In two counterfactuals, I compare the LIHTC to its natural policy alternative, vouchers. The first counterfactual replaces the LIHTC with a tenant-based subsidy that achieves the same aggregate benefit to households as the LIHTC, while allowing the fiscal cost to change. The second holds the fiscal cost fixed and allows the transfer to households to change.

Both comparisons consider a stylized voucher program. That is, I abstract away from many differences between these two programs and from institutional details of vouchers. The main purpose of the comparison is therefore to interpret, within the context of the model, the magnitudes of the incidence estimates. After presenting these results, I discuss considerations beyond the model that are likely relevant to comparisons of the LIHTC to real-world vouchers.

Figure 11 reports results for the first counterfactual. Holding household utility fixed, vouchers reduce present-value fiscal costs by 25 percent, or by $56,000 per LIHTC unit. In the fixed-budget counterfactual, the switch to vouchers expands the housing stock and raises household welfare. Holding the budget fixed, vouchers' effect on the aggregate housing stock is about 40 percent larger than that of the LIHTC. The total welfare benefit to households would be about 60 percent larger, rising from from \$87,000 per LIHTC unit to \$138,000.

I then use the model to examine how this cost difference arises. Vouchers and the LIHTC impose opposite-signed pecuniary externalities on market rents. All else equal, the government must therefore make a larger outlay under vouchers to achieve the same overall welfare gain for households. In particular, offsetting vouchers' pecuniary impacts on unassisted households requires a transfer to assisted tenants that is more than twice the per-unit present value of the rent discounts received by LIHTC tenants. Weighing in the other direction, vouchers by assumption avoid the LIHTC's entry costs, and they incur less incidence on developer profits. On balance, these latter forces leave vouchers with a modest fiscal advantage over the LIHTC. My estimates of the compensated difference in fiscal costs are close to accounting-style comparisons of housing policies, which also find moderate advantages for vouchers (Weicher, 2012).

My analysis omits many interesting differences between real-world housing programs, some of which are project-specific in nature. First, as noted above, project-based programs can simplify the provision of other supports, particularly healthcare, alongside housing. This aspect is challenging to value and could apply to the half of LIHTC properties that provide elderly or supportive housing (Table 1). Second, the LIHTC's federal tax expenditure understates its total cost: The developments often also receive subsidized debt, state-level assistance, and local property tax abatements. Recent estimates by Lang and Olsen (2023) suggest that, in California, these "hidden subsidies" may amount to twice the federal tax credit. Third, vouchers have administrative costs ignored here. Finally, the differences in average income between assisted and non-assisted households imply the household benefits from vouchers are more progressively distributed than those of project-based assistance. On balance, these factors are likely to enlarge LIHTC's cost advantage on average, but more significantly, they create project-level heterogeneity. Such heterogeneity could plausibly justify the occasional use of project-based assistance but hardly overturns a presumption that it is inferior to vouchers in general.

9 Conclusion

In the U.S. and other countries, limitations on the supply of housing in desirable cities have driven up housing costs, creating demand for policy answers to the "housing crisis." At the same time, poor households have long faced severe housing-cost burdens, reflecting a denominator problem of low incomes more so than a numerator problem of high rents. These twin social challenges motivate many low-income housing policies, including public housing, rent vouchers, and project-based subsidies. Are project-based subsidies effective answers to either of these challenges?

I examine this question in the context of the LIHTC, the largest U.S. project-based subsidy. Using newly-collected data on developer applications and three quasi-experimental research designs, I estimate a dynamic model of developer behavior, and then I use the model to evaluate the LIHTC's impacts and incidence. I find the LIHTC has little impact on the overall size of the housing stock, as the LIHTC heavily displaces private development that would have otherwise soon occurred. Project-based subsidies are therefore of limited value as an answer to the first challenge.

I also find project-based subsidies ultimately benefit lower-income households, making them relevant to the second social challenge. LIHTC residents receive a transfer through below-market rents, albeit modest on average, and the transfer also reallocates new housing to lower-income households than would otherwise live in it. Allocating capital in housing markets via a complex governmental process, however, does create two significant problems: high entry costs and significant incidence on a subset of developers with a LIHTC-specific cost advantage. This conclusion has particular relevance to recent U.S. expansions of "supply-side" subsidies in domains such as energy and manufacturing. An important concern for policy design will be how difficult it is to compete for these subsidies, both because of the intrinsic costs of policy complexity and because such complexity may shift incidence to those firms most adept at navigating it.

References

Aaron, Henry J, *Shelter and Subsidies: Who Benefits from Federal Housing Policies?*, Brookings Institution, 1972.

Abdulkadiroğlu, Atila, Joshua D Angrist, Yusuke Narita, and Parag A Pathak, "Research Design Meets Market Design: Using Centralized Assignment for Impact Evaluation," *Econometrica*, 2017, *85* (5), 1373–1432.

Albouy, David, Gabriel Ehrlich, and Yingyi Liu, "Housing Demand, Cost-of-Living Inequality, and the Affordability Crisis," Working Paper 22816, National Bureau of Economic Research 2016.

Almagro, Milena and Tomás Domínguez-Iino, "Location Sorting and Endogenous Amenities: Evidence from Amsterdam," Working Paper 2022.

Anagol, Santosh, Fernando V Ferreira, and Jonah M Rexer, "Estimating the Economic Value of Zoning Reform," Working Paper 29440, National Bureau of Economic Research 2021.

Andrews, Isaiah, Matthew Gentzkow, and Jesse M Shapiro, "Measuring the Sensitivity of Parameter Estimates to Estimation Moments," *Quarterly Journal of Economics*, 2017, *132* (4), 1553–1592.

Baum-Snow, Nathaniel and Justin Marion, "The Effects of Low Income Housing Tax Credit Developments on Neighborhoods," *Journal of Public Economics*, 2009, *93* (5-6), 654–666.

Berger, David, Nicholas Turner, and Eric Zwick, "Stimulating Housing Markets," *The Journal of Finance*, 2020, *75* (1), 277–321.

Bergman, Peter, Raj Chetty, Stefanie DeLuca, Nathaniel Hendren, Lawrence F Katz, and Christopher Palmer, "Creating Moves to Opportunity: Experimental Evidence on Barriers to Neighborhood Choice," *American Economic Review*, forthcoming.

Burge, Gregory S, "Do Tenants Capture the Benefits from the Low-Income Housing Tax Credit Program?," *Real Estate Economics*, 2011, *39* (1), 71–96.

Calder-Wang, Sophie, "The Distributional Impact of the Sharing Economy on the Housing Market," Working Paper 2022.

Chodorow-Reich, Gabriel, John Coglianese, and Loukas Karabarbounis, "The Macro Effects of Unemployment Benefit Extensions: A Measurement Error Approach," *Quarterly Journal of Economics*, 2019, *134* (1), 227–279.

Christensen, Peter and Christopher Timmins, "The Damages and Distortions from Discrimination in the Rental Housing Market," *Quarterly Journal of Economics*, 2023, *138* (4), 2505–2557.

Collinson, Robert, Ingrid Gould Ellen, and Jens Ludwig, "Low-Income Housing Policy," in "Economics of Means-Tested Transfer Programs in the United States, Volume 2," University of Chicago Press, 2015, pp. 59–126.

Cook, Cody, Pearl Z Li, and Ariel J Binder, "Where to Build Affordable Housing? Evaluating the Tradeoffs of Location," Working Paper 2023.

Davis, Morris A, Jesse Gregory, and Daniel A Hartley, "The Long-Run Effects of Low-Income Housing on Neighborhood Composition," Working Paper 2019.

Derby, Elena, "Does Growing Up in Tax-Subsidized Housing Lead to Higher Earnings and Educational Attainment?," Working Paper 2021.

Desai, Mihir, Dhammika Dharmapala, and Monica Singhal, "Tax Incentives for Affordable Housing: The Low Income Housing Tax Credit," *Tax Policy and the Economy*, 2010, *24* (1), 181–205.

Diamond, Rebecca and Tim McQuade, "Who Wants Affordable Housing in Their Backyard? An Equilibrium Analysis of Low-Income Property Development," *Journal of Political Economy*, 2019, *127* (3), 1063–1117.

_ , _ , and Franklin Qian, "The Effects of Rent Control Expansion on Tenants, Landlords, and Inequality: Evidence from San Francisco," *American Economic Review*, 2019, *109* (9), 3365–3394.

Ellen, Ingrid G, Keren M Horn, and Katherine M O'Regan, "Poverty Concentration and the Low Income Housing Tax Credit: Effects of Siting and Tenant Composition," *Journal of Housing Economics*, 2016, *34*, 49–59.

Ellen, Ingrid Gould and Keren Mertens Horn, "Points for Place: Can State Governments Shape Siting Patterns of Low-Income Housing Tax Credit Developments?," *Housing Policy Debate*, 2018, *28* (5), 727–745.

Eriksen, Michael D and Stuart S Rosenthal, "Crowd Out Effects of Place-Based Subsidized Rental Housing: New Evidence from the LIHTC Program," *Journal of Public Economics*, 2010, *94* (11-12), 953–966.

Favilukis, Jack, Pierre Mabille, and Stijn Van Nieuwerburgh, "Affordable Housing and City Welfare," *Review of Economic Studies*, 2023, *90* (1), 293–330.

Feiveson, Laura, "General Revenue Sharing and Public Sector Unions," *Journal of Public Economics*, 2015, *125*, 28–45.

Freedman, Matthew and Emily G Owens, "Low-Income Housing Development and Crime," *Journal of Urban Economics*, 2011, *70* (2-3), 115–131.

Glaeser, Edward L and Joseph Gyourko, *Rethinking Federal Housing Policy*, AEI: American Enterprise Institute for Public Policy Research, 2008.

Gourieroux, Christian, Alain Monfort, and Eric Renault, "Indirect Inference," *Journal of Applied Econometrics*, 1993, *8* (S1), S85–S118.

Gyourko, Joseph, "Housing Supply," *Annual Review of Economics*, 2009, *1* (1), 295–318.

Holmes, Thomas J and Holger Sieg, "Structural Estimation in Urban Economics," in "Handbook of Regional and Urban Economics," Vol. 5, Elsevier, 2015, pp. 69–114.

Hopenhayn, Hugo A, "Entry, Exit, and Firm Dynamics in Long Run Equilibrium," *Econometrica*, 1992, pp. 1127–1150.

Hortaçsu, Ali and David McAdams, "Mechanism Choice and Strategic Bidding in Divisible Good Auctions: An Empirical Analysis of the Turkish Treasury Auction Market," *Journal of Political Economy*, 2010, *118* (5), 833–865.

House, Christopher L and Matthew D Shapiro, "Temporary Investment Tax Incentives: Theory with Evidence from Bonus Depreciation," *American Economic Review*, 2008, *98* (3), 737–768.

Hsiao, Allan, "Sea Level Rise and Urban Adaptation in Jakarta," Working Paper 2022.

Jacob, Brian A and Lars Lefgren, "The Impact of Research Grant Funding on Scientific Productivity," *Journal of Public Economics*, 2011, *95* (9-10), 1168–1177.

Joint Committee on Taxation, "Present Law And Background Relating To Residential Real Estate," Technical Report JCX-4-23 2023.

Kazis, Noah M, "The Failed Federalism of Affordable Housing: Why States don't Use Housing Vouchers," *Mich. L. Rev.*, 2022, *121*, 221.

Khadduri, Jill and Charles Wilkins, "Designing Subsidized Rental Housing Programs: What Have We Learned?," in "Revisiting Rental Housing: Policies, Programs, and Priorities," Brookings Institution Press, 2008, pp. 161–190.

Lang, Bree J and Edgar O Olsen, "The Hidden Subsidies of the Low-Income Housing Tax Credit Program," Working Paper 2023.

Levy, Antoine, "Housing Policy with Home-Biased Landlords: Evidence from French Rental Markets," Working Paper 2021.

_ , "Regulating Housing Quality: Evidence from France," Working Paper 2023.

Mian, Atif and Amir Sufi, "The Effects of Fiscal Stimulus: Evidence from the 2009 Cash for Clunkers Program," *The Quarterly Journal of Economics*, 2012, *127* (3), 1107–1142.

Murphy, Alvin, "A Dynamic Model of Housing Supply," *American Economic Journal: Economic Policy*, 2018, *10* (4), 243–267.

Olsen, Edgar O, "Housing Programs for Low-Income Households," in "Means-Tested Transfer Programs in the United States," University of Chicago Press, 2003, pp. 365–442.

_ **and Jeffrey E Zabel**, "US Housing Policy," in "Handbook of Regional and Urban Economics," Vol. 5, Elsevier, 2015, pp. 887–986.

Orlebeke, Charles J, "The Evolution of Low-Income Housing Policy, 1949 to 1999," *Housing Policy Debate*, 2000, *11* (2), 489–520.

Pakes, Ariel, "Alternative Models for Moment Inequalities," *Econometrica*, 2010, *78* (6), 1783–1822.

Quigley, John M, "Rental Housing Assistance," *Cityscape*, 2011, pp. 147–158.

Rosenbaum, Paul R and Donald B Rubin, "The Central Role of the Propensity Score in Observational Studies for Causal Effects," *Biometrika*, 1983, *70* (1), 41–55.

Rosenthal, Stuart S, "Are Private Markets and Filtering a Viable Source of Low-Income Housing? Estimates from a "Repeat Income" Model," *American Economic Review*, 2014, *104* (2), 687–706.

Rust, John, "Parametric Policy Iteration: An Efficient Algorithm for Solving Multidimensional DP Problems?," Working Paper 2000.

Shelburne, Mark, "LIHTC Qualified Allocation Plans, Explained (A Series)," https://www.novoco.com/notes-from-novogradac/lihtc-qualified-allocation-plans-explained-series 2021.

Sinai, Todd and Joel Waldfogel, "Do Low-Income Housing Subsidies Increase the Occupied Housing Stock?," *Journal of Public Economics*, 2005, *89* (11-12), 2137–2164.

Soltas, Evan J, "The Price of Inclusion: Evidence from Housing Developer Behavior," *Review of Economics and Statistics*, forthcoming.

Sportiche, Noemie, "Early-Life Health Impacts of Affordable Housing in Higher-Income Areas: Evidence from Massachusetts Chapter 40B," Working Paper 2022.

Suárez-Serrato, Juan Carlos and Philippe Wingender, "Estimating Local Fiscal Multipliers," Working Paper 22425, National Bureau of Economic Research 2016.

Sweeting, Andrew, "Dynamic Product Positioning in Differentiated Product Markets: The Effect of Fees for Musical Performance Rights on the Commercial Radio Industry," *Econometrica*, 2013, *81* (5), 1763–1803.

U.S. Department of Housing and Urban Development, *Housing in the Seventies: A Report of the National Housing Policy Review* 1973.

Vale, Lawrence J and Yonah Freemark, "From Public Housing to Public–Private Housing: 75 Years of American Social Experimentation," *Journal of the American Planning Association*, 2012, *78* (4), 379–402.

Waldinger, Daniel, "Targeting In-Kind Transfers Through Market Design: A Revealed Preference Analysis of Public Housing Allocation," *American Economic Review*, 2021, *111* (8), 2660–96.

Weicher, John C, *Housing Policy at a Crossroads: The Why, How, and Who of Assistance Programs*, AEI Press, 2012.

Winnick, Louis, "The Triumph of Housing Allowance Programs: How a Fundamental Policy Conflict Was Resolved," *Cityscape*, 1995, pp. 95–121.

Appendices for Online Publication

A Additional Tables and Figures **2**

B Supplemental Information **36**

C Theoretical and Structural Appendix **46**

A Additional Tables and Figures

Figure A1: Data Coverage

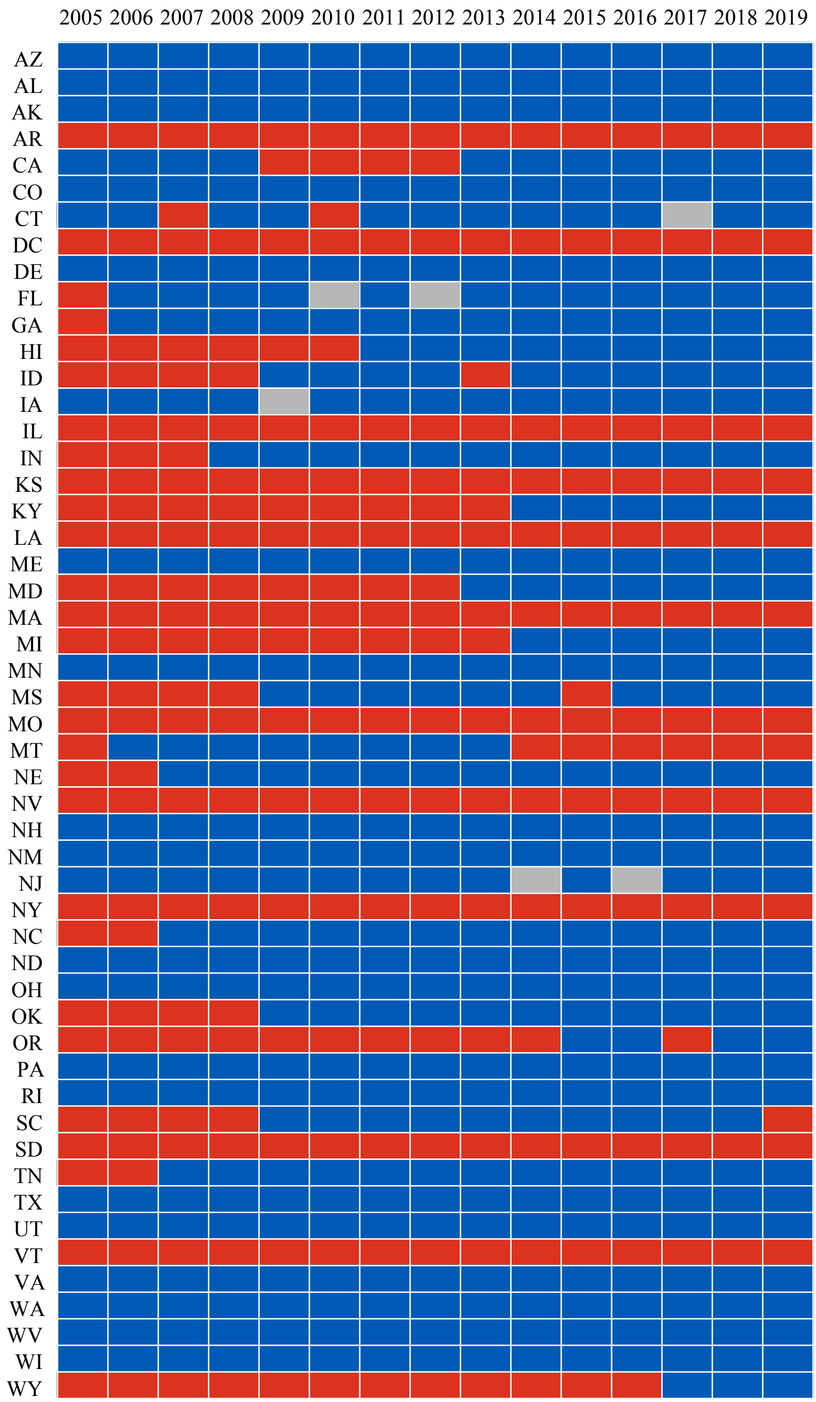

Notes: This figure shows the coverage of the LIHTC application data. Blue indicates coverage, red indicates non-coverage, and gray indicates that the state did not hold a LIHTC round in that year.

Figure A2: Neighborhood Effects of Tax Credit Awards:
Controlling for Application Characteristics

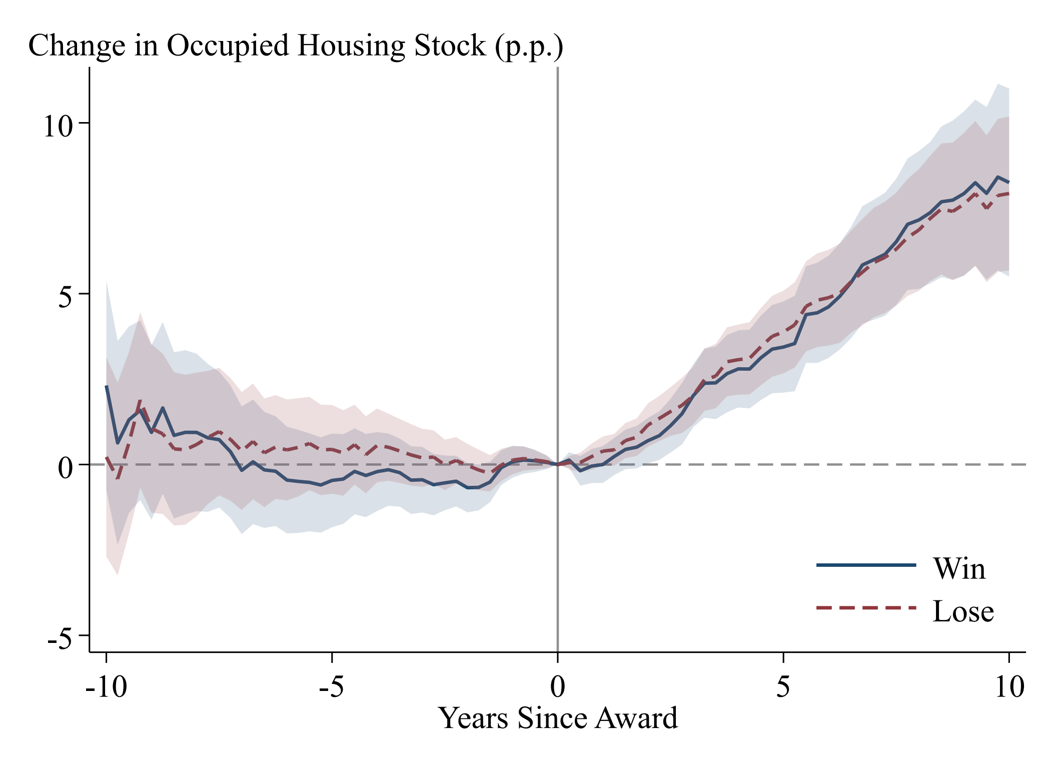

Notes: The figure plots win and lose effects on the occupied housing stock (change in percentage points), as in Figure 7. The specification is an event-study analog of Equation 9 that includes non-applicant tracts: $\Delta \log \mathrm{E}[Y_{i,t+k}] = \alpha_k \mathrm{Win}_{it} + \beta_k \mathrm{Lose}_{it} + \boldsymbol{X}_{it} \boldsymbol{\gamma}_k$, for k quarters or years from the award. It is estimated via Poisson regression. The key distinction from Figure 7 is that it includes in controls X_{it} the two application characteristics on which there is significant imbalance: whether the application is entirely new-construction and the leave-out win rate of the developer. The color bands depict pointwise 95-percent confidence intervals, and standard errors are clustered at the tract level.

