Rationed, Not Repriced: Mortgage Access and House Prices*

Wenchuan Zhao† Masashi Takahashi‡

November 16, 2026

Abstract

Government support for mortgage credit often results in higher house prices rather than in more homeowners. We study the Duty to Serve mandate, which requires the government-sponsored enterprises (GSEs) to buy more mortgages in eligible rural areas, a target they meet by making qualifying loans more profitable for lenders to originate. Combining spatial difference-in-differences and structural estimation, we show that lenders keep the higher margin instead of lowering borrowers' interest rates, and use it to approve creditworthy applicants they had previously denied. The concession costs only a few basis points per loan, while the approval is worth several thousand dollars a year to the borrower who receives it. In equilibrium, marginal homebuyers keep the surplus and incumbent owners gain little because of elastic housing supply.

JEL classification: G21, G28, G51, H22, R31, R58

Keywords: Credit Rationing, Mortgage Policy, Housing Supply Elasticity, Welfare Incidence, Spatial Equilibrium

*We thank our advisors Martin Schmalz, Joel Shapiro, Shumiao Ouyang, Brent Ambrose, and Jiro Yoshida for their continued support. We thank Milena Almagro, Francisco Amaral, Elliot Anenberg, Arun Barman, Eirik Brandsaas, Yann Joël Cerasi (discussant), Konhee Chang, Justin Contat, William Doerner, Christian Hilber, Tran Huynh (discussant), Michael LaCour-Little, Raven Molloy, Bernardo Morais, Nuno Mota, Ken Okamura, Daniel Ringo, Tess Scharlemann, Stijn Van Nieuwerburgh, Ansgar Walther, and seminar participants at the Federal Housing Finance Agency (FHFA), the University of Reading, and the University of Oxford for helpful comments. We also thank participants at the European Meeting of the Econometric Society (ESEM), the FMA European Conference, the Freddie Mac Rural Housing Symposium, the Lugano Real Estate and Urban Economics Conference, the UEA European Meeting, and the UEA Summer School. The paper was previously circulated as "Mortgage without Capitalization: Evidence from GSE Lending."

†Saïd Business School, University of Oxford; email: wenchuan.zhao.dphil@said.oxford.edu

‡Smeal College of Business, Pennsylvania State University; email: mft5548@psu.edu

1 Introduction

Housing affordability has become a national concern that extends well beyond large cities. In 2026, Congress passed the first major federal housing legislation in decades in response to this concern (United States Congress, 2026). Because housing supply is relatively fixed in the short run, the traditional way to make housing affordable runs through the mortgage market, such as cheaper rates, public guarantees and waived fees. Cheaper credit, however, raises every buyer's bid at once, so its gains surface in house prices and accrue to the people who already own (Hilber and Turner, 2014; Carozzi et al., 2024; Greenwald and Guren, 2025). Why does support that travels through mortgages so rarely reach the households it targets? This paper argues that credit policy need not make loans cheaper to take effect. When loans are rationed at posted terms, the marginal applicant is creditworthy at the price everyone else pays but falls short of the lender's required margin by a small amount, so a concession that lenders keep as margin still changes which applications are worth approving, and transmission runs through the extensive margin of credit without repricing anyone. Whether the gain stays with the new borrowers depends on whether the housing supply can meet the demand they bring.

We study the Duty to Serve (DTS) program, a concession of this kind. Implemented by final rule in December 2016, DTS directs the government-sponsored enterprises (GSEs), Fannie Mae and Freddie Mac, to expand purchases of mortgages in underserved rural tracts against annual loan-count targets. The GSEs meet those targets by improving what a qualifying loan is worth to the lender who delivers it, through purchase commitments, execution options, and pricing at delivery, rather than by changing the terms borrowers see.[1] Eligibility follows Census delineations that predate the program, drawing a sharp tract-level boundary. The rural markets at the boundary also let us observe the housing-market adjustment: land is abundant and building lightly regulated, so supply can respond if demand arrives (Saiz, 2010; Gyourko et al., 2021).

[1] Section 2 describes the policy in detail.

Using a spatial difference-in-differences design, we find that the concession takes effect without touching the price of credit. The GSEs shift their purchases toward small-balance loans, and conventional approval rates rise by 1.35 percentage points, with the approval gain largest for applicants just below the debt-to-income limit. Contract rates do not move, implying that lenders keep the execution premium as margin, and the concession reaches borrowers as approvals. The marginal borrowers are low- and moderate-income owner-occupants who take smaller loans relative to income, and delinquency stays flat through the COVID period. Together, these patterns say the concession admitted applicants who were creditworthy at posted terms and lacked only a lender for whom the loan was worth making, not that underwriting loosened across the board.

While the policy delivers benefits to the new borrowers, the housing market determines how much of the gain they retain. When supply is inelastic, house prices rise to clear the new demand, and the gain accrues to incumbent homeowners. In the rural setting of this policy, the opposite holds. The stock of owner-occupied units grows by about 4%, through new construction and conversion of rental units; house prices do not move, and rents even fall by roughly 2.4% as renters shift into ownership.

The reduced form cannot value the access: contract terms do not move, so no price exists from which to infer what a loan is worth to an applicant who used to be denied. Nor can it divide the gain across households, since welfare is an equilibrium object. And it cannot say how the same concession would behave in other housing markets, or how a different instrument, such as cheaper credit, would fare in this rural market.

We therefore build a structural model of the two decisions that produced these facts. On the household side, workers who differ in credit access choose jointly where to live, where to work, and whether to own or rent, in a nested discrete-choice system (McFadden, 1978; Berry et al., 1995). On the lender side, a profit-maximizing lender approves an application when its expected return is positive, and DTS raises what a qualifying rural loan is worth to the lender, which enters as a lower effective funding cost and relaxes the approval threshold

for marginal applicants. The two sides meet in the cost of owning a home: the approval a lender grants reaches the household as a reduction in that cost, the object the model prices. Housing supply responds to prices with an elasticity that we allow to differ across the eligibility boundary, so how the demand shock splits between quantities and prices depends on where it lands. We estimate the two sides jointly by simulated method of moments (SMM), matching the model to the difference-in-differences responses above.

We first use the model to price the rationing relief for households. For a credit-constrained buyer whose origination the policy induces, the approval is worth more than a third of the annual cost of owning the home, while the execution wedge that produces the approval is worth a few tenths of one percent of that same cost. The gap means that mortgage credit in these areas is severely rationed, and it identifies the margin that matters: a policy that merely lowered mortgage rates could transmit at most the wedge itself.

Next, the model divides the relief across households. The gains concentrate on the credit-constrained buyers the policy targets, about $2,800 per subsidized origination in the year of purchase. Renters gain modestly from lower rents and the new option to buy, while incumbent owners are roughly unaffected, because elastic supply adds units instead of bidding up the homes they own. Two counterfactuals then vary how the support is delivered and where it lands. Delivered through contract rates at the same concession per loan, the policy is nearly inert: a rate cut must pay every borrower to reach the few near the threshold, and it does not reach the applicant whose constraint is the approval itself. Delivered into a supply-constrained urban market, the same concession capitalizes into prices: the capitalized share rises from under 2% to 17%, the marginal buyer's gain erodes, and the gains rotate from new buyers to incumbent owners.

Our first contribution is to the literature on how credit supply passes through lenders and into housing markets. On the price margin, pass-through of the GSEs' funding advantage to contract rates is modest and incomplete, with originators retaining much of the advantage as margin (Passmore, 2005; Sherlund, 2008; Lehnert et al., 2008; Scharfstein and

Sunderam, 2016; Fuster et al., 2024). On the quantity margin, which applications lenders approve responds to funding and securitization conditions (Loutskina and Strahan, 2009; Keys et al., 2010; Puri et al., 2011; Jiménez et al., 2012; Fuster and Vickery, 2015), and funding shocks in consumer credit pass through as limits rather than prices (Agarwal et al., 2018). Downstream, credit expansions capitalize into house prices where supply or the rental margin cannot absorb them (Mian and Sufi, 2009; Favara and Imbs, 2015; Adelino et al., 2012; Landvoigt et al., 2015; Di Maggio and Kermani, 2017; Orlando and Welke, 2025), propagate into household leverage, and end in distress (Mian and Sufi, 2011; Mian et al., 2017); Greenwald and Guren (2025) trace the division between prices and homeownership to the margin that absorbs the demand. We document a new regime in which a lender-side concession with no rate pass-through transmits through approvals, and an elastic-supply market absorbs the shock on the quantity margin at stable prices and with no detectable deterioration in credit quality. A credit expansion thus has two regime variables: transmission at the lender's approval margin, which decides whether a concession becomes access, and the elasticity of housing supply, which decides whether access shows up as homeownership or as prices, and to whom the gains accrue.

Our second contribution is an estimate of the shadow price of mortgage credit rationing, the user-cost reduction per origination that the policy induces, from policy variation that moves the approval margin alone. Lenders ration credit in equilibrium, denying loans at posted terms, whether because a higher rate would worsen the applicant pool (Stiglitz and Weiss, 1981) or because underwriting thresholds bind at the lender's margin (Duca and Rosenthal, 1991; Ambrose et al., 2002). The DTS margin is of the second kind, which is why relieving it requires no repricing. Moreover, our result is consistent with the results that access to credit matters more for housing demand than its price (Acolin et al., 2016; DeFusco and Paciorek, 2017; Gete and Reher, 2018; Anenberg et al., 2019; Fuster and Zafar, 2021), and regulating the same underwriting margin moves quantities, not rates (DeFusco et al., 2020). Yet the welfare cost of rationing has been quantified only through calibrated models,

for want of policy variation at that margin (Campbell and Cocco, 2015; Greenwald, 2018; Kaplan et al., 2020). The policy record reflects the same gap: guarantees expanded at the conforming limit buy little homeownership (Grundl and Kim, 2021), the Affordable Housing Goals left targeted ownership unchanged (Bostic and Gabriel, 2006; An et al., 2007; Bhutta, 2012), the Community Reinvestment Act raised lending to lower-income borrowers without raising income-conditional default while its broader effects remain contested (Ringo, 2023; Bhutta, 2011; Agarwal et al., 2012; Saadi, 2020; Brevoort, 2024; Conway et al., 2023), and demand-side subsidies capitalize into prices precisely where housing supply constraints bind (Hilber and Turner, 2014; Carozzi et al., 2024; Sommer and Sullivan, 2018), the constraints a large literature identifies as the root of the affordability problem (Quigley and Raphael, 2004; Gyourko and Molloy, 2015; Glaeser and Gyourko, 2018; Molloy, 2020; Baum-Snow and Duranton, 2025). The DTS boundary provides the missing variation, and the joint lender-household model converts it into a price, roughly two orders of magnitude above the execution wedge that produces the relief, and into an incidence, the household-level counterpart of aggregate welfare analyses of GSE guarantees (Elenev et al., 2016; Jeske et al., 2013; Gete and Zecchetto, 2018). Mid-century access expansions did build ownership (Fetter, 2013); the practical lesson of our estimates is to condition both the design and the expected incidence of credit policy on the local supply environment.

2 The Duty to Serve Program

Congress created the duty in 1992,[2] and the Federal Housing Finance Agency (FHFA) implemented it by final rule on December 13, 2016 (Federal Housing Finance Agency, 2016). The rule requires the government-sponsored enterprises (GSEs), Fannie Mae and Freddie Mac,

[2]Section 1335 of the Federal Housing Enterprises Financial Safety and Soundness Act of 1992 directed the GSEs to provide leadership in facilitating a secondary market for mortgages serving very low-, low-, and moderate-income families in underserved markets. The provision lay largely dormant until the Housing and Economic Recovery Act of 2008, which established FHFA and gave it authority to issue implementing regulations. The duty is statutory rather than an artifact of conservatorship, so the GSEs would remain subject to it after an exit from conservatorship. Appendix A gives the fuller history.

to serve three underserved markets: manufactured housing, affordable housing preservation, and rural housing, the market we study. Each Enterprise operates under a three-year Underserved Markets Plan, approved by FHFA, with the first plans taking effect on January 1, 2018.[3] For the rural market, the plans commit the GSEs to purchase targets for qualifying loans: loans in high-needs rural regions, loans from small rural lenders and community development financial institutions, and products for high-needs rural populations. Targets are measured in loan counts rather than dollar volumes, so the yardstick is the number of borrowers served. FHFA evaluates performance annually and publishes a rating,[4] and the duty binds the Enterprises alone: lenders face no obligation under the program, so any change in lender behavior must come from a change in the economics of lending.

The GSEs do not originate loans, so meeting a count target requires making qualifying loans worth originating and delivering. In our sample period their instrument is purchase execution rather than posted pricing. The plans commit the Enterprises to expand the lenders they buy from and to purchase through cash, flow, bulk, and seasoned-loan channels (Freddie Mac, 2017; Fannie Mae, 2019b), and the Enterprises' reports to FHFA show the lever in use: in 2018, short of its target for rural loans from small financial institutions, Fannie Mae reported that it would "continue to execute an outreach and pricing strategy with our small financial institution lenders, resulting in an increase in loan purchase volume" (Fannie Mae, 2018b).[5] Alongside execution, the Enterprises lowered the cost of originating a qualifying loan. Fannie Mae introduced an appraisal waiver for high-needs rural purchase loans in December 2018 (Fannie Mae, 2018a); Freddie Mac clarified rural property eligibility, income

[3]FHFA reviews and approves each plan before its effective date, and proposed amendments require FHFA approval. The one-year gap between the rule and the plans' effective date gave market participants certainty about the geographic definitions and the targeting structure. Our account of program operation draws on program documents and discussions with industry participants.

[4]The evaluation combines a quantitative assessment of the plan objectives, a qualitative assessment of activities, and potential extra credit for certain impactful activities, and culminates in a public rating on a five-point scale from "Fails" to "Exceeds", disclosed in FHFA's annual report to Congress. Non-compliance triggers an FHFA notice and a remediation requirement; there are no automatic financial penalties. The regime resembles supervisory bank-examination ratings rather than litigation-based enforcement.

[5]Industry participants describe qualifying rural small-balance loans as purchased above par, partly to meet the targets and partly because small loans prepay slowly. Program documents do not disclose the size of the premium.

calculation, and valuation requirements in Guide Bulletins 2019-20 and 2019-22 and trained nearly 1,000 lenders in its rural property flexibilities (Freddie Mac, 2019); and both delivered down-payment and closing-cost assistance through lenders, Fannie Mae to nineteen small rural lenders in 2019 and Freddie Mac through its BorrowSmart program with 27 lenders (Fannie Mae, 2019a; Freddie Mac, 2019).[6] Nothing in the concession sets the borrower's rate or fees; it changes what a qualifying loan is worth to the lender that originates and delivers it.[7]

Eligibility, by contrast, is mechanical, and only the rural market carries a precise geographic definition. A census tract qualifies as rural if it lies outside a Metropolitan Statistical Area, or within one but outside the Census Bureau's urbanized areas, and with a housing density at or below 64 housing units per square mile. These delineations exist independently of the program and predate it, so the formula draws a sharp tract-level eligibility boundary.

3 Data

3.1 Mortgage Data

We use HMDA loan-level data as the primary source. Reporting requirements changed in 2018, altering variable definitions and formats. We harmonize the pre- and post-2018 files and retain variables on geography, loan characteristics, property attributes, and borrower demographics. We restrict the sample to originated, conventional, first-lien, single-family mortgages that were purchased by Fannie Mae or Freddie Mac, including both purchase

[6] None of the 2018–2021 plan objectives for rural single-family lending contains a fee waiver or guarantee-fee change. The objectives are loan-count targets for high-needs rural regions and for loans from small financial institutions, evaluated annually by FHFA. Appendix A describes the non-pricing channels in more detail.

[7] The GSEs' posted pricing instrument is the loan-level price adjustment (LLPA), an upfront fee in percent of the loan amount charged on delivery according to credit score, loan-to-value ratio, and loan features; waiving it raises the price the lender receives by the same amount. FHFA eliminated LLPAs on single-family loans supporting Duty to Serve on October 24, 2022, effective for deliveries settling on or after December 1, 2022 (Federal Housing Finance Agency, 2022; Freddie Mac, 2022), after our 2013–2021 estimation window. Appendix Figure A1 reproduces the waiver table from the current matrix (Fannie Mae, 2026). The waiver is not part of the variation we use.

and refinance loans. Panel A of Table 1 reports summary statistics. Monetary variables are deflated to 2010 dollars using the GDP deflator and winsorized at the 1st and 99th percentiles.

We complement HMDA with the Black Knight McDash Mortgage Performance dataset for more detailed loan-level information on active and terminated mortgages, including loan amount, interest rate, LTV, credit score, ZIP code, and default outcomes. Panel B of Table 1 reports summary statistics for conventional loans within five miles of the treatment boundary. The average loan is a 30-year fixed-rate mortgage with a balance of roughly \$315K, a 3.8% interest rate, a 72% LTV ratio, and a credit score near 757. McDash identifies location at the ZIP-code level only, so we define treatment exposure continuously using ZIP–tract crosswalks rather than tract-level eligibility.[8] The McDash sample begins in 2015 rather than 2013 as elsewhere in the paper. Coverage of exposed ZIP codes differs in those two years. A single acquisition channel, wholesale lending, supplies 60 percent of the loans there in 2013 and 30 percent in 2014, against about 6 percent in 2015–2016 and under 2 percent afterward, and its share in unexposed ZIP codes never exceeds about 6 percent. The 2013 cohort is also more seasoned at entry: 61 percent of its loans board within twelve months of origination, against about 70 percent in each of the next three cohorts. Estimates for those two years therefore reflect which servicers report, unevenly across the two sides of the boundary. Throughout the McDash analysis we keep only loans that board within twelve months of origination. Loan amount, property value, and the interest rate are winsorized at the 1st and 99th percentiles. The credit score is not, and the regressions use the loan-to-value ratio as recorded: winsorizing the score at the 1st percentile censors the left tail of the score distribution, which is where the composition results lie.

[8] In most cases, a ZIP code covers multiple census tracts. We weight each census tract by its share of the ZIP code's residential addresses to compute the treatment exposure of the ZIP code. In our sample, 34% of ZIP codes have no exposure at all, which form a clean control group.

Table 1: **Summary Statistics – Mortgage**

	Mean	Median	SD	Min	Max
Panel A: HMDA Data					
Log Real Loan Amount	12.14	12.17	0.46	11.14	12.96
Log Real Income	11.32	11.33	0.47	10.33	12.32
Log Real Property Value	12.62	12.62	0.45	11.67	13.53
Loan-to-Value Ratio	71.69	74.47	14.86	36.23	95.00
Debt-to-Income Ratio	0.37	0.41	0.13	0.10	0.49
Rate Spread	0.28	0.24	0.32	−0.28	1.17
Panel B: McDash Data					
Loan Term (years)	28.73	30	4.29	5	30
Loan Amount ($K)	314.84	287	187.78	5	726
Fixed-Rate (%)	88.41				
Interest Rate (%)	3.77	3.75	0.88	1.75	8
Loan-to-Value Ratio (%)	71.91	78.87	23.35	3.38	100
Front-End Debt-to-Income Ratio (%)	31.90	33	9.73	7	50
Credit Score	757.42	770	45.17	300	850
Share of loans ever reaching (%):					
30 Days Delinquent	7.35				
60 Days Delinquent	3.33				
90 Days Delinquent	2.48				
Foreclosure	0.20				
Bankruptcy	0.21				
Loans	205,544				
ZIP Codes	3,662				

Notes: Panel A reports summary statistics for the HMDA loan-level data, restricted to originated, conventional, first-lien, single-family loans purchased by GSEs over the 2013–2021 period. Panel B reports summary statistics for the McDash loan-level data: conventional purchase loans originated over 2015–2021 in ZIP codes within five miles of the eligibility boundary and boarded within twelve months of origination. The delinquency, foreclosure, and bankruptcy rows report the share of loans that ever reach that status in the servicing record, so later cohorts are observed for a shorter period. Panel A variables, and the loan amount, interest rate, loan-to-value ratio, and debt-to-income ratio in Panel B, are winsorized at the 1st and 99th percentiles; the credit score is not.

3.2 Housing Market

We use CoreLogic data on U.S. real estate transactions, mortgage originations, and tax assessor records from 2000 to 2022. The analysis relies on two linked files identified by a unique property identifier (CLIP). The Property database contains property characteristics and assessed values, and the Owner Transfer database records transaction prices, sale dates, and buyer characteristics.

The data distinguish investor, corporate, and owner-occupant purchases, as well as mortgage-financed versus cash transactions and new construction versus resale properties.[9] We use repeated sales of the same property to study housing-price dynamics. Panel A of Table 2 reports summary statistics for the repeated-sales sample over 2013–2021. The average transaction is a roughly $300K property with six rooms and a building age of about 40 years. Variables are winsorized at the 1st and 99th percentiles.

We also aggregate transactions to the tract-year level to measure local housing-market activity. These measures include transaction volume and ownership composition within each tract. Panel B of Table 2 reports tract-level statistics for areas within five miles of the treatment boundary, separately for rural and non-rural tracts. Transaction volumes are similar across the boundary, but non-rural tracts have substantially higher average transaction prices.

3.3 Census Characteristics

We use the American Community Survey (ACS) five-year files from the National Historical Geographic Information System (NHGIS) to measure labor and housing market conditions at the census-tract level. Because ACS five-year files pool observations over rolling windows, we assign the midpoint year of each window as the reference period, yielding coverage from 2007 to 2021. We harmonize all observations to 2010 census tract boundaries to construct a

[9] Buyer and seller names are anonymized in our extract, so individual buyers cannot be linked across transactions.

Table 2: **Summary Statistics: Housing Market (CoreLogic)**

Variable	Mean	Median	SD	Min	Max
Panel A: Transaction-Level Repeated-Sales Sample					
Transaction Price ($K)	292.79	192.07	375.28	10.99	2737.6
Building Size (sqft)	1840.27	1624	896.56	616	5529
Number of Rooms	6.41	6	2	3	14
Property Age (years)	42.49	35	31.94	1	125
Rural ({0,1})	0.11	0	0.31	0	1
Distance to Boundary (miles)	3.59	3.82	1.07	0.1	5
Observations	2,557,701				
Unique Properties	1,216,104				
Panel B: Tract-Level Transaction Measures					
Non-rural Tracts					
# Transactions	57	39	64	0	498
# Owner-occupied Transactions	44	28	54	0	417
Average Price ($K)	486	313	533	22	3,286
Observations	134,090				
Rural Tracts					
# Transactions	53	38	60	0	345
# Owner-occupied Transactions	43	31	49	0	256
Average Price ($K)	280	219	214	27	1,350
Observations	21,208				

Notes: Panel A reports summary statistics for transaction-level CoreLogic data, restricted to properties with repeated sales observed during the 2013–2021 period. Panel B reports summary statistics for tract-level housing transaction measures constructed from CoreLogic over the 2013–2021 period. Variables are winsorized at the 1st and 99th percentiles.

consistent panel of tract-level demographic, housing, and employment characteristics. Appendix Table A2 reports summary statistics.

In the structural model, households decide where to live and work, and therefore we also use commuting and migration data. At the aggregate level, the Longitudinal Employer-Household Dynamics Origin-Destination Employment Statistics (LODES) from the U.S. Census Bureau helps us to construct tract-to-tract commuting flows that enter the structural model in Section 5. At the individual level, we use Infutor (previously Verisk) to distinguish homeowners from renters. Appendix C describes the construction procedure.

We classify workers as high-skill if they hold a bachelor's degree or higher. We use education rather than observed credit access because education is predetermined and unaffected by the DTS policy. Credit constraints are endogenous to treatment: DTS relaxes borrowing constraints for some households, changing the composition of the credit-constrained group after treatment. In pre-treatment data, low-skill tracts report lower applicant incomes, smaller loan amounts, and higher mortgage denial rates. Their loan pool also contains a larger share of small-balance and FHA loans than the pool of high-skill tracts (Table A3). Individual records from Infutor show a strong correlation between educational attainment and income (Table A4). We interpret skill heterogeneity in Section 5 as capturing differences in credit access rather than education itself.

4 Reduced-form Empirical Results

This section presents the reduced-form evidence. We first describe the spatial difference-in-differences design built on the DTS eligibility boundary, and then follow the concession from the GSEs to lenders, to borrowers, and into the housing market.

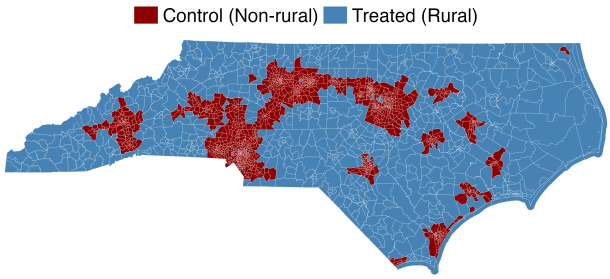
4.1 Identification Strategy

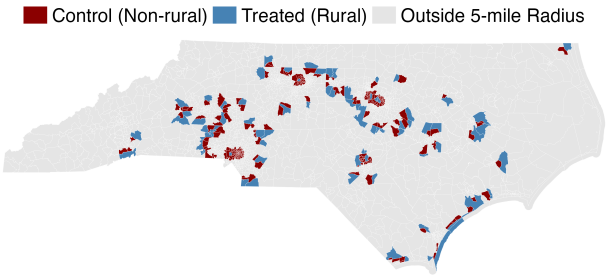
The primary identification challenge is that the rural tracts targeted by the program differ systematically from non-rural tracts, so a simple comparison would conflate the program's effects with pre-existing economic and demographic disparities. We address this selection problem with the program's own eligibility formula. Because eligibility follows Census delineations, the formula draws a boundary between tracts that face the same local economy but different program status, and we restrict the sample to its neighborhood. The treatment group consists of rural tracts within five miles of the boundary, and the control group comprises non-rural tracts within the same distance. All specifications compare treated and control tracts that share the same boundary segment within the same year, through boundary-by-year fixed effects, so unobserved local shocks common to a boundary neighborhood do not confound the estimates. Figure 1 illustrates the design for three states: North Carolina, Ohio, and Pennsylvania. The left panel maps rural and non-rural areas, and the right panel shows the treated and control tracts that remain after excluding those more than five miles from the boundary.

The treatment period begins on December 13, 2016, when the final DTS rule was published. Although the program officially became effective on January 1, 2018, we date treatment from the announcement so that any anticipation during 2017 counts as treatment and does not contaminate the pre-period. We estimate the following spatial DiD specification, restricting the sample to tracts located within five miles of the DTS eligibility boundary:

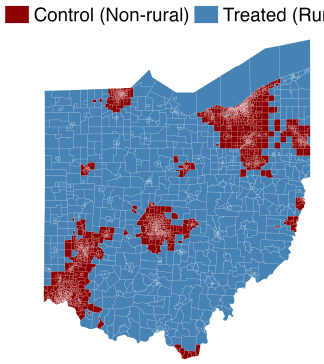
$$Y_{ct} = \beta\big(\text{Rural}_c \times \text{Post}_t\big) + X'_{ct}\delta + \gamma_c + \gamma_{b(c),t} + \varepsilon_{ct}, \tag{1}$$

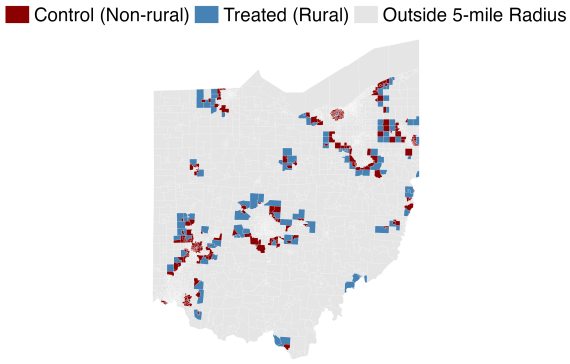
where Y_{ct} is the outcome for tract c in year t, Rural_c is an indicator equal to one if tract c lies within the rural areas, and Post_t is an indicator for the post-announcement period. The vector X_{ct} collects time-varying tract-level demographic controls, such as the racial and age distributions. The term γ_c denotes tract fixed effects that absorb time-invariant local



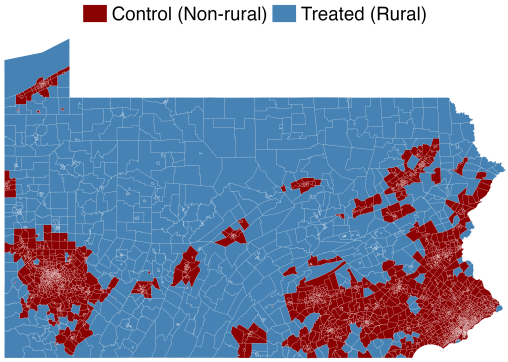


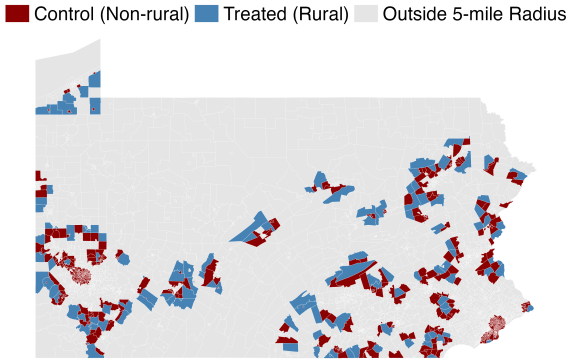
(a) North Carolina





(b) Ohio





(c) Pennsylvania

Figure 1: **Illustration of Identification Strategy**

Notes: These figures illustrate the spatial difference-in-differences design. The left panel for each state shows the complete division of census tracts into rural (treated) and non-rural (control) areas. The right panel zooms in on the tracts within a 5-mile radius of the rural/non-rural boundary, which constitute the sample for our analysis.

characteristics, and $\gamma_{b(c),t}$ represents boundary-by-year fixed effects that flexibly control for time-varying shocks common to tracts adjacent to the same eligibility boundary segment. Standard errors are clustered at the state level throughout. Identification of β therefore comes from within-boundary, within-year comparisons of treated and nearby control tracts over time.

