Refinancing Advertising, Inequality, and the Last Mile of Monetary Policy

Zhenling Jiang
University of Pennsylvania

Donggwan Kim
Boston College

Yesim Orhun*
University of Michigan

Abstract

Many mortgage borrowers fail to refinance even when doing so would yield substantial savings, representing a key inefficiency in household finance and a barrier to monetary policy transmission. We study whether advertising can reduce this friction. Using a comprehensive dataset that links television advertising to mortgage origination and performance data from HMDA, Freddie Mac, and Equifax, we study how exposure to advertising from lenders influences borrowers' refinancing decisions. We find that advertising increases mortgage borrowers' refinancing hazard. Almost all the causal impact of advertising on refinancing uptake is focused on cases where refinancing yields monetary savings. We further investigate how the advertising-induced increases in refinancing affect demographic disparities in refinancing propensities across income, gender, and race. Overall, our results suggest that advertising can be a relatively efficient policy lever to improve household financial decisions and address the "last-mile" problem of monetary policy.

Keywords: Mortgage refinancing, Television advertising, Monetary policy transmission, Income inequality, Household finance.

*zhenling@wharton.upenn.edu, donggwan.kim@bc.edu, aorhun@umich.edu.

1 Introduction

Mortgages represent the largest category of household debt in the United States. When interest rates fall, households with fixed-rate mortgages can lower their monthly payments and increase disposable income by refinancing. Mortgage refinancing is therefore a key channel through which monetary policy affects household balance sheets and consumption (e.g., Auclert, 2019; Beraja et al., 2019; Di Maggio et al., 2017; Wong, 2019). Yet many borrowers fail to refinance even when doing so would yield substantial savings, leaving thousands of dollars on the table (e.g., Campbell, 2006; Johnson et al., 2019; Keys et al., 2016). Moreover, these refinancing shortfalls are not evenly distributed: low-income households and racial minorities refinance at systematically lower rates, especially when refinancing incentives are strongest (Agarwal et al., 2024; Gerardi et al., 2023). Because lower-income households tend to have higher marginal propensities to consume, weaker refinancing pass-through among these households can further dampen the aggregate stimulus from rate cuts while widening economic disparities.

Therefore, from a policy perspective, understanding and ultimately reducing the frictions that prevent borrowers from refinancing is central to both strengthening monetary-policy pass-through and limiting the unequal wealth consequences of refinancing shortfalls. Prior work has suggested that refinancing failures often reflect demand-side frictions, including inattention, procrastination, or misperceptions about eligibility and potential gains (Agarwal et al., 2016; Andersen et al., 2020; Berger et al., 2025; Johnson et al., 2019; Keys et al., 2016; Maxted et al., 2025). These demand-side frictions are natural in this setting. Even when refinancing is financially attractive, it does not occur automatically: borrowers must recognize it and then actively begin, and persist through, a multi-stage application process. To the extent frictions arise from attention and follow-through, they are difficult to remedy through regulation alone. Prior work therefore evaluates interventions that seek to trigger borrower attention and engagement, such as mail reminders and direct offers (Byrne et al., 2023; Grundl and Kim, 2019; Johnson et al., 2019; Keys et al., 2016). These studies show mixed evidence and are largely confined to small-scale, targeted interventions. As a result, whether large-scale interventions can meaningfully reduce refinancing frictions remains an open question.

In this paper, we examine whether mass-market television advertising reduces refinancing frictions and improves the transmission of monetary policy to household balance sheets. Television advertising is an economically significant feature of the refinancing market, with mortgage lenders spending millions of dollars per year over the past decade and blanketing entire media markets, particularly during periods of lower interest rates. Television advertising operates precisely at the attention and engagement stages where the demand-side refinancing frictions arise and does so at a much larger scale compared to targeted interventions.

Even though lender advertising is pervasive and economically important, its effect on refinancing is not obvious a priori. In a competitive market, advertising may be highly effective from a lender's perspective (i.e., shifting borrower demand toward the advertiser), yet have little impact on aggregate refinancing if it mainly reallocates borrowers across lenders rather than inducing additional refinancing. More subtly, even when advertising does increase refinancing on average, it need not do so where it matters most for policy: its usefulness for monetary transmission and consumer welfare depends critically on which borrowers respond. If advertising disproportionately encourages borrowers with little to gain from interest-rate reductions, its relevance for monetary transmission may be limited. If responses differ systematically across demographic groups, advertising may have important distributional implications. Identifying which borrowers advertising primarily influences is therefore central to assessing its policy relevance: the extent to which advertising strengthens interest-rate pass-through depends on whether it increases refinancing among borrowers for whom failing to refinance is most costly.

A distinctive feature of the mortgage refinancing context, unlike most settings studied in the advertising literature, is that the economic value of refinancing is largely observable. Given a borrower's loan terms and prevailing interest rates, we can directly quantify the potential savings refinancing offers using established approaches from the prior literature. This enables us to test whether advertising responses are concentrated among consumers who stand to benefit the most, and to translate the causal impact of advertising exposure into dollar-valued gains. This is particularly important in the context of refinancing, where the stakes can be large and vary sharply across borrowers depending on their contracts and the interest-rate environment. Moreover, because we observe borrower demographics, we can examine how advertising interacts with existing disparities in refinancing behavior. Together, these features make mortgage refinancing a rare setting in which the efficiency and distributional consequences of mass-market advertising can be studied directly.

We construct a unique loan-level panel dataset to study the effect of TV advertising on borrowers' refinancing decisions, rich enough to allow for heterogeneity by potential financial gains and by demographic group. We link mortgage lender television advertising across 129 designated market areas (DMA) to comprehensive mortgage records for 30-year fixed-rate loans, combining data from

HMDA, Freddie Mac, and Equifax. The Equifax records enable us to distinguish refinancing from other prepayments events (e.g., moves) and to observe mortgage-related credit inquiries even when they do not result in a completed refinance. Together, these sources provide origination terms and borrower demographics, monthly updates to credit scores and payment histories, and measures of both refinancing activity and refinancing intent. The resulting panel comprises approximately 945,000 loans and 39 million loan-month observations from January 2016 through December 2021. We document that refinancing activity during this period mirrors patterns documented in prior work: refinancing responds to interest-rate savings potential, but substantially less so among low-income households and among Black or Hispanic borrowers.

Estimating the causal impact of advertising on refinancing choices is complicated by the fact that advertising is not random — lenders choose where and when to advertise in response to expected market demand. To address this concern, we follow the approach made popular by Shapiro et al. (2021), which leverages institutional features of the television ad-buying process. In particular, we absorb predictable variation in refinancing demand and advertising targeting through DMA, origination-cohort, and year-month fixed effects. Remaining variation in advertising intensity reflects both limits to advertisers' ability to forecast demand and institutional constraints in ad purchasing and slot availability, providing plausibly exogenous variation for identification.

We find that television advertising increases refinancing activity, but only among borrowers for whom refinancing is financially attractive. A one standard deviation increase in advertising exposure raises borrowers' monthly refinancing probability by 10-15 percent when potential interest-rate savings are substantial, while having no detectable effect for those with weaker incentives. This pattern is consistent with advertising alleviating attention and engagement frictions, rather than inducing refinancing with little economic benefit. The implied savings are large relative to the cost of advertising. A one–standard-deviation increase in advertising exposure costs approximately \$0.30 per household and generates expected savings of \$13–\$35 on average, and \$32–\$88 during low rate environments. These returns are comparable to those documented for highly targeted, smaller-scale interventions that successfully used reminders or offers to overcome behavioral frictions. Our estimates suggest that mass-market television advertising can be an effective and cost-efficient intervention for addressing demand-side frictions and improving interest-rate pass-through.

We next examine how advertising affects disparities in refinancing outcomes. We find that advertising has a neutral or modestly offsetting effect on racial disparities in refinancing. In contrast, advertising widens income disparities. Advertising exposure increases refinancing substantially more among higher-income households than among those in the lowest income quartile, conditional on similar refinancing incentives and observable eligibility criteria. Quantitatively, a one–standard-deviation increase in advertising exposure increases the existing refinancing gap between borrowers in the highest and lowest income quartiles by about 18 percent. We provide suggestive evidence that this gap does not arise from differences in advertising exposure or initial interest (proxied by credit inquiries), but instead reflects differences in the borrowers' likelihood of successfully completing the refinancing process. As a result, advertising ultimately amplifies income disparities in refinancing outcomes by interacting with downstream frictions in the refinancing process.

Our paper sits at the intersection of three literatures: refinancing frictions and monetary policy transmission, the impact of advertising on consumer choices, and demographic disparities in household finance. For the refinancing frictions and monetary policy transmission literatures, our results highlight the role of mass-market advertising in facilitating monetary pass-through. By concentrating its effects where potential gains are largest, television advertising improves the transmission of rate cuts to household balance sheets and mitigates the last-mile frictions that often limit the effectiveness of monetary stimulus. These findings point to communication-based interventions as a natural complement to traditional monetary tools. More broadly, such communication need not be confined to private lenders. Johnson et al. (2019) highlight the lack of trust in lending institutions as a barrier to refinancing, implying that the effects documented in this paper may represent lower bounds on the impact of potential communications that come from government agencies or not-for-profit entities.

From the perspective of the advertising effectiveness literature, a distinctive feature of our setting is that we can observe individual-level heterogeneity in the potential savings from taking the advertised action. In most consumer goods markets, consumer surplus depends on unobserved preferences, making it difficult to assess whether advertising disproportionately influences those who stand to gain the most. In the mortgage refinancing context, by contrast, potential savings are largely determined by observable loan terms and prevailing interest rates. This allows us to test directly whether advertising-induced responses are concentrated among borrowers with large potential interest savings and to quantify how mass-market advertising shapes the distribution of refinancing activity and the transmission of interest-rate cuts to households.

Finally, for the literature on demographic disparities in household finance, our distributional results suggest that mass-market informational interventions can widen income disparities in refinancing if other constraints are not addressed. For example, if other downstream frictions (e.g., time or liquidity constraints) disproportionately hinder low-income households, then equal increases in awareness or interest in refinancing need not translate into equal refinancing outcomes. This suggests that policies aimed at facilitating refinancing should pair communication strategies with efforts to alleviate downstream bottlenecks, particularly during periods of high refinancing demand.

2 Institutional Background and Prior Literature

Understanding why borrowers often fail to refinance, despite substantial potential savings, is informed by the institutional features of the refinancing process. In this section, we first describe the refinancing decision and the multi-stage process through which borrowers pursue refinancing opportunities, highlighting the role of regulation in shaping prices and information. We then review prior evidence on failures to refinance, the behavioral and supply-side frictions that contribute to these failures, and persistent disparities across demographic groups. Finally, we discuss the burgeoning empirical literature on interventions designed to prompt refinancing, which motivates our focus on mass-market advertising and heterogeneity in its effects across borrowers with different financial incentives and household characteristics.

2.1 Refinancing Decision and Process

Most borrowers refinance to obtain a lower mortgage rate, and even modest rate reductions can generate substantial savings, particularly for borrowers with larger balances or longer remaining terms.[1] By lowering rates and increasing disposable income, refinancing is an important channel through which monetary policy affects household consumption and the broader economy (Agarwal et al., 2023; Beraja et al., 2019; Berger et al., 2021; Di Maggio et al., 2017, 2020; Eichenbaum et al., 2022). The benefits of refinancing depend on the interest rate differential, remaining loan balance, and expected tenure, while costs include explicit refinancing fees (typically \$2,000 to \$5,000 depending on loan size (Keys et al., 2016; Zhang, 2023), as well as the option value of waiting for potentially lower rates. Prior literature introduces normative benchmarks for identifying refinancing incentives that take all these factors into account (Agarwal et al., 2013; Deng et al., 2000).

Observed refinancing behavior, however, departs sharply from these normative benchmarks (Agarwal et al., 2016; Andersen et al., 2020). Understanding failures to refinance and potential

[1] Other reasons include changing the term structure of the mortgage, converting from an adjustable-rate to a fixed-rate mortgage, extracting home equity, or changing the identity of the borrower.

interventions requires recognizing the potential frictions that may exist in the multi-stage process of refinancing. Even when refinancing is financially attractive, borrowers must first recognize the opportunity and initiate contact with a lender. Borrowers then progress through subsequent stages of the refinancing funnel, including authorizing a credit check and submitting a formal application, undergoing underwriting review with associated documentation requirements, and ultimately closing. At each stage, borrowers may exit voluntarily or involuntarily due to time costs, inattention, or failure to meet underwriting criteria. As a result, potential savings need not ignite engagement, and early-stage engagement need not translate into completed refinancing.

As borrowers progress through the refinancing funnel, the terms they face are shaped by institutional features of the mortgage market. For conforming mortgages, which we focus on in this study, pricing is largely rule-based, determined by market conditions and observable borrower risk characteristics through the Loan-Level Price Adjustment (LLPA) schedule set by the Federal Housing Finance Agency.[2] While lenders retain limited discretion, discrimination based on protected characteristics is prohibited under federal law, notably the Equal Credit Opportunity Act. In addition, post-crisis disclosure reforms under the TILA–RESPA Integrated Disclosure (TRID) rules standardize the presentation of key loan terms, facilitating comparison across lenders. Together, these institutional features imply that prices and loan information are relatively transparent and standardized. These regulations do not, however, address frictions that prevent borrowers from recognizing refinancing opportunities or ease the procedural burdens of progressing through the refinancing funnel.

2.2 Failure to Refinance

A large empirical literature documents that many borrowers fail to refinance despite facing refinancing incentives that appear large relative to reasonable estimates of refinancing costs and having no obvious eligibility constraints (Agarwal et al., 2016, 2024; Andersen et al., 2020; Johnson et al., 2019; Keys et al., 2016). These failures are economically meaningful. For example, Keys et al. (2016) find that approximately 20% of U.S. homeowners who would benefit from refinancing fail to do so, leaving a median of $11,500 in net present value and $160 per month in foregone savings. Similarly, Johnson et al. (2019) find that roughly half of in-the-money borrowers at a large financial institution do not refinance, leaving an average of $8,700 in expected savings unclaimed. Andersen

[2] Base rates are linked to yields on the 10-year Treasury note, with guarantee fees added according to the LLPA schedule.

et al. (2020) document a similar widespread refinancing inertia in Denmark.

Refinancing inertia creates a "last-mile problem" for monetary policy. Reductions in interest rates do not automatically translate into lower household borrowing costs. Sluggish mortgage refinancing limits the extent to which monetary easing reaches household balance sheets, dampening consumption responses and weakening aggregate stimulus (Calza et al., 2013). Beraja et al. (2019) show that regions with more active refinancing exhibit substantially larger consumption responses to interest rate cuts. Consequently, refinancing failures not only affect individual borrowers' finances but, when pervasive, also constitute a first-order impediment to the pass-through of monetary policy.

Refinancing failures reflect frictions operating in multiple stages of the refinancing funnel. The literature identifies important demand-side frictions. At the upper funnel, borrowers may fail to initiate refinancing due to inattention (Andersen et al., 2020; Byrne et al., 2023), limited financial sophistication (Agarwal et al., 2016; Andersen et al., 2020), procrastination (Maxted et al., 2025), and distrust of financial institutions (Johnson et al., 2019). Documentation requirements and liquidity constraints can create additional barriers, either by impeding completion later in the funnel or by deterring borrowers from applying in anticipation of these barriers (e.g., DeFusco and Mondragon, 2020). Supply-side frictions may also matter, especially during refinancing booms. Sharp declines in interest rates trigger large refinancing waves that can overwhelm lender capacity, potentially resulting in rationed attention, prioritization of more profitable applications, and longer processing times (Agarwal et al., 2024; Beraja et al., 2019; Fuster et al., 2024).

Failures to refinance are not evenly distributed across borrowers. Low-income and minority households refinance at substantially lower rates, even after controlling for refinancing incentives and loan and credit characteristics (Keys et al., 2016). Agarwal et al. (2024) show that refinancing inequality across income groups widened during the COVID-19 pandemic, with disadvantaged borrowers less likely to take advantage of historically low rates. Prior work has examined several mechanisms that may contribute to disparities across borrowers, including differences in financial sophistication (Agarwal et al., 2016; Andersen et al., 2020), access to information (Maturana and Nickerson, 2019; McCartney and Shah, 2022), search and access frictions (Kim, 2024), exposure to labor-market shocks during downturns (DeFusco and Mondragon, 2020), and frictions at the application stage (Frame et al., 2025; Li and Zhang, 2024).

The sluggishness of refinancing and disparities in that sluggishness naturally raise the question of whether policy or market-based interventions can meaningfully increase refinancing, particularly for borrowers who are least likely to refinance. A small literature examines several targeted outreach campaigns designed to prompt borrower attention and engagement, with mixed results. Byrne et al. (2023) show that low-cost information and reminder letters sent to approximately 12,000 variable-rate mortgage holders at a large Irish retail bank substantially increase refinancing take-up. Similarly, Grundl and Kim (2019) find that mail-in offers raise refinancing demand among about 9,000 U.S. borrowers. However, other studies find limited effects. Keys et al. (2016) report that mailed offers sent to approximately 200 low-income Chicago borrowers had little impact, with many recipients failing to open the letters or follow through. In three field experiments involving pre-approved Home Affordable Refinance Program (HARP) offers to 80,000–110,000 borrowers per wave, Johnson et al. (2019) find no meaningful increase in application rates and highlight the role of trust-related frictions.[3] Taken together, these studies provide mixed evidence on the effectiveness of mail-based outreach campaigns. It is unclear whether the mixed evidence reflects differences in refinancing incentives, differences in borrower populations, or both.

We contribute to this literature by studying the effectiveness of mass-market TV advertising as a potential alternative tool to encourage refinancing, focusing on its ability to strengthen monetary pass-through as well as its distributional consequences. We assess pass-through relevance by testing whether ad effectiveness scales with borrowers' savings potential, and we examine distributional consequences by testing whether ad responsiveness differs across demographic groups, conditional on savings potential. Distinguishing incentive-based from demographic heterogeneity has been difficult in prior work, likely due to the smaller-scale and targeted nature of the prior interventions.

2.3 TV Advertising

Mortgage lenders devote substantial resources to direct-to-consumer advertising, particularly during periods of declining interest rates. In particular, the majority of the spending is devoted to mass-market television campaigns. TV ads typically emphasize the prevailing low-rate environment, highlight potential monthly savings, and stress the convenience of the application process, often featuring explicit calls to action inviting viewers to check their eligibility. In this sense, TV ads can potentially mitigate frictions at the top of the refinancing decision funnel, as they aim to shape borrower awareness and engagement.

From a policy-perspective, whether television advertising has desirable effects on refinancing

[3] The Home Affordable Refinance Program (HARP) was a federal program introduced during the financial crisis that allowed borrowers with government-backed mortgages to refinance despite high LTV ratios.

is ex ante ambiguous for five reasons. First, it is not a foregone conclusion that advertising increases aggregate refinancing. While reminder and attention effects are well documented in other contexts (e.g., He and Klein, 2023), and Hu et al. (2024) show that the entry of financial news programming increases refinancing activity, lender advertising in a competitive market may primarily reallocate borrowers across lenders rather than expand the total refinancing volume. Second, even if advertising raises refinancing on average, it may be a blunt instrument: television advertising is untargeted and reaches many borrowers with weak refinancing incentives, so its effectiveness could be substantially lower than that of targeted mail-based reminders and offers. Third, for monetary policy transmission, it is important to assess whether it increases refinancing among borrowers who stand to save. Advertising could instead induce refinancing among borrowers with limited potential savings (e.g., inducing cash-out refinancing), contributing little to interest-rate pass-through. Fourth, advertising that benefits lenders need not benefit borrowers. For example, Célérier and Tak (2025) show that under weak regulation, bank advertising steered minority borrowers toward harmful products, and Gurun et al. (2016) document that origination advertising during the pre-crisis period was associated with more expensive mortgages. Whether similar concerns extend to refinancing advertising's impact on who refinances remains an open empirical question. Fifth, because advertising's impact on refinancing may differ across demographic groups, mass-market advertising could either mitigate or exacerbate existing disparities in who benefits from interest-rate declines. In this sense, studying lender advertising advances the household-finance literature by connecting a major market institution, i.e., mass-market advertising, to refinancing behavior, monetary-policy pass-through, and the distribution of refinancing gains.

Our paper also contributes to the broader literature on advertising effectiveness. A large body of work studies how advertising influences consumer behavior, documenting effects on demand, market shares, and consumer attention (e.g., Hristakeva and Mortimer, 2021; Shapiro, 2018; Shapiro et al., 2021). In banking, Honka et al. (2017) find that advertising primarily increases consumer awareness rather than shifting preferences, with the awareness elasticity roughly five times larger than the choice elasticity, consistent with an informative role for advertising. In the mortgage refinancing context, Kim et al. (2025) show that the racial identity of actors featured in television ads affects lender choice, suggesting that ad content, not just volume, shapes borrower responses.

We complement this work by examining advertising effects in a setting where the economic stakes of responding are unusually transparent. Although there is growing interest in the societal consequences of advertising (e.g., Ghosh Dastidar et al., 2023), it is typically difficult to assess whether advertising primarily influences consumers who stand to gain the most from taking the advertised action. In most consumer goods markets, consumer surplus depends on unobserved preferences, making it difficult to determine whether advertising is more effective among those who would benefit the most. In the mortgage refinancing context, by contrast, potential gains are largely determined by observable loan terms, interest rates, and borrower characteristics.[4] This allows us to test directly whether advertising effects are concentrated among borrowers with large savings potential from refinancing, which is central to interpreting both the distributional effects and policy relevance of advertising.

3 Data and Refinancing Incentives

3.1 Data Sources

Our analysis draws on a novel, large-scale dataset that integrates four complementary data sources: Home Mortgage Disclosure Act (HMDA) records, Freddie Mac Single-Family Loan-Level data, proprietary consumer credit data from Equifax, Inc., and Nielsen AdIntel TV advertising data.

No single data source contains all the information needed to study refinancing incentives, advertising exposure, and refinancing decisions (including search behavior) within the same set of loans over time. We therefore construct a loan-level monthly panel through a series of merges described below. The Online Appendix documents each merge step in detail, including match algorithms, match rates by year, and representativeness checks.

3.1.1 Loan-level data

The Freddie Mac data cover conforming, fixed-rate mortgage loans and provide the backbone of our panel: origination characteristics (loan amount, mortgage rate, property location), monthly performance outcomes (prepayment, delinquency, modification), and time-varying eligibility metrics (remaining balance, delinquency status).[5] The HMDA data supply borrower demographics (e.g., race/ethnicity, income) that are essential for studying disparities. We match HMDA originations to Freddie Mac loans for the years 2010–2017 using a waterfall procedure based on overlapping loan characteristics at origination (loan amount, origination year, loan purpose, property location, and,

[4] We recognize that some borrowers refinance for reasons beyond interest-rate savings (e.g., cash-out refinancing or changes in loan terms); therefore, we abstain from making strong welfare statements.

[5] We define a loan as delinquent if it is 90 days or more past due, consistent with earlier work.

when available, lender identity). Close to 40% of loans are matched one-to-one, with match rates improving from approximately 30% in 2010–2012 to over 45% in 2016–2017 (Online Appendix, Table A1).

A key limitation of the merged HMDA–Freddie Mac data is that they do not distinguish whether a prepayment reflects refinancing or a home sale. To classify prepayment reasons, we match the HMDA–Freddie Mac sample to Equifax consumer credit records using characteristics observed both at origination and over the loan's performance history (loan amount, origination timing, property location, and prepayment timing). We classify a prepayment as refinancing if the borrower (identified via encrypted SSN) originates a new mortgage loan within a short window around the closing date of the original loan at the same property location.[6] The Equifax match rate ranges from 70% to 82% across yearly cohorts (Online Appendix, Table A2). The Equifax data also provide updated credit scores (VantageScore), which are critical for constructing refinancing incentives conditional on time-varying credit risk. In addition, we obtain mortgage-related hard credit inquiries that require consumer consent. We use these credit inquiries to proxy for borrower interest in refinancing even in the absence of a completed refinancing loan. Additional details on the two sequential matching procedures, including the waterfall algorithms and representativeness checks, are provided in Online Appendix, Section D.

Finally, we merge the Zillow Home Value Index data at the county-month level, which allows us to track home price appreciation since origination and impute updated loan-to-value (LTV) ratios over the life of each loan.[7]

3.1.2 Advertising data

We obtain television advertising data from Nielsen AdIntel, which records every national and local TV ad airing across U.S. designated market areas (DMAs). For each ad airing, AdIntel reports the advertiser, ad creative, program, media channel, airing time, DMA, and the Gross Rating Points (GRPs) delivered. GRPs measure the volume of advertising exposure delivered to households in a DMA, defined as the product of *reach* (the share of households exposed) and *frequency* (the average number of exposures among those reached). For example, reaching 25 percent of households four

[6] To account for reporting delays, we allow up to a three-month window between the closing of the original loan and the origination of the new loan. Prepayments not classified as refinancing are treated as moves; we retain these observations up to six months before the move date.

[7] The current home value is estimated by multiplying the reported original home price by the change in the Zillow Home Value Index at the county level since the year and month of origination. The updated LTV is the outstanding loan balance divided by this estimated home value.

times corresponds to 100 GRPs. Unlike advertising spending, which reflects what advertisers pay for airtime, GRPs measure what audiences actually see, making them the standard metric for quantifying advertising exposure. Because the mapping from dollars to impressions varies across markets and over time due to differences in slot availability, programming, audience size, and competition from other advertisers, GRPs provide a more accurate and comparable measure of advertising intensity than expenditures.

We restrict the sample to mortgage ads, excluding reverse mortgage lenders, VA-specialized lenders, and mortgage modification and relief service providers firms based on advertiser and product names.[8] We aggregate mortgage-lender ad airings (measured in GRPs) to the DMA×month level, yielding a monthly panel of advertising intensity across markets. Because advertising effects reflect cumulative recent exposure rather than contemporaneous exposure alone, we follow the standard approach in the advertising effectiveness literature (Shapiro et al., 2021) and construct an advertising stock measure. Specifically, advertising stock is defined as an exponentially weighted moving average of past GRPs:

$$AdStock_{td(i)} = \sum_{\tau=t-L}^{\tau=t} \delta^{t-\tau} \mathrm{GRP}_{\tau d(i)},$$

where $\mathrm{GRP}_{\tau d(i)}$ is the total GRP in month τ and DMA $d(i)$, $\delta \in (0,1)$ is the decay parameter governing how quickly past exposure depreciates, and L is the look-back window. Consistent with prior work on durable advertising effects, we set weekly $\delta = 0.90$ and use a 12-month look-back window.[9] This structure is well suited to refinancing, a low-frequency deliberative decision for which television advertising builds salience and awareness over time.