Figure A3: Neighborhood Effects of Tax Credit Awards,
Instrumenting for Actual Wins with Simulated Wins

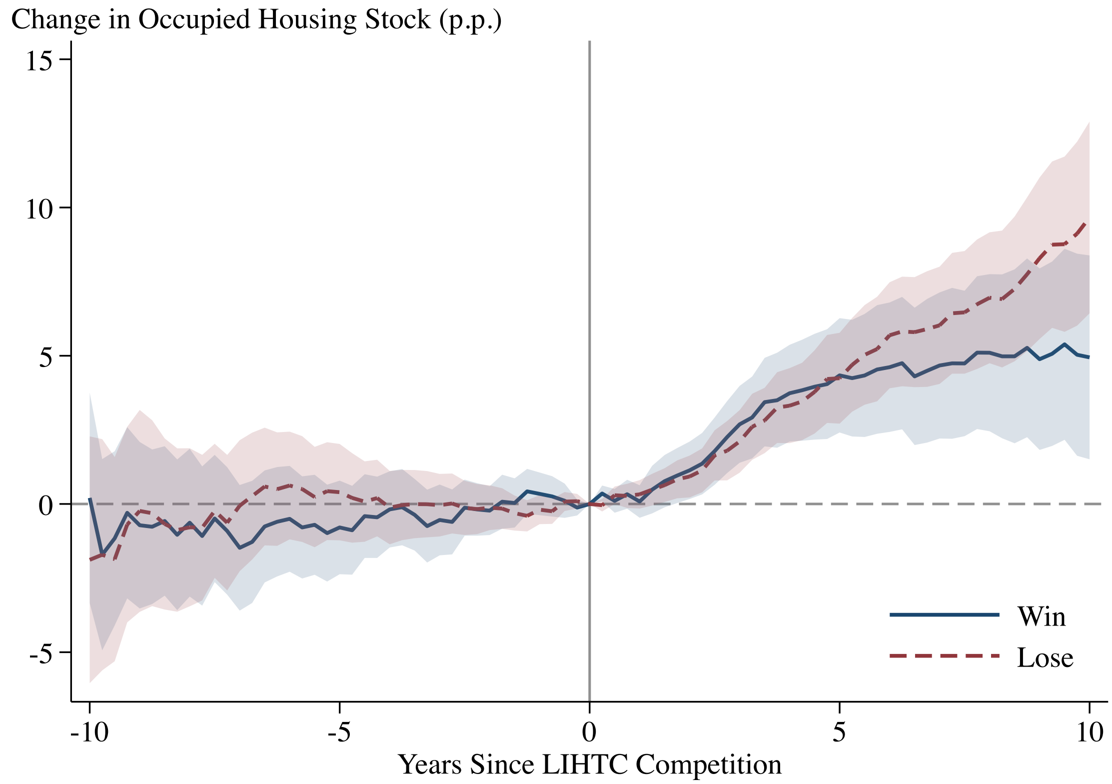

Notes: This figure plots win and lose effects on the occupied housing stock (change in percentage points). The specification is an instrumental-variables (IV) version of Figure 7, where I instrument for the actual tax-credit assignment of an application with its simulated assignment. I implement the IV within the Poisson regression using a control function. The color bands depict pointwise 95-percent confidence intervals, and standard errors are clustered at the tract level.

Figure A4: Parcel-Level Effects of Tax Credit Awards (Weighted)

Panel A: Total Development

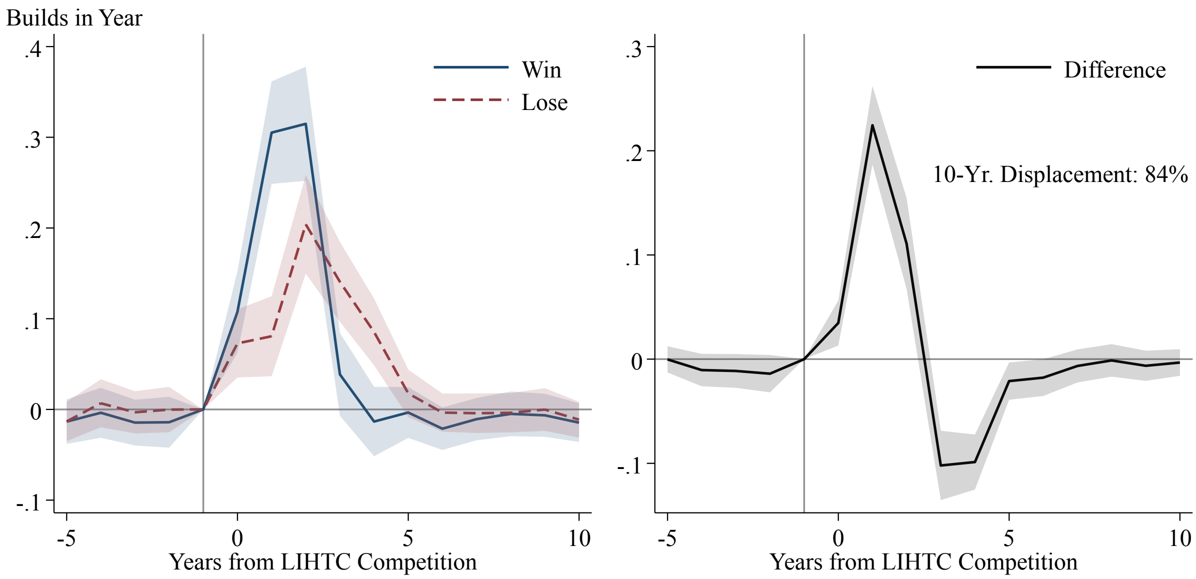

Panel B: Development by Type

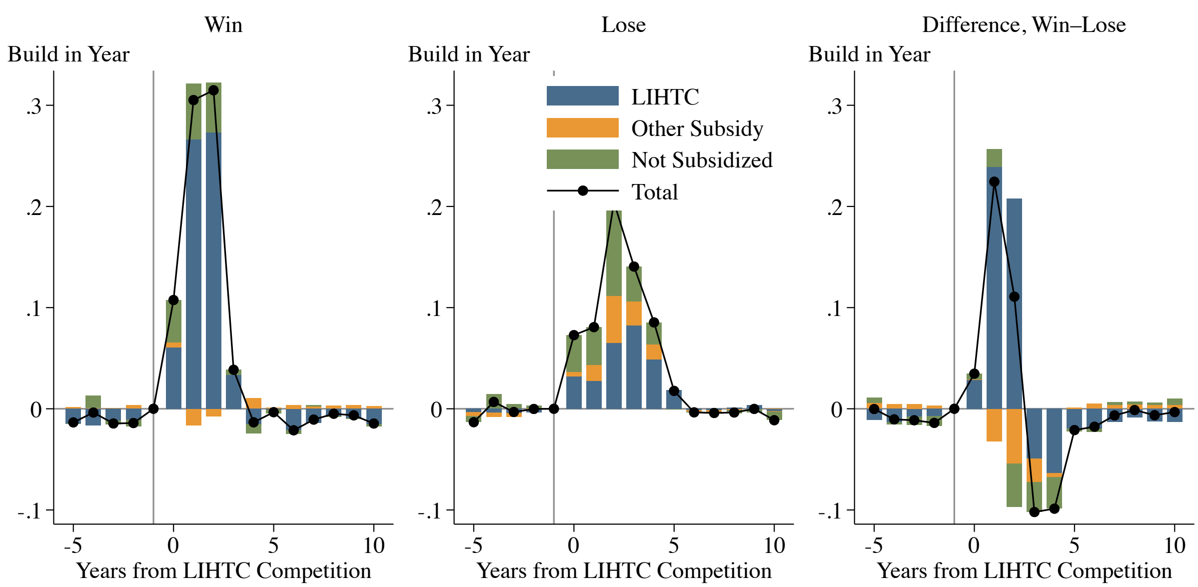

Notes: Panel A plots win and lose effects (on left) and the win–lose difference (on right), of tax credit awards on the probability that construction is recorded as completed in a given year. Panel B decomposes these development impacts into LIHTC, other-subsidy, and unsubsidized development. Multiple parcels associated with a single application are aggregated using floor-space weighting. Standard errors in Panel A are clustered by the consolidated parcel.

Figure A5: Parcel-Level Effects of Tax Credit Awards: Heterogeneity by Win Probability

Panel A: Event Study

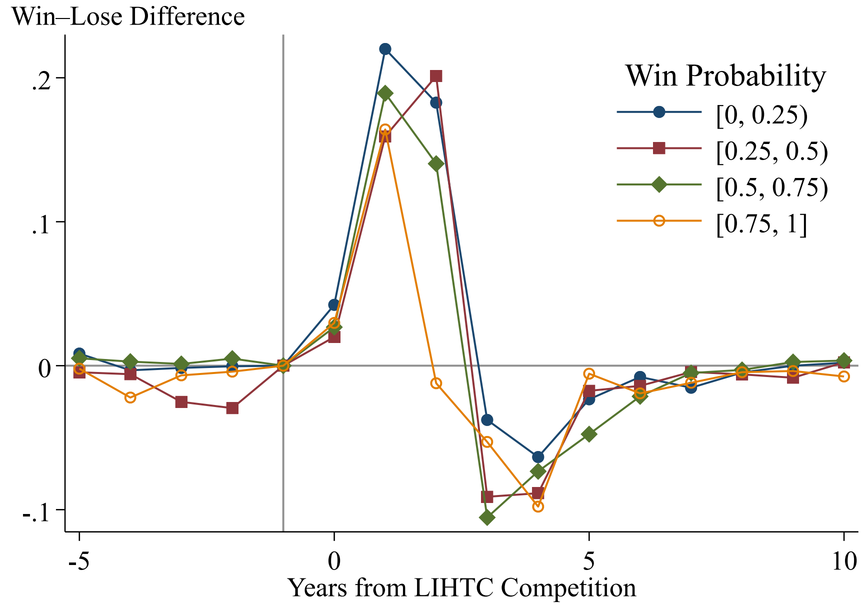

Panel B: Displacement Rate Summary

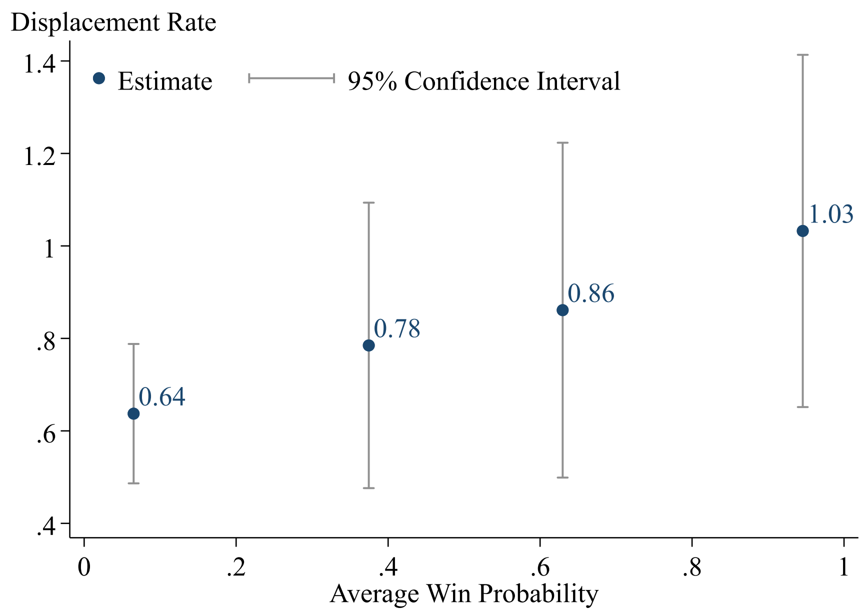

Notes: Panel A plots effects of winning versus losing a tax credit award on the probability that construction is recorded as completed in a given year. The figure reports these estimates separately by splitting the sample into quarter intervals of the application's win probability. Panel B reports the 10-year displacement rate by interval of the win probability. Standard errors in Panel A are clustered by the consolidated parcel.

Figure A6: Winners-Versus-Losers Comparison in Neighborhood Event Study

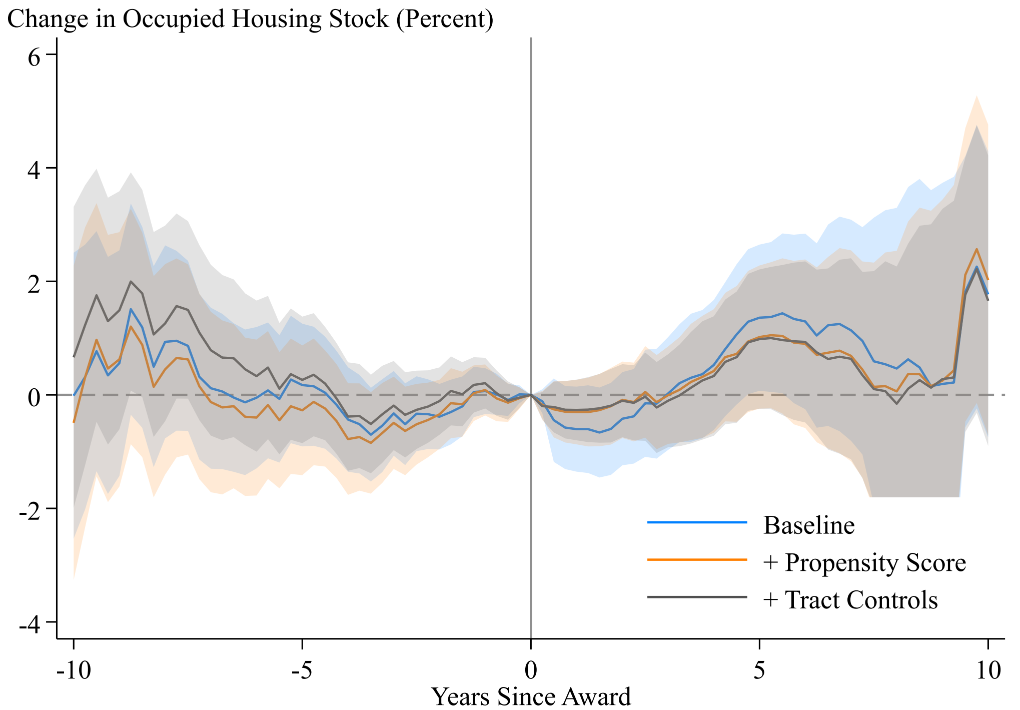

Notes: This figure plots win–lose difference coefficients in an event-study comparison of Census tracts with LIHTC applications. The blue line shows baseline estimates, which only include county–year fixed effects as controls. The orange line augments this specification with the flexible control for win probability. The black line adds controls for Census tract pre-award characteristics, as listed in Section 4. The color bands depict pointwise 95-percent confidence intervals, and standard errors are clustered at the tract level.

Figure A7: Detailed Land-Use Effects of Tax Credit Awards

Panel A: Single- Versus Multi-Family

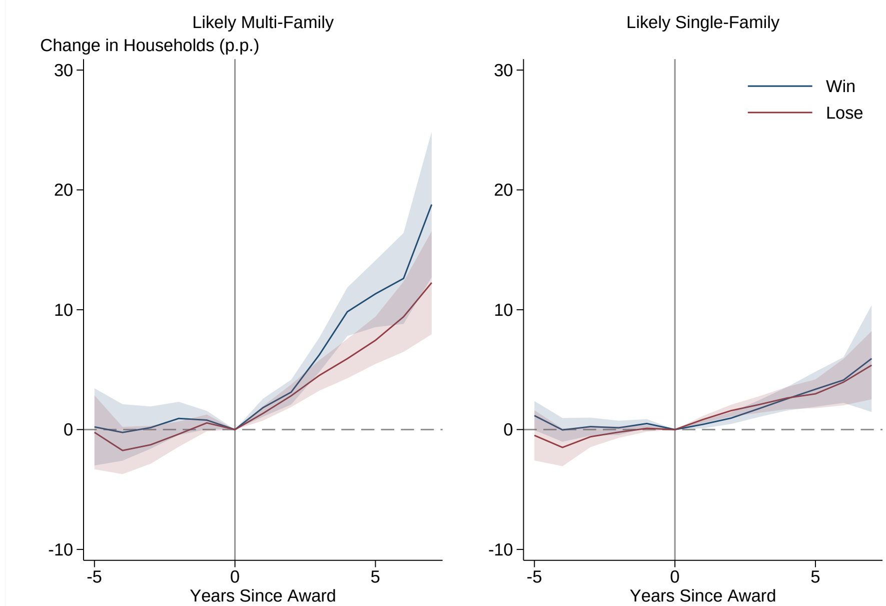

Panel B: Rental Versus Owner-Occupied

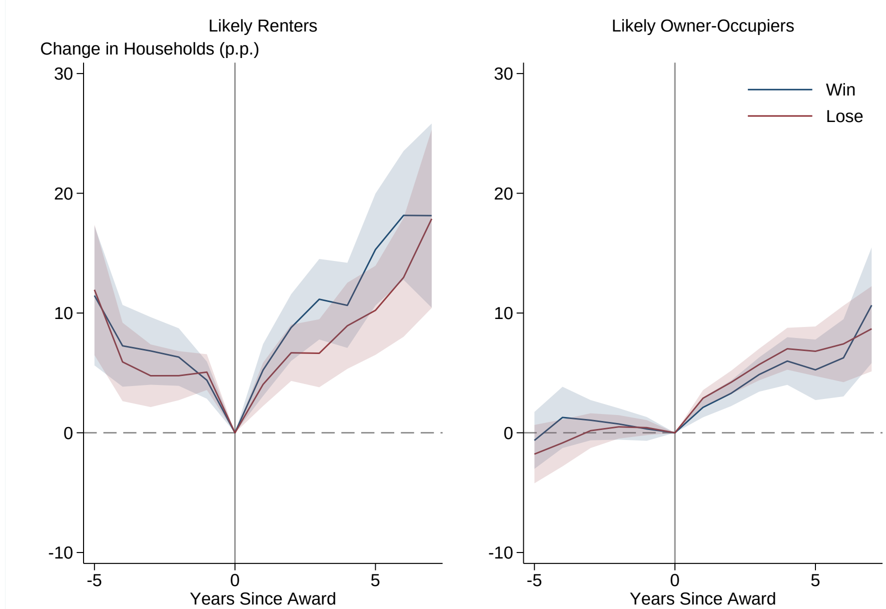

Notes: Panel A plots win and lose effects on the number of households (change in percentage points) in the same Census block group as the LIHTC application who reside in multi-family or single-family residences. Panel B reports the same effects but for households who are likely renters versus likely owner-occupants. In both, the specification is an event-study analog of Equation 9 that includes non-applicant tracts: $\Delta \log \mathrm{E}[Y_{i,t+k}] = \alpha_k \mathrm{Win}_{it} + \beta_k \mathrm{Lose}_{it} + \boldsymbol{X}_{it}\boldsymbol{\gamma}_k$, for k quarters or years from the award. The specification is estimated via Poisson regression, and it includes controls X_{it} for the win probability, pre-award tract characteristics, and the baseline level of the outcome. In both panels, the color bands depict pointwise 95-percent confidence intervals, and standard errors are clustered at the tract level.

