

NBER WORKING PAPER SERIES

DO LENDERS STILL DISCRIMINATE?
A ROBUST APPROACH FOR ASSESSING DIFFERENCES IN MENUS

David Hao Zhang
Paul S. Willen

Working Paper 29142
<http://www.nber.org/papers/w29142>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2021

We thank Isaiah Andrews, John Campbell, Edward Glaeser, Robin Lee, Ariel Pakes, Adi Sunderam, and Elie Tamer for their continuous advice on this paper. We are further indebted to Alex Bell, Debopam Bhattacharya, Neil Bhutta, Gary Chamberlain, Chris Foote, Andreas Fuster, Camilo Garcia-Jimeno, Aurel Hizmo, Jun Ishii, Larry Katz, Xinwei Ma, Daniel Ringo, Jonathan Roth, Alexander Ufier, Andrei Shleifer, Jialan Wang, attendees of the 2020 System Applied Micro Conference, 2020 System Econometrics Conference, 2020 Econometric Society European Winter Meetings, 2021 Federal Reserve Day-Ahead Conference on Financial Markets and Institutions, 2021 Southwest Finance Association Meetings, 2021 Royal Economic Society Annual Conference, 2021 OSU PhD Real Estate Conference, SOLE 2021 Annual Meeting, 2021 IAEE Conference, and 2021 Financial Intermediation Research Society (FIRS) Conference for their valuable comments. All errors are our own. The views expressed in this paper are solely those of the authors and not necessarily those of the Federal Reserve Bank of Boston nor the Federal Reserve System. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2021 by David Hao Zhang and Paul S. Willen. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Do Lenders Still Discriminate? A Robust Approach for Assessing Differences in Menus

David Hao Zhang and Paul S. Willen

NBER Working Paper No. 29142

August 2021

JEL No. C12,G21,G51

ABSTRACT

Motivated by the assessment of racial discrimination in mortgage pricing, we introduce a new methodology for comparing the menus of options borrowers face based on their choices. First, we show how standard regression-based approaches for assessing discrimination in the menus context can lead to misleading and contradictory results. Second, we propose a new methodology that is robust to these problems based on relatively weak economic assumptions. More specifically, we use pairwise dominance relationships in choices supplemented by restrictions on the range of plausible menus to define (1) a test statistic for equality in menus and (2) a difference in menus (DIM) metric for assessing whether one group of borrowers would prefer to switch to another group's menus. Our statistics are robust to arbitrary heterogeneity in borrower preferences across racial groups, are sharp in terms of identification, and can be efficiently computed using Optimal Transport methods. Third, we devise a new approach for inference on the value of Optimal Transport problems based on directional differentiation. Fourth, we use our methodology to estimate mortgage pricing differentials by race on a novel data set linking 2018–2019 Home Mortgage Disclosure Act (HMDA) data to Optimal Blue rate locks. We find robust evidence for mortgage pricing differentials by race, particularly among Conforming mortgage borrowers who are relatively creditworthy.

David Hao Zhang

Harvard University

dzhang@hbs.edu

Paul S. Willen

Federal Reserve Bank of Boston

Research Department

600 Atlantic Avenue

Boston, MA 02210-2204

and NBER

paul.willen@bos.frb.org

1 Introduction

Whether mortgage lenders discriminate against minority borrowers is an important question both in terms of academic research and in regard to its policy relevance.¹ However, the task of assessing whether lenders discriminate by offering minority borrowers worse prices is complicated by the fact there are two dimensions to mortgage pricing in the US: the interest rate and the upfront fees charged by the lender. In particular, borrowers can choose to pay higher upfront fees (in the industry referred to as paying discount points) in return for lower interest rates. Or, conversely, they can get the lender to pay some of their closing costs in exchange for a higher interest rate. We show that the availability of this choice between a higher upfront fee and a higher interest rate creates a “menu problem” that makes the detection of lender discrimination nontrivial, with the methods implemented in the literature susceptible to false and contradictory results. The menu problem is also broadly applicable to empirical analyses of differences in opportunity under multi-dimensional choice. We propose a novel identification argument and a new procedure for inference to deal with this menu problem, and apply it to a new dataset on mortgage pricing by race.

Many studies find that minority consumers pay higher interest rates compared with observationally similar white consumers in the mortgage market.² While this can be interpreted as evidence that lenders systematically discriminated against minority borrowers by offering them worse pricing on their mortgages, the existence of a rate-point trade-off leads to an alternative explanation: minority consumers were offered the same menu of choices but ended up with higher rate/lower point combinations because, for example, they could not afford to pay points.³ Explaining rate gaps with offsetting point gaps does not rule out structural

¹Since the financial crisis, the Department of Justice has reached settlements of well over \$500 million with lenders that overcharged Black and Hispanic borrowers in violation of the Fair Lending Act, as explained in Bhutta and Hizmo (2020). These settlements include \$335 million with Bank of America (on behalf of Countrywide), \$175 million with Wells Fargo, and \$55 million with JP Morgan Chase. On June 12, 2019, Sen. Elizabeth Warren wrote on Twitter, “For generations, lenders have given African American & Latino families fewer loans at worse terms than similar white borrowers. Tech alone won’t fix the problem. A new analysis found that discrimination is hardwired into lending algorithms. I want answers.” <https://twitter.com/senwarren/status/1138909674781237253>.

²See, for example, Black and Schweitzer (1985), Boehm, Thistle, and Schlottmann (2006), Bocian, Ernst, and Li (2008), Kau, Keenan, and Munneke (2012), Ghent, Hernández-Murillo, and Owyang (2014), Cheng, Lin, and Liu (2015), Bartlett, Morse, Stanton, and Wallace (2021). Relatedly, Munnell, Tootell, Browne, and McEneaney (1996) and Tootell (1996) find that minority borrowers are more likely to be rejected for mortgages; Black, Boehm, and DeGennaro (2003) find that minority borrowers pay higher yield spreads when refinancing their mortgage; and Ambrose, Conklin, and Lopez (2020) find that minority borrowers pay more in broker fees particularly when faced with a white broker.

³The logic the different rates could reflect different options on the same menu was well understood in the industry. Skanderson, Darius, and Butler (2014), for example, write that, “[v]arious interest rates are available to any given borrower, at any given point in time, depending upon his or her needs and preferences [regarding points and closing costs]. The fact that similar borrowers selected different rates on this basis does

disparities between racial groups, but it has policy implications that are very different from those of one in which the lenders themselves are systematically offering minority consumers inferior menus of rates and discount point options. Given data on the borrowers' chosen mortgage rates and points (but not the menus borrowers faced, which are not typically observable), our objective is to examine whether lenders discriminated against Black borrowers in the sense of offering them a distribution of menus that was worse than the one offered to observationally similar white consumers, a practice we call *discrimination in menus*.

Our first contribution is to point out that there exists a nontrivial econometric problem involved with assessing differences in the distribution of menus offered to minority and non-Hispanic white borrowers. Researchers might think we can address heterogeneity in preferences over rates and discount points by racial group by simply controlling for the discount points in their regressions. However, we show that the approach of controlling for rates and discount points can lead to contradictory results. For example, it is sometimes observed that controlling for rate, minority borrowers pay the same closing costs as white borrowers (so it looks like there is no mortgage pricing differential by race), but controlling for closing costs, minority borrowers pay a higher rate (so it looks like there is a mortgage pricing differential by race). This is related to the “reverse regression” problem of Goldberger (1984) which appears naturally in the menu setting. The menu problem also goes beyond this “reverse regression” problem. Even when forward and reverse regressions are consistent, we show that false positives in the sense of detecting discrimination when none exists and false negatives in the sense of failing to detect discrimination when it does exist can still appear. Finally, we show that even a seemingly foolproof comparisons of means—that is, checking if minority consumers on average pay both a higher interest rate and more in lender fees—can still lead to false positives and false negatives if interpreted as evidence of discrimination in menus.⁴ These issues emerge when minority and non-Hispanic white borrowers are likely to make different decisions (that is, have different unobserved preferences) over menu items, which is particularly relevant because the possibility of heterogeneity in preferences across racial groups is usually the reason researchers would seek to control for the choices of discount points in the first place.

The menu problem is important for the mortgage pricing discrimination literature. There are two main methods by which a large literature assesses discrimination in mortgage pricing

not mean that they were treated differently in a sense that should be relevant for fair lending compliance. Nevertheless, such differences in borrower choices can create the appearance of pricing disparities if they happen to be correlated with a prohibited basis and are not captured in the data and accounted for in a fair lending analysis.”

⁴It follows that adjusting by a known range of rate and lender fee trade-offs and then comparing means can be similarly problematic.

given the rate and discount point trade-off. First, Courchane and Nickerson (1997), Bhutta and Hizmo (2020), and Bartlett et al. (2021) look at whether Black borrowers paid more in points conditional on rate.⁵ Second, Woodward (2008) and Bartlett, Morse, Stanton, and Wallace (November 2019) compare the interest rates of minority and white borrowers after adjusting for points using a known range of rate-point trade-offs, and find that minorities paid higher point-adjusted-rates for mortgages. As we discussed, both of the existing methods used in the literature, (1) controlling for rates and (2) comparing mean levels of rates and discount points after adjusting by a known range of slopes, can lead to false positives and false negatives and can contradict one another as an assessment of mortgage pricing discrimination. Indeed, we show that in our FHA sample the choice of either (1) or (2) does lead to apparently contradictory results on whether mortgage pricing discrimination exists. Our robust solution to the menu problem would therefore allow researchers and regulators to assess discrimination in mortgage markets in a more internally consistent and theoretically sound way.

The menu problem also extends well beyond the mortgage setting. Generally speaking, the problem is relevant whenever a researcher wishes to assess disparities in opportunity given data on choices while allowing for heterogeneous preferences across groups. For example, when workers make decisions that trade off wages and hours worked, researchers may wish to assess the extent to which the gender gap in pay may be explained by the choice of hours, as in the model of Goldin (2014). The menu problem implies that popular measures of gender inequality, such as the gender pay gap conditional on hours or even the gender pay gap after adjusting for all relevant average compensating differentials, are not necessarily informative about whether the data on wages can be explained by heterogeneous preferences over hours worked. Our robust metrics can also be useful for this type of problems depending on the institutional details the researcher has access to. Therefore, the problems we point out and the solution we propose may be of broad interest.⁶

As a solution to the menu problem, we propose (1) a new metric for detecting whether

⁵In samples of FHA mortgages, Courchane and Nickerson (1997) find a differential in points paid by race conditional on rate, while Bhutta and Hizmo (2020) do not. One possible reason for this difference is that Bhutta and Hizmo (2020) improved on the earlier literature by using a much larger sample and constructing a more uniform sample of loans, which are improvements we largely adopt. Bhutta and Hizmo (2020)'s study also includes an analysis of the mean levels of rates and points paid conditional on covariates, which is similar in spirit to the second method we discuss. Finally, beyond comparing prices, they find that lenders received more revenue from loans that were made to minorities once points and secondary marketing revenue are added together.

⁶The labor literature has also analyzed differences in worker productivity and wage inequality within the context of structurally specified models of labor supply (e.g. Hwang, Reed, and Hubbard (1992) and Bell (2019)). Our illustrations highlight the role of seemingly innocuous assumptions, such as the restriction of productivity differences to a single component of unobserved heterogeneity, that is typically used in these models.

there exists a difference in the distribution of menus offered to two groups and (2) a new lower bound measure for assessing differences in menus (DIM) for the extent to which one group of consumers would like to switch to another group’s menus. Both metrics are based on pairwise dominance relationships in the data (that is, a mortgage with a lower rate and paying fewer points dominates a mortgage with a higher rate and more points) that can be supplemented by industry knowledge. Based on these pairwise relationships, we ask the question of whether the data *can* be rationalized by a model of equality in menus but heterogeneity in preferences, and if not, we compute an average difference in menus perceived by one group of consumers when switching to another group’s menus. Unlike the existing methodology used in the literature, our metrics are robust to any form of unobserved differences in preferences across borrower groups.

The sample counterparts to both of our metrics can be computed as solutions to optimal transport problems, which are computationally well understood and can be efficiently computed through linear programming. As a technical contribution, we also derive a new approach to uniformly valid inference for the value of optimal transport problems, which we implement for our metrics to distinguish between statistical noise and actual differences in menus. Conventional approaches to inference, such as bootstrapping, fail for optimal transport problems, because the objective function can be non-differentiable (Fang and Santos, 2018). We prove that optimal transport problems are directionally differentiable in the sense of Shapiro (1991) and Fang and Santos (2018). We then apply the asymptotic results of Fang and Santos (2018), which we combine with a Bonferroni correction following Romano, Shaikh, and Wolf (2014) and McCloskey (2017) to address sampling error in the directional derivatives. We show that this approach leads to asymptotically uniformly valid size control for hypothesis testing in the value of optimal transport problems, and test it in a Monte Carlo simulation. Our new approach to inference in optimal transport may be useful for other researchers who wish to conduct inference on the value of optimal transport problems, many of which are described in Galichon (2016).

Empirically, we use our metrics to assess mortgage pricing discrimination in the 2018–2019 Home Mortgage Disclosure Act (HMDA) data matched to Optimal Blue rate locks. We show that we can detect a difference in menus offered by the same lender in the same county and within narrow covariate groups for conforming mortgages for both Black and Hispanic borrowers relative to non-Hispanic white borrowers. Furthermore, we show that on average Black borrowers getting conforming mortgages would be willing to increase their interest rate by at least 1.7 basis points in order to switch to the menus of non-Hispanic white borrowers. Similarly, Hispanic borrowers are on average willing to pay 1.5 basis points more in interest rate in order to switch menus with non-Hispanic white borrowers. Our finding that

racial differences in lender pricing remains relevant for conforming mortgages is consistent with Bartlett et al. (2021), although the amount of interest rate discrimination we detect is smaller in magnitude. On the other hand, we do not detect interest rate discrimination in FHA mortgages, which is consistent with Bhutta and Hizmo (2020). Within conforming mortgages, the differences in menus we detect are particularly concentrated among borrowers with lower loan-to-value (LTV) ratios and are not explained by multi-product discounts or loan originator compensation.⁷ The fact that the mortgage pricing discrimination we detect is concentrated among the more creditworthy conforming mortgage borrowers is consistent with the less risky non-Hispanic white borrowers being more likely to be offered discounts during the search and negotiation process.⁸

In our empirical application, we focus on the important and basic question of whether we can detect differences in the rate and discount point menus faced by observably similar non-Hispanic white and minority mortgage borrowers. The answer to this question is important because one needs to ascertain whether there *exists* disparate treatment before considering its causes. Our use of the word “discrimination” to refer to this disparate treatment conditional on observables is in line with the recent literature and is justified by the unusual institutional details and the regulatory framework developed around this market. First, as explained in Bartlett et al. (2021), the US fair-lending law imposes a requirement of no pricing differentials by race conditional on observables in our setting, so our results are naturally interesting from a regulatory perspective. Second, for the types of mortgages we focus on, lenders are insured from the risk of default by either the GSEs or the FHA, who compensate investors for any losses of principal but allows them to benefit from the more favorable prepayment risk of minorities. Therefore, in contrast to some stereotypes, mortgages from Black and Hispanic borrowers are likely significantly more valuable than those of observably similar white borrowers due to their lower prepayment risk, as shown in the simulations of Kau, Fang, and Munneke (2019) and in analyses of the actual mortgage-backed securities prices in Gerardi, Willen, and Zhang (2020). As a result, even if lenders did use unobservables that are correlated with race to price for expected loan performance in a possibly illegal manner,

⁷Relatedly, Ambrose, Conklin, and Lopez (2020) find in a pre-2008 sample period that the more credit-worthy minority borrowers pay higher broker fees than observationally similar non-Hispanic white borrowers, even though their default risks are similar. Regulation Z of 2011 now forbids mortgage brokers from varying the fees they charge across borrowers except as a function of the loan amount, shutting down this particular channel of disparity, though similar forces may be play in the disparities we uncover.

⁸Negotiation in mortgage markets is common. Studies that look at the search and negotiation process in mortgage markets include Allen, Clark, and Houde (2014), Allen, Clark, and Houde (2019), and Bhutta, Fuster, and Hizmo (2019). A reason why minority borrowers may behave as if they have higher search costs is in Agarwal, Grigsby, Hortaçsu, Matvos, Seru, and Yao (2020), where borrowers internalize a higher probability of rejection. Price discrimination by such effective search costs would violate the US fair-lending law, to the extent it results in disparate impact by race within lender, according to Bartlett et al. (2021).

it would be unlikely to justify the unfavorable pricing to minorities we find.

The rest of this paper is structured as follows. Section 2 explains the motivation of our paper by exploring why heuristic approaches to analyzing discrimination in menus may be misleading. It also provides intuition for our approach. Section 3 formally defines our metrics for assessing discrimination in menus. Section 4 describes a methodology for conducting inference on our metrics. Section 5 shows our data and empirical results. Section 6 concludes.

2 The menu problem

2.1 Problem description

In this section we introduce the menu problem and discuss why intuitively appealing approaches for assessing discrimination in menus may be misleading. By way of background, there are two dimensions of pricing for mortgages in the United States: an upfront fee/discount points and the interest rate, where each point is customarily worth 1 percent of the loan amount. Consumers can have the option of picking a particular rate and discount point combination that best suits their preferences and financial constraints. We plot those choices from an example rate sheet in Figure 1. We also present a screenshot illustrating this trade-off from an online mortgage price comparison service in Appendix Figure A.1. In particular, borrowers can pay discount points to reduce their interest rate or receive money from the lender to help cover their closing costs by getting lender credit (paying negative points). The sense in which we think about lender discrimination in menus, then, is for minority borrowers to receive a worse rate-point schedule than white borrowers.

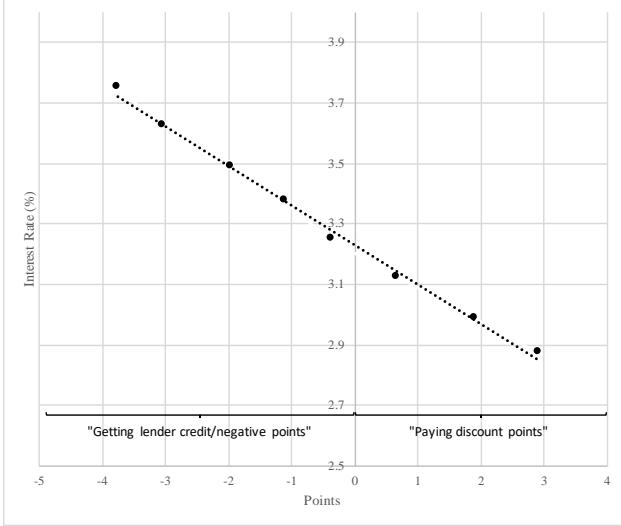


Figure 1: An example set of menu items from a lender rate sheet.

Note: the data for this figure is taken from a Loan Sifter rate sheet for a \$300k conforming, new purchase mortgage in the Boston area with a loan-to-value ratio of 90% and a FICO score of 750.

In practice, the researcher often observes the distribution of borrower choices x but not the underlying menus \mathbf{m} where $x \in \mathbf{m}$ is chosen from. The menu problem then emerges as the problem of inference on the extent to which a matched group of borrowers who are by construction observationally similar in terms of covariates faced the same distributions of menus. More specifically, testing for equality of menus can be written as testing the null hypothesis \mathcal{H}_0 that the distributions of menus being offered to both groups are equal. That is, suppose $\mathbf{m}_1 \sim \mathbf{M}_1$ for borrowers in group 1 and $\mathbf{m}_2 \sim \mathbf{M}_2$ for borrowers in group 2, the menu problem is the hypothesis testing problem where:

$$\mathcal{H}_0 : \mathbf{M}_1 = \mathbf{M}_2, \quad (1)$$

$$\mathcal{H}_1 : \neg \mathcal{H}_0. \quad (2)$$

An important stated objective of the applied literature on mortgage pricing discrimination, and indeed the literature on differences in opportunity more broadly, is assessing equality in menus as specified in Equation (1). Nevertheless, existing methodology used in this literature tends to fall short of being able to credibly assess inequality in menus. We point out a series of difficulties associated with using existing methodology to test $\mathcal{H}_0 : \mathbf{M}_1 = \mathbf{M}_2$, which we call the “menu problem.”