3.2 Descriptive Statistics

Our analysis sample consists of 944,485 loans and 39,104,968 loan-month observations spanning January 2016 through December 2021, covering the top 129 DMAs for which AdIntel measures all TV advertising airings using monitoring devices. We restrict attention to conventional, first-lien, 30-year fixed-rate mortgages on owner-occupied, single-family properties originated between 2010

[8] Online Appendix, Section E details our lender inclusion and exclusion criteria, the construction of DMA-level GRPs from national and local ads, and the distribution of ad exposure across racial and income groups (which is broadly similar across groups for cable television ads).

[9] Our conclusions are robust to using contemporaneous GRPs instead of the stock measure. The residual variation in ad stock after absorbing DMA and year-month fixed effects remains substantial (coefficient of variation = 0.136; see Online Appendix, Section E, Figure B3).

and 2017, excluding loans with exotic features (e.g., negative amortization, prepayment penalties). We further require complete demographic information (race/ethnicity, gender, income), a loan balance above $60,000, and no foreclosure during the sample period. The AdIntel merge restricts the sample to the top 129 DMAs where advertising exposure is reliably measured. Tables A3 and A4 in the Online Appendix document attrition at each restriction step.[10] We track each loan from January 2016 (or from origination, if later) until prepayment or the end of the sample period. To ensure that refinancing propensity is not contaminated by an impending home sale, loans identified as prepaid due to a sale are retained only up to six months before the inferred move date (based on the closing date of the existing loan).[11]

Table 1 displays summary statistics for the key variables used in our analysis. The top panel reports loan and borrower characteristics at origination (one observation per loan); the bottom panel reports time-varying characteristics at the loan-month level. The sample of borrowers is 84% White, 4% Black, 7% Hispanic, the rest being mostly Asian. At origination, 52% of the loans are jointly held, and 44% of single-borrower loans are held by female borrowers.

The average interest rate at origination in our sample is 4.16%, with the prevailing rate being on average 0.35% lower than the borrower's origination rate. The average FICO score at origination, as well as the current credit score, is 747, suggesting no general trend overall. The average loan size is approximately $242K at origination, with $215k balance and 311 months remaining on average at any given time in the data period across households. Congruently, LTV at origination is 79.6% but its current value, at any given time in the sample period, is 58.5% on average. We also track the months since the most recent delinquency or modification, as these events could affect refinancing eligibility, especially when they occur within the past 12 months.

Our key outcome variable is the refinancing indicator, with an average monthly hazard of 0.91% over the sample period. We also track mortgage-related credit inquiries, with an average monthly inquiry probability of roughly 2.5%. Our key treatment variable is the advertising GRP stock, which has a sample mean of 3,501 and a standard deviation of 1,571.

[10] To focus on borrowers who are plausibly eligible to refinance, so that failure to refinance reflects behavioral frictions rather than credit constraints, we apply dynamic eligibility screens: we exclude loan-months in which the borrower's credit score falls below 620, the LTV ratio exceeds 80%, or the loan is currently delinquent or modified. Loans less than 12 months old are also excluded.

[11] Because our sample is limited to loans securitized by Freddie Mac, it excludes loans held by Fannie Mae, government-insured (FHA/VA) loans, and non-conforming (jumbo or non-QM) mortgages. FHA and VA borrowers tend to have lower incomes and are disproportionately Black and Hispanic, so our sample likely understates the true extent of demographic disparities in refinancing. Our estimates of advertising effects should therefore be interpreted as applying to the conforming conventional mortgage market.

Table 1: Descriptive statistics for the estimation sample

Variable	Mean	Std. Dev.	Percentiles 25th	Median	75th
Loan and borrower characteristics at origination (loan level, N= 944,485)					
White (d)	0.84	0.37	1	1	1
Black (d)	0.04	0.18	0	0	1
Hispanic (d)	0.07	0.25	0	0	1
Asian (d)	0.06	0.24	0	0	1
Single account, male (d)	0.27	0.45	0	0	1
Single account, female (d)	0.21	0.41	0	0	0
Joint account (d)	0.52	0.50	0	1	1
Income ($1,000)	98.0	68.2	58.0	85.0	121.0
Interest rate at origination (%)	4.16	0.43	3.88	4.13	4.49
Credit score (FICO)	747.1	47.3	716	756	785
Loan size ($1,000)	241.8	113.8	152	224	315
Loan-to-value (LTV)	79.6	18.0	72	80	90
Refinance at origination (d)	0.48	0.50	0	0	1
Time-varying factors impacting returns to refinancing (loan-month, N= 39M)					
Outstanding - prevailing rate	0.35	0.72	-0.17	0.33	0.86
Loan balance ($1,000)	215.1	104.6	132.6	196.6	280.1
Months remaining	311.3	27.0	294	315	331
Time-varying characteristics impacting refinance eligibility (loan-month, N= 39M)					
Credit score (Vantage, current)	747.1	66.3	712	763	800
Loan-to-value (LTV, current)	58.5	17.2	47.1	59.7	70.5
Months since delinquency†	12.0	13.7	1	7	18
Months since modification†	21.4	16.9	8	18	31
Refinancing and mortgage-related credit inquiries (loan-month, N= 39M)					
Refinancing indicator (×100)	0.91	9.50	0	0	0
Inquiry indicator (×100)	2.52	15.67	0	0	0
TV mortgage ad exposure, measured in GRPs (loan-month; N = 39M)					
Ad GRP flow	1,890	973	1,242	1,624	2,223
Ad GRP stock	3,501	1,571	2,452	3,124	3,993

Notes: † Conditional on having experienced a delinquency or modification.

The refinancing hazard and ad exposure intensity are correlated over time, presumably driven by the interest rate environment. Figure 1 displays the monthly refinancing hazard alongside the average 30-year fixed mortgage rate and monthly GRPs (in thousands) over the sample period. From 2016 to late 2018, mortgage rates hovered around 4–5%, the refinancing hazard remained below 1%, and advertising averaged roughly 1,500–2,500 GRPs per month. Beginning in mid-2019 and accelerating through the COVID-19 period in 2020–2021, the mortgage rates fell sharply, dropping below 3%. This decline coincided with a surge in refinancing activity (peaking above 3%) and in advertising intensity, with GRPs exceeding 4,000/month (roughly 40 ad exposures per capita per month). Notably, although advertising and refinancing both respond to the interest rate environment, they do not move in lockstep: advertising intensity spikes with different timing and magnitude across rate episodes.

Figure 1: Refinancing Hazard, Mortgage Rate, and Advertising GRPs over Time

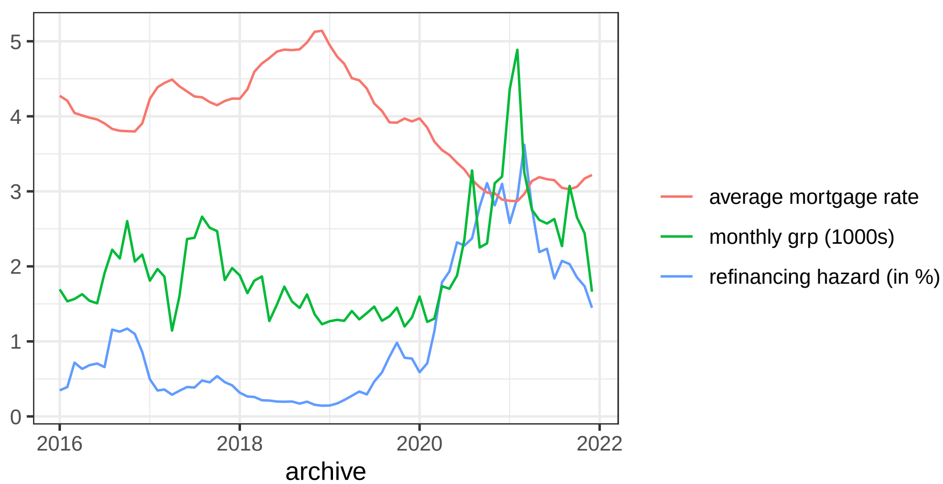

The left axis indicates the monthly refinancing hazard (in %), the prevailing average 30-year fixed mortgage rate (in %), and the total mortgage advertising GRPs (in thousands), January 2016–December 2021. GRPs measure gross rating points delivered to households across 129 DMAs.

Group disparities in refinancing are readily apparent in raw data. Figure 2 plots the monthly refinancing hazard by race/ethnicity (top panel) and income quartile (bottom panel), with the 30-year fixed rate shown on the right axis for reference. During the low-rate period of 2020–2021, refinancing activity increased across all groups, but the magnitude of the responses varied substantially. At the peak, White and Asian borrowers' monthly hazards exceeded 3%, while Black borrowers' hazard reached only about 2%. Disparities by income are even more pronounced: borrowers in the top income quartile (Q4) peaked near 5%, while those in the bottom income quartile (Q1) barely reached 2%. Across both panels, the gaps across different groups widen sharply when interest rates fall and refinancing becomes most attractive, and then compress as rates rise. These disparities motivate our heterogeneous analyses, which examine responses to advertising across demographic groups. In doing so, we assess whether advertising amplifies or mitigates existing disparities after accounting for differences in refinancing incentives and other covariates.

Figure 2: Monthly Refinancing Hazard by Race/Ethnicity and Income Quartile

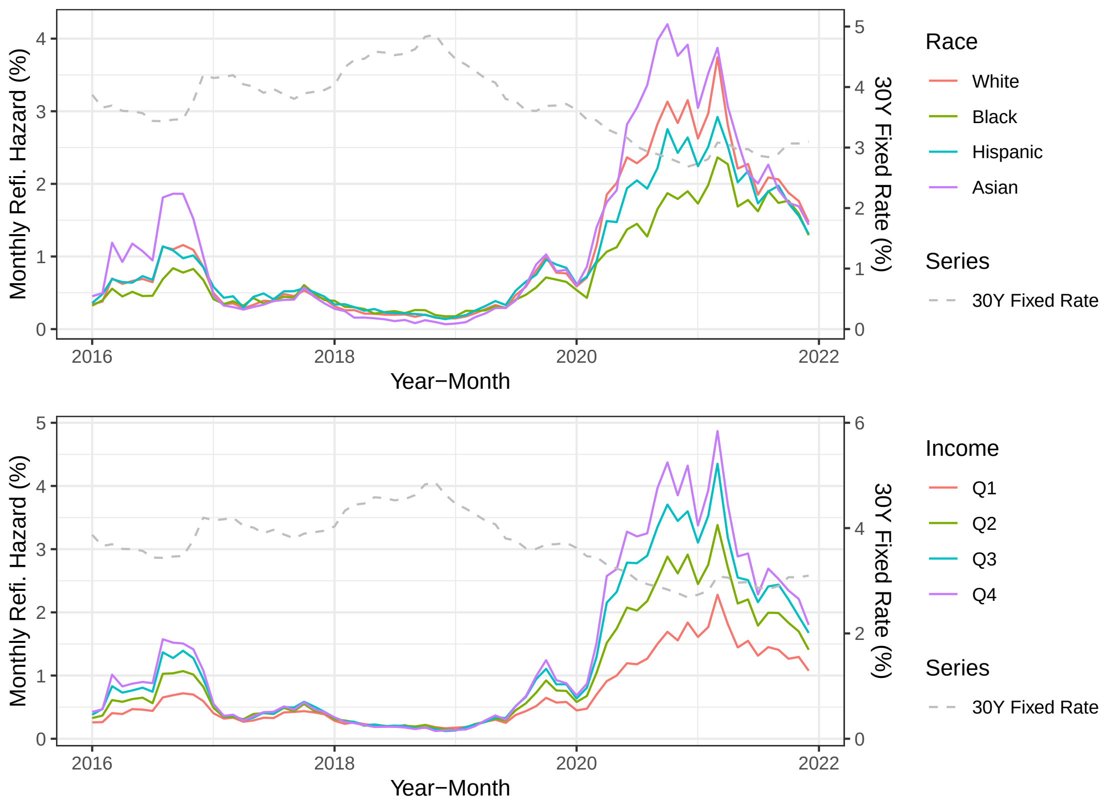

Top panel: monthly refinancing hazard by race/ethnicity. Bottom panel: monthly refinancing hazard by household income quartile at origination. The 30-year fixed mortgage rate is plotted on the right axis for reference. Unconditional on incentives or borrower characteristics.

3.3 Refinancing Incentives

To study refinancing behavior, we measure the economic incentive to refinance for each borrower over time. Standard theory frames mortgage refinancing as the exercise of a prepayment option: the borrower should refinance when the present value of interest savings exceeds transaction costs plus the option value of waiting (Agarwal et al., 2013; Deng et al., 2000). Empirically, the literature has operationalized refinancing incentives using three related but distinct metrics, each with different strengths. We construct all three measures for each borrower-month in our panel, which allows us to assess robustness across alternative definitions of refinancing value and permits comparability with prior work.

The simplest measure is the *rate gap*, defined as the difference between the borrower's current contract rate and the rate the borrower is expected to receive upon refinancing. While commonly used (e.g., Gerardi et al., 2023; Keys et al., 2016), the rate gap abstracts from important loan features, such as loan size and remaining term, that impact potential gains from refinancing, as well as refinancing fees. Option-value measures, such as the *ADL threshold* proposed by Agarwal et al. (2013), solve a dynamic choice problem that incorporates the rate gap, expectations about future interest rates, transaction costs, and other exogenous shocks, such as expected mobility, to derive a conservative threshold for when refinancing is optimal. While theoretically grounded, such measures rely on strong assumptions, and can be difficult to interpret. *Monthly savings* measure the net reduction in the borrower's monthly mortgage payment from refinancing. This metric incorporates loan balance, remaining term, and refinancing fees, providing a transparent measure of refinancing value that captures the cash-flow gains most salient to households (e.g., Johnson et al., 2019). In practice, refinance advertising and lender communications routinely frame refinancing benefits in terms of monthly payment reductions, consistent with the idea that borrowers evaluate refinancing primarily through its impact on monthly affordability and household cash flow.

All three measures require predicting the interest rate a borrower would receive if they were refinancing at a given point in time. In practice, refinancing rates for conforming loans vary systematically with borrower risk characteristics through the Loan-Level Price Adjustment (LLPA) schedule, which maps combinations of credit scores and LTV ratios into rate adjustments added to the benchmark base rate.[12] Because both credit scores and LTV ratios evolve over time, meaningful incentive measures must be borrower-specific and time-varying. We leverage the Equifax match, which provides updated credit scores, together with our imputed LTV ratios to predict individualized refinancing rates at a monthly frequency.

We adopt a two-stage hedonic rate prediction model (details in Appendix F). In the first stage, we regress observed refinancing rates on a full set of year-month $\times$ LLPA cell fixed effects, capturing the nonlinear pricing schedule over credit scores and LTV ratios interacted with time, along with state $\times$ 3-digit ZIP code fixed effects and loan size decile fixed effects. This flexible specification

[12]For an example of an LLPA schedule, see Figure F1 in the Online Appendix F.

absorbs both granular borrower risk pricing and local market conditions. The model explains 83% of the variation in observed refinancing rates.[13] In the second stage, we apply the estimated pricing function to the full loan-month panel, using each borrower's updated credit score and LTV ratio to predict the interest rate they would have faced in each month, had they refinanced. We then use the predicted interest rate for each loan-month to calculate *rate gap*, the *ADL threshold*, and *monthly savings* (details in Appendix G.4).

3.4 Refinancing Patterns

We now document two sets of baseline patterns that motivate our empirical analysis of advertising: the nonlinear relationship between incentives and refinancing, and the demographic disparities in refinancing conditional on incentives.

3.4.1 Refinancing Hazard with Incentives

Figure 3 Panel (a) plots the raw monthly refinancing hazard (solid line) against the distribution of each incentive metric (histogram bars). Across all three measures, the pattern is strikingly consistent. When incentives are negative (e.g., when refinancing would raise the borrower's payments) the refinancing hazard is low and nearly flat below 0.5%. The hazard then rises sharply: for monthly savings, the monthly refinancing hazard climbs from roughly 0.3% at zero to over 3% at \$200; for the rate gap, from roughly 0.2% at zero to nearly 3% at a 1.5 percentage points gap; and for the ADL threshold, from near zero at −1 to over 3% once the threshold crosses zero. At the highest incentive levels, however, the response flattens or even slightly declines, producing the S-shaped pattern visible in all three panels.

To characterize this relationship flexibly, we estimate a linear probability model of refinancing on a piecewise linear spline of the incentive variable:

$$\text{refi}_{it} = \sum_j \alpha_j \, S_j(\text{Incentive}_{it}, \boldsymbol{\kappa}) + \gamma D_i + \lambda X_{it} + \eta_{q(i)} + \mu_{d(i)} + \delta_t + e_{it}, \tag{1}$$

[13] Augmenting the pricing model with borrower demographics yields only a negligible improvement in fit; predicted rates with and without demographics have a correlation exceeding 0.99. We use the specification without demographics in the main analysis, as it can be estimated on the full Freddie Mac refinancing sample (N = 4,974,143). Adding race, gender, and income as regressors shows that Black borrowers pay approximately 1.4 basis points more than comparable White borrowers, consistent with prior findings. See Table F1 in Appendix F.

Figure 3: Refinancing Hazard with Incentives

Panel (a): Raw Refinance Hazard

(a) Monthly Saving

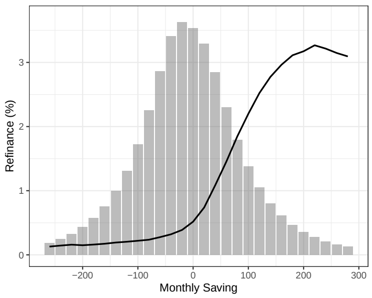

(b) Rate Gap

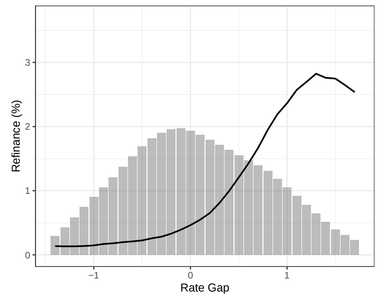

(c) ADL Incentive

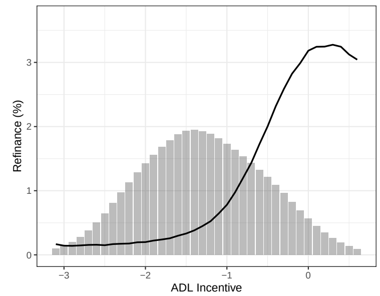

Note: Solid lines plot the raw monthly refinancing hazard against each incentive measure. Histogram bars show the distribution of incentives across loan-month observations.

Panel (b): Estimated Refinance Hazard

(d) Monthly Saving

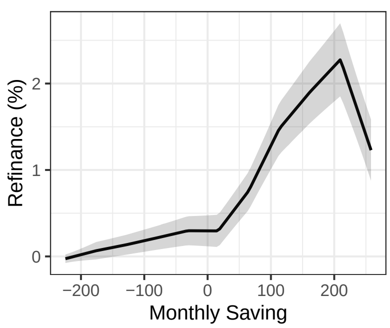

(e) Rate Gap

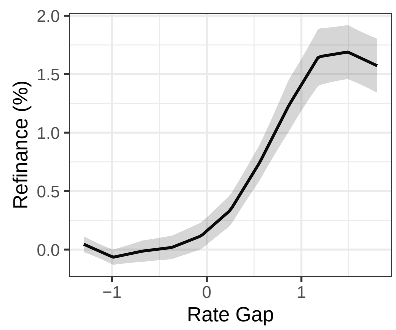

(f) ADL Incentive

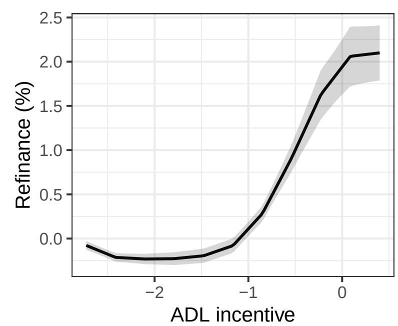

Note: Solid lines plot the estimated monthly refinancing hazard, with shaded areas indicating 95% confidence bands. Controls include borrower demographics, time-varying credit characteristics, and DMA, cohort, and year-month fixed effects. Standard errors are double-clustered by DMA and origination year-quarter.

where refi_{it} is an indicator of whether the borrower i refinances in month t, conditional on the loan remaining active. The spline basis functions S_j are defined over knots $\boldsymbol{\kappa}$ placed at the 25th, 50th, 60th, 70th, 80th, 90th, and 95th percentiles of the incentive distribution.[14] The vector D_i includes indicators for race/ethnicity, household income quartile, and gender/joint-filing status. Time-varying controls X_{it} include credit score bins, indicators for LTV above 80% and 90%, a quadratic in loan age, and delinquency and modification history. We include DMA fixed effects $\mu_{d(i)}$, year-month fixed effects δ_t, and origination cohort (year-quarter) fixed effects $\eta_{q(i)}$. Standard errors are double-clustered by DMA and origination year-quarter.

[14]Specifically, $S_1(x;\boldsymbol{\kappa}) = \min\{x, \kappa_1\}$, $S_j(x;\boldsymbol{\kappa}) = \max\{0, \min\{x, \kappa_j\} - \kappa_{j-1}\}$ for interior segments, and $S_8(x;\boldsymbol{\kappa}) = \max\{0, x - \kappa_7\}$.

Panel (b) of Figure 3 reports the estimated refinancing hazard as a function of each incentive measure, with 95% confidence bands. The nonlinear pattern observed in the raw data persists after conditioning on observables and fixed effects. For monthly savings, the estimated hazard is very low when potential savings are negative and rises sharply once savings turn positive, reaching approximately 2.3% at around \$200 in monthly savings before declining. The rate gap and ADL figures display a similar pattern: a flat region at low incentives, followed by a steep increase, and a flattening or slight decline at the upper end of the incentive distribution. The consistency of this shape across three distinct incentive measures, each with different assumptions and units, confirms that the nonlinearity is a robust feature of refinancing behavior. The flattening at very high incentive levels is consistent with a selection interpretation: borrowers who remain unrefinanced despite large potential gains are likely those most affected by inattention, procrastination, or other frictions that dampen their responsiveness to strong economic incentives.

3.4.2 Refinancing Disparities

We next examine whether the refinancing disparities observed in the raw data persist after conditioning on individualized incentives. We interact the spline bases with D_i, which includes borrower income quartile, race/ethnicity, and gender/joint-filing indicators:

$$\text{refi}_{it} = \sum_j (\alpha_{1j} + \alpha_{2j} D_i)\, S_j(\text{Incentive}_{it}, \boldsymbol{\kappa}) + \lambda X_{it} + \mu_{d(i)} + \eta_{q(i)} + \delta_t + e_{it} \tag{2}$$

All other variables are defined as before. This specification allows the refinancing response to incentives to differ flexibly across groups while controlling the underlying refinance incentive.

Figure 4 plots the estimated group-specific refinancing hazards as a function of monthly savings. Figure 4 (a) shows that income disparities are large and widen dramatically at high incentive levels. At \$200 in potential monthly savings, borrowers in the top income quartile (Q4) refinance at roughly 2.3% per month, compared with only about 0.5% for borrowers in the bottom quartile. This nearly 1.8 percentage points gap is roughly four times the refinancing hazard of Q1 borrowers. The middle quartiles (Q2 and Q3) fall in between, with Q3 closely tracking Q4. For non-positive monthly savings, refinancing hazards across all four groups cluster near zero with negligible differences.

Figure 4 (b) shows similar pattern for racial disparities. At high incentive levels, White and Asian borrowers refinance at approximately 2.0–2.3% per month, while Black borrowers peak at roughly 1.5% and Hispanic borrowers at about 1.8%. As with income, the gaps are negligible in the negative-savings region and emerge only once refinancing becomes financially attractive.

Figure 4: Refinancing Decisions by Income and Race/Ethnicity

(a) Income Category

(b) Race and Ethnicity

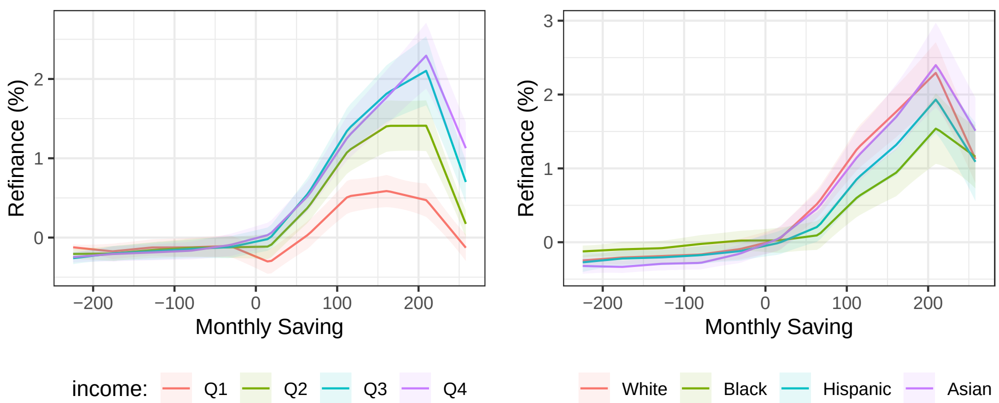

Figure 5 reports the pairwise group differences along with their 95% confidence intervals. The income gap (Q4 minus Q1) rises steadily from near zero at negative savings to approximately 1.8 percentage points at \$200 in monthly savings. The White–Black gap follows a hump-shaped pattern, rising from zero to a peak of roughly 0.8 percentage points around \$150 in savings before narrowing at the very highest incentive levels. The White–Hispanic gap is smaller in magnitude, peaking at approximately 0.4 percentage points. All gaps are estimated precisely in the positive-savings region. Results using the rate gap and ADL measures are reported in the lower panels of Figure 5 and are qualitatively similar.

These patterns are consistent with earlier work by Agarwal et al. (2024), Gerardi et al. (2023) and Keys et al. (2016). Disparities in refinancing have important implications. Refinancing frictions disproportionately affect groups for whom refinancing would generate the largest financial gains. Since lower-income households typically have higher marginal propensities to consume, such disparities weaken the pass-through of interest rate declines to aggregate consumption. More broadly, the concentration of refinancing gaps at high incentive levels highlights a last-mile problem in household finance: even when refinancing is clearly beneficial, many households, particularly those already disadvantaged, fail to translate incentives into realized financial benefits.

