The Effect of Advertising on Household Mortgage Decisions*

Vikram Jambulapati† Jack Liebersohn‡

October 2026

Abstract

Using discontinuities in local television market borders, we study how mortgage advertising affects household refinancing decisions. We estimate that $10,000 of additional advertising generates 6.27 additional refinances across all lenders and 8.42 additional loans for the advertiser, with an imprecisely estimated net decline at competing lenders. Under our benchmark assumptions about households' rate of refinancing, each advertising dollar generates $6.78 in mortgage value for borrowers over five years. We calculate $1.97 in pre-tax production income per advertising dollar for advertisers using industry average mortgage origination revenues and expenses. The benefits for borrowers come almost entirely from refinancing sooner rather than from changing lenders, and we calculate positive benefits under a variety of alternative assumptions about borrowers' rate of refinancing.

JEL Classification: D12, G21, L13, L15, M37

Keywords: Mortgages, advertising, refinancing, consumer search, market expansion

*Draft date: October 5, 2026. This paper is based on Vikram Jambulapati's job market paper, which appeared as Chapter 1 of his 2025 dissertation, *Microeconomics of Consumer Financial Decisions*. We thank Chuqing Jin, Yufeng Wu, Julie Cullen, Gordon Dahl, Michelle White, Daniel Green, and David Genesove for helpful comments and suggestions, as well as participants in the AREUEA Virtual Seminar and seminars at UC Irvine.

†UC San Diego.

‡UC Irvine.

1 Introduction

Many households find financial decisions both difficult and confusing. Across a variety of financial products, households leave money on the table by paying unnecessarily high prices and neglecting valuable opportunities. For example, in the market for mortgage refinancing, high costs and inattention cost households billions of dollars per year. These mistakes raise the natural concern that households can be misled by financial advertising, which could steer them towards high-cost products rather than helping them make good choices. Consistent with this, Gurun, Matvos, and Seru (2016) show that subprime lenders with worse rates advertise more, especially to unsophisticated households. At the same time, advertising could draw attention to opportunities that households would otherwise miss.

This paper studies the effects of advertising on refinancing decisions in the U.S. mortgage market. Using discontinuities at advertising borders, we estimate how advertising affects both the overall rate of refinancing and household choice of lender. Both choices—when to refinance and which lender to choose—can affect households' mortgage costs. For example, advertising can make the possibility of a beneficial refinancing more salient, but it could also steer households towards more expensive lenders.

Our empirical strategy follows recent research on consumer-product advertising (e.g., Shapiro (2018)) by exploiting sharp variation at the borders of television media markets. Households on opposite sides of a Designated Market Area (DMA) border receive different local television advertisements. The border changes both the total amount of mortgage advertising and the set of lenders advertising locally. We use these differences to estimate the effect of advertising on total refinancing and on lender choice among households that refinance. We then combine the estimates with observed mortgage terms and assumptions about when households would otherwise refinance to calculate borrower mortgage value and advertiser returns.

Our first finding is that advertising expands the refinancing market. Using deeds data from 2007 through 2021, we estimate local projections to trace how advertising changes refinancing in ZIP codes near the border. An additional $10,000 of local television advertising across all lenders generates 6.27 additional refinances within three months under our calculations. The response is largest in the first month after an increase in advertising, and

most of the measured response occurs within three months. The estimates per advertising dollar over six and twelve months are similar. We also show that the estimated responses to advertising are larger among households with more to gain from refinancing.

Our second finding is that advertising shifts refinancing borrowers toward the advertising lender. Because deeds data do not reliably identify every lender, we study lender choice using data from the Home Mortgage Disclosure Act (HMDA). HMDA reports the year but not the month of origination, so we estimate this effect annually. Using these data, we study effects during advertising campaigns and the lender advantage outside those campaigns. An advertising campaign might directly nudge households towards refinancing right away, and could also create a persistent advantage for the advertiser that continues after the campaign ends. We find a lender share gap before campaigns begin, consistent with persistence from previous advertising. During ongoing campaigns, the gap is 0.182 percentage points larger than before the campaign, or 0.065 percentage points per additional cent of annual advertising per resident. Combining this response with estimated effects on overall refinancing gives 8.42 additional advertiser loans per $10,000 and an imprecisely estimated net decline of 2.15 loans at competitors.

What are the net costs and benefits of advertising for households? To answer this question, we use a stylized account of household attention to refinancing that allows a household to refinance later if it does not respond to an advertisement. We use the attention rate estimated by Berger et al. (2025): we assume households reconsider their mortgage once every 4.3 years on average and refinance when the available rate makes refinancing worthwhile. Applying these assumptions to our estimates, we calculate that borrowers save a total of $6.78 over the subsequent five years for each dollar spent on mortgage advertising. This calculation uses borrowers' actual estimated remaining mortgage balances and combines information on mortgage rates, loan terms, and refinancing costs from the data.

We consider a variety of alternative assumptions to understand why the estimated benefits for households are so large. Borrowers gain mostly because advertising causes them to refinance sooner: many households in the sample stand to benefit a great deal from refinancing, and the ones that benefit more are precisely the ones that are more responsive. The correlation between responsiveness and benefits is an important reason that the benefits are so high: assigning the same response to every eligible mortgage reduces the estimated benefit

to $1.21. Changes in lender choice due to advertising have relatively little effect on average mortgage costs, because advertisers offer rates and fees close to those of non-advertisers in our data. Even when we apply the higher interest and fee costs in the published example in Gurun, Matvos, and Seru (2016), who study the subprime period, we calculate $2.92 in gains for households, again coming largely from earlier refinancing.

Next, we ask how much advertisers themselves benefit from mortgage advertisements. Rather than modeling mortgage production structurally, we apply the Mortgage Bankers Association's 2021 average revenue and production cost per loan to the advertiser's additional originations. This calculation applies average production income to the additional loans. It does not require advertising and profits to be uncorrelated, but it assumes these loans earn the industry average. Production expenses include marketing and overhead, so the result is not marginal advertising profit. Our calculation yields $1.97 in pre-tax production income per advertising dollar. However, returns to advertising vary across borrowers and periods: advertising when many households would benefit from refinancing can yield higher returns because households respond more strongly when refinancing is more beneficial.

The identifying assumption for the border strategy is that mortgage demand and lender access otherwise vary smoothly across a DMA border where advertising changes discontinuously. Our ZIP-level data allow us to estimate treatment effects and covariate balance in a narrow band around the border, making this a plausible assumption. We show that the treatment effects occur sharply at DMA borders, and that the results are robust to a variety of alternate samples, control variables, and alternative estimation approaches. We also show that households across DMA borders are comparable using data on demographic characteristics, deposit quantities, branch presence, other lending, and potential refinancing gains. We also examine refinancing before advertising rises and purchase mortgages as checks on the design. The monthly estimates using deeds data show little change in refinancing before advertising rises, and purchase mortgages show little response, as we would predict. The campaign comparison further assumes that the pre-campaign share gap is a valid comparison for the gap during ongoing advertising.

The paper contributes to three areas of existing research. We add to work on mistakes and limited search in household finance (Gomes, Haliassos, and Ramadorai, 2021; Keys, Pope, and Pope, 2016; Johnson, Meier, and Toubia, 2019; Bhutta, Fuster, and Hizmo, 2026).

Previous research has shown that households fail to refinance and pay more than they need to. Our contribution is to ask whether advertising by lenders corrects those mistakes or compounds them, and how it does so. We show that the benefits to advertising depend not just on how many households respond, but also on their mortgage terms and how soon they would have refinanced anyway.

Second, we contribute to research on financial product advertising and competition. The most important precedent to this paper is Gurun, Matvos, and Seru (2016), who study advertising and mortgage expensiveness during the subprime boom. They find that lenders that advertise more originate more expensive mortgages. We therefore examine how applying the higher costs in their published example would change our results. Grundl and Kim (2019) study the targeting and effects of direct-mail refinancing advertisements. In retail banking, Honka, Hortaçsu, and Vitorino (2017) find that advertising increases awareness and search, while Chen, Hu, and Ma (2026) distinguish pricing, service, and trust-related content in bank advertising videos. We study local mortgage television advertising after 2007 and estimate its effects on total refinancing and lender choice. These effects allow us to calculate net financial costs and benefits to households as well as the effects to competitors in the mortgage market.

Finally, we contribute to research on the value of advertising for consumers. A huge literature studies how advertising can change product demand, but little is known about whether consumers actually benefit from it. Previous papers have used a variety of approaches to measure consumer benefits. Brynjolfsson et al. (2025) use an experiment to value Facebook with and without advertisements. Jeziorski and Segal (2015) estimate a dynamic model of search advertisement clicks and find that better ad placement and targeting can raise consumer welfare. Aizawa and Kim (2025) combine demand estimates with a model of advertising in health insurance marketplaces; they find that government advertising raises enrollment and welfare, while private advertising has at most modest effects on market expansion. Shapiro (2022) shows that workplace absenteeism falls due to antidepressant advertising. By studying mortgages, we are able to directly calculate the financial implications of advertising for households on a per-dollar basis. We do so by calculating the net change in payments, fees, and remaining debt, in present value terms. One benefit of our calculations is that they require minimal theoretical structure (we do not estimate a model of consumer choice), but

this means that we are not able to make calculations for households' overall welfare.

The remainder of the paper proceeds as follows. Section 2 describes the mortgage setting, local television markets, and data. Section 3 presents the two border designs. Section 4 reports the reduced-form results. Section 5 calculates borrower value and advertiser returns. Section 6 concludes.

2 Setting and Data

2.1 Advertising and household financial decisions

Households exhibit substantial inattention and inertia across financial markets. They leave savings in low-paying accounts even when better-paying accounts are available from the same bank (Adams et al., 2021), and their retirement saving depends strongly on whether employers enroll them automatically and which contributions they select as defaults (Madrian and Shea, 2001). Mortgage refinancing is another prominent example: households can leave their mortgages unchanged even when they have substantial savings available (Keys, Pope, and Pope, 2016; Andersen et al., 2020; Berger et al., 2025).

Households also often do not choose the financial product with the highest financial payoff, even when they make an active choice. Investors choose high-fee index funds even when fees are made clear, and credit-card borrowers who shop less tend to pay higher rates (Choi, Laibson, and Madrian, 2010; Stango and Zinman, 2016). Rates on observationally similar mortgages likewise vary substantially across lenders, and borrowers with less mortgage knowledge or who compare fewer offers tend to pay more (Bhutta, Fuster, and Hizmo, 2026).

Recent research has shown that reminders can lead households to refinance their mortgages, at least in some cases (Byrne et al., 2023). Mortgage advertising may perform the same function, supplied by the private market. The profits from new originations provide a private incentive for lenders to remind borrowers about the opportunity to refinance. At the same time, profitability and borrower financial gains may not be perfectly aligned. A lender benefits when a borrower chooses its mortgage, even if another lender offers better terms or refinancing is not yet worthwhile. Borrowers could nevertheless gain from an advertisement that prompts a worthwhile refinance, even if the advertiser does not offer the cheapest loan.

The mortgage market is a useful setting for valuing advertising because households' old and new financial choices have specific contractual features that can be observed in some years. To calculate the present value of both old and new mortgages, we use data on rates, fees, balances, and remaining maturity. The main missing piece is the service, convenience, and relationships associated with each mortgage, which may also influence household choice. We first ask whether advertising causes more households to refinance and then whether it changes which lender households choose. We combine both estimated responses with borrowers' potential savings and lenders' mortgage terms to calculate the financial costs and benefits of advertising.

2.2 Local television markets and DMA borders

A Designated Market Area, or DMA, is a geographic region used to define a local television market in the United States. Nielsen assigns every county to one of 210 mutually exclusive DMAs, based largely on television viewing patterns. DMAs often cross state lines but generally follow county boundaries.

Broadcast networks deliver programming through local affiliates, so DMA borders generate variation in local television advertising. Under the Federal Communications Commission's cable and satellite carriage rules, Nielsen DMA assignments determine which local stations a household receives (Federal Communications Commission, 2022). Nearby households on opposite sides of a border may therefore watch different NBC or CBS affiliates and see different local advertisements. DMA borders were created with the structure of television markets (and the placement of television antennas) in mind; while they sometimes follow administrative borders, they often do not. Researchers have used the sharp and often unusual discontinuities created by DMA borders to study television advertising in several product markets (Shapiro, 2018, 2022; Hitsch, Shapiro, and Tuchman, 2021; Sinkinson and Starc, 2019). In our setting, the main feature of borders is a change in local mortgage advertising per resident.

2.3 Data

We combine television advertising from Kantar Media, recorded deeds and mortgages, and HMDA. Throughout, we deflate spending and dollar amounts to 2021 dollars using CPI-U.

Advertising. Our advertising measure comes from Kantar Media's Ad$pender data. We aggregate local advertising expenditures to the advertiser-DMA-month level and divide by DMA population, yielding advertising dollars per resident. The total-refinancing analysis combines monthly Kantar and ZTRAX data from January 2007 through December 2021, when ZTRAX has stable national coverage. The main lender-choice sample combines annual Kantar and HMDA data from 2007 through 2021. Ad$pender disaggregates financial advertising by category, and we keep mortgage-market advertisements whenever possible.[1]

The data distinguish national advertisements, which run across media markets, from local DMA-specific advertisements. Because national advertisements reach both sides of a DMA border, we exclude them from the treatment definition. We report national and local spending separately in descriptive statistics. The estimates and calculations of returns per advertising dollar use local spending only.

We observe expenditures but do not have data on the content of individual advertisements. We therefore cannot distinguish messages about refinancing opportunities from messages about a lender's rates, service, or reputation.[2]

After removing duplicates and clearly nonmortgage advertisers, we observe $929.1 million of local mortgage advertising across 202 media markets from 2007 through 2021. Appendix A describes which advertisers we exclude.

Matching advertisers to lenders. We match advertising brands and parent companies to lenders as reported in HMDA. Most matches between advertisers and lenders are unambiguous based on names alone; when a name is ambiguous (for example, "First Bank"), we use the geography of the local advertisements to match the advertiser to the relevant local

1. The source categories are mortgage loans; home equity loans and lines of credit; and general promotion and multi-category consumer loan and credit products. The broader categories matter especially in early years. We remove documented nonmortgage advertisers, but cannot classify the content of every advertisement.

2. Chen, Hu, and Ma (2026) study bank advertising videos and the types of messages included in the advertisements to ask how lenders market themselves.

institution. We check identities by year, and account for mergers, subsidiaries, and changes in brand name. For the lender choice estimates we require a usable HMDA reporting identity and positive lending somewhere nationally (but not locally) in that year. When we study aggregate refinancing using deeds data, we do not limit advertising to matched lenders.

Two mortgage panels. We create two panels of mortgage origination data to estimate the effects of advertising on total refinancing and lender market shares. The first panel is based on deeds data, and identifies the origination month of each mortgage, but does not reliably identify the lender. The second dataset comes from HMDA, and identifies lender identities but is only available on an annual basis. We use the first panel to estimate the response of total refinancing and the second to estimate changes in lender choice.

Deeds and mortgage histories. We organize each property's history into spells. A spell begins when a mortgage is recorded on the property and ends at the next recorded purchase or qualifying refinance, which in turn begins a new spell; if neither occurs, the spell remains open until the end of the sample. In each month, the eligible stock of properties consists of the properties in an open spell, i.e., those that have an active mortgage. A single property, even without a change in ownership, can contribute more than one refinancing over the sample period.

For monthly estimates, we record whether a property refinances in each month. For cumulative estimates, we record whether it refinances at least once within the next three, six, or twelve months, so each property counts at most once at a given horizon. The main advertising exposure variable ends in December 2021, since 2021 is the last full year of Ad$pender coverage available to us. Advertising is observed through September 2022 and deeds through December 2022. We retain exposures through December 2021 for the three- and six-month spending-adjusted estimates and through September 2021 for the twelve-month estimate.

Potential refinancing gains. To estimate what a household could save by refinancing, we begin by considering its outstanding mortgage, including mortgages recorded before the advertising sample begins. We estimate the balance remaining in the month of advertising by amortizing the recorded amount at the national mortgage rate when the loan originated.

We use a consistent reported term or maturity date where available and otherwise assume a 30-year original mortgage. We value a new mortgage with the same remaining balance and maturity at the national rate observed one month before advertising. Over the following five years, or until maturity if sooner, we compare the present value of payments and remaining debt under the old and new mortgages and subtract estimated closing costs. These costs come from 2021 HMDA refinance loans of similar size, net of lender credits and counting points once. We leave potential gains unclassified when a recorded contract is nonstandard or the required information is missing. Appendix A.2 gives the calculation and explains how we allow for later refinancing when valuing the effects of advertising.

HMDA. HMDA contains loan-level mortgage data reported by covered financial institutions and is publicly available on an annual basis. We observe application information, loan purpose and amount, occupancy, property location, lender identity, and borrower demographics. The coverage of other variables varies over time; for example, interest rates are available only beginning in 2018. We use 2007–2021 to match the period of the deeds analysis. The earlier advertising files include broader consumer finance categories; we do not treat the start year alone as a guarantee that every advertisement is mortgage specific.

For each ZIP-year, we calculate each matched advertiser's share of all HMDA refinance originations in that ZIP. The panel includes lenders and DMA borders where the lender advertises on at least one side of the border at any time from 2007-2021. Our main ongoing campaign sample requires positive local advertising on the same side and zero on the other side in both the current and preceding year. To understand the persistent effects of historical advertising, we additionally estimate how advertisers' market share changes at borders in the years before a campaign begins and after it ends. The before and after comparisons use years when local advertising is zero on both sides, but households may still be influenced by past advertisements.

To assess coverage, we compare our two main data sources, ZTRAX and HMDA. We compare refinance counts from the two sources in Appendix Figure 12. From 2007 through 2021, their national annual counts have a correlation of 0.964, with qualifying ZTRAX events exceeding HMDA refinance originations by about 10 percent. The county-year correlation is 0.966, and the within-year log correlation across common ZIP-years is 0.858. Differences in

reporting coverage and the classification of refinancing events account for the gap in levels.

Rates and fees. We create a lender-level average rate measure using loan-level data on interest rates. HMDA added interest rates and detailed loan-cost fields in 2018; we construct the rate measures for 2018–2021. For each year, we regress the reported interest rate on log property value, log borrower income, log loan amount, combined loan-to-value ratio, and debt-to-income ratio, with lender fixed effects. We center the estimated lender component on that year's loan weighted mean among matched lender portfolios and add it to the monthly national mortgage rate. This gives each lender an annual relative rate while allowing national rates to change over the year. For fees, we use 2021 net loan costs as a share of the loan amount, adjusted for borrower and loan characteristics. Appendix A.1 provides additional construction details.

In 2021, lenders with usable rate estimates account for 70.5 percent of local advertising spending. We use the reference fee when a lender's fee estimate is unavailable, affecting 0.4 percent of this linked spending. Section 5.4 shows how borrower value changes when advertisers' mortgage costs differ from this benchmark.

Geography and other covariates. We designate the location of each ZIP code as its Census internal point and require at least 99 percent of its area to lie in the same DMA. We also require the DMA containing the internal point, the DMA containing most of the ZIP's area, and the DMA used to assign advertising to agree. We limit the sample to ZIP codes whose internal points lie within ten miles of the DMA border. ZIP codes can appear more than once, if their distance is less than ten miles of multiple DMA boundaries. Section 3 describes the clustering used in each design.

We also use household and demographic data from the 2015–2019 American Community Survey. For balance tests, we add prior-year income and wages from IRS ZIP files and prior-year ZIP-level employment, establishments, and payrolls.

2.4 Samples and descriptive patterns

Table 1 reports the two estimation samples; Appendix A describes their construction.

