

Bank Regulation and Affordable Housing Supply*

Yuwen Ji[†]

October 22, 2026

Abstract

This paper studies the financing cost of affordable housing arising from the policy interaction between the Community Reinvestment Act (CRA) and the Low-Income Housing Tax Credit (LIHTC) program. Large banks provide 85 percent of equity investment in the LIHTC market, motivated by CRA incentives. However, CRA evaluations tie banks' investment to their local deposit footprints, which creates spatially fragmented demand for tax credits. I assemble a novel dataset that unpacks the capital stack and development costs of LIHTC projects obtained through FOIA requests. Using bank mergers and acquisitions to generate plausibly exogenous variation in local CRA demand, I find that increases in bank deposits significantly raise LIHTC investment in affected counties. Tax credit prices are higher where banks have stronger CRA incentives, while syndication further depresses prices in thin-deposit counties. Markets with lower tax credit prices see less entry and smaller projects. To quantify the efficiency consequences, I develop and estimate a structural model of affordable housing supply and tax credit demand. Counterfactual simulations show that removing the spatial constraints embedded in CRA incentives, as proposed in the 2023 CRA modernization, would equalize developers' financing costs, reallocate equity across space, reduce syndication costs, increase affordable housing supply and improve the efficiency of the LIHTC program.

Keywords: Affordable Housing, Spatial Fragmentation, Bank Regulation, Tax Credit Finance

*I am deeply grateful to my advisors, Lu Han, Christopher Timmins, Anthony DeFusco, and Dayin Zhang for their guidance and support. I also thank Ben Couillard (discussant), Kurt Paulsen, Richard Green, Kerry Vandell, Pierre Mabillet, Timothy Riddiough, Konstantin Milbradt, Erica Jiang, Christian Hilber, Francisco Amaral, Sean Flynn, Martin Schmalz, Luming Chen, Evan Soltas, Jason Allen, Yongheng Deng, Heejin Yoon, Alina Arefeva, and seminar participants at the University of Wisconsin–Madison Real Estate Seminar, Zurich–Bern–Wisconsin Conference on Real Estate and Urban Economics, Zurich–Oxford Doctoral Symposium on Real Estate Markets, and Wisconsin Real Estate PhD Alumni Conference for their helpful comments and feedback. This project has also benefited from industry insights shared by Paul Aylesworth and staff at WHEDA and Cinnaire, as well as the cooperation of state housing agencies in fulfilling my data requests. Financial support from the Russell Sage Foundation Doctoral Dissertation Grants (R-2502-52038) and Real Estate Research Institute Dissertation Grants are gratefully acknowledged. All errors are my own.

[†]Ji: University of Wisconsin-Madison. Email: yuwen.ji@wisc.edu.

1 Introduction

The housing affordability crisis has emerged as one of the most pressing social challenges in the United States. Rising construction costs, restrictive land-use regulation, and other constraints on housing supply have made it increasingly difficult to expand the housing stock ([Baum-Snow and Duranton, 2025](#); [Glaeser and Gyourko, 2025](#)). Expanding the supply of affordable housing requires substantial capital, yet the revenues generated by income-restricted rents are often insufficient to support the cost of development. Governments therefore rely on subsidies and regulatory incentives to make affordable housing financially feasible and attract private investment ([Clapp, 1981](#); [Olsen, 2003](#); [Collinson et al., 2015](#); [Soltas, 2024b](#); [Krimmel and Wang, 2026](#)). In the United States, the primary federal program is the Low-Income Housing Tax Credit (LIHTC), which provides developers with tax credits that can be sold to investors in exchange for equity financing. Since its creation in 1986, the LIHTC has become the largest federal program supporting the development of affordable rental housing.

A distinctive feature of the LIHTC is that the government does not directly provide most of the capital used to finance these projects. Instead, the program relies on private investors to purchase tax credits, and banks provide roughly 85 percent of LIHTC equity investment.¹ Banks have an additional incentive to make these investments because affordable housing investments can receive regulatory credit under the Community Reinvestment Act (CRA). The CRA evaluates banks on their provision of credit and investment to the communities in which they collect deposits, linking the regulatory value of an affordable housing investment to a bank's geographic footprint ([Desai et al., 2010b](#); [U.S. Government Accountability Office, 2012](#); [Marquis et al., 2012](#); [Kneebone and Reid, 2021](#); [Cook et al., 2023](#)).

This interaction between housing subsidies and bank regulation creates a spatially fragmented market for affordable housing finance. A bank has stronger regulatory incentives to purchase tax credits from projects located in areas where it has a larger deposit presence. Consequently, otherwise similar tax credits can attract different investor demand depending on where the housing project is located. Stronger local demand raises the amount investors are willing to pay for tax credits, increasing the equity available to developers and potentially changing both where affordable housing is developed and how much housing is built. Thus, a regulation intended to encourage community investment may shape both the spatial distribution of affordable housing and the efficiency of public subsidies.

In this paper, I quantify these effects by combining new data on LIHTC financing with a model of affordable housing supply and CRA-driven investor demand. I first show that banks' LIHTC investments respond to their local deposit footprints. Banks invest more in counties where they hold

¹See industry estimates in comments submitted for the 2023 Regulatory Capital Rule, Source: [Federal Reserve](#).

more deposits, and these markets have higher tax credit prices and lower investor returns. Banks are more likely to invest through syndicators in thin-deposit markets, where pooling fragmented capital imposes additional costs. These differences in financing conditions translate into housing supply, as markets with higher tax credit prices see more entry and larger projects. I develop and estimate a model of CRA-driven investor demand, banks' choice between direct and syndicated investment, and affordable housing supply to quantify these mechanisms and their equilibrium implications. I use the estimated model and counterfactual to evaluate how changing the geographic structure of CRA incentives affects financing costs, the spatial distribution of affordable housing, and the efficiency of LIHTC subsidies.

I combine several new datasets linking banks' local CRA incentives to LIHTC investment and developer behavior. I first assemble detailed LIHTC applications obtained through public-records requests to 15 state housing finance agencies. The applications report project financing, investor and syndicator identities, developer characteristics, project size and costs, scoring, and 15-year operating pro formas for both awarded and unsuccessful applications.² I next construct a bank-by-county panel of community-development investments from OCC Part 24 filings and link it to the FDIC Summary of Deposits and bank merger histories. These data allow me to measure banks' local LIHTC investment and deposit exposure and to exploit merger-induced changes in local deposits. To estimate the supply model, I use competitive 9 percent LIHTC applications in Texas from 2010 to 2024, which provide detailed information on developers' project choices and the allocation process. Finally, I supplement these data with digitized CRA performance evaluations, bank subsidiary records, and county characteristics.

Using these data, I document four empirical facts. First, banks invest more in LIHTC where they hold more deposits. To address the endogenous location of deposits, I exploit bank mergers that shift deposits across banks within counties. The IV estimates imply that an additional \$1 billion of local deposits increases LIHTC investment by \$1.37 million, providing direct evidence that CRA incentives tie investment to banks' deposit footprints.

Second, this localized incentive is capitalized into tax credit prices. Prices are higher in counties with more large bank deposits and, within investing banks, prices are also higher where the bank holds more local deposits. Third, fragmented CRA incentives create an additional financing friction. Projects in thin-deposit counties rely more on syndicators to pool bank commitments, and syndication further depresses tax credit prices. Thin-deposit markets therefore receive less equity due to both weaker local CRA demand and higher costs of pooling fragmented capital.

Finally, developers' supply decisions vary systematically with these spatially heterogeneous financing costs. A 10-cent higher tax credit price is associated with 8 additional units per project and 0.24 additional applications per county-year. Together, these findings show how CRA incentives

²Appendix Figure B1 reports data coverage by state and year.

affect not only where banks invest but also tax credit prices, syndication costs, and ultimately the spatial distribution of affordable housing supply.

Motivated by these facts, I develop an equilibrium model of LIHTC supply and bank demand. On the supply side, developers choose project size to maximize expected profits, trading off tax-credit equity and rental income against development costs and the probability of receiving an allocation. The model incorporates the institutional features that shape this choice, including the competitive allocation process, credit-request caps, zoning and program size constraints, and the use of standardized architectural templates.

On the demand side, banks choose LIHTC investment given a CRA benchmark tied to their local deposits. Banks can originate investments directly, incurring a fixed underwriting cost, or invest through syndicators that charge a proportional fee. Banks with larger commitments are therefore more likely to originate directly, while smaller, geographically fragmented commitments are more likely to be syndicated. County-level tax credit prices clear the market between developer credit supply and bank demand, linking the spatial distribution of deposits to financing conditions and affordable housing supply.

I estimate the supply model using Texas LIHTC applications from 2010 to 2024. I first estimate the award process from comparisons of winners and losers within the same competition and recover construction costs from observed project expenditures. I then estimate developers' size choices by Simulated Method of Moments (SMM). For each application, the model draws a profit shock and solves for optimal project size under the nonlinear tax-credit schedule, including the size-dependent developer-fee allowance and credit-request cap. Delivered size additionally reflects the minimum project size, rural project cap, parcel-specific zoning limits, and the option to adopt a standardized architectural template. The estimator matches the resulting distribution of project sizes and bunching at these institutional thresholds to the data. On the demand side, I estimate the cost of deviating from the CRA benchmark, the syndication cost, and heterogeneity in direct-origination costs by matching tax credit prices and banks' origination choices across local markets.

The CRA therefore has two distinct effects: it increases aggregate investor demand for LIHTC, but its fragmented evaluation framework directs that demand toward deposit-rich markets. My counterfactuals separate the aggregate benefit of CRA-induced demand from the efficiency consequences of its spatial variation. I first compare the observed equilibrium with a counterfactual that removes the CRA-induced price premium. CRA substantially increases equity available to LIHTC projects, but directs these gains mainly toward deposit-rich markets. Counties in the top deposit quartile account for 94 percent of the associated increase in housing supply.

I next equalize tax credit prices across counties and allow developers to re-optimize project size. Holding both the government credit budget and total investor equity fixed, equalization reallocates equity toward thin-deposit markets, where project size is less constrained by credit-request caps,

and zoning, making housing supply more responsive to additional equity. Expected awarded units increase by 840, or 1.16 percent, with most of the gains concentrated in counties with fewer large-bank deposits.

Finally, I use the full equilibrium model to simulate banks' origination choices under statewide assessment areas. I solve for the resulting equilibrium tax credit prices and take them into the supply model to quantify developers' responses. Holding banks' total commitments and the aggregate tax credit budget fixed, relaxing the spatial constraint would equalize CRA incentives across counties, allow banks to consolidate direct-origination desks and reduce the syndication costs created by fragmented CRA incentives. The combined effect increases expected awarded units by 939, or 1.30 percent, of which 840 units arise from spatial reallocation and the remaining 99 from lower origination costs.

This paper first contributes to the literature on the LIHTC program. Existing work studies how LIHTC affects housing markets, neighborhood outcomes, and the spatial distribution of affordable housing (Cummings and DiPasquale, 1999; Sinai and Waldfogel, 2005; Baum-Snow and Marion, 2009; Ellen et al., 2009; Eriksen and Rosenthal, 2010; Burge, 2011; Diamond and McQuade, 2019; Brennan et al., 2022; Favilukis et al., 2023), as well as how program rules shape where subsidized projects are built (Ellen et al., 2015; Ellen and Horn, 2018). A related literature studies the cost and incidence of subsidized housing and the value of the tax credits used to finance it (Eriksen, 2009; Desai et al., 2010a; Soltas, 2024a; Lang and Olsen, 2025). I shift attention upstream to the investor market that finances LIHTC projects and study how financing frictions in this market affect the efficiency and spatial distribution of LIHTC supply.

Second, the paper relates to the literature on housing supply and the production decisions of developers. Existing work shows that geography, land-use regulation, permitting, and construction technology shape both the level and elasticity of housing supply (Saiz, 2010; Paciorek, 2013; Combes et al., 2021; Baum-Snow and Han, 2024; D'Amico et al., 2024; Gabriel and Kung, 2026; Rollet, 2025; Soltas and Gruber, 2026). More closely related to my model, developers choose project scale by trading off the returns to additional units against increasing construction and land costs (Chau et al., 2007; Ahlfeldt and Barr, 2022; Eriksen and Orlando, 2022). I extend this framework by incorporating financing conditions and constraints directly into developers' project-size decisions. Because tax-credit equity is a major source of project financing, variation in its price changes the return to additional units. This links financing frictions to the intensive margin of housing supply.

Third, the paper contributes to the literature on the CRA and the geography of banking. Existing work on the CRA focuses primarily on mortgage and consumer lending (Bostic and Robinson, 2003; Agarwal et al., 2012; Avery and Brevoort, 2015; Ghent et al., 2015; Saadi, 2020; Ding and Nakamura, 2021; Ringo, 2023; Conway et al., 2023; Cespedes et al., 2024), while much less is

known about banks' CRA investment test. This margin is difficult to study because transaction-level data linking banks to individual investments are scarce and the institutional structure of these investments is complex. I study this investment margin using newly assembled regulatory filings and project-level LIHTC data, showing that deposit-based CRA incentives shape where banks supply tax-credit equity, whether they invest directly or through syndicators, and the prices developers receive. More broadly, the paper connects to evidence that bank funding, lending, and entry remain spatially segmented (Hurst et al., 2016; Nguyen, 2019; Oberfield et al., 2024; Aguirregabiria et al., 2024) and that banks specialize across markets and activities (Dick, 2007; Berger and Dick, 2007; Blickle et al., 2026; Jiang et al., 2026). The paper extends these insights from lending to equity investment and links the geography of bank regulation to the financing and supply of affordable housing.

Finally, the paper contributes to a growing literature on public-private financing of socially targeted investment and infrastructure financing. A longstanding literature studies how governments rely on private firms and private capital to provide publicly subsidized goods and infrastructure (Hart et al., 1997; Hart, 2003; Iossa and Martimort, 2015). Tax-credit programs provide a distinct form of such public-private partnership. Because project sponsors often lack sufficient tax capacity to monetize these benefits directly, specialized investors intermediate the subsidy, making investor demand and the organization of the tax-equity market important determinants of how much public support ultimately reaches projects (Desai et al., 2010a; Johnston, 2019; Garrett and Shive, 2025; Darmouni et al., 2026). This mechanism also connects to the literature on sustainable investing and green-asset pricing, which shows that investor preferences and mandates can affect asset prices and firms' cost of capital (Zerbib, 2019; Pástor et al., 2021; Pedersen et al., 2021; Baker et al., 2022). I identify a regulatory source of investor-specific demand and exploit its spatial variation to study how investor incentives affect the allocation of capital. When these incentives are geographically fragmented, costly syndication reduces the pass-through of public subsidies to projects and shapes the spatial distribution of affordable housing supply.

2 Institutional Background

This section describes the key institutional features of the Community Reinvestment Act (CRA) and the Low-Income Housing Tax Credit (LIHTC) relevant to my study.

2.1 Community Reinvestment Act

The Community Reinvestment Act (CRA), enacted in 1977, requires federally insured banks to help meet the credit needs of the communities in which they operate. Federal banking regulators,

including the OCC, Federal Reserve, and FDIC, periodically evaluate banks' CRA performance, with ratings from Outstanding to Substantial Noncompliance. Poor CRA performance can delay or prevent regulatory approval of mergers and acquisitions, branch openings, and relocations.³

CRA investment test. Large banks are evaluated under separate lending, investment, and service tests, with the investment test accounting for approximately one quarter of the overall rating.⁴ The investment test evaluates the geographic distribution of banks' total qualified community development investments relative to their deposits across assessment areas.⁵ This geographic structure motivates a constant investment-to-deposit ratio across markets as a benchmark, with a median annualized ratio of 0.20 percent as discussed in Section 3.4. In practice, LIHTC investments constitute the majority of CRA-qualified investments.

Syndication. Banks can underwrite LIHTC investments directly or through specialized syndicators. Direct origination entails costs that are easier to justify at larger commitment scales, particularly in high-deposit markets. In thin-deposit markets, by contrast, CRA incentives are fragmented into smaller commitments across banks, requiring syndicators to pool commitments into project-level equity funds, creating an additional syndication cost.

Geographic reform in 2023 CRA modernization. The geographic fragmentation of CRA incentives has long generated concerns about “hot spots” and “deserts” in community development. One of the proposals in the 2023 CRA modernization sought to reduce these geographic disparities by broadening the areas in which community development activities could receive CRA consideration beyond banks' traditional assessment areas.⁶ The rule did not take effect following litigation but motivates the policy counterfactual in this paper.

2.2 Low-Income Housing Tax Credit

The Low-Income Housing Tax Credit (LIHTC), created by the Tax Reform Act of 1986 and made permanent in 1993, is the largest federal program supporting the production of affordable rental

³See examples of [branch expansion](#), [bank merger](#), and [transaction](#).

⁴The asset threshold defining a large bank is adjusted annually. ([Source](#))

⁵Appendix Tables [B1](#) and [B2](#) illustrate this evaluation framework. Section 3.4 provides supporting evidence.

⁶This proposal was also supported by the comment letters from [National Council of State Housing Agencies \(2022\)](#), [Affordable Housing Tax Credit Coalition \(2021\)](#), and [American Bankers Association \(2022\)](#). The agencies described one objective as “eliminating barriers in the current framework that have resulted in a mismatch in the supply and demand of community development financing activities in certain geographic areas.” FFIEC (2023), *Community Reinvestment Act Final Rule*, p. 801. ([Source](#))

housing. Since its inception, the program has financed the construction or rehabilitation of approximately 3.7 million rental units.⁷

State allocation. LIHTC provides competitive 9 percent credits covering roughly 70 percent of qualified basis and bond-financed 4 percent credits covering roughly 30 percent. States allocate 9 percent credits through Qualified Allocation Plans (QAPs), which specify eligibility, scoring criteria, set-asides, and geographic competition pools based on factors including income targeting, project feasibility, subsidy efficiency, and location (Ellen et al., 2009). The allocation process therefore determines which projects and locations receive credits while incorporating affordable housing needs, equity considerations, and other state policy objectives. I take the state’s QAP priorities and geographic competition structure as given and study the efficiency of financing within this allocation system, allowing award probabilities across competing projects to adjust endogenously with project choices.

Developers’ application. Developers apply with a proposed site, development plan, and financing package. Applications report project size and unit mix, income and rent restrictions, development costs, acreage and zoning restrictions, and operating pro formas. Developers typically secure site control before applying and provide evidence of financing, including debt commitments and equity letters from prospective tax-credit investors. The amount of tax credits a project can receive depends on its qualified basis, which consists primarily of eligible development costs excluding land and may receive a 30 percent basis boost for projects located in Qualified Census Tracts (QCTs) or Difficult Development Areas (DDAs). Credit requests are further constrained by project-level credit caps and program-specific size limits. Once awarded, LIHTC projects must maintain affordability for at least 30 years.

Policy interaction. LIHTC became closely linked to bank regulation following the 1995 CRA reform, which introduced a separate investment test for large banks. LIHTC investments can count as CRA-qualified investments when banks finance projects within their assessment areas.⁸ This interaction provides banks with an additional regulatory benefit from LIHTC investments, but also generates an unintended spatial consequence by increasing banks’ willingness to pay for tax credits in markets where their CRA incentives are stronger. Developers sell the 10-year stream of awarded credits for upfront equity. A price of \$0.90 raises 90 cents of equity per dollar of cred-

⁷U.S. Department of Housing and Urban Development, [Low-Income Housing Tax Credit Database](#).

⁸The finance literature also documents an additional source of mismatch that branch-level deposits do not perfectly reflect the geographic residence of depositors (Dick, 2008; Aguirregabiria et al., 2017; Sakong and Zentefis, 2023; Kundu et al., 2021). FDIC guidance allows banks to assign deposits based on account-holder location, account activity, where the account was opened, or branch-management purposes.

its, while investors receive one-tenth of the credits annually for 10 years. Higher prices increase project equity and returns to scale, translating geographically heterogeneous CRA incentives into affordable housing supply.

3 Data

I combine data from LIHTC applications obtained through FOIA requests, matching LIHTC property with bank investment and subsidiary data, CRA performance evaluations, and the FDIC Summary of Deposits.

3.1 Bank Community-Development Housing Investment

I measure banks' community-development housing investment using public welfare investment filings submitted to the Office of the Comptroller of the Currency (OCC) under 12 CFR Part 24.⁹ The filings identify both the investing bank and the geographic area targeted by each investment, allowing me to link banks' local investment activity to their local deposit exposure. Data Appendix E provide additional details on the Part 24 reporting and classification, bank crosswalk and panel construction, and merger sample construction and timing conventions.

Part 24 investment filings. Each filing reports the investing institution, year, investment vehicle, dollar amount, activity financed, and target geography. Affordable housing is the dominant use of Part 24: it accounts for 85.2 percent of investment dollars over 2004–2022. Of these housing dollars, \$132.7 billion (81.7%) is explicitly identified as LIHTC transactions.¹⁰

Because the OCC is the only federal banking regulator that releases these filings at the individual-investment level, the sample is restricted to OCC-supervised institutions. It nevertheless captures the six largest bank LIHTC investors. I aggregate filings to the bank holding company $\times$ county $\times$ year level and match them to county deposits from the FDIC Summary of Deposits. For each bank–county pair with at least one filing, I complete the panel over 2004–2022 and assign zero investment in years with no filing.

Merger sample. The identification strategy below uses acquisitions to move deposits into counties a bank already serves, so the merger sample is built alongside the investment panel. I construct

⁹12 CFR Part 24 permits national banks to make real-estate investments otherwise restricted under the National Bank Act when designed primarily to promote the public welfare, including investments in affordable housing.

¹⁰Part 24 filings do not distinguish between 9% and 4% LIHTC investments, so the baseline outcome captures both credit types. As a robustness check, I link investment LLCs to the LIHTC database to identify a subset of 9% and 4% projects separately and find results consistent with the baseline estimates.

merger-induced deposit shocks from FDIC BankFind institution-history records matched to the SOD. I restrict attention to unassisted, arm’s-length acquisitions in which the acquirer and target had overlapping pre-merger footprints spanning at least two states, and measure exposure using the target’s last pre-merger county deposits. Where a bank–county pair sees more than one acquisition, I date the event from the first and censor the panel from the second onward.

Table 1: Summary Statistics, OCC Part 24 LIHTC Sample

	Obs.	Mean	S.D.	Min	Median	Max
Panel A. LIHTC investment, all bank $\times$ county $\times$ year cells						
LIHTC Investment (M\$)	23,366	2.622	10.89	0.000	0.000	359.64
Number of LIHTC Investments	23,366	0.171	0.548	0.000	0.000	10.00
Panel B. LIHTC investment, cells with at least one LIHTC investment						
LIHTC Investment (M\$)	2,861	21.41	23.79	0.001	15.33	359.64
Number of LIHTC Investments	2,861	1.400	0.859	1.000	1.000	10.00
Panel C. Deposits and county characteristics, all cells						
Bank Deposits in County (\$bn)	23,366	3.461	17.93	0.000	0.602	822.53
Inherited Deposits (\$bn, winsorized)	23,366	0.179	1.432	0.000	0.000	27.59
County Population (M)	23,366	0.810	1.149	0.037	0.495	10.11
Median Monthly Rent (K\$)	23,366	1.322	0.326	0.703	1.259	3.009
Median Household Income (K\$)	23,366	83.56	20.86	44.08	79.40	178.04
Poverty Rate (%)	23,366	13.07	4.920	2.482	12.82	34.96
Share Rental (%)	23,366	36.51	10.78	13.41	35.32	81.03
Share White (%)	23,366	72.35	15.57	13.91	75.05	97.79
Unemployment (%)	23,366	7.275	2.417	2.116	6.905	18.41

Notes: Bank $\times$ county $\times$ year observations, 2004–2022, for the sample estimated in Table 4; county-years in which only one sample bank appears are absorbed by the county $\times$ year effects and excluded. Pairs in which the bank invests but never operates a branch are excluded. The panel is balanced, so Panel A includes bank–county–year cells with no LIHTC investment. Panel B restricts to cells with at least one LIHTC investment. Number of LIHTC Investments counts projects. Bank Deposits in County are June 30 Summary of Deposits branch deposits of the holding company in that county, zero where it has no branch there. Inherited Deposits are the target’s last pre-merger deposits in the county, zero for bank–counties never exposed to a qualifying merger. County Population is in millions of residents, Median Monthly Rent and Median Household Income in thousands of 2024 dollars.

Summary statistics. Table 1 summarizes the estimation sample. It contains 23,366 bank–county–year observations covering 1,416 bank–county pairs, 78 banks, and 363 counties over 2004–2022. Unconditionally, the average bank–county–year contains \$2.62 million of LIHTC investment, with a median of zero. Conditional on positive investment, the average cell contains 1.40 investments totaling \$21.4 million, with a median of \$15.3 million. Banks hold about \$3.5 billion of deposits in the average county in the sample, while 17.8 percent of bank–county–year cells have zero de-

posits.¹¹ Among the 200 treated bank–county pairs, the deposits inherited through a merger average \$2.0 billion, with a median of \$0.6 billion.

3.2 15-state LIHTC Application Data

Table 2 reports summary statistics for the 15-state LIHTC application sample. I obtained application packages from state housing finance agencies through public-records requests. Figure B1 summarizes the geographic and temporal coverage of these data across the 15 states. Sample coverage varies across states due to differences in record-retention practices. For example, some states delete unsuccessful applications after specified retention periods. Variable coverage also varies because state agencies may redact financing information deemed confidential.