Building on this baseline DiD framework, we additionally estimate event-study specifications to examine the dynamic effects of the policy:

$$Y_{ct} = \sum_{k \neq -1} \beta_k \big(\text{Rural}_c \times \mathbf{1}\{t - t_0 = k\} \big) + X_{ct}'\delta + \gamma_c + \gamma_{b(c),t} + \varepsilon_{ct}. \tag{2}$$

Here, t_0 is the year 2017 and k indexes event time. The coefficient for the year prior to the announcement is omitted as the reference period. For outcomes that are discrete or non-negative counts, including loan applications and originations by occupancy type, we estimate the same specifications by Poisson regression, so the reported coefficients are interpreted as log differences. In nearly all event-study figures, the pre-treatment coefficients are close to zero and statistically insignificant, supporting the parallel trends assumption underlying our identification strategy.

The post-treatment window overlaps with the COVID-19 pandemic, which could confound the estimates if it hit treated rural tracts and control tracts differently. Because the comparison is between adjacent tracts on the two sides of a boundary segment, the pandemic's local effects enter as a common shock that the boundary-by-year effects absorb. Excluding 2020 and after leaves all results qualitatively the same; Appendix Table A10 reports the main estimates on the 2013–2019 sample.

A second concern is the violation of the Stable Unit Treatment Value Assumption (SUTVA), as economic effects and migration could spill over from treated tracts to adjacent control tracts. We vary the bandwidth around the boundary (3-mile and 10-mile windows) and estimate "donut" specifications that exclude tracts immediately adjacent to the boundary (3–5-mile and 5–10-mile windows). Our coefficients remain robust under the alternative

specifications.[10] Our individual address histories also allow us to test the migration channel directly. The share of treated-tract in-movers originating on the control side of the boundary band is stable at 16% before and after the program, and a difference-in-differences on cross-boundary inflows is statistically zero as shown in Appendix Table A13.

4.2 Mortgage Credit Supply

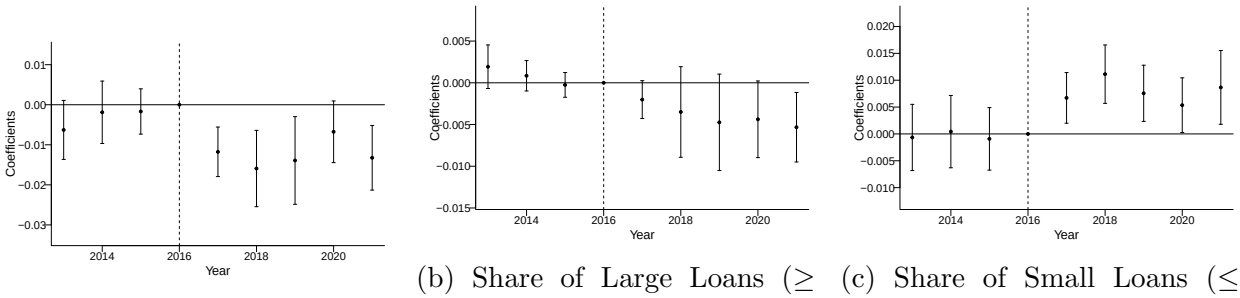
DTS commits the GSEs to purchase more mortgages in eligible rural tracts. This part verifies that the program reached this instrument, by tracing the volume and composition of the loans the GSEs purchased, and tests whether lending outside the GSE channel responded.

Figure 2 reports dynamic difference-in-differences estimates for the composition of GSE loan purchases in treated tracts.[11] GSE purchases in treated tracts rise by about 7% relative to control tracts, and their composition tilts toward smaller loans. Pre-treatment coefficients are flat in all three panels. After the mandate, the average size of purchased loans falls, as the share of large loans declines, though estimated less precisely, and the share of small loans rises steadily. In the static specification, average loan size falls by roughly 0.9% and the small-loan share rises by 0.8 percentage points, although the tilt is modest.

We also study the substitution from non-GSE into GSE lending. Table 3 Panel A shows that rejection rates among GSE-eligible applications fall in treated tracts by 1.3 percentage points, mirroring the approval response in Section 4.3. Panel B shows that the average size of non-GSE loans does not change overall and falls by 1.8% among loans below $100,000. Together, the program broadened GSE purchases of rural mortgages across the size distribution.

[10] Results are available upon request.

[11] This figure uses fixed-dollar thresholds ($100,000 and $500,000 in real 2010 dollars) so that the binary share series are consistent across years. The lender-margin analysis in Section 4.3 uses the FHFA Underserved Markets Plans bin definitions (DTS small-balance, conforming, jumbo) that match the policy's literal small-balance threshold.



(a) Average Loan Size

(b) Share of Large Loans (≥ $500k)

(c) Share of Small Loans (≤ $100k)

Figure 2: **Dynamic Effects on Mortgage Composition**

Notes: These figures present the dynamic difference-in-differences estimates of the effect of the Duty to Serve program on the composition of GSE-purchased mortgages. The sample includes all census tracts within a 5-mile radius of a rural/non-rural boundary. The vertical dashed line indicates the program's announcement year (2016). Each point represents the estimated coefficient for a given year relative to 2016, the announcement year (omitted), with 95% confidence intervals. The sample runs from 2013 to 2021 and ends before the December 2022 elimination of loan-level price adjustments on Duty to Serve loans. The figures show a significant decrease in the average loan size, driven by a decline in the share of large loans and a corresponding sharp increase in the share of small loans following the program's implementation.

Table 3: **Effects on Loan Rejection Rates and Non-GSE Loan Amounts**

	Panel A: Rejection Rates			Panel B: Non-GSE Amounts		
	All Data	Large Loan	Small Loan	All Data	Large Loan	Small Loan
	(1)	(2)	(3)	(4)	(5)	(6)
Rural × Post	-0.0134***	-0.0316	0.0002	-0.0026	-0.0185	-0.0178***
	(0.0022)	(0.0212)	(0.0031)	(0.0050)	(0.0137)	(0.0053)
Observations	7,468,826	357,950	1,084,393	4,688,994	791,679	723,044
R^2	0.08409	0.30724	0.09109	0.59787	0.36539	0.18065
Tract fixed effects	✓	✓	✓	✓	✓	✓
Boundary × Year fixed effects	✓	✓	✓			
Boundary × Year × Lender fixed effects				✓	✓	✓

Notes: This table combines results on loan rejection rates (Panel A) and non-GSE loan amounts (Panel B). Panel A estimates the treatment effect on loan application rejection (*is_rejected*) using a linear probability model on GSE-eligible HMDA applications. Panel B presents a placebo test using only non-GSE loan originations with the log of the real loan amount as the dependent variable. Within each panel, columns report the All-data sample, the Large-loan subsample (≥ $500,000), and the Small-loan subsample (≤ $100,000). Panel A uses tract and boundary × year fixed effects. Panel B uses tract and boundary × year × non-GSE channel fixed effects, where channel denotes the type of non-GSE financing (Affiliate, Bank, Credit Union/Finance, Other, Portfolio, Private Securitization). Standard errors are clustered at the state level and reported in parentheses. ***p<0.01, **p<0.05, *p<0.1.

4.3 Lender Response

Having established that the concession arrived, we ask what lenders did with it. A lender whose execution economics improve has two ways to respond: pass the gain to borrowers through the price of credit, or retain it as margin, which makes marginal applications profitable to approve and sell. The institutional design predicts the second, because GSE execution gives lenders a willing buyer for a previously thin segment. We test this prediction with a four-stage decomposition of the credit pipeline, from application through approval and origination to GSE sale, and then ask whether the concession reached the contract rate.

Table 4 reports difference-in-differences estimates for the four stages of the conventional first-lien single-family credit pipeline. The top-down analysis starts from total applications received, a demand-side count that lenders do not directly control. The remaining three stages are loan-level conditional probabilities that reflect lender discretion: approved given application, originated given approval, and GSE-sold given origination. Each stage is estimated for the full sample and within loan-size bins.[12]

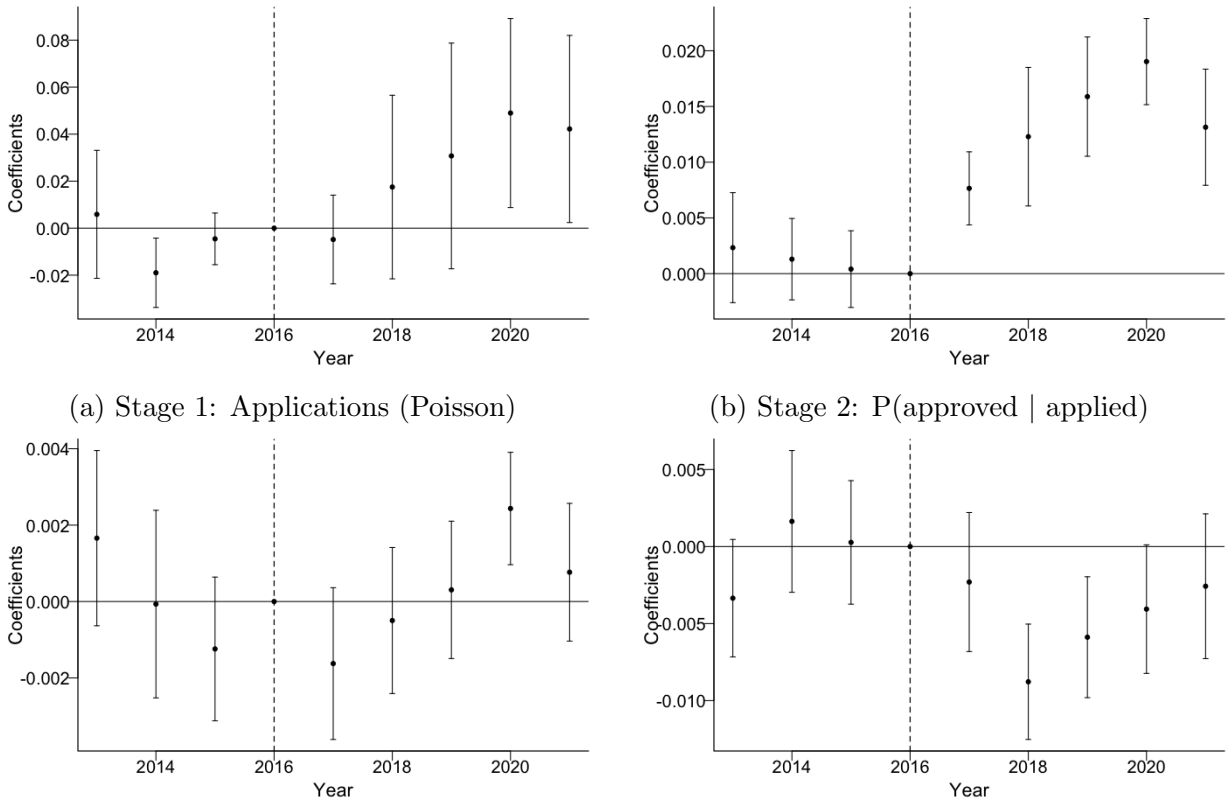
Among all stages, the second stage is the binding margin. Conventional approval rates rise by 1.35 percentage points on the full sample, with gains of 1.09 pp in the DTS small-balance segment and 1.92 pp in the conforming segment.[13] By contrast, applications do not surge, and origination conditional on approval does not change. The final stage shows a small decline in the share of originations sold to the GSEs within the year, but the decline is dilution. Also, the within-lender share routed to the GSE channel barely moves, so the mandate changes which loans are originated rather than how originations are routed.[14] Figure 3

[12] We define loan-size bins following the FHFA Underserved Markets Plans: *DTS small-balance* ($\leq$ \$200,000 for 2013–2020 and $\leq$ \$250,000 for 2021), *conforming* (above DTS small-balance and at or below the annual FHFA one-unit conforming loan limit), and *jumbo* (above the conforming limit).

[13] Approval also rises in the jumbo segment, but we do not read this estimate alongside the targeted segments. Loans above the standard conforming limit are mostly ineligible for GSE sale, although about a fifth of the bin's originations in treated tracts are sold to the GSEs under high-cost-area limits, and the estimate is the noisiest of the three bins. More importantly, jumbo applications rise by 10.5%, the only segment with a demand response, so the jumbo approval margin conflates a changing applicant pool with lender discretion.

[14] Appendix Table A5 splits the approval and routing responses by lender type over 2013–2017, when HMDA still reported agency codes: the approval gain is concentrated at independent mortgage companies

shows the corresponding event study for all four stages.



(a) Stage 1: Applications (Poisson)

(b) Stage 2: P(approved | applied)

(c) Stage 3: P(originated | approved)

(d) Stage 4: P(GSE-sold | originated)

Figure 3: **Dynamic Effects on Credit Outcomes across Lending Pipeline**

Notes: Each panel reports the event-study coefficients β_k from a separate difference-in-differences regression of the indicated funnel-stage outcome on Rural $\times$ $\mathbf{1}\{\text{year} = k\}$. Panel (a) is a Poisson regression on application counts at the (tract, year) level; Panels (b)–(d) are linear-probability regressions on the indicated conditional binary outcomes at the loan level. Panel (d) additionally includes lender (HMDA respondent) $\times$ year fixed effects. The sample includes HMDA conventional first-lien single-family applications in census tracts within five miles of the DTS rural eligibility boundary, 2013–2021. The vertical dashed line indicates the program's announcement year (2016), which is the omitted reference year. Each point represents the estimated coefficient for a given year relative to 2016, with 95% confidence intervals constructed from standard errors clustered at the state level.

Looking into denied applications with a recorded primary reason helps to understand the binding margin of mortgage approval. The composition tilts away from debt-to-income, as shown by Figure 4. The DTI share of recorded denials falls in 2017 and 2018 and fades by 2020, right after the announcement of the mandate, while the collateral and credit-history

and credit unions, the lenders whose originate-to-sell business is most exposed to GSE execution, and the decline in the GSE-sold share is confined to Freddie Mac deliveries and to purchase loans.

Table 4: **Effects on Credit Outcomes across Lending Pipeline**

	All	DTS small-balance (≤$200K)	Conforming	Jumbo
Panel A: Stage 1 – Poisson regression of application counts (tract, year)				
Rural × Post	0.0365	−0.0046	0.0284	0.1051**
	(0.0235)	(0.0195)	(0.0275)	(0.0522)
Tract FE	✓	✓	✓	✓
Boundary × Year FE	✓	✓	✓	✓
Pseudo-R^2	0.9264	0.8762	0.9133	0.8942
Observations	122,343	118,158	109,520	74,621
Panel B: Stage 2 – P(approved | applied), loan level				
Rural × Post	0.0135***	0.0109***	0.0192***	0.0164**
	(0.0020)	(0.0014)	(0.0032)	(0.0072)
Tract FE	✓	✓	✓	✓
Boundary × Year FE	✓	✓	✓	✓
R^2	0.0216	0.0287	0.0233	0.0218
Observations	15,570,522	6,447,642	6,945,935	2,176,945
Panel C: Stage 3 – P(originated | approved), loan level				
Rural × Post	0.0005	0.0006	0.0009	0.0003
	(0.0007)	(0.0007)	(0.0009)	(0.0031)
Tract FE	✓	✓	✓	✓
Boundary × Year FE	✓	✓	✓	✓
R^2	0.0107	0.0142	0.0125	0.0198
Observations	10,080,904	4,109,245	4,558,931	1,412,728
Panel D: Stage 4 – P(GSE-sold | originated), loan level				
Rural × Post	−0.0036***	−0.0017	−0.0020	0.0079
	(0.0013)	(0.0010)	(0.0019)	(0.0060)
Tract FE	✓	✓	✓	✓
Boundary × Year FE	✓	✓	✓	✓
Lender × Year FE	✓	✓	✓	✓
R^2	0.5137	0.5448	0.5574	0.5117
Observations	9,631,555	3,909,284	4,375,076	1,347,195

Notes: Each panel reports a separate difference-in-differences regression of the corresponding pipeline stage on Rural × Post. Sample: HMDA conventional first-lien single-family applications within 5 miles of the DTS rural eligibility boundary, 2013–2021. Stage 1 is a Poisson regression on application counts at the (tract, year) level; Stages 2–4 are linear-probability regressions at the loan level on the indicated conditional binary outcomes. Stage 4 adds lender (HMDA respondent) × year fixed effects to identify within-lender variation in GSE channel routing. The Pseudo-R^2 row in Panel A is the McFadden pseudo-R^2 for Poisson; the R^2 rows in Panels B–D include the variance explained by the fixed effects. The DTS small-balance threshold is $200,000 through 2020 and $250,000 in 2021. Standard errors are clustered at the state level. ***p<0.01, **p<0.05, *p<0.1.

shares do not move.[15] Better execution can relieve a debt-to-income constraint through loan size. Because loan principal is a choice variable while income and outside debt are fixed, a premium on delivery makes the smaller loan that brings an applicant below an unchanged DTI threshold profitable for the lender to originate and sell.

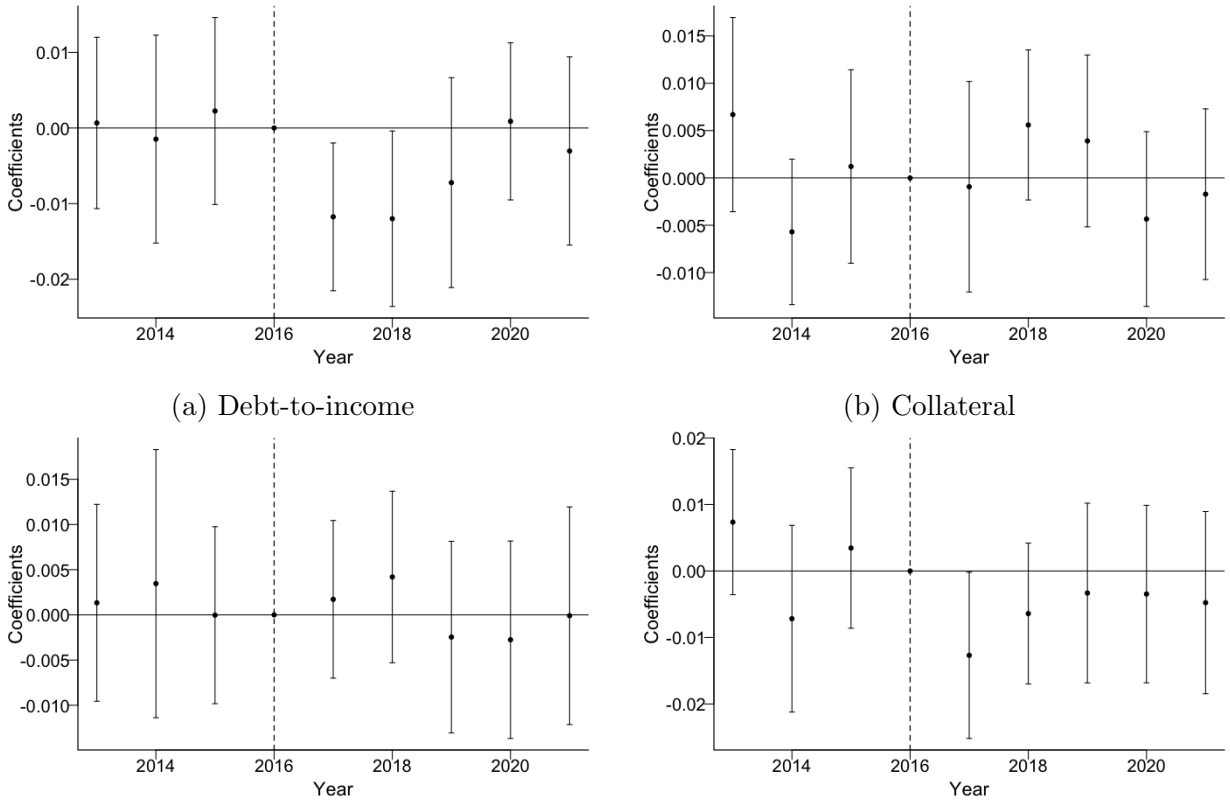
Panel A of Figure 5 locates the approval gain along the ratio, re-estimating the response within bins of the loan-to-income ratio at application, the loan-side component of the debt-to-income ratio that HMDA reports in every year. The response is an inverted U. It is smallest where the ratio is low and approval was never in doubt, rises through the range in which a typical applicant's debt-service ratio approaches the underwriting limit, peaks at loan-to-income ratios of 3.5 to 4, and is small and imprecise above 5, where reaching the limit would require cutting the loan by a fifth or more.[16] Binning instead on a pre-period predicted denial probability, a summary of everything HMDA reports about an applicant's risk of denial (loan amount, income, the loan-to-income ratio, purpose, and occupancy), produces no gradient at all (Panel b): the approval gain stays close to the pooled 1.2 percentage points in every decile of that risk. Applicants who are likely to be denied for any reason gain no more than applicants who are not, while applicants near the debt-to-income limit gain the most, so the relief follows the ratio, not overall riskiness.[17]

As the above test locates the response in quantities, the remaining question is whether the concession reached the contract rate. Figure 6 reports dynamic estimates for the note rate in

[15] Appendix Table A6 reports the within-denied composition DiD, 2013–2021. The composition is suggestive, for two reasons that Appendix Table A7 makes explicit. First, denial reasons were optional in HMDA before 2018 and mandatory after, and three-quarters of the decline in the denial rate per application falls on denials with no recorded reason. Second, in levels the decline is spread across debt-to-income, collateral, and credit-history denials, whether the unrecorded denials are dropped or allocated by each lender's recorded shares.

[16] Loan-to-income maps into the debt-to-income ratio through repayment terms and other obligations. At the 2017–2021 average 30-year rate of 3.7 percent, principal and interest are about 5.5 percent of the balance per year, taxes and insurance add roughly two percent, and non-housing obligations add eight to ten points of income for a typical applicant, so a loan-to-income ratio of 3.5 to 4 corresponds to a back-end ratio in the high 30s and a ratio above 5 to one near 50. The automated-underwriting limit was 45 percent until July 2017 and 50 percent thereafter.

[17] The pattern is not a change in who applies either. Appendix Table A8 shows that the pooled approval response falls only from 1.35 to between 1.1 and 1.2 percentage points when boundary-by-year effects are interacted with loan-to-income bins, with predicted-denial deciles, or with purpose, income-decile, and occupancy cells.



(a) Debt-to-income

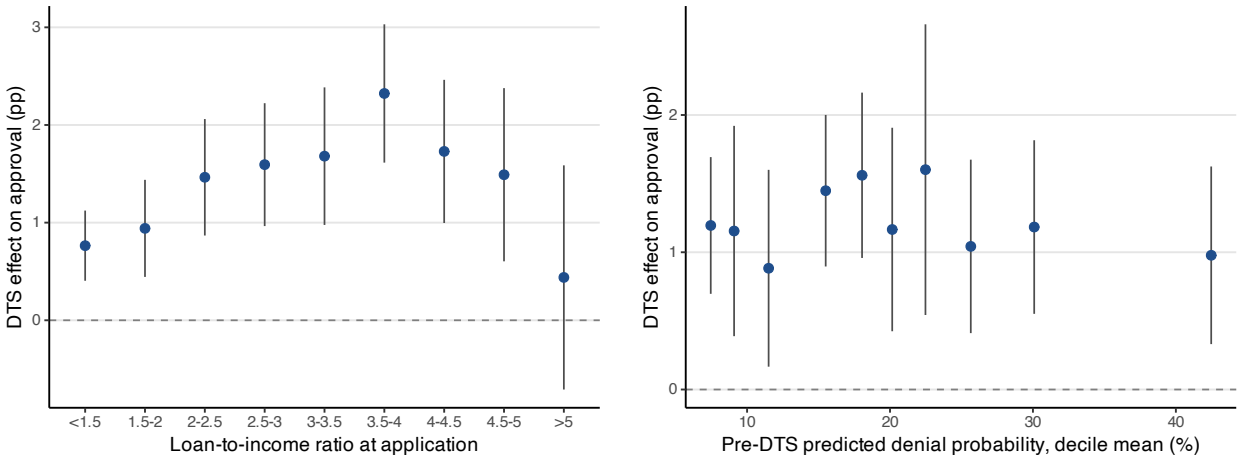
(b) Collateral

(c) Credit history

(d) Collateral or DTI

Figure 4: **Dynamic Effects on Denial Reason Composition**

Notes: Each panel reports the event-study coefficients β_k from a separate within-denied-application composition regression. The dependent variable is an indicator for the primary denial reason equaling the panel category (DTI, collateral, credit history, or the union of collateral and DTI). The sample includes denied applications with a recorded primary denial reason in HMDA conventional first-lien single-family applications in census tracts within five miles of the DTS rural eligibility boundary, 2013–2021. The vertical dashed line indicates the program's announcement year (2016), which is the omitted reference year. Each point represents the estimated coefficient for a given year relative to 2016, with 95% confidence intervals constructed from standard errors clustered at the state level. Pre-2018 denial-reason reporting is optional under Regulation C and post-2018 reporting is mandatory; the specification conditions on having a recorded reason, so it measures composition among recorded denials and is exposed to differential reporting across lenders and arms, which Appendix Table A7 addresses in levels.



(a) By loan-to-income ratio at application

(b) By pre-period predicted denial probability

Figure 5: **Approval Response by Loan-to-Income Ratio and Predicted Denial Probability**

Notes: The figure plots the treatment effect on loan approval by bin. Each point is the coefficient of Rural × Post from the Stage 2 approval specification of Table 4, estimated separately within the indicated bin, and the error bars represent 95% confidence intervals. The sample consists of HMDA conventional first-lien single-family applications with non-missing applicant income within five miles of the DTS rural eligibility boundary over 2013–2021. In Panel (a), applications are binned by the loan-to-income ratio at application. In Panel (b), applications are binned by deciles of a predicted denial probability, estimated by a logit of denial on within-year deciles of loan amount and income, the loan-to-income ratio and its square, a purchase indicator, and owner occupancy, using decided applications in 2013–2016; each point is placed at the decile mean. All specifications include census tract and boundary × year fixed effects. Standard errors are clustered at the state level. Appendix Table A8 reports the estimates, the pre-period denial rate in each bin, and the pooled coefficient within bins.

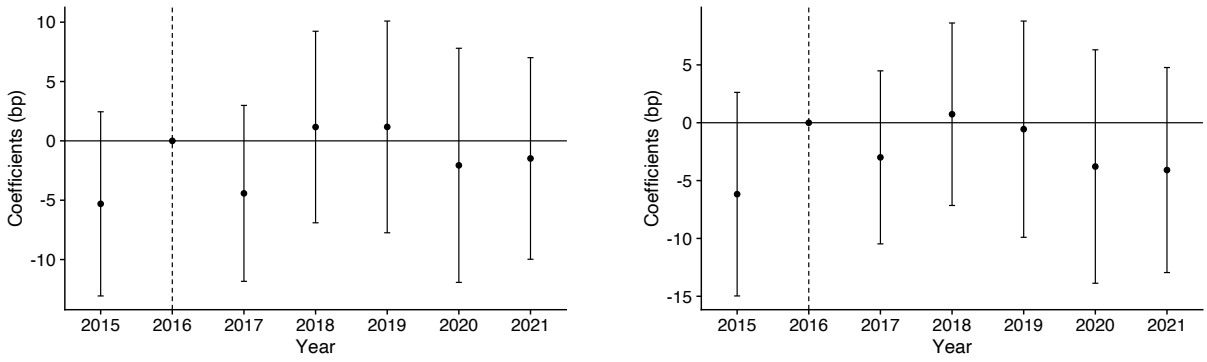
the McDash servicing data. Rates on loans originated in exposed ZIP codes do not move after 2016. No post-period coefficient in Panel (a) is distinguishable from zero, and the pooled estimate is 1.7 basis points at full exposure, with a 95% confidence interval from −4.7 to 8.1. Panel (b) holds the credit score, the loan-to-value ratio, the term, and mortgage insurance fixed and gives the same answer, a pooled estimate of 0.8 basis points.[18] The institutional record settles the channel. No posted price to the borrower changed during the estimation window (Section 2), the premium is received by the lender at delivery, and lenders retain it as margin, so its value surfaces as additional approvals rather than cheaper loans. The GSE literature documents the same retention, with modest or negligible consumer-rate benefits from GSE eligibility (Passmore, 2005; Sherlund, 2008; Lehnert et al., 2008) and substantial capture of the funding advantage by originators (Scharfstein and Sunderam, 2016; Fuster et al., 2024); weaker lender competition in rural markets plausibly reinforces it, though we do not identify market structure directly.