Lenders advertise in months with both large and small refinancing opportunities. Panel A of Table 2 compares households' potential rate savings in months with more and less local advertising. The rate saving is the national mortgage rate when the existing loan originated minus the current national rate; it is not a comparison of advertiser rates with their own borrowers' rates. The share of eligible mortgages with at least 50 basis points of potential rate savings is 54.2 percent when months are weighted by advertising dollars and 54.8 percent when each month receives equal weight. Panel B compares the rates offered by advertisers in 2021. The advertising weighted adjusted rate is 2.944 percent, close to the national reference of 2.945 percent.

Appendix Figure 11 shows annual real spending on the same 76 markets with explicit records in every sample month. Annual expenditure averages $63.7 million in 2007–2013, $45.6 million in 2014–2019, and $58.0 million in 2020–2021. Following the literature on television advertising, our main unit of study is dollar spending on advertising. Ideally we would also have data on the particular television audiences watching and the content of the advertisements. Spending reflects both the quantity of advertising and airtime prices.

Table 1: Estimation Samples

	(1) Total refinancing Deeds	(2) Lender choice HMDA
Years	2007–2021	2007–2021
Border pairs	413	309
ZIP codes	6,743	6,245
ZIP–border assignments	7,636	6,917
Border-pair months	70,577	–
Lender–border–year assignments	–	15,702
Lenders	–	507
Estimation observations	1,279,367	367,102

Notes: This table shows summary statistics for the deeds and HMDA panels over 2007–2021. Column (1) reports the ZIP, border, and month observations used for the three- and six-month estimates in Table 3, based on deeds data. Column (2) reports the lender, ZIP, border, and year observations used for the campaign estimates in Table 5, based on HMDA. Both samples require the ZIP's Census internal point and at least 99 percent of its area to lie in the same DMA, with the internal point more than one and no more than ten miles from the border. A ZIP can appear in several comparisons if it is near several DMA borders. Campaign assignment counts distinguish observations before and during campaigns.

Table 2: When Mortgage Advertising Occurs and Advertisers' Mortgage Rates

	Calendar-month mean	Ad-dollar-weighted mean	Difference
Panel A. National advertising timing, 2007–2021			
National mortgage rate (percent)	4.294	4.376	8.2 bp
Eligible mortgages with at least 50 bp to gain (percent)	54.8	54.2	-0.6 pp
Mean rate saving opportunity (basis points)	82.1	81.6	-0.5 bp

Panel B. Advertisers' mortgage rates, 2021	
Advertising-weighted standardized lender rate	2.944%
Advertising-weighted national mortgage rate	2.945%
Advertising-weighted lender rate minus national rate	-0.1 bp
Unweighted annual lender rate minus national reference	2.9 bp
Correlation of log annual advertising with relative lender rate	−0.127
Lenders in the annual correlation	283
Share of local advertising spending linked to a lender rate	70.5%

Notes: Advertising expenditures are in 2021 dollars. Panel A compares the simple average over the 180 months from 2007 through 2021 with the average weighted by each month's local mortgage television spending. The rate-saving opportunity for an eligible ZTRAX mortgage is the national rate when it was originated minus the current national rate. Total local advertising spending is $929.1 million. Panel B uses year-specific advertiser–lender matches and HMDA rates adjusted for borrower and loan characteristics. Each lender has one relative rate for 2021, which is added to the monthly national rate. The first three rows weight by local advertising dollars. The next two rows give each of the 283 lenders with positive annual spending and a usable rate equal weight; spending is summed over the year before taking logs. These lenders account for $50.1 million of $71.1 million in local spot spending. Appendix Table 12 reports annual correlation intervals and the corresponding 2018–2020 comparisons.

3 Empirical Strategy

3.1 Overview

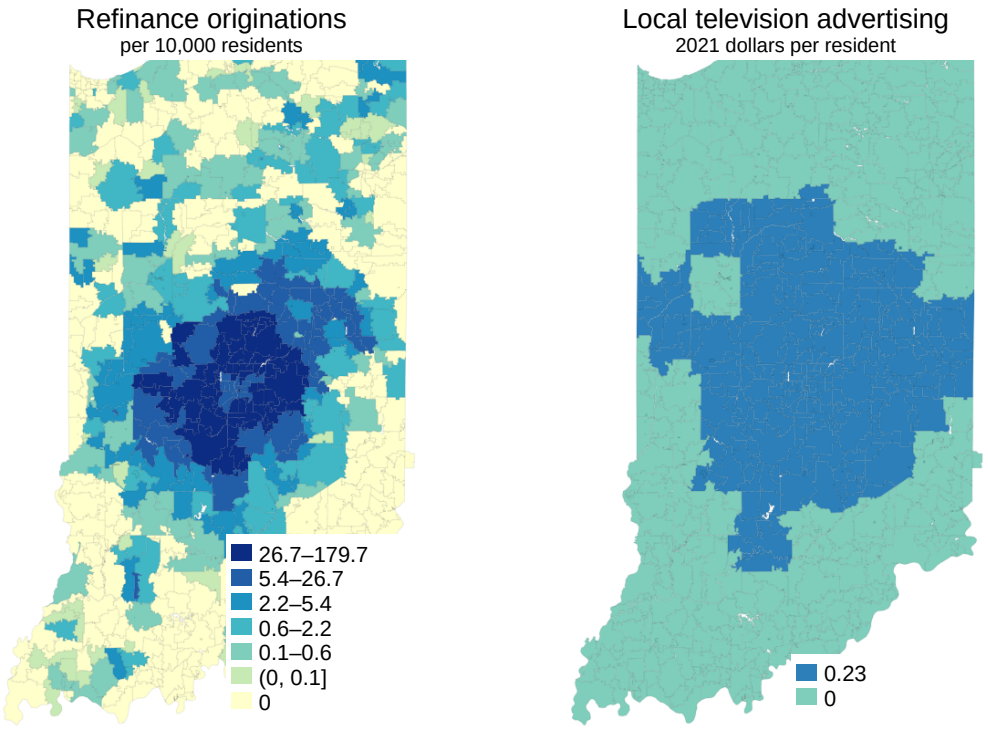
Lenders advertise the most in the regions where they do the most lending. While this association might reflect the causal impact of advertising, it could also result from branch networks or preexisting customers. A naive comparison of high- and low-advertising regions risks conflating these reasons. Following earlier DMA-border designs (Shapiro, 2018, 2022; Hitsch, Shapiro, and Tuchman, 2021; Sinkinson and Starc, 2019), we compare nearby ZIP codes on opposite sides of a DMA border. Compared to many previous studies, our data provides finer geographic coverage, which allows us to study treatment effects and covariate balance in a narrow band around DMA borders.

Figure 1 shows how mortgage lending and advertising boundaries look in practice. Both lenders illustrated in the figure, Forum Credit Union and Eastman Credit Union, lend and advertise most heavily near their bases of operation. But within the area where they originate mortgages, the DMA border creates a sharp discontinuity in exposure to advertisements. We pool these comparisons across borders, lenders, and years.

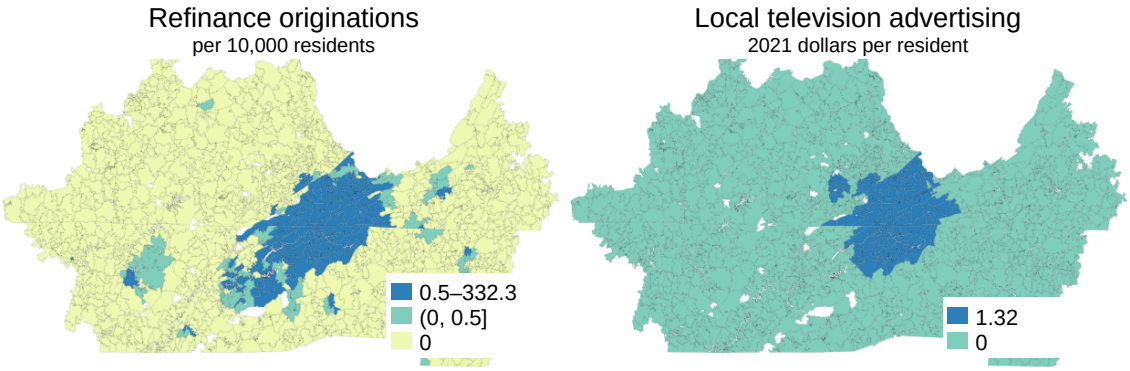
First, to study whether advertisements cause households to refinance at all, we sum refinancing and advertising across lenders within each DMA-border and month. We then use local projections to trace the effect of advertising on refinancing over the subsequent months.

Second, for each lender, border, and year, we compare the lender's refinance share in nearby ZIP codes receiving more and less of its advertising across the DMA border. Repeating this comparison across lenders, borders, and years gives us a very large number of contrasts and the statistical power to estimate lender choice.

In both samples we include ZIP codes whose internal point lies within ten miles of a DMA boundary. We exclude ZIP–border observations whose internal point is within one mile of the boundary, where advertising may spill across it. Including these observations leaves the total-refinancing result similar (Appendix Table 13). Sections 4.2 and 4.4 also report estimates from the five-mile samples.



(a) Forum Credit Union, Indianapolis, Indiana



(b) Eastman Credit Union, Tri-Cities, Tennessee–Virginia

Figure 1: Examples of Local Television Advertising at Media-Market Borders

Notes: The maps illustrate the geographic comparison for two mortgage advertisers over 2005–2019. In each example, the left map shades ZIP codes by the lender's refinance originations per 10,000 residents and the right map shades media markets by its local television advertising. In 2021 dollars, spending over this period was $0.23 per resident in Forum's advertised market and $1.32 in Eastman's; spending was zero in the surrounding markets. Lines mark DMA borders. The empirical analysis uses all qualifying borders and advertisers.

3.2 Total refinancing

Let A_{dt} denote local mortgage television advertising per resident in DMA d and month t. Let y_{dbt} denote refinances per 10,000 eligible properties on side d of border b. We estimate the monthly response to advertising at all horizons using the following specification:

$$y_{db,t+h} = \alpha_{dbh} + \lambda_{bth} + \beta_h A_{dt} + \sum_{\ell=1}^{3} \rho_{h\ell} A_{d,t-\ell} + \sum_{\ell=1}^{3} \phi_{h\ell} y_{db,t-\ell} + \varepsilon_{dbth}, \qquad h \in \mathcal{H}. \tag{1}$$

Here α_{dbh} are border-side-by-horizon fixed effects, and λ_{bth} are border-month-by-horizon fixed effects. Border-side fixed effects absorb permanent differences between the two sides, and border-month fixed effects absorb shocks common to both sides in a given month. The coefficient β_h therefore measures the relative change in refinancing h months after advertising rises, controlling for three months of prior advertising and refinancing in that area.

We scale β_h to report the change in refinances per 10,000 eligible properties caused by an additional cent of advertising per DMA resident. Regressions weight each border side by its number of eligible properties and cluster standard errors two ways by DMA and calendar month. ZIP codes in the same DMA receive the same advertising, including when they enter different border comparisons. DMA clustering allows their errors to be related across locations and over time. Month clustering allows errors across the country to be related in the same month. We also estimate separate coefficients for months -6 through -4 to test whether refinancing begins to change before advertising rises, and omit months -3 through -1, whose refinancing outcomes enter as controls.

The cumulative outcome variables measure whether an eligible property refinances within three, six, or twelve months after advertising rises. For $H \in \{3, 6, 12\}$, let E_{idbt} be the number of properties in ZIP code i on side d of border b that are eligible for refinancing in month t; and let R^H_{idbt} be the number of those properties that refinance in months $t+1$ through $t+H$. Define

$$Y^H_{idbt} = 10{,}000 \times \frac{R^H_{idbt}}{E_{idbt}}. \tag{2}$$

We estimate the three-, six-, and twelve-month outcomes at the ZIP-month level, using the same lags and fixed effects as equation (1) and weighting each observation by its stock of eligible properties.

Advertising campaigns typically last several months, so increases in advertising are predictably followed by high advertising over subsequent months as well. Appendix Figure 13 shows how innovations to advertising are associated with future advertising. An increase in advertising in month t is therefore followed by both the initial advertisement and additional spending. To estimate refinancing per dollar spent, we account for both, following Ramey and Zubairy (2018). At each horizon, we divide the total estimated refinancing response by the total estimated increase in spending over the same period:

$$\theta_H = \frac{\beta_H^R}{\beta_H^A}, \qquad S_{dt}^H = \sum_{h=0}^{H} A_{d,t+h}. \tag{3}$$

Here β_H^R is the coefficient on current advertising in the regression for Y_{idbt}^H, and β_H^A is the corresponding coefficient when total spending S_{dt}^H is the outcome. Both regressions use the same observations, controls, fixed effects, and weights. Equivalently, we regress cumulative refinancing on cumulative spending, instrumenting spending with advertising in the initial month. This gives the refinancing response per unit of additional spending, including spending after that month. Standard errors account for estimation of both coefficients. We use this calculation for the cumulative refinancing estimates and borrower savings below.

3.3 Lender choice

The lender choice analysis uses annual HMDA refinance originations. Rather than tracking mortgages over time, as we are able to with deeds data, we compare lenders' market shares on either side of a DMA border. Since HMDA does not let us track individual mortgages or properties over time, each observation is a lender, border, ZIP code, and year, including zero shares where a lender makes no refinances in that ZIP. The main dependent variable is the lender's refinance market share.

The local projections trace the response of total refinancing over the following months. Lender choice need not operate on the same timescale. A short campaign could make a lender salient to households that refinance years later. A household could also refinance with a lender in response to an advertisement and return to that lender later. These are possible persistent effects of past advertising, particularly relevant for lenders with a continuing

regional presence.

For the basic border comparison, we define the side with more advertising as having positive distance. Let D_{ijbt} indicate that side, r_{ijbt} be signed distance from the ZIP's Census internal point to the boundary, and s_{ijbt} the lender's refinance share. We estimate

$$s_{ijbt} = \alpha_{jbt} + \delta D_{ijbt} + \rho^{-}_{jbt} r^{-}_{ijbt} + \rho^{+}_{jbt} r^{+}_{ijbt} + u_{ijbt}, \tag{4}$$

where $r^{-} = \min(r, 0)$ and $r^{+} = \max(r, 0)$. Intercepts and the two distance slopes differ by lender, border, and year. We also estimate this specification with advertising per resident as the dependent variable, which gives the average jump in advertising at the border.

The jump in market share can reflect both current advertisements and historical advertisements run by the same lender in that DMA. To distinguish them, we estimate border discontinuities during ongoing campaigns, excluding the launch year, and in the year before an observed campaign begins, when local advertising is zero on both sides. The pre-campaign gap is consistent with an advantage left by past advertising. Subtracting it from the ongoing gap measures the additional effect associated with the current campaign if that earlier gap is a valid comparison.

The additional effect of an ongoing campaign can be estimated using a single regression. Let a index the lender, border, advertising side, and observed start used for the comparison, and let O_{at} indicate an ongoing year. We estimate

$$s_{iat} = \alpha_{at} + \delta_0 D_{ia} + \delta_1 D_{ia} O_{at} + \rho^{-}_{at} r^{-}_{ia} + \rho^{+}_{at} r^{+}_{ia} + u_{iat}. \tag{5}$$

The pre-campaign discontinuity is δ_0, the ongoing discontinuity is $\delta_0 + \delta_1$, and their difference is δ_1. We use the first observed start with a usable pre-campaign year and the same ZIPs before and during ongoing advertising. Ongoing years have positive advertising on the same side in the current and preceding year and zero on the other side in both; they may include later spells. The pre-campaign year enters once and each qualifying ongoing year enters once. Appendix B gives the sample rules.

We are interested in the effect of an ongoing advertising campaign per dollar of spending. We divide δ_1 by the estimated discontinuity in advertising at the same borders during ongoing campaigns. This gives the change in annual market share per additional cent of spending

per resident. The ratio is equivalently the IV coefficient from instrumenting advertising with DO while including D. We use IV standard errors.

All specifications use 2010 ZIP population weights. A ZIP can enter more than one border comparison. We cluster standard errors by ZIP and by lender–border pair across all years. ZIP clustering allows observations of the same location to move together; lender–border clustering allows observations of the same lender at the same boundary to move together over time. Neither dimension treats a repeated ZIP or campaign as independent.

Geographic plots visually show how advertising and market shares change at DMA borders. For these plots, one-mile bin indicators replace the side indicator and distance slopes. The reference bin is the nearest retained bin on the left, from two to one miles below zero, fixed at zero without a confidence interval. For the campaign plots, we also interact the bins with the ongoing indicator and display the two periods and their difference. The plots retain the same fixed effects, weights, and clustering.

3.4 Heterogeneity

The value of market expansion depends on which households refinance. We divide eligible mortgages into three groups according to their estimated five-year gains after financial refinancing costs: nonpositive gains, positive gains up to \$5,000, and gains above \$5,000. We assign each mortgage to a group before the advertisement and follow its refinancing over the next three months. For each group, we estimate the advertising response using the fraction of its eligible mortgages that refinance, weighting by that group's eligible stock. We use the same border orientation, advertising lags, outcome-history controls, and two-way clustering by DMA and month as in the aggregate regression. The coefficients on the advertising and refinancing lags, and the fixed effects, are estimated separately for each group. We compare the advertising coefficients in a joint regression on common border-months with observations on both sides.

To compare the effectiveness of different advertisers, we estimate the lender-choice regression separately for banks, independent mortgage companies, and credit unions. We use a joint version of the three regressions to test whether their effects differ. Finally, we use HMDA income data to compare lender-choice responses across borrower income groups.

3.5 Identification and interpretation

Both designs compare nearby ZIP codes that receive different local television advertising because they lie on opposite sides of a DMA border. The identifying assumption in both empirical strategies is that local mortgage demand and lender access otherwise change smoothly at the border. For the local projections using deeds data, the main concern is that some other local factor both raises mortgage demand and is correlated with advertising. For the border discontinuity estimates using HMDA data, the main concern is that a lender is more attractive or easier to reach on the side of the border where it advertises more. We conduct a series of robustness and placebo tests to support these identifying assumptions.

4 Results

4.1 Evidence at the border

Appendix Figure 6 shows how advertising and lender choice vary across the border in the broader sample with different advertising on the two sides. As in the regressions, we exclude ZIP codes whose Census internal points lie within one mile of the border because advertising may spill across it. We show the main campaign comparison in Section 4.4.

Figure 2 tests for discontinuities in predetermined local characteristics at DMA borders. These tests use equation (4) but replace lender share with the listed characteristic. The figure reports 15 ZIP characteristics, 11 measures of bank presence and lending, and borrowers' mean potential gain from refinancing. Bank measures use the lenders with verified historical banking portfolios; missing bank information is not assigned to zero. Of these 27 characteristics, one is distinguishable from zero at the 10 percent level and none at the 5 percent level.

We also test whether refinancing changes before advertising rises. To do this, we run local projections on data before innovations to advertising, dropping months -3 through -1 because they are included as controls. Figure 3 shows little change in months -6 through -4. The joint test and the corresponding sample are reported in the figure note. This check does not rule out lenders advertising in anticipation of future refinancing demand.

