Testing Models of Economic Discrimination Using the Discretionary Markup of Indirect Auto Loans

Jonathan A. Lanning [1]

I am grateful for helpful comments from Robert Adams, Brian Bucks, Benjamin Keys, Brian Kovak, Christopher Palmer, Shanthi Ramnath, and seminar/conference participants at the NBER conferences on Innovative Data in Household Finance and Discrimination in the Financial System, the Consumer Financial Protection Bureau, the Schar School of Policy and Government at George Mason University, the Society of Government Economists annual conference, the Allied Social Sciences Association annual meeting, the Federal Reserve Bank of Chicago, and the Philadelphia Federal Reserve's biennial conference on Auto Lending. Deepra Yusuf and Christopher Stocks provided valuable assistance in assembling the data.

[1] Federal Reserve Bank of Chicago. e-mail jonathan.lanning@chi.frb.org. Note: much of the work on this project was conducted while the author was employed by the Consumer Financial Protection Bureau, and the data describing auto loans that this paper relies on are proprietary to the CFPB. However, this project is the result of the author's independent research and does not necessarily represent the views of the Federal Reserve Bank of Chicago, or the Consumer Financial Protection Bureau. **This work is a preliminary draft, and should not be circulated or cited without the author's permission.**

1. Introduction

The economics literature has been concerned with discrimination since at least 1957, but few studies have directly tested various competing models of economic discrimination.[1] However, understanding which models of economic models of discrimination are at play in the market is quite important, as models of taste-based, statistical, and search-based discrimination each imply different conclusions about the sources and persistence of discrimination.[2] For example, if observed disparities are due to taste-based discrimination, then it is puzzling why and how they persist, and this persistence likely implies diffuse prejudice or a lack of competition. If the disparities are due to statistical discrimination, then the disparities themselves are potentially efficient, but raise the question as to why there has not been investment in a technology to reduce the informational asymmetry.

Effective antidiscrimination policy also depends on understanding the types of discrimination affecting market outcomes; in fact, the same anti-discrimination policy could either mitigate or exacerbate disparity, depending on that disparity's root cause.[3] For example, affirmative action-type policies may effectively combat taste-based discrimination if they impose a sufficiently high cost on indulging prejudice. But these same policies could mute signals sent by the disfavored group, widening the informational asymmetries at the heart of a statistical discrimination model (Coate and Loury 1993). As such, distinguishing between these alternative theories of discrimination is an important task for those interested in understanding and combating discrimination in markets.

---

[1]Becker (1957) is generally considered to be the seminal work on the economics of discrimination. Lang and Lehmann (2012) provides an excellent overview of the studies that test empirical evidence of discrimination in labor markets to assess consistency with a single model (or develop a model of discrimination that is generally consistent with the evidence). There is a smaller literature exploring multiple potential sources of discrimination, though most of these papers rely on data from idiosyncratic sources, like game shows (e.g. (Levitt 2004), (Antonovics et al. 2005)) or the labor markets in sports (e.g. (Lanning 2010)), to obtain the necessary data on productivity.

[2]See Altonji and Blank (1999), Cain (1986), Lang and Lehmann (2012) for surveys of the literature on models of economic discrimination and their policy implications.

[3]See Lang and Lehmann (2012) at 8.1.

That said, the gap that currently exists in the literature is understandable. To clearly differentiate between (or decomposing the relative effects of) these models requires data describing, at the very least: i) individuals' true productivity or skill; ii) clear and predictive signals of individuals' productivity or skill; iii) the distribution of prejudice across agents in the market; and iv) individuals' search costs (which in turn depend on difficult to observe parameters, such as the share of prejudiced or discriminating agents in a market). To date, empirical studies of discrimination have generally lacked data that simultaneously encompass all of these dimensions. In the absence of these data, most empirical tests of discrimination have been "indirect," in that they assess consistency with predictions of the models, as opposed to directly linking the hypothesized source of discrimination to the observed disparities.

In this paper I conduct direct empirical tests of the theoretical roots of observed racial disparities using a newly available administrative data set of indirect auto loans. Auto loans are typically the largest loans for non-home buyers, and most consumers will take out many more auto loans than home loans in their lives. As such, it has been noted that "the magnitude and relative frequency of vehicle purchases suggest that differential treatment by race in the vehicle market may have important implications for differences in wealth and financial well-being by race" (Charles et al. 2008, p. 315). Yet there have been surprisingly few studies of this market.[4] This also means that any differential experiences with these loans can lead to substantial differences in the financial outcomes experienced by consumers.[5] The market for indirect auto loans in particular is a large and complex one. Indirect auto loans are those arranged by a dealer on a borrower's behalf and estimates suggest they comprise more than 80% of all auto loans (Grunewald et al. 2020). At its most simple, an auto transaction progresses as follows: consumers first search for a make/model of vehicle, then search for a dealer, then bargain over a price for that vehicle, and then (and only

---

[4]Charles et al. (2008), Attanasio et al. (2008), Cohen (2012), and Grunewald et al. (2020) are among the limited number of studies focusing on auto loans.

[5]See Blau and Graham (1990), Barsky et al. (2002), and Charles and Hurst (2002) for examples of how race-based disparities in financial markets can impact the Black-White wealth gap.

then) negotiate financing terms. These terms are largely based on information such as risked-based interest rate, rate/term structure options, etc. that are observable only by the dealer. In other words, customers are often “locked in” to a deal once the financing stage has been reached, and are at an informational disadvantage in these interactions. This gives the dealer substantial leverage in these negotiations. In this last stage of the transaction, a dealer may “markup” an indirect loans by as much as 250 basis points (2.5 percentage points) over the risk-based “buy rate” at which the lender is willing to extend credit. This markup is levied solely at the discretion of the dealer, and the dealer receives additional compensation from the lender in exchange for the increase in rate. Moreover, the buyer has no information about whther his or her loan is marked up, let alone by how much. The discretionary nature of this dealer markup, the relatively light competition ensured by the asymmetric information, and the high search costs, make for a fascinating, yet underexplored market.[6] These factors also make this an ideal market for testing the classic theories of economic discrimination.

Using data from the General Social Survey (GSS) and proprietary supervisory auto data from the Consumer Financial Protection Bureau, I perform direct tests of a Becker-style taste-based discrimination, a standard model of statistical discrimination, and a random search model with discrimination. Because the data provide detailed information on the entire transaction, I can use the negotiated price of the vehicle to reject the hypothesis that differences in negotiation skill is behind the disparity in markup. Then, following Charles and Guryan (2008) I use measures of prejudice derived from the GSS to test the sharp prediction of the Becker model that racial disparity in markup will be closely related to the marginal level of prejudice in a market, but will not be substantially impacted by the average level of prejudice. I next test for statistical discrimination by using administrative data to evaluate the relationship between the signals of financial sophistication available to the dealer (e.g., credit scores) to see whether the predictions of

[6] To my knowledge, only Cohen (2012) and Grunewald et al. (2020) have been able to directly explore markup previously, the former with a more limited set of data obtained for litigation purposes, and the latter with the same data used here.

statistical discrimination models, i.e. that consumers receive lower markup with positive signals, and Black consumers receive relatively lower returns to a positive signal, are supported by the data. Finally, I test the prediction of a basic search model of discrimination, that the share of prejudiced (unprejudiced) individuals in a population will be associated with worse (better) average outcomes for disfavored types. In each of these tests I am able to nest the predictions of taste-based models, to see if the key percentiles of the distribution of prejudice remain predictive of disparities even when signaling and search costs are accounted for.

Ultimately I find that taste-based discrimination of the form proposed by Becker is strongly consistent with the data in every specification and across all robustness checks (for example, using alternative measures of prejudice and including alternate percentiles of prejudice). I find no compelling evidence suggesting that statistical discrimination by dealers or differential search values contribute to the observed markup disparities; in fact, I find evidence that is highly *inconsistent* with each of those models. As such, it appears that taste-based discrimination on the part of dealers contributes substantially to the observed disparities in average markups paid by Black and White borrowers.

The remainder of the paper progresses as follows: in Section 2 I present some background information on the market and related literature; in Section 3 I present a simple theoretical construct for the analysis; in Section 4 I describe the data; in Section 5 I describe the empirical approach I employ to test the various models of economic discrimination; in Section 6 I present the results; and I conclude in Section 7.

2. Background

Most of the economics literature focused on discrimination falls into one of two camps: theories of economic discrimination that may (or may not) explain persistent demographic-based disparities, and empirical studies of disparities that can be attributed to discrimination, either by implication or residual disparities from regressions with many controls. The most common examples of the former are taste-based models (the seminal work being Becker (1957)) and models of statistical discrimination (e.g. Arrow (1972), Aigner and Cain (1977)), though more recent work has focused on search (e.g. Black (1995), Sasaki (1999), Lanning (2014)), wage posting (e.g. Lang et al. (2005)), and other, often more technical models. Examples of empirical studies are so numerous and varied that any attempt to list representative works here would be folly (though summaries and examples can be found in Cain (1986), Altonji and Blank (1999), Lang and Lehmann (2012), and other survey articles).

Despite the large literature on economic discrimination, few studies have tried to directly link empirical evidence of discrimination to specific theoretical underpinnings. Attempts have been made to test various forms of statistical discrimination (e.g. Altonji and Blank (1999), Arcidiacono et al. (2010), Lang and Manove (2011)) with some results consistent with various formulations of these models. However the results do not point to a clear and consistent story about differences in the initial observation of productivity vs. differential rates of learning, and struggle to distinguish between differences in productivity predictions based on information vs. prejudiced priors (which may be related). Additionally, whenever discussing statistical models of discrimination, one must question why the market would have coordinated on race as a reasonable measure along which to divide individuals. Or, as William Spriggs put in in his open letter to economists: “How does a model assume that an entire set of actors, observing the infinite diversity of human beings, all settle on race as a meaningful marker independent of history, laws, and social norms? And, miraculously, those same rational actors use statistical methods to find only negative attributes highly correlated with race.” (Spriggs 2020).

Most tests of taste-based models fall into the “indirect” category, testing consistency of outcomes with the predictions of models. For example, these tests will assess the relative returns observed for loans or job offers made to Black and White individuals, and conclude that taste-based discrimination exists if Black individuals are systematically more profitable than their White counterparts (Becker 1957). There are at least two difficulties with this approach. First, it is unclear if this test differentiates between taste-based models with racial animus and statistical models with inaccurate priors. Some recent studies have shown that biased beliefs can be eliminated from the market over time, in much the same way as Becker predicts prejudice may be cleared from the market (Bohren et al. 2019). However, Coate and Loury (1993) show that inaccurate beliefs can manifest in self-fulfilling disparities over time. As such, outcome tests may be capturing the results of racial animus that has yet to be cleared from the market, or inaccurate beliefs that are in the process of either leaving the market, or actually altering market outcomes. The second issue with outcome tests is they struggle to deal with the infra-marginality problem, where even absent discrimination the profitability of outcomes could differ. Accounting for this issue has shown that outcome tests may yield misleading results (Ayers 2002; Simoiu et al. 2017).

One of the only papers to directly test a taste-based model of discrimination is Charles and Guryan (2008).[7] In that article the authors exploit a key facet of Becker's classical taste-based model of discrimination to evaluate whether patterns in labor market data are related to inputs of the model. Their results are clear: the Black-White wage gap is closely related to the marginal level of prejudice against Blacks in a market, but not the average level of this prejudice. These findings are strongly consistent with the specific predictions of the Becker model.

But the Charles and Guryan study is not without its weaknesses. Lang and Lehmann (2012) notes that the Charles and Guryan approach assumes "that all firms are the same size, that black and white workers are perfectly segregated, that there is no consumer prejudice (or at least that the market can avoid it), that the distribution of prejudice is the same among employers as among the population as a whole and that the labor force participation rates of black and white workers are the same," assumptions which are "unreasonably strong" (Lang and Lehmann 2012, p.972). Additionally, there are historical trends in labor market disparities that are difficult to reconcile

[7] Some other exceptions are papers exploiting data from game shows (e.g. Levitt (2004), Antonovics et al. (2005), Anwar (2012)), or sports (e.g. Szymanski (2000), Lanning (2010), which consider the labor market for athletes, and List (2004) which considers the market for sports memorabilia). However, the former utilize data from non-market transactions, and the latter explore markets that are small, and may be highly idiosyncratic. This leaves Charles and Guryan (2008) as arguably the most comprehensive and general test of a specific model of taste-based discrimination.

with the trends in Charles and Guryan's prejudice index (although the latter measure might be changing in measured value rather than underlying prejudice, if GSS respondents have become more uncomfortable truthfully revealing racial animus over time). Still, the Charles and Guryan study stands as the sharpest test of Becker-type discrimination in the literature, and supports that prejudice plays a significant role in the Black-White wage gap.