One natural approach to assessing whether lenders offered minority and non-Hispanic white borrowers different menus is to control for one dimension of the menu, that is, conditional on the distribution of covariates, to estimate whether minority borrowers who received

the same interest rate as white borrowers paid more in discount points. This approach was used in Courchane and Nickerson (1997), Bhutta and Hizmo (2020), and Bartlett et al. (2021). A problem with this approach, however, is that it can lead to contradictory estimates depending on whether the researcher chooses to control for rates or points.⁹ The situation in Figure 2 shows that it is possible for a regression of points on rate to show no discrimination against minorities while a regression of rate on points shows discrimination with the same example data. In this figure, we represent example data from minority and white borrowers using black and white dots, respectively; regression line by the dashed line; and the difference to the regression line by the solid arrows. Figure 2a shows that a regression of points on rate and borrower race would show a zero coefficient for minorities, with the two arrows balancing each other out. On the other hand, Figure 2b shows that, using the same data, a regression of rate on points would instead give a positive coefficient for minorities. This sort of contradiction is not particular to the linear regression case, and as we show in Appendix Figure A.2, it can appear with general conditional expectations.

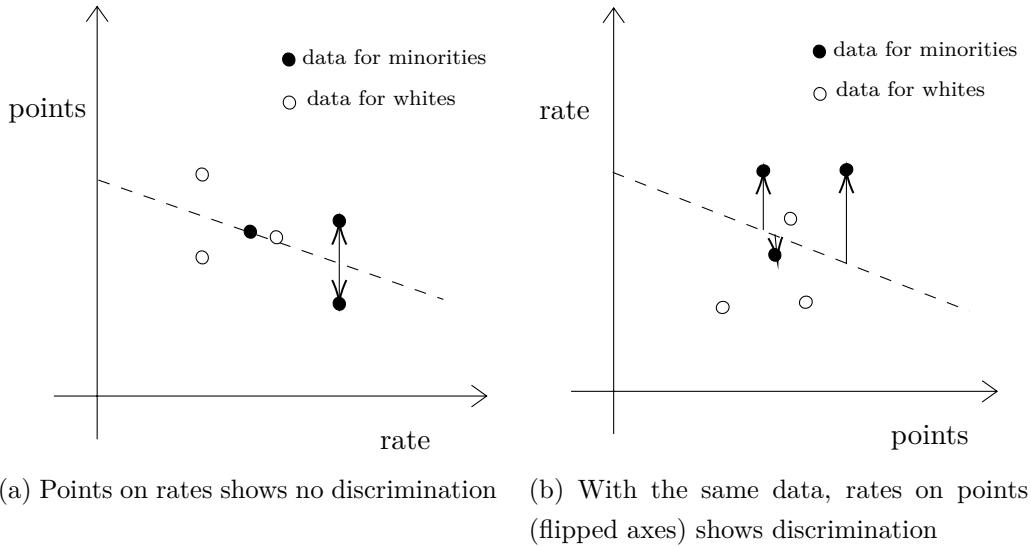


Figure 2: How the choice of which menu dimension to control for can lead to contradictory findings of discrimination

Note: Figures 2a and 2b illustrate example situations where the forward regression of rates on points and the reverse regression of points on rate can lead to different results. The dashed lines represent the regression line, and the arrowed lines represent the distance of minority borrowers' choices to the regression line.

This possibility is not merely a theoretical curiosity. Figure 3a plots the predicted values of discount points paid in a regression with race dummies interacted with levels of interest

⁹Bhutta and Hizmo (2020) also runs lender-specific regressions, and Bartlett et al. (2021) also controls for lender-specific slopes in a robustness check. However, the problems identified here exists as long as there are within-lender heterogeneity in menus, which we find is prevalent in Section 2.2.

rate spread rounded to the nearest quarter of a percent for the FHA mortgages in our sample. It shows that, conditional conditional on the rate paid, Black and white borrowers pay similar numbers of points. On the other hand, using the same data, Figure 3b shows the predicted values of interest rate paid in a regression with race dummies interacted with levels of discount points rounded to the nearest half a point. It suggests that, conditional conditional on the number of points paid, Black borrower pays higher rates than white borrowers. Thus, the regression approach can give strikingly clear, though ultimately misleading, results.

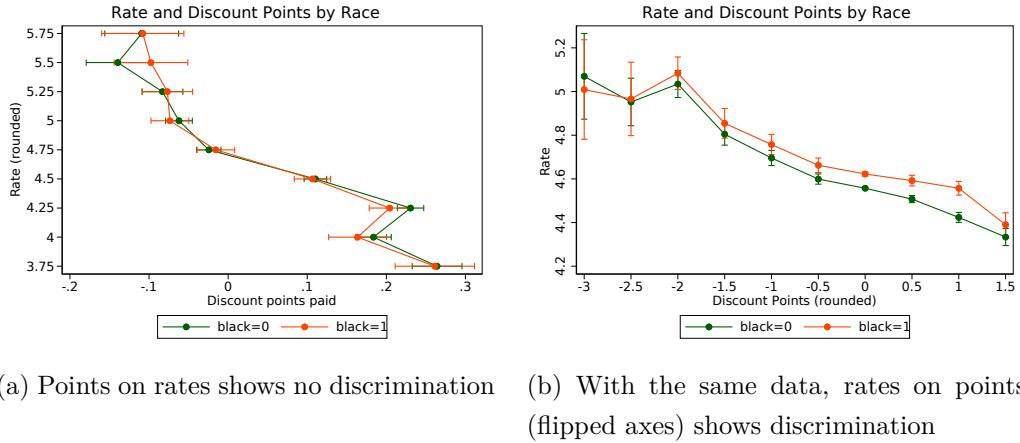


Figure 3: An empirical example for the choice of which menu dimension to control for can lead to contradictory findings of discrimination

Note: Figure 3a plots the predicted values from forward regression of discount points on rates rounded to the nearest quarter interacted with race, with 95% confidence intervals. Figure 3b plots the reverse regression of rate on discount points rounded to the nearest one-half interacted with race, with 95% confidence intervals. The data used is our covariate matched FHA sample as described in Section 5.

This situation in Figures 2 and 3 can be viewed as a version of the reverse regression problem of Goldberger (1984). This problem emerge naturally in the menus setting, since when borrowers are choosing from a menu of options it is in theory not clear which direction the regression should be run. The problem with testing equality in menus goes beyond the reverse regression problem, as there can be false positives and false negatives even when the forward and reverse regressions are consistent (and even when a simple comparison of means is consistent) which we will discuss in the rest of this section.

Figure 4 shows how even when forward and reverse regressions consistently detect discrimination or no discrimination, the heuristic of controlling for one dimension of the menu can lead to false positives and false negatives if interpreted as a test for equality in menus as in Equation (1). In this figure, we represent example data from minority and white borrowers using black and white dots, respectively, and menus by dotted lines where each dotted

line is one potential menu from which a borrower may draw. The left panel of Figure 4a shows a false positive situation in which minority borrowers paid more in rate, controlling for points, and more points, controlling for rate, even though lenders offered both minority and white consumers the same distribution of menus. The only difference in borrower behavior by group is that minority consumers chose to pay fewer points on every menu. That is, even though $\mathbf{M}_1 = \mathbf{M}_2$ in reality, as represented by a common set of dotted lines facing both groups, it appears as if minorities are worse off controlling for either dimension of the menu. In the right panel of Figure 4b, we illustrate a false negative situation in which minority consumers paid the same rate conditional on points but faced a worse distribution of menus, since the bottom menu (the most advantageous menu) was offered only to white borrowers while the second-to-bottom menu was offered only to minority borrowers. Furthermore, Figure 4b can be constructed in a way such that the variance of the rate and discount points paid are equal, such that forward and reverse regressions both give the same false negative but there does exist discrimination in menus. That is, in the situation of Figure 4b, $\mathbf{M}_1 \neq \mathbf{M}_2$ in reality, but controlling for either dimension of the menu the researcher would find no difference between the two groups.

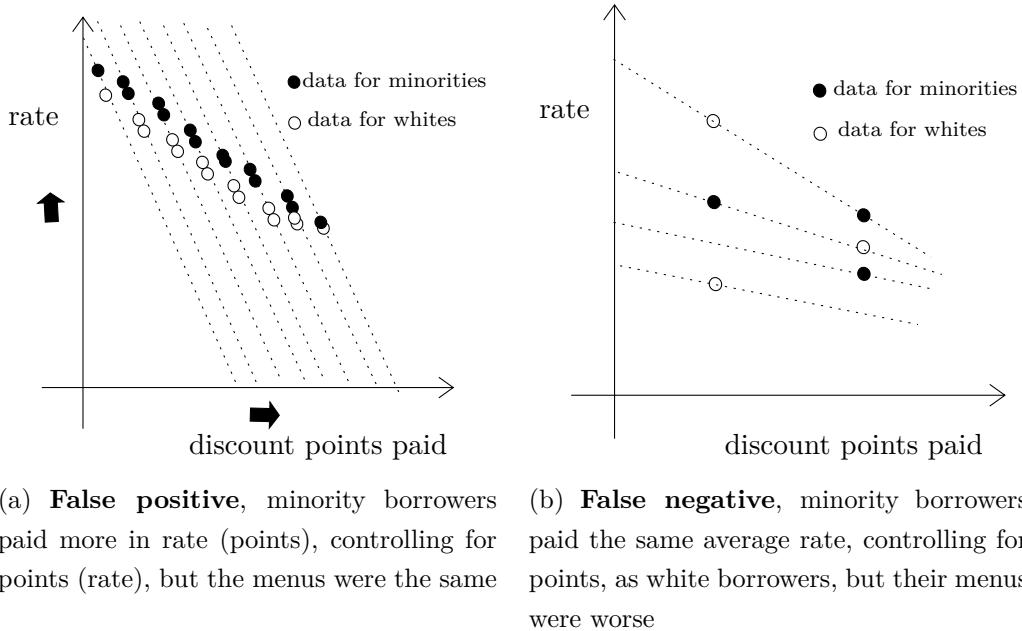
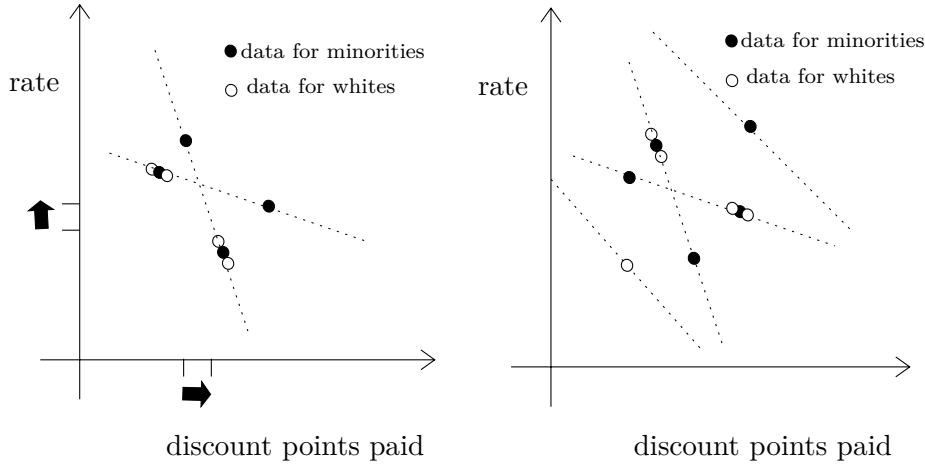


Figure 4: False positives and false negatives from controlling for one direction of the menu

Note: Figures 4a and 4b illustrate example situations where even when forward and reverse regressions are consistent, there can still be false positive and false negative estimates of mortgage pricing differentials by race. Figure 4a represents a situation where minority borrowers consistently choose to pay fewer points than white borrowers. A regression analysis would show that minority borrowers paid more in rate controlling for points, and more in points conditional on rate. However, in this example they faced the same distribution of menus as white borrowers. Figure 4b represents a situation where forward and reverse regressions can both show no difference, but where one white borrower's menu is better than one minority borrower's.

A second intuitively appealing approach for assessing discrimination in menus is to compare means, thus avoiding the problem of having to estimate menu slopes from the data. In other words, the researcher may wish to check if minority consumers paid more on average in both rates and discount points, such that they are disadvantaged in both dimensions compared with observationally similar white borrowers. A variation of this approach is to take a pre-defined range of rate-point trade-offs as the slope estimated from external sources, which is done in Woodward (2008) and Bartlett et al. (November 2019). While this avoids the problem that regressions may incorrectly estimate menu slopes, it can still lead to false positives and false negatives when slopes are not constant across menus, thus breaking the assumption (1) that all menus share the same shape. This is a realistic problem, because we know from rate sheet data that an unobserved heterogeneity in slopes does exist in the mortgage setting, as the rate-point trade-offs do vary substantially across lenders and over time (Figure 9). Figure 5a illustrates how, when slopes differ across menus, a false positive in which minority consumers pay more on average in terms of both rates and points but faced the same distribution of menus as white borrowers can occur ($\mathbf{M}_1 = \mathbf{M}_2$). And Figure 5b illustrates a false negative possibility in which minority borrowers paid the same average rates and points as white borrowers but did face worse menus ($\mathbf{M}_1 \neq \mathbf{M}_2$).



(a) **False positive**, minority consumers paid more on average in both rates and points, but their menus are the same

(b) **False negative**, minority consumers pay the same average rate and points as white borrowers, but their menus were worse

Figure 5: False positives and false negatives from checking if minority borrowers paid more on average in both rates and points

Note: Figure 5a illustrates an example situation where minority borrowers pay more in mean rates and points compared to white borrowers, but both groups of borrowers faced the same distributions of menus. Figure 5b illustrates an example situation where minority borrowers pay the same mean rates and points compared to white borrowers, but faced worse menus than white borrowers.

The mechanism for when a comparison of means would lead to false positives in the case of Figure 5a is that in the example, minority borrowers respond less to differences in the slopes of the rate-point menus compared with white borrowers. This possibility is empirically relevant, because there are two directions in which constraints can drive borrower choices of rates and closing costs. First, if borrowers are cash constrained when getting the loan, they may need to get a lower closing cost mortgage (pay fewer discount points) regardless of what rate-point trade-offs lenders offer. Second, if borrowers are debt-to-income (DTI) constrained, they may need to pay more points to buy down the rate to increase their borrowing limit.¹⁰ Therefore, the finding that minority borrowers on average pay more than white borrowers may simply reflect the fact that minority borrowers are more constrained in their choices, which is a form of “disadvantage” that is not necessarily due to lenders discriminating against them by offering them different menus.

While simple in hindsight, the situation of Figure 5a also illustrates how a seemingly innocuous assumption that differences in menus faced by different agents can be summarized by a single monotonically additive term is actually a strong assumption when there are

¹⁰See, for example: <https://www.thetruthaboutmortgage.com/dti-debt-to-income-ratio/>.

differences in both menu intercepts and slopes in reality. Perhaps due to tractability, such an assumption is popular in models of labor productivity (e.g. Hwang, Reed, and Hubbard (1992), Bell (2019)), where agents are assumed to face structurally defined menus of trade-offs indexed by a single productivity term, and estimation proceeds by devising a method to consistently estimate the distribution of that term. Our illustration implies that by ignoring the multidimensional way unobserved heterogeneity in trade-offs can enter into the model (for example, in terms of both levels and slopes), a researcher can be led to conclude that, for example, one group of workers is more productive than another group when one group is simply more flexible than another group in terms of their preferences over wage-amenities trade-offs. Therefore, whether a single index assumption is reasonable or not depends on how much of a concern the situation illustrated in our Figure 5a would be in the specific empirical context in which they are applied.¹¹

Fundamentally, the shortcomings of existing empirical methodologies, when applied to the menu problem, can be summarized as a combination of omitted variables bias and misspecification bias. Roughly speaking, the assumptions underlying the heuristic approaches of assessing discrimination in menus are that (1) all menus share the same shape, with unobserved heterogeneity in menus being due to an additive error term, and (2) that this shape can be correctly estimated/inferred from other data. The omitted variables problem stems from the fact that the preferences of consumers, since they are unobserved, can lead to bias when the slopes of menus are estimated from data by controlling for them, thus breaking implicit assumption (2). This is the main problem with simply controlling for some dimensions of the menu, as illustrated in Figures 2 and 4. While it is less recognized in the mortgage pricing discrimination literature, much work in labor productivity estimation is indeed aimed at estimating the slopes of trade-offs and addressing the omitted variables problem in (2). A more subtle problem is that correctly estimating the average slopes of trade-offs is not sufficient for assessing equality in menus. Models that assume away heterogeneity in slopes and interpret residuals to average trade-offs as differences in menus can generate a mis-specification error when heterogeneity in slopes does exist in reality, which can lead to misleading inference by breaking assumption (1). This is the problem illustrated in Figure 5.

¹¹Another setting where a single (log-)additive productivity assumption is popular is in the production function literature (e.g. Olley and Pakes (1996), Levinsohn and Petrin (2003), Ackerberg, Caves, and Frazer (2015)). Therefore, while methods from the production estimation literature are useful for estimating average returns to capital and labor, they are also problematic if applied as a test of equality in menus as in our mortgage context.

2.2 The relevancy of the “menu problem” in the mortgage setting

As we illustrated in Figure 3, the possibility for contradiction in regression estimates of differences in menus is an empirically relevant problem in the mortgage discrimination setting in the sense that it is present in our data. More fundamentally, the problems illustrated in Section 2.1 depend on (1) the existence of unobserved heterogeneity in menus and (2) in the differences in preferences between borrower racial groups. We argue that these issues are first order in the mortgage setting.

Empirically, we find that there exists large unobserved heterogeneity in menus across borrowers within narrow covariate groups and from the same lender, suggesting that existing methods that condition on these observables cannot eliminate the problems we point out. In particular, we use the data described in Section 5 to sample pairs of borrowers who got a mortgage with the same lender, on the same day, with the same mortgage type (Conforming or FHA) and the same loan-level price adjustment (LLPA) and loan amount group, and compute the implied slopes given by their rate and point choices, which we plot in Figure A.3. Within this sample and among borrowers whose rate and point choices were not identical, 51% of Conforming pairs and 52% of FHA pairs had choices that dominated one another. Furthermore, Figure A.3 also shows that the implied exchange rate between them have large dispersion. Finally, R^2 of our regressions of points on rate in Tables A.3, as well as those of the literature, are well below 1. These observations all imply that there exists large unobserved heterogeneity in the menus, even conditional on narrow covariate and lender groups.

The reason that researchers have emphasized the role of discount points in studies of mortgage pricing discrimination is due to a presumption that Black borrowers have different preferences over menus items compared to White borrowers, perhaps due to differences in liquidity or borrowing constraints. Indeed, we observe Black and White borrowers making different choices in the data. In particular, Table 3 of our empirical application shows that Black borrowers pay more points than similar White borrowers for Conforming mortgages and fewer points than similar White borrowers for FHA mortgages. However, as we point out, there is a difficulty in attributing this difference in choices to a difference in the menus presented to Black and White borrowers or the heterogeneity in preferences between the racial groups. The advantage of our methodology is that it allows us to say that, under weak assumptions about preferences and menus, some of the racial difference in choices can *only* be rationalized by a difference in the distribution of menus facing Black and White borrowers among Conforming mortgage borrowers.

3 Robust Metrics for Assessing Discrimination in Menus

In this section, we define our robust metrics for assessing discrimination in menus. We consider equality in menus to mean equality in the distribution of menus offered to Black and non-Hispanic white borrowers, conditional on their observables. First, we present intuition for our approach in Section 3.1. We then specify our model of borrower choice more formally in Section 3.2. We keep our model fairly simple; menus are treated simply as a collection of items. Then, we define a direct test metric for equality in menus in Section 3.3 and a more welfare-relevant differences in menus metric for whether one group of consumers would like to switch to another distribution of menus in Section 3.4. We discuss the power of our identification results in Section 3.5. We leave inference on these metrics to Section 4. To increase clarity, we focus our exposition on the mortgage setting with two choice dimensions. Nevertheless, our general approach can be applied to contexts with more than two dimensions and may be useful in other settings where similar restrictions can be made.