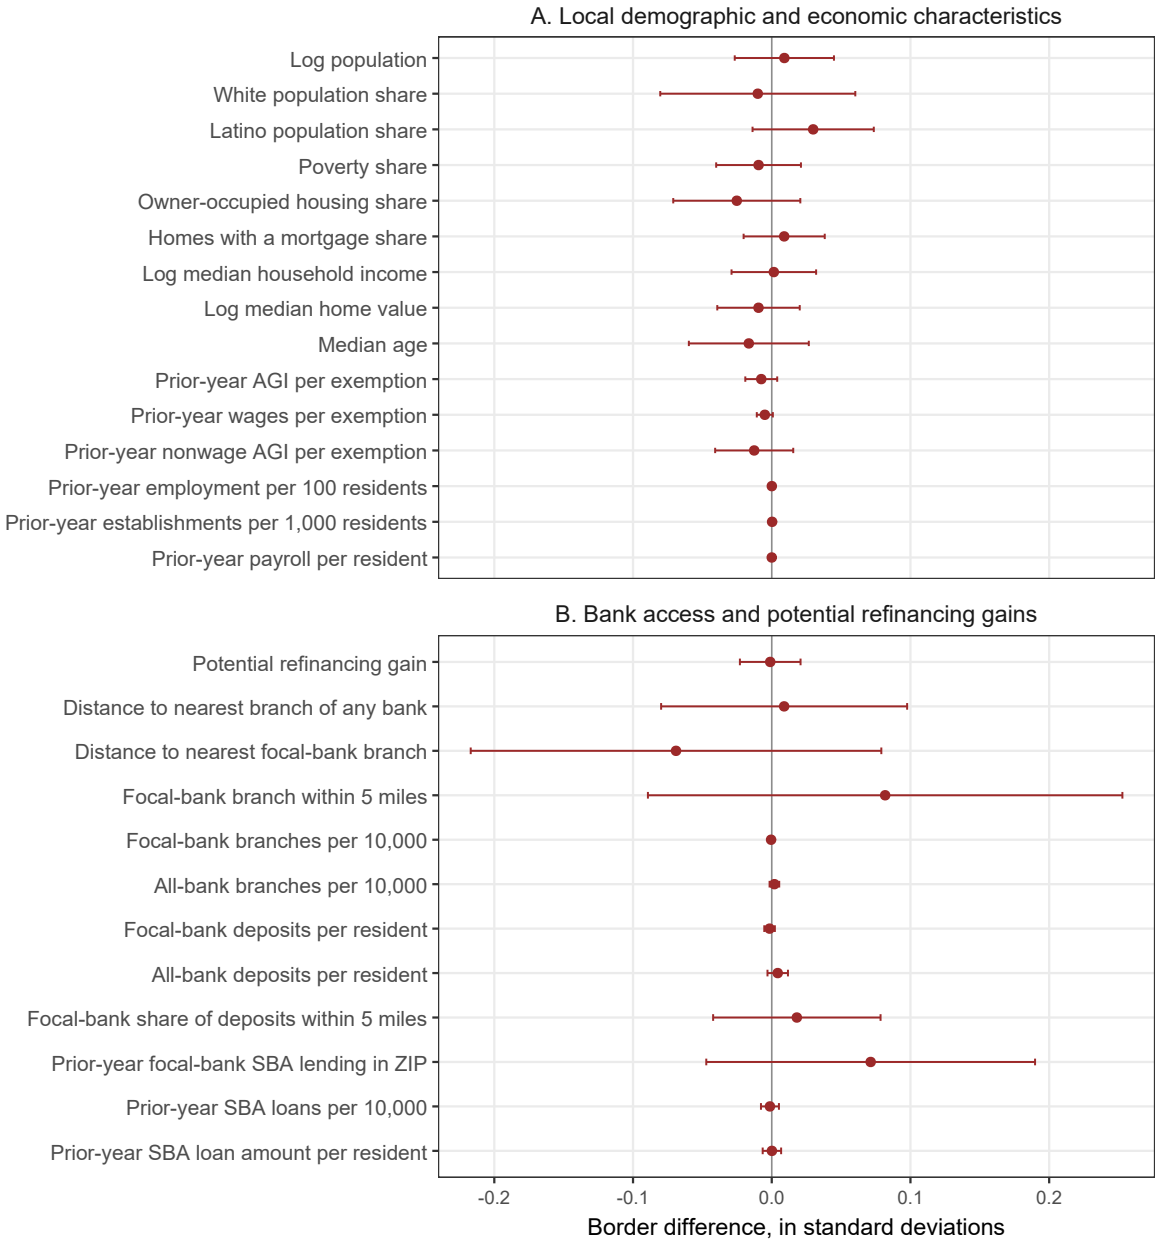


Figure 2: Covariate Balance at Media-Market Borders

Notes: Points show border discontinuities in the listed variables, in standard deviations, estimated using equation (4). Bars are pointwise 95 percent confidence intervals. Demographic measures use available ZIP and year observations in the broader lender sample. Bank measures use the subset with verified historical banking information; missing information is not set to zero. Potential gains use the previous December's mortgages. Standard deviations are calculated before duplicating observations across borders. Regressions use 2010 ZIP population weights, ZIPs between one and ten miles from the border, and standard errors clustered by ZIP and lender–border pair across years.

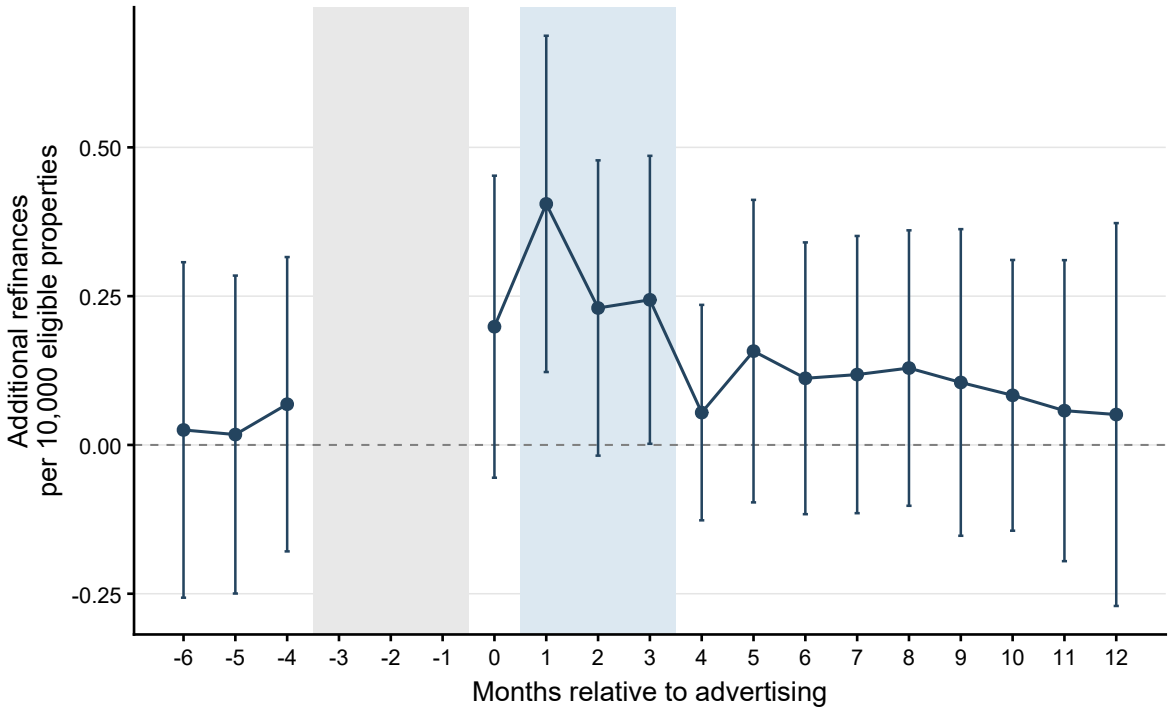


Figure 3: Effect of Local Mortgage Advertising on Monthly Refinancing

Notes: Points show the monthly refinancing response to an additional cent of advertising per DMA resident in month zero, estimated using equation (1). Refinancing is measured per 10,000 eligible properties and advertising in 2021 dollars. Responses include the effects of subsequent advertising associated with the initial increase. Months -3 through -1 enter as controls and are omitted from the plot. Blue shading marks months 1–3; gray shading marks the omitted months. The sample uses ZIP codes whose internal points are between one and ten miles from the DMA border. Estimates are weighted by the number of eligible properties. Bars are pointwise 95 percent confidence intervals, clustered by DMA and calendar month. The joint test for months -6 through -4 has $p = 0.919$.

We also examine purchase mortgages for evidence of local housing or credit shocks. We expect purchase mortgages to respond less to advertising because a buyer receives the relevant advertising at their previous residence, which may be in another DMA. Figure 4 shows small and imprecise effects for purchase mortgages. Tests of zero purchase effects have $p = 0.249$ before exposure, $p = 0.363$ through month three, and $p = 0.526$ through month six.

4.2 Total refinancing effects

Figure 3 shows the monthly refinancing response to advertising. Refinancing rises after advertising increases, with the largest effects occurring one month later. Table 3 shows an increase of 0.419 refinances per 10,000 eligible properties for each additional cent of advertising per resident during the exposure month and following three months. The estimates per dollar at six and twelve months are similar. Splitting the results by mortgage type, borrowers take out 0.390 additional conforming-size loans and 0.029 jumbo-size loans per 10,000 eligible properties within three months.

We think the timing of the response makes sense given that refinancing is not immediate. ICE Mortgage Technology reports that a refinance took an average of 59 days to close in December 2020 and 48 days in June 2021 (ICE Mortgage Technology, 2021b, 2021a). The fact that our effects peak in the month after advertising and persist for a second month is consistent with this timing. Households originating in the same month as the advertising may already have been considering refinancing or may have had preapproved offers.

A new loan is conforming-size if its original amount is at or below the Federal Housing Finance Agency's one-unit conforming loan limit for its county at the time of recording, and jumbo-size otherwise.

Robustness. In Appendix D, we show robustness by reporting estimates that use alternative geographic definitions and regression weights. The three-month effect is 0.405 (0.161) when we limit ZIP codes to be within five miles of the border and 0.411 (0.161) when we include ZIP codes within one mile. Reweighting observations so that ZIP codes assigned to more than one border receive the same total weight gives 0.431 (0.174). We also compare population weights with the eligible mortgage weights used in the main estimates.

We also show robustness to using a different method of accounting for future advertising.

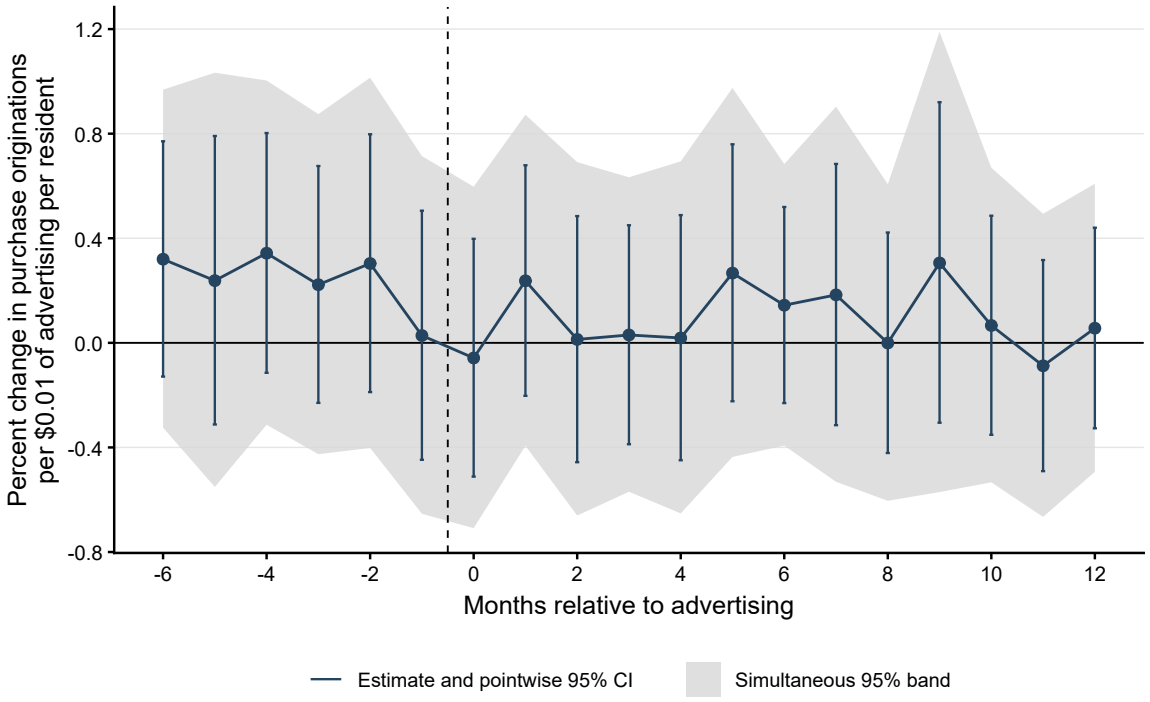


Figure 4: Effect of Local Mortgage Advertising on Purchase-Mortgage Originations

Notes: Advertising expenditures are in 2021 dollars. The figure estimates the percent change in monthly purchase-mortgage originations caused by an additional $0.01 of local mortgage television advertising per media-market resident. The purchase counts come from the deeds data. This purchase-mortgage check uses advertising exposures from 2006 through 2021, with originations followed in subsequent months. The regressions use the one-to-ten-mile border-side sample, border-side and border-pair-by-month fixed effects, and three lags of advertising. Points and bars are estimates and pointwise 95 percent confidence intervals; the gray region is a simultaneous 95 percent band. Tests of zero effects have $p = 0.249$ before advertising, $p = 0.363$ through month 3, and $p = 0.526$ through month 6. Standard errors are clustered by DMA and calendar month.

Rather than the cumulative-spending approach described in Section 3, we estimate how this month's advertising and refinancing carry into the next month, then use those relationships to calculate the response to a one-month increase in advertising. This gives an estimate of 0.555 mortgages, compared with 0.447 from local projections on the same monthly outcomes and observations (Appendix Table 13).

Finally, Appendix Table 13 reports Poisson estimates using the same controls and fixed effects. Poisson describes advertising as changing refinancing proportionally, whereas the linear regression allows the same number of additional refinances per eligible property regardless of the initial refinancing rate. The estimates need not coincide if responses differ across markets. The spending curves in Appendix F provide a separate check on the linear specification, although they are noisy.

Table 3: Effect of Local Mortgage Advertising on Total Refinancing

	Months after advertising		
	3	6	12
	(1)	(2)	(3)
All refinances	0.419***	0.451**	0.572*
	(0.162)	(0.201)	(0.298)
Conforming size	0.390**	0.413**	0.547*
	(0.162)	(0.201)	(0.298)
Jumbo size	0.029*	0.038*	0.025
	(0.015)	(0.021)	(0.029)
Observations	1,279,367	1,279,367	1,256,629

Notes: Estimates report additional refinances per 10,000 eligible properties for each additional cent of advertising per DMA resident, in 2021 dollars. Columns follow properties for three, six, or twelve months after advertising. We divide the refinancing response by the spending response over the outcome window, as in equation (3). Both loan size categories use all eligible properties as their denominator. Regressions use the controls and fixed effects in equation (1). Estimates are weighted by the number of eligible properties and use ZIP codes located between one and ten miles from the relevant DMA border. The three- and six-month columns use advertising months from January 2007 through December 2021; the twelve-month column ends in September 2021. Standard errors in parentheses account for estimation error in both advertising and refinancing and are clustered by DMA and calendar month. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

4.3 Who responds to advertising

An advertisement that prompts refinancing is more valuable when the household has more to gain. Table 4 compares borrowers by their estimated five-year savings after financial

refinancing costs. We use the three groups defined in Section 3.4: nonpositive gains, positive gains up to $5,000, and gains above $5,000. Section 2.3 and Appendix A.2 describe how we measure these gains from the existing mortgage and prevailing rates.

An additional cent of advertising per resident raises refinancing by 0.234 per 10,000 eligible properties in the lowest-gain group, 0.240 in the middle group, and 0.501 in the highest-gain group. Their observed three-month refinancing rates are 1.18, 1.70, and 2.57 percent, respectively. The increase in response rates is what we would expect if households with greater benefits from refinancing respond more to advertising. We see a similar pattern in baseline refinancing rates. It is a bit surprising that households in the "no benefit" group respond at all, and indeed that they refinance at all. We think the likely reason is that our estimates of refinancing benefits are approximations, which can blur the relationship between benefits and response rates. In Section 5.2, we use a different approach to take account of refinancing benefits by putting refinancing gains on the left side of the regression specification. This lets us estimate savings directly without first dividing households into gain groups; it does not eliminate error in the measured gains.

Table 4: Refinancing Responses by Potential Gain

	Nonpositive gains	Positive gains up to $5,000	Above $5,000
Additional refinances within three months			
Additional refinances	0.234*	0.240**	0.501*
	(0.122)	(0.117)	(0.299)
Baseline refinancing rate (percent)	1.18	1.70	2.57
High minus nonpositive gains		0.267 (0.334)	
Equal responses: p-value		0.684	

Notes: Estimates report additional refinances per 10,000 eligible mortgages for each additional cent of advertising per resident over the initial month and following three months. Groups use potential five-year savings after financial refinancing costs, measured before advertising and expressed in 2021 dollars. The middle group has positive gains up to $5,000. Each group uses its own eligible mortgage count as denominator and weight. Regressions use the controls and fixed effects in equation (1), estimated separately by group on common observations with ZIPs on both sides of the border. Baseline rates are observed three-month refinancing percentages. The final rows compare the group responses, retaining their covariance. Standard errors in parentheses are clustered by DMA and calendar month. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

4.4 Effect of advertising on lender choice

Next, we turn to the effects of advertising on lender market shares. Table 5 reports the results.

We first estimate the jump in market share during ongoing advertising campaigns. The estimated border discontinuity is 0.334 percentage points for the sample where we observe a usable year before the campaign begins.

Next, we estimate the discontinuity at the same borders in the year before the observed campaign starts. The gap is 0.152 percentage points. Market shares therefore rise by 0.182 percentage points during a campaign relative to one year before, with standard error 0.069. We use this difference as the additional effect associated with the current campaign. We also report a valuation calculation using the full ongoing gap.

We convert this change in market share to an effect per advertising dollar. At the same borders, the ongoing advertising gap is 2.808 cents per resident. Dividing the share difference by the spending difference gives 0.065 percentage points of lender share per additional cent, with standard error 0.022. Section 5 uses this estimate to calculate borrower value and advertiser returns.

Table 5: Lender Shares Before and During Advertising Campaigns

Comparison	Estimate	Standard error
Panel A. Share discontinuities (percentage points)		
Before campaign	0.152	0.053
During ongoing campaign	0.334	0.075
Ongoing minus before	0.182	0.069
Panel B. Response per cent of advertising		
Advertising gap (cents per resident)	2.808	0.487
Share response (percentage points per cent)	0.065	0.022

Notes: Panel A reports discontinuities in annual lender refinance shares, in percentage points, estimated jointly before and during campaigns using equation (5). "Before" is the year preceding the first usable observed campaign start, when local advertising is zero on both sides. "Ongoing" requires positive advertising on the campaign side in the current and preceding year and zero on the other side in both years. Launch years are excluded. Each pre year enters once and each qualifying ongoing year enters once. Panel B divides the ongoing minus before share difference by the advertising difference at the same borders during ongoing campaigns. Regressions use 2010 ZIP population weights and ZIPs between one and ten miles from the border. Standard errors are clustered by ZIP and lender–border pair across years and include estimation error in both the share and spending responses.