The strong nonlinearity in refinancing response and the concentration of large demographic disparities at high incentive levels are baseline patterns that set the stage for our central empirical questions. The nonlinear response suggests that the effect of any intervention, including advertising,

Figure 5: Demographic Gaps in Refinancing Conditional on Incentives

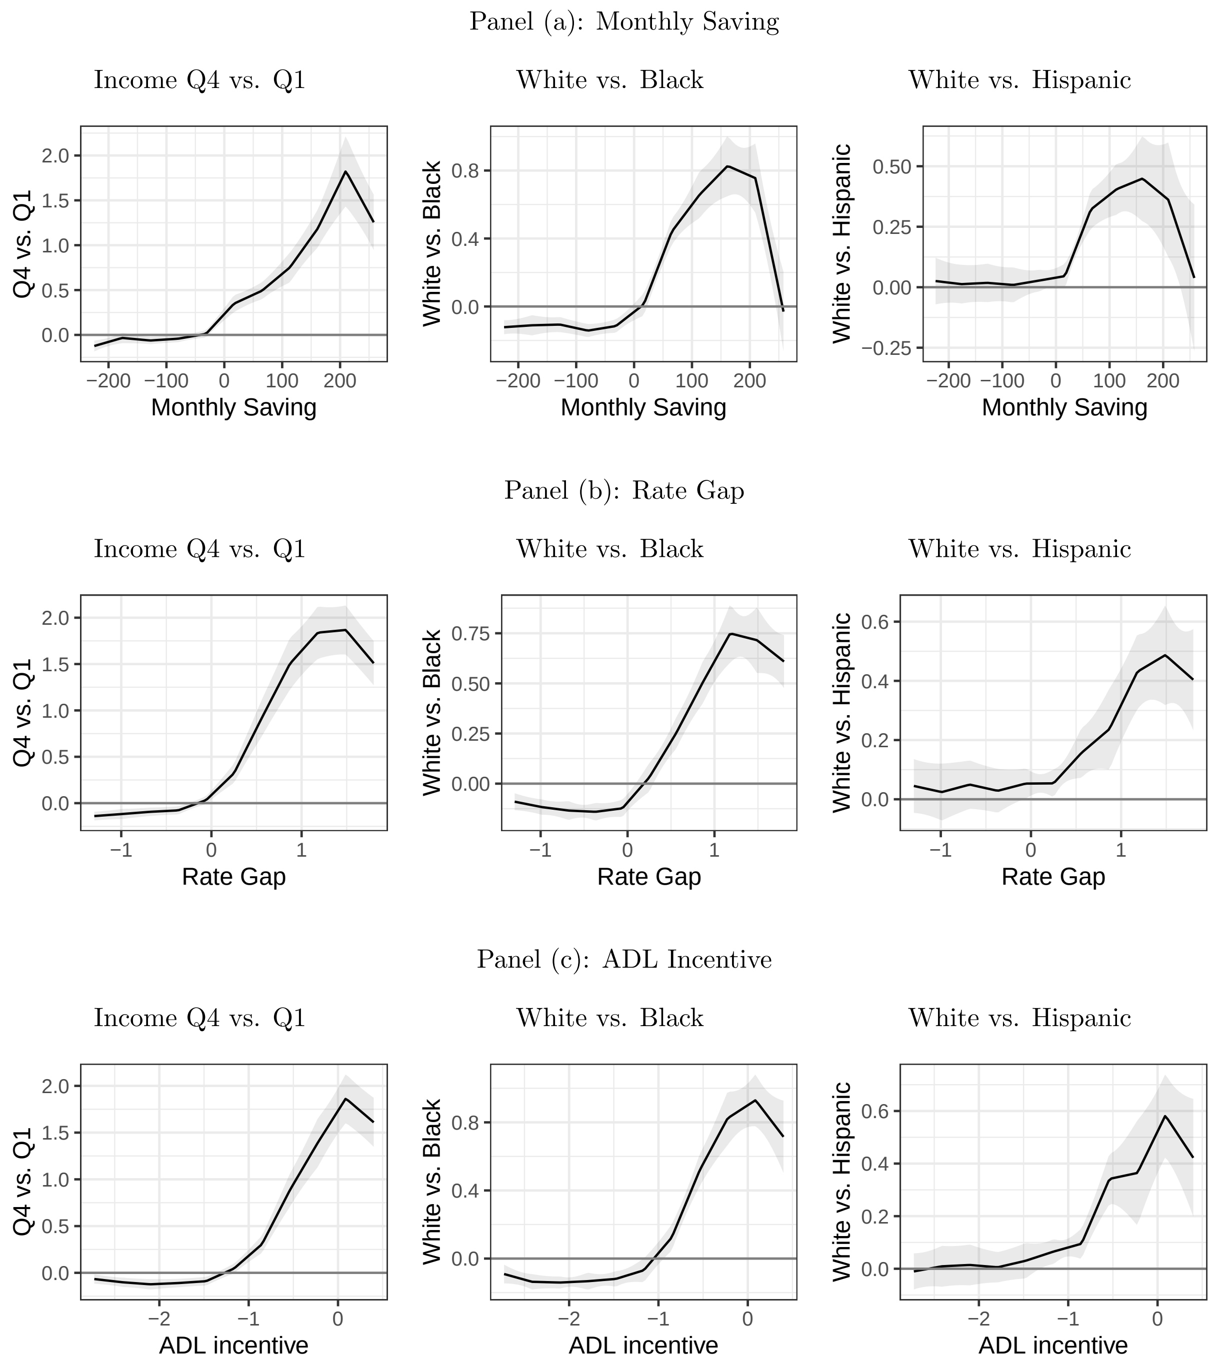

Note: Solid lines plot pairwise differences in estimated refinancing hazard between demographic groups as a function of refinancing incentives. Shaded bands show 95% confidence intervals. Positive values indicate higher refinancing rates for the first-named group. Controls and fixed effects as in Figure 3 Panel (b).

may depend on where in the incentive distribution it operates. The demographic gaps, occurring precisely where refinancing is most valuable, raise the question of whether advertising narrows or widens existing disparities. We turn to these questions next.

4 Impact of TV Advertising on Refinancing Behavior

We now turn to the central question of the paper: does television advertising affect refinancing behavior, and if so, where in the incentive distribution does it operate? The baseline patterns documented in Section 3.4 showed that refinancing responds nonlinearly to financial incentives, with the steepest response in the region where savings first become positive. This nonlinearity is relevant to consider for policy implications of advertising. If advertising primarily operates in regions where refinancing is financially beneficial but underutilized, it will strength monetary policy pass-through. Conversely, if its effects are concentrated at low or negative incentive levels, advertising may raise concerns about inefficient refinancing activity.

4.1 Empirical Specification

To estimate how advertising effects vary across the incentive distribution, we augment the baseline spline specification from Section 3.4 by interacting advertising stock with indicators for each incentive region:

$$\begin{aligned} \text{refi}_{it} = &\sum_{j=1}^{10} \alpha_j \, S_j(\text{Incentive}_{it}, \boldsymbol{\kappa}) + \sum_{j=1}^{10} \beta_j \text{AdStock}_{it} \cdot \mathbf{1}(\kappa_{j-1} < \text{Incentive}_{it} < \kappa_j) + \\ &\gamma D_i + \lambda X_{it} + \eta_{q(i)} + \mu_{d(i)} + \delta_t + e_{it}. \end{aligned} \tag{3}$$

S_j are the spline basis functions defined in Section 3.4, and the indicator functions $\mathbf{1}(\cdot)$ partition the incentive distribution into regions using the same knots $\boldsymbol{\kappa}$. The coefficients β_j capture the effect of a one-standard-deviation increase in advertising stock on the refinancing hazard within each incentive region, allowing advertising to shift the *level* of refinancing within regions. All other controls and fixed effects, demographics (D_i), time-varying borrower characteristics (X_{it}), DMA fixed effects ($\mu_{d(i)}$), year-month fixed effects (δ_t), and origination-cohort fixed effects ($\eta_{q(i)}$), are as described previously. Standard errors are double-clustered by DMA and origination year-quarter.

Our identification strategy follows Shapiro et al. (2021) and leverages institutional features of the television advertising market that limit advertisers' ability to target short-run, market-specific refinancing demand. Television advertising is purchased across four media types—Cable, Network, Syndicated, and Spot—that differ in how ads are bought and distributed across designated market areas (DMAs). Cable ads are bought and aired nationally; Network and Syndicated ads are purchased nationally but cleared and distributed locally, resulting in differential impressions across DMAs; and Spot ads are bought and aired within specific DMAs.

Crucially, the majority of TV advertising (roughly 80%) is purchased months in advance in the "upfront" market, where advertisers commit to blocks of ad inventory before observing the precise timing of macroeconomic shocks or local demand conditions. The remaining inventory is sold in the "scatter" market, where supply is limited and prices fluctuate with factors unrelated to mortgage demand, such as political advertising, sports programming, and network scheduling changes. Because of these institutional constraints (i.e., bulk buying, limited scatter-market availability, bundled inventory, and the local distribution of nationally purchased ads), the realized distribution of advertising across DMAs and months is only loosely linked to contemporaneous mortgage refinancing demand, even when aggregate advertising budgets respond to broad interest-rate environments.

Our specification flexibly controls for predictable differences in refinancing across media markets, borrower cohorts, and aggregate interest-rate environments: DMA fixed effects absorb time-invariant differences in mortgage market competitiveness, lender mix, and media-market characteristics, year-month fixed effects absorb nationwide shifts in refinancing demand driven by macroeconomic conditions, and origination-cohort fixed effects absorb systematic differences in refinancing propensity across loans originated under different credit environments.

After removing these predictable components, the remaining within-DMA, within-month variation in advertising exposure reflects quasi-random fluctuations induced by the institutional frictions in the TV advertising market (i.e.,unpredictable slot availability, bundling, pricing shocks in local media markets, and networks' distribution of unsold national inventory across DMAs), and are plausibly orthogonal to local refinancing demand. After absorbing fixed effects, the residual variation in advertising stock is substantial, providing the identifying variation needed to estimate causal advertising effects (see Figure E3 and the discussion in Online Appendix E).

4.2 Results

Figure 6 presents the main results. For each incentive measure, the left panel plots the estimated refinancing hazard at mean advertising exposure (normalized ad GRP stock = 0) and at one standard deviation above the mean (ad GRP stock = 1). The right panel plots the advertising lift (i.e., the difference between these two curves) with 95% confidence intervals.

The results reveal a clear pattern that is consistent across all three incentive measures. When financial incentives are low or negative, advertising has no detectable effect on refinancing: the two curves in each left panel overlap almost exactly in this region, and the corresponding lift estimates in the right panels are indistinguishable from zero. The effect of advertising emerges only once refinancing becomes financially attractive and grows larger as incentives increase.

For monthly savings (panel a), the two curves begin to diverge around \$50 in potential savings. The advertising lift rises from approximately zero at negative savings levels to about 0.10 percentage points at \$50 in monthly savings, 0.17 at \$100, and approximately 0.25–0.30 percentage points at \$150–\$250 in monthly savings. The 95% confidence interval excludes zero for all positive-savings bins. At the highest incentive levels, this lift represents roughly a 10–15% increase relative to the baseline refinancing hazard—a substantial effect for a mass-market intervention. The rate gap (panel b) shows the same pattern: negligible effects at negative rate gaps, with the lift rising to approximately 0.15–0.23 percentage points at rate gaps above one percentage point. For the ADL measure (panel c), the pattern is similar, with close to zero effects at lower incentives and reaching approximately 0.15–0.22 percentage points in the region where refinancing is unambiguously optimal according to the option-value criterion.[15]

As a robustness check, we replace the step-function specification with a continuous spline specification by interacting spline basis functions, rather than incentive region indicators, with advertising stock. This allows the advertising effect to vary smoothly across the incentive distribution. All other controls, fixed effects, and clustering remain the same. Figure A1 in Appendix A.1 reports the results. The magnitude and overall pattern of the advertising effect across incentives closely mirror the main analysis. For the rest of the analysis, we use the step-function approach to preserve a more interpretable set of region-specific advertising effects.

Downstream real effects of refinancing are economically meaningful. A growing body of evidence documents sizable consumption responses to mortgage payment reductions and refinancing activity. Di Maggio et al. (2017) shows that adjustable-rate mortgages resets that lower monthly payments increase car purchases by 35 percent. In the post-crisis low-rate refinancing boom, Di Maggio et al. (2020) find that increased refinancing of conforming loans generated substantial payment reductions

[15]Table C1 in Appendix C reports the underlying regression coefficients. The point estimates confirm the visual pattern: coefficients on advertising × incentive-bin interactions are small and statistically insignificant in the low-incentive region and become larger, positive, and significant in the high-incentive regions.

Figure 6: Overall Ad Effect on Refinance

Panel (a): Monthly Saving

Refinance by Ad Levels

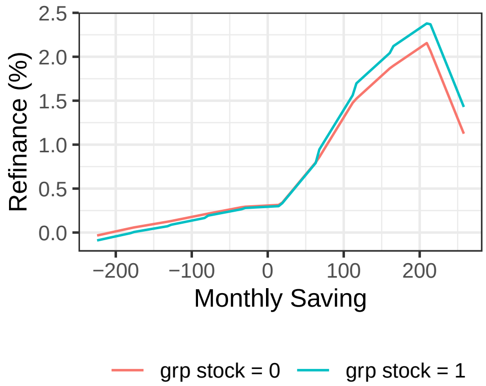

(a) Ad Effect on Refinance

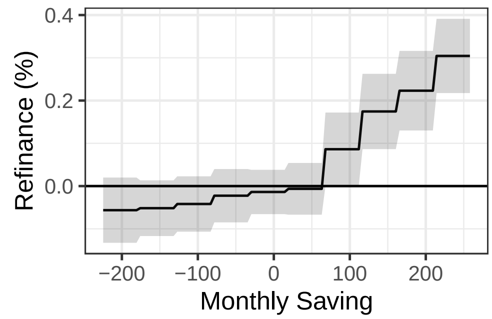

Panel (b): Rate Gap

Refinance by Ad Levels

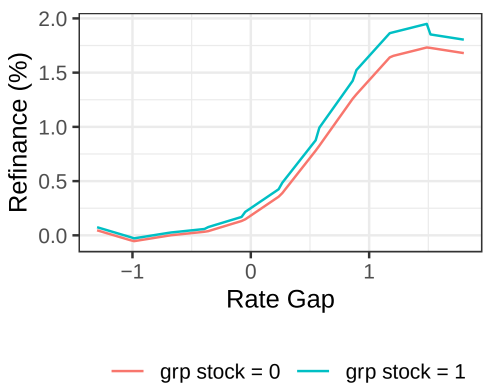

(b) Ad Effect on Refinance

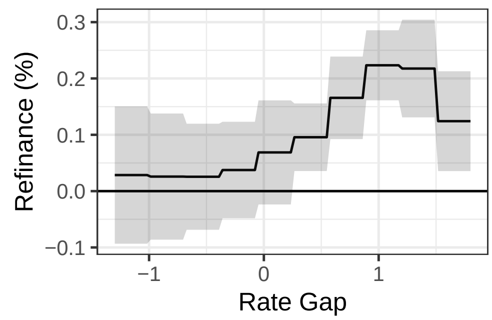

Panel (c): ADL Incentive

Refinance by Ad Levels

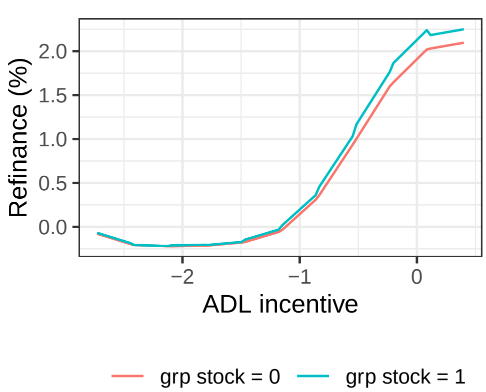

(c) Ad Effect on Refinance

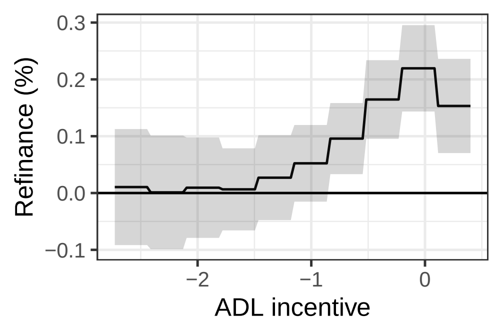

and raised durable spending by 12% among refinancing households. Exploiting eligibility rules in the Home Affordable Refinancing Program (HARP), Agarwal et al. (2023) find an increase of about $1,600 in durable expenditures over two years following refinancing, with larger increases among more indebted borrowers. In a structural life-cycle framework, Eichenbaum et al. (2022) estimate that a 50 basis point mortgage rate decline raises consumption by 1-1.4% depending on the distribution of potential refinancing savings.

Our results have two important implications for evaluating advertising as a potential policy tool. First, the absence of advertising effects at low or negative incentive levels indicates that TV advertising does not meaningfully induce refinancing for motives other than lowering rates. This null effect is reassuring in light of recent evidence that the cash-out channel of monetary policy has limited real effects: Anenberg et al. (2025) show that most cash-out proceeds substitute for other forms of borrowing, yielding a marginal propensity to consume of only $0.04 per dollar of equity extracted. Second, because advertising operates precisely where potential rate-reduction savings are largest, it strengthens the pass-through of monetary policy to household balance sheets. By prompting refinancing among borrowers with substantial unrealized gains from lower rates, advertising increases the effective take-up of monetary stimulus, and narrows the gap between the theoretical cash-flow benefits of rate cuts and their realized impact on household finances.

4.3 Discussion: Cost-Effectiveness and Magnitudes

To quantify the economic benefit generated by TV advertising, we compute the expected financial savings induced by a one-standard-deviation increase in advertising exposure. Since advertising effects decay gradually, a one-time increase in exposure at time t increases the refinancing hazard over several subsequent months. We summarize this dynamic effect by the six-month cumulative lift in refinancing. Specifically, let Δp_{it} denote the increase in the probability that household i refinances within six months when exposure at t is increased by one standard deviation, relative to the observed level of advertising.

We translate the increase in refinance probability into economic value by scaling it by the borrower's potential monthly savings. For each loan-month, the expected incremental savings are $\Delta p_{it} \times \text{MonthSave}_{i,t}$. We then compute the net present value of the stream of monthly payment reductions that refinancing would generate over the remaining loan term, denoted NPV_{it}. The expected long-term financial savings attributable to the advertising lift is therefore:

$$\Delta p_{it} \times \text{NPV}_{it}.$$

We aggregate these expected savings across loan-months and divide by the number of households to obtain per-household expected benefits. We also report bounds: the upper bound assumes savings are realized over the full remaining loan term, while the lower bound incorporates a prepayment hazard that shortens the effective horizon over which refinancing savings are realized (see Appendix G.4).

A one-standard-deviation increase in advertising exposure generates an NPV of savings between \$13–\$35 per household, depending on the discounting assumption. Restricting the calculations to borrowers with positive savings raises the implied NPV to \$21-\$57. The effect is substantially larger in low-rate environments, when both refinancing responsiveness and potential payment savings are higher. For example, in early 2020 (mortgage rates 3.13%-3.72%), the implied NPV of savings is \$14-\$38 per household, whereas in late 2020 (rates 2.66%-3.07%), it increases to \$32-\$88 per household.

These gains come at a modest cost. A one-standard-deviation increase in advertising exposure costs approximately \$0.30 per household.[16] Comparing expected savings-NPV to this incremental advertising cost implies a benefit-cost ratio of approximately 40:1 to 110:1 in the entire sample (or 70:1 to 190:1 among those with positive savings). In lower rate environments, this ratio is even higher: 45:1 to 125:1 in early 2020 and 107:1 to 294:1 in late 2020. Overall, these magnitudes suggest that television advertising operates as a highly efficient channel for generating household savings.

To benchmark these magnitudes, it is useful to compare TV advertising to the targeted mail-based interventions studied in the recent work. Byrne et al. (2023) send letters and reminders to 12,000 Irish mortgage holders. The most effective treatment, two letters costing €1 per household, generated average savings with an NPV of about €84 per household. Similarly, Grundl and Kim (2019) study direct mail refinancing offers sent to 10,000 US households and find that such mailings increase savings by \$36 - \$52 per household. However, Keys et al. (2016) find that mailed offers

[16] To compute the per-household cost of a one-standard-deviation increase in advertising exposure, we first calculate the average cost per GRP as mean spending divided by GRPs across all ad airings. We convert this measure to a per-household basis by dividing by the total number of U.S. TV households, and then multiply by the standard deviation of ad exposure.

to roughly 200 low-income Chicago households had little effect, and Johnson et al. (2019) find no meaningful increase in applications from pre-approved HARP offers targeted to 80,000–110,000 eligible borrowers per wave.

Direct mail is typically targeted using borrower-level information, concentrating outreach on households with the strongest refinancing incentives. Television advertising, by contrast, trades precision for scale. Our results show that it can still produce sizable savings at a favorable cost.

Because television advertising is untargeted, it reaches many households with little or no refinancing incentives at a given time. Yet we do not observe an increase in refinancing among borrowers in the low-incentive region. This pattern likely reflects the institutional environment governing the refinancing process. In particular, regulatory guardrails such as the TILA-RESPA Integrated Disclosures (TRID) require standardized and transparent disclosure of loan terms, payments, and closing costs. As a result, while advertising may prompt borrowers to explore refinancing, it cannot circumvent the regulated information environment that makes the costs and limited benefits salient.

To further probe this idea, we examine whether advertising increases *search* even when it does not culminate in refinancing. We proxy for active search using hard credit inquiries.[17] We estimate an analogue of Equation (3), replacing the refinancing outcome with an indicator for a mortgage-related credit inquiry indicator and the using a one-month lag of ad stock, since inquiries typically occur about a month prior to origination. Figure 7 plots the estimated advertising effect on inquiry behavior across the incentive distribution. In contrast to the refinancing results, advertising increases credit inquiries across a broader range of incentives, including levels at which the households would not realize any savings.

This pattern is consistent with advertising broadly stimulating search activity, even among households with little savings potential. Yet, in the low-incentive region, this increased search does not translate into completed refinancing. Instead, the refinancing process and its disclosures appears to discipline the final decision, screening out marginal or unprofitable refinances as borrowers confront standardized information on rates, payments, and closing costs during the origination stage.

[17] A hard inquiry (recorded in Equifax credit panel) occurs when a borrower authorizes a lender to pull their credit file, typically at the application stage. Unlike soft inquiries, which can be initiated without consumer consent, hard inquiries reflect verified borrower-initiated activity. Although mortgage lenders typically pull consumer credit reports from all three major credit bureaus, some inquiries made by lenders can be misclassified as non–mortgage inquiries. To address this measurement issue, we flag borrowers who refinance but have no recorded mortgage inquiry in the three months prior to refinancing. For these cases, we impute a single inquiry dated one month prior to refinancing, which increases the total from approximately 975 thousand to one million loan-month observations with an inquiry history.

Figure 7: Overall Ad Effect on Credit Inquiry

Panel (a): Monthly Saving

Credit Inquiry by Ad Levels (a) Ad Effect on Credit Inquiry

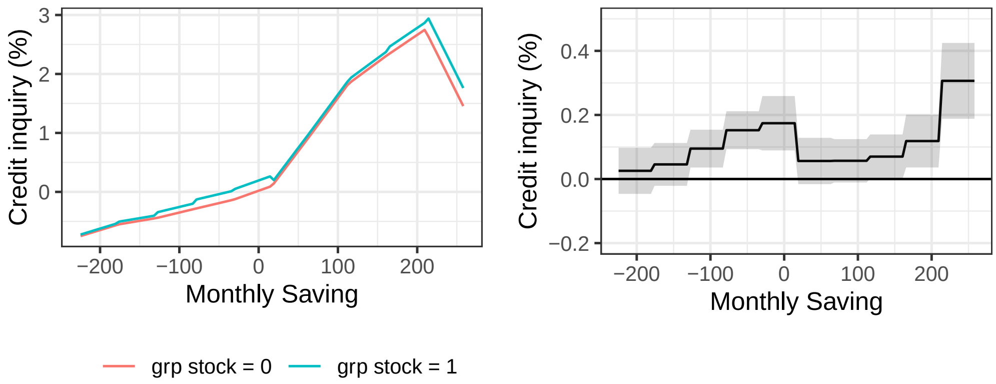

Panel (b): Rate Gap

Credit Inquiry by Ad Levels (b) Ad Effect on Credit Inquiry

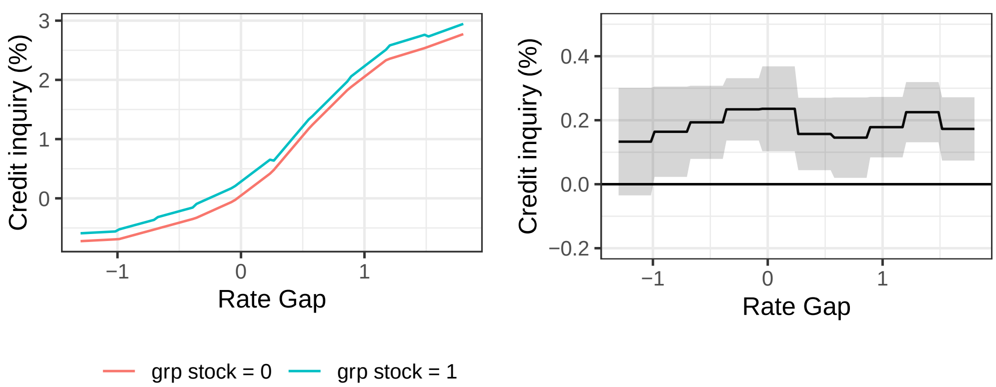

Panel (c): ADL Incentive

Credit Inquiry by Ad Levels (c) Ad Effect on Credit Inquiry

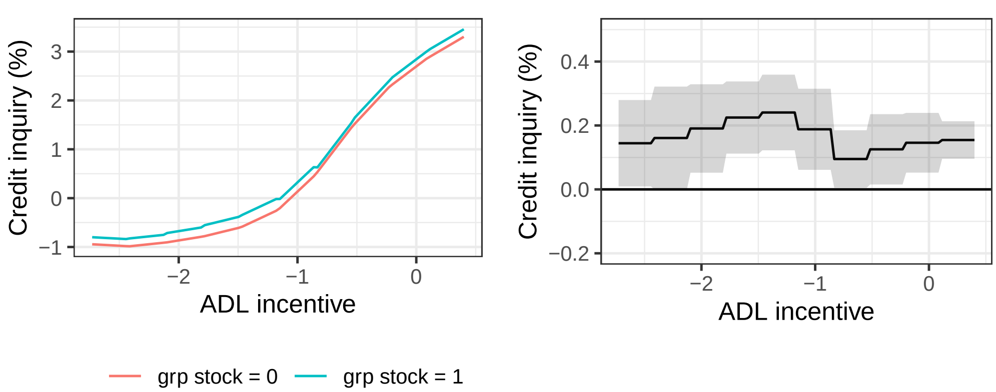

5 Impact of TV Advertising on Refinancing Disparities

Section 4 establishes that television advertising increases refinancing precisely where financial incentives are strongest, thereby strengthening the pass-through of monetary policy on average. Yet, the increase in the aggregate can mask important distributional consequences. As documented in Section 3.4, there are substantial refinancing disparities across income and racial groups, particularly when potential savings are large. We now examine whether advertising narrows or widens the disparities across demographic groups.