4.4 Borrower Response

The next empirical question is who received the access. Previous credit expansions reached investors and stretched borrowers and ended in leverage and default (Mian and Sufi, 2009, 2011; Mian et al., 2017), and speculative demand amplified both the boom and the bust (DeFusco et al., 2022). We find no support for the speculation reading along four dimensions, namely income composition, buyer type and leverage, risk at origination, and post-origination distress. On each margin, the marginal borrowers look like the statute's intended reach: low- and moderate-income owner-occupants entering with small loans, no leverage buildup, and delinquency in line with their observable risk.

Table 5 summarizes the borrower margins in three panels. Panel A shows that the approval response is broad across the income distribution; we sharpen this margin below

[18] The window contains a single pre-period year, whose coefficient is negative and not distinguishable from zero, so we read the panels as a bound on pass-through: the unconditional estimate rules out a decline in the contract rate larger than about five basis points.



(a) Unconditional

(b) Conditional on loan characteristics

Figure 6: **Dynamic Effects on Mortgage Interest Rates**

Notes: The figure plots dynamic difference-in-differences estimates for the note rate, in basis points, at the loan level. The sample consists of conventional, non-jumbo purchase loans in the McDash servicing panel within five miles of the DTS rural eligibility boundary, originated over 2015–2021 and boarded within twelve months of origination; 2013–2014 are excluded because panel entry breaks in those years. Exposure is defined as the share of a ZIP code's residential addresses in eligible tracts, fixed at its 2016 value, so coefficients are at full exposure. The coefficient of Exposure $\times$ year identifies the differential change in the note rate for a fully exposed ZIP code relative to the 2016 origination cohort, which is normalized to zero; the dashed line marks 2016. All specifications include ZIP code, state $\times$ year, boundary $\times$ year, acquisition channel, structure type, and loan type fixed effects and control for distance to the boundary and loan age at boarding. Panel (a) is unconditional on loan characteristics, so its coefficients include any change in the mix of loans originated. Panel (b) is conditional on the credit score, the loan-to-value ratio, the term, and mortgage insurance status, so its coefficients compare loans with the same observable risk. The error bars represent 95% confidence intervals. Standard errors are clustered at the state level.

using the statutory income bands the GSEs apply. Panel B runs Poisson regressions on tract-year origination counts split by occupancy status. Principal-residence counts rise by 5.2% in treated tracts, with no evidence of a disproportionate investor response, as the non-principal estimate is similar in size but too imprecise to interpret.[19] Panel C shows that the loan-to-income ratio (LTI) falls by 1.4 percent for the low- and middle-income terciles and is unchanged at the top, so new entrants buy smaller homes for a given income without levering up, a pattern the leverage measures below confirm.

Figure 7 replaces the sample terciles with the area-median-income (AMI) categories the GSEs apply under DTS: very low ($\leq$ 50% of AMI), low (50–80%), and moderate (80–100%). The approval response concentrates in the statutory target bands, rising by 1.8 percentage points for low-income and 2.1 percentage points for moderate-income applicants, against 1.3 percentage points for applicants above 100% of AMI, the band with the highest pre-period approval rate. The very-low-income band shows no approval gain. From a pre-period approval rate of 52%, these applicants remain rationed, consistent with an execution wedge of a few basis points relaxing the approval threshold only for applicants near it. Loan-size adjustment shows the mirror-image gradient, as LTI falls by 2.6 percent in the very-low band and 1.3 percent in the low band, and is statistically indistinguishable from zero at and above 80% of AMI. The mandate thus reaches the households the statute names, with credit expanding for low- and moderate-income families while the very-low-income segment adjusts on loan size instead of approval.[20]

Correspondingly, unlike the canonical credit expansion, leverage falls rather than rises. Table 6 reports DiD estimates on three LTI measures for the GSE-purchased conventional

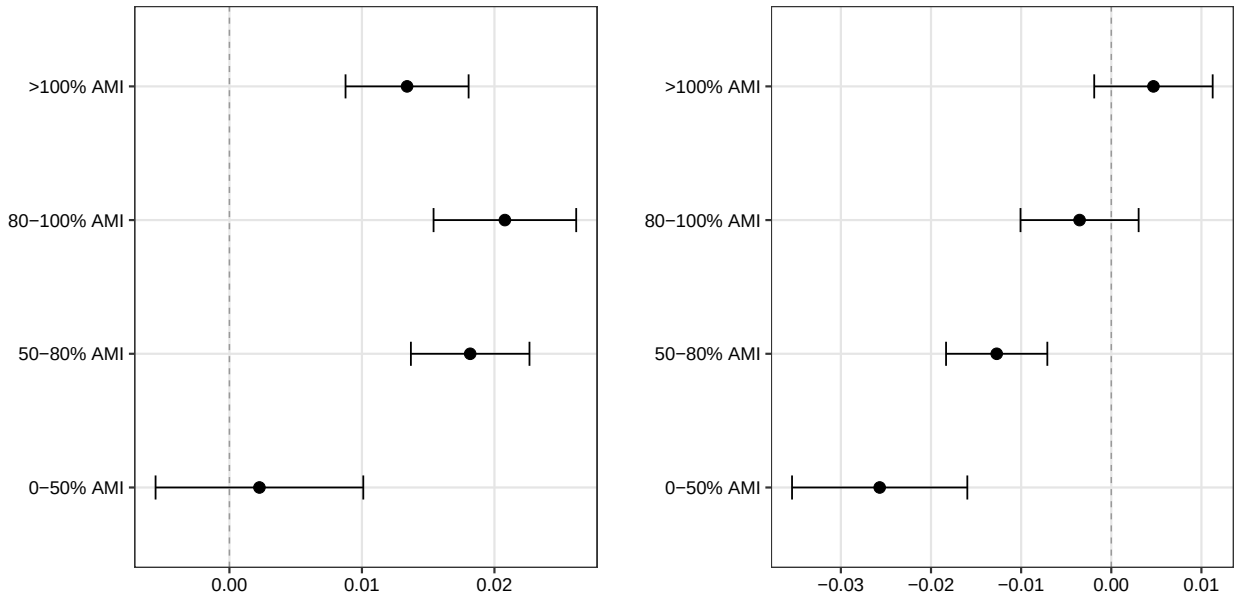
[19] HMDA redefined its occupancy field in 2018. Principal residence is the same category before and after the change, and the non-principal category pools second homes and investment properties on both sides, so the two counts are defined consistently across the break; the boundary-by-year fixed effects absorb any common shift in reporting, and the treated-control gap in the principal-residence share of originations is stable at about five percentage points in every year of the sample.

[20] Appendix Figure A2 reports the same margins by applicant income tercile and by applicant race; the small-balance gains concentrate in the lower terciles, and the small-loan share rises for White, Black, and Asian borrowers, by the most, 2.5 percentage points, for Black borrowers.

Table 5: **Borrower Margin: Approval, Origination Counts, and Loan-to-Income**

Panel A: Stage 2 – P(approved | applied) by applicant income tercile			
	Low income	Middle income	High income
Rural × Post	0.0119***	0.0149***	0.0118***
	(0.0016)	(0.0024)	(0.0026)
Tract FE	✓	✓	✓
Boundary × Year FE	✓	✓	✓
R^2	0.0309	0.0246	0.0187
Observations	4,393,279	4,806,518	5,417,923
Panel B: Poisson DiD on tract-year origination counts by occupancy			
	Principal residence	Non-principal residence	
Rural × Post	0.0520**	0.0506	
	(0.0224)	(0.0509)	
Tract FE	✓	✓	
Boundary × Year FE	✓	✓	
Pseudo-R^2	0.9096	0.6953	
Observations	120,577	120,083	
Panel C: $\log(LTI)$ *on originated loans by applicant income tercile*			
	Low income	Middle income	High income
Rural × Post	−0.0144***	−0.0140***	0.0077
	(0.0022)	(0.0032)	(0.0059)
Tract FE	✓	✓	✓
Boundary × Year FE	✓	✓	✓
R^2	0.2958	0.4203	0.2984
Observations	2,559,472	3,139,048	3,607,751

Notes: The table presents the treatment effect on the borrower margin in three panels. The sample consists of HMDA conventional first-lien single-family applications within five miles of the DTS rural eligibility boundary over 2013–2021. Panel A reports linear-probability estimates of the approval rate conditional on application, by applicant-income tercile; terciles are defined on 2013–2016 income. Panel B reports Poisson estimates on the tract-year counts of principal-residence and non-principal-residence originations, so the coefficients are log differences in expected counts between treated and control tracts; the pseudo-R^2 is McFadden's. Panel C reports estimates on the log loan-to-income ratio of originated loans with non-missing income, with the ratio winsorized at the 1% and 99% levels. The coefficient of Rural × Post identifies the treatment effect of the program. All specifications include census tract and boundary × year fixed effects. Standard errors, clustered at the state level, are reported in parentheses. ***p<0.01, **p<0.05, *p<0.1.



(a) Loan Approval

(b) Loan-to-Income (log)

Figure 7: Borrower Margin by Area Median Income (AMI) Category

Notes: The figure plots the treatment effect on loan approval and loan size by area-median-income (AMI) category. Each point is the coefficient of Rural × Post estimated separately for the indicated category, and the error bars represent 95% confidence intervals; the vertical dashed line marks zero. The sample consists of HMDA conventional first-lien single-family records with non-missing income within five miles of the DTS rural eligibility boundary over 2013–2021. AMI categories follow the statutory income definitions the GSEs apply under Duty to Serve: borrower income is measured relative to the FFIEC median family income of the property's metropolitan area or statewide non-metropolitan area, matched by tract and year. Panel (a) reports approval conditional on application, and Panel (b) reports the log loan-to-income ratio of originated loans. Pre-period approval rates in control tracts are 0.52, 0.63, 0.66, and 0.70, and pre-period mean loan-to-income ratios are 3.87, 3.09, 2.82, and 2.20, from the very-low-income to the above-100% category. All specifications include census tract and boundary × year fixed effects. Standard errors are clustered at the state level.

first-lien sample, 2013–2021.[21] All three coefficients are negative. The decline is what the resizing mechanism of Section 4.3 implies, as new entrants take small-balance loans that pull average LTI down. It also rules out a leverage buildup on the mortgage itself, because LTI is the loan-side component of the debt-to-income ratio.

Table 6: **Effects on Leverage Outcomes**

	log(LTI)	LTI (level)	Share LTI ≥ 4
Rural × Post	-0.0163^{***}	-0.0512^{***}	-0.0076^{***}
	(0.0034)	(0.0093)	(0.0017)
Tract FE	Yes	Yes	Yes
Boundary × Year × GSE FE	Yes	Yes	Yes
N (loans)	4,888,315	4,888,315	4,888,315

Notes: The table presents the treatment effect on three measures of borrower leverage. The sample consists of HMDA conventional first-lien single-family originations purchased by Fannie Mae or Freddie Mac in census tracts within five miles of the DTS rural eligibility boundary over 2013–2021. The loan-to-income ratio (LTI) is defined as the loan amount divided by applicant income, both in 2010 dollars. Rural is an indicator for tracts inside the rural eligibility boundary, and Post is an indicator for 2017 and later years. The coefficient of Rural × Post identifies the treatment effect of the program on the outcome in each column. All specifications include census tract and boundary × year × GSE fixed effects. Standard errors, clustered at the state level, are reported in parentheses. ***p<0.01, **p<0.05, *p<0.1.

What the program admitted also shows up in what the loans look like at origination. Figure 8 reports composition estimates on the same McDash sample. At full exposure, the average credit score falls by 3.4 points and the share originated at or below 660 rises by 1.6 percentage points; neither estimate is distinguishable from zero. The front-end debt-to-income ratio rises by 2.3 percentage points while the loan-to-value ratio does not move.[22] The shift that does appear is on a margin the GSEs do not price. Their loan-level price adjustments, the upfront fees charged at delivery, are set by the credit score and the loan-

[21] LTV is available in HMDA across this window but is sparsely populated and inconsistently measured before the 2018 reporting reform, which precludes a clean DiD. We use LTI, computed from HMDA loan amount and applicant income (both deflated to real 2010 dollars), as a standard leverage proxy.

[22] Both Enterprises relaxed underwriting at this margin in July 2017. Fannie Mae's Desktop Underwriter version 10.1 removed the overlays that had made mortgages at the maximum allowable debt-to-income ratio ineligible, and Freddie Mac adopted less restrictive criteria at roughly the same time (Federal Housing Finance Agency Office of Inspector General, 2019). The change is national, so the boundary comparison differences out its common component; what it cannot difference out is a differential response, which would require the two sides of a segment to have differed in how many applicants sat near the limit.

to-value ratio (Section 2), and neither moves. The debt-to-income ratio enters underwriting only as a ceiling and sets no fee. That margin carries little default risk.[23]



(a) Mean credit score

(b) Share at or below 660

(c) Loan-to-value ratio

(d) Front-end debt-to-income ratio

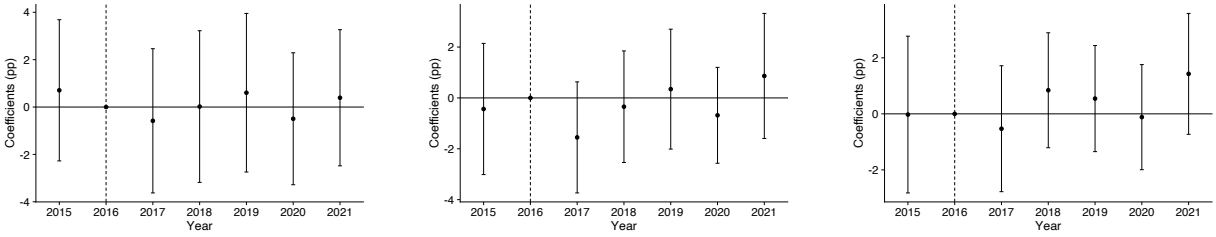
Figure 8: **Ex Ante Composition of Originated Loans**

Notes: The figure plots dynamic difference-in-differences estimates for the characteristics of loans at origination. The outcomes are the mean credit score, in points, and, in percentage points, the share of loans with a credit score at or below 660, the loan-to-value ratio, and the front-end debt-to-income ratio. The front-end debt-to-income ratio is defined as housing payments over income; the back-end ratio used in underwriting is not observed in McDash. The sample, the definition of exposure, the fixed effects, and the controls follow Figure 6. The coefficient of Exposure × year identifies the differential change for a fully exposed ZIP code relative to the 2016 origination cohort, which is normalized to zero. The error bars represent 95% confidence intervals. Standard errors are clustered at the state level.

Finally, Figure 9 shows no evidence that DTS raises delinquency. It reports event-study coefficients for whether a loan ever reaches 30, 60, or 90 days past due, on the same McDash sample. No coefficient is distinguishable from zero at any severity in any year. Because

[23] In the 2016Q1 Freddie Mac vintage, adding the debt-to-income ratio to a default model with the credit score and the loan-to-value ratio raises the pseudo-R^2 by less than one percentage point, and loans with ratios above 45 percent, approved only with compensating factors, default no more often within three years than loans at 41 to 45 percent (Appendix Table A9).

the flags record the whole servicing history, the 2017–2019 cohorts are the sharper test: each is observed through the end of the panel and so past the forbearance window, and the estimates there are −0.31, −0.14, and 0.43 percentage points, none significant, against pre-period delinquency rates of 9.8, 4.4, and 3.5 percent in unexposed ZIP codes. The 95% confidence intervals exclude increases above 1.7, 1.2, and 1.5 percentage points at full exposure, and the average loan in an exposed ZIP code has an exposure of 0.22.[24]



(a) 30 days past due (b) 60 days past due (c) 90 days past due

Figure 9: **Dynamic Effects on Mortgage Delinquency**

Notes: The figure plots dynamic difference-in-differences estimates for mortgage delinquency at the loan level, so that every performing loan stays in the denominator. The outcome in each panel is an indicator, in percentage points, for the loan ever reaching 30, 60, or 90 days past due through the end of the servicing panel. The sample, the definition of exposure, the fixed effects, and the controls follow Figure 6. The coefficient of Exposure × year identifies the differential change for a fully exposed ZIP code relative to the 2016 origination cohort, which is normalized to zero. Foreclosure and bankruptcy are reported in Appendix Figure A5. The error bars represent 95% confidence intervals. Standard errors are clustered at the state level.

The approval relief thus reaches low- and moderate-income owner-occupants, their loans are smaller relative to income, and their performance through the COVID period matches their observable risk. Two margins remain outside our data, as HMDA does not consistently report second liens or cash-out refinances in the boundary sample, so home-equity extraction is unobserved, and the size of the shock does not support a test of macro-financial propagation. On the margins we observe, the expansion delivers access, not speculation. The remaining question is how the housing market absorbs it.

[24]Foreclosure and bankruptcy, which add the servicer's and the court's decisions to the borrower's and which the 2020–2021 moratoria suppressed, are reported in Appendix Figure A5; no coefficient there is distinguishable from zero in any year.

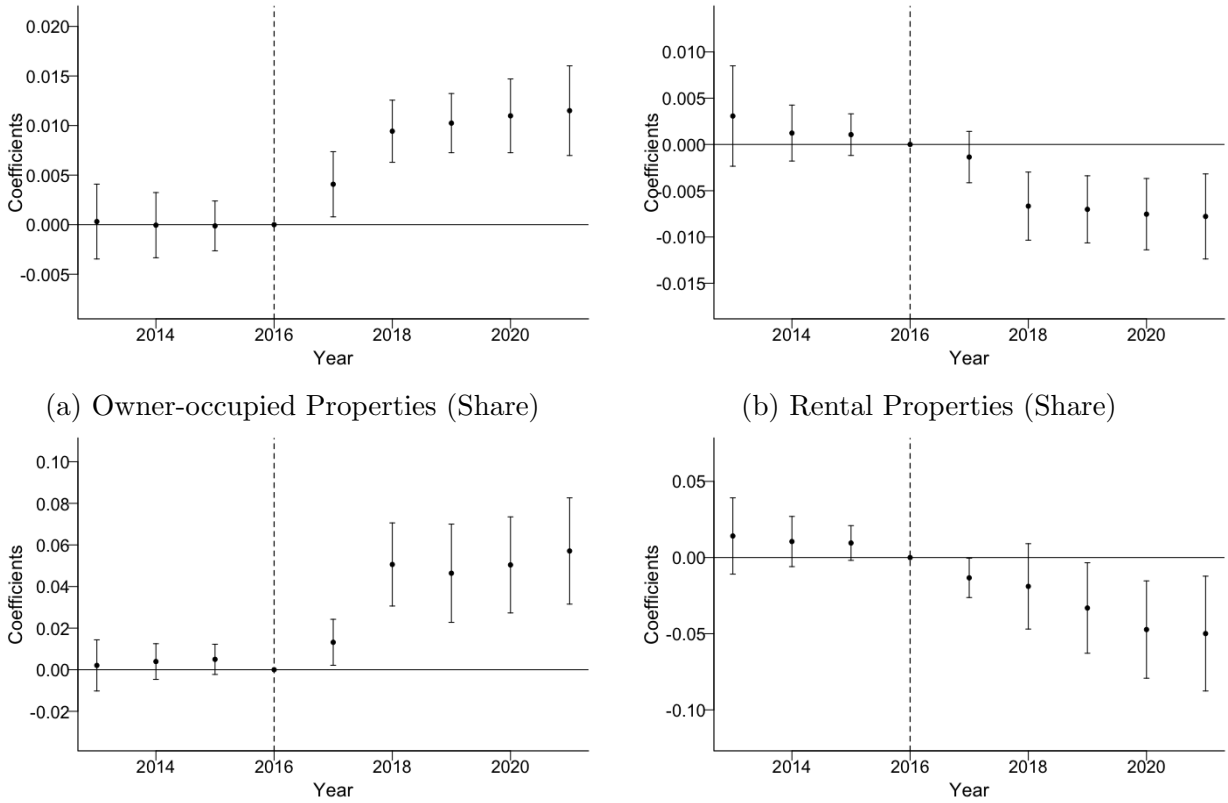
4.5 Housing Market Outcomes

The final question is where the gain comes to rest. The new buyers value their access, but whether they keep that value depends on how the housing market absorbs their demand, because any price response transfers part of the surplus to the people who already own. The credit-housing literature documents the first case as a strong regularity, in which increases in mortgage availability raise local house prices and incumbent owners capture much of the value in supply-constrained metropolitan markets (Mian and Sufi, 2009; Favara and Imbs, 2015; Adelino et al., 2025). By contrast, the elastic-supply rural markets where DTS operates deliver the second case, as the local market absorbs the demand shock through quantity expansion at stable prices.

Owner-occupancy rises and rental units fall in treated tracts. Figure 10 shows dynamic treatment effects on owner-occupied and renter-occupied housing in share form and log-count form. The owner-occupied share rises to 0.4 pp above controls in 2017 and reaches 1.2 pp by 2021 (Panel a), and the count of owner-occupied units climbs from 1.3% to roughly 5.7% above controls over the same window, averaging 4.06% over the post period. Rental properties move in mirror image, with the renter-occupied share down 0.8 pp and the rental-unit count down 5% by 2021 (Panels b and d).

The new owner-occupied stock comes from two margins, new construction and tenure conversion of existing rental units. Building-permit records corroborate the construction margin, as new residential permits rise by roughly 30% in eligible tracts, entirely in single-family units, with a monotone build-up starting in 2018, the year the first DTS plans took effect (Appendix Figure A3). The mirror-image rental decline is the conversion margin, and the rental-market evidence below examines the household flows on that margin. Combining the two count series, the total occupied housing stock is roughly 2% larger in treated tracts by 2021.[25]

[25] The baseline owner-occupied share of occupied units is 0.65 (Appendix Table A2), so the 2021 estimates for the two series imply total growth of $0.65 \times 5.7\% - 0.35 \times 5\% \approx 2\%$.



(a) Owner-occupied Properties (Share)

(b) Rental Properties (Share)

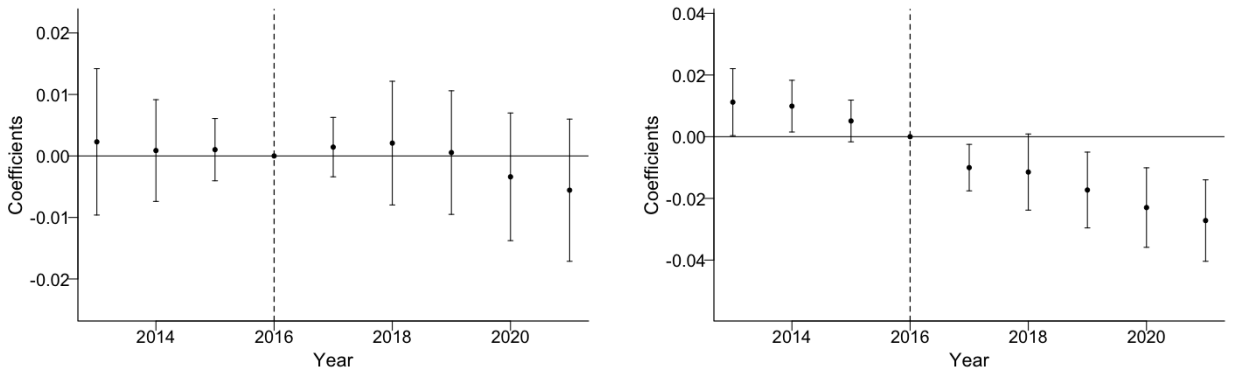
(c) Owner-occupied Properties (Units)

(d) Rental Properties (Units)

Figure 10: **Dynamic Effects on Homeownership**

Notes: This figure presents the dynamic difference-in-differences estimates of the effect of the Duty to Serve program on properties by tenure. The sample includes all census tracts within a 5-mile radius of a rural/non-rural boundary. The vertical dashed line indicates the program's announcement year (2016). Each point represents the estimated coefficient for a given year relative to 2016, the announcement year (omitted), with 95% confidence intervals.

Prices show no statistically detectable response. Panel A of Figure 11 shows flat tract-average home prices, and Figure 12 re-estimates the response within properties in the CoreLogic repeat-sales sample, with property fixed effects holding the dwelling constant. The within-property treatment effect is economically and statistically indistinguishable from zero throughout.



(a) Average Housing Price (b) Average Rent

Figure 11: **Dynamic Effects on the Housing Market**

Notes: These figures present the dynamic difference-in-differences estimates of the effect of the Duty to Serve program on the housing market. The sample includes all census tracts within a 5-mile radius of a rural/non-rural boundary. The vertical dashed line indicates the program's announcement year (2016). Each point represents the estimated coefficient for a given year relative to 2016, the announcement year (omitted), with 95% confidence intervals.

By contrast, rents fall by roughly 2.4% on average. Market rents may fall as rental demand shrinks, or the average may fall mechanically, because the units converting to owner-occupancy sit at the higher-quality end of the rental stock while new affordable units enter at the bottom. Aggregate data cannot separate the two, but the demand side is directly observable in household address histories from Infutor. Figure 13 reports event studies of the three margins of rental demand: fewer renters arriving, incumbent renters leaving, or incumbent renters switching into ownership. The margin that moves is entry, as renter arrivals into treated tracts fall by roughly 11% of their pre-period rate per year, a decline that builds monotonically after 2017 and holds in levels as well as rates.[26] By contrast,

[26]Because the household stock of treated tracts grows over the post period, a declining arrival rate could in principle coincide with a stable or rising number of arriving renters. Re-estimating arrivals in levels by

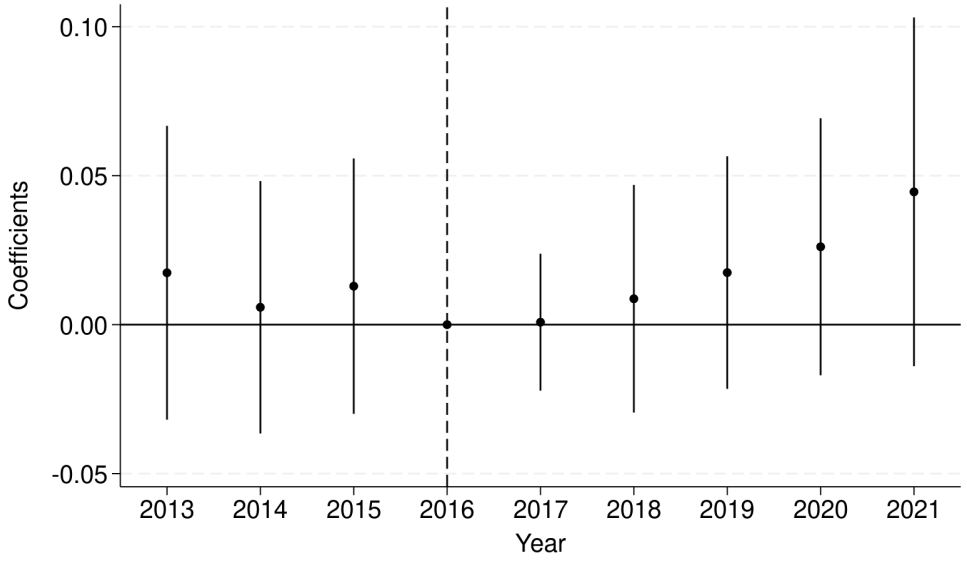
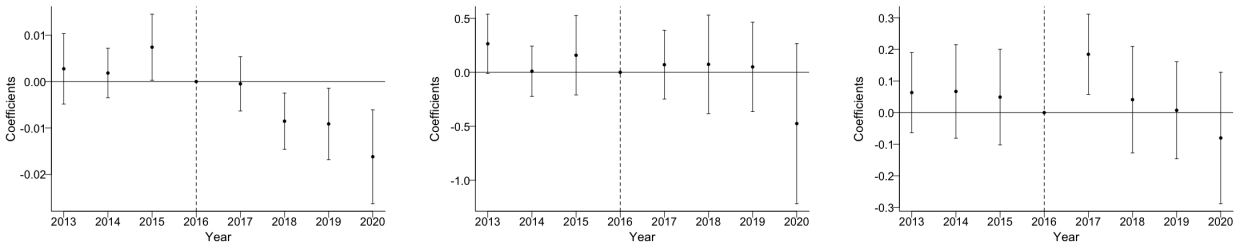


Figure 12: **Dynamic Effects on House Prices in the Repeated-Sales Sample**
Notes: This figure presents dynamic difference-in-differences estimates of the effect of the DTS program on housing prices using repeated-sales samples from CoreLogic. The outcome variable and fixed effects are defined at the repeated-sale property level. The sample includes all census tracts within a five-mile radius of the rural–non-rural boundary. The vertical dashed line marks the program announcement year (2016). Each point represents the estimated coefficient for a given year relative to the year immediately preceding the announcement, with 95% confidence intervals.

renter-to-owner transitions or renter departures offer no clear evidence in either direction. Rental demand therefore falls at the entry margin, regardless of how much of the measured rent decline is compositional.



(a) Renter arrival rate (b) Renter departure rate (c) Local conversion

Figure 13: **Rental-Market Adjustment Margins**

Notes: The figure plots the dynamic treatment effects on three margins of rental demand, measured from Infutor address histories. Panel (a) shows the renter arrival rate, in percentage points of the tract's household stock; Panel (b) shows the renter departure rate and Panel (c) the rate at which renters become owners through a local move, both in percentage points of the tract's renter households. The sample consists of census tracts within five miles of the DTS rural eligibility boundary over 2013–2020, with transitions dated by their origin year. Each point is the coefficient of Rural interacted with the indicated year relative to 2016, which is normalized to zero; the vertical dashed line marks the announcement year. All specifications include census tract and boundary × year fixed effects and weight tracts by the outcome's denominator. The error bars represent 95% confidence intervals. Standard errors are clustered at the state level.