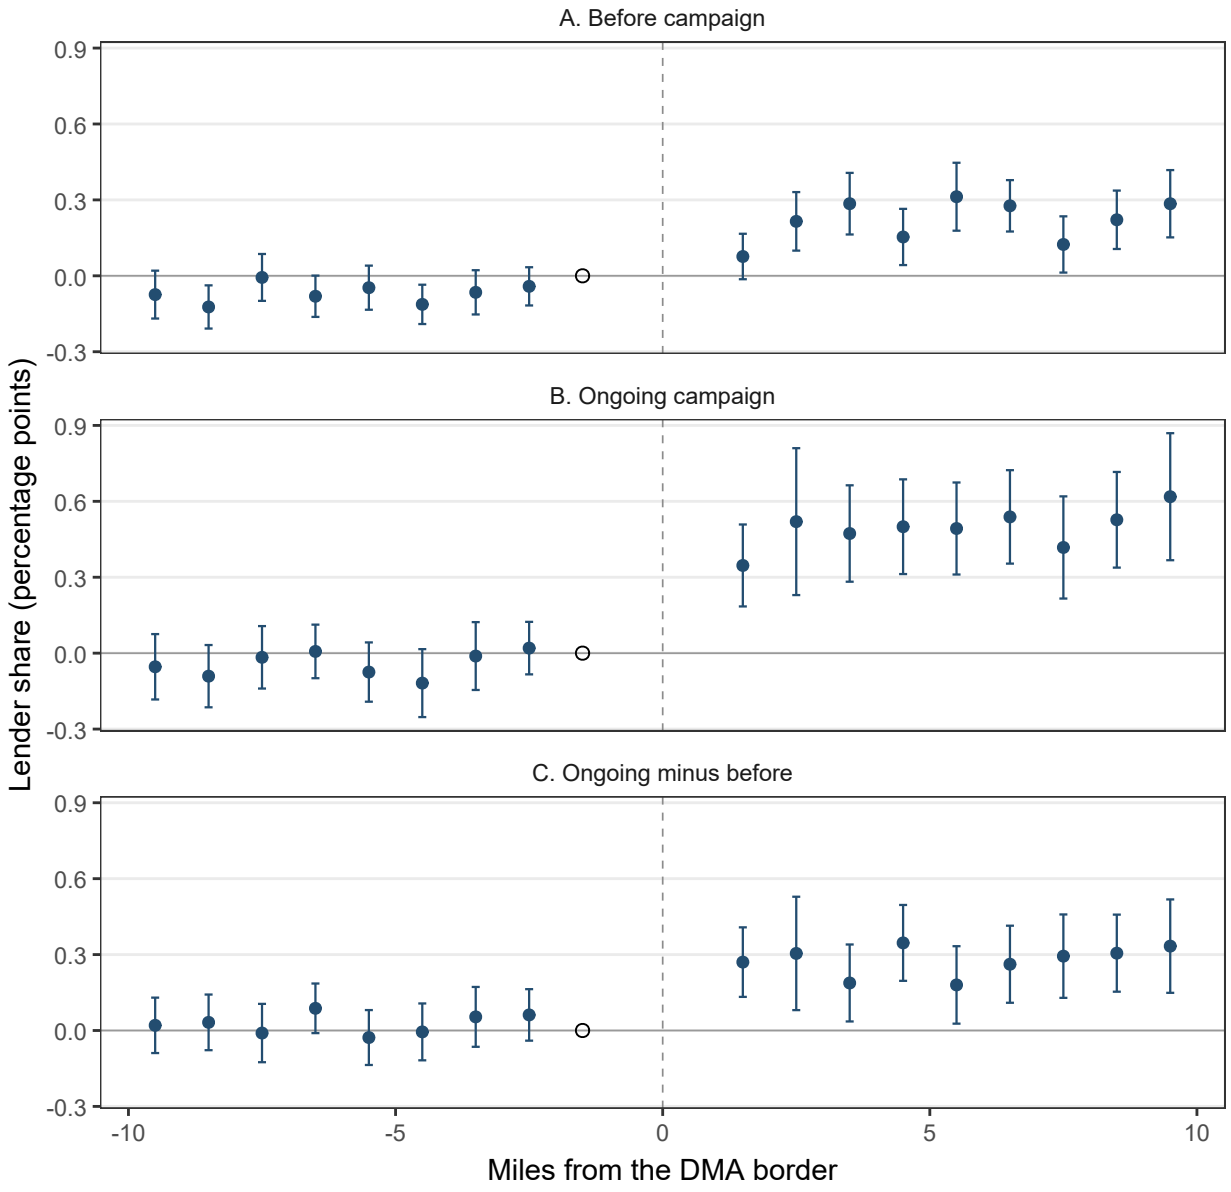


Figure 5: Lender Shares Before and During Campaigns at the Border

Notes: Points show estimated lender share differences in percentage points by one mile distance bin. Panels show the period before campaigns, ongoing advertising, and their difference, using the sample in Table 5. Distance is measured from the ZIP's Census internal point and is positive on the campaign's advertising side. Bin indicators and their interactions with the ongoing indicator replace the side indicators and distance slopes in equation (5), keeping the other fixed effects and weights the same. ZIPs within one mile are excluded because advertising may spill across the border. The reference bin, from -2 to -1 miles, is fixed at zero without an interval. Bars are pointwise 90 percent confidence intervals, clustered by ZIP and lender–border pair.

Persistence and robustness. Appendix B reports share gaps one to two, three to five, and at least six years after an observed campaign ends, when local advertising is zero on both sides. The gaps are 0.151, 0.101, 0.134 percentage points, respectively. These comparisons include campaigns without an observed start. They are consistent with persistent effects of past advertising, but do not track the same campaigns for the full period. The appendix also reports all ongoing campaigns, including those without a usable pre-campaign year.

Table 8 considers alternative campaign definitions, population, housing and mortgage market weights, within-state comparisons, polygon centroids, a narrower distance band, clustering, and Poisson regressions. The ongoing gaps remain positive across these checks. Most changes leave the campaign difference similar, although mortgage market weights and within-state comparisons give smaller estimates. When we give the pre and ongoing periods equal total weight, the pre-campaign gap moves closer to the ongoing gap. The Poisson difference is also imprecise. Table 10 carries the alternative linear share responses into the financial calculations.

5 Borrower Value and Advertiser Returns

Our goal is to calculate the average financial costs and benefits of advertising for households. We proceed in two steps. First, we convert the estimated aggregate refinancing and lender-share responses into comparable units—additional loans per advertising dollar. Second, we value the financial consequences for borrowers and lenders as a result of those new mortgages.

5.1 Additional loans per advertising dollar

An advertiser gains loans when more households refinance and when it receives a larger share of refinance originations. Let E be the increase in total refinancing per additional advertising dollar and s_0 the advertiser's initial share of originations. If its share stayed unchanged and only total originations increased, the advertiser would receive $s_0 E$ additional loans.

The estimates in Section 4.4 give the increase in market share for each additional cent of annual spending per resident. Denote this response by β_s, expressed as a fraction rather than in percentage points. In a DMA with population P, one additional advertising dollar is equivalent to $1/(0.01P)$ cents per resident. If annual refinancing volume is Q_y, the adver-

tiser's higher share adds $Q_y\beta_s/(0.01P)$ loans per advertising dollar. The net changes at the advertiser, B, and its competitors, C, are

$$B = s_0 E + \frac{Q_y \beta_s}{0.01P}, \qquad C = E - B. \tag{6}$$

The share response multiplies annual refinancing volume because both the HMDA share and advertising expenditure are measured by year. These are marginal changes per dollar; we multiply them by 10,000 for the lending figures reported in the text.

For E, we use the three-month refinancing response in column (1) of Table 3. Call this coefficient e_3: it is additional refinances per 10,000 eligible mortgages for an additional cent of advertising per resident. One dollar buys $100/P$ cents per resident; if the DMA contains M eligible mortgages, this produces $e_3 M/(100P)$ additional loans. We average $M/(100P)$ across DMA months using their shares of local advertising spending. Calling this average D_M, the conversion is simply $E = e_3 D_M$.

For β_s, we use the ongoing-minus-before estimate in Panel B of Table 5, divided by 100 to convert percentage points to a fraction. We average $Q_y/(0.01P)$ using the same advertising weights and call the result D_Q. Thus the national loan calculation is $B = s_0 E + \beta_s D_Q$. The annual volume is counted once through each year's total weight, not multiplied by twelve.

We use the campaign difference because it nets out the share advantage present before the campaign. It is not an estimate of the full long-run return to advertising. Combining it with the three-month refinancing response assumes that refinancing outside that window makes little difference to the response per dollar and that the annual and monthly estimates describe comparable changes in advertising. We also assume the border responses apply to the full DMA mortgage populations used in the calculation. Appendix A.4 reports the mortgage counts and advertising weights and a check using the twelve-month response.[3]

These are net lending changes, not identified counts of borrowers switching between particular lenders. Newly induced refinancers can borrow from the advertiser or its competitors.

3. The deeds outcomes follow each advertising month, whereas HMDA shares cover calendar years. Extending the deeds window does not make the advertising dates identical across the two designs.

5.2 Financial gains from refinancing

Advertising that nudges households towards sooner and more valuable refinancing benefits them. But the benefit is smaller if they would soon have refinanced anyway, and it can be negative if refinancing costs exceed the gain or they refinance too early. Let G_i be the financial gain from refinancing at reference mortgage terms rather than following the household's mortgage path without the initial refinance. We write

$$G_i = G_i^{\text{gross}} - F_i - L_{0i} + L_{1i}. \tag{7}$$

Here G_i^{gross} is the saving before fees assuming no later refinance, F_i is the initial refinancing fee, L_{0i} is the expected net saving from later refinancing without the initial refinance, and L_{1i} is the expected net saving from refinancing again after it. Both later savings are discounted to the advertising month and net of their refinancing fees.

We calculate each object using each mortgage's estimated remaining balance, rate, and maturity. We assume that any new mortgage begins with the remaining principal balance and matures on the same date. We then compare payments over five years and remaining debt at the end of five years (or until maturity if sooner), discounted to the advertising date. We let refinancing fees vary with the mortgage balance. Appendix A.2 provides more details on our calculation.

Following Berger et al. (2025), we assume that households reconsider their mortgages once every 4.3 years on average, an annual attention intensity of 0.234 in their model. When they reconsider, they refinance if the available rate meets the fee-dependent threshold and reduces borrowing costs after fees. We allow for one later refinance for both households that refinance as a result of advertising and households that do not. A household that has already refinanced will have less to gain from doing so again, even though it reconsiders equally often. The calculation uses the national interest rates that subsequently occurred.

To estimate the savings generated by advertising, we calculate G_i for every eligible mortgage in the advertising month. Our outcome is the sum of these gains among mortgages that refinance within three months, divided by the eligible stock and expressed per 10,000 mortgages. We use the same border regression, controls, fixed effects, weights, and clustering as for the refinancing count. We divide its response by the cumulative spending response on

the same observations. Negative gains enter with their signs. The appendix explains how we assign gains when mortgage information is missing.[4]

Let g denote the savings response divided by the spending response on the same observations, as in equation (3). Its units are dollars of savings per 10,000 eligible mortgages for an additional cent of advertising per resident. To express it per advertising dollar, we multiply by the same D_M used for the refinancing count:

$$T = gD_M. \tag{8}$$

Our calculation allows households with different gains to respond differently. For comparison, we multiply the number of additional refinances per advertising dollar on the valuation sample by the mortgage count weighted mean gain, $\bar{G}$. This gives $E_{\text{val}}\bar{G}$ and assumes the same refinancing response regardless of potential gains.

We also account for differences in advertisers' rates and fees. Let c be the advertiser's additional mortgage cost per loan relative to the reference terms offered by other lenders, including interest and fees in present value. Since the advertiser originates B additional loans per advertising dollar, total borrower financial value per dollar is

$$V = T - Bc. \tag{9}$$

The cost adjustment applies to all additional advertiser loans, including loans to households that would otherwise have refinanced with a competitor. A cheaper advertiser has $c < 0$ and adds to borrower gains.

5.3 Borrower savings and advertiser returns

Table 6 shows the estimated responses and each step of the calculation. Under the assumptions above, $10,000 of additional advertising generates 6.27 refinances overall. At the advertiser's initial share of 1.57 percent, an unchanged share would give it 0.098 additional loans.

4. The two calculations use the same 1,211,212 ZIP–border–month observations in the 171-month valuation window. We cannot calculate individual gains for 16.7 percent of eligible mortgage–month observations because mortgage information is missing or unsuitable. We assign these mortgages the average gain among eligible mortgages with usable information in the same DMA, month, and original rate-gap group. Observed refinances with negative measured gains enter with their signed values. The difference from $E\bar{G}$ depends on which borrowers respond and how the regression weights locations and months.

Its higher annual share adds 8.321 loans. Together these give 8.42 additional advertiser loans and -2.15 loans at competitors on net.

The savings regression gives $4,483.67 per 10,000 eligible mortgages for an additional cent of advertising per resident. Its standard error is $2,553.05. Multiplying by D_M, the eligible mortgage count per resident divided by 100 and averaged with advertising spending weights, gives

$$T = 4,483.67 \times 0.00149609 \simeq \$6.708$$

per advertising dollar. The savings regression uses the valuation sample with sufficient mortgage and rate information. Its refinancing-count coefficient differs from the full-period estimate and is reported in the note to Table 7.

Among advertisers with matched 2021 rates and the fee estimates described above, a $300,000 mortgage costs $87 less than the national reference on average. These lenders account for 70.5 percent of 2021 local advertising expenditure. We apply their advertising-weighted mean cost difference to each additional advertiser loan. At 8.420 additional advertiser loans per $10,000, this contributes $0.074 to borrower gains per advertising dollar. It is an average-cost comparison, not an estimate of the mortgage terms received by each marginal borrower.

Putting the two amounts together gives

$$V = T - Bc = \$6.708 + \$0.074 \simeq \$6.78.$$

The 90 percent confidence interval is [$0.50, $13.06]. It includes estimation uncertainty in the savings and spending responses while holding the financial assumptions, conversion factors, and lender-cost adjustment fixed. These are financial gains to borrowers, not a measure of total social surplus.

Advertiser returns. We calculate advertiser production income by multiplying additional loans by the Mortgage Bankers Association's 2021 average revenue and expenses per loan. Average revenue was $11,003 and production expenses were $8,664, leaving $2,339 in pre-tax production income (Mortgage Bankers Association, 2022). Applied to the additional advertiser loans, these averages give $9.26 in revenue and $1.97 in production income per

Table 6: From Advertising Responses to Borrower Savings and Advertiser Returns

Quantity	Source or calculation	Estimate
Panel A. Estimated advertising responses		
Refinancing response, e_3	Table 3, col. (1)	0.419 (0.162)
Lender share response, $100\beta_s$	Table 5, Panel B	0.065 (0.022)
Savings response, g	Section 5.2	4,483.672 (2,553.051)
Panel B. Additional loans per $10,000 of advertising		
All lenders	0.419073×14.960926	6.270
Advertiser, share held fixed	1.566875% of 6.269720	0.098
Advertiser, higher share	0.064680×128.652074	8.321
Advertising lender, total	0.098+8.321	8.420
Competing lenders, net	6.270−8.420	-2.150
	90% confidence interval	[-8.283, 3.984]
Panel C. Borrower savings per advertising dollar		
Savings from earlier refinancing	4,483.67×0.00149609	$6.708
Advertiser cost adjustment	$\frac{8.419508}{10{,}000}$×$87.39	$0.074
Total borrower savings	$6.708+$0.074	$6.782
	90% confidence interval	[$0.499, $13.064]
Panel D. Advertiser returns per advertising dollar		
Gross production revenue	$\frac{8.419508}{10{,}000}$×$11,003	$9.264
Production expenses	$\frac{8.419508}{10{,}000}$×$8,664	$7.295
Pre-tax production income at industry average costs	$\frac{8.419508}{10{,}000}$×$2,339	$1.969

Notes: Panel A reports estimated responses and standard errors in parentheses. The refinancing and savings responses are per 10,000 eligible mortgages for an additional cent of advertising per resident, including subsequent spending. The lender share response is percentage points per additional annual cent per resident. In Panel B, the first multiplier converts eligible mortgage counts per resident into loans per $10,000 of spending; the second uses annual refinance counts per resident. The advertiser's initial share is 1.57 percent. Section 5.1 explains these conversions and their timing assumptions. Panel C subtracts the advertiser's additional mortgage cost from refinancing savings, as in equation (9); the benchmark advertiser costs $87 less per loan. Panel D applies industry average revenue and expenses, including marketing and overhead, and does not measure marginal advertising profit. All amounts are in 2021 dollars. Calculations use unrounded inputs in Appendix A.4. The competitor interval assumes zero covariance between the two empirical designs. The borrower interval includes savings and spending uncertainty, holding financial assumptions, conversions, and the lender cost adjustment fixed.

advertising dollar. These calculations assume the additional loans earn industry average revenue and incur industry average costs. Marketing and overhead are included in those costs, so this is not marginal advertising profit. Marginal costs could be lower if additional loans do not require more fixed spending. Repeat business and other customer relationships may also provide income (Jin, Liebersohn, and Wu, 2026).

5.4 Alternative valuation assumptions

Table 7 shows how the benefits of advertising change under different assumptions. First, assigning the same refinancing response regardless of potential gains yields $1.21 in borrower value per advertising dollar. We multiply the estimated number of additional refinances on the valuation sample by the mortgage count weighted average gain on that sample, then retain the advertiser cost adjustment.

Next, we ask how the benefits change when we use a different assumption about later refinancing. When households reconsider their mortgages more often, more would have refinanced without the advertisement, so the benefits of advertising are lower. Reconsideration every 3.0 years gives $6.29 per advertising dollar, and every 8.5 years gives $7.47, compared with $6.78 at 4.3 years in the benchmark. We recalculate both future refinancing paths when changing this assumption. With all remaining balances set to $300,000, the gains are slightly higher at $7.39. Without refinancing fees, they are $8.15.

We separately change the advertiser's mortgage cost. The published example in Gurun, Matvos, and Seru (2016), scaled to a $300,000 mortgage, gives an additional present-value cost of $4,500 per advertiser loan. Appendix A.3 gives the conversion. Applying this cost to the additional advertiser loans subtracts $3.79 per advertising dollar from the gains due to earlier refinancing. Borrower value is then $2.92, with a conditional 90 percent interval of [-$3.36, $9.20]. This interval holds the advertiser loan count fixed. At the point estimates, an excess cost of about $8,000 per additional advertiser loan would eliminate the aggregate borrower gain.

Table 7: Borrower Value under Alternative Assumptions

Assumption	Borrower value per advertising dollar
Benchmark	$6.78
Same refinancing response regardless of potential gain	$1.21
Households reconsider every 3.0 years	$6.29
Households reconsider every 8.5 years	$7.47
No refinancing fees	$8.15
All remaining balances set to $300,000	$7.39
Advertiser charges $4,500 more per loan	$2.92

Notes: Values are in 2021 dollars. Each row changes one benchmark assumption. The benchmark uses the estimated savings response, remaining mortgage balances, fees that vary with balance, and annual attention intensity 0.234. The common-response row uses the valuation-sample refinancing coefficient 0.4741649424 and mortgage-count-weighted mean gain $1,606.60; Appendix A.4 gives the calculation. The $300,000 row changes remaining balances and recalculates fees and later refinancing, keeping rates, terms, dates, and regression observations unchanged. The attention rows use annual intensities of 0.330 and 0.118, so households reconsider every 3.0 and 8.5 years on average. The attention and no-fee rows recalculate savings with and without the initial refinance, including which later refinances are worthwhile, on the same regression sample and set of mortgages. The no-fee row sets refinancing fees to zero and keeps benchmark attention. The final row applies the $4,500 excess cost from the published example in Gurun, Matvos, and Seru (2016) and is the only row that changes the lender-cost adjustment. All rows use the same full-period estimate of additional advertiser loans and the same conversion to advertising dollars. Table 6 reports the benchmark confidence interval.

6 Conclusion

Households can leave valuable refinancing opportunities unused and may compare only a small number of lenders when they do refinance. The effects of advertising can come through both channels. On the one hand, it could steer households toward whichever lender advertises, even if that lender does not offer a particularly attractive rate; on the other hand, it could raise awareness of refinancing among borrowers who might not otherwise think of it.

We estimate that $10,000 of advertising leads to about 6.27 additional refinances within three months. Combining this response with the annual campaign-based lender-share estimate and annual refinancing volume gives 8.42 additional advertiser loans and a net decline of 2.15 loans at other lenders. The competitor effect is imprecisely estimated; we cannot reject no net change in competitors' lending. These net changes do not identify which borrowers switched lenders. We calculate the financial effects using the estimated refinancing and lender-choice responses, mortgage cash flows, and by making assumptions about later refinancing. Overall, $1 of advertising leads to $6.78 in financial savings under our benchmark assumptions, after accounting for refinancing fees. Nearly all the benefit comes from households refinancing sooner than they otherwise would. Among the lenders we can match to mortgage rates, advertisers charge rates close to the national average, so changing households' lender choice plays a quantitatively small role. Even if we assume that advertisers charge higher rates and fees than other lenders, as in the published example in Gurun, Matvos, and Seru (2016), the point estimates still imply gains for households, although the gains are smaller and less precisely estimated.

