

TV Advertising Effectiveness with Racial Minority Representation: Evidence from the Mortgage Market

[Click here for the most recent version.](#)

Donggwan Kim
Washington University

Zhenling Jiang
University of Pennsylvania

Raphael Thomadsen*
Washington University

July 25, 2023

Abstract

This paper examines the impact of racial minority representation on advertising effectiveness. We do this by first assembling data on 10 million mortgage refinance loans, along with data on TV advertisements for mortgage refinance. We construct a measure of minority representation from video ads using computer vision techniques, and extract additional video and transcript features from the advertisements using a variational autoencoder and a text embedding model. We then apply a Double Machine Learning model to estimate how the minority representation in ads affects which lender consumers choose for their refinancing, while controlling for high-dimensional image and text features, as well as a rich set of fixed effects. We find that ads with higher minority representation are more effective in driving consumer choices: as the minority share in ads increases from 15% to 25%, the advertising elasticity increases from 0.037 to 0.042 (a relative increase of 14%). This effect is more pronounced among minority borrowers but is also positive among White borrowers. Across the political spectrum, minority representation has a larger impact among liberal-leaning consumers. In addition to our observational study, we conduct a pre-registered lab experiment ($N = 2,796$) where we manipulate the race of the actors using generative AI technology. The results are consistent with those from our observational study, providing further causal evidence for our findings. We discuss potential mechanisms driving these results, as well as the implications of our findings.

Keywords: minority representation; TV advertising; double machine learning; image and text embeddings

*donggwan.kim@wustl.edu, zhenling@wharton.upenn.edu, thomadsen@wustl.edu.

1 Introduction

In recent years, there has been a growing emphasis on promoting diversity, equity, and inclusion (DEI), which has prompted companies to make commitments to advance DEI. One visible manifestation of this commitment is the increasing number of chief diversity officers (CDOs) tasked with driving DEI initiatives at the organizational level. According to an article by McKinsey & Company, over 50% of Fortune 500 firms have appointed CDOs as of 2022.¹ Efforts to promote DEI can also be observed in many other areas of society, such as inclusive hiring practices in the workplace, fostering diverse student bodies in education, and the inclusion of diverse characters and narratives in media.

From a marketing perspective, promoting diversity and minority representation in advertising holds significant importance for companies for a number of reasons. Advertising is a powerful tool for marketers to connect with consumers and convey brand values. By including minority actors in advertisements, companies can better engage with their minority customer base, as these consumers respond more positively to ads featuring actors from their own racial background (e.g., [Deshpandé and Stayman, 1994](#); [Aaker, Brumbaugh, and Grier, 2000](#)). Beyond the racial fit between the consumers and actors, companies can also signal their commitment to DEI initiatives through minority representation in their ads. Recent surveys conducted by Microsoft and Facebook find that consumers are more trusting of brands that represent diversity in their ads, and ads featuring more diverse actors are associated with higher ad recall, both suggesting positive consumer attitudes towards minority representation in advertising.²

This paper investigates the impact of racial minority representation on the effectiveness of TV advertising in the empirical context of mortgage refinancing. Refinancing a mortgage is one of the most financially consequential decisions that a household can make. In 2020 alone, \$2.6 trillion worth of mortgage loans were refinanced.³ Prior studies have shown that minority consumers are less likely to refinance their mortgages compared to White consumers with similar characteristics, forgoing substantial potential savings ([Gerardi, Lambie-Hanson, and Willen, 2021](#); [Gerardi, Willen, and Zhang, 2023](#)). The lower refinance take-up rate among minorities contributes to the well-documented racial disparities in the mortgage market (e.g., [Bartlett et al., 2022](#); [Bhutta, Hizmo, and Ringo, 2022](#)). From this perspective, understanding the impact of minority representation in ads in this market holds particular social significance.

More specifically, we aim to answer the following research questions. First, how does the effectiveness of TV advertising change with varying degrees of minority representation? After showing the main effect that higher minority representation increases advertising effectiveness, we further investigate potential heterogeneous effects. In particular, how do the effects differ based on the borrower's own race and political

¹<https://www.mckinsey.com/capabilities/people-and-organizational-performance/our-insights/unlocking-the-potential-of-chief-diversity-officers>. Accessed July 12, 2023

²<https://about.ads.microsoft.com/en-us/insights/inclusive-marketing-whitepaper>;
<https://www.facebook.com/business/news/insights/the-difference-diversity-makes-in-online-advertising>. Accessed July 12, 2023

³<https://www.statista.com/statistics/205946/us-refinance-mortgage-originations-since-1990/>. Accessed July 12, 2023

leanings? Finally, taking together results from both observational and experimental studies, we discuss several potential mechanisms that are consistent with our findings.

To answer these research questions, we obtain loan-level mortgage origination data from 2018 to 2021. This dataset includes information on the borrower’s race and census tract of the property, which allows us to link the home location to census tract-level political voting data. We then merge this loan origination data with TV mortgage advertising data obtained from Kantar Media, covering the same time period. In addition to advertising spending data, we also obtain the video files of the TV ad creatives. Using these videos, we utilize computer vision techniques to determine the race of each actor and construct a measure of minority representation in advertising.

There are two main challenges in estimating the impact of minority representation on TV advertising effectiveness using observational data. First, lenders may advertise more towards individuals who are more likely to be responsive or during time periods when the return on advertising is expected to be high. Second, the level of minority representation in ads may be correlated with video features, such as visual elements or advertising messages, which can also affect the effectiveness of the ads. Ignoring these potential correlations can introduce omitted variable bias in the estimates.

We address these concerns using a Double Machine Learning (Double ML) model (Chernozhukov et al., 2018). In the Double ML model, we explicitly allow the advertising levels and minority representation in ads to depend on a very flexible functional form on high-level interactions between lender, location (Designated Market Area, or DMA), and time (year) fixed effects, as well as high-dimensional video attributes. The video attributes include image embeddings which we obtain by training a variational autoencoder (VAE), as well as ad transcript embeddings, which we obtain from a pre-trained embedding model from OpenAI. Because the Double ML model allows for interactions between each of the control variables, it includes the benefits one would get from a fixed effects regression with a rich set of interactive fixed effects. As a benchmark, we also present the results of a fixed effects regression with lender-DMA and lender-time fixed effects. This model, however, does not control for the possibility that advertisements with more minority actors might also have different messages or video imagery. The Double ML approach controls for this correlation by estimating how the advertising levels and minority representation are correlated with the video and message attributes, and then using only the residual variation to estimate the causal effect. We find that the estimates from the fixed effect model are mostly similar to results from the Double ML model.

Our results show that increased minority representation leads to an increase in the effectiveness of the ads. Specifically, we find that as the minority share in ads increases from 15% (representing the median value at the lender-DMA level in our data) to 25%, the advertising elasticity increases from 0.037 to 0.042 (a 14% relative increase). Moreover, we observe that the impact of minority representation is more pronounced among minority borrowers than White borrowers, but is also positive among White borrowers. Across the political spectrum, we find a stronger effect among liberal-leaning consumers compared to conservative-leaning consumers, suggesting that support for diversity and minority representation, or attitudes about

race in general, may play an important role.

We further complement our observational study with a pre-registered lab experiment ($N = 2,796$) where we directly manipulate the race of the families in advertisements using generative AI technology. The results from our experimental study confirm the main effect observed in our observational study: Participants who are randomly assigned to advertisements featuring minority families report a higher likelihood of applying for refinancing from the advertised lender and recommending the lender compared to those randomly assigned to ads featuring White families. Furthermore, the experimental results show consistent patterns on the relative impact of minority presence in ads based on the participants' self-reported race and political orientation. The converging results from the experiment with random assignment help assure us that our estimated effects from our observational study are not simply the result of a subtle endogeneity or omitted variable biases.

There are several possible mechanisms that are consistent with our results. First, consistent with our heterogeneous effects, prior results have documented that minority consumers have a stronger preference for racial homophily than White consumers (e.g., [Deshpandé and Stayman, 1994](#); [Aaker, Brumbaugh, and Grier, 2000](#); [Mollica, Gray, and Trevino, 2003](#)), and liberal-leaning individuals have greater support for racial diversity and equity than conservatives (e.g., [Agarwal and Sen, 2022](#); [Babar, Adeli, and Burtch, 2022](#); [Aneja, Luca, and Reshef, 2023](#)). By asking follow-up questions in the experiment, we find that ads featuring minority actors lead to favorable brand perceptions, such as perceptions of broad loan options, fair lending practices, and inclusiveness toward individuals of all backgrounds, all of which can contribute to higher ad effectiveness. Lastly, ads featuring minority actors can be more effective simply because they are less common, making them stand out to consumers ([Pieters, Warlop, and Wedel, 2002](#); [Rosengren et al., 2020](#)). Indeed, our experimental participants perceive ads featuring minority actors as new, fresh, and attention-grabbing compared to ads with only White actors.

Our findings offer valuable insights for brands in shaping their advertising strategies. By featuring minority actors, brands can not only signal their commitment to DEI, but can also increase the effectiveness of their advertising efforts. Our results also have important policy implications. Given that featuring minority actors is particularly effective in reaching minority consumers, it has the potential to be a useful strategy to provide these consumers with information about refinancing opportunities and encourage them to refinance, especially in times when interest rates are low. This, in turn, could help reduce racial disparities in the refinance take-up rate.

Our research contributes to several streams of literature. First, our paper is closely related to the nascent literature on the impact of DEI initiatives and social equity movements on both society and business. In terms of societal impact, [Agarwal and Sen \(2022\)](#) find a significant increase in demand for anti-racist books requested by public school teachers following the killing of George Floyd. From a managerial perspective, [Balakrishnan, Nam, and Buell \(2022\)](#) and [Khan and Kalra \(2022\)](#) demonstrate that signals of diversity at the corporate level have a positive impact on consumer attitudes. Furthermore, signaling racial identity can increase demand for minority-owned businesses on platforms like Yelp and improve the success rate

of requests for help (Babar, Adeli, and Burtch, 2022; Kirgios et al., 2022; Aneja, Luca, and Reshef, 2023). Beyond racial diversity, Goli and Mummalaneni (2023) find that an increase in women’s screen time positively impacts the viewership of cable news shows. However, it is important to note that DEI initiatives may not always receive favorable responses, underscoring the need for careful assessment of the potential benefits and risks associated with such initiatives. For example, Wang et al. (2022) find that firms’ social media posts related to the Black Lives Matter movement reduce consumer engagement on social media platforms. Our paper contributes to this literature by studying how minority representation in TV ads impacts advertising effectiveness.

Within this domain, the studies closest to ours are two concurrent working papers by Hartmann, Netzer, and Zalta (2023) and Overgoor et al. (2023). Hartmann, Netzer, and Zalta (2023) find that online display advertisements featuring minority actors achieve higher click-through rates than those with White actors. Overgoor et al. (2023) study the impact of Black actor share in TV ads on consumers’ purchase intentions and find that the effect depends on the processing route. Our study differs from these sets of papers in a few important ways. First, we study advertising effectiveness with actual consumer demand rather than relying on clicking behaviors or self-reported purchase intentions. Second, we account for the potential correlation between the race of the actors and visual and text features in ads in order to minimize omitted variable biases. Thus, we have a much richer set of controls in our study. Third, by leveraging detailed information on each consumer’s race and (census tract level) political leaning, we estimate heterogeneous effects along these dimensions. This, in conjunction with the experimental results, allows us to discuss potential mechanisms at play.

Our paper also adds to the literature on the impact of advertising content. Beyond lab experiments, ad content has received less attention than the examination of ad quantity in economics and marketing. Bertrand et al. (2010) measure the effect of informational content (e.g., interest rates) and non-informational content (e.g., a photo featuring an attractive woman) in direct mail ads for consumer loans through a large-scale field experiment, and find strong evidence of the significant impact of ad content. Since then, marketers have utilized either field experiments (e.g., Sudhir, Roy, and Cherian, 2016; Sahni, Wheeler, and Chintagunta, 2018; Morozov and Tuchman, 2022) or observational data (e.g., Liaukonyte, Teixeira, and Wilbur 2015; Lee, Hosanagar, and Nair 2018; Tsai and Honka 2021; Fossen et al. 2022) to gain deeper insights into the impacts of both informational and non-informational ad content. Building on this literature, our study explores how racial representation in ads, as a form of non-informational content, can influence consumer demand for the advertised lender.

The rest of the paper is organized as follows. Section 2 describes our main data and provides descriptive statistics. Section 3 describes the multi-modal features we extract from the video data. Section 4 presents our empirical strategies. Section 5 documents the empirical results. Section 6 describes the online experiment and documents the experimental results. Section 7 discusses potential mechanisms. Finally, Section 8 concludes.

2 Data Source and Descriptive Statistics

In this section, we describe the two main datasets used in our study: the mortgage origination data (Section 2.1) and the TV advertising data (Section 2.2). We also discuss how we extract race information from advertising video files and present descriptive statistics in Section 2.3.

2.1 Mortgage Origination Data

We obtain loan-level mortgage origination data from the Home Mortgage Disclosure Act (HMDA) database for the period from 2018 to 2021. The HMDA is a U.S. federal law that mandates mortgage lenders to disclose detailed information about their mortgage lending activities. The data collected under this law cover approximately 90% of total mortgage originations with reporting exemptions for small lenders (Bhutta, Laufer, and Ringo, 2017). This dataset provides comprehensive loan-level information, including the originating lender, year of origination, loan size, and loan type. Additionally, the HMDA data include borrower characteristics, including race, which allow us to study how consumers from different racial backgrounds respond to minority representation in ads. We also use the property’s census tract to collect information about political voting patterns.⁴

A new mortgage loan can be originated for the purpose of a home purchase or refinancing. We focus on refinances because consumers often rely on real estate agents or mortgage brokers when choosing a lender during the home purchase process, which can lead to limited advertising effectiveness at that stage. In contrast, refinancing decisions are typically made independently by consumers.

We focus on conventional mortgage loans that follow standard underwriting guidelines, such as having a minimum credit score of 620 or above and a debt-to-income ratio below 50%. Consequently, we exclude non-conventional, government-backed loans, including VA loans for veterans or active-duty service members, FHA loans for low-income and low-credit score consumers, and USDA loans for rural areas. Conventional mortgage loans represent approximately 80% of the total mortgage originations during our sample period (Liu, Jo, and Chen, 2022). Following previous studies on household finance, we impose additional inclusion criteria (e.g., Bartlett et al., 2022; Bhutta, Hizmo, and Ringo, 2022; Gerardi, Willen, and Zhang, 2023): These loans must be first-lien mortgages for owner-occupied, site-built, single-family residential homes with a minimum loan size of \$100,000. Additionally, we exclude jumbo loans that exceed conventional loan size limits and other unconventional loan types, such as reverse mortgages, interest-only loans, balloon payment loans, and negatively amortizing loans.