Figure A8: Effects of Tax Credit Awards on the Subsidized Housing Stock

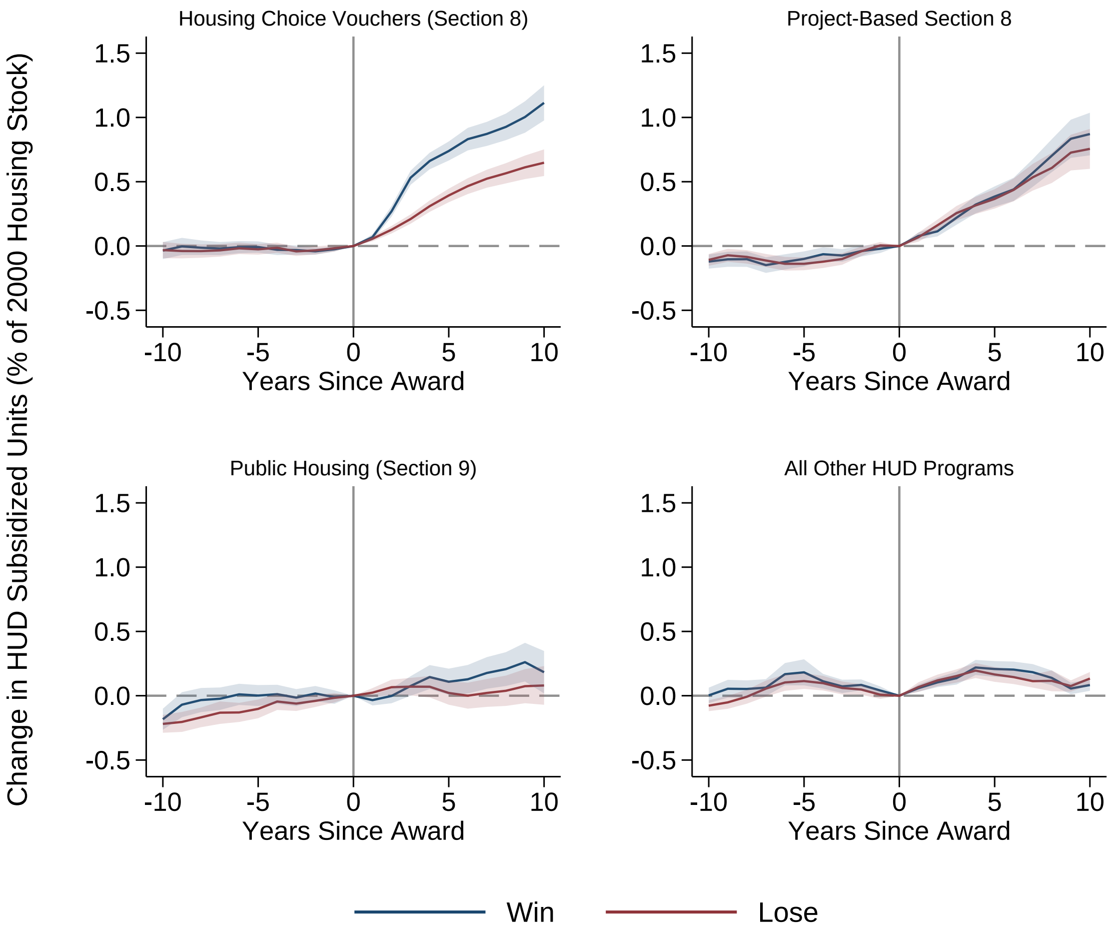

Notes: This figure plots win and lose effects on four categories of the subsidized housing stock. The outcome is defined by taking annual HUD counts of subsidized units by Census tract and dividing it by the count of households in 2000. The numerator is from the Picture of Subsidized Households database, and the denominator is from the 2000 Census. The specification is an event-study analog of Equation 9 that includes non-applicant tracts: $\Delta \log \mathrm{E}[Y_{i,t+k}] = \alpha_k \mathrm{Win}_{it} + \beta_k \mathrm{Lose}_{it} + \boldsymbol{X}_{it}\boldsymbol{\gamma}_k$, for k quarters or years from the award. The specification is estimated via Poisson regression, and it includes controls X_{it} for the win probability, pre-award tract characteristics, and the baseline level of the outcome. In both panels, the color bands depict pointwise 95-percent confidence intervals, and standard errors are clustered at the tract level.

Figure A9: Effects of Tax Credit Awards on Population by Income Decile

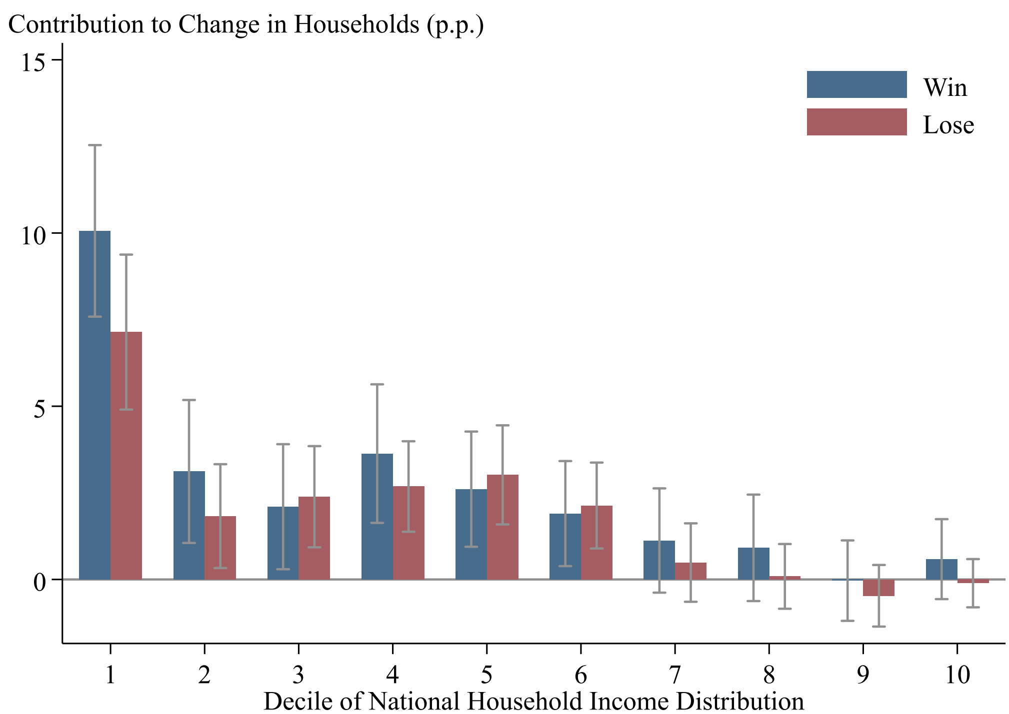

Notes: This figure plots win and lose effects on the number of households by income decile in the same Census block group as the LIHTC application. The time horizon is 10 years after the award, and I rescale the effects so that they represent contributions to the aggregate percentage change in the block-group household count. The specification is an event-study analog of Equation 9 that includes non-applicant tracts: $\Delta \log \mathrm{E}[Y_{i,t+k}] = \alpha_k \mathrm{Win}_{it} + \beta_k \mathrm{Lose}_{it} + \boldsymbol{X}_{it} \boldsymbol{\gamma}_k$, for k quarters or years from the award. The specification is estimated via Poisson regression, and it includes controls X_{it} for the win probability, pre-award tract characteristics, and the baseline levels of the outcomes. That is, to address measurement error in household income, all specifications includes control for all baseline income-decile household counts. The gray bars depict 95-percent confidence intervals, and standard errors are clustered at the tract level.

Figure A10: Replication of Tract-Level Event Study (Data Axle)

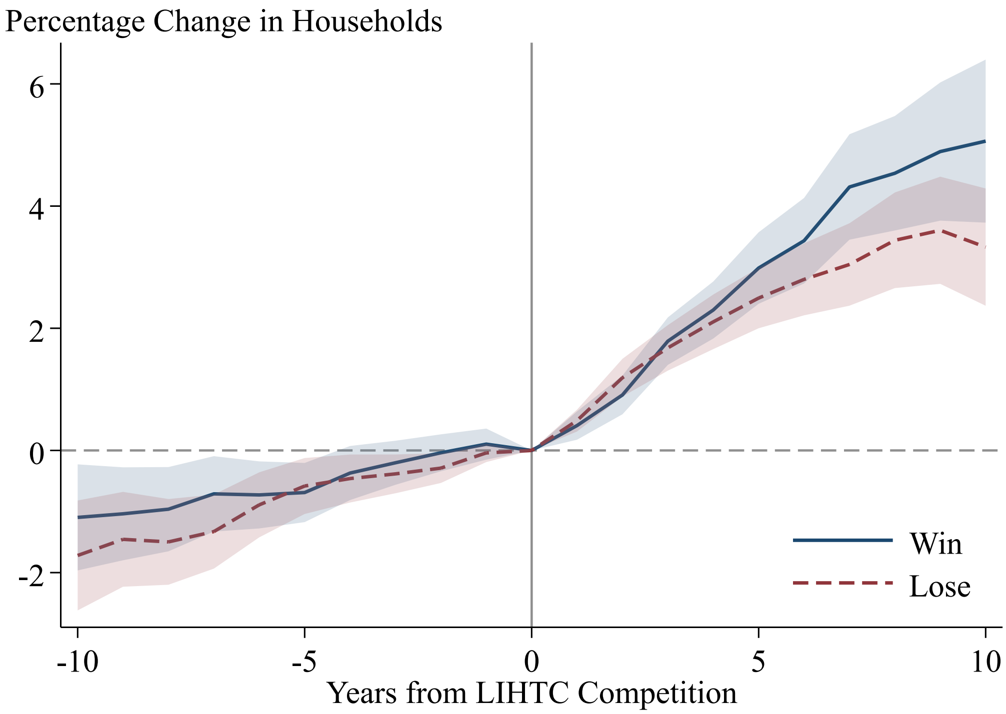

Notes: This plots win and lose effects effects on the change (in percentage points) in the household count in the LIHTC application's Census tract. The specification is intended for comparison to Figure 7, and it includes controls for the win probability and pre-award tract characteristics. The color bands depict pointwise 95-percent confidence intervals, and standard errors are clustered at the tract level.

Figure A11: Neighborhood Effects of Tax Credit Awards, No Win-Probability Control

Panel A: Tract (USPS)

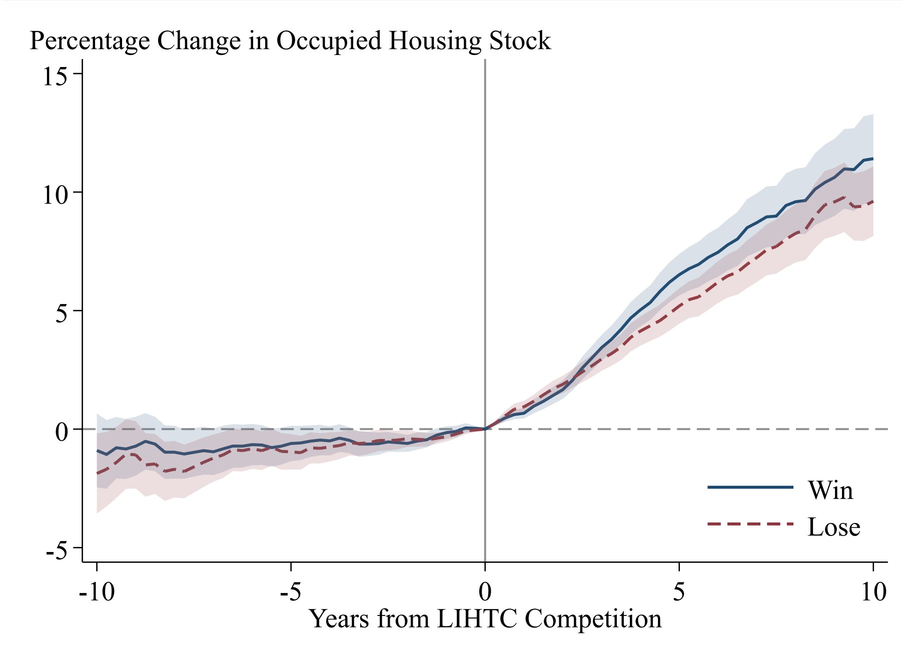

Panel B: Block Group (Data Axle)

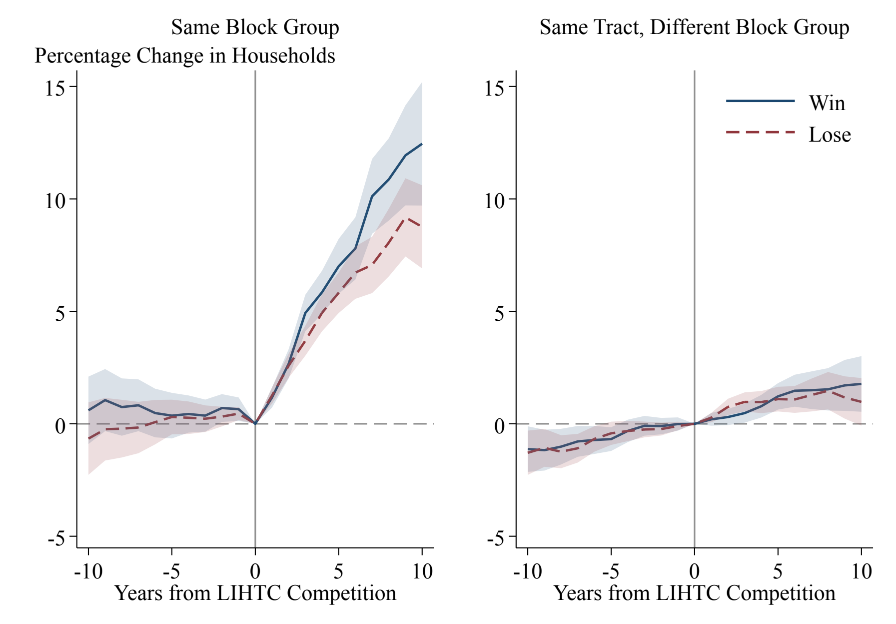

Notes: Panel A plots win and lose effects on the occupied housing stock (change in percentage points) from a quarterly event-study analog of Equation 9 that includes non-applicant tracts. Panel B plots win and lose effects effects on the change (in percentage points) in the household count in the same Census block group as the LIHTC application, or in the same tract but other block groups. All specifications include controls for pre-award tract characteristics but exclude the win-probability control. See Figure 7 for comparison. In both panels, the color bands depict pointwise 95-percent confidence intervals, and standard errors are clustered at the tract level.

Figure A12: Cohort-Specific Effects of LIHTC Application Wins and Losses

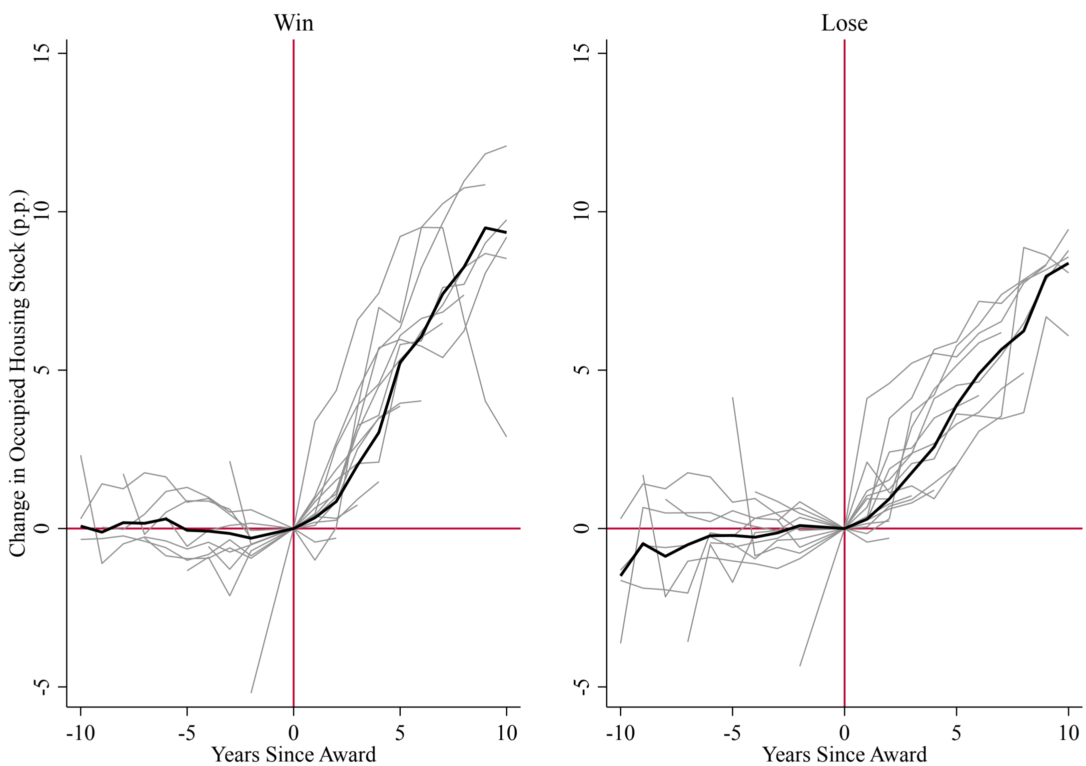

Notes: This figure plots estimates of cohort-specific quarterly coefficients $\beta_{k,y}$ from an event-study specification that is estimated on the subsample of tracts that have application in year y or that never apply. Estimated via a Poisson regression, the specification is $\Delta \log \mathrm{E}[Y_{i,t+k}] = \alpha_{k,y}\mathrm{Win}_{it} + \beta_{k,y}\mathrm{Lose}_{it} + \boldsymbol{X}_{it}\boldsymbol{\gamma}_{k,y}$ for each time horizon k and cohort y. Each grey line traces a path of coefficients for the same cohort y over the quarters k relative to the quarter of application. The solid black line displays the precision-weighted averages of these cohort-specific coefficients at each event time horizon.

Figure A13: Application Supply Responses to Changes in Eligibility for Basis Boost: Falsification Test

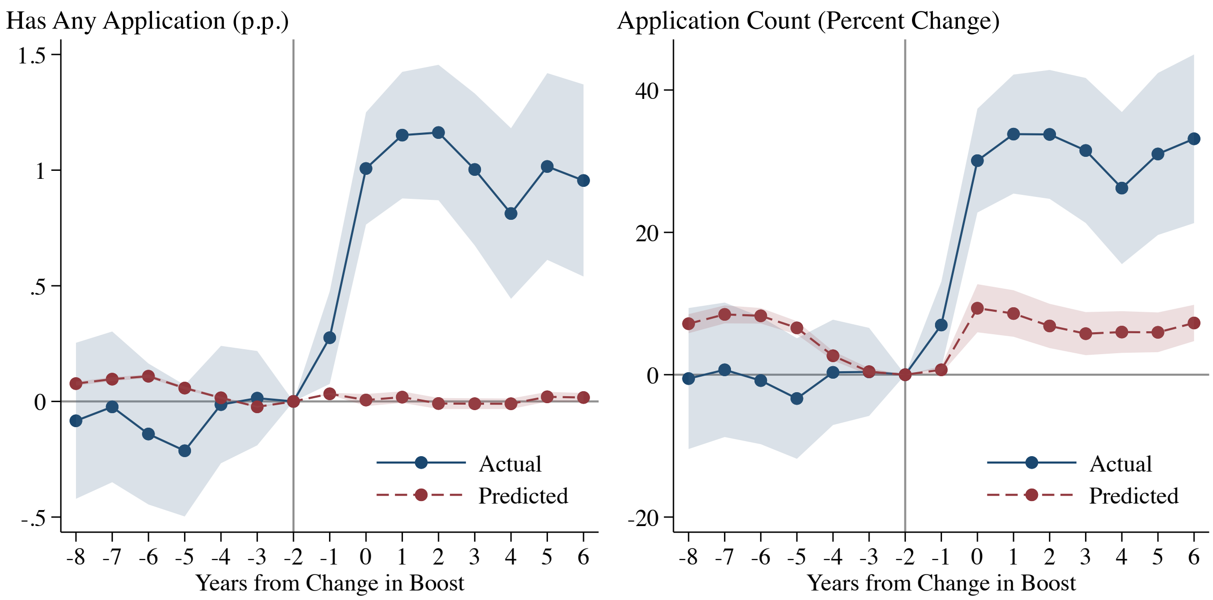

Notes: This figure plots event-study coefficients from Equation 10 for various application supply outcomes. The "actual" coefficients plot responses for actual application-supply outcomes. The "predicted" coefficients plot responses for an outcome constructed as the predicted values of year-specific regressions on the QCT running variables. Entry and exit from boost are assumed to have symmetric effects. All specifications are estimated by Poisson regression and include state–year fixed effects as controls. The bands show pointwise 95-percent confidence intervals, with standard errors clustered by tract.

Figure A14: Application Supply Responses to Changes in Eligibility for Basis Boost: Secondary Outcomes

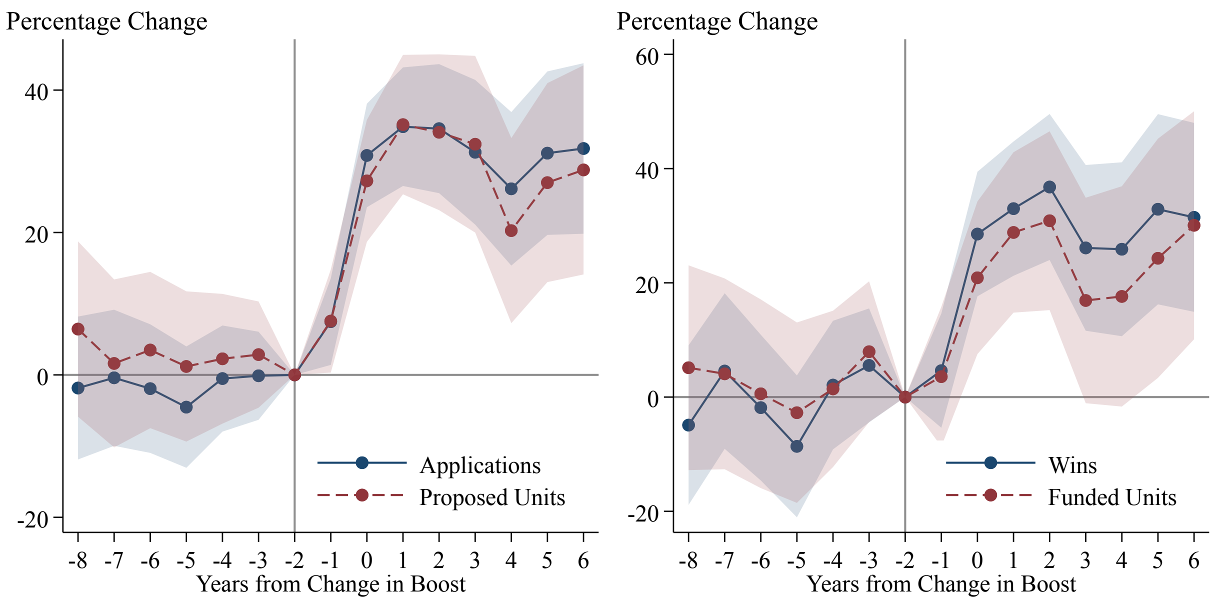

Notes: This figure plots event-study coefficients from Equation 10 for various application supply outcomes. Entry and exit from boost are assumed to have symmetric effects. All specifications are estimated by Poisson regression and include state–year fixed effects as controls. The bands show pointwise 95-percent confidence intervals, with standard errors clustered by tract.