Lenders' returns also depend on when they advertise and whom they reach. Advertising to households with more to gain can be more profitable when their stronger response brings enough additional business to the advertiser. Using industry average production income from mortgage origination, we calculate $1.97 in pre-tax production income per advertising dollar.

References

Adams, Paul, Stefan Hunt, Christopher Palmer, and Redis Zaliauskas. 2021. "Testing the Effectiveness of Consumer Financial Disclosure: Experimental Evidence from Savings Accounts." *Journal of Financial Economics* 141 (1): 122–147. https://doi.org/10.1016/j.jfineco.2020.05.009.

Aizawa, Naoki, and You Suk Kim. 2025. "Public and Private Provision of Information in Market-Based Public Programs: Evidence from Advertising in Health Insurance Marketplaces." *American Economic Journal: Economic Policy* 17 (2): 84–121. https://doi.org/10.1257/pol.20220352.

Andersen, Steffen, John Y. Campbell, Kasper Meisner Nielsen, and Tarun Ramadorai. 2020. "Sources of Inaction in Household Finance: Evidence from the Danish Mortgage Market." *American Economic Review* 110 (10): 3184–3230. https://doi.org/10.1257/aer.20180865.

Berger, David, Konstantin Milbradt, Fabrice Tourre, and Joseph Vavra. 2025. "Optimal Mortgage Refinancing with Inattention." *American Economic Review: Insights* 7 (4): 497–515. https://doi.org/10.1257/aeri.20240296.

Bhutta, Neil, Andreas Fuster, and Aurel Hizmo. 2026. "Paying Too Much? Borrower Sophistication and Overpayment in the U.S. Mortgage Market." *Journal of Finance* 81 (1): 49–90. https://doi.org/10.1111/jofi.70001.

Brynjolfsson, Erik, Avinash Collis, Daniel Deisenroth, Haritz Garro, Daley Kutzman, Asad Liaqat, and Nils Wernerfelt. 2025. "The Consumer Welfare Effects of Online Ads: Evidence from a Nine-Year Experiment." *American Economic Review: Insights* 7 (4): 447–462. https://doi.org/10.1257/aeri.20240452.

Byrne, Shane, Kenneth Devine, Michael King, Yvonne McCarthy, and Christopher Palmer. 2023. *The Last Mile of Monetary Policy: Inattention, Reminders, and the Refinancing Channel.* Working Paper 31043. National Bureau of Economic Research. https://doi.org/10.3386/w31043.

Chen, Xugan, Allen Hu, and Song Ma. 2026. "How Do Banks Compete? Evidence from Advertising Videos." Working paper, February 2026. https://songma.github.io/files/chenhuma_bankimages.pdf.

Choi, James J., David Laibson, and Brigitte C. Madrian. 2010. "Why Does the Law of One Price Fail? An Experiment on Index Mutual Funds." *Review of Financial Studies* 23 (4): 1405–1432. https://doi.org/10.1093/rfs/hhp097.

Federal Communications Commission. 2022. *Definition of Television Markets for Cable and Satellite Signal Carriage.* Rules defining local television markets using Nielsen Designated Market Areas, FCC 22-89. https://docs.fcc.gov/public/attachments/FCC-22-89A1.pdf.

Gomes, Francisco, Michael Haliassos, and Tarun Ramadorai. 2021. "Household Finance." *Journal of Economic Literature* 59 (3): 919–1000. https://doi.org/10.1257/jel.20201461.

Grundl, Serafin, and You Suk Kim. 2019. "Consumer Mistakes and Advertising: The Case of Mortgage Refinancing." *Quantitative Marketing and Economics* 17 (2): 161–213. https://doi.org/10.1007/s11129-018-9207-3.

Gurun, Umit G., Gregor Matvos, and Amit Seru. 2016. "Advertising Expensive Mortgages." *Journal of Finance* 71 (5): 2371–2416. https://doi.org/10.1111/jofi.12423.

Hitsch, Günter J., Bradley T. Shapiro, and Anna E. Tuchman. 2021. "TV Advertising Effectiveness and Profitability: Generalizable Results from 288 Brands." *Econometrica* 89 (4): 1855–1879.

Honka, Elisabeth, Ali Hortaçsu, and Maria Ana Vitorino. 2017. "Advertising, Consumer Awareness, and Choice: Evidence from the U.S. Banking Industry." *RAND Journal of Economics* 48 (3): 611–646. https://doi.org/10.1111/1756-2171.12188.

ICE Mortgage Technology. 2021a. *Latest Origination Insight Report Shows Purchases Eclipse Refinances for First Time in 18 Months,* August 9, 2021. https://ir.theice.com/press/news-details/2021/ICE-Mortgage-Technologys-Latest-Origination-Insight-Report-Shows-Purchases-Eclipse-Refinances-for-First-Time-in-18-Months/default.aspx.

———. 2021b. *Origination Insight Report Sees Interest Rates Reach Historical Low in December as Refinances Continue to Surge,* January 21, 2021. https://www.businesswire.com/news/home/20210121005250/en/ICE-Mortgage-Technology-Origination-Insight-Report-Sees-Interest-Rates-Reach-Historical-Low-in-December-as-Refinances-Continue-to-Surge.

Jeziorski, Przemyslaw, and Ilya Segal. 2015. "What Makes Them Click: Empirical Analysis of Consumer Demand for Search Advertising." *American Economic Journal: Microeconomics* 7 (3): 24–53. https://doi.org/10.1257/mic.20100119.

Jin, Chuqing, Jack Liebersohn, and Yufeng Wu. 2026. "Cross-Selling Mortgages and Deposits." Working paper.

Johnson, Eric J., Stephan Meier, and Olivier Toubia. 2019. "What's the Catch? Suspicion of Bank Motives and Sluggish Refinancing." *Review of Financial Studies* 32 (2): 467–495. https://doi.org/10.1093/rfs/hhy061.

Keys, Benjamin J., Devin G. Pope, and Jaren C. Pope. 2016. "Failure to Refinance." *Journal of Financial Economics* 122 (3): 482–499. https://doi.org/10.1016/j.jfineco.2016.01.031.

Madrian, Brigitte C., and Dennis F. Shea. 2001. "The Power of Suggestion: Inertia in 401(k) Participation and Savings Behavior." *The Quarterly Journal of Economics* 116 (4): 1149–1187. https://doi.org/10.1162/003355301753265543.

Mortgage Bankers Association. 2022. *IMB Production Profits Fall in 2021 from Record 2020.* Mortgage production revenue, expense, and pre-tax net production income per originated loan, April. https://newslink.mba.org/mba-newslinks/2022/april/mba-newslink-wednesday-apr-13-2022/mba-imb-production-profits-fall-in-2021-from-record-2020/.

Ramey, Valerie A., and Sarah Zubairy. 2018. "Government Spending Multipliers in Good Times and in Bad: Evidence from US Historical Data." *Journal of Political Economy* 126 (2): 850–901. https://doi.org/10.1086/696277.

Shapiro, Bradley T. 2018. "Positive Spillovers and Free Riding in Advertising of Prescription Pharmaceuticals: The Case of Antidepressants." *Journal of Political Economy* 126 (1): 381–437. https://doi.org/10.1086/695475.

———. 2022. "Promoting Wellness or Waste? Evidence from Antidepressant Advertising." *American Economic Journal: Microeconomics* 14 (2): 439–477. https://doi.org/10.1257/mic.20190277.

Sinkinson, Michael, and Amanda Starc. 2019. "Ask Your Doctor? Direct-to-Consumer Advertising of Pharmaceuticals." *Review of Economic Studies* 86 (2): 836–881. https://doi.org/10.1093/restud/rdy016.

Stango, Victor, and Jonathan Zinman. 2016. "Borrowing High versus Borrowing Higher: Price Dispersion and Shopping Behavior in the U.S. Credit Card Market." *Review of Financial Studies* 29 (4): 979–1006. https://doi.org/10.1093/rfs/hhv072.

A Data and Financial Calculations

Early advertising records often group mortgage promotions with other consumer credit. We retain those categories but exclude sixteen advertiser groups whose documented businesses provide only nonmortgage products, such as payday loans, automobile-title loans, debt settlement, or litigation finance. We retain mixed-product firms, mortgage intermediaries, and unresolved brands. The exclusions include Freedom Financial, a debt relief advertiser, not the separate mortgage lender Freedom Mortgage. Removing an advertiser does not remove its loans from the market totals.

Advertiser matches are year-specific and distinguish a lender with no originations from an advertiser whose reporting identity is unknown. We check reporting identifiers, historical ownership, and geographic footprints, remove incorrect links, and omit unresolved acquisition years. The current panel applies the reviewed dated identities, including repaired quiet years, to every eligible border. We require positive national refinance lending in each retained lender-year, so a zero ZIP share cannot arise merely because a lender did not yet exist or had closed.

Table 1 reports the main campaign sample. The broader overall comparison has 872,388 lender–ZIP–border–year observations, 985 lenders, and 6,248 ZIP codes. The campaign sample has 367,102 observations: 95,402 before and 271,700 during ongoing advertising. Before constructing these samples, we replicate all overlapping outcomes from the dated-repair audit and preserve every eligible geographic comparison. Missing advertiser identities and missing lender-year records are not treated as zero advertising or zero lending.

Figure 12 compares deeds-based refinancing events with HMDA refinance originations. From 2007 to 2021, national annual counts have a correlation of 0.964, and deeds events number 1.099 times HMDA originations. Across county-years observed in both sources, the correlation is 0.966; within years, the correlation of log counts across ZIP codes is 0.858. Counting only deeds events with an observed earlier mortgage lowers the ratio to 0.757. The figure also shows 2022, after the advertising sample ends, but these statistics use 2007–2021 only.

For one-unit first liens in 2018–2021, we use HMDA's conforming-loan flag where available and consistent with the reported amount. Otherwise, we classify a loan only when rounding

leaves its status unambiguous. We drop loans we cannot classify, including 2011 loans with amounts between the two limits in effect that year.

A.1 Lender rates and fees

We construct standardized lender rates separately for each year from 2018 through 2021. The rate sample consists of originated 360-month, conventional, first-lien mortgages on site-built one-to-four-family properties, excluding open-end and business-purpose loans. We winsorize reported rates at the first and ninety-ninth percentiles and estimate

$$r^{w}_{i\ell t} = \alpha_{\ell t} + \beta_{1t} \log V_i + \beta_{2t} \log Y_i + \beta_{3t} \log L_i + \beta_{4t} \, \mathrm{CLTV}_i + \beta_{5t} \, \mathrm{DTI}_i + \varepsilon_{i\ell t}, \tag{10}$$

where V_i is property value, Y_i is borrower income, L_i is loan amount, CLTV is the combined loan-to-value ratio, and DTI is the debt-to-income ratio. The lender fixed effect $\alpha_{\ell t}$ is the standardized rate for legal entity ℓ. For parent companies with several legal entities, we average the available rates across the parent's matched entities using regression loan counts. Matches follow ownership in the year of the loan. We then subtract the loan weighted mean across these matched portfolios in each year. We add the parent's annual relative rate to the monthly national rate.

For lender fee comparisons, we regress 2021 loan costs net of lender credits, divided by the loan amount, on the same borrower and loan characteristics and an indicator for cash-out refinancing, with lender fixed effects. We winsorize cost shares at the first and ninety-ninth percentiles. We center the lender effects on their loan-weighted mean and add the national mean cost share. We combine legal entities within a parent using regression loan counts. Among the 283 lenders with rate estimates, 270 have adjusted fee estimates; the other 13 receive the national reference fee. These 13 lenders account for 0.4 percent of spending linked to rates. The calculation applies the matched advertisers' advertising-weighted mean excess mortgage cost to all additional advertiser loans.

Table 12 reports annual advertising and rate comparisons for 2018–2021. We sum spending within lender and year before taking logs and give each matched lender equal weight. The 2021 comparison contains 283 lenders and has a correlation of -0.127. Coverage is measured against all local advertising spending in the same year. These correlations do not

measure the mortgage cost difference for borrowers whose lender choice changes because of advertising. The financial calculation instead uses advertisers' average present value cost relative to the reference terms, weighted by advertising expenditure.

A.2 Potential gains and later refinancing

A borrower gains from refinancing when the reduction in payments and debt exceeds the cost of the new loan. Let P_s^o and P_s^n be the old and new scheduled payments, and B_H^o and B_H^n the remaining balances after H months. Potential net savings are

$$S = \sum_{s=1}^{H} \frac{P_s^o - P_s^n}{(1+d)^s} + \frac{B_H^o - B_H^n}{(1+d)^H} - F(B), \tag{11}$$

where B is remaining principal and d is the new monthly mortgage rate. Both loans start with the same principal and mature on the same date. We set H to 60 months or the remaining term, whichever is shorter. We include the difference in remaining balances so that a loan that repays principal more slowly does not appear to save money.

We approximate the old mortgage rate with the national rate when the mortgage was recorded. To assign borrowers to gain groups, we use the national rate one month before advertising as the new mortgage rate. To value refinancing in response to advertising, we use the rate in the advertising month and the rates in later months. We use a consistent reported term or infer one from a maturity date near a standard 10-, 15-, 20-, 25-, or 30-year contract; otherwise we assume a 30-year term. We do not directly calculate gains for known nonstandard contracts or for mortgages whose required information is uncertain. Group membership is set before advertising and does not change during the three-month outcome window. We do not observe the household's actual rate, credit eligibility, or refinancing offers.

The fee schedule comes from 3,632,947 conventional, non-cash-out HMDA refinance originations in 2021 with usable loan costs. We subtract lender credits from total loan costs, counting discount points only once, and interpolate mean net costs across loan-amount groups. We calculate amortization and discounted cash flows in nominal dollars, convert the results to 2021 dollars, and apply the 2021 fee schedule in real terms to every year. The fees do not

include the household's time and effort.

A household reconsiders its mortgage in a given month with probability $1 - \exp(-a/12)$, where a is the annual attention intensity. A later refinance is worthwhile in a month when the available rate meets the fee-dependent threshold and, after fees, lowers borrowing costs over the rest of the valuation window. In such a month, the household refinances with probability equal to the chance that it reconsiders that month times the chance that it has not already used its one later refinance in an earlier month. We allow this one later refinance whether or not the household first refinances in response to advertising, and every refinance pays the applicable fee. With zero fees, a later refinance must strictly lower borrowing costs, so the household does not use up its later refinance on a loan that saves nothing.

When we cannot calculate a mortgage's savings, we assign it the average gain among eligible mortgages with usable information in the same DMA, month, and original rate-gap group. These assigned gains enter only the dollar-valued outcome; the gain-group regressions in Table 4 use only mortgages whose savings we can calculate. The benchmark and common-response calculations use the same observations. When we change the attention rate or fees, we recalculate savings with and without advertising but keep the observations and the set of mortgages with usable information fixed.

A.3 The higher advertiser-cost comparison

Gurun, Matvos, and Seru (2016) study subprime purchase mortgages in 2002–2006. In their adjustable-rate example, the reported estimate implies a 7.11-basis-point increase in the reset rate per \$1,000 of quarterly advertising, or $7.11 \times 25.46 = 181.02$ basis points at average spending of \$25,460. They calculate that this premium costs a borrower about \$3,000 in present value on a \$200,000 mortgage, assuming the borrower pays it for one year after the two-year reset and discounting at 10 percent. We apply the implied cost of 1.5 percent of the balance to a \$300,000 loan, which gives \$4,500 per advertiser loan.

When we change the advertiser's excess cost, we hold T and B fixed, so only the term Bc in equation (9) changes.

A.4 Annual lender shares and returns per advertising dollar

We convert the border estimates to effects per advertising dollar using eligible mortgages and annual refinancing flows throughout each DMA. Let M_{dt} be eligible properties in month t, Q_{dy} be refinances during calendar year y, and P_d be the DMA's 2010 population. Let w_{dt} be the share of 2007–2021 local mortgage advertising spending in that DMA and month, so that $\sum_{dt} w_{dt} = 1$. If e_H is the estimated refinancing response per 10,000 properties per additional cent per resident, the national calculation is

$$D_M = \sum_{dt} w_{dt} \frac{M_{dt}}{100 P_d}, \qquad D_Q = \sum_{dt} w_{dt} \frac{Q_{d,y(t)}}{0.01 P_d}, \tag{12}$$
$$E_H = e_H D_M, \qquad B_H = s_0 E_H + \beta_s D_Q, \qquad C_H = E_H - B_H.$$

One dollar buys $1/(0.01 P_d)$ cents per resident, and the refinancing coefficient is per 10,000 properties. This gives the factor $M_{dt}/(100 P_d)$. The annual HMDA advertising measure sums spending over the year, so its share coefficient multiplies the twelve-month flow Q_{dy}. Summing the monthly spending weights within a year gives that DMA-year's share of spending. The counts cover each DMA, with no duplication for ZIP codes near several borders. Both conversion factors use the population denominator used for local advertising per resident. We obtain $D_M = 0.001496092551$ and $D_Q = 1.286520736307$, equivalent to an annual baseline of 859.92 refinances per 10,000 eligible properties under these weights. We apply the pooled border responses to this observed national stock and flow.[5]

The loan counts use the following inputs:

$$e_3 = 0.4190730143 \qquad (\text{SE} = 0.1617172245),$$
$$\beta_s = 0.000646804172 \quad (\text{SE} = 0.000223039501),$$
$$s_0 = 0.015668749801.$$

We retain the reference s_0 calculated from unique lender–ZIP–year observations with known prior-year shares, weighting by ZIP population. This calculation uses 2008–2021 because the

5. The counts use the same amount, maturity, and refinancing definitions as the regressions. The retained deeds files cover the 50 states and District of Columbia, but Vermont has no qualifying events. We map 99.755 percent of recorded eligible mortgage-months to DMAs; the rest are excluded. These are registered properties covered by the data, rather than a census of outstanding mortgages.

reviewed prior-year lender records begin in 2007; unavailable prior shares are not assigned zero. The savings coefficient is $g = 4,483.67222144$ dollars per 10,000 eligible mortgages per additional cent per resident, with standard error 2,553.05094466. Multiplying by D_M gives \$6.708 in savings per advertising dollar. The lender-cost adjustment uses $c = -87.39192547$ dollars per additional advertiser loan. Calculations use the full-precision saved estimates.

The common-response row uses $e_3 = 0.4741649424$ on the valuation sample and an average gain of \$1,606.60, giving $e_3 D_M \bar{G} - Bc = \$1.21$. This mean divides total observed and imputed gains by eligible mortgage counts in the same ZIP–border–month sample.

In the border sample, the three-month baseline is 163.26 refinances per 10,000 properties. Multiplying it by four gives 653.04, close to the observed annual flow of 662.31 on the same sample. Using the twelve-month spending-adjusted response in the national calculation gives 8.56 additional refinances per \$10,000 of advertising, 8.46 additional advertiser loans, and a net competitor change of 0.11. The borrower-value calculation uses the three-month savings outcome throughout, including its joint covariance with spending. We do not infer a twelve-month savings response from the refinancing count.

B Additional Lender Choice Results

The main comparison includes each pre-campaign observation once and every qualifying ongoing year once, using ZIP population weights. ZIPs near several borders enter each eligible comparison. Table 8 changes one part of this design at a time. All estimated variants remain available in the replication files.

The all-ongoing check adds campaigns without a usable pre year. The recent-advertising check removes pre years immediately following a campaign end. Equal-period weights repeat each pre observation for every ongoing year; they do not give each region equal weight or remove repeated borders. Within-state comparisons exclude locations within one mile of state-line segments and estimate separate intercepts and slopes within states. Housing weights exclude 15 ZIPs with zero housing units. Population weights on that support give a difference of 0.182 (0.069).