Although most of the studies mentioned to this point have been focused (explicitly or implicitly) on labor markets, these markets may not be ideal for understanding the disparities that face disfavored groups in the broader economy. Labor interactions are often repeated games where agents are able to learn, strategically signal, and reoptimize over time. Additionally, there are myriad considerations outside the direct employer/employee interactions that may impact hiring decisions (e.g. coworker or customer prejudice). Finally, direct signals of productivity (e.g. education) might matter less to potential employers than non-quantifiable "soft skills" that signal the laborer would be enjoyable to work with; such signals are problematic for quantitative analysis, as they are observed by the employer, but not by the econometrician. Many of the assumptions underpinning static models of discrimination (e.g. difficulty in avoiding prejudice actors, minimal learning, no unobserved signals, etc.) may be more reasonable in a market for a single, discrete transaction (like many of those in consumer finance markets), rather than a repeated, long-term economic relationship like a labor contract.

Not coincidentally, bias and discrimination have also been explored in household finance markets (e.g. Morse and Pence (2020), Bhutta and Hizmo (2021), Bartlett et al. (2021), Dobbie et al. (2020)). Of particular interest to this paper are the studies of discrimination in the market for automobiles, and the financing thereof. This is actually quite a complex market. Purchasing an automobile requires multiple stages of search (Busse and Silva-Risso 2010; Grunewald et al. 2020). First, a buyer must decide on his/her desired vehicle model (or models) and ideal options, then shop across dealers to optimize the price, then, in most transactions, negotiate over financing, supplementary insurance, and add-ons (such as warranties, service packages, etc.). Faltering at a late stage in this process (e.g. walking away from a deal due to interest rate) can lead to the customer having to repeat the earlier stages and bear all of the associated costs again.

One component of this transaction that is particularly opaque to consumers is the practice of "dealer markup" of interest rates on indirect loans. Indirect loans are those secured from a financial institution by the dealer on behalf of the buyer. This saves the buyer the time and effort of securing his/her own financing, and allows the dealer to serve as a "one-stop shop" for the buyer. This is a very popular option for consumers, and estimates suggest that more than 80% of cars financed in the U.S. have that financing arranged through a dealer (Davis 2012). In exchange for this service, financial institutions compensate dealers in two ways: they pay a flat fee to the dealer, and/or they allow the dealer to increase the interest rate offered to the buyer and to share in the additional revenue this produces. Markup is typically capped at between 200 and 250 basis points (2.0 and 2.5 percentage points), and dealers are free to impose any allowable amount of markup at their discretion.[8] Note that this markup is completely hidden from the consumer; the only number reported to the consumer by the dealer is the interest rate inclusive of any added markup. In fact, only 21% of buyers claim they are aware that a dealer can adjust rates without consent, and it is unclear if they are aware of how the markup process actually works (Davis 2012). Additionally, consumers may lack information about the potential interest rates available to them. Unlike with mortgages, there are relatively few online resources that allow borrowers to aggregate and compare rates across many lenders. It also isn't clear how the distribution of rate available in the market for direct loans (where the customer contracts directly with a lender outside of the transaction with he dealer) compares to those available in the indirect market, as dealers may be able to more easily search across multiple lenders, benefit from volume discounts if they source many loans to select lenders, etc. In addition, factors related to the price of a vehicle like MSRP, invoice cost, etc., which are fixed for each vehicle can be found via numerous resources on the internet, while

[8]The specific markup caps are often dealer and/or deal-specific, though none of the lenders in the sample permit markups in excess of 250 basis points. For transactions lenders consider more "risky," markup can also be capped at lower levels (e.g. 100 or 150 basis points) to mitigate the increased risk of default that may accompany the imposition of higher markups on less creditworthy borrowers.

consumers can only learn about the market interest rates available to him/her by going through a formal application process. This is likely to result in consumers focusing on the pricing of a vehicle prior to negotiations, rather than the financing terms. Finally, the interest rate may not even be the aspect of the actual loan upon which a buyer is likely to focus, as term and loan amount (a function of price, down payment, and add-ons) are decided upon simultaneously with the rate, and may each impact rate in ways that are not obvious to the consumer. In such a complex process, consumers may focus on a more simple “index” measure of the deal, like monthly payment. Argyle et al. (2020a) provide direct evidence of exactly this fact, and the general lack of salience is further supported by the fact that 61% of consumers report being unaware of the APR on their auto loans (Davis and Frank 2011). As such, the dealer likely has very few impediments to marking a loan up, even if the buyer is financially savvy.

What is particularly intriguing about this practice from an economic discrimination standpoint is that the markup is purely discretionary. Dealers may (or may not) vary the markup based on any factors they see fit, e.g. the buyer’s perceived financial sophistication or negotiation skills, the rents already extracted on the deal, the dealer’s personal feelings toward the buyer, etc. This type of opaque, purely discretionary pricing in a market where it is difficult to shop is rare, and well-positioned to reveal the underpinnings of economic discrimination that may manifest less (or differently) in more competitive and fully informed markets. However, these features have not piqued the interest of the economics literature nearly as much as has discrimination in the negotiated prices of vehicles.

In general, the literature focusing on discrimination in the prices paid for vehicles focuses relatively more on the economic models underlying disparities than do studies of labor markets. This is perhaps due to the fact that the key prices in labor markets are a function of productivity, which is much more difficult to estimate/observe than is the price of a vehicle. Regardless, many of the studies of car purchases have explicit implications for the economic models driving the disparity (though few have specific tests for those models). For example, Ayers and Siegelman (1995) conduct a field experiment showing that auto dealers quote lower prices to White males compared to Black males, but the resulting data do not strongly support a particular model of discrimination. Goldberg (1996) uses the Consumer Expenditure Survey to show that, conditional on other factors (e.g. vehicle model, dealer financing, first-time buyers, etc.), there does not seem to be racial disparity in auto prices. However, the author determines that there is greater dispersion in these other factors for Black customers, implying that statistical discrimination that might rationalize the disparities documented in Ayers and Siegelman (1995). Morton et al. (2003) shows that minority buyers who shop for vehicles on the internet pay nearly the same prices as do Whites, controlling for consumers' income, education, and neighborhood characteristics. This is consistent with either a theory of taste-based discrimination or a theory of statistical discrimination (as dealers cannot observe a buyer's race/ethnicity in order to manifest any distaste, nor discern any demographic-related signals about willingness to pay). But, again, these studies focus on the prices negotiated for particular vehicles. These prices are dependent on easier to observe factors compared to financing, especially since buyers are almost certainly more likely to use the internet to research these purchases now compared to the years in which these studies were conducted.

There has been comparatively less work done on exploring the financing of vehicles, and the potential economic disparities arising in these transactions. Some notable exceptions are Attanasio et al. (2008), Cohen (2012), Charles et al. (2008), and Butler et al. (2021). Attanasio et al. (2008) uses Consumer Expenditure Survey data to explore the effects of borrowing constraints on auto loans, but do not explore racial disparities or the impact of markup. Cohen (2012) analyzes data from captive auto lenders (financing companies associated with a particular manufacturer) to show that there is a great deal of dispersion in the markup of auto loans, and that Black and Hispanic borrowers were systematically charged higher markups compared to their non-Hispanic White counterparts. Charles et al. (2008) use the Survey of Consumer Finances and quantile regressions to show that most of the racial difference in interest rates in the vehicle purchase market occurs at the percentiles above the median. However, they also show that, conditional on financing a loan through a "traditional" bank, there is no discernible disparity in the rates paid by Black borrowers compared to White borrowers. They do, however, find differential treatment of borrowers at the types of finance companies analyzed by Cohen (2012), and ask the question "why [do] Blacks finance their loans at vehicle finance companies at all, given that they pay higher rates of interest there?" (Charles et al. 2008, p. 319). The authors posit that it could be financial literacy or differential probabilities of denial, but can perform no direct tests. Additionally, as their data contain only coarse measure of interest rates, they are unable to determine if disparities in markup exist between the Black and White borrowers at either type of lending institution.

Most similar to this paper is the work of Butler et al. (2021). That paper links consumer credit record to HMDA data in order to more precisely identify the race of individuals taking auto loans, and observe factors not available in credit records (e.g. income). They then assess underwriting and pricing disparities, as well as the performance of those loans in a Becker outcome test framework, and conclude that tast-based discrimination is likely. However, their sample is constrained to borrowers who appear in HMDA, in which Black individuals are systematically unrepresented. They are also unable to differentiate between algorithmic rate setting and discretionary components of pricing, direct vs. indirect loans (which typically have different pricing models), and are unable to observe key factors in pricing and underwriting decisions (e.g. vehicle characteristics).

In contrast to previous studies, I am able to use broad administrative data from many financial institutions. These data include all the objective variables used by the institutions to underwrite and price the loan, the financial characteristics of borrower observable to the dealer (e.g. FICO, risk-based interest, income), vehicle information (e.g. make, model, year, new/used), and negotiated terms of the transaction (e.g. price paid, add-ons, etc.). I also separately observe buy rate and markup for each deal. This allows for a more comprehensive and complete assessment of this under-explored market, as well as facilitating tests for specific forms of economic discrimination that may be driving disparity in the market.

3. Economic Theories of Markup and Discrimination

It is important to be specific about the models of economic discrimination being tested, prior to developing and conducting tests of those models. First, I develop a simple theoretical framework for markup. Then I offer key predictions from three representative models of economic discrimination: taste-based discrimination, statistical discrimination, and search with discrimination. These models of discrimination are more carefully developed in Appendix A.

3.1. A Theory of Markup

In effect, a dealer bears no direct cost when adding markup to a loan, and derives constant, linear revenues from each additional basis point of markup added. However, most lenders pay a relatively small flat fee to the dealer for each deal that does not include markup. As such, the dealer has a strong incentive to markup each loan, subject only to the limits on markup imposed by a lender, and the fear that too much markup will cause the buyer to either walk away from the entire deal, or seek financing from outside sources (in which case the dealer loses both the markup, and the flat fee). Notably, the rents earned from the negotiated price of the vehicle do not directly factor in to the dealer's desired markup (although, for especially low rent deals the dealer may actually prefer to lose the entire deal, rather than receiving only the flat fee or the compensation from a small amount of markup).

Let m_i be the markup imposed by the dealer on buyer i, M be the maximum allowed markup by the lender, $\bar{m}_i$ be the borrower's reservation markup (i.e. the maximum markup buyer i is willing to pay), ϕ_i be the value of the flat fee paid to the dealer for buyer i's deal absent any markup, and r_k be the rents earned by dealer k from the non-financing aspects of the deal (e.g. from the negotiated price of the vehicle). Then, based on the markup imposed and the borrower's reservation level of markup, the dealer receives the following levels of utility:

$$\begin{aligned} m_i > \bar{m}_i \quad &\Longrightarrow \quad U_k = \{0, r_k\} \\ m_i = 0 < \bar{m}_i \quad &\Longrightarrow \quad U_k = r_k + \phi_i \\ 0 < m_i \leq \bar{m}_i \quad &\Longrightarrow \quad U_k = r_k + m_i \end{aligned} \tag{1}$$

.

That is, if a dealer imposes too high a markup, then the dealer either loses the entire deal (earning 0), or only those rents achieved from the non-financing aspects of the deal; if a dealer imposes no markup (and the buyer is willing to pay the non-marked up rate), the dealer earns the non-financing rents plus a flat fee; and if a dealer imposes a markup between 0 and the buyer's reservation markup, the dealer earns the non-financing rents plus the value of the markup. However, the dealer cannot directly observe the buyer's reservation level of markup. As such, the dealer dealer's expectation of $\bar{m}_i$ will determine the optimal markup to be imposed, such that:

$$\begin{aligned} E[\bar{m}_i] \leq \phi_i \quad &\implies \quad m_i = 0 \\ E[\bar{m}_i] \geq M \quad &\implies \quad m_i = M \\ \phi_i < E[\bar{m}_i] < M \quad &\implies \quad m_i = E[\bar{m}_i]. \end{aligned} \tag{2}$$

In words: if a dealer expects the buyer's reservation markup is below the value of the flat fee, the dealer will impose no markup; if a dealer expects a buyer's reservation markup is greater than or equal to the markup allowable markup, the dealer will impose the maximum markup; and if the dealer believes the buyer's reservation markup is between these levels, then the dealer will set markup equal to the buyer's expected reservation markup.