Figure A15: Application Supply Responses to Changes in Eligibility for Basis Boost

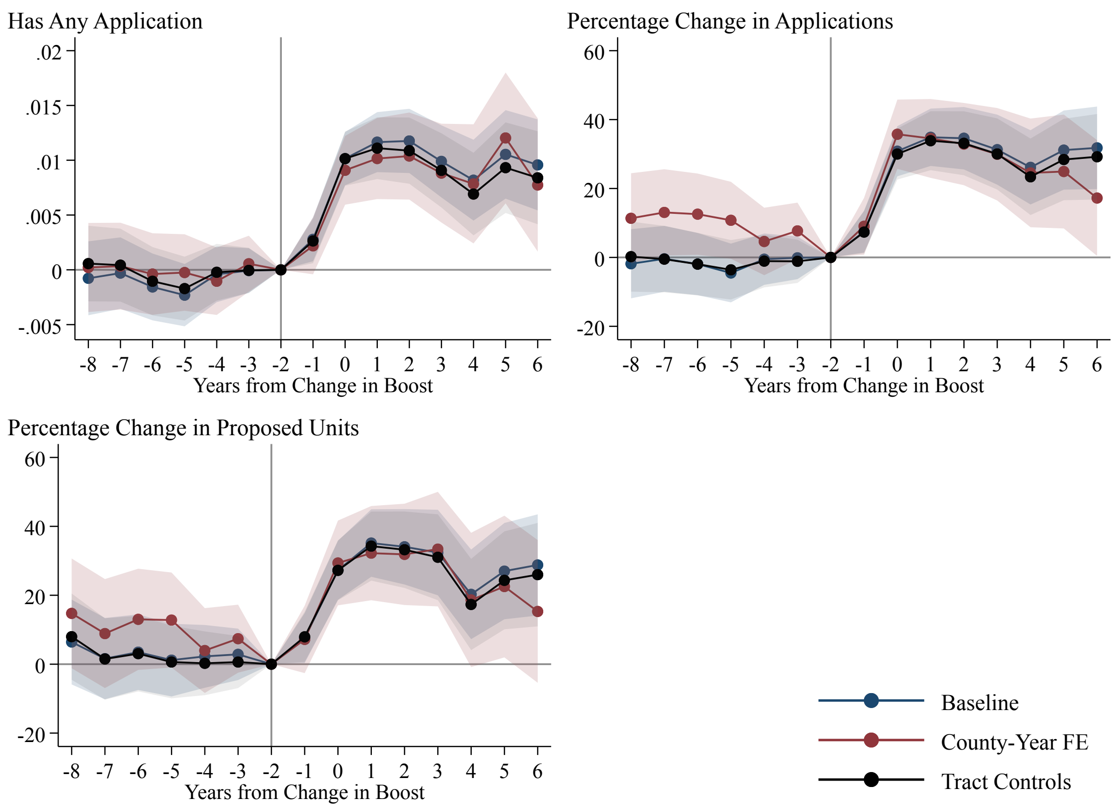

Notes: This figure plots event-study coefficients from Equation 10 for various application supply outcomes. Entry and exit from boost are assumed to have symmetric effects. In the baseline specification, I include state–year fixed effects. The other specifications introduce county–year fixed effects or tract controls interacted with year indicators. The tract controls are decile-group indicators for the tract's poverty rate, population density, and cumulative change in population from 1990 to 2000. The bands show pointwise 95-percent confidence intervals, with standard errors clustered by tract.

Figure A16: Application Supply Responses to Changes in Eligibility for Basis Boost: Include Always-Boosted Tracts

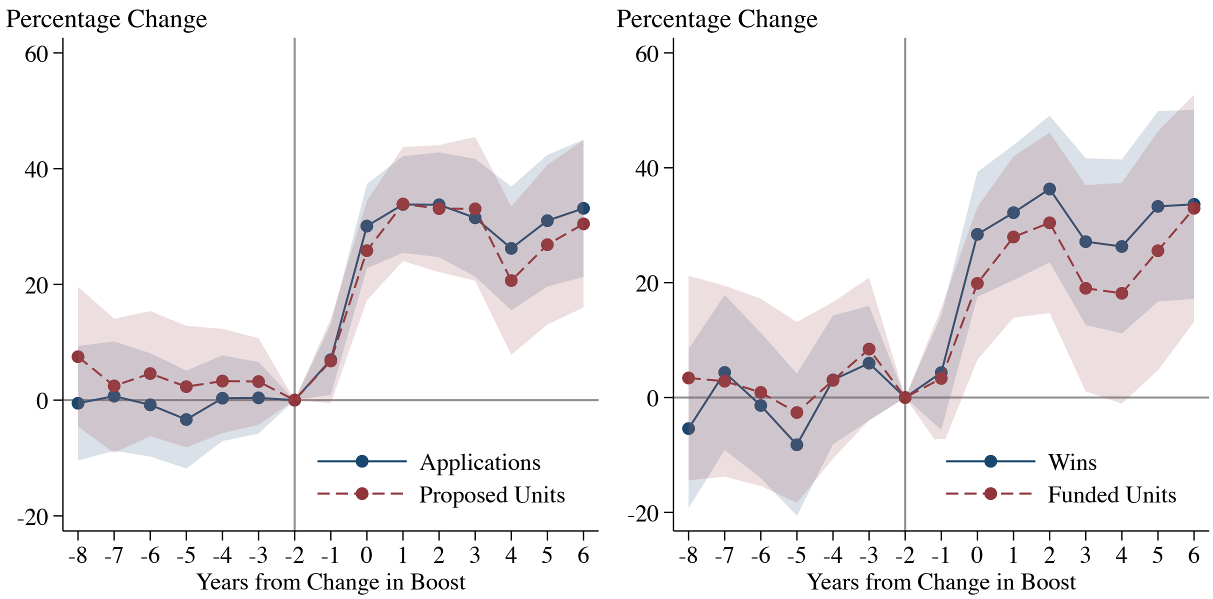

Notes: This figure plots event-study coefficients from Equation 10 for various application supply outcomes. The sole difference from Appendix Figure A14 is that here I include always–bosted tracts in the sample. Entry and exit from boost are assumed to have symmetric effects. All specifications are estimated by Poisson regression and include state–year fixed effects as controls. The bands show pointwise 95-percent confidence intervals, with standard errors clustered by tract.

Figure A17: Application Supply Responses to Changes in Eligibility for Basis Boost: QCT and DDA Effects

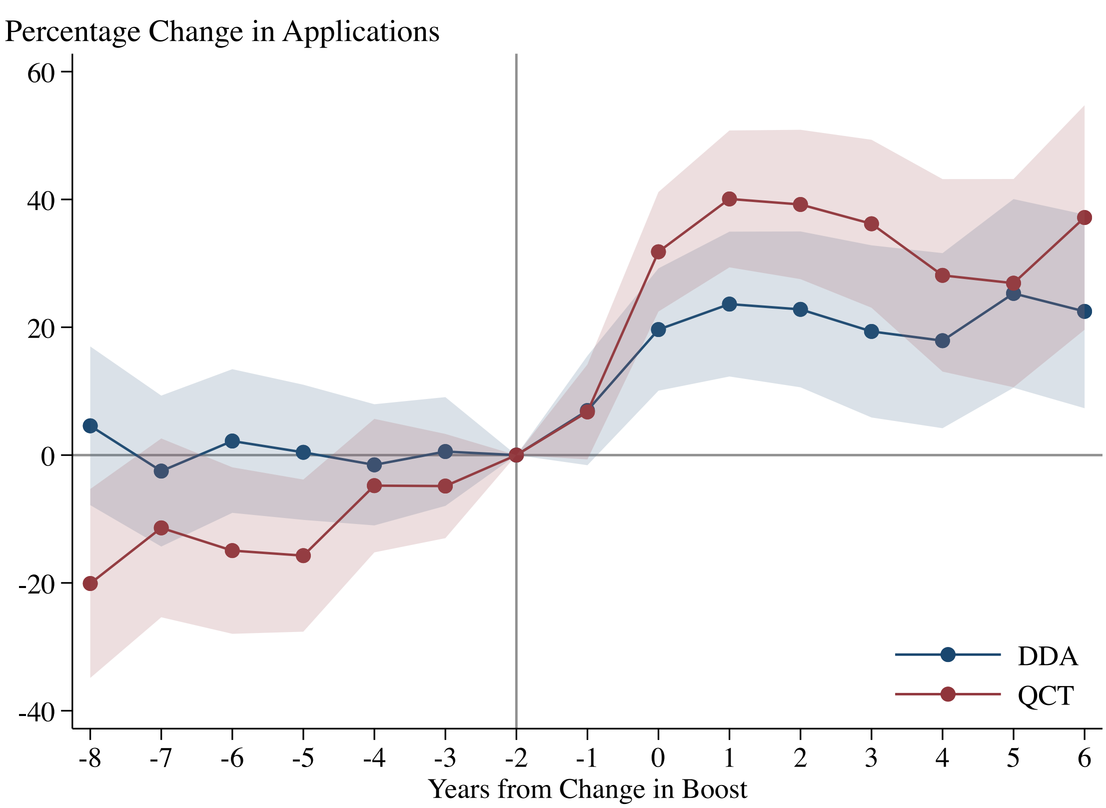

Notes: This figure plots event-study coefficients from Equation 10. Entry and exit from boost are assumed to have symmetric effects, but I allow the two causes of basis boosts (Qualified Census Tract [QCT] and Difficult Development Area [DDA]) to have different effects on application supply. The specification is estimated by Poisson regression and includes state–year fixed effects as controls. The bands show pointwise 95-percent confidence intervals, with standard errors clustered by tract.

Figure A18: LIHTC Application Volume Around the Qualified Census Tract Threshold: Alternative Outcomes

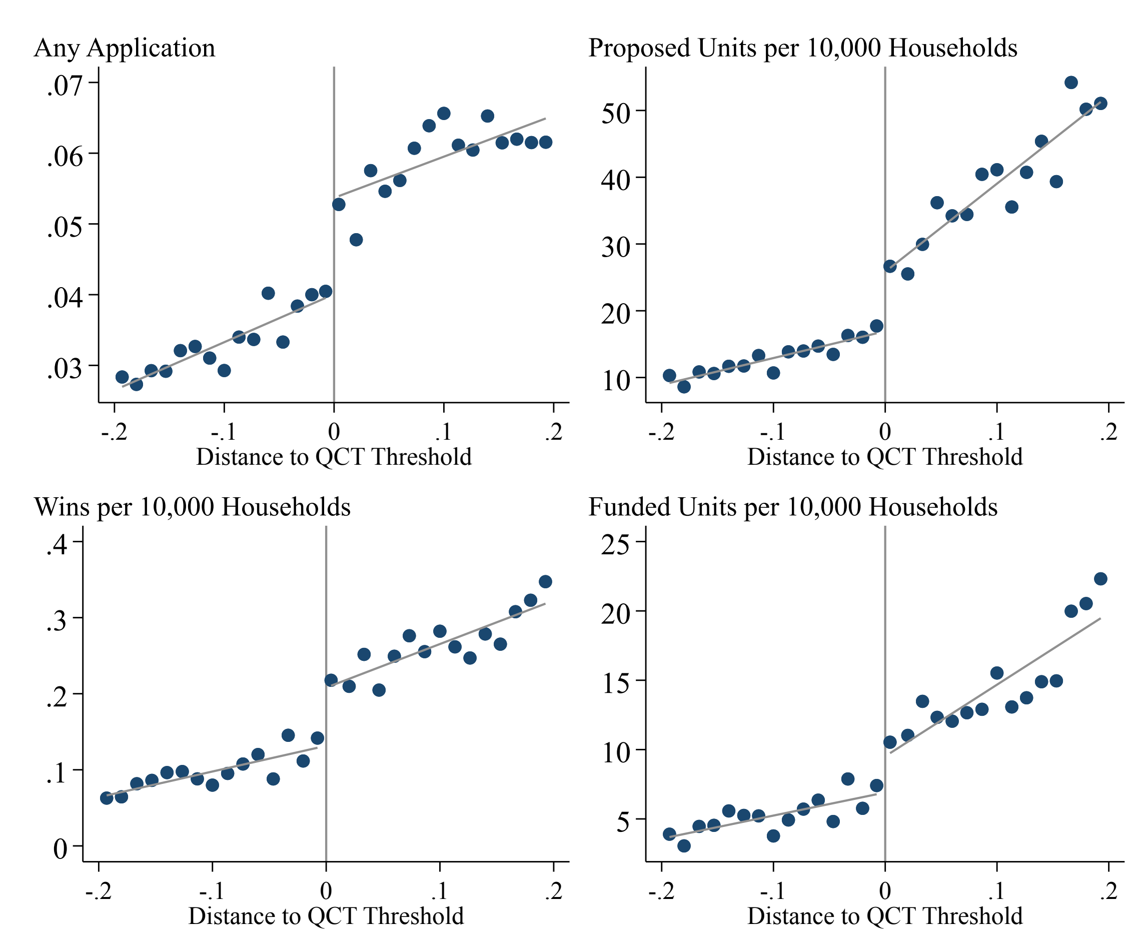

Notes: This figure shows the impact of the Qualified Census Tract (QCT) threshold on four measures of tract-level LIHTC application volume. In all panels, the running variable is a distance defined with respect to tract rank within its metropolitan area. All panels split the data into 15 equal-interval bins on either side of the QCT threshold.

Figure A19: Covariate Smoothness Around the Qualified Census Tract Threshold

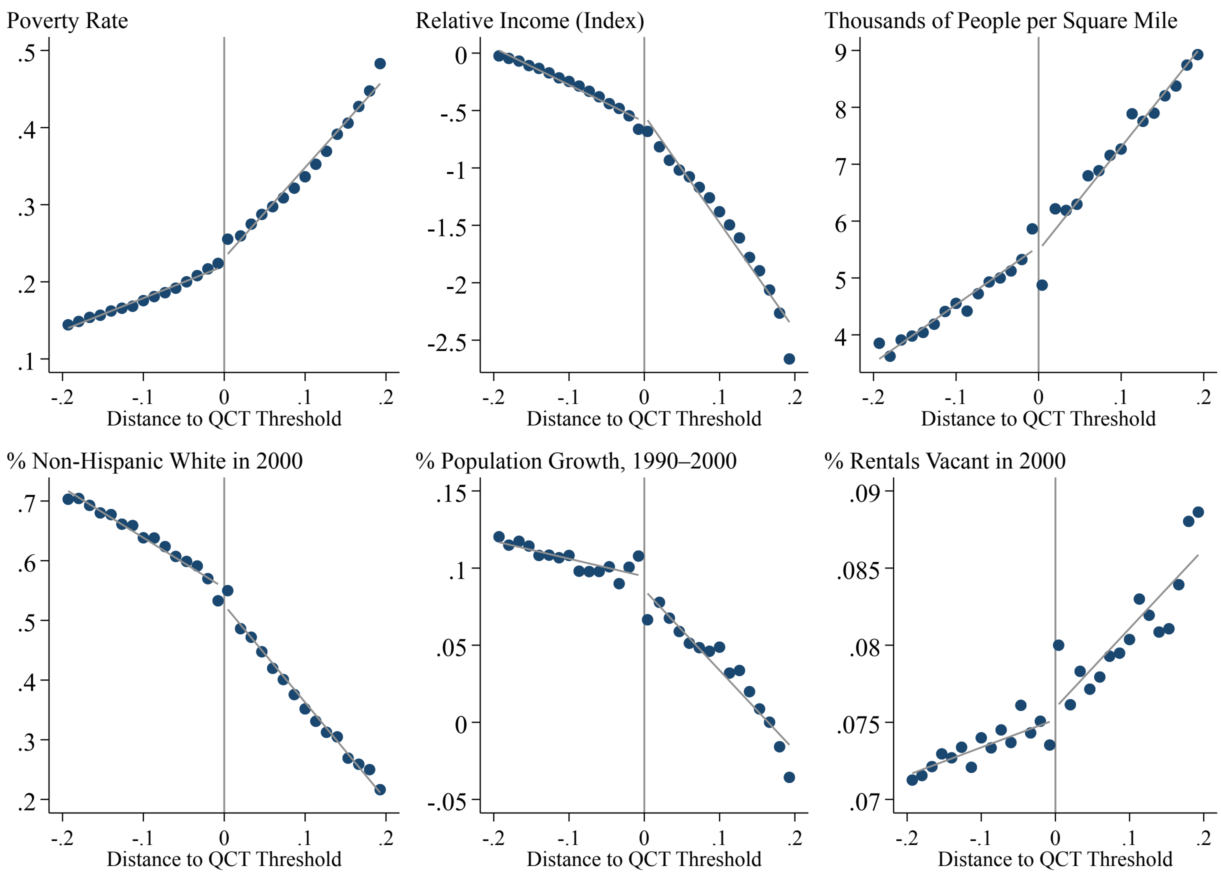

Notes: This figure examines the conditional expectation function for six covariates around the Qualified Census Tract (QCT) threshold. In all panels, the running variable is a distance defined with respect to tract rank within its metropolitan area. All panels split the data into 15 equal-interval bins on either side of the QCT threshold. The "relative income" measure is an index, standardized to zero mean and unit standard deviation, used in QCT assignment that compares a tract's median household income to its metro-area median.

Figure A20: Application Characteristics Around the Qualified Census Tract Threshold

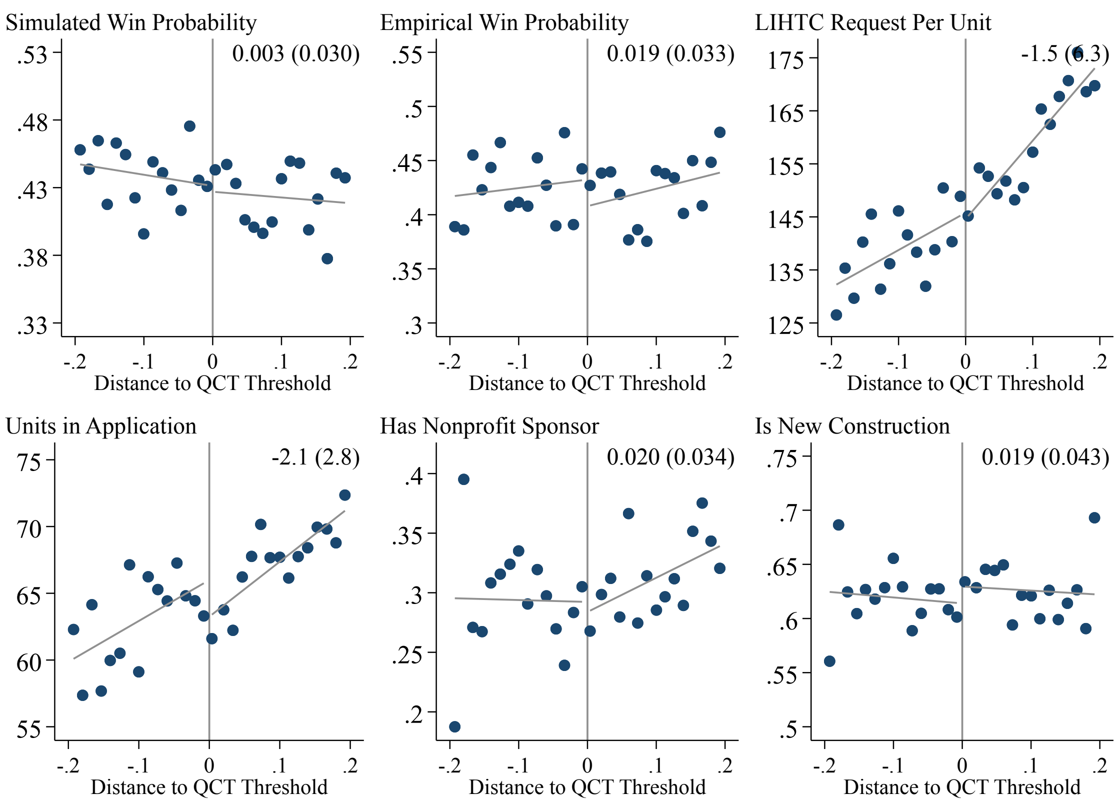

Notes: This figure shows average application characteristics around the threshold of Qualified Census Tract (QCT). In all panels, the running variable is a distance defined with respect to tract rank within its metropolitan area, and the data is split into 15 equal-interval bins on either side of the QCT threshold. The RD estimate and standard error corresponding to the plotted outcome is reported in the upper-right corner of each panel. LIHTC request per unit is in thousands of dollars.