Table 8: Robustness of Lender Choice Estimates

Specification	Before	Ongoing	Difference	Share per cent	N
Panel A. Linear estimates (percentage points)					
Main comparison	0.152 (0.053)	0.334 (0.075)	0.182 (0.069)	0.065 (0.022)	367,102
All ongoing campaigns	0.152 (0.053)	0.292 (0.070)	0.139 (0.068)	0.051 (0.023)	451,851
Drop recently advertised pre years	0.171 (0.061)	0.334 (0.075)	0.163 (0.076)	0.058 (0.026)	346,309
Omit first ongoing year	0.169 (0.074)	0.417 (0.107)	0.248 (0.101)	0.070 (0.025)	210,326
Two quiet pre years	0.170 (0.055)	0.292 (0.085)	0.122 (0.071)	0.047 (0.021)	316,144
Use each current spell	0.167 (0.057)	0.334 (0.076)	0.167 (0.068)	0.059 (0.021)	380,262
Equal total weights across periods	0.253 (0.068)	0.334 (0.075)	0.081 (0.051)	0.029 (0.018)	543,400
Compare within states	0.105 (0.062)	0.198 (0.068)	0.093 (0.062)	0.034 (0.022)	265,901
Weight by refinance market size	0.167 (0.052)	0.258 (0.053)	0.091 (0.056)	0.030 (0.018)	367,102
Weight by housing units	0.136 (0.052)	0.321 (0.069)	0.185 (0.065)	0.067 (0.023)	366,452
Polygon centroid geography	0.146 (0.052)	0.323 (0.068)	0.177 (0.064)	0.063 (0.022)	370,970
Within five miles	0.037 (0.075)	0.197 (0.120)	0.160 (0.113)	0.051 (0.029)	139,406
Cluster by DMA and year	0.152 (0.052)	0.334 (0.070)	0.182 (0.043)	0.065 (0.015)	367,102
Panel B. Poisson estimates (percent gaps)					
Main comparison	15.51 (7.47)	25.24 (8.73)	9.73 (7.95)	–	237,787
Equal total weights across periods	24.73 (11.33)	25.24 (8.73)	0.51 (7.42)	–	352,065
Compare within states	6.39 (7.43)	11.33 (7.82)	4.94 (7.27)	–	163,671

Notes: Panel A estimates equation (5). Entries are share gaps before and during campaigns, their difference, and the difference divided by the ongoing spending gap. Samples, controls and weights follow Table 5 except as indicated. Parentheses contain standard errors clustered by ZIP and lender–border pair; the DMA and year row uses 171 DMAs and 15 years. Panel B estimates $E[s_{ic}|X] = \exp(F_{ic}+\eta H_{ic}+\delta H_{ic}O_c)$ with the same fixed effects and slopes. Before and ongoing entries are $100(e^{\eta}-1)$ and $100(e^{\eta+\delta}-1)$; their difference uses the joint covariance. Poisson retains population weights but drops all-zero groups and separated observations. Its percent effects are not percentage points.

Table 9: Lender Shares During and After Advertising Campaigns

Comparison	OLS (pp)	N	Poisson (%)	N
All active advertising comparisons	0.161 (0.039)	872,388	5.47 (2.02)	590,523
All ongoing campaigns	0.292 (0.070)	356,449	15.73 (5.59)	238,864
1–2 years after campaign ends	0.151 (0.041)	460,502	9.87 (3.88)	277,595
3–5 years after campaign ends	0.101 (0.034)	337,580	8.04 (3.86)	200,919
6+ years after campaign ends	0.134 (0.038)	261,637	15.17 (5.20)	165,932

Notes: Each row estimates $s_{ic} = \alpha_c + \rho_c^- r_{ic}^- + \rho_c^+ r_{ic}^+ + \eta H_{ic} + \epsilon_{ic}$ by OLS, and the corresponding exponential conditional mean by Poisson. Here c denotes a lender–border–year comparison, H the advertised side, and r signed distance. OLS entries are percentage points; Poisson entries are $100(e^{\eta} - 1)$ percent. The overall row uses all active lender assignments with different advertising across the border. Ongoing campaigns require advertising on the same side in the current and preceding year. Post-campaign rows require zero local advertising on both sides and measure time since the most recent observed one-sided campaign end. These rows pool available campaigns; they need not observe a start or follow the same campaigns over time. Regressions use 2010 ZIP population weights, the one-mile exclusion, and ZIP and lender–border clustered standard errors in parentheses.

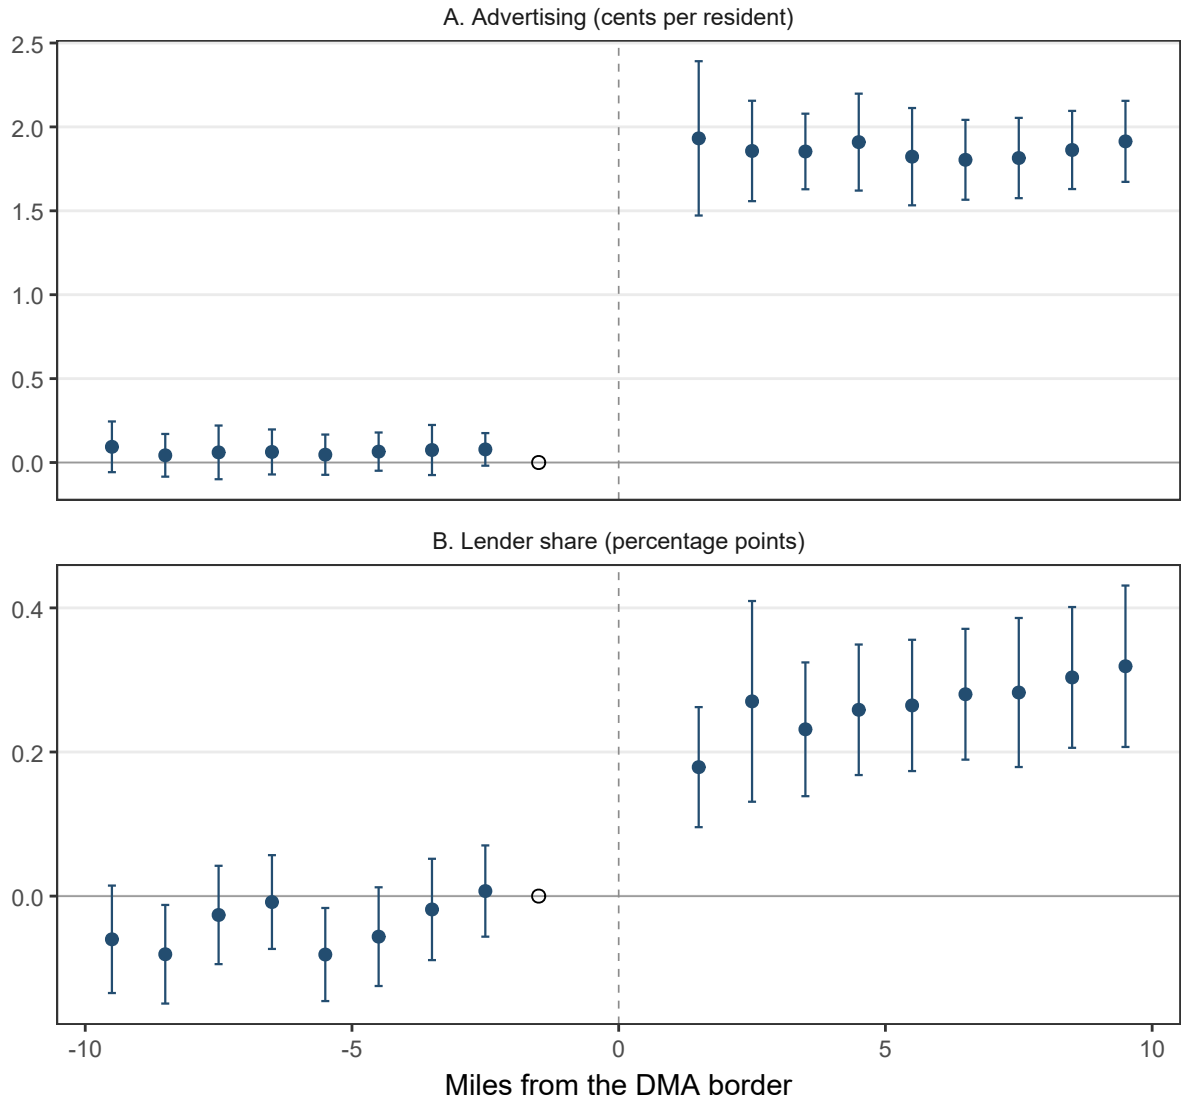


Figure 6: Lender Advertising and Lender Choice at Media-Market Borders

Notes: Advertising expenditures are in 2021 dollars. The figure pools lender–border–year assignments in the primary one-to-ten-mile sample. Distance is signed so that the side receiving more advertising from the focal lender is positive. We divide distance into one-mile bins and regress each outcome on the bin indicators, omitting the nearest retained bin on the left, from two to one miles below zero. The open marker at -1.5 miles shows this reference bin, fixed at zero without a confidence interval. We exclude ZIP codes whose Census internal points lie within one mile of the border because advertising may spill across it. Panel A reports focal-lender television advertising in cents per resident, and Panel B reports the focal lender's percentage share of local HMDA refinance originations. The regressions include lender–border–year fixed effects and use 2010 ZIP-population weights. Bars are pointwise 90 percent confidence intervals with two-way clustering by ZIP code and lender–border across years.

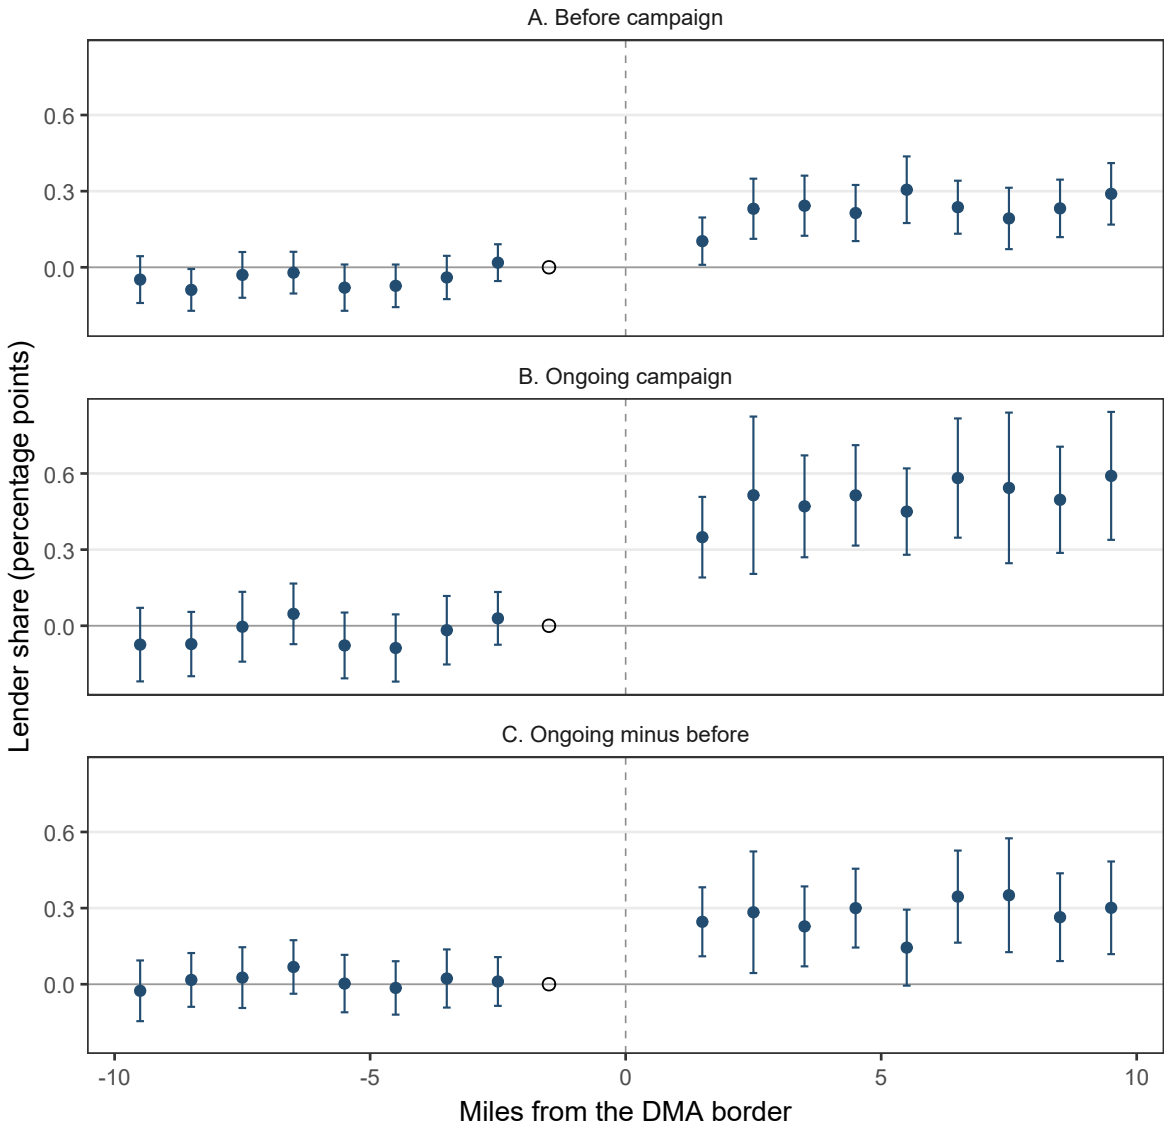


Figure 7: Campaign Share Differences Using Polygon Centroids

Notes: Points show one-mile-bin coefficients before campaigns, during ongoing advertising, and their difference, as in Figure 5. The sample is rebuilt using each ZIP polygon's area centroid rather than its Census internal point, including ZIPs that enter the one-to-ten-mile band. The DMA assignment and 99 percent area requirement are unchanged. The reference bin from −2 to −1 miles is fixed at zero without an interval. Bars are pointwise 90 percent confidence intervals, clustered by ZIP and lender–border pair.

B.1 Advertising spending bins

We replace advertising spending with indicators for its ten positive-spending deciles, using zero spending as the reference. Figure 8 plots the estimated coefficients during ongoing campaigns. The estimated share differences generally rise with spending, although the lower deciles are noisy.

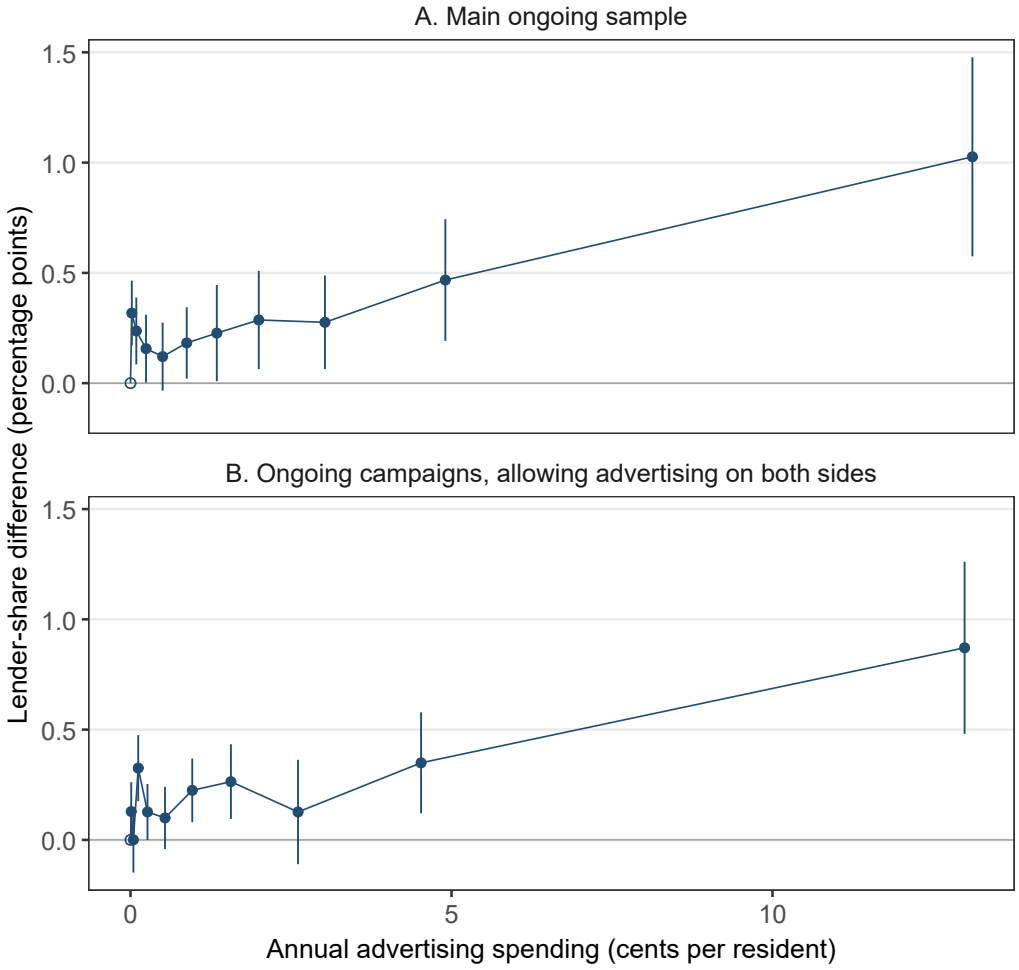


Figure 8: Lender Shares by Advertising Spending

Notes: The regression replaces continuous advertising with spending-bin indicators:

$$s_{ic} = \alpha_c + \xi_c^+ r_{ic}^+ + \xi_c^- r_{ic}^- + \sum_{k=1}^{10} \beta_k \mathbf{1}\{A_{ic} \in I_k\} + \epsilon_{ic}.$$

Here s is lender share in percentage points, A annual advertising in cents per resident (2021 dollars), i a ZIP, and c a lender–border–campaign–year comparison. Signed distance r enters as $r^+ = \max(r, 0)$ and $r^- = \min(r, 0)$, with separate slopes for every comparison. Each side enters its own spending bin. The intervals I_k are positive-spending deciles, formed with 2010 ZIP population weights; zero spending is omitted. Points show $\hat{\beta}_k$ at each bin's weighted mean spending; the zero reference has no interval. Panel A uses the main ongoing sample. Panel B also includes campaigns without a usable pre year and allows advertising on both sides, requiring the higher-spending side to have positive, higher spending in the preceding year. Both panels retain the one-mile donut and population weights, with ZIP and lender–border clustering across years. Bars are pointwise 90 percent intervals. These are ongoing share differences, without subtracting the pre-campaign gap.

B.2 Lender shares and advertising responses

The basic logit model predicts that a lender's share response to advertising is proportional to $s(1-s)$, where s is its initial share, if advertising enters each lender's utility with the same coefficient. For small shares, this is approximately proportional to the share itself. Figure 9 shows larger ongoing share differences for lenders with larger pre-campaign shares. These lenders also tend to spend more on advertising.

To account for spending, Figure 10 divides each group's ongoing-minus-before share response by its advertising response. The point estimates show no clear proportional increase with initial share. This pattern is consistent with advertising making a lender salient. The estimates are imprecise, however, and do not establish a rejection of logit. Since the second figure also subtracts the pre-campaign gap, spending alone need not explain the difference between the two figures.