3.1 Intuition

Our metrics for assessing discrimination in menus are based on whether the data *can* be rationalized by a model in which all groups of borrowers faced the same distribution of menus. Ignoring for now sampling error to build intuition, the common thread in the false positive situations of Figures 4a and 5a is that there does exist a possible common distribution of menus that rationalizes the choice distributions of both minority and white borrowers, in the sense of there being a possible one-to-one match between minority and white borrowers where within each match both of the borrowers' choices could have come from the same menu. This is the criteria we use for assessing equality in menus.

By construction, our metrics are robust to false positives, since we will detect discrimination in menus only when there is no way to rationalize the data under the assumption of a common distribution of menus, regardless of the nature of any preference heterogeneity between the two groups of borrowers. Our methodology can also detect discrimination in menus where the existing heuristic approaches to the menu problem fail to do so. In particular, we illustrate in Figure 6 how the situation of Figures 4b, in which regressions controlling for either rates or points, would show a false negative, but the data fail our one-to-one matching condition under the assumption that borrowers with pairwise strictly dominated choices (that is, paying more in terms of both rates and points) could not have shared menus with one another. In other words, there is no way to construct a common distribution of menus for minority and white borrowers that explains the data. Analogously, the false negative example from comparing means in Figure 5b also fails our criterion.

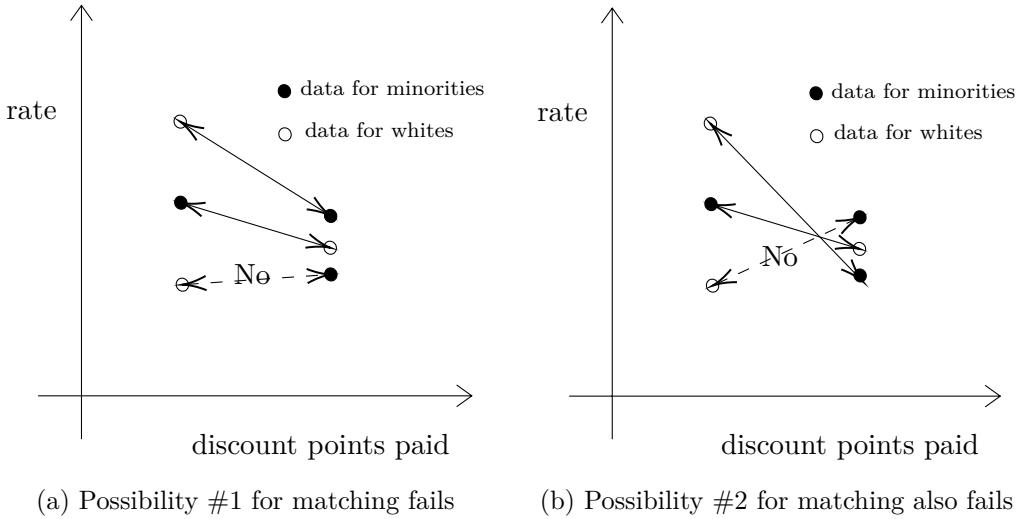


Figure 6: How data from Figure 4b fail a “perfect matching” condition

Note: Figure 6 shows how data from Figure 4b, which can produce false negatives in linear regressions, fails our “perfect matching” condition that each minority borrower should be matchable to a white borrower one-to-one.

To summarize, we define new metrics based on whether a set of preferences can rationalize the data under equality in menus, which are robust to the false positives. Furthermore, in some situations, such as those in Figures 4b and 5b, our metrics can detect discrimination when existing, heuristic approaches fail to do so. Nevertheless, a drawback of our approach is that it still leaves some possibility for false negatives, because the mere existence of a set of preferences that explains the data under equality in menus does not mean that it is the true set of preferences. This is a weakness compared with experimental approaches that may allow the researcher to directly observe menus, but the advantage of our approach is that it requires only data on outcomes and few assumptions on the data-generating process. Furthermore, our metrics are *sharp* in the partial identification sense that false negatives can only occur when there exists a common distribution of menus rationalizing the data that cannot be ruled out under the assumptions given.

3.2 Model

We now define our model more precisely. A menu item has values over k dimensions of attributes,¹² which we encode by $x \in \mathbb{X} \subset \mathbb{R}^k$. A menu $\mathbf{m} \subseteq \mathbb{X}$ is a set of such menu items that are presented to the borrowers. When borrower i is presented with a menu \mathbf{m} , we observe them making a choice that maximizes their utility over menu items $u_i(x)$. That

¹²In our context, the two dimensions of mortgage pricing are interest rates and discount points.

is, we observe choices x_i where:

$$x_i \in \arg \max_x \{u_i(x) : x \in \mathbf{m}\}. \quad (3)$$

To keep the distribution of menus Lebesgue measurable and to implement the inference procedure of Section 4, we make the simplifying assumption that the set of items available to choose from is finite:

Assumption 1. (*Finiteness*) *The set of possible menu items, \mathbb{X} , is finite.*

Under Assumption 1, we can consider a probability distribution over possible menus $\mathbf{m} \sim \mathbf{M}$. This setup is fairly general and follows from consumers having standard (that is, complete and transitive) preferences over menu items.¹³

3.3 A robust test for inequality in menus

Suppose borrowers with a similar distribution of covariates in groups 1 and 2 face menus $\mathbf{m}_1 \sim \mathbf{M}_1$ for borrowers in group 1 and $\mathbf{m}_2 \sim \mathbf{M}_2$ for borrowers in group 2. The researcher wishes to compare the distribution menus across two groups of borrowers. More specifically, testing for equality of menus can be written as testing the null hypothesis that the distributions of menus being offered to both groups are equal. That is:

$$\mathcal{H}_0 : \mathbf{M}_1 = \mathbf{M}_2, \quad (4)$$

$$\mathcal{H}_1 : \neg \mathcal{H}_0. \quad (5)$$

To go from data on choices to statements about menus, we place restrictions on the choices that could have been plausibly made from the same menus in terms of borrower preferences. For mortgages, it is plausible to assume that paying more in both interest rates and discount points is a dominated choice (and indeed would not be offered as a choice by the loan originator), which is the intuition we use in Figure 6 to reject equality in menus in that situation. We formalize this as Assumption 2:

Assumption 2. (*Dominance*) *Paying more in rates and points is dominated. More formally, let $x_1 = [r_1, y_1]$, $x_2 = [r_2, y_2]$, where r_1, r_2 represents rates and y_1, y_2 represents points. Then if $r_1 > r_2, y_1 \geq y_2$, or $r_1 \geq r_2, y_1 > y_2$, $u_i(x_1) < u_i(x_2), \forall i$.*

¹³Our Equation (3) that consumers maximize utility over menu items does rule out more behavioral representations of preferences over menus, such as in Gul and Pesendorfer (2001) and Ellis and Masatlioglu (2021) where the existence of some menu items may “tempt” consumers to change their rankings of other menu items. In that case, our statistical test of inequality in menus would still be valid, but the interpretation of our more welfare-relevant differences in menus (DIM) metric would be nuanced.

We illustrate in Figure 7 the restrictions on the observed choices that may come from the same menu under Assumption 2. The observed choice of the borrower, shown as the black dot, implies that they did not have the lower-left dashed quadrant available on their menu, since otherwise they would have chosen it. Similarly, any choice in the upper-right quadrant could not have been from the same menu as the choice indicated, as that agent would have an incentive to switch to that choice.

Note that while Assumption 2 is defined in the form of preferences of borrowers, it could have been alternatively formulated in the sense of dominated choices being unlikely to be *offered* to borrowers, which would have led to the same restrictions. Importantly, this implies that our test of equality in menus would hold even if borrowers are behavioral such that Assumption 2 fails, as long as lenders' menus fall within the range we specify in Assumption 3. The main purpose of Assumption 2 is to allow us to compute a welfare-relevant metric for assessing differences in menus (DIM) later on.

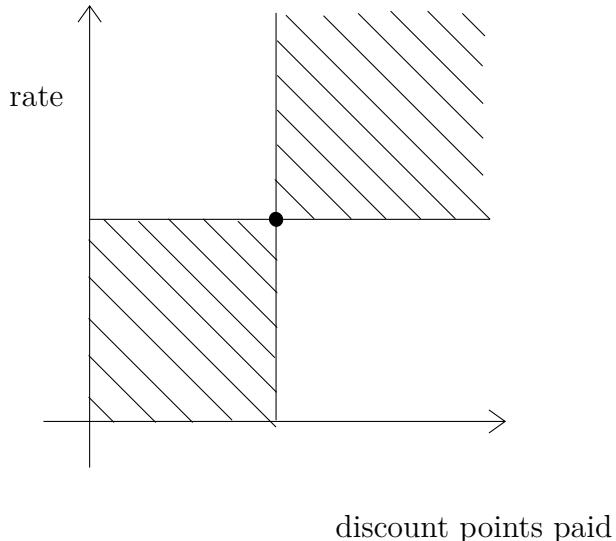


Figure 7: Restriction on what cannot lie on the same menus from dominance.

Note: this figure illustrates the impact of a restriction on what cannot lie on the same menus based on a assumption that dominated choices are not taken/offered. The dashed region represents what could not have lied on the menu of the borrower making the choice illustrated by the black dot.

While Assumption 2 is sufficient for rejecting equality in menus in the example situation of Figure 6, in our empirical application it is too weak of a restriction to be informative by itself. For our empirical analyses, we further adopt the industry rule of thumb of Bartlett et al. (November 2019) that each point paid reduces the interest rate on a mortgage by one-eighth to one-fourth for conforming mortgages, with an expanded range for FHA mortgages. This is an assumption about menus rather than about preferences, which we formalize as

Assumption 3:

Assumption 3. (Restriction on Menus) *In menus, each point paid reduces the rate by between $[a, b]$. More formally, $x_1 = [r_1, y_1]$, and $x_2 = [r_2, y_2]$ can lie on the same menu $\{x_1, x_2\} \subseteq \mathbf{m}$ only if:*

$$a \leq \frac{r_2 - r_1}{y_2 - y_1} \leq b \text{ or } x_1 = x_2, \quad (6)$$

where $0 \leq a \leq b \leq \infty$, and $x_1 = [r_1, y_1], x_2 = [r_2, y_2]$, with r_1, r_2 representing rates and y_1, y_2 representing points.

We illustrate in Figure 8 the effect of defining a menu set based on Assumption 3. As Figure 8 indicates, the range of possible choices that could have come from the same menu as that of the consumer with the choice illustrated by the black dot is more restricted under this assumption, compared with using only dominance relationships in terms of preferences as in Assumption 2. Thus, this improves our ability to detect discrimination in menus. Since Assumption 3 is weakly stricter than Assumption 2 for our test statistic for equality in menus, the only place where Assumption 2 is used in our empirical application is to make welfare comparisons of menu distributions.

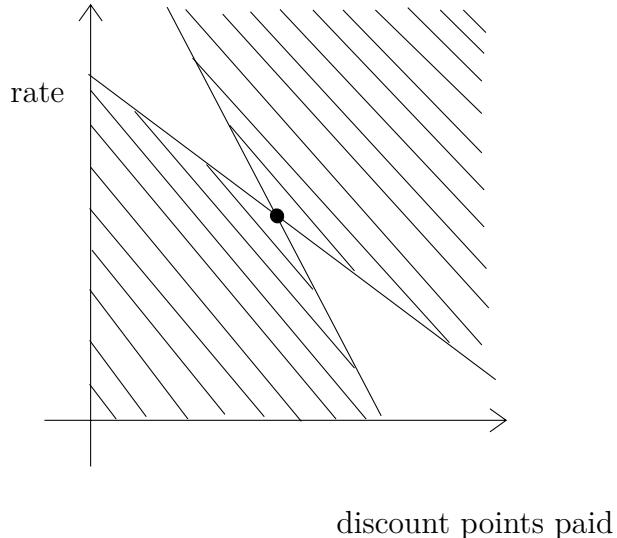


Figure 8: Restriction on what cannot lie on the same menus, from industry knowledge.

Note: this figure illustrates the impact of a restriction on menus from industry knowledge. The dashed region represents what could not have lied on the menu of the borrower making the choice illustrated by the black dot.

Empirically, we find that the industry rule of thumb from Bartlett et al. (November 2019), which motivated Assumption 3, covers a vast majority of menus based on data from

a sample of lender rate sheets that enumerates the rate and discount point menus. Using the 2014 LoanSifter data from Fuster, Lo, and Willen (2019), we estimate slopes of menus within their sample of 30-year purchase mortgages across seven different MSAs (Chicago, Houston, Los Angeles, Miami, New York City, Seattle, and San Francisco) and a range of loan amounts, FICO scores, and LTVs. The sample construction is discussed in more detail in Fuster, Lo, and Willen (2019). We estimate the slopes of the rate-point trade-off by taking the difference in the interpolated rate from 0 points to 2 points and dividing by 2. In this sample, the rule of thumb that each point paid is worth 1/8 to 1/4 of a point covers 94.4 percent of all rate sheet observations for conforming mortgages. For FHA mortgages, we use an expanded rule that each point is worth 1/32 to 1/4 in rate, which covers 97.6 percent of all observations. We illustrate this in Figure 9. Thus, we believe that our Assumption 3 is reasonable.

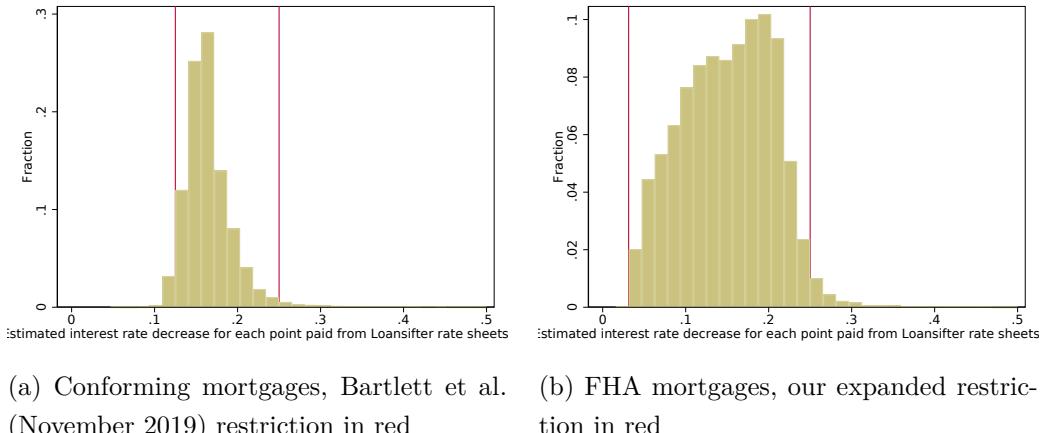


Figure 9: Rate sheet evidence for our menu slopes Assumption 3.

Note: these figures are constructed using 2014 LoanSifter rate sheet data from Fuster, Lo, and Willen (2019) for conforming and FHA mortgages. The slope of the rate-point trade-off is estimated by taking the interpolated rate from 0 to 2 points and dividing by 2. The red lines in Figure 9a represent 1/8 and 1/4, and the red lines in Figure 9b represent 1/32 and 1/4.

In addition, we see substantial heterogeneity between the menu slopes across lender-weeks in Figure 9, perhaps reflecting market power or lender- and time-specific costs. As we explained earlier, the existence of this heterogeneity interacted with possible differences in preferences between the two groups makes an approach like ours necessary for computing robust metrics of discrimination in menus.

Under Assumption 3, let $x_1 = [r_1, y_1], x_2 = [r_2, y_2]$. We define an indicator function for

whether choices x_1, x_2 could have come from the same menu:

$$\phi(x_1, x_2) = \begin{cases} 1, & \text{if } a \leq \frac{r_2-r_1}{y_2-y_1} \leq b \text{ or } x_1 = x_2 \\ 0, & \text{otherwise.} \end{cases} \quad (7)$$

After defining this function, we have the following identification result for when vectors of choice probabilities $\mathbf{p}_1 = [p_1(x_1), p_1(x_2), \dots], \mathbf{p}_2 = [p_2(x_1), p_2(x_2), \dots]$ from observationally similar groups of borrowers can be rationalized under the null hypothesis of equality in the distribution of menus $\mathcal{H}_0 : M_1 = M_2$:

Theorem 1. *Under Assumptions 1 and 3, choice probabilities $\mathbf{p}_1, \mathbf{p}_2$ can be generated from the same underlying distribution of menus $\mathbf{M}_1 = \mathbf{M}_2$ if and only if there exists a coupling with probability mass function $\boldsymbol{\pi}(x_1, x_2) : \mathbb{X} \times \mathbb{X} \rightarrow [0, 1]$ with implied marginal densities $\sum_{x_2} \boldsymbol{\pi} = \mathbf{p}_1, \sum_{x_1} \boldsymbol{\pi} = \mathbf{p}_2$ such that:*

$$T \equiv 1 - E_{\boldsymbol{\pi}} \phi(x_1, x_2) = 0. \quad (8)$$

Proof. See Appendix A.1.1. \square

Theorem 1 formalizes the intuition from Section 2 that equality in menus should imply a way to “match” observations to one another such that each pair can come from the same set of menus. In particular, $\boldsymbol{\pi}$ serves as a coupling or “matching function,” where each of its entries $\boldsymbol{\pi}(x_1, x_2)$ represents the extent to which choice probabilities $p_1(x_1)$ and $p_2(x_2)$ came from the same menu, and the requirement of Equation (8) is that this coupling should lie entirely in the area where $\phi(x_1, x_2) = 1$. Based on Theorem 1, we are ready to define a test statistic that looks at the extent to which this matching is deficient:

Definition 1. *Our test statistic for equality in menus, \hat{T} , given sample choice probabilities $\hat{\mathbf{p}}_1, \hat{\mathbf{p}}_2$, is:*

$$\hat{T} = \min_{\boldsymbol{\pi}(x_1, x_2)} 1 - E_{\boldsymbol{\pi}} \phi(x_1, x_2), \text{ s.t. } \sum_{x_2} \boldsymbol{\pi} = \hat{\mathbf{p}}_1, \sum_{x_1} \boldsymbol{\pi} = \hat{\mathbf{p}}_2, \boldsymbol{\pi} \geq 0, \quad (9)$$

where the extent to which $\hat{T} > 0$ indicates a failure of the perfect matching condition in Theorem 1 in sample, which is evidence against equality in menus. It measures the fraction of each sample that cannot be matched to one another. The statistic \hat{T} in Equation (9) is a finite dimensional optimal transport objective that can be efficiently computed using a linear program. To deal with sampling error inherent in \hat{T} , we discuss in Section 4 how hypothesis

testing for \hat{T} from Equation (9), as an optimal transport objective, can be consistently simulated.