5.1 Empirical Specification

Whether advertising narrows or widens disparities depends on where the gaps arise in the refinancing process. If disparities primarily reflect differences in awareness or inattention, advertising may reduce the gaps by equalizing access to information and increasing engagement in the upper funnel. If, instead, disparities are driven by downstream barriers (e.g., documentation burdens, differential lender screening and capacity constraints), then increased awareness in the upper funnel from advertising may translate to unequal completion rates, potentially widening the gaps in realized refinancing outcomes.

To test for heterogeneous advertising effects across demographic groups, we extend the baseline specification in Section 4.1 to allow the advertising effect to vary jointly by financial incentives and borrower characteristics. Specifically, we estimate:

$$\begin{aligned}\text{refi}_{it} = &\sum_{j=1}^{10} \alpha_j\, S_j(\text{Incentive}_{it}, \boldsymbol{\kappa}) + \sum_{j=1}^{10} \alpha_j^D\, S_j(\text{Incentive}_{it}, \boldsymbol{\kappa}) \cdot\, D_i + \\ &\sum_{j=1}^{10} \beta_j \text{AdStock}_{it} \cdot\, \mathbf{1}(\kappa_{j-1} < \text{Incentive}_{it} < \kappa_j) + \sum_{j=1}^{10} \beta_j^D\, \text{AdStock}_{it} \cdot\, \mathbf{1}(\kappa_{j-1} < \text{Incentive}_{it} < \kappa_j) \cdot\, D_i + \\ &\gamma D_i + \lambda X_{it} + \eta_{q(i)} + \mu_{d(i)} + \delta_t + e_{it}.\end{aligned} \tag{4}$$

As in Equation (3), the dependent variable refi_{it} equals 1 if borrower i refinances in month t. The spline terms $\sum_{j=1}^{10} \alpha_j S_j(\text{Incentive}_{it}, \boldsymbol{\kappa})$ flexibly capture the nonlinear relationship between refinancing incentives and the refinancing hazard. Importantly, $\sum_{j=1}^{10} \alpha_j^D S_j(\text{Incentive}_{it}, \boldsymbol{\kappa}) \cdot D_i$ allows the baseline sensitivity to incentives to vary across demographic groups, so that any heterogeneity in advertising effects is not mechanically driven by pre-existing differences in how groups respond to incentives.

Heterogeneity in advertising effects enters through the interaction between AdStock_{it} and incentive-region indicators. AdStock_{it} denotes standardized advertising stock. The coefficients β_j capture the effect of a one-standard-deviation increase in advertising exposure within incentive region j for the omitted demographic group, while β_j^D measures the differential advertising lift for group D_i relative to the omitted category within the same incentive region. The indicators $\mathbf{1}(\kappa_{j-1} < \text{Incentive}_{it} < \kappa_j)$ partition the incentive distribution using knots $\boldsymbol{\kappa}$.

Borrower demographic characteristics D_i, controls X_{it} and fixed effects follow the baseline specifications. The vector X_{it} includes updated credit score bins, loan-to-value indicators, delinquency and modification history, and loan age. We include DMA fixed effects $\mu_{d(i)}$, origination cohort (year-quarter) fixed effects $\eta_{q(i)}$, and year-month fixed effects δ_t. Standard errors are double-clustered by DMA and origination year-quarter.

5.2 Results

Figure 8: Heterogeneous Ad Effect on Refinancing

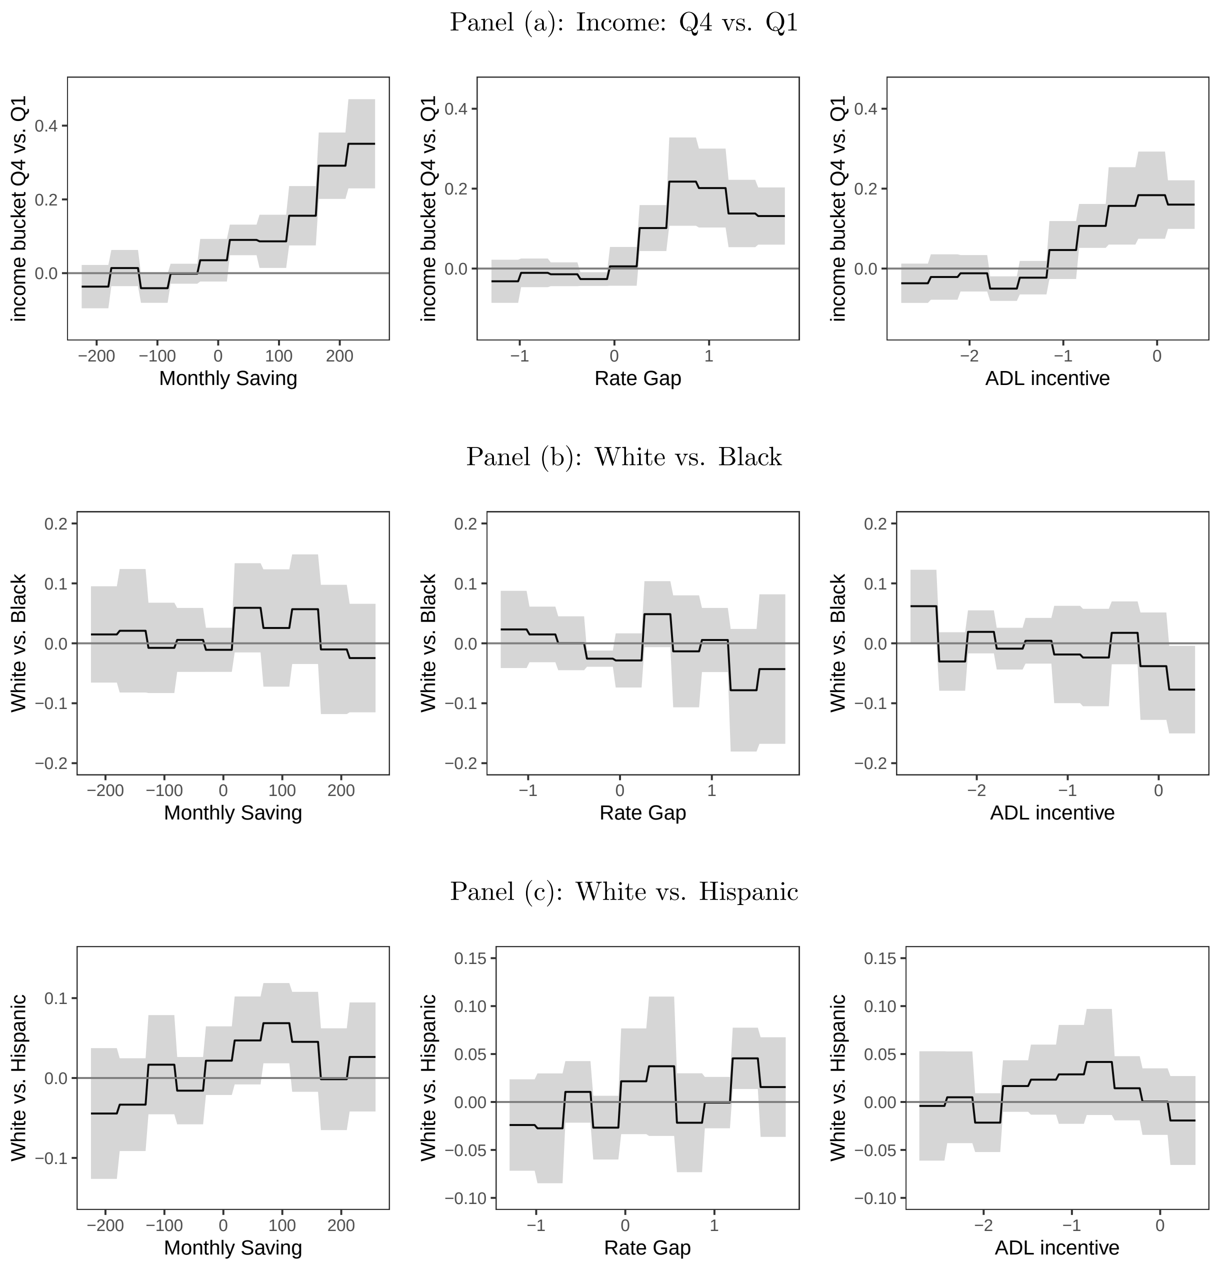

Figure 8 plots the estimated differential advertising lift across demographic groups as a function of refinancing incentives. Each panel reports the differential effect of advertising between the indicated groups, based on the triple-interaction coefficients β_j^D from Equation (4). Shaded bands denote 95% confidence intervals.

We begin with income disparities. Figure 8 Panel (a) plots the differential advertising lift between borrowers in the highest income quartile (Q4) and those in the lowest quartile (Q1). When refinancing incentives are weak or negative, the differential lift is close to zero and statistically insignificant, consistent with our broader finding that advertising has little effect when financial gains are small. In the high-incentive regions, however, an income differential emerges. For borrowers with large potential savings (roughly \$150–\$250 in monthly payment reductions), the advertising lift for Q4 exceeds that for Q1 by about 0.30–0.35 percentage points per month. The 95% confidence interval excludes zero for all bins above approximately \$50 in monthly savings. In proportional terms, this differential represents a widening of the income-based refinancing savings gap of roughly 18% relative to the baseline gap documented in Section 3.4 in the regions where refinancing is most financially attractive.

This pattern is robust to alternative measures of incentives. Using the rate gap or the ADL threshold yields the same qualitative conclusion: once refinancing becomes clearly beneficial, advertising generates a larger increase in refinancing among higher-income borrowers. Appendix C.2 reports the regression coefficients.

In contrast, we find little evidence of differential effects by race or ethnicity. Panels (b) and (c) show that the estimated White–Black and White–Hispanic differentials fluctuate around zero and are statistically indistinguishable from zero in nearly all incentive regions. The confidence intervals are sufficiently tight to rule out economically large racial differences in advertising lift, especially in the high-incentive regions where refinancing activity is most responsive. Thus, while substantial baseline racial disparities in refinancing persist (Section 5.2), advertising does not appear to amplify them.

Taken together, these results indicate that television advertising does not widen racial disparities in refinancing but does increase income-based disparities at high incentive levels. While advertising lifts refinancing for all groups where incentives are strong, it differentially increases the realized gains from refinancing for the higher income group.

5.3 Discussion

To translate these heterogeneous effects into dollar terms, we compute group-specific expected savings following the approach in Section 4.3. For Q4 households, the advertising-induced increase in expected monthly savings is approximately \$27–74 per household, depending on the discount window. For Q1 households, the corresponding figure is much lower at about \$2–4 per household.

Both groups gain from advertising relative to a no-advertising counterfactual, but the gain is substantially larger for high-income borrowers.

The larger advertising lift observed for higher-income borrowers can arise through several distinct channels: differences in advertising exposure, differences in initial engagement in response to advertising, or differences in the likelihood that initial engagements convert into completed refinances. Distinguishing among these mechanisms is important for interpreting whether advertising interacts with downstream frictions. We examine evidence along each margin in turn.

The first possibility is differential exposure: low-income households see fewer ads, either because lenders concentrate advertising in higher-income DMAs or because viewing patterns differ across income groups. This channel would generate a mechanical difference in advertising "dose" without differential advertising responsiveness. Our specification estimates effects conditional on DMA-level advertising exposure and includes DMA fixed effects, absorbing systematic targeting across DMAs.

Furthermore, Appendix J provides direct evidence on advertising exposure by demographic group using cable television viewership data. If anything, Black and lower-income households receive slightly *more* advertising exposure than White and higher-income households. We therefore view differential exposure as an unlikely primary driver of the income differential in advertising lift, although we cannot rule it out completely.

The second possibility is that advertising induces different levels of initial engagement across income groups. Even with the same exposure, higher-income borrowers may be more likely to engage with the refinancing process after seeing ads, so heterogeneity in completed refinancing could reflect differences in responsiveness rather than differences in dose. To assess this margin, we use mortgage-related hard credit inquiries as a proxy for active engagement that requires borrower authorization. We estimate the analogue of Equation (4), replacing the refinancing outcome with an inquiry indicator and using a one-month lag of ad stock to reflect the typical gap between inquiry and loan origination.

Figure 9 reports the results. Panel (a) shows little evidence of systematic income differences in the inquiry response: the estimated Q4–Q1 differential in inquiry lift is small, noisy, and statistically indistinguishable from zero across virtually the entire incentive distribution. These results suggest that the income gap in refinancing lift does not appear to be primarily driven by the differential ad response at the awareness or inquiry stage.

Figure 9: Heterogeneous Ad Effect on Credit Inquiry

Panel (a): Income: Q4 vs. Q1

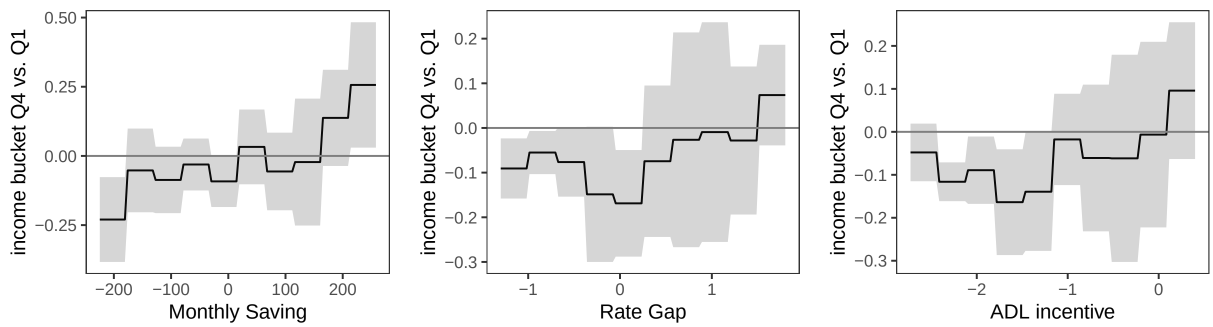

Panel (b): White vs. Black

Income: Q4 vs. Q1 White vs. Black White vs. Hispanic

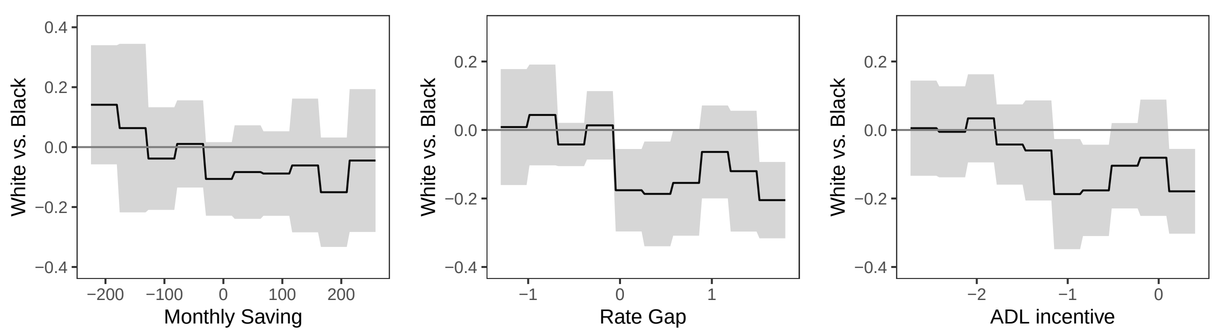

Panel (c): White vs. Hispanic

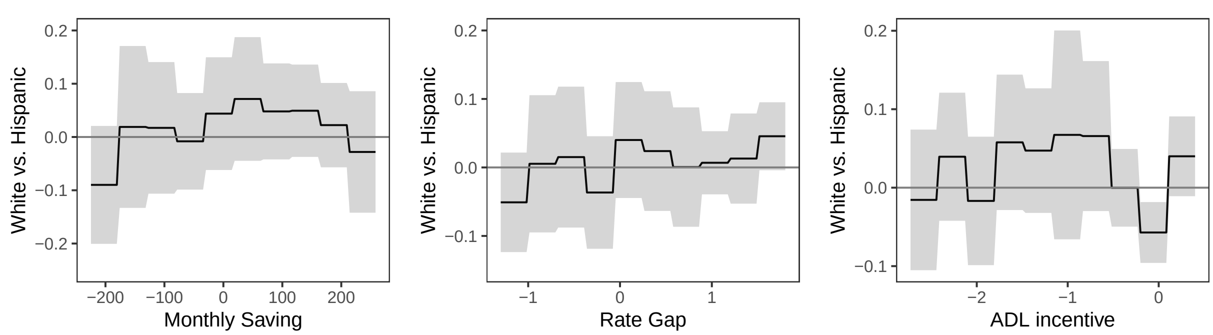

Note: The figure plots differential advertising effects on refinancing across demographic groups over the incentive distribution. Panel (a) shows Q4−Q1 (income), Panel (b) and Panel (c) show White−Black and White−Hispanic differences. Shaded bands denote 95% confidence intervals. Positive values indicate a larger advertising effect for the first-listed group.

.

A similar null result holds for race and ethnicity. Panel (b) reports the differential advertising lift on credit inquiries between White and Black borrowers, and Panel (c) reports the analogous White–Hispanic comparison. Across all three incentive measures, the estimated gaps hover around zero and are statistically indistinguishable from zero in nearly every incentive region. The confidence intervals are sufficiently tight to rule out economically meaningful racial or ethnic differences in inquiry responses to advertising.

Appendix B.1 reports the baseline disparities in credit inquiries, which are substantially smaller than refinancing disparities, consistent with the funnel interpretation that gaps widen between the search and closing stages. This pattern also aligns with Frame et al. (2025), who show that racial disparities in mortgage markets arise primarily at the loan officer interaction stage.

Taken together, the results suggest that the income disparity in completed refinancing does not appear to originate at the inquiry stage. The combination of similar inquiry responses but divergent refinancing responses points to differences in conversion from application to closing. After initiating contact with a lender, borrowers need to complete documentation, clear underwriting, and finalize the transaction – stages at which barriers may systematically differ across income groups.

On the borrower side, low-income households may have a higher opportunity cost of time, greater difficulty assembling documentation, or more volatile income and employment, reducing the likelihood that an initiated inquiry culminates in a completed refinance. Consistent with this interpretation, Agarwal et al. (2020) argue that screening amplifies the costs of refinancing in ways that differ across borrower types.

On the lender side, capacity constraints and profitability considerations may lead lenders to prioritize larger or more profitable loans, which are more likely to be held by higher income borrowers. During the 2020–2021 refinancing boom, supply constraints appear to have rationed marginal applicants: Frazier and Goodstein (2025) show that lenders screened out borrowers with low balances, low incomes, and low credit scores, while Agarwal et al. (2015) find that competitive frictions in the HARP market reduced take-up by 10–20%, with amplified effects among more indebted borrowers. Related evidence indicates that capacity constraints during the pandemic-era refinancing wave disproportionately affected lower-income households (Agarwal et al., 2024; Fuster et al., 2024). While our findings are consistent with downstream frictions of this kind, we also find that the heterogeneous advertising effect is not confined to the COVID-19 period, suggesting that differential conversion is not solely driven by pandemic-specific labor-market shocks or lender capacity constraints.

Taken together, the evidence suggests that advertising lifts all inattentive boats: it broadens awareness and prompts similar increases in initial search across groups. Yet the gains from this search are disproportionately realized by higher-income households, who are more likely to convert an inquiry into a completed refinancing. The resulting income gap in refinancing thus appears to arise not mainly from differential advertising responsiveness at the search stage, but from differential conversion between application and closing. If the marginal, low-gain applications are screened out in this process, this would simply reflect the institutional discipline of a regulated transaction. However, our results indicate that completions among high-gain borrowers are being impeded, implying that downstream burdens must be disproportionately borne by lower-income households. These frictions amplify income gaps precisely where refinancing is most valuable.

6 Conclusion

Mortgage refinancing is a central channel through which interest-rate cuts translate into higher household cash flow and consumption, yet many borrowers fail to refinance even when doing so would yield substantial savings. We study whether mass-market television advertising by mortgage lenders can help overcome "last-mile" frictions in this transmission mechanism, and we assess both its effectiveness for aggregate pass-through and its distributional consequences.

We document four main findings. First, television advertising increases refinancing, and only where it matters for pass-through: effects are concentrated among borrowers with positive and sizable savings potential. Second, these effects are economically meaningful. Translating advertising-induced refinancing into payment reductions, a one–standard-deviation increase in advertising exposure generates substantial expected savings per household, with particularly large effects in low-rate environments. Third, advertising can have uneven distributional implications. We find little evidence that advertising widens racial disparities in refinancing, but it increases income-based disparities at high incentive levels: higher-income households capture a larger share of the advertising-induced gains when refinancing is most valuable. Fourth, the income gap appears to arise downstream in the refinancing funnel: advertising increases mortgage-related credit inquiries similarly across income groups, but higher-income borrowers are more likely to convert inquiries into completed refinances.

Our findings have two broad policy implications. First, they speak to the usefulness of mass-market advertising as a scalable and relatively low-cost policy lever for strengthening the pass-through of interest-rate declines to household balance sheets. Our results show that lender advertising appears to spur engagement precisely when potential savings are large, generating positive spillovers for monetary transmission even when advertisers are competing for market share. At the same time, evidence of trust-related frictions in targeted outreach (Johnson et al., 2019) suggests that our estimated effects may be a lower bound on what broad communication can achieve: messages delivered by trusted public agencies or nonprofits could plausibly have larger impacts than for-profit lender ads, especially among disadvantaged borrowers. Advertising also offers a scalable tool to policy-makers to deploy geographic targeting (e.g., concentrating media buys in markets where refinancing is unusually sluggish) to bolster aggregate stimulus where and when it is most needed. Yet our results also caution that increasing awareness and engagement is not enough. When downstream frictions exist and impede certain groups more than others, interventions that raise search can translate into unequal realized gains and may widen inequality. Effective communication strategies should therefore be paired with efforts that reduce closing-stage bottlenecks so that households who have the most to gain are also able to capture those gains.

Second, we offer a new lens for evaluating advertising's consequences, an object of first-order importance for policy in domains where advertising-induced mistakes can be very costly. The key question is whether advertising-induced behavior is concentrated among consumers who have the most to gain from taking the advertised action. In many consumer markets, the welfare implications of ad-induced actions are hard to assess because surplus is driven by unobserved preferences. In mortgage refinancing, potential gains are largely observable, allowing us to evaluate whether advertising moves decisions in economically meaningful directions. Reassuringly, despite being untargeted and increasing initial search even among low-gain households, advertising does not translate into excess refinancing when gains are minimal. This result may be largely driven by the disciplining role of regulated, standardized disclosures in the closing process. In this setting, mass-market advertising generates substantial savings per dollar spent. Our work illustrates a novel approach for evaluating the impact of advertising through an incentive-alignment lens. We hope this approach motivates further work on advertising's welfare consequences in other settings, where informational guardrails may be weaker.

References

Agarwal, Sumit, Gene Amromin, Souphala Chomsisengphet, Tim Landvoigt, Tomasz Piskorski, Amit Seru, and Vincent Yao (2015) "Mortgage refinancing, consumer spending, and competition: Evidence from the Home Affordable Refinancing Program," *NBER Working Paper*, 21512.

——— (2023) "Mortgage Refinancing, Consumer Spending, and Competition: Evidence from the Home Affordable Refinance Program," *The Review of Economic Studies*, 90 (2), 499–537.

Agarwal, Sumit, Yongheng Deng, Chenxi Gu, Jia He, Wenlan Qian, and Yuan Ren (2024) "Refinancing inequality during the COVID-19 pandemic," *Journal of Financial and Quantitative Analysis*, 59 (1), 121–152.

Agarwal, Sumit, John C Driscoll, and David I Laibson (2013) "Optimal mortgage refinancing: A closed-form solution," *Journal of Money, Credit and Banking*, 45 (4), 591–622.

Agarwal, Sumit, John Grigsby, Ali Hortacsu, Gregor Mazer, Paolo Pazzona, Maurizio Tumer, and Vincent Wenig (2020) "Searching for approval," *Working Paper*.

Agarwal, Sumit, Richard J Rosen, and Vincent Yao (2016) "Why do borrowers make mortgage refinancing mistakes?" *Management Science*, 62 (12), 3494–3509.

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.

Anenberg, Elliot, Tess Scharlemann, and Eileen Van Straelen (2025) "Borrowing and spending in the money: Debt substitution and the cash-out refinance channel of monetary policy," *American Economic Review*, 115 (11), 3909–3940.

Auclert, Adrien (2019) "Monetary policy and the redistribution channel," *American Economic Review*, 109 (6), 2333–2367.

Bartlett, Robert, Adair Morse, Richard Stanton, and Nancy Wallace (2022) "Consumer-lending discrimination in the FinTech era," *Journal of Financial Economics*, 143 (1), 30–56.

Beraja, Martin, Andreas Fuster, Erik Hurst, and Joseph Vavra (2019) "Regional heterogeneity and the refinancing channel of monetary policy," *Quarterly Journal of Economics*, 134 (1), 109–183.

Berger, David, Konstantin Milbradt, Fabrice Tourre, and Joseph Vavra (2021) "Mortgage prepayment and path-dependent effects of monetary policy," *American Economic Review*, 111 (9), 2829–2878.

——— (2025) "Optimal mortgage refinancing with inattention," *American Economic Review: Insights*, 7 (4), 497–515.

Byrne, David, Robert Kelly, and Conor O'Toole (2023) "Attention, refinancing, and monetary policy pass-through," *Journal of Finance*, 78 (6), 3265–3308.

Calza, Alessandro, Tommaso Monacelli, and Livio Stracca (2013) "Housing finance and monetary policy," *Journal of the European Economic Association*, 11 (suppl_1), 101–122.

Campbell, John Y (2006) "Household finance," *The Journal of Finance*, 61 (4), 1553–1604.

Célérier, Claire and Purnoor Tak (2025) "Finance, advertising, and race," *The Review of Financial Studies*, 38 (11), 3149–3204.

DeFusco, Anthony A and John A Mondragon (2020) "No job, no money, no refi: Frictions to refinancing in a recession," *Journal of Finance*, 75 (5), 2327–2376.

Deng, Yongheng, John M Quigley, and Robert Van Order (2000) "Mortgage terminations, heterogeneity and the exercise of mortgage options," *Econometrica*, 68 (2), 275–308, 10.1111/1468-0262.00118.

Di Maggio, Marco, Amir Kermani, Benjamin J Keys, Tomasz Piskorski, Rodney Ramcharan, Amit Seru, and Vincent Yao (2017) "Interest rate pass-through: Mortgage rates, household consumption, and voluntary deleveraging," *American Economic Review*, 107 (11), 3550–3588.