Figure A21: Nonparametric Estimates of Bunching at Rent Score Kink

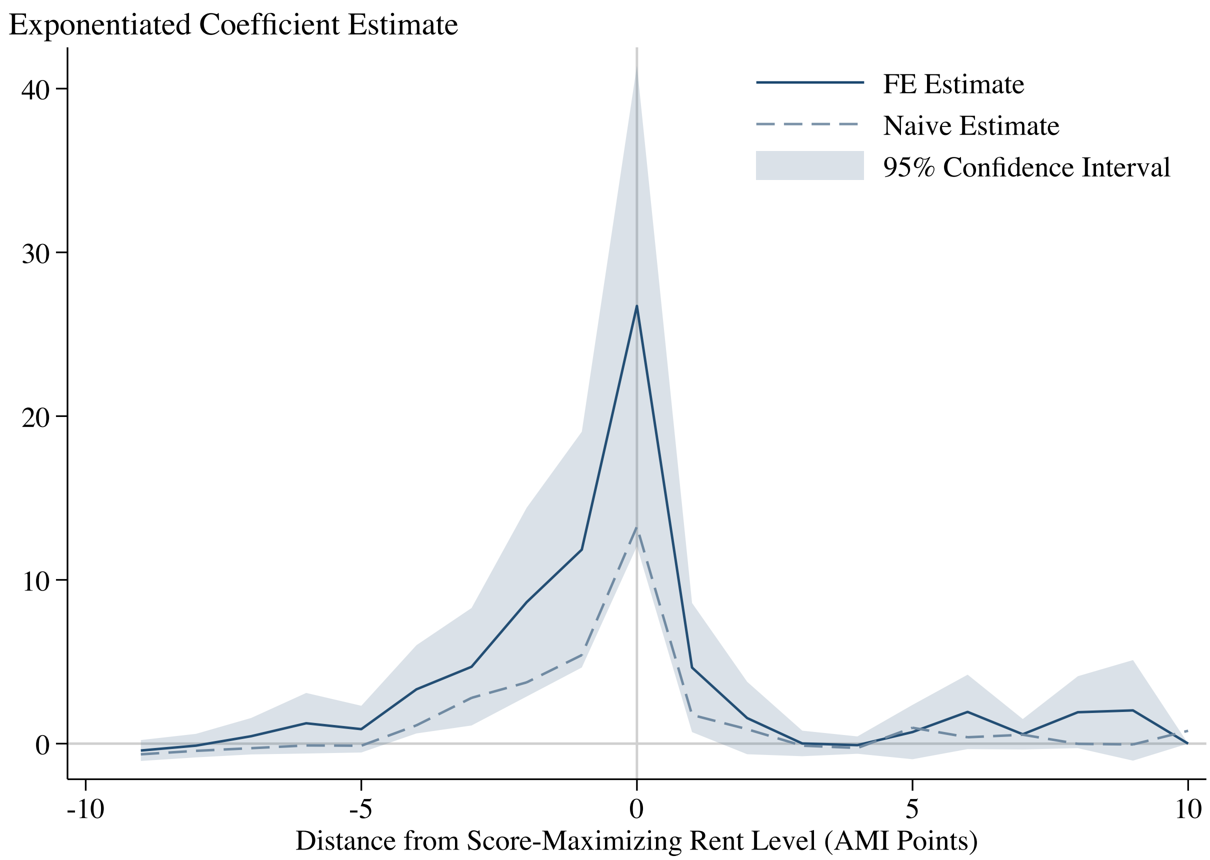

Notes: This figure depicts estimated effects of kinked incentives in Qualified Allocation Plans for setting rents in LIHTC units. The coefficients come from a Poisson regression $\log \mathrm{E}[A_{rm}] = \alpha_r + \alpha_m + \sum_s \beta_s 1[r - m = s]$, where A_{rm} is the count of applications with rent level r given that the score-maximizing rent level is m. The coefficients α_r and α_m are respectively fixed effects for an application's rent level and for an application round's score-maximizing rent level. "FE" bunching coefficients β_s are identified from heterogeneity in the score-maximizing rent level. The "naive" specification omits the fixed effects α_r and α_m, pooling variation. An application's rent level is defined relative to the local annual median income (AMI) of the household to whom its units would be "affordable," that is, no greater than 30 percent of income. To obtain counts, I bin the application data by rounding to full percentage points of AMI. To facilitate interpretation, the Poisson coefficients are exponentiated to $\exp(\beta_s) - 1$.

Figure A22: Heterogeneity in the Trade-Off Between Win Probability and Rental Income

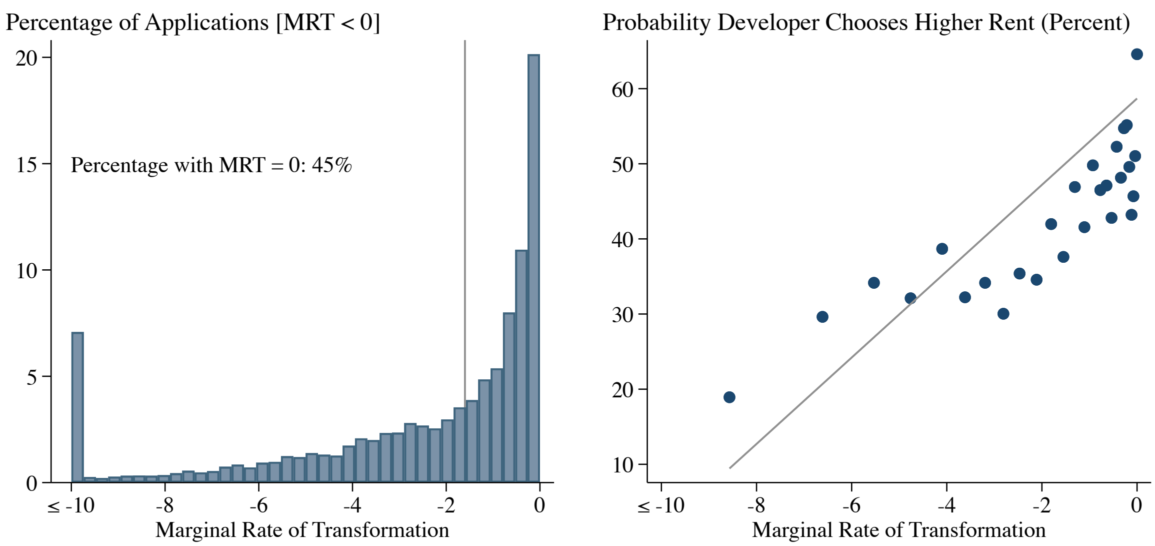

Notes: The left panel is a histogram of the marginal rate of transformation (MRT), local to the developer's actual choice. The MRT is defined as a ratio of the percentage-point change in win probability to the change in the present discounted value of rental income per unit, measured in thousands of dollars. To aid visualization, I winsorize the distribution at MRT < -10, and I drop the all observations with MRT $= 0$. The shares reported on the vertical axis are thus conditional on a non-zero MRT. The right panel is a binned scatterplot of the probability the developer accepts a choice giving them a higher rental income but lower win probability, scattered with respect to the MRT associated with this choice.

Figure A23: Income Comparison of LIHTC Residents to Other Locals

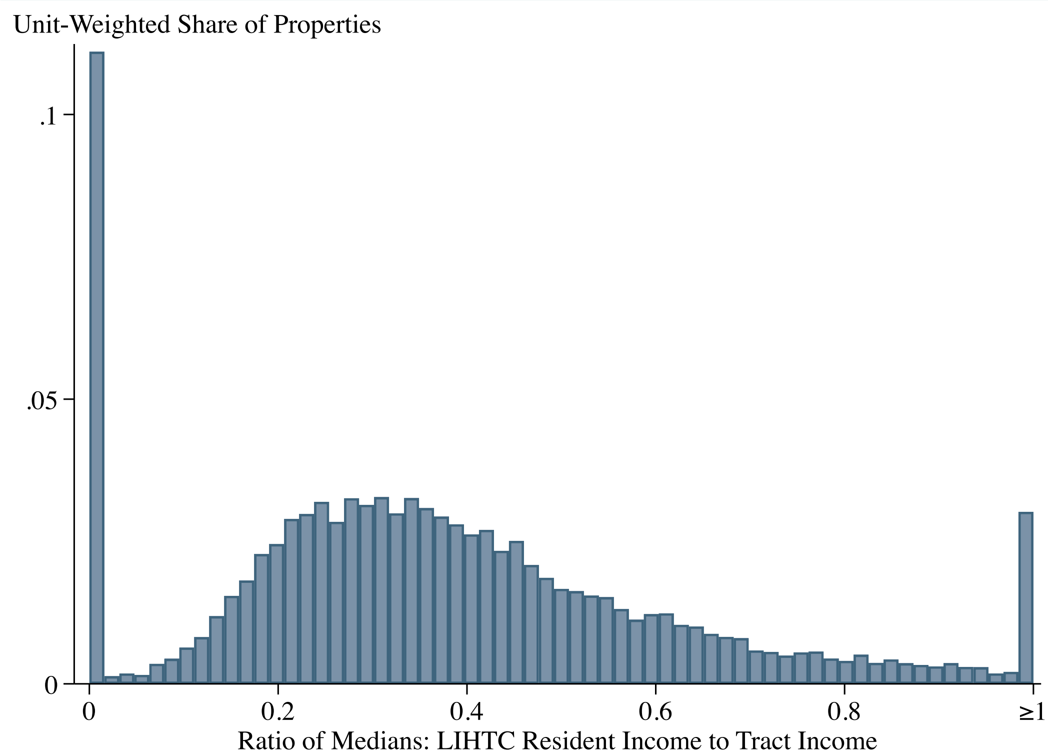

Notes: This figure plots the distribution of a ratio of median incomes. The ratio compares the median household income of LIHTC residents at the property level to the median in the corresponding Census tract. Median-income data for LIHTC residents is as of 2019; the tract data come from the ACS centered on 2019. The histogram is weighted by each LIHTC property's total unit count.

Figure A24: Map of LIHTC Applications

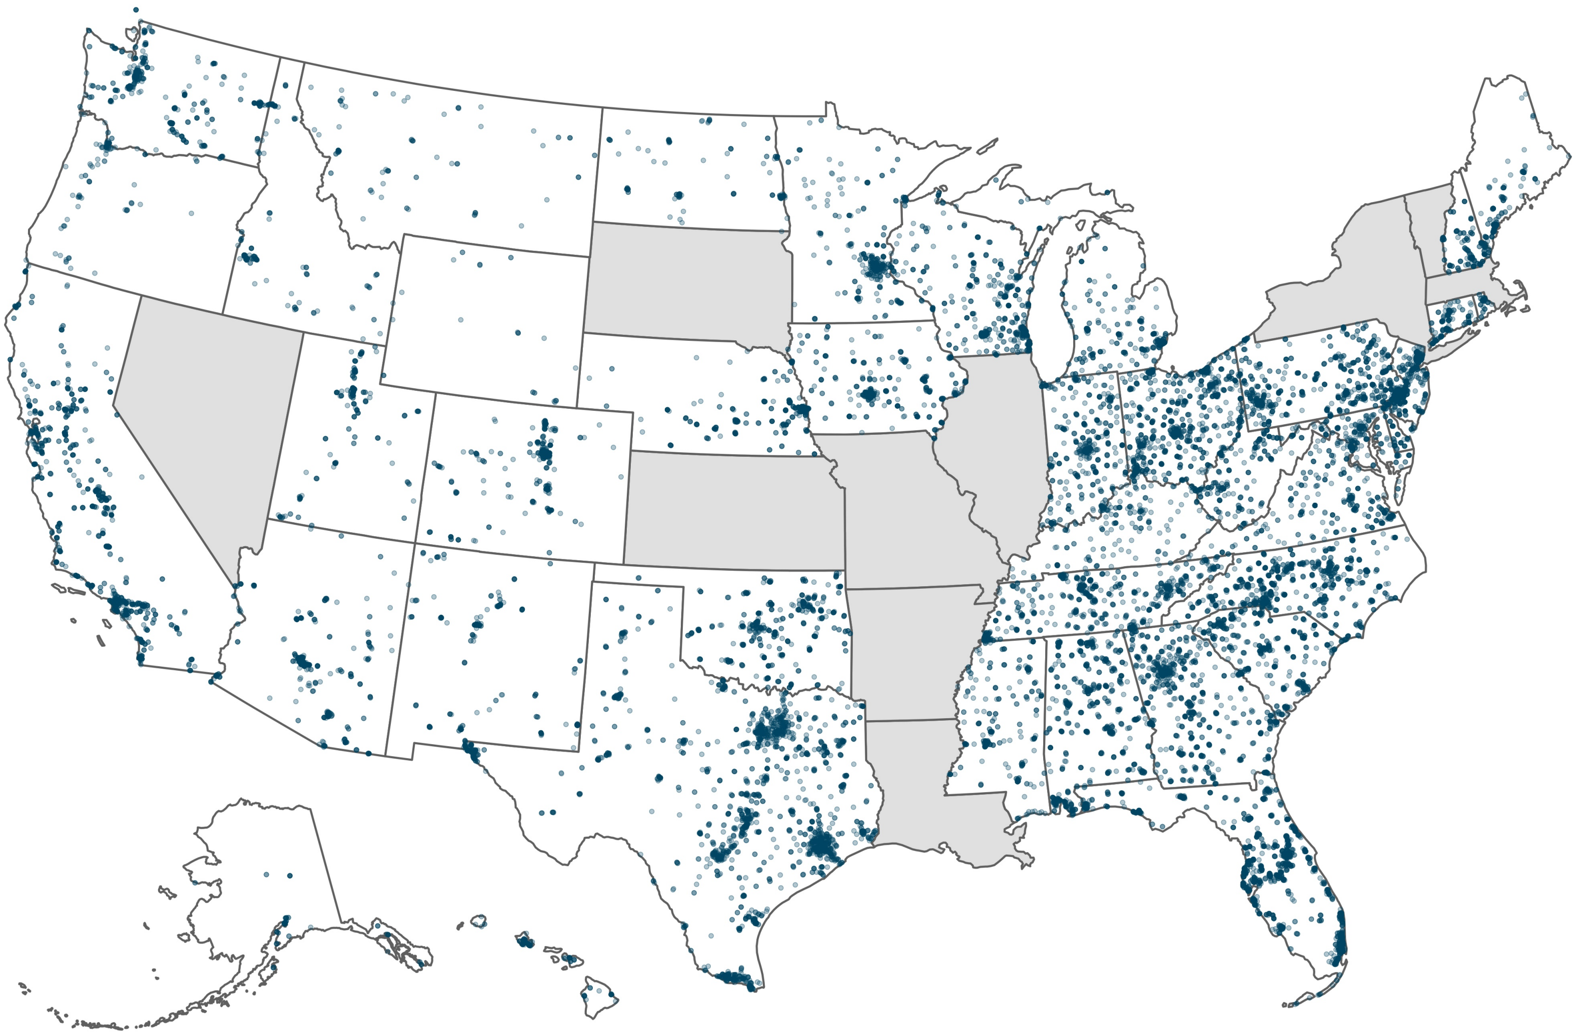

Notes: This figure displays a map of the tax-credit applications. States are shaded in grey if no data are available.

Figure A25: Counts of Applications and Proposed Units by Year and Outcome

Panel A: Applications

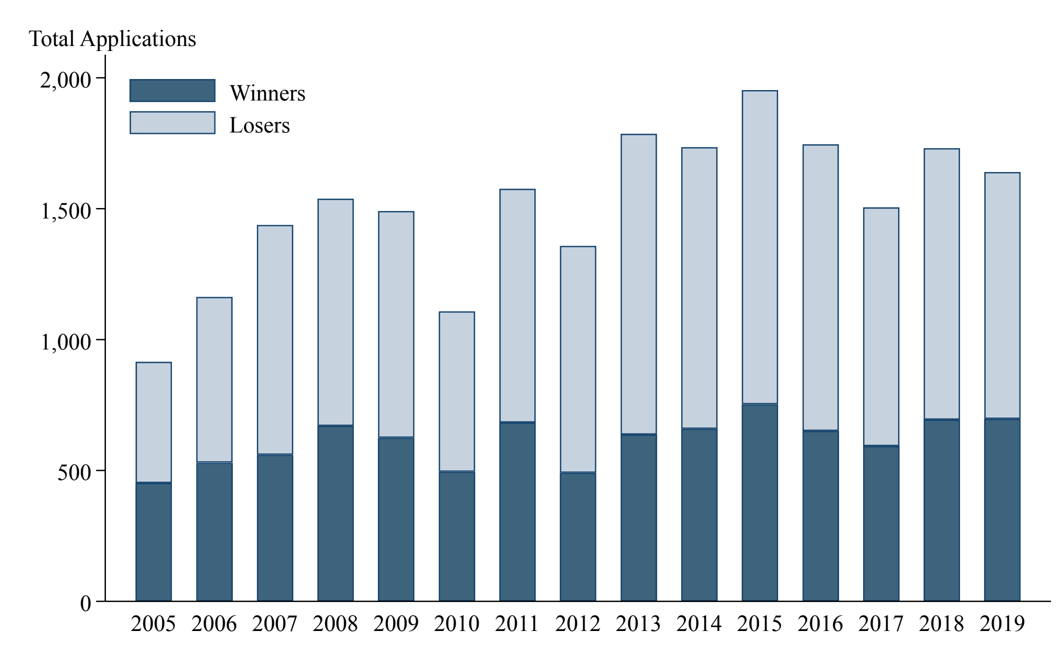

Panel B: Proposed Units

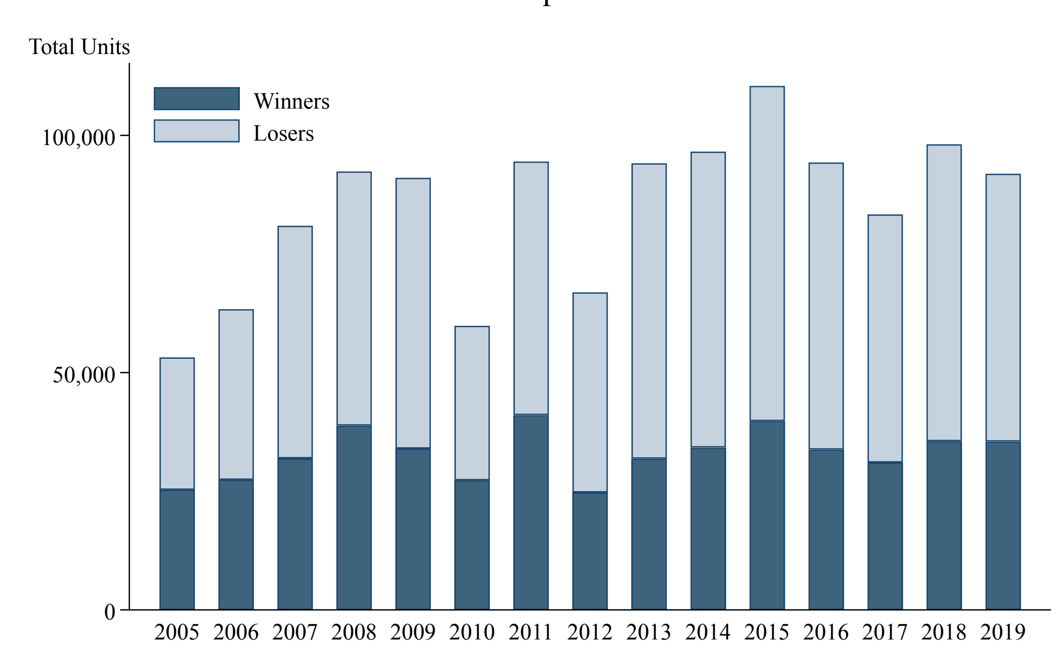

Notes: This figure plots, in Panels A and B respectively, the counts of applications and proposed total units (both income-restricted and unrestricted) in the data for each year and application outcome.