Financing. The application packages provide unusually detailed information on project financing arrangements. The median tax credit price is \$0.90 per dollar of federal tax credit, with substantial dispersion across projects, ranging from \$0.51 to \$1.26. The median project raises \$8.78 million in LIHTC equity against \$1 million in annual tax credits. Roughly three quarters of classified projects reach investors through a third-party syndicator, while 18.7 percent receive equity directly from banks. In only 23.2 percent of projects does the equity investor also serve as a construction or permanent lender, suggesting limited overlap between debt and equity providers, potentially reflecting investors’ diverging incentives to diversify their exposure across projects (Kroszner and Strahan, 2001; Jiang et al., 2010; Fecht et al., 2021; Chu et al., 2024).

The pay-in variables capture the timing and duration of the investor’s capital commitment. On average, a quarter of equity is committed at or before construction starts; earlier contributions provide developers with capital sooner but expose investors to the project for a longer investment duration and therefore may be reflected in a discounted tax credit price. The average construction term is about one and a half years. A construction term of zero can arise particularly for acquisition and rehabilitation projects, which often have shorter or no construction-related vacancy period and can deliver credits relatively quickly.

The remaining financing variables capture project leverage and operating capacity. Four percent of projects have no strict repayment obligations from conventional amortizing hard debt and may instead rely on soft or subordinate financing with limited required payments during the affordability period and a balloon payment at maturity. The average projected Year-1 debt-service coverage ratio is 1.26, with unlevered projects entering at zero (1.31 among levered projects); negative values arise mechanically when projected Year-1 net operating income is negative.

¹¹Most of the remaining zero-deposit observations occur in years before a bank enters the county and after the bank exits.

Table 2: Summary Statistics, 15-state LIHTC Application Sample

	Mean	S.D.	Min	Median	Max
Financing					
Tax credit price (\$/credit)	0.894	0.080	0.510	0.900	1.261
Tax credit equity (\$m)	10.60	6.62	0.11	8.78	53.87
Annual credit (\$m)	1.17	0.71	0.02	1.00	5.00
Bank self-syndicated	0.187	0.390	0.000	0.000	1.000
Equity investor is also lender	0.232	0.422	0.000	0.000	1.000
Equity committed early (share)	0.250	0.249	0.000	0.200	1.000
Construction term (years)	1.56	0.56	0.00	1.67	5.00
No hard debt	0.040	0.195	0.000	0.000	1.000
DSCR	1.26	0.41	-2.31	1.26	6.59
Year-1 NOI per unit (\$k)	3.09	2.09	-7.41	2.68	31.85
Cost					
Total units	72.3	38.6	6.0	65.0	500.0
Acreage	5.7	5.8	0.0	4.3	86.6
Soft cost (\$m)	3.07	3.28	0.00	1.88	67.98
Total development cost (\$m)	16.2	10.5	0.8	13.5	140.1
Developer experience (# prior)	8.0	15.6	0.0	2.0	158.0

Notes: This table reports summary statistics for the 15-state application sample: All variables are extracted from the application packages. Tax credit price is the equity paid per dollar of 10-year federal credit. Tax credit equity is the total LIHTC equity committed in the equity letter, and annual credit is the annual LIHTC credit requested. Bank self-syndicated equals one when the named equity provider is a bank investing directly rather than through a third-party syndicator. Equity investor is also lender equals one when the construction or permanent lender belongs to the same institution as the equity provider. Equity committed early is the share of total equity paid in at or before construction closing, from the equity pay-in schedule. Construction term is the length of the construction period from the construction financing commitment, in years. No hard debt equals one for projects with no hard (amortizing) debt; DSCR is the Year-1 debt-service coverage ratio from the application pro forma, set to zero for projects with no hard debt. Year-1 NOI per unit is the first-year projected net operating income divided by total units. Total units, acreage and total development cost are extracted from the cost schedule in the application package. Soft costs are costs that cannot be included in eligible basis for claiming tax credits. Developer experience counts the developer's prior applications observed in the sample before the current application year.

Project characteristics. The average project contains 72 units on 5.7 acres and has a total development cost of \$16.2 million. The average developer has eight prior applications observed in the sample, compared with a median of two, indicating that the applicant pool includes both repeat developers and first-time entrants.

3.3 Texas LIHTC Application Data

The unit of analysis on the supply side is a 9 percent competitive LIHTC application submitted to the Texas Department of Housing and Community Affairs (TDHCA) between 2010 and 2024.

I restrict the sample to applications with a positive QAP score and non-missing weighted AMI, acreage, eligible basis, tax credit price, annual credits, and subsidy ratio. The resulting sample contains 1,910 applications, of which 961 received an allocation. Table 3 reports summary statistics separately for awarded and non-awarded applications.

Table 3: Summary Statistics, Texas LIHTC Application Sample

	Awarded (N=961)		Non-awarded (N=949)			
	Mean	S.D.	Mean	S.D.	Diff.	(s.e.)
Award						
QAP score	165.8	22.5	167.7	26.1	-1.9	(1.3)
Residualized QAP score	1.8	10.7	-1.8	12.5	3.7***	(0.8)
Weighted AMI	0.577	0.059	0.574	0.055	0.003	(0.003)
BDR-adj. weighted AMI	0.583	0.061	0.580	0.059	0.003	(0.003)
Subsidy ratio	1.526	0.381	1.480	0.330	0.046**	(0.020)
USDA loan	0.162	0.369	0.071	0.256	0.092***	(0.022)
Rural	0.376	0.485	0.293	0.455	0.083***	(0.023)
At-risk	0.231	0.422	0.104	0.306	0.127***	(0.022)
Nonprofit developer	0.176	0.381	0.177	0.382	-0.001	(0.018)
Payoff						
Tax credit price (\$/credit)	0.881	0.076	0.876	0.076	0.005	(0.004)
Total units	87.9	44.2	94.4	41.8	-6.5***	(2.0)
Annual credit (\$m)	1.12	0.50	1.19	0.47	-0.07***	(0.02)
Basis boost (share)	0.988	0.111	0.978	0.147	0.010*	(0.005)
Total development cost (\$m)	15.2	9.2	15.3	7.2	-0.1	(0.5)
Soft cost (\$m)	3.10	4.04	2.78	2.70	0.33	(0.21)
Acreage	6.4	4.8	7.4	5.4	-1.0***	(0.2)
Year-1 NOI (\$m)	0.30	0.26	0.32	0.22	-0.01	(0.01)
Rehabilitation	0.226	0.418	0.095	0.293	0.131***	(0.022)
Senior / elderly	0.348	0.476	0.336	0.473	0.011	(0.024)
Supportive housing	0.027	0.162	0.020	0.140	0.007	(0.006)

Notes: This table reports summary statistics for the Texas supply-estimation sample. QAP score is the application's total competitive score. Residualized QAP score is the score net of weighted AMI, the subsidy ratio, and application-year fixed effects, so the residual primarily captures scoring advantages from project amenities and community support. Weighted AMI is the unit-weighted average income limit of the proposed units, as a share of area median income, with market-rate units measured at 100% AMI. BDR-adj. weighted AMI is the weighted AMI adjusted by bedroom mix. Subsidy ratio is total development cost per dollar of tax credit equity. USDA equals one for projects with USDA (Rural Development) financing; at-risk for existing subsidized developments whose affordability restrictions expire within two years; nonprofit for nonprofit developers. Basis boost (share) equals one for applications getting a 30% bonus for basis (up to 130 percent). Rehabilitation, senior / elderly, and supportive housing are indicators for the proposed construction type and target population. Differences are awarded minus non-awarded means, with standard errors clustered by county; *** p<0.01, ** p<0.05, * p<0.10.

Award. The first group contains variables governing the probability of receiving an allocation. Because the QAP is rewritten annually, raw scores are comparable only within an application year. I therefore also construct a residualized QAP score, as in the supply estimation, to capture an application’s scoring advantage beyond its affordability targeting and subsidy intensity. The remaining variables capture dimensions explicitly favored by the allocation system, including affordability targeting, rural status, USDA financing, at-risk developments, and nonprofit sponsorship.

Payoff. The second group contains variables determining the developer’s payoff conditional on receiving an allocation. These include project scale and costs, the amount and price of tax credits, the basis boost, projected operating income, and project type. Together, these variables determine the development costs, equity proceeds, and expected operating returns that enter the supply model.

Awarded and non-awarded applications are strikingly similar on the fundamentals that determine underwriting quality: raw QAP scores, tax credit prices, weighted AMI, total development costs, soft costs, and projected net operating income are all statistically indistinguishable. Within a year, however, scoring does select winners: the residualized QAP score which mainly reflects amenities and community support is 3.7 points higher for awarded applications (significant at 1%).

The remaining differences line up with the allocation rules rather than with project fundamentals. Awarded projects are smaller (6.5 fewer units and \$70,000 less in annual credits); they are disproportionately rural (37.6 versus 29.3 percent), rehabilitation (22.6 versus 9.5 percent), and disproportionately carry USDA financing and at-risk status (differences of 9 and 13 percentage points). Awarded projects also obtain a slightly higher subsidy ratio and are marginally more likely to claim the basis boost. These patterns are consistent with the fact that award probabilities are shaped by set-asides, regional allocation targets, and within-year scoring, which the supply model in section 5.1 incorporates directly into the probability of winning.

3.4 CRA Rating and FDIC Summary of Deposits

Banks are evaluated for CRA compliance approximately every three to four years, with large banks receiving separate investment-test ratings for each state and multistate MSA in which they operate. I digitize these evaluations for 18 banks observed providing direct LIHTC commitment letters, recovering investment-test ratings and the geographic distribution of qualified community development investments across assessment areas. The resulting sample covers 81 evaluations from 2003 to 2024 and 4,724 assessment-area observations. All banks pass the institution-level investment test: 59 are Outstanding (72.8%), 18 High Satisfactory (22.2%), and 4 Low Satisfactory (4.9%), with none rated Needs to Improve or Substantial Noncompliance.

Regulators evaluate total qualified investments, including both new investments made during

the evaluation period and prior-period investments that remain outstanding. New investments account for 72 percent of reported qualified investments, and the median evaluation spans 3.75 years. I measure local deposit exposure using branch-level deposits from the FDIC Summary of Deposits, excluding wholesale-sourced deposits reported at headquarters and main offices. The median assessment area receives new qualified investments equal to 0.71 percent of local deposits over an evaluation period, or 0.20 percent annually.

3.5 County Economic and Demographic Data

County-level economic and demographic controls come from the American Community Survey (ACS) five-year estimates (2006–2010 through 2019–2023). The controls include population, median gross rent, the poverty rate, the renter share of occupied housing units, the White population share, the unemployment rate, and median household income. Because ACS five-year estimates are available for all counties regardless of population size, these measures cover the full sample, including small and rural counties. I then linearly interpolate within counties between Census and ACS reference years to construct an annual county-year panel.

4 Reduced-form Evidence

In this section, I document four reduced-form patterns. First, CRA incentives direct banks to invest in areas where their deposits are concentrated. Second, tax credit prices are higher where banks hold more deposits. Third, greater reliance on costly syndication further lowers tax credit prices in thin-deposit counties. Fourth, developers are more likely to enter the LIHTC competition and build larger projects where tax credit prices are higher.

4.1 CRA directs banks to invest where deposits concentrate

Deposits and investment in the cross section. Panel A of Table 4 reports the relationship between a bank’s county deposits and its LIHTC investment in the OCC Part 24 filing sample. Importantly, the investment outcome measures new Part 24 investments made by the bank in a given year, rather than its stock of outstanding investments; investments previously made by an acquired bank therefore do not mechanically enter the acquirer’s post-merger outcome.

A trillion dollars of additional deposits is associated with 2.4 additional LIHTC projects and \$57 million more LIHTC equity investment. The specification compares changes in a bank’s deposits and investment within the same county over time, while absorbing substantial potential confounding variation. County $\times$ year fixed effects absorb time-varying local shocks, including housing

demand and credit conditions, while bank–county fixed effects absorb persistent bank-specific investment patterns across counties. The remaining concern is time-varying bank–county shocks that jointly affect deposits and LIHTC investment. For example, a bank may expand its local business relationships in ways that simultaneously attract deposits and generate new relationships with affordable-housing developers. Such shocks would induce a correlation between deposits and LIHTC investment even absent a causal effect of deposits. I therefore turn to merger-induced changes in local deposits, which isolate variation in a bank’s deposit that is plausibly unrelated to contemporaneous changes in its local LIHTC investment opportunities.

Table 4: Bank deposits and LIHTC investment

	(1) Number of LIHTC Projects	(2) LIHTC investment (M\$)
Panel A. Least squares: investment on deposits held		
Deposits (\$bn)	0.0024* (0.0014)	0.057*** (0.013)
Panel B. Instrumental variables: investment on deposits held		
Instrumented deposits (\$bn)	0.0831*** (0.0299)	1.367*** (0.378)
First-stage F	19.7	19.7
Bank $\times$ county effects	✓	✓
County $\times$ year effects	✓	✓
Observations	23,366	23,366
Clustered on	Bank	Bank

Notes: This table presents estimates of the relationship between bank deposits and LIHTC investment at the bank $\times$ county $\times$ year level from 2004–2022. Column (1) measures investment by the number of OCC Part 24 filings identifying LIHTC investments, and column (2) measures investment by their total dollar value in millions of 2024 dollars. Panel A regresses LIHTC investment on the bank’s Summary of Deposits deposits in the county, measured in billions of dollars. Panel B instruments county deposits with the post-merger indicator interacted with the inherited pre-merger deposits and reports the corresponding instrumental-variable estimates. Inherited deposits are winsorized at the 99th percentile. Standard errors are clustered at the bank level and presented in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Identification. To address the endogeneity of the deposit distribution, I exploit variation in local deposits generated by bank mergers and acquisitions (M&As). Consider a merger in which acquirer i purchases target j in year t_m . Bank i inherits the deposits held by j in each county where the two banks previously overlapped. Because CRA incentives are tied to banks’ local deposit footprints, the merger generates a differential increase in local CRA exposure across counties according to the target’s pre-merger deposit holdings.

I construct the merger sample from FDIC BankFind institution-history records matched to the SOD and impose the following restrictions. First, I focus on counties where i and j had overlapping deposit footprints prior to the merger, ensuring that inherited deposits represent an increase in exposure to an existing market rather than geographic expansion (Nguyen, 2019).¹² Second, I require the overlapping footprint to span at least two states, so that identification comes from the geographic distribution of deposits inherited through mergers spanning multiple local markets rather than from a transaction concentrated in a single market (Garmaise and Moskowitz, 2006). Third, I measure exposure using target j 's last observed pre-merger deposits in county c , denoted d_{c,t_m-1}^j , so that the magnitude of the merger shock is fixed prior to the transaction. These restrictions leave 27 M&A events by 14 acquiring banks affecting 200 bank–county pairs; 1,216 bank–county pairs are never treated and serve as controls.

Estimation. I instrument changes in banks' local deposits with the interaction between the post-merger indicator and inherited pre-merger deposits. The first stage is

$$D_{ct}^b = \pi \left(\text{Post}_{ct}^b \times d_{c,t_m-1}^j \right) + \xi_{b,c}^1 + \xi_{c,t}^1 + \varepsilon_{1,ct}^b, \quad (1)$$

where D_{ct}^b denotes bank b 's deposits in county c , $\text{Post}_{ct}^b = \mathbf{1}\{t \geq t_m\}$ equals one beginning in the merger year, and d_{c,t_m-1}^j is the target's deposits in county c in the last pre-merger SOD vintage. The fixed effects $\xi_{b,c}$ absorb time-invariant differences across bank–county pairs, including the level of inherited deposits, while $\xi_{c,t}$ absorb common county–year shocks. The instrument therefore exploits the differential post-merger increase in deposits across counties according to the amount of deposits inherited from the target.

I then estimate

$$Y_{ct}^b = \beta \hat{D}_{ct}^b + \xi_{b,c}^2 + \xi_{c,t}^2 + \varepsilon_{2,ct}^b, \quad (2)$$

where Y_{ct}^b is LIHTC investment. The coefficient β therefore captures the effect of merger-induced changes in local deposits on LIHTC investment.

I examine the dynamics of the merger shock by interacting event time with inherited deposits:

$$Y_{ct}^b = \sum_{k \neq -1} \beta_k \left[\mathbf{1}\{t - t_m = k\} \times d_{c,t_m-1}^j \right] + \xi_{b,c} + \xi_{c,t} + \varepsilon_{3,ct}^b, \quad (3)$$


where $k = -1$ is the omitted period and the event window is $k \in [-6, 6]$. The coefficients therefore measure the change in the outcome per dollar of inherited deposits relative to the year before the merger. Because merger timing is staggered, conventional TWFE event-study estimates may use

¹²For bank–county pairs with multiple events, I use the first merger and censor observations from the second onward.

already-treated units as controls for later-treated units and can be biased under heterogeneous treatment effects (Goodman-Bacon, 2021). I therefore report TWFE alongside the interaction-weighted estimator of Sun and Abraham (2021), the two-stage estimator of Gardner (2021), and a stacked design following Cengiz et al. (2019) that uses never-treated units as controls in Appendix A.1 with details of sample construction for each estimator. Standard errors are clustered at the bank level.

Results. Figure 1a shows a strong first stage ($F = 19.7$). Each billion dollars of inherited deposits increases the acquirer’s local deposits by about \$0.86 billion in the merger year, with the increase remaining positive thereafter. Figure 1b shows the corresponding reduced-form response of LIHTC investment. Investment shows no differential pre-trends and rises with a two-year delay, reaching about \$1 million per billion dollars of inherited deposits in subsequent years. This delay is consistent with banks adjusting LIHTC investment as they approach their next CRA evaluation following the M&A event.

Figure 1: Merger-Induced Deposit Shocks and Bank LIHTC Investment



Notes: Panel (a) reports the effect on county deposits per billion dollars of inherited pre-merger deposits. Panel (b) reports the corresponding effect on LIHTC investment per billion dollars of inherited deposits using three staggered difference-in-differences estimators (Sun-Abraham estimates reported in A.1). In both panels, $k = -1$ is omitted, bars and shaded areas report 95% confidence intervals, and standard errors are clustered at the bank level.

Panel B of Table 4 instruments bank–county deposits with post-merger inherited deposits, $\text{Post}_{ct}^b \times d_{c,tm-1}^j$. The IV estimate implies that an additional billion dollars of merger-induced local deposits increases LIHTC equity investment by \$1.367 million. The estimates are similar across the four staggered-treatment estimators reported in Appendix Table A1, ranging from \$1.033 to \$1.38 million in equity per billion dollars of deposits.

The IV estimate is substantially larger than the OLS estimate in Panel A. One potential explanation is measurement of the local deposit footprint. Many large bank–county observations are head-office locations with substantial non-retail deposits, for which deposits reported in the Summary of Deposits may not closely measure the deposit base relevant for local CRA incentives. Consistent with this interpretation, the deposit–investment relationship becomes substantially flatter among very large bank–county pairs. The merger instrument mitigates this concern by isolating deposit changes generated by acquisitions of existing local deposit footprints rather than using all variation in reported deposits.

The magnitude of the IV estimate is broadly consistent with the CRA performance evaluations discussed in Section 3.4. Median new qualified investments equal 0.20 percent of local deposits per year, and LIHTC accounts for approximately 69.6 percent of public welfare investments, which yields an examination-based benchmark of about 0.139 percent of deposits, which is close to the IV estimate of 0.137 percent.

4.2 Tax credit prices are higher where banks hold more deposits.

Pricing tax credits as an asset. A LIHTC allocation generates a stream of federal tax credits that investors can claim over a 10-year credit period. To interpret observed tax-credit prices, I translate each price into the credit-implied return by solving the following equation:

$$p_{it} = \frac{1}{10} \sum_{k=0}^9 \frac{1}{(1+r)^k} \quad (4)$$


where p_{it} denotes the price paid per dollar of tax credit and r is the credit-implied return.

Figure 2 plots the conversion curve and the distribution of LIHTC tax credit prices. Most prices cluster around 90 cents per dollar of credit, corresponding to implied returns of roughly 2.5–3 percent, well below the 10-year Treasury yield. Notably, some projects transact at or above one dollar per dollar of credit, implying returns near zero. Panel (b) makes this comparison explicit by subtracting from each observed price the risk-free benchmark price implied by the contemporaneous 10-year Treasury yield. Appendix Figure A1 shows the corresponding comparison over years.

The substantial positive premium suggests that investors potentially value LIHTC investments for benefits, such as CRA regulatory consideration, beyond the tax credits themselves. This resembles the “greenium” in sustainable asset markets, where investors with regulatory or nonpecuniary preferences accept lower financial returns on assets carrying a preferred attribute (Zerbib, 2019; Pástor et al., 2021; Pedersen et al., 2021; Baker et al., 2022). A key distinction, however, is that the CRA premium is inherently spatial, arising from regulatory obligations rather than solely from investors’ heterogeneous preferences, and its value is tied to banks’ local deposit footprints. Con-

sequently, the same underlying tax credit can carry different regulatory value to investors across locations. This location-specific regulatory value generates geographic heterogeneity in willingness to pay and translated into LIHTC prices in equilibrium. Thus, unlike a conventional greenium” attached primarily to an asset’s characteristics, the CRA premium reflects the interaction between an asset’s location and the investor’s regulatory exposure.

Figure 2: IRR from LIHTC Tax Credit and Histogram



Notes: Panel (a) plots the histogram of tax credit prices and the conversion curve between the tax credit price and the investor’s internal rate of return, $p_{it} = \frac{1}{10} \sum_{k=0}^9 (1+r)^{-k}$; the dashed line marks the 10-year Treasury yield. Panel (b) plots the histogram of the difference between the observed tax credit price and the risk-free price implied by the contemporaneous 10-year Treasury yield. The dashed line marks a zero premium.

Estimation. The empirical analysis exploits this spatial variation by testing whether the same bank pays higher tax-credit prices where it has a larger deposit footprint and hence stronger CRA incentives.

To isolate this regulatory premium, I control for other pricing factors. Financing structure can affect valuation through the timing and risk of project cash flows¹³, as in renewable energy project finance (Johnston, 2019; Garrett and Shive, 2025; Darmouni et al., 2026). I therefore control for the timing of equity contributions and construction term from equity commitment letters, as well as the debt-service coverage ratio from the 15-year projected pro forma. I also include syndicator fixed effects to absorb persistent differences in fees and pricing practices across intermediaries. Year fixed effects absorb aggregate market conditions affecting LIHTC demand, including changes

¹³The foreclosure rates and recapture rates for LIHTC properties are extremely low. See Appendix Figure B4 for the historical foreclosure rate from 2002 to 2022. The low realized risk of LIHTC investments has also been emphasized in regulatory comment letters by the [Affordable Housing Tax Credit Coalition](#) and other 15 associations.

in investors' tax capacity and the after-tax opportunity cost of capital, much as the value of tax exemption is capitalized into municipal bond prices (Miller, 1977; Green, 1993; Ang et al., 2010; Longstaff, 2011). Finally, because banks may provide both debt and tax-credit equity to the same project and price the two jointly, I control for whether the equity investor is also the lender. This occurs in 23.2 percent of projects, consistent with banks limiting their joint debt and equity exposure to individual projects.

Table 5: Bank Deposits and Tax Credit Price

	Tax Credit Price (Cents Per \$1 Tax Credit)				
	(1)	(2)	(3)	(4)	(5)
log(Cnty Large-Bank Dep)	0.6465*** (0.2359)	0.6019*** (0.2173)			
log(Investor Bank Deposits)			0.1064*** (0.0403)	0.1049*** (0.0319)	0.1746*** (0.0646)
Pay-in Schedule	✓	✓	✓	✓	✓
No-Debt & Levered DSCR	✓	✓	✓	✓	✓
Same Lender	✓	✓	✓	✓	✓
County Controls	✓	✓	✓		
Credit Supply Controls	✓	✓	✓		
Syndicator FE	✓	✓	✓	✓	✓
County FE, Year FE	✓	✓	✓		
County × Year FE				✓	✓
Bank × County FE					✓
Sample	All	Excl. booking	All	All	All
Observations	11,481	9,822	1,456	1,456	1,456
Counties	635	617	101	101	101
Dep. Var Mean	89.771	89.778	94.711	94.711	94.711
Adjusted R-squared	0.788	0.795	0.781	0.841	0.825

Notes: This table reports the relationship between bank deposits and LIHTC tax credit prices, winsorized at the 1st and 99th percentiles. Columns (1)–(2) use total large-bank deposits in the project county, net of non-retail deposits reported in CRA evaluations; column (2) further excludes counties with non-retail booking deposits. Columns (3)–(5) use the investing bank's own county deposits for projects matched to bank investors. County controls include rent, income, and population; credit supply controls include county and state 9 and 4 percent allocations; project controls include equity pay-in timing, construction term, debt structure, debt-service coverage, and whether the equity investor is also the lender. Standard errors are clustered by county. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Results. Table 5 shows a positive relationship between local bank deposits and tax credit prices. Columns (1)–(2) use county-level large-bank deposits and yield similar estimates whether non-retail booking deposits are netted out or booking counties are excluded. A one log-point increase in deposits is associated with a 0.60–0.65 cent increase in the tax credit price.


Columns (3)–(4) compare prices across projects financed by banks with different local deposit footprints, with column (4) restricting the comparison to banks investing in the same county and year. Column (5) further adds bank-by-county fixed effects, identifying from changes in a given bank’s local deposits over time relative to other banks investing in the same county-year. Across these specifications, a one log-point increase in the investor’s local deposits is associated with a 0.10–0.17 cent higher price, equivalent to a 0.03–0.06 percentage point reduction in the investor’s implied return, as shown in Appendix Table A2. Banks pay higher prices for LIHTC tax credits where they have larger local deposit footprints.

4.3 Syndication concentrates where local bank capital is thin

LIHTC syndication. Banks can invest in LIHTC projects through two channels. A bank can invest directly in a project, sourcing and underwriting the deal itself, or it can invest through a syndicated fund that pools equity from multiple investors, performs the sourcing, underwriting, and monitoring functions and charges for this intermediation. In the application sample, 77 percent of the 10,329 projects that identify an equity provider obtain their equity letter through a third-party syndicator.