Di Maggio, Marco, Amir Kermani, and Christopher J Palmer (2020) "How quantitative easing works: Evidence on the refinancing channel," *Review of Economic Studies*, 87 (3), 1498–1528.

Eichenbaum, Martin, Sergio Rebelo, and Arlene Wong (2022) "State-dependent effects of monetary policy: The refinancing channel," *American Economic Review*, 112 (3), 721–761.

Frame, W Scott, Ruidi Huang, Erica Xuewei Jiang, Yeonjoon Lee, Will Shuo Liu, Erik J Mayer, and Adi Sunderam (2025) "The impact of minority representation at mortgage lenders," *The Journal of Finance*, 80 (2), 1209–1260.

Frazier, Nick and Ryan Goodstein (2025) "Supply constraints and failure to refinance," FDIC CFR Working Paper FDIC CFR WP 2023-01, Federal Deposit Insurance Corporation.

Fuster, Andreas, Stephanie H Lo, and Paul S Willen (2024) "The time-varying price of financial intermediation in the mortgage market," *Journal of Finance*, 79 (4), 2553–2601.

Gerardi, Kristopher, Paul S Willen, and David Hao Zhang (2023) "Mortgage prepayment, race, and monetary policy," *Journal of Financial Economics*, 147 (3), 498–524.

Ghosh Dastidar, Samanway, Jean-Pierre Dubé, and Navdeep S Sahni (2023) "Societal spillovers of TV advertising: A natural experiment," *Working Paper*.

Grundl, Serafin and You Suk Kim (2019) "Advertising expensive mortgages," *Journal of Finance*, 74 (5), 2371–2416.

Gurun, Umit G, Gregor Matvos, and Amit Seru (2016) "Advertising expensive mortgages," *Journal of Finance*, 71 (5), 2371–2416.

He, Alex and Tobias J Klein (2023) "Advertising as a reminder: Evidence from the Dutch State Lottery," *Management Science*, 69 (4), 2015–2038.

Honka, Elisabeth, Ali Hortacsu, and Maria Ana Vitorino (2017) "Advertising, consumer awareness, and choice: evidence from the US banking industry," *RAND Journal of Economics*, 48 (3), 611–646.

Hristakeva, Sylvia and Julie Holland Mortimer (2021) "Price dispersion and legacy discounts in the national television advertising market," *Management Science*, 67 (12), 7632–7654.

Hu, Sheng, Jiasun Li, Huan Tang, and Zhi Zhen (2024) "Financial media as a money doctor," *Review of Financial Studies*, 37 (2), 442–489.

Johnson, Eric, 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.

Keys, Benjamin J, Devin G Pope, and Jaren C Pope (2016) "Failure to refinance," *Journal of Financial Economics*, 122 (3), 482–499.

Kim, Donggwan, Zhenling Jiang, and Raphael Thomadsen (2025) "Tv advertising effectiveness with racial minority representation: Evidence from the mortgage market," *The Wharton School Research Paper Forthcoming*.

Kim, You Suk (2024) "Internet access and mortgage refinancing," *Working Paper*.

Li, Keyuan and Yiran Zhang (2024) "Does monetary policy narrow or widen the racial gap in the US housing market?" *Working Paper*.

Maturana, Gonzalo and Jordan Nickerson (2019) "Teachers teaching teachers: The role of workplace peer effects in financial decisions," *Review of Financial Studies*, 32 (10), 3920–3957.

Maxted, Peter, David Laibson, and Benjamin Moll (2025) "Present bias amplifies the household balance-sheet channels of macroeconomic policy," *Quarterly Journal of Economics*, 140 (1), 691–743.

McCartney, W Ben and Avni M Shah (2022) "Household mortgage refinancing decisions are neighbor influenced, especially along racial lines," *Journal of Urban Economics*, 128, 103306.

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.

Shapiro, Bradley T, Günter J Hitsch, and Anna E Tuchman (2021) "TV advertising effectiveness and profitability: Generalizable results from 288 brands," *Econometrica*, 89 (4), 1855–1879.

Wong, Arlene (2019) "Refinancing and the transmission of monetary policy to consumption," *Unpublished manuscript*, 20 (4).

Zhang, David (2023) "Closing costs, refinancing, and inefficiencies in the mortgage market," *Working Paper*.

A Robustness Checks

A.1 Overall Ad Effect Using Splines

Figure A1: Overall Ad Effect Using Splines

Panel (a): Monthly Saving

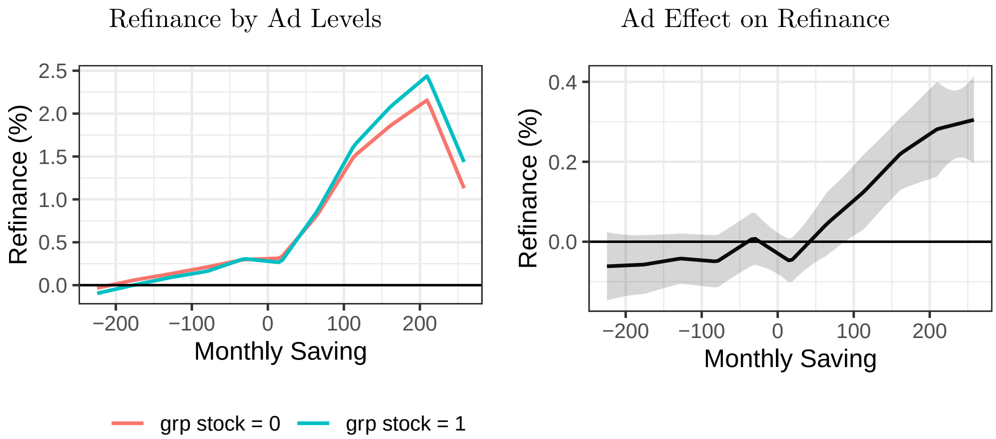

Panel (b): Rate Gap

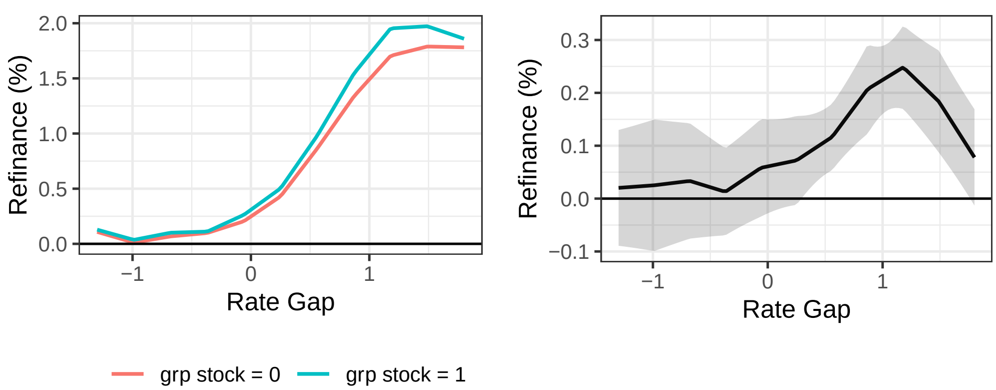

Panel (c): ADL Incentive

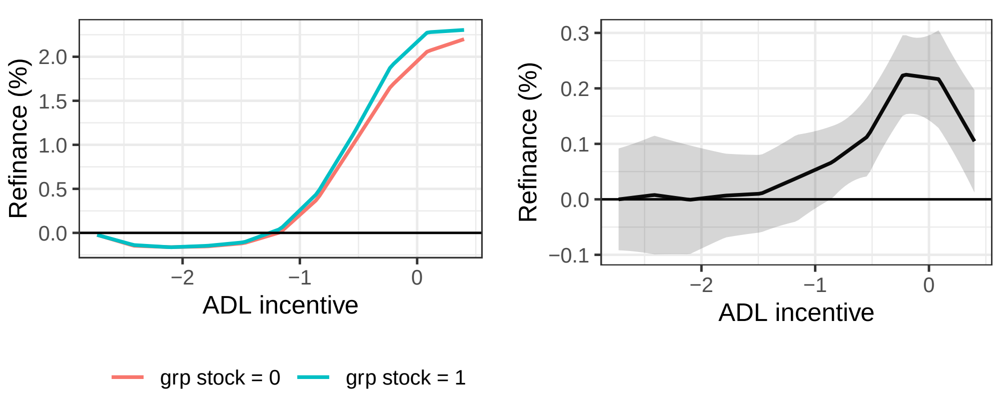

B Results for Credit Inquiry

B.1 Disparities in Credit Inquiries

Smaller income and racial disparity for credit inquiry relative to refinance.

Figure B1: Credit Inquiry Disparity

(a) Income Category (b) Race and Ethnicity

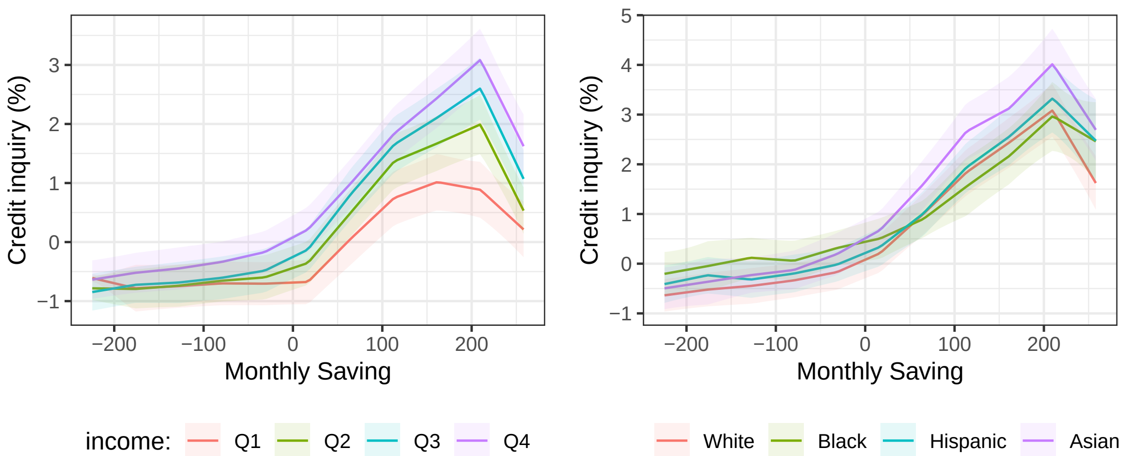

Figure B2: Gap in Credit Inquiry

Panel (a): Monthly Saving

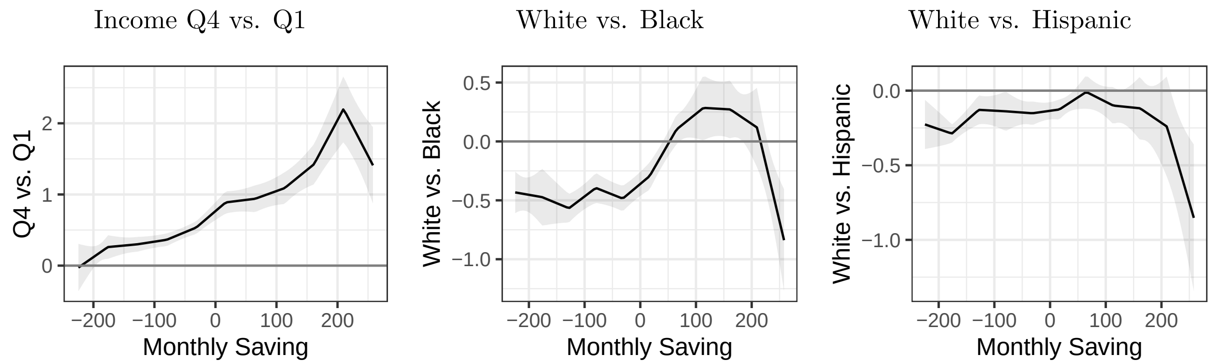

Panel (b): Rate Gap

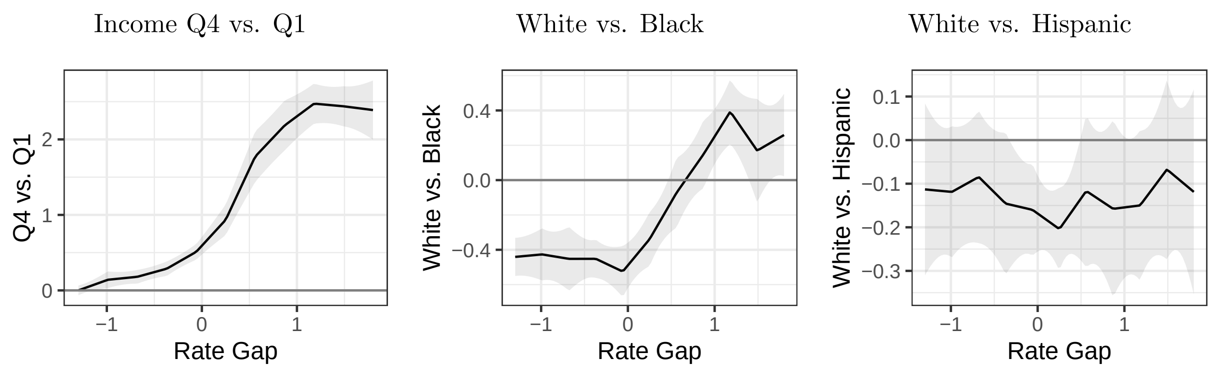

Panel (c): ADL Incentive

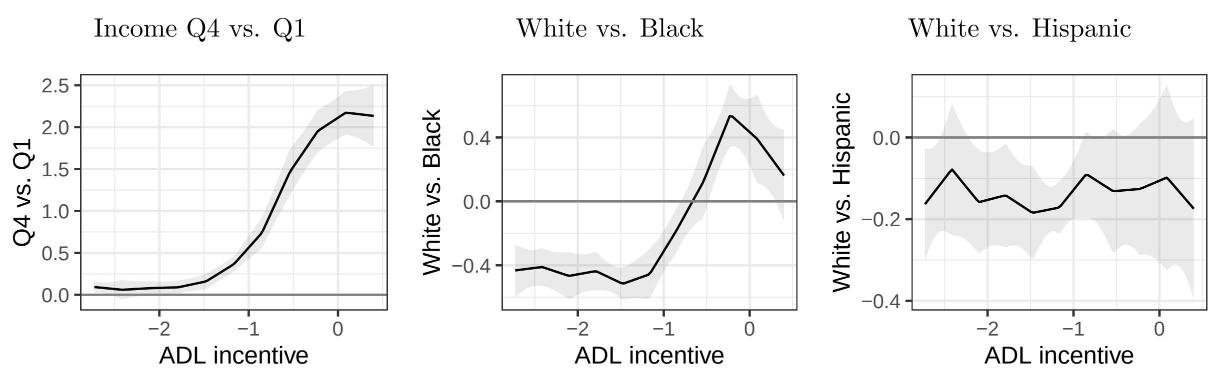

C Regression Results

C.1 Overall Ad Effect

Table C1: Overall Ad Effect

	Monthly Saving		Rate Gap		ADL Incentive	
	Coefficient	(Std. Err.)	Coefficient	(Std. Err.)	Coefficient	(Std. Err.)
β_j: AdStock × Incentive bin						
Decile 1	-0.0564	(0.0389)	0.0285	(0.0622)	0.0105	(0.0521)
Decile 2	-0.0517	(0.0332)	0.0258	(0.0571)	0.0011	(0.0510)
Decile 3	-0.0419	(0.0330)	0.0255	(0.0480)	0.0094	(0.0451)
Decile 4	-0.0225	(0.0318)	0.0375	(0.0437)	0.0065	(0.0369)
Decile 5	-0.0138	(0.0265)	0.0687	(0.0471)	0.0270	(0.0381)
Decile 6	-0.0064	(0.0309)	0.0957**	(0.0306)	0.0523	(0.0344)
Decile 7	0.0862*	(0.0437)	0.1656***	(0.0374)	0.0958**	(0.0319)
Decile 8	0.1745***	(0.0449)	0.2234***	(0.0317)	0.1646***	(0.0353)
Decile 9	0.2231***	(0.0475)	0.2176***	(0.0442)	0.2195***	(0.0388)
Decile 10	0.3044***	(0.0443)	0.1242**	(0.0452)	0.1532**	(0.0423)
α_j: Splines	✓					
$\gamma : D_i$						
Income Q2	0.1414***	(0.0133)	0.2295***	(0.0209)	0.1348***	(0.0145)
Income Q3	0.2482***	(0.0230)	0.3939***	(0.0353)	0.2406***	(0.0256)
Income Q4	0.2785***	(0.0304)	0.4833***	(0.0457)	0.2788***	(0.0347)
Black	-0.1249***	(0.0165)	-0.1351***	(0.0165)	-0.1399***	(0.0162)
Hispanic	-0.1428***	(0.0148)	-0.1424***	(0.0138)	-0.1471***	(0.0148)
Asian	-0.0588*	(0.0239)	0.0035	(0.0238)	-0.0272	(0.0234)
Male	-0.0263***	(0.0068)	-0.0374***	(0.0073)	-0.0243**	(0.0069)
Female	0.1141***	(0.0093)	0.1072***	(0.0093)	0.1126***	(0.0093)
First mortgage	-0.0604***	(0.0074)	-0.0619***	(0.0088)	-0.0555***	(0.0078)
$\lambda : X_{it}$						
LTV>80%	-0.0466*	(0.0202)	-0.0746***	(0.0189)	-0.0918***	(0.0185)
LTV>90%	-0.2298***	(0.0399)	-0.2242***	(0.0405)	-0.2250***	(0.0389)
Delinq within 12 mons	-0.6958***	(0.0594)	-0.7194***	(0.0579)	-0.7587***	(0.0618)
Delinq over 12 mons	-0.2185***	(0.0234)	-0.2251***	(0.0222)	-0.2517***	(0.0238)
Modified within 12 mons	-0.0745	(0.2061)	0.4059**	(0.1294)	0.4186**	(0.1385)
Modified over 12 mons	-0.6707***	(0.1483)	-0.1424	(0.0961)	-0.1683	(0.1005)
Credit score $(650, 750]$	0.0706***	(0.0157)	0.0525**	(0.0164)	0.0368*	(0.0144)
Credit score >750	-0.0498.	(0.0289)	-0.0380	(0.0293)	-0.0470.	(0.0266)
Refinanced Loan	-0.0020	(0.0268)	-0.0452	(0.0267)	-0.0515.	(0.0277)
Loan age	0.1583*	(0.0618)	0.0728	(0.0540)	0.0857	(0.0580)
Loan age^2	-0.1041***	(0.0102)	-0.1058***	(0.0111)	-0.1114***	(0.0115)
DMA Fixed Effects	✓		✓		✓	
YYYYQQ at Origination FE	✓		✓		✓	
Year-Month FE	✓		✓		✓	
Observations	39,104,968		39,104,968		39,104,968	
R^2	0.01280		0.01190		0.01280	

Dependent variable: monthly refinancing indicator (×100). Ad stock is standardized (mean zero, unit standard deviation). β_j coefficients capture the effect of a one-standard-deviation increase in ad stock on the refinancing hazard within each incentive decile. Standard errors double-clustered by DMA and origination year-quarter. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, · $p < 0.1$.

C.2 Heterogeneous Ad Effect

Table C2: Heterogeneous Ad Effect

	Monthly Saving		Rate Gap		ADL Incentive	
	Coefficient	(Std. Err.)	Coefficient	(Std. Err.)	Coefficient	(Std. Err.)
β_{1j}: AdStock × Incentive bin						
Decile 1	-0.0190	(0.0442)	0.0124	(0.0581)	0.0070	(0.0499)
Decile 2	-0.0504	(0.0422)	0.0104	(0.0572)	-0.0032	(0.0495)
Decile 3	0.0004	(0.0368)	0.0124	(0.0456)	0.0010	(0.0466)
Decile 4	-0.0173	(0.0342)	0.0256	(0.0423)	0.0240	(0.0339)
Decile 5	-0.0277	(0.0268)	0.0406	(0.0441)	0.0088	(0.0331)
Decile 6	-0.0287	(0.0290)	0.0372	(0.0324)	0.0074	(0.0291)
Decile 7	0.0633	(0.0400)	0.0204	(0.0299)	0.0498	(0.0345)
Decile 8	0.0929*	(0.0454)	0.0814.	(0.0401)	0.0794*	(0.0337)
Decile 9	0.0032	(0.0354)	0.1254**	(0.0455)	0.0956	(0.0607)
Decile 10	0.0339	(0.0612)	0.0540	(0.0562)	0.0584	(0.0474)
β_{2j}: Income × AdStock × Incentive bin						
Income Q2 x Decile 1	-0.0217	(0.0236)	-0.0095	(0.0106)	-0.0139	(0.0128)
Income Q3 X Decile 1	-0.0151	(0.0266)	-0.0066	(0.0168)	-0.0116	(0.0196)
Income Q4 x Decile 1	-0.0365	(0.0300)	-0.0319	(0.0275)	-0.0368	(0.0251)
Income Q2 x Decile 2	0.0199	(0.0167)	-0.0176*	(0.0076)	-0.0132	(0.0120)
Income Q3 x Decile 2	0.0059	(0.0209)	-0.0059	(0.0118)	-0.0107	(0.0157)
Income Q4 x Decile 2	0.0138	(0.0250)	-0.0107	(0.0183)	-0.0212	(0.0290)
Income Q2 x Decile 3	-0.0284*	(0.0138)	-0.0043	(0.0084)	-0.0029	(0.0108)
Income Q3 x Decile 3	-0.0287.	(0.0156)	-0.0017	(0.0079)	0.0020	(0.0157)
Income Q4 x Decile 3	-0.0408.	(0.0202)	-0.0143	(0.0152)	-0.0119	(0.0233)
Income Q2 x Decile 4	0.0004	(0.0100)	-0.0052	(0.0088)	-0.0188	(0.0136)
Income Q3 x Decile 4	0.0174	(0.0182)	0.0003	(0.0139)	-0.0309*	(0.0136)
Income Q4 x Decile 4	-0.0015	(0.0138)	-0.0264**	(0.0088)	-0.0502**	(0.0157)
Income Q2 x Decile 5	0.0392*	(0.0146)	0.0259.	(0.0141)	0.0271	(0.0170)
Income Q3 x Decile 5	0.0366	(0.0218)	0.0191	(0.0172)	0.0226	(0.0197)
Income Q4 x Decile 5	0.0350	(0.0295)	0.0055	(0.0247)	-0.0228	(0.0214)
Income Q2 x Decile 6	0.0603***	(0.0158)	0.0566***	(0.0095)	0.0541**	(0.0174)
Income Q3 x Decile 6	0.0591**	(0.0178)	0.0574*	(0.0239)	0.0456*	(0.0216)
Income Q4 x Decile 6	0.0902***	(0.0211)	0.1014**	(0.0294)	0.0464	(0.0370)
Income Q2 x Decile 7	0.0536.	(0.0265)	0.1114**	(0.0314)	0.0407.	(0.0219)
Income Q3 x Decile 7	0.0722.	(0.0366)	0.1737***	(0.0429)	0.0639*	(0.0258)
Income Q4 x Decile 7	0.0863*	(0.0368)	0.2176***	(0.0564)	0.1067***	(0.0280)
Income Q2 x Decile 8	0.1046**	(0.0289)	0.1265***	(0.0283)	0.0965**	(0.0299)
Income Q3 x Decile 8	0.1640***	(0.0302)	0.1733***	(0.0346)	0.1115**	(0.0379)
Income Q4 x Decile 8	0.1557***	(0.0411)	0.2015***	(0.0504)	0.1568**	(0.0494)
Income Q2 x Decile 9	0.2083***	(0.0563)	0.1062***	(0.0226)	0.1130***	(0.0295)
Income Q3 x Decile 9	0.2844***	(0.0473)	0.1649***	(0.0295)	0.1931***	(0.0374)
Income Q4 x Decile 9	0.2913***	(0.0459)	0.1378**	(0.0430)	0.1837**	(0.0557)

Continued on next page

	Coefficient	(Std. Err.)	Coefficient	(Std. Err.)	Coefficient	(Std. Err.)
Income Q2 x Decile 10	0.1613***	(0.0352)	0.0608**	(0.0190)	0.0697***	(0.0187)
Income Q3 x Decile 10	0.2787***	(0.0597)	0.1214***	(0.0259)	0.1457***	(0.0280)
Income Q4 x Decile 10	0.3508***	(0.0617)	0.1314**	(0.0365)	0.1601***	(0.0310)
β_{3j}: Race $\times$AdStock $\times$ Incentive bin						
Black x Decile 1	-0.0148	(0.0410)	-0.0232	(0.0329)	-0.0620.	(0.0310)
Hispanic x Decile 1	0.0444	(0.0417)	0.0240	(0.0243)	0.0041	(0.0291)
Asian x Decile 1	-0.0252	(0.0350)	-0.0240	(0.0316)	-0.0251	(0.0270)
Black x Decile 2	-0.0209	(0.0526)	-0.0148	(0.0237)	0.0303	(0.0249)
Hispanic x Decile 2	0.0333	(0.0296)	0.0274	(0.0292)	-0.0049	(0.0244)
Asian x Decile 2	0.0031	(0.0312)	-0.0182	(0.0242)	-0.0239	(0.0187)
Black x Decile 3	0.0076	(0.0384)	-0.0001	(0.0229)	-0.0191	(0.0183)
Hispanic x Decile 3	-0.0167	(0.0317)	-0.0105	(0.0164)	0.0215	(0.0156)
Asian x Decile 3	0.0123	(0.0229)	0.0155	(0.0149)	-0.0071	(0.0225)
Black x Decile 4	-0.0057	(0.0272)	0.0256***	(0.0069)	0.0089	(0.0178)
Hispanic x Decile 4	0.0158	(0.0215)	0.0268	(0.0170)	-0.0166	(0.0137)
Asian x Decile 4	0.0099	(0.0235)	0.0180	(0.0208)	0.0045	(0.0264)
Black x Decile 5	0.0108	(0.0188)	0.0286	(0.0231)	-0.0042	(0.0195)
Hispanic x Decile 5	-0.0217	(0.0219)	-0.0216	(0.0281)	-0.0233	(0.0186)
Asian x Decile 5	-0.0081	(0.0259)	-0.0042	(0.0207)	0.0075	(0.0219)
Black x Decile 6	-0.0592	(0.0380)	-0.0487.	(0.0281)	0.0187	(0.0414)
Hispanic x Decile 6	-0.0470	(0.0281)	-0.0373	(0.0371)	-0.0288	(0.0263)
Asian x Decile 6	-0.0508*	(0.0247)	-0.0023	(0.0226)	-0.0264	(0.0228)
Black x Decile 7	-0.0256	(0.0499)	0.0134	(0.0477)	0.0237	(0.0415)
Hispanic x Decile 7	-0.0687*	(0.0256)	0.0216	(0.0263)	-0.0417	(0.0282)
Asian x Decile 7	-0.0222	(0.0269)	0.0684*	(0.0275)	-0.0012	(0.0236)
Black x Decile 8	-0.0569	(0.0467)	-0.0055	(0.0273)	-0.0175	(0.0267)
Hispanic x Decile 8	-0.0453	(0.0319)	0.0006	(0.0137)	-0.0143	(0.0171)
Asian x Decile 8	-0.0402	(0.0476)	0.0744.	(0.0408)	0.0361	(0.0325)
Black x Decile 9	0.0102	(0.0550)	0.0783	(0.0522)	0.0381	(0.0457)
Hispanic x Decile 9	0.0015	(0.0325)	-0.0455**	(0.0163)	-0.0004	(0.0177)
Asian x Decile 9	0.0816*	(0.0391)	-0.0534	(0.0347)	0.0382	(0.0493)
Black x Decile 10	0.0246	(0.0462)	0.0430	(0.0636)	0.0773*	(0.0373)
Hispanic x Decile 10	-0.0264	(0.0348)	-0.0156	(0.0265)	0.0193	(0.0237)
Asian x Decile 10	0.0669*	(0.0296)	0.0149	(0.0393)	0.0109	(0.0415)
α_{1j}: Splines	✓					
α_{2j}: Income $\times$ Splines	✓					
α_{3j}: Race $\times$ Splines	✓					
$\gamma : D_i$						
Income Q2	-0.0984**	(0.0350)	-0.0653***	(0.0146)	-0.0378*	(0.0164)
Income Q3	-0.1494***	(0.0393)	-0.1048***	(0.0153)	-0.0562***	(0.0126)
Income Q4	-0.1509***	(0.0341)	-0.1534***	(0.0142)	-0.0847**	(0.0254)
Black	0.1014***	(0.0258)	0.0746**	(0.0261)	0.0554	(0.0334)
Hispanic	0.0142	(0.0382)	-0.0215	(0.0330)	0.0220	(0.0259)
Asian	-0.0803**	(0.0282)	-0.1289***	(0.0259)	-0.0847*	(0.0319)