Since the TV advertising data (discussed in Section 2.2) covers the top 101 Designated Market Areas (DMAs), we include loans originating from these markets. Further, as we seek to analyze the heterogeneous responses of consumers from different racial backgrounds, we exclude loans with missing or mixed joint

⁴We obtain census block group-level estimates for the 2020 presidential election from Bryan (2022), which uses the methods described in Amos, McDonald, and Watkins (2017). This data has been used in academic studies, such as Babar, Adeli, and Burtch (2022).

race information.⁵ These data cleaning procedures lead to a sample of 9.7 million loans for our analysis. We provide further details of the data cleaning process and the number of excluded loans at each step in Appendix A. For ease of estimation, we randomly select 2.89 million borrowers, which account for 30% of the full sample, as our estimation sample.

Table 1: Market Share, Advertising Spending, and Minority Share in Ads by Lender per Year

Lender	Market Share	Ad Spending (per 1,000 Capita)	Minority Share in Ads
Rocket Mortgage	10.87%	\$1147.90	34.69%
United Wholesale Mortgage	5.69%	\$7.62	39.04%
Wells Fargo	4.03%	\$341.13	20.48%
JP Morgan Chase	3.77%	\$0.67	21.95%
LoanDepot	2.78%	\$98.17	13.98%
Nationstar	1.90%	\$26.97	13.33%
Bank of America	1.72%	\$0.00	—
Caliber Home Loans	1.52%	\$0.00	—
US Bank	1.30%	\$6.68	18.25%
Fairway Independent	1.28%	\$197.05	8.08%
PennyMac	1.18%	\$0.00	—
Guaranteed Rate	1.14%	\$80.50	23.53%
Flagstar Bank	1.12%	\$16.81	0.00%
Home Point Financial	1.08%	\$0.00	—
Freedom Mortgage	1.01%	\$1.45	25.15%
Newrez Mortgage	0.86%	\$0.00	—
Provident Funding	0.85%	\$0.00	—
AmeriSave Mortgage	0.77%	\$105.42	4.29%
Citizens Bank	0.74%	\$6.07	34.88%
Better Mortgage	0.74%	\$0.20	0.00%
CrossCountry Mortgage	0.74%	\$75.59	23.52%
PNC Bank	0.72%	\$0.06	36.28%
Broker Solution Bank	0.70%	\$0.31	32.52%
Cardinal Financial	0.69%	\$6.08	8.02%
Finance of America	0.68%	\$10.62	6.42%
Guild Mortgage	0.60%	\$13.94	6.03%
Fifth Third Bank	0.52%	\$0.00	—
Huntington Natl. Bank	0.49%	\$0.00	—
Movement Mortgage	0.49%	\$0.26	25.00%
American Financing Corp.	0.42%	\$941.54	5.59%

Notes: Ad spending denotes the total ad spending per 1,000 capita, including both national ads and local ads across the top 101 DMAs.

When studying the choice of lenders, we narrow our focus to the top 30 lenders, which collectively represent over 50% of all refinancing mortgage originations, and categorize the remaining smaller lenders as

⁵We compare borrowers with missing or joint ethnic/racial information to those with complete information on loan size, income, and age and find similar distributions. Further details are provided in Appendix B.

the “outside” option. Column 1 of Table 1 presents the top 30 lenders ranked by their average market share per year, as reported in column 2.⁶ Rocket Mortgage (formerly known as Quicken Loans) has the largest market share, followed by United Wholesale Mortgage, Wells Fargo, and JP Morgan Chase. These top four lenders account for 24% of the market share. Mortgage lending is a much less concentrated market than many other markets, such as airlines or breakfast cereals, where the top 4 companies have market shares of 67% and 85%, respectively.⁷ We define the variables in columns 3 and 4 of Table 1 in the sections below.

2.2 Mortgage TV Advertising Data

We obtain TV advertising data from Kantar Media for the same sample period as in the mortgage data. This data includes monthly advertising spending at the lender-DMA-ad creative level for both national and local ads, covering the top 101 Designated Market Areas (DMAs).⁸ To account for population differences across DMAs, we scale the local advertising spending using the population of the corresponding DMA to obtain ad spending per capita, following previous research in TV advertising (e.g., Shapiro 2018; Tsai and Honka 2021). Similarly, we scale the national advertising spending using the national population. The total ad spending for a specific lender within a specific DMA is defined as the lender’s national ad spending per capita plus their local ad spending per capita within the DMA.

Using the ad spending data, we show the average ad spending per 1,000 capita per year for each lender across all of the 101 DMAs in Column 3 of Table 1. We observe significant variations in total advertising spending among lenders. Rocket Mortgage is the largest advertising spender during the sample period, followed by American Financing and Wells Fargo. However, some major lenders, such as JP Morgan Chase and Bank of America, allocate little or no budget to TV advertising.

Besides the advertising spending data, we also collect the video files of the TV ad creatives. In our data, there are a total of 1,441 unique ad creatives aired by the top 30 lenders.⁹ For each ad creative, we observe the total ad spending at the DMA-month level. We utilize these ad videos to determine if and to what extent they feature minority actors. In addition to race, we extract visual and textual features from these video ads, as detailed in Section 3.

⁶In this list, we exclude BB&T and SunTrust, which merged into Truist in December 2019. Because of the merger, there was a time period where they had separate advertising campaigns but reported to HMDA under the new name Truist, which creates challenges in matching the advertising data with the loan origination data.

⁷<https://www.statista.com/statistics/250577/domestic-market-share-of-leading-us-airlines/>;
<https://www.statista.com/statistics/858562/cereal-company-market-share-us/>. Accessed July 12, 2023

⁸In the U.S., TV markets, known as DMAs, are defined by the Nielsen Company to measure ratings across different geographic regions. Each DMA typically consists of multiple counties, with a major city at its center, along with surrounding smaller counties. Advertisers have the option to purchase national ads that are broadcasted across all DMAs (a total of 210) or local ads that are limited to specific DMAs (e.g., Chicago DMA).

⁹This number is based on the ad creative names reported in the Kantar data. We exclude a small number of ads specifically targeting reverse mortgage loans.

2.3 Race Detection and Minority Representation Measure

To determine the race of each actor in advertisements, we leverage pre-trained computer vision algorithms from Clarifai Inc. rather than training our own model. Pre-trained models, which are trained on large datasets, generally outperform models trained by researchers on smaller data. Several prior studies in marketing and management have also used Clarifai’s pre-trained models (e.g., [Dzyabura and Peres, 2021](#); [Khern-am nuai et al., 2021](#); [Zhang et al., 2022](#); [Hartmann, Netzer, and Zalta, 2023](#)). We build a customized “workflow” on Clarifai: We first detect any faces in a given frame, then crop an image for each detected face, and finally predict the probability that the face belongs to each racial group. We apply this workflow to frame-level image data, where we sample one frame per second from each ad video. Figure 1 shows two examples. On the top panel, Clarifai detects two faces and predicts that the male actor is Black with a probability of 0.86 and the female actor is Black with a 0.99 probability. On the bottom panel, one face is detected, and the actor is predicted to be East Asian with a probability of 0.80. We have manually checked a number of predictions and found the Clarifai algorithms to be highly accurate.

Before constructing our measure of racial representation using the predicted race information, we conduct two additional data processing steps. First, we group certain racial categories from Clarifai to align with the categories in the HMDA data. Specifically, we combine “White” and “Middle Eastern” into the White category and group “East Asian”, “Southeast Asian”, and “Indian” into the Asian category. This results in four racial categories: White, Black, Hispanic, and Asian.¹⁰ Second, we exclude a small number of predictions where the probability of the most likely race falls below 70%, similar to previous studies (e.g., [An and Kwak, 2019](#); [Guitart and Stremersch, 2021](#)). This is to ensure that the detected race variable contains minimal measurement errors.

To measure the level of racial representation in videos, we take into account both the duration of time that each race appears on the screen and the extent of screen sharing when a video features multiple actors. Suppressing the subscript for each ad video for brevity, let $f = 1, \dots, F$ denote the frame with human faces in the video. Let J^f denote the number of actors in frame f and $R_j^f \in \{W, B, H, A\}$ denote whether the race of the j^{th} actor in frame f is White (W), Black (B), Hispanic (H) or Asian (A). When a frame contains multiple actors ($J^f \geq 2$), we divide the screen share for each actor by the number of actors present in the frame ($\frac{1}{J^f}$). The measure of racial representation for each video is calculated as follows:

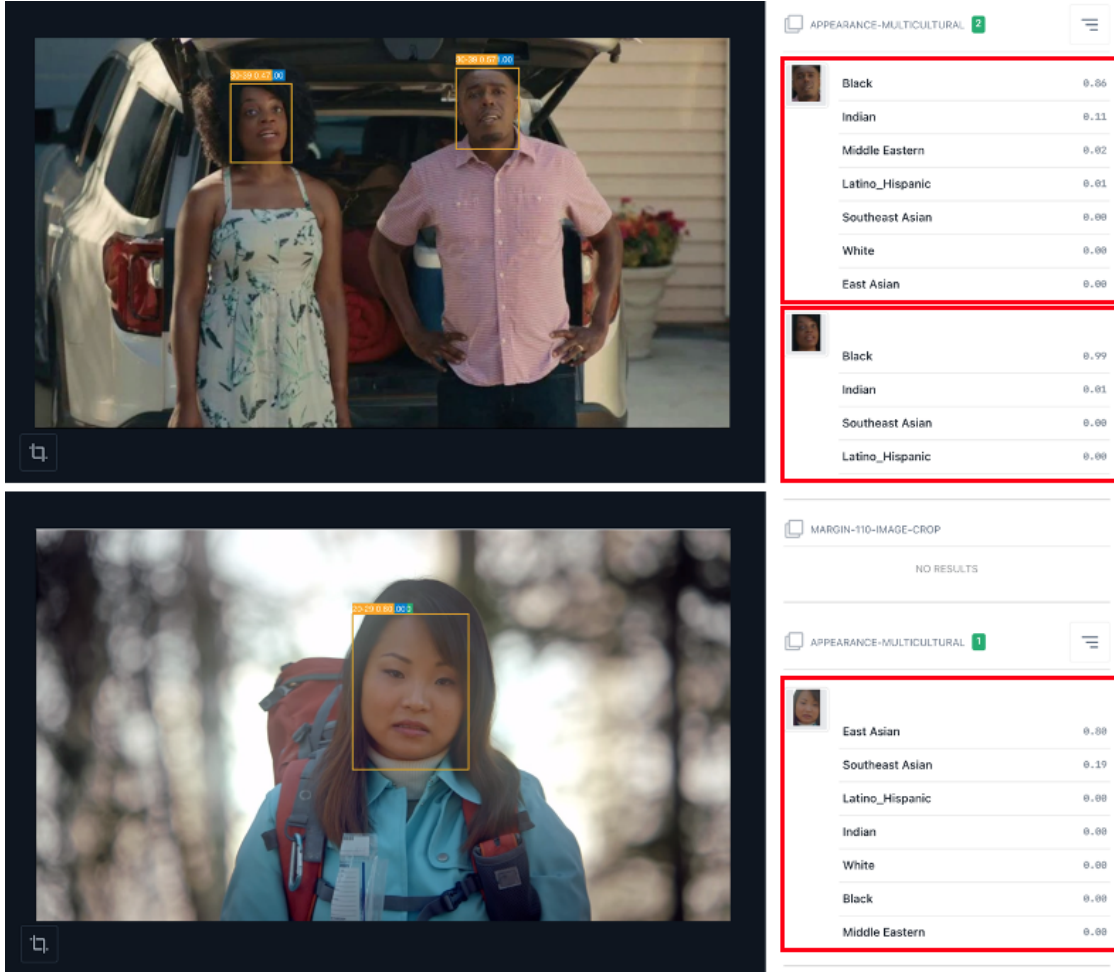
$$Share^{Race} = \frac{1}{F} \sum_{f=1}^F \left(\frac{1}{J^f} \sum_{j=1}^{J^f} \mathbb{1}\{R_j^f = Race\} \right), \quad (1)$$

where $Race \in \{W, B, H, A\}$ denotes each racial group, and $\mathbb{1}\{R_j^f = Race\}$ is an indicator function that takes the value of 1 if the race of the j^{th} actor in frame f matches the given $Race$, and 0 otherwise.

Let’s consider a 30-second ad as an example. Suppose a White actor appears in the first 14 seconds, and

¹⁰In our study, we use race to refer to both ethnicity and race, and these classifications are based on the appearance of actors in the ads. Thus, we treat Hispanic as a racial category.

Figure 1: Clarifai Examples



then two actors, one White and one Black, appear together during the next 14 seconds, and the lender’s logo is displayed in the last 2 seconds. In this example, the total number of frames with human faces F is 28. The share of White actors is calculated as $Share^W = \frac{1}{28} (14 \cdot 1 + 14 \cdot 0.5) = 0.75$ because the first 14 seconds only have a White actor (screen share of 1), and the next 14 seconds have both a White actor and a Black actor (screen share of $\frac{1}{2}$ for each). Similarly, the share of Black actors is $Share^B = \frac{1}{28} (14 \cdot 0 + 14 \cdot 0.5) = 0.25$ because the Black actor appears in the latter 14 seconds together with the White actor.

Table 2 presents the average share for each racial group at the ad creative level. We observe that the average share of White representation is approximately 0.8. Among the minority groups, the share of Black representation is the highest at 0.13, indicating that a significant portion of the minority representation in our data comes from Black actors. In contrast, the shares of Hispanic and Asian representation are relatively low at 0.01 and 0.06, respectively.¹¹


¹¹We observe that the Hispanic representation is the lowest. Clarifai’s race detection algorithm can only identify

Table 2: Racial Representation in Ads at the Ad Creative Level ($N = 1,441$)

Race	Mean (SD)	Min	Median	Max
White	0.797 (0.262)	0	0.902	1
Black	0.130 (0.229)	0	0	1
Hispanic	0.009 (0.042)	0	0	0.545
Asian	0.064 (0.124)	0	0	1

Figure 2 plots the monthly shares of minority representation for Black, Hispanic, and Asian actors from 2018 to 2021. We calculate the weighted average share of racial representation for each minority group within a given month, considering both national and local ads across all of the 101 DMAs and using ad spending per capita as the weight. The shares of Hispanic and Asian representation stay relatively stable over time. However, there is a significant increase in the share of Black representation in ads in the second half of 2020. The dashed line in the figure corresponds to the time of George Floyd’s murder. While it is difficult to pin down the exact reasons, the increase in Black representation in ads aligns with lenders responding to the social movement advocating for greater diversity, equity, and inclusion in ads in response to George Floyd’s murder. A similar upward trend in the representation of minority actors, primarily driven by an increase in Black actors, is also observed in online display advertising (Hartmann, Netzer, and Zalta, 2023).

Figure 2: Racial Representation in Ads over Time



Given that the majority of minority representation in ads comes from Black actors, we combine the shares of Black, Hispanic, and Asian representations to obtain the overall share of minority representation in each advertisement, a :

individuals with more indigenous Hispanic features. As a result, many Hispanic individuals may be classified as White. A similar challenge is discussed in a related study by Davis et al. (2019).

$$Share_a^{Minority} = Share_a^B + Share_a^H + Share_a^A. \quad (2)$$

Because the HMDA data is only available at the annual level, we ultimately aggregate the advertising spending and minority share variables to the annual level as well by lender and DMA. The advertising spending is calculated by taking the sum of the spending per capita, and the minority share is the weighted average of minority share across different advertising creatives, where the weight is the spending for each advertisement in the DMA.