The four margins trace a single causal chain. The GSEs shifted purchases toward small-balance loans; lenders converted the concession into approvals instead of lower contract rates; low- and moderate-income owner-occupants entered with small loans; and the housing market met them with new units without price increases, while rental demand thinned at the entry margin. The pattern has a simple explanation. In these rural tracts land is abundant and building is lightly regulated, so builders answer new demand with new units. If that is the right explanation, the pattern should weaken where building is harder, and it does: when we split the boundary tracts at the median of the Saiz (2010) housing supply elasticity, areas where supply responds easily absorb the shock through tenure shifts at flat prices, while more constrained areas show ownership gains less than a third as large (Appendix Section F.3). The quantity-without-price pattern is the elastic-supply counterpart of Adelino et al. (2012),

Poisson pseudo-maximum-likelihood on tract-year counts rules this out, as the number of arriving renters falls by roughly 14.5% with clean pre-trends.

whose price effects of conforming-loan eligibility are concentrated in regions with low housing supply elasticity, and it matches Bhutta and Ringo (2021), who find a large home-buying response to a small cut in FHA insurance premiums, concentrated among lower-income borrowers. What the reduced form cannot deliver is the value of this access to the households who receive it, the division of the gains among new buyers, incumbent owners, and renters, and the behavior of the same policy where supply cannot respond. Section 5 builds the model that delivers all three.

5 Structural Model

We construct a structural model of mortgage credit, household tenure choice, and housing-market equilibrium that delivers the three objects the reduced form cannot produce: the value of an approval to the household who receives it, the division of the surplus among new buyers, renters, and incumbent owners, and the behavior of the same concession under other delivery margins and in other housing markets. The model has two interacting blocks. First, a spatial-equilibrium block, building on canonical spatial-equilibrium models (Rosen, 1979; Roback, 1982; Diamond, 2016) and the discrete-choice framework of McFadden (1974), models heterogeneous households' joint choice of location, workplace, and tenure; only a share of households re-optimizes in a period, and the credit subsidy accrues to new owner originations alone. Second, a lender threshold block determines which marginal applications are approved as a function of the GSE-mediated funding cost (Campbell and Cocco, 2015; Loutskina and Strahan, 2009). Estimating the two blocks together yields the shadow price of an approval alongside the execution wedge that produces it, and the welfare incidence across household groups. Counterfactuals then deliver the same concession through contract rates and move it into a supply-constrained housing market. The model is silent on leverage, equity extraction, and default, which require longer time-series data and are handled with suggestive reduced-form evidence in Section 4.4.

5.1 Model Setup

The economy consists of I census tracts indexed by $i, j \in \{1, \ldots, I\}$, with i denoting residential location and j workplace location. Each tract has a binary DTS treatment status $\lambda_i \in \{0, 1\}$. Households differ by skill $s \in \{H, L\}$, defined by educational attainment as in Section 3, where H proxies for high credit access and L for constrained credit access, and by their origin location g_0 and origin tenure $k_0 \in \{O, R\}$, owner or renter.[27] The economy is in equilibrium in period $t = 0$, the DTS policy is introduced immediately before period $t = 1$, and we solve for the new equilibrium given the policy shock.

Among all households, a share λ_{s,k_0} of each type draws mover status and re-optimizes over the full location-workplace-tenure menu. The remaining $1 - \lambda_{s,k_0}$ stayers keep their residential market and tenure, while workplace choice remains free. For the λ_{s,k_0} households, they choose a residential and workplace location pair (i, j) in the upper nest, and within that pair they choose a tenure $k \in \{O, R\}$ in the lower nest.[28] Specifically, a household of skill s choosing (i, j, k) in period t has indirect utility

$$\tilde{V}^{s,h}_{i,j,k,t} = \beta^s \left(w^s_{j,t} - \chi^s_{i,j,t} - \theta^{s,k} \tilde{r}_{i,k,t} + \Gamma_{i,k,t} \right) + \gamma^s a^s_{i,t} + c_s \mathbf{1}\{k = k_0\} + \tilde{\varepsilon}^{s,h}_{i,j,k,t}, \tag{3}$$

where β^s is the income coefficient, $w^s_{j,t}$ the real wage at workplace j, $\chi^s_{i,j,t}$ the commuting cost, $\tilde{r}_{i,k,t}$ the housing service cost for a unit of tenure k, $\theta^{s,k}$ the housing expenditure share, $\Gamma_{i,k,t}$ the policy impact defined below, $a^s_{i,t}$ the log amenity, γ^s the amenity weight in utility, c_s the tenure-continuation premium applied by movers to cells that keep the origin tenure, and $\tilde{\varepsilon}^{s,h}_{i,j,k,t}$ an idiosyncratic preference shock indexed by household h. The shocks are

[27] We collect these into a household type $z = (s, g_0, k_0)$. Skill is predetermined and exogenous to DTS (Section 3), and the mapping between skill and credit access in our HMDA sample is documented in Table A3.

[28] A flat conditional logit would impose the independence of irrelevant alternatives (IIA). That is implausible here: a household's tenure options within a given location are far closer substitutes than options in other locations, so a renter induced into ownership is much more likely to buy in the same location than to relocate. The nest keeps IIA between the two tenure alternatives within a location-workplace pair and relaxes it across pairs: with the dissimilarity parameter $\rho < 1$, the unobserved tastes for owning and renting in the same pair are correlated, so a gain in the value of owning in a tract draws mainly from renting in that tract rather than from other locations; $\rho \to 1$ recovers the flat logit.

jointly distributed across (i, j, k) as a nested generalized extreme value (GEV) distribution (McFadden, 1978), with the location-workplace pair (i, j) as the upper nest and tenure k as the lower nest. The dissimilarity parameter $\rho \in (0, 1]$ governs the within-nest correlation: smaller ρ implies that tenure choices within a given (i, j) are more substitutable than across-location moves. The marginal distribution of each $\tilde{\varepsilon}^{s,h}_{i,j,k,t}$ is Type-I extreme value (Gumbel), and the joint distribution reduces to independent Gumbel as $\rho \to 1$, in which case the model collapses to a standard conditional logit (McFadden, 1974).

The housing service cost depends on tenure. Renters pay rent, and owners pay a cost proportional to the house price scaled by the user cost of housing κ_i:

$$\tilde{r}_{i,k,t} = \begin{cases} r_{i,t} & \text{if } k = R, \\ p_{i,t} + \kappa_i & \text{if } k = O, \end{cases} \tag{4}$$

where $p_{i,t}$ and $r_{i,t}$ are the log price and log rent in tract i, and $\kappa_i = \ln(\iota_i + d_i + \omega_i - g_i)$ combines the market interest rate, depreciation rate, property tax rate, and expected capital appreciation rate (Poterba, 1984).[29] The DTS policy shock $\Gamma_{i,k,t}$ enters utility only for new owner originations in eligible tracts.

We define Γ as the policy benefit and τ as its size, which is the user-cost reduction that eligibility delivers on each new owner origination. The benefit goes to households that take out a new purchase mortgage in an eligible market, and the program can buy such a loan in exactly two situations: a renter buys a home, in any market, or an owner buys a home outside her origin market.[30] Such a household receives $\Gamma := \theta^{s,O}\tau$, the reduction scaled by its

[29] Throughout the model, lowercase $p_{i,t}$, $r_{i,t}$, $w^s_{j,t}$, $\tilde{r}_{i,k,t}$, and $a^s_{i,t}$ denote log prices, log rents, log wages, the log housing service cost, and the log amenity, while uppercase $P_{i,t}$, $R_{i,t}$, $\tilde{R}_{i,k,t}$, $W^s_{j,t}$, $A^s_{i,t}$, and $H_{i,k,t}$ denote the corresponding levels used in the housing-supply, equilibrium, and hat-algebra equations of Sections 5.2 and 5.4.

[30] Our household-level data allow us to observe directly who stayed, who moved, and whether a mover ended up owning or renting in each period. These household records enter the estimation as transition rates by origin group, the stay and conversion moments, and the model produces the same rates from the aggregate flows. However, the model cannot capture an owner who sells her home and buys a different one within the same market, as it tracks markets, not individual properties. This ignores a small number of genuine new originations.

housing expenditure share. Everyone else receives nothing. An owner who stays in her origin market keeps her existing mortgage, so there is no loan for the program to buy; renters who keep renting are outside the program by definition; and ineligible markets carry no benefit at all.

Movers solve the joint location-tenure problem under the nested-logit structure (McFadden, 1978). Let $V^s_{i,j,k,t}$ denote the deterministic component of utility. For a mover, the joint choice probability factorizes into the product of an upper-nest probability over location-workplace pairs and a lower-nest conditional probability over tenures within a pair:

$$\pi^s_{i,j,k,t} = \pi^s_{i,j,t} \cdot \pi^s_{k|i,j,t}, \tag{5}$$

where the upper-nest probability over location-workplace pairs is

$$\pi^s_{i,j,t} = \frac{(I^s_{i,j,t})^\rho}{\sum_{i',j'}(I^s_{i',j',t})^\rho}, \tag{6}$$

the lower-nest conditional probability over tenures within nest (i, j) is

$$\pi^s_{k|i,j,t} = \frac{\exp(V^s_{i,j,k,t}/\rho)}{I^s_{i,j,t}}, \tag{7}$$

and $I^s_{i,j,t} = \sum_k \exp(V^s_{i,j,k,t}/\rho)$ is the inclusive value of the lower nest. In estimation, the baseline mean utilities in these expressions are recovered by Berry inversion.

On the labor market side, firms in workplace j produce a homogeneous tradable good (the numeraire) using high- and low-skilled labor in a Cobb-Douglas technology with location-specific productivity:

$$Y_{j,t} = Z_{j,t}(N^H_{j,t})^\alpha (N^L_{j,t})^{1-\alpha}, \tag{8}$$

where $\alpha \in (0, 1)$ is the high-skilled labor share. To allow the model to reproduce the empirical wage stability under population inflows, we let productivity be endogenous to

local employment density through agglomeration economies parameterized by $\sigma \geq 0$:

$$Z_{j,t} = \bar{Z}_{j,t} \left(N_{j,t}^{H} + N_{j,t}^{L} \right)^{\sigma}, \tag{9}$$

with $\bar{Z}_{j,t}$ the exogenous fundamental productivity. Firms hire labor up to the point where marginal product equals the wage, so the equilibrium wages combine an agglomeration force (positive in total density) with a relative-supply force (negative in own-skill density).

Housing supply is segmented by tenure. Supply of tenure k in tract i is

$$H_{i,k,t} = \bar{H}_{i,k} \, P_{i,k,t}^{\mu_{i,k}}, \tag{10}$$

with the owner-occupied supply elasticity

$$\mu_{i,O} = \mu_{i,\text{base}} \times \begin{cases} \xi & \text{if } \lambda_i = 1, \\ 1 & \text{otherwise,} \end{cases} \tag{11}$$

where $\mu_{i,\text{base}}$ is the baseline elasticity calibrated from Saiz (2010) and $\xi > 1$ is the multiplier capturing the additional supply responsiveness in DTS-eligible rural tracts. Rental supply elasticity is $\mu_{i,R} = \mu_{i,\text{base}}$ at the baseline, unaffected by the policy. The parameter ξ is the model's central regime variable: it governs whether the credit shock is absorbed through price capitalization (low ξ) or quantity expansion (high ξ). The decline in rents documented in Section 4.5 arises from the demand side: the subsidy makes ownership more attractive for the marginal entrant, reducing rental demand and pulling rents down.

Local amenities depend on the scale of the tract's population relative to baseline. Let $N_{i,t}^{s} = \sum_{j,k} F_{i,j,k,t}^{s}$ denote the resident population of skill s in market i at time t. The amenity level perceived by a resident of skill s is

$$A_{i,t}^{s} = \bar{A}_{i}^{s} \cdot \left(\frac{1}{2} \left(\frac{N_{i,t}^{H}}{N_{i,0}^{H}} + \frac{N_{i,t}^{L}}{N_{i,0}^{L}} \right) \right)^{-\delta}, \tag{12}$$

where $\bar{A}_i^s > 0$ is the exogenous amenity level and $\delta \geq 0$ is a congestion parameter: growth in total population strains shared infrastructure and reduces amenities uniformly across skill groups. The congestion factor is the unweighted average of the two skill groups' population ratios relative to baseline; it equals one at baseline and coincides with total population growth when the two groups grow in proportion, so amenities respond to scale and population inflows are self-damping.[31]

The structural parameters governing the spatial-equilibrium block are the policy subsidy τ, the supply multiplier ξ, the nesting parameter ρ, the four mover shares λ_{s,k_0}, and the two tenure-continuation premia c_s. Next, we formally define the spatial equilibrium in Section 5.2, and introduce the lender problem in Section 5.3.

5.2 Spatial Equilibrium

A spatial equilibrium is a vector of wages $\{w_{j,t}^{s*}\}$, housing prices $\{p_{i,t}^*\}$, rents $\{r_{i,t}^*\}$, amenities $\{A_{i,t}^{s*}\}$, population flows $\{F_{i,j,k,t}^{s*}\}$, labor allocations $\{N_{j,t}^{s*}\}$, and housing stocks $\{H_{i,k,t}^*\}$ such that the following conditions hold for all locations i, j, skill types s, and tenure choices $k \in \{O, R\}$:

- Population flows aggregate optimal stayer and mover choices across types:

$$F_{i,j,k,t}^{s*} = \sum_{z:s(z)=s} N_z \left[(1 - \lambda_{s,k_0}) \, \pi_{j,t}^{S,z} \mathbf{1}\{i = g_0, k = k_0\} + \lambda_{s,k_0} \, \pi_{i,j,k,t}^{s,z} \right]. \tag{13}$$

[31] We also considered the general specification in which amenities respond to the skill composition of the population through spillover elasticities, but restrict them to zero. First, spillovers are weakly identified in this setting. The policy moves both skill populations nearly proportionally. When composition is stable, the spillover terms are almost collinear with the scale term that congestion already carries, so the data can separate them only from the small residual variation in skill mix. Second, the amenity channels this parameter represents, density-driven retail variety and service quality of the kind Diamond (2016) documents across cities, operate through density thresholds that the low-density tracts of the rural boundary sample do not reach.

- High-skilled labor demand equals high-skilled labor supply:

$$N_{j,t}^{H*} = \sum_{i,k} F_{i,j,k,t}^{H*}, \qquad w_{j,t}^{H*} = \ln \alpha + \ln Z_{j,t}^{*} + (\alpha - 1) \ln N_{j,t}^{H*} + (1 - \alpha) \ln N_{j,t}^{L*}. \quad (14)$$

- Low-skilled labor demand equals low-skilled labor supply:

$$N_{j,t}^{L*} = \sum_{i,k} F_{i,j,k,t}^{L*}, \qquad w_{j,t}^{L*} = \ln(1 - \alpha) + \ln Z_{j,t}^{*} + \alpha \ln N_{j,t}^{H*} - \alpha \ln N_{j,t}^{L*}. \quad (15)$$

- Housing supply equals demand for each tenure cell:

$$H_{i,k,t}^{*} = \sum_{s,j} F_{i,j,k,t}^{s*}. \quad (16)$$

- Endogenous amenities are determined by the population composition:

$$A_{i,t}^{s*} = \bar{A}_{i}^{s} \cdot \left(N_{i,t}^{H*}\right)^{\alpha_{H|s}} \cdot \left(N_{i,t}^{L*}\right)^{\alpha_{L|s}} \cdot \left(\frac{1}{2}\left(\frac{N_{i,t}^{H*}}{N_{i,0}^{H}} + \frac{N_{i,t}^{L*}}{N_{i,0}^{L}}\right)\right)^{-\delta}. \quad (17)$$

In the spatial equilibrium, we take the user-cost reduction τ as a primitive and treat it as the household-side benefit of DTS eligibility per new origination. The reduced-form evidence in Section 4.3 shows that the policy operates through the lender's accept-or-reject decision. The structural counterpart of this finding is a lender problem in which DTS raises the net proceeds and lowers the delivery cost of a qualifying rural loan, which we summarize as a reduction in the lender's effective funding cost, relaxing the approval threshold for marginal applications.

5.3 Lender Problem

A representative lender evaluating an application from borrower b in market m approves the loan if expected per-unit profit is positive:

$$P_a(b,m) = \mathbf{1}\{\,[r(b,m) - c(m)] - q(b)\cdot \mathrm{LGD} - \kappa_L > 0\,\}\,, \tag{18}$$

where $r(b,m)$ is the contract rate, $c(m)$ the lender's funding cost, $q(b)$ the expected default probability, LGD the loss given default, and κ_L the lender's required margin per unit of loan principal (covering servicing, capital, and the risk premium beyond expected default loss). The default probability is a probit in the borrower's credit score, loan-to-value ratio, and debt-to-income ratio, the determinants that drive default in Campbell and Cocco (2015):

$$q(b) = \Phi(\beta_0 + \beta_F \cdot \mathrm{FICO}_b + \beta_L \cdot \mathrm{LTV}_b + \beta_D \cdot \mathrm{DTI}_b)\,, \tag{19}$$

DTS raises the value of a qualifying loan to the lender that originates and delivers it, through better purchase execution on delivery and a lower cost of originating an eligible loan (Section 2). Both enter the lender's problem in the same place, as a reduction in the effective funding cost $c(m)$ in eligible tracts after the policy is announced, so we model the concession as a wedge in $c(m)$ and let it differ by the loan-size bins defined in Section 4.3:

$$c(m) = \begin{cases} c_0 - \Delta_{\mathrm{size}(m)}, & \text{if } \lambda_{i(m)} = 1 \text{ and } t = 1, \\ c_0, & \text{otherwise,} \end{cases} \tag{20}$$

where Δ_{size} is the bin-specific wedge with size $\in \{\mathrm{Small}, \mathrm{Conforming}, \mathrm{Jumbo}\}$. A fraction γ_{pass} of the wedge passes through to the borrower's contract rate as a lower offer; the remaining fraction stays as additional lender margin and contributes to relaxation of the approval threshold:

$$r(b,m) = r_0(b) - \gamma_{\mathrm{pass}} \cdot \Delta_{\mathrm{size}(m)} \cdot \mathbf{1}\{\lambda_{i(m)} = 1, t = 1\}\,. \tag{21}$$

where $r_0(b)$ is the base rate schedule.

The user-cost subsidy τ in the spatial-equilibrium block is the household-side counterpart of the lender wedge Δ_{size}. A fraction γ_{pass} of Δ reaches the borrower as a lower contract rate: on a loan of size L it saves $\gamma_{\text{pass}}\bar{\Delta}L$ per year, or $\gamma_{\text{pass}}\bar{\Delta}(L/y)$ of income for a loan-to-income ratio L/y, which with annual housing expenditure $\theta^{s,O}y$ is a reduction of $\gamma_{\text{pass}}\bar{\Delta}(L/y)/\theta^{s,O}$ in the user cost κ_i of equation (4). The remaining $(1-\gamma_{\text{pass}})\Delta$ widens the lender's margin and lowers the approval threshold. This access channel leaves contract rates unchanged, so it is invisible in the rate-spread moment, and its value to the household is the user-cost equivalent of an approval that the baseline threshold would have denied. Section 5.6 compares the two channels in the welfare results.

The structural parameters of the lender problem are the three wedges $(\Delta_{\text{Small}}, \Delta_{\text{Conforming}}, \Delta_{\text{Jumbo}})$, estimated jointly with the spatial parameters in Section 5.4.

5.4 Estimation

Estimation of the model makes use of both market- and individual-level data. The market-level data consist of the boundary responses of Section 4, which record how prices, rents, owner and rental stocks, and approval rates moved. The address histories record who stayed, who moved, and who switched tenure, and these transition rates discipline how many households face each choice margin in a given period. We use the model to fit them jointly, because they scale each other in equilibrium. A subsidy estimated on market responses alone would be spread across households who never faced the choice, and transition rates alone would say nothing about what access was worth.

Concretely, we estimate twelve structural parameters jointly: the spatial and mobility parameters

$$\Theta_{\text{sm}} = (\tau, \xi, \rho, \lambda_{O,H}, \lambda_{O,L}, \lambda_{R,H}, \lambda_{R,L}, c_H, c_L)$$

and the lender-side funding-cost wedges

$$\Theta_{\text{le}} = (\Delta_{\text{Small}}, \Delta_{\text{Conforming}}, \Delta_{\text{Jumbo}}),$$

on eighteen moments, listed with their values and sources in Table 7. They fall into the two families above: the market-level difference-in-differences responses of Section 4, and the household-level transition moments constructed from anonymized Infutor address histories in the boundary sample.[32]

The remaining parameters are calibrated (Table 7, Panel C). On the household side, the income coefficient is $\beta = 2.5$ and the amenity weights are $\gamma^H = 1.0$ and $\gamma^L = 0.5$ (Diamond, 2016; Monte et al., 2018); the housing expenditure shares are 0.24 and 0.30 for high-skilled owners and renters and 0.43 and 0.55 for low-skilled owners and renters, from debt-to-income limits and cost-burden data; the congestion parameter is $\delta = 0.20$ (Almagro and Domínguez-Iino, 2025); and the realized share of the ownership dividend is $\eta = 1$, its upper bound. On the production side, the agglomeration elasticity is $\sigma = 0.04$ (Combes et al., 2012), and the high-skilled labor share $\alpha = 0.35$ matches the high-skilled wage-bill share in the estimation sample. The amenity spillovers $\alpha_{H|H}$, $\alpha_{H|L}$, $\alpha_{L|L}$, and $\alpha_{L|H}$ are set to zero: in the thin and dispersed populations of the rural boundary sample the unrestricted cross-skill estimate is small, negative, and unstable across specifications, so the congestion term carries the scale effects. In the lender block, the default probit has slopes of -0.005 per credit-score point, 2.5 per unit of loan-to-value, and 1.5 per unit of debt-to-income, with the intercept anchored so that a reference borrower with a credit score of 720, a loan-to-value ratio of 0.85, and a debt-to-income ratio of 0.35 defaults with probability 1.5%, the aggregate rate over the relevant Freddie Mac vintages;[33] the base rate $r_0(b)$ is the Freddie Mac Primary Mortgage Market Survey 30-year rate plus the public loan-level price adjustments for credit score and

[32] We target arrival contrasts rather than resident-population levels because, with amenity spillovers restricted to zero, the composition of arrivals carries the borrower-composition information directly.

[33] For the reference borrower the function rises from 0.8% at a debt-to-income ratio of 20 percent to 2.1% at 45 percent, a factor of 2.6, against a factor of 2.4 in observed 90-day delinquency across the same bins in the 2016Q1 vintage after adjustment for credit score and loan-to-value (Appendix Table A9, Panel C).

loan-to-value; the loss given default is 0.30; and the required margin $\kappa_L = 0.008$ anchors the integrated approval probability to the 67.2% pre-2017 conventional approval rate of the analysis sample, which also absorbs the level of the funding cost c_0, so only the change Δ, identified from the approval difference-in-differences, matters for the results. Finally, $\gamma_{\text{pass}} = 0$, because the within-post rural–control gap in HMDA rate spreads on the post-2018 conventional sample is indistinguishable from zero (Section 4.3).[34]

To make estimation computationally tractable and to address the sparseness of bilateral commuting flows, we aggregate census tracts into groups for the structural estimation. Tracts that share the same boundary segment and treatment status are combined into a single group, and tracts more than five miles from any boundary are pooled into a single outside option for which all change variables are normalized to one. We then estimate the model on commuting flows between groups rather than between individual tracts, with location-level variables averaged across the constituent tracts using resident or worker populations as weights. Baseline mean utilities are Berry-inverted on this grouped system, so the model reproduces the baseline flow panel exactly and estimation moves only the policy response.

Given the baseline allocation, we solve for the post-policy equilibrium in proportional changes from the baseline, the exact hat algebra of Dekle et al. (2007): for every level variable X, $\hat{X} = X_1/X_0$. Working in changes lets the unobserved fundamentals (baseline productivities, amenity levels, and housing-stock scales) cancel, so the post-policy equilibrium follows from the observed baseline, the calibrated parameters, and the structural parameters alone. Wages, prices, rents, and amenities update through the four equilibrium blocks of Section 5.2, with the owner-supply elasticity $\mu_{i,O} = \mu_{i,\text{base}} \times \xi$ in treated tracts and $\mu_{i,\text{base}}$ otherwise; Appendix B.1 states the equations. The change in the systematic utility of a mover cell is

$$\Delta V^s_{i,j,k} = \beta^s \left(\ln \hat{W}^s_j - \theta^{s,k} \ln \hat{P}_{i,k} \right) + \beta^s \Gamma^s_{i,k,1} + \gamma^s \ln \hat{A}^s_i, \tag{22}$$

[34]The gap is about one basis point, statistically significant but negligible next to the wedges, in the DTS-eligible small-balance subsample.

where the policy term Γ enters every eligible owner cell except the own-location owner cell of origin-owner types. Post-policy flows aggregate the updated stayer and mover choices across household types with the baseline probabilities as weights, and we iterate on the blocks from $\hat{X} = 1$ until convergence.

The address histories add household-level evidence. For each household we see whether it stayed in its origin market over the transition window and, if it began as a renter, whether it owned a home at the end. Averaging within boundary arms by skill and origin tenure gives four stay rates and four renter-to-owner conversion rates, and the model produces the same rates. In the model, a household stays either because it is not among the movers or because it is, considers every market, and still chooses home. A renter, in turn, becomes an owner only by moving, so the conversion rate is the share of renter-origin households that move and buy within the boundary band. Because a household that reconsiders every market rarely lands back home, each stay rate pins down its mover share almost directly: a stay rate of 90 percent means a mover share of about 10 percent. The conversion rates then pin down the tenure-continuation premium c_s, since a larger premium means fewer conversions.

The remaining parameters follow the same moment-to-parameter logic; Appendix B.5 tabulates the mapping. The subsidy τ works on the new-origination margin and is identified by the arrival contrasts and the owner-quantity response. The supply multiplier ξ is identified by the coexistence of a flat price response with a strong quantity response. The nesting parameter ρ is pinned by the rent response together with the tenure tilt of arrivals, as stronger within-location tenure substitution converts arriving demand into ownership and pulls rents down. Each wedge Δ_{size} is identified by the approval change in its own loan-size bin, the only place it enters the lender's profit. In all, eighteen moments identify twelve parameters, leaving six overidentifying restrictions, and the estimator minimizes the quadratic SMM criterion with a diagonal weighting matrix that weights each moment by the inverse of its squared standard error (Appendix B.4).

5.5 Parameter Estimates

Table 7 presents the target moments, estimated parameters, and calibrated parameters. We discuss each estimated parameter in turn.

The policy subsidy $\tau = 0.456$ is the user-cost reduction per new origination and thus the shadow price of an induced approval. For a borrower whose origination the policy induces, the loan is worth a 37% reduction in the user cost of owning in the origination year.[35] The magnitude reflects two features of the estimation. First, it prices the relief of a binding constraint for the household at the margin, not a discount available to every owner. Second, the approval-margin relief is substantial, consistent with the evidence that housing demand responds far more to financing access than to its price (Fuster and Zafar, 2021).

With a nesting parameter $\rho = 0.454$ well below one, movers treat the owner-renter choice within a given location as far more substitutable than choices across locations. This is the entry-margin adjustment the address histories document, where renter arrivals fall while local renter-to-owner conversion stays flat (Section 4.5), and it is how the model generates the rent decline, as the rental channel of arrivals thins.

The supply elasticity multiplier $\xi = 2.50$ governs how the credit shock is split between the access channel (quantity expansion reaching new buyers) and the capitalization channel (price appreciation benefiting incumbents). The point estimate implies that owner-occupied housing supply in DTS-eligible tracts is roughly two and a half times as elastic as in adjacent non-rural areas. The magnitude is plausible for this setting, as the literature finds the highest supply elasticities where developable land is abundant and regulation is light (Saiz, 2010; Gyourko et al., 2021), and manufactured-housing collateral as a DTS focus area adds a production-technology channel, since factory-built housing enters the stock at lower marginal cost than site-built construction.

[35] Income-equivalent conversions multiply τ by the housing expenditure share $\theta^{s,k}$, which is calibrated rather than estimated, and are therefore illustrative rather than additional estimates: for the origination year, a credit-constrained new buyer devoting 43% of income to housing values the approval at $\tau \times \theta^{L,O} \approx 19.6\%$ of annual income, and a high-credit-access buyer at $\approx 10.9\%$.