3.4 Metric for assessing differences in menus

The direct test of inequality in the distribution of menus in Section 3.3, while indicative of discrimination in menus, has the drawback that it may not be the object of interest for researchers. The fact that the distribution of menus presented to one group is different in some aspect from the distribution of menus presented to another group may not be of welfare consequence, a problem that Abadie (2020) discusses in more detail. Rather, for the purposes of comparing menus across two distributions, we want a metric for assessing whether the distribution of menus from one group is meaningfully “better” than that of another group. For this purpose, we ask the question: *If Black consumers were instead assigned white menus, how much better off would they be?*

Conceptually, we consider the object of interest to be the change in welfare when Black consumers were instead assigned white menus, under an assignment rule $\pi(i, j)$ that maps each consumer $i \in \mathcal{I}_1$ from group 1 to the menu of consumer $j \in \mathcal{I}_2$ from group 2. Giving all consumers the same welfare weight, this objective can be represented by Equation (10):

$$\Delta W_{\mathcal{I}_1 \rightarrow \mathcal{I}_2, \pi} = \sum_{i \in \mathcal{I}_1, j \in \mathcal{I}_2} \pi(i, j)(u_i(\mathbf{m}_j) - u_i(\mathbf{m}_i)). \quad (10)$$

To get at $\Delta W_{\mathcal{I}_1 \rightarrow \mathcal{I}_2, \pi}$, we proxy for the utility difference $u_i(\mathbf{m}_j) - u_i(\mathbf{m}_i)$ through a metric $d_{i \rightarrow j}(\mathbf{m}_i, \mathbf{m}_j)$ that measures the extent to which consumer i would be willing to increase the interest rates on their loan in order to switch from menu \mathbf{m}_i to menu \mathbf{m}_j :

$$u_i(\mathbf{m}_j) - u_i(\mathbf{m}_i) = d_{i \rightarrow j}(\mathbf{m}_i, \mathbf{m}_j) \equiv \sup\{\delta \in \mathbb{R} : u_i(\{x + \delta e^r, x \in \mathbf{m}_j\}) \geq u_i(\mathbf{m}_i)\}, \quad (11)$$

where e^r represents a basis vector that is equal to 1 at the location indexing interest rates. If consumers have constant marginal utility over interest rates such that utility can be represented as $u_i(x = [r, y]) = r + f(y)$, then $d_{i \rightarrow j}(\mathbf{m}_i, \mathbf{m}_j)$ is directly proportional to the utility change for consumer i after switching to menu j . Even if consumers do not have constant marginal utility over interest rates, it is still meaningful as a “willingness to pay” metric, since $d_{i \rightarrow j}(\mathbf{m}_i, \mathbf{m}_j)$ measures the extent to which consumer i would be willing to increase the interest rates to switch from \mathbf{m}_i to \mathbf{m}_j . In the rest of this section we will show how we can compute an informative lower bound for this metric given the data, $\underline{d}_{i \rightarrow j}(x_i, x_j) \leq d_{i \rightarrow j}(\mathbf{m}_i, \mathbf{m}_j), x_i \in \mathbf{m}_i, x_j \in \mathbf{m}_j$, which then leads our differences in menus measure.¹⁴

¹⁴We note that while we define our $\underline{DIM}_{1 \rightarrow 2}$ over “willingness to pay” in terms of interest rates, we could

To define our lower bound, we make an additional assumption that menus are complete in points, such that all choices of points are available to borrowers, which we formalize as Assumption 4. This is an approximation, since lenders may limit the choices of points to certain decimals (for example, 0.134, 0.266, ...) rather than literally the full range, but the practical implications of such small gaps in menus are likely small. Furthermore, there may be information constraints on the part of borrowers such that they do not “see” their full choice set (that is, some borrowers may not know that they can pay/receive points). Our metric is robust to this information constraints problem since as long as the subsets of choices borrowers “see” are held constant, our lower bound metric would remain valid.

Assumption 4. (*Completeness*) *The menus are complete in discount points. More specifically, $\forall \mathbf{m}, \forall y', \exists x = [r, y'] \in \mathbf{m}$.*

The effect of Assumption 4 is illustrated in Figure 10. Under the assumption that the mortgage menus are complete in discount points, we can meaningfully say that the minority borrower whose choice is represented by the black dot would have preferred the menu of the white borrower whose choice is represented by the white dot, because there exists a level of discount points such that all possible choices in the white borrower’s menu dominate the minority borrower’s choice. Otherwise, the minority borrower might not have preferred the white borrower’s menu because the white borrower’s menu could have been a singleton that the minority borrower dislikes. Therefore, adding the assumption of menu completeness in points sharpens the comparison of menus.

have also defined it using points. However, due to the possible existence of cash on hand constraints that are likely binding for many consumers, constant marginal utility is very unlikely to hold for points, which makes welfare aggregation less credible.

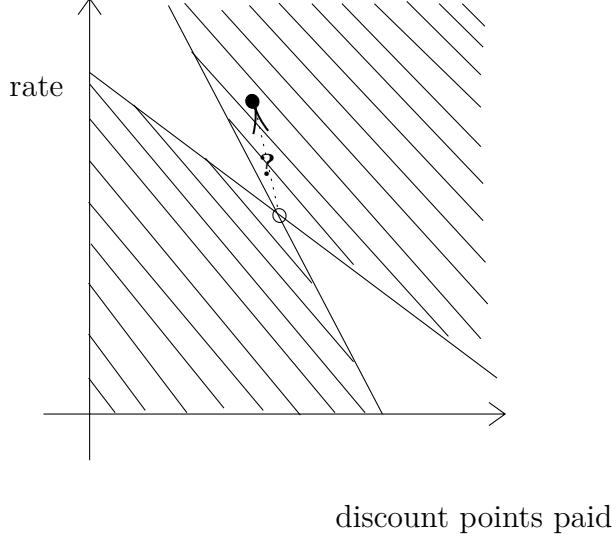


Figure 10: Impact of assuming that menus are complete in either rates or points

Note: this figure illustrates the impact of an assumption on the completeness of menus. The dashed region represents what could not have lied on the menu of the borrower represented by the white dot. While the black dot falls within the dashed region, without assuming completeness it would be unknown if the borrower making the choice represented by the black dot would have preferred the menu containing the white dot. Under completeness, we know that the borrower making the choice represented by the black dot would have preferred the menu containing the white dot by revealed preference.

We illustrate in Figure 11 how we can construct a lower bound for the willingness to pay in terms of interest rates $\underline{d}_{i \rightarrow j}(x_i, x_j)$ under Assumption 4. There, borrower 1, who made a choice x_1 from an unobserved menu \mathbf{m}_1 , has made a choice that is dominated (in terms of paying a higher rate at the same level of points) by any possible menu of borrower 2, who made a choice x_2 from a menu \mathbf{m}_2 . This implies that, by revealed preference of borrower 1, the menu that borrower 2 faced is better than borrower 1's menu, or $\mathbf{m}_2 \succ_1 \mathbf{m}_1$. For borrower 1 to possibly become indifferent between \mathbf{m}_1 and \mathbf{m}_2 , \mathbf{m}_2 needs to be shifted up by at least the amount indicated in the figure in the dimension of interest rates. In other words, a lower (sharp) lower bound for $d_{1 \rightarrow 2}(\mathbf{m}_1, \mathbf{m}_2)$ is $\underline{d}_{1 \rightarrow 2}(x_1, x_2)$, in the sense that borrower 1 is willing to pay at least $\underline{d}_{1 \rightarrow 2}(x_1, x_2)$ more in interest rate in order to get borrower 2's menu. Similarly, the menu faced by borrower 2 would need to be shifted downward by at most the negative $\underline{d}_{3 \rightarrow 2}$ before it dominates x_3 's choice. Therefore, borrower 3, with choice x_3 , would need to pay at least $-\underline{d}_{3 \rightarrow 2}$, or receive at most $\underline{d}_{3 \rightarrow 2}$, before being willing to switch to borrower 2's menu.

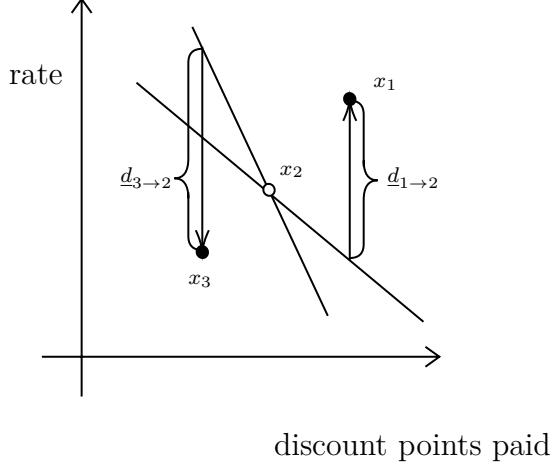


Figure 11: Lower bound for borrower 1's willingness to pay a higher interest rate in order to get borrower 2's menu.

Note: this figure illustrates the computation of our lower bound measure on differences in menus. The borrower making the choice x_1 would pay at least $d_{1 \rightarrow 2}$ in additional interest rate in order to switch to x_2 's menu. The borrower making the choice x_3 would require at most a reduction of $-d_{3 \rightarrow 2}$ in interest rate in order to switch to x_2 's menu.

Formalizing the intuition from Figure 11, we define our lower bound for the willingness of borrower 1 to switch to borrower 2's menu under Assumptions 1 through 4 as follows:

$$d_{1 \rightarrow 2}(x_1, x_2) \equiv r_1 - r_2 + a \max(y_1 - y_2, 0) + b \min(y_1 - y_2, 0) \leq d_{1 \rightarrow 2}(x_1, x_2), \quad (12)$$

where in the first line j indexes points, and in the second line we take it to the mortgage setting and let $x_1 = [r_1, y_1], x_2 = [r_2, y_2]$, where r_1, r_2 are rates and y_1, y_2 are points. The lower bound for how much borrower 1 would be willing to pay to switch to the menus of borrower 2, $d_{1 \rightarrow 2}(x_1, x_2)$, then allows us to define our differences in menus (DIM) metric:

Theorem 2. *Under Assumptions 1 through 4, choice probabilities $\mathbf{p}_1, \mathbf{p}_2$ implies a DIM measure:*

$$\underline{DIM}_{1 \rightarrow 2} = \min_{\boldsymbol{\pi}(x_1, x_2)} E_{\boldsymbol{\pi}} d_{1 \rightarrow 2}(x_1, x_2), \text{ s.t. } \sum_{x_2} \boldsymbol{\pi} = \mathbf{p}_1, \sum_{x_1} \boldsymbol{\pi} = \mathbf{p}_2, \boldsymbol{\pi} \geq 0 \quad (13)$$

where $\underline{DIM}_{1 \rightarrow 2}$ serves as a lower bound for the average willingness to pay in terms of interest rates for borrowers in group 1 to switch menus with borrowers in group 2. If borrowers have constant marginal utility in interest rate, then:

$$\underline{DIM}_{1 \rightarrow 2} \leq \Delta W_{\mathcal{I}_1 \rightarrow \mathcal{I}_2, \boldsymbol{\pi}} \quad (14)$$

where utility is measured in terms of interest rate paid.

Proof. See Appendix A.1.2. □

Theorem 2 shows that, when all consumers have the same constant marginal utility over interest rates (normalized to 1), our DIM metric is as a lower bound for the change in welfare for when consumers in group 1 are instead assigned menus from group 2 in an arbitrary way $\Delta W_{I_1, 1 \rightarrow 2, \pi_{1 \rightarrow 2}}$. If instead consumers do not have constant marginal utility over interest rates, then the $\underline{DIM}_{1 \rightarrow 2}$ metric could still be interpreted as the average increase in interest rates consumers in group 1 would be willing to pay in order to switch to menus from group 2. Furthermore, by Theorem 1, equality in menus would imply that $\underline{DIM}_{1 \rightarrow 2} \leq 0$, so a finding that $\underline{DIM}_{1 \rightarrow 2} > 0$ is also rejection of equality in menus in a welfare relevant way.

The sample analogue of the DIM metric follows immediately from Definition 2.

Definition 2. *Our empirical differences in menus metric, $\hat{DIM}_{1 \rightarrow 2}$, given choice probabilities $\hat{\mathbf{p}}_1, \hat{\mathbf{p}}_2$, is:*

$$\hat{DIM}_{1 \rightarrow 2} = \min_{\boldsymbol{\pi}(x_1, x_2)} E_{\boldsymbol{\pi}} \underline{d}_{1 \rightarrow 2}(x_1, x_2), \text{ s.t. } \sum_{x_2} \boldsymbol{\pi} = \hat{\mathbf{p}}_1, \sum_{x_1} \boldsymbol{\pi} = \hat{\mathbf{p}}_2, \boldsymbol{\pi} \geq 0 \quad (15)$$

In terms of inference, the sample DIM metric in Definition 2 is also the value of a finite dimensional optimal transport problem. We discuss hypothesis testing in this class of problems in Section 4.

3.4.1 Extension: a “simple DIM”

An simple extension of our DIM metric is:¹⁵

$$\underline{DIM}_{1 \rightarrow 2}^s = E_{x_1 \sim f_1(x_1), x_2 \sim f_2(x_2)} \underline{d}_{1 \rightarrow 2}(x_1, x_2) \quad (16)$$

Where $f_1(x_1)$ is the distribution of borrower choices in group 1, and $f_2(x_2)$ is the distribution of borrower choices in group 2.

The $\underline{DIM}_{1 \rightarrow 2}^s$ metric as defined in Equation (16) has a straightforward interpretation: it is a lower bound for the expected willingness to pay of borrowers in group 1 after switching to *random* group 2 borrowers’ menus. This is compared to the worst-possible assignment as in our Definition 2. The drawback of this metric is it requires an additional assumption to be a valid test of differences in menus, such as random assignment of menus conditional on the covariates groups within which it is computed. Such an assumption may be violated. For

¹⁵This section owes its existence to Daniel Ringo. The authors have omitted this possibility from the initial version of the paper in favor of the more robust alternative as presented in Definition 2. Dr. Ringo has independently arrived at this metric and used it to replicate our main results during his discussion of our paper at the 2021 Federal Reserve Day-Ahead Conference on Financial Markets and Institutions. In doing so, he illustrated the computational advantages of this simpler metric which may be useful in some contexts.

example, if, for borrowers in both groups, the heterogeneity in menus is due to unobservables that may be correlated with preferences, it is possible that the $\underline{DIM}_{1 \rightarrow 2}^s$ metric would produce a false positive. Our DIM metric in Definition 2, on the other hand, is robust to such a violation.

Nevertheless, the “simple DIM” metric has the advantage of being easy to compute and conduct inference on. Its sample counter-part is:

$$\hat{DIM}_{1 \rightarrow 2}^s = E_{x_1 \sim \hat{f}_1(x_1), x_2 \sim \hat{f}_2(x_2)} \underline{d}_{1 \rightarrow 2}(x_1, x_2) \quad (17)$$

which may be computed by simply sampling random pairs of choices x_1, x_2 , conditional on covariates, and then averaging their $\underline{d}_{1 \rightarrow 2}(x_1, x_2)$ metric. Standard techniques such as bootstrapping may then be used for inference.

In our empirical application, the “simple DIM” gives similar qualitative results to our DIM metric. Thus, researchers may find such a metric useful for data exploration before moving on to our more robust metric. Furthermore, if the researcher believes that random assignment of menus conditional on the covariates groups is a valid assumption to make, the “simple DIM” is also a valid metric for assessing difference in menus.

3.5 When does our method have power?

Our metrics for differences in menus are robust in the sense that they are immune to the false positives problem from which the existing methods suffer, but they may still generate false negatives in that there exist scenarios where the data are rationalizable under equality in the distribution of menus, but in fact the distribution is different. Generally speaking, our method has power only to the extent that the borrowers’ choices cannot be rationalized by an equal distribution of menus under the restrictions on preferences and menus that we have made. In this subsection, we give some examples of scenarios where this would or would not occur. In our later empirical analysis, we show that our methodology does have enough power to be useful in detecting discrimination in mortgage markets.

Figure 12 illustrates a scenario in which white and minority borrowers face the same default menu, but some white borrowers are offered a discretionary discount in terms of points. In the figure, the menu represented by the dashed line is shifted leftward for a white borrower, but all minority borrowers face the original menu. This shift makes the bottom-left choice by the white borrower not matchable to any of the minority borrowers’ choices, so our one-to-one matching condition for equality in menus in Theorem 1 is broken. Our Theorem 1 would therefore have power to detect a difference in the distribution of menus offered to white and minority borrowers in this case. Furthermore, when the discretionary

discount being given to white borrowers is large enough relative to the range of permissible menus, our DIM metric defined in Theorem 2 would also show that minority borrowers would be willing to pay to switch to the white borrowers' menus.

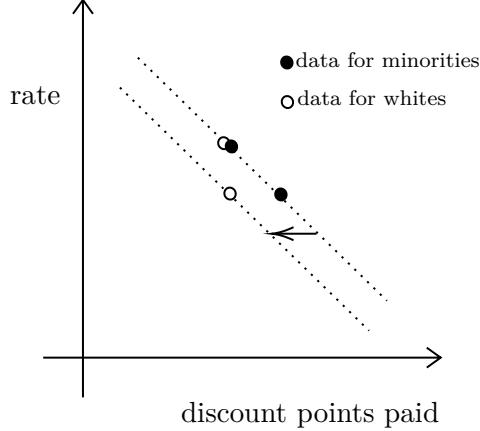


Figure 12: Power to detect discrimination when discretionary discounts offered to only some white borrowers make the data not rationalizable under equality in menus

Note: this figure illustrates how, if the loan originator offered only a white borrower a discount on points paid as represented by a leftward shift of the dashed line, and the borrower maintained their choice of rate, our methodology would have power to detect this pricing differential.

Figure 13 gives an example situation in which minority and white borrowers were presented with different menus as illustrated in the dotted lines, with white borrowers on average paying less in both rates and points, but the data are rationalizable under equality in menus. More specifically, while the data were generated by menus represented by the dotted lines which differ in distribution between white and minority borrowers, the data can be rationalized with one minority and one white borrower both choosing from the hypothetical menu represented by the dashed line and the other two borrowers choosing from the remaining dotted line. In this scenario, we cannot rule out that the true distribution of menus is represented by the dashed line plus a dotted line rather than the two dotted lines, and therefore our metrics from both Definition 1 and Definition 2 would fail to reject equality in menus.

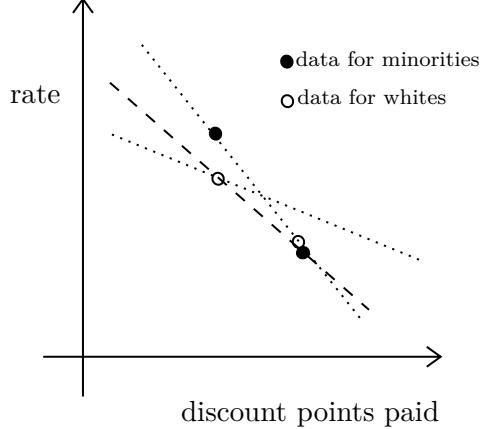


Figure 13: No power when the data can be rationalized by (incorrect but plausible) distribution of menus that are equal across the racial groups

Note: this figure illustrates how, if the data is as illustrated in black and white dots, it would be unknown to the researcher whether the true underlying distributions of menus are represented by the two dotted lines (in which case there would be an inequality in the distribution of menus) or the dashed line and a dotted line (in which case there would be equality in the distribution of menus).

While our approach cannot detect differences in menus in the situation of Figure 13, there is a good reason for this: It is possible that the borrowers' choices were truly generated by the dashed line and the dotted line such that the racial groups did in fact face the same distribution of menus. It is indeed impossible to rule out that possibility given the data presented. This is the sense in which our metrics are sharp. Empirically in Section 5, we show that we are able to reject equality in menus for conforming mortgages, which shows that our method does have enough power to be useful in our mortgage setting.

4 Inference for Optimal Transport

In this section, we devise a new procedure for conducting hypothesis testing on the values of optimal transport problems that includes our metrics derived in Section 3. This is needed because the objective values of our model can be non-differentiable with respect to $\mathbf{p}_1, \mathbf{p}_2$ which implies that simple bootstrap methods may not be consistent (Fang and Santos, 2018). While some existing methods in the literature can be applied in more general contexts, they are either too conservative or converge too slowly based on our simulations, which we discuss at the end of this section. Our new procedure may be of independent interest for other researchers who wish to apply optimal transport methods to economics; some of those applications are listed in Galichon (2016).