We present the minority share in ads for each lender in Table 1 column 3. While most lenders that advertise use minority actors in at least some of their ads, the minority share in ads varies significantly across lenders. For example, Rocket Mortgage and United Wholesale Mortgage have relatively higher minority shares, while other lenders, such as Fairway Independent, AmeriSave Mortgage, and American Financing, are significantly less likely to feature minority actors during our sample period, despite their significant advertising expenditures.

Table 3 presents summary statistics of the key variables at the borrower-lender level. In Panel A, we see that, on average, borrowers have access to 28 lenders in their respective DMA, including the outside option. While the top 30 lenders are generally available in most of the 101 DMAs, there are some exceptions, such as Fifth Third Bank, which primarily focuses its loan originations in the Midwest. We also find that 23% of borrowers belong to racial minority groups, with Black, Hispanic, and Asian borrowers accounting for 4.52%, 8.44%, and 9.85% of the market, respectively. The difference between Democratic and Republic vote shares is calculated as the difference between the number of votes for the Democratic and Republican candidates divided by the sum of their votes in the 2020 presidential election at the census tract level. Panel B of Table 3 presents the summary statistics for the ad spending per capita and minority share variables. On average, 16.8% of the screen time is occupied by minority actors in mortgage ads, which is slightly lower than the fraction of the minority consumers in this market, as shown in Panel A.

3 Other Video Ad Features

While our main goal is to estimate the impact of minority representation in advertising on consumers' choice of lenders, the presence of different races in ads may be correlated with other confounding factors, such as visual elements or advertising messages. For instance, advertisements featuring minority actors might hypothetically place a greater emphasis on the ease of loan application. This potential difference in messaging could introduce omitted variables bias if it is not properly addressed. Therefore, we seek to account for the possibility of such correlation by controlling for a large number of video features. In this section, we describe how we extract high-dimensional visual and textual features from ads.

Table 3: Summary Statistics at the Borrower-Lender Level ($N = 81.2\text{M}$)

Panel A: Borrower Characteristics				
	Mean (SD)	Min	Median	Max
Number of Available Lenders	28.06 (3.52)	8	29	31
Minority Borrower	0.23 (0.42)	0	0	1
Dem.–Rep. Vote Shares	0.08 (0.36)	-0.97	0.07	0.97

Panel B: Advertising Characteristics				
	Mean (SD)	Min	Median	Max
Ad Spending per 1,000 Capita	31.11 (125.09)	0	0	1,071.33
Minority Share in Ads	0.168 (0.149)	0	0.151	1
Black Share in Ads	0.108 (0.128)	0	0.040	1
Hispanic Share in Ads	0.012 (0.025)	0	0	0.500
Asian Share in Ads	0.048 (0.046)	0	0.028	0.413

Notes: Minority share in ads, as well as the racial breakouts, are conditional on positive ad spending.

3.1 Visual Features

To extract visual features from the video data, we first pre-process the video data and select a smaller number of images per video. A typical video in our data has 1.37 billion pixel values (960 pixels for width \times 540 pixels for height \times 30 seconds \times 30 frames per second \times 3 color channels).¹² To reduce the computational burden, we sample one frame every 5 seconds. This results in an average of 6 frames per video since most video ads are 30 seconds long. Following the standard practices in computer vision, we then resize the frame-level data to have 150,528 pixel values per frame (224p for width \times 224p for height \times 3 color channels).

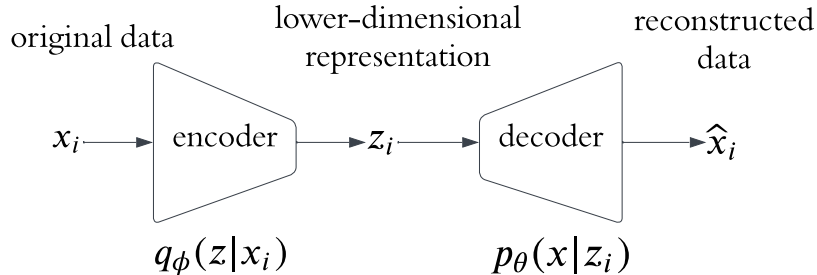
There are two broad approaches that one can use to extract features from image data. The first approach is to extract a number of researcher-defined features, which are often interpretable. For example, [Zhang et al. \(2022\)](#) examine how 12 image features, such as brightness and situation, impact the demand for Airbnb properties. This approach can offer interpretable insights into what exactly is captured in the image data. However, it is often not straightforward to determine which features to extract. Furthermore, unless pre-trained algorithms are available, researchers would need to train their own models to extract the desired features, which are more likely to be subject to measurement errors.

The second approach is to obtain embeddings that represent image data. These embeddings are numeric vectors that may not be easily interpretable. However, unlike the first approach, researchers do not need to manually define a list of features to extract. Instead, they rely on a data-driven approach to capture as much relevant information as possible from the data. Since the primary purpose of extracting these features is to use them as control variables to address potential confounding factors in our application, we prioritize comprehensiveness over interpretability. Therefore, we use the second approach to extract image features.

¹²The vast majority of the video ads have a resolution of 960p \times 540p with a frame rate of 30 frames per second.

We use a variational autoencoder (VAE) to obtain a lower-dimensional representation of the high-dimensional image data (Kingma and Welling, 2013; Rezende, Mohamed, and Wierstra, 2014). VAEs have been applied in several recent marketing studies (e.g., Dew, Ansari, and Toubia 2022; Burnap, Hauser, and Timoshenko 2023; Tian, Dew, and Iyengar 2023). The structure of a VAE is illustrated in Figure 3. In our context, x_i denotes the i -th image from our frame-level video data, represented by high-dimensional pixel-level values ($x_i \in \mathbb{R}^d$), where d is equal to 150,528. This input x_i is passed through an encoder network denoted as $q_\phi(z|x_i)$, which generates a lower-dimensional latent representation denoted as z_i . More specifically, the encoder network generates a stochastic representation by outputting the parameters (mean and variance) of the distribution $q_\phi(z|x_i)$, which is a Gaussian probability density. $z_i \in \mathbb{R}^k$ is sampled from this distribution and has a significantly smaller dimension compared to x_i , with $k \ll d$. In our case, we set $k = 100$. The low-dimensional representation z_i is then fed into the decoder network denoted as $p_\theta(x|z_i)$, which generates the reconstructed image $\hat{x}_i \in \mathbb{R}^d$.

Figure 3: Illustration of Variational Autoencoder



Intuitively, a VAE is a semi-supervised machine learning model that aims to reconstruct the original image using a low-dimensional latent representation that captures the necessary information to minimize the reconstruction error. More formally, a VAE is trained by minimizing the following loss function:

$$L_{\phi, \theta}(x) = -E_{z \sim q_\phi(z|x)} [\log(p_\theta(x|z))] + D_{KL}(q_\phi(z|x) || p(z)), \quad (3)$$

where the first term is the expected log-likelihood of the data x given the latent representation z . The expectation is taken over the encoder’s distribution over the representation z , so that it depends on both the encoder parameters ϕ and the decoder parameters θ . This term is typically known as the reconstruction loss since it encourages the decoder to accurately reconstruct the original data by maximizing the likelihood. The D_{KL} term is the regularization loss. This is the Kullback–Leibler divergence between the encoder’s distribution $q_\phi(z|x)$, which is the variational approximation to the posterior distribution, and the prior distribution $p(z)$, which is assumed to be a standard normal distribution. This regularization term ensures that the latent representation follows a smooth distribution. We refer interested readers to Kingma, Welling

et al. (2019) for further details on the model and estimation.

We implement a VAE using the pre-processed frame-level video data. We train a convolutional encoder-decoder neural network that includes fully connected layers in between to generate a 100-dimensional latent representation z . For more detailed information about the network structure and the training process, please refer to Appendix C.

Once the model is trained, we use the trained encoder to generate a lower-dimensional representation z for each frame. We then aggregate the z 's at the video level by taking the average across all frames within the video. To align with the unit of observation in our mortgage origination data, we further aggregate these video-level image embeddings to the lender-DMA-year level by calculating the weighted average across both national and local ads, where the ad spending per capita is the weight.

3.2 Text Features

To extract features that capture the messages conveyed in the ads, we first obtain the video transcripts using the Amazon Transcribe API, a speech recognition service that converts the audio content of the ad into text. Similar to image analysis, there are two broad approaches to extracting features from textual data. One could seek to identify a set of interpretable topics or sentiments within the ads (e.g., ease of application or a low mortgage rate). Alternatively, one could use a pre-trained large language model to obtain text embeddings that may not be directly interpretable but contain more comprehensive information. As the text features will serve as control variables in our application, we choose the second approach and use a text embedding model to represent the content of ads.

To represent the transcript of each ad in a low-dimensional vector, we utilize a pre-trained embedding model from OpenAI. Specifically, we use the “text-embedding-ada-002” model from OpenAI. While OpenAI offers multiple embedding models, such as “davinci”, “curie”, and “baggage” that are better suited for different tasks (e.g., clustering or search), they recommend “ada-002” for most use cases.¹³ The ada-002 model generates 1,536-dimensional embeddings per document. Since this dimensionality is still relatively high to be used as control variables in our causal inference model, we further compress the embeddings into 100 dimensions using another VAE model. Further details on the network structure and training process can be found in Appendix C. Similarly as the image embeddings, we aggregate these text embeddings to the lender-DMA-year level by taking the weighted average across both national and local ads, where the ad spending per capita is the weight.

4 Empirical Strategy

After establishing that lenders have increased the minority share in their ads, we seek to understand how consumers respond to minority representation in ads. In this section, we describe our empirical strategy for

¹³For more information, see <https://platform.openai.com/docs/guides/embeddings>. Accessed July 13, 2023

estimating the impact of minority share in ads on consumers’ choice of lenders. We begin by describing a benchmark regression model with fixed effects in Section 4.1, which can account for many concerns related to advertisers targeting ads with greater minority representation over specific DMAs and time periods. However, these fixed effects regressions could be vulnerable to potential confounding variables. To address this concern, we use the double machine learning (Double ML) estimator as our main empirical strategy, which is described in Section 4.2. The Double ML estimator allows us to account for the high-dimensional video features described in Section 3 in a flexible functional form and obtain consistent estimates for the main parameters.

4.1 Benchmark: Regression with Fixed Effects

We start by describing a benchmark regression model with a large number of fixed effects to estimate the impact of ad spending and minority representation in the ads on consumer choices.¹⁴

$$y_{i,j} = \beta_1 \cdot \log(1 + Ad_{i,j}) + \beta_2 \cdot \log(1 + Ad_{i,j}) \cdot MS_{i,j} + \delta_{j,m(i)} + T_{j,t(i)} + e_{i,j}, \quad (4)$$

where $y_{i,j}$ is a binary variable that equals 1 if consumer i chooses lender j and 0 otherwise. $Ad_{i,j}$ represents the total ad spending per capita by lender j in DMA $m(i)$ in year $t(i)$. Here, $m(i)$ denotes the DMA where consumer i resides, and $t(i)$ denotes the year when the consumer obtains a refinance loan. Note that the $m(i)$ and $t(i)$ subscripts are suppressed whenever the i subscript is included because each customer only considers a refinance in a specific market during a specific year, making the extra two subscripts superfluous. $MS_{i,j}$ indicates the corresponding minority share in the ads lender j has in market $m(i)$ in year $t(i)$, as defined in Section 2.3.

This regression includes lender-DMA fixed effects $\delta_{j,m(i)}$ and lender-year fixed effects $T_{j,t(i)}$. The lender-market fixed effects account for local, time-invariant confounding factors, such as lenders consistently advertising more in certain DMAs with higher demand or certain consumer characteristics. These fixed effects can also account for higher demand due to a particular lender having more offices or a longer history in a particular market. The lender-year fixed effects capture global, time-varying confounding factors, such as lenders choosing to advertise more, or including more minorities in advertising, in certain time periods, such as the increase in minority representation in advertising that occurred after the murder of George Floyd. The two key parameters of interest, β_1 (which measures the baseline level of advertising effectiveness) and β_2 (which captures how advertising effectiveness varies based on different levels of minority share), would then be identified from the variation in advertising by specific lenders within given markets across time. We show in Appendix D that there is sufficient residual variation in advertising with fixed effects that allows

¹⁴Because multinomial logit or probit models do not handle large numbers of fixed effects well, for the ease of computation, we use a linear probability model as a reasonable approximation for the more micro-founded multinomial logit or probit models, similar to Tsai and Honka (2021) and Wang (2022). This also allows us to be parallel to our Double ML model, where it is challenging to run a multinomial logit or probit. To account for potential correlations within individuals, we bootstrap the standard errors, and further details are provided in Section 4.2.

for the estimation of these parameters. If this level of variation were to be approximately random, then the fixed effects model would yield causal estimates.

Besides the main effects, we could also estimate the results vary based on each consumer’s race. Let $\mathbb{1}\{M_i\}$ be an indicator variable that takes the value of 1 if individual i is from a racial minority group and 0 otherwise (i.e., non-Hispanic White). We extend the model in Equation 4 as follows:

$$y_{i,j} = (\beta_1 + \beta_2 \cdot \mathbb{1}\{M_i\}) \cdot \log(1 + Ad_{i,j}) + (\beta_3 + \beta_4 \cdot \mathbb{1}\{M_i\}) \cdot \log(1 + Ad_{i,j}) \cdot MS_{i,j} + \delta_{j,m(i)} + T_{j,t(i)} + e_{i,j}, \quad (5)$$

where β_2 captures the difference in baseline advertising effectiveness for minority consumers compared to White consumers, and β_4 estimates the difference in the impact of minority share in ads on minority consumers compared to White consumers. If the estimated β_4 is positive and statistically significant, it indicates that minority representation in ads has a stronger impact on minority consumers compared to White consumers.

Similarly, we examine the heterogeneous effects based on borrowers’ political ideology. We use the census tract-level voting outcomes from the 2020 presidential election to represent the political ideology of each consumer. We create a new variable DEM_i that represents the difference in the vote shares between Biden and Trump in 2020 within the census tract where consumer i ’s property is located. More specifically, DEM_i is defined as $\frac{(\text{number of Biden votes}_i - \text{number of Trump votes}_i)}{(\text{number of Biden votes}_i + \text{number of Trump votes}_i)}$. Third party votes are discarded in this calculation. DEM_i is bounded between 1 and -1, where 1 (-1) means that 100% of votes went to Biden (Trump). We then estimate the following regression:

$$y_{i,j} = (\beta_1 + \beta_2 \cdot DEM_i) \cdot \log(1 + Ad_{i,j}) + (\beta_3 + \beta_4 \cdot DEM_i) \cdot \log(1 + Ad_{i,j}) \cdot MS_{i,j} + \delta_{j,m(i)} + T_{j,t(i)} + e_{i,j}. \quad (6)$$

Similar to Equation 5, β_2 captures the different baseline advertising effectiveness based on consumers’ political leanings and β_4 captures how the impact of minority share in ads varies with political leanings. If the estimated β_4 is positive and statistically significant, it indicates that ads featuring minority actors have a stronger impact on liberal-leaning consumers compared to conservative-leaning consumers.