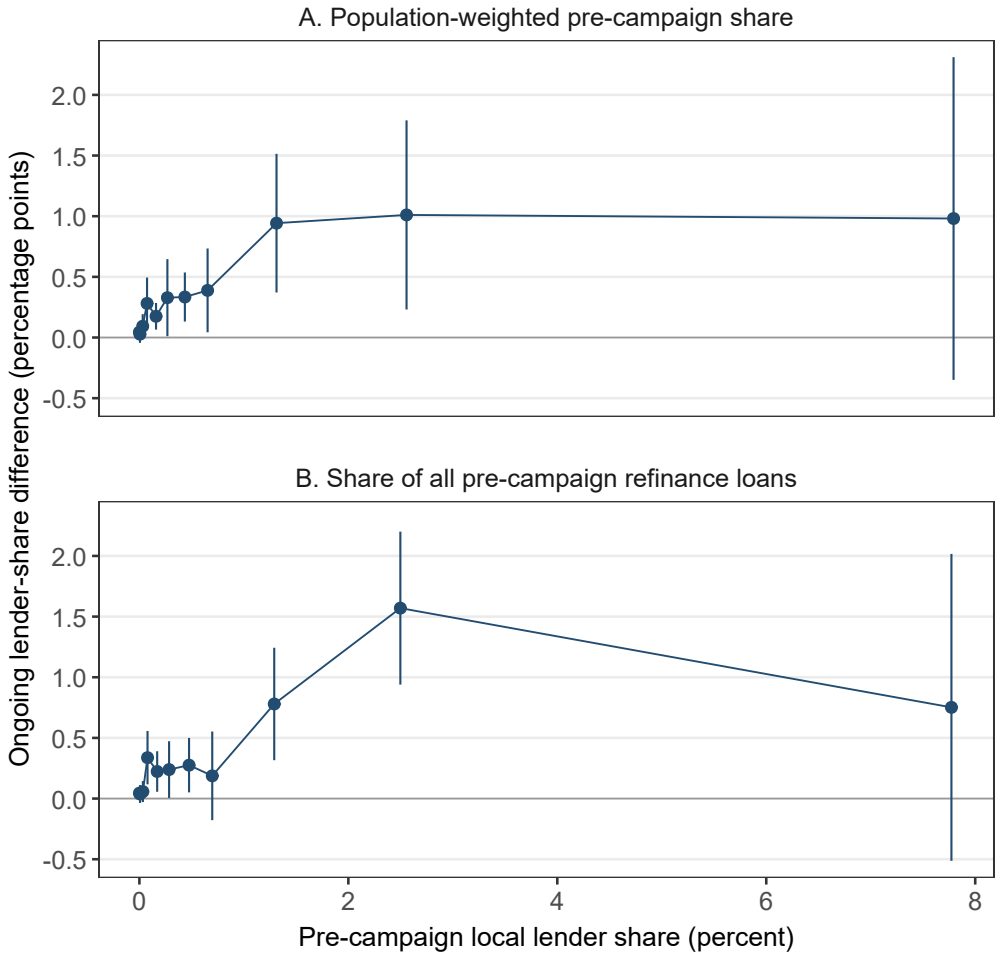


Figure 9: Ongoing Advertising Responses by Initial Lender Share

Notes: The regression gives each initial-share group its own before and ongoing gaps:

$$s_{ic} = F^s_{ic} + \sum_{k=0}^{10} (\eta_k H_{ic} + \delta_k H_{ic} O_c) \mathbf{1}\{g_c = k\} + \epsilon_{ic}.$$

Share s is in percentage points. Here i is a ZIP and c a lender–border–campaign–year comparison; H_{ic} marks the campaign's advertised side and O_c an ongoing year. Writing r for signed distance in miles, $r^+ = \max(r, 0)$ and $r^- = \min(r, 0)$. Each comparison has its own intercept and two slopes, $F^W_{ic} = \alpha^W_c + \xi^{W,+}_c r^+_{ic} + \xi^{W,-}_c r^-_{ic}$. Within-state comparisons also separate these terms by state. The group g_c uses the lender's refinance share in the year before the campaign, over the same ZIPs retained on both sides. Zero observed loans form group 0; positive shares form ten deciles using ongoing population weights. Panel A uses the population-weighted average of ZIP shares; Panel B uses lender loans divided by total refinance loans in those ZIPs. Only the grouping changes: all regressions use 2010 ZIP population weights. Points show $\hat{\eta}_k + \hat{\delta}_k$ at each group's weighted mean initial share; bars are pointwise 90 percent intervals. The one-mile donut and ZIP and lender–border clusters across years are retained. Freedom Financial is excluded; the joint sample has 367,102 observations. An observed zero loan count is not assumed to be a zero choice probability.

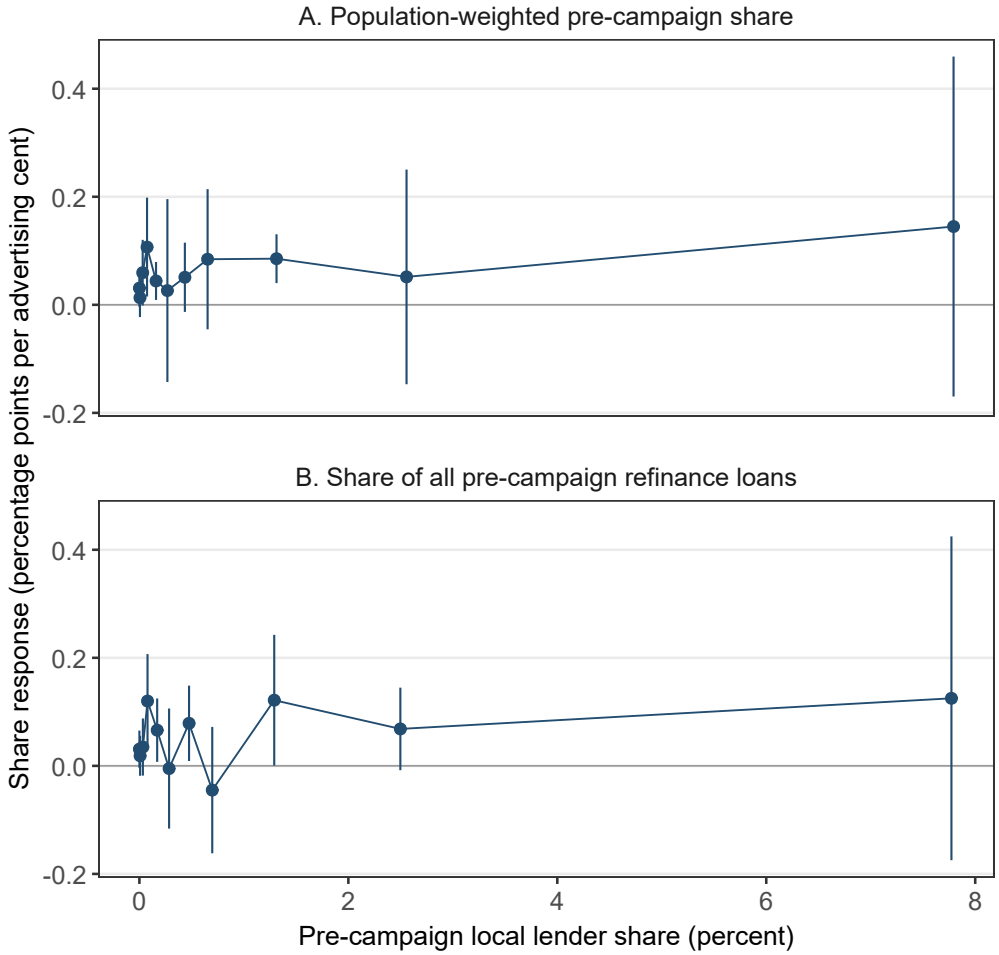


Figure 10: Advertising Responses per Cent by Initial Lender Share

Notes: The share regression and the advertising regression during ongoing years are

$$s_{ic} = F^s_{ic} + \sum_{k=0}^{10} (\eta_k H_{ic} + \delta_k H_{ic} O_c) \mathbf{1}\{g_c = k\} + \epsilon_{ic},$$

$$A_{ic} = F^A_{ic} + \sum_{k=0}^{10} \pi_k H_{ic} \mathbf{1}\{g_c = k\} + u_{ic} \qquad (O_c = 1).$$

Share s is in percentage points and annual advertising A in cents per resident. Here i is a ZIP and c a lender–border–campaign–year comparison; H_{ic} marks the campaign's advertised side and O_c an ongoing year. Writing r for signed distance in miles, $r^+ = \max(r, 0)$ and $r^- = \min(r, 0)$. Each comparison has its own intercept and two slopes, $F^W_{ic} = \alpha^W_c + \xi^{W,+}_c r^+_{ic} + \xi^{W,-}_c r^-_{ic}$. Within-state comparisons also separate these terms by state. Group g_c is the pre-campaign lender-share group: zero observed loans or one of ten positive-share deciles. The groups and joint sample are those in Figure 9. Points show $\hat{\delta}_k/\hat{\pi}_k$, the ongoing-minus-before share gap divided by the ongoing advertising gap, at each group's weighted mean initial share. Intervals account for both regressions and their covariance using the equivalent IV regression. All regressions use 2010 ZIP population weights, the one-mile donut, and ZIP and lender–border clusters across years. Bars are pointwise 90 percent intervals. These are campaign comparisons, not directly estimated derivatives of individual choice probabilities.

C Additional Financial and Descriptive Results

Table 10: Financial Calculations with Alternative Lender Responses

Lender choice estimate	Advertiser loans	Competitor loans	Borrower value	Production income	High-cost value
Main comparison	8.42	-2.15	$6.78	$1.97	$2.92
Equal total weights across periods	3.79	2.48	$6.74	$0.89	$5.00
Within-state comparisons	4.50	1.77	$6.75	$1.05	$4.68
Refinance market weights	4.02	2.25	$6.74	$0.94	$4.90
Housing unit weights	8.67	-2.40	$6.78	$2.03	$2.81
Polygon centroid lender estimate	8.25	-1.98	$6.78	$1.93	$2.99
Full ongoing share gap	15.39	-9.12	$6.84	$3.60	-$0.22

Notes: Loan counts are per $10,000 of advertising; values are per advertising dollar. Each row replaces only the lender share response in equation (6). Total refinancing, the estimated savings response, initial advertiser share and financial assumptions remain at their benchmark values in Table 6. The centroid row therefore changes lender geography only. The high-cost column assumes the advertiser charges $4,500 more per loan. Production income applies industry average costs, not marginal costs. The final row uses the full ongoing gap without subtracting the pre-campaign gap; it is an alternative calculation, not an identified lifetime return to advertising. Poisson percent gaps are not inserted as percentage-point coefficients.

Rebuilding the total-refinancing sample with polygon centroids gives 0.408 additional refinances per 10,000 eligible properties per cent of advertising (standard error 0.165). Saved mortgage simulations do not cover 16,271 of the newly required ZIP months. Among the centroid observations with simulations, the savings response is $4,946.69 per 10,000 eligible mortgages per cent. This is a restricted-sample check, not a full-centroid estimate of borrower value.

Table 11: Supplementary Lender Choice Comparisons

Group or comparison	Estimate	SE	N
Panel A. Lender type, 2019–2021			
Banks	0.166	0.063	33,973
Independent mortgage companies	0.029	0.023	70,438
Credit unions	0.099	0.039	48,805
Panel B. Borrower income, 2007–2021			
Income below $50,000	0.073	0.033	274,304
Income at least $50,000	0.088	0.030	281,515
Higher minus lower income	0.015	0.025	555,819

Notes: Entries are percentage points of lender share per additional cent of advertising. For each group g, the share regression is $s_{ic}^{g} = \alpha_c^g + \rho_c^{g,-} r_{ic}^{-} + \rho_c^{g,+} r_{ic}^{+} + \eta_g H_{ic} + \epsilon_{ic}^{g}$; the advertising regression replaces s^g with annual spending per resident. Entries divide the share gap by the spending gap. These are overall advertising comparisons, not ongoing-minus-before campaign estimates. Lender type uses the preceding year's HMDA classification. Income is in 2021 dollars; missing income is omitted. Income groups are estimated jointly and their difference retains the covariance. Only verified unchanged legal portfolios with available classifications enter these checks. Regressions use 2010 ZIP population weights and ZIP and lender–border clustered standard errors, accounting for both regressions.

C.1 Local television advertising expenditure

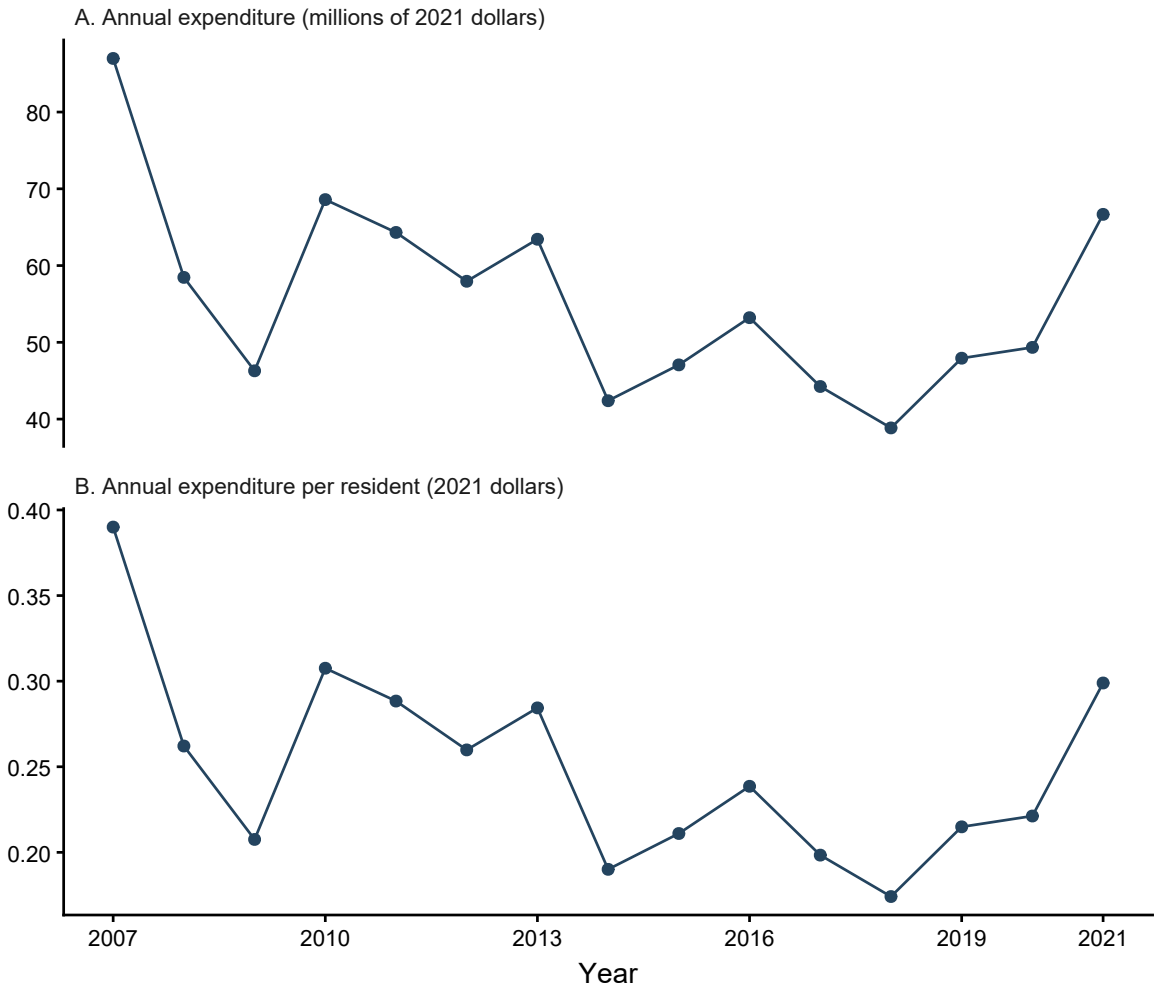


Figure 11: Local Mortgage Television Advertising Expenditure

Notes: Both panels use the same 76 markets with explicit finite source records in every month from 2007 through 2021. Recorded zeros are included; absent market-months are not assumed zero. Expenditure is in 2021 dollars and population is the 2010 DMA population. Expenditure combines the amount of advertising and airtime prices, not audience impressions or online advertising.

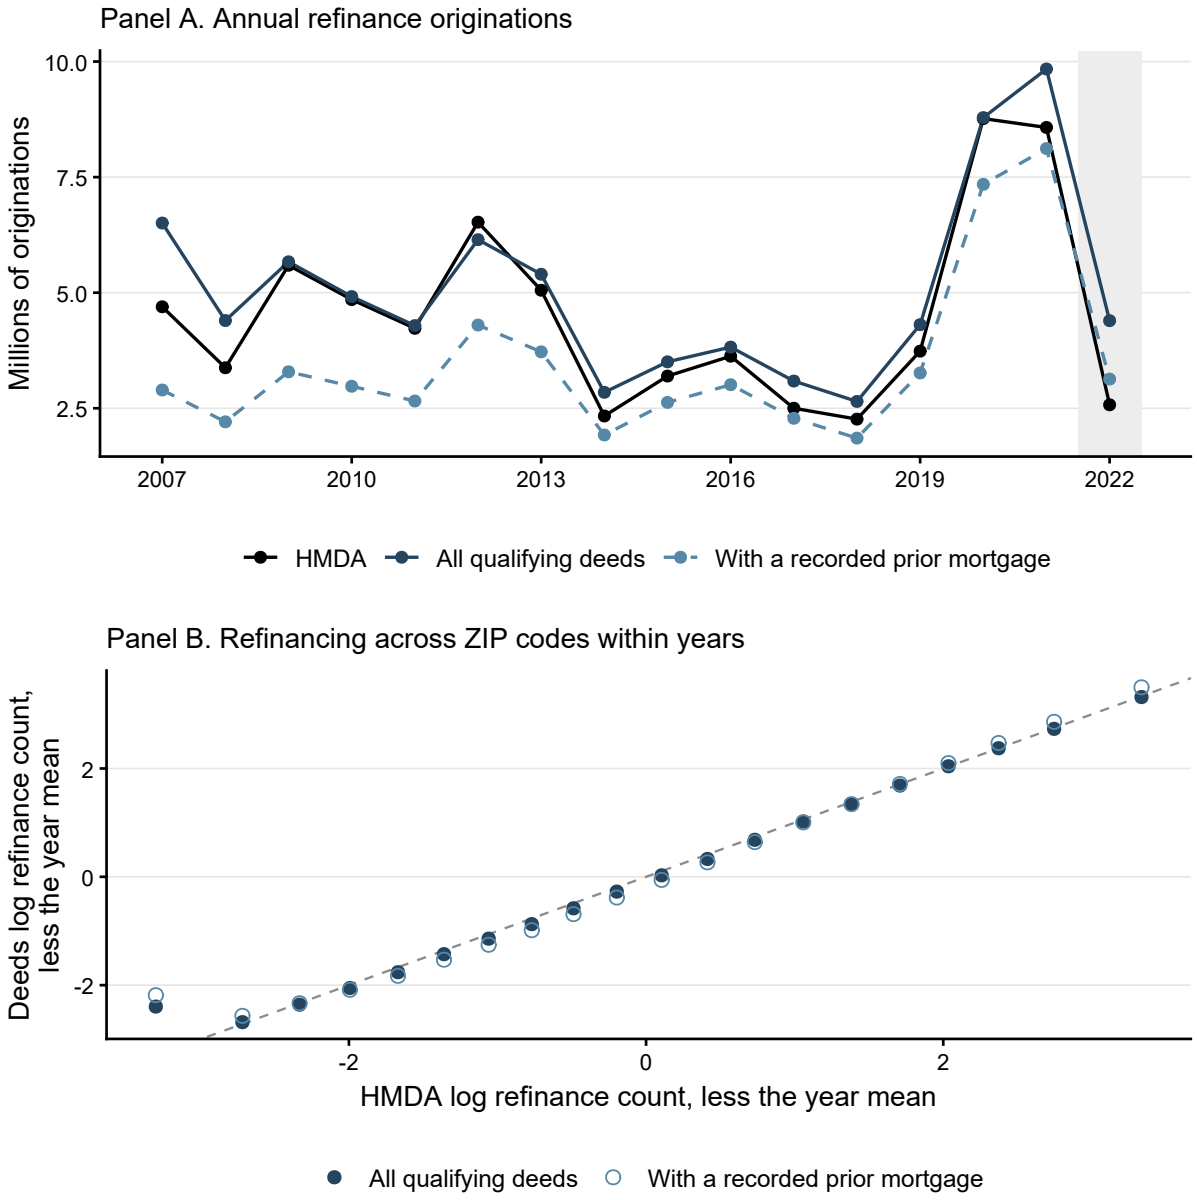


Figure 12: Validation of the Deeds-Based Refinancing Measure

Notes: Panel A compares national annual HMDA refinance originations with qualifying refinancing events in the deeds data, both for all events and for events with a preceding recorded mortgage. HMDA originations are allocated to 2010 ZCTAs. Qualifying deeds events are mortgages of at least $10,000 in nominal dollars and no more than $5 million in 2021 dollars, excluding loan-type codes CL, CM, FA, and RM and records within 30 days of a deed. We remove duplicates. Panel B compares counts across ZIP codes observed in both sources within the same year. It groups 451,066 ZIP–year observations into 20 bins after subtracting each year's mean log count. National annual correlations with HMDA are 0.964 for all qualifying deeds and 0.937 for events with a recorded prior mortgage; within-year log-count correlations are 0.858 and 0.864. The shaded 2022 observations fall after the advertising sample and are excluded from these correlations.