We consider hypothesis testing on the value φ of a finite dimensional optimal transport

problem with cost function $\phi(x_1, x_2)$, $x_1, x_2 \in \mathbb{X}$ and marginal distributions $\mathbf{p}_1, \mathbf{p}_2$:

$$\hat{\varphi}(\hat{\mathbf{p}}_1, \hat{\mathbf{p}}_2) = \min_{\boldsymbol{\pi} \in \mathcal{P}_{(x_1, x_2)}} E_{\boldsymbol{\pi}} \phi \text{ s.t. } \sum_{x_2} \boldsymbol{\pi} = \hat{\mathbf{p}}_1, \sum_{x_1} \boldsymbol{\pi} = \hat{\mathbf{p}}_2, \boldsymbol{\pi} \geq 0, \quad (18)$$

where the hypothesis is in the form of the value of the optimal transport $\varphi(\mathbf{p}_1, \mathbf{p}_2)$ as a function of the true marginal distributions $\mathbf{p}_1, \mathbf{p}_2$ being less than or equal to some value φ_0 :

$$\mathcal{H}_0 : \varphi(\mathbf{p}_1, \mathbf{p}_2) \leq \varphi_0, \quad (19)$$

$$\mathcal{H}_a : \varphi(\mathbf{p}_1, \mathbf{p}_2) > \varphi_0. \quad (20)$$

The form of the null hypothesis in Equation (19) is especially relevant to us because both our test of equality in menus (that is, whether $T \leq 0$) and our lower bound DIM metric (that is, whether $DIM \leq DIM_0$) can be expressed in terms of it. We provide a methodology to conduct this hypothesis test by looking at the asymptotic distribution of $\hat{\varphi}$ and finding a critical value to compare the observed $\hat{\varphi}$ to under \mathcal{H}_0 . This can then be inverted into a confidence interval for the true value of φ .

As an overview, we combine the directional derivatives approach of Fang and Santos (2018) and Shapiro (1990) with a size correction of Romano, Shaikh, and Wolf (2014) and McCloskey (2017), which allows us to conduct hypothesis testing for optimal transport with uniform size control. To do so, we prove the directional differentiability for optimal transport problems on finite domains, and show how the general approach can be implemented as a linear program with complementarity constraints (LPCC).

We use the definition of Hadamard directional differentiation from Fang and Santos (2018), with some notational differences tailored to the optimal transport setting. Here, the value of an optimal transport represents a map $\varphi : \mathbb{D}_\varphi \rightarrow \mathbb{R}$, where $\mathbb{D}_\varphi = \mathcal{P}_{\mathbb{X}} \times \mathcal{P}_{\mathbb{X}}$, $\mathcal{P}_{\mathbb{X}}$ is the set of probability measures on \mathbb{X} . Let $\mathbb{D}_0 = \{P_1 - P_2 : P_1 \in \mathcal{P}_{\mathbb{X}}, P_2 \in \mathcal{P}_{\mathbb{X}}\}$ be the set of possible differences in probability measures, and $\theta = \{p_1, p_2\}$ be the marginal distributions, then:

Definition 3. (Fang and Santos, 2018) A map $\varphi : \mathbb{D}_\varphi \rightarrow \mathbb{R}$ is said to be Hadamard directionally differentiable at $\theta \in \mathbb{D}_\varphi$ tangentially to the set \mathbb{D}_0 , if there is a continuous linear map $\varphi'_\theta : \mathbb{D}_0 \rightarrow \mathbb{R}$ such that:

$$\lim_{n \rightarrow \infty} \left\| \frac{\varphi(\theta + t_n \mathbf{h}_n) - \varphi(\theta)}{t_n} - \varphi'_\theta(\mathbf{h}) \right\| = 0 \quad (21)$$

for all sequences $\{\mathbf{h}_n\} \subset \mathbb{D}_0$ and $\{t_n\} \subset \mathbb{R}^+$, such that $t_n \rightarrow^+ 0$, $\mathbf{h}_n \rightarrow \mathbf{h} \in \mathbb{D}_0$ as $n \rightarrow \infty$ and $\theta + t_n \mathbf{h}_n \in \mathbb{D}_\varphi$ for all n .

The main difference between the Hadamard directional differentiability and the typical notion of differentiability is that t_n approaches zero from the positive direction in Definition 3. Loosely speaking, the directional derivatives represent the change in the value of the function for a small change in its inputs “in the direction \mathbf{h} ” for each \mathbf{h} .

We show in Theorem 3 that the value of all Monge-Kantorovich optimal transport problems with bounded cost functions on finite spaces is Hadamard directionally differentiable in the sense of Definition 3. In particular, Theorem 3 is a generalization of Sommerfeld and Munk (2018), which shows that the Wasserstein metric (the value of an optimal transport problem with the cost function restricted to distance metrics) on finite spaces is directionally differentiable.¹⁶

Theorem 3. *The value φ of an optimal transport problem with cost function $\phi(x_1, x_2)$, $x_1, x_2 \in \mathbb{X}$, where $M = \sup |\phi| < \infty$ and $\dim(\mathbb{X}) < \infty$, is Hadamard directionally differentiable, with derivative equal to:*

$$\varphi'_{\mathbf{p}_1, \mathbf{p}_2}(\mathbf{h}_1, \mathbf{h}_2) = \max_{\mathbf{u}, \mathbf{v} \in \Psi^*(\mathbf{p}_1, \mathbf{p}_2)} \mathbf{h}_1^T \mathbf{u} + \mathbf{h}_2^T \mathbf{v} \quad (22)$$

where $\Psi^*(\mathbf{p}_1, \mathbf{p}_2) = \{\mathbf{u}, \mathbf{v} : \mathbf{p}_1^T \mathbf{u} + \mathbf{p}_2^T \mathbf{v} = \varphi(\mathbf{p}_1, \mathbf{p}_2), u(x_1) + v(x_2) \leq \phi(x_1, x_2) \forall x_1, x_2\}$ is the set of dual solutions to the linear programming problem, for all $\{\mathbf{p}_1, \mathbf{p}_2\} \in \mathbb{D}_\varphi$, tangentially to the set \mathbb{D}_0 .

Proof. See Appendix A.1.3. □

Under i.i.d. sampling, we know that $\hat{\mathbf{p}}_1 - \mathbf{p}_1$ and $\hat{\mathbf{p}}_2 - \mathbf{p}_2$ approach a multivariate Normal distribution:

$$\hat{\mathbf{p}}_1 - \mathbf{p}_1 \xrightarrow{d} N \left(0, \begin{bmatrix} p_{1,1}(1-p_{1,1}) & -p_{1,1}p_{1,2} & \dots \\ -p_{1,2}p_{1,1} & p_{1,2}(1-p_{1,2}) & \dots \\ \dots & \dots & \dots \end{bmatrix} \right), \quad (23)$$

and likewise for $\hat{\mathbf{p}}_2 - \mathbf{p}_2$, such that by construction, Assumptions 2.1 and 2.2 of Fang and Santos (2018) are satisfied. Then, Theorem 2.1 of Fang and Santos (2018) immediately implies that:

$$r_n[\varphi(\hat{\mathbf{p}}_1, \hat{\mathbf{p}}_2) - \varphi(\mathbf{p}_1, \mathbf{p}_2)] = \varphi'_{\mathbf{p}_1, \mathbf{p}_2}(r_n[(\hat{\mathbf{p}}_1, \hat{\mathbf{p}}_2) - (\mathbf{p}_1, \mathbf{p}_2)]) + o_p(1), \quad (24)$$

¹⁶It is also related to Tameling, Sommerfeld, and Munk (2019), who prove that the Wasserstein distance on countable metric spaces is directionally differentiable. Our Theorem 3 can be similarly extended to countable metric spaces under the assumption that the cost function ϕ is continuous.

such that the asymptotic distribution of $\varphi(\{\hat{\mathbf{p}}_1, \hat{\mathbf{p}}_2\})$ can be obtained via the directional Delta method. The remaining challenge for inference is that the true $\mathbf{p}_1, \mathbf{p}_2$ used in $\varphi'_{\{\mathbf{p}_1, \mathbf{p}_2\}}$ in Equation 24 are not known, and thus must be estimated. While we could have used the plug-in analogue $\varphi'_{\{\hat{\mathbf{p}}_1, \hat{\mathbf{p}}_2\}}$, that would converge only pointwise and not uniformly, which is known in the moment inequalities literature to be a poor approximation to the finite sample properties of estimators for which there are discontinuities in the pointwise asymptotic distribution (Andrews and Soares, 2010). We approach this problem following the logic of Romano, Shaikh, and Wolf (2014) and McCloskey (2017). More specifically, we suppose there exists confidence bands for $[\mathbf{p}_1, \mathbf{p}_2]$ at level β such that:

$$\limsup_{n \rightarrow \infty} \Pr([\mathbf{p}_1, \mathbf{p}_2] \in \hat{\mathcal{P}}_{n, \beta}) \geq 1 - \beta. \quad (25)$$

Many uniform confidence bands satisfying Equation (25) are available, for example from Montiel Olea and Plagborg-Møller (2019). Then, we take our estimate of the directional derivative as the maximum directional derivative within this confidence band:

$$\hat{\varphi}'_\beta(\mathbf{h}_1, \mathbf{h}_2) = \max_{\mathbf{u}, \mathbf{v} \in \Psi(\mathbf{p}_1, \mathbf{p}_2) : [\mathbf{p}_1, \mathbf{p}_2] \in \hat{\mathcal{P}}_\beta} \mathbf{h}_1^T \mathbf{u} + \mathbf{h}_2^T \mathbf{v}, \quad (26)$$

where $\Psi = \{\mathbf{u}, \mathbf{v} : \mathbf{p}_1^T \mathbf{u} + \mathbf{p}_2^T \mathbf{v} \leq \varphi_0, u(x_1) + v(x_2) \leq \phi(x_1, x_2) \forall x_1, x_2\}$ are the set of dual solutions under the null hypothesis $\mathcal{H}_0 : \varphi \leq \varphi_0$.

Next, we define the critical value for φ by using our estimated maximum directional derivative from Equation (26) at level $1 - \alpha + \beta$:

$$\hat{c}_{1-\alpha+\beta} = \inf\{c \in \mathbb{R} : \Pr(\hat{\varphi}'_\beta(\mathbf{h}_1, \mathbf{h}_2) \leq c) \geq 1 - \alpha + \beta\}, \quad (27)$$

where the distribution of $\mathbf{h}_1, \mathbf{h}_2$ is the asymptotic distribution of $\hat{\mathbf{p}}_1 - \mathbf{p}_1, \hat{\mathbf{p}}_2 - \mathbf{p}_2$. In the following Corollary 1, we will prove uniform coverage for when the observed value $\hat{\varphi}$ is less than the critical value $\hat{c}_{1-\alpha+\beta}$, and suggest a computationally tractable version of it as a linear program with complementarity constraints (LPCC).¹⁷

Corollary 1. *Suppose we have uniform confidence bands for $[\mathbf{p}_1, \mathbf{p}_2] \in \hat{\mathcal{P}}_\beta$ that provide uniform coverage as in Equation (25), then under $\mathcal{H}_0 : \varphi \leq \varphi_0$:*

$$\limsup_{n \rightarrow \infty} \Pr(\hat{\varphi} - \varphi_0 \geq \hat{c}_{1-\alpha+\beta}) \leq \alpha, \quad (28)$$

¹⁷As explained in Hsieh, Shi, and Shum (2021), LPCCs are well understood computationally and are implemented in software such as Knitro: https://www.artelys.com/docs/knitro/2_userGuide/complementarity.html.

where $\hat{c}_{n,1-\alpha+\beta}$ is computed as in Equation (27).

Proof. See Appendix A.1.4. \square

Corollary 1 implies that uniformly valid hypothesis testing for the value of φ can be conducted by first computing a set of uniform confidence bands $[\mathbf{p}_1, \mathbf{p}_2] \in \hat{\mathcal{P}}_\beta$, and then maximizing over all directional derivatives within these bands as in Equation (26). Directly maximizing over the directional derivative defined in Equation (26) is computationally intensive, because it involves optimizing over a nonlinear dual value constraint $\mathbf{p}_1^T \mathbf{u} + \mathbf{p}_2^T \mathbf{v} \leq \varphi_0$. To deal with this, we replace it with an equivalent complementary slackness formulation $\boldsymbol{\pi}^T \mathbf{s} = 0, \boldsymbol{\pi} \geq 0, \mathbf{s} \geq 0$ where $u(x_1) + v(x_2) + s(x_1, x_2) = \phi(x_1, x_2)$, which implies that the elements of $\boldsymbol{\pi}$ and \mathbf{s} cannot be positive simultaneously. Following the operations research shorthand, we represent this constraint by $\boldsymbol{\pi} \leq 0 \perp \mathbf{s} \geq 0$. The derivation of this equivalence can be found in standard texts on optimal transport/linear programming. In particular, Hsieh, Shi, and Shum (2021) use a similar set of conditions for their projection method. Based on this equivalency, the problem of finding critical values for the null hypothesis $\mathcal{H}_0 : \varphi \leq \varphi_0$ versus the alternative $\mathcal{H}_a : \varphi > \varphi_0$ can be computed through the following LPCC which we implement:

$$\hat{\varphi}'_\beta(\mathbf{h}_1, \mathbf{h}_2) = \max_{\mathbf{u}, \mathbf{v}, \mathbf{p}_1, \mathbf{p}_2, \mathbf{s}, \boldsymbol{\pi}} \mathbf{h}_1^T \mathbf{u} + \mathbf{h}_2^T \mathbf{v}, \quad (29)$$

$$\sum_{x_2} \boldsymbol{\pi} = \mathbf{p}_1, \quad (30)$$

$$\sum_{x_1} \boldsymbol{\pi} = \mathbf{p}_2, \quad (31)$$

$$E_{\boldsymbol{\pi}} \phi \leq \varphi_0, \quad (32)$$

$$u(x_1) + v(x_2) + s(x_1, x_2) = \phi(x_1, x_2), \quad (33)$$

$$[\mathbf{p}_1, \mathbf{p}_2] \in \hat{\mathcal{P}}_\beta, \quad (34)$$

$$\boldsymbol{\pi}, \mathbf{s} \geq 0, \quad (35)$$

$$\boldsymbol{\pi} \leq 0 \perp \mathbf{s} \geq 0. \quad (36)$$

To test our econometric approach, we conduct a Monte Carlo simulation with two possibilities for points $\{0, 1\}$ and five possibilities for rate $\{3, 3.25, 3.5, 3.75, 4\}$. Furthermore, Black and white borrowers choose each of the rate-point options with probability $\frac{1}{10}$ such that the null discrimination of no discrimination in menus is satisfied. We compute $\hat{\mathcal{P}}_\beta$ using the plug-in sup-t band of Montiel Olea and Plagborg-Møller (2019). Let $\beta = \frac{1}{10}\alpha$ following Romano, Shaikh, and Wolf (2014), and show the simulated probability that we reject equality in menus $\mathcal{H}_0 : \varphi_0 = 0$ at the $\alpha = 0.01, 0.025, 0.05, 0.1$ levels in Table 1. As Table 1

shows, our approach has the approximately correct size across a wide range of sample sizes and significance levels.

Table 1: Control of our size-corrected directional derivative approaches to inference

Sample size	Significance level			
	1%	2.5%	5%	10%
$n_1 = n_2 = 500$	0.6	2.3	4.7	9.3
$n_1 = n_2 = 1000$	0.8	2.9	5.3	9.3
$n_1 = n_2 = 5000$	1.1	2.3	4.7	9.1
$n_1 = n_2 = 10000$	0.7	2.0	3.8	7.6
$n_1 = n_2 = 50000$	0.8	1.9	4.1	8.6

Note: Computed via 2000 sample draws and 500 draws of $\mathbf{h}_1, \mathbf{h}_2$ from the estimated asymptotic multivariate Normal distribution for $\mathbf{p}_1, \mathbf{p}_2$ within each sample draw.

Compared with existing methodology that can be applied to the optimal transport context, the advantage of our procedure is that it achieves uniform coverage without being overly conservative. In particular, Hsieh, Shi, and Shum (2021) have a novel projection method for parameter inference in mathematical programming problems, which is a broader set of problems than optimal transport, but their approach is conservative. In our empirical context, this conservativeness tends to make the confidence intervals uninformative. Another approach that is theoretically valid in this setting is the general m -out-of- n subsampling method of Politis and Romano (1994), but in addition to requiring the researcher to choose a subsample size m , it can require very large samples for convergence.

We look at the control of these competing approaches at the 5 percent significance level under our simulation setting in Appendix Table A.1. In that table, HSS (2021) refers to the projection method of Hsieh, Shi, and Shum (2021), the m -out-of- n subsampling approach refers to the method of Politis and Romano (1994), and in the final column, the size-corrected directional derivatives approach is replicated from Table 1 for comparison. Table A.1 shows that the HSS (2021) approach rejects the null with probability close to 0 percent, consistent with its conservativeness. On the other hand, the m -out-of- n subsampling approach tends to reject at rates greater than 5 percent for all values of m we tried, and appears to require more than 50,000 observations in order to converge to the correct rejection rate, which is significantly larger than our available sample size.¹⁸

¹⁸A method related to m -out-of- n subsampling is the numerical bootstrap of Hong and Li (2020), which can be more data efficient than the subsampling method. But, we were not able to find a suitable choice of ϵ_n that converges to the correct coverage in our simulations using that method.

In Section 5, we also report one-sided confidence intervals from the inversion of our hypothesis test. Nothing in our econometric theory precludes us from also testing the other direction and reporting two-sided confidence intervals instead.¹⁹ However, since our economic theory is focused on getting a *lower bound* for the existence and welfare effects of discrimination, one-sided confidence intervals are particularly suitable for our purposes.

In summary, in this section we devised a new procedure for inference in optimal transport that is uniformly valid and not overly conservative for our empirical context. Our empirical analysis in Section 5 shows that we are able to strongly reject equality in menus for conforming mortgages using our methodology.

5 Empirical Estimates of Mortgage Discrimination

5.1 Data

We apply our methodology to a new data set constructed via matching the 2018–2019 public Home Mortgage Disclosure Act (HMDA) data to Optimal Blue rate locks. The public HMDA data contains loan level information along with indicators for borrower race and ethnicity, and we take the borrowers with an HMDA-derived race of “Black or African American” as our sample of Black borrowers, borrowers with a derived ethnicity of “Hispanic or Latino” as our sample of Hispanic borrowers, and borrowers with an HMDA-derived race of “White” along with a HMDA derived ethnicity of “Not Hispanic or Latino” as our sample of non-Hispanic white borrowers.

The HMDA data does not contain information on credit scores, and so we supplement it using data from Optimal Blue. Optimal Blue is a rate locking platform used by lenders that comprise of 40% of the US mortgage market, and has been used to study many questions regarding mortgage pricing.²⁰ The data construction process proceeds in two steps. First, we use year, loan type, loan purpose, occupancy, loan amount, interest rate, LTV (allowing for a 5% difference), and zip code to match each loan in Optimal Blue to originated loans in the 2018-2019 public HMDA data. For FHA mortgages, we further allow the mortgage insurance premium to be added to the loan amount. Second, we use the output of the first step to assign a de-identified set of lender IDs from HMDA to de-identified Optimal Blue lender IDs. Overall, we were able to uniquely match 55.4% of Optimal Blue rate locks to HMDA, almost all of which are unique. Our match rate is comparable to a 66% “lock pull-through rate,”

¹⁹The other direction may require more computational finesse, however, since it would involve taking the minimum of a maximum.

²⁰For example, Bhutta, Fuster, and Hizmo (2019), Bhutta and Hizmo (2020), Bhutta and Ringo (2021), and Fuster, Hizmo, Lambie-Hanson, Vickery, and Willen (2021).

which is the rate at which locks turn into originated loans, that we understand is reasonable based on industry sources.

Starting from the uniquely matched HMDA-Optimal Blue matched data set, we further restrict our analysis to standard 30-year, new-purchase, fixed-rate, first-lien mortgages on owner-occupied, site-built properties without prepayment-penalties, balloon, interest-only, negative-amortization, or non-amortizing features. Table 2 compares the HMDA data, which include the complete set of mortgages originated in the United States with such characteristics, to our matched sample. We find that our matched sample has very similar average loan sizes, LTVs, rates, points, and the percentage composition of Black and Hispanic borrowers compared with the HMDA data, as can be seen from Table 2. One known caveat to this is that lenders using the Optimal Blue platform tend to be smaller lenders, with the larger lenders being more likely to have their own platform for rate locking.