The question then becomes how a dealer sets its expectation of a buyer's reservation markup. A general theory of how a dealer forms its expectation of $\bar{m}_i$ could take the form:

$$E[\bar{m}_i] = f(p_i, s_i, \epsilon_i), \tag{3}$$

where p_i is a proxy of buyer i's negotiation skill as measured by the relative price paid for the vehicle, s_i is a proxy of buyer i's financial sophistication (σ_i), and ϵ_i is an idiosyncratic term specific to the interaction between buyer i and the dealer. Note that this allows negotiation skill and financial sophistication to vary independently from one another, as many skilled negotiators may not realize the extent to which they can negotiate their interest rate.[9] Given the dealer's discretion to impose markup, and the lack of transparency into the markup process, an important question remains:

[9] It is also conceivable that some individuals would have high financial sophistication–e.g. know about the practice of dealer markup–but not be strong enough negotiators to take full advantage of that information. However, as few consumers are aware of the existence of markup, and even those who are aware are unlikely to have substantial negotiating leverage, this possibility seems relatively remote.

what would keep the dealer from assigning the maximum markup to all customers? In this model, only the risk of losing all rents from the transaction if the customer rejects the entire deal due to the high interest rate (i.e. $U = 0$), or losing the returns from financing the deal if the customer secures alternative financing due to the high interest rate (i.e. not realizing m_i or ϕ).[10]

3.2. Key Predictions from Common Models of Economic Discrimination

Becker-style taste-based models of discrimination rely on one agent's distaste or animus for another. In the case of markup, this would arise if the dealers assigning markup were prejudiced against Black customers. The predictions here are straightforward: prejudiced dealers require additional rents from deals with Black customers compared to deals with White customers, in order to compensate the dealers for their distaste. But an oft-overlooked feature of these models is that it is the marginal level of prejudice in a market that determines the equilibrium level of disparity. That is, neither the extreme levels of prejudice, nor the average level of prejudice, would affect the observed (average) disparity. The intuition for this result is familiar. In effect, the more extreme values of prejudice in the market are avoided by disfavored groups, leaving only the most prejudiced agents with whom a transaction is closed to set the market level of disparity. This is directly analogous to any observed equilibrium market price, which is related to the marginal

[10] In practice, there may be other considerations as well. For example, dealers are often under pressure from sales quotas, and may wish to streamline negotiations in order to complete deals with as little risk or negotiation as possible. Additionally, it could be that compensation for F&I departments/employees is non-linear with respect to markup (e.g. once certain aggregate levels are achieved, the incentives drop). It is also the case that some lenders limit the available markup for particular deals, based on borrower, dealer, and/or deal characteristics. Lastly, there are numerous potential strategic and behavioral considerations outside the scope of the model (e.g. expectations about repeat customers, reservation rents that lead to lower markups on already profitable deals, etc.). However, these factors are mostly idiosyncratic, and would be hard pressed to explain any persistent patterns of disparity in the market.

willingness to pay, rather than extreme or average values.[11] This result–which will be crucial in the empirical approach I employ–is more formally developed both by Charles and Guryan (2008), and in Appendix A.

Models of statistical discrimination are myriad and varied, but each rely on an information asymmetry to explain disparities. These are typically straightforward signalling models where one decision maker a prior about some important characteristic of an agent, and updates that prior based on some signal of that agent's true value for the characteristic. Disparities arise when the priors vary for different types (e.g. favored and disfavored types), and/or the perceived quality of the signal is different between these types.[12] Often (though not always) this results in a prediction that a positive signal for a member of the "low prior" type will have a larger return compared to that same signal for a "high prior" type. However, a net disparity typically manifests despite the higher return, as the "low prior" type is still penalized for their group's perceived characteristics.

In the case of markup, statistical discrimination could most easily arise if dealers have different priors about the financial sophistication of its Black and White customers. This could be reasonable if a dealer believes that Black and White customers have systematically different levels of experience with the financial system, informed by lower average credit characteristics, credit scores, etc. It seems unlikely dealers would differentially perceive the signals of financial sophistication they

[11] Note that in both cases, the avoidance of extreme values need not be strategic or costless. That is, Black customers would not need to know the identities of the most prejudiced actor and specifically avoid them, any more so than a supplier of milk would need to know the identities of buyers with the lowest willingness to pay. In each case, were a "bad match" to result, there would be no transaction, and the parties would either exit the market or move on to the next available agent with whom to transact. In a market without frictions, this would converge to the described equilibrium instantaneously. One specific and important friction that may affect this prediction, search costs, will be discussed shortly.

[12] Note that the priors need not be accurate, ex ante. Coate and Loury (1993) show that a self-fulfilling prophecy can result if noisy or inaccurate perceptions about the distribution of skills result in under-investment in those skills, which would validate the inaccurate priors in the general equilibrium. Although this is just one potential equilibrium, it highlights that the priors need not be accurate for the model to converge to a "patronizing equilibrium."

observe–specifically FICO and credit characteristics–as these signals come from a common source and and consistent in their calculation.[13] In this case, dealers may perceive a higher value of $E[\bar{m}_i]$ for their Black customers, and seek systematically higher markups accordingly. However, a signal of a high FICO should have a larger effect on the markup for Black customers, as it would force an update to the prior for those customers that is larger in magnitude than that same signal would require for White customers. This result is also formalized in Appendix A.

The final model considered here is a search model with discrimination. One way to conceptualize search models with discrimination is that they formalize the frictions referenced in the Becker model. If a "bad match" occurs, the simplest form of the Becker model suggests the parties would simply move on to the next option without cost or consequence. In a search model, the existence of potential "bad matches" changes the reservation markups and equilibrium markup offers for searchers facing discrimination. The intuition is straightforward: Black customers face some portion of the market that will not offer them viable levels of markup due to high levels of prejudice.[14] As such, their expected returns to searching a random dealer are lower relative to White customers, meaning they would be willing to accept a higher level of markup in order to avoid having to search again. Dealers willing to lend to Black customers are aware of these search costs, and in equilibrium offer higher average markup to Black customers compared to those that are White. The testable implications of this model are twofold. First, higher levels of "extreme" prejudice should result in

[13] Note that this does not imply that credit scores are unbiased, or that their signal value is necessarily consistent across types. It is possible that the information from the "traditional" financial upon which FICO scores are based would differentially measure Black and White customers' creditworthiness. However, as navigating the traditional financial system is often harder for racial and ethnic minorities (due to their over-representation in under-banked communities, racial disparities in access to credit, etc.), the obvious ways in which a credit score could carry different weight for these groups would seem to imply that a high credit score for a Black customer could, if anything, be a stronger signal of financial sophistication.

[14] Prejudice is just one source of the disparity in offers. Search models of discrimination often do not delineate a specific source of the disparity in offer probabilities, and this disparity could arise for extreme information asymmetry, hysteresis, or any other motivation for not lending (at reasonable prices) to a particular group of potential borrowers. However, prejudice is a natural, and commonly assumed, source for this disparity.

higher markup differentials, as the expected benefits of searching a highly prejudiced market are lower for Black customers. Note that this contrasts the Beck prediction in that it focuses on the prevalence of extreme prejudice, rather than the marginal level of prejudice. However, these are not mutually exclusive, as it could be that a significant portion of the market is highly prejudiced, but the marginal level of prejudice is low. The second testable implication is that the proportion of a population that is Black should lead to lower differentials, as the cost to dealers of indulging in such extreme prejudice in a market with more Black customers would be higher (thereby forcing a less discriminatory market through competition). I present a search model with discrimination that more formally develops these implications in Appendix A.

It is an important, though often overlooked, feature of models of economic discrimination that they need not exist separately from one another. That is, it could be that multiple sources of discrimination are present in any given market transaction. For example, a statistical discriminator may weight both the customer's signals of financial sophistication and his/her outside options (derived from a search model) when assigning markup. Alternatively, a taste-based discriminator may indulge his/her preferences differently conditional on the perceived risk that a customer would walk away from the deal, a perception that could be informed by a signal of financial sophistication or outside options. In other words, these models may well work in concert to result in the observed disparities. Given the unique data I am able to utilize, I am able to test many of theses models' implications not simply in isolation, but in conjunction with one another.

4. Data

The data come from two sources, a proprietary data set composed of administrative data from a number of indirect auto lenders compiled from the Consumer Financial Protection Bureau's supervisory examinations, and the General Social Survey (GSS). The administrative data provide detailed, transaction-level information on borrowers, vehicles, financing, and markup. The latter provides information on social attitudes which allows me to construct measures of prejudice. Together, these data are comprehensive enough to enable tests of the specific predictions of economics models of discrimination in a more comprehensive and complete way than has been previously attempted.

The data describing indirect auto loans comes from the Consumer Financial Protection Bureau (Bureau). These are administrative data collected in the course of the Bureau's efforts to evaluate lenders' compliance with federal consumer financial laws, including the Equal Credit Opportunity Act (ECOA).[15] The data comprise millions indirect auto loans booked between 2008 and 2013, and contain all the information used by the institutions for underwriting and assigning the risk-based buy rates (interest rates absent markup) to the loans.[16] These data contain all the income, credit, and similar information that is likely to be available to the dealer.

To calculate the discretionary dealer markup amounts, the buy rates (the minimum interest rates at which the lenders will initiate the loans) are subtracted from the contract rates (the interest rates that are ultimately paid by the borrowers). In addition to the contract and buy rates, these data include measures of the characteristics upon which a lender would underwrite and price a loan (e.g. credit history, down payment, vehicle characteristics, etc.). These characteristics are useful for understanding if there is some correlation between the factors affecting buy rate, and the dealer imposed markup levied on top of that buy rate. Summary statistics for some of the key measures in the auto loan data are presented in Table 1.

Noticeably absent from these data are a borrower's race, ethnicity, and gender. In fact, collecting this information for an auto loan applicant is prohibited by law. However, the data available

[15]The lenders in this sample include both traditional banks and finance companies that specialize in auto loans (captives and non-captives). The specific count, identities, and types of lenders must remain masked for privacy purposes; however the data set encompasses a sizable share of the market for indirect auto loans (the included lenders that typically comprise more than 20% of the indirect market in any given year).

[16]Note that I exclude subvented loans (those subsidized and advertised by specific, often captive, lenders), as these typically disallow markup. I also exclude loans with markup caps below 200 basis points to ensure consistent markup potential. Finally, I exclude loans with FICO scores below 620, as the subprime auto loan market differs substantially in its structure from the prime/near-prime markets. Including these data does not have a substantial effect on the results reported later.

to the CFPB do contain a borrower's name and address, enabling construction of race/ethnicity and gender proxies. In order to estimate each borrower's race, I employ the Bayesian Improved Surname Geocoding proxy (BISG) first developed by Elliott et al. (2008). This approach and its relative performance in the context of indirect auto lending are described in detail by CFPB (2014).

As a summary, I employ the BISG approach by using the 2010 Census list of frequently occurring surnames to establish six mutually exclusive baseline probabilities that a borrower would self-identify as Hispanic, White, Black, Asian and Pacific Islander (API), American Indian/Alaska Native (AIAN), and Multiracial.[17] These probabilities, $s(r|j)$, are simply the proportion of individuals with surname j who self-report being in race/ethnicity group r. I then update these prior probabilities with geocoded race/ethnicity data from the 2010 Census SF1 file. Here I employ a hierarchical matching technique, attempting to match a borrower's address to the smallest Census geographic entity for which a usable match is achieved. For the majority of borrowers this results in a match to a Census block group; for the remaining borrowers matches are achieved to a Census tract or a zip code. For each geographic entity k I then calculate $g(k|r)$, the proportion of all people who self report being race r who reside in the geographic entity. This is simply the probability of a selected geographic entity of residence given an individual's race/ethnicity group.

To construct the proxy I make the standard (and necessary) assumption that, given an individual's race, the probability of residing in the matched geographic entity is independent of his or her surname. This allows for the application of Bayes' Rule to yield a posterior distribution given by

$$p(r|j,k) = \frac{c(r,j,k)}{c(White,j,k) + c(Black,j,k) + c(Hisp.,j,k) + c(API,j,k) + c(AIAN,j,k) + c(Multi.,j,k)}, \tag{4}$$

where $c(r,j,k) = p(r|j) * g(k|r)$. Table 2 displays some key descriptive statistics for these proxies. It is important to note that each individual has a set of proxied race/ethnicity probabilities,

[17] The Census surname data are available via https://www.census.gov/genealogy/www/data/2010surnames/.

and is not assigned to a particular group by the BISG process. As discussed in Elliott et al. (2008) and CFPB (2014), these estimated probabilities act as relatively efficient and accurate proxies of a borrower's race and ethnicity. However, in some instances it might be useful or important to assign the borrower to a particular race/ethnicity, e.g. for analyses typically constrained to only members of certain groups. In these instances I employ two alternative strategies: a maximum a posteriori (MAP) approach that categorizes borrowers into the group for which they have the largest proxy value, and a repeated imputation approach where a race is randomly assigned to each observation based on the proxy probabilities, and the process is repeated to generate a distribution of results.[18] Summary statistics by assigned race for borrowers assigned to the Black and White categories are presented in Table 7.

It should be noted that there have been criticisms of the use of the BISG proxy to estimate race/ethnicity in consumer fiance. In an odd quirk of consumer finance law, lenders are required to attempt to collect and report race/ethnicity information for home mortgages that they originate, but forbidden from collecting this information for other loans. This means that borrowers' demographic characteristics are known for most mortgages, allowing for an evaluation of the performance of the BISG proxy in predicting the race/ethnicity of mortgage borrowers. Employing this strategy, studies have shown that the BISG over-predicts the count of minorities when applied to home mortgage data (e.g. CFPB (2014) and Voicu (2016)). This result is unsurprising. By using the Census the BISG proxy is, in effect, calibrated to the demographics of the overall population. To the extent that home buyers are not a representative sample of households, the BISG should not accurately predict the races/ethnicities of this group. More specifically, if racial/ethnic minorities are relatively underrepresented in the population of home buyers, the BISG would overstate their numbers.