Banks differ substantially in their channel choices. Appendix Figure B2 shows that some banks predominantly originate LIHTC investments directly, while others invest almost exclusively through syndicated funds. Such specialization is also documented in local banking and lending patterns (Dick, 2007; Berger and Dick, 2007; Blickle et al., 2026; Jiang et al., 2026). Consequently, spatial variation in syndication reflects both bank composition across markets and variation in CRA incentives within banks across markets.

Figure 3: Fragmented CRA incentives and syndication



Notes: The figure illustrates how geographically fragmented CRA incentives shape LIHTC investment. Banks split fixed CRA commitments across counties according to their local deposit footprints. While Bank A, with a large local deposit base in County 1, can finance projects in County 1 directly, projects in Counties 2 and 3 must pool commitments from multiple banks and therefore rely on syndication.

My focus is on the latter source of spatial variation. Because CRA incentives are tied to banks’

local deposit footprints, a bank's aggregate CRA commitment is effectively split across the counties in which it operates. Figure 3 illustrates the mechanism. Bank A has sufficient local CRA demand to finance the project in County 1 directly. In Counties 2 and 3, individual banks have smaller local investment commitment, making direct origination less economical relative to syndication, which pool investments from multiple investors but incurring a syndication cost. Thus, thin-deposit counties rely more heavily on syndicated investment, raising financing costs. This county-level pattern can reflect both differences in bank composition across markets and within-bank channel choice.

Table 6: Local Bank Deposits, Bank Composition, and Syndication

	Syndicated share		Tax credit price (cents)	
	(1)	(2)	(3)	(4)
log(Large-Bank Deposits)	-0.0265*** (0.0061)	-0.0251*** (0.0061)	0.1635*** (0.0563)	0.1571*** (0.0560)
County syndicated share			-1.4641*** (0.2037)	-1.4443*** (0.2035)
Share _{Direct-Capable Bank Deposits}		-0.0462 (0.0373)		0.1561 (0.3202)
Share _{Direct} × log(Large Bank Dep)		-0.0252** (0.0110)		0.2795** (0.1247)
County controls, State × Year FE, CSA FE	✓	✓	✓	✓
Observations	3,575	3,575	3,575	3,575
Dep. var. mean	0.808	0.808	89.32	89.32

Notes: The unit of observation is a county × year in the 15-state application sample, 2004–2023. Columns (1)–(2) use the share of applications financed through a third-party syndicator; columns (3)–(4) use the county-year mean of deal-level tax credit price residuals (+mean) after controlling for project characteristics, CSA and state × year fixed effects. The direct-capable deposit share is the fraction held by banks observed making direct LIHTC investments elsewhere in the sample, excluding the county's own applications. Interacted variables are mean-centered. Standard errors are clustered by county. *** p<0.01, ** p<0.05, * p<0.10.

Results. Table 6 provides evidence consistent with this mechanism. Column (1) shows that a one log-point increase in local large-bank deposits is associated with a 2.7 percentage point decline in the syndicated share. Column (2) exploits heterogeneity in banks' propensity to invest directly. A one-log-point increase in deposits is associated with a 1.3 percentage point lower syndicated share when no local deposits are held by direct-capable banks, compared with 3.8 percentage points when all local deposits are held by direct-capable banks. Together, these patterns are consistent with syndication being more prevalent where local CRA demand is thin, with the deposit–syndication gradient particularly strong where a larger share of local deposits is held by direct-capable banks.

Columns (3) and (4) show how this financing margin is capitalized into tax credit prices. Tax credit prices are low in markets with fewer local deposits and decline further with higher share of syndication. Thin-deposit counties therefore may face a double disadvantage: weaker local CRA demand is associated with lower tax credit prices, while fragmented CRA incentives require projects to rely more heavily on syndication to pool investment, further increasing developers' financing costs.

4.4 Counties with higher tax credit prices have more entry and larger projects.

Developers sell allocated tax credits to investors in exchange for equity financing. As in the standard housing supply framework, project scale depends on the return to adding housing relative to its development cost (Chau et al., 2007; Eriksen and Orlando, 2022; Ahlfeldt and Barr, 2022). In the LIHTC market, adding units increases development costs but also expands eligible basis and the amount of tax credits the project can claim. A higher tax credit price thus increases the equity raised from those additional credits, effectively lowering the net cost of expanding the project. Developers therefore have an incentive to propose larger projects when tax credit prices are higher. Higher prices also increase the expected payoff of development more generally, making entry attractive to projects that would otherwise not be worthwhile. Tax credit prices can therefore affect affordable housing supply along both the intensive margin of project size and the extensive margin of developer entry.

Results. Columns (1) and (2) of Table 7 examine project size. Comparing projects within the same county and year, a 10-cent higher tax credit price is associated with 7.9 additional units, relative to a mean project size of 75 units. Excluding applications that bunch at unit caps increases the estimate slightly to 8.0 units (column 2), suggesting that binding size constraints slightly attenuate the baseline correlation.¹⁴

Columns (3)–(5) examine the extensive margin. A 10-cent higher tax credit price is associated with 0.24 additional applications filed in a county–year, relative to a mean of 3.05 applications. In contrast, higher prices are not associated with an increase in the number of projects awarded in column (4). Consequently, the additional entry increases competition with a lower award probability in column (5).

These correlations motivate, but do not identify, the supply response needed for the counterfactuals. Project size is a constrained choice subject to parcel constraints and institutional kinks and caps in the LIHTC program. As a result, the response to tax credit prices is inherently nonlinear: the same price change may induce substantial expansion for an interior project but no change for

¹⁴A full list of unit caps and share of bunching is listed in B.7.

Table 7: Tax Credit Price, Entry and Project Size

	Project Size		Entry		
	Total Units (1)	(2)	App filed (3)	Awarded proj (4)	Win Prob (5)
Tax Credit Price	78.92*** (14.13)	79.91*** (14.33)	2.446** (1.200)	-0.022 (0.510)	-0.484** (0.237)
County controls			✓	✓	✓
Acreage, Zoning, Developer FE	✓	✓			
County FE			✓	✓	✓
State $\times$ year FE			✓	✓	✓
County $\times$ year FE	✓	✓			
Excluding Bunching		✓			
Observations	8,904	8,738	3,795	3,795	3,795
Dependent var mean	74.96	74.93	3.05	1.15	0.43
Adjusted R^2	0.498	0.501	0.755	0.538	0.132

Notes: This table reports the relationship between tax credit prices and project size (total units) in columns (1)–(2), and developer entry and award outcomes in columns (3)–(5). Column (2) excludes applications whose unit counts bunch at project-size ceilings imposed by the state’s Qualified Allocation Plan. Column (3)–(5) excludes state–years in which administrative records contain only awarded projects. The dependent variable is the number of LIHTC applications filed in the county–year in column (3), the number of projects awarded the credit in column (4), and the share of applications awarded in column (5). Tax Credit Price is the price per dollar of 10-year credit, residualized on pay-in schedule (percent of early commitment before construction starts, construction term, and their interaction), with syndicator FE absorbed. County controls include median gross rent, median household income, and log population. Entry regressions additionally exclude states with incomplete application records and state–years in which only awarded applications are kept. Standard errors are clustered by county. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

one at a binding cap. Cross-project regressions therefore do not identify how a given project would resize under a counterfactual tax credit price. The supply model in Section 5 addresses this limitation by allowing developers to re-optimize project size as prices change, subject to parcel-specific and institutional constraints and the competitive award process.

5 Model

This section develops an equilibrium model linking CRA-induced demand for tax credits to affordable housing supply. Developers choose project size, trading off the return to tax-credit equity and rental income against development costs and the probability of receiving an allocation. Banks choose LIHTC investment and whether to originate directly or through syndicators, with CRA incentives tied to their local deposit footprints. Equilibrium tax-credit prices clear local tax credit

markets, linking the spatial distribution of bank deposits to developer financing and affordable housing supply.

5.1 Affordable Housing Supply from LIHTC Developers

I model developer’s problem using the institutional setting of the Texas LIHTC program and application data from 2010 to 2024. Texas provides a particularly clean setting for modeling developers’ choices for two reasons. First, federal tax credits are the primary funding source of LIHTC projects in Texas, without layering additional state tax credits. This allows the return to project size to be mapped directly from federal tax-credit rules and observed tax-credit prices. In states such as California and Georgia, additional state credits introduce another subsidy margin that changes the return to project scale. Appendix Figure B5 illustrates this distinction across states. Second, the Texas data provide a large sample with detailed project-level information, including zoning constraints, pay-in schedules, pro formas, and scoring variables, that disciplines both the award process and the project size choice.

Developer’s problem. Each developer signs a contingent contract prior to applying for the LIHTC subsidy and holds the exclusive right to develop parcel i , incurring an entry cost κ_c .¹⁵ Parcels differ in fixed development costs and other characteristics, denoted f_i and Z_i . In principle, the developer chooses two margins: project size q and an affordability profile ϕ_i , defined as the unit-weighted average income limit relative to AMI.¹⁶ The developer solves

$$\max_{q, \phi_i} \left\{ \Pi_i^D(q, \phi_i) = \underbrace{\omega_i(q, \phi_i)}_{\text{award probability}} \underbrace{\pi_i(q, \phi_i)}_{\text{profit if awarded}} - \underbrace{\kappa_c}_{\text{entry cost}} \right\}. \quad (5)$$

I model the size choice conditional on the observed affordability profile. In practice, ϕ_i may reflect commitments made before the LIHTC application, through land lease agreements, funding commitments, or other project-specific arrangements that I do not observe. I therefore treat ϕ_i as predetermined and hold it fixed in the counterfactuals. The developer’s problem can consequently be written as

$$\max_q \Pi_i^D(q) = \omega_i(q) \pi_i(q) - \kappa_c. \quad (6)$$

¹⁵The entry cost is included for completeness but is not estimated because potential entrants are not observed. The estimated supply model and counterfactuals therefore condition on the observed applicant pool and do not model endogenous entry.

¹⁶I assign market-rate units an income limit of 100% of AMI. Alternative AMI assignments yield nearly identical results, as both mixed-income projects and the share of market-rate units are small in the estimation sample; see Appendix Figure A3.

Profit. Profit conditional on receiving an allocation consists of tax-credit equity, discounted rental cash flow, and a size-proportional profit shock, net of development costs:

$$\pi_i(q) = \underbrace{\pi_i^E(q)}_{\text{tax-credit equity}} + \underbrace{\pi_i^R(q)}_{\text{discounted rental cash flow}} + \underbrace{\sigma^D \varepsilon_i q}_{\text{profit shock}} - \underbrace{[f_i + vc_i(q)]}_{\text{development cost}}, \quad \varepsilon_i \sim N(0, 1). \quad (7)$$

Tax Credit Equity. Developers finance the project by selling tax credits to investors at the county-year price p_{ct} . The amount of 10-year tax credit which developers can request is proportional to project construction cost, adjusted for the basis boost and developer-fee allowance, subject to the QAP credit-request cap:

$$\begin{aligned} \pi_i^E(q) &= p_{ct} \min \{ \theta_{1t} \beta_i (1 + \delta_i) vc_i(q), \bar{c}_t \}, \\ \beta_i &\in [1, 1.3], \\ \delta_i &= \begin{cases} 0.20, & q \leq 49, \\ 0.15, & q > 49, \end{cases} \\ \bar{c}_t &= \begin{cases} \$20 \text{ million}, & t \leq 2012 \text{ or } t \geq 2022, \\ \$15 \text{ million}, & 2013 \leq t \leq 2021. \end{cases} \end{aligned} \quad (8)$$

Here, θ_{1t} converts eligible basis into the 10-year tax-credit stream.¹⁷ β_i is up to 130% if projects are located in a QCT, DDA, or TDHCA-designated high-opportunity area. δ_i is the developer-fee allowance as a share of construction cost—20 percent for projects of 49 units or fewer and 15 percent for larger projects—which developers typically exhaust, as shown in Appendix Figure A6. The credit cap $\bar{c}_t$ equals \$20 million through 2012 and from 2022 onward, and \$15 million during 2013–2021.

The credit-request cap in equation (8) therefore introduces a kink in the return to project size. Once the request cap binds, an additional unit generates no additional tax-credit equity. In the sample, 44.9 percent of urban applications are at the cap. This kink is important for both model fit and the counterfactuals: it generates bunching at cap-implied project sizes and dampens the size response of high-price urban projects to changes in tax-credit prices.

Rental Cash Flow. Developers collect net rental cash flows anchored to projected pro formas. I assume that LIHTC properties depreciate over 30 years. θ_2 is calibrated as the median ratio across pro formas of the discounted present value of projected 30-year net operating income relative to

¹⁷The 9 percent program allows projects to receive up to \$0.90 in ten-year credits per dollar of eligible basis, but in practice additional government subsidies, grants, and other gap financing may lower this conversion ratio. I use year-specific ratio to capture year-level variation in the availability of complementary financing.

$$\begin{aligned}
\phi_i \bar{r}_{ct} q_i \\
\pi_i^R(q) = \theta_2 \phi_i \bar{r}_{ct} (q + \rho_i q^2), \\
\rho_i = \rho^0 + \rho^R \mathbf{1}_i\{\text{Rural}\} + \rho^{SH} \mathbf{1}_i\{\text{SH}\} + \rho^S \mathbf{1}_i\{\text{Senior}\} + \rho^t.
\end{aligned} \tag{9}$$

Here, $\bar{r}_{ct}$ is the two-bedroom HUD rent limit implied by the county's area median income.¹⁸ The project's affordability profile ϕ_i is defined as the unit-weighted income limit as a fraction of AMI, so $\phi_i \bar{r}_{ct}$ captures the average rent permitted by the project's income limits. The factor θ_2 maps this rent into discounted pro-forma cash flows.

Concavity in size, $\rho_i < 0$, captures local eligible housing demand. As project size increases, the marginal unit draws from an increasingly thin pool of eligible tenants. The shifters ρ^R , ρ^{SH} , and ρ^S allow market depth to differ for projects targeting rural, supportive-housing, and senior population, respectively. The cohort effects ρ^t allow the curvature of rental cash flows to vary across application years.

Cost. Fixed cost f_i in equation (7) consists of project-specific land and non-basis soft costs. Construction cost vary with project size, density, and construction type:

$$vc_i(q) = \alpha_0 + \alpha_1 q + \alpha_2 \mathbf{1}_i\{\text{Rehab}\}q + \alpha_3 \frac{q}{\ell_i} + u_i q, \tag{10}$$

where α_0 and α_1 capture baseline construction cost and α_2 captures the cost difference for rehabilitation projects. ℓ_i is site acreage and α_3 captures the additional per-unit cost associated with developing on smaller sites. u_i is a project-specific per-unit cost component observed by the developer.

Winning probability. TDHCA allocates 9% credits through year-round competitions. In each application year, the QAP reserves an at-risk set-aside statewide and divides the remaining allocation across 13 service regions, each split into urban and rural subregions. Within each competition, applications are ranked and funded until the available credits are exhausted.

I model the award process as a conditional logit within these competition pools. Appendix C develops the conditional likelihood, the recursion used to compute inclusion probabilities, and the offset representation. Conditional on pool c funding K_c of its n_c applications, the probability that

¹⁸The sample mean of unit size is 1.82 bedroom. As a robustness check, I alternatively measure project size in bedrooms, which yields a comparable fit as reported in Appendix D.2.

project i in the winner set is

$$\Pr(i \in W \mid q) = \frac{\overbrace{\sum_{A \subseteq I_c, |A|=K_c, i \in A} \exp\left(\sum_{j \in A} v_j\right)}^{\text{feasible winner sets containing } i}}{\underbrace{\sum_{A \subseteq I_c, |A|=K_c} \exp\left(\sum_{j \in A} v_j\right)}_{\text{all feasible sets of } K_c \text{ winners}}}.$$

$$v_i(q) = \underbrace{\lambda_1 \mu_i(q)}_{\text{efficient use of subsidy}} + \underbrace{\lambda_2 \phi_i}_{\text{affordability target}} + \underbrace{Z_i' \gamma}_{\text{amenities and community support}}.$$

where the subsidy ratio μ_i , defined as total development cost relative to tax-credit equity, measures the leverage of other funding sources and enters the QAP as a measure of subsidy efficiency. Z_i includes the residualized QAP score, net of scoring components related to affordable profile ϕ_i and subsidy ratio μ_i . The residualized score mainly captures the remaining dimensions in the QAP scoring rules, primarily neighborhood amenities and measures of local community support. Z_i also includes indicators for the USDA and nonprofit set-asides.¹⁹ Pool-level factors are differenced out by the within-pool comparison.

From the developer's perspective, the winning probability admits an exact offset-logistic representation. The offset $\log \frac{e_{K_c-1}(w_{-i})}{e_{K_c}(w_{-i})}$ summarizes the strength of competing applications in the pool, and the resulting inclusion probabilities satisfy

$$\omega_i(q) = \Lambda \left(v_i(q) + \log \frac{e_{K_c-1}(w_{-i})}{e_{K_c}(w_{-i})} \right),$$

$$\sum_{i \in c} \omega_i = K_c.$$

where $\Lambda(x) = (1 + e^{-x})^{-1}$ denotes the logistic CDF. The affordability profile ϕ_i enters the award index directly, reflecting the QAP's preference for deeper income targeting. Project size q affects the award probability through the subsidy ratio $\mu_i(q)$. Expanding a project may increase its reliance on tax-credit equity, lowering its subsidy efficiency and, in turn, its probability of receiving an allocation. Appendix C.5 discusses this potential concern and evaluates the validity of holding the number of winners fixed in the counterfactual.

¹⁹TDHCA designates USDA and nonprofit set-asides, but unlike the At-Risk set-aside and geographic service regions, these two set-asides do not constitute separate competition pools. I therefore treat USDA and nonprofit status as project attributes that receive preferential treatment directly in the allocation process.

Optimal project size. Given the profit shock, the developer chooses

$$\max_q \omega_i(q) [\bar{\pi}_i(q) + \sigma^D \varepsilon_i q], \quad (13)$$

where $\bar{\pi}_i(q)$ denotes profit excluding the idiosyncratic profit shock. An interior optimal project size q^* satisfies

$$\omega_i(q) [\bar{\pi}'_i(q) + \sigma^D \varepsilon_i] + \omega'_i(q) [\bar{\pi}_i(q) + \sigma^D \varepsilon_i q] = 0. \quad (14)$$

The first term captures the direct marginal return to expanding the project, while the second captures the effect of project size on the probability of receiving an allocation. Below the credit-request cap, a higher tax-credit price raises the marginal return to size and increase the optimal project size q^* . A higher price can also affect the award probability through the subsidy ratio, although this channel is quantitatively smaller in the estimated model.

Delivered project size. Five institutional features map the interior optimum q_i^* into delivered project size q_i . Projects are subject to a 16-unit minimum, an 80-unit cap for rural new construction (NC) and adaptive reuse (AR) projects, a 252-unit cap for urban new construction and adaptive reuse projects before 2013, and zoning-imposed maximum unit limits. In addition, developers may adopt an existing architectural template given similar parcel attributes, generating bunching at a discrete set of project sizes $\mathcal{M}$ that accounts for 38.8% of observed projects.²⁰

$$\begin{aligned} q_i &\geq 16, \\ q_i &\leq 80, \text{ if } \mathbf{1}_i\{\text{Rural}\} = 1 \text{ and } i \text{ is NC or AR,} \\ q_i &\leq 252, \text{ if } \mathbf{1}_i\{\text{Urban}\} = 1, i \text{ is NC or AR, and } t \leq 2013, \\ q_i &\leq \bar{q}_{\text{zoning}}, \\ \mathcal{M} &= \{24, 36, 48, 60, 72, 80, 96, 100, 120, 152, 180, 200\}, \end{aligned} \quad (15)$$

$$\Pr(q_i \in \mathcal{M} \mid q_i^*) = \Lambda \left(\kappa_0 - \kappa_1 \min_{m \in \mathcal{M}} |q_i^* - m| \right).$$

Template adoption is more likely when the interior optimum q_i^* is close to a standardized size m , capturing the cost savings from reusing existing designs. Otherwise, the developer delivers the custom size implied by the constrained optimum.

²⁰As an example of template adoption, Kaplan Thompson Architects describes its work: “we designed a prototype for a Passive House multi-family building that could be deployed at multiple sites and achieve certification at different orientations and locations.” The architect explicitly attributes cost and design efficiencies to repetition. ([Source](#))

5.2 Tax Credit Demand from Bank Investor

Banks demand LIHTC equity both for its tax benefits and for its value toward CRA obligations. I model the CRA incentive as a deposit-based investment benchmark, so a bank's demand for LIHTC equity depends on its local deposit footprint. Banks can meet this demand either by originating investments directly, subject to a fixed underwriting cost, or through a syndicator that charges a proportional fee. The choice between these channels depends on the size of the bank's equity commitment and its cost of direct origination. County-level tax-credit prices are determined by aggregating these bank-level demand schedules and equating total bank demand with the equity value of tax credits supplied by developers.²¹

Bank's problem. Let a_{ct}^b denote the dollar amount of LIHTC equity invested by bank b in county c over evaluation period t . Because CRA exams cover multiple years, I aggregate the model to four-year periods, summing flows, averaging prices, and defining direct origination and desk costs over each period.

$$\pi_{ct}^b(a) = \underbrace{\nu_t \frac{a}{p_{ct}}}_{\text{tax benefit}} - \underbrace{a}_{\text{equity investment}} - \underbrace{\min\{F_s^b, \tau a\}}_{\text{investment cost}} - \underbrace{\frac{\eta}{2} D_{ct}^b \left(\frac{\tilde{a}}{D_{ct}^b} - \zeta_c + \sigma^\zeta \varepsilon_{ct}^\zeta \right)^2}_{\text{cost of CRA deviation}}, \quad (16)$$

$$\tilde{a} = (1 + \tau \mathbf{1}\{d_{ct}^b = 0\}) a + F_s^b \mathbf{1}\{d_{ct}^b = 1\}, \quad \varepsilon_{ct}^\zeta \sim N(0, 1).$$

Here ν_t is the fundamental value of the tax credits a/p_{ct} , measured by the price paid in counties with no large-bank deposits and averaged over the years in t , and p_{ct} is their market price.²² Banks can originate LIHTC investments either directly by operating an internal underwriting desk at fixed cost F_s^b per evaluation period, or indirectly through a syndicator, which charges a proportional fee τ .²³

The final term captures the CRA incentive to invest in proportion to the bank's local deposit footprint. The baseline benchmark for 9% LIHTC investment is $\zeta_c D_{ct}^b$, where ζ_c scales a common target by the county's 9% share of total LIHTC credits.²⁴ The shock $\sigma^\zeta \varepsilon_{ct}^\zeta$ captures unobserved

²¹I classify a bank as direct-capable if I observe at least two applications in which an equity commitment letter is issued directly by the bank or its community-investment subsidiary. This classification captures banks with observed in-house LIHTC origination capacity. It is distinct from a proprietary fund, which is a single-investor fund managed by a third-party syndicator on behalf of the bank; investments through such funds are classified as syndicated.

²²I do not explicitly model other tax-credit investors, whose demand reflects the tax benefits of LIHTC investment and should not exhibit the geographic preferences induced by the CRA. Their valuation is therefore absorbed into the common non-CRA value ν_t .

²³I treat the syndication fee as constant across markets, so its effect on equilibrium tax-credit prices operates through the share of syndicated deals. Appendix Table A4 finds limited variation in syndication costs with the fragmentation of large-bank deposits, particularly in thin-deposit counties where only one or two large banks are typically present.

²⁴Counties are an approximation to CRA assessment-area geography. In practice, banks may combine neighboring counties, but assessment areas vary across banks and over time, while counties provide the finest common unit for

county-period factors that shift banks' demand for 9% LIHTC relative to other CRA-qualified activities, so that the effective benchmark is $(\zeta_c - \sigma^\zeta \varepsilon_{ct}^\zeta) D_{ct}^b$. The parameter η governs the shadow cost of deviating from this benchmark, reweighted by local deposits D_{ct}^b . The examined commitment $\tilde{a}$ is the bank's total qualified investment: for syndicated investments it includes the fee τa , and for direct investments it includes the desk cost F_s^b . Commitments are bounded below by zero.

The first-order condition gives bank b 's marginal willingness to pay in county c ,

$$p_{ct}^b(a) = \frac{\nu_t}{1 + \tau_d - \eta(1 + \tau_d) \left(\zeta_c - \sigma^\zeta \varepsilon_{ct}^\zeta - \frac{\tilde{a}}{D_{ct}^b} \right)}, \quad \tau_d \equiv \tau \mathbf{1}\{d_{ct}^b = 0\}, \quad (17)$$

so that direct origination is the case $\tau_d = 0$ with $\tilde{a} = a + F_s^b$. A bank whose LIHTC investment is low relative to its preferred allocation is willing to pay a premium for tax credits. This premium declines as investment approaches the preferred allocation, generating downward-sloping demand.

Inverting equation (17) gives bank b 's demand,

$$a_{ct}^b(p) = \max \left\{ \frac{\zeta_c - \sigma^\zeta \varepsilon_{ct}^\zeta}{1 + \tau_d} D_{ct}^b + \frac{D_{ct}^b}{\eta(1 + \tau_d)^2} \left(\frac{\nu_t}{p} - 1 - \tau_d \right) - F_s^b \mathbf{1}\{d_{ct}^b = 1\}, 0 \right\}. \quad (18)$$

Demand is increasing in local deposits and decreasing in price. A bank that originates directly buys F_s^b less equity than its examined commitment, because the desk cost itself earns examiner credit. As $\eta \rightarrow \infty$, the examined commitment converges to $(\zeta_c - \sigma^\zeta \varepsilon_{ct}^\zeta) D_{ct}^b$.