Table 12: Annual Advertising and Standardized Lender Rates

Year	Matched lenders	Share of local advertising spending linked to a rate	Correlation	95% interval
2018	255	76.6%	-0.064	[-0.185, 0.060]
2019	251	78.4%	0.094	[-0.031, 0.215]
2020	244	66.5%	-0.010	[-0.136, 0.116]
2021	283	70.5%	-0.127	[-0.240, -0.011]

Notes: The table reports the correlation between log annual local television spending and the standardized lender rate, in 2021 dollars. Each observation is a parent company with positive local spending and an estimated rate in the same year, and each receives equal weight. Monthly spending is summed before taking logs. Rates are adjusted for the characteristics in equation (10) and centered on the loan weighted mean among matched portfolios. Advertiser matches are year specific. Spending coverage uses all local advertisements in the same year. Intervals use the Fisher transformation, treating parent companies as independent and estimated rates as given.

D Additional Total Refinancing Results

Appendix Table 13 reports alternative samples, weights, clustering, and specifications for total refinancing. The centroid check rebuilds the sample using each ZIP polygon's area centroid instead of its Census internal point, including ZIPs that enter the distance band. Other checks change one choice at a time. The underlying estimates and additional combinations are retained in the replication files.

Table 13: Robustness of Total Refinancing Estimates

Specification	Estimate	SE	N
Main specification	0.419	0.162	1,279,367
Within five miles	0.405	0.161	494,332
Include ZIPs within one mile	0.411	0.161	1,313,327
Downweight repeated borders	0.431	0.174	1,279,367
Polygon centroid geography	0.408	0.165	1,277,659
ZIP population weights	0.394	0.152	1,278,542
Cluster by ZIP and month	0.419	0.139	1,279,367
Poisson, scaled to sample mean	0.198	0.114	1,259,108
Control for subsequent advertising	0.334	0.272	1,279,367
Monthly-event outcome, common sample			
Local projections	0.447	0.160	–
First-order time-series model	0.555	0.158	–

Notes: The table reports the three-month refinancing response under the alternatives listed in the first column. Entries are additional refinances per 10,000 eligible properties per additional cent of advertising per DMA resident, in 2021 dollars. Except as indicated, regressions use the controls, fixed effects and weights in column (1) of Table 3. Standard errors account for the refinancing and spending responses and are clustered by DMA and calendar month. Population weights omit ZIPs with zero or missing population. The Poisson model is $E[Y^3_{ibdt}|X] = \exp(\alpha_{bd} + \lambda_{bt} + \beta A_{dt} + \sum_{\ell=1}^{3}[\rho_\ell A_{d,t-\ell} + \phi_\ell \bar{y}^{(\ell)}_{bdt}])$. It retains mortgage weights and removes all-zero groups. The displayed level response multiplies its spending-normalized log response by its weighted outcome mean; the corresponding percent response is 0.121 percent. This local scaling is not nonlinear IV. The subsequent-advertising row adds $A_{d,t+1}, A_{d,t+2}, A_{d,t+3}$ as controls and reports the direct current-spending coefficient. The final rows count monthly refinancing events and may count repeat refinances, unlike the main unique-property outcome. They use identical observations and weights; their time-series equations are given in Appendix E.

E Advertising over Time

Figure 13 shows how advertising evolves around an increase in spending. For an alternative estimate of refinancing per cent spent, we fit a first-order model to advertising and monthly refinancing: $A_{dt} = \alpha_{db} + \mu_{bt} + \rho A_{d,t-1} + u_{dbt}$ and $y_{dbt} = \eta_{db} + \tau_{bt} + \phi y_{db,t-1} + b_0 A_{dt} + b_1 A_{d,t-1} + v_{dbt}$. The estimated persistence coefficients are $\rho = 0.599$ for advertising and $\phi = 0.526$ for refinancing. We use these equations to trace both responses following an initial increase in advertising and divide cumulative refinancing by cumulative spending. This assumes first-order dynamics and that advertising is unrelated to the current refinancing error conditional on the included variables.

Table 13 compares the three-month calculation with local projections using the same observations and initial month mortgage weights. The estimates are close. The table also reports a regression controlling directly for advertising in the following three months. That comparison holds subsequent spending fixed and requires it to be unrelated to unobserved refinancing shocks conditional on the other controls.

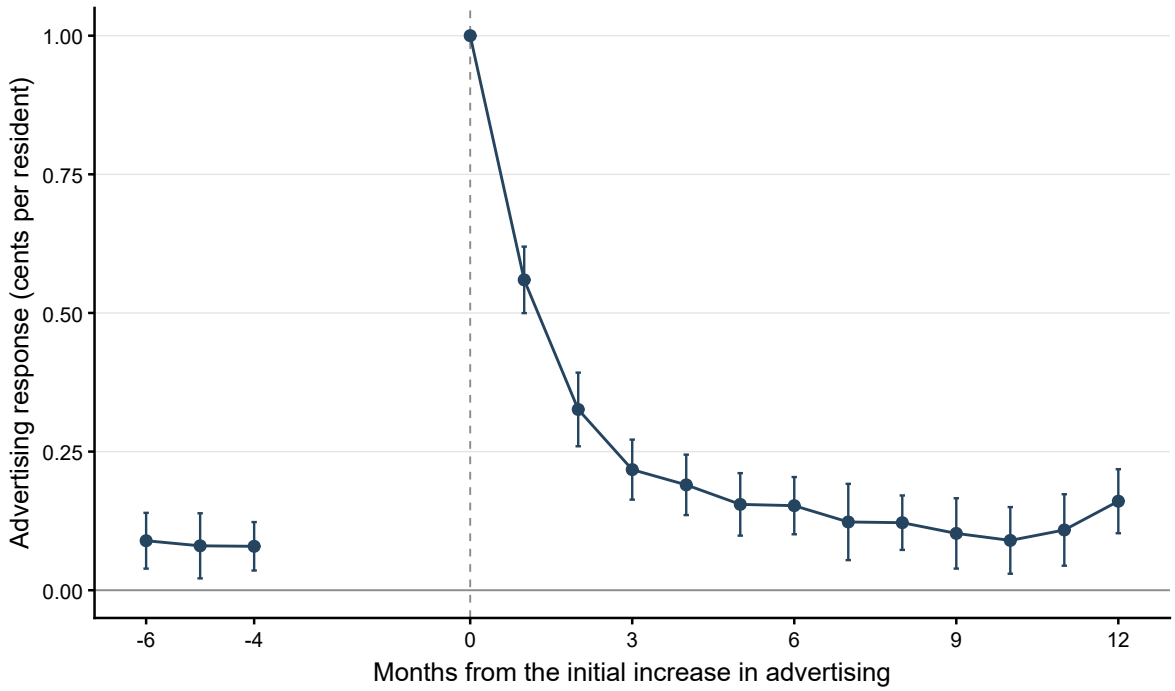


Figure 13: Advertising before and after an Increase in Spending

Notes: For each horizon h, the figure estimates

$$A_{d,t+h} = \alpha^h_{db} + \mu^h_{bt} + \beta_h A_{dt} + \sum_{\ell=1}^{3} (\rho^h_\ell A_{d,t-\ell} + \phi^h_\ell y_{db,t-\ell}) + u^h_{dbt}.$$

Here A is advertising in cents per DMA resident, y the monthly refinancing rate per 10,000 eligible properties, b a border, d its DMA side, and t a month. The terms α_{db} and μ_{bt} are border-side and border-month fixed effects. Points show $\hat{\beta}_h$ for a one-cent increase in the initial month; $\beta_0 = 1$ mechanically. The regression uses horizon-specific eligible-mortgage weights and complete advertising follow-up through month twelve, with exposure months from January 2007 through September 2021. Bars are pointwise 90 percent intervals, clustered by DMA and calendar month. Months -3 through -1 enter as controls and are omitted.

F Refinancing Responses by Advertising Spending

We replace current advertising in the three-month refinancing regression with indicators for five or ten positive-spending bins. Zero spending is the reference. The sample, lags, fixed effects, weights, and clustering are unchanged. We also estimate each bin's cumulative-spending response, so the refinancing response can be expressed per cent of advertising.

The refinancing response generally rises with spending. The point estimates per advertising cent are roughly flat, but noisy. Adding the bins to the original continuous-spending regression gives joint p-values of 0.066 with five bins and 0.336 with ten. These estimates do not rule out meaningful differences across spending levels. The lowest decile has a weak spending response and no bounded ratio interval.

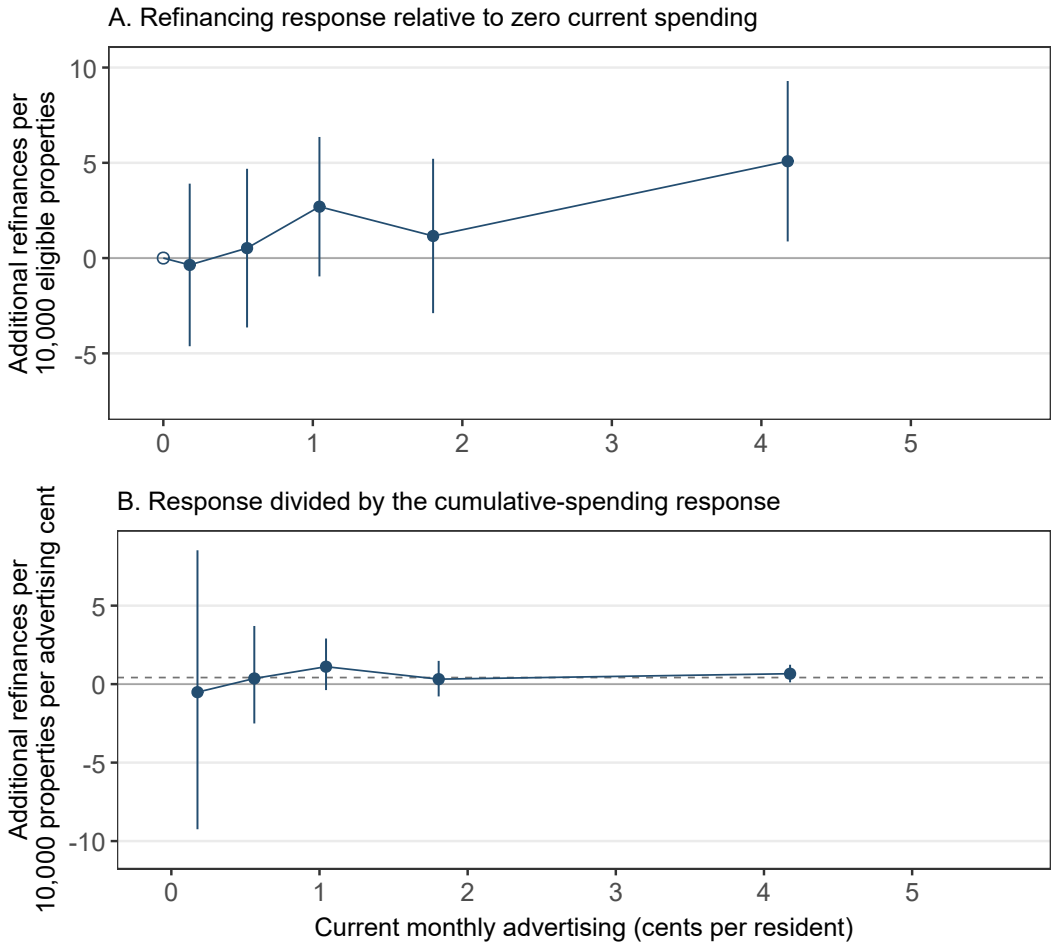


Figure 14: Refinancing Responses by Advertising Spending: Five Bins

Notes: We estimate

$$W^m_{ibdt} = \alpha^m_{bd} + \lambda^m_{bt} + \sum_{k=1}^{K} \beta^m_k \mathbf{1}\{A_{dt} \in I_k\} + \sum_{\ell=1}^{3} \left(\rho^m_\ell A_{d,t-\ell} + \phi^m_\ell \bar{y}^{(\ell)}_{bdt}\right) + \epsilon^m_{ibdt}, \qquad m \in \{R, S\}.$$

Here i is a ZIP, b a border, d its DMA side, and t a month. W^R counts eligible properties refinancing in months $t+1$ through $t+3$ per 10,000 eligible properties; $W^S = \sum_{h=0}^{3} A_{d,t+h}$. Monthly advertising A is in cents per resident (2021 dollars). The monthly refinancing control $\bar{y}^{(\ell)}_{bdt}$ pools counts and eligible stock at $t-\ell$ over the ZIPs included at t. α_{bd} and λ_{bt} are border–side and border–month fixed effects. The $K = 5$ positive-spending bins I_k use eligible-mortgage weights; zero spending is omitted. Panels A and B plot $\hat{\beta}^R_k$ and $\hat{\beta}^R_k/\hat{\beta}^S_k$ at each bin's weighted mean current spending. The ratio is an average contrast, not a marginal slope. Regressions use 1,279,367 stacked ZIP–border–month observations in 2007–2021, eligible-mortgage weights, and the one-mile donut. Bars are pointwise 90 percent intervals, clustered by DMA and month. Ratio intervals invert the joint test of $\beta^R_k - r\beta^S_k = 0$ for candidate ratios r. The dashed line is the linear estimate, 0.419. Zero has no interval in Panel A and an undefined ratio in Panel B.

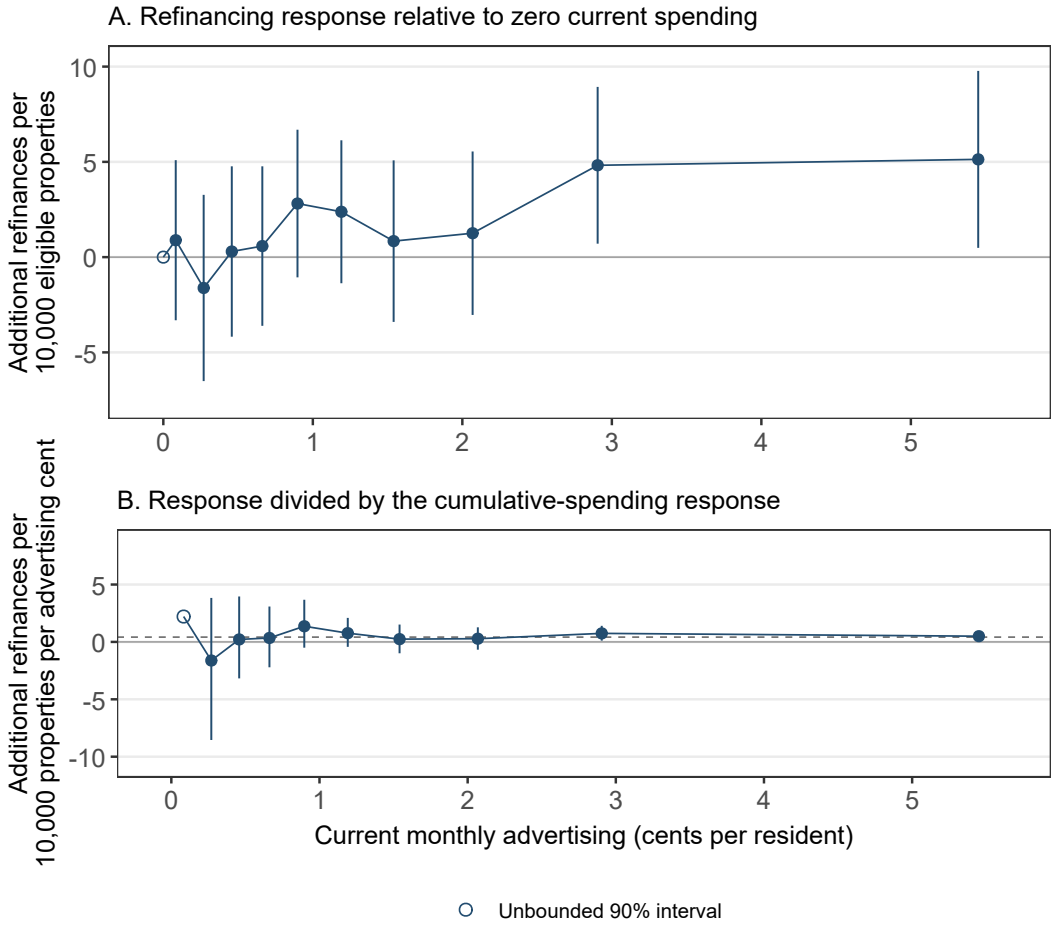


Figure 15: Refinancing Responses by Advertising Spending: Ten Bins

Notes: We estimate

$$W^m_{ibdt} = \alpha^m_{bd} + \lambda^m_{bt} + \sum_{k=1}^{K} \beta^m_k \mathbf{1}\{A_{dt} \in I_k\} + \sum_{\ell=1}^{3} \left(\rho^m_\ell A_{d,t-\ell} + \phi^m_\ell \bar{y}^{(\ell)}_{bdt}\right) + \epsilon^m_{ibdt}, \qquad m \in \{R, S\}.$$

Here i is a ZIP, b a border, d its DMA side, and t a month. W^R counts eligible properties refinancing in months $t+1$ through $t+3$ per 10,000 eligible properties; $W^S = \sum_{h=0}^{3} A_{d,t+h}$. Monthly advertising A is in cents per resident (2021 dollars). The monthly refinancing control $\bar{y}^{(\ell)}_{bdt}$ pools counts and eligible stock at $t-\ell$ over the ZIPs included at t. α_{bd} and λ_{bt} are border–side and border–month fixed effects. The $K = 10$ positive-spending bins I_k use eligible-mortgage weights; zero spending is omitted. Panels A and B plot $\hat{\beta}^R_k$ and $\hat{\beta}^R_k/\hat{\beta}^S_k$ at each bin's weighted mean current spending. The ratio is an average contrast, not a marginal slope. Regressions use 1,279,367 stacked ZIP–border–month observations in 2007–2021, eligible-mortgage weights, and the one-mile donut. Bars are pointwise 90 percent intervals, clustered by DMA and month. Ratio intervals invert the joint test of $\beta^R_k - r\beta^S_k = 0$ for candidate ratios r. The dashed line is the linear estimate, 0.419. Zero has no interval in Panel A and an undefined ratio in Panel B. The open point in Panel B has an unbounded 90 percent interval because its spending response is weak.