[18] As a robustness check, I also used an alternative "threshold" approach where borrowers were assigned to a single race/ethnicity group if their proxy values exceeded .5 and .8 for a particular category. However this type of assignment leaves a number of borrowers uncategorized and making this a less preferred approach. That said, all results from this alternative classification were qualitatively similar to the preferred specifications.

On the other hand, if the population of auto loan borrowers more closely tracks the overall population, then the concerns about potential bias in the BISG estimates would be at least somewhat allayed. Table 4 presents summaries from GSS questions related to home and auto purchases. Responses indicate that non-White respondents are much less likely to have ever purchased a home compared to the overall population (and especially in the last five years). This is true in both the measures of counts, and the measures of propensity. However, the share of auto purchases from dealers in the last five years (the most relevant group for this paper) looks much more representative of the overall population. Specifically the share of auto buyers who are Black is less than 1.6 percentage points lower than the overall population, and the share of White only differs by 0.2 percentage points. As such, the concerns surrounding bias in the proxy due to a non-representative population of borrowers seem minimal here.[19]

The General Social Survey also provides the information that allows me to proxy for the levels of prejudice preenst in a market. The GSS is a (typically) biennial survey consisting of myriad and varied questions, including many related to demographics and attitudes. Many of these attitudinal questions relate to issues of race, and allow for a comparison of racial sentiment across geography and time. These data form the basis from which indices of racial prejudice are constructed. Following the general approach of Charles and Guryan (2008) I create a "summary indices" of racial prejudice based on the GSS questions that relate to racial attitudes but do not contain strong policy implications.Excluding questions with a strong policy implication is meant to avoid the possibility that an individual may have strong feelings toward government programs, but that those feelings may not indicate any strong feelings regarding race.[20] These indices are the

[19]It should also be noted that the MAP assignment process will likely lead to an under-counting of minorities. For example, individuals in relatively integrated communities with surnames assigned at similar rates by minority and majority populations will be systematically assigned to the majority category. Put differently, all individuals with proxy values indicating a 25% of being Black, and a 75% chance of being White will be assigned to White, even though a quarter of these individuals are expected to be Black. Consistent with this notion is that $Pr(Black|White/Black)$ is approximately 12.6% using the direct probabilities, while it is approximately 8.5% using the MAP assignment.

[20]For example, a response indicating agreement with statement "We are spending too much money improving the

mean of the (normalized) standard deviations away from a baseline average, where the baseline is the average in the year the question first appeared. As such, a one unit increase in the index value can be interpreted as being an average of one standard deviation above the baseline level of racial prejudice across all questions included in that respondent's index.

As in Charles and Guryan (2008) I create two indices. The "full index" encompasses all of the questions that relate to racial prejudice (but do not have direct implications for government policy).[21] The "core index" is constructed using only four key questions used to a create measure of the marginal discriminator's prejudice, percentile measure of prejudice, and the within-regional variation of prejudice over time. For reasons Charles and Guryan (2008) describe in detail, the results from the analyses using the core index are strongly preferred.[22] These "core" questions are:

- racpres: "If your party nominated a Black person for President, would you vote for him if he were qualified for the job?" [yes/no] [23]

- racfew/rachalf/racmost: "Would you yourself have any objection to sending your children to a school where [a few of/half/more than half] the children are [Whites/Blacks]? [yes/no][24]

condition of blacks" may indicate opposition the funding of programs to improve the conditions of Blacks, as opposed to disagreement with the mission of such programs.

[21]Note that one of these 21 variables is a composite of the response to three related questions regarding the level of integration desired in a school which a child will attend.

[22]Charles and Guryan preferred the core index over the full index when constructing percentiles, as the variance structure of the full index makes it a less reliable measure when constructing distributions. This is because the number of questions varies by year, making year to year comparisons of index values inconsistent in their variance. As such, the resulting percentiles may be biased, and any estimates based on those percentiles may also be biased, and will be difficult, if not impossible, to assess for statistical significance. That said, the analyses to be performed here will report results for both the full and core indices, as the results from the full index may still be instructive, even if they are less precise and/ore reliable.

[23]The gender specificity is imported from the original question.

[24]From these three questions I generate a composite indicator that takes the value 1 if the response to any of the questions is "yes."

- racmar : “Do you think there should be laws against marriages between Blacks and Whites?” [yes/no]

- racseg : “White people have a right to keep Blacks out of their neighborhoods if they want to, and Blacks should respect that right” [SD/D/N/A/SA].

Of particular importance to this analysis are the “marginal” and average levels of prejudice in a market. As in Charles and Guryan (2008) I define the “marginal” level of prejudice as the b^{th} percentile of the prejudice index, where b is the percent of that area’s population that is Black. The “average” level is simply the mean of the prejudice index. Table 5 shows the proportion White, proportion Black, average prejudice index, and marginal prejudice index according to the GSS. All results are presented by census region, in ascending order of the marginal level of prejudice measured in the full index to highlight the variation in this key measure.[25]

In combination, the expansive data on discretionary markup in indirect auto loans, race/ethnicity proxies, and the indices of prejudice should allow for a direct and nuanced exploration of models of discrimination.

5. Empirical Approach

The first requirement in assessing the impact of economic discrimination on market disparities is to document that such disparities exist. Table 6 displays the coefficients from simple regressions of markup on the race proxy probabilities. Of note, the first column shows that the estimated overall disparity facing Black customers is more than 21 basis points. The remaining columns of Table 6 show the coefficients by region. As additional evidence of disparity Table 7 shows the means of key measures for Black and White customers as determined by the MAP assignment.[26] Again,

[25]Due to restrictions on the GSS sensitive data files, all prejudice-related measures are described in the paper at the region level, where they are publicly available. However, all estimates rely on state-level measures.

[26]Throughout the MAP analyses, the sample is constrained to only observations assigned to non-Hispanic White or Black.

disparities are apparent, with the average markup experienced by Black customers exceeding that of White customers by 14 basis points. Finally, figure 7 is a histogram showing the distributions of markup faced by Black and White customers as determined by MAP. Of note is that Black customers appear to experience 200 basis points of markup more frequently than do White customers, while white customers experience no markup more frequently than do Black customers. This implies that much of the disparity is likely due to relatively prevalence of extreme markups in these populations. In each case there are economically meaningful and statistically significant disparities facing Black borrowers. Regardless, these simple disparities (in both the frequency and magnitude of markup) imply an analysis of the sources of markup is warranted.

5.1. Exploring Non-Discriminatory Explanations

In the hypothesized model described in Equation 3, disparities in markup may be related to many factors other than prejudice. These include differences in buyers' negotiation skill and financial sophistication, as well as their potential interaction with a buyer's race/ethnicity.[27] Unfortunately, these skills are difficult, if not impossible, to observe directly. However, the richness of the available data allows for the observation of outcomes that are likely directly related to these skills.

For example, negotiation skill should also be related to the relative price at which the buyer is able to settle the deal. In fact, it should be the case that negotiation skill is better illustrated by the vehicle price than the interest rate, as there is little doubt this is a more salient factor to the average

---

[27]It should be noted at this point that the analysis presented here is not a legal analysis of what factors should be considered when estimating disparities. That is, the specifications presented are developed to explore why markup occurs, and the portions of the disparities associated with differences in the various controls are not necessarily justified from a legal standpoint. See CFPB (2015) for a further discussion of the use of controls in legal analyses of markup disparities. Also note that "explaining away" disparity in this way may understate the true impact of prejudice (Spriggs 2020). However, as this is a test of discrimination in this specific market, this exercise seems justified.

buyer. Table 7 shows summary statistics–notably including average markup–by price quartile. These quartiles are calculated controlling for make, model, new/used status, region (to account for geographic differences in demand), and vehicle year. The table shows that, as the relative price for a given deal increases (indicating less negotiation skill on the part of a buyer) that likelihood of being marked up monotonically increases. However, the magnitude of this change is quite small, with the proportion marked up increasing 4 percentage points, from .745 to .786. The average markup amounts do not follow a similar pattern–they decrease from quartile 1 through quartile 3, then increase in quartile four–but are similarly small in the magnitude of their differences. Given this, it appears that a consumer's skill in negotiating for a lower price of the vehicle may be unrelated to markup. While it is true buyers who pay a relatively high price for a given vehicle experience slightly higher markup rates, the amounts of markup paid do not follow a pattern consistent with a hypothesis of negotiation skill. Additionally, the differences in markup experiences by price quartile are small enough that any relationship seems at best economically trivial, and potentially illusory.[28]

That said, a vehicle's price may have ambiguous influence on a buyer's markup. Conditional on a buyer having approximately equal information about a vehicle's "fair" price and a "fair" interest rate, negotiation skill should serve a buyer equally well when settling on a price and interest rate for the vehicle. However, price is a far more salient feature of the deal to most buyers. As such, dealers may look to add additional markup to low price deals, to make up for rents lost during the price negotiation. As such, we may not expect the relationship between price and markup to be monotonic, and the observed differences could be the net effect of competing influences (e.g. lower priced deals have consumers who are better negotiators but dealers who have stronger desires for markup, and vice versa).[29]

---

[28]Regressions of markup on relative price paid for the vehicle yield economically trivial coefficients, and including indicators of price quartile in the subsequent analyses has no meaningful effect on the results.

[29]This could also be true if relatively wealthy buyers are not particularly price sensitive, and expend little effort to negotiate the price of a vehicle down. If these buyers are willing or able to pay cash for the vehicle, dealers may offer low markup–or potentially even marked down–rates, to entice these buyers into financing.

To proxy for a dealer's assessment of the buyer's financial sophistication, I use the buy rate offered by the lender.[30] Since the buy rate is a risk-based rate assigned by a financial institution based on the buyer's full credit report, it can serve as an index of the buyer's overall financial history.[31] However, there may be a competing influence of buy rate on markup beyond financial sophistication. Namely, as the buy rate increases, it may be more difficult for a dealer to add markup to the transaction. This is due to concerns ranging from practical (e.g. a buyer paying an already high rate may be more likely to balk at any increase to that rate), to programmatic (e.g. many lenders place stricter limits on the allowable markup for high buy rate deals), to legal (e.g. many states have usury caps, although these are more often binding on subprime borrowers who are excluded from this analysis). Still, if markup is related to a buyer's financial sophistication, it should exhibit some relationship to the buy rate assigned to a particular deal.

Table 8 shows summary statistics–notably including average markup–by buy rate quartile. The results show that the average markup imposed on buyers is higher as they move from buy rate quartile 1 through buy rate quartile 3. In buy rate quartile 4, however, this trend is reversed, with buyers paying (slightly) less markup than those in quartile 3. The fourth quartile of buy rate shows a somewhat similar anomaly in markup likelihood. Buyers in the first three quartiles experience average markup chances within 1.2 percentage points of one another (though the pattern is not monotonic). However, in the fourth quartile, the likelihood of markup drops by more than 4 percentage points relative to the next lowest quartile's likelihood. It is possible this is a result of

[30] In the more substantive analysis I use both the buy rate and a borrower's FICO score, but felt a single measure was sufficient for the descriptive analysis presented here.

[31] While the buy rate captures much of the information suggestive of financial sophistication, it is still an imperfect proxy. For example, well-educated young buyers may be quite financially savvy, yet have a "thin" credit file leading to higher buy rates, while a long and spotless credit history may be more closely related to a buyer's wealth, rather than his/her sophistication. Still, the buy rate is a simple, consolidated measure of many factors related to participation in financial markets, and therefore a plausible proxy for the buyer's financial sophistication. More importantly, it is directly observed by the dealer in each and every indirect financing transaction, and therefore constitutes a more than reasonable proxy for the dealer's assessment of the buyer's financial sophistication.

markup capacity (as usury rates may bind on some high buy rate borrowers) or tolerance (as "sticker shock" may dissuade some borrowers from taking very high interest loans) being somewhat different in the fourth quartile of buy rate. Regardless, the differences across quartiles here are small, and do not imply a substantial and systematic relationship between buy rate and markup (though this relationship will be explored more in the analysis that follows).

It should be noted that the relationship between buy rate quartiles is not independent of race. For example, Black and Hispanic customers are more likely to be in the third and fourth quartiles of buy rates, indicating that the markup disparities facing them may be related to financial sophistication. This is in contrast to the price quartile results, where there does not appear to be a corresponding association between the quartile and race/ethnicity. To more flexibly account for this, I further assess the relationship between markup and the interaction of buy rate quartiles with race/ethnicity. Specifically, estimate probits for the probability of being marked up and regressions for amount of markup by buy rate quartiles. The results show, generally, that the coefficients on the proxies of minority status trace the pattern of markup; that is, as the average risk/amount of markup increases, so too do the coefficients on minority status. In other words, as the overall risks associated with markup increase, the additional risks faced by minority borrowers increases as well.[32]

5.2. Test for Becker-Type Discrimination

To assess the impact of Becker-type taste-based discrimination I employ a similar identification approach as in Charles and Guryan (2008). Specifically, I rely on the sharp prediction of the Becker model noted previously: that disparities will be related to the marginal level of prejudice in a market, but generally unaffected by the average level of prejudice in that same market.