Table 7: **Spatial-Equilibrium Estimation: Target Moments, Model Fit, and Parameter Estimates**

Panel A: Target Moments and Model Fit

Moment	*Description*	*Data*	*S.E.*	*Model*
Spatial moments (%)				
ΔPrice_{Avg}	Change in tract-average owner-occupied price	−0.20	0.54	0.86
ΔRent	Change in rents	−2.43	0.61	−1.37
ΔQty_{Own}	Change in owner-occupied units	4.06	0.99	3.70
ΔQty_{Rent}	Change in renter-occupied units	−4.09	1.30	−2.60
Lender approval-rate moments (pp)				
$\Delta\text{Stage2}_{\text{Small}}$	DTS-small approval rate (≤$200k)	1.09	0.14	1.08
$\Delta\text{Stage2}_{\text{Conforming}}$	Conforming approval rate	1.92	0.32	1.91
$\Delta\text{Stage2}_{\text{Jumbo}}$	Jumbo approval rate	1.64	0.72	1.64
Infutor mobility moments (%)				
$\text{Stay}_{O,H}$	Stay rate, high-skilled owners	91.56	0.24	91.54
$\text{Stay}_{O,L}$	Stay rate, low-skilled owners	91.30	0.23	91.27
$\text{Stay}_{R,H}$	Stay rate, high-skilled renters	88.20	0.52	88.24
$\text{Stay}_{R,L}$	Stay rate, low-skilled renters	87.14	0.28	87.13
Conv^{C}_{H}	Renter-to-owner conversion, control, high-skilled	3.94	0.14	3.82
Conv^{T}_{H}	Renter-to-owner conversion, treated, high-skilled	3.64	0.17	3.82
Conv^{C}_{L}	Renter-to-owner conversion, control, low-skilled	3.58	0.14	3.65
Conv^{T}_{L}	Renter-to-owner conversion, treated, low-skilled	3.73	0.14	3.65
Arrival contrasts (ratio)				
Inflow^{rel}_{H}	In-migration into eligible vs. ineligible tracts, high-skilled (% gap)	36.5	8.3	20.2
Inflow^{rel}_{L}	In-migration into eligible vs. ineligible tracts, low-skilled (% gap)	24.2	7.2	28.5
Arr^{rel}_{Rent}	Arrivals entering as renters, eligible vs. ineligible tracts (% gap)	−14.5	4.5	−19.7

Panel B: Estimated Structural Parameters

Parameter	*Description*	*Estimate*	*Boot. S.E.*
τ	Policy subsidy (user-cost reduction)	0.456^{***}	0.070
ξ	Rural owner-supply elasticity multiplier	2.499^{***}	0.922
ρ	Nesting parameter (tenure substitution)	0.454^{***}	0.074
$\lambda_{O,H}$	Mover share, high-skilled owner origins	0.085^{***}	0.002
$\lambda_{O,L}$	Mover share, low-skilled owner origins	0.088^{***}	0.002
$\lambda_{R,H}$	Mover share, high-skilled renter origins	0.118^{***}	0.005
$\lambda_{R,L}$	Mover share, low-skilled renter origins	0.129^{***}	0.003
c_H	Tenure-continuation premium, high-skilled movers	0.445^{***}	0.054
c_L	Tenure-continuation premium, low-skilled movers	0.001	0.020
Lender-Layer Wedges (basis points)			
Δ_{Small}	GSE wedge on DTS-eligible small-balance (see notes)	2.94^{***}	0.39
$\Delta_{\text{Conforming}}$	GSE wedge on conforming-balance loans	4.58^{***}	0.73
Δ_{Jumbo}	Effective margin shift, above-limit bin	3.66^{**}	1.59

Panel C: Calibrated Parameters

Parameter	*Description*	*Value*	*Source*
Preference Parameters			
γ^H	High-skilled sensitivity to amenities	1.00	Diamond (2016)
γ^L	Low-skilled sensitivity to amenities	0.50	Diamond (2016)
β	Migration / income elasticity	2.50	Diamond (2016); Monte et al. (2018)
Housing Expenditure Shares (θ)			
$\theta_{O,H}$	High-skilled owners	0.24	DTI limits, Cost Burden data
$\theta_{R,H}$	High-skilled renters	0.30	DTI limits, Cost Burden data
$\theta_{O,L}$	Low-skilled owners	0.43	DTI limits, Cost Burden data
$\theta_{R,L}$	Low-skilled renters	0.55	DTI limits, Cost Burden data
Production and Amenity Parameters			
α	High-skilled labor share in production	0.35	H-skill wage bill share from data
σ	Agglomeration elasticity	0.04	Combes et al. (2012)
δ	Level-based congestion in amenity	0.20	Almagro and Domínguez-Iino (2025), Saiz-range translation
η	HELOC realization fraction	1.00	Upper bound (full extraction)
γ_{pass}	GSE-to-borrower rate pass-through	0.00	Calibrated from HMDA rate-spread DiD null

Notes: The table presents the target moments, the model fit, and the parameter estimates of the joint SMM estimation. Panel A lists the eighteen targeted moments and their sources. The four spatial moments are the difference-in-differences estimates of Section 4.5 for treated relative to control tracts within five miles of the DTS rural eligibility boundary; the three approval moments are the bin-specific Stage 2 coefficients of Table 4; and the eleven mobility moments are stay rates and renter-to-owner conversion rates by tenure and skill, arrival contrasts by skill, and the renter-channel arrival contrast, all measured from Infutor address histories. The Model column reports the model-implied moments at the estimated parameters. Panel B reports the twelve parameters estimated jointly: the spatial, mobility, and tenure-continuation parameters, identified from the spatial and mobility moments, and the three loan-size-specific wedges, identified from the approval moments. Standard errors are from a parametric bootstrap with 100 draws, in which the moment targets are perturbed by independent normal draws at their standard errors and the model is re-estimated on each draw. Panel C reports the calibrated parameters. The DTS small-balance threshold is $200,000 through 2020 and $250,000 in 2021. $^{***}p<0.01$, $^{**}p<0.05$, $^{*}p<0.1$.

The mobility parameters are the exposure margin that scales the per-origination subsidy into aggregate responses. The mover shares are 8.5% and 8.8% over the transition window for high- and low-skilled owner origins and 11.8% and 12.9% for renter origins: renters move more than owners, and the skill differential is small. The estimates $c_H = 0.445$ and $c_L = 0$ imply an asymmetry in tenure persistence. High-skilled households persist in their origin tenure beyond what prices, wages, and amenities explain, while the tenure outcomes of credit-constrained households are fully accounted for by the observable market terms and the approval margin.

The lender wedges, the second group of estimates, are uniformly small. The above-limit bin returns $\Delta_{\text{Jumbo}} \approx 3.7$ basis points, which we read as an effective shift in the lender's margin in that bin rather than as a GSE concession, since most of its loans cannot be delivered to the GSEs and its applicant pool changes with the program (Section 4.3). The smallness also reflects the high density of marginal applicants at the lender's profit threshold. Per percentage point of approval gained, the three bins require similar relief, so the conforming wedge is the largest because its approval response is the largest, not because its applicants are harder to flip.[36]

Placing the funding wedge and user cost reduction side by side yields the paper's central magnitude, because the two are the cost and the value of the same approval. An execution wedge of a few basis points accompanies a user-cost reduction of roughly 37% for the marginal borrower it admits, roughly two orders of magnitude larger. The gap has two implications. First, mortgage credit in these markets was substantially rationed: previously rejected applicants gain far more from approval at unchanged terms than the premium required to induce the approval. Second, the policy operates through the access channel, so essentially all of τ reflects the value of approval itself rather than cheaper credit. Both im-

[36] In institutional terms, a wedge of three to five basis points per year is what a delivery premium of a quarter of a point on the loan amount, amortized over a five-year expected life, is worth to the lender. The premium industry participants describe for qualifying rural loans (Section 2) is of this order, whereas a waiver of the posted loan-level price adjustments, which did not apply in our window, would be one to two points upfront and several times larger.

plications are testable inside the estimated model. Section 5.6 first prices who collects this value across household groups, and the counterfactuals of Section 5.7 then deliver the same concession through contract rates instead of the approval margin, at the same concession per loan and at the same per-buyer benefit, to test the access-channel reading directly.

5.6 Welfare Analysis

Our next step is to study how the gains are distributed among new buyers, renters, and incumbent homeowners. For each household type we compute the equivalent variation implied by the mover-stayer structure:

$$EV_z = \frac{1}{\beta_s}\left[(1-\lambda_{s,k_0})\Delta \log D_z^S + \lambda_{s,k_0}\Delta \log D_z^M\right] + \mathbf{1}\{k_0 = O\}\,\eta\,\theta^{s,O}\,\Delta \log P_{g_0}^O, \tag{23}$$

where D_z^S and D_z^M are the stayer and mover logsums, so a type's welfare change mixes its stayer and mover exposure at the mover share.[37] The final term is the ownership dividend: the capital gain on an incumbent owner's baseline dwelling, expressed as an annual flow by annuitizing it at the user-cost rate so that it sits on the same footing as the other terms. Annual owner housing expenditure is $\theta^{s,O}y = uP$, so the annuity of a log price change $\Delta \log P$ is $\theta^{s,O}\Delta \log P$ of income per year, and the lump-sum gain is $1/u$, here 25, times that amount. The fraction $\eta \in [0,1]$ is the share of the annuitized gain the owner realizes through refinancing, extraction, or sale; $\eta = 1$ is the upper bound at which the owner captures the full annuity, and for an owner who stays put it is also the case in which the gain offsets the higher cost of the owner's own housing services.

Table 8 decomposes each group's gain into its channels, the subsidy, housing costs, wages,

[37] Because wages enter utility in logs, the logsum term is a log-income equivalent: a value G corresponds to a proportional income change of $e^G - 1$, which the tables' percent-of-income entries approximate to first order. The approximation is immaterial below a few percent and understates the largest entries by about two points; the per-induced-buyer relief of 19.17 and 10.58 log points is 21.1 and 11.2 percent of income in levels, or about \$6,100 and \$5,700 in the origination year.

congestion, the ownership dividend, and the reallocation gain from re-optimizing, so every mechanism discussed below is a column of the table. At the calibrated $\gamma_{\text{pass}} = 0$ baseline the rate channel is zero and the Subsidy column is entirely the access channel.

Table 8: **Welfare Incidence by Group, with Channel Decomposition**

Group	**Total** *EV (%)*	**Subsidy**	**Housing** *Cost*	**Wage**	**Congestion**	**Ownership** *Dividend*	**Reallocation**	**EV** *($/yr)*
Low-skilled New buyers	9.72	10.16	-0.36	0.00	-0.09	0.00	0.00	2,821
High-skilled New buyers	2.27	2.64	-0.18	-0.01	-0.18	0.00	0.00	1,152
Low-skilled Incumbent renters	0.73	0.11	0.36	-0.00	-0.15	0.00	0.41	211
High-skilled Incumbent renters	0.31	0.06	0.37	-0.00	-0.16	0.00	0.03	157
Low-skilled Incumbent owners	0.05	0.07	-0.18	-0.01	-0.16	0.36	-0.03	14
High-skilled Incumbent owners	-0.03	0.09	-0.16	-0.01	-0.16	0.18	0.03	-17
Incumbent aggregate	0.13	0.09	-0.02	-0.01	-0.16	0.15	0.08	44

Notes: The table presents the equivalent variation by household group, in percent of annual income, decomposed into channels. Reallocation is the exact change in the logsum minus the sum of the first-order channel terms, so the channels add up to the Total column. Incumbent rows are defined by origin and cover households that began the period in a treated tract, whether they stayed or moved. New-buyer rows are defined by destination and cover movers who take a subsidized purchase mortgage in a treated tract. For new buyers, the Subsidy entry is the full relief $\theta^{s,O}\tau$ multiplied by the induced share, defined as one minus the ratio of baseline to post-policy new-buyer flow in treated tracts; the induced share is 52% for low-skilled and 24% for high-skilled buyers, and purchasers who would have originated at baseline receive none of the relief. Per induced buyer, the gain is 19.17% (low-skilled) and 10.58% (high-skilled) of income in the origination year, the full relief net of the general-equilibrium terms. The Ownership Dividend is the annuity, at the user-cost rate of 0.04, of the capital gain on an incumbent owner's baseline dwelling, equal to $\theta^{s,O}\Delta \log P$ of income per year and realized at fraction $\eta = 1$; the lump-sum gain is 25 times the annual entry. At $\eta = 0.5$ the dividend halves and both incumbent-owner totals turn slightly negative; no other group moves. The small positive Subsidy entries for incumbent groups are the option value of a subsidized purchase for the minority who move. The aggregate row is the population-weighted incumbent total; new buyers are excluded from it because the two groups are defined on different bases and overlap. Dollar columns convert totals at mean annual earnings by skill in treated tracts over 2013–2016, in 2010 dollars: $50,743 for high-skilled and $29,030 for low-skilled workers.

Table 8 shows that the surplus goes to the households the policy admits. A buyer whose origination is induced by the policy gains close to a fifth of a year's income if credit-constrained and about a tenth if not; averaged over all subsidized originations, with the induced share as the weight, the gain is 9.7% and 2.3% of income in the origination year.[38] Nearly all of it is the subsidy itself, and the offsets from housing costs and congestion are small. Renters gain a fraction of a percent of income, from lower rents as the rental margin of arrivals thins and from the new option to buy with the subsidy, which for credit-constrained

[38]The induced share, the fraction of post-policy new originations that would not have occurred at baseline flows, is 52% for low-skilled and 24% for high-skilled buyers; purchasers who would have originated at baseline are assigned none of the relief.

renters is the larger part. Incumbent owners are approximately unaffected: higher prices raise the value of the dwelling they hold and, for an owner who stays put, the cost of the housing services they consume by about as much, so the annuitized dividend nets out against housing costs and congestion. With supply absorbing the demand shock in quantities, the asset-side transfer to incumbents that capitalization produces in supply-constrained settings is largely absent.

In dollars, a subsidized origination is worth about $2,800 in the origination year to a low-skilled buyer and $1,200 to a high-skilled buyer, and the full relief per induced buyer is about $5,600 and $5,400; the aggregate incumbent gain is $47 million per year across treated tracts, about $44 per household.

5.7 Counterfactual Analysis

We conduct four counterfactuals to ask how the welfare a credit-supply policy delivers depends on how the credit is delivered and where it lands, holding the estimated credit shock fixed. The first two deliver the same support through the price of credit rather than through approval, at the same concession per loan and at the same per-buyer benefit; the last two move the program into less elastic housing markets, removing the estimated rural supply premium ($\xi = 1$) and imposing the supply conditions of a constrained urban market ($\xi = 0.2$). Table 9 reports the market responses and the welfare incidence by household group; Figure 14 traces the incidence continuously across the supply regime.

The first counterfactual changes the delivery margin at the same concession per loan. We hold the execution premium Δ fixed and pass it entirely to the contract rate ($\gamma_{\text{pass}} = 1$) instead of the approval threshold, on every mortgage in treated tracts, since a rate cut cannot be aimed at the applicants near the threshold.[39] The owner-stock response disappears, and

[39]Paying the same concession on every mortgage costs nearly nine times the access channel's outlay: the concession reaches the roughly 96,000 qualifying new originations per period, while the rate cut reaches the roughly 840,000 mortgages held by treated owners and new buyers. Holding the outlay equal instead spreads it to about 0.01% of income per household. The rate version is inert under either accounting. The horizon cancels, since the execution premium and the rate cut are both annualized per unit of principal over the life

Table 9: **Counterfactual Policy Experiments**

		Counterfactual Scenarios			
	Baseline	Rate, same concession ($\gamma_{\text{pass}} = 1$)	Rate, equal benefit ($\psi = 1$)	No rural premium ($\xi = 1$)	Inelastic supply ($\xi = 0.2$)
Panel A: Market outcomes (DiD, %)					
Per-unit price	0.87	0.00	0.87	1.90	7.68
Rent	−1.37	0.00	−1.37	−1.33	−1.13
Owner-occupied units	3.70	0.00	3.70	3.60	3.09
Panel B: Welfare incidence (EV, % of income)					
Low-skilled New Buyer	9.72	0.11	9.72	9.06	5.55
High-skilled New Buyer	2.27	0.11	2.27	1.97	0.34
Low-skilled Incumbent	0.05	0.11	17.95	0.09	0.32
High-skilled Incumbent	−0.03	0.11	9.99	−0.01	0.12
Low-skilled Renter	0.73	0.00	0.73	0.70	0.58
High-skilled Renter	0.31	0.00	0.31	0.30	0.25

Notes: The table presents model-implied difference-in-differences responses under the estimated parameters (Baseline) and four counterfactuals. Panel A reports market outcomes in percent. Panel B reports the equivalent variation in percent of annual income by household group, defined by skill (high-skilled for high credit access, low-skilled for credit-constrained), incumbency, and tenure; new-buyer rows are credited to policy-induced originations as in Table 8, and valued per induced buyer the baseline entries are 19.17 (low-skilled) and 10.58 (high-skilled) percent. The two rate counterfactuals deliver the policy through the borrower's contract rate instead of the approval margin. "Rate, same concession" holds the GSE concession per loan Δ fixed and passes it fully to the contract rate on every mortgage in treated tracts ($\gamma_{\text{pass}} = 1$); the transfer is worth $\bar{\Delta}\, L/y \approx 0.11$ percent of income to each treated owner and new buyer and costs about nine times the outlay of the access channel, and at equal outlay it is about 0.01 percent of income per household. "Rate, equal benefit" holds the household user-cost reduction τ fixed and delivers it as a rate cut to all owners ($\psi = 1$); market outcomes are unchanged to the reported precision, and the incumbent gains are a rate transfer on existing mortgages, of which about 88% accrues to inframarginal owners who would have owned regardless. "No rural premium" sets the rural owner-supply elasticity multiplier to one ($\xi = 1$), so treated tracts keep the Saiz baseline elasticity of about 2.4. "Inelastic supply" sets the multiplier to $\xi = 0.2$, an effective elasticity of about 0.5, comparable to the most supply-constrained metropolitan areas.

the gain to a low-skilled new buyer collapses from 9.7% of income to about ten basis points, which is the rate transfer itself. In magnitude, the rate channel can transmit at most the concession, about ten basis points of income or 25 to 45 basis points of user cost, while approval delivers roughly 37% of user cost to the rationed marginal borrower. However, rationing binds on quantity rather than price, so a rate cut helps borrowers who already clear the threshold but never reaches the applicant who is denied at any rate.

The second counterfactual changes the delivery margin at equal benefit. We ask what happens if the policy delivers the same user-cost reduction τ as a rate cut. Because the policy enters the model as a user-cost reduction on owning, the equilibrium is essentially invariant to the delivery margin, but what moves is the distribution of welfare. A rate cut is a subsidy that incumbents collect on their existing mortgages: low-skilled incumbent-owner welfare jumps from near zero to 18.0% of income and high-skilled to 10.0%, new buyers and renters are unchanged, and 88% of the spending leaks to inframarginal incumbents. Matching the access channel through the rate margin would require a contract-rate cut of about 5 percentage points for low-skilled and 3 percentage points for high-skilled buyers against the 3.5 basis points the concession can pass, which is unrealistic. Therefore, rate cuts are more expensive and poorly targeted.

The last two counterfactuals begin moving the program across the supply regime. Setting $\xi = 1$ removes the estimated rural supply premium, so treated tracts keep the Saiz baseline elasticity ($\mu \approx 2.4$), still an elastic market rather than a supply-constrained one. Prices rise by twice as much and the per-unit capitalization share doubles from 1.9% to 4.2%, yet the incidence barely moves. Capitalization responds to the supply regime immediately, but within elastic markets the access result remains robust. By contrast, setting $\xi = 0.2$, an effective elasticity of about 0.5, the same program stops producing access and starts producing capitalization. The per-unit capitalization share rises from 1.9% to 16.8%, the owner-stock response falls, and the incidence rotates toward incumbents. The causal gain of

of the loan.

a high-skilled new buyer is nearly wiped out, and the low-skilled gain roughly halves. Renter gains also shrink, and incumbent owners flip from approximately zero to clear gains. Aimed at elastic rural markets, the mandate buys access for rationed buyers; aimed at a constrained urban market, the same mandate would buy asset appreciation for incumbents.

Figure 14 traces the rotation continuously across eligible-area supply elasticities and locates the incidence reversal by skill. For high-skilled households the causal new-buyer and incumbent-owner gains cross at an eligible-area elasticity of about 0.43, inside the supply-constrained range in which the mortgage interest deduction fails to raise homeownership (Hilber and Turner, 2014). For the low-skilled group there is no crossing anywhere on the grid, as the causal new-buyer gain exceeds the incumbent gain at every elasticity down to the grid floor. The asymmetry suggests that even in supply-constrained markets the approval margin delivers more to the rationed entrant than capitalization delivers to the incumbent, while for unconstrained households the ranking flips.

The counterfactuals answer when a mortgage-credit policy reaches the credit-constrained. It should relieve the rationing constraint through the extensive margin, or the support is inert or leaks to incumbents. And it needs to land where housing supply is elastic, or the access it provides capitalizes into prices. The estimated DTS policy clears both conditions, which is why the gains concentrate on the credit-constrained margin the program was written for.

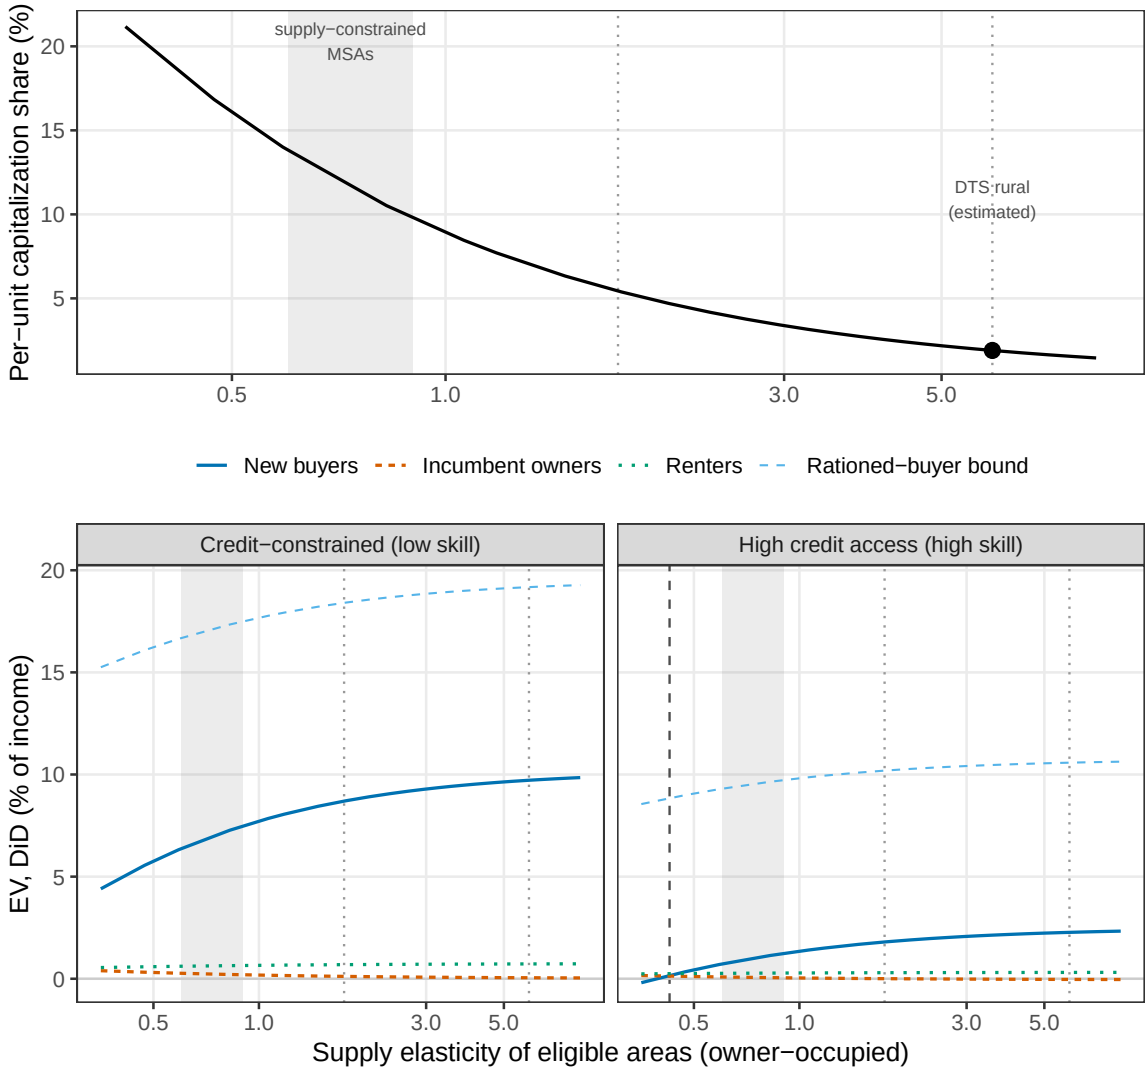


Figure 14: Supply Elasticity as the Regime Variable for Welfare Incidence

Notes: The figure plots the estimated model across the housing supply elasticity of eligible areas, holding the credit shock and all other parameters at their estimated values. The horizontal axis is the mean Saiz baseline elasticity of treated tracts multiplied by the rural premium ξ, on a log scale. The top panel plots the capitalization share, defined as the model-implied price change relative to the effective subsidy. The bottom panels plot the equivalent variation by household group, separately for low-skilled (credit-constrained) and high-skilled (high credit access) households; the solid new-buyer line is the average over policy-induced originations, and the light dashed line is the full relief per induced buyer, its upper bound. The dashed vertical line in the high-skilled panel marks the elasticity of 0.43 at which the new-buyer and incumbent-owner gains are equal; in the low-skilled panel the new-buyer gain exceeds the incumbent-owner gain at every elasticity on the grid, so no such point exists. The gray band marks the 0.6–0.9 range of supply-constrained metropolitan areas, and the dotted vertical lines mark the metropolitan median of 1.75 in Saiz (2010) and the elasticity of the program as estimated. New-buyer gains remain positive throughout, so the crossing marks where incidence tilts toward incumbents, not where entrants lose.

6 Conclusion

When the economics of originating a loan improve, borrowers need not see a lower price; they may instead see loans they would otherwise have been denied. The Duty to Serve concession never reached the price of credit: contract rates did not move, and lenders kept the execution premium as margin. Yet the margin bought access. Approvals rose, low- and moderate-income owner-occupants entered with smaller loans, and distress did not rise. As a result, homeownership grew while house prices stayed flat. The surplus is large for buyers induced into ownership, modest for renters, and approximately zero for incumbent owners.

The delivery margin is half of the reason. Mortgage credit in these markets is rationed on quantity rather than price, and the shadow price of relieving the rationing is roughly two orders of magnitude above the lender-side execution wedge that accompanies it. Delivered through contract rates instead, the same support is nearly inert at the same concession per loan, and at the same per-buyer benefit it leaks overwhelmingly to inframarginal owners who would have owned regardless. Within the model, a rate cut reaches only borrowers who already clear the approval threshold, never the applicant whose constraint is the approval itself. DTS works because it improves the economics of the approval decision directly.

Elastic housing supply is the other half. The new demand became owner-occupied homes rather than higher prices. Moved to the supply conditions of a constrained urban market in our counterfactuals, the program changes character: the capitalization share rises from under 2% to 17% of the per-origination subsidy, the marginal buyer's gain erodes, and incumbent owners become clear beneficiaries. Evidence pooled across supply regimes therefore mixes two qualitatively different outcomes, access and capitalization, and mischaracterizes what the policy achieves in each.

Three open questions follow. Our estimates compare steady states over a medium-run window, so the long-run incidence, including default behavior over a full credit cycle, remains open. The shadow price we estimate is specific to mortgage markets at the DTS boundary, so whether rationing relief is comparably valuable in other rationed credit markets, from

small-business lending to the GSEs' other underserved-market duties in manufactured housing and affordable-housing preservation, is an empirical question the joint lender-household framework can take on. And translating the supply-conditioning principle into an eligibility formula, anticipating how lenders and developers would respond to one, is a policy-design problem beyond this paper. As policy returns to housing affordability, whether a credit expansion buys homeowners or house prices is decided by design: by the margin through which the support is delivered, and by the markets where it lands.

References

Acolin, Arthur, Jesse Bricker, Paul Calem, and Susan Wachter, "Borrowing Constraints and Homeownership," *American Economic Review*, 2016, *106* (5), 625–629.

Adelino, Manuel, Antoinette Schoar, and Felipe Severino, "Credit Supply and House Prices: Evidence from Mortgage Market Segmentation," NBER Working Paper 17832, National Bureau of Economic Research 2012.

_ , _ , and _ , "Credit Supply and House Prices: Evidence from Mortgage Market Segmentation," *Journal of Financial Economics*, 2025, *163*, 103958.

Agarwal, Sumit, Efraim Benmelech, Nittai Bergman, Bernardo Morais, and Amit Seru, "Did the Community Reinvestment Act (CRA) Lead to Risky Lending?," NBER Working Paper 18609, National Bureau of Economic Research 2012.

_ , Souphala Chomsisengphet, Neale Mahoney, and Johannes Stroebel, "Do Banks Pass through Credit Expansions to Consumers Who Want to Borrow?," *The Quarterly Journal of Economics*, 2018, *133* (1), 129–190.

Almagro, Milena and Tomás Domínguez-Iino, "Location Sorting and Endogenous Amenities: Evidence From Amsterdam," *Econometrica*, 2025, *93* (3), 1031–1071.

Ambrose, Brent W., Anthony Pennington-Cross, and Anthony M. Yezer, "Credit Rationing in the U.S. Mortgage Market: Evidence from Variation in FHA Market Shares," *Journal of Urban Economics*, 2002, *51* (2), 272–294.