4.2 Double Machine Learning

While the benchmark fixed effects regressions are likely to account for the largest sources of endogeneity, it is still possible that advertisements with higher minority representation may also differ in terms of their messaging or other video features. We can account for these effects using the Double ML estimator (Belloni, Chernozhukov, and Hansen 2014; Chernozhukov et al. 2018). Double ML allows us to estimate causal effects in the presence of high-dimensional covariates. It has recently gained increasing popularity in economics and marketing for causal inference using observational data with high-dimensional control variables (e.g.,

Dube et al., 2020; Ellickson, Kar, and Reeder III, 2022; Gershon and Jiang, 2022; Gordon, Moakler, and Zettelmeyer, 2022). The high-level intuition behind Double ML is to leverage machine learning models to remove or “partial out” the influences of high-dimensional control variables from both the outcome and treatment variables. By doing so, we obtain orthogonalized residuals, which are then used to estimate the causal parameters.

We start by describing the Double ML estimator in our application. We specify the outcome model as a partial linear model:

$$y_{i,j} = \beta D_{i,j} + g(X_{i,j}) + e_{i,j}, \quad (7)$$

where $y_{i,j}$ denotes the binary choice variable as defined previously, and $D_{i,j}$ denote the key causal variables of interest: ad spending per capita, $\log(1 + Ad_{i,j})$, and the interaction term of ad spending and minority share, $\log(1 + Ad_{i,j}) \cdot MS_{i,j}$.¹⁵ $X_{i,j}$ denotes our high-dimensional control variables: the visual and textual features of the ads as described in Section 3, as well as lender, DMA, and year fixed effects. All the variables in $X_{i,j}$ can be thought of as nuisance variables that need to be accounted for in the model but are not the main variables of interest. The impact of the high-dimensional control variables $X_{i,j}$ on the outcome $y_{i,j}$ is captured through a flexible function denoted by $g(\cdot)$.

One naive approach to estimating Equation 7 would be to simply fit a machine learning model to obtain the estimate of the flexible function $\widehat{g(X)}$ and plug it into the regression model in Equation 7. However, under such an approach, the estimates for the main coefficients of interest $\widehat{\beta}$ will be biased. The reason behind this bias can be understood through the regularization in machine learning models, which results in $\mathbb{E}(\widehat{g(X)}) \neq g(X)$, and introduces regularization bias. While this bias diminishes as the sample size (n) increases, it does so at a rate slower than $n^{-1/2}$. Furthermore, the machine learning model can also overfit the training data because of the flexible functional form, resulting in overfitting bias.

The Double ML estimator solves both issues of regularization bias and overfitting bias through orthogonalization and sample-splitting. For orthogonalization, we estimate a second equation to predict the key causal variables $D_{i,j}$ given control variables $X_{i,j}$ (video features and fixed effects) through another flexible function denoted by $h(\cdot)$:

$$D_{i,j} = h(X_{i,j}) + \epsilon_{i,j} \quad (8)$$

To implement Double ML in our application, we use a Random Forest as the machine learning model to obtain the estimates of the conditional expectations $\widehat{l(X)} = \widehat{E(y | X)}$ and $\widehat{h(X)} = \widehat{E(D | X)}$. Using neural networks as the machine learning model gives similar results. Since we have multiple treatment variables, we fit a separate machine learning model for each treatment variable. After fitting the model, the Double ML

¹⁵To estimate heterogeneous effects, $D_{i,j}$ can further include the interaction terms with consumers’ own race $\mathbb{1}\{M_i\}$ as in Equation 5, and interaction terms with political leaning DEM_i as in Equation 6.

estimate is obtained through a residuals-on-residuals regression. Using vector notation, we use the residuals of the outcome variable as $\widehat{\boldsymbol{\epsilon}} = \mathbf{y} - \widehat{\boldsymbol{l}}(\mathbf{X})$, and the residuals of the treatment variables as $\widehat{\boldsymbol{\epsilon}} = \mathbf{D} - \widehat{\boldsymbol{h}}(\mathbf{X})$. These residuals can be thought of as the variations in the dependent and key causal variables (e.g., advertising spending, minority share in ads) after controlling for, or “partialling-out” the effects of the control variables. The main parameters of interest can then be estimated as:

$$\widehat{\boldsymbol{\beta}} = \left(\widehat{\boldsymbol{\epsilon}}' \widehat{\boldsymbol{\epsilon}} \right)^{-1} \left(\widehat{\boldsymbol{\epsilon}}' \widehat{\boldsymbol{\epsilon}} \right) \quad (9)$$

There is one more consideration. The procedure above deals with regularization bias, but the estimated $\widehat{\boldsymbol{\beta}}$ may still be biased due to overfitting bias. This issue is solved by sample-splitting, where we randomly partition the data into K subsets, called folds. For each fold k , we fit the machine learning models to obtain $\widehat{\boldsymbol{l}}(\cdot)$ and $\widehat{\boldsymbol{h}}(\cdot)$ using all folds except the k -th fold, take the fitted models, and estimate $\widehat{\boldsymbol{\beta}}^k$ using the k -th fold. The key is that the observations used to estimate $\widehat{\boldsymbol{\beta}}^k$ are different from those used to fit the machine learning models. Doing so avoids bias that can arise due to overfitting. After iterating through all K folds, we compute the final Double ML estimate by averaging the K estimates. In our application, we have experimented with different numbers of folds, ranging from two to four, and obtained similar results. We opt for two folds for computational efficiency. We refer interested readers to [Chernozhukov et al. \(2018\)](#) and references therein for more technical details.

The conventional approach of obtaining standard errors from the Double ML estimator does not directly apply in our setting. This is because the conventional approach assumes that the error terms are independent and identically distributed across all observations. This assumption is violated in our multinomial linear probability model because the observations for each individual are correlated since each consumer chooses one of the lenders for refinancing. To properly account for the correlation structure, we calculate the standard errors at the individual level using the block bootstrapping technique ([Cameron and Miller, 2015](#)). More specifically, we resample data on an individual level with replacement and compute the parameter estimate for each bootstrap sample. The bootstrap standard error is then calculated as the standard deviation of the 200 bootstrap estimates.

5 Empirical Results

In this section, we present results from both the fixed effects model and the Double ML model. We start by discussing the main effect of including more minority actors on the overall effectiveness of advertising in [Section 5.1](#). We then describe the heterogeneous effects based on consumer characteristics, including race and political leanings, in [Section 5.2](#).

5.1 Main Effect

As discussed in Section 2.1, we estimate both models on a random subset of 30% of consumers (or 2.89 million consumers). Table 4 presents the results of estimating the impact of minority representation in ads on consumers’ choice of lenders. The fixed effects model corresponds to Equation 4, while the Double ML model corresponds to Equations 7 and 8, where the treatment variables D include ad spending and the interaction term of ad spending and minority share in ads. The results from the two models are not statistically different. The baseline effect of ad spending on lender choice (β_1) is positive and statistically significant. Moreover, the main parameter of interest, the interaction term of ad spending and minority share (β_2), is positive and statistically significant. These results indicate that a higher representation of minority actors increases the overall effectiveness of advertising.

Table 4: Effects of Minority Share in Ads on Consumer Choices

Model	Lender choice	
	F.E. Reg. (1)	Double ML (2)
$\beta_1 : \log(1 + Ad)$	0.026*** (0.006)	0.035*** (0.007)
$\beta_2 : \log(1 + Ad) \cdot MS$	0.056** (0.022)	0.060*** (0.020)
N	81,203,548	81,203,548

Notes: MS denotes the minority share in ads. Standard errors, clustered at individual, in parentheses. ** $p < 0.05$; *** $p < 0.01$.

To interpret the effect size of the estimates, we calculate the implied advertising elasticity of demand, which is commonly used as a measure of advertising effectiveness. With simple algebra, the advertising elasticity can be derived as: $\frac{\partial y}{\partial Ad} \frac{Ad}{y} = \beta_1 \cdot \frac{Ad}{1+Ad} \frac{1}{y} + \beta_2 \cdot MS \cdot \frac{Ad}{1+Ad} \frac{1}{y}$, where the notations are defined the same way as in Section 4. The first term represents the elasticity of ad spending in the absence of minority representation, while the second term represents the incremental effect of the minority share in ads on ad elasticity. Using the sample averages for all the variables, we calculate the average advertising elasticity of demand to be 0.030 in the fixed effects model and 0.038 in the Double ML model. The effect sizes are broadly in-line with previous studies: the average elasticity is 0.023 for 288 consumer packaged goods (Shapiro, Hitsch, and Tuchman, 2021), 0.026 for cigarette product placement on TV (Goli et al., 2022), 0.030 for auto insurance (Tsai and Honka, 2021), 0.031 for antidepressant (Shapiro, 2022), 0.05–0.06 for satellite TV operators (Yang, Lee, and Chintagunta, 2021), and 0.08 for e-cigarettes (Tuchman, 2019).

To examine the impact of minority share on advertising effectiveness, we calculate the average advertising elasticity under different levels of minority share in ads while keeping the total level of TV advertising spending constant. Specifically, we present the advertising elasticity for two levels of minority share: 15% (close to the median minority share at the lender-DMA level) and 25%, using our preferred Double ML

model. The results are presented in Table 5. As the minority share increases from 15% to 25%, the estimated elasticity increases from 0.037 to 0.042, representing a 13.6% increase in relative terms. This result suggests that increasing the minority share in advertising can increase the effectiveness of advertising at an economically meaningful level.

Table 5: Advertising Elasticities with Minority Representation

Minority share	Advertising elasticity	
	Mean	95% C.I.
	(1)	(2)
15%	0.0373	[0.0207, 0.0534]
25%	0.0424	[0.0224, 0.0623]

Notes: The range in brackets [] denotes the 95% confidence interval.

5.2 Heterogeneous Effects

In this section, we investigate how the impact of including more minorities in advertising varies with the borrowers' characteristics. We start by examining how the results vary with the borrower's race. The results are presented in Table 6. Across both models, the coefficient of minority share in ads for minority borrowers (β_4) is positive and significant. This indicates that higher minority representation in ads has a larger impact on minority borrowers compared to White borrowers. The effect of minority share for White borrowers (β_3) is also positive and significant, although smaller than that for minority borrowers.

Table 6: Heterogeneous Effects based on Consumers' Race

Model	Lender choice	
	F.E. Reg.	Double ML
	(1)	(2)
$\beta_1 : \log(1 + Ad)$	0.030*** (0.006)	0.032*** (0.007)
$\beta_2 : \log(1 + Ad) \cdot \mathbb{1}\{M\}$	-0.014*** (0.002)	0.0003 (0.002)
$\beta_3 : \log(1 + Ad) \cdot MS$	0.043* (0.022)	0.043** (0.020)
$\beta_4 : \log(1 + Ad) \cdot MS \cdot \mathbb{1}\{M\}$	0.108*** (0.006)	0.061*** (0.004)
N	81,203,548	81,203,548

Notes: MS denotes the minority share in ads; $\mathbb{1}\{M\} = 1$ for minority borrowers. Standard errors, clustered at individual, in parentheses. * $p < 0.1$; ** $p < 0.05$; *** $p < 0.01$.

Further comparing the results from the fixed effects model and the Double ML model, we observe that although the results are generally in-line with each other, there are some differences. In particular, we see

that the baseline impact of ad spending with no minority actors ($\beta_1+\beta_2$) is much smaller in the fixed effects regression than in the Double ML model, while the impact of including minority actors in the advertising ($\beta_3+\beta_4$) is larger for minority customers in the fixed effects regression. These differences suggest that ads featuring minority actors may contain some video features or messaging attributes that can particularly influence minority borrowers’ choice of lenders, and that a failure to account for these features in a simply fixed effects regression may overstate the impact of featuring minority actors in ads on minority borrowers.

To interpret the effect size of the estimates, we examine the advertising elasticities for both White and minority borrowers at two different levels of minority share (15% and 25%), using our preferred model of Double ML. The results are presented in Table 7. Among White borrowers, when the minority share increases from 15% to 25%, the advertising elasticity increases by 11.0%. The effect is even more pronounced among minority borrowers. With the same change in minority share, the advertising elasticity increases by 21.8%. These results confirm that increased minority representation in ads has a larger impact on minority borrowers compared to White borrowers.

Table 7: Advertising Elasticities based on Consumers’ Race

Minority share	White consumers		Minority consumers	
	Mean	95% C.I.	Mean	95% C.I.
	(1)	(2)	(3)	(4)
15%	0.0326	[0.0171, 0.0481]	0.0404	[0.0196, 0.0615]
25%	0.0362	[0.0171, 0.0554]	0.0492	[0.0244, 0.0743]

Notes: The range in brackets [] denotes the 95% confidence interval.

The stronger impact observed on minority borrowers aligns with previous behavioral research indicating that minority borrowers are more sensitive to racial cues and show stronger racial homophily effects (e.g., [Deshpandé and Stayman, 1994](#); [Aaker, Brumbaugh, and Grier, 2000](#); [Mollica, Gray, and Trevino, 2003](#)). However, the positive impact on White borrowers suggests that other mechanisms beyond racial homophily are likely contributing to the observed effects.

We also examine the heterogeneous effects based on the borrower’s political ideology. The results are presented in Table 8. Recall that *DEM* is defined as the difference in the vote shares between the Democratic and Republican candidates in the 2020 election, divided by the sum of the vote shares for the Democratic and Republican candidates. Both models show that the coefficient on the interaction term between the minority share in ads and the vote share difference (β_4) is positive and significant, indicating that increased minority representation in ads has a larger impact on liberal-leaning borrowers than on conservative-leaning borrowers. However, the magnitudes are again quite different between the two models. Ultimately, we believe that the Double ML estimates represent the more robust estimates because this model accounts for the higher-level video and transcript attributes.

To interpret the effect size of the estimates, we examine the advertising elasticities for two groups of borrowers: moderately liberal and moderately conservative. Specifically, we consider a moderately liberal-

Table 8: Heterogeneous Effects based on Political Leaning

Model	Lender choice	
	F.E. Reg. (1)	Double ML (2)
$\beta_1 : \log(1 + Ad)$	0.032*** (0.006)	0.035*** (0.007)
$\beta_2 : \log(1 + Ad) \cdot DEM$	-0.035*** (0.003)	-0.020*** (0.002)
$\beta_3 : \log(1 + Ad) \cdot MS$	0.044** (0.022)	0.067*** (0.020)
$\beta_4 : \log(1 + Ad) \cdot DEM \cdot MS$	0.050*** (0.007)	0.020** (0.005)
N	81,203,548	81,203,548

Notes: MS denotes the minority share in ads; DEM denotes the difference in vote shares between the Democratic and Republican candidates. Standard errors, clustered at individual, in parentheses. ** $p < 0.05$; *** $p < 0.01$.

leaning group with a 25% advantage for the Democratic candidate (i.e., $DEM = 0.25$) and a moderately conservative-leaning group with a 25% advantage for the Republican candidate (i.e., $DEM = -0.25$). For each group, we consider the impact of changing the minority share in ads from 15% to 25%. Results are presented in Table 9. Among liberal-leaning borrowers, the advertising elasticity increases by 17.6%, while the advertising elasticity increases by 15.7% among conservative-leaning borrowers, which is slightly smaller.