Continued on next page

	Coefficient	(Std. Err.)	Coefficient	(Std. Err.)	Coefficient	(Std. Err.)
Male	-0.0313**	(0.0088)	-0.0393***	(0.0074)	-0.0278***	(0.0070)
Female	0.1180***	(0.0111)	0.1097***	(0.0094)	0.1179***	(0.0095)
First mortgage	-0.0617***	(0.0090)	-0.0674***	(0.0084)	-0.0603***	(0.0071)
$\lambda : X_{it}$						
LTV>80%	-0.0550*	(0.0209)	-0.0769***	(0.0189)	-0.0902***	(0.0183)
LTV>90%	-0.2264***	(0.0409)	-0.2464***	(0.0372)	-0.2500***	(0.0369)
Delinq within 12 mons	-0.6975***	(0.0618)	-0.7249***	(0.0599)	-0.7561***	(0.0628)
Delinq over 12 mons	-0.2115***	(0.0260)	-0.2210***	(0.0221)	-0.2406***	(0.0234)
Modified within 12 mons	0.0010	(0.1878)	0.3739**	(0.1299)	0.3743**	(0.1344)
Modified over 12 mons	-0.5324***	(0.1339)	-0.1458	(0.0972)	-0.1705.	(0.0999)
Credit score (650,750]	0.0985***	(0.0174)	0.0598**	(0.0170)	0.0351*	(0.0145)
Credit score > 750	-0.0118	(0.0303)	-0.0281	(0.0298)	-0.0506.	(0.0267)
Refinanced Loan	0.0063	(0.0289)	-0.0412	(0.0266)	-0.0486.	(0.0271)
Loan age	0.1166.	(0.0598)	0.0609	(0.0532)	0.0702	(0.0561)
Loan age^2	-0.1035***	(0.0119)	-0.1058***	(0.0112)	-0.1050***	(0.0095)
DMA Fixed Effects	✓		✓		✓	
YYYYQQ at Origination FE	✓		✓		✓	
Year-Month FE	✓		✓		✓	
Observations	39,104,968		39,104,968		39,104,968	
R^2	0.01314		0.01277		0.01342	

Notes: Dependent variable: monthly refinancing indicator (×100). β_{1j}: baseline advertising effect by incentive decile. β_{2j}: differential advertising effect by income quartile (relative to Q1). β_{3j}: differential advertising effect by race (relative to White). Standard errors double-clustered by DMA and origination year-quarter. *** $p < 0.001$, ** $p < 0.01$, * $p < 0.05$, $^{\cdot}$ $p < 0.1$.

D Loan-level Panel Data Construction

Since addressing our research questions requires information that is not available in any single data source to our best knowledge, we construct a novel loan-level dataset by combining three sources of mortgage data: (a) Home Mortgage Disclosure Act (HMDA) data, (b) Freddie Mac (FHLMC) Single-Family Loan-Level data, and (c) consumer credit bureau data from Equifax, Inc. HMDA provides the most comprehensive coverage of mortgage originations, including rich borrower and loan characteristics such as race and income. Freddie Mac data track the subsequent performance of these loans, allowing us to observe detailed information on monthly payment history, delinquency status, and prepayment events. Finally, Equifax data allow us to distinguish whether a prepayment reflects refinancing or moving, while also providing time-varying credit scores. Taken together, this tri-merged dataset allows us to follow each loan from origination through termination and to identify the underlying reasons for loan prepayment (if any).

Although these datasets do not share a common unique identifier, they have substantial overlap in loan-level characteristics, such as loan size, origination timing, loan purpose, and property location, which we use to perform one-to-one matches across the three sources. Specifically, we merge the data in two sequential steps: first matching HMDA records to Freddie Mac loans (hereafter, HMDA–FHLMC), and then matching the resulting matched sample to Equifax credit records (hereafter, HMDA–FHLMC–EFX).

The remainder of this appendix documents the pre-merge data cleaning procedures, the matching algorithms used at each stage, and the resulting one-to-one match rates to demonstrate the quality of the merged dataset. We then describe the final sample construction and post-merge cleaning steps applied prior to analysis.

D.1 Loan-Level Matches between HMDA and Freddie Mac

While both the HMDA and Freddie Mac datasets provide detailed information with significant overlap, including loan size, origination timing, loan purpose, and property location, HMDA contains more comprehensive information at origination, including borrower race and income, whereas Freddie Mac provides richer data on loan performance over time. This difference motivates merging the two datasets.

We merge the HMDA and Freddie Mac data for the years 2010 through 2017 (inclusive). Following prior literature (e.g., Bartlett et al., 2022), we restrict both samples to first-lien, site-built, owner-occupied loans originated either for home purchase or refinancing. We further require that property location is reported (i.e., census tract, county code, and state code). Finally, for HMDA records only, we additionally restrict the sample to loans reported as sold to Freddie Mac. Table D1 reports the resulting number of loans in HMDA and Freddie Mac. The counts follow similar time trends, but the HMDA counts are consistently smaller, likely due to HMDA's earlier reporting deadline (March 1 of the following year), which leaves limited time for some loans (especially, those originated late in the year) to be sold to Freddie Mac.

Before discussing the data merging procedures, we first examine the similarity of key variables used for matching. Figure D1 compares the distributions of loan size, refinance share at origination (defined as the fraction of newly originated loans that are refinances), and loan origination share

Table D1: Number of loans and match rate by year

	Pre-Matching		Post-Matching	
Year	HMDA	FHLMC	Matched	Match rate*
2010	1,318,009	1,663,806	405,569	30.77%
2011	1,085,586	1,247,138	380,744	35.07%
2012	1,598,836	1,811,897	438,844	27.45%
2013	1,474,392	1,588,320	550,754	37.35%
2014	881,841	986,236	429,482	48.70%
2015	1,139,828	1,294,937	509,648	44.71%
2016	1,309,932	1,493,619	573,451	43.78%
2017	1,083,642	1,191,765	508,042	46.88%
Total	9,892,066	11,277,718	3,796,534	38.39%

Notes: * is calcualted as $\frac{\text{Number of matched loans}}{\text{Number of HMDA loans}}$.

by state and by MSA (based on loan counts). The distributions of loan size, refinance share, and loan share by state are closely aligned across the two datasets. In contrast, we observe some discrepancies in the loan share by MSA. Further inspection suggests two explanations. First, Freddie Mac contains more loans located in non-MSA regions, as some rural areas are not assigned an MSA code. Second, although infrequent, MSA codes can change over time, and such changes are reflected in the Freddie Mac data but not in HMDA, leading to potential mismatches. (We describe how we address this issue in the merging algorithm below.) Overall, these descriptive patterns provide supporting evidence that the remaining loans are suitable for one-to-one matching.

Figure D1: Comparison of key loan characteristics between HMDA and Freddie Mac

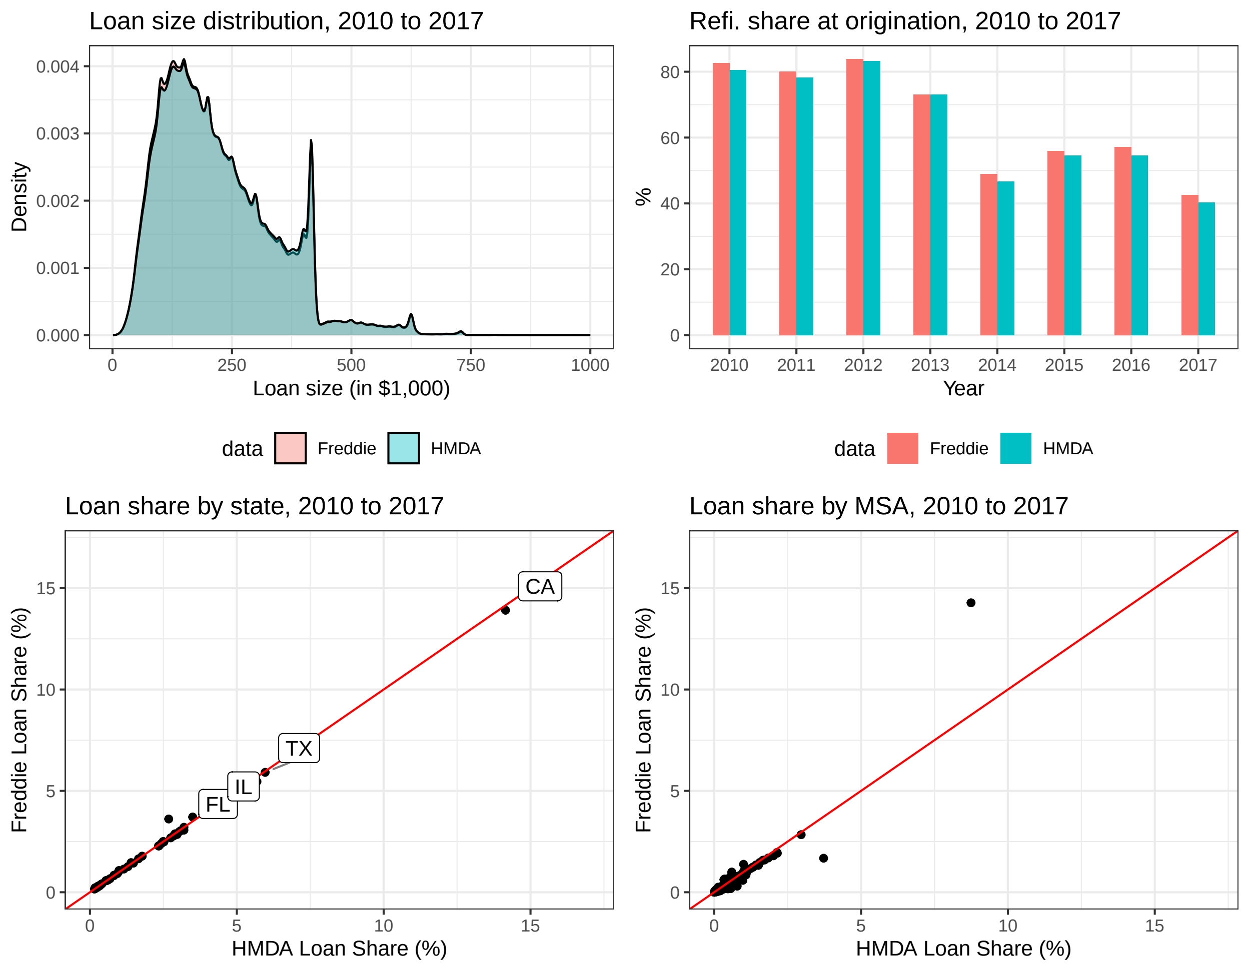

Next, we compare market shares by lender, separately for home purchase and refinance loans at origination. While HMDA reports the originating lender for all loans, regardless of origination volume, Freddie Mac reports lenders only if they account for at least 1% of total transfer volume (in dollars) in a given quarter. As a result, lender-level comparisons across the two datasets are possible only for the subset of lenders that meet this 1% threshold. Figure D2 reports lender-level market shares for home purchase and refinance loans in 2016.[18] Market shares for lenders in this selected group are highly similar across the two datasets, suggesting that the two samples are comparable prior to matching.

[18]We focus on 2016 because the set of lenders meeting the 1% market-share threshold varies across years. However, we find that within-lender market shares are highly similar across all sample years.

Figure D2: Market share by lender

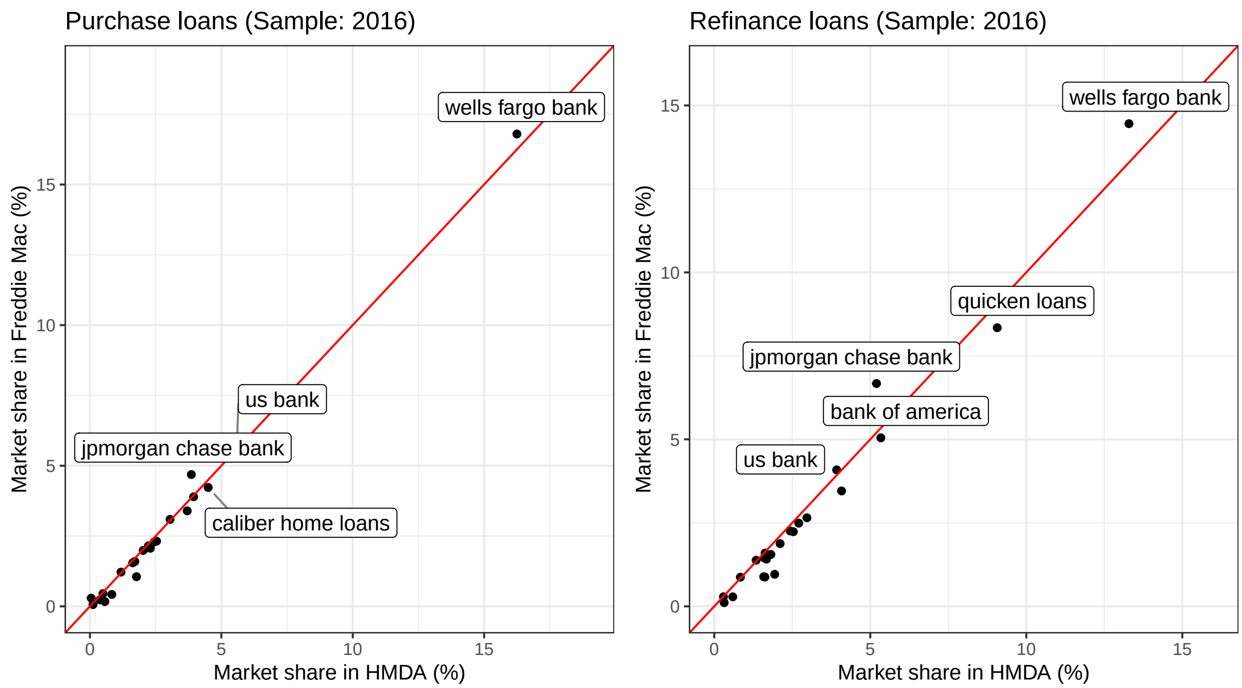

Given the substantial overlap in key variables, and the evidence that their distributions are highly similar across datasets (with the exception of MSA and lender names), we implement a waterfall-style matching procedure as follows:

- Step 1: Match using full geographic information
  - Step 1(a): Match on year of origination, loan size (in \$1,000s), loan purpose, 3-digit ZIP code, MSA, state, and lender.
  - Step 1(b): Remove one-to-one matched loans from Step 1(a), then match remaining loans using all variables above except for lender.
- Step 2: Relax geographic conditions
  - Step 2(a): Repeat Step 1(a), excluding MSA from the matching criteria.
  - Step 2(b): Remove one-to-one matched loans from Step 2(a), then repeat Step 1(b), excluding both MSA and lender.

We report the number of one-to-one matched loans and match rates by year in Table D1. Match rates are lower in the earlier four years (typically around 30%) but increase to about 45% in the later four years.

To address potential concerns that the merging procedure may introduce selection bias, we examine whether matched loans differ systematically from either the full sample or the unmatched loans. Figure D3 presents the results. First, loan sizes are broadly similar across groups, although match rates decline at the conforming loan limits (indicated by vertical lines), where one-to-one matching is less common. Second, FICO score distributions are nearly identical across samples. Third, match rates tend to decrease in counties with larger populations. Lastly, lenders whose identities are disclosed in both HMDA and Freddie Mac (i.e., those meeting the 1% market-share threshold) are more likely to be matched, resulting in slightly higher lender-level market shares among matched loans. Overall, despite these modest differences, the matched sample appears broadly representative of the underlying borrower population, suggesting that any selection bias is likely limited.

Figure D3: Comparison of key variables across matched and unmatched/full samples

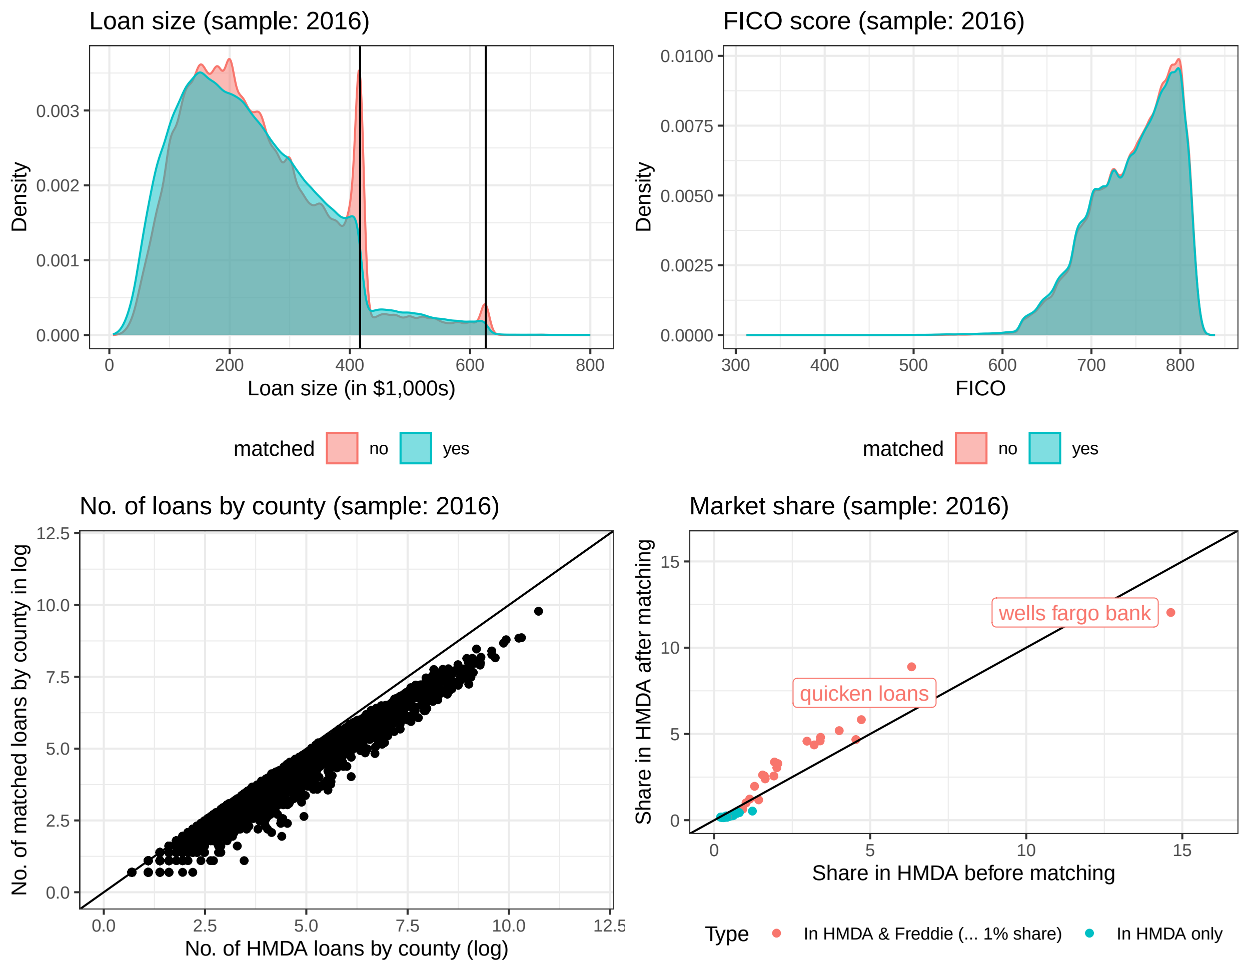

D.2 Loan-Level Matches between HMDA–Freddie Mac and Equifax

While the matched HMDA–FHLMC data provide detailed information on loan and borrower characteristics (primarily at origination) and subsequent loan performance (including prepayments), they do not allow us to identify the reason for prepayment, which is central to our study. To determine whether a prepayment is due to refinancing, we merge the matched HMDA–FHLMC data with consumer credit bureau data from Equifax, Inc., which contain information on almost all consumer credit products held by each individual, along with updated credit scores (VantageScore) and credit inquiries (including mortgage-related inquiries).

Following this merge, we classify a prepayment as a refinancing event if the borrower (identified via an encrypted SSN) obtains a new mortgage shortly after the original loan closes and both loans are associated with the same ZIP code; otherwise, we classify the prepayment as non-refinancing.[19]

We also obtain updated credit scores and mortgage-related consumer credit inquiries for the matched loans. For credit scores, we use the lower of the two scores if a loan has joint borrowers. For credit inquiries, although lenders typically pull consumer credit reports from all three major credit bureaus, some inquiries can be misclassified as non–mortgage inquiries. To address this measurement issue, we identify borrowers whorefinance but have no recorded mortgage inquiry within the three months preceding refinancing. For these cases, we impute a single inquiry dated one month prior to refinancing.

Before merging, we select all conventional loans reported as sold to Freddie Mac in the Equifax data. Figure D4 shows that the number of new loans observed in Equifax closely aligns with the unmatched full sample of Freddie Mac loans, indicating very similar coverage across the two datasets.

Figure D4: Comparison of new loan counts between Freddie Mac and Equifax

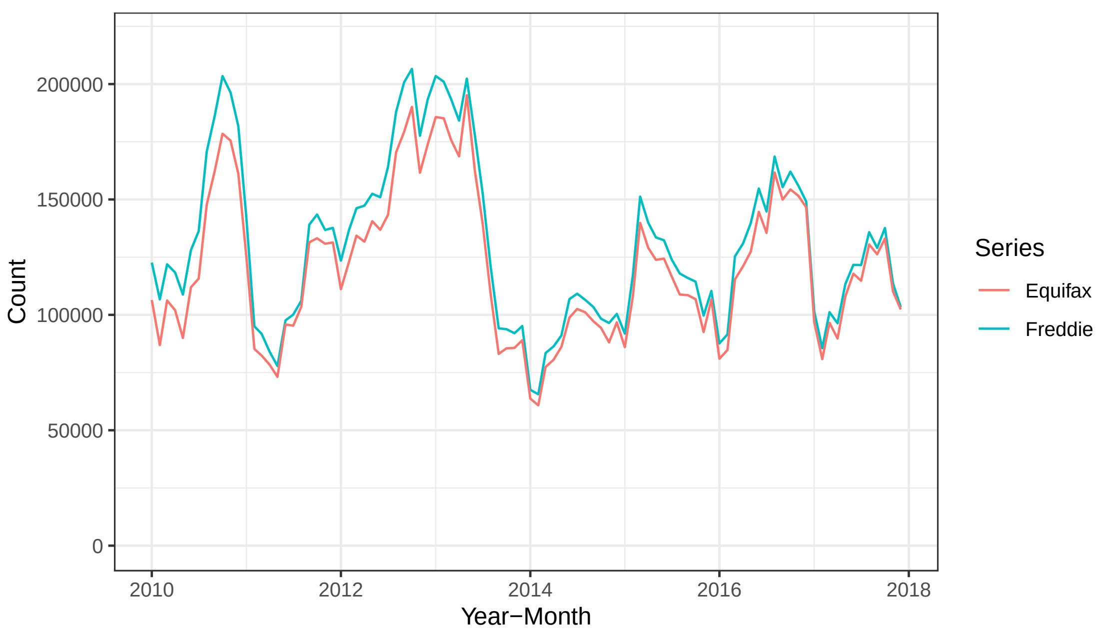

We also find that loan performance is highly consistent across the two datasets. As an illustrative example, Figure D5 plots the monthly count of prepaid loans for the 2016 origination cohort. The

[19]To account for reporting delays, we allow up to a three-month window between the closing of the original loan and the origination of the new loan.

two datasets show nearly identical patterns, including sharp increases during 2020–2021, which coincide with the historically low interest rate environment.[20]

Figure D5: Monthly prepayments among loans originated in 2016

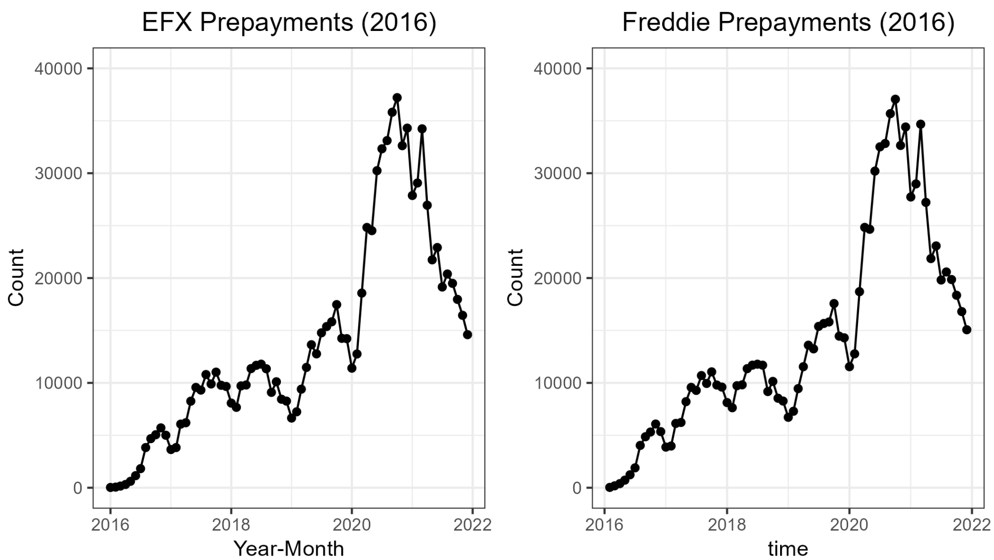

Having established the comparability of the Equifax and Freddie Mac data, we next apply a waterfall-style matching procedure using the set of overlapping variables between the HMDA–FHLMC and Equifax datasets:

- Step 1: Match loans based on loan size (in \$1,000s), 3-digit ZIP code, state, origination year–month, closing year–month (if applicable), and single/joint account dummy.
- Step 2: For remaining unmatched loans, allow for a modest reporting delay:
  - Match on loan size, 3-digit ZIP, state, and origination year–month, singe/joint account dummy; and
  - Allow the closing year–month to differ by up to one month.