Figure A26: Round-Level Distributions of Average Simulated Win Probability and Explained Share of Variance in Wins

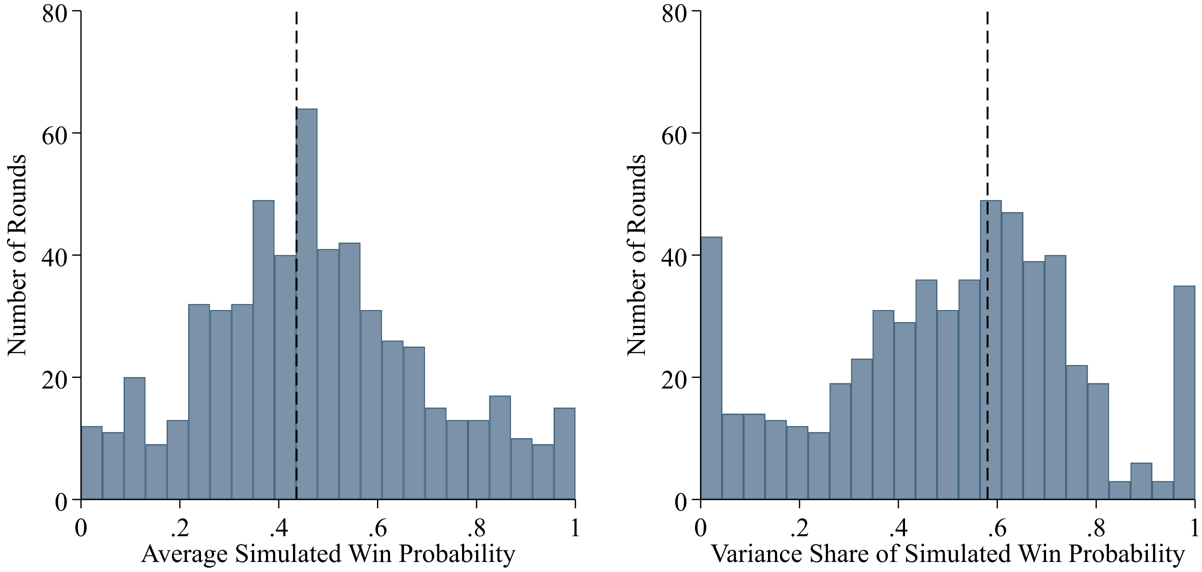

Notes: This figure displays, in the left panel, the distribution of average simulated win probabilities at the level of tax-credit competition round. In the right panel, the figure displays the round-level distribution of the share of variance of wins that is explained by variance in simulated win probabilities, that is, $\text{Var}(\hat{p}_i)/\text{Var}(\text{Win}_i)$. The vertical lines denote unweighted averages of the average simulated win probability (left) and the explained share of variance in wins (right).

Figure A27: Analysis of Self-Scores

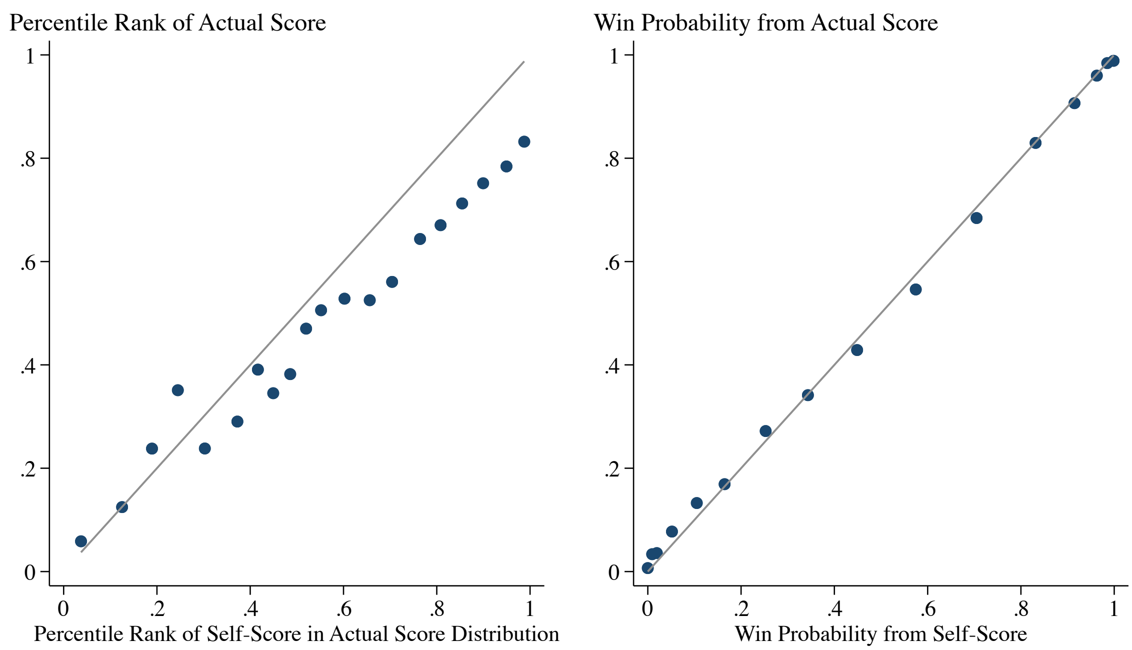

Notes: This figure displays, in the left panel, a binned scatterplot of actual application scores versus self-scores. I transform scores into percentile ranks within the distribution of actual scores. In the right panel, the figure displays a binned scatterplot comparing estimates of the win probability using actual and self-scores. The procedure to compute win probabilities using self-scores is described in Appendix B.

Table A1: Heterogeneous Application Supply Responses to Qualified Census Tract Threshold

	Main Effect		Interaction Effect		
	Est. (1)	SE (2)	Est. (3)	SE (4)	Median (5)
Panel A: High Population Growth					
QCT Threshold	0.592***	(0.084)	-0.291***	(0.083)	0.043
Panel B: High Rental Vacancy					
QCT Threshold	0.435***	(0.100)	0.019	(0.088)	0.063
Panel C: High Population Density					
QCT Threshold	0.467***	(0.084)	-0.062	(0.083)	3.85
Panel D: High Poverty Rate					
QCT Threshold	0.501***	(0.116)	-0.156	(0.124)	0.230
Panel E: High Non-Hispanic White Share					
QCT Threshold	0.383***	(0.090)	0.123	(0.087)	0.539

Notes: This table estimates heterogeneous effects of the Qualified Census Tract (QCT) threshold on application supply according to Census tract characteristics. In all panels, the outcome measure is the application count in that tract–year. Effects are estimated using a local-linear Poisson regression with a triangular kernel of bandwidth 0.2 around the QCT threshold. Main effects of the QCT threshold are reported in Columns 1 and 2, and interaction effects are reported in Columns 3 and 4. In estimating heterogeneous effects, I split Census tracts at the median of the named variable, with the median reported in Column 5. $^{*} = p < 0.10$, $^{**} = p < 0.05$, $^{***} = p < 0.01$.

Table A2: Estimating Developer Win Valuations from Bidding Behavior: Binding Subsample

	(1) OLS	(2) IV	(3) Cond. Logit.	(4) + Ctrl. Funct.
Win Probability	0.407***	1.801***	1.199***	6.277***
	(0.029)	(0.057)	(0.091)	(0.240)
Log Average Rent	1.029***	3.667***	2.970***	11.956***
	(0.073)	(0.144)	(0.216)	(0.548)
Applications	4,274	4,274	4,268	4,268
Marg. Rate of Substitution	1.590	1.282	1.558	1.199
	(0.059)	(0.029)	(0.057)	(0.028)
Mean Win Value Per Unit	$32,426	$40,235	$33,090	$43,019
	(1,244)	(971)	(1,247)	(931)
Developer Incidence Share	0.255	0.317	0.260	0.338
	(0.010)	(0.009)	(0.011)	(0.009)

Notes: This table reports estimates of coefficients in Equation 13. I use the subsample of the application data in which I estimate that rent regulations bind. In Column 4, I take a control-function approach to instrument for the win probability in the conditional logit. Marginal rates of substitution are calculated as a ratio of coefficients; see the text for detail on the other calculations. Standard errors are clustered by application. * = $p < 0.10$, ** = $p < 0.05$, *** = $p < 0.01$.

Table A3: Structural Parameter Estimates

Group	Parameter	Estimate	Standard Error
Outside Option	Intercept	-0.064	0.002
	Poverty Rate	-0.029	0.000
	Log Population Density	0.197	0.002
	Log Unit Count	0.277	0.003
	Log Credits Per Unit	-0.347	0.002
	Permanent Unobs.	0.924	0.010
	Reapplicant	0.396	0.003
Net Value of Win	Intercept	0.388	0.005
	Poverty Rate	-0.465	0.014
	Log Population Density	0.171	0.002
	Log Unit Count	-0.263	0.003
	Log Credits Per Unit	-0.574	0.006
	Permanent Unobs.	0.731	0.010
	Reapplicant	0.513	0.007
Entry Cost	Intercept	0.584	0.003
	Poverty Rate	-0.212	0.003
	Log Population Density	-0.015	0.000
	Log Unit Count	-0.203	0.003
	Log Credits Per Unit	-0.103	0.001
	Permanent Unobs.	-0.073	0.001
	Reapplicant	-0.386	0.004
Win Probability	Parameter 1	1.270	0.004
	Parameter 2	0.686	0.003
	Parameter 3	0.759	0.008
Temp. Unobs. Dispersion	Application Cost	0.005	0.000
	Outside Option	0.011	0.000
	Bidding	0.097	0.007
Log Units	Mean	3.804	0.020
	Standard Deviation	0.497	0.006
Log Credits Per Unit	Mean	12.142	0.018
	Standard Deviation	0.492	0.009

Notes: This table reports estimates of the model structural parameters, along with standard errors. $^{*} = p < 0.10$, $^{**} = p < 0.05$, $^{***} = p < 0.01$.

Table A4: Model Fit

		Data		Simulation		
Group	Moment	Estimate	SE	Estimate	SE	$\|t\|$
Application: Cross-Sectional Regression	Intercept	0.030	0.001	0.036	0.003	2.07
	Poverty Rate	0.146	0.007	0.000	0.022	6.37
	Log Population Density	-0.000	0.000	0.007	0.001	5.30
Application: Average Derivative	Application Probability	0.014	0.002	0.033	0.002	6.17
	Win Probability	0.003	0.030	-0.053	0.003	1.84
	Log Request	-0.010	0.043	-0.008	0.007	0.04
	Log Units	-0.030	0.044	0.079	0.009	2.43
Post-Loss Behavior	Intercept	0.254	0.004	0.235	0.005	2.88
	Win Probability	0.024	0.013	0.130	0.037	2.72
	Poverty Rate	-0.031	0.029	0.110	0.042	2.77
	Log Population Density	0.030	0.003	0.023	0.004	1.70
	Log Units	-0.022	0.009	0.042	0.015	3.62
	Log Credits Per Unit	-0.048	0.009	-0.040	0.013	0.48
	Reapplicant	-0.027	0.009	-0.028	0.014	0.03
Reapplication Choice	Intercept	0.367	0.004	0.408	0.006	5.36
	Win Probability	0.040	0.015	0.252	0.042	4.75
	Poverty Rate	-0.017	0.032	0.072	0.048	1.53
	Log Population Density	0.013	0.003	0.032	0.004	3.65
	Log Units	-0.074	0.010	-0.043	0.017	1.56
	Log Credits Per Unit	0.020	0.010	0.032	0.015	0.68
	Reapplicant	0.117	0.010	-0.006	0.014	6.51
Bid Choice	Win Probability	0.386	0.025	0.353	0.011	1.19
	Log Average Rent	0.755	0.079	0.817	0.028	0.74
	Win Prob. × Poverty	-0.361	0.171	-0.464	0.067	0.56
	Log Avg. Rent × Poverty	-1.092	0.553	0.008	0.206	1.86
	Win Prob. × Density	-0.059	0.014	-0.021	0.007	2.43
	Log Avg. Rent × Density	-0.097	0.043	-0.168	0.019	1.53
Distributional Parameters	Win Probability Mean	0.436	0.003	0.412	0.026	0.91
	Win Probability SD	0.387	0.002	0.347	0.019	2.11
	Log Units Mean	4.005	0.004	4.011	0.047	0.13
	Log Units SD	0.511	0.002	0.617	0.033	3.16
	Log Request Mean	12.280	0.004	12.282	0.039	0.04
	Log Request SD	0.573	0.003	0.508	0.027	2.35
	Win Probability AR1 Intercept	0.274	0.007	0.273	0.007	0.14
	Win Probability AR1 Slope	0.393	0.018	0.394	0.018	0.07

Notes: This table reports estimates and standard errors for the data moments to be matched in the structural model, alongside their simulation analogs. The final column reports the absolute value of the t-statistic testing equality of each data and simulation moment. $^{*} = p < 0.10$, $^{**} = p < 0.05$, $^{***} = p < 0.01$.

Table A5: Sensitivity Analysis

	(1) App. Supply: Intercept	(2) App. Supply: Derivative	(3) Win–Lose: Intercept	(4) Reapplication: Intercept	(5) Bidding: Win Prob. Coef.	(6) Bidding: Rent Coef.
Panel A: Model Primitives						
Outside Option	-0.264	0.013	0.009	-0.561	-0.169	0.045
Win Value	-0.162	-0.331	-0.045	0.081	0.150	0.074
Entry Cost	-0.006	-0.083	-0.007	0.095	0.026	-0.003
Panel B: Incidence						
Household Share	-0.041	0.035	-0.018	-0.229	-0.032	-0.009
Developer Share	0.021	-0.003	0.017	0.122	-0.005	0.013
Entry Cost Share	0.019	-0.032	0.002	0.108	0.037	-0.003
Panel C: Impacts						
Displacement Rate	-0.016	0.007	0.002	-0.039	-0.001	-0.010
Marginal Cost Per Unit	-0.5	1.4	-3.0	-1.1	8.9	0.2

Notes: This table reports the results of a sensitivity analysis following Andrews et al. (2017). I consider centered one-percentage-point perturbations to the regression coefficients used as empirical moments. I re-estimate the structural model with 0.5-p.p. upward and downward perturbations to each moment. This centers the sensitivity analysis on the actual estimates and circumvents non-smoothness concerns. Marginal costs per unit are in thousands of dollars. $^{*} = p < 0.10$, $^{**} = p < 0.05$, $^{***} = p < 0.01$.

Table A6: More Estimates of Model Primitives and Potential-Applicant Characteristics

	(1)	(2)	(3)	(4)	(5)	(6)
	Winners			Losers		
	Median	P25	P75	Median	P25	P75
Panel A: Model Primitives as a Share of Basis						
Ex-Ante Value	1.51	1.16	1.80	1.34	1.05	1.71
Outside Option	1.35	0.93	1.67	1.22	0.82	1.62
Win Value	1.72	1.43	2.21	1.23	0.64	1.62
Application Cost	0.17	0.10	0.25	0.25	0.03	0.33
Panel B: Model Primitives in Thousands of Dollars Per Unit						
Ex-Ante Value	178	121	264	211	144	311
Outside Option	149	98	218	186	104	276
Win Value	211	154	288	178	112	248
Application Cost	19	9	27	30	4	57

Notes: This table reports estimates of the model primitives, the ex-ante value of the application value function $\mathrm{E}[V^A(\boldsymbol{s}_{it}, \boldsymbol{\varepsilon}_{it}) \mid \boldsymbol{s}_{it}]$, the outside option $\pi_0(\boldsymbol{s}_{it})$, the win value $\pi_1(\boldsymbol{s}_{it})$, and the application cost $\kappa(\boldsymbol{s}_{it})$.

Table A7: Selectivity Correction for New-Unit Rent Hedonic Estimation

	(1)	(2)	(3)	(4)
Log Median Rent, All Units	0.967 (0.007)	0.997 (0.008)	0.887 (0.013)	0.864 (0.016)
Selection Correction		0.076 (0.008)	0.069 (0.008)	0.055 (0.009)
Controls			✓	✓
State FE				✓
Clusters	14,615	14,579	14,544	14,544
First-Stage F Stat.		1,111,658	69,777	51,145
Mean Selectivity Bias		0.099	0.083	0.066

Notes: This table examines the bias from sample selection in the estimation of willingness to pay for new rental units by Census tract. The specification estimated is Equation 17. Bootstrap standard errors account for first-stage estimation of the selectivity correction. $^{*} = p < 0.10$, $^{**} = p < 0.05$, $^{***} = p < 0.01$.

Table A8: Application Data Coverage: By Property and State–Year

	(1)	(2)	(3)	(4)	(5)	(6)
	Property Coverage			State–Year Coverage		
	Yes	No	Diff. (SE)	Yes	No	Diff. (SE)
Panel A: Property Characteristics						
Unit Count	60.85	66.45	-5.60*** (0.90)	65.67	61.10	4.56*** (0.93)
% Low-Income Units	89.7	89.9	-0.3 (0.6)	88.7	91.8	-3.1*** (0.6)
Monthly Rent	1,672.97	1,693.44	-20.48** (8.70)	1,684.57	1,684.08	0.49 (9.08)
Rents Below Fed. Max.	75.8	54.9	20.9*** (0.9)	70.5	53.2	17.3*** (0.9)
% New Construction	68.9	60.4	8.5*** (0.8)	67.3	58.5	8.8*** (0.9)
Family	61.7	55.6	6.1*** (1.0)	58.9	56.8	2.1** (1.0)
Elderly	44.3	33.2	11.2*** (1.1)	38.9	35.8	3.1*** (1.1)
Other	15.1	18.0	-2.9*** (0.6)	17.3	15.8	1.6** (0.6)
PDV Tax Credits Per Unit	278,327	161,329	116,997*** (9,014)	239,807	165,749	74,059*** (9,241)
Nonprofit	27.6	33.5	-5.9*** (0.8)	29.2	34.3	-5.1*** (0.9)
Panel B: Location Characteristics						
Mean Income Per Capita	27,539	27,554	-15 (330)	27,366	27,886	-520 (344)
% Poor	77.1	75.3	1.8*** (0.3)	76.3	75.8	0.5 (0.4)
% Less than HS	15.2	15.7	-0.4* (0.2)	15.5	15.3	0.2 (0.3)
% HS Graduate	32.0	31.5	0.5** (0.3)	32.1	30.9	1.2*** (0.3)
% Some College	29.5	28.7	0.8*** (0.2)	29.3	28.6	0.6*** (0.2)
% College Graduate	14.8	14.9	-0.1 (0.2)	14.6	15.3	-0.7*** (0.2)
% More than College	8.5	9.3	-0.8*** (0.2)	8.5	9.8	-1.3*** (0.2)
% Non-Hispanic White	54.7	51.5	3.2*** (0.8)	52.6	53.5	-0.9 (0.8)
% Non-Hispanic Black	23.1	27.7	-4.6*** (0.7)	25.6	25.8	-0.2 (0.8)
% Hispanic	14.8	13.8	1.0** (0.5)	14.6	13.5	1.1** (0.5)
% Asian	2.4	2.4	-0.0 (0.1)	2.4	2.4	0.0 (0.1)
Pop. Density (per sq. mi.)	3,708	5,252	-1,544*** (193)	4,155	5,349	-1,194*** (202)
% Rentals Vacant	6.01	6.00	0.00 (0.15)	5.97	6.07	-0.11 (0.16)
Observations	6,334	8,528		9,568	5,294	
P-val of Balance Test			0.000			0.000

Notes: This table uses HUD's LIHTC property database to compare properties with matches in my application data to unmatched properties. I use HUD data from 2005 to 2019 for all U.S. states. Location characteristics are from the five-year ACS centered on 2019. Columns 1 and 2 report means of the matched and unmatched samples at the property level, and Column 3 reports the difference in means. Columns 4–6 proceed in parallel, but defining the match at the state–year level, so as to distinguish between match failures and sample coverage. $^{*} = p < 0.10$, $^{**} = p < 0.05$, $^{***} = p < 0.01$.

B Supplemental Information

B.1 Data Construction

This appendix section explains the construction of data used in the analysis.

Data Coverage. Appendix Figure A24 shows a map of applications in my data. The key populous states my data miss are Illinois (IL), Massachusetts (MA), and New York (NY).[40] Appendix Figure A25 shows application counts and their total proposed units by the year of application and by their grant assignment. Nationwide, the actual number of funded LIHTC units has slightly declined over this period. Data coverage therefore improves considerably over time.

Appendix Table A8 examines the coverage of my application data by matching winning applications to HUD's LIHTC property database. This table excludes winning applications that are not successfully matched into the HUD data. The HUD data suffers from significant flaws: Its own coverage is far from complete, and properties are often miscoded in terms of the type of credit they receive (4% or 9%), or are allocated funding in years other than the one in which they apply. With these caveats in mind, however, it still provides a useful way to examine the consequences of the incomplete state–year coverage of my application data. Columns 4–6 show that state–years covered in my application data do differ on observables from all LIHTC state–years from 2005 to 2019. They tend to have larger projects in unit count and in tax credits per unit, and they are more likely to set rents below the federal maximum. The differences on location characteristics are smaller, though my application data tends to overweight state–years with lower population density.

Appendix Figure A26 provides some histograms at the application-round level. In particular, its left panel shows that the win probability in the median round is about 40 percent, and there is substantial variation in win probabilities across rounds. Its right panel shows the distribution of a measure of a round's ex-ante probability, the "explained share" $\text{Var}(\hat{p}_i)/\text{Var}(\text{W}_i)$. When this explained share is one, then applicants face no uncertainty as to whether they would win or lose. When this explained share is zero, then the competition is equivalent to a uniform lottery. Intermediate cases are weighted lotteries. The histogram finds that applicants in most LIHTC rounds feature substantial ex-ante uncertainty about their grant assignment.

Standardizing Geography Definitions. This paper almost exclusively uses the 2000 definition of Census tracts, block groups, and blocks. The rationale for this choice is that 2010 and 2020

[40]Data were unavailable in IL due to the state's public-records law protecting the contact information of developers. The MA and NY housing finance agencies informed me they do not store some key variables for my analysis outside of the paper applications submitted by developers. Among smaller states, Nevada, Kansas, and Louisiana all cited record-keeping issues and were unable to entirely meet my records requests. A state law in Arkansas explicitly forbids their agency to release LIHTC applications. Missouri and Vermont do not use numerical rules to score applications. In the District of Columbia and South Dakota, the agencies refused my records requests, and my submissions to their state appeal boards were unsuccessful.

geography definitions are potentially endogenous to the LIHTC, in that developments can require these geographies to be redrawn. The exceptions regarding 2000 Census geographies are the RDD, rent discount, and counterfactual tenant income analyses.