Table 9: Advertising Elasticities based on Political Leaning

Minority share	Liberal consumers		Conservative consumers	
	Mean (1)	95% C.I. (2)	Mean (3)	95% C.I. (4)
15%	0.0346	[0.0164, 0.0528]	0.0418	[0.0267, 0.0569]
25%	0.0407	[0.0190, 0.0624]	0.0471	[0.0289, 0.0652]

Notes: The range in brackets [] denotes the 95% confidence interval.

These results are consistent with related studies that indicate that individuals with liberal-leaning political ideologies are more likely to be more supportive of racial diversity and related social movements in multiple contexts (e.g., Agarwal and Sen, 2022; Aneja, Luca, and Reshef, 2023; Babar, Adeli, and Burtch, 2022). We do not find a strong negative impact of minority representation on conservative-leaning borrowers. Even when we extrapolate DEM to an extreme value of -1 (i.e., 100% of votes going for the Republican candidate), the total impact of minority share in ads is positive in the double ML model and close to 0 in the fixed effects model. One caveat with this analysis is that the political-leaning data is only observed at the aggregate census tract level. We will revisit this point when discussing results from our experimental study in Section 6.

6 Experimental Study

We complement the analysis with observational data with an experimental study, where we directly manipulate the race of actors using generative AI technology. We describe the experimental design in Section 6.1 and discuss the results and implications in Section 6.2. The experimental study serves two main purposes. First, while we believe that the Double Machine Learning estimator provides causal estimates, there is always the hypothetical concern of endogeneity or omitted variable biases with observational data. With the experiment, we are able to measure a clean causal relationship with random assignment of ads with different racial compositions. The fact that our results from the observational study match those from the experiment adds confidence that our empirical findings are not driven by some subtle endogeneity story. Second, we use the experiment to inform us about the potential mechanisms by asking participants a number of attitudinal questions about the ads they see, which we discuss in Section 7.

6.1 Experimental Design

As outlined in our preregistered research plan (https://aspredicted.org/B7C_P49), we aim to recruit a total of 2,800 participants from CloudResearch. We plan to recruit 2,000 participants who are either mortgage borrowers or homeowners, and due to the limited available pool of participants, supplement with 800 general population participants. The purpose of prioritizing homeowners is to ensure that the sample is comparable to the borrowers in the observational data. We ended up with a sample of 1,903 participants who were either mortgage borrowers or homeowners and 902 general population participants. After excluding participants who failed the attention check ($n = 9$), our final sample consists of 2,796 participants with an average age of 42 years and 51.3% female.¹⁶

Participants were presented with the following text: “In this survey, we would like you to imagine that you currently have a mortgage loan on your home and you are considering refinancing the mortgage to reduce interest rates. You come across a refinance advertisement from AnchorPoint Refi. Please consider the ad as you are thinking about your refinancing decisions.” Participants were then shown an advertisement featuring customers who had recently refinanced with the advertised lender. Participants were randomly assigned to one of seven conditions, with each condition featuring actors from different racial groups in the ad.

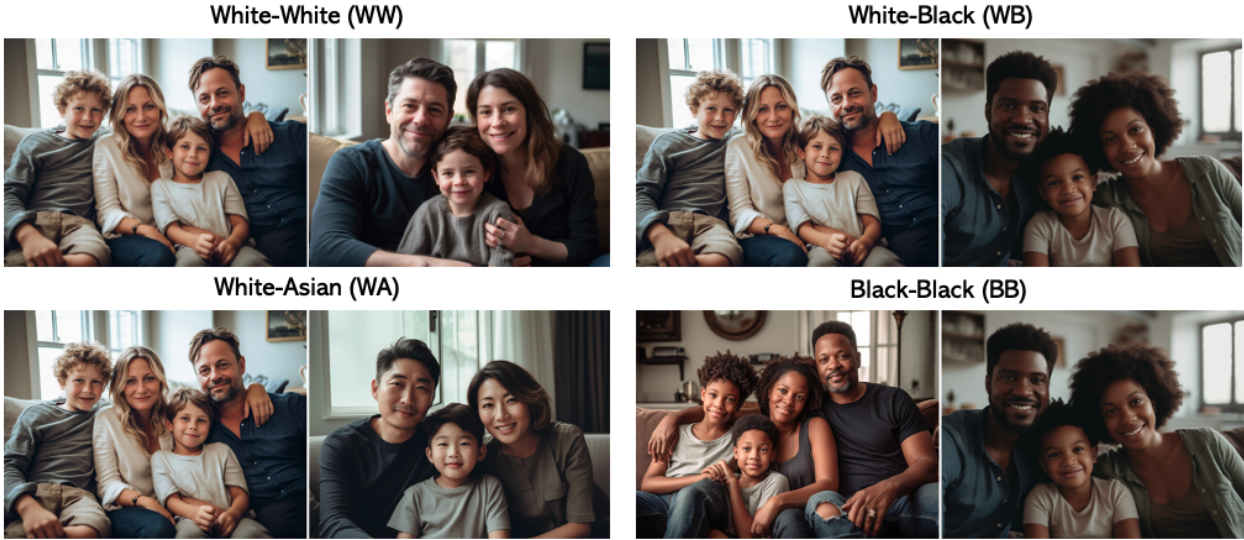
We use Midjourney V5, a generative AI technology that generates highly realistic images based on text prompts, to experimentally manipulate the race of the actors in the ads. To ensure similarity among the generated images, we provided explicit and detailed instructions to Midjourney, including specifications for the number of children and their genders. This is to ensure that other aspects of the images do not differ systematically across conditions. In four out of the seven conditions, participants were presented with an image featuring two families. As shown in Figure 4(a), these four conditions include two White families (WW), one White and one Black family (WB), one White and one Asian family (WA), or two Black families

¹⁶We exclude two outliers in the reported age (660 and 677) when calculating the average age.


(BB). The remaining three conditions, shown in Figure 4(b), featured an image with one White family (W), one Black family (B), or one Asian family (A). The purpose of having both the two-family conditions and the single-family conditions is to explore whether the positive consumer responses observed in our observational study hold when ads feature racially diverse representation (the two-family conditions) as well as minority representation (the single-family conditions).

Figure 4: Advertising Images Featuring Different Races

(a) Two-family conditions



(b) Single-family conditions



After showing one of the seven advertisements, we measure two key dependent variables (DVs): the likelihood of submitting a loan application with the advertised lender and the likelihood of recommending the advertised lender to a friend who is looking to refinance. Both variables are measured on a scale of 1 to 7, where 1 indicates “Not at all likely” and 7 indicates “Very likely.” While the DV of the likelihood to submit an application is better aligned with our dependent variable in the observational analysis, we also included the DV of the likelihood to recommend the lender, which may better capture the overall brand attitude or

impression. The two measures are positively correlated with the correlation coefficient $\rho = 0.771$.

To explore the potential mechanisms behind our results, we also ask participants six attitudinal questions after they respond to the key outcome variables: “To what extent would you agree or disagree with the following statements?” on a 1 to 7 scale where 1 indicates “Strongly disagree,” and 7 indicates “Strongly agree.” The statements cover the breadth of product offerings (“The advertised lender has broad, flexible loan options that fit different financial situations and needs”), whether the respondent felt included (“The advertised lender caters to people like me”), financial inclusiveness (“The advertised lender is inclusive towards individuals of all backgrounds”), fair lending practices (“The advertised lender has fair lending practices without predatory interest rates and hidden fees”), the freshness of the ad (“The advertisement feels new and fresh”), and whether the ad garners attention (“The advertisement is attention-grabbing”).

Lastly, participants were asked to provide their demographic information, including age, gender identity, and race and ethnicity (White/Caucasian, Black/African American, Hispanic/Latinx, Asian, Native American/Alaska Native, Mixed race/multiracial, Others, and Prefer not to disclose). They were also asked to describe their political orientation on a 1 to 7 scale, where 1 indicates “Very liberal” and 7 indicates “Very conservative.”

6.2 Experimental Results

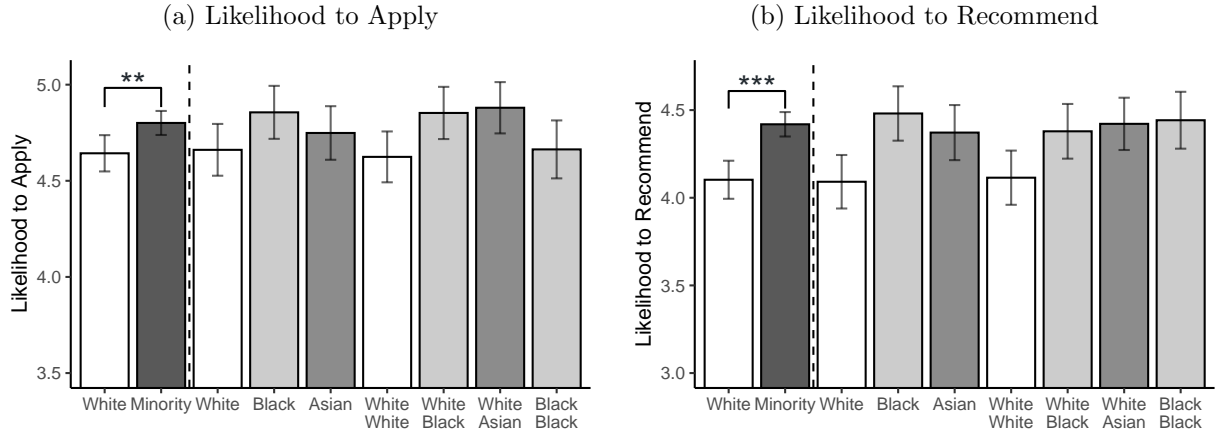
Main Effect

We start by presenting the results of the key outcome measures. Figure 5(a) shows the likelihood-to-apply measure and Figure 5(b) shows the likelihood-to-recommend measure. We first compare the conditions featuring minority families (WB, WA, BB, B, and A) with those featuring only White families (WW and W). Overall, participants who are randomly assigned to advertisements featuring minority families report a higher likelihood of submitting an application with the advertised lender compared to those who see White families (4.80 vs. 4.64, $p = 0.006$). The gap is larger for the likelihood to recommend: Participants who see minority families report a higher likelihood to recommend the advertised lender compared to those who see White families (4.42 vs. 4.10, $p < 0.001$). The experimental results are consistent with the observational study where we find that ads with a higher minority share are more effective.

Figure 5 further shows the results separately for each of the 7 experimental conditions. In the single-family conditions, participants in the Black condition report a higher likelihood of applying with the lender (4.86 vs. 4.66, $p = 0.048$) and recommending the lender (4.48 vs. 4.09, $p < 0.001$) compared to those in the White condition. The Asian condition shows a slightly higher likelihood to apply (4.75 vs. 4.66, $p = 0.376$) and a higher likelihood to recommend (4.37 vs. 4.09, $p = 0.012$) compared to the White condition.

Among the two-family conditions, we use the condition with two White families as the benchmark. Participants in the White-Black condition are more likely to apply to the lender (4.85 vs. 4.62, $p = 0.018$) and recommend the lender (4.38 vs. 4.11, $p = 0.018$). Similarly, participants in the White-Asian condition

Figure 5: Likelihood to Apply and Recommend



Notes: Error bars denote the 95% confidence interval. $**p < 0.01$; $***p < 0.001$.

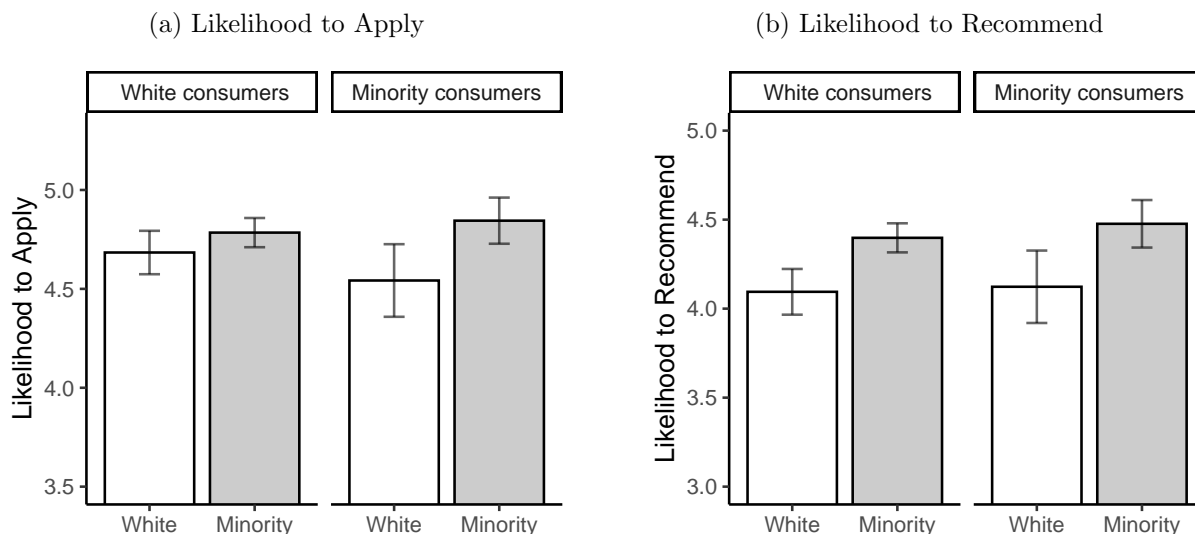
are more likely to apply (4.88 vs. 4.62, $p = 0.008$) and recommend (4.42 vs. 4.11, $p = 0.005$). When participants see an ad featuring two Black families, the likelihood of applying is just slightly higher (4.66 vs. 4.62, $p = 0.703$) and the likelihood of recommending is significantly higher (4.44 vs. 4.11, $p = 0.004$). The results with two Black families point to the possibility of a boundary condition in our observational study: The measured effect may not fully extrapolate to a region where the minority share is substantially higher than what was observed in our observational data. With that said, despite a possible boundary condition, the minority share in ads observed in the data is still lower than the actual share of minority borrowers in the mortgage refinancing market. Overall, while not all comparisons show statistical differences, we find converging results that featuring minority actors increases advertising effectiveness.

Heterogeneous Effects

We examine the heterogeneous effects based on the self-reported racial/ethnic information of the participants. After excluding 10 participants who chose not to disclose their race, there are 2,025 participants who identified as “White/Caucasian,” and we group the rest 761 participants as “minority consumers.” Figure 6 presents the likelihood of making a loan application and a recommendation for these two groups under the conditions that displayed only White families (WW or W) and those with minority families (the other 5 conditions). Figure 6(a) shows that the presence of minorities in the ads increases the likelihood of applying much more for minority consumers (4.84 vs. 4.53, $p = 0.007$) than White consumers (4.78 vs. 4.68, $p = 0.136$). Figure 6(b) shows that both White consumers (4.40 vs. 4.09, $p < 0.001$) and minority consumers (4.48 vs. 4.12, $p = 0.005$) report a significantly higher likelihood of recommending the lender when they are assigned to conditions featuring minority families compared to conditions featuring White families. We report the results for each condition separately in Appendix E. These results are broadly consistent with those from our observational study, where we find that while ads with a higher minority share are more effective for

both White and minority groups, the impact of minority share is stronger among minority consumers.