Channel choice. Each bank chooses between direct investment and investing through a syndicator. Direct origination requires a fixed underwriting cost F_s^b per evaluation period, whereas syndication costs a fraction τ of the equity commitment. For direct-capable banks I parameterize the fixed cost as

$$F_s^b = \exp(\alpha_b + \alpha_s), \quad \sum_s \omega_s \alpha_s = 0, \quad (19)$$

where α_b captures persistent differences in internal underwriting capacity and origination strategy, and α_s captures differences in state-specific fixed cost and deal-sourcing relationships. The LIHTC market is decentralized: credits are allocated through state housing agencies that differ in QAP rules, application procedures, complementary subsidies, and developer networks, and these costs are easier to amortize in states with thicker LIHTC markets. The weights ω_s are state shares of the bank–county–period observations in which a channel choice is observed. I normalize the weighted average state effect to zero, so α_b captures the bank-specific component of log origination costs.

deposits, tax-credit supply, and investment. This approximation may overstate fragmentation if banks substitute investments across counties within an assessment area, although competitive pricing should limit persistent price differences across such counties.

Because the desk cost is itself credited, a bank that originates directly buys $\tilde{a} - F_s^b$ of equity and avoids the fee on that amount, so the comparison $F_s^b < \tau a$ becomes $F_s^b < \tau \tilde{a}/(1 + \tau)$. A bank originates directly when the fee it avoids, together with an idiosyncratic benefit of originating in house, exceeds the fixed cost:

$$d_{ct}^b = \mathbf{1} \left\{ \frac{\tau}{1 + \tau} \tilde{a}_{ct}^b - F_s^b - \xi_{ct}^b > 0 \right\}, \quad \xi_{ct}^b \sim \text{Logistic}(0, \sigma^d), \quad (20)$$

where $\tilde{a}_{ct}^b$ is the bank's examined commitment under direct origination, from equation (17) with $\tau_d = 0$. Mixing the two channels within a county-year is rare (2.0%) even over a full CRA evaluation cycle (5.4%), as illustrated in Appendix Figure B3. The shock ξ_{ct}^b captures deal-specific benefits of in-house origination, such as a relationship with the sponsor or economies with the bank's other lending on the project, and σ^d , measured in dollars, governs their dispersion. The realized desk cost $F_s^b + \xi_{ct}^b$ is floored at zero. The probability of direct origination is

$$\Pr(d_{ct}^b = 1 \mid \tilde{a}_{ct}^b, \alpha_b, \alpha_s) = \Lambda \left(\frac{\tau \tilde{a}_{ct}^b / (1 + \tau) - F_s^b}{\sigma^d} \right), \quad \Lambda(x) = \frac{e^x}{1 + e^x}. \quad (21)$$

Banks are thus more likely to originate directly when their fixed cost of direct investment is lower and when their commitment, and with it the fee they avoid, is larger. Lengthening the evaluation period raises the commitment behind a single desk payment without raising the cost proportionally, so desks become more attractive the longer the cycle. Banks never observed originating directly are assumed to invest exclusively through syndicators. For these banks, origination costs are only set-identified as sufficiently high to make syndication optimal, and I assume they continue to invest through syndicators in the counterfactuals.

County price determination. Bank commitments, channel choices, and equilibrium prices are jointly determined, since commitments depend on prices while prices depend on the share of investment that is syndicated. Market clearing equates the equity banks actually place—their commitments net of fees and desk costs—with the equity the awarded credits must raise. In counties with no large-bank deposits, the CRA deviation cost is absent and outside investors clear the market at the fundamental value ν_t , adjusted for project-specific pay-in and risk characteristics. These counties therefore do not enter the large-bank demand equilibrium and are used to calibrate ν_t .²⁵

This aggregation highlights three sources of spatial price variation. First, a larger local deposit base shifts aggregate demand outward. Second, syndicated investments incur the proportional fee

²⁵The model abstracts from within-county-period price dispersion across investors in Table 5. County-year fixed effects alone explain 89 percent of the variation in project-level tax credit prices (Figure A2). I instead focus on cross-county price variation generated by local CRA demand and credit supply.

τ , lowering banks' willingness to pay relative to direct origination, while direct investments place F_s^b less equity per dollar of commitment. Finally, ε_{ct}^ζ captures county-period variation in banks' preferred allocation to 9% LIHTC relative to other CRA-qualified activities. Conditional on deposits and credit supply, a higher ε_{ct}^ζ lowers LIHTC demand and the equilibrium price.

5.3 Equilibrium

Given project characteristics, local bank deposits, and model parameters, an equilibrium is defined by developer project sizes q_i , award probabilities $\{\omega_i\}$, bank equity investments and origination channels $\{a_{ct}^b, d_{ct}^b\}$, and county-year tax-credit prices $\{p_{ct}\}$ such that:

1. Developers optimize: Given tax-credit prices, developers choose project sizes to maximize expected profits according to equation (13), subject to the institutional constraints in equation (15). Credit requests are determined by equation (8).
2. Banks optimize: Given tax-credit prices, banks choose LIHTC equity investment according to equation (18) and origination channels according to equation (21).
3. Tax-credit markets clear: In each county-year, aggregate bank demand for LIHTC equity equals the equity value of expected allocated credit supply,²⁶

$$\sum_b a_{ct}^b(p_{ct}) = p_{ct} \sum_{i \in (c,t)} \omega_i(q_i) c_i(q_i). \quad (22)$$

6 Estimation

6.1 Affordable Housing Supply

Estimation proceeds in three blocks. The award parameters are estimated by conditional maximum likelihood; the cost parameters by least squares on per-unit costs; and the revenue-curvature, dispersion, and template parameters by SMM on the delivered-size distribution. The award process is directly observed from QAP scores and allocation outcomes, while institutional kinks, caps, and template bunching distort observed project sizes and could contaminate award estimates in a joint system. I therefore estimate the award process first and the size parameters conditional on these estimates.

²⁶The supply and demand model operates at two time scales. Bank commitments, origination choices, and market clearing are determined over CRA evaluation periods, while developers choose project size annually; the resulting county-period CRA premium is mapped into annual prices through the year-specific fundamental value ν_t .

Award block. I estimate the award parameters $(\lambda_1, \lambda_2, \gamma)$ by maximizing the conditional likelihood of observed winner sets, $\sum_c \log \Pr(W_c)$, from equation (11). Competition pools are defined by the statewide At-Risk set-aside and the urban and rural subregions within each service region, yielding 394 pool–year cells. Conditioning on the number of awards K_c in each pool, identification comes entirely from comparisons of winners and non-winners within the same competition, and all pool-level heterogeneity is differenced out. Standard errors are clustered at the pool–year level.

Cost block. I estimate $(\alpha_0, \alpha_1, \alpha_2, \alpha_3)$ by regressing per-unit construction cost, vc_i/q_i , on inverse project size, a constant, the rehabilitation indicator, and inverse acreage. I treat the resulting residual $\hat{u}_i$ as a project-specific cost component and hold it fixed when evaluating alternative project sizes.

Size block. I calibrate the credit-to-basis conversion θ_{1t} to the median ratio of 10-year credit requests to developer fee-grossed boosted basis among year- t applications. θ_2 is calibrated to the median ratio of the present value of 30-year projected net operating income to $\phi_i \bar{r}_{ct} q_i$.

Given the award and cost estimates and the calibrated (θ_{1t}, θ_2) , I estimate $\Theta^S = (\rho^0, \rho^{SH}, \rho^S, \rho^R, \sigma^D, \kappa_0, \kappa_1)$ by SMM. For each application, I simulate profit and template-adoption shocks, solve the developer’s size problem on a grid, and impose the rural cap, 16-unit floor, zoning, and template-adoption rule. I match 21 moments of the delivered size distribution: the mean, standard deviation, and the 10th, 25th, 50th, 75th, and 90th percentiles; the shares above 120 units and at or below 49 units; the share delivered exactly at a template size and the shares at 1–2, 3–5, 6–10, and 11 or more units from the nearest template size; the rural mean and standard deviation, the share of rural new-construction projects at the 76–80 unit cap, and the rural share at or below 49 units; the senior and supportive-housing means; and the share of urban applications requesting the statutory credit cap. Cohort effects ρ^t are calibrated to match mean project size within each application year. Standard errors are computed from the SMM sandwich, accounting for simulation error and sampling uncertainty in the first-stage award estimates.

Results. Table 8 reports the estimates. In the award block, the residualized QAP score, which primarily captures project amenities and local community support, strongly predicts allocation within an award pool. Consistent with the QAP’s set-aside provisions, USDA participation receives preferential treatment in the allocation process, whereas the nonprofit set-aside does not translate into a detectable advantage in award outcomes conditional on the other components of the QAP. The estimates also show that projects leveraging more external financing are more likely to receive an award, consistent with the QAP’s emphasis on efficient use of tax-credit subsidies.

Conditional on these award determinants, the estimated coefficient on the affordability profile is negative, consistent with deeper affordability being favored, but plays a limited role in deter-

Table 8: Supply-Side Key Parameter Estimates

Coefficient	Parameter	Estimate	SE
Panel A: Award stage (MLE)			
Residualized QAP score	γ_1	0.168***	(0.046)
USDA financing	γ_2	1.049***	(0.353)
Nonprofit set-aside	γ_3	0.142	(0.149)
Subsidy ratio μ_i	λ_1	0.430**	(0.189)
Affordability profile ϕ_i	λ_2	-0.180	(1.616)
Panel B: Per-unit cost (OLS)			
Fixed eligible basis (\$M)	α_0	0.967***	(0.141)
Per-unit cost, new construction	α_1	11.061***	(0.174)
Rehabilitation shifter	α_2	-4.353***	(0.242)
Density (inverse acreage)	α_3	1.219***	(0.302)
Panel C: Size stage (SMM)			
Revenue curvature	ρ^0	-0.317***	(0.016)
Supportive-housing shifter	ρ^{SH}	-0.277***	(0.066)
Senior shifter	ρ^S	-0.112***	(0.033)
Rural shifter	ρ^R	-0.460***	(0.021)
Profit-shock dispersion	σ	1.930***	(0.071)
Template adoption: level	κ_0	-0.195	(0.226)
Template adoption: distance decay	κ_1	0.111**	(0.048)
Applications		1,910	
Awards		961	

Notes: Table estimates are based on 1,910 Texas LIHTC applications from 2010–2024, of which 961 received an award. Panel A reports conditional maximum-likelihood estimates of the award block over 394 pool–year cells, with standard errors clustered at the pool–year level. Panel B reports the per-unit cost projection with heteroskedasticity-robust standard errors. Panel C reports the SMM estimates of the revenue-curvature, dispersion, and template parameters, with sandwich standard errors that account for the sampling error in Panel A and condition on Panel B. For estimation, costs and land values are measured in millions of dollars, the two-bedroom rent limit at 100 percent of AMI in thousands of dollars per month, and project size in hundreds of units. * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$


mining award outcomes. Winners and losers differ in mean affordability by only 0.003. This is also consistent with affordability commitments being largely pre-determined through the project’s financing arrangements, land-use agreements, and other funding requirements.

On the cost side, the estimates imply a baseline construction cost of approximately \$110,600 per unit for new construction, with acquisition–rehabilitation projects costing about \$43,500 less per unit. Construction costs are also higher on smaller parcels with higher density.

The size-choice estimates capture several institutional features of LIHTC development. Rural, supportive-housing and senior projects exhibit stronger diminishing returns to project scale, con-

sistent with the smaller potential tenant base. The model also captures bunching at standardized project sizes. When the unconstrained optimum coincides with a template size in $\mathcal{M}$, the developer adopts the template with probability of about 45 percent. This probability declines with distance from the template menu.

Figure 4: Model Fit of Supply Estimates



Notes: Panel (a) compares the distribution of delivered project sizes in the data and model simulations; the dotted line denotes the 80-unit cap for rural new construction and adaptive-reuse projects. Panel (b) plots binned averages of realized award rates against model-implied inclusion probabilities by decile; the dashed line denotes the 45-degree line.

Model fit. The model closely reproduces both the distribution of project size and the allocation of tax credits. On the size margin, it matches the main percentiles, the upper tail, and the substantial bunching at standardized project sizes. The rural cap, template adoption, and credit-request cap account for substantial concentrations in the observed size distribution. On the allocation margin, realized award outcome closely track predicted award probabilities along the 45-degree line.

6.2 Tax Credit Demand

I estimate the demand model by SMM. For each candidate parameter vector

$$\Theta^D = (\tau, \eta, \sigma^\zeta, \sigma^d),$$

I solve each bank's optimal investment under direct origination and syndication and determine its channel choice by comparing the resulting optimized payoffs. I then jointly simulate banks'

channel choices, county equilibrium prices, and investment quantities on the observed distribution of deposits and tax-credit supply. The four parameters are estimated from nine moments of prices and origination choices.

For each candidate parameter vector, the bank and state components of the direct-origination cost $\{\alpha_b\}$ and $\{\alpha_s\}$ are set so that each bank's and each state's simulated direct-origination rate equals its observed rate. I calibrate $\zeta_c = \zeta \chi_c$, where ζ is the merger-based IV estimate of LIHTC investment per dollar of deposits from Table 4 and χ_c is the county's 9% share of total LIHTC credits. I set ν_t to the tax-credit price paid in counties with no large-bank deposits, averaged over the years in period t .

The estimation sample contains 1,548 county-periods from 2010–2024, formed as four-year evaluation windows, and 14,965 bank-county-period observations. Deposits, credit supply and equity are summed within a window, and the tax-credit price is total equity divided by the total face value of credits, so it is the price of the window's bundle. I classify a named bank as direct-capable if it issues at least two direct equity commitments over the sample, and a bank counts as originating directly in a window if it issues at least one such commitment during it. Each of the 13 direct-capable banks receives its own estimated bank effect, while the effects of banks without information on direct origination are only set-identified.

Moments and identification. For each candidate Θ^D , I draw county-period benchmark shocks and jointly solve for banks' channel choices and equilibrium prices. The SMM estimator matches moments relating prices and channel choices to local deposits, credit supply, and syndication, as well as residual price dispersion and the deal-level hedonic syndication discount. I instrument local credit supply using QCT and DDA designations and a shift-share instrument based on counties' pre-2005 shares of state credit awards.

These moments separately discipline the four parameters. τ is identified from the price block, principally the slope of the county price on the syndicated share, with the deal-level hedonic discount imposed as an additional target that currently supplies most of the precision. The within-bank relationship between direct origination and own-county deposits identifies σ^d ; the response of prices to deposits and $\tilde{S}_{ct}/D_{ct}$ identifies η ; and residual within-state-period price dispersion identifies σ^ζ . Appendix D.3 provides additional estimation details.

Results. Table 9 reports the estimates. In the baseline SMM specification, the syndication cost is \$0.0149 per dollar of investment. The cost of deviating from the CRA benchmark is 0.664, while the dispersion in the CRA benchmark shock is 0.0600 and the origination-shock scale is 0.720. Incorporating the IV supply-to-deposit price gradient lowers the syndication cost slightly to \$0.0142 and the origination-shock scale to 0.653, raises the cost of CRA deviation to 0.765, and

Table 9: Key Parameter Estimates of Demand Model

	Parameter	SMM	SMM (+IV)
Syndication fee	τ	0.0149*** (0.0025)	0.0142*** (0.0024)
Cost of CRA deviation	η	0.664*** (0.107)	0.765*** (0.111)
CRA benchmark shock (s.d.)	σ^ζ	0.0600*** (0.0102)	0.0495*** (0.0087)
Origination-shock scale (\$M)	σ^d	0.720*** (0.240)	0.653*** (0.232)
First-step 2SLS slope, price on $\tilde{S}/D$	β^{IV}		-0.758*** (0.161)
First-stage F (clustered by county)			14.9
Criterion J (standard-error weighted)		16.62	18.66
Bank–county–period cells		14,965	

Notes: This table reports SMM estimates of the demand model, in which the unit of observation is the bank–county–evaluation period and the underwriting desk is a fixed cost paid once per period. The SMM specification targets eight moments of prices and origination choices; SMM (+IV) additionally targets the within-state–period price gradient with respect to the supply-to-deposit ratio. τ denotes the proportional syndication cost, η the shadow cost of deviations from the CRA benchmark, σ^ζ the standard deviation of the county–period benchmark shock, and σ^d the scale of the direct-origination shock, measured in millions of dollars. Bank and state direct-origination cost effects are concentrated out to match their respective origination rates; their estimates are shown in Figure 5. Moments are weighted by the inverse of their state-block bootstrap variances. Standard errors are bootstrapped from 50 random draws. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.


lowers the dispersion in the CRA benchmark shock to 0.0495. The targeted IV estimate of the supply-to-deposit price gradient is -0.758, with a first-stage F -statistic of 14.9.

Out-of-sample validation. I validate the estimated direct-origination costs using OCC Part 24 filings, an independent data source that records banks’ community-development investments, including investments in 4 percent LIHTC projects. Figure 5 compares the estimated bank and state cost effects with composition-adjusted syndication shares in the OCC data. The estimated costs capture persistent differences in direct-origination cost and higher costs predict less direct investment in the OCC filings.

7 Counterfactual

The counterfactual analysis quantifies how fragmented CRA demand affects the allocation and efficiency of the LIHTC program, holding government tax credits and investor equity fixed. I proceed

Figure 5: Out-of-Sample Validation



Notes: This figure compares estimated direct-origination costs with composition-adjusted direct-investment shares in OCC Part 24 filings, which are not used in estimation. Panels (a) and (b) report bank and state effects, respectively, evaluated at the average effect of the other dimension. Dashed lines show linear fits. Four state-chartered banks are excluded because they do not file Part 24 reports with the OCC. JPMorgan Chase is also excluded because its estimated cost is based on only two direct commitments. The OCC measure is an imperfect proxy for investment channel based on the identity of the reported investment vehicle as discussed in B.4.

in two parts. The first three exercises use the supply model to isolate the redistributive consequences of spatial variation in tax credit prices. I first remove the CRA-induced price premium, then equalize prices across counties, and finally impose budget neutrality through alternative adjustments to prices and credit sizing. These exercises trace how the geography of CRA demand redistributes investor equity and housing production across markets. The final exercise incorporates the demand model to study the syndication cost of spatially fragmented CRA incentives. I replace county assessment areas with statewide assessment areas and allow banks to re-optimize their origination channels, quantifying both the gains from pooling CRA demand across counties and the associated savings in syndication and origination costs. Appendix B.9 provides additional discussion of the welfare implications of these geographic redistribution.


7.1 Redistribution

Non-CRA and equalized tax credit prices. I first use the estimated supply model to isolate how the spatial pattern of tax credit prices affects housing supply, holding the set of awarded projects fixed. For each application, I replace its observed tax credit price with an alternative price and allow the developer to re-optimize project size. These exercises therefore capture the intensive-margin

supply response to tax credit prices, rather than changes in award outcomes.

The first exercise examines how CRA incentives provide an additional source of demand for LIHTC investments and shape the spatial distribution of housing supply through their effect on tax credit prices. I construct a non-CRA-incentive price from projects located in counties where large banks hold no deposits at all, the closest empirical analogue to a market without CRA incentives over 2012-2024.²⁷ I adjust these prices for pay-in timing, deal risk, lender structure, and state composition, and assign the resulting price to Texas applications. As shown in Panel (a) of Figure 6, the tax credit price is 4.5 cents per credit dollar below the Texas credit-weighted average and re-optimizing project size at this price reduces the units of awarded projects by 3,164 (-4.32 percent) in total. Panel (b) in Figure 6 shows that CRA incentives are concentrated in the top deposit quartile with substantially higher premium in tax credit prices, and removing the incentive reduces expected supply by 2,971 units (4.93 percent), accounting for 94 percent of the statewide loss. By contrast, expected supply in the bottom deposit quartile falls by only 0.10 percent.

Figure 6: CRA Incentives, Tax Credit Prices, and Housing Supply



Notes: Panel (a) plots the annual credit-weighted average Texas tax credit price and the non-CRA-incentive price, estimated from projects in counties with no large-bank deposits after adjusting for project characteristics and state composition. Error bars report 95 percent confidence intervals. Panel (b) reports changes in expected awarded units when projects are assigned either the non-CRA-incentive price or the annual Texas average price and developers re-optimize project size, holding award outcome fixed. Changes are relative to baseline.

The second exercise retains the CRA incentives but removes its spatial constraint, allowing banks to invest statewide until tax credit prices equalize within each year. I therefore assign

²⁷The reporting window starts in 2012 to exclude the post-crisis recovery of the investor market. In 2009–2011, bank taxable income and direct LIHTC investment had collapsed and were rebuilding, tax credit prices cleared at 0.69–0.77 dollars per credit dollar in Texas.

each project the credit-weighted average Texas price in the same year and allow developers to re-optimize project size, holding awards fixed. As shown in Figure 6(b), equalizing prices increases expected units by 7.18 percent in the bottom deposit quartile, 7.21 percent in the second, and 6.18 percent in the third, compared with only -0.09 percent in the top quartile. The gains are concentrated in low-deposit counties, where removing the spatial constraint of CRA demand raises tax credit prices and expands housing supply.

These two exercises isolate the supply response and spatial dispersion of tax credit prices without imposing budget neutrality or market clearing. Developers facing higher counterfactual tax credit prices optimally choose larger projects and request more tax credits, so part of the increase in housing supply is financed by higher government tax credit expenditures. The next subsection removes the spatial constraint on CRA investment while holding the government tax credit budget fixed.


Budget-neutral price equalization. I next remove the spatial constraint on CRA investment, allowing banks to invest statewide until tax credit prices equalize within each year, while holding the government's tax credit budget fixed. In each year, tax credits trade at a common statewide price, and developers re-optimize project size and credit requests, with award probabilities adjusting through the subsidy ratio. The common price adjusts so that expected awarded tax credits remain at their baseline level in each year, removing the spatial price variation while holding the government's credit allocation fixed.

As shown in Figure 7(a), the budget-neutral common price is on average 1.5 cents below the observed credit-weighted average price. Raising prices in low-price markets induces developers to choose larger projects and request more credits, so keeping the same credit budget fixed requires a lower statewide price. Equalizing prices increases expected awarded units by 168 (0.23 percent), from 72,382 to 72,550, while reducing investor equity by \$148 million (1.69 percent).

The aggregate effect masks substantial spatial redistribution. As shown in Figure 7(b), expected units increase by 104 (4.33 percent), 134 (3.96 percent), and 207 (2.81 percent) across the bottom three deposit quartiles but fall by 277 (0.47 percent) in the top quartile. Investor equity is redistributed similarly: an additional \$61 million in the bottom three quartiles supports 444 more units (3.38 percent), while a \$208 million reduction in the top quartile is associated with only 277 fewer units (0.47 percent). Thus, the common price reallocates equity toward markets where housing supply is more elastic to tax credit equity.

This asymmetry partly reflects the per-project credit cap. 42% of applications in the top deposit quartile are at the credit cap, compared with essentially none in the lower quartiles. For capped projects, higher prices increase equity on existing credits without increasing the return to project size, making housing supply less elastic to tax credit equity. Equalizing prices therefore shifts

Figure 7: Budget-Neutral Equalization of Tax Credit Prices



Notes: Panel (a) compares the annual credit-weighted average Texas tax credit price with the budget-neutral common price that holds expected awarded tax credits at their baseline level. Panel (b) reports percentage changes in expected awarded units and investor equity by county large-bank deposit quartile and statewide. Developers re-optimize project size and credit requests under the common price, with award probabilities adjusting through the subsidy ratio.

equity toward the more elastic lower-deposit markets, allowing the same credit budget to support more housing. The \$148 million reduction in investor equity is also large relative to the resulting losses: at the top-quartile average development cost of \$186,900 per unit, replacing the 277 lost units would cost about \$52 million. Thus, the interaction between spatially dispersed CRA-driven tax credit prices and the LIHTC credit cap creates a misallocation large enough that equalizing prices increases housing supply and generates sufficient savings to offset the losses in high-deposit markets, while leaving a net surplus.


Budget-neutral price equalization with credit sizing. The price-adjustment closure holds the government’s credit budget fixed but allows total investor equity to fall by \$148 million, so translating its equity savings into additional housing requires an external calculation. I therefore consider an alternative closure that holds both the government credit budget and total investor equity fixed.

I implement this closure by allowing the agency to resize each project’s tax-credit allocation. Credit sizing means adjusting the amount of credits awarded per dollar of eligible basis to reflect the project’s financing need: when a given amount of credits raises more equity with higher tax credit prices, or when other funding sources fill more of the financing gap, fewer tax credits are needed. In the model, I approximate this re-underwriting by setting a common tax credit price within each year equal to baseline investor equity per awarded credit and scaling θ_{1t} , the conversion from boosted eligible basis into tax credits, by a common annual factor that holds expected awarded ten-year

credits at baseline. This scaling does not change the statutory credit rate; it is a reduced-form representation of project-level credit sizing under the equalized price. The statutory per-project credit cap remains unchanged, developers re-optimize project size, and award probabilities adjust through the subsidy ratio.

The required adjustment in credit sizing is small. Figure 8(a) shows that the scaling factor averages 0.992 and reduces credits per dollar of eligible basis by at most 2.1 percent. The resulting common price differs from the credit-weighted average baseline price by at most 0.7 cents. Figure 8(b) shows the spatial incidence. Under credit sizing, expected units increase by 125 (5.21 percent), 172 (5.11 percent), and 289 (3.92 percent) in the bottom three deposit quartiles and by 254 (0.43 percent) in the top quartile, compared with 104 (4.33 percent), 134 (3.96 percent), 207 (2.81 percent), and -277 (-0.47 percent) under price adjustment. Under price adjustment, the bottom three quartiles gain \$61 million in investor equity, while the top quartile loses \$208 million and statewide investor equity falls by \$148 million. Under credit sizing, which holds statewide investor equity fixed, approximately \$86 million is reallocated from the top quartile to the bottom three quartiles.