An, Xudong, Raphael W. Bostic, Yongheng Deng, and Stuart A. Gabriel, "GSE Loan Purchases, the FHA, and Housing Outcomes in Targeted, Low-Income Neighborhoods," in "Brookings-Wharton Papers on Urban Affairs" 2007, pp. 205–256.

Anderson, Donald G. M., "Iterative Procedures for Nonlinear Integral Equations," *Journal of the ACM*, 1965, *12* (4), 547–560.

Anenberg, Elliot, Aurel Hizmo, Edward Kung, and Raven Molloy, "Measuring Mortgage Credit Availability: A Frontier Estimation Approach," *Journal of Applied Econometrics*, 2019, *34* (6), 865–882.

Baum-Snow, Nathaniel and Gilles Duranton, "Housing supply and housing affordability," Technical Report, National Bureau of Economic Research 2025.

Berry, Steven, James Levinsohn, and Ariel Pakes, "Automobile Prices in Market Equilibrium," *Econometrica*, 1995, *63* (4), 841–890.

Bhutta, Neil, "The Community Reinvestment Act and Mortgage Lending to Lower Income Borrowers and Neighborhoods," *Journal of Law and Economics*, 2011, *54* (4), 953–983.

_ , "GSE Activity and Mortgage Supply in Lower-Income and Minority Neighborhoods: The Effect of the Affordable Housing Goals," *Journal of Real Estate Finance and Economics*, 2012, *45* (1), 238–261.

_ **and Daniel Ringo**, "The effect of interest rates on home buying: Evidence from a shock to mortgage insurance premiums," *Journal of Monetary Economics*, 2021, *118*, 195–211.

Bostic, Raphael W. and Stuart A. Gabriel, "Do the GSEs Matter to Low-Income Housing Markets? An Assessment of the Effects of the GSE Loan Purchase Goals on California Housing Outcomes," *Journal of Urban Economics*, 2006, *59* (3), 458–475.

Brevoort, Kenneth P., "Reexamining the 'Role of the Community Reinvestment Act in Mortgage Supply and the U.S. Housing Boom'," Finance and Economics Discussion Series 2024-009, Board of Governors of the Federal Reserve System 2024.

Campbell, John Y. and João F. Cocco, "A Model of Mortgage Default," *Journal of Finance*, 2015, *70* (4), 1495–1554.

Carozzi, Felipe, Christian A. L. Hilber, and Xiaolun Yu, "On the Economic Impacts of Mortgage Credit Expansion Policies: Evidence from Help to Buy," *Journal of Urban Economics*, 2024, *139*, 103611.

Combes, Pierre-Philippe, Gilles Duranton, Laurent Gobillon, Diego Puga, and Sébastien Roux, "The Productivity Advantages of Large Cities: Distinguishing Agglomeration from Firm Selection," *Econometrica*, 2012, *80* (6), 2543–2594.

Conway, Jacob, Jack N. Glaser, and Matthew C. Plosser, "Does the Community Reinvestment Act Improve Consumers' Access to Credit?," Staff Report 1048, Federal Reserve Bank of New York 2023.

DeFusco, Anthony A. and Andrew Paciorek, "The Interest Rate Elasticity of Mortgage Demand: Evidence from Bunching at the Conforming Loan Limit," *American Economic Journal: Economic Policy*, 2017, *9* (1), 210–240.

_ **, Charles G. Nathanson, and Eric Zwick**, "Speculative dynamics of prices and volume," *Journal of Financial Economics*, 2022, *146* (1), 205–229.

_ **, Stephanie Johnson, and John Mondragon**, "Regulating Household Leverage," *Review of Economic Studies*, 2020, *87* (2), 914–958.

Dekle, Robert, Jonathan Eaton, and Samuel Kortum, "Unbalanced trade," *American Economic Review*, 2007, *97* (2), 351–355.

Di Maggio, Marco and Amir Kermani, "Credit-Induced Boom and Bust," *Review of Financial Studies*, 2017, *30* (11), 3711–3758.

Diamond, Rebecca, "The Determinants and Welfare Implications of US Workers' Diverging Location Choices by Skill: 1980–2000," *American Economic Review*, 2016, *106* (3), 479–524.

Duca, John V. and Stuart S. Rosenthal, "An Empirical Test of Credit Rationing in the Mortgage Market," *Journal of Urban Economics*, 1991, *29* (2), 218–234.

Elenev, Vadim, Tim Landvoigt, and Stijn Van Nieuwerburgh, "Phasing Out the GSEs," *Journal of Monetary Economics*, 2016, *81*, 111–132.

Fannie Mae, "Duty to Serve Rural Housing Loan Product Report, Housing in High-Needs Rural Regions: Annual 2018," Narrative report filed with the Federal Housing Finance Agency, 2018. https://www.fhfa.gov/sites/default/files/reports_11_23/Enterprise%20quarterly%20and%20annual%20reports/2018-DTS-Reports/FannieMae/FNM_Rural_HN_Regions_2_2018.pdf.

_ , "Duty to Serve Rural Housing Loan Purchase Reports, Financing by Small Financial Institutions: Second Quarter, Third Quarter, and Annual 2018," Narrative report filed with the Federal Housing Finance Agency, 2018. https://www.fhfa.gov/sites/default/files/reports_11_23/Enterprise%20quarterly%20and%20annual%20reports/2018-DTS-Reports/FannieMae/FNM_Rural_Small_Fin_Inst_1_2018.pdf.

_ , "Duty to Serve Rural Housing Loan Purchase Report, Financing by Small Financial Institutions: Annual 2019," Narrative report filed with the Federal Housing Finance Agency, 2019. https://www.fhfa.gov/sites/default/files/reports_11_23/Enterprise%20quarterly%20and%20annual%20reports/2019-DTS-Reports/Fannie%20Mae/RH_Small_Financial_Institutions_Loan_Purchase_2019.pdf.

_ , "Duty to Serve Underserved Markets Plan 2018–2021," Modified plan, Federal Housing Finance Agency redline copy, effective December 20, 2019. https://www.fhfa.gov/sites/default/files/documents/FannieMaeDTS%20Plan_2018-2021_Redline.pdf.

_ , "Loan-Level Price Adjustment (LLPA) Matrix," Fannie Mae Single-Family Selling Guide resource, version effective January 28, 2026. https://singlefamily.fanniemae.com/media/9391/display.

Favara, Giovanni and Jean Imbs, "Credit Supply and the Price of Housing," *American Economic Review*, 2015, *105* (3), 958–992.

Federal Housing Finance Agency, "Enterprise Duty to Serve Underserved Markets; Final Rule," 81 Fed. Reg. 96242 (December 29, 2016), codified at 12 C.F.R. Part 1282 2016.

_ , "FHFA Announces Targeted Pricing Changes to Enterprise Pricing Framework," Press release, October 24, 2022. https://www.fhfa.gov/news/news-release/fhfa-announces-targeted-pricing-changes-to-enterprise-pricing-framework.

Federal Housing Finance Agency Office of Inspector General, "An Overview of Enterprise Debt-to-Income Ratios," White paper WPR-2019-002, March 27, 2019. https://www.fhfaoig.gov/sites/default/files/WPR-2019-002.pdf.

Fetter, Daniel K., "How Do Mortgage Subsidies Affect Home Ownership? Evidence from the Mid-century GI Bills," *American Economic Journal: Economic Policy*, 2013, *5* (2), 111–147.

Freddie Mac, "Duty to Serve Underserved Markets Plan for 2018–2020," Underserved Markets Plan submitted to the Federal Housing Finance Agency, December 2017. https://ruralhome.org/wp-content/uploads/storage/documents/other_docs/freddie-dts-2017.pdf.

_ , "Duty to Serve Rural Housing Loan Product Report, Housing in High-Needs Rural Regions: Annual 2019," Narrative report filed with the Federal Housing Finance Agency, 2019. https://www.fhfa.gov/sites/default/files/reports_11_23/Enterprise%20quarterly%20and%20annual%20reports/2019-DTS-Reports/Freddie%20Mac/Rural_HN_Regions_Loan_Product_2019.pdf.

_ , "Single-Family Seller/Servicer Guide Bulletin 2022-22," Bulletin dated October 31, 2022. https://guide.freddiemac.com/app/servicing/bulletin/2022-22.

Fuster, Andreas and Basit Zafar, "The Sensitivity of Housing Demand to Financing Conditions: Evidence from a Survey," *American Economic Journal: Economic Policy*, 2021, *13* (1), 231–265.

_ **and James Vickery**, "Securitization and the Fixed-Rate Mortgage," *Review of Financial Studies*, 2015, *28* (1), 176–211.

_ **, Stephanie H. Lo, and Paul S. Willen**, "The Time-Varying Price of Financial Intermediation in the Mortgage Market," *Journal of Finance*, 2024, *79* (4), 2553–2602.

Gete, Pedro and Franco Zecchetto, "Distributional Implications of Government Guarantees in Mortgage Markets," *Review of Financial Studies*, 2018, *31* (3), 1064–1097.

_ **and Michael Reher**, "Mortgage Supply and Housing Rents," *The Review of Financial Studies*, 2018, *31* (12), 4884–4911.

Glaeser, Edward and Joseph Gyourko, "The Economic Implications of Housing Supply," *Journal of Economic Perspectives*, 2018, *32* (1), 3–30.

Greenwald, Daniel L., "The Mortgage Credit Channel of Macroeconomic Transmission," 2018. MIT Sloan Research Paper No. 5184-16.

_ **and Adam Guren**, "Do Credit Conditions Move House Prices?," *American Economic Review*, 2025, *115* (10), 3559–3596.

Grundl, Serafin and You Suk Kim, "The Marginal Effect of Government Mortgage Guarantees on Homeownership," *Journal of Monetary Economics*, 2021, *119*, 75–89.

Gyourko, Joseph and Raven Molloy, "Regulation and Housing Supply," in Gilles Duranton, J. Vernon Henderson, and William C. Strange, eds., *Handbook of Regional and Urban Economics*, Vol. 5, Elsevier, 2015, pp. 1289–1337.

_ **, Jonathan S. Hartley, and Jacob Krimmel**, "The Local Residential Land Use Regulatory Environment across US Housing Markets: Evidence from a New Wharton Index," *Journal of Urban Economics*, 2021, *124*, 103337.

Hilber, Christian A. L. and Tracy M. Turner, "The Mortgage Interest Deduction and Its Impact on Homeownership Decisions," *Review of Economics and Statistics*, 2014, *96* (4), 618–637.

Jeske, Karsten, Dirk Krueger, and Kurt Mitman, "Housing, Mortgage Bailout Guarantees and the Macro Economy," *Journal of Monetary Economics*, 2013, *60* (8), 917–935.

Jiménez, Gabriel, Steven Ongena, José-Luis Peydró, and Jesús Saurina, "Credit Supply and Monetary Policy: Identifying the Bank Balance-Sheet Channel with Loan Applications," *American Economic Review*, 2012, *102* (5), 2301–2326.

Kaplan, Greg, Kurt Mitman, and Giovanni L. Violante, "The Housing Boom and Bust: Model Meets Evidence," *Journal of Political Economy*, 2020, *128* (9), 3285–3345.

Keys, Benjamin J., Tanmoy Mukherjee, Amit Seru, and Vikrant Vig, "Did Securitization Lead to Lax Screening? Evidence from Subprime Loans," *The Quarterly Journal of Economics*, 2010, *125* (1), 307–362.

Landvoigt, Tim, Monika Piazzesi, and Martin Schneider, "The Housing Market(s) of San Diego," *American Economic Review*, 2015, *105* (4), 1371–1407.

Lehnert, Andreas, S. Wayne Passmore, and Shane M. Sherlund, "GSEs, Mortgage Rates, and Secondary Market Activities," *Journal of Real Estate Finance and Economics*, 2008, *36* (3), 343–363.

Loutskina, Elena and Philip E. Strahan, "Securitization and the Declining Impact of Bank Finance on Loan Supply: Evidence from Mortgage Originations," *Journal of Finance*, 2009, *64* (2), 861–889.

McFadden, Daniel, "Conditional logit analysis of qualitative choice behavior," in Paul Zarembka, ed., *Frontiers in Econometrics*, Academic Press, 1974, pp. 105–142.

_ , "Modelling the Choice of Residential Location," in A. Karlqvist, F. Snickars, and J. Weibull, eds., *Spatial Interaction Theory and Planning Models*, Amsterdam: North-Holland, 1978, pp. 75–96.

Mian, Atif, Amir Sufi, and Emil Verner, "Household Debt and Business Cycles Worldwide," *Quarterly Journal of Economics*, 2017, *132* (4), 1755–1817.

_ **and** _ , "The Consequences of Mortgage Credit Expansion: Evidence from the U.S. Mortgage Default Crisis," *Quarterly Journal of Economics*, 2009, *124* (4), 1449–1496.

_ **and** _ , "House Prices, Home Equity-Based Borrowing, and the US Household Leverage Crisis," *American Economic Review*, 2011, *101* (5), 2132–2156.

Molloy, Raven, "The Effect of Housing Supply Regulation on Housing Affordability: A Review," *Regional Science and Urban Economics*, 2020, *80*, 103350.

Monte, Ferdinando, Stephen J. Redding, and Esteban Rossi-Hansberg, "Commuting, Migration, and Local Employment Elasticities," *American Economic Review*, December 2018, *108* (12), 3855–3890.

Mota, Nuno, "An Analysis of Small Balance Loan Origination in Rural and High-Needs Rural Areas," Technical Report, Fannie Mae, Economic and Strategic Research December 2019.

Orlando, Anthony W. and Gerd Welke, "Borrowing on the Wrong Side of the Tracks: Evidence from Mortgage Loan Discontinuities," *Journal of Banking and Finance*, 2025, *176*, 107438.

Passmore, Wayne, "The GSE Implicit Subsidy and the Value of Government Ambiguity," *Real Estate Economics*, 2005, *33* (3), 465–486.

Poterba, James M., "Tax Subsidies to Owner-Occupied Housing: An Asset-Market Approach," *The Quarterly Journal of Economics*, 1984, *99* (4), 729–752.

Puri, Manju, Jörg Rocholl, and Sascha Steffen, "Global Retail Lending in the Aftermath of the US Financial Crisis: Distinguishing between Supply and Demand Effects," *Journal of Financial Economics*, 2011, *100* (3), 556–578.

Quigley, John M. and Steven Raphael, "Is Housing Unaffordable? Why Isn't It More Affordable?," *Journal of Economic Perspectives*, 2004, *18* (1), 191–214.

Ringo, Daniel, "Mortgage Lending, Default, and the Community Reinvestment Act," *Journal of Money, Credit and Banking*, 2023, *55* (1), 77–102.

Roback, Jennifer, "Wages, Rents, and the Quality of Life," *Journal of Political Economy*, 1982, *90* (6), 1257–1278.

Rosen, Sherwin, "Wage-Based Indexes of Urban Quality of Life," in "Current Issues in Urban Economics" 1979, pp. 74–104.

Saadi, Vahid, "Role of the Community Reinvestment Act in Mortgage Supply and the U.S. Housing Boom," *Review of Financial Studies*, 2020, *33* (11), 5288–5332.

Saiz, Albert, "The Geographic Determinants of Housing Supply," *Quarterly Journal of Economics*, 2010, *125* (3), 1253–1296.

Scharfstein, David S. and Adi Sunderam, "Market Power in Mortgage Lending and the Transmission of Monetary Policy," *Harvard Business School Working Paper*, 2016.

Sherlund, Shane M., "The Jumbo-Conforming Spread: A Semiparametric Approach," Finance and Economics Discussion Series 2008-01, Board of Governors of the Federal Reserve System 2008.

Sommer, Kamila and Paul Sullivan, "Implications of US Tax Policy for House Prices, Rents, and Homeownership," *American Economic Review*, 2018, *108* (2), 241–274.

Stiglitz, Joseph E. and Andrew Weiss, "Credit Rationing in Markets with Imperfect Information," *American Economic Review*, 1981, *71* (3), 393–410.

United States Congress, "21st Century ROAD to Housing Act," H.R. 6644, 119th Congress, enacted July 11, 2026.

Walker, Homer F. and Peng Ni, "Anderson Acceleration for Fixed-Point Iterations," *SIAM Journal on Numerical Analysis*, 2011, *49* (4), 1715–1735.

A Online Appendix: Institutional Detail and Other Federal Mortgage-Credit Programs

DTS is among the newest federal programs directing mortgage credit toward underserved households and among the least studied. The final rule dates to December 2016, and the prior record consists of FHFA's annual evaluations of the Enterprises' plans and Fannie Mae's descriptive analysis of small-balance rural lending (Mota, 2019). Three design features set it apart from the programs below. Eligibility follows a published tract-level formula that creates a sharp geographic boundary; the targets are loan counts delivered through the GSEs' purchase execution, so the instrument works on the number of households served; and the program operates in rural markets with elastic housing supply, whereas most evidence on credit expansions comes from supply-constrained metropolitan settings.

A.1 Statutory history and program administration

The DTS program originates in Section 1335 of the Federal Housing Enterprises Financial Safety and Soundness Act of 1992, which directed the government-sponsored enterprises (GSEs) Fannie Mae and Freddie Mac to provide leadership in facilitating a secondary market for mortgages serving very low-, low-, and moderate-income families in underserved markets. The statute remained largely dormant for sixteen years and was substantively activated by the Housing and Economic Recovery Act of 2008 (HERA). HERA placed the GSEs in conservatorship, established the Federal Housing Finance Agency (FHFA) as their independent regulator and conservator, and gave FHFA the authority to issue implementing regulations for DTS. The duty itself is statutory rather than an artifact of conservatorship: it attaches to the Enterprises under Section 1335 of the Safety and Soundness Act, so the GSEs would remain subject to DTS even after an exit from conservatorship. FHFA issued the final rule implementing DTS on December 13, 2016 (Federal Housing Finance Agency, 2016), and the first three-year Underserved Markets Plans took formal effect on January 1, 2018.

Each plan lists objectives under the rural activities the rule names: housing in high-needs rural regions, housing for high-needs rural populations, financing by small financial institutions, and small multifamily rental properties. Every objective carries an evaluation area (outreach, loan product, loan purchase, or investment), and loan-purchase objectives set annual loan-count targets against a baseline of prior purchases (Freddie Mac, 2017; Fannie Mae, 2019b). The Enterprises report to FHFA quarterly and annually, and FHFA rates each market on a quantitative and a qualitative assessment, with extra credit for designated activities, on a five-step scale from Fails to Exceeds; a rating of Fails constitutes noncompliance (Federal Housing Finance Agency, 2016). The duty is enforceable under the same provisions as the housing goals: FHFA may require a housing plan, and only refusal to submit one, an unacceptable plan, or failure to follow it triggers cease-and-desist orders or civil money penalties, so a low rating carries no automatic financial penalty. FHFA found both Enterprises compliant in every year from 2018 to 2021. The count targets bound early in the cycle, with Fannie Mae short of its small-institution target at mid-year 2018 (Fannie Mae, 2018b), and were exceeded by wide margins in 2020 and 2021 as refinancing surged.

Beyond targets and execution, the plans commit the GSEs to direct outreach: engaging small rural lenders, community banks, and community development financial institutions to encourage the origination and delivery of qualifying loans. This outreach channel complements the execution premium by familiarizing lenders with GSE delivery for loan types they had previously retained in portfolio or declined to originate, and outreach activities receive credit in FHFA's evaluation of each plan.

A.2 Other federal mortgage-credit programs

We compare DTS with other federal programs that direct mortgage credit toward underserved households. Each differs from DTS on at least one of the design dimensions discussed in Section 2: the tract-level eligibility formula, the count-based targets delivered through GSE purchase execution, and the elastic-supply rural setting.

GSE Affordable Housing Goals (AHG). The AHG, established by HUD in the early 1990s and continued under FHFA, set percentage targets for the share of each GSE's annual acquisitions serving specified borrower-income, geographic, or property-type categories. Evaluations find little effect on targeted lending or homeownership (Bostic and Gabriel, 2006; An et al., 2007; Bhutta, 2012). A binding share target, like a binding count target, gives a qualifying acquisition a positive shadow value, so the form of the target does not by itself explain the difference in findings. Three features of DTS do. First, the AHG categories were broad and were largely met with loans the Enterprises would have acquired anyway, so the constraint carried little value at the margin (Bhutta, 2012); the DTS counts are set on a thin segment the Enterprises were not otherwise serving at the required volume, and they bound in 2018 and 2019 (Section 2). Second, a share target can be met by reducing non-qualifying acquisitions, whereas a count target is met only by adding qualifying loans. Third, the DTS plans commit the Enterprises to named execution actions that FHFA rates each year, whereas the AHG left the means of compliance to the Enterprises.

Community Reinvestment Act (CRA). The CRA, enacted in 1977, evaluates depository institutions' lending records in the communities where they accept deposits; examination ratings affect approvals of branch openings, mergers, and acquisitions. The CRA therefore works through institution-level supervisory pressure on banks, and evaluations find that this pressure shifts lending volume at the examined institutions, with mixed evidence on loan risk (Agarwal et al., 2012; Saadi, 2020; Ringo, 2023). DTS involves no deposit-side supervision and applies to qualifying loans regardless of the originating institution.

Federal Housing Administration (FHA). The FHA, established in 1934 and now a unit of HUD, insures lenders against borrower default in exchange for a borrower-paid mortgage insurance premium, with underwriting standards (down payment, loan-to-value, debt-to-income) set by HUD. It is a borrower-side insurance program that changes the price of credit for eligible borrowers (Bhutta and Ringo, 2021). DTS charges borrowers nothing and

insures nothing; it changes the lender's secondary-market execution.

USDA Section 502 and rural FHA / VA activity. The U.S. Department of Agriculture lends directly to, and guarantees loans for, qualifying rural borrowers under its Section 502 programs, and the FHA and VA insure or guarantee rural originations; all of these operate outside the GSE secondary market that DTS works through.

B Online Appendix: Estimation Details

This appendix collects the numerical details underlying the joint SMM estimation in Section 5.4. We describe the hat-algebra formulation of the equilibrium with the typed mover-stayer aggregation, the fixed-point solver and the specification gates, the integration of the lender problem over the borrower distribution, the SMM objective, moment construction, and inference, and the moment-to-parameter identification mapping.

B.1 Equilibrium in Hat Algebra

We solve the post-policy equilibrium of Section 5.2 in changes from the observed baseline using the exact hat algebra of Dekle et al. (2007). For every level variable X, let $\hat{X} = X_1/X_0$ denote its multiplicative change. The change of variables substitutes out the unobserved fundamentals (the baseline productivities $\bar{Z}_j$, exogenous amenity levels $\bar{A}_i^s$, and housing stock scale parameters $\bar{H}_{i,k}$), so the post-policy equilibrium can be computed from the observed baseline allocation, the calibrated parameters, and the structural parameters Θ alone. The baseline itself is anchored by Berry inversion: mean utilities δ_n on the grouped cell space are inverted so that the type-aggregated choice system reproduces the observed baseline flow panel exactly, and the inversion is cached and reused across parameter evaluations that hold the mobility block fixed.

The hat-form equilibrium consists of four blocks. Wages adjust to clear segmented labor markets with agglomeration:

$$\hat{W}_j^H = (\hat{N}_j^{\text{Total}})^{\sigma}(\hat{N}_j^L/\hat{N}_j^H)^{1-\alpha}, \qquad \hat{W}_j^L = (\hat{N}_j^{\text{Total}})^{\sigma}(\hat{N}_j^H/\hat{N}_j^L)^{\alpha}. \tag{24}$$

House prices and rents clear tenure-specific housing markets,

$$\hat{P}_{i,k} = (\hat{D}_{i,k})^{1/\mu_{i,k}}, \tag{25}$$

with $\mu_{i,O} = \mu_{i,\text{base}} \times \xi$ in treated tracts and $\mu_{i,\text{base}}$ otherwise. Amenities update with the new population composition,

$$\hat{A}_i^s = (\hat{N}_i^H)^{\alpha_{H|s}} (\hat{N}_i^L)^{\alpha_{L|s}} \big((\hat{N}_i^H + \hat{N}_i^L)/2\big)^{-\delta}, \tag{26}$$

which is the hat-form counterpart of equation (17); with the amenity spillovers restricted to zero the expression reduces to the congestion factor $-\delta$ that captures the strain of population growth on shared infrastructure. Finally, population flows update by the locked mixture aggregation. Movers of each type re-optimize over the full menu via the nested-logit update

$$F_{i,j,k,1}^{M,z} = F_{i,j,k,0}^{M,z} \cdot \frac{\exp(\Delta V_{i,j,k}^s/\rho) \cdot (\hat{I}_{ij}^s)^{\rho-1}}{\sum_{i',j'} \pi_{i',j'|0}^{s,z} (\hat{I}_{i'j'}^s)^{\rho}}, \tag{27}$$

where $\Delta V_{i,j,k}^s$ is the change in the systematic utility component, $\hat{I}_{ij}^s$ the change in the inclusive value of the tenure nest at (i,j), and the policy factor inside ΔV is carved from the own-location owner cells of origin-owner types under new-origination incidence. Stayers update only the workplace dimension on the origin-location, origin-tenure slice, and no stayer menu contains the policy term. Cell-level flows sum the stayer and mover systems across types with weights $1 - \lambda_{s,k_0}$ and λ_{s,k_0}, exactly as in the equilibrium definition of Section 5.2.

B.2 Fixed-Point Solver

Given a candidate parameter vector Θ, the four-block system (24)–(27) defines a map on the state vector $\mathbf{x} = (\hat{L}_j^s, \hat{N}_i^s, \hat{N}_{i,k})$. We solve its fixed point by damped iteration, $\mathbf{x} \leftarrow 0.9\,\mathbf{x} + 0.1\,g(\mathbf{x})$, with Anderson acceleration (Anderson, 1965; Walker and Ni, 2011) available after 1,000 iterations; at the reported estimates every solve converges within about 150 iterations, before the acceleration phase. Following Almagro and Domínguez-Iino (2025), we apply a Picard-Mann damping step to the residential population block after each iteration:

$$(\hat{N}_i^s)^{t+1} = (1 - \lambda_{\text{amen}}) \cdot (\hat{N}_i^s)^t + \lambda_{\text{amen}} \cdot (\hat{N}_i^s)^{t+1,\text{trial}}, \tag{28}$$

with $\lambda_{\text{amen}} = 0.5$. The damping does not alter the fixed point. Convergence is declared at $\max_i |\mathbf{x}_i^{t+1} - \mathbf{x}_i^t| < 10^{-5}$ in the final-precision evaluation, and at 10^{-3} inside the SMM inner loop.

The equilibrium of Section 5.2 need not be globally unique in models with endogenous amenities. Following Almagro and Domínguez-Iino (2025), we initialize the solver at the observed pre-policy allocation as a selection rule, which lets the counterfactual continuation track the observed equilibrium; every scenario and grid point of the counterfactual analysis in Section 5.7 converges under this rule at the final-precision tolerance.

Four specification gates validate the locked mover-stayer implementation against its nesting benchmarks; all four pass at the estimation build.

- *Mover-limit equivalence.* With all mover shares set to one the lock is vacuous and the typed solver must reproduce the unrestricted benchmark cell by cell; the maximum gap is 6.9×10^{-18} against a tolerance of 10^{-7}.

- *Stayer-limit freeze.* As the mover shares approach zero, no menu contains the policy term, so every equilibrium change must equal one and conversions must vanish; the maximum gap is 1.3×10^{-9} against a tolerance of 10^{-5}.

- *Margin mix.* The audited sensitivity of the conversion moments to τ is 2.9 moment standard errors per unit of τ, against roughly 160 in a specification without the tenure lock, in which in-place tenure flips drove the estimator to the boundary. The margins that identify τ under the lock are the arrival contrasts (5.9 standard errors per unit), the owner-quantity DiD (7.5), and the renter-arrival contrast (-3.9, the correct sign).

- *Inversion.* The Berry contraction residual is 10^{-8} and the baseline flow panel is reproduced to 10^{-8} relative error.

B.3 Lender Problem: Numerical Integration

The lender problem in Section 5.3 produces a model-implied Stage 2 approval rate by integrating the per-application approval probability over the borrower-characteristic distribution. For computational tractability, we replace the discrete profit-positive indicator with a logistic-smoothed approximation $P_a(b,m) = (1+\exp(-s\cdot E[\Pi(b,m)/L(b)]))^{-1}$, with $s = 200$. The discrete indicator is recovered as $s \to \infty$; the bin-specific execution wedges Δ_{size} and their cross-bin ordering are stable across $s \in \{100, 200, 500, 1000\}$.

Borrower characteristics (FICO, LTV, DTI) are drawn as independent truncated normals within each (skill, rural, size bin) cell. The LTV and DTI moments of each cell are calibrated from HMDA applicants in 2018 and later, the years in which HMDA reports both ratios, and the FICO moments by loan-size bin from the public Freddie Mac single-family loan-level files. The approval integral within each cell is evaluated by Monte Carlo with 200,000 draws per skill-by-arm cell, using the same draws across parameter evaluations and across the three loan-size bins so that simulation noise does not enter the optimizer's view. Because the lender moments depend only on the execution wedges and the spatial and mobility moments do not, the two blocks separate in the numerical work: spatial perturbations reuse the base lender block, and wedge perturbations skip the equilibrium solver.