Figure 6: Heterogeneous Effects By Race

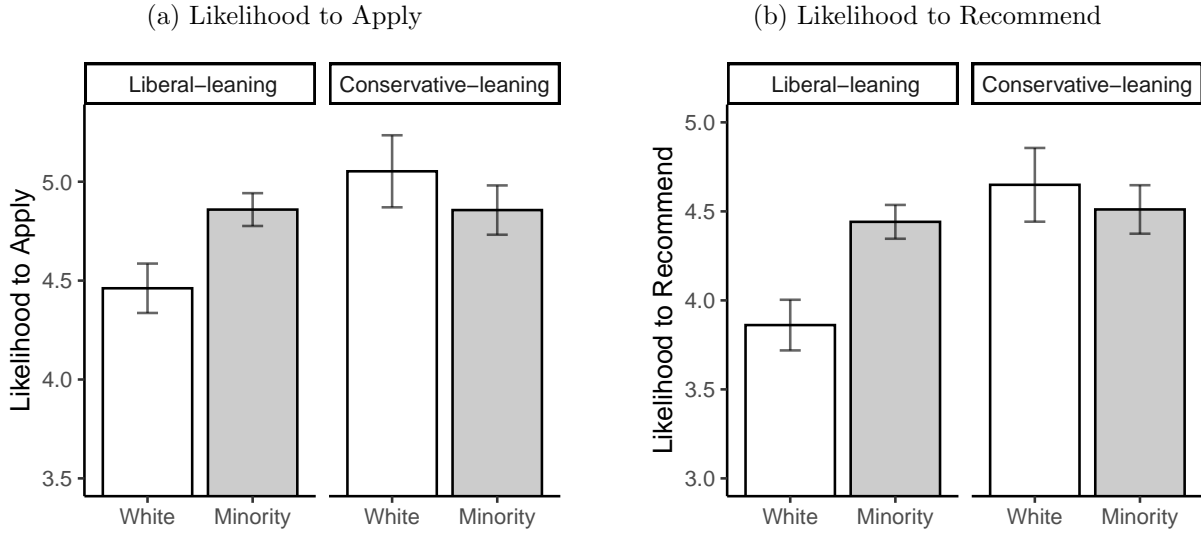


Notes: Error bars denote the 95% confidence interval

Next, we investigate how the results vary based on the self-reported political leanings of participants. As a reminder, our political scale ranged from 1 to 7, where 1 represents “Very liberal,” and 7 represents “Very conservative.” There are 1,454 “liberal-leaning” participants who chose 1 to 3 on this scale, and 786 “conservative-leaning” participants who chose 5 to 7. As shown in Figure 7(a), liberal-leaning consumers report a significantly higher likelihood to submit an application when they are assigned to a condition featuring minority families compared to a condition featuring only White families (4.86 vs. 4.46, $p < 0.001$). Conservative-leaning consumers, on the other hand, are less likely to apply when assigned to conditions that feature minority families compared to only White families (4.86 vs. 5.05, $p = 0.083$), although this effect is borderline insignificant. This pattern also holds when we look at the likelihood to recommend, as shown in Figure 7(b). Liberal-leaning consumers are significantly more likely to recommend the advertised lender when assigned to the conditions featuring minority families (4.44 vs 3.86, $p < 0.001$), while conservative-leaning consumers do not have a significant difference in their likelihood of recommending (4.51 vs 4.65, $p = 0.275$). These experimental results are broadly consistent with the observational study, where we find that the impact of minority representation is larger for liberal-leaning consumers compared to conservative-leaning consumers, although they cast doubt as to whether conservative customers respond positively or negatively to minority representation in an ad.

Comparing Figures 6 and 7, the difference in consumer responses to ads featuring White versus minority families appears even larger when participants are grouped based on their political leaning rather than their own race. This suggests that the effectiveness of minority representation in ads may be better predicted by political ideology than race. Recall that our observational study indicates that race plays a larger role than

Figure 7: Heterogeneous Effects By Political Leaning



Notes: Error bars denote the 95% confidence interval

political leaning (Tables 7 and 9). This could be because while we observe individual-level race information in the observational data, we rely on the census tract level data to proxy political leaning for each consumer, which inherently introduces measurement errors. The experimental results suggest that we would likely expect an even greater difference between liberal- and conservative-leaning consumers based on individual-level political ideology.

7 Potential Mechanisms

So far, we have shown converging evidence from both the observational and experimental studies. In this section, we discuss several potential mechanisms that could explain our findings. To do so, we draw on related prior literature, our empirical results, and the follow-up questions in the experiment that measure participants’ perceptions of the advertised lender and the advertisement (see Section 6.1). Our goal here is not to pin down a single definitive mechanism; rather we show evidence of several possible explanations that are consistent with our findings. Indeed, the effect that ads with minority representation are more effective is likely to be multi-determined.

One potential explanation for our findings is that customers care about the racial match between themselves and the race shown in the ads, with the importance of match being especially high for minority customers. Past behavioral literature has found that racial fit plays a significant role in advertising effectiveness, particularly for minority consumers (e.g., Deshpandé and Stayman, 1994; Aaker, Brumbaugh, and Grier, 2000). In our study, we find that minority consumers are significantly more likely to believe that “the advertised lender caters to people like me” when they see ads featuring minority actors, as opposed

to ads with White actors (4.79 vs. 4.13, $p < 0.001$). White consumers, on the other hand, show only a marginal increase in the belief that ads featuring White actors cater to people like them compared to ads with minority actors (4.79 vs. 4.73, $p = 0.354$). Taking these two findings together, placing minority actors in advertisements should be more effective because of the positive effect this has on minority consumers and non-negative response it generates from White consumers in terms of racial fit. This is consistent with our finding that minority representation has a stronger impact on minority consumers from both observational and experimental studies (Table 6 and Figure 6).

Racial fit is unlikely to be the only mechanism, however, since we find ads featuring minority consumers are also more effective among White consumers, although at a smaller magnitude. There are a number of mechanisms that are consistent with this result, as well.

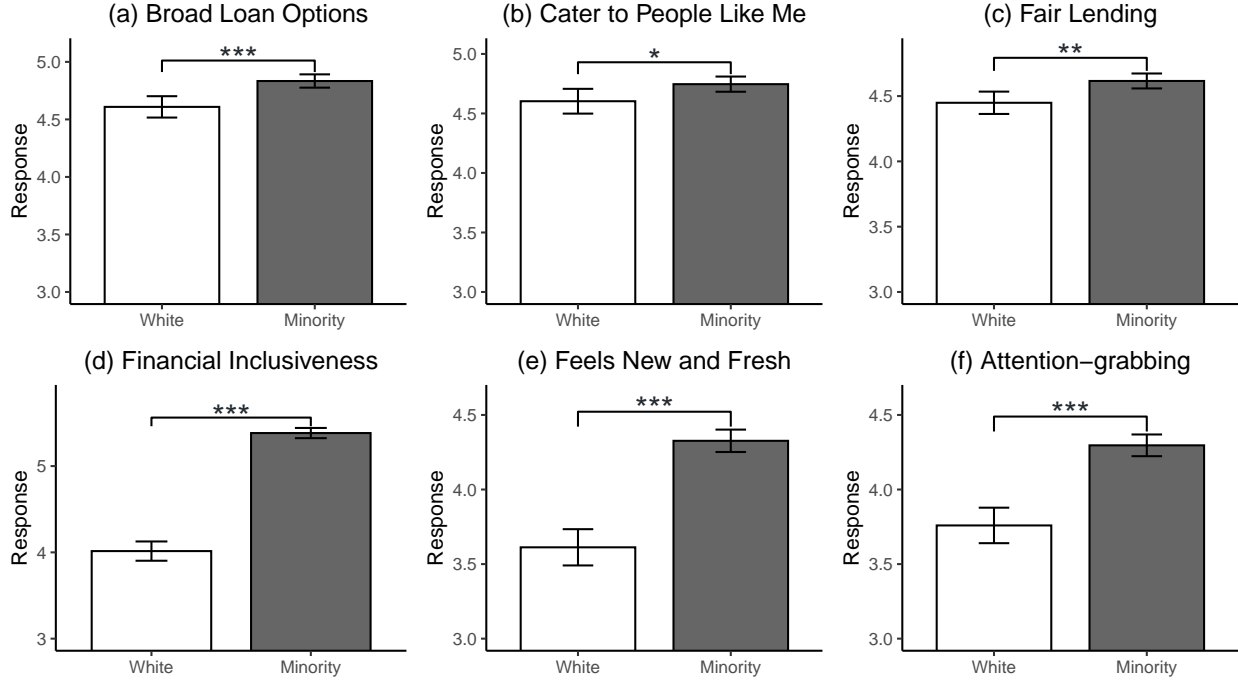
One possible mechanism is that the presence of minority actors in advertisements can speak to consumers' support for diversity and minority representation. Several recent studies have documented greater support for DEI initiatives among liberal-leaning individuals (Agarwal and Sen 2022; Babar, Adeli, and Burtch 2022; Aneja, Luca, and Reshef 2023). Consistent with these studies, we find that liberal-leaning consumers, including White customers, respond more positively toward ads featuring minority actors compared to conservative-leaning consumers. The heterogeneity of the effect across the political spectrum suggests that the extent of support for diversity and minority representation, or attitude about race in general, likely plays an important role in determining how consumers react to ads featuring minority consumers.

Moreover, the presence of minority actors in ads can affect consumers' perceptions of the advertised brand. In other words, even when presented with the same ad copy for a fictional brand, consumers may perceive the brand differently based on the race of the actors. This aligns with prior research that finds that firms with diverse workforces are perceived to be more moral (Khan and Kalra, 2022). We examine brand perceptions using the attitudinal measures collected in the experimental study. Figure 8 compares the results for participants who are randomly assigned to see ads featuring White vs. minority families.¹⁷ After seeing ads featuring minority consumers, participants are more likely to perceive the advertised lender to have broad and flexible loan options (4.83 vs. 4.61, $p < 0.001$), have fair lending practices without predatory pricing and hidden fees (4.62 vs. 4.45, $p = 0.001$), and be inclusive towards individuals of all backgrounds (5.38 vs. 4.01, $p < 0.001$). As all these brand perceptions are positively correlated with the likelihood of applying for a loan and recommending the lender (Table E4 in Appendix E), minority representation can increase the effectiveness of the ads through these favorable brand perceptions.

Lastly, ads featuring minority actors can be more effective simply because they are less common, making them stand out and appear more salient to consumers. Having minority representation in ads, therefore, can be one way where firms differentiate their advertisements from others and potentially increase their effectiveness (Pieters, Warlop, and Wedel, 2002; Rosengren et al., 2020). In our experimental study, we find that participants are more likely to perceive the ads featuring minority actors as new and fresh (4.33 vs. 3.61,

¹⁷Detailed results by each condition can be found in Appendix E.

Figure 8: Consumer Perceptions of the Advertised Lender and Advertisement



Note: Error bars denote the 95% confidence interval; * $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$.

$p < 0.001$) and attention-grabbing (3.88 vs. 3.64, $p < 0.001$) compared to ads with White actors. Therefore, ads with minority actors are more effective because of these visual cues, which are positively correlated with the likelihood of loan application and recommendation (Table E4 in Appendix E).

8 Conclusion

Given the growing emphasis on diversity and minority representation, it is crucial for brands to understand the impact of including racial minority actors in ads on advertising effectiveness. In this paper, we find that greater minority representation in TV ads increases the effectiveness of advertising in the mortgage refinancing market. The impact of minority representation is stronger among minority consumers as well as liberal-leaning consumers. A pre-registered experimental study where we directly manipulate the race of the actors in ads shows consistent results with the observational study. Leveraging the heterogeneous results from our observational study and attitudinal questions from the experimental study, we discuss several potential mechanisms that are consistent with our findings.

Our research offers valuable insights for brands seeking to promote racially diverse and inclusive representation in their advertising strategies. Our results suggest that featuring minority actors not only achieves the social goal of increasing minority representation but also leads to higher advertising effectiveness, given the current level of advertising spending. Our results also have important policy implications. Since

featuring minority actors in advertising is an effective strategy for reaching minority consumers, it has the potential to contribute to improving financial inclusion in the mortgage market. In particular, minority consumers have been shown to be less likely to refinance when it is beneficial to do so (Gerardi, Lambie-Hanson, and Willen, 2021; Gerardi, Willen, and Zhang, 2023), resulting in missed interest savings. Increasing minority representation in TV advertising can encourage refinancing among minority consumers and help reduce racial disparities in the mortgage market.

References

- Aaker, Jennifer L, Anne M Brumbaugh, and Sonya A Grier (2000), “Nontarget markets and viewer distinctiveness: The impact of target marketing on advertising attitudes,” *Journal of Consumer Psychology*, 9 (3), 127–140.
- Agarwal, Saharsh and Ananya Sen (2022), “Antiracist curriculum and digital platforms: Evidence from Black Lives Matter,” *Management Science*, 68 (4), 2932–2948.
- Amos, Brian, Michael P McDonald, and Russell Watkins (2017), “When boundaries collide: Constructing a national database of demographic and voting statistics,” *Public Opinion Quarterly*, 81 (S1), 385–400.
- An, Jisun and Haewoon Kwak “Gender and racial diversity in commercial brands’ advertising images on social media,” “Social Informatics: 11th International Conference, SocInfo 2019, Doha, Qatar, November 18–21, 2019, Proceedings 11,” pages 79–94, Springer (2019).
- Aneja, Abhay, Michael Luca, and Oren Reshef “Black Ownership Matters: Does Revealing Race Increase Demand For Minority-Owned Businesses?,” Technical report, National Bureau of Economic Research (2023).
- Babar, Yash, Ali Adeli, and Gordon Burtch (2022), “The Effects of Online Social Identity Signals on Retailer Demand,” *Available at SSRN 4292001*.
- Balakrishnan, Maya, Jimin Nam, and Ryan W Buell (2022), “Differentiating on Diversity: How Disclosing Workforce Diversity Improves Brand Attitudes,” *Available at SSRN 4369758*.
- 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.
- Belloni, Alexandre, Victor Chernozhukov, and Christian Hansen (2014), “Inference on treatment effects after selection among high-dimensional controls,” *The Review of Economic Studies*, 81 (2), 608–650.
- Bertrand, Marianne, Dean Karlan, Sendhil Mullainathan, Eldar Shafir, and Jonathan Zinman (2010), “What’s advertising content worth? Evidence from a consumer credit marketing field experiment,” *The Quarterly Journal of Economics*, 125 (1), 263–306.
- Bhutta, Neil, Aurel Hizmo, and Daniel Ringo (2022), “How much does racial bias affect mortgage lending? Evidence from human and algorithmic credit decisions,” *Available at SSRN 3887663*.
- Bhutta, Neil, Steven Laufer, and Daniel R Ringo (2017), “Residential mortgage lending in 2016: Evidence from the Home Mortgage Disclosure Act data,” *Fed. Res. Bull.*, 103, 1.
- Bryan, Michael “U.S. Voting by Census Block Groups,” (2022) <https://doi.org/10.7910/DVN/NKNWBX>.