Figure 8: Budget-Neutral Equalization of Tax Credit Prices



Notes: Panel (a) plots the estimated credit-to-basis ratio $\theta_{1,t}$, defined as the median ratio of requested ten-year credits to boosted, fee-grossed eligible basis, with 95 percent bootstrap confidence intervals. The dashed series reports the counterfactual ratio under the credit-sizing closure, which scales $\theta_{1,t}$ annually to hold expected awarded credits at baseline while equalizing prices. The scaling factor averages 0.992, with a maximum reduction of 2.1 percent. Panel (b) reports changes in expected awarded units by county large-bank deposit quartile and statewide. Dark bars equalize prices while holding the government credit budget fixed; light bars additionally adjust credit sizing to hold total investor equity fixed. The price closure reduces investor equity by \$148 million. In both exercises, developers re-optimize project size and award probabilities adjust through the subsidy ratio. Expected awarded units equal the sum of award probability times project size across 1,654 applications from 2012–2024. Counties are grouped by 2015 large-bank deposits.

The aggregate effects differ because the credit-sizing closure keeps the equity savings inside the program. Expected awarded units increase by 840 (1.16 percent), compared with 168 (0.23 percent) with a small adjustment in θ_{1t} . The additional 672 units therefore measure, within the model, the housing produced when the price closure's equity savings remain in the program rather than accruing to investors. For comparison, a simple calculation using average development costs implies that the \$148 million reduction in investor equity under the price closure could finance roughly 790 additional units.

The difference between the two closures reflects the interaction between credit sizing and the per-project credit cap. A marginal reduction in θ_{1t} affects only projects below the cap and has little effect on their optimal size, while capped metropolitan projects retain their credit allocations rather than absorbing the price reduction required under the price-adjustment closure. Because θ_{1t} is a year-specific conversion rate calibrated from the median ratio of requested credits to eligible basis, the exercise should be interpreted as a reduced-form approximation to project-level underwriting rather than a literal uniform change in agency policy.

7.2 Efficiency Gains from Statewide Assessment Areas


I use the estimated model to quantify the efficiency gains from re-optimizing banks' investment channels. The counterfactual holds fixed each bank's total commitment and the aggregate tax credit budget, allows banks to re-optimize their investment channels, and lets a common tax credit price clear the statewide market.

Counterfactual equilibrium. For each state and evaluation period, banks choose where to operate direct-origination desks to minimize total origination costs, subject to bank commitments and county credit supply. Banks consolidate desks where fixed costs exceed the syndication fees they avoid and concentrate direct origination where it is most cost effective; the remaining commitment is routed through syndicators.

Computationally, I jointly solve for the cost-minimizing match between banks and projects and the market-clearing tax credit price. Channel choices determine origination costs and bank demand, while the statewide price adjusts to clear the tax credit market. I solve this problem using the same simulation draws as in the baseline.

Results. Under statewide evaluation, banks substantially consolidate origination desks and save syndication costs. The number of active county-desks per 4-year evaluation window falls from 27.2 to 19.5 (-28%), and annualized realized desk costs decline from \$2.36 million to \$1.17 million (-50%). Syndication costs fall from \$1.16 million to \$0.60 million per year (-48%). Capitalizing these savings raises the tax credit price by 0.21 cents per dollar of tax credit.

Figure 9: Re-optimizing Origination Choice



Notes: Texas counties are grouped by quartiles of large-bank deposits in 2015 (horizontal axis) and total allocated tax credits over 2012–2024 (vertical axis). Counties with no large-bank deposits are included in Q1; gray cells contain no counties. Quartiles are defined over counties. Panel (a) reports the credit-weighted change in tax credit prices under the statewide counterfactual, in cents per credit dollar, for 1,654 applications in 167 counties over 2012–2024. Panel (b) reports the change in the share of county LIHTC equity directly originated by direct-capable banks, in percentage points.

Figure 9(a) reports the heterogeneous price response. Relative to today’s observed prices, tax credit prices rise by 4.67 cents in the lowest-deposit quartile, by 3.78 and 3.04 cents in the second and third quartiles, and fall by 0.43 cents in the highest-deposit quartile. Figure 9(b) shows the corresponding channel choices. Across counties, the saved desk and syndication costs are concentrated in high-deposit markets. Limited tax credit supply in bottom-quartile counties cannot justify a desk, while larger markets support greater consolidation into direct origination, with the largest increases in the top quartiles.

I then take the counterfactual prices into the developer model and adjust credit sizing to hold the government’s aggregate credit budget fixed. The counterfactual increases equity reaching projects by \$20.1 million over 2012–2024 and raises expected awarded units by 939, or 1.30 percent. The gains are highly heterogeneous across space. Expected units increase by 127 units (5.3 percent) in counties in the lowest deposit quartile, compared with 336 units (0.6 percent) in the highest quartile.

Discussion. Re-optimized origination generates 99 of the 939 additional expected units, and the remaining 840 come from price equalization alone. This estimated channel gain is conservative for three reasons. First, banks never observed originating directly are constrained to continue syndicating in the counterfactual, as their fixed origination costs are identified only from below. The

reported savings therefore provide a lower bound on the gains from re-optimizing channel choice, and they accrue to only 7 direct-capable banks in Texas. Second, I treat the syndication fee as fixed rather than modeling it endogenously. If a higher share of direct origination lowers syndication fees by reducing syndication demand, the model understates the resulting cost savings. Finally, equity commitment letters identify the source of capital but not the full intermediation structure. Banks classified as direct may still use single-investor funds, while fragmented markets may rely on more costly multi-investor funds. Because the composition of bank capital behind syndicated funds is not observed, the model may further understate the cost of fragmented CRA incentives.

8 Conclusion

This paper studies how bank regulation shapes the financing and spatial distribution of affordable housing. The interaction between the CRA and LIHTC program channels substantial private capital into affordable housing, but ties banks' demand for tax credits to their local deposit footprints. These differences in local demand translate into higher tax credit prices in deposit-rich counties. In counties with fewer deposits, projects instead rely more heavily on syndicators to pool capital across banks, adding another financing cost for developers.

I provide evidence for both channels. Using bank mergers as shocks to local CRA exposure, I find that banks increase LIHTC investment when they inherit local deposits. Stronger local CRA demand is associated with higher tax credit prices, and markets with higher prices observe larger projects and more entry. At the same time, projects in thin-deposit markets rely more heavily on syndication to pool fragmented bank demand, further reducing the equity that reaches developers. These differences in local CRA demand affect both the price developers receive for tax credits and their reliance on syndication.

I quantify these effects using an equilibrium model of developer project-size choice and bank tax-credit demand. Counterfactuals show that the spatial distribution of CRA incentives matters for the efficiency of LIHTC financing. Holding both the government credit budget and investor equity fixed, equalizing tax credit prices across counties increases expected awarded units by 840, or 1.16 percent, by reallocating equity toward markets where housing supply is more responsive to additional financing. Broadening CRA assessment areas statewide also allows banks to consolidate direct-origination desks and reduce syndication costs. These origination-cost savings increase expected awarded units by an additional 99 units, bringing the total increase to 939 units, or 1.30 percent.

The results also speak to the geographic reforms in the 2023 CRA modernization, which proposed to broaden consideration of community development financing beyond banks' traditional assessment areas. The counterfactuals show that relaxing these geographic constraints can real-

locate tax-credit equity toward markets with weaker local bank demand and reduce the costs of pooling fragmented investment, increasing affordable housing production without increasing tax-credit budget or investor outlays. More generally, when public subsidies are financed by private investors whose regulatory incentives vary across space, the geography of those incentives can affect both where capital flows and how much of it reaches developers, and ultimately housing supply.

References

- Agarwal, S., Benmelech, E., Bergman, N., and Seru, A. (2012). Did the community reinvestment act (cra) lead to risky lending? Technical report, National Bureau of Economic Research.
- Aguirregabiria, V., Clark, R., and Wang, H. (2017). The geographic flow of bank funding and access to credit: Branch networks and local-market competition.
- Aguirregabiria, V., Clark, R., and Wang, H. (2024). The geographic flow of bank funding and access to credit: Branch networks, local synergies and competition. *arXiv preprint arXiv:2407.03517*.
- Ahlfeldt, G. M. and Barr, J. (2022). The economics of skyscrapers: A synthesis. *Journal of Urban Economics*, 129:103419.
- Ang, A., Bhansali, V., and Xing, Y. (2010). Taxes on tax-exempt bonds. *Journal of Finance*, 65(2):565–601.
- Avery, R. B. and Brevoort, K. P. (2015). The subprime crisis: Is government housing policy to blame? *Review of Economics and Statistics*, 97(2):352–363.
- Baker, M., Bergstresser, D., Serafeim, G., and Wurgler, J. (2022). The pricing and ownership of US green bonds. *Annual Review of Financial Economics*, 14:415–437.
- Baum-Snow, N. and Duranton, G. (2025). Housing supply and housing affordability. Technical report, National Bureau of Economic Research.
- Baum-Snow, N. and Han, L. (2024). The microgeography of housing supply. *Journal of Political Economy*, 132(6):1897–1946.
- Baum-Snow, N. and Marion, J. (2009). The effects of low income housing tax credit developments on neighborhoods. *Journal of Public Economics*, 93(5-6):654–666.
- Berger, A. N. and Dick, A. A. (2007). Entry into banking markets and the early-mover advantage. *Journal of Money, Credit and Banking*, 39(4):775–807.
- Blickle, K., Parlatore, C., and Saunders, A. (2026). Specialization in banking. *The Journal of Finance*, 81(3):1531–1572.
- Bostic, R. W. and Robinson, B. L. (2003). Do cra agreements influence lending patterns? *Real Estate Economics*, 31(1):23–51.

- Brennan, M., Mehta, A., and Steil, J. (2022). In harm's way? the effect of disasters on the magnitude and location of low-income housing tax credit allocations. *Journal of Policy Analysis and Management*, 41(2):486–514.
- Burge, G. S. (2011). Do tenants capture the benefits from the low-income housing tax credit program? *Real Estate Economics*, 39(1):71–96.
- Cengiz, D., Dube, A., Lindner, A., and Zipperer, B. (2019). The effect of minimum wages on low-wage jobs. *Quarterly Journal of Economics*, 134(3):1405–1454.
- Cespedes, J., Jiang, E. X., Parra, C., and Zhang, J. (2024). Branching out inequality: The impact of credit equality policies. *Available at SSRN*, 4702125.
- Chamberlain, G. (1980). Analysis of covariance with qualitative data. *The Review of Economic Studies*, 47(1):225–238.
- Chau, K.-W., Wong, S. K., Yau, Y., and Yeung, A. K. C. (2007). Determining optimal building height. *Urban Studies*, 44(3):591–607.
- Chetty, R., Hendren, N., and Katz, L. F. (2016). The effects of exposure to better neighborhoods on children: New evidence from the moving to opportunity experiment. *American Economic Review*, 106(4):855–902.
- Chu, Y., Lin, L. X., and Xiao, Z. (2024). Agree to disagree: Lender equity holdings, within-syndicate conflicts, and covenant design. *Journal of Financial Intermediation*, 57:101065.
- Clapp, J. M. (1981). The impact of inclusionary zoning on the location and type of construction activity. *Real Estate Economics*, 9(4):436–456.
- Collinson, R., Ellen, I. G., and Ludwig, J. (2015). Low-income housing policy. Working Paper 21071, National Bureau of Economic Research.
- Combes, P.-P., Duranton, G., and Gobillon, L. (2021). The production function for housing: Evidence from france. *Journal of Political Economy*, 129(10):2766–2816.
- Conway, J., Glaser, J. N., and Plosser, M. C. (2023). Does the community reinvestment act improve consumers' access to credit? *Available at SSRN* 4321265.
- Cook, C., Li, P. Z., and Binder, A. J. (2023). *Where to build affordable housing?: Evaluating the tradeoffs of location*. US Census Bureau, Center for Economic Studies Rochester, NY.
- Cummings, J. L. and DiPasquale, D. (1999). The low-income housing tax credit: An analysis of the first ten years. *Housing Policy Debate*, 10(2):251–307.

- Currie, J. and Van Parys, J. (2025). Affordable housing during childhood improves long-term outcomes of women and their children. Working Paper 34464, National Bureau of Economic Research.
- Currie, J. and Yelowitz, A. (2000). Are public housing projects good for kids? *Journal of Public Economics*, 75(1):99–124.
- D’Amico, L., Glaeser, E. L., Gyourko, J., Kerr, W. R., and Ponzetto, G. A. M. (2024). Why has construction productivity stagnated? the role of land-use regulation. Working Paper 33188, National Bureau of Economic Research.
- Darmouni, O., Lehner, F., and Sastry, P. (2026). Financing the energy transition. *SSRN Working Paper*. SSRN 6871419.
- Desai, M., Dharmapala, D., and Singhal, M. (2010a). Investable tax credits: The case of the low-income housing tax credit. In *Tax Policy and the Economy*, volume 24, pages 181–205. University of Chicago Press.
- Desai, M., Dharmapala, D., and Singhal, M. (2010b). Tax incentives for affordable housing: the low income housing tax credit. *Tax policy and the economy*, 24(1):181–205.
- Diamond, R. and McQuade, T. (2019). Who wants affordable housing in their backyard? an equilibrium analysis of low-income property development. *Journal of Political Economy*, 127(3):1063–1117.
- Dick, A. A. (2007). Market size, service quality, and competition in banking. *Journal of Money, Credit and Banking*, 39(1):49–81.
- Dick, A. A. (2008). Demand estimation and consumer welfare in the banking industry. *Journal of Banking & Finance*, 32(8):1661–1676.
- Ding, L. and Nakamura, L. (2021). “don’t know what you got till it’s gone”: The community reinvestment act in a changing financial landscape. *Journal of Real Estate Research*, 43(1).
- Ellen, I. G., Horn, K., Kuai, Y., Pazuniak, R., and Williams, M. D. (2015). Effect of qap incentives on the location of lihtc properties.
- Ellen, I. G. and Horn, K. M. (2018). Points for place: Can state governments shape siting patterns of low-income housing tax credit developments? *Housing Policy Debate*, 28(5):727–745.

- Ellen, I. G., O'Regan, K., and Voicu, I. (2009). Siting, spillovers, and segregation: A reexamination of the low income housing tax credit program. *Housing markets and the economy: Risk, regulation, and policy*, pages 233–67.
- Eriksen, M. D. (2009). The market price of low-income housing tax credits. *Journal of Urban Economics*, 66(2):141–149.
- Eriksen, M. D. and Orlando, A. W. (2022). Returns to scale in residential construction: The marginal impact of building height. *Real Estate Economics*, 50(2):534–564.
- Eriksen, M. D. and Rosenthal, S. S. (2010). Crowd out effects of place-based subsidized rental housing: New evidence from the lihtc program. *Journal of Public Economics*, 94(11-12).
- Favilukis, J., Mabilie, P., and Van Nieuwerburgh, S. (2023). Affordable housing and city welfare. *The Review of Economic Studies*, 90(1):293–330.
- Fecht, F., Peydro, J.-L., and Tümer-Alkan, G. (2021). Banks' equity stakes in firms: A blessing or curse in credit markets? Technical report, Barcelona School of Economics.
- Gabriel, S. A. and Kung, E. (2026). Development approval timelines, approval uncertainty, and new housing supply: Evidence from los angeles. *Journal of Urban Economics*. Forthcoming.
- Gardner, J. (2021). Two-stage differences in differences. Working paper, University of Mississippi.
- Garmaise, M. J. and Moskowitz, T. J. (2006). Bank mergers and crime: The real and social effects of credit market competition. *the Journal of Finance*, 61(2):495–538.
- Garrett, D. and Shive, S. (2025). The power of tax equity investors. *SSRN Working Paper*. SSRN 4344966.
- Ghent, A. C., Hernández-Murillo, R., and Owyang, M. T. (2015). Did affordable housing legislation contribute to the subprime securities boom? *Real Estate Economics*, 43(4):820–854.
- Glaeser, E. and Gyourko, J. (2025). America's housing supply problem: The closing of the suburban frontier? *Brookings Papers on Economic Activity*, 2025(1):375–425.
- Goodman-Bacon, A. (2021). Difference-in-differences with variation in treatment timing. *Journal of Econometrics*, 225(2):254–277.
- Green, R. C. (1993). A simple model of the taxable and tax-exempt yield curves. *Review of Financial Studies*, 6(2):233–264.

- Hart, O. (2003). Incomplete contracts and public ownership: Remarks, and an application to public-private partnerships. *Economic Journal*, 113(486):C69–C76.
- Hart, O., Shleifer, A., and Vishny, R. W. (1997). The proper scope of government: Theory and an application to prisons. *Quarterly Journal of Economics*, 112(4):1127–1161.
- Hurst, E., Keys, B. J., Seru, A., and Vavra, J. (2016). Regional redistribution through the us mortgage market. *American Economic Review*, 106(10):2982–3028.
- Iossa, E. and Martimort, D. (2015). The simple microeconomics of public-private partnerships. *Journal of Public Economic Theory*, 17(1):4–48.
- Jacob, B. A., Kapustin, M., and Ludwig, J. (2015). The impact of housing assistance on child outcomes: Evidence from a randomized housing lottery. *The Quarterly Journal of Economics*, 130(1):465–506.
- Jiang, E. X., Yu, G. Y., and Zhang, J. (2026). Bank competition amid digital disruption: Implications for financial inclusion. *The Journal of Finance*, 81(4):1951–2004.
- Jiang, W., Li, K., and Shao, P. (2010). When shareholders are creditors: Effects of the simultaneous holding of equity and debt by non-commercial banking institutions. *The Review of Financial Studies*, 23(10):3595–3637.
- Johnston, S. (2019). Nonrefundable tax credits versus grants: The impact of subsidy form on the effectiveness of subsidies for renewable energy. *Journal of the Association of Environmental and Resource Economists*, 6(3):433–460.
- Kling, J. R., Liebman, J. B., and Katz, L. F. (2007). Experimental analysis of neighborhood effects. *Econometrica*, 75(1):83–119.
- Kneebone, E. and Reid, C. K. (2021). The complexity of financing low-income housing tax credit housing in the united states. *Terner Center for Housing Innovation, UC Berkeley*. <https://ternercenter.berkeley.edu/research-andpolicy/lihtc-complexity>.
- Krimmel, J. and Wang, B. X. (2026). Upzoning with strings attached: Evidence from seattle’s affordable housing mandate. *Regional Science and Urban Economics*, page 104219.
- Kroszner, R. S. and Strahan, P. E. (2001). Bankers on boards:: monitoring, conflicts of interest, and lender liability. *Journal of financial Economics*, 62(3):415–452.
- Kundu, S., Park, S., and Vats, N. (2021). The geography of bank deposits and the origins of aggregate fluctuations. *Available at SSRN 3883605*.

- Lang, B. J. and Olsen, E. O. (2025). The hidden subsidies of the low-income housing tax credit program. *National Tax Journal*, 78(3):577–608.
- Longstaff, F. A. (2011). Municipal debt and marginal tax rates: Is there a tax premium in asset prices? *Journal of Finance*, 66(3):721–751.
- Marquis, C., Guthrie, D., and Almandoz, J. (2012). State activism and the hidden incentives behind bank acquisitions. *Social science research*, 41(1):130–145.
- Miller, M. H. (1977). Debt and taxes. *Journal of Finance*, 32(2):261–275.
- Nguyen, H.-L. Q. (2019). Are credit markets still local? evidence from bank branch closings. *American Economic Journal: Applied Economics*, 11(1):1–32.
- Oberfield, E., Rossi-Hansberg, E., Trachter, N., and Wenning, D. T. (2024). Banks in space: Nber working papers 32256.
- Olsen, E. O. (2003). Housing programs for low-income households. In *Means-tested transfer programs in the United States*, pages 365–442. University of Chicago Press.
- Paciorek, A. (2013). Supply constraints and housing market dynamics. *Journal of Urban Economics*, 77:11–26.
- Pástor, L., Stambaugh, R. F., and Taylor, L. A. (2021). Sustainable investing in equilibrium. *Journal of Financial Economics*, 142(2):550–571.
- Pedersen, L. H., Fitzgibbons, S., and Pomorski, L. (2021). Responsible investing: The esg-efficient frontier. *Journal of Financial Economics*, 142(2):572–597.
- Ringo, D. (2023). Mortgage lending, default, and the community reinvestment act. *Journal of Money, Credit and Banking*, 55(1):77–102.
- Rollet, V. (2025). Can we rebuild a city? the dynamics of urban redevelopment. Working Paper.
- Saadi, V. (2020). Role of the community reinvestment act in mortgage supply and the us housing boom. *The Review of Financial Studies*, 33(11):5288–5332.
- Saiz, A. (2010). The geographic determinants of housing supply. *Quarterly Journal of Economics*, 125(3):1253–1296.
- Sakong, J. and Zentefis, A. K. (2023). Bank branch access: Evidence from geolocation data.

- Sinai, T. and Waldfoegel, J. (2005). Do low-income housing subsidies increase the occupied housing stock? *Journal of public Economics*, 89(11-12):2137–2164.
- Soltas, E. (2024a). Tax incentives and the supply of low-income housing. *Unpublished manuscript*.
- Soltas, E. and Gruber, J. (2026). How costly is permitting in housing development? Working Paper.
- Soltas, E. J. (2024b). The price of inclusion: Evidence from housing developer behavior. *Review of Economics and Statistics*, 106(6):1588–1606.
- Sun, L. and Abraham, S. (2021). Estimating dynamic treatment effects in event studies with heterogeneous treatment effects. *Journal of Econometrics*, 225(2):175–199.
- U.S. Department of Housing and Urban Development (2017). Defining housing affordability. PD&R Edge. Office of Policy Development and Research.
- U.S. Government Accountability Office (2012). Community reinvestment act: Challenges in quantifying its effect on low-income housing tax credit investment. Technical Report GAO-12-869R, U.S. Government Accountability Office, Washington, DC.
- Zerbib, O. D. (2019). The effect of pro-environmental preferences on bond prices: Evidence from green bonds. *Journal of Banking and Finance*, 98:39–60.

Appendix

A	Additional Tables and Figures	1
A.1	LIHTC investments, four DID estimators	1
A.2	Bank Deposits, Supervisory Agency and Rate of Return	2
A.3	LIHTC Prices, Risk-Free Benchmark, and Market Condition	3
A.4	Decompose Variation in Tax Credit Price	4
A.5	Fixed Syndication Costs	5
A.6	Income Targeting and Low-Income Unit Composition	6
A.7	Units Bunching and Developer Overhead Allowance	7
B	Additional Institutional Details	1
B.1	Data Coverage by State and Year	1
B.2	Example of CRA Investment Test	2
B.3	OCC National Bank Community Development Investment Annual Charts	3
B.4	Bank Heterogeneity in Direct vs. Syndicated Investment	4
B.5	Mixed Channel Choice	5
B.6	Foreclosure Risk of LIHTC Property	6
B.7	Unit Caps in 15-state QAPs	7
B.8	Heterogeneity in Project Funding Sources by State	8
B.9	Credit Allocation, Housing Needs, and CRA Incentives	9
C	The Conditional-Logit Award Process	10
C.1	Conditional Likelihood	10
C.2	Inclusion Probabilities and Adding Up	11
C.3	Winning Probability in the Developer's Problem	12
C.4	Estimation and Inference	13
C.5	Discussion	13
D	Additional Details in Model Estimation	15
D.1	Supply Estimation	15
D.2	Supply Model Extension: Unit Size as Alternative Measure	22
D.3	Demand Estimation	25
E	Data Appendix: Bank Public Welfare Investments	30
E.1	OCC Public Welfare Investments	30
E.2	Local Deposits and the Bank–County Panel	32
E.3	Bank Mergers and Acquisitions	34

A Additional Tables and Figures

A.1 LIHTC investments, four DID estimators

All four estimators use the same post-merger horizon but differ in how treatment effects are aggregated across merger cohorts. TWFE uses the full panel; Sun–Abraham estimates cohort-specific effects and averages them using treated-pair shares; Gardner estimates fixed effects from untreated observations and gives equal weight to treated post-merger observations; and the stacked estimator constructs separate event-window samples for each merger cohort using never-treated pairs as controls. Sun–Abraham and stacked exclude cohorts with fewer than five treated pairs. The IV estimate is constructed as the ratio of the estimated post-merger effect on LIHTC investment to the corresponding effect on deposits, and the reported first-stage F is the squared t statistic on the deposit effect.