As noted in Section 4 I define the "marginal" level of prejudice as the b^{th} percentile of the prejudice index, where b is the percent of that state's population that is Black. The "average" level

[32]These results are not reported here, but are available in an appendix, upon request.

is simply the mean of the prejudice index in that state. In effect, this approach exploits geographic variation in prejudice levels (summarized in Table 5) to explain racial disparities. This obviously requires there to be some geographic variation in markup disparities by race. To motivate this approach, Table 6 presents the results from region-specific regressions of markup on the various race, ethnicity, and gender proxies.[33] The coefficients for Pr(Back) range from -.122 in the Pacific region, to .399 in the East South Central region, exhibiting a span of more than 52 basis points. Additionally, the coefficients are each statistically significant at the 5 percent level or better, with the exceptions of the coefficient in the Mountain region. This finding implies that there is indeed substantial, region-specific variation in the average markup disparities between Black and White borrowers, and that these differences are statistically and economically relevant.

To assess the impact of the marginal and average levels of prejudice on markup disparities I estimate the following regression

$$m_i = \alpha + \beta C_i + \gamma S_j + \eta I_{i,j}, \tag{5}$$

where C_i is a vector of individual-specific characteristics (including the demographic proxies), S_j is a vector of state-specific characteristics (including the marginal and average prejudice levels), and $I_{i,j}$ is a vector of interactions between the proxied probability of being Black and state characteristics. The Becker prediction is that the interaction coefficients will show a relatively large and significant relationship between the $pr_blackxmarginal$ interaction and markup, and a comparatively trivial relationship between $pr_blackxaverage$ and markup.

As mentioned previously, Lang and Lehmann (2012) note that in the context of the Charles and Guryan paper this approach assumed “that all firms are the same size, that black and white workers are perfectly segregated, that there is no consumer prejudice (or at least that the market can avoid it), that the distribution of prejudice is the same among employers as among the population as a whole and that the labor force participation rates of black and white workers are the

[33]The more substantive analysis to follow relies on state-level variation. However, the state-level results rely on the GSS sensitive data files, and cannot be reported here. Instead, I present the region-level data that rely only on the public-use GSS data.

same," assumptions which they deemed "unreasonably strong" (Lang and Lehmann 2012, p.972). However, these assumptions seem more reasonable here. While auto dealers certainly vary in size, ultimately the indirect lending transaction is atomistic and therefore idiosyncratic enough that market power issues should not be of much concern (as interactions between a borrower and an F&I agent results in a stand alone transaction regardless of employer size, rather than placement into a more commoditized job at a large employer). There is no opportunity consumer prejudice to arise in this market, as the lenders who effectively play the rol eof consumer (by "purchasing" the loan from the dealer) do not know the race or ethnicity of the borrower. Finally, specific selection of prejudiced/unprejudiced types into working at an auto dealership seems unlikely, as more than 1.3 million people are employed by automotive and motor vehicle dealerships (F&I agents are often selected and promoted from within auto dealerships, e.g. former sales people, operations associates, business managers, etc.), and this employment is spread fairly proportionally across geography.[34] As such, there is no reason to believe the distribution of prejudice within auto dealerships is different from that of the general population in a particular geography.[35] So, while these points are still important to consider as potential threats to identification, they appear to be more reasonable here than in the context of employment.

Lastly, it should be noted this is a "direct" test of the model, in that is specifically assesses the impact of the model's assumed source of disparity on outcomes. This approach should provide a more compelling test of the model compared to an outcome test, which assess the data's consistency with a models predictions, rather than linking the evidence directly to the model's inputs.

---

[34]Employment numbers from BLS as of March 2018.

[35]Another important difference between F&I managers and employers when considering the potential impact of prejudice is that hiring and compensation practices are generally subject to EEOC oversight and hence HR managers often have internal policies intended to identify and avoid potential disparities. By contrast, F&I transactions are subject to comparatively less oversight at the dealers level (e.g. the CFPB's examinations of markup disparities focused on lenders rather than dealers), so one might expect correspondingly less proactive attention to be paid to identifying and avoiding potential disparities.

5.3. Test for Statistical Discrimination

To assess whether asymmetric returns to information explain some of the observed disparity, I test the relationship between signals of a borrower's financial sophistication (FICO score and buy rate) and his or her markup. Importantly, it is unlikely there are signals of a buyer's true financial sophistication available to the dealer that are not captured in the data. As dealers do not have to interact with buyers after the transaction, "soft skills" and personality traits should be less of a consideration relative to labor markets. Additionally, any signals of sophistication not recorded in the data (e.g. showing an awareness of the lending process or typical market rates) are highly likely to be correlated with the observed signals (e.g. FICO), as savvy borrowers are likely to have worked to improve these signals in the lead up to the transaction. As such, this approach should provide a more direct test of the economic model, as it links outcomes to signals we know to be observed, rather than assessing dynamic learning patterns consistent with the predictions of statistical models (as in Dobbie et al. (2020), Altonji and Pierret (2001)).

It should be noted that few borrowers are aware of the practice of dealer markup, and even fewer (perhaps none) observe such markup. As such, financial sophistication likely plays only a small role in directly determining markup, though the signal of sophistication may well affect a dealer's markup offer.[36] Since markup is purely discretionary and not a function of credit risk, neither FICO score nor buy rate should directly impact the markup in any systematic way. However, dealers may assign markup in proportion to their perceptions of a buyer's financial sophistication (which the dealer could reasonably perceive as being related to a buyer's tolerance for markup), and may use FICO scores or buy rates as a signal of a borrower's financial sophistication.[37] As

[36]The most obvious way in which financial sophistication may affect markup is that savvy customers may be more likely to obtain a financing offer directly from their bank prior to negotiating with a dealer. However, this may actually have an ambiguous effect on final markup. As indirect offers are often lower than direct offers (due in part to the dealer's ability to shop, volume of business, etc.), a direct offer may inadvertently serve as a floor on the rate a customer can negotiate, and may actually result in greater markup than a dealer would have assigned to a customer perceived as being financially sophisticated who did not present a competing offer.

[37]It should be noted that, in practice, neither FICO nor buy rate is not a particularly comprehensive measure

such, these measures should be exactly the type of signal assumed by statistical discrimination models. In a model of statistical discrimination as described in Equation A7, dealers would apply less weight to a Black borrower's individual signal of sophistication, while weighting more heavily the disfavored group's lower average characteristics (e.g. lower average FICOs, higher average buy rates, as shown in Tables 1 and 7).

While there are a number of ways one could test for statistical discrimination in this market, I employ a straightforward approach that credibly allows for decomposition of estimated disparity into taste-based and statistical discrimination components. To achieve this, I simply regress each borrower's markup on race, the marginal and average levels of prejudice, the signal of financial sophistication, and the non-race factors interacted with race. If the financial sophistication signals do impact markup, the expectation is that higher FICOs would be associated with lower markups, and lower buy rates would be associated with lower markups (i.e. the coefficient on FICO would be negative, and the coefficient on buy rate would be positive). Throughout the analysis I use a dummy variable equal to one if the borrower's FICO score is 720 or above to denote a "high FICO" signal.[38] If there are differential returns to financial sophistication that disfavor Black borrowers, then the coefficient on the Black x high FICO interaction term should be positive and significant (both statistically and economically), indicating that better signals result in relatively higher markups for Black customers. Similarly, the coefficient on Black x buy rate should be negative, indicating better (lower) buy rate signals are associated with relatively higher markups for Black customers. In both cases, if these asymmetric returns supplant the Becker-type discrimination implied by the previous findings, then the coefficients on the interaction between Black and marginal level of

---

of financial sophistication. However, all that matters for it to be a credible signal of sophistication in a statistical discrimination model and/or in a time-constrained negotiation is that it be observable to the dealer at a relatively low cost, and that the dealer believes it is correlated with a buyer's financial sophistication. I believe these conditions are likely met by FICO and buy rate.

[38] 720 is used here as it is often the highest FICO delineation present on the rate sheets that determine a borrower's buy rate. In alternative specifications I use different thresholds, as well as FICO as a linear term. All results are qualitatively consistent with those presented here, though specific point estimates and significance levels do vary.

prejudice should be attenuated.

5.4. Test for Search with Discrimination

To test the implications of search models of discrimination I evaluate the relationships between the share of Black searchers, prejudiced firms, and unprejudiced firms with the observed disparities in markup. I use a similar approach as I did with the Becker model, using the share of "very prejudiced" (or "very unprejudiced") respondents in a state to proxy for the share of "very prejudiced" (or "very unprejudiced") firms.[39] Specifically, I use the share of respondents within a state whose prejudiced index values are two or more standard deviations away from the national average to denote the share of the market that is prejudiced/unprejudiced in a state.[40] The search models in Section 3 show that the markup gap should decrease as the share of very prejudiced firms in a state decreases (i.e. $\frac{d(m^w - m^b)}{dtheta_p} > 0$), as the share of very unprejudiced firms increases (i.e. $\frac{d(m^w - m^b)}{d\theta_n} < 0$), and/or as the share of a state's population that is Black increases ($\frac{d(m^w - m^b)}{dBlack} < 0$). To test these relationships I run simple regressions with these values as covariates, to see if they behave in a manner consistent with a search model's predictions.

Note the implication that the share of prejudiced/unprejudiced firms will impact disparities bears some similarity to the implication of the Becker model. In both cases, the models show that the average level of prejudice should not have a direct impact on the disparity. In search, the share of "very prejudiced" agents should affect the market outcome, while in the Becker model it is the prejudice level at a specific point in the distribution that matters. In effect this allows this test of

[39] It is important to keep in mind that search models are not truly independent models of discrimination. Rather, they are models of propagation and/or amplification. That is, these models assume some type of disparity in opportunities, then show these disparities affect the behavior of unprejudiced actors, and describe how frictions impede competition from clearing the disparities from the market (Black 1995; Lanning 2014). These disparities can come from prejudice levels–as assumed here–information asymmetries, or other sources. What matters most is that some portion of the market is inaccessible to one group of potential searchers.

[40] Alternative approaches using one or three standard deviations yield similar results.

the search model's implications to serve as a robustness check on, or refinement of, the tests of the Becker model. That is, if the share of prejudiced firms does not seem to impact the gap while the marginal level of prejudice does, that is further evidence that the specific point in the distribution of prejudice identified by the Becker model is what is affecting the gap (rather than the average or extreme levels, or the share of extreme values).

6. Results and Discussion

6.1. Becker Test Results

Column I of Table 9 shows the relevant results for the direct test of the Becker predictions using the core index of prejudice, and using the proxy probabilities directly in the regression.[41] The regression estimates a large coefficient on the marginal level of prejudice of 98 basis points, which is significant at the 5% level despite the limitation to 51 state-level clusters when estimating standard errors. Putting these numbers in a more practical scale, moving from the least prejudiced census region (at the margin) to the most prejudiced census region would lead to an expected increase in markup for Black borrowers of 71 basis points, or approximately 63 percent of the average amount of markup charged to borrowers. The coefficient for the average level of prejudice is much smaller at 12.3 basis points, and is statistically insignificant at any commonly accepted level. Putting this into practical terms, moving from the least prejudiced region (measured at the mean) to the most prejudice region would result in an expected change to markup of 4 basis points, or less than 3.6 percent of the average markup.

These results are strongly consistent with the Becker theory of taste-based discrimination. The large coefficient on the Black x marginal interaction is a sharp prediction of the model, as is the small coefficient on the Black x average interaction. That the former coefficient is statistically significant and the latter is statistically indistinguishable from zero is also generally consistent with

[41] All regressions also include controls for the term of the loan.

the model's prediction (as it implies not just the magnitude, but the strength of the association between outcomes and marginal levels of prejudice is stronger than the association between outcomes and average levels of prejudice).[42]

Since most analyses of racial disparities in economic markets focus on one "test" and one "control" group, I employ two strategies to constrain the analysis to only Black and White borrowers. The first is an application of a maximum a posteriori (MAP) assignment. This simply assigns each borrower into a single, mutually exclusive race/ethnicity category based on the maximum estimated proxy.[43] While MAP assignment results in distinct "test" and "control' group assignments, this approach may increase measurement error due to the "false certainty" of the classifications.[44] I also employ a multiple imputation approach where each borrower's race is assigned by drawing a random value from a uniform distribution and using that value along with the proxy probabilities to determine the borrower's race, then estimating the markup regression. This process is repeated 1,000 times, and the distribution of coefficients is captured.