Application data cover 2005 through 2019, and outcome data continue through 2021. I therefore harmonize geographies across the 2000, 2010, and 2020 Censuses. I use probabilistic crosswalks from IPUMS NHGIS which map from 2000 to 2010 and from 2010 to 2020 (Manson, 2017). I reverse the direction of the crosswalks, and I combine the crosswalks in mapping data from the 2020 Census geography.

Another instance where I use geographic crosswalks is to reconcile HUD's Difficult Development Area (DDA) boundaries to lower-level Census geographies. First, I digitized lists of DDAs for the application years 2002 through 2016.[41] Over time, these DDAs have been defined at varying combinations of geographic levels: county, county subdivision / town / place, Zip Code Tabulation Area (ZCTA), and several metropolitan area concepts (MSA/PMSA/CBSA). Once digitized, I used probabilistic crosswalks from Geocorr to map DDAs to the tract level.[42] Finally, I harmonize the tracts to the 2000 Census definition as above.

Geocoding Addresses. I used a combination of the Google Maps API, Geocodio, and OpenStreetMap to geocode street addresses in applications. Results were first checked for basic consistency with rules. For instance, the geocoded state should match the state in which the application is submitted. Similarly, the geocoded county, town, and zip code should generally match information submitted on the application. Inconsistencies were then hand-reviewed.[43] Geocoding results were also carefully hand-reviewed whenever these services flagged the geocode as inexact.

Some applications provide the address information in formats that do not lend themselves to easy geocoding. This often occurs when a project is proposed in newer neighborhoods, or low-density areas, where exact street numbers have not been assigned to every parcel. For instance, the street-address field provided is sometimes an intersection, or with reference to an intersection (e.g. "A St, 500 ft S of A St and B Ave" or "SEC [southeast corner] of A St and B Ave"). Such addresses required extensive hand-review to assign geocodes.

A notable co-benefit of the data linkages to the NHPD and to CoreLogic was an opportunity to detect and resolve other geocoding errors. Inconsistencies were also hand-reviewed. Through these steps, I concluded my geocoding of LIHTC applications is likely to be highly accurate. Wilson et al. (2022) document that the HUD LIHTC database has significant geocoding errors, motivating my attention to this issue.

[41]The originals are available in PDFs on the HUD website: https://www.huduser.gov/portal/datasets/qct.html.

[42]Geocorr is maintained by the Missouri Census Data Center and is available at https://mcdc.missouri.edu/applications/geocorr.html.

[43]In surprisingly many cases, the application submission is incorrect, and then I keep the geocoding output version (e.g., the town is misspelled, the zip code is incompatible with other information).

Detecting Reapplicants. To identify applications that are potentially reapplications, I first use rule-based record-linkage methods (`dtalink` in Stata) to link my data to *itself* (i.e., duplicating the database and treating the duplicate as if it were another file). After dropping cases where an observation links to itself, using a unique application identifier that I assigned at data entry, the data contain applications linked to their potential reapplications. Due to symmetry in the record-linkage rules, all pairs necessarily appear in duplicate (A links to B, so B also links to A), which is easily resolved by dropping one of each such pair.

Variables used in the record linkage included substrings of the project name and address, geography (city, county, geographic coordinates in degree-wise Manhattan distance), developer contact information, and project variables (unit count, funding request). Record-linkage programs often assign "scores" based on the extent of matching criteria. I used this score to divide potential links into extremely likely or unlikely reapplications and marginally-likely reapplications. I hand-reviewed all applications of the marginal group and auto-coded those at the extremes.

Data Axle (Infogroup). The Data Axle files are annual cross-sections from their "Consumer Historical" database. I recode several variables for my analysis. First, the variable `owner_renter_status` provides a scale of 0 to 9 for the likelihood the resident is a renter or owner-occupant. I split this scale at the midpoint in identifying likely renters and owner-occupants. Second, the land-use variable is `location_type`. I exclude any households coded as living in nursing homes, retirement homes, trailers, or undefined location types in the single-family versus multi-family analysis. Third, the variable `find_div_1000` is the household's predicted income in thousands of current dollars. I map this to national household deciles in each year using Table A-4a of Semega and Kollar (2022).

Competition Dates. Analyses in Section 4 use quarterly data on the award dates of LIHTC competitions. I collected these by hand in several steps. First, I looked for related documents that had datestamps on them, such as award announcement press releases, the lists of winners, or schedules included in the QAP. Second, I used the metadata on spreadsheets or other public records shared with me, which was often not wiped in the public-records release process. These spreadsheets often contain a "last saved" or "last printed" date which is the award date or close to this date. Third, states typically follow the same annual or biannual schedule, so I interpolated missing years.

National Housing Preservation Database (NHPD) Linkage. I use the NHPD to establish whether, in the absence of a LIHTC award, a losing application builds subsidized housing via another subsidy. I use Stata's `dtalink` to match the application data to the NHPD, using the same variables listed in "Detecting Reapplicants." I also then hand-review marginal potential matches.

CoreLogic Linkage. I link both winning and losing applicants to the CoreLogic data. For losers, this is to establish whether any development occurs. For winners, this is both to confirm development and to obtain additional data (for use in potential future analyses).

The challenge in matching the CoreLogic data is its scale, requiring me to carefully control false positives, as most parcels are not LIHTC parcels. I begin the match process by identifying a subset of near-sure matches by (1) matching substrings of the LIHTC project name to the parcel's first-listed LLC owner, (2) the exact street address, or (3) a very small distance between my and CoreLogic's geocoded latitudes and longitudes. I hand-reviewed all such matches to confirm they were not false positives.

I then use a regression model in the CoreLogic to predict likely matches based on a larger set of variables. I do this iteratively in rounds of matching: After hand-reviewing the latest set of matches, I add them to a file of confirmed matches used in predicting other likely matches. Over many rounds of matching, I varied the set of variables, but they broadly fall into two classes. First are common variables across the two datasets, like those used to identify reapplicants or perform the NHPD match. Second are CoreLogic-only variables only measured in CoreLogic but that I find to be highly effective in avoiding false positives, such as lot size and land-use classification.

B.2 Running Variable Definition in QCT RDD

This appendix section explains the differences from Baum-Snow and Marion (2009) and Davis et al. (2019) in defining the running variable for the Qualified Census Tract (QCT) regression discontinuity design.

The basic criteria to be a QCT are (1) a tract poverty rate above 25 percent and (2) a low tract median household income. The definition of the second criterion is a ratio of tract median income to its metropolitan area's median income, adjusted for household size. A tract may qualify if more than 50 percent of the tract earns less than 60 percent of the adjusted metro-area median. Baum-Snow and Marion (2009) use only the latter threshold, whereas Davis et al. (2019) uses the minimum of the distances to each threshold to collapse the two-dimensional discontinuity into a single running variable.

HUD's implementation of these thresholds is more complex than this summary. Among the complexities are: (1) an adjustment to ensure that no more than 20 percent of a metropolitan area is classified as a QCT, (2) disqualifications for tracts with high sampling error in the ACS, (3) unusual methods of averaging over multiple American Community Survey years, and (4) a tiered ranking based on whether the tract meets one or both of the criteria. These aspects all otherwise reduce the first-stage coefficient in the QCT RDD.

From this complexity, a single priority ranking of tracts within each metro area nevertheless does arise. That is, each tract i in each metro area m has a well-defined rank r_i^m, with a threshold rank $\bar{r}^m$ such that—with one exception of which I am aware—a tract's QCT assignment is given by $D_i = 1[r_i^m < \bar{r}^m]$. I use r_i^m as my running variable, rescaled into a metro-area percentile rank.

The exception is that HUD classifies small tracts as QCTs that are beyond the metro-area threshold rank $\bar{r}^m$, so that it can use up any "spare" capacity below the threshold of 20 percent of metro-area population. It does this in rank order, explaining why my running variable still has some tracts on the wrong side of the threshold still being QCTs.

Table 2 shows a first-stage coefficient of 0.62—that is, being just to the right of the QCT threshold on the running variable raises the probability of being a QCT by 62 percentage points. Baum-Snow and Marion (2009) do not provide an equivalent estimate, and the first stage in Davis et al. (2019) is approximately 0.4. There is thus a considerable power gain to be extracted from precisely replicating HUD's QCT assignment formula.

B.3 Sampling Variation in Basis Boost

My analysis of the role of sampling variation in the basis boost is greatly facilitated by the HUD Qualified Census Tract data files.[44] These files include the official Census estimates of the margin of error for each observation. HUD includes these values in the data files because the QCT rules exclude tracts from QCT eligibility if these margins of error are too large relative to the estimate.

Census reports margins of error in the American Community Survey (ACS), the source of the QCT variables since 2013, by multiplying estimated standard errors by exactly 1.645.[45] For each observation i, I took two random draws from the normal distribution $N(\mu_i, \sigma_i)$, where μ_i was the Census estimate for that observation and σ_i is its implied standard error.

Using the same programs as to measure the QCT running variable, I then calculate whether a tract would have been assigned as a QCT. I hold DDA status as given, and so tracts that are already in DDAs will be boosted irrespective of the random draws.

B.4 Construction of Hypothetical Applications

The following states are included in the subsample for which I can analyze the trade-off between win probability and rental income: AZ, CA, CO, GA, IA, IN, NM, OH, OK, TX, UT, WA, WI. In these states, I have data on proposed rents in applications, and the QAP scoring rule features some bonus for lower rents.

After reviewing these states' QAPs in each year, I calculated how many rent-related points each application received, given their current application. I then simulated alternative applications by recalculating the rent-related points and swapping out their actual score points for this alternative.

For applications at the federal maximum rent (60 percent of area median income), I did not simulate an "up" deviation to a higher rent, as this would disqualify their application. For

[44]Available at https://www.huduser.gov/portal/datasets/qct.html.

[45]See "Worked Examples for Approximating Standard Errors Using American Community Survey Data," https://www2.census.gov/programs-surveys/acs/tech$_{docs/accuracy/2018_{A}CS_{A}ccuracy_{D}ocument_{W}orked_{E}xamples.pdf}$.

applications at the score-maximizing kink, I simulated a "down" application with a modestly lower rent but which received no additional QAP points.

I follow Section 2.3 to map from scores to win probabilities. I only allow one of the K "folds" of the applications to deviate to an alternative rent ("up" or "down") at a time. Alternative win probabilities are calculated when applications are in the deviation fold. Within the deviation fold, it is randomly assigned in each bootstrap whether an application will deviate up or down.

B.5 Robustness Checks

This appendix summarizes robustness checks for analyses presented in the main text of the paper.

Section 4. Appendix Figure A10 replicates the USPS tract-level event study in the Data Axle data, finding considerable attenuation, which explains why the block-group and tract impacts in Figure 7 are of roughly the same magnitude. Appendix Figure A11 shows that the results of Figure 7 change little when I omit the win-probability control. This implies winners are not highly selected on their outside option relative to losers. Appendix Figure A12 presents event studies estimated separately for each treatment-year cohort, ruling out the possibility my results are driven by improper comparisons between early-treated and late-treated cohorts.

Section 5. Appendix Figure A14 shows estimated boost effects on a broader set of outcomes, such as the probability of at least one application from a Census tract in that year, counts of winning and losing applications, and counts of funded units. Appendix Figure A15 shows the estimates are robust to county–year fixed effects, as well as to controlling for tract characteristics with time-varying coefficients. Appendix Figure A16 includes always-boosted tracts. Appendix Figure A17 separately estimates impacts for the two boost triggers (QCT and DDA).

B.6 Analysis of Self-Scores

A feature of the LIHTC is that, in some states, developer applicants are required or recommended to submit "self-scores," wherein they fill out the QAP scoring rubric before their application is reviewed by the government. The self-score data allow me to explore developer's elicited expectations, similar to the survey in Kapor et al. (2020). The main takeaway from the self-score data is that the typical developer essentially knows their application's score when they apply.

I begin by comparing actual and self-scores in rank terms, ranking applications within their rounds. Self-scores are highly informative about actual scores (Spearman rank correlation coefficient of 0.73) and in fact understate actual scores on average. The left panel of Appendix Figure A27 displays a binned scatterplot of actual scores ranks versus self-scores. I transform both self-scores and actual scores into ranks in the distribution of actual scores. On average, an applicant whose self-score implies the top-rank score in their round has a 80th-percentile actual score, whereas a bottom self-scoring applicant has a 5th-percentile actual score.

I then re-estimate applicants' win probabilities using their self-scores. Following Hendren (2013), my approach allows for developers to strategically report self-scores that are not their true beliefs. To summarize the approach, I estimate the conditional distribution of actual scores given self-scores, and then I compute the developer's win probability by sampling from this conditional distribution. The key implementation steps are: (1) pool scores across rounds by transforming them into percentile ranks, and (2) use kernel-density methods to nonparametrically estimate the copula between actual score percentile and self-score percentile. This approach assumes that applicants know the distribution of their potential rivals but not the draw.

Similar to the results for self-scores, the new estimates of win probabilities show developers are highly informed. The R^2 of the win-probability measures constructed using actual scores and self-scores is 0.94. The right panel of Appendix Figure A27 shows the binned scatterplot of these two estimates of win probabilities. Perhaps surprisingly, there is substantial mass of the self-score-based win probabilities near zero and one. A lack of knowledge about one's own score is thus unlikely to explain why I observe so many developer applications with little chance of winning. A remaining potential explanation is that applicants are less informed about their potential rivals than my simulations assume.

B.7 Estimating the Marginal Willingness to Pay for New Rental Housing

This section introduces an econometric strategy for estimating the marginal willingness to pay (MWTP) for new rental housing by location. It addresses a sample-selection problem as in Heckman (1979): Rents on new units are only observed when new units are built, and construction occurs when potential rents exceed construction costs.

Such sample selection introduces a bias in hedonic regressions of new-unit rents on location characteristics. Failing to account for sample selection will cause hedonic regressions to systematically overstate potential new-unit rents in neighborhoods where no new rental units were built.

Addressing such sample-selection concerns is of potentially substantial importance, as new construction of rental units occurs in only 17 percent of Census tracts from 2010 to 2019. By implication, the researcher must impute new-unit rents in 83 percent of tracts.

Selection Model. Following Heckman (1979), I derive a bivariate-normal sample-selection model. I then discuss the use of lagged construction activity as a proxy for local construction costs, motivating its use a instrument for the selection correction. Consider the following model of the

MWTP for rental housing of age a in location i:

$$R_i(a) = \boldsymbol{x}_i \boldsymbol{\beta}(a) + \delta(a)\xi_i + \varepsilon_i(a),$$

where x_i contains observable location characteristics, ξ_i is a time-invariant scalar unobservable for the location, and $\varepsilon_i(a)$ is an i.i.d. normal error term across locations and unit ages. The outcome $R_i(a)$ is the (log) rent for units of age a in location i. The terms $\boldsymbol{\beta}(a)$ and $\delta(a)$ are age-specific coefficients on location characteristics and the time-invariant unobservable respectively.

Suppose there are two ages of units, new ($a = 1$) and old ($a = 0$). I implicitly remove the time-invariant location unobservable term ξ_i by controlling for old-unit rents in the same location.[46] This yields a MWTP for new units of

$$R_i(1) = \boldsymbol{x}_i \tilde{\boldsymbol{\beta}} + \delta R_i(0) + \tilde{\varepsilon}_i. \tag{15}$$

I now show the selection problem. The MWTP for new rental housing $R_i(1)$ is observed if and only if it exceeds construction cost:

$$B_i(1) = 1[R_i(1) \geq C_i(1)],$$

as otherwise there is no new rental development in location i. I model construction costs by age and location as:

$$C_i(a) = \boldsymbol{x}_i \boldsymbol{\gamma} + \rho(a)\eta_i + u_i(a),$$

where η_i is similarly a time-invariant cost of the location and $u_i(a)$ is an i.i.d. normal error term across locations and unit ages. This can be similarly rewritten as

$$C_i(1) = \boldsymbol{x}_i \tilde{\boldsymbol{\gamma}} + \rho C_i(0) + \tilde{u}_i. \tag{16}$$

Combining Equations 15 and 16, I obtain that the gap between rents and costs is

$$R_i(1) - C_i(1) = \boldsymbol{x}_i \boldsymbol{\theta}_1 + \delta R_i(0) - \rho C_i(0) - \nu_i.$$

Using the normality of the error terms, the expected value of this gap, conditional on new construc-

[46] This selection-correction approach will therefore only yield estimates of the MWTP in locations with some old rental housing, as such rents are otherwise unobserved. This covers the vast majority of locations of interest. I impute rents from nearby Census block groups for remaining missing observations.

tion occurring in the location, is

$$\mathrm{E}[\nu_i \mid B_i(1) = 1] = \frac{\phi(\tilde{\rho}C_i(0) + \boldsymbol{x}_i\tilde{\boldsymbol{\theta}}_1 + \tilde{\delta}R_i(0))}{\Phi(\tilde{\rho}C_i(0) + \boldsymbol{x}_i\tilde{\boldsymbol{\theta}}_1 + \tilde{\delta}R_i(0))},$$

where coefficients with tildes indicate normalization by the selection model's scale coefficient.

This model predicts that, observed new-unit rents are positively selected on the error term when new rental units are observed. Such sample-selection bias will be strongest when a location has characteristics that predict the absence of construction. It is thus of particular relevance for estimating the MWTP for subsidized development, which is less constrained by ex-ante profitability. Using bivariate normality, we obtain that

$$\mathrm{E}[R_i(1)|x_i, R_i(0), C_i(0)] = \boldsymbol{x}_i\tilde{\boldsymbol{\beta}} + \delta R_i(0) + \lambda\frac{\phi(\tilde{\rho}C_i(0) + \boldsymbol{x}_i\tilde{\boldsymbol{\theta}}_1 + \tilde{\delta}R_i(0))}{\Phi(\tilde{\rho}C_i(0) + \boldsymbol{x}_i\tilde{\boldsymbol{\theta}}_1 + \tilde{\delta}R_i(0))}, \tag{17}$$

where λ is a coefficient on the selection-correction term.

Results. Nonparametric identification of the parameters $(\boldsymbol{\beta}, \delta, \lambda)$ requires a selection instrument. Supply-side variables are natural instruments in this context, and the model above motivates lagged construction costs $C_i(0)$ for this purpose. In equilibrium, such costs should affect rents on new units only as a proxy for current construction costs. Formally, the instrument requires the exogeneity assumption that

$$\mathrm{E}[\tilde{\varepsilon}_i(1)u_i(0)|\boldsymbol{x}_i, R_i(0)] = 0,$$

implying that lagged construction costs are unrelated to current MWTP for new units conditional on location characteristics x_i and current rents $R_i(0)$ on old units. This instrument will be relevant to the extent that there is a persistent component of construction cost, as in the model above. In particular, I use the log count of rental units built before 1989 in the locality.

I estimate Equation 17 by the Heckman (1979) two-step approach, bootstrapping over both steps so as to obtain standard errors of correct size. I use data at the Census tract level. For greater data availability, I use log median rents across all units in the tract, rather than exclusively old units as in $R_i(0)$. I instrument for the selection-correction term with the following four variables: the shares of rental units built in the tract from 1990 to 1999 and from 2000 to 2009, along with binary indicators for whether any rental units in the tract were built in these intervals.

Appendix Table A7 reports the results. Column 1 shows that a tract's new unit rents are closely predicted by the rents on older units in the same tract. Column 2 includes the selection correction, instrumented using lagged construction. Columns 3 and 4 augment this specification with tract characteristics and with state fixed effects. In particular, the characteristics are: log population density, age shares of population (less than 18, 18–34, 35–64, 65 and over), racial/ethnic shares of population (non-Hispanic white, non-Hispanic black, Hispanic, Asian, other), and educational-attainment shares of population (less than high school, high-school graduate and some college, bachelor's degree and up), log median household income, and the housing-unit vacancy rate.

Across specifications, the coefficient on selection-correction term is statistically significant, and the instruments are strong. The estimates imply an average selectivity bias of 6 to 10 percentage points in extrapolating from tracts with new rental units to tracts without them. The selectivity bias falls sharply when I include a powerful predictor of selection as a control: the log count of rental units built in the tract before 1989. Overall, the results suggest that sample-selection concerns are of considerable importance in estimating MWTP for new rental housing. Failing to account for sample selection would have otherwise led to overestimates of MWTP and therefore of tenant incidence.

Regulated Rents. To calculate LIHTC rent discounts, I also require regulated rents. Federal regulations specify maximum rents in terms of fractions of AMI: For instance, the federal maximum rent for LIHTC units is that they are "affordable" (i.e., no more than 30 percent of income) for a household making 60 percent of the area median. The maximum monthly rent is thus 1.5 percent ($0.6 \cdot 0.3/12 = 0.015$) of AMI. I use HUD data on Area Median Incomes (AMI) in 2019, aligning with the central ACS year for market rents.

There is a specific AMI for each household size, whereas rents are set in terms of a unit's number of bedrooms. HUD's conversion rule is 1.5 persons per bedroom: A one-bedroom unit's regulated rent is computed by averaging the AMI levels of a one- and two-person household, for instance, whereas a two-bedroom unit's rent reflects the AMI level of a three-person household.[47] These steps yield rent levels for units by numbers of bedrooms in each HUD geography. I project the data to the tract level using ACS data on the distribution of numbers of bedrooms in rental units by tract. I use this approach to calculate an average federal maximum rent in LIHTC units by tract that would be allowed at 100% AMI, and then I adjust using the actual limits in my application data, which are specified as fractions of AMI.