Table 2: Comparison of means between the 2018–2019 HMDA data and our HMDA-Optimal Blue matched data

	Loan Size	LTV	Rate	Points	% Black	% Hispanic	<i>N</i>
<i>Panel A: Conforming mortgages</i>							
HMDA data	\$261,566	84.4	4.50	0.08	4.5	9.0	3,730,152
Matched sample	\$258,205	83.5	4.59	0.12	3.9	8.6	817,588
<i>Panel B: FHA mortgages</i>							
HMDA data	\$215,144	96.1	4.57	0.07	14.1	19.7	1,437,088
Matched sample	\$220,031	95.7	4.63	0.13	13.7	18.8	360,202

Note: This table compares the 2018–2019 HMDA data and our HMDA-Optimal Blue matched sample for 30-year, new-purchase, fixed-rate, first-lien mortgages on owner-occupied, site-built properties without prepayment-penalties, balloon, interest-only, negative-amortization, or non-amortization features. Outliers for points above 4 and below -4 and outliers for rates below 2 and above 10.25 were excluded.

To control for the impact of lender and borrower characteristics within each loan program, we exactly match observations from Black and non-Hispanic white borrowers without replacement on groups of covariates, taking a random observation when multiple white borrowers can be matched to a Black borrower. This creates a sample containing equal numbers of Black and non-Hispanic white borrowers that are exactly matched on their covariates. If lenders offered Black and non-Hispanic white borrowers the same distribution of menus conditional on covariates, our covariate-matched sample of Black and non-Hispanic white borrowers should then have faced the same distribution of menus. The covariates that we used to match are lender, county, month of lock, four categories of loan amount, eight categories of FICO scores and nine categories of LTVs as defined in the GSE Loan-Level Price Adjustment (LLPA) matrix. Thus, we control for the effects of the interactions of all these

Table 3: Summary statistics on the covariate matched sample

Panel A: Black and Non-Hispanic White Covariate Matched Sample						
	Conforming			FHA		
	Black	White	Difference	Black	White	Difference
Raw						
Rate Spread (bps)	36.2	31.4	4.8	48.3	45.4	2.9
Points (bps)	12.6	10.1	2.5	10.6	13.6	-3.0
De-meanned & rounded						
Rate Spread (bps)	3.1	-1.7	4.8	2.2	-0.7	2.9
Points	1.5	-0.9	2.4	-1.6	1.6	-3.2
Sample size	6,398	6,398		4,711	4,711	
Panel B: Hispanic and Non-Hispanic White Covariate Matched Sample						
	Conforming			FHA		
	Hispanic	White	Difference	Hispanic	White	Difference
Raw						
Rate Spread (bps)	38.3	34.5	3.8	49.3	45.4	3.9
Points (bps)	15.6	11.8	3.8	13.5	13.3	-0.2
De-meanned & rounded						
Rate Spread (bps)	2.5	-1.2	3.7	2.6	-1.2	3.8
Points	2.3	-1.3	3.6	-0.1	0.2	-0.3
Sample size	14,758	14,758		6,156	6,156	

Note: This table lists the summary statistics for our Black and white as well as Hispanic and non-Hispanic white lender-county-month and covariate matched sample of 30-year, new-purchase, fixed-rate, first-lien mortgages on owner-occupied, site-built properties without prepayment-penalties, balloon, interest-only, negative-amortization, or non-amortization features. Outliers for points above 4 and below -4 were dropped, and all rate spreads were between -1.25 and 3. De-meanned and rounded rates and points were demeaned by lender-county-month and covariate group and then rounded into the nearest eighths for rate and halves for points.

covariates in our assessment of equality in menus. This set of controls is similar to what is used in Bartlett et al. (2021).

Summary statistics for our matched sample are in Table 3. Mortgage interest rates were measured based on a spread to the Freddie Mac Weekly Survey rate during the week of the rate lock. For our empirical analysis, we de-mean within each lender-county-month covariate group and round rates to the nearest eighths and points to the nearest halves, and we show in Table 3 that this step does not substantively change the mean differences in rates or points.²¹

²¹Both the HMDA data and the Optimal Blue data contain information about discount points paid that sometimes disagrees with one another. While the literature has used both data sources, we focus on the HMDA information, because it is more precisely defined as the discount points to reduce the interest rate in Line A.01 of the Closing Cost Details page of the Closing Disclosure, according to CFPB's Regulation C. Furthermore, we find, as shown in Appendix Table A.2, that regression of the HMDA information on points on origination charges and total loan costs has a much stronger R^2 , with a coefficient closer to 1, compared with the Optimal Blue information on points. In a regression controlling for the HMDA points, the effect of Optimal Blue points has minimal additional explanatory power for origination charges and total loan costs. We interpret this as suggestive evidence for there being more measurement error in the Optimal Blue definition of points. Our results using the Optimal Blue data on points are qualitatively similar.

As shown in Table 3, Black borrowers paid 4.8 basis points more in interest rate and 2.5 basis points more in points for conforming mortgages. On the other hand, they paid only 2.9 basis points more in interest rate and -3.2 basis points fewer points for FHA mortgages. While this comparison of means cannot be interpreted as evidence for or against discrimination in menus (as we noted in Section 2), it does show that the distribution of data underlying conforming mortgages is different from the distribution of data for FHA mortgages. Appendix A.3 shows that the choice of heuristic used in analysis can lead to contradictory results in the FHA sample.

Note that because our data is limited to the set of mortgages that originated, our results should be interpreted as assessing differences in menus faced by the set of borrowers who took out a mortgage. The mortgage pricing discrimination literature including Bartlett et al. (2021) and Bhutta and Hizmo (2020) has argued that the pricing difference conditional on origination is policy relevant conditional on our observables. If borrowers that are offered worse prices are more likely to not get a mortgage at all, the extent of mortgage pricing differences at the application stage may be larger than what we detect at the origination stage. Our general methodology could be used to compare differences in menus at both stages if we had data on the chosen rates and points in the Loan Estimates and Closing Disclosures as well as at origination. HMDA does not yet collect such data on non-originated mortgage applications, and so we leave the assessment of mortgage pricing differences at the application stage for future research.

5.2 Analysis of differences in menus using our metrics

In this section we present our assessment of whether lenders offered minority borrowers worse menus of rates and points. Table 4 presents results from testing for equality between menus for Black borrowers and menus for white borrowers using our Definition 1. We find, as shown in columns (1) and (3), that our test statistic for inequality is positive and highly significant for both Black and Hispanic borrowers for conforming mortgages. More specifically, for conforming mortgages in the Black versus non-Hispanic white matched sample, our test statistic in column (1) is $\hat{T} = 2.69$, which indicates that 2.69 percent of Black borrowers' choices in the data could not have been matched to those of non-Hispanic white borrowers' choices that could have been on the same menu. Using our inference approach described in Section 4, we find that statistic is different from zero with $p < 0.01$. For Hispanic borrowers, our test statistic is $\hat{T} = 1.60$ in column (3), with $p < 0.05$. For FHA mortgages, on the other hand, we are unable to reject equality between the menus for Black and non-Hispanic white borrowers, with a test statistic of $\hat{T} = 0$ in column (2), but we are able to reject it at the 10

percent level for Hispanic and non-Hispanic white borrowers with a test statistic of $\hat{T} = 1.90$ in column (4).

Table 4: Results from our test of equality in menus (\hat{T}).

	Black vs Non-Hispanic White		Hispanic vs Non-Hispanic White	
	(1)	(2)	(3)	(4)
	Conforming	FHA	Conforming	FHA
Test statistic (\hat{T})	2.69***	0.00	1.60**	1.90*
95% CI	[0.67, ∞)	[0.00, ∞)	[0.21, ∞)	[0.00, ∞)
N	12,796	9,422	29,516	12,312

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

Note: This table shows results from our test for equality in menus based on our Definition 1, with \hat{T} in units of percentage points. We use the Black and non-Hispanic white and Hispanic and non-Hispanic white lender-county-month and covariate-matched samples of 30-year, new-purchase, fixed-rate, first-lien mortgages on owner-occupied, site-built properties without prepayment-penalties, balloon, interest-only, negative-amortization, or non-amortization features. Outliers for points above 4 and below -4 were excluded. Confidence intervals are computed through inversion of the hypothesis test as described in Section 4, using 2,000 draws from the asymptotic Normal distribution implied by $\hat{\mathbf{p}}_1, \hat{\mathbf{p}}_2$.

Table 5 presents our results for our differences in menus metric as in Definition 2. Column (1) shows that for conforming mortgages, Black borrowers on average are willing to pay at least $\hat{DIM}_{1 \rightarrow 2} = 1.66$ basis points more in interest rates in order to get the non-Hispanic white borrowers' menus for conforming mortgages. This again rejects equality in menus and indicates that the distribution of menus faced by black borrowers is worse than that faced by white borrowers. Similarly, Column (3) shows that Hispanic borrowers are willing to pay at least $\hat{DIM}_{1 \rightarrow 2} = 1.46$ basis points more in order to get the non-Hispanic white borrowers' menus. Both of these statistics are significant at the 1% level. Our lower bound for how much more in interest rates non-Hispanic white borrowers would be willing to pay to switch to minority menus, on the other hand, are consistently negative. The magnitudes of these point estimates are small, even a small difference in interest rate at origination leads to a large difference in payments over the lifetime of the mortgage, as explained in Bartlett et al. (November 2019).

Table 5: Results for our lower bound for the average interest rate increase (bps) needed for consumers to remain indifferent after switching to another group's menus (\hat{DIM})

	Black vs Non-Hispanic White		Hispanic vs Non-Hispanic White	
	Conforming (1)	FHA (2)	Conforming (3)	FHA (4)
Minority to white ($\hat{DIM}_{1 \rightarrow 2}$)	1.66*** [1.09, ∞)	-1.83 [-2.60, ∞)	1.46*** [1.03, ∞)	-0.91 [-1.89, ∞)
95% CI				
White to minority ($\hat{DIM}_{2 \rightarrow 1}$)	-6.16 [-6.82, ∞)	-7.03 [-7.85, ∞)	-6.04 [-6.40, ∞)	-7.57 [-8.24, ∞)
95% CI				
<i>N</i>	12,796	9,422	29,516	12,312

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

Note: This table shows results for our metric for differences in menus (DIM) based on our Definition 2, with \hat{DIM} in units of basis points. We use the Black and non-Hispanic white and Hispanic and non-Hispanic white lender-county-month and covariate group matched samples of 30-year, new-purchase, fixed-rate, first-lien mortgages on owner-occupied, site-built properties without prepayment-penalties, balloon, interest-only, negative-amortization, or non-amortization features. Outliers for points above 4 and below -4 were excluded. P-values were computed using 2,000 draws from the asymptotic Normal distribution implied by \hat{p}_1, \hat{p}_2 using our procedure in Section 4, and confidence intervals are computed through inversion of the hypothesis test. The sample sizes for the conforming and FHA Black versus non-Hispanic white matched samples are equal by coincidence.

In summary, we find that lenders offered Black and Hispanic borrowers a distribution of menus that was different from the distribution they offered non-Hispanic white borrowers for conforming mortgages. In particular, 2.69 percent of Black borrowers' choices were unable to be matched to non-Hispanic white borrowers based on column (1) of Table 4, and for those unmatched Black borrowers their minimum willingness to pay to switch to unmatched non-Hispanic white borrowers' menus is at least 62 basis points (1.66/0.0269, combining with information in column (1) of Table 5). Similarly, column (3) of Table 4 shows that 1.46 percent of Hispanic borrowers' choices were unable to be matched to those of non-Hispanic white borrowers, with an average willingness to pay to switch to unmatched non-Hispanic white borrowers' menus being at least 91 basis points among those borrowers (1.46/0.0160, combining with information in column (3) of Table 5).

On the other hand, our results for FHA mortgages is more mixed: Column (2) of Table 4 shows that we cannot reject equality between the menus for Black and non-Hispanic white borrowers for FHA mortgages, and while column (4) of Table 4 shows that we are able to do so for Hispanic and non-Hispanic white borrowers at the 10 percent level, we are unable to reject a zero DIM metric in column (4) of Table 5 in terms of the average increase in rate that Hispanic borrowers would be willing to pay in order to receive the menus of non-Hispanic white borrowers.

5.3 Further analyses of differences in menus

To further explore where differences in menus occurs, we divide the sample into different LTV and FICO buckets. We detect more discrimination among the conforming, lower LTV and higher FICO (that is, more creditworthy) borrowers. In particular, focusing on the DIM metric, columns (1) and (3) of Table 6 show that we detect large differences in menus for conforming mortgages in the LTV under 75 and LTV of 75 to 80 categories, such that Black and Hispanic borrowers would be willing to pay 6-7 basis points to switch to non-Hispanic white menus. Over a 80 LTV, we can no longer say that Black and Hispanic borrowers would be willing to switch to non-Hispanic white menus. The results by FICO shown in Table 7 is more mixed, but suggests that our DIM metric is in general increasing in the FICO scores of conforming Black and Hispanic borrowers.

An internally consistent explanation for the fact that we primarily detect mortgage pricing discrimination among conforming, low-LTV borrowers is that lenders are more willing to offer discretionary discounts to the creditworthy non-Hispanic white borrowers, and less willing to do so for minority borrowers who are similarly creditworthy in terms of their underwriting variables. In particular, as shown in Ambrose, Conklin, and Lopez (2020), racial differences in default risk are very low among lower-LTV and higher-FICO borrowers, and in any case they are insured. Furthermore, it seems unlikely given the strict regulatory environment surrounding the mortgage market that lenders would condition their rate sheets and first offers based on race. Therefore, by process of elimination, we believe that the search and negotiation process, particularly for the more creditworthy borrowers, may play an important role in the within-lender disparate outcomes we find.

Table 6: Analysis of differences in menus comparing borrowers across categories of loan-to-value (LTV) ratio

	Black vs Non-Hispanic White		Hispanic vs Non-Hispanic White	
	(1) Conforming	(2) FHA	(3) Conforming	(4) FHA
<i>Panel A: Test of Equality in Menus (\hat{T})</i>				
LTV ≤ 75 (\hat{T})	11.01**	-	8.26***	-
95% CI	[0.34, ∞)	-	[2.45, ∞)	-
$75 < \text{LTV} \leq 80$ (\hat{T})	7.82***	-	7.30***	-
95% CI	[2.55, ∞)	-	[4.08, ∞)	-
$80 < \text{LTV} \leq 90$ (\hat{T})	1.52	-	0.67	-
95% CI	[0, ∞)	-	[0, ∞)	-
$90 < \text{LTV} \leq 95$ (\hat{T})	0.93	-	0.80	-
95% CI	[0, ∞)	-	[0, ∞)	-
LTV > 95 (\hat{T})	1.82	0.17	0.90	1.77
95% CI	[0, ∞)	[0, ∞)	[0, ∞)	[0, ∞)
<i>Panel B: Difference in Menus (\hat{DIM}) Metric</i>				
LTV ≤ 75 (\hat{T})	7.37***	-	5.80***	-
95% CI	[4.91, ∞)	-	[4.16, ∞)	-
$75 < \text{LTV} \leq 80$ ($\hat{DIM}_{1 \rightarrow 2}$)	6.13***	-	5.16***	-
95% CI	[4.54, ∞)	-	[4.33, ∞)	-
$80 < \text{LTV} \leq 90$ ($\hat{DIM}_{1 \rightarrow 2}$)	0.18	-	-1.30	-
95% CI	[-1.04, ∞)	-	[-2.25, ∞)	-
$90 < \text{LTV} \leq 95$ ($\hat{DIM}_{1 \rightarrow 2}$)	-0.29	-	0.22	-
95% CI	[-0.91, ∞)	-	[-0.32, ∞)	-
LTV > 95 ($\hat{DIM}_{1 \rightarrow 2}$)	-0.53	-1.76	-2.19	-1.18
95% CI	[-1.58, ∞)	[-2.45, ∞)	[-3.93, ∞)	[-1.73, ∞)
$N_{\text{LTV} \leq 75}$	840	0	3,142	0
$N_{75 < \text{LTV} \leq 80}$	3,002	0	8,492	0
$N_{80 < \text{LTV} \leq 90}$	1,714	0	3,988	88
$N_{90 < \text{LTV} \leq 95}$	4,844	86	9,742	162
$N_{\text{LTV} > 95}$	2,396	9,302	4,152	12,060

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

Note: Panel A shows results for our test of equality in menu (\hat{T}) in units of percentage points based on our Definition 1, and Panel B shows our metric for differences in menus (\hat{DIM}) in units of basis points based on our Definition 2, with \hat{DIM} in units of basis points. We match on lender-county-month and covariate groups for 30-year, new-purchase, fixed-rate, first-lien mortgages on owner-occupied, site-built properties without prepayment-penalties, balloon, interest-only, negative-amortization, or non-amortization features. Outliers for points above 4 and below -4 were excluded. Confidence intervals are computed through inversion of the hypothesis test as described in Section 4, using 2,000 draws from the asymptotic Normal distribution implied by \hat{p}_1, \hat{p}_2 . Categories with under 500 observations are omitted.

Table 7: Analysis of differences in menus comparing borrowers across categories of FICO scores

	Black vs Non-Hispanic White		Hispanic vs Non-Hispanic White	
	(1) Conforming	(2) FHA	(3) Conforming	(4) FHA
<i>Panel A: Test of Equality in Menus (\hat{T})</i>				
FICO ≥ 740 (\hat{T})	3.08 ***	9.23	2.59 ***	1.78
95% CI	[0.37, ∞)	[0, ∞)	[1.04, ∞)	[0, ∞)
700 \leq FICO < 740 (\hat{T})	3.15	7.79	0.91	9.35 **
95% CI	[0, ∞)	[0, ∞)	[0, ∞)	[0.87, ∞)
660 \leq FICO < 700 (\hat{T})	4.38	2.22	3.99	6.53 ***
95% CI	[0, ∞)	[0, ∞)	[0, ∞)	[1.93, ∞)
620 \leq FICO < 660 (\hat{T})	-	1.57	-	1.17
95% CI		[0, ∞)		[0, ∞)
FICO < 620 (\hat{T})	-	1.67	-	0.17
95% CI		[0, ∞)		[0, ∞)
<i>Panel B: Difference in Menus (\hat{DIM}) Metric</i>				
FICO ≥ 740 ($\hat{DIM}_{1 \rightarrow 2}$)	1.84 ***	-0.59	2.16 ***	-1.33
95% CI	[1.10, ∞)	[-3.13, ∞)	[1.71, ∞)	[-3.84, ∞)
700 \leq FICO < 740 ($\hat{DIM}_{1 \rightarrow 2}$)	2.19 **	-2.97	-0.74	-1.30
95% CI	[0.90, ∞)	[-5.22, ∞)	[-1.64, ∞)	[-2.83, ∞)
660 \leq FICO < 700 ($\hat{DIM}_{1 \rightarrow 2}$)	1.53	-2.09	1.16	0.72
95% CI	[-0.70, ∞)	[-3.30, ∞)	[-1.31, ∞)	[-0.23, ∞)
620 \leq FICO < 660 ($\hat{DIM}_{1 \rightarrow 2}$)	-	-	-	-1.06
95% CI		[-2.36, ∞)		[-2.14, ∞)
FICO < 620 ($\hat{DIM}_{1 \rightarrow 2}$)	-	-2.82	-	-5.66
95% CI		[-3.89, ∞)		[-7.81, ∞)
$N_{\text{FICO} \geq 740}$	9,914	354	22,964	662
$N_{700 \leq \text{FICO} < 740}$	2,002	612	4,590	1,108
$N_{660 \leq \text{FICO} < 700}$	738	2,884	1,680	4,008
$N_{620 \leq \text{FICO} < 660}$	142	4,372	282	5,188
$N_{\text{FICO} < 620}$	0	1,200	0	1,346

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

Note: Panel A shows results for our test of equality in menu (\hat{T}) in units of percentage points based on our Definition 1 and Panel B shows our metric for differences in menus (\hat{DIM}) in units of basis points based on our Definition 2, with \hat{DIM} in units of basis points. We match on lender-county-month and covariate groups for 30-year, new-purchase, fixed-rate, first-lien mortgages on owner-occupied, site-built properties without prepayment-penalties, balloon, interest-only, negative-amortization, or non-amortization features. Outliers for points above 4 and below -4 were excluded. Confidence intervals are computed through inversion of the hypothesis test as described in Section 4, using 2,000 draws from the asymptotic Normal distribution implied by \hat{p}_1, \hat{p}_2 . Categories with under 500 observations are omitted.