- Burnap, Alex, John R Hauser, and Artem Timoshenko (2023), “Product aesthetic design: A machine learning augmentation,” *Marketing Science*.
- Cameron, A Colin and Douglas L Miller (2015), “A practitioner’s guide to cluster-robust inference,” *Journal of Human Resources*, 50 (2), 317–372.
- Chernozhukov, Victor, Denis Chetverikov, Mert Demirer, Esther Duflo, Christian Hansen, Whitney Newey, and James Robins (2018), “Double/debiased machine learning for treatment and structural parameters,” *The Econometric Journal*, 21 (1), C1–C68.
- Davis, Donald R, Jonathan I Dingel, Joan Monras, and Eduardo Morales (2019), “How segregated is urban consumption?,” *Journal of Political Economy*, 127 (4), 1684–1738.
- Deshpandé, Rohit and Douglas M Stayman (1994), “A tale of two cities: Distinctiveness theory and advertising effectiveness,” *Journal of Marketing Research*, 31 (1), 57–64.
- Dew, Ryan, Asim Ansari, and Olivier Toubia (2022), “Letting logos speak: Leveraging multiview representation learning for data-driven branding and logo design,” *Marketing Science*, 41 (2), 401–425.
- Dube, Arindrajit, Jeff Jacobs, Suresh Naidu, and Siddharth Suri (2020), “Monopsony in online labor markets,” *American Economic Review: Insights*, 2 (1), 33–46.
- Dzyabura, Daria and Renana Peres (2021), “Visual elicitation of brand perception,” *Journal of Marketing*, 85 (4), 44–66.
- Ellickson, Paul B, Wreetabrata Kar, and James C Reeder III (2022), “Estimating marketing component effects: Double machine learning from targeted digital promotions,” *Marketing Science*.
- Fossen, Beth L, Donggwan Kim, David A Schweidel, and Raphael Thomadsen (2022), “The role of slant and message consistency in political advertising effectiveness: evidence from the 2016 presidential election,” *Quantitative Marketing and Economics*, 20 (1), 1–37.
- Gerardi, Kristopher, Lauren Lambie-Hanson, and Paul Willen “Racial differences in mortgage refinancing, distress, and housing wealth accumulation during covid-19,” Discussion Paper 21-02, Federal Reserve Bank of Philadelphia Consumer Finance Institute (2021).
- Gerardi, Kristopher, Paul S Willen, and David Hao Zhang (2023), “Mortgage prepayment, race, and monetary policy,” *Journal of Financial Economics*, 147 (3), 498–524.
- Gershon, Rachel and Zhenling Jiang (2022), “Referral Contagion: Downstream Benefits of Customer Referrals,” *Available at SSRN 4214781*.

- Goli, Ali and Simha Mummalaneni (2023), “Gender Diversity on Cable News: An Analysis of On-Screen Talent and Viewership,” *Available at SSRN 4462592*.
- Goli, Ali, Simha Mummalaneni, Pradeep K Chintagunta, and Sanjay K Dhar (2022), “Show and Sell: Studying the Effects of Branded Cigarette Product Placement in TV Shows on Cigarette Sales,” *Marketing Science*, 41 (6), 1163–1180.
- Gordon, Brett R, Robert Moakler, and Florian Zettelmeyer (2022), “Close enough? a large-scale exploration of non-experimental approaches to advertising measurement,” *Marketing Science*.
- Guitart, Ivan A and Stefan Stremersch (2021), “The impact of informational and emotional television ad content on online search and sales,” *Journal of Marketing Research*, 58 (2), 299–320.
- Hartmann, Jochen, Oded Netzer, and Rachel Zalta (2023), “Diversity in Advertising in Times of Racial Unrest,” *Available at SSRN 4466378*.
- Khan, Uzma and Ajay Kalra (2022), “It’s good to be different: How diversity impacts judgments of moral behavior,” *Journal of Consumer Research*, 49 (2), 177–201.
- Khern-am nuai, Warut, Hyunji So, Maxime C Cohen, and Yossiri Adulyasak (2021), “Selecting Cover Images for Restaurant Reviews: AI vs. Wisdom of the Crowd,” *Available at SSRN 3911469*.
- Kingma, Diederik P and Max Welling (2013), “Auto-encoding variational bayes,” *arXiv preprint arXiv:1312.6114*.
- Kingma, Diederik P, Max Welling et al. (2019), “An introduction to variational autoencoders,” *Foundations and Trends® in Machine Learning*, 12 (4), 307–392.
- Kirgios, Erika L, Aneesh Rai, Edward H Chang, and Katherine L Milkman (2022), “When seeking help, women and racial/ethnic minorities benefit from explicitly stating their identity,” *Nature Human Behaviour*, 6 (3), 383–391.
- Lee, Dokyun, Kartik Hosanagar, and Harikesh S Nair (2018), “Advertising content and consumer engagement on social media: Evidence from Facebook,” *Management Science*, 64 (11), 5105–5131.
- Liaukonyte, Jura, Thales Teixeira, and Kenneth C Wilbur (2015), “Television advertising and online shopping,” *Marketing Science*, 34 (3), 311–330.
- Liu, Feng, Young Jo, and Eileen Chen (2022), “Data Point: 2021 Mortgage Market Activity and Trends,” *Consumer Financial Protection Bureau Office of Research Reports Series*, (7).
- Mollica, Kelly A, Barbara Gray, and Linda K Trevino (2003), “Racial homophily and its persistence in newcomers’ social networks,” *Organization Science*, 14 (2), 123–136.

- Morozov, Ilya and Anna Tuchman (2022), “Where Does Advertising Content Lead You? We Created a Bookstore to Find Out,” *Available at SSRN 4308400*.
- Overgoor, Gijs, Gokhan Yildirim, Yakov Bart, and Koen Pauwels (2023), “From Representation to Reception: Evaluating the Impact of Diversity in TV Advertising on Consumer Purchase Intention,” *Available at SSRN 4471248*.
- Pieters, Rik, Luk Warlop, and Michel Wedel (2002), “Breaking through the clutter: Benefits of advertisement originality and familiarity for brand attention and memory,” *Management science*, 48 (6), 765–781.
- Rezende, Danilo Jimenez, Shakir Mohamed, and Daan Wierstra “Stochastic backpropagation and approximate inference in deep generative models,” “International conference on machine learning,” pages 1278–1286, PMLR (2014).
- Rosengren, Sara, Martin Eisend, Scott Koslow, and Micael Dahlen (2020), “A meta-analysis of when and how advertising creativity works,” *Journal of Marketing*, 84 (6), 39–56.
- Sahni, Navdeep S, S Christian Wheeler, and Pradeep Chintagunta (2018), “Personalization in email marketing: The role of noninformative advertising content,” *Marketing Science*, 37 (2), 236–258.
- 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 (2022), “Promoting wellness or waste? evidence from antidepressant advertising,” *American Economic Journal: Microeconomics*, 14 (2), 439–77.
- 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.
- Sudhir, K, Subroto Roy, and Mathew Cherian (2016), “Do sympathy biases induce charitable giving? The effects of advertising content,” *Marketing Science*, 35 (6), 849–869.
- Tian, Zijun, Ryan Dew, and Raghuram Iyengar (2023), “Mega or Micro? Influencer Selection Using Follower Elasticity,” *Available at SSRN 4173421*.
- Tsai, Yi-Lin and Elisabeth Honka (2021), “Informational and noninformational advertising content,” *Marketing Science*, 40 (6), 1030–1058.
- Tuchman, Anna E (2019), “Advertising and demand for addictive goods: The effects of e-cigarette advertising,” *Marketing science*, 38 (6), 994–1022.
- Wang, Yanfei (2022), “Competition and multilevel technology adoption: A dynamic analysis of electronic medical records adoption in US hospitals,” *International Economic Review*, 63 (3), 1357–1395.

- Wang, Yang, Marco Shaojun Qin, Xueming Luo, and Yu Kou (2022), “Frontiers: How support for Black Lives Matter impacts consumer responses on social media,” *Marketing Science*, 41 (6), 1029–1044.
- Yang, Joonhyuk, Jung Youn Lee, and Pradeep K Chintagunta (2021), “Commercial success through commercials? Advertising and pay-tv operators,” *Journal of Marketing Research*, 58 (5), 925–947.
- Zhang, Shunyuan, Dokyun Lee, Param Vir Singh, and Kannan Srinivasan (2022), “What makes a good image? Airbnb demand analytics leveraging interpretable image features,” *Management Science*, 68 (8), 5644–5666.

Appendix A: Sample Selection Criteria

In this appendix, we outline our sample selection criteria and provide the resulting number of observations.

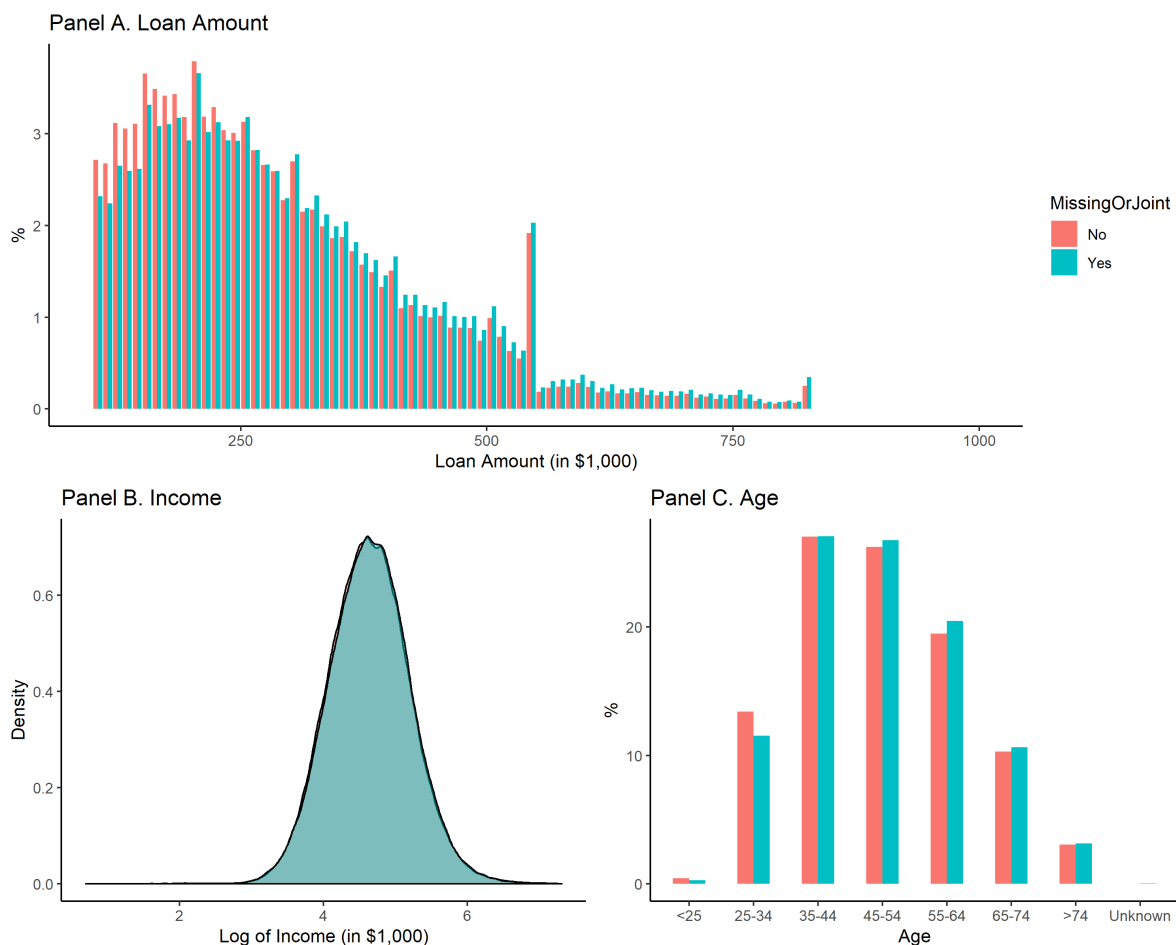
The selection process is as follows:

1. We collect loans that originated directly from lenders for home purchase or refinance purposes during our sample period, resulting in a total of 43,878,666 loans.
2. Among these loans, we select 34,510,763 conventional loans.
3. We then select 19,523,098 refinance loans, which account for about 57% of all conventional loans.
4. Our selection criteria focus on first-lien mortgage loans for owner-occupied, single-family, site-built residential homes. Additionally, we apply two conditions: the loan length must fall within the range of 10 to 30 years, and the loan size should exceed \$100,000 but be less than \$1,000,000. After applying these filters, we are left with 14,849,323 loans.
5. Next, we remove jumbo loans, which are loans where the loan amount exceeds the limit for conforming loans. This step results in a remaining total of 14,408,940 loans.
6. We exclude “exotic loans” with the following characteristics: reverse, open-end, interest-only, non- or negatively amortizing, balloon payment, or those with prepayment penalties. Additionally, we remove a small number of loans with zero or over 20% APRs. After applying these additional filters, we end up with 14,075,414 loans.
7. We further refine our selection by choosing loans whose underlying property locations belong to the top 101 Designated Market Areas (DMAs) we consider, resulting in a remaining sample of 12,814,253 loans.
8. Finally, we exclude loans with missing ethnic/racial information or cases where the borrowers on the same loan belong to different ethnic or racial backgrounds (i.e., joint). Additionally, we remove a small number of loans originating to American Indian or Alaska Native and Native Hawaiian or Other Pacific Islander borrowers. After applying these criteria, we are left with a total of 9,781,509 loans.

Appendix B: Selection on Observables

In this appendix, we present suggestive evidence of similar observable characteristics between individuals with missing or joint ethnic/racial information (Group A) and those with complete or single ethnic/racial information (Group B) using the 2021 data. Figure B1 shows these comparisons. In Panel A, we observe a slight difference in loan amounts, with borrowers from Group A borrowing, on average, \$15,266 more than those from Group B. However, considering the average loan amount of \$290,854, the difference appears relatively small. Panel B shows nearly identical income distributions, indicating that the disparity in loan size is not driven by income. In Panel C, borrowers from Group A have a slightly higher average age than those from Group B. Overall, the observed differences on these characteristics are relatively small, suggesting that any potential selection issue is likely small.