Table D2 reports the number of matched loans and the corresponding match rates by year, along with the initial number of loans in the HMDA–FHLMC sample. Match rates are generally high, ranging from 70.5% to 81.8%.

[20]We observe similar patterns in other origination cohorts. We focus on 2016, rather than pooling across years, because prepayment timing differs across cohorts and aggregation would mask these dynamics.

Table D2: Match rate between HMDA-FHLMC and Equifax (EFX) by year

Year	HMDA-FHLMC	EFX-HMDA-FHLMC	Match rate
2010	405,569	285,872	70.5%
2011	380,744	281,106	73.8%
2012	438,844	318,791	72.6%
2013	550,754	410,657	74.6%
2014	429,482	332,910	77.5%
2015	509,648	395,327	77.6%
2016	573,451	439,157	76.6%
2017	508,042	415,774	81.8%

D.3 Panel data construction

So far, we have described two loan-level matching procedures, which are cross-sectional, whereas our analysis ultimately relies on panel data. To this end, we construct a panel dataset that tracks each loan in the tri-merged sample on a monthly basis until either prepayment—whether voluntary or involuntary, along with its reason—or the end of the sample period.

We begin by applying additional sample restrictions to the HMDA–FHLMC data. Importantly, the loan-level matches between the HMDA–FHLMC and Equifax data described above are conducted prior to applying these restrictions, so as not to exclude potentially matchable loans. Specifically, we restrict the HMDA–FHLMC sample to loans that satisfy the following criteria: (a) race/ethnicity information is populated; (b) gender information is populated; (c) FICO score, loan-to-value ratio (LTV), and income are populated; (d) the loan is a standard 30-year fixed-rate mortgage; (e) the loan does not experience foreclosure by the end of the sample period; and (f) the loan is originated within the top 129 DMAs, where advertising exposure is measured more reliably. Table D3 reports the number of HMDA–FHLMC loans in the full sample and the number remaining after each of these filters is applied.

Table D3: Number of loans in the HMDA-FHLMC panel data by year

Year	Matched	Race filled	Gender filled	FICO/LTV/Income	30y FRM	No Foreclosure	Top DMAs
2010	405,569	282,285	282,030	279,268	186,562	182,955	166,159
2011	380,744	260,810	260,612	248,291	148,562	145,867	133,992
2012	438,844	310,125	309,908	299,687	193,372	190,277	174,307
2013	550,754	383,873	383,521	376,198	246,609	243,269	224,880
2014	429,482	303,593	303,284	299,624	217,958	215,857	200,024
2015	509,648	356,064	355,715	344,195	247,591	246,448	228,664
2016	573,451	400,393	400,072	390,800	287,769	286,678	265,966
2017	508,042	354,524	354,160	346,842	269,335	268,026	248,573

Next, we select the subset of these filtered HMDA–FHLMC loans that are successfully matched to Equifax in order to determine whether a prepayment reflects refinancing or moving. Table D4 the total number of one-to-one matched loans between HMDA and Freddie Mac that meet our selection criteria (corresponding to the last column in Table D3), the number of these loans successfully matched to Equifax, the number of loans that remain active at the start of the sample period (January 2016), and the number of loans with an outstanding balance above $60,000, since loans with very small balances are generally difficult to refinance profitably given closing costs.

Table D4: Number of loans in the EFX-HMDA-FHLMC panel data by year

Year	HMDA-FHLMC	Matched with EFX	Active 2016-01	Balance above $60k
2010	166,159	113,607	44,202	38,742
2011	133,992	96,904	48,030	43,574
2012	174,307	124,674	96,497	89,874
2013	224,880	168,553	137,284	129,377
2014	200,024	154,921	127,530	121,320
2015	228,664	176,522	171,103	165,114
2016	265,966	200,907	198,816	193,287
2017	248,573	202,969	200,593	195,587

Finally, to focus on prepayments due to refinancing rather than moving, we allow loans that are ultimately prepaid due to moving to exit the panel, beginning six months prior to the move. We assume that borrowers remain potential refinancers until six months before the moving date, but once a move is anticipated, they exit the pool of potential refinancing borrowers.

Our final panel contains 39,104,968 loan-month observations covering 944,485 loans. Table 1 presents key loan and borrower characteristics at origination, followed by refinancing and credit inquiry behavior over time, time-varying characteristics that affect refinancing eligibility (i.e., key underwriting variables that determine whether a borrower can refinance), and returns to refinancing (i.e., the potential gains from refinancing). The updated loan-to-value (LTV) ratio is calculated as the outstanding loan balance at the beginning of a given month divided by the estimated current home value, multiplied by 100. We estimate the current home value by multiplying the original home price by the change in the Zillow Home Value Index at the county level since the year and month of origination. We describe the construction of advertising GRPs in Appendix E.

E TV Mortgage Advertising GRP Construction

For our advertising exposure measures, we use Nielsen Ad Intel data, which cover television advertising occurrences and impressions at the ad-airing level across four types of ads: Cable, Network, Syndicated, and Spot. Ads in the first three categories (Cable, Network, and Syndicated) are purchased and aired nationally, whereas Spot TV ads are sold and aired locally at the DMA (Designated Market Area) level. Although Network and Syndicated ads are nationally purchased, the data typically include DMA-level viewership measures for these ads. In contrast, such locally measured viewership is not available for Cable ads.

Following Shapiro et al. (2021) and their replication code, we construct local, DMA-level viewership measures for Network and Syndicated TV ads. We then define gross rating points (GRPs) for each ad airing as the number of viewers divided by the total number of TV viewers in the DMA, multiplied by 100. GRPs are calculated similarly for Spot TV ads. For Cable TV ads, we use the total number of TV viewers nationwide as the denominator.

Before aggregating the data to the appropriate level, we exclude a set of lenders and ads that correspond to reverse mortgage lenders, VA-specialized lenders, and forbearance firms, based on advertiser and product names. Specially, we automate this exclusion process as much as possible by removing advertisers or ads containing keywords such as "reverse," "senior," "foreclosure," "modification." In addition, we manually exclude a small number of large, well-known lenders in these categories (e.g., American Advisors Group for reverse mortgages and Veterans United for VA lending).

For the remaining ads, we aggregate GRPs to the DMA–year–month level across all ads to obtain local GRPs for mortgage advertising. We then add the GRP measures from Cable ads uniformly across all DMAs. As a result, variation in advertising exposure across DMAs and over time comes from differences in Spot TV advertising purchases and variation in local viewership of Network and Syndicated TV ads.

Figure E1 presents average monthly ad GRPs by year and month, along with the corresponding 30-year fixed mortgage rates from Freddie Mac's Primary Mortgage Market Survey. The average monthly ad GRP is calculated as the mean of monthly GRPs across the top 129 DMAs. The two series show a negative correlation. Given that TV advertising placements are often purchased well in advance of airing to obtain price discounts (Hristakeva and Mortimer, 2021), this pattern suggests that lenders increase advertising commitments when interest rates are expected to decline and refinancing demand is anticipated to rise, although it is also possible that some advertising placements are purchased closer to the airing date, as realized interest rates turn out to be lower than previously expected.

Figure E1: Monthly ad GRP and 30-year fixed rate

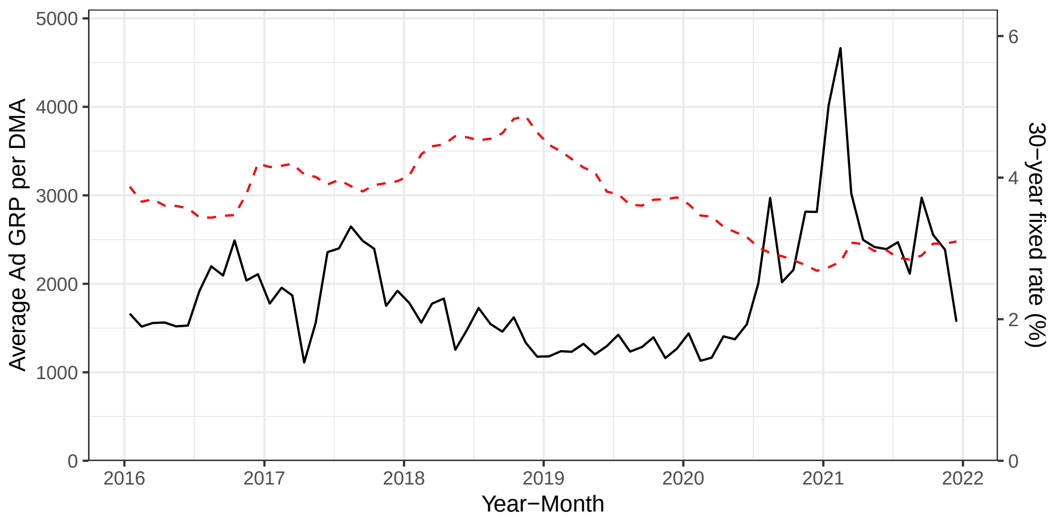

Notes: Solid and dashed lines represent ad GRP and 30-year fixed mortgage rates, respectively.

Although limited to cable television ads, our data provide detailed information on viewership by demographic characteristics, including race and income, which allows us to examine whether advertising exposure evolves similarly across different groups. Figure E2 presents monthly advertising GRPs by race (specifically, White and Black viewers) in Panel (A) and by income category in Panel (B). While we observe some differences in exposure levels, with Black and lower-income consumers exposed to more ads, the time-series patterns are broadly comparable across both racial and income groups.

Figure E2: Cable TV Advertising Exposure Across Racial and Income Groups

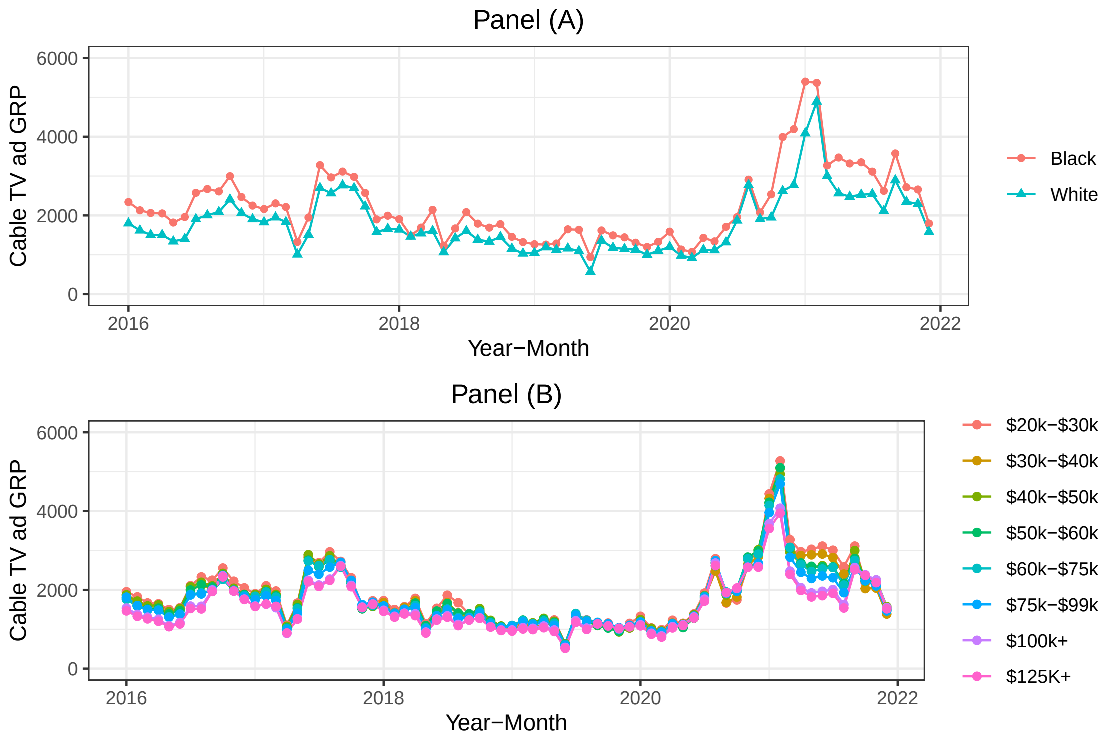

Next, we define the measure of advertising exposure used in our analyses. Let $d(i)$ denote the DMA in which consumer i resides, and let t denote the year-month. We construct advertising exposure as a stock variable:

$$AdStock_{td(i)} = \sum_{\tau=t-12}^{\tau=t-1} \delta^{t-\tau-1} Ad_{\tau d(i)},$$

where $Ad_{\tau d(i)}$ is the total GRP (including both national and local ads) in month τ and DMA $d(i)$. We consider the past 12 months of ad GRPs, starting with the month immediately preceding t, to account for the fact that there is typically a lag of 30 to 45 days between loan application, which occurs close to ad exposure, and loan closing for conventional loans. Following Shapiro et al. (2021), we set $\delta = 0.9$.

We present summary statistics for the ad GRP stock, along with ad GRP flow, at the bottom of Table 1. These statistics indicate substantial variation.

However, for identification, what matters is not the raw variation in advertising exposure, but the residual variation in ad GRP stock after controlling for the rich set of fixed effects used in our analysis. Accordingly, following common practice in the TV advertising literature (e.g., Shapiro et al., 2021), we examine the distribution of residual ad GRP stock after controlling for DMA and year–month fixed effects. Figure E3 presents this residual distribution, which shows that substantial variation remains. We further quantify this variation using the coefficient of variation, defined as the standard deviation of the residual ad GRP stock divided by the unconditional mean of the ad GRP stock. This measure captures the extent of unexplained variation relative to the overall level of advertising exposure. We find a coefficient of variation of 0.136, indicating sizable residual dispersion relative to the mean.

Figure E3: Residual Variation in Ad GRP Stock

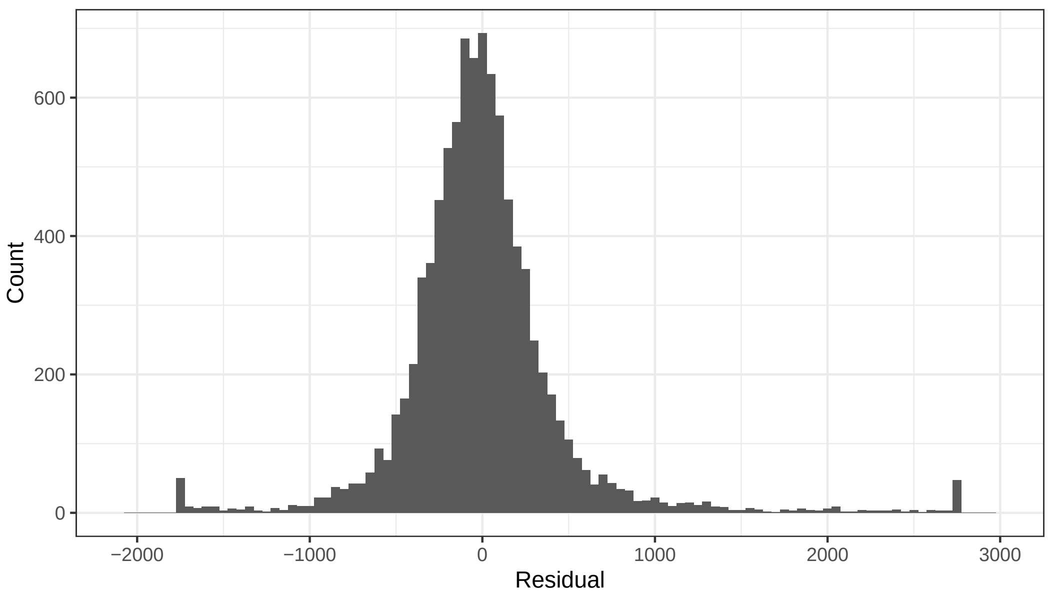

Notes: Residuals are winsorized at the 0.5th and 99.5th percentile values.

F Predicting Counterfactual Mortgage Rates

In this appendix, we describe how we predict the counterfactual mortgage rate that a borrower would obtain if they were to refinance at any point after the initial loan origination.

As described in Section 3.3, we construct the counterfactual mortgage rate using a two-stage approach. In the first stage, we estimate the relationship between observed mortgage rates (i.e., rates actually received by borrowers) and loan and borrower characteristics, following a procedure similar to Bartlett et al. (2022) that exploits the institutional structure of mortgage pricing for conforming loans. Specifically, pricing for conforming loans is largely determined by a base funding rate and a schedule of loan-level price adjustments (LLPAs) that depend on default risk, determined by discrete FICO score and loan-to-value (LTV) buckets. As shown in Figure F1, lenders add rate adjustments, determined by FICO score and LTV ranges, to the base funding rate. At the same time, lenders retain some discretion in setting final rates, for example by passing through marketing or advertising costs to borrowers. In the second stage, we predict counterfactual mortgage rates by plugging updated loan characteristics into the estimated relationship from the first stage.

Figure F1: Loan-Level Price Adjustment (LLPA) Matrix

Credit Score	LTV Range									
	Applicable for all loans with terms greater than 15 years									
	≤ 30.00%	30.01 – 60.00%	60.01 – 70.00%	70.01 – 75.00%	75.01 – 80.00%	80.01 – 85.00%	85.01 – 90.00%	90.01 – 95.00%	>95.00%	SFC
≥ = 780	0.000%	0.000%	0.000%	0.125%	0.500%	0.625%	0.500%	0.375%	0.375%	007
760 – 779	0.000%	0.000%	0.125%	0.375%	0.875%	1.000%	0.750%	0.625%	0.625%	007
740 – 759	0.000%	0.000%	0.250%	0.750%	1.125%	1.375%	1.125%	1.000%	1.000%	007
720 – 739	0.000%	0.000%	0.500%	1.000%	1.625%	1.750%	1.500%	1.250%	1.250%	007
700 – 719	0.000%	0.000%	0.625%	1.250%	1.875%	2.125%	1.750%	1.625%	1.625%	007
680 – 699	0.000%	0.000%	0.875%	1.625%	2.250%	2.500%	2.125%	1.750%	1.750%	007
660 – 679	0.000%	0.125 %	1.125%	1.875%	2.500%	3.000%	2.375%	2.125%	2.125%	007
640 - 659	0.000%	0.250%	1.375%	2.125%	2.875%	3.375%	2.875%	2.500%	2.500%	007
≤ 639	0.000%	0.375%	1.750%	2.500%	3.500%	3.875%	3.625%	2.500%	2.500%	007

Notes: The LLPA Matrix was obtained from Fannie Mae on Jan. 26, 2026.

Let i index loans and t index year-months. In the first stage, we estimate the following regression:

$$r_{it} = \delta_{t \times LLPA(FICO_i, LTV_i)} + \mu_{State_i \cdot Zip_i} + \mu_{SizeDecile_{it}} + \epsilon_{it},$$

where the term $\delta_{t \times LLPA(FICO_i, LTV_i)}$ denotes fixed effects for the interaction between year-month and the loan-level pricing adjustment (LLPA) cell defined by the borrower's FICO score and LTV ranges, as in Figure F1. The term $\mu_{State_i \cdot Zip_i}$ denotes fixed effects for the interaction between state and 3-digit ZIP code, while $\mu_{SizeDecile_{it}}$ denotes loan size decile fixed effects.

The estimation sample consists of full (unmatched) Freddie Mac loans that refinance between October 2015 and December 2021.

We report model performance in Panel A of Table F1. Column (1) shows that LLPA cells interacted with year–month fixed effects alone explain approximately 84% of the variation in mortgage rates. Columns (2) and (3) show that adding state fixed effects or state-by-ZIP fixed effects yields a modest increase in R^2, suggesting that geographic controls provide additional, but limited, predictive power.

As an alternative first-stage specification, we also consider a model that estimates the relationship between mortgage rates actually received by borrowers and loan- and borrower-level observables, including race, gender, and income. Specifically, we estimate the following regression:

$$r_{it} = \beta_{Race_i} + \beta_{Gender_i} + \beta_{IncDecile_{it}} + \delta_{t \times LLPA(FICO_i, LTV_i)} + \mu_{State_i \cdot Zip_i} + \mu_{SizeDecile_{it}} + \epsilon_{it},$$

where β_{Race_i}, β_{Gender_i}, and $\beta_{IncDecile_{it}}$ denote coefficients on the borrower's race, gender, and income decile, respectively. All other variables are defined as in the previous specification.

Because these additional demographic variables are available only in the HMDA data, we estimate this regression using the merged HMDA–Freddie Mac sample, covering our sample period through December 2021.[21]

Panel B of Table F1 reports the coefficients on the demographic variables, along with the corresponding R^2 values. Column (1) shows that LLPA cells interacted with year–month fixed effects and state-by-ZIP fixed effects explain about 83.1% of the variation in mortgage rates, which is very similar to Column (3) of Panel A. Column (2) shows that adding racial information increases R^2 only marginally, while Column (3) shows that additionally controlling for gender and income yields little further improvement in explanatory power. Although demographic variables contribute little to overall model fit, we find that Black borrowers pay mortgage rates that are, on average, 1.4 basis points (0.014%) higher than comparable White borrowers.

Having estimated the two fixed-effects regression models, with and without demographic controls, we use the estimated coefficients to predict counterfactual mortgage rates for each loan in our sample at each year–month, up to the point at which the loan exits the dataset, either through prepayment or at the end of the sample period. Specifically, we construct the predicted rate by

[21] Starting in 2018, the set of variables reported in HMDA expanded, requiring modifications to our loan-level matching procedure between HMDA and Freddie Mac. In particular, we add loan term (in months) and mortgage rates into the waterfall-style matching procedure. As a result, the match rates are higher than those reported in Table D1.

Table F1: Model fit for predicting rates: with and without demographics

Panel A: Without Demographics			
	(1)	(2)	(3)
R^2	0.8388	0.8402	0.8416
Loan Deciles Controlled	Y	Y	Y
LLPA × Year-Month FE	Y	Y	Y
State FE		Y	
State × Zip FE			Y
N	4,974,143	4,974,143	4,974,143

Panel B: With Demographics			
	(1)	(2)	(3)
Black		0.0145***	0.0138***
		(0.0014)	(0.0014)
Hispanic		-0.0041***	-0.0036***
		(0.0011)	(0.0011)
Asian		-0.0682***	-0.0675***
		(0.0012)	(0.0012)
Single Acct. Female			0.0128***
			(0.0008)
Joint Acct.			0.0079***
			(0.0007)
Income decile 2			-0.0093***
			(0.0013)
Income decile 3			-0.0068***
			(0.0013)
Income decile 4			-0.0049***
			(0.0013)
Income decile 5			-0.0043***
			(0.0014)
Income decile 6			-0.0036***
			(0.0014)
Income decile 7			-0.0041***
			(0.0014)
Income decile 8			-0.0041***
			(0.0015)
Income decile 9			-0.0048***
			(0.0015)
Income decile 10			-0.0076***
			(0.0016)
R^2	0.8309	0.8313	0.8314
Loan Deciles Controlled	Y	Y	Y
LLPA × Year-Month FE	Y	Y	Y
State × Zip FE	Y	Y	Y
N	1,249,721	1,249,721	1,249,721

plugging in either (a) the borrower's FICO score at origination and the updated loan-to-value ratio, or (b) the quarterly updated VantageScore and the updated loan-to-value ratio, together with other time-invariant demographic characteristics.

We find that the predicted rates based on FICO at origination and the rates based on the quarterly updated VantageScore are highly correlated (correlation = 0.97). Moreover, the correlation between the predicted rates from the models with and without demographics is 0.99 (both based on the updated VantageScore). Accordingly, in our main analysis, we use the quarterly updated VantageScore and the first-stage model without demographic controls, which is trained on a larger sample.

Using the predicted counterfactual rates, Figure F2 presents the fraction of mortgage borrowers with a rate gap, defined as the difference between the predicted counterfactual rate and the outstanding mortgage rate, that is greater than or equal to either 0 or 50 basis points. Overlaid in the 30-year fixed mortgage rate. The figure shows that the fraction of borrowers with a positive rate gap increases when the 30-year fixed rate is low, such as during 2016–2017 and from mid-2019 through the end of 2021, indicating that a sizable fraction of borrowers would benefit from refinancing during most of our sample period, except for 2018 through mid-2019.

Figure F2: Fraction of Borrowers with Positive Refinancing Rate Gaps

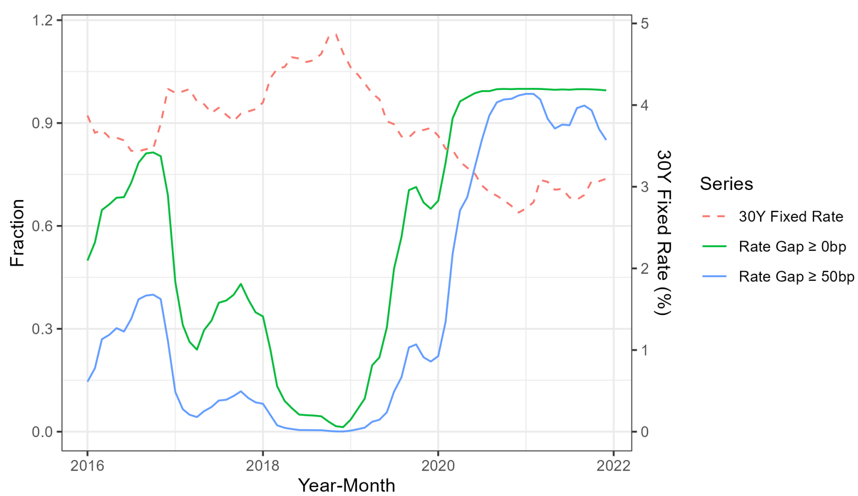

Note: The refinancing rate gap is defined as the difference between the predicted counterfactual rate and the outstanding mortgage rate.