Results from the MAP estimation of Equation A1 are presented in column IV of Table 9. The regression again estimates a substantial coefficient on the marginal level of prejudice of 49 basis points, which is again significant at the 5% level despite the limitation to 51 state-level clusters when estimating standard errors. Putting these numbers in a more practical scale, moving from the least prejudiced census region (at the margin) to the most prejudiced census region would lead to an expected increase in markup for Black borrowers of around 36 basis points, or approximately 32% percent of the average amount of markup charged to borrowers. The coefficient for the average level of prejudice is once again smaller and statistically indistinguishable from zero. Even taken at face

[42]Given that the data must be clustered at the state level, I am not overly concerned with the specific statistical significance of the estimates. While the data are extensive, prejudice is only identified at the state level, effectively limiting the analysis to 51 observations. As such, statistical significance is relatively difficult to achieve. However, the p-values remain instructive, especially when they are particularly large or small.

[43]In the unlikely event there is no distinct maximum proxy value, the observation is dropped from the analysis.

[44]For more discussion of this approach, and specifically its performance in a predicting the race/ethnicity of moertgage borrowers, see Voicu (2016).

value the coefficient of 8.8 basis points implies moving from the least prejudiced region (measured at the mean) to the most prejudice region would result in an expected change to markup of less than 3 basis points, or 2.6% percent of the average markup.

The means and 95% confidence intervals of the coefficients from the multiple imputation simulations are displayed in Table 10. They show the mean of the estimated coefficients is 56 basis points for the interaction between Black and the marginal level of prejudice, and 6 basis points for the interaction between Black and the average level of prejudice. As with the previous results, these findings are strongly consistent with the predictions of taste-based models of discrimination. Importantly, these coefficients are also tightly distributed, with 95% confidence interval ranges of less than 7 basis points for the marginal interaction, and 3.5 basis points for the average interaction. This implies that the estimates are not particularly sensitive to random variation in the assignment of race based on the BISG proxy.

In summary, in each specification there is evidence highly consistent with key predictions of the Becker model of discrimination. These results suggest that prejudice may well be responsible for the observed disparity in markup facing Black customers relative to their White counterparts. That the findings are so consistent in direction and relative magnitude suggests this finding is quite robust. However, the analysis thus far considers only one potentially explanatory model; the subsequent analyses will nest the test of Becker-type discrimination while considering alternative explanations as well.

6.2. Statistical Discrimination Test Results

Column II in Table 9 shows the results when markup is regressed on the interaction between Black and key percentiles of the prejudice index as well as an indicator for high FICO. The coefficient on high FICO is negative and significant, indicating that higher FICO customers indeed experience less average markup. This is generally consistent with the notion that FICO may signal financial sophistication to a dealer, who in turn would assign lower markup. However, the coefficient on

Black x FICO is negative and significant. This implies that as borrowers signal greater financial sophistication, the markup gap between Black and White customers decreases. This is in stark contrast to the prediction of the standard model of statistical discrimination presented in Section 3. Additionally, the coefficient is small enough (-4.1 basis points) so as to imply almost no impact on the markup gap. As such, these findings seem highly inconsistent with a hypothesis of statistical discrimination.

Additional suggestive evidence against the existence of this type of statistical discrimination can be found in the residuals from a simple regression of markup on FICO. The model in Section 3 suggests that markups for Black customers will have a tighter distribution conditional on signal than will markups for White customers. This implies that the residuals from a regression of markup on FICO would have more weight in the center of the distribution for White customers than for Black customers. Figure 7 displays the residuals from a regression of markup on FICO, and shows just the opposite. Specifically, there is greater relative weight in the center of the distribution of residuals for White customers relative to Black customers (as determined by the MAP assignment). Again, this is evidence that asymmetric returns to FICO as a siganl of financial sophistication do not appear to be a meaningful source of discrimination in this market.

Column III in Table 9 shows results from a regression using buy rate in place of FICO as the signal of financial sophistication. The coefficient on buy rate is negative and significant, indicating that customers with lower buy rates experience less average markup. This is inconsistent with the notion that buy rate may signal financial sophistication to a dealer, as the lower buy rate indicates greater sophistication, but also higher markup. However, the coefficient on Black x buy rate is also negative and significant. This implies that as borrowers signal greater financial sophistication (through a lower buy rate), the markup gap between Black and White customers does increase. This is generally consistent with the prediction of the standard model of statistical discrimination presented in Section 3. However, the coefficient is quite small. As a sense of the magnitude of the coefficient, a one standard deviation decrease in the buy rate would result in an increase in the markup gap of less than 4 basis points. As such, these findings offer little support for a hypothesis of statistical discrimination.[45]

Columns V and VI of Table 9 show the results from similar signaling specifications using the MAP assignment approach. In each case the results are similar in both direction and magnitude to those using the race probability proxy directly. If anything these results show even less support for a hypothesis of statistical discrimination.

Of particular interest in each of these specifications is that the coefficients on Black x marginal prejudice remain relatively large, while the coefficients on Black x average prejudice remain comparatively small. This is again consistent with the sharp predictions of the Becker model.[46] In effect, the specifications in columns II, III, V, and VI of Table 9 nest the hypotheses of taste-based and statistical discrimination models, and find much stronger evidence for the former.

Table 10 presents additional evidence to this effect. The "signaling specification" results display the mean, 5th, and 95th percentile coefficients for prejudice interactions and high FICO indicator from regressions based on 1,000 samples of the data using random assignment of race based on the proxy values. The results here tell a similar story: there is a strong evidence of taste-based discrimination (as seen in the strong relationship between the markup gap and marginal prejudice levels and much weaker relationship with the average prejudice level), and a relationship with FICO that is inconsistent with the predictions of a statistical discrimination model. Compared to

---

[45]There are significant concerns about assuming a linear relationship between buy rate and markup. For example, very low buy rate customers may have greater tolerance for markup, as they may still achieve lower rates than they would in the direct auto lending market. Similarly, very high buy rate customers may have limited capacity for markup before running into usury caps or rates that would provide enough "sticker shock" so as to discourage them from originating a loan. To account for this, I ran alternative specification using buy rate as a quartic, substituting buy rate quartile for "raw" buy rate, and using a dummy indicator for "low buy rate." Each of these resulted in qualitatively similar results to those presented here.

[46]While the significance of the estimates do not quite achieve the 10% level, this is in large part due to the limited number of clusters in the data(51). Still, the significance levels are non-trivial given the limited number of effective observations and, coupled with the economically meaningful magnitudes, imply some impact of taste-based discrimination.

the MAP results, the Black x marginal prejudice coefficient is slightly larger, the Black x average prejudice coefficient is slightly smaller, the FICO $\geq$ 720 coefficient is very similar, and the Black x FICO $\geq$ 720 coefficient is a larger negative value. If anything, these results are more consistent with the Becker predictions, and less consistent with statistical discrimination predictions than their MAP counterparts.

In short: the findings in each specification of are quite stable and consistent in their implications that Black/White markup disparities appear to be heavily influenced by the marginal level of prejudice, and much less so by differing returns to signals of financial sophistication.

6.3. Search With Discrimination Test Results

Table 11 shows the results of regressions of markup on the key components of the search model with discrimination described above. Column I displays the results of test of the specification where all unprejudiced firms set type-dependent markups, meaning the model does not depend on the share of highly unprejudiced firms. The results show a large, significant coefficient on the proxy for Black interacted with the share Black in a state. This is the opposite prediction of the model described in Section 3. That is, as the share of a state's population that is Black increases, the disparity in relative markup should decrease. As such, the large and significant positive coefficient is inconsistent with the predictions of the search model.[47] Additionally, the coefficient on the proxy for Black interacted with the share prejudiced in a state is also inconsistent with such a model. Here the model predicts the coefficient should be positive and significant, indicating that the higher the share of prejudiced firms in the market, the higher the markup on Black customers relative to their non-Hispanic White counterparts. However, the coefficient in Column I is negative, large in

[47]While the coefficient is large, it is important to keep the scale in mind; moving from the state with the highest share Black to the state with the lowest share Black would result in an estimated decrease in markup of 67 basis point for a Black customer.

magnitude, and statistically significant.[48]

Not only are these finding in stark contrast to the predictions of the model, they prove to be robust to including highly unprejudiced firms. Column II of Table 11 shows the results of a specification where the share of highly unprejudiced respondents is added to the specification in Column I. Here the coefficient on the proxy for Black interacted with the share Black remains positive, and the coefficient on the race proxy interacted with prejudiced becomes even more negative. It should be noted that the coefficient on the interaction between the proxy for Black and the share unprejudiced does have the expected sign and a substantial magnitude (moving from the region with the lowest share unprejudiced to the region with the highest share is associated with a 37 basis point decline in relative markup on Black customers). However, this finding does little to overcome the inconsistencies with a search model presented by the other key coefficients.

It is perhaps not surprising that search disparities fail to explain the observed differences in outcomes. Heterogeneous outcomes require heterogeneous search costs, or inelastic demand and seller cost heterogeneity (Argyle et al. 2020b). But while the costs of search may well differ for Black and White customers, it is likely their effective search costs are quite similar, in that both appear to be high enough to curtail any real search behavior. Butler et al. (2021) show suggestive evidence that search over loans is relatively uncommon, and the multi-stage process of auto bargaining and the intermediation of the loan process could discourage search once the financing stage has been reached. If search costs are high enough that neither group regularly crosses the extensive margin, then the potential differences in search costs would not result in disparate outcomes.

Regardless, adding the marginal and average levels of the prejudice index to the the specifications still proves instructive, as they once again imply that Becker-type discrimination is consistent with the observed disparities. Columns III and IV show the results of the regression in Columns I and II when the interactions of the race proxy with the marginal and average indices are included.

[48] Again, the magnitude of the coefficient may be initially misleading; scaling the coefficient for the relevant range of values shows that a Black customer moving from the region with the highest share of prejudiced respondents to the region with the lowest share would expect an increase in markup of nearly 40 basis points.

Consistent with the model's predictions, the coefficient on the race proxy interacted with the share Black does adopt the expected sign in both specification III and IV. However, in each case the coefficients on the proxy for Black interacted with the share prejudiced in a state remains large and negative. This is, once again, highly inconsistent with the search model's predictions. Additionally, the sign on the race proxy interacted with the share unprejudiced is positive and large in specification IV, which is also inconsistent with the model's predictions. Finally, the coefficient on the marginal index interaction is positive and large, and the coefficient on the average index interaction is comparatively smaller. While this is somewhat consistent with the Becker predictions, the magnitudes of the coefficients are quite large (probably unreasonably so), and the coefficient on the average index is certainly not trivial, as the Becker model would predict. However, it should be noted that there are a number of econometric issues with specifications III and IV that should give pause before putting much weight on their findings. For example, the share of a state's population that is Black is directly included in the specification, and also influences the marginal index. Additionally, the shares prejudiced/unprejudiced are necessarily related to the marginal and average indices. As such the regressions in III and IV should be seen as a follow up to those in I and II. That is, I and II provide strong evidence that the empirical reality is inconsistent with the search model's predictions, while III and IV offer suggestive evidence that, even when controlling for the factors of import to the search model, the Becker predictions seem more consistent with the evidence.

7. Conclusion

This paper empirically tested the theoretical roots of observed racial disparities in a market uniquely suited to doing so. By using a newly available administrative data set on indirect auto loans and focusing on a purely discretionary component of the interest rate, I was able to compare the predictions of classical taste-based, statistical, and search models of discrimination. These data are expansive, and the market for indirect auto loans is both an ideal one in which to explore discrimination, yet largely unexplored to this point. However, the analysis I am able to conduct is not without its caveats. Most notably, the actual race of the borrowers in the data are unknown.

In place of using the borrower's known race, I employed three strategies: I estimated Bayesian-improved surname and geography (BISG) race proxies for each observation and used these directly; I employed a maximum a posteriori (MAP) assignment that attributes a race to each observation based on the maximum of that observation's BISG proxies; and I conducted a repeated imputation simulation where races are assigned at random based on the BISG, and the analysis is repeated to generate an empirical distribution of estimates. Each of these approaches yield similar qualitative results, though the specific magnitudes of the estimates relying on the different approaches vary.

Caveats aside, the evidence appears strongly consistent with the predictions of a Becker-type model of taste-based discrimination. Specifically, the sharp predictions that racial disparities in outcomes are closely associated with the marginal, but not the average, level of prejudice in a market are repeatedly borne out in the data. This is true for every specification, regardless of how race is proxied for. Even in specifications where I nest models of statistical discrimination or search as competing hypotheses, the results prove consistent with the Becker predictions. Moreover, the data show substantial and significant inconsistencies with the predictions of the statistical and search models of discrimination. As such, this paper both provides strong evidence that Becker-type discrimination drives a substantial portion of the gap in markup added to the indirect auto loans made to Black customers, and evidence that competing models of economic discrimination do not.

Table 1. Summary statistics for loan characteristics in the supervisory auto data.