5.4 Robustness

In Section 5.2 we showed that lenders offered Black and Hispanic borrowers a less advantageous distribution of menus compared to observationally similar non-Hispanic white borrowers for conforming mortgages. In this section we conduct a series of robustness checks that rule out some explanations for this finding.

5.4.1 Multi-product discounts

One possible reason why minorities may receive worse menus of rates and points is that they are less likely to be offered multi-product discounts set by banks. For example, banks may offer discounts on rates and points based the amount of deposits a borrower hold with them. To check whether this explains our findings, Table 8 replicates our analysis for a sample of non-bank lenders. These are lenders that do not offer the depository services that banks do, and focus primarily on mortgage lending. As Table 8 shows, our results are robust to this sample restriction. Indeed, the point estimates for both our test of equality in menu (\hat{T}) and our differences in menus (\hat{DIM}) increased under this specification. These results suggests that multi-product discounts offered by banks is not an explanation for the racial differences in menus.

Table 8: Analysis of differences in menus after restricting to non-bank lenders

	Black vs Non-Hispanic White		Hispanic vs Non-Hispanic White	
	(1) Conforming	(2) FHA	(3) Conforming	(4) FHA
<i>Panel A: Test of Equality in Menus (\hat{T})</i>				
\hat{T}	2.38***	0.87	2.46***	1.84
95% CI	[0.77, ∞)	[0, ∞)	[0.86, ∞)	[0, ∞)
<i>Panel B: Difference in Menus ($\hat{D}\hat{M}$) Metric</i>				
Minority to white ($\hat{D}\hat{M}_{1 \rightarrow 2}$)	2.04***	-1.37	1.96**	-1.34
95% CI	[1.37, ∞)	[-2.27, ∞)	[1.37, ∞)	[-2.03, ∞)
White to minority ($\hat{D}\hat{M}_{2 \rightarrow 1}$)	-6.57	-7.56	-6.67	-7.66
95% CI	[-7.21, ∞)	[-8.36, ∞)	[-7.13, ∞)	[-8.36, ∞)
<i>N</i>	10,128	8,106	24,358	10,852

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

Note: Panel A shows results for our test of equality in menu (\hat{T}) in units of percentage points based on our Definition 1, and Panel B shows our metric for differences in menus ($\hat{D}\hat{M}$) in units of basis points based on our Definition 2, with $\hat{D}\hat{M}$ in units of basis points. We match on lender-county-month and covariate groups plus loan originator compensation for 30-year, new-purchase, fixed-rate, first-lien mortgages on owner-occupied, site-built properties without prepayment-penalties, balloon, interest-only, negative-amortization, or non-amortization features. The sample is further restricted to non-bank lenders as defined in the HMDA lender file constructed by Robert Avery. Outliers for points above 4 and below -4 were excluded. Confidence intervals are computed through inversion of the hypothesis test as described in Section 4, using 2,000 draws from the asymptotic Normal distribution implied by \hat{p}_1, \hat{p}_2 .

5.4.2 Loan originator compensation

Another possibility that may explain the differences in menus that we detect is that minority borrowers may be more likely to get mortgages from loan originators that charge more for their services. Within a mortgage bank, some loan originators may provide more services in exchange for a higher compensation, and such differences in compensation would be mechanically reflected in the menus they present to borrowers. Our Optimal Blue data contains information on loan originator compensation for a subsample of lenders. To check whether this explains our findings, we conduct a robustness check by matching on levels of loan originator compensation, rounded to the nearest 1 percent of the loan amount, in addition to our usual covariates. The results are shown in Table 9. In particular, Table 9 shows that while our \hat{T} results lost significance likely due to the impact of smaller samples, the level and significance of our differences in menus $\hat{D}\hat{M}$ results are qualitatively unchanged. In particular, the point estimate for our minority-to-white $\hat{D}\hat{M}_{1 \rightarrow 2}$ measure for conforming mortgages increased from 1.66 to 2.18 for Black borrowers and from 1.46 to 1.54 for Hispanic borrowers. This suggests that differences in loan originator compensation does not explain our results.

Table 9: Analysis of differences in menus after further matching on loan originator compensation

	Black vs Non-Hispanic White		Hispanic vs Non-Hispanic White	
	(1) Conforming	(2) FHA	(3) Conforming	(4) FHA
\hat{T}	4.62	3.48	2.13	4.83*
95% CI	$[0, \infty)$	$[0, \infty)$	$[0, \infty)$	$[0, \infty)$
Minority to white ($\hat{D}\hat{M}_{1 \rightarrow 2}$)	2.18***	0.49	1.54**	-1.09
95% CI	$[0.81, \infty)$	$[-0.94, \infty)$	$[0.55, \infty)$	$[-2.21, \infty)$
White to minority ($\hat{D}\hat{M}_{2 \rightarrow 1}$)	-4.86	-8.28	-5.87	-7.56
95% CI	$[-6.11, \infty)$	$[-9.90, \infty)$	$[-6.83, \infty)$	$[-8.51, \infty)$
N	1,818	1,818	5,290	2,438

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

Note: Panel A shows results for our test of equality in menu (\hat{T}) in units of percentage points based on our Definition 1, and Panel B shows our metric for differences in menus ($\hat{D}\hat{M}$) in units of basis points based on our Definition 2, with $\hat{D}\hat{M}$ in units of basis points. We match on lender-county-month and covariate groups plus loan originator compensation for 30-year, new-purchase, fixed-rate, first-lien mortgages on owner-occupied, site-built properties without prepayment-penalties, balloon, interest-only, negative-amortization, or non-amortization features. Outliers for points above 4 and below -4 were excluded. Confidence intervals are computed through inversion of the hypothesis test as described in Section 4, using 2,000 draws from the asymptotic Normal distribution implied by $\hat{\mathbf{p}}_1, \hat{\mathbf{p}}_2$.

5.4.3 Endogeneity of LTV

Finally, we explore the possibility that the borrower may simultaneously choose their loan-to-value (LTV) ratio along with their rate/points choice. Our focus on the rate and points choice takes as given a plausible setting in which borrowers make their mortgage size choice before shopping for rates and points. Indeed, most price comparison websites require borrowers to enter their property and loan amount information before showing offers with different rate and points choices, which supports such an assumption. Nevertheless, it is possible that some borrowers may be able to jointly optimize on points and LTV. To investigate this possibility, we show in Table 10 the results of running our algorithm on only a sample of borrowers with an LTV of exactly 80. These borrowers are more likely to be constrained in their LTV choices. Results of Table 10 show that we detect large differences in menus among this sub-population, with magnitudes that are consistent with the 75-80 LTV sample in Table 6. Thus, we conclude that joint optimization of LTV and points is likely not a driver of the racial differences in menus that we detect.

Table 10: Analysis of differences in menus after restricting to borrowers with exactly 80 LTV

	Black vs Non-Hispanic White		Hispanic vs Non-Hispanic White	
	(1) Conforming	(2) FHA	(3) Conforming	(4) FHA
\hat{T}	8.27***	-	9.08***	-
95% CI	[2.42, ∞)	-	[5.25, ∞)	-
Minority to white ($\hat{DIM}_{1 \rightarrow 2}$)	5.81***	-	6.09***	-
95% CI	[4.21, ∞)	-	[5.13, ∞)	-
White to minority ($\hat{DIM}_{2 \rightarrow 1}$)	-10.71	-	-5.87	-
95% CI	[-11.90, ∞)	-	[-11.70, ∞)	-
<i>N</i>	2,600	-	7,318	-

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

Note: Panel A shows results for our test of equality in menu (\hat{T}) in units of percentage points based on our Definition 1, and Panel B shows our metric for differences in menus (\hat{DIM}) in units of basis points based on our Definition 2, with \hat{DIM} in units of basis points. We match on lender-county-month and covariate groups for 30-year, new-purchase, fixed-rate, first-lien mortgages on owner-occupied, site-built properties without prepayment-penalties, balloon, interest-only, negative-amortization, or non-amortization features, and restrict our sample to borrowers with a loan-to-value ratio of exactly 80. Outliers for points above 4 and below -4 were excluded. Confidence intervals are computed through inversion of the hypothesis test as described in Section 4, using 2,000 draws from the asymptotic Normal distribution implied by \hat{p}_1, \hat{p}_2 .

6 Discussion

We identify a “menu problem” that confounds the estimation of mortgage pricing discrimination, which stems from potential unobserved preference heterogeneity across borrower groups. We also devise a new methodology for assessing differences in menus that is robust to such preference heterogeneity, and use it to produce new estimates of mortgage pricing discrimination. Empirically, we find that mortgage pricing differentials by race still exists, particularly among lower LTV conforming borrowers.

While the menu problem is broadly relevant, our specific methodology for addressing it is unlikely to be the final word on the problem. The main benefit of our approach is that it requires relatively few assumptions to be valid. One promising path for future methodological research is the exploration of alternative avenues of identification for assessing differences in menus.

Finally, the “menu problem” we identify suggests that the typical administrative data on borrower choices falls short of ideal in terms of assessing discrimination in menus. As such, collecting actual data on menus faced by borrowers of different races, perhaps in a field-experiment/audit study setting, would be another valuable contribution.

References

Abadie, Alberto. 2020. “Statistical Nonsignificance in Empirical Economics.” *American Economic Review: Insights* 2(2): 193–208.

Ackerberg, Daniel A., Kevin Caves, and Garth Frazer. 2015. “Identification Properties of Recent Production Function Estimators.” *Econometrica* 83(6): 2411–2451.

Agarwal, Sumit, John Grigsby, Ali Hortaçsu, Gregor Matvos, Amit Seru, and Vincent Yao. 2020. “Searching for Approval.” Working Paper 27341. National Bureau of Economic Research.

Allen, Jason, Robert Clark, and Jean-François Houde. 2014. “The Effect of Mergers in Search Markets: Evidence from the Canadian Mortgage Industry.” *American Economic Review* 104(10): 3365–96.

Allen, Jason, Robert Clark, and Jean-François Houde. 2019. “Search Frictions and Market Power in Negotiated-Price Markets.” *Journal of Political Economy* 127(4): 1550–1598.

Ambrose, Brent W, James N Conklin, and Luis A Lopez. 2020. “Does Borrower and Broker Race Affect the Cost of Mortgage Credit?” *The Review of Financial Studies* 34(2): 790–826.

Andrews, Donald W. K., and Gustavo Soares. 2010. “Inference for Parameters Defined by Moment Inequalities Using Generalized Moment Selection.” *Econometrica* 78(1): 119–157.

Bartlett, Robert, Adair Morse, Richard Stanton, and Nancy Wallace. 2021. “Consumer-lending discrimination in the FinTech Era.” *Journal of Financial Economics*.

Bartlett, Robert, Adair Morse, Richard Stanton, and Nancy Wallace. November 2019. “Consumer-lending discrimination in the FinTech Era.” *Working Paper*.

Bell, Alex. 2019. “Job Amenities & Earnings Inequality.” *Working Paper*.

Bhutta, Neil, Andreas Fuster, and Aurel Hizmo. 2019. “Paying Too Much? Price Dispersion in the US Mortgage Market.” *Working Paper*.

Bhutta, Neil, and Aurel Hizmo. 2020. “Do Minorities Pay More for Mortgages?” *The Review of Financial Studies* 34(2): 763–789.

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

Black, Harold A., Thomas P. Boehm, and Ramon P. DeGennaro. 2003. “Is there discrimination in mortgage pricing? The case of overages.” *Journal of Banking & Finance* 27(6): 1139 – 1165.

Black, Harold A, and Robert L Schweitzer. 1985. “A Canonical Analysis of Mortgage Lending Terms: Testing for Lending Discrimination at a Commercial Bank.” *Urban Studies* 22(1): 13–19.

Bocian, Debbie Gruenstein, Keith S. Ernst, and Wei Li. 2008. “Race, ethnicity and subprime home loan pricing.” *Journal of Economics and Business* 60(1): 110 – 124.

Boehm, Thomas P., Paul D. Thistle, and Alan Schlottmann. 2006. “Rates and race: An analysis of racial disparities in mortgage rates.” *Housing Policy Debate* 17(1): 109–149.

Cheng, Ping, Zhenguo Lin, and Yingchun Liu. 2015. “Racial discrepancy in mortgage interest rates.” *The Journal of Real Estate Finance and Economics* 51(1): 101–120.

Courchane, Marsha, and David Nickerson. 1997. “Discrimination Resulting from Overage Practices.” *Journal of Financial Services Research* 11(1): 133–151.

Ellis, Andrew, and Yusufcan Masatlioglu. 2021. “Choice with Endogenous Categorization.” *The Review of Economic Studies*.

Fang, Zheng, and Andres Santos. 2018. “Inference on Directionally Differentiable Functions.” *The Review of Economic Studies* 86(1): 377–412.

Fuster, Andreas, Aurel Hizmo, Lauren Lambie-Hanson, James Vickery, and Paul Willen. 2021. “How Resilient is Mortgage Credit Supply? Evidence from the COVID-19 Pandemic.” *Working Paper*.

Fuster, Andreas, Stephanie H. Lo, and Paul S. Willen. 2019. “The Time-Varying Price of Financial Intermediation in the Mortgage Market.” *Working Paper*.

Gal, Tomas, and Harvey J Greenberg. 2012. *Advances in sensitivity analysis and parametric programming*, vol. 6. Springer Science & Business Media.

Galichon, Alfred. 2016. *Optimal transport methods in economics*. Princeton University Press.

Gerardi, Kristopher, Paul Willen, and David Zhang. 2020. “Mortgage Prepayment, Race, and Monetary Policy.” *Working paper*.

Ghent, Andra C., Rubén Hernández-Murillo, and Michael T. Owyang. 2014. “Differences in subprime loan pricing across races and neighborhoods.” *Regional Science and Urban Economics* 48: 199 – 215.

Goldberger, Arthur S. 1984. “Reverse Regression and Salary Discrimination.” *The Journal of Human Resources* 19(3): 293–318.

Goldin, Claudia. 2014. “A Grand Gender Convergence: Its Last Chapter.” *American Economic Review* 104(4): 1091–1119.

Gul, Faruk, and Wolfgang Pesendorfer. 2001. “Temptation and Self-Control.” *Econometrica* 69(6): 1403–1435.

Hong, Han, and Jessie Li. 2020. “The numerical bootstrap.” *Annals of Statistics* 48(1): 397–412.

Hsieh, Yu-Wei, Xiaoxia Shi, and Matthew Shum. 2021. “Inference on estimators defined by mathematical programming.” *Journal of Econometrics*.

Hwang, Hae-shin, W Robert Reed, and Carlton Hubbard. 1992. “Compensating wage differentials and unobserved productivity.” *Journal of Political Economy* 100(4): 835–858.

Kau, James B, Lu Fang, and Henry J Munneke. 2019. “An unintended consequence of mortgage financing regulation—a racial disparity.” *The Journal of Real Estate Finance and Economics* 59(4): 549–588.

Kau, James B, Donald C Keenan, and Henry J Munneke. 2012. “Racial discrimination and mortgage lending.” *The Journal of Real Estate Finance and Economics* 45(2): 289–304.

Levinsohn, James, and Amil Petrin. 2003. “Estimating Production Functions Using Inputs to Control for Unobservables.” *The Review of Economic Studies* 70(2): 317–341.

McCloskey, Adam. 2017. “Bonferroni-based size-correction for nonstandard testing problems.” *Journal of Econometrics* 200(1): 17 – 35.

Montiel Olea, José Luis, and Mikkel Plagborg-Møller. 2019. “Simultaneous confidence bands: Theory, implementation, and an application to SVARs.” *Journal of Applied Econometrics* 34(1): 1–17.

Munnell, Alicia H., Geoffrey M. B. Tootell, Lynn E. Browne, and James McEneaney. 1996. “Mortgage Lending in Boston: Interpreting HMDA Data.” *The American Economic Review* 86(1): 25–53.

Olley, G. Steven, and Ariel Pakes. 1996. “The Dynamics of Productivity in the Telecommunications Equipment Industry.” *Econometrica* 64(6): 1263–1297.

Politis, Dimitris N., and Joseph P. Romano. 1994. “Large Sample Confidence Regions Based on Subsamples under Minimal Assumptions.” *Annals of Statistics* 22(4): 2031–2050.

Romano, Joseph P., Azeem M. Shaikh, and Michael Wolf. 2014. “A Practical Two-Step Method for Testing Moment Inequalities.” *Econometrica* 82(5): 1979–2002.

Santambrogio, Filippo. 2015. “Optimal transport for applied mathematicians.” *Birkhäuser, NY* 55(58-63): 94.

Shapiro, Alexander. 1990. “On concepts of directional differentiability.” *Journal of Optimization Theory and Applications* 66: 477–487.

Shapiro, Alexander. 1991. “Asymptotic analysis of stochastic programs.” *Annals of Operations Research* 30: 169–186.

Skanderson, David, Ivan Darius, and Tammy Butler. 2014. “Containing Fair Lending Risk.” *Mortgage Banking*.

Sommerfeld, Max, and Axel Munk. 2018. “Inference for empirical Wasserstein distances on finite spaces.” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 80(1): 219–238.

Tameling, Carla, Max Sommerfeld, and Axel Munk. 2019. “Empirical optimal transport on countable metric spaces: Distributional limits and statistical applications.” *The Annals of Applied Probability* 29(5): 2744–2781.

Tootell, Geoffrey M. B. 1996. “Redlining in Boston: Do Mortgage Lenders Discriminate Against Neighborhoods?” *The Quarterly Journal of Economics* 111(4): 1049–1079.

Villani, Cédric. 2008. *Optimal transport: old and new*, vol. 338. Springer Science & Business Media.

Woodward, Susan E. 2008. “A Study of Closing Costs for FHA Mortgages.” *U.S. Department of Housing and Urban Development Report*.

Do Lenders Still Discriminate? A Robust Approach for Assessing Differences in Menus

Appendix

This appendix provides proofs and supplements the empirical analysis of Zhang and Willen (2021). Below is a list of the sections contained in this appendix.