G Calculating Refinancing Incentives

In this appendix, we provide a detailed description of the construction of the three refinancing incentive measures.

G.1 Rate Gap

In Appendix F, we describe the calculation of the borrower-specific counterfactual refinancing rates $r^{\text{new}}_{i,t}$, assuming borrower i refinances at time t. The refinancing rate gap is defined as the difference between the original interest rate r_i^{orig} and the counterfactual rate $r_{i,t}^{new}$:

$$RateGap_{i,t} = r_i^{orig} - r_{i,t}^{new}.$$

G.2 ADL Incentive

Agarwal et al. (2013) defines a (conservative) optimal refinancing threshold $O_{i,t}$ that takes into account closing cost and marginal tax rate, remaining balance and term, and expectation of future events (such as rate changes, moving, inflation). More specifically, the optimal refinancing threshold for loan i at time t is

$$O_{i,t} = \sqrt{\frac{\sigma \kappa_{i,t}}{M_{i,t}(1-\tau)} \sqrt{2(\rho + \lambda_{i,t})}}.$$

For parameters that are not loan- or time-specific, we use suggested values from Agarwal et al. (2013): $\sigma = 0.0109$ denotes the annual standard deviation of mortgage rates, $\tau = 0.28$ is the marginal tax rate, and $\rho = 0.05$ is the discount rate. The other parameters are loan or time specific: $M_{i,t}$ is the remaining balance (in dollars), $\kappa_{i,t}$ denotes the closing costs (\$2,000 plus 1% of $M_{i,t}$). The prepayment hazard $\lambda_{i,t}$ depends on the annual probability of moving, $\mu = 0.1$, a term that depends on the original interest rate r_i^{orig}, the remaining terms (in years) $T_{i,t}$, as well as the inflation rate $\pi = 0.03$.

$$\lambda_{i,t} = \mu + \frac{r_i^{orig.}}{exp(r_i^{orig.} T_{i,t}) - 1} + \pi.$$

After computing $O_{i,t}$ for each loan i at time t, we define the ADL incentive as the difference between the $RateGap_{i,t}$ and the optimal refinancing threshold:

$$ADL_{i,t} = (r_i^{orig} - r_{i,t}^{new}) - O_{i,t}.$$

G.3 Monthly Savings

We compute the potential monthly savings from refinancing at each point in time, net of closing costs.

Let P, r, and T denote the loan principal, annual interest rate, and loan term (in months), respectively. For a fixed-rate mortgage, the monthly payment is

$$M = P \cdot \frac{\frac{r}{12}\,(1+\frac{r}{12})^T}{(1+\frac{r}{12})^T - 1}.$$

For borrower i, the original monthly payment M_i^{Orig} is calculated using the initial loan balance P, the original interest rate r, and a 30-year maturity ($T = 360$).

To compute the counterfactual monthly payment $M_{i,t}^{\text{New}}$, we use the borrower-specific counterfactual refinancing rates $r_{i,t}^{\text{new}}$. The counterfactual payment is computed using the formula above with two additional assumptions. First, the principal is set equal to the remaining loan balance plus refinancing closing costs, which we assume to be \$2,000 plus 1% of the remaining balance. This construction assumes no cash-out refinancing. In practice, cash-out refinancing may reflect additional liquidity motives—such as substituting for higher-cost personal loans or home equity lines of credit—that are not captured by our measure of payment savings. Second, we assume that the refinanced loan maintains the remaining maturity, so that $T^{new} = 360 - age_{it}$, where $age_{i,t}$ denotes the number of months since origination. Holding the original maturity isolates the effect of interest rate reductions on monthly payments, rather than mechanically lowering payments through extending the loan term.

Accordingly, our approach provides a benchmark measure of refinancing incentives driven exclusively by interest rate-induced payment reductions, which abstracts from other potential refinancing motives. Potential monthly savings are defined as the difference between the original and counterfactual monthly payments:

$$M_{i,t}^{\text{Save}} = M_i^{\text{Orig}} - M_{i,t}^{\text{New}}.$$

The three measures are positively correlated. The pooled correlation between monthly savings and the rate gap is 0.81, between monthly savings and ADL is 0.77, and between the rate gap and ADL is 0.96. Figure G1 illustrates these relationships using pairwise heatmaps.

Figure G1: Three Measures of Refinancing Incentives

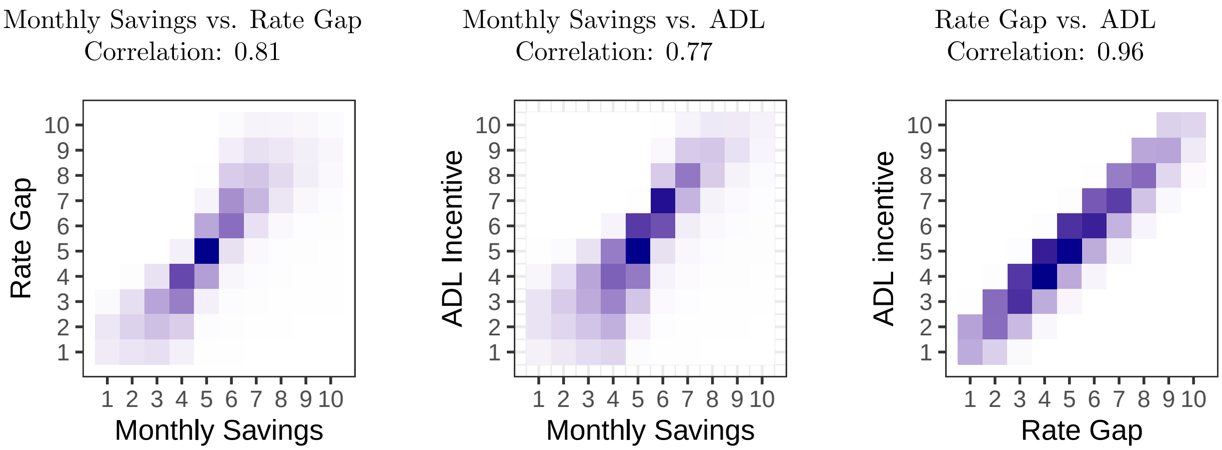

Notes: Values 1 to 10 correspond to the incentive regions as defined in Section 4.1. Darker shading corresponds to areas with a higher number of observations. The reported correlations are computed using all observations.

G.4 NPV of Refinance Savings

Next, we compute the net present value (NPV) of refinancing savings by discounting the stream of monthly payment savings over the remaining life of the loan. Let ρ denote the annual discount rate. If borrower i refinances in month t, the NPV of savings is given by

$$NPV_{i,t}(\rho) = M_{i,t}^{\text{Save}} \cdot \frac{1 - \left(1 + \frac{\rho}{12}\right)^{-T_{i,t}}}{\frac{\rho}{12}},$$

where $M_{i,t}^{\text{Save}}$ denotes monthly payment savings from refinancing and $T_{i,t}$ is the remaining loan term in months.

We also consider an alternative NPV measure that accounts for the possibility of early loan termination. Specifically, we augment the annual discount rate by an expected repayment hazard $\lambda_{i,t}$, which is calculated in the same way as described in the ADL incentives in Appendix G.2. Under this adjustment, the effective annual discount rate becomes $\rho + \lambda_{i,t}$. This adjustment captures the possibility that borrowers may repay or terminate their loans prior to maturity, thereby shortening the horizon over which refinancing savings are realized.

In our analysis, we set $\rho = 0.05$ to compute an upper bound on NPV savings. Using $\rho + \lambda_{i,t}$, which is close to 0.2 on average, yields a lower bound, since the higher effective discount rate reduces the present value of future savings.

H Ad Effect in Pre-COVID and COVID Periods

Our sample spans January 2016 through December 2021, covering both a period of relatively stable interest rates (2016–2019) and the dramatic rate decline during the COVID-19 pandemic (2020–2021). To assess whether the advertising effects documented in Section 4 are driven by the unusual pandemic environment, we estimate the heterogeneous advertising effect model separately for the pre-COVID period (January 2016 through February 2020) and the COVID period (March 2020 through December 2021).

Figure H1: Ad Effect on Refinancing: Pre-COVID vs. COVID Period

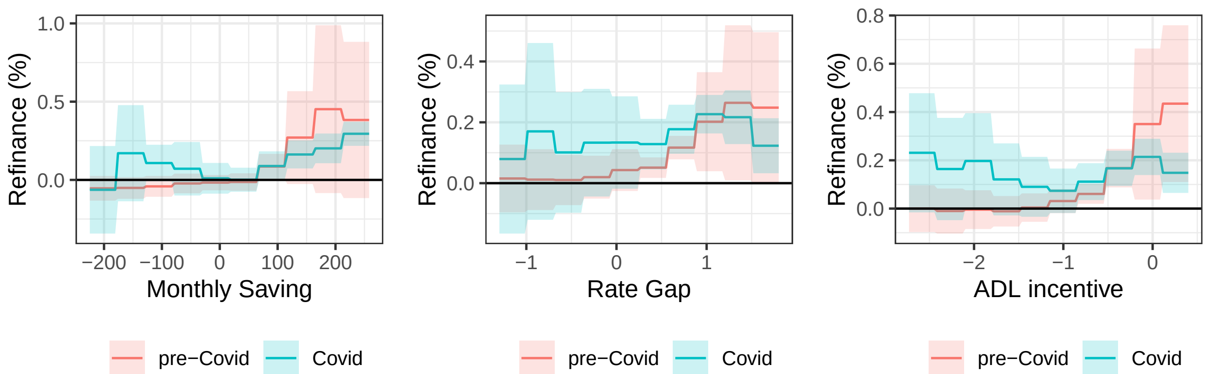

Notes: Estimates are from a joint model with three-way interaction (incentive bin × ad stock × COVID-period indicator). The pre-COVID period is January 2016 through March 2020; the COVID period is April 2020 through December 2021.

Figure H2 reports the heterogeneous advertising lift on refinancing—the Q4–Q1 income differential and the White–Black and White–Hispanic racial differentials—separately for the two sub-periods. The income differential is present in both periods: the Q4–Q1 gap in advertising lift widens at high incentive levels regardless of the rate environment. Similarly, the racial null result is stable: the White–Black and White–Hispanic differentials remain statistically indistinguishable from zero in both sub-periods. These findings confirm that the patterns documented in Section 4 are not artifacts of the pandemic.

Figure H2: Heterogeneous Ad Effect on Refinancing: Pre-COVID vs. COVID Period

Panel (a): Income: Q4 vs. Q1

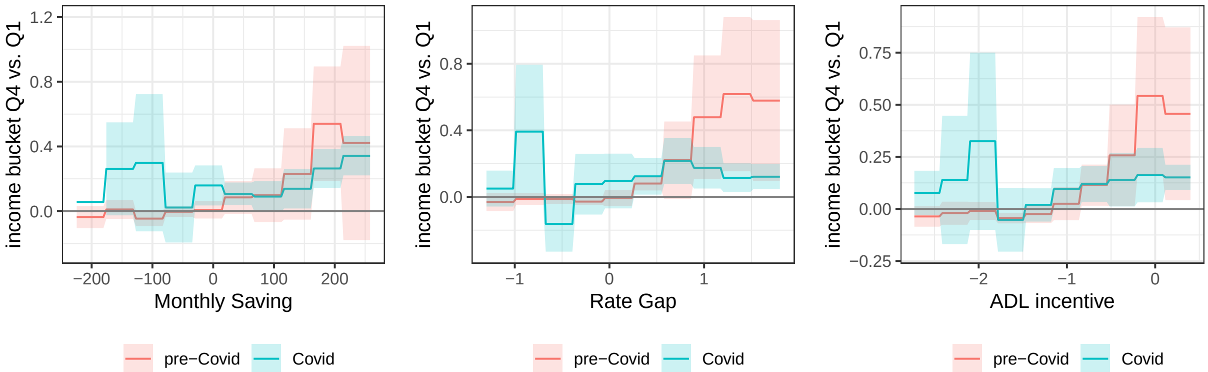

Panel (b): White vs. Black

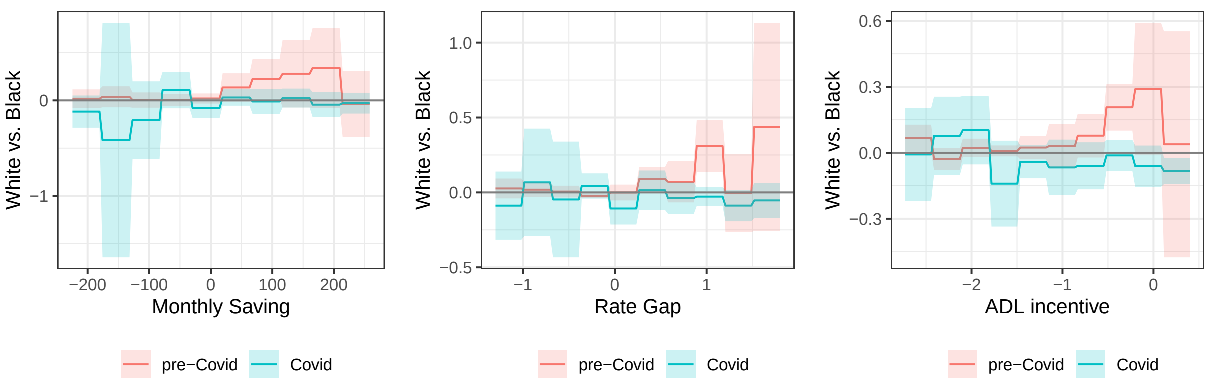

Panel (c): White vs. Hispanic

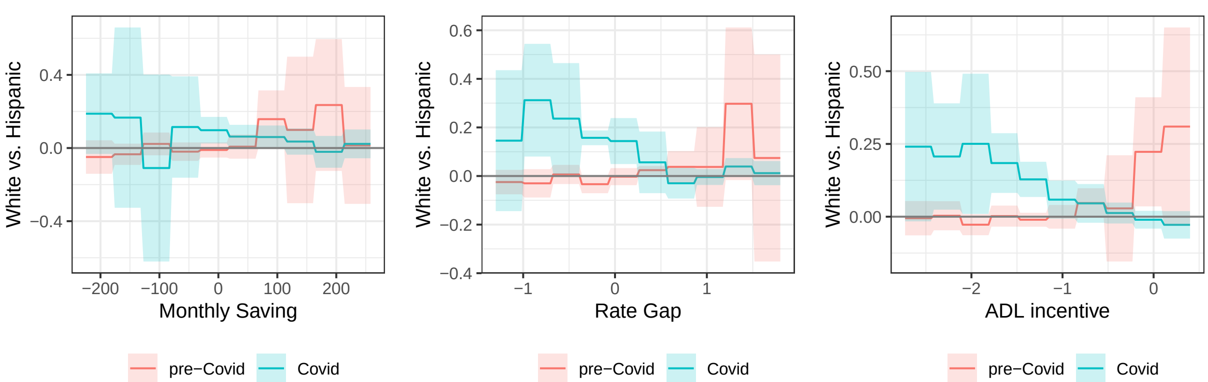

Notes: Estimates are from a joint model with four-way interaction (incentive bin × ad stock × demographics × COVID-period indicator). The pre-COVID period is January 2016 through February 2020; the COVID period is March 2020 through December 2021.

Figure H3 reports the corresponding results for credit inquiries. The absence of an income differential in inquiry response to advertising is likewise stable across sub-periods, reinforcing the interpretation that the income gap in refinancing arises from barriers to completion rather than from differential interest.

Figure H3: Heterogeneous Ad Effect on Credit Inquiry: Pre-COVID vs. COVID Period

Panel (a): Income: Q4 vs. Q1

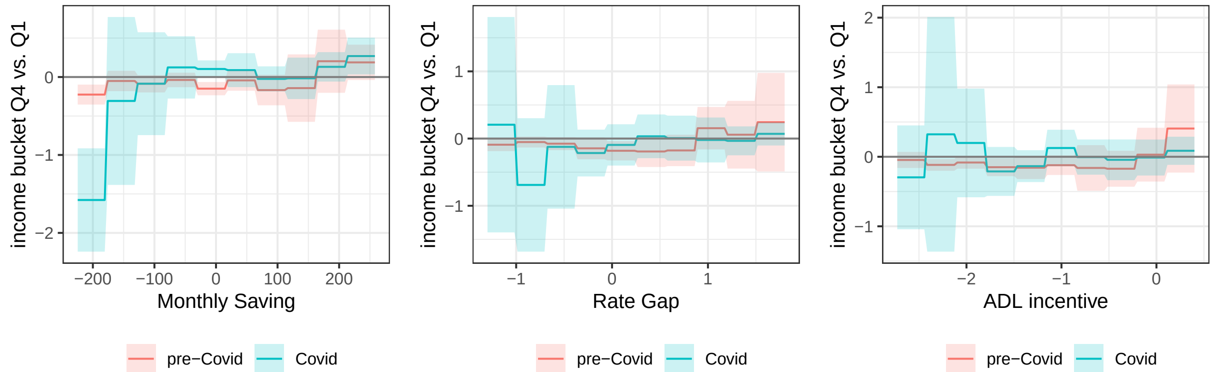

Panel (b): White vs. Black

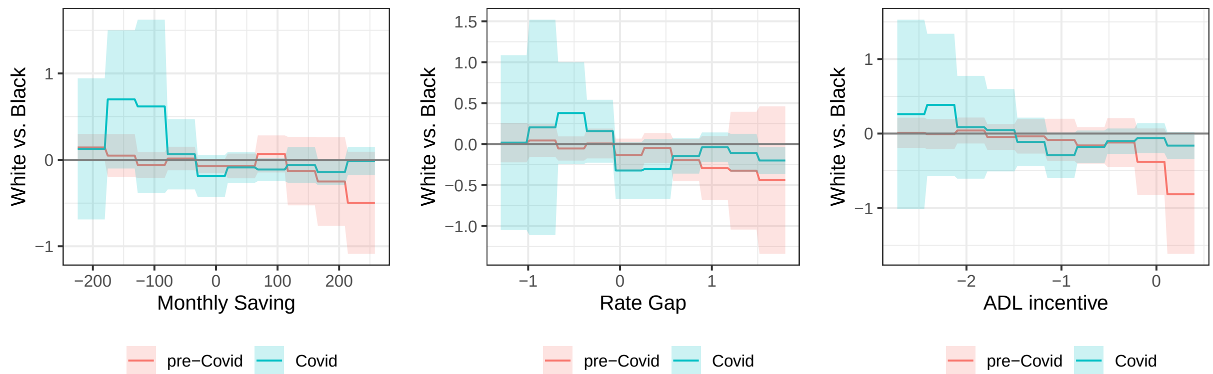

Panel (c): White vs. Hispanic

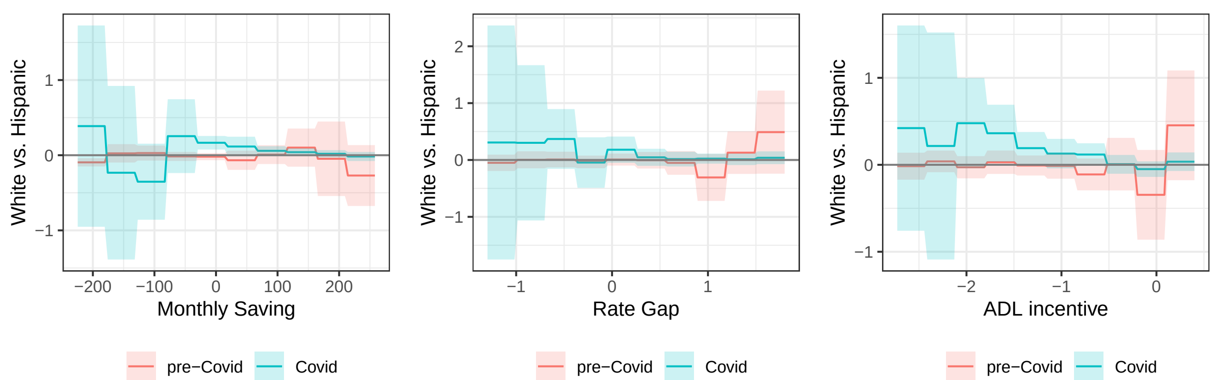

Notes: Estimates are from a joint model with four-way interaction (incentive bin × ad stock × demographics × COVID-period indicator).

I Gender Heterogeneity in Advertising Effects

In the main text, we focus on income and racial heterogeneity in advertising effects. Here we extend the analysis to gender. Figure I1 plots the monthly refinancing hazard by gender as a function of incentives, while Figure I2 reports the heterogeneous advertising lift on refinancing by gender.

Figure I1: Refinancing Hazard by Gender

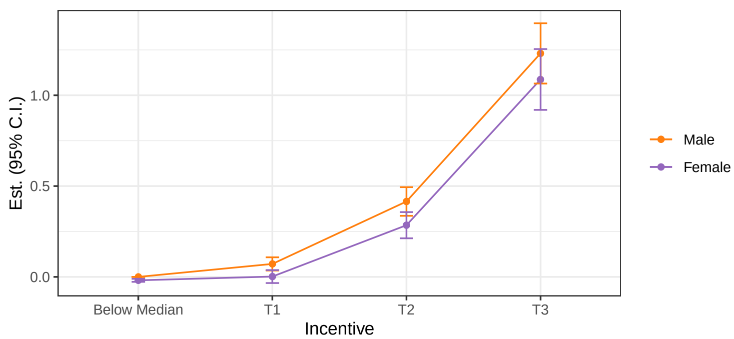

Figure I2: Heterogeneous Ad Effect on Refinancing by Gender

(a) Ad Effect by Gender (Rate Gap)

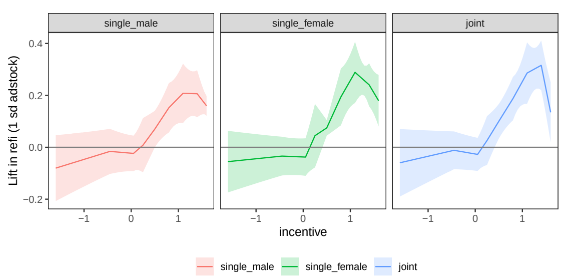

(b) Gender Differential in Ad Lift (Rate Gap)

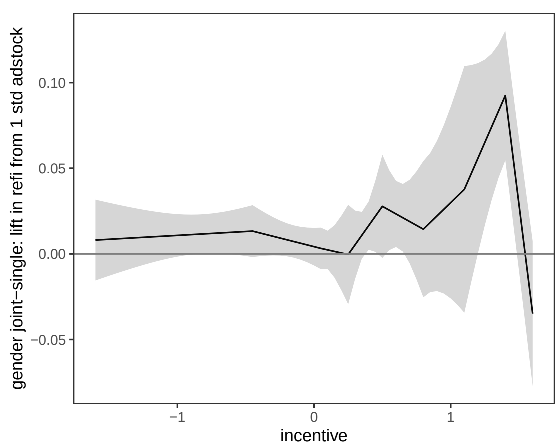

Notes: Panel (a) plots the advertising lift on refinancing separately for male and female borrowers. Panel (b) plots the male–female differential in advertising lift with 95% confidence intervals. The specification follows Specification 4 the main text, replacing income/race interactions with a gender indicator.

Figure I3 reports the corresponding results for credit inquiries by gender.

Figure I3: Heterogeneous Ad Effect on Credit Inquiry by Gender

(a) Ad Effect by Gender (Rate Gap)

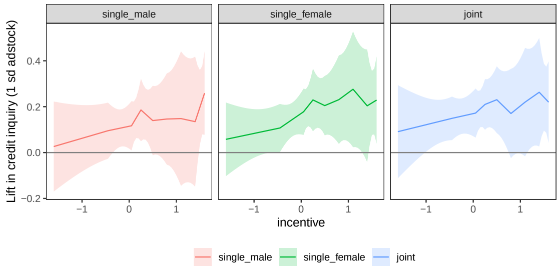

(b) Gender Differential in Inquiry Lift (Rate Gap)

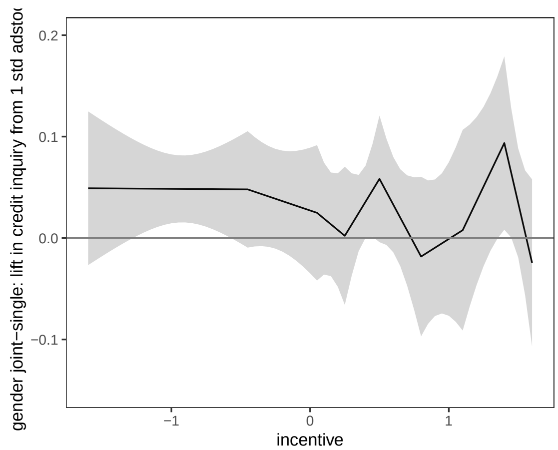

J Advertising Exposure by Demographics

A potential concern is that advertising effects differ across groups because of differential advertising exposure rather than differential responsiveness. While our specification controls for DMA-level exposure, within-DMA differences in television viewership could generate differential effective exposure across demographic groups. Figure J1 provides evidence on this channel by plotting advertising exposure (measured in GRPs from cable television data, which provide demographic breakdowns) by racial group. The time-series patterns are broadly similar across groups, though Black and lower-income viewers exhibit slightly higher overall exposure levels, likely reflecting differences in television viewing habits. The key point is that exposure does not systematically favor the groups that exhibit larger advertising responses (high-income, White borrowers), suggesting that differential exposure is unlikely to drive the heterogeneous refinancing effects we document.

Figure J1: TV Advertising Exposure by Racial Group

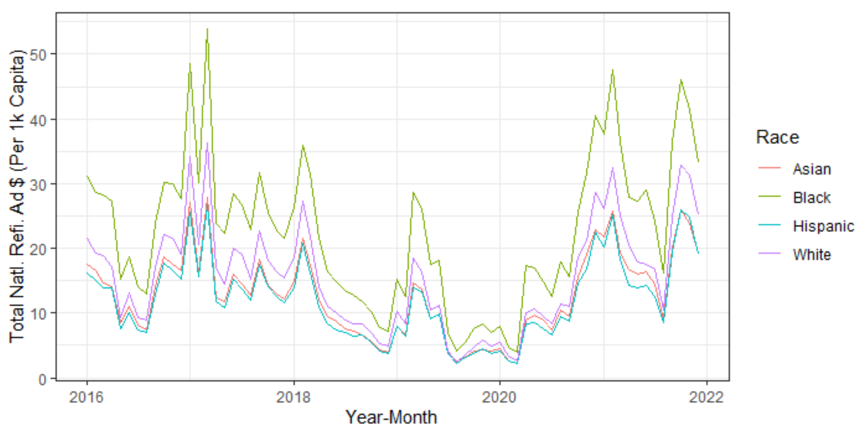

Notes: Monthly cable TV advertising GRPs by racial group. Data from Nielsen AdIntel viewer-level demographic breakdowns for cable television.