		Region								
	Overall	ESC	SATL	WSC	MATL	ENC	WNC	PAC	NE	MOUNT
Buy Rate	4.66	5.14	4.66	4.75	4.21	4.83	4.10	5.05	5.05	4.96
	(3.61)	(3.74)	(3.77)	(3.79)	(3.19)	(3.12)	(3.02)	(4.35)	(4.35)	(4.01)
Markup	1.12	1.10	1.18	1.23	1.18	0.92	1.20	1.07	1.07	1.19
	(0.87)	(0.91)	(0.84)	(0.86)	(0.87)	(0.91)	(0.79)	(0.86)	(0.86)	(0.83)
Marked Up?	0.73	0.69	0.76	0.76	0.75	0.63	0.80	0.73	0.73	0.79
	(0.44)	(0.46)	(0.42)	(0.42)	(0.43)	(0.48)	(0.40)	(0.45)	(0.45)	(0.41)
Loan Amount	22,670	22,070	22,996	26,000	22,388	20,368	21,650	22,776	22,776	23,445
	11,384	10,195	11,390	12,529	11,478	10,100	9,982	12,315	12,315	11,555
FICO	743	736	738	737	753	749	755	736	736	736
	(66.22)	(64.56)	(65.50)	(65.00)	(66.50)	(66.89)	(63.75)	(66.56)	(66.56)	(66.53)
Loan to Value	45.64	46.77	49.43	56.48	40.39	40.52	43.20	40.66	40.66	44.67
	(49.65)	(51.40)	(49.87)	(50.59)	(46.49)	(47.81)	(52.29)	(48.90)	(48.90)	(49.96)
Used?	0.49	0.45	0.52	0.47	0.46	0.38	0.54	0.59	0.59	0.56
	(0.50)	(0.50)	(0.50)	(0.50)	(0.50)	(0.49)	(0.50)	(0.49)	(0.49)	(0.50)
N=	>7.5M	>500k	>1.5M	>1M	>750k	>1.25M	>500k	>750k	>300k	>500k

Note. — Standard deviations in parentheses. Census regions presented in ascending order of prejudice index at the percentile corresponding to the percent Black in the region. Counts are approximated to mitigate re-identification risk.

Table 2. Summary statistics for the full set of race/ethnicity/gender proxies calculated from supervisory auto data.

		Region								
	Overall	ESC	SATL	WSC	MATL	ENC	WNC	PAC	NE	MOUNT
Pr(White)	0.736	0.775	0.682	0.621	0.794	0.838	0.885	0.558	0.558	0.740
	(0.318)	(0.269)	(0.317)	(0.352)	(0.305)	(0.259)	(0.198)	(0.362)	(0.362)	(0.308)
Pr(Black)	0.106	0.176	0.192	0.126	0.084	0.080	0.047	0.059	0.059	0.039
	(0.204)	(0.253)	(0.255)	(0.210)	(0.193)	(0.188)	(0.122)	(0.129)	(0.129)	(0.090)
Pr(Hispanic)	0.088	0.015	0.072	0.179	0.061	0.037	0.019	0.211	0.211	0.130
	(0.232)	(0.074)	(0.209)	(0.328)	(0.189)	(0.146)	(0.093)	(0.327)	(0.327)	(0.265)
Pr(API)	0.039	0.012	0.030	0.037	0.044	0.024	0.020	0.118	0.118	0.031
	(0.141)	(0.070)	(0.122)	(0.139)	(0.153)	(0.109)	(0.099)	(0.236)	(0.236)	(0.110)
Pr(AIAN)	0.011	0.009	0.007	0.020	0.003	0.006	0.013	0.014	0.014	0.037
	(0.053)	(0.040)	(0.037)	(0.068)	(0.021)	(0.028)	(0.058)	(0.048)	(0.048)	(0.125)
Pr(Multi./Other)	0.019	0.013	0.017	0.019	0.015	0.015	0.015	0.040	0.040	0.023
	(0.030)	(0.019)	(0.025)	(0.027)	(0.026)	(0.022)	(0.022)	(0.051)	(0.051)	(0.028)
Pr(Female (only))	0.284	0.279	0.302	0.275	0.293	0.283	0.249	0.286	0.286	0.265
	(0.436)	(0.431)	(0.443)	(0.428)	(0.443)	(0.436)	(0.415)	(0.436)	(0.436)	(0.425)
N=	>7.5M	>500k	>1.5M	>1M	>750k	>1.25M	>500k	>750k	>300k	>500k

Note. — Standard deviations in parentheses. Census regions presented in ascending order of prejudice index at the percentile corresponding to the percent Black in the region. Counts are approximated to mitigate re-identification risk.

Table 3. Selected summary statistics from superviosry auto data using MAP assignment, by census region.

	Overall	ESC	SATL	WSC	MATL	ENC	WNC	PAC	NE	MOUNT
Black										
Marked Up?	0.757	0.777	0.768	0.753	0.808	0.700	0.791	0.668	0.791	0.749
	(0.429)	(0.416)	(0.422)	(0.431)	(0.394)	(0.458)	(0.407)	(0.471)	(0.406)	(0.433)
Markup Amt.	1.24	1.32	1.23	1.25	1.39	1.15	1.26	1.02	1.33	1.18
	(0.867)	(0.861)	(0.844)	(0.875)	(0.837)	(0.923)	(0.815)	(0.889)	(0.843)	(0.867)
FICO	717.3	718.9	719.2	717.2	714.5	709.9	722.2	715.8	721.5	715.3
	(65.5)	(64.9)	(65.1)	(65.2)	(65.7)	(66.9)	(67.0)	(65.9)	(67.0)	(64.7)
Buy Rate	6.27	6.42	5.86	6.35	6.43	6.87	6.19	7.22	5.79	6.76
	(5.01)	(5.09)	(4.81)	(5.06)	(4.94)	(4.86)	(5.05)	(6.23)	(4.39)	(5.50)
White										
Marked Up?	0.727	0.678	0.759	0.764	0.736	0.617	0.801	0.738	0.757	0.787
	(0.446)	(0.467)	(0.427)	(0.425)	(0.441)	(0.486)	(0.399)	(0.440)	(0.429)	(0.410)
Markup Amt.	1.10	1.06	1.16	1.21	1.14	0.89	1.19	1.07	1.16	1.18
	(0.865)	(0.912)	(0.841)	(0.855)	(0.872)	(0.900)	(0.787)	(0.844)	(0.848)	(0.821)
FICO	748.1	738.7	743.2	743.1	757.8	753.0	756.8	742.1	758.5	739.8
	(65.7)	(64.1)	(65.1)	(64.7)	(65.5)	(66.1)	(63.4)	(66.7)	(65.1)	(66.6)
Buy Rate	4.41	4.94	4.37	4.39	3.96	4.70	4.01	4.67	3.75	4.71
	(3.29)	(3.43)	(3.40)	(3.46)	(2.82)	(2.91)	(2.90)	(4.05)	(2.74)	(3.80)

Note. — Black and White are estimated from maximum a posteriori (MAP) assignment (i.e. borrower is assigned to the race/ethnicity category corresponding to that borrower's largest estimated race/ethnicity proxy). This technique likely under-counts Black customers in the data, e.g. surnames assigned at similar rates by minority and majority populations in a given area will systematically lead to greater-than-representative rates of assignments to the majority category. Marked up shows the proportion of borrowers marked up. Markup is the average amount of markup in interest rate points(unconditional on whether a borrower is marked up). FICO is the credit score used for pricing the application, and buy rate is the risk-based interest rate determined by the lender based on the borrower's application.

Table 4. Responses to GSS Questions About Home and Auto Purchases by Race.

	White	Black	Other
Share of respondents who have made various purchases who are White/Black/Other:			
Overall percent of respondents	80.6%	14.1%	5.3%
% of respondents who ever purchased a home	88.5%	8.6%	2.9%
% of respondents who purchased home in last 5 years	90.5%	7.0%	2.5%
% of respondents who purchased a car from a dealership in last 5 years	81.7%	12.5%	5.8%
% of resopndents who purchased any used car in last 5 years	80.8%	12.4%	6.8%
Share within respondents who are White/Black/Other that have made such purchases:			
Overall percent of respondents	100.0%	100.0%	100.0%
% of respondents in group who ever purchased a home	69.5%	41.5%	35.1%
% of respondents in group who purchased home in last 5 years	31.4%	14.9%	13.5%
% of respondents in group who purchased a car from a dealership in last 5 years	55.8%	47.6%	59.0%
% of resopndents in group who purchased any used car in last 5 years	34.4%	29.0%	42.5%

Note. — Share of respondents making various purchases show the relative demographics of all respondents, those who purchased homes, and those who purchased cars. Share within demographic groups show the propensity to make various purchases. Both indicate that a sample of car purchasers is more likely to be similar to the overall population than would be a sample of home purchasers.

Table 5. Summary statistics from General Social Survey data.

		Region								
	Overall	ESC	SATL	WSC	MATL	ENC	WNC	PAC	NE	MOUNT
Proportion Black	0.139	0.210	0.222	0.183	0.152	0.125	0.083	0.071	0.049	0.024
	0.346	0.407	0.416	0.387	0.359	0.331	0.276	0.256	0.217	0.152
Proportion White	0.812	0.782	0.742	0.749	0.797	0.851	0.893	0.817	0.919	0.896
	0.390	0.413	0.438	0.434	0.402	0.356	0.309	0.387	0.272	0.305
Mean Distaste Index	-0.214	-0.021	-0.130	-0.146	-0.263	-0.203	-0.223	-0.336	-0.349	-0.326
	0.628	0.691	0.683	0.656	0.604	0.605	0.597	0.562	0.562	0.567
Distaste Index at Bth Percentile	-0.858	-0.596	-0.672	-0.700	-0.858	-0.861	-1.025	-1.031	-1.159	-1.321
N=	53,608	3,545	10,202	4,996	7,969	9,956	4,009	7,170	2,563	3,198

Note. — Standard deviations in parentheses. Census regions presented in ascending order of prejudice index at the b^{th} percentile, where b is the percent Black in the region.

Table 6. Simple markup regressions by census region.

	Overall	ESC	SATL	WSC	MATL	ENC	WNC	PAC	NE	MOUNT
Black	0.213	0.399	0.103	0.064	0.304	0.326	0.055	-0.122	0.207	-0.006
	(0.002)	(0.005)	(0.003)	(0.004)	(0.005)	(0.004)	(0.009)	(0.008)	(0.012)	(0.014)
Hispanic	0.183	0.396	0.188	0.082	0.441	0.493	0.195	-0.006	0.283	0.105
	(0.001)	(0.016)	(0.003)	(0.003)	(0.004)	(0.005)	(0.011)	(0.003)	(0.010)	(0.005)
API	0.105	0.124	0.088	-0.003	0.263	0.244	0.097	0.039	-0.014	-0.017
	(0.002)	(0.019)	(0.006)	(0.006)	(0.006)	(0.007)	(0.011)	(0.004)	(0.012)	(0.011)
AIAN	0.202	0.094	-0.070	0.118	0.027	-0.303	0.207	-0.033	0.063	0.153
	(0.006)	(0.032)	(0.018)	(0.012)	(0.041)	(0.029)	(0.017)	(0.019)	(0.068)	(0.009)
Multi./Other	-0.320	-1.366	-0.458	-0.677	0.463	-0.870	0.449	-0.136	0.266	-0.512
	(0.011)	(0.077)	(0.029)	(0.034)	(0.035)	(0.039)	(0.049)	(0.019)	(0.055)	(0.045)
N	>7.5M	>500k	>1.5M	>1M	>750k	>1M	>500k	>750k	>250k	>500k
R^2	0.014	0.030	0.017	0.013	0.023	0.019	0.010	0.011	0.010	0.017

Note. — Robust standard errors in parentheses. Dependent variable is the amount marked up. Census regions presented in ascending order of prejudice index at the percentile corresponding to the percent Black in the region (both calculated from the GSS; see Table 5.

Table 7. Summary statistics for auto transactions, by quartile of price paid for vehicle.

	Overall	Price Quartile 1	Price Quartile 2	Price Quartile 3	Price Quartile 4
Markup Amount	1.125	1.189	1.184	1.174	1.180
	(0.868)	(0.882)	(0.842)	(0.825)	(0.809)
Prop. Marked Up	0.733	0.745	0.764	0.772	0.786
	(0.442)	(0.436)	(0.425)	(0.419)	(0.410)
Black	0.044	0.108	0.109	0.109	0.107
	(0.001)	(0.204)	(0.206)	(0.206)	(0.206)
Hispanic	0.039	0.097	0.096	0.094	0.092
	(0.001)	(0.241)	(0.241)	(0.238)	(0.237)
API	0.065	0.046	0.043	0.040	0.038
	(0.001)	(0.156)	(0.148)	(0.141)	(0.137)
AIAN	0.081	0.012	0.012	0.012	0.012
	(0.003)	(0.056)	(0.056)	(0.056)	(0.055)
Muli./Other	-0.132	0.020	0.020	0.019	0.019
	(0.005)	(0.032)	(0.030)	(0.030)	(0.030)
N	> 7.5M	>1.5M	>1.5M	>1.5M	>1.5M

Note. — Price quartiles are determined controlling for a vehicle's make, model, year, new/used status, and region (to account for geographic differences in demand). Note that more than 500,000 transactions do not have disaggregated price data (instead reporting the overall loan amount, inclusive of add-ons and exclusive of down payment). As such, the weighted averages of the quartile averages will not necessarily equate to the overall average.