B.8 Imputing Counterfactual Tenant Incomes

I estimate the distribution of LIHTC tenant incomes using property-level data from HUD (Table 8 of the 2019 LIHTC Tenant Data release). These data contain property-level medians as well as household shares with annual incomes in the following ranges: \$0 to \$5,000; \$5,000 to \$10,000; \$10,000 to \$15,000; \$15,000 to \$20,000; \$20,000 and up. The income shares are extensively censored for privacy.

[47] See https://www.huduser.gov/portal/datasets/il.html#faq_2023.

I model the property-level income distributions as normal: $y_{ib} \sim \mathcal{N}(\mu_b, \sigma_b)$. In particular, I use median incomes to estimate the two parameters:

$$\mu_b = \alpha + \beta \cdot \text{MedInc}_b$$
$$\sigma_b = \gamma + \delta \cdot \text{MedInc}_b.$$

I obtain estimates of the parameters by method of moments, in particular by minimizing the sum of square residuals between actual and predicted shares when reported.

For counterfactual tenants, I take the estimated market rents from above, and I find the household income distribution in ACS data that corresponds to that rent. In particular, I use ACS Table B25122 from the 2021 5-year ACS tabulation, which is centered on 2019. This table contains the joint distribution of household income and gross rent by Census tract. I am therefore able to leverage geography as well as the rent level to construct a counterfactual.

C Theoretical and Structural Appendix

C.1 Identification

Nonparametric identification of the developer primitives in the model poses three issues. First, winners and losers may be selected by the government using information that is not directly observable to the economist but may be known by the developers themselves. Second, applicants self-select into participation from a pool of potential applicants. Third, there are both permanent and transitory unobservable characteristics of developers that must be disentangled.

Let s denote the observable component of the developer's state, and let ξ and ε respectively denote the permanent and transitory unobservables. The primitives to be identified are: (1) application flow payoffs $\Pi^A(a, s, \xi)$, (2) building flow payoffs $\Pi^B(b, s, \xi)$, (3) the distribution of permanent unobservables $G_\xi(\cdot|s)$, and (4) the distribution of transitory unobservables $F_\varepsilon(\cdot|s, \xi)$. The unobservable distributions may depend flexibly on the observable state s and, for the transitory unobservables, the permanent unobservable ξ. As applying and building are binary choices, let $\Delta\Pi^A(s, \xi) = \Pi^A(1, s, \xi) - \Pi^A(0, s, \xi)$ and $\Delta\Pi^B(s, \xi) = \Pi^B(1, s, \xi) - \Pi^B(0, s, \xi)$.

The data for each application are the grant assignment W_{it}, the win probability p_{it}, observables s_{it}, and the development choice B_{it} for losers. Non-applicants are not observed at the level of individual (potential) applications, but instead the economist has counts of applications by characteristics.

Step 1: Selection of Winners. The first step of the identification argument is to use the random assignment of grants conditional on p_{it} to address selection into winning. In particular, consider the probability of development after applying and losing:

$$\Pr(B_{it} = 1 \mid s, A_{it} = 1, W_{it} = 0) = \int F_\varepsilon(\Delta\Pi^B(s,\xi) \mid s, \xi)\, dG_\xi(s \mid A_{it} = 1, W_{it} = 0). \qquad (18)$$

A naive attempt to estimate $\Pi^B(s,\xi)$ using the building probabilities of losing applicants would be confounded by both selection of winners and self-selection into application. In particular, selection generates covariance between s and $\mathrm{E}[\xi|s, A_{it} = 1, W_{it} = 0]$.

The win probability p_{it} balances winners and losers on the unobservables (ξ, ε). Letting $s' = (s, p)$, Equation 19 becomes

$$\Pr(B_{it} = 1 \mid s', A_{it} = 1) = \int F_\varepsilon(\Delta\Pi^B(s',\xi) \mid s', \xi)\, dG_\xi(s' \mid A_{it} = 1). \qquad (19)$$

Within sets of s', winners and losers are balanced on ξ. However, applicants remain unrepresentative of the population distribution of $G_\xi(s')$. Insofar as the distribution of ξ depends upon s', self-selection still generates covariance between s' and $\mathrm{E}[\xi|s', A_{it} = 1, W_{it} = 0]$, confounding estimation of $\Pi^B(s,\xi)$.

Step 2: Self-Selection of Applicants. I overcome the self-selection issue through an "identification-at-infinity" argument as in Heckman (1990). I use shocks to the grant generosity as a selection instrument. In particular, I rely on an argument z of $s'(z)$ such that $\lim_{z\to\infty} P(A_{it} = 1 \mid s(z)) = 1$ for all z. Within such limit sets of $s'(z)$, I have that

$$\lim_{z\to\infty} \Pr(B_{it} = 1 \mid s'(z)) = \int F_\varepsilon(\Delta\Pi^B(s',\xi) \mid s', \xi)\, dG_\xi(s'). \qquad (20)$$

Winners and losers are therefore balanced on ξ, and applicants are representative of the population $G_\xi(s)$.

Step 3: Permanent and Transitory Unobservables. It remains to distinguish the respective distributions F_ε and G_ξ of permanent and transitory unobservables. To do so, I first note that we can similarly express the application probability as

$$\Pr(A_{it} = 1 \mid s') = \int F_\varepsilon(\Delta\Pi^A(s',\xi) \mid s', \xi)\, dG_\xi(s'). \qquad (21)$$

I use the panel structure of the data, along with a support restriction on ξ, to obtain identification. Following Bonhomme (2012) and Arellano and Bonhomme (2011), I assume that $G_\xi(s')$ has finite points of support. To apply their arguments, I define several objects.

Let $P_{y|s'}$ be a matrix with columns of application probabilities $\Pr(A_{it} = 1 \mid s')$ and of building probabilities $\Pr(B_{it} = 1 \mid s')$ at each state s'. In addition, let $F_{(s',\xi)}(\theta)$ be a matrix of probabilities

$\Pr(A_{it} = 1 \mid s', \xi)$ and $\Pr(B_{it} = 1 \mid s', \xi)$ that also condition on the persistent unobservable ξ, given some value of the structural parameters θ. Finally, let $G_\xi(s')$ be the probabilities of each value of the permanent unobservable ξ given the observable state s', as above.

Then, for all s', we have that

$$P_{y|s'} = F_{(s',\xi)}(\theta) G_\xi(s'),$$

which, following Bonhomme (2012), generates the restrictions on θ:

$$\{I - F_{(s',\xi)}(\theta)[F_{(s',\xi)}(\theta)' F_{(s',\xi)}(\theta)]^{-1} F_{(s',\xi)}(\theta)\} P_{y|s'} = 0.$$

An alternative path to identification is to require that the distribution of permanent unobservables G_ξ does not depend on s'.

Further Points. I now discuss three remaining issues for identification.

The first issue is distinguishing win values $\pi_1(s, \xi)$ from entry costs $\kappa(s, \xi)$, which are both contained within $\Pi^A(a, s, \xi)$. I use the bid-inversion moments for this purpose. Another route to identification, however, would be an instrument for the win probability that is excluded from either $\pi(s, \xi)$ or $\kappa(s, \xi)$ or from both.

The second issue is the joint distribution of the application and building errors within F_ε, or similarly for the components of G_ξ. In principle, these could be identified by the response of private development to the generosity instrument. Intuitively, the errors in F_ε are positively correlated if marginal applicants are less likely to develop privately upon losing than observably-similar inframarginal applicants. The cross-sectional relationship between win probabilities and development choices is also informative about these joint distributions.

The third issue is the distinction between state dependence and unobservable heterogeneity. "Excess support" from a longer panel dimension can provide the necessary identifying variation, under an assumption that the process governing state dependence is shorter than the panel. Another path to identification is identification-at-infinity argument. When a developer will never build without subsidy, the impact on the value function from the shift in the outside option, $\pi_0(\boldsymbol{s}_{it}(h_{it} = 1)) - \pi_0(\boldsymbol{s}_{it}(h_{it} = 0))$, is zero. Outside of this limit set, developers can benefit from the effect of applying on their post-application outside option.

C.2 Estimation Details

Parametric Policy Iteration. I present my implementation of PPI in detail, following Sweeting (2013).

Let $\widehat{P}^A(a, \boldsymbol{s}_{it})$ be an initial estimate of the developer's application probability from the state $\boldsymbol{s}_{it}$,

$\Pr(A_{it} = a \mid \boldsymbol{s}_{it})$. Similarly, let $\widehat{P}^B(b, \boldsymbol{s}_{it})$ be an initial estimate of the building probability.[48] Using Equation 2 and these initial choice probabilities, I will compute the unconditional expectation of the developer's application value function as

$$\bar{V}^A(\boldsymbol{s}_{it}) = \sum_a \widehat{P}^A(a, \boldsymbol{s}_{it}) \left[\tilde{\Pi}^A(a, \boldsymbol{s}_{it}) + \beta \mathrm{E}[\bar{V}^A(\boldsymbol{s}_{it+1}) \,\middle|\, \boldsymbol{s}_{it}, A_{it} = a]\right], \tag{22}$$

where $\tilde{\Pi}^A(a, \boldsymbol{s}_{it}) = \Pi^A(a, \boldsymbol{s}_{it}) + \mathrm{E}[\varepsilon^A_{it}(a) \mid \boldsymbol{s}_{it}, A_{it} = a]$. Under the distributional assumptions on errors, $\mathrm{E}[\varepsilon^A_{it}(a) \mid \boldsymbol{s}_{it}, A_{it} = a] = \sigma_a(\gamma - \log \Pr(A_{it} = a | \boldsymbol{s}_{it}))$, where γ is Euler's constant.

The first step of PPI, termed "policy valuation" in Rust (2000), computes the expected value function $\bar{V}^A(\boldsymbol{s}_{it})$ under a given set of choice probabilities through a regression-based approximation. To set up the regression, let the vector $\bar{\boldsymbol{\Pi}}_P$ collect the expected payoffs at each state: $\bar{\boldsymbol{\Pi}}_P = \mathrm{E}_P[\tilde{\Pi}^A(a, \boldsymbol{s}_{it})] = \sum_a \widehat{P}^A(a, \boldsymbol{s}_{it}) \tilde{\Pi}^A(a, \boldsymbol{s}_{it})$, integrating over application choices. I then assume the application value function can be approximated by basis functions $\phi_k(\cdot)$ of the state:

$$\bar{V}^A(\boldsymbol{s}_{it}) \approx \sum_{k=1}^{K} \lambda_k \phi_k(\boldsymbol{s}_{it}) = \boldsymbol{\Phi}_{it} \boldsymbol{\lambda},$$

where I let $\boldsymbol{\Phi}_{it} = \boldsymbol{\Phi}(\boldsymbol{s}_{it})$, and where λ_k is a coefficient on the kth basis function. I use quadratic polynomials of the state variables as basis functions. Using this approximation and Equation 22, the application value function can be represented by

$$\boldsymbol{\Phi}_{it} \boldsymbol{\lambda} = \bar{\boldsymbol{\Pi}}_P + \beta \mathrm{E}_P[\boldsymbol{\Phi}_{it+1}] \boldsymbol{\lambda},$$

with $\mathrm{E}_P[\boldsymbol{\Phi}_{it+1}]$ collecting the K entries of $\mathrm{E}_P[\phi_k(\boldsymbol{s}_{it+1}) \mid \boldsymbol{s}_{it}]$.

Let the matrix $\boldsymbol{\Phi}$ stack the basis-function vector $\boldsymbol{\Phi}_{it}$ across individuals and periods. If there were exactly K unique states, equal to the number of basis functions, one could obtain coefficients $\boldsymbol{\lambda}$ on the basis functions by $(\boldsymbol{\Phi} - \beta \mathrm{E}_P[\boldsymbol{\Phi}])^{-1} \boldsymbol{\Pi}_P$. In the typical (overidentified) case, the number of unique states exceeds the number of basis vectors, and estimates of $\boldsymbol{\lambda}$ can be obtained by

$$\widehat{\boldsymbol{\lambda}} = ((\boldsymbol{\Phi} - \beta \mathrm{E}_P[\boldsymbol{\Phi}])'(\boldsymbol{\Phi} - \beta \mathrm{E}_P[\boldsymbol{\Phi}]))^{-1} (\boldsymbol{\Phi} - \beta \mathrm{E}_P[\boldsymbol{\Phi}])' \boldsymbol{\Pi}_P$$

which are the coefficients from an OLS regression of $\bar{\boldsymbol{\Pi}}_P$ on the matrix $\boldsymbol{\Phi} - \beta \mathrm{E}_P[\boldsymbol{\Phi}]$.

To complete the policy-valuation step, I use the estimates $\widehat{\boldsymbol{\lambda}}$ to approximate the continuation value functions by $\widehat{\mathrm{FV}}(a, \boldsymbol{s}_{it}) \approx \mathrm{E}[V^A(\boldsymbol{s}_{it+1}) \mid \boldsymbol{s}_{it}, A_{it} = a]$. In particular,

$$\widehat{\mathrm{FV}}(a, \boldsymbol{s}_{it}) = \mathrm{E}[\boldsymbol{\Phi}_{it+1} \mid \boldsymbol{s}_{it}, A_{it} = a] \widehat{\boldsymbol{\lambda}}.$$

[48] I initialize the choice probabilities using the predicted values of a flexible logit model on tract-level observables.

Using these continuation values, I form the choice-specific values of applying and not applying:

$$V^A(a=1,\boldsymbol{s}_{it}) = \Pi^A(a,\boldsymbol{s}_{it}) + \beta(1-p_{it})\widehat{\mathrm{FV}}(1,\boldsymbol{s}_{it})$$
$$V^A(a=0,\boldsymbol{s}_{it}) = \widehat{P}^B(b,\boldsymbol{s}_{it+1}))[\pi_0(s_{it+1}) + E[\varepsilon^B_{it}(1)\,|b=1]] + [1-\widehat{P}^B(b,\boldsymbol{s}_{it+1})]\widehat{\mathrm{FV}}(b=0,\boldsymbol{s}_{it}),$$

where, with some misuse of notation, $\widehat{\mathrm{FV}}(b=0,\boldsymbol{s}_{it}) = \mathrm{E}[\boldsymbol{\Phi}_{it+1}\mid \boldsymbol{s}_{it}, B_{it}=0]\widehat{\boldsymbol{\lambda}}$. By Equations 4 and 5, I can also form choice-specific values for the building decision:

$$V^B(b=1,\boldsymbol{s}_{it}) = \pi_0, \quad V^B(b=0,\boldsymbol{s}_{it}) = \widehat{\mathrm{FV}}(b=0,\boldsymbol{s}_{it}).$$

In the second step of PPI, termed "policy improvement," I use the choice-specific values to update the choice probabilities. In particular,

$$\widehat{P}^A(a,\boldsymbol{s}_{it}) = \Lambda\left(\sigma_a^{-1}[V^A(1,\boldsymbol{s}_{it}) - V^A(0,\boldsymbol{s}_{it})]\right) \text{ and } \widehat{P}^B(b,\boldsymbol{s}_{it})) = \Lambda\left(\sigma_b^{-1}[V^B(1,\boldsymbol{s}_{it}) - V^B(0,\boldsymbol{s}_{it})]\right)$$

where the logistic function is $\Lambda(\cdot) = \exp(\cdot)/(1+\exp(\cdot))$.

The iteration can now proceed. I use the probabilities from the latest policy-improvement step to recompute the policy-valuation step, iterating until the choice probabilities converge.

Sampling. A challenge in estimation is that applying for the LIHTC is rare. Each year, approximately three percent of tracts have at least one application. To raise precision without an excessive number of observations, either simulated or actual sdata, I over-sample tracts with applications and then re-weight the sample.

For the actual data, I use choice-based sampling as in Manski and Lerman (1977). In particular, I configure the sample so that I retain all tract–years with applications, and I randomly sample tract–years without applications so that such observations represent one-fourth of the sample. In estimating the model, I always match data moments weighted to reflect the choice-based sample.

The same problem also arises in the simulated data. Here I use importance sampling to raise the application probability in the unweighted simulated data. To motivate my approach, note that applications with a strongly positive value for the persistent unobservable ξ_i are likelier to apply (Appendix Table A3). Instead of the standard normal, I therefore use $\xi_i \sim \mathcal{N}(2,1)$ and construct simulation weights $\omega_i = \phi(\xi_i)/\phi(\xi_i - 2)$, where $\phi(\cdot)$ denotes the standard normal density. In practice, I find this roughly triples the unweighted application probability in the simulation, improving model precision at lower counts of simulated observations.

Drawing Potential Applicants. In the simulation, each tract draws a potential applicant each year. I first sample with replacement from the empirical distribution of tract–years, providing me with a simulated joint distribution of tract characteristics.

For each simulated potential application, I then take i.i.d. lognormal draws for the number of units and the tax credit amount per unit. The parameters of the lognormal distribution are estimated. This procedure abstracts away from correlation between application characteristics and tract characteristics. In my context, however, a negligible share of variation in application characteristics is explained by tract characteristics. This simplification is likely inconsequential.

I also randomly sample the persistent unobservable ξ_i and an indicator for whether the potential applicant is a reapplicant. I set the share of potential reapplicants to 0.22, consistent with the empirical share of applicants that are reapplicants. As I do not target this share as a moment, this choice merely controls the relative precision of the data moments for application and reapplication. I thusignore the "initial conditions" problem of Heckman (1981) in drawing independent samples of potential-reapplication status, the persistent unobservable, and application and tract observable characteristics.

Counterfactual 1: No LIHTC. I calculate the incidence of the LIHTC by computing a counterfactual in which the LIHTC does not exist. I do so by setting win probabilities to zero and by changing the structural parameters so that entry costs are arbitrarily large and win values are zero.

I find the new equilibrium rents and housing quantity as follows. From the model's demand side, I impose that $\Delta \log \bar{r} = \frac{1}{\varepsilon} \Delta \log H$, where $\bar{r}$ is the average rent, inclusive of the LIHTC rent savings, and H is the total housing stock. The log change $\Delta \log \bar{r}$ therefore accounts for both rent savings and general-equilibrium effects.

On the supply side, I increment the intercept coefficient on the outside option until equilibrium is reached. An important detail in this approach is that the market rent r^m prevails in the outside option, both before and after the LIHTC is eliminated. I therefore net out the value of the rent savings from $\Delta \log \bar{r}$ in adjusting outside options. This adjustment results in the market rent r^m rising less in logarithmic terms than the average rent $\bar{r}$.

Counterfactual 2: Stylized Voucher. I eliminate the LIHTC as above. I then introduce the stylized voucher as a subsidy wedge between the rents facing the supply and demand sides of the model. This is accomplished by incrementing pre-voucher rents until enough housing is produced that, at the post-voucher rent, the model's demand side wants to consume as much housing as it did under the LIHTC. This approach also restores the same average rent $\bar{r}$ as under the LIHTC. In the budget-balanced voucher counterfactual, I increment pre-voucher rents as above but stop when the fiscal cost of the voucher equals the fiscal cost of the LIHTC.

References for Appendices

Andrews, Isaiah, Matthew Gentzkow, and Jesse M Shapiro, "Measuring the Sensitivity of Parameter Estimates to Estimation Moments," *Quarterly Journal of Economics*, 2017, *132* (4), 1553–1592.

Arellano, Manuel and Stéphane Bonhomme, "Nonlinear Panel Data Analysis," *Annu. Rev. Econ.*, 2011, *3* (1), 395–424.

Baum-Snow, Nathaniel and Justin Marion, "The Effects of Low Income Housing Tax Credit Developments on Neighborhoods," *Journal of Public Economics*, 2009, *93* (5-6), 654–666.

Bonhomme, Stéphane, "Functional Differencing," *Econometrica*, 2012, *80* (4), 1337–1385.

Davis, Morris A, Jesse Gregory, and Daniel A Hartley, "The Long-Run Effects of Low-Income Housing on Neighborhood Composition," Working Paper 2019.

Heckman, James, "Sample Selection Bias as a Specification Error," *Econometrica*, 1979, *47* (1), 153–161.

_ , "An Incidental Parameters Problem and the Problem of Initial Conditions in Estimating a Discrete Time-Discrete Data Stochastic Process," *The Structural Analysis of Discrete Data*, 1981, pp. 179–195.

_ , "Varieties of Selection Bias," *American Economic Review*, 1990, *80* (2), 313–318.

Hendren, Nathaniel, "Private Information and Insurance Rejections," *Econometrica*, 2013, *81* (5), 1713–1762.

Kapor, Adam J, Christopher A Neilson, and Seth D Zimmerman, "Heterogeneous Beliefs and School Choice Mechanisms," *American Economic Review*, 2020, *110* (5), 1274–1315.

Manski, Charles F and Steven R Lerman, "The Estimation of Choice Probabilities from Choice Based Samples," *Econometrica*, 1977, *45* (8), 1977–1988.

Manson, Steven M, "IPUMS National Historical Geographic Information System: Version 12.0," 2017.

Rust, John, "Parametric Policy Iteration: An Efficient Algorithm for Solving Multidimensional DP Problems?," Working Paper 2000.

Semega, Jessica and Melissa Kollar, "Income in the United States: 2021," *Census Bureau Report P60-276*, 2022.

Sweeting, Andrew, "Dynamic Product Positioning in Differentiated Product Markets: The Effect of Fees for Musical Performance Rights on the Commercial Radio Industry," *Econometrica*, 2013, *81* (5), 1763–1803.

Wilson, Nicole E, Michael Hankinson, Asya Magazinnik, and Melissa Sands, "Inaccuracies in Low Income Housing Geocodes: When and Why They Matter," *Urban Affairs Review*, 2022.