Table of Contents

A.1	Proofs of results	2
A.1.1	Proof of Theorem 1	2
A.1.2	Proof of Theorem 2	2
A.1.3	Proof of Theorem 3	2
A.1.4	Proof of Corollary 1	4
A.2	Additional Tables and Figures	5

A.1 Proofs of results

A.1.1 Proof of Theorem 1

Proof. In the forward direction, if such a $\pi(x_1, x_2)$ exists, then it is possible to construct a series of menus $\mathbf{m} = \{x_1, x_2\}$ for each $\{x_1, x_2\}$ such that $\phi(x_1, x_2) = 1$, each appearing with probability $g_1(\mathbf{m}) = g_2(\mathbf{m}) = \pi(x_1, x_2)$, from which group 1 consumers chose x_1 and group 2 consumers chose x_2 . Under this construction, the distributions of menus across the two groups are equal such that $\mathbf{M}_1 = \mathbf{M}_2$, and the choice probabilities are rationalized. The reverse direction follows from the fact that, denoting $c(x_1, x_2|\mathbf{m})$ by the probability that group 1 consumers choose $x_1 \in \mathbf{m}$ and group 2 consumers choose $x_2 \in \mathbf{m}$ given a menu \mathbf{m} , we can construct the required $\pi(x_1, x_2) = \sum_{\mathbf{m}} c(x_1, x_2|\mathbf{m})g(\mathbf{m})$ for any $c(x_1, x_2|\mathbf{m})$, $g(\mathbf{m})$ where $g(\mathbf{m}) = g_1(\mathbf{m}) = g_2(\mathbf{m})$. \square

A.1.2 Proof of Theorem 2

Proof. Under constant marginal utility in interest rate, Equations (11) and (12) imply that $u_i(\mathbf{m}_j) - u_i(\mathbf{m}_i) = \beta d_{i \rightarrow j}(\mathbf{m}_i, \mathbf{m}_j) \geq \beta \underline{d}_{1 \rightarrow 2}(x_1, x_2)$ for some constant β . Then, by property of minimum we must have $\beta \underline{DIM}_{1 \rightarrow 2} \leq \Delta W_{\mathcal{I}_1 \rightarrow \mathcal{I}_2, \boldsymbol{\pi}}$, $\forall \boldsymbol{\pi}$, with $\beta \equiv 1$ if utility were measured in terms of interest rate paid. \square

A.1.3 Proof of Theorem 3

Proof. First, we show that φ is Gâteaux directionally differentiable in the sense that the limit:

$$\varphi'_{\mathbf{p}_1, \mathbf{p}_2}(\mathbf{h}_1, \mathbf{h}_2) = \lim_{t \rightarrow 0^+} \frac{\varphi(\mathbf{p}_1 + t\mathbf{h}_1, \mathbf{p}_2 + t\mathbf{h}_2) - \varphi(\mathbf{p}_1, \mathbf{p}_2)}{t} \quad (37)$$

exists for all and is equal to that given by Equation (22) for $\{\mathbf{p}_1, \mathbf{p}_2\} \in \mathbb{D}_\varphi$ and $\{\mathbf{h}_1, \mathbf{h}_2\} \in \mathbb{D}_0$. To do this, without loss of generality letting $\phi^* = \phi + M$, $M = \sup |\phi|$ such that $\phi^* \geq 0$ and $\varphi = \varphi^* - M$, we transform the problem to standard linear programming form:

$$\varphi^* = \min_{\boldsymbol{\pi}} E_{\boldsymbol{\pi}} \phi^* \text{ s.t. } E_{x_2} \boldsymbol{\pi} \geq \mathbf{p}_1, E_{x_1} \boldsymbol{\pi} \geq \mathbf{p}_2, \boldsymbol{\pi} \geq 0 \quad (38)$$

and then use Theorem 3.1 of Gal and Greenberg (2012), who set out conditions for the Gâteaux directional differentiability of standard linear programs with inequality constraints. In particular, we need to check primal and dual stability in the sense that, if we let $\Pi^*(p_1, p_2)$ be the set of primal solutions to the linear programming problem in Equation (38), then the

set of primal solutions reachable from perturbations in the direction $\{h_1, h_2\}$ is non-empty, such that:

$$\Pi^\infty(\{\mathbf{p}_1, \mathbf{p}_2\}, \{\mathbf{h}_1, \mathbf{h}_2\}) = \{\boldsymbol{\pi} : \{\boldsymbol{\pi}^k\} \rightarrow \boldsymbol{\pi} \text{ for some } \{\epsilon_k \rightarrow 0^+\}\}, \quad (39)$$

$$\text{with } \boldsymbol{\pi}^k \in \boldsymbol{\Pi}^*(\mathbf{p}_1 + \epsilon_k \mathbf{h}_1, \mathbf{p}_2 + \epsilon_k \mathbf{h}_2) \neq \emptyset, \quad (40)$$

and analogously for dual solutions. This can be done by referencing existing results:

1. Since $\{\mathbf{p}_1 + \epsilon_k \mathbf{h}_1, \mathbf{p}_2 + \epsilon_k \mathbf{h}_2\}$ are a series of probability measures for $\epsilon_k \leq 1$, primal stability in the sense of Equation (39) is guaranteed by Theorem 5.19 in Villani (2008).
2. Similarly, dual stability is guaranteed by Theorem 1.52 in Santambrogio (2015), in particular by taking the sequence of c-concave Kantorovich potentials corresponding to $(\mathbf{p}_1 + \epsilon_k \mathbf{h}_1, \mathbf{p}_2 + \epsilon_k \mathbf{h}_2)\}$.

Therefore, φ^* is Gâteaux differentiable. Second, we show that φ^* is Lipschitz, such that the Gâteaux directionally differentiability of φ^* is equivalent to Hadamard directionally differentiability following Shapiro (1990). For the l_1 norm $l_1(\{\mathbf{p}_{1,1}, \mathbf{p}_{1,2}\}, \{\mathbf{p}_{2,1}, \mathbf{p}_{2,2}\}) = \sum_x |\mathbf{p}_{1,1}(x) - \mathbf{p}_{1,2}(x)| + |\mathbf{p}_{2,1}(x) - \mathbf{p}_{2,2}(x)|$, we will show that:

$$|\varphi^*(\mathbf{p}_{1,1}, \mathbf{p}_{1,2}) - \varphi^*(\mathbf{p}_{2,1}, \mathbf{p}_{2,2})| \leq Ml_1(\{\mathbf{p}_{1,1}, \mathbf{p}_{1,2}\}, \{\mathbf{p}_{2,1}, \mathbf{p}_{2,2}\}) \quad (41)$$

More specifically, let $\mathbf{p}_1^- = \min\{\mathbf{p}_{1,1}, \mathbf{p}_{1,2}\}$, $\mathbf{p}_1^+ = \max\{\mathbf{p}_{1,1}, \mathbf{p}_{1,2}\}$, and analogously for $\mathbf{p}_2^-, \mathbf{p}_2^+$. By construction we know that $\varphi^*(\mathbf{p}_1^-, \mathbf{p}_2^-) \leq \varphi^*(\mathbf{p}_{1,1}, \mathbf{p}_{1,2}), \varphi^*(\mathbf{p}_{2,1}, \mathbf{p}_{2,2})$ and $\varphi^*(\mathbf{p}_1^+, \mathbf{p}_2^+) \geq \varphi^*(\mathbf{p}_{1,1}, \mathbf{p}_{1,2}), \varphi^*(\mathbf{p}_{2,1}, \mathbf{p}_{2,2})$, and therefore:

$$\varphi^*(\mathbf{p}_{1,1}, \mathbf{p}_{1,2}) - \varphi^*(\mathbf{p}_1^+, \mathbf{p}_2^+) \leq \varphi^*(\mathbf{p}_{1,1}, \mathbf{p}_{1,2}) - \varphi^*(\mathbf{p}_{2,1}, \mathbf{p}_{2,2}) \leq \varphi^*(\mathbf{p}_{1,1}, \mathbf{p}_{1,2}) - \varphi^*(\mathbf{p}_1^-, \mathbf{p}_2^-) \quad (42)$$

Furthermore, we know that $\varphi^*(\mathbf{p}_{1,1}, \mathbf{p}_{1,2}) - \varphi^*(\mathbf{p}_1^-, \mathbf{p}_2^-) \leq Ml_1$ since taking the optimal plan from $\pi^- = \varphi^*(\{\mathbf{p}_1^-, \mathbf{p}_2^-\})$ and then constructing a plan $\pi^{-,*} = \pi^- + (\mathbf{p}_{1,1} - \mathbf{p}_1^-)(\mathbf{p}_{1,2} - \mathbf{p}_2^-)$ yields an upper bound for the value $\varphi^*(\mathbf{p}_{1,1}, \mathbf{p}_{1,2}) \leq E_{\pi^{-,*}}\phi = \varphi^*(\{\mathbf{p}_1^-, \mathbf{p}_2^-\}) + E_{(\mathbf{p}_{1,1}-\mathbf{p}_1^-)(\mathbf{p}_{1,2}-\mathbf{p}_2^-)}\phi \leq \varphi^*(\{\mathbf{p}_1^-, \mathbf{p}_2^-\}) + M \sum (\mathbf{p}_{1,1} - \mathbf{p}_1^-)(\mathbf{p}_{1,2} - \mathbf{p}_2^-) \leq \varphi^*(\{\mathbf{p}_1^-, \mathbf{p}_2^-\}) + Ml_1$. Similarly, we have that $\varphi^*(\mathbf{p}_{1,1}, \mathbf{p}_{1,2}) - \varphi^*(\mathbf{p}_1^+, \mathbf{p}_2^+) \geq -Ml_1$. Substituting into Equation (42) yields:

$$-Ml_1 \leq \varphi^*(\mathbf{p}_{1,1}, \mathbf{p}_{1,2}) - \varphi^*(\mathbf{p}_{2,1}, \mathbf{p}_{2,2}) \leq Ml_1 \quad (43)$$

which implies Equation (41) and that the mapping φ^* is Lipschitz.

□

A.1.4 Proof of Corollary 1

Proof. Suppose we knew the true $\mathbf{p}_1, \mathbf{p}_2$. Then, the “oracle” analogue of the critical value at level $1 - \alpha + \beta$ conditional on the true $\mathbf{p}_1, \mathbf{p}_2$ is:

$$\hat{c}_{1-\alpha+\beta}^o = \inf\{c \in \mathbb{R} : \Pr(\hat{\psi}'_{\mathbf{p}_1, \mathbf{p}_2}(\mathbf{h}_1, \mathbf{h}_2) \leq c) \geq 1 - \alpha + \beta\} \quad (44)$$

By construction, $\hat{c}_{1-\alpha+\beta} \geq \hat{c}_{1-\alpha+\beta}^o$ whenever $[\mathbf{p}_1, \mathbf{p}_2] \in \hat{\mathcal{P}}_\beta$. Therefore, if we let the event $[\mathbf{p}_1, \mathbf{p}_2] \in \hat{\mathcal{P}}_{n,\beta}$ be E and $[\mathbf{p}_1, \mathbf{p}_2] \notin \hat{\mathcal{P}}_{n,\beta}$ be $\neg E$, it follows that:

$$\Pr(\hat{\psi} \geq \hat{c}_{1-\alpha+\beta}) = \Pr(\hat{\psi} \geq \hat{c}_{1-\alpha+\beta} | E) \Pr(E) + \Pr(\hat{\psi} \geq \hat{c}_{1-\alpha+\beta} | \neg E) \Pr(\neg E) \quad (45)$$

$$\leq \Pr(\hat{\psi} \geq \hat{c}_{1-\alpha+\beta}^o) \Pr(E) + \Pr(\hat{\psi} \geq \hat{c}_{1-\alpha+\beta} | \neg E) \Pr(\neg E) \quad (46)$$

$$\leq \Pr(\hat{\psi} \geq \hat{c}_{1-\alpha+\beta}^o) + \Pr(\neg E) \quad (47)$$

By property of limsup, we know that:

$$\limsup_{n \rightarrow \infty} \Pr_{\mathbf{p}_1, \mathbf{p}_2}(\hat{\varphi}_n \geq \hat{c}_{n,1-\alpha+\beta}) \leq \limsup_{n \rightarrow \infty} \Pr(\hat{\psi}_n \geq \hat{c}_{n,1-\alpha+\beta}^o) + \limsup_{n \rightarrow \infty} \Pr([\mathbf{p}_1, \mathbf{p}_2] \notin \hat{\mathcal{P}}_{n,\beta}) \quad (48)$$

By Equation (24) and the Portmanteau Theorem, we have:

$$\limsup_{n \rightarrow \infty} \Pr_{\mathbf{p}_1, \mathbf{p}_2}((\hat{\varphi}_n \geq \hat{c}_{n,1-\alpha+\beta}^o)) \leq \alpha - \beta \quad (49)$$

By Equation (25) for the uniform confidence band, we know that:

$$\limsup_{n \rightarrow \infty} \Pr_{\mathbf{p}_1, \mathbf{p}_2}([\mathbf{p}_1, \mathbf{p}_2] \notin \hat{\mathcal{P}}_{n,\beta}) \leq \beta \quad (50)$$

Combining the two parts, we have:

$$\limsup_{n \rightarrow \infty} \Pr_{\mathbf{p}_1, \mathbf{p}_2}(\hat{\varphi}_n \geq \hat{c}_{n,1-\alpha+\beta}) \leq \beta + (\alpha - \beta) = \alpha \quad (51)$$

as required. □

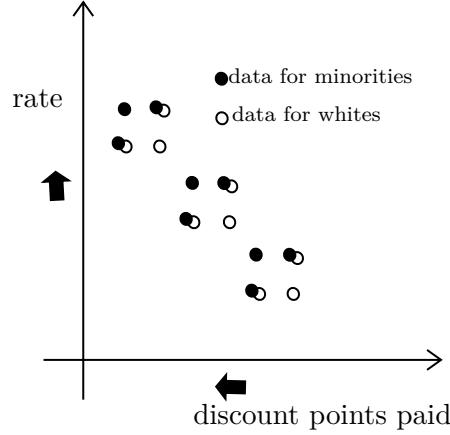
A.2 Additional Tables and Figures

Figure A.1: Screenshot of advertised rate and closing cost table

Rate	Upfront fees	
2.500% 30 year fixed refinance	\$0 Points: 0 5 year cost: \$35,348	Next Offer details
2.375% 30 year fixed refinance	\$2,510 Points: 0.67 5 year cost: \$36,052	Next Offer details
2.250% 30 year fixed refinance	\$4,712 Points: 1.404 5 year cost: \$36,451	Next Offer details

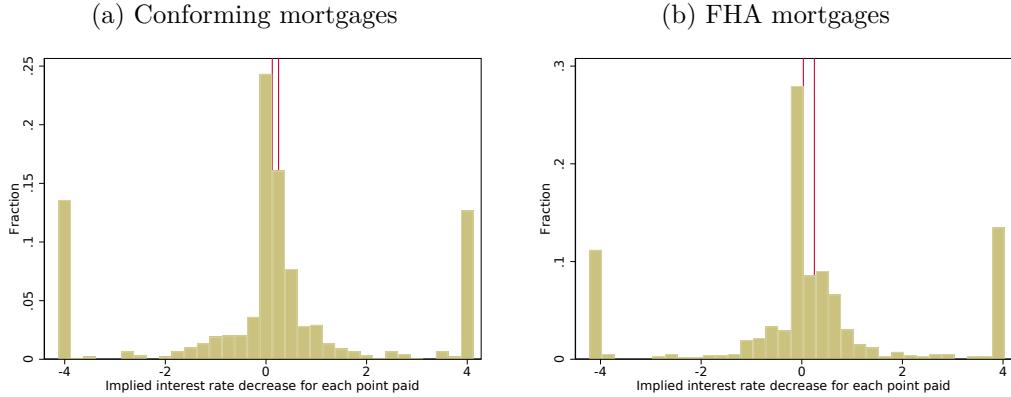
Note: This figure shows a screenshot obtained by the authors from Bankrate.com for a \$300,000 refinancing mortgage with baseline characteristics on December 19, 2020. It shows how a borrower may choose to pay 0 points for a 2.500% interest rate mortgage, 0.67 points for a 2.375% interest rate mortgage, or 1.404 points for a 2.250% interest rate mortgage. Also, lender review websites such as NerdWallet and MortgageWaldo evaluate lenders based on both interest rates and discount points/closing costs.

Figure A.2: How the choice of which menu dimension to condition on can yield contradictory results



Note: This Figure shows that a “contradictory” assessment of discrimination can appear nonlinearly, and it complements the linear regression case of Figure 2.

Figure A.3: Implied interest rate decrease per point paid, based on randomly sampled pairs of borrowers



Note: these figures are constructed using HMDA-Optimal Blue matched data discussed in Section 5. Pairs of borrowers are sampled conditional on lender, lock date, type of mortgage, LLPA group, and loan amount group. The implied interest rate decrease for each point paid is estimated by taking the pairs of borrower rate and point choices and computing $-\frac{r_2 - r_1}{y_2 - y_1}$, where r_1, r_2 are the interest rates paid by the borrowers and y_1, y_2 are the discount points paid by the borrowers. Values below -4 and above 4 are replaced with -4 and 4, respectively. The red lines in Figure A.3a represent 1/8 and 1/4, which is the plausible range of menu slopes discussed in Bartlett et al. (November 2019). The red lines in Figure A.3b represent 1/32 and 1/4, which is the plausible range of menu slopes we find based on rate sheet data, as shown in Figure 9b.

Table A.1: Other approaches to inference in optimal transport, compared with our size-corrected directional derivatives approach

	HSS (2021)	m out of n subsampling			Size-corrected directional derivatives
		$m = n^{2/3}$	$m = n^{1/2}$	$m = n^{1/3}$	
$n_1 = n_2 = 500$	0.0	31.8	21.9	16.2	4.4
$n_1 = n_2 = 1000$	0.0	28.3	18.4	12.5	4.7
$n_1 = n_2 = 5000$	0.0	20.0	11.3	8.6	4.8
$n_1 = n_2 = 10000$	0.1	18.3	10.3	8.4	5.8
$n_1 = n_2 = 50000$	0.0	13.7	8.6	6.5	5.0

Note: Entries represent the probability of rejecting the null at the 5 [percent level. The control of the subsampling approach is taken with subsample size m as indicated, with $n = \frac{n_1 n_2}{n_1 + n_2} = \frac{1}{2} n_1$. The control of our size-corrected directional derivatives approach was computed via 2000 sample draws and 500 draws of h_1, h_2 from the estimated asymptotic multivariate Normal distribution for p_1, p_2 within each sample draw.

Table A.2: Regressions of origination costs and total loan costs as a percentage of the loan amount on HMDA's information on points (hmda_points) versus Optimal Blue's information on points (ob_points)

	Origination costs			Total loan costs		
	(1)	(2)	(3)	(4)	(5)	(6)
hmda_points	0.915*** (0.00277)		0.902*** (0.00308)	0.940*** (0.00372)		0.941*** (0.00463)
ob_points		0.469*** (0.00377)	0.0200*** (0.00220)		0.474*** (0.00434)	0.00531 (0.00361)
_cons	0.580*** (0.000432)	0.560*** (0.00129)	0.575*** (0.000591)	2.194*** (0.000577)	2.174*** (0.00149)	2.190*** (0.000927)
N	1224911	1221338	1221208	1224417	1220921	1220715
R^2	0.553	0.231	0.553	0.233	0.093	0.232

Standard errors in parentheses

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

Note: The sample consists of the conforming and FHA purchase mortgages originated within the retail channel within our 2018–2019 HMDA–Optimal Blue matched sample. In each regression, we excluded observations with extreme outliers for points (below -4 or above 4) and for origination costs and total loan costs as a percentage of the loan amount below -3 percent or above 10 percent. All regressions include lender by county by year by product type fixed effects. Standard errors were also clustered at the lender by county by year by product type level.

Table A.3: Assessments of lender discrimination using two heuristic approaches in the Black and non-Hispanic white matched sample

	Heuristic 1		Heuristic 2				Alternate Heuristic 1	
	Conforming	FHA	Conforming		FHA		Conforming	FHA
	(1) points	(2) points	(3) rate_1/8	(4) rate_1/4	(5) rate_1/32	(6) rate_1/4	(7) rate	(8) rate
black	4.482*** (0.994)	-0.908 (1.319)	2.905*** (0.327)	3.309*** (0.376)	3.032*** (0.436)	2.519*** (0.509)	2.627*** (0.322)	2.983*** (0.437)
Rate Decile FE	Yes	Yes	No	No	No	No	No	No
Points Decile FE	No	No	No	No	No	No	Yes	Yes
<i>N</i>	12278	9208	12278	12278	9208	9208	12278	9208
<i>R</i> ²	0.041	0.025	0.006	0.006	0.005	0.003	0.023	0.025

Standard errors in parentheses

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

Note: This table presents heuristic analyses of data for our Black and non-Hispanic white lender-county-month and covariate-matched sample of 30-year, new-purchase, fixed-rate, first-lien mortgages on owner-occupied, site-built properties without prepayment-penalties, balloon, interest-only, negative-amortization, or non-amortization features. Outliers for points above 4 and below -4 were excluded, and rate spreads below -55 basis points and above 90 basis points were excluded. Columns (1) and (2) presents OLS regressions of points on a dummy variable indicating if a borrower is black along with deciles of rate, for conforming and FHA mortgages. Columns (7) and (8) presents an alternative version of this heuristic by running a OLS regression of rate on a dummy variable indicating if a borrower is black along with deciles of points. Columns (3)-(6) shows results of OLS regressions on the “points-adjusted rate” where the adjustment factor is shown in the underscore.