Table A1: Bank Deposits and LIHTC Investment, by Staggered Estimator

	(1) TWFE	(2) Sun–Abraham [†]	(3) Gardner	(4) Stacked
Panel A. Number of LIHTC Projects				
Instrumented deposits (\$bn)	0.0831*** (0.0299)	0.0648** (0.0272)	0.0816*** (0.0241)	0.0657** (0.0263)
Panel B. LIHTC investment (M\$)				
Instrumented deposits (\$bn)	1.367*** (0.378)	1.033** (0.473)	1.377*** (0.357)	1.176** (0.495)
First-stage F	19.7	10.7	125.5	25.4
Bank $\times$ county effects	✓	✓	✓	✓
County $\times$ year effects	✓	✓	✓	✓
Observations	23,366	23,277	22,916	116,879
Clustered on	Bank	Bank	Bank	Bank

Notes: This table examines the robustness of the IV estimates in Panel B of Table 4 to alternative estimators that account for staggered merger timing. Panel A reports the number of LIHTC projects and Panel B their dollar value in millions of 2024 dollars. Deposits are instrumented with the post-merger indicator interacted with inherited deposits, winsorized at the 99th percentile. [†]Column (2) instead instruments deposits with the post-merger indicator alone: the Sun–Abraham estimator saturates the design in cohort $\times$ event-time cells, and with a continuous shock most cells are identified by only a few bank–county pairs, so the scaled version is not reliably estimated. The coefficient remains the effect per billion dollars of deposits held and is comparable across columns. Sample sizes vary across estimators, with column (4) reporting stacked observations. Standard errors are clustered at the bank level and reported in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

A.2 Bank Deposits, Supervisory Agency and Rate of Return


Table A2: Bank deposits and implied internal rate of return

	Tax Credit Price	Implied Internal Rate of Return (%)				
	(1)	(2)	(3)	(4)	(5)	(6)
log(Cnty Large-Bank Dep)	0.6140* (0.3204)	-0.1714** (0.0707)	-0.1723*** (0.0641)			
× FED Share of Deposits	-0.0873 (0.2949)					
× FDIC Share of Deposits	-0.1730 (0.3229)					
FED Share of Deposits	2.7089 (6.0417)					
FDIC Share of Deposits	4.8745 (6.4382)					
log(Investor Bank Deposits)				-0.0263** (0.0107)	-0.0295*** (0.0082)	-0.0559*** (0.0199)
Pay-in Schedule	✓	✓	✓	✓	✓	✓
No-Debt & Levered DSCR	✓	✓	✓	✓	✓	✓
Same Lender	✓	✓	✓	✓	✓	✓
County Controls	✓	✓	✓	✓		
Credit Supply Controls	✓	✓	✓	✓		
Syndicator FE	✓	✓	✓	✓	✓	✓
County FE, Year FE	✓	✓	✓	✓		
County × Year FE					✓	✓
Bank × County FE						✓
Observations	11,481	11,481	9,822	1,456	1,456	1,456
Counties	635	635	617	101	101	101
Dep. Var Mean	89.771	2.650	2.642	1.350	1.350	1.350
Adjusted R-squared	0.788	0.806	0.815	0.805	0.860	0.846

Notes: This table provides additional robustness checks. Column (1) tests whether the relationship between large-bank deposits and tax-credit prices varies with the local supervisory composition of deposits, measured by the county deposit shares held by Federal Reserve- and FDIC-supervised banks. Columns (2)–(6) reproduce the specifications in Table 5 using the implied return as the dependent variable. Standard errors are clustered by county. *** p<0.01, ** p<0.05, * p<0.10.

A.3 LIHTC Prices, Risk-Free Benchmark, and Market Condition

Figure A1: LIHTC Tax Credit Prices and the Risk-Free Benchmark



Notes: The figure compares average tax credit prices by application year with the price implied by the 10-year Treasury yield. The solid line reports the mean LIHTC price in the 15-state sample, adjusted for changes in state composition, with bars denoting 95% confidence intervals. The adjustment residualizes prices on state and application-year fixed effects and removes each deal's estimated state effect, centered at its deal-weighted mean, thereby holding the pooled mean fixed while removing changes due to states entering or leaving the sample. State effects are estimated using all priced applications from 2003–2025. The dashed line reports the risk-free price, $p^{rft} = \frac{1}{10} \sum_{k=0}^9 (1 + r_t)^{-k}$, where r_t is the annual average 10-year Treasury constant-maturity yield (FRED DGS10).


A.4 Decompose Variation in Tax Credit Price

Table A3: Decompose Variation in Tax Credit Price, Texas

Variable	County-Year FE	Full Model
Percent of Early Commitment		-0.0011 (0.0255)
Construction Term		0.0059 (0.0063)
Percent of Early Commitment $\times$ Construction Term		0.0005 (0.0128)
Developer Experience		0.0000 (0.0001)
County-Year FE	Yes	Yes
R-squared	0.893	0.894
Adjusted R-squared	0.819	0.819
Observations	1,962	1,962

Notes: Table A3 reports OLS regressions of project-level tax credit prices on county-year fixed effects and project-level controls. Column (1) includes only county-year fixed effects. Column (2) adds pay-in schedule and developer experience. The pay-in schedule variables include the share of equity committed early, the construction term, and their interaction. Developer experience is measured as the number of LIHTC projects the developer applied for before the application year. Developers are identified using the developer name and the associated email domain. Heteroskedasticity-robust standard errors are reported in parentheses.

Figure A2: Price Variation in Tax Credit Price, Texas



Notes: Figure A2 plots the R^2 from OLS regressions of project-level tax credit price on alternative covariate sets (See Table A3). Bars are sorted by R^2 of each model. All regressions are estimated on the same sample of projects with the non-missing dependent variable and covariates.

A.5 Fixed Syndication Costs

This subsection tests the assumption of a fixed syndication fee across markets in the demand model. Table A4 examines this restriction by testing whether the syndicated–direct price gap varies with local bank fragmentation in thin-deposit counties, where syndication is most prevalent. The price gap remains approximately 1.1 cents across specifications, with no significant interaction with alternative measures of fragmentation. This provides little evidence of systematic cross-market variation in syndication costs.

Table A4: Fragmentation and the Syndication Price Gap in Thin-Deposit Counties

	Tax credit price (cents)		
	(1)	(2)	(3)
County syndicated share	-1.114*** (0.317)	-1.116*** (0.330)	-1.106*** (0.332)
ln(Effective number of large banks)		0.327 (0.323)	0.339 (0.322)
Syndicated share $\times$ ln(Effective number)		0.139 (0.614)	-0.535 (0.757)
Syndicated share $\times$ log(Large-Bank Deposits)			0.600 (0.532)
log(Large-Bank County Deposits)	0.174 (0.140)	0.093 (0.157)	0.076 (0.157)
County controls	✓	✓	✓
State $\times$ Year FE	✓	✓	✓
CSA FE	✓	✓	✓
Observations	1,392	1,392	1,392
Dep. var. mean	88.88	88.88	88.88
Syndicated share, mean	0.868	0.868	0.868
Effective large banks, median	2.2	2.2	2.2

Notes: The sample consists of county $\times$ year observations in the bottom two quintiles of local large-bank deposits in the 15-state application sample, 2004–2023. The dependent variable, controls, and fixed effects are as in Table 6. The effective number of large banks is the inverse Herfindahl index of county large-bank deposit shares. Interacted variables are centered at their sample means, so the coefficient on syndicated share measures the price gap between fully syndicated and fully direct counties at mean fragmentation. Standard errors are clustered by county. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

A.6 Income Targeting and Low-Income Unit Composition

Figure A3: Fraction of LI Units





Figure A4: Weighted Income Limits



Notes: Figure A3 plots the distribution of the share of low-income (LI) units across LIHTC applications in the sample. The figure shows substantial bunching near 100 percent, indicating that many projects allocate nearly all units as low-income units. In some cases, projects reserve a small number of units (e.g., one to four units) for managers or employees, causing the LI share to fall slightly below 100 percent. Figure A4 compares the distribution of the weighted average of developer-specified income limits across awarded and non-awarded LIHTC applications. The weighted income limit is constructed using unit shares across different AMI restrictions within a project. Manager units are excluded from the calculation, and market-rate units are assigned a 100 percent income limit. Lower values indicate deeper affordability targeting. The figure suggests substantial overlap between awarded and non-awarded projects.

A.7 Units Bunching and Developer Overhead Allowance

Figure A5: Distribution of Total Units





Figure A6: Developer Fee



Notes: Figure A5 plots the distribution of project size for Texas LIHTC applications. The bunching at common sizes such as 48, 60, 72, and 96 units suggests that project size is partly shaped by standardized development templates and architectural design increments. The mass below 50 units reflects the Texas developer-fee threshold, which allows a higher developer-fee cap for developments with 49 or fewer units. Figure A6 plots the distribution of developer fees as a share of eligible basis for awarded and non-awarded Texas LIHTC applications. The measure is constructed as Developer Fee divided by eligible basis net of Developer Fee. Both awarded and non-awarded applications exhibit substantial bunching near the statutory caps of 15 and 20 percent. This pattern suggests that developers frequently set subsidizable overhead at the maximum amount allowed by underwriting rules, rather than reducing developer fees to potentially increase the probability of receiving an allocation.

B Additional Institutional Details

B.1 Data Coverage by State and Year

Figure B1: LIHTC Application Data Coverage by State and Year

State	2003	2004	2005	2006	2007	2008	2009	2010	2011	2012	2013	2014	2015	2016	2017	2018	2019	2020	2021	2022	2023	2024	2025	Applications
CA																								2,814
FL																								2,868
GA																								1,404
IA																								360
ID																								53
KS																								123
MN																								346
NC																								912
NE																								137
NH																								273
NM																								38
OH																								795
OK																								903
TX																								2,059
WI																								79
All	63	73	93	135	148	274	373	558	469	549	725	757	892	960	933	988	994	1,020	967	703	825	446	219	13,164

Notes: This figure shows the availability of LIHTC application data by state and year. Each blue cell indicates that application-level data were collected and digitized for that state-year, while blank cells indicate that application records were not available. Data availability varies across states depending on the capacity of state housing agency to maintain records. The rest of states are missing because (1) 20 states did not respond to or proceed with my requests, (2) 6 states cited that fulfilling the request would require substantial time or fees, and (3) 9 states rejected the request because the author is not a state resident or the data is confidential.

B.2 Example of CRA Investment Test

Table B1: CRA Investment Test for Wells Fargo in Arizona (2023), rated “High Satisfactory”

Assessment Area	Total Community Development Investment				Deposit
	Volume	% of Total	Amount \$(000's)	% of Total	% of Deposits \$
Full-Scope Review:					
Phoenix CSA	458	71.3	217,518	74.2	76.0
Tucson CSA	112	17.4	72,586	24.8	12.9
Limited-Scope Review:					
Flagstaff MSA	32	5.0	969	0.3	2.0
Lake Havasu City MSA	4	0.6	92	0.0	2.6
Prescott Valley MSA	19	3.0	715	0.2	2.6
Sierra Vista MSA	3	0.5	459	0.2	1.0
Yuma MSA	5	0.8	53	0.0	1.8
Arizona Combined Non-MSA	9	1.4	581	0.2	1.0

Notes: This table shows an example from Wells Fargo CRA Performance Evaluation in state of Arizona, as reported by regulators in 2023. The table compares the geographic distribution of a bank’s qualified community development investments with the distribution of its deposits across assessment areas. “Volume” denotes the number of qualified community development investments, while “Dollar Amount” reports the total investment value in thousands of dollars. The “% of Total” columns report each assessment area’s share of the bank’s total community development investment activity. The final column reports the share of the bank’s total deposits located in each rating area. “Full Review” areas receive a larger weight on statewide rating than “Limited Review”.

Table B2: CRA Investment Test for Wells Fargo in Arizona (2019), rated “Needs to Improve”

Assessment Area	Total Community Development Investment				Deposit
	Volume	% of Total	Amount \$(000's)	% of Total	% of Deposits \$
Full-Scope Review:					
Phoenix MSA	724	71.4	91,837	47.3	72.5
Tucson CSA	185	18.2	95,971	49.4	15.0
Limit-Scope Review:					
Flagstaff MSA	19	1.9	188	0.1	2.2
Lake Havasu City MSA	11	1.1	41	0.0	2.8
Prescott MSA	22	2.2	249	0.1	3.0
Sierra Vista MSA	2	0.2	1,057	0.5	1.1
Yuma MSA	23	2.3	130	0.1	1.8
AZ Combined NonMetro	28	2.8	4,807	2.5	1.6

Notes: This table provides an example of Wells Fargo’s poor CRA Investment Test performance in Arizona in 2019. The evaluation states that “the institution has a poor level of qualified CD investments and grants, particularly those that are not routinely provided by private investors.”

B.3 OCC National Bank Community Development Investment Annual Charts

Table B3: Example of CD Investments in WI (2023)


Investment Bank	PWI Apartment and Fund	Target Area	Amount
Associated Bank	Bay City Lofts	Green Bay, WI	\$4,022,358
Associated Bank	Shield Apartments	Madison, WI	\$7,199,280
Bremer Bank	200 Rittenhouse Avenue	Bayfield, WI	\$919,009
Old National Bank	City East Apartments	Green Bay, WI	\$6,273,031
U.S. Bank	New Hampton Apartments	Milwaukee, WI	\$21,266,802
U.S. Bank	Belle City Square Prime Tenant	Racine, WI	\$4,806,054
U.S. Bank	Kenosha Uptown Lofts	Kenosha, WI	\$13,334,746
Wells Fargo Bank	Tosa MC Apartments	Wauwatosa, WI	\$12,438,366
Bremer Bank	Bremer Community Invest Fund II	Chippewa Falls	\$6,949,800
Bremer Bank	Bremer Community Invest Fund II	Milwaukee	\$5,370,300
JPMorgan Chase Bank	RJ Housing Fund 73 L.P.	Milwaukee, WI	\$12,830,886
Town Bank	R4 Capital	Beloit, WI	\$1,000,000
U.S. Bank	USBCDC Investment Fund 398	Antigo, WI	\$4,680,000
U.S. Bank	Twain Investment Fund 4	Racine, WI	\$4,806,054
U.S. Bank	GenCap Emem Sheboygan 101	Sheboygan, WI	\$19,313,161
U.S. Bank	Twain Investment Fund 720	Milwaukee, WI	\$9,913,800
U.S. Bank	Twain Investment Fund 247	West Allis, WI	\$4,106,700
U.S. Bank	Twain Investment Fund 724	Greenwood, WI	\$10,374,000

Notes: This table presents examples of community development (CD) investments made by banks in Wisconsin in 2023. The table lists the investing bank, the investment vehicle or project, the geographic target area, and the investment amount. Investments include both direct financing of affordable housing developments (e.g., apartment projects) and participation in specialized housing investment funds. Large national banks such as U.S. Bank, JPMorgan Chase, and Wells Fargo appear alongside regional institutions such as Associated Bank and Bremer Bank. In the data construction process, I use apartment name, target area and investment year as merge key to match investors.

B.4 Bank Heterogeneity in Direct vs. Syndicated Investment

Projects are classified as direct when the investment name indicates a proprietary bank vehicle, wholly owned subsidiary, or single-property partnership without a recognized syndicator, and as syndicated when it identifies a known third-party LIHTC syndicator (e.g., Alliant, Boston Capital, Enterprise/NEF, Hudson, Hunt, RBC, Red Stone, Raymond James, WNC, Richman, OCCH, and similar). Because the filings do not directly report the underwriting channel, these classifications are imperfect name-based proxies used only to illustrate cross-bank differences; the demand model instead identifies channels from equity commitment letters.

Figure B2: Bank Heterogeneity in Direct vs. Syndicated Investment




Notes: Each row counts unique LIHTC/housing investment reports filed by the bank in the FFIEC Community Development Investment Annual Charts, 2003–2023, with the top 30 banks by project count shown and bank-name variants collapsed into a single entity.

B.5 Mixed Channel Choice

Using the Investment Test review period stated in each bank's OCC performance evaluation, 53.9 percent of the 2,703 bank–county–evaluation period cells are entirely direct, 40.7 percent are entirely syndicated and 5.4 percent combine the two channels. Of the 4,684 bank–county–year cells, 59.0 percent are entirely direct, 39.0 percent are entirely syndicated and 2.0 percent combine the two channels.

Figure B3: Channel Choice in Bank-County-Year




Notes: Each row counts bank–county–year observations. An observation is Direct only if every reported investment in county-year is a direct investment, Syndicated only if every report is syndicated, and Mixed Channel if it contains both.

B.6 Foreclosure Risk of LIHTC Property

Figure B4: Low Foreclosure Rate in LIHTC Market

Annual LIHTC Foreclosure Rate vs. Conventional Multifamily Delinquency Rate



Notes: Figure B4 compares the annual foreclosure rate for LIHTC properties with the delinquency rate for conventional multifamily properties from 2002 to 2022. The dark blue area reports the LIHTC foreclosure ratio, while the light blue area reports the FDIC-insured multifamily delinquency ratio. While conventional multifamily loan delinquency rose sharply during the Great Recession—peaking at over 4 percent around 2010—the foreclosure rate among LIHTC properties remained close to zero throughout the period. The figure illustrates the historically low default risk associated with LIHTC properties. ([Source](#))

B.7 Unit Caps in 15-state QAPs


Table B4: QAP Unit Cap

State	Cap	Scope	Years	<i>N</i>	At cap
TX	80	rural, new construction or adaptive reuse	2010–	237	24.9%
NC	80	Central, East, West Regions, 9% new construction	2015–2018	245	23.3%
NC	84	Central, East, West Regions, 9% new construction	2019–	410	8.5%
CA	80	rural set-aside (low-income units)	2001–2024	193	4.1%
NC	120	Metro Region, 9% new construction	2015–	257	2.3%
TX	252	urban, new construction or adaptive reuse	2005–2013	396	0.0%
CA	150	other than rural set-aside (low-income units)	2008–2017	1,799	0.0%
TX	120	rural, tax-exempt bond	2020–	112	0.0%
NC	200	tax-exempt bond, new construction	2015–	912	0.0%
IA	144	Polk County (agency reserves the right)	2005–	100	0.0%
IA	96	other MSA central counties (reserved right)	2005–	111	0.0%

Notes: This table summarizes project-size caps imposed by state Qualified Allocation Plans (QAPs) in the 15-state application sample. *N* is the number of applications subject to each ceiling (the state, region or rural status, development type and years in scope), and “At cap” is the share of those applications whose unit count equals the ceiling exactly. Applications that bunch at these ceilings are excluded from the column (2) in Table 7. The table reports only explicit upper bounds on project size and does not include other QAP-induced size kinks, such as changes in developer-fee allowances at specified unit thresholds or additional amenity requirements for developments exceeding specified sizes.

B.8 Heterogeneity in Project Funding Sources by State

Figure B5: Distribution of Efficient Use of Subsidy by State



Notes: Figure B5 plots the distribution of the subsidy ratio across LIHTC applications by state, where the subsidy ratio is defined as the ratio of total project cost to equity raised from federal LIHTC credits. Higher values indicate that projects rely more heavily on non-LIHTC financing sources relative to tax credit equity. The figure highlights substantial heterogeneity across states in the efficient use of LIHTC subsidies, reflecting differences in project costs, state allocation rules, local government support, and the availability of rental-income-backed commercial loans. For example, states which allow LIHTC developers to apply federal and state LIHTCs to the same eligible basis, mechanically lowers reliance on federal LIHTCs. Because developers also endogenize the amount of state credits layered onto each project, this financing margin generates additional within-state dispersion in reliance on federal LIHTCs.

B.9 Credit Allocation, Housing Needs, and CRA Incentives

To provide distributional context, I compare the geographic allocation of tax credits in Texas with low-income housing needs. This comparison is not intended as a welfare analysis: the social value of affordable housing depends not only on the number of low-income households served, but also on housing stability and affordability, neighborhood spillovers, access to neighborhood opportunity, and long-run effects on children (Currie and Yelowitz, 2000; Kling et al., 2007; Jacob et al., 2015; Chetty et al., 2016; Diamond and McQuade, 2019; Cook et al., 2023; Currie and Van Parys, 2025). Rather, it illustrates how the spatial incentives created by the CRA interact with a tax credit allocation that is already broadly targeted toward low-income households.

Figure B6: HH Income < 30%





Figure B7: HH Income < 30%, $\frac{\text{Rents}}{\text{Income}} > 30\%$



Notes: Quartiles cover all 254 Texas counties and are based on 2015 large-bank deposits. Credits are annual 9% awards from TDHCA, averaged over 2010–2024. Renter counts come from the 2015–2019 ACS five-year county tables. Eligible renters live in households below 30% of county median income; the rent-burdened measure additionally requires gross rent of at least 30% of income. Household counts are converted to persons using renter household size. Because the ACS does not identify LIHTC tenants, these measures may include existing LIHTC residents.

Figures B6 and B7 compare annual 9% tax credits with the number of low-income renters across quartiles of 2015 large-bank deposits. Figure B6 counts people in renter households below 30% of county median income, while Figure B7 additionally restricts the sample to rent-burdened households, defined as those spending more than 30% of income on housing costs (U.S. Department of Housing and Urban Development, 2017). Credit allocations broadly track these measures of housing needs. Instead, CRA incentives introduce an additional spatial force on the financing cost. Conditional on credit allocations, banks have stronger incentives to invest where their deposit footprints are larger, raising investor demand and equity prices. The counterfactuals therefore focus on whether reducing geographic differences in financing conditions can increase affordable housing supply under a fixed tax credit budget, without taking a stand on the socially optimal geographic allocation of affordable housing.

C The Conditional-Logit Award Process

This appendix develops the award model in Section 5.1. Applications compete within allocation pools, and each pool funds a fixed number of projects. I condition on the number of awards in each pool and estimate how those awards are distributed across applications. I first derive the conditional likelihood and the associated inclusion probabilities. I then show how these probabilities enter the developer's problem and describe estimation and inference.

C.1 Conditional Likelihood

Applications are indexed by i and allocation pools by c . A pool is defined by application year and the relevant TDHCA competition: the statewide At-Risk set-aside or one of the urban and rural subregions of the thirteen Uniform State Service Regions. Let I_c denote the applications competing in pool c , let $n_c = |I_c|$, and let

$$K_c = \sum_{i \in I_c} A_i$$

denote the observed number of awards, where $A_i = 1$ if application i receives an allocation.

Within a pool, application i has award index

$$v_i = Z_i' \gamma + \lambda_1 \mu_i + \lambda_2 \phi_i, \quad w_i = e^{v_i}, \quad (C.1)$$

as in equation (11). The latent award probability is

$$\Pr(A_i = 1 \mid \alpha_c, v_i) = \Lambda(\alpha_c + v_i), \quad i \in I_c, \quad (C.2)$$

where α_c is an unrestricted pool effect. Conditional on α_c and the application indices, award shocks are independent across applications.

The key feature of the model is that estimation conditions on K_c . Because every feasible outcome under this conditioning contains exactly K_c winners, the common pool effect α_c drops from the likelihood (Chamberlain, 1980). For any set $W \subseteq I_c$ containing K_c applications,

$$\Pr \left(\{i : A_i = 1\} = W \mid \sum_{i \in I_c} A_i = K_c \right) = \frac{\prod_{i \in W} w_i}{\sum_{\substack{S \subseteq I_c \\ |S| = K_c}} \prod_{j \in S} w_j}. \quad (C.3)$$

To see this, let $p_i = \Lambda(\alpha_c + v_i)$. The unconditional probability associated with winner set W

can be written as

$$\prod_{i \in W} p_i \prod_{i \notin W} (1 - p_i) = \left[\prod_{i \in I_c} (1 - p_i) \right] e^{\alpha_c K_c} \prod_{i \in W} w_i.$$

Conditional on K_c , the first two terms are common to every feasible winner set and therefore cancel. Equation (C.3) follows.

Thus, identification comes entirely from comparisons among applications competing in the same pool. Any factor common to all applications in a pool—including the size of the pool’s allocation and other common pool–year shocks—is absorbed by the conditioning. Pools in which all applications win or all applications lose contain no within-pool outcome variation and do not contribute to the likelihood.

C.2 Inclusion Probabilities and Adding Up

The structural model requires the probability that a particular application receives one of the K_c awards, rather than the probability of a particular winner set. I refer to this object as the application’s inclusion probability:

$$\omega_i \equiv \Pr \left(i \in W \left| \sum_{j \in I_c} A_j = K_c \right. \right). \quad (\text{C.4})$$

For compactness, define

$$e_k(w) = \sum_{\substack{S \subseteq I_c \\ |S|=k}} \prod_{j \in S} w_j, \quad e_0(w) = 1,$$

and let w_{-i} denote the vector of indices excluding application i . Summing equation (C.3) over all winner sets containing i gives

$$\omega_i = \frac{w_i e_{K_c-1}(w_{-i})}{e_{K_c}(w)}. \quad (\text{C.5})$$

These inclusion probabilities satisfy

$$\sum_{i \in I_c} \omega_i = K_c. \quad (\text{C.6})$$

The adding-up restriction follows directly from conditioning on exactly K_c winners:

$$\sum_{i \in I_c} \omega_i = \mathbb{E} \left[\sum_{i \in I_c} \mathbf{1}\{i \in W\} \left| |W| = K_c \right. \right] = K_c.$$

Equation (C.6) distinguishes the inclusion probabilities from choice probabilities in a standard

multinomial logit. They do not sum to one unless the pool has a single award. Instead, the model distributes the pool's K_c awards across its n_c applicants. This property is useful in the structural model because predicted awards therefore respect the observed capacity of every competition pool exactly.

C.3 Winning Probability in the Developer's Problem

The developer's problem requires the winning probability as a function of its own project characteristics. Holding the other applications in the pool fixed, equation (C.5) has an exact logistic representation. Using

$$e_{K_c}(w) = w_i e_{K_c-1}(w_{-i}) + e_{K_c}(w_{-i}),$$

the inclusion probability can be written as

$$\begin{aligned} \omega_i &= \frac{w_i e_{K_c-1}(w_{-i})}{w_i e_{K_c-1}(w_{-i}) + e_{K_c}(w_{-i})} \\ &= \Lambda \left(v_i + \log \frac{e_{K_c-1}(w_{-i})}{e_{K_c}(w_{-i})} \right). \end{aligned} \tag{C.7}$$

The second term in equation (C.7) is the rival offset in equation (12). It summarizes two features of the developer's competitive environment: the number of awards available in the pool and the award indices of all competing applications. Conditional on these objects, the developer faces an ordinary logistic winning probability in its own index.

Because the offset does not vary with application i 's own choice when rival applications are held fixed,

$$\frac{\partial \omega_i}{\partial v_i} = \omega_i(1 - \omega_i). \tag{C.8}$$

Project size affects v_i through the subsidy ratio $\mu_i(q)$, so

$$\frac{\partial \omega_i(q)}{\partial q} = \omega_i(q) [1 - \omega_i(q)] \lambda_1 \frac{\partial \mu_i(q)}{\partial q}. \tag{C.9}$$

This is the derivative entering the developer's first-order condition in equation (14).