Figure B1: Distribution of Loan Size, Income, and Age by Ethnic/Racial Information Status

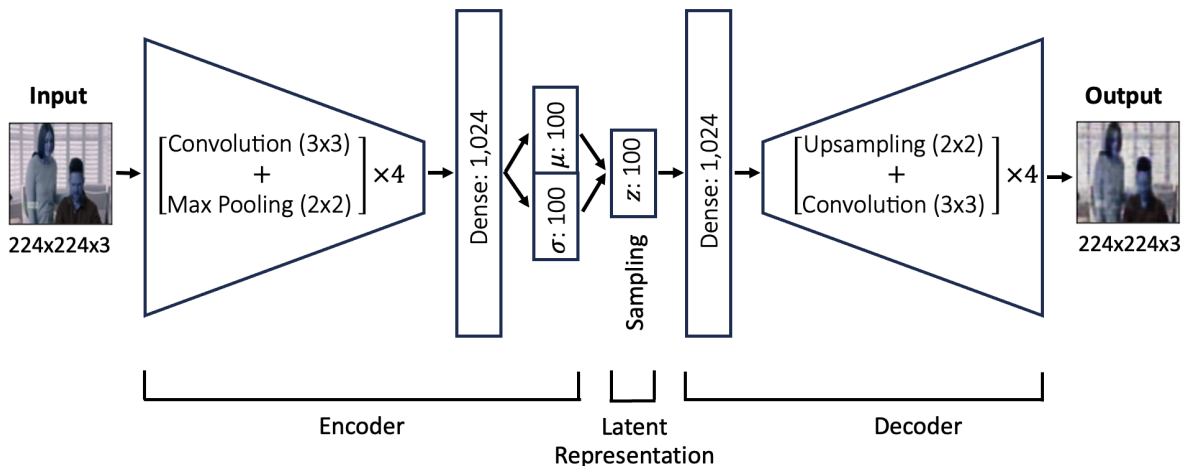


Appendix C: VAE Architecture and Estimation Details

In this appendix, we describe the network structure and the estimation procedure used to estimate Variational Autoencoders (VAEs). Figure C1 shows the network structure employed to represent image data in a lower-dimensional vector space, along with an example of an input image and the corresponding reconstructed image. As discussed in Section 3.1, an input image is represented as a 150,528-dimensional vector ($224 \times 224 \times 3$). This input is passed through the encoder network, which consists of four convolutional layers and two fully-connected/dense layers. In the convolutional layers, we apply commonly used 3×3 filters and 2×2 max pooling operations, gradually increasing the number of channels to 32, 32, 64, and 64. The purpose of this process is to extract meaningful features from the image while reducing its dimensionality. The resulting output from the convolutional layers is then flattened to a 12,544-dimensional vector ($14 \times 14 \times 64$) and fed into the two fully-connected/dense layers. We use the Rectified Linear Unit (ReLU) as the activation function in these two layers. The output of the encoder network is the latent vector, denoted as z , which is sampled from a multivariate Gaussian distribution with parameters μ and σ . This latent vector captures a compact representation of the input image and has a dimensionality of 100. The decoder network is simply the inverse of the encoder network. It takes the latent vector z as input and gradually increases its dimensionality until it matches the original vector space of the input image. The purpose of the decoder network is to reconstruct an image close to the input image based on the compact representation z .

The model is trained with a batch size of 64 for 100 epochs. We employ adaptive learning rates, starting with an initial rate of 0.005. To prevent overfitting, we also implement an early stopping rule.

Figure C1: VAE Network Structure



We implement another VAE to represent the 1,536-dimensional output of the OpenAI’s text embedding model (“text-embedding-ada-002”) into a 100-dimensional latent vector space. Since the input data does not contain spatial information like images, we use a simpler network structure that does not require convolutional

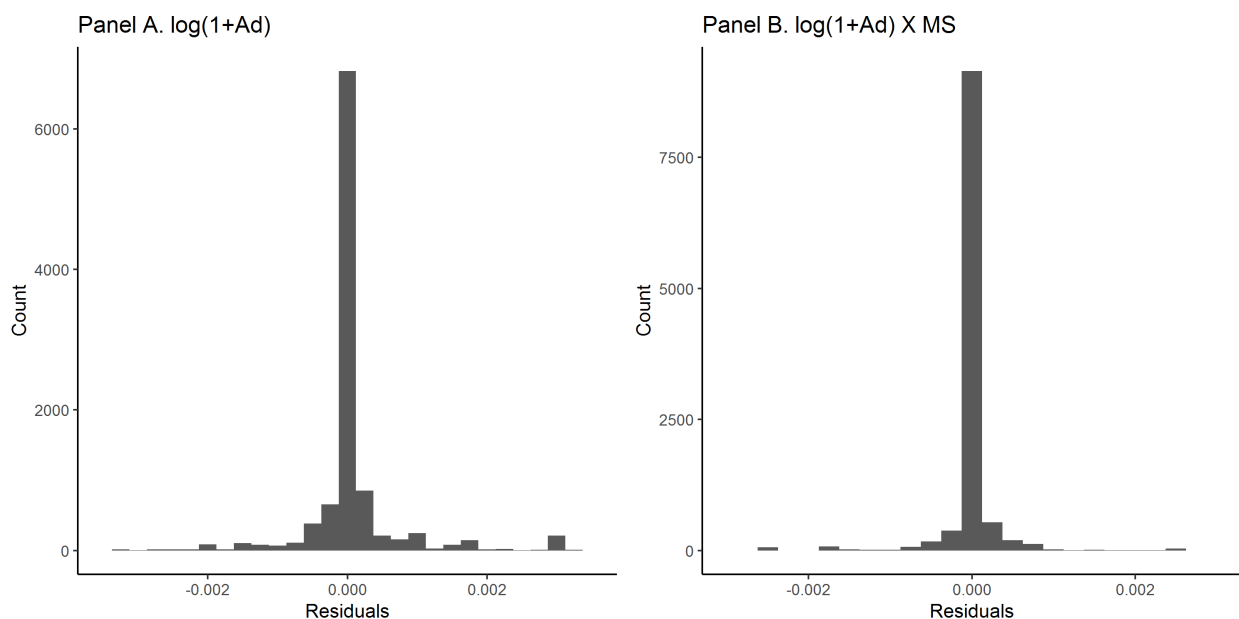
neural networks. Specifically, we use the fully connected layers for training. In the encoder network, we use a sequence of fully-connected layers with dimensions of 1024, 512, 256, and 100. In the decoder network, we perform the inverse operations, gradually increasing the dimensionality of the latent vector until it matches the original input space. The estimation procedure follows a similar approach as before. We train the model using a batch size of 64 for 100 epochs with adaptive learning rates that start with an initial rate of 0.005. To prevent overfitting, we implement an early stopping rule.

Appendix D: Residual Variation in Advertising

One potential concern about controlling for a large set of fixed effects is that there might be little residual variation in advertising. To address this concern, we explore whether we have sufficient variation in advertising after accounting for the fixed effects. Following previous studies (e.g., [Shapiro, Hitsch, and Tuchman, 2021](#); [Tsai and Honka, 2021](#)), we regress $\log(1 + Ad)$ on the lender-DMA and lender-year fixed effects, where Ad represents the total ad spending per capita, including both national and local ad spending. The unit of observation is the lender-DMA-year level. To assess the extent of variation in advertising not explained by the fixed effects, we calculate the ratio of the standard deviation of the residuals to the unconditional mean of ad spending. Additionally, we regress the interaction between $\log(1 + Ad)$ and MS on the same set of fixed effects, where MS denotes the corresponding minority share in ads at the lender-DMA-year level. We then calculate the ratio of the standard deviation of the residuals to the unconditional mean of the interaction term.

Figure D1 shows the distributions of the residual variations. We observe a significant level of residual variation in both Panel A and B. Moreover, the calculated ratios of the standard deviation to the unconditional mean are 0.179 for Panel A and 0.171 for Panel B, suggesting sufficient variation in the data.

Figure D1: Distributions of Residual Variations in Advertising Variables



Appendix E: Additional Experimental Findings

In this appendix, we present additional findings from the experiment. We first examine the heterogeneous effects by race. Specifically, we regress the likelihood of loan application and recommendation on the seven conditions, using the WW condition as the baseline (intercept), within each of the two groups: White consumers and minority consumers. Table E1 presents the results. In columns 1 and 2, we observe that the conditions featuring minority families generally have positive impacts on both DVs among White consumers, although some of the effects are statistically insignificant. On the other hand, in columns 3 and 4, we find that the impacts are greater among minority consumers, and the coefficients are mostly significant. These results align with the findings from our observational study.

Table E1: Heterogeneous Effects based on Each Consumer’s Race

	White Consumers		Minority Consumers	
	(1)	(2)	(3)	(4)
	Application	Recommend	Application	Recommend
White-Black (WB)	0.157 (0.118)	0.200 (0.132)	0.421** (0.184)	0.437** (0.209)
White-Asian (WA)	0.230** (0.138)	0.310** (0.153)	0.300 (0.187)	0.295 (0.212)
Black-Black (BB)	-0.074 (0.117)	0.275** (0.131)	0.324* (0.184)	0.468** (0.209)
Single White (W)	-0.019 (0.117)	-0.070 (0.131)	0.168 (0.182)	0.078 (0.207)
Single Black (B)	0.103 (0.117)	0.341*** (0.132)	0.560*** (0.193)	0.407* (0.219)
Single Asian (A)	0.025 (0.119)	0.205 (0.134)	0.389** (0.186)	0.418** (0.211)
Intercept	4.694*** (0.084)	4.130*** (0.094)	4.462*** (0.129)	4.077*** (0.146)
<i>N</i>	2,025	2,025	761	761
Adj. <i>R</i> ²	0.002	0.006	0.007	0.005

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Next, we investigate the heterogeneous effects by each consumer’s political ideology. To do so, we divide the sample based on the self-reported political leanings, as described in Section 6.2. The first group consists of individuals who lean toward a liberal ideology, while the second group consists of those who lean toward a conservative ideology. Within each group, we regress the likelihood of loan application and recommendation on the seven conditions, using the WW condition as the baseline (intercept). The results are presented in Table E2. In columns 1 and 2, we observe that the conditions featuring minority families consistently have positive and statistically significant effects among liberal consumers. In contrast, when considering columns 3 and 4, which correspond to conservative-leaning consumers, the effects are generally negative, although not statistically significant in many cases. Among other conditions, the BB condition stands out with a large and statistically significant negative impact. Overall, these results show a stronger impact of minority representation among liberal consumers, which aligns with the findings from our observational study.

Table E2: Heterogeneous Effects based on Each Consumer’s Political Leaning

	Liberal Consumers		Conservative Consumers	
	(1)	(2)	(3)	(4)
	Application	Recommend	Application	Recommend
White-Black (WB)	0.344*** (0.132)	0.436*** (0.151)	-0.061 (0.194)	-0.074 (0.214)
White-Asian (WA)	0.481*** (0.128)	0.547*** (0.147)	-0.096 (0.194)	0.096 (0.215)
Black-Black (BB)	0.516*** (0.133)	0.824*** (0.151)	-0.517*** (0.194)	-0.354 (0.215)
Single White (W)	0.044 (0.131)	0.003 (0.149)	-0.069 (0.195)	0.029 (0.216)
Single Black (B)	0.381*** (0.133)	0.586*** (0.152)	-0.148 (0.194)	-0.117 (0.215)
Single Asian (A)	0.371*** (0.131)	0.524*** (0.150)	-0.350* (0.202)	-0.180 (0.223)
Intercept	4.439*** (0.092)	3.860*** (0.105)	5.087*** (0.137)	4.635*** (0.152)
<i>N</i>	1,454	1,454	786	786
Adj. <i>R</i> ²	0.016	0.029	0.006	-0.0003

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

We further examine consumers’ perceptions of the advertised lender and advertisements across the seven conditions. Specifically, we regress each of the six attitudinal measures collected in our experiment on the seven experimental conditions, using the WW condition as the baseline (intercept). The results are presented in Table E3. Overall, we observe that conditions featuring minority families lead to more positive perceptions across all dimensions we consider.

Lastly, we consider how the six attitudinal measures towards the lender and the advertisement correlates with the likelihood of loan application and recommendation. The results are presented in Table E4. In both Panel A and B, we observe that each of the six attitudinal measures is positively correlated with the likelihood to apply for the advertised lender as well as recommend it.

Table E3: Perceptions of the Advertised Lender and Advertisement

	(1)	(2)	(3)	(4)	(5)	(6)
	Broad Options	Cater to Me	Fair Lending	Inclusive	Fresh&New	Attention
White-Black (WB)	0.278*** (0.094)	0.258** (0.104)	0.200** (0.090)	1.857*** (0.100)	0.774*** (0.122)	0.658*** (0.118)
White-Asian (WA)	0.320*** (0.093)	0.350*** (0.103)	0.276*** (0.090)	1.716*** (0.099)	0.708*** (0.121)	0.610*** (0.118)
Black-Black (BB)	0.184** (0.094)	-0.030 (0.104)	0.099 (0.091)	1.421*** (0.100)	0.864*** (0.122)	0.612*** (0.118)
Single White (W)	0.126 (0.093)	0.080 (0.103)	0.158* (0.090)	0.694*** (0.099)	0.187 (0.121)	0.182 (0.117)
Single Asian (A)	0.272*** (0.095)	0.189* (0.105)	0.365*** (0.092)	1.823*** (0.101)	0.723*** (0.123)	0.451*** (0.120)
Single Black (B)	0.391*** (0.094)	0.147 (0.104)	0.303*** (0.091)	1.793*** (0.101)	0.979*** (0.123)	0.814*** (0.119)
Intercept	4.545*** (0.066)	4.562*** (0.073)	4.368*** (0.064)	3.662*** (0.071)	3.517*** (0.086)	3.667*** (0.084)
<i>N</i>	2,796	2,796	2,796	2,796	2,796	2,796
Adj. <i>R</i> ²	0.006	0.005	0.006	0.179	0.035	0.023

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$

Table E4: Impact of Consumer Perceptions on the Likelihood of Application and Recommendation

Panel A: Likelihood of Loan Application						
	(1)	(2)	(3)	(4)	(5)	(6)
	Application					
Broad Options	0.598*** (0.017)					
Cater to Me		0.591*** (0.014)				
Fair Lending			0.585*** (0.017)			
Inclusive				0.405*** (0.015)		
Fresh & New					0.453*** (0.013)	
Attention						0.452*** (0.013)
Intercept	1.906*** (0.082)	1.975*** (0.070)	2.081*** (0.083)	2.738*** (0.080)	2.888*** (0.056)	2.883*** (0.059)
<i>N</i>	2,796	2,796	2,796	2,796	2,796	2,796
Adj. <i>R</i> ²	0.318	0.381	0.286	0.202	0.320	0.296
Panel B: Likelihood of Recommendation						
	(1)	(2)	(3)	(4)	(5)	(6)
	Recommendation					
Broad Options	0.651*** (0.019)					
Cater to Me		0.632*** (0.017)				
Fair Lending			0.671*** (0.020)			
Inclusive				0.456*** (0.017)		
Fresh & New					0.538*** (0.014)	
Attention						0.552*** (0.014)
Intercept	1.223*** (0.094)	1.355*** (0.082)	1.263*** (0.093)	2.057*** (0.090)	2.111*** (0.062)	2.042*** (0.064)
<i>N</i>	2,796	2,796	2,796	2,796	2,796	2,796
Adj. <i>R</i> ²	0.297	0.342	0.295	0.200	0.353	0.347

Notes: * $p < 0.1$, ** $p < 0.05$, *** $p < 0.01$