B.4 SMM Objective and Moments

Given Θ, the model produces a vector $m_{\text{model}}(\Theta)$ matched to the empirical vector m_{data} by minimizing

$$Q(\Theta) = [m_{\text{data}} - m_{\text{model}}(\Theta)]' W [m_{\text{data}} - m_{\text{model}}(\Theta)], \tag{29}$$

where W is a diagonal weighting matrix with weights $1/\text{SE}^2$ for each moment. The eighteen targeted moments and their sources are as follows.

- *Spatial DiDs* (four; Section 4.5): tract-average owner-occupied log price −0.20 (SE 0.54), log rent −2.43 (0.61), owner-occupied units +4.06 (0.99), rental units −4.10

(1.30), all in percent. Skill-specific resident-population DiDs are not targeted; with amenity spillovers restricted to zero, the arrival contrasts below carry the borrower-composition information.

- *Lender Stage 2 approval DiDs* (three; Section 4.3): +1.09 (0.14), +1.92 (0.32), and +1.64 (0.72) percentage points for the DTS-small, conforming, and jumbo bins.

- *Mobility moments* (eight; Infutor address histories in the boundary sample): treated-arm stay rates by origin tenure and skill, 91.56 (0.24), 91.30 (0.23), 88.20 (0.52), and 87.14 (0.28) percent for owner-H, owner-L, renter-H, and renter-L; and in-band renter-to-owner conversion rates by boundary arm and skill, 3.94 (0.14) and 3.58 (0.14) percent in the control arm and 3.65 (0.17) and 3.73 (0.14) percent in the treated arm for H and L. Both sets are computed from the canonical transition kernel of the address-history panel (2014–2016 origin states to 2018–2020 outcomes), using counts restricted to records with valid observation dates and reweighted for the difference in household-profile coverage between arms (63.6 percent of treated-side and 54.4 percent of control-side addresses carry a profile); standard errors are group-clustered bootstrap standard errors over boundary clusters (500 draws). A conversion is a household of renter origin that ends the period as an owner anywhere in the boundary band, the same event the locked model generates through movers only.

- *Arrival contrasts* (three): relative in-migration contrasts $(T - C)/C$ by skill, +0.365 (0.083) and +0.242 (0.072), from the same transition kernel; and the renter-channel arrival contrast −0.145 (0.045), constructed as $\exp(b) - 1$ from a PPML level DiD of tract-year renter-arrival counts on a balanced tract panel including zeros, with tract and boundary-by-year fixed effects, state-clustered standard errors, and 2016 as the reference year, with the standard error mapped through the delta method; a joint pre-trend test on the pre-period coefficients is clean.

A first stage warm-starts the six mobility parameters (the four mover shares and the

two continuation premia) on the mobility moments alone, holding the spatial parameters and lender wedges fixed, from the empirical mover shares implied by the stay rates. The joint stage then estimates all twelve parameters on all eighteen moments by Nelder-Mead, warm-started at the unrestricted optimum with perturbed restarts; the restarts converge to the same neighborhood and the reported optimum. Standard errors are from a parametric bootstrap with 100 draws: each draw perturbs the eighteen moment targets by independent normal draws at their estimated standard errors and re-estimates all twelve parameters.

B.5 Identification

Each parameter is identified by one or more moments in the model.

- The four mover shares λ_{s,k_0} are pinned nearly one for one by the four treated-arm stay rates: a stayer keeps location and tenure, so the stay rate of each origin tenure and skill cell is one minus the mover share times the movers' away-share.

- The tenure-continuation premia c_s are identified by the conversion rates given the mover shares: the rate at which renter-origin households end the period as in-band owners is governed by how strongly movers persist in their origin tenure.

- The user-cost reduction τ is identified by the arrival contrasts and the owner-quantity DiD: under new-origination incidence the subsidy moves arrivals into treated-arm ownership and the owner stock, while the conversion response stays muted under the lock (the margin-mix gate above).

- The supply-elasticity multiplier ξ is identified by the price-versus-quantity split on the owner side: the flat price DiD against the strong owner-quantity DiD bounds ξ from below, and the bound is one-sided because beyond a threshold additional elasticity produces the same near-zero capitalization.

- The nesting parameter ρ is identified by the rent DiD together with the tenure tilt of

arrivals: smaller ρ produces larger rent declines for a given owner-side expansion as arriving demand substitutes from renting toward owning.

- The three loan-size-specific execution wedges Δ_{size} are identified by the corresponding Stage 2 approval-rate DiDs: each Δ enters the lender's profit function only for applications in its own bin, so cross-bin spillover in the model-implied approval rate is small by construction.

The following table summarizes the mapping.

Parameter	*Primary identifying moments*
$\lambda_{O,H}, \lambda_{O,L}, \lambda_{R,H}, \lambda_{R,L}$	Treated-arm stay rates by tenure $\times$ skill
c_H, c_L	Renter-to-owner conversion rates by arm $\times$ skill
τ	Arrival contrasts, renter-arrival contrast, owner-quantity DiD
ξ	Price DiD vs. owner-quantity DiD (lower bound)
ρ	Rent DiD, tenure tilt of arrivals
$\Delta_{\text{Small}}, \Delta_{\text{Conforming}}, \Delta_{\text{Jumbo}}$	Stage 2 approval DiD in the own loan-size bin

C Online Appendix: Data Construction

This appendix details the construction of key datasets used in our analysis, including an annual panel of commuting flows between U.S. census tracts, disaggregated by worker skill and housing tenure. For the construction of the commuting data, the primary objective is to generate a measure of commuting flows, J_{odtk}, representing the number of workers in skill-tenure category k traveling from origin tract o to destination tract d in year t (in this appendix, the index k runs over the four skill–tenure categories defined below and is distinct from the tenure index of Section 5). This requires a data fusion methodology that integrates three main sources: the Longitudinal Employer-Household Dynamics (LODES) dataset, the American Community Survey (ACS) Public Use Microdata Sample (PUMS), and the National Household Travel Survey (NHTS). The core challenge lies in the fact that LODES provides geographically granular commuting flows but lacks detailed worker characteristics, whereas the ACS PUMS contains rich individual-level data but is only available at the coarser Public Use Microdata Area (PUMA) level.

C.1 Baseline Commuting Matrix and Worker Characteristics

The foundation of our dataset is the tract-to-tract commuting matrix derived from the LODES Origin-Destination (OD) files. For each year, we aggregate the raw census block-level flows to the 2010 census tract level to ensure geographic consistency over time. This process yields a baseline matrix of total commuters, J_{odt}, for every origin-destination-year triplet.

We define worker characteristics using the ACS PUMS data. Worker skill is a binary classification based on educational attainment: individuals with a bachelor's degree or higher are classified as *High-Skill*, and all others are classified as *Low-Skill*. Housing tenure is also binary: workers are classified as *Owners* or *Renters* based on the tenure status of their household. Combining these yields four mutually exclusive skill-tenure categories (k): (1)

High-Skill Owner, (2) High-Skill Renter, (3) Low-Skill Owner, and (4) Low-Skill Renter.

C.2 Imputing Worker Characteristics with a Gradient Boosting Model

To disaggregate the total commuting flows from LODES, we impute the skill-tenure composition of workers residing in each origin tract. This is accomplished by training a predictive model on the ACS PUMS microdata and NHTS data, and then applying it to the tract-level demographic profiles available in LODES. The 2009 NHTS version provides origin and destination census tracts for each traveler, so we rely on the 2009 data.

We employ an XGBoost (Extreme Gradient Boosting) model, a machine learning algorithm well suited for capturing complex, non-linear relationships and high-order interactions between variables. The model's objective is to predict the probability that a worker belongs to one of the four skill-tenure categories, conditional on a set of characteristics available in both the ACS and LODES datasets. These predictor variables include coarse but consistently defined bins for worker age and earnings, as provided in the LODES data files.

The trained XGBoost model outputs a probability distribution over the four skill-tenure categories for any given combination of input features (e.g., for a worker in the lowest age bin and highest earnings bin). We then use the LODES Residence Area Characteristics (RAC) files to determine the population count for each combination of these features within every origin census tract o. Table A1 lists the predictor variables. The estimated proportion of workers in skill-tenure category k for tract o in year t, denoted $\hat{P}_{otk}$, is calculated as a weighted average of the model's predicted probabilities, where the weights are the shares of the tract's resident worker population in each demographic cell.

Finally, we disaggregate the total commuting flows under the assumption that the skill-tenure distribution of commuters from a tract mirrors that of its resident population. We jointly predict skill and tenure instead of training separate models, in order to account for the correlation between the two (for example, high-skilled workers are more likely to own). This

procedure is applied to every commuting pair for each year, yielding the final disaggregated panel.

Table A1: Model Variables and Descriptions

Variable Group	*Count*	*Prefix*	*Description*
Residential Area Characteristics (RAC)			
Worker age shares	3	`share_CA*_rac`	Age distribution of jobs (29 or younger, 30–54, 55+)
Worker earnings shares	3	`share_CE*_rac`	Earnings distribution (low, medium, high)
Industry shares	20	`share_CNS*_rac`	Jobs by NAICS sector (agriculture to public admin.)
Workplace Area Characteristics (WAC)			
Worker age shares	3	`share_CA*_wac`	Age distribution at workplace
Worker earnings shares	3	`share_CE*_wac`	Earnings distribution at workplace
Industry shares	20	`share_CNS*_wac`	Industry distribution at workplace
Origin-Destination Commute Flow Characteristics (OD)			
Worker age shares	3	`share_SA*_od`	Age distribution in commute flow
Worker earnings shares	3	`share_SE*_od`	Earnings distribution in commute flow
Worker industry shares	3	`share_SI*_od`	Industry groupings in commute flow
Home-Work Tract Differences			
Demographics	8	`diff_p*`	Race/ethnicity and age distribution differences
Housing	5	`diff_*`	Tenure, home value, rent, vacancy differences
Employment	4	`diff_l*`	Wages and employment differences by skill level
Socioeconomics	3	`diff_p*`	Poverty, education, and unemployment differences
Total Predictors	*81*		
Other Variables			
Target: Skill	1	`y_skill`	Binary: 1 = high-skilled, 0 = low-skilled
Target: Ownership	1	`y_ownership`	Binary: 1 = homeowner, 0 = renter
Weight	1	`wtperfin`	Person-level survey weight

D Online Appendix: Additional Figures

Purchase Money Loans – LLPA by Credit Score/LTV Ratio										
Credit Score	LTV Range									
	Applicable for all loans with terms greater than 15 years									
	≤ 30.00%	30.01 – 60.00%	60.01 – 70.00%	70.01 – 75.00%	75.01 – 80.00%	80.01 – 85.00%	85.01 – 90.00%	90.01 – 95.00%	>95.00%	SFC
≥ = 780	0.000%	0.000%	0.000%	0.000%	0.375%	0.375%	0.250%	0.250%	0.125%	N/A
760 – 779	0.000%	0.000%	0.000%	0.250%	0.625%	0.625%	0.500%	0.500%	0.250%	N/A
740 – 759	0.000%	0.000%	0.125%	0.375%	0.875%	1.000%	0.750%	0.625%	0.500%	N/A
720 – 739	0.000%	0.000%	0.250%	0.750%	1.250%	1.250%	1.000%	0.875%	0.750%	N/A
700 – 719	0.000%	0.000%	0.375%	0.875%	1.375%	1.500%	1.250%	1.125%	0.875%	N/A
680 – 699	0.000%	0.000%	0.625%	1.125%	1.750%	1.875%	1.500%	1.375%	1.125%	N/A
660 – 679	0.000%	0.000%	0.750%	1.375%	1.875%	2.125%	1.750%	1.625%	1.250%	N/A
640 - 659	0.000%	0.000%	1.125%	1.500%	2.250%	2.500%	2.000%	1.875%	1.500%	N/A
≤ 639	0.000%	0.125%	1.500%	2.125%	2.750%	2.875%	2.625%	2.250%	1.750%	N/A

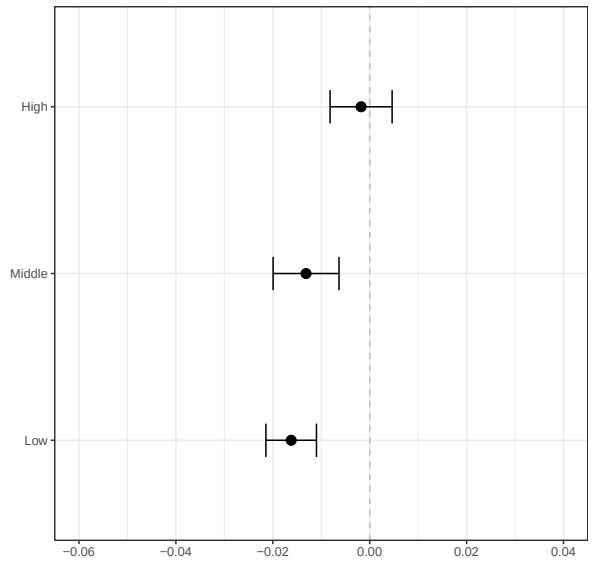
(a) Loan-Level Price Adjustment (LLPA) Matrix

LLPA Waivers All LLPAs will be waived for the following loans (Exception: LLPAs for Minimum Mortgage Insurance Coverage Option will be charged if applicable)	
Product Feature	**SFC**
HomeReady® loans	900
Loans to first-time homebuyers with qualifying income ≤100% area median income (AMI) or 120% AMI in high-cost areas	NA
Loans meeting Duty to Serve requirements (Purchase and limited cash-out refinances, principal residences with total qualifying income ≤100% AMI) • Manufactured housing including MH Advantage® • Rural housing - loans in high needs rural regions • Loans to Native Americans on tribal lands • Loans originated by "small financial institutions" • Affordable housing preservation loans – loans financing ENERGY STAR® certified improvements, loans with shared equity	874

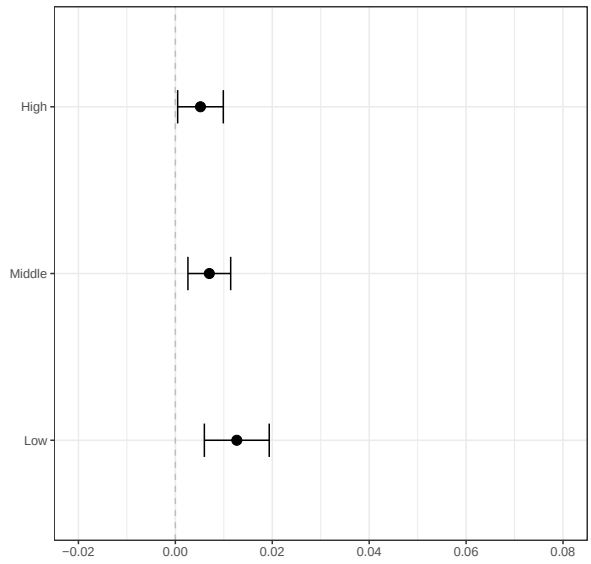
(b) LLPA waiver for loans meeting Duty to Serve requirements

Figure A1: **The Post-2022 DTS Fee Waiver (Fannie Mae LLPA Matrix)**

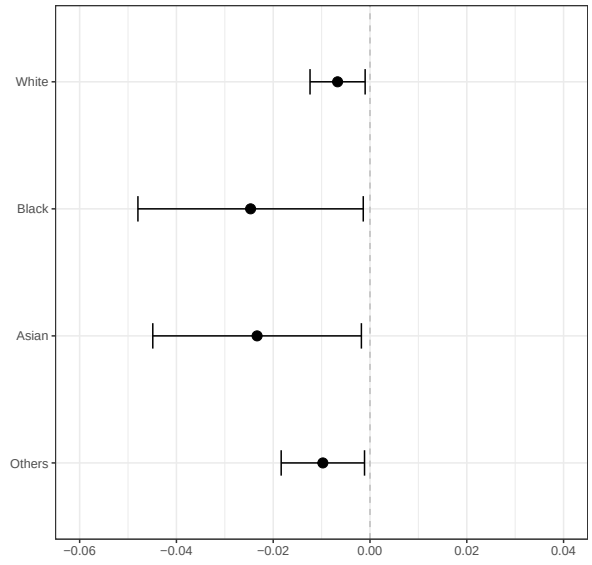
Notes: This figure reproduces excerpts from Fannie Mae's Loan-Level Price Adjustment (LLPA) Matrix, version effective January 28, 2026 (Fannie Mae, 2026). Panel (a) shows the baseline LLPA grid: price adjustments, in percentage points of the loan amount, charged on loans sold to Fannie Mae as a function of credit score, loan-to-value ratio, and other loan features. Panel (b) shows the waiver table, under which LLPAs are waived for "loans meeting Duty to Serve requirements".



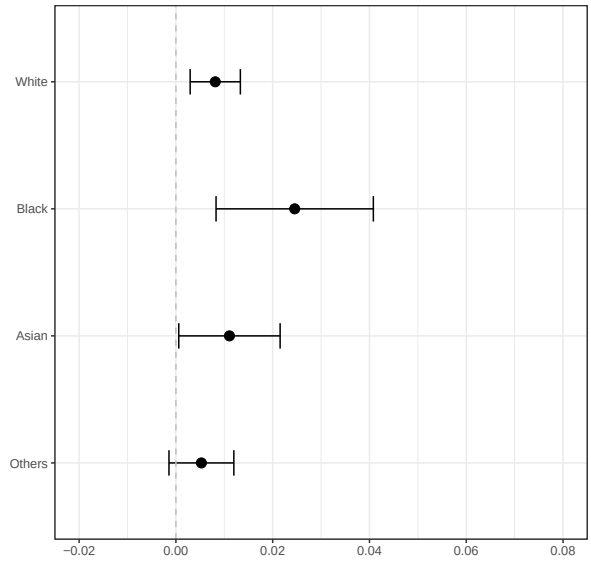
(a) Loan Amount by Income



(b) Small Loans by Income



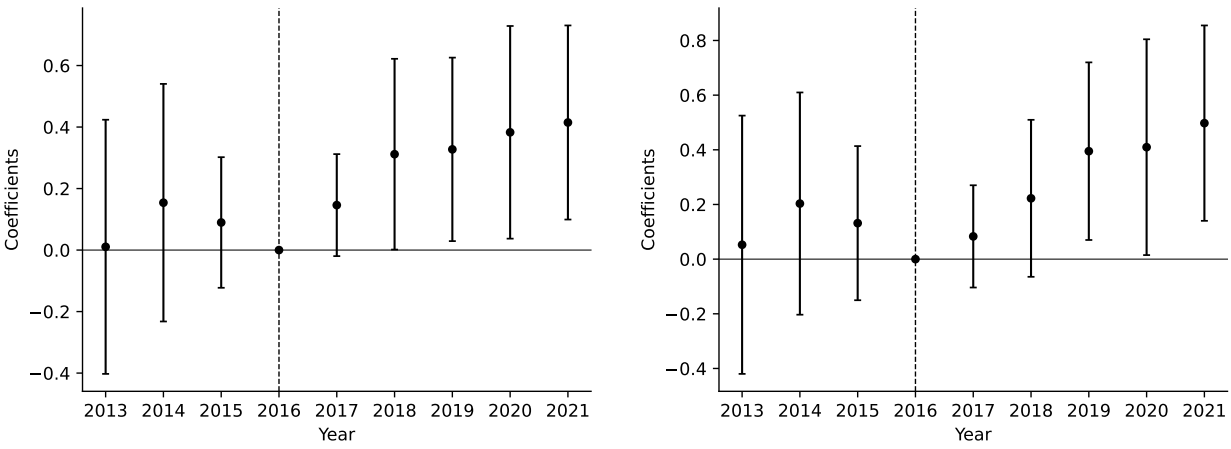
(c) Loan Amount by Race



(d) Small Loans by Race

Figure A2: Heterogeneous Effects on Loan Outcomes by Borrower Characteristics

Notes: These figures present heterogeneous treatment effects of the Duty to Serve program, examining how its impact on mortgage composition varies by borrower income and race. Income and race categories are based on census-tract-level classifications. Each point is the Rural × Post difference-in-differences coefficient for the indicated group, with a 95% confidence interval based on standard errors clustered at the state level; the vertical dashed line marks zero. The sample runs from 2013 to 2021.

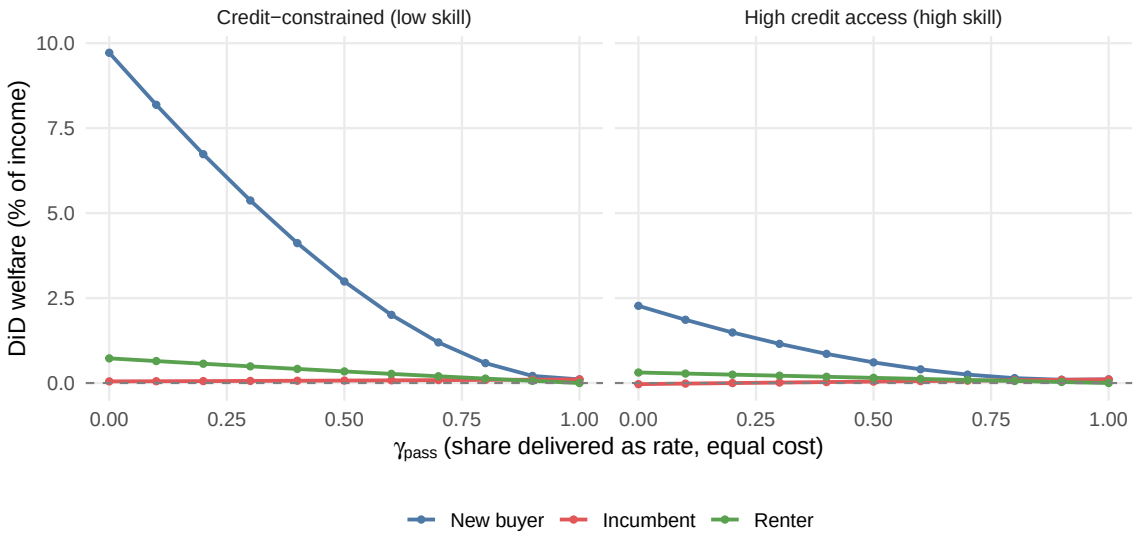


(a) New Residential Permits (Total)

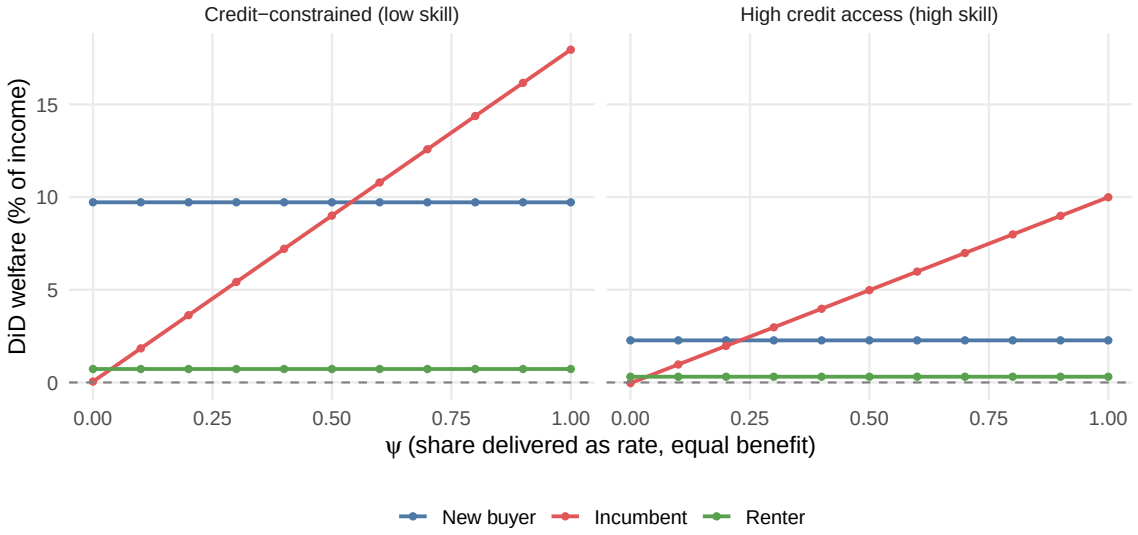
(b) New Single-Family Permits

Figure A3: New Construction at the DTS Boundary: Building-Permit Event Study

Notes: Each point is the Rural × year difference-in-differences coefficient (reference year 2016) from a Poisson pseudo-maximum-likelihood regression of the tract-level count of new residential building permits on tract and boundary × year fixed effects, with a 95% confidence interval based on standard errors clustered at the state level. The sample is a balanced tract-year panel of building permits from Builty, 2013–2021, restricted to tracts within 5 mi of the DTS boundary in boundary pairs with permit coverage on both sides and to reporting localities observed in every sample year; panel (a) covers 1,776 tracts (216 treated) in 31 states and panel (b) 1,681 tracts (210 treated). Outcome counts are winsorized at 1% / 99%. Panel (a) counts all new residential permits; panel (b) counts new single-family permits. The permit year is the issuance date (submission date if missing). The corresponding pooled Rural × Post coefficients are 0.251 (SE 0.108) in panel (a) and 0.229 (SE 0.120) in panel (b). Estimates are local to Builty-covered jurisdictions, and coverage is thinner on the rural side of the boundary.



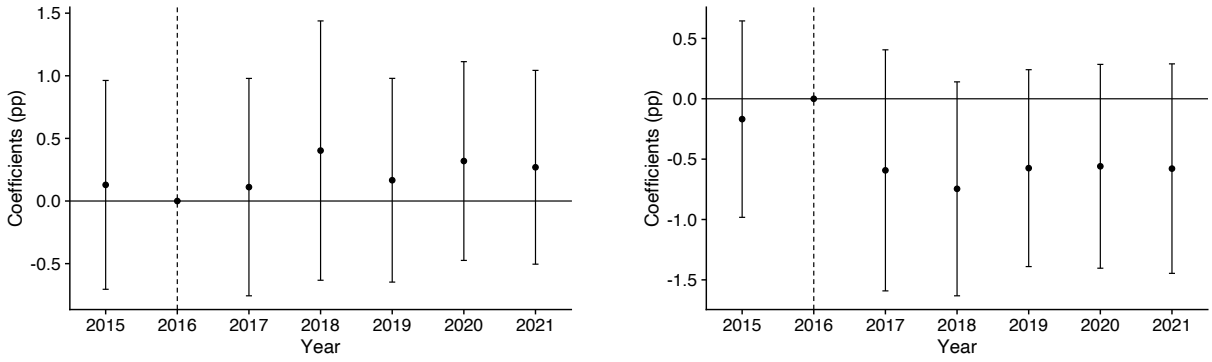
(a) Rate delivery, equal cost (γ_{pass} from access to rate)



(b) Rate delivery, equal benefit (ψ from access to rate)

Figure A4: **Counterfactual welfare incidence by household group**

Notes: Each panel plots the model-implied difference-in-differences equivalent variation (EV, percent of annual income) for the six household groups, faceted by skill (credit-constrained low skill on the left, high credit access on the right), as a continuous function of the counterfactual parameter, holding the other structural parameters at their joint-estimation point values. Panel (a) varies the rate pass-through fraction γ_{pass} holding the GSE concession fixed; Panel (b) varies the rate share ψ holding the household benefit fixed (the equilibrium is invariant, so new-buyer and renter welfare is flat).



(a) Foreclosure

(b) Bankruptcy

Figure A5: **Dynamic Effects on Foreclosure and Bankruptcy**

Notes: The figure plots dynamic difference-in-differences estimates for foreclosure and bankruptcy at the loan level. The outcome in each panel is an indicator, in percentage points, for the loan ever entering foreclosure or bankruptcy through the end of the servicing panel; both outcomes are rare in the pre-period. The sample, the definition of exposure, the fixed effects, and the controls follow Figure 6. The coefficient of Exposure $\times$ year identifies the differential change for a fully exposed ZIP code relative to the 2016 origination cohort, which is normalized to zero. Foreclosure was suspended for federally backed mortgages by the CARES Act moratorium over 2020–2021, and bankruptcy filings were curtailed by court closures over the same period. The error bars represent 95% confidence intervals. Standard errors are clustered at the state level.

E Online Appendix: Additional Tables

Table A2: **Descriptive Statistics for ACS Data**

	Mean	Median	SD	Min	Max
Occupied Units (share)	0.89	0.91	0.07	0.67	0.99
Owner-Occupied (share)	0.65	0.68	0.19	0.17	0.93
Log Home Value	11.98	11.94	0.59	10.87	13.26
Log Rent	6.70	6.68	0.32	6.12	7.42
Rent-to-Home-Value Ratio	0.07	0.06	0.03	0.02	0.14
Log Wage (High Skill)	10.65	10.65	0.25	10.06	11.18
Log Wage (Low Skill)	10.13	10.13	0.22	9.66	10.60
Log Employment (High Skill)	6.07	6.12	0.81	4.13	7.51
Log Employment (Low Skill)	6.81	6.87	0.51	5.45	7.70
College Educated (share)	0.27	0.24	0.13	0.08	0.61
Unemployment Rate (share)	0.02	0.02	0.01	0.00	0.05
Poverty Rate (share)	0.07	0.06	0.04	0.02	0.18
Non-Hispanic Black (share)	0.13	0.06	0.15	0.00	0.70
Asian (share)	0.04	0.02	0.05	0.00	0.24
Non-Hispanic White (share)	0.74	0.81	0.22	0.14	0.99

Notes: The table presents summary statistics for the census-tract panel constructed from ACS five-year files (NHGIS). The sample consists of all U.S. census tracts over 2012–2021, with each five-year window assigned to its midpoint year and tracts harmonized to 2010 boundaries. Shares are relative to the tract population or, for tenure and occupancy, to occupied housing units; home values, rents, wages, and employment enter in logs.