The offset representation applies to a unilateral change in project i , holding its competitors fixed. When a counterfactual changes multiple applications in the same pool, I recompute the inclusion probabilities using equation (C.5) at the counterfactual indices while holding the number of awards K_c fixed. The resulting probabilities therefore continue to satisfy equation (C.6) within every pool.

C.4 Estimation and Inference

I estimate $(\gamma, \lambda_1, \lambda_2)$ by conditional maximum likelihood. Let W_c denote the observed winner set in pool c . Summing equation (C.3) across pools with within-pool outcome variation gives

$$\mathcal{L}(\gamma, \lambda_1, \lambda_2) = \sum_{c: 0 < K_c < n_c} \left[\sum_{i \in W_c} v_i - \log e_{K_c}(w) \right]. \quad (\text{C.10})$$

Let x_i collect the regressors entering v_i . Differentiating equation (C.10) gives the score

$$\frac{\partial \mathcal{L}}{\partial (\gamma', \lambda_1, \lambda_2)'} = \sum_c \sum_{i \in I_c} x_i (A_i - \omega_i). \quad (\text{C.11})$$

The fitted value in the score is therefore the inclusion probability ω_i , and equation (C.6) implies that fitted and observed awards add up to the same K_c within every pool.

I compute the likelihood and inclusion probabilities without enumerating the $\binom{n_c}{K_c}$ possible winner sets. The denominator follows from the recursion

$$e_k(w_{1:j}) = e_k(w_{1:j-1}) + w_j e_{k-1}(w_{1:j-1}), \quad (\text{C.12})$$

which evaluates the likelihood in $O(n_c K_c)$ operations per pool. The leave-one-out terms required for the inclusion probabilities satisfy

$$e_k(w_{-i}) = e_k(w) - w_i e_{k-1}(w_{-i}), \quad e_0(w_{-i}) = 1.$$

Thus neither estimation nor prediction requires enumeration of feasible winner sets.

For inference, I use a sandwich covariance estimator constructed from the conditional-likelihood scores and cluster the scores at the competition-pool level. This allows award outcomes within the same pool to be correlated beyond the dependence induced mechanically by conditioning on K_c . The resulting covariance matrix is used for the award-stage standard errors reported in the main text and is propagated through the simulated project-size model as described in Appendix D.1.

C.5 Discussion

I hold the number of winners (K_c) fixed within each competition pool, so when counterfactual project sizes change credit requests, total credits awarded within a pool may also change, potentially generating implicit reallocations across pools. In Texas, however, sub-region credit allocations are not binding budgets. The Regional Allocation Formula (RAF) determines only the amount initially available, an individual application may request up to 150 percent of its sub-region's initial

allocation, and unused credits are subsequently redistributed through rural and statewide collapse procedures. Consistent with these rules, 49 percent of sub-region-years award more than their initial RAF allocation. Moreover, the counterfactual generates relatively small reallocations. median pool-level credit use changes by only 0.1 percent, and only 2.5 percent of pools change by more than 10 percent.

Table C1: Sub-region Pool Budget

Year	Credits (\$M)		Awards K_c		Share of sub-regions (%)		Awarded / RAF amount		
	RAF	Awarded	Total	Median	No award	Above RAF	p10	Median	p90
2010	46.7	65.2	51	1	19	65	0.00	1.36	2.53
2011	48.0	39.7	33	1	19	38	0.00	0.80	2.48
2012	43.7	40.3	35	1	12	46	0.16	0.98	1.44
2013	51.6	53.5	55	1	0	35	0.92	0.99	1.49
2014	52.6	52.8	53	1	4	38	0.80	0.98	1.44
2015	51.7	52.2	47	1	0	42	0.88	1.00	1.50
2016	56.1	57.3	52	1	0	27	0.90	0.97	1.47
2017	56.9	58.6	53	1	0	38	0.95	0.99	1.49
2018	65.6	69.7	63	2	0	65	0.88	1.30	1.83
2019	67.5	67.6	55	1	8	42	0.84	1.00	1.50
2020	69.5	69.8	54	1	4	54	0.91	1.09	1.50
2021	70.2	70.7	55	1	0	54	0.75	1.01	1.50
2022	65.3	60.1	42	1	4	54	0.20	1.11	1.50
2023	70.2	65.0	42	1	4	62	0.72	1.47	1.50
2024	74.0	85.5	50	1	4	73	0.00	1.27	1.50
All	889.6	907.8	740	1	5	49	0.64	1.00	1.50

D Additional Details in Model Estimation

This appendix provides additional details on the model estimation in Section 6.

D.1 Supply Estimation

In the supply estimation, I first estimate the award process by conditional maximum likelihood and the construction-cost parameters by least squares. Conditional on these estimates, I estimate the remaining parameters governing project size by simulated method of moments. I describe the simulation procedure, moments, inference, identification, and model fit below.

For each candidate parameter vector, I simulate delivered project size for every application using $S = 20$ draws of the profit shock $\varepsilon_i \sim N(0, 1)$. I use common random numbers throughout the estimation. Table D1 summarizes the procedure.

Table D1: Simulation of Delivered Project Size

#	Step	Procedure
1	Construct model objects	Compute $\pi_i^E(q)$, $\pi_i^R(q)$, $vc_i(q)$, $\omega_i(q)$, and their derivatives on 240 grid points spanning 5 to 500 units
2	Set revenue curvature	Set $\rho_i = \rho^0 + \rho^R \mathbf{1}_i\{\text{Rural}\} + \rho^{SH} \mathbf{1}_i\{\text{SH}\} + \rho^S \mathbf{1}_i\{\text{Senior}\} + \rho^t$
3	Draw profit shocks	Draw $\varepsilon_i^{(s)} \sim N(0, 1)$, $s = 1, \dots, 20$, using a fixed seed
4	Solve for interior size	Solve equation (14) on the size grid for each application and simulation draw, with linear interpolation between adjacent grid points
5	Impose institutional constraints	Apply the 16-unit minimum, zoning, and, where applicable, the 80-unit cap
6	Apply template adoption	Map the constrained optimum to the nearest $m \in \mathcal{M}$ with the probability in equation (15), using matched uniform draws
7	Concentrate cohort effects	Adjust ρ^t so that simulated and observed mean project size coincide within each application year
8	Construct moments	Compute the 21 simulated moments used in estimation

Notes: Profit shocks and uniform draws are held fixed across parameter vectors. The size grid contains 240 equally spaced points between 5 and 500 units.

Solving for project size. For each application and simulation draw, I solve the developer's first-order condition from equation (14),

$$\omega_i(q) \left[\bar{\pi}'_i(q) + \sigma^D \varepsilon_i^{(s)} \right] + \omega'_i(q) \left[\bar{\pi}_i(q) + \sigma^D \varepsilon_i^{(s)} q \right] = 0, \quad (\text{D.1})$$

where

$$\bar{\pi}_i(q) = \pi_i^E(q) + \pi_i^R(q) - f_i - vc_i(q). \quad (\text{D.2})$$

I evaluate equation (D.1) on a 240-point grid spanning 5 to 500 units and linearly interpolate at the first sign change. If the first-order condition does not cross zero, I use the grid point at which its absolute value is minimized. At the estimated parameters, 99.4% of application–draw pairs have a unique root.

Institutional constraints and template adoption. I next impose the 16-unit minimum and, for rural new construction and adaptive-reuse projects, the 80-unit cap, zoning and cap for urban projects. The site’s zoning code limits the number of units that can be developed on the parcel. When the zoning verification letter or municipal ordinance specifies a maximum residential density, I set the zoning capacity $\bar{z}_i$ equal to the permitted units per acre times the site’s acreage. I set $\bar{z}_i = \infty$ when no binding density limit is observed, including sites in Houston and its extraterritorial jurisdiction, districts without a density standard, and sites with rezoning pending at the time of application. Finally, developers may then adopt one of the standardized project sizes

$$\mathcal{M} = \{24, 36, 48, 60, 72, 80, 96, 100, 120, 152, 180, 200\}$$

according to equation (15). The adoption decision uses matched uniform draws held fixed across parameter vectors.

Cohort Effects. Rental-cash-flow curvature is

$$\rho_i = \rho^0 + \rho^R \mathbf{1}_i\{\text{Rural}\} + \rho^{SH} \mathbf{1}_i\{\text{SH}\} + \rho^S \mathbf{1}_i\{\text{Senior}\} + \rho^t,$$

as in equation (9). Rather than estimate the fifteen cohort effects ρ^t jointly with the remaining size parameters, I concentrate them out. At each evaluation of the objective, I choose the cohort effects so that simulated mean project size equals observed mean project size in each application year:

$$\frac{1}{S n_t} \sum_{s=1}^S \sum_{i: t(i)=t} q_i^{(s)} = \frac{1}{n_t} \sum_{i: t(i)=t} q_i, \quad t = 2010, \dots, 2024. \quad (\text{D.3})$$

Because the cohort effects absorb changes in average project size across years, I normalize

$$\sum_t \rho^t = 0. \quad (\text{D.4})$$

The intercept ρ^0 therefore measures average rental-cash-flow curvature, while ρ^t captures deviations from that average.

Moments and Weighting. I use 21 moments to estimate the project-size parameters. The moments summarize the distribution of project size, bunching at standardized sizes, the rural size constraint, differences across target populations, and the tax-credit cap. For the template moments, distance is measured relative to the nearest standardized size,

$$\min_{m \in \mathcal{M}} |q - m|,$$

after rounding project size to the nearest unit.

Location and spread (7): mean and standard deviation of q ;

10th, 25th, 50th, 75th, and 90th percentiles of q ;

Tails (2): $\Pr(q > 120)$, $\Pr(q \leq 49)$;

Templates (5): $\Pr\left(\min_{m \in \mathcal{M}} |q - m| = 0\right)$,
 $\Pr\left(\min_{m \in \mathcal{M}} |q - m| \in \{1, 2\}\right)$,
 $\Pr\left(3 \leq \min_{m \in \mathcal{M}} |q - m| \leq 5\right)$,
 $\Pr\left(6 \leq \min_{m \in \mathcal{M}} |q - m| \leq 10\right)$,
 $\Pr\left(\min_{m \in \mathcal{M}} |q - m| \geq 11\right)$;

Rural (4): $\mathbb{E}[q \mid \text{Rural}]$, $\text{SD}(q \mid \text{Rural})$,

$\Pr(76 \leq q \leq 80 \mid \text{80-unit cap})$, $\Pr(q \leq 49 \mid \text{Rural})$;

Segments (2): $\mathbb{E}[q \mid \text{Senior}]$, $\mathbb{E}[q \mid \text{SH}]$;

Credit cap (1): $\Pr(\text{credit cap binds} \mid \text{Urban})$.

For the final moment, the simulated credit request follows equation (8). I classify a simulated project as being at the cap when its credit request reaches at least 97% of $\bar{c}_t$. In the data, I apply the same 97% threshold to the reported credit request.

I minimize the weighted distance between the simulated and observed moments:

$$Q = \sum_{k=1}^{21} \left(\frac{\hat{m}_k - m_k}{\hat{s}_k} \right)^2, \quad (\text{D.5})$$

where m_k and $\hat{m}_k$ denote the observed and simulated moments, respectively, and

$$\hat{s}_k = \max \{ \text{SE}_{\text{boot}}(m_k), \underline{s}_k \}. \quad (\text{D.6})$$

Bootstrap standard errors are computed from 500 resamples of applications. I set the lower bound $\underline{s}_k$ to one unit for size moments and 0.01 for shares. The lower bound prevents moments at empirical mass points from receiving disproportionate weight because of mechanically small sampling variation.

Optimization. I minimize equation (D.5) over

$$(\rho^{SH}, \rho^S, \rho^R, \sigma^D, \kappa_0, \kappa_1),$$

with σ^D and κ_1 optimized in logs to impose positivity. I use Nelder–Mead because the institutional constraints and template adoption make the simulated objective only piecewise smooth. The search uses twelve starting values, eleven drawn from a Latin-hypercube design over the economically relevant parameter region, and retains the solution with the lowest objective value. Common random numbers are held fixed throughout.

The objective has a second local minimum in which greater profit-shock dispersion substitutes for template adoption. Its objective value is two to three times that of the reported solution. I therefore report estimates from the lower-objective basin.

Standard Errors. Award-stage standard errors use the pool–year clustered sandwich estimator described in Appendix C. For the project-size parameters, I compute the Jacobian of the 21 simulated moments with respect to $(\rho^{SH}, \rho^S, \rho^R, \sigma^D, \kappa_0, \kappa_1)$ numerically by central differences, re-concentrating the cohort effects at each perturbation. The covariance matrix uses the bootstrap covariance of the data moments, adjusted for simulation noise, and propagates uncertainty from the estimated award parameters through the simulated size moments.

Specifically, the covariance matrix is

$$\hat{V} = (G'WG)^{-1}G'W \left[\left(1 + \frac{1}{S}\right) \hat{\Omega}_m + G_b \hat{V}_{\text{aw}} G_b' \right] WG(G'WG)^{-1}, \quad (\text{D.7})$$

where G is the Jacobian of the simulated moments with respect to the project-size parameters, W is the diagonal weighting matrix implied by equation (D.6), $\hat{\Omega}_m$ is the bootstrap covariance matrix of the data moments, $\hat{V}_{\text{aw}}$ is the award-stage covariance matrix, and G_b is the Jacobian of the simulated moments with respect to the award parameters. I compute G_b by adjusting each award coefficient and re-solving the project-size problem. The second term therefore propagates uncertainty from

the award stage into the project-size estimates.

Standard errors for ρ^0 and the cohort effects account for the sampling variation in the cohort mean sizes used to concentrate them out. For each cohort, I divide the sampling error of mean project size by the numerical response of simulated mean size to the corresponding cohort effect.

Identification. The award and construction-cost parameters are identified as described in the main text. In the project-size block, the cohort effects ρ^t match mean project size within application year. Conditional on these means, ρ^R , ρ^{SH} , and ρ^S are disciplined by differences in the size distributions of rural, supportive-housing, and senior projects. The profit-shock dispersion σ^D is disciplined primarily by the within-cohort dispersion and tails of the project-size distribution.

The template parameters are disciplined by bunching around the standardized sizes in $\mathcal{M}$. The parameter κ_0 shifts the overall probability of adopting a standardized design, while κ_1 determines how quickly adoption declines with the distance between the interior optimum and the nearest standardized size. The on-template share and the distribution of distances from $\mathcal{M}$ jointly distinguish these margins. The data are most informative about adoption near a standardized size and less informative about the adoption probability far from the menu.

Model Fit. Table D2 reports the cohort effects, and Table D3 compares the observed and simulated moments. The model reproduces the main features of the size distribution, including its dispersion, the rural cap, and bunching at standardized project sizes. The largest discrepancies occur around several off-menu multiples of four and in the upper tail. In particular, many observed projects at 40, 56, 64, 76, 84, and 104 units lie four units from a standardized size. Because the model allows either adoption of a size in $\mathcal{M}$ or the continuous constrained optimum, it does not reproduce all of this additional discreteness. The model also modestly overpredicts the urban share of applications at the tax-credit cap.

Table D2: Cohort Effects in Rental Cash-Flow Curvature

Application year	ρ^t
2010	+0.131
2011	-0.006
2012	+0.017
2013	+0.138
2014	+0.087
2015	+0.125
2016	+0.016
2017	+0.005
2018	+0.101
2019	-0.056
2020	-0.050
2021	+0.014
2022	-0.328
2023	-0.218
2024	+0.024

Notes: This table reports the cohort effects ρ^t in equation (9). The effects are normalized to sum to zero and are concentrated out by matching simulated and observed mean project size within each application year, as in equation (D.3). A positive value implies less concavity in rental cash flow and therefore a higher marginal return to project size.

Table D3: Model Fit: Data and Simulated Moments

	Data	Model
Mean units	91.126	91.105
SD units	43.086	41.777
10th percentile	40.000	39.639
25th percentile	60.000	60.000
Median	80.000	84.317
75th percentile	120.000	120.000
90th percentile	150.000	152.000
Share >120 units	0.219	0.210
Share ≤ 49 units	0.201	0.169
Share on QAP menu	0.388	0.412
Rural mean units	57.551	57.781
Rural SD units	21.985	23.870
Rural New Construction share 76–80	0.300	0.262
Rural share ≤ 49	0.454	0.387
Urban share at cap	0.449	0.540
Senior mean units	86.792	85.831
Supportive-housing mean units	88.978	83.262
Share 1–2 units from a template	0.147	0.173
Share 3–5 units from a template	0.232	0.193
Share 6–10 units from a template	0.171	0.167
Share ≥ 11 units from a template	0.062	0.055
KS statistic (size dist.)	–	0.038
Corr(sim, obs) size	–	0.630

Notes: The first 21 rows report the moments entering equation (D.5), computed from the 1,910 observed applications and the corresponding simulated project sizes. The urban credit-cap moment is the share of applications whose credit request reaches at least 97% of the applicable cap. Template-distance moments measure distance in units from the nearest size in $\mathcal{M}$. The KS statistic is the maximum distance between the observed and simulated project-size distributions. The correlation is between observed project size and each application’s mean simulated size across draws.

D.2 Supply Model Extension: Unit Size as Alternative Measure


Estimation. The baseline model measures project size by the number of units, consistent with how the QAP defines the minimum project size, rural cap, developer-fee threshold, and other institutional constraints. Unit counts, however, do not capture differences in unit size across projects. I therefore re-estimate the supply model using the number of bedrooms as the measure of project size. The results are quantitatively similar to the baseline specification.

I construct each project's bedroom mix from the Bedrooms panel of the Populations Served sheet and project-level rent schedule. I code efficiency units as one bedroom. The sample mean is 1.82 bedrooms per unit, with larger units in family projects (2.05) than in supportive-housing (1.51) and senior projects (1.41). Average bedrooms per unit are similar across rural and urban projects and do not systematically vary with unit count.

I re-estimate the supply model using bedrooms B_i as the size choice, holding each project's bedroom mix $\bar{b}_i$ fixed. Unit-based constraints are scaled by $\bar{b}_i$, while the dollar-denominated tax-credit cap remains unchanged. Construction costs and rental cash flows are re-estimated in bedrooms, while the award model remains unchanged. I use the same SMM procedure as in the baseline, expressing size moments in bedrooms and QAP thresholds in units-equivalent, $B_i/\bar{b}_i$.

Table D4 reports the estimates. The award estimates are unchanged, while the cost and size-choice parameters rescale with the bedroom denomination. The main heterogeneity parameters retain similar signs and relative magnitudes across specifications.

Figure D1: Model Fit with Project Size Measured in Bedrooms



Notes: Panel (a) compares the distribution of project size in the data and model simulations when size is measured in bedrooms. Panel (b) compares observed winning rates with predicted award probabilities; the dashed 45-degree line indicates perfect calibration.

Table D4: Supply Estimates with Project Size Measured in Bedrooms

	Parameter	Estimate	SE
Panel A: Award stage (MLE)			
Residualized QAP score	γ_1	0.168***	(0.046)
USDA financing	γ_2	1.049***	(0.353)
Nonprofit set-aside	γ_3	0.142	(0.149)
Subsidy ratio μ_i	λ_1	0.430**	(0.189)
Affordability profile ϕ_i	λ_2	-0.180	(1.616)
Panel B: Per-unit cost (OLS)			
Fixed eligible basis (\$M)	α_0	1.872***	(0.171)
Per-bedroom cost, new construction	α_1	5.560***	(0.127)
Rehabilitation shifter	α_2	-2.762***	(0.197)
Density (inverse acreage)	α_3	1.264***	(0.351)
Panel C: Size stage (SMM)			
Revenue curvature	ρ^0	-0.193***	(0.009)
Supportive-housing shifter	ρ^{SH}	-0.167***	(0.048)
Senior shifter	ρ^S	-0.121***	(0.017)
Rural shifter	ρ^R	-0.233***	(0.014)
Profit-shock dispersion	σ	1.143***	(0.047)
Template adoption: level	κ_0	-0.239	(0.177)
Template adoption: distance decay	κ_1	0.070***	(0.023)
Applications		1,910	
Awards		961	

Notes: This table re-estimates the supply model measuring project size in bedrooms. Panel A reports the award estimates, Panel B the construction-cost estimates, and Panel C the size-choice estimates. Standard errors are clustered by competition pool in Panel A and are heteroskedasticity-robust in Panel B. Panel C reports SMM standard errors incorporating uncertainty from the award stage.

Figure D1 shows that the model reproduces the bedroom-size distribution about as well as it reproduces the unit-size distribution in the baseline specification. The model also closely matches the mean, median, and 25th percentile of the bedroom distribution, as well as the shares of projects adopting a template and reaching the tax-credit request cap. The remaining discrepancies occur in the same parts of the distribution as in the baseline specification, primarily the upper tail and the mass near the rural cap. Overall, measuring project size in bedrooms rather than units mainly rescales the cost and size-choice parameters.

Counterfactual. Table D5 shows similar supply responses across the two specifications. Under budget-neutral price equalization, expected awarded housing rises by 169 units-equivalent (0.23

percent) in the bedroom model, compared with 168 units (0.23 percent) in the baseline, with similar gains across deposit quartiles. Under credit sizing, the corresponding gains are 585 (0.80 percent) and 840 (1.16 percent). The smaller credit-sizing response in the bedroom model reflects a weaker response of credit requests and hence a smaller adjustment in θ_{1t} .

The underlying price elasticity is nearly identical: a 10-cent increase in tax-credit prices raises mean optimal project size by 9.3 percent in the baseline and 9.2 percent in the bedroom model, and raises expected awarded housing by 6.8 percent in both. Thus, measuring project size in bedrooms leaves the estimated price responsiveness of supply and the qualitative counterfactual conclusions largely unchanged.

Table D5: Counterfactual supply response: baseline versus bedroom model

	Baseline model	Bedroom model	
	(units)	(units-equivalent)	(bedrooms)
Panel A: budget-neutral price equalization, statewide 2012–2024			
Baseline expected awarded housing	72,382	73,330	126,442
Budget-neutral price (δ_t)	+167 (+0.23%)	+169 (+0.23%)	+270 (+0.21%)
Credit sizing (θ_{1t} scaled)	+840 (+1.16%)	+585 (+0.80%)	+948 (+0.75%)
Panel B: budget-neutral price closure by 2015 large-bank deposit quartile, percent change			
Q1 (low deposits)	+4.33%	+2.36%	+2.53%
Q2	+3.96%	+3.67%	+3.40%
Q3	+2.81%	+2.74%	+2.66%
Q4 (high deposits)	-0.47%	-0.42%	-0.37%
Panel C: credit sizing closure by 2015 large-bank deposit quartile, percent change			
Q1 (low deposits)	+5.21%	+2.68%	+2.87%
Q2	+5.11%	+4.46%	+4.11%
Q3	+3.92%	+3.39%	+3.30%
Q4 (high deposits)	+0.43%	+0.13%	+0.15%
Panel D: a 10-cent increase in every application's tax credit price			
Mean optimal project size, status quo	88.7	88.8	157.3
Change in mean optimal project size	+8.27 (+9.33%)	+8.21 (+9.24%)	+14.44 (+9.19%)
Change in expected awarded housing	+4,943 (+6.83%)	+4,950 (+6.75%)	+8,414 (+6.65%)
Change in expected number of awards	-20.6	-20.3	

Notes. Expected awarded housing over the 2012–2024 cohorts (1,654 applications in 167 counties) under the counterfactuals of Section 7. Bedroom-model outcomes are converted to units project by project ($B_i/\bar{b}_i$) and also reported in bedrooms. Panels A–C compare the three budget-neutral closures and their effects across quartiles of 2015 large-bank deposits. Panel D increases every application's tax credit price by 10 cents. Simulation draws are common across regimes and models.

D.3 Demand Estimation

This appendix provides additional details on the estimation of the demand model in Section 6.2. I estimate

$$\Theta^D = (\tau, \eta, \sigma^\zeta, \sigma^d)$$

by simulated method of moments (SMM), taking the county benchmark ζ_c and the fundamental value ν_t as calibrated inputs. For each candidate Θ^D , I solve for equilibrium prices and banks' origination choices, concentrate out the bank and state components of direct-origination costs, and simulate the model using common random numbers.

Estimation sample and calibration. The unit of observation is a bank–county–evaluation period. I partition 2010–2024 into four windows, 2010–13, 2014–17, 2018–21, and 2022–24, and treat the underwriting desk as a fixed cost incurred once per window. Within each window, I aggregate credits, equity, and deposits; define the county price as total equity divided by total credits; and construct the syndicated share as the equity-weighted share of syndicated investments. A bank is classified as originating directly if it is named as the equity investor on at least two applications during the window. The resulting panel covers 15 states, 1,548 county–periods, and 14,965 bank–county–period observations. Table D6 summarizes the procedure.

Moments. I use nine moments to identify the four structural parameters. County-level variables are demeaned within state–period before computing slopes and standard deviations. The moments are:

Syndication fee (2): slope of p_{ct} on the syndicated share;
deal-level hedonic discount on syndicated deals;

Prices (3): slope of p_{ct} on $\log D_{ct}$;
slope of p_{ct} on $\tilde{S}_{ct}/D_{ct}$;
s.d. of p_{ct} ;

Channel composition (2): named banks' direct share of county equity;
slope of the syndicated share on $\log D_{ct}$;

Origination (2): within-bank slope of d_{ct}^b on $\log D_{ct}^b$;
share of bank–state pairs that never originate directly.