Table A3: **Skill and Credit Access in the Pre-Treatment Boundary Sample**

	High-skill tracts (1)	Low-skill tracts (2)	Difference (3)	t-stat (4)
Mean applicant income ($)	124 749	68 745	56 003***	49.04
Mean loan amount ($)	288 897	163 117	125 779***	42.89
Mortgage denial rate (%)	14.95	21.00	−6.05***	-57.24
Share of small-balance loans, ≤ $100k (%)	8.88	26.75	−17.87***	-46.75
Share of FHA loans (%)	12.12	26.23	−14.11***	-55.08
Number of tracts	6778	6778		

Notes: The sample is HMDA loan applications from 2013–2016 in census tracts within 5 miles of the DTS rural eligibility boundary. Tracts are split at the median ACS share of population aged 25 and over with a bachelor's degree or higher. Columns (1) and (2) report group means; column (3) is the difference; column (4) is the t-statistic from a regression of the outcome on the high-skill indicator with standard errors clustered at the census-tract level. The denial-rate row uses the full application sample; all other rows are restricted to originated loans. Real monetary values are in 2010 dollars. ***p<0.01, **p<0.05, *p<0.1.

Table A4: **Household Education and Income in the Boundary Sample**

	Household income band (row share, %)			
Education of household head	<$50k	$50–100k	>$100k	Households
Some high school	71.5	21.3	7.2	9,843
High school	47.8	30.3	21.9	2,742,852
Some college	30.0	34.8	35.2	1,713,616
Vocational/technical	28.5	34.6	36.9	99,718
College	16.2	25.9	57.9	3,644,222
Graduate school	17.3	28.4	54.3	1,802,612
High-skill household (college+)	16.5	26.7	56.7	5,446,834
Low-skill household (below college)	40.8	32.1	27.1	4,566,029

Notes: This table uses Infutor individual records to cross-tabulate the education and estimated household income band of the household profiled at each address in the five-mile boundary sample. Income bands are the modeled estimates by Infutor. The household-level gradient corroborates the tract-level skill-credit bridge of Table A3, suggesting that education predicts the income that governs loan size and approval throughout the boundary sample.

Table A5: **Lender Margin: Stage 2 (Approval) and Stage 4 (GSE-Sale) Heterogeneity**

Sample	Coef	SE	N
Panel A: Stage 2 P(approved | applied) by lender type, 2013–2017			
Bank	0.0042	(0.0032)	1,141,159
Credit Union	0.0060*	(0.0030)	505,381
Independent mortgage company	0.0070***	(0.0026)	2,668,649
Large bank (CFPB-supervised)	0.0037	(0.0037)	2,355,885
Pooled (all types)	0.0053***	(0.0015)	6,671,074
Panel B: Stage 4 P(GSE-sold | originated)			
By lender type, 2013–2017			
Bank	−0.0068*	(0.0034)	833,251
Credit Union	0.0014	(0.0055)	330,058
Independent mortgage company	−0.0027	(0.0027)	1,598,209
Large bank (CFPB-supervised)	0.0011	(0.0036)	1,477,531
By GSE, 2013–2021			
Fannie sold	0.0002	(0.0016)	9,631,555
Freddie sold	−0.0038***	(0.0014)	9,631,555
By pre-treatment HHI tercile			
Low HHI	−0.0029*	(0.0016)	4,980,130
Mid HHI	−0.0049***	(0.0016)	3,061,350
High HHI	−0.0022	(0.0024)	1,589,491
Panel C: Stage 4 by loan purpose, 2013–2021			
Purchase	−0.0063***	(0.0017)	4,239,484
Refinance	−0.0001	(0.0013)	5,392,071
Panel D: Stage 2 by lender type × loan-size bin, 2013–2017			
Bank, DTS small (≤$200k)	0.0079*	(0.0045)	585,048
Bank, Conforming	0.0019	(0.0052)	431,735
Bank, Jumbo	0.0026	(0.0115)	124,376
Credit Union, DTS small (≤$200k)	−0.0001	(0.0048)	323,287
Credit Union, Conforming	0.0240***	(0.0074)	149,409
Credit Union, Jumbo	−0.0125	(0.0195)	32,685
Independent mortgage company, DTS small (≤$200k)	0.0024	(0.0041)	1,141,807
Independent mortgage company, Conforming	0.0150***	(0.0038)	1,245,331
Independent mortgage company, Jumbo	0.0059	(0.0115)	281,511
Large bank (CFPB-supervised), DTS small (≤$200k)	0.0037	(0.0042)	1,141,875
Large bank (CFPB-supervised), Conforming	0.0037	(0.0054)	743,797
Large bank (CFPB-supervised), Jumbo	0.0153*	(0.0083)	470,213

Notes: The table presents the approval and GSE-sale responses by lender type, purchaser, market concentration, and loan-size bin. The sample consists of HMDA conventional first-lien single-family loans within five miles of the DTS rural eligibility boundary; the outcome is an indicator for approval on the application sample (Panels A and D) and an indicator for sale to a GSE on originated loans (Panels B and C). Lender type is defined from the HMDA agency code: codes 1, 2, and 3 are depositories supervised by the OCC, the Federal Reserve, and the FDIC; code 5 is credit unions (NCUA); code 7 is nondepository mortgage companies (HUD); and code 9 is depositories with assets above $10 billion and their affiliates (CFPB). Agency codes are available for 2013–2017 only, because HMDA moved to LEI identifiers in 2018, so the lender-type panels use that window. Tract HHI is computed from 2013–2016 origination shares by HMDA respondent. The DTS small-balance threshold is $200,000 through 2020 and $250,000 in 2021. All specifications include census tract and boundary × year fixed effects, and the GSE-sale specifications add lender × year fixed effects. Standard errors are clustered at the state level. ***p<0.01, **p<0.05, *p<0.1.

Table A6: Within-Denied-Application Composition Shift, by Denial Reason

	Debt-to-income	Collateral	Credit history	Collateral or DTI
Panel A: Full sample (within-denied composition)				
Rural × Post	−0.0057**	−0.0017	−0.0014	−0.0074**
	(0.0023)	(0.0028)	(0.0025)	(0.0032)
Panel B: By loan-size bin				
DTS-small (≤$200K)	−0.0087***	0.0039	−0.0023	−0.0048
	(0.0027)	(0.0039)	(0.0035)	(0.0036)
Conforming	−0.0011	−0.0113**	0.0013	−0.0124***
	(0.0045)	(0.0044)	(0.0034)	(0.0040)
Jumbo	−0.0175	−0.0002	−0.0011	−0.0178*
	(0.0114)	(0.0060)	(0.0042)	(0.0102)
Panel C: By borrower-income tercile				
Low income	−0.0041	0.0011	−0.0028	−0.0030
	(0.0035)	(0.0040)	(0.0031)	(0.0033)
Mid income	−0.0034	−0.0050	0.0004	−0.0085
	(0.0046)	(0.0049)	(0.0042)	(0.0073)
High income	−0.0043	−0.0063	−0.0040	−0.0107**
	(0.0044)	(0.0040)	(0.0049)	(0.0051)
Tract FE	✓	✓	✓	✓
Boundary × Year FE	✓	✓	✓	✓
R^2 (Panel A)	0.0334	0.0494	0.0495	0.0278
Observations (Panel A)	1,560,590	1,560,590	1,560,590	1,560,590

Notes: The table presents the treatment effect on the composition of denial reasons among denied applications. The sample consists of HMDA conventional first-lien single-family purchase and refinance applications within five miles of the DTS rural eligibility boundary over 2013–2021, restricted to denied applications with a recorded primary denial reason. In each column the outcome is an indicator equal to one if the primary reason matches the column heading, and each cell reports the coefficient of Rural × Post from a separate regression. Loan-size bins follow Table 4: DTS small-balance is at or below $200,000 through 2020 and $250,000 in 2021, conforming is above that and at or below the FHFA one-unit conforming limit of the year, and jumbo is above the conforming limit. Income terciles are defined year by year within the boundary sample. All specifications include census tract and boundary × year fixed effects. Standard errors are clustered at the state level. ***p<0.01, **p<0.05, *p<0.1.

Table A7: Denial Response by Recorded Reason: Levels (Share of Applications)

	Denied (any)	Debt-to-income	Collateral	Credit history	Other recorded	No reason recorded
Panel A: Full sample, 2013–2021						
Rural × Post	−0.0068***	−0.0008	−0.0013***	−0.0007*	0.0012**	−0.0052***
	(0.0011)	(0.0005)	(0.0004)	(0.0004)	(0.0005)	(0.0012)
Pre-period mean, treated	0.1773	0.0281	0.0284	0.0261	0.0351	0.0596
Observations				15,570,522		
Panel B: Missing-reason denials allocated by lender-year recorded shares						
Rural × Post	−0.0068***	−0.0020***	−0.0026***	−0.0018***	−0.0004	–
	(0.0011)	(0.0003)	(0.0004)	(0.0004)	(0.0005)	
Panel C: By loan-size bin						
DTS-small (≤$200K)	−0.0044***	−0.0006*	0.0004	−0.0005	0.0017***	−0.0055***
	(0.0010)	(0.0004)	(0.0006)	(0.0006)	(0.0004)	(0.0008)
Conforming	−0.0108***	−0.0012*	−0.0031***	−0.0008**	0.0003	−0.0060***
	(0.0018)	(0.0007)	(0.0004)	(0.0004)	(0.0006)	(0.0012)
Jumbo	−0.0117**	−0.0026	−0.0044***	−0.0012*	0.0014	−0.0048*
	(0.0056)	(0.0026)	(0.0009)	(0.0007)	(0.0009)	(0.0024)
Panel D: Pre-COVID window, 2013–2019						
Rural × Post	−0.0037***	−0.0008	−0.0002	−0.0000	0.0017***	−0.0043***
	(0.0011)	(0.0005)	(0.0005)	(0.0004)	(0.0005)	(0.0011)
Panel E: By loan purpose						
Purchase	−0.0034***	−0.0000	−0.0004	−0.0008**	0.0009	−0.0030***
	(0.0010)	(0.0004)	(0.0003)	(0.0003)	(0.0006)	(0.0008)
Refinance	−0.0072***	−0.0014**	−0.0015**	−0.0009	0.0012**	−0.0046***
	(0.0017)	(0.0006)	(0.0006)	(0.0006)	(0.0006)	(0.0014)
Tract FE				✓		
Boundary × Year FE				✓		

Notes: The table presents the treatment effect on denial rates by recorded reason, measured as a share of all applications, so the five reason columns sum to the first column. The sample consists of HMDA conventional first-lien single-family purchase and refinance applications within five miles of the DTS rural eligibility boundary over 2013–2021. In each column the outcome is an indicator equal to one if the application was denied for the column's reason, and each cell reports the coefficient of Rural × Post from a separate regression. Other recorded reasons are employment history, insufficient cash, unverifiable information, incomplete application, mortgage insurance denied, and other. Reporting a denial reason was optional in HMDA before 2018 and mandatory after, so the last column captures any differential change in reporting; Panel B allocates denials with no recorded reason across reasons in proportion to the recorded-reason shares of the same lender and year, using boundary-year shares when the lender records fewer than five denials. Loan-size bins follow Table 4. All specifications include census tract and boundary × year fixed effects. Standard errors are clustered at the state level. ***p<0.01, **p<0.05, *p<0.1.

Table A8: **Approval Response by Distance from the Debt-to-Income Limit**

	Rural × Post (pp)	S.E.	Observations	Pre-period denial rate, treated (%)	Denials flipped (% of pre-period)
Panel A: Stage 2 approval DiD by loan-to-income ratio at application					
<1.5	0.76***	(0.18)	2,999,518	18.2	4.2
1.5-2	0.94***	(0.25)	2,154,725	17.9	5.3
2-2.5	1.46***	(0.30)	2,241,079	17.9	8.2
2.5-3	1.59***	(0.32)	1,932,326	18.8	8.5
3-3.5	1.68***	(0.36)	1,547,846	20.3	8.3
3.5-4	2.32***	(0.36)	1,174,585	22.4	10.4
4-4.5	1.73***	(0.37)	847,080	25.3	6.8
4.5-5	1.49***	(0.45)	608,785	29.0	5.1
>5	0.44	(0.59)	1,117,519	47.9	0.9
Panel B: Pooled Stage 2 approval DiD, with and without bin × boundary × year fixed effects					
All applications, prediction sample	1.35***	(0.18)	14,623,463	20.5	6.6
Within predicted-denial deciles	1.19***	(0.20)	14,623,463		
Within predetermined-feature deciles	1.24***	(0.20)	14,623,463		
Within loan-to-income bins	1.21***	(0.17)	14,623,463		
Within purchase vs. refinance	1.13***	(0.21)	14,623,463		
Within purpose × income decile × occupancy	1.12***	(0.21)	14,623,463		

Notes: The table presents the treatment effect on loan approval by distance from the debt-to-income limit. The sample consists of HMDA conventional first-lien single-family applications with non-missing applicant income within five miles of the DTS rural eligibility boundary over 2013–2021. Panel A reports the Stage 2 approval coefficient of Table 4 estimated separately within bins of the loan-to-income ratio at application, defined as the loan amount divided by applicant income; the last two columns report the 2013–2016 denial rate among decided applications in treated tracts and the coefficient as a share of that rate. Panel B reports the pooled coefficient on the same sample, and then the coefficient when the boundary × year fixed effects are interacted with a bin: deciles of a predicted denial probability, estimated by a logit of denial on within-year deciles of loan amount and income, the loan-to-income ratio and its square, a purchase indicator, and owner occupancy, using decided applications in 2013–2016; deciles of a predetermined-feature version of that probability that omits loan amount and the loan-to-income ratio; loan-to-income bins; loan purpose; and purpose × income decile × occupancy cells. All specifications include census tract and boundary × year fixed effects. Standard errors are clustered at the state level. ***p<0.01, **p<0.05, *p<0.1.

Table A9: DTI and Early Default: Predictive Content Conditional on FICO and Collateral

	(1)	(2)	(3)	(4)
Panel A: Logit, D90+ within 36 months (full sample)				
FICO	−0.0192***	−0.0197***	−0.0193***	
	(0.0004)	(0.0004)	(0.0004)	
CLTV		0.0240***	0.0236***	
		(0.0015)	(0.0015)	
DTI			0.0356***	0.0444***
			(0.0026)	(0.0024)
Pseudo-R^2	0.0769	0.0872	0.0942	0.0132
Observations	248,688	248,688	248,688	248,688
Panel B: Non-metro subsample				
FICO	−0.0181***	−0.0184***	−0.0182***	
	(0.0013)	(0.0014)	(0.0014)	
CLTV		0.0228***	0.0222***	
		(0.0049)	(0.0049)	
DTI			0.0263***	0.0354***
			(0.0078)	(0.0073)
Pseudo-R^2	0.0709	0.0800	0.0844	0.0092
Observations	21,205	21,205	21,205	21,205

Panel C: D90+ within 36 months by origination DTI bin (full sample)

			D90+ (%)	
DTI bin	Loans	Share (%)	Raw	Adjusted
≤25	49,469	19.9	0.48	0.75
26–30	35,934	14.4	0.69	0.77
31–35	41,403	16.6	1.01	0.96
36–40	46,543	18.7	1.18	1.02
41–43	31,031	12.5	1.64	1.43
44–45	23,456	9.4	2.12	1.79
46–50	20,852	8.4	0.93	1.24

Notes: The table presents the predictive content of the debt-to-income ratio for early default, conditional on credit score and collateral. The sample consists of one-unit primary-residence loans originated in 2016Q1 in the Freddie Mac Single-Family Loan-Level Dataset, the last quarter before the December 2016 final rule, with non-missing FICO score, combined loan-to-value ratio (CLTV), and debt-to-income ratio (DTI). The outcome is an indicator equal to one if the loan reaches 90-day delinquency or worse within 36 months of origination, a window that closes in 2019 and so predates pandemic-era forbearance. Panels A and B report logit coefficients; Panel B restricts the sample to properties outside metropolitan statistical areas, the segment DTS targets. Panel C sorts loans into origination-DTI bins and reports each bin's share of loans, its raw 90-day delinquency rate, and a rate adjusted for FICO × CLTV composition, computed by demeaning the outcome within FICO × CLTV cells and recentering on the sample mean. Loans with DTI above 45 percent require compensating factors under GSE underwriting. The dataset identifies location only at the three-digit ZIP level, so the table describes market-wide underwriting patterns rather than boundary estimates. Standard errors are reported in parentheses. ***p<0.01, **p<0.05, *p<0.1.

F Online Appendix: Additional Robustness Checks

F.1 Pre-COVID Subsample

Our post-treatment window (2017–2021) overlaps with the COVID-19 pandemic. Although our spatial DiD design compares geographically adjacent tracts on either side of the same eligibility boundary, so that pandemic-era shocks enter as a common local disturbance absorbed by the boundary-by-year fixed effects, we verify that our results do not depend on the inclusion of the pandemic period. Table A10 replicates the main static DiD specifications restricting the sample to 2013–2019.

Table A10: **Pre-COVID Robustness: Main Results Excluding 2020–2021**

	Full Sample		Pre-COVID (2013–2019)		
Outcome	Estimate	SE	Estimate	SE	N
Panel A: Credit Expansion (HMDA)					
Log Loan Amount	-0.0086^{**}	(0.0036)	-0.0129^{***}	(0.0042)	2,810,916
Small Loan Share	0.0078^{**}	(0.0033)	0.0098^{***}	(0.0029)	2,810,916
Panel B: Housing Market (ACS)					
Housing Price	−0.0020	(0.0054)	−0.0002	(0.0053)	94,953
Rent	-0.0243^{***}	(0.0061)	-0.0195^{***}	(0.0057)	95,580
Owner-Occupied Share	0.0092^{***}	(0.0019)	0.0075^{***}	(0.0019)	95,778
Renter-Occupied Share	-0.0074^{***}	(0.0020)	-0.0059^{***}	(0.0021)	95,778
Panel C: Employment by Skill (ACS)					
Log Low-Skilled Emp.	0.0400^{***}	(0.0098)	0.0314^{***}	(0.0083)	95,778
Log High-Skilled Emp.	0.0393^{***}	(0.0137)	0.0298^{***}	(0.0112)	95,778

Notes: This table compares the baseline DiD estimates (full sample, 2013–2021) with estimates restricting the post-treatment window to 2017–2019. The reported N is for the pre-COVID (2013–2019) specification; the full-sample specifications use the 2013–2021 panel with correspondingly larger samples (e.g., $N = 123{,}257$ for the occupancy-share outcomes). All ACS specifications include tract and boundary-by-year fixed effects, as well as demographic controls (racial composition, age distribution, poverty rate, unemployment rate, and college share). The HMDA specifications use boundary-by-year-by-GSE fixed effects in place of boundary-by-year. Standard errors clustered at the state level are in parentheses. ***p<0.01, **p<0.05, *p<0.1.

All results are quantitatively stable. The credit expansion is more pronounced in the pre-COVID window: average loan sizes decline by 1.3% pre-COVID versus 0.9% in the full sample, and the small-loan share increases by 1.0 percentage point pre-COVID versus 0.8 pp

in the full sample, both significant at the 1% level in the pre-COVID window. Housing prices remain indistinguishable from zero in both samples, so the absence of price capitalization is not a pandemic artifact. The increase in the owner-occupied share (0.75 pp pre-COVID versus 0.92 pp full sample) and the decline in rents (2.0% pre-COVID versus 2.4% full sample) are smaller by approximately 18–20%, consistent with the monotone build-up in the event studies and two fewer post-treatment years. Employment by skill is likewise preserved: low-skilled employment rises by 3.1% (versus 4.0%) and high-skilled employment by 3.0% (versus 3.9%), each significant at the 1% level.

F.2 Placebo in Time

A further threat to identification is the possibility that treated and control tracts were already diverging prior to the DTS announcement. While the event-study figures in the main text show pre-treatment coefficients close to zero for most outcomes, we provide a more formal test by estimating the main DiD specification with a placebo treatment date. We restrict the sample to the pre-DTS period (2013–2016) and assign a placebo treatment at the end of 2014, so that "pre" corresponds to 2013–2014 and "post" to 2015–2016. Table A11 reports the results.

Seven of the eight placebo estimates are statistically insignificant, with magnitudes close to zero. The sole significant coefficient is on rent ($t = -2.08$, $p = 0.04$), but its magnitude (-0.008) is one-third of the actual treatment effect (-0.024). The key outcomes that drive the paper's main conclusions all produce precise zeros in the placebo test, confirming that the treatment effects documented in the main text are attributable to the DTS program rather than to pre-existing differential trends.

F.3 Heterogeneity by Baseline Supply Elasticity

Our model predicts that the transmission of credit expansion into quantities versus prices depends on the elasticity of local housing supply. To test this prediction, we split the sample

Table A11: **Placebo-in-Time Test: Placebo Treatment at End-2014 (Sample: 2013–2016)**

Outcome	Estimate	SE	t	p	N
Panel A: Credit Expansion (HMDA)					
Log Loan Amount	0.0020	(0.0027)	0.73	0.47	1,718,078
Small Loan Share	0.0001	(0.0025)	0.05	0.96	1,718,078
Panel B: Housing Market (ACS)					
Housing Price	−0.0014	(0.0039)	−0.36	0.72	54,175
Rent	−0.0080	(0.0039)	−2.08	0.04	54,536
Owner-Occupied Share	−0.0008	(0.0013)	−0.59	0.56	54,651
Renter-Occupied Share	−0.0010	(0.0016)	−0.60	0.55	54,651
Panel C: Employment by Skill (ACS)					
Log Low-Skilled Emp.	0.0012	(0.0032)	0.38	0.70	54,651
Log High-Skilled Emp.	−0.0027	(0.0078)	−0.35	0.73	54,651

Notes: This table reports DiD estimates using only pre-DTS data (2013–2016) with a placebo treatment date at the end of 2014. The "post" indicator equals one for 2015–2016. All ACS specifications include tract and boundary-by-year fixed effects, as well as demographic controls. The HMDA specifications use boundary-by-year-by-GSE fixed effects in place of boundary-by-year. Standard errors clustered at the state level are in parentheses. ***p<0.01, **p<0.05, *p<0.1.

at the median tract-level housing supply elasticity from Saiz (2010), mapped to census tracts via MSA assignment. Table A12 reports the main DiD estimates separately for tracts in low-elasticity and high-elasticity areas.

The results are consistent with the model's mechanism. In high-elasticity areas, the treatment effect operates primarily through tenure shifts: the owner-occupied share increases by 1.4 percentage points while housing prices are indistinguishable from zero. In low-elasticity areas, the owner-occupied share rises by a smaller 0.4 percentage points, and housing prices do not rise, although the sample's median elasticity of 1.83 leaves the low-elasticity half elastic by metropolitan standards, so the split is a weak test of capitalization. Both subsamples show significant rent declines, though the magnitude is larger in low-elasticity areas (−2.9% versus −2.0%), consistent with a compositional shift toward more affordable rentals where supply is constrained. Low-skilled employment rises in both subsamples, by more in low-elasticity areas (5.9% versus 1.9%). High-skilled employment rises sharply in high-elasticity

Table A12: **Heterogeneity by Baseline Housing Supply Elasticity**

	Low Elasticity		High Elasticity	
Outcome	Estimate	SE	Estimate	SE
Panel A: Credit Expansion (HMDA)				
Log Loan Amount	−0.0161***	(0.0061)	−0.0071	(0.0054)
Small Loan Share	0.0064*	(0.0038)	0.0150***	(0.0053)
Panel B: Housing Market (ACS)				
Housing Price	−0.0096	(0.0077)	0.0055	(0.0061)
Rent	−0.0289***	(0.0097)	−0.0202***	(0.0072)
Owner-Occupied Share	0.0044**	(0.0021)	0.0136***	(0.0035)
Renter-Occupied Share	−0.0043	(0.0027)	−0.0103***	(0.0035)
Panel C: Neighborhood Sorting (ACS)				
Log Low-Skilled Emp.	0.0586***	(0.0156)	0.0189**	(0.0075)
Log High-Skilled Emp.	0.0110	(0.0198)	0.0631***	(0.0156)

Notes: This table reports DiD estimates separately for tracts below ("Low") and above ("High") the median Saiz (2010) housing supply elasticity. The median elasticity in our sample is 1.83. All ACS specifications include tract and boundary-by-year fixed effects with demographic controls. The HMDA specifications use boundary-by-year-by-GSE fixed effects in place of boundary-by-year. Standard errors clustered at the state level are in parentheses. ***p<0.01, **p<0.05, *p<0.1.

areas (6.3%, significant at the 1% level) but is statistically indistinguishable from zero in low-elasticity areas, indicating that elastic-supply markets absorb a broader cross-section of new entrants. The credit expansion effects are present in both subsamples, with high-elasticity areas showing a larger shift toward small loans. These patterns are consistent with the supply regime shaping how the credit expansion divides between quantities and prices.

F.4 Cross-Boundary Migration and SUTVA

The boundary design assumes that control tracts are unaffected by the program. The most direct threat to this assumption is policy-induced migration: if DTS eligibility draws households out of adjacent ineligible tracts, control populations, rents, and tenure counts fall for policy-induced reasons, and the boundary estimates overstate the true effects. We test this channel directly using individual address histories from Infutor Data Solutions, which record each mover's origin and destination address. To this end, we assign addresses to 2010 census

tracts and classify every residential move into a boundary-band tract by the location of its origin. Table A13 reports the origin composition of inflows into treated and control tracts before and after the program.

Table A13: **Origin Composition of Residential Inflows at the Eligibility Boundary**

Origin of in-mover	2014–2016	2018–2019	2018–2020
Panel A: Inflows into treated (rural) tracts			
Outside five-mile band	57.8	63.2	64.1
Opposite side of boundary (band)	17.1	16.2	16.2
Same tract	11.0	7.9	7.2
Same side of boundary (band)	4.5	3.5	3.4
Origin unobserved	9.6	9.2	9.1
Panel B: Inflows into control (non-rural) tracts			
Outside five-mile band	47.0	56.9	58.2
Opposite side of boundary (band)	1.8	1.9	1.9
Same tract	9.6	7.3	6.6
Same side of boundary (band)	32.1	25.0	24.4
Origin unobserved	9.5	9.0	8.9
Panel C: Cross-arm flows, persons per year			
Control to treated (cross-arm)	61,739	88,935	93,449
Treated to control (cross-arm)	50,681	73,794	76,204
Net, control to treated	11,058	15,142	17,245
From outside the band, into treated	208,691	346,993	369,418
From outside the band, into control	1,301,942	2,235,226	2,340,312

Notes: This table reports column shares (percent) of residential moves into boundary-band census tracts by origin location, constructed from Infutor address histories matched to 2010 census tracts. A move is dated by the month its destination address becomes effective; the pre-period is 2014–2016 and the post-period excludes the 2017 transition year. Opposite side of boundary denotes origins in the other treatment arm within the five-mile band, the channel through which policy-induced reallocation would contaminate control outcomes. Moves include renters and owners. The larger opposite-side share in Panel A than Panel B reflects the relative population mass of the two arms. Panel C reports the cross-arm flows by direction and, as a benchmark for the growth of recorded moves over the period, inflows from outside the band, all in average persons per year over the window; the reallocation regression in the text identifies the differential between the two cross-arm directions.

The contamination channel is empirically negligible. The share of treated-tract in-movers originating on the control side of the band is stable at 16–17% before and after the program, and the placebo composition for control-tract inflows is likewise flat. Cross-arm flows rose in both directions after the program, by 51 percent into treated tracts and 50 percent into

control tracts, alongside a 77 to 80 percent rise in all recorded moves into the band from outside it (Panel C), so the common component reflects churn and record coverage rather than the policy. The policy-induced reallocation is the differential between the two directions. A difference-in-differences regression of log cross-arm inflows on destination arm and post, with tract and boundary-by-year fixed effects and state-clustered standard errors, yields 0.021 for the 2018–2020 window and 0.011 for 2018–2019, statistically indistinguishable from zero. The coefficient is the difference between the proportional responses of the two directions; it measures the net pull on the control arm under the assumption that the reverse flow carries no policy response beyond the common growth, which the equal growth of the two directions supports. On that reading, the point estimate applied to the baseline flow of 61,739 persons per year implies a net excess flow of roughly 1,300 persons per year, 0.004 percent of the control-arm population, and the upper end of the 95 percent confidence interval caps it at 0.014 percent per year. Without the adjustment, the raw net change in the control arm's loss to the treated arm is about 6,200 persons per year, 0.018 percent of its population. Both bounds are about one to five percent of the annualized owner-occupancy effect documented in Section 4.5. Two caveats bound the scope of this test. The address histories record gross residential moves by owners and renters alike, so the estimate speaks to total cross-boundary migration and not to owner or renter flows separately; and the test addresses the migration channel of interference only, since credit-supply reallocation by lenders across the boundary is a distinct channel and is examined through the lender-margin results in the main text.