The price-pressure moment—the slope of p_{ct} on $\tilde{S}_{ct}/D_{ct}$ —is targeted to its 2SLS estimate of -0.758. I instrument supply relative to deposits using county QCT and DDA tract counts and a

Table D6: Simulation of Bank Demand and Channel Choice

	Step	Procedure
1	Construct model objects	For each candidate $\Theta^D = (\tau, \eta, \sigma^\zeta, \sigma^d)$, construct bank deposits, county credit supply, the CRA benchmark, and the fundamental value for each bank–county–period
2	Draw shocks	Draw county–period benchmark shocks $\varepsilon_{ct}^{\zeta(r)}$ and bank–county–period origination shocks $\xi_{ct}^{b(r)}$, using common random numbers across parameter vectors
3	Guess county prices	Initialize the county price p_{ct} and compute the corresponding CRA premium $x_{ct} = 1 - \nu_t/p_{ct}$
4	Optimize within each channel	For each bank, solve the investment problem separately under direct origination and syndication, obtaining optimal quantities $a_{ct}^{D,b*}$ and $a_{ct}^{S,b*}$ and the associated optimized payoffs $V_{ct}^{D,b*}$ and $V_{ct}^{S,b*}$
5	Choose origination channel	Compare the optimized payoffs and choose direct origination when $V_{ct}^{D,b*} - V_{ct}^{S,b*} + \xi_{ct}^{b(r)} > 0$; banks without direct origination capacity are restricted to syndication
6	Clear county markets	Aggregate banks' optimal investments under their chosen channels and update p_{ct} until aggregate bank investment equals county tax-credit equity supply
7	Solve channel–price equilibrium	Repeat Steps 4–6 until bank channel choices, investment quantities, and county prices jointly converge
8	Concentrate cost effects	Adjust the bank and state components of direct-origination costs so that simulated direct-origination rates match their observed counterparts, resolving the county equilibria after each update
9	Construct moments	Compute the nine simulated moments of prices, channel composition, and origination choices used in estimation
10	Evaluate SMM criterion	Average moments across simulation draws, evaluate equation (D.8), and update Θ^D until the SMM objective converges

Notes: Benchmark and origination shocks are held fixed across parameter vectors. At each candidate parameter vector, investment quantities are optimized separately under direct origination and syndication before the bank chooses its origination channel, so the channel decision compares optimized payoffs rather than costs evaluated at a common investment quantity.

shift-share instrument that interacts pre-2005 county award shares with leave-own-out state credit totals, averaged over each evaluation window. The county-clustered first-stage F -statistic is 14.9. To ensure that the empirical and simulated moments isolate the same variation, I apply the same 2SLS specification and instruments to each simulated dataset and use the resulting IV coefficient as the simulated counterpart. Column (1) of Table 9 excludes this moment.

The within-bank origination gradient is computed after demeaning d_{ct}^b and $\log D_{ct}^b$ within each of the 18 named banks. The never-originate moment uses the 64 named bank–state pairs with at least 15 channel cells.

SMM criterion. For each candidate Θ^D , I simulate the model using $R = 60$ draws of the benchmark and origination shocks. The same draws are held fixed across objective evaluations so that changes in simulated moments reflect parameter changes rather than simulation noise. The criterion is

$$Q(\Theta^D) = \sum_{k=1}^9 \left(\frac{\hat{m}_k(\Theta^D) - m_k}{\hat{s}_k} \right)^2, \quad \hat{m}_k(\Theta^D) = \frac{1}{R} \sum_{r=1}^R m_k^{(r)}(\Theta^D), \quad (\text{D.8})$$

where $\hat{s}_k$ is the standard deviation of moment k across 200 state-block bootstrap samples. I re-estimate the IV target in each bootstrap sample and use the regression standard error for the hedonic moment. Resampling at the state level preserves dependence in prices and origination within a state's LIHTC allocation system.

Optimization and inference. I minimize equation (D.8) over the logarithms of the four structural parameters. I use Nelder–Mead because discrete realized channel choices and numerically concentrated cost effects make the simulated criterion nonsmooth. The initial simplex perturbs each log parameter by 0.1, and the algorithm terminates when the simplex differs by less than 10^{-3} in parameters and 0.02 in the objective.

Table 9 reports stratified county-block bootstrap standard errors based on 50 replications per specification. Within each state, I resample counties with replacement, preserving the state's number of counties and moving all periods of a sampled county together. Each replication re-estimates the IV target, draws the hedonic target from $\mathcal{N}(0.0130, 0.0019^2)$, reclassifies direct-capable banks using the resampled commitments, and re-estimates the model starting from the full-sample estimates while holding the weighting matrix fixed. For computational tractability, bootstrap replications use 20 simulation draws, 12 bisection steps, four cost-effect updates per objective evaluation, and at most 200 optimization iterations. Resampling counties within states keeps all state effects identified in every replication.

As a comparison, I also compute the asymptotic covariance matrix

$$\hat{V} = (G'WG)^{-1}G'W \left[\left(1 + \frac{1}{R} \right) \hat{\Omega}_m \right] WG(G'WG)^{-1}, \quad (\text{D.9})$$

where $W = \text{diag}(\hat{s}_k^{(-2)})$, $\hat{\Omega}_m$ is the state-block bootstrap covariance matrix of the data moments, and G is the Jacobian of simulated moments with respect to the parameters. I compute G using central differences with 5 percent perturbations and 200 simulation draws, re-solving the concentrated cost effects from common starting values at each perturbation.

The asymptotic and bootstrap standard errors are similar. In column (2), the asymptotic standard errors for τ , η , σ^ζ , and σ^d are 0.0021, 0.129, 0.0082, and 0.312, compared with bootstrap standard errors of 0.0024, 0.111, 0.0087, and 0.232. I report the bootstrap estimates because the asymptotic

calculation conditions on the observed origination rates targeted by the concentration step.

Identification. The parameters are identified by distinct features of prices and banks' origination choices. The syndication fee τ is disciplined primarily by the deal-level hedonic discount. The price gradient with respect to the syndicated share provides additional information because a larger τ makes equilibrium prices more sensitive to the composition of investment channels. Consistent with this interpretation, dropping the hedonic moment moves τ toward 0.02 along a relatively flat ridge with σ^d .

Conditional on τ , the dispersion in direct-origination costs, σ^d , is identified primarily by the within-bank relationship between direct origination and local deposits. Equation (21) implies that the sensitivity of the origination probability to a bank's local commitment depends on the ratio τ/σ^d . The CRA curvature parameter η is disciplined by the response of prices to deposits and to tax-credit supply relative to deposits. Finally, σ^ζ is identified by residual price dispersion not explained by these gradients. Because the price moments are constructed within state-period, their identifying variation is largely insensitive to the calibrated level of the CRA benchmark. The local sensitivity matrix $\Lambda = (G'WG)^{-1}G'W$, evaluated at the column (2) estimates, confirms these patterns.

Model fit. Table D7 compares the observed and simulated moments. Because the weighting matrix is diagonal rather than the inverse of the full moment covariance matrix, I use the criterion value J to compare specifications rather than as a formal specification test.

In column (2), the criterion is $J = 18.6$. The largest contributions come from the syndicated-share gradient with respect to county deposits (6.4), the within-bank deposit gradient (4.0), the price gradient with respect to the syndicated share (2.9), the price-deposit gradient (2.6), the price-pressure gradient (1.5), and named banks' direct share of equity (1.0); the remaining moments contribute less than 0.1 in total. The two largest discrepancies have opposite implications for the source of spatial variation in origination. The model predicts that the syndicated share declines more strongly with county deposits than in the data (-0.041 versus -0.027), while an individual bank's direct-origination probability responds less strongly to its own local deposits (0.039 versus 0.054).

By construction, the concentration step reproduces each bank's and state's direct-origination rate. In column (1), where the price-pressure moment is excluded from estimation, the model predicts a price-pressure slope of -0.50, compared with the IV estimate of -0.758, approximately 1.6 standard errors away.

Table D7: Demand Model Fit: Data and Simulated Moments

	Data	Model (1)	Model (2)
Price on syndicated share	-0.0355	-0.0236	-0.0256
Price on log large-bank deposits	0.0057	0.0066	0.0074
Price on $\tilde{S}/D$ (instrumented)	-0.758	-0.502 [†]	-0.571
S.d. of price	0.0382	0.0379	0.0380
Named-bank direct share of equity	0.153	0.132	0.125
Syndicated share on log deposits	-0.0265	-0.0411	-0.0412
Direct-origination rate	0.132	0.132 [‡]	0.132 [‡]
S.d. of bank direct-origination rates	0.103	0.104 [‡]	0.104 [‡]
Direct origination on log own deposits (within bank)	0.0538	0.0383	0.0396
Syndication discount (hedonic)	0.0132	0.0134	0.0128
Bank–state pairs never originating directly	0.344	0.359	0.357
Criterion J (standard-error weighted)	–	16.62	18.66
Overidentifying restrictions	–	4	5

Notes: Model moments are averages across $R = 60$ simulation draws at the estimates in Table 9. County-level moments are computed after demeaning within state–period. [†] Not included in the criterion in column (1); the reported value is the model’s out-of-criterion prediction. [‡] Not included in the criterion; bank and state cost effects are concentrated out to match these rates.

E Data Appendix: Bank Public Welfare Investments

This appendix describes the construction of the bank-level data used in the analysis. The procedure has three steps. First, I use OCC public welfare investment filings to measure banks' investments in LIHTC and other affordable housing. Second, I combine these investments with the FDIC Summary of Deposits (SOD) to construct the bank $\times$ county $\times$ year panel used to relate local deposit exposure to CRA-motivated investment. Third, I use FDIC institution-history records to identify mergers that generate plausibly exogenous changes in an acquiring bank's deposits within its existing geographic footprint.

E.1 OCC Public Welfare Investments

Part 24 Investment Filings. My measure of bank investment comes from filings made under 12 CFR Part 24. Part 24 permits national banks and federal savings associations to make equity investments designed primarily to promote the public welfare, principally by benefiting low- and moderate-income individuals and communities. These investments include equity stakes in LIHTC partnerships and constitute an important channel through which banks make CRA-eligible community development investments. A bank may invest up to 5 percent of capital and surplus under Part 24, or up to 15 percent with supervisory approval.

For each investment, the filing identifies the investing institution, investment vehicle, dollar amount, geographic target area, and, for most of the sample period, the activity being financed. The OCC compiles these filings in its *Community Development Investments At-A-Glance* releases. A key advantage is that individual investments are observed directly rather than inferred from aggregate measures of CRA activity.

The OCC is the only federal banking regulator that publicly reports these investments at the individual-investment level. The Federal Reserve and FDIC administer parallel authorities for state-chartered banks but do not release comparable investment-level records. I therefore observe only OCC-supervised institutions. This restriction nevertheless captures the major bank investors in the LIHTC market: Bank of America, JPMorgan Chase, Wells Fargo, U.S. Bank, PNC, and Citibank all operate national bank charters and appear in the Part 24 data.

Geographic assignment. Part 24 filings report the area targeted by an investment rather than the address of a particular property. I therefore use counties as the geographic unit of analysis, matching the level at which I measure banks' deposit exposure. Among the target areas named in housing records, 69.6 percent are a single city and 16.3 percent a single state, while 29.7 percent of dollars are targeted "Nationwide" and cannot be localized. I map reported cities to counties using a crosswalk constructed from the HUD LIHTC database's city and county identifiers; 98.6 percent

of city references resolve. This procedure assigns \$92.7 billion, or 56.7 percent, of housing dollars to counties.

A small share of filings report one investment against several target areas without specifying how the amount is divided. These combined-area filings make up 7.6 percent of housing records and 11.7 percent of housing dollars. In the baseline, I split the investment equally across all listed areas, including state-level and “Nationwide” targets that subsequently fall outside the county panel. This prevents the full value of a multi-area fund from being mechanically assigned to the subset of locations that can be resolved to counties. Results are similar when I instead exclude combined-area filings altogether.

Sample Period and Investment Type. I compile all 28 annual OCC releases from 1997 through 2024. After removing filings that a release lists more than once, they contain 18,311 investments totaling \$249.1 billion. The analysis uses investments made between 2004 and 2022. Before 2003, the releases do not report either the geographic target or the activity financed. These fields first appear in 2003, but geographic information in that year is reported only at the state or national level and no investment is explicitly identified as LIHTC. I therefore begin the analysis in 2004. After 2022, the OCC stops reporting the activity field needed to distinguish LIHTC and other affordable-housing investments.

Affordable housing is the dominant use of Part 24 investment. As shown in Table E1, housing accounts for \$162.4 billion over 2004–2022, compared with \$28.3 billion devoted to other activities such as economic development and community development entities. Among housing investments, \$132.7 billion is explicitly identified as LIHTC.

I classify an investment as LIHTC when its OCC activity field explicitly refers to the LIHTC. The remaining housing investments are classified as non-LIHTC affordable housing. The distinction generally appears to reflect meaningful reporting rather than inconsistent labeling. The two categories coexist throughout the sample; 38.5 percent of bank–years contain both types of filings, and 82.2 percent of dollars reported simply as affordable housing come from banks that separately identify LIHTC investments in the same year. My baseline measure uses explicitly identified LIHTC investment.²⁸

²⁸The main exception is 2008–2009 and, to a lesser degree, 2014, when the share of housing investment explicitly labeled LIHTC falls sharply despite little change in total housing investment. Matching filing names against the HUD LIHTC property database confirms this is a labeling artifact: requiring a LIHTC property of the same name in a state the filing targets re-identifies filings worth about \$1.1 billion, raising the LIHTC share from 7 to 21 percent in 2008, from 23 to 35 percent in 2009, and from 56 to 63 percent in 2014. The recovery is partial because the unlabeled filings in those years are mostly syndicated funds, which have no property name to match and estimating on the recovered sample does not change the results.

Investing Bank. I map each OCC filing to the top-tier bank holding company reported in the SOD in the year of the investment.²⁹ Where the identifier recorded against a filing is not one the SOD recognises, I fall back on matching the institution by name within the year, restricting candidates to OCC-chartered institutions. Nineteen holding companies would otherwise be dropped from the panel entirely, because an unmatched identifier resolves to a holding company with no deposits anywhere. This maps the 121 investing institutions in the OCC data to 113 banking organizations.

I use the holding company as the bank unit throughout the analysis because the empirical variation of interest is a banking organization’s local deposit footprint and the associated CRA incentive to invest in a county. Using individual charters would additionally introduce artificial entry and exit when a holding company consolidates subsidiary charters without changing the underlying banking organization or its local deposit base.³⁰

I also consolidate duplicate or partial filings before aggregation. Multiple filings by the same bank for the same investment vehicle in the same year are combined. Filings for the same vehicle in different years are retained because LIHTC investments are frequently funded in multiple tranches. Downward revisions and residual amounts below \$1,000 are folded into the principal filing for the same vehicle, matched on bank and a canonical form of the vehicle name so that a revision filed as “Frenchtown Square Partners, LLC” is matched to the commitment filed as “Frenchtown Square Partners”. There are 144 such partial filings among 10,701 housing records; the 19 with no principal filing anywhere in the data are dropped rather than counted as investments in their own right.

E.2 Local Deposits and the Bank–County Panel

Local Deposits. I measure a bank’s local deposit exposure using the FDIC Summary of Deposits. The SOD reports deposits at every branch of an FDIC-insured institution as of June 30 each year. I aggregate branch deposits to the bank holding company $\times$ county $\times$ year level.³¹

Because the SOD assigns deposits to the branch at which they are booked, it can overstate a bank’s local deposit base when a main office books non-retail deposits raised outside the surrounding market. CRA examiners identify and exclude such deposits when evaluating a bank’s local deposit footprint. As a robustness check, I recover these exclusions from banks’ CRA performance evaluations and subtract them from the SOD measure, applying each reported non-retail deposit share between consecutive evaluations. Netting out these deposits reduces the measured

²⁹For institutions with no bank holding company, I use the charter’s own RSSD identifier. For charters that no longer appear in the SOD when an investment is reported, I use the nearest available SOD year.

³⁰In constructing the holding-company crosswalk, I recode the SOD value RSSDHCR=0, which denotes an institution with no holding company, to the institution’s own identifier. Otherwise all independent banks would incorrectly be treated as a single banking organization.

³¹If the bank appears in the SOD but has no branch in a county, deposits are coded as zero; bank–years absent from the SOD are dropped because deposits are then unobserved rather than zero.

deposit base by 16.2 percent.

This measure captures the local deposit exposure relevant for the CRA incentive studied throughout the paper. A bank with branches and a larger deposit base in a county has greater exposure to that local market for CRA purposes. Accordingly, a bank–county–year in which the banking organization appears in the SOD but has no branch in a particular county is coded as having zero deposits in that county. Bank–years in which the organization does not appear in the SOD are instead treated as missing, since in that case deposits cannot be distinguished from an unobserved institution.

As a validation exercise, I compare the resulting deposit measure with county deposits independently recorded in the CRA performance evaluation. Among the 1,992 bank–county–years observed in both sources, the measures agree almost exactly: their log correlation is 1.0000, the median ratio is 1.0000, and 99.8 percent differ by less than one percent. All investment and deposit amounts used in the analysis are expressed in 2024 dollars using the CPI-U.

Assigning Investments to Counties. Part 24 filings identify the geographic area targeted by an investment rather than the address of an individual LIHTC property. I therefore assign investments to counties using the geographic descriptions in the filings. County is the finest geography that can be constructed consistently over the sample and also corresponds to the local deposit measure used in the analysis.

When a filing reports a city, I map the city to a county using a crosswalk constructed from city and county identifiers in HUD’s LIHTC database. This procedure resolves most city references; I hand-code common neighborhoods, boroughs, and obvious spelling errors and leave genuinely ambiguous geographic references unresolved.

Geographic precision improves considerably over the sample period. City references account for 34 percent of the target areas named in LIHTC filings in 2004, exceed one-half from 2007, and run between 66 and 85 percent after 2010. Over 2004–2022, 40.3 percent of LIHTC dollars cannot be attributed to a county, because they are reported only at the state or national level or name a place the crosswalk cannot resolve. This limitation is concentrated in the early years: between 82 and 93 percent of dollars are unassignable in 2004–2007, compared with 23 to 42 percent thereafter.

Some filings list multiple geographic target areas without reporting how the investment is allocated across them. This occurs for 609 LIHTC filings, representing 11.8 percent of LIHTC dollars. In the baseline construction, I divide the reported amount equally across all listed areas. Areas identified only as a state or as “Nationwide” receive their share of the allocation but subsequently fall outside the county panel. This procedure prevents the full value of a multi-area or nationwide fund from being mechanically assigned to the subset of locations that can be resolved to counties. Because the allocation across areas is not observed, I also report specifications that exclude these

filings altogether.

Bank–County Investment Panel. After assigning investments to banks, counties, and years, I aggregate the Part 24 records to the bank holding company $\times$ county $\times$ year level. The LIHTC sample contains 4,840 positive bank–county–year cells, covering 78 banking organizations and 877 counties between 2004 and 2022 and representing \$79.2 billion of geographically attributable LIHTC investment.

For estimation, I complete the panel over the bank–county pairs that appear in the investment data: a pair enters if the bank filed at least one Part 24 LIHTC investment in that county during the window, and is then observed in every year of it. A bank that makes no investment in such a county in a particular year enters with zero investment rather than disappearing. These zeros are 88.2 percent of the resulting cells and identify the timing of investment within a bank–county relationship. Counties in which a bank never invests are absent from the panel rather than present as zeros, so the design speaks to when a bank invests where it is active, not to whether it becomes active somewhere new.

Table E1: Composition of Part 24 Public Welfare Investment, 2004–2022

	Records		Amount		Excl. 2008–09	
	Number	Share (%)	\$B	Share (%)	Records (%)	Amount (%)
All Part 24 investment	14,829	100.0	190.7	100.0	100.0	100.0
Affordable housing	10,556	71.2	162.4	85.2	71.0	85.3
LIHTC named	8,198	55.3	132.7	69.6	58.5	72.8
Other affordable housing	2,358	15.9	29.8	15.6	12.5	12.4
Other (NMTC, CEDEs)	4,273	28.8	28.3	14.8	29.0	14.7

Notes: OCC *Community Development Investments At-A-Glance* annual releases, 2004–2022, after removing filings the releases list twice. Classification is the OCC activity field. Shares are of all Part 24 investment, so the two indented levels each sum to the line above them; investments naming the LIHTC are 77.7 percent of affordable housing records and 81.7 percent of its dollars. The last two columns repeat the shares dropping 2008 and 2009, the two years in which the LIHTC label is applied unusually sparsely: the housing share is 85.3 percent of dollars, while the LIHTC share of housing dollars rises to 85.4 percent, so the dip is a labelling discontinuity rather than a shift out of the credit. Those two years hold 1,040 investments worth \$10.0 billion.

E.3 Bank Mergers and Acquisitions

The identification strategy uses bank mergers to isolate changes in a bank’s local deposit exposure that arise from the acquisition of another banking organization. I construct these shocks by combining the FDIC institution-history file with pre-merger branch deposits from the SOD.

Identifying Mergers Across Banking Organizations. The FDIC institution-history file records changes in the status of FDIC-insured charters, including the absorbed charter, surviving charter, transaction date, and type of transaction. The raw file contains 23,299 charter-level changes between 1969 and 2026. Many of these records, however, are reorganizations within an existing bank holding company rather than acquisitions between distinct banking organizations. Because these transactions do not create new deposit exposure for the holding company, they provide no identifying variation.

I therefore retain transactions only when the absorbed and acquiring charters belong to different top-tier holding companies before the transaction. I establish pre-merger ownership using SOD records in the six years preceding the FDIC event. This restriction removes 18,707 charter-level events, or 80.3 percent of the raw institution-history file, and leaves 4,592 changes in ownership between banking organizations.³²

Among these transactions, I retain unassisted mergers, absorptions, and consolidations. I exclude transactions associated with bank failures or FDIC assistance because their timing may respond directly to local economic conditions. I also exclude affiliated-institution poolings, which generally represent corporate reorganizations rather than acquisitions generating new exposure.

Merger Timing and Pre-Merger Deposits. The FDIC effective date refers to the legal merger of the underlying charter, which can occur after ownership has already consolidated at the holding-company level. Because the relevant CRA exposure changes when the banking organizations consolidate, I use the SOD to identify the first year in which the target and acquirer appear under the same top-tier holding company. For each transaction, I then measure the target's county-level deposits using the last SOD observation before this consolidation.

Using pre-merger deposits fixes the size and geographic distribution of the acquired deposit shock before the acquirer's subsequent decisions. Post-merger branch closures, sales, or deposit reallocations are therefore not incorporated into the exposure measure.

Because the inherited amount is a nominal figure recorded in the last SOD vintage before the consolidation, I deflate it to 2024 dollars at that vintage year rather than at the year of the observation it is attached to. The exposure is then constant within a bank-county pair, and is absorbed by the bank-county effects, so that only its interaction with the post-merger indicator is estimated. Deflating at the observation year instead makes the same nominal amount drift by roughly 23 percent across a six-year event window, and the drift is monotone in event time, so the interaction picks up the price level as well as the merger.

When an acquirer completes multiple transactions in the same year, I sum the inherited de-

³²This restriction also removes institutions that never appear in the SOD, such as some trust and credit-card banks, because their pre-merger local deposit footprint cannot be measured.

posits across targets within each county. This procedure yields 20,989 acquirer $\times$ year $\times$ county observations associated with 3,934 merger events.

Constructing the Merger-Induced Deposit Shock. The empirical design focuses on mergers that increase an acquirer’s deposits within counties in which it was already operating. I therefore retain only target deposits in counties where the acquiring banking organization had positive deposits before the merger. This restriction distinguishes an increase in exposure to an existing local market from geographic expansion into a new market, which could simultaneously alter both the bank’s CRA geography and its investment opportunities (Nguyen, 2019).

I additionally require the pre-merger overlap between the target and acquirer to span at least two states. The resulting variation therefore comes from the geographic distribution of deposits acquired across the acquirer’s existing footprint rather than from a transaction concentrated in a single local market.

Table E2 summarizes these restrictions. The resulting merger-shock sample contains 1,850 acquirer–year–county observations associated with 171 merger events, 114 acquirers, and 780 counties. The median event transfers \$1.1 billion of deposits, with an interquartile range of \$0.4 billion to \$3.6 billion. Where a bank–county pair experiences more than one qualifying merger I date the event from the first and censor the panel from the second onward, so that no post-merger window contains a second, separately identified shock. After restricting to banks observed in the Part 24 investment data and bank–county pairs entering the estimation sample, the identifying variation comes from 27 merger events involving 14 acquiring banking organizations and 200 bank–county pairs.


Table E2: Construction of the Merger Sample

	Observations	Events	Acquirers	Counties
All transactions with a measurable footprint	20,989	3,934	1,972	2,819
excluding affiliated poolings	20,555	3,831	1,937	2,806
excluding failures and assisted mergers	18,798	3,474	1,841	2,745
counties with a pre-merger overlap	4,785	1,879	1,096	1,359
overlap spanning at least two states	1,850	171	114	780

Notes: FDIC institution-history records matched to the Summary of Deposits. An observation is an acquirer $\times$ year $\times$ county and an event is an acquirer $\times$ year. The starting row already excludes charter consolidations within the same top-tier holding company, in which no deposits change hands between institutions, and charters that never appear in the Summary of Deposits, for which a pre-merger footprint cannot be measured.

Timing of the Merger Events. Figure E1 plots the merger events by the year in which the deposits changed hands. Qualifying mergers occur steadily across the sample, between two and twelve a year, with a trough in 2009 and 2010, when merger activity outside failure resolution largely stopped.

Figure E1: Merger Events by Year



Notes: Light bars are merger events passing the restrictions in Table E2: arms-length, unassisted, an overlapping pre-merger deposit footprint, and overlap spanning at least two states. Dark bars are the subset that enters the estimation sample: the acquirer files Part 24 investments, and the treated bank–county pair holds deposits in the county and survives the county $\times$ year effects. The figure covers the 2003–2022 panel window with 150 qualifying events. An event is an acquirer $\times$ year, so an acquirer closing several transactions in one year counts once.