Table 8. Summary statistics for auto transaction by quartile of buy rate assigned to borrower.

	Overall	Buy Rate Quartile 1	Buy Rate Quartile 2	Buy Rate Quartile 3	Buy Rate Quartile 4
Markup Amount	1.125	1.073	1.102	1.169	1.164
	(0.868)	(0.821)	(0.859)	(0.873)	(0.927)
Prop. Marked Up	0.733	0.749	0.737	0.741	0.694
	(0.442)	(0.433)	(0.440)	(0.438)	(0.461)
Black	0.106	0.086	0.081	0.110	0.143
	(0.106)	(0.171)	(0.222)	(0.207)	(0.249)
Hispanic	0.088	0.064	0.041	0.098	0.117
	(0.232)	(0.191)	(0.144)	(0.247)	(0.269)
API	0.039	0.043	0.000	0.038	0.034
	(0.141)	(0.152)	(0.000)	(0.138)	(0.130)
AIAN	0.011	0.011	0.019	0.012	0.013
	(0.053)	(0.047)	(0.030)	(0.054)	(0.062)
Muli./Other	0.019	0.019	0.000	0.019	0.020
	(0.030)	(0.029)	(0.000)	(0.030)	(0.031)
N	> 7.5M	>1.75M	>1.75M	>1.75M	>1.75M

Note. — Buy rate is the risk-based pricing of the loan as determined by the lender. In addition to establishing the "baseline price" of the loan, it may also serve as an index of the consumer's financial sophistication, as they are observed by the dealers, and constitute a standalone indicator of all relevant information in a consumer's credit history(as determined by lender).

Table 9. Markup Regressions Evaluating Predictions of Taste-Based and Statistical Models of Discrimination.

	Race Probability Proxy			MAP Assignment to Race		
	I	II	III	IV	V	VI
Black x Marginal Index	0.980	0.847	0.754	0.491	0.390	0.342
	(0.046)	(0.107)	(0.120)	(0.040)	(0.115)	(0.178)
Black x Average Index	0.123	0.140	0.150	0.088	0.125	0.097
	(0.704)	(0.661)	(0.641)	(0.584)	(0.433)	(0.550)
Black x FICO $\geq$ 720		-0.041			-0.016	
		(0.036)			(0.204)	
Black x Buy Rate			-0.020			-0.017
			(0.000)			(0.000)
FICO $\geq$ 720		-0.294			-0.316	
		(0.000)			(0.000)	
Buy Rate			-0.011			-0.008
			(0.000)			(0.000)
N	$> 7.5M$	$> 7.5M$	$> 7.5M$	$> 6.5M$	$> 6.5M$	$> 6.5M$

Note. — Independent variable is the discretionary markup imposed on the loan. In the race probability specifications the BISG probabilities of a borrower's race are used to proxy for race (with all complementary proxies included). In the MAP assignment to race specifications race is assigned based on the maximum proxy value, and only those borrowers assigned to Black or White are kept. The construction of the prejudice indices are described in section 4. P-values based on standard errors clustered at the state level in parentheses. Marginal and average prejudice levels test implications of the Becker model. FICO scores and buy rate are signals of the customer's financial sophistication that serve as a test for the prediction of statistical discrimination models. Each specification includes all components of the interacted terms as covariates.

Table 10. Coefficient Estimates from Repeaded Imputation Markup Regressions.

	Mean Coefficient	Fifth Percentile	Ninety-Fifth Percentile
Basic Specification			
Black x Marginal Index	0.560	0.526	0.595
Black x Average Index	0.060	0.043	0.078
Signaling Specification			
Black x Marginal Index	0.421	0.388	0.458
Black x Average Index	0.085	0.067	0.103
Black x FICO $\geq$ 720	-0.027	-0.029	-0.024
FICO $\geq$ 720	-0.3008	-0.3012	-0.3005

Note. — Coefficients are from 1,000 regressions where the observations were randomly assigned a race based on the BISG proxy values. As in Table 9, the independent variable is the discretionary markup imposed on the loan, and the indices are determined based on responses to relevant questions in the GSS reported in the GSS. Each specification also includes all components of the interaction terms as covariates.

Table 11. Coefficient Estimates from Search Regressions.

	I	II	III	IV
Pr(Black) x Share Black	2.541	1.598	-2.343	-2.675
	(0.000)	(0.000)	(0.006)	(0.003)
Pr(Black) x Share Prejudiced	-3.957	-5.188	-5.163	-5.262
	(0.000)	(0.000)	(0.000)	(0.000)
Pr(Black) x Share Unprejudiced		-8.041		8.903
		(0.000)		(0.084)
Pr(Black) x Marginal Index			4.674	4.832
			(0.000)	(0.000)
Pr(Black) x Average Index			0.939	2.512
			(0.002)	(0.000)
N	>7.5M	>7.5M	>7.5M	>7.5M
R^2	0.015	0.016	0.018	0.023

Note. — Independent variable is the discretionary markup imposed on the loan. Pr(Black) is the BISG probability of a borrower being Black. Share Black is the share of a state's population that is Black. Shares prejudiced/unprejudiced are the shares of GSS respondent's whose prejudice indices are two or more standard deviations above/below the national average. Marginal index is the b^{th} percentile of the prejudice index, where the b is the percent Black in a given state; average index is the mean of the prejudice index in the borrower's state. P-values based on standard errors clustered at the state level in parentheses. All specifications include the non-interacted components of the interaction terms, and the full set of BISG race probabilities (with non-Hispanic White as the omitted group). Search models predict positive coefficients on Pr(Black) x Share Prejudice and negative coefficients on Pr(Black) x Share Black and Pr(Black) x Share Prejudiced (see Section 5.4).

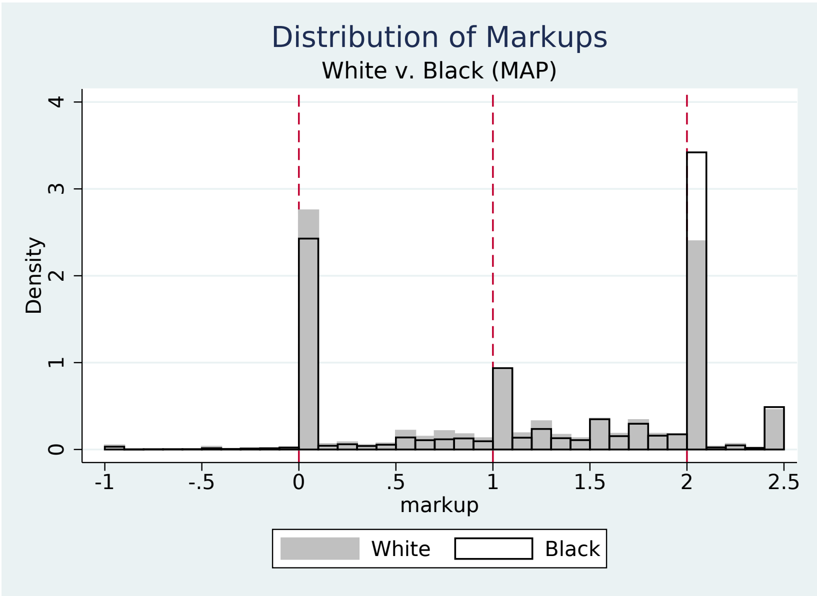

Fig. 1.— Distribution of markups imposed on loans for Black and White customers, with race approximated by MAP assignment. See Section 4 for more details.

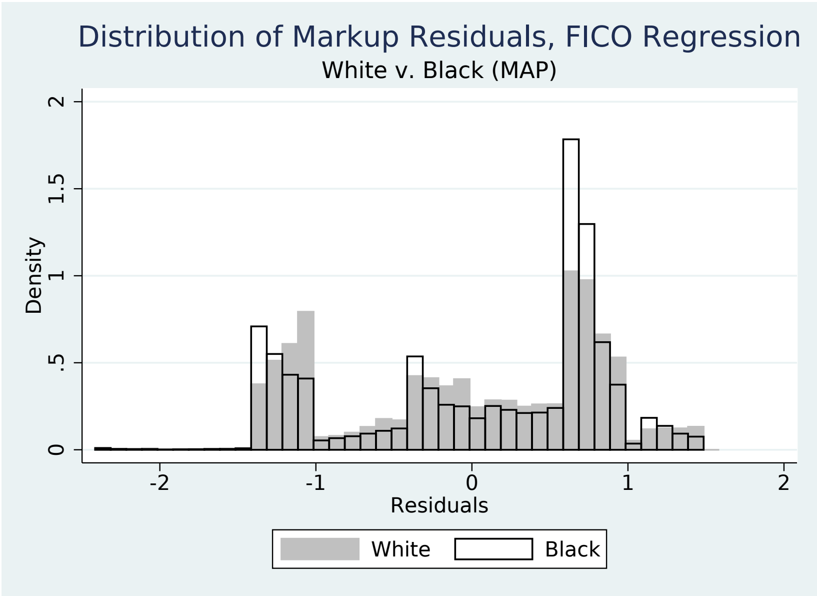

Fig. 2.— Distribution of residuals from regression of markup on FICO for Black and White customers, with race approximated by MAP assignment. See Section 5 for more details.

REFERENCES

Aigner, Dennis J. and Glen G. Cain, "Statistical Discrimination in the Labor Market," *Industrial and Labor Relations Review*, 1977, *30*, 175–187.

Altonji, Joseph G. and Charles R. Pierret, "Employer Learning and Statistical Discrimination," *Quarterly Journal of Economics*, February 2001, *16* (1), 313–350.

——— **and Rebecca M. Blank**, "Race and Gender in the Labor Market," in Orley Ashenfelter, Richard Laynard, and David Card, eds., *Handbook of Labor Economics*, Vol. 3C, Elsevier Science B.V., 1999, chapter 48, pp. 3143–3259.

Antonovics, Kate, Peter Arcidiacono, and Randall Walsh, "Games and Discrimination: Lessons from 'The Weakest Link',," *Journal of Human Resources*, Fall 2005, *40* (4), 918–947.

Anwar, Shamena, "Testing for discrimination: Evidence from the game show Street Smarts," *Journal of Economic Behavior & Organization*, 2012, *81* (1), 268–285.

Arcidiacono, Peter, Patrick Bayer, and Aurel Hizmo, "Beyond Signaling and Human Capital: Education and the Revelation of Ability," *American Economics Journal: Applied Economics*, 2010, *2* (4), 76–104.

Argyle, Bronson, Taylor D. Nadauld, and Christopher Palmer, "Monthly Payment Targeting and the Demand for Maturity," *Review of Financial Studies*, 2020, *33* (11), 5416–5462.

——— **,** ——— **, and** ——— , "Real Effects of Search Frictions in Consumer Credit Markets," *NBER Working Paper 26645*, 2020.

Arrow, Kenneth J., "Some Mathematical Models of Race in the Labor Market," in A.H. Pascal, ed., *Racial Discrimination in Economics Life*, Lexington Books, 1972, pp. 187–204.

Attanasio, Orazio P., Pinelopi Koujianou Goldberg, and Ekaterini Kyriazidou, "Credit Constraints In The Market For Consumer Durables: Evidence From Micro Data On Car Loans," *International Economics Review*, 2008, *49* (2), 401–436.

Ayers, Ian, "Outcome Tests of Racial Disparities in Police Practices," *Justice Research and Policy*, 2002, *4*, 131–142.

——— **and Peter Siegelman**, "Race and Gender Discrimination in Bargaining for a New Car," *American Economic Review*, 1995, *85* (3), 304–321.

Barsky, Robert, John Bound, Kerwin Kofi Charles, and Joseph P. Lupton, "Accounting for the Black-White Wealth Gap: A Nonparametric Approach," *Journal of the American Statistical Association*, 2002, *97* (459), 663–673.

Bartlett, Robert, Adair Morse, Richard Stanton, and Nancy Wallace, "Consumer-Lending Discrimination in the FinTech Era," *Journal of Financial Economics*, 2021, *In Press.*

Becker, Gary S., *The Economics of Discrimination*, Chicago: University of Chicago Press, 1957.

Bhutta, Neil and Aurel Hizmo, "Do Minorities Pay More for Mortgages?," *The Review of Financial Studies*, 2021, *34* (2), 63789.

Black, Dan A., "Discrimination in an Equilibrium Search Model," *Journal of Labor Economics*, April 1995, *13* (2), 309–334.

Blau, Francine D. and John W. Graham, "Black-White Differences in Wealth and Asset Composition," *Quarterly Journal of Economics*, May 1990, *105* (2), 321–339.

Bohren, J. Aislinn, Alex Imas, and Michael Rosenberg, "Outcome Tests of Racial Disparities in Police Practices," *American Economic Review*, 2019, *109* (10), 3395–3436.

Busse, Meghan R. and Jorge M. Silva-Risso, ""One Discriminatory Rent" or "Double Jeopardy": Multicomponent Negotiation for New Car Purchases," *American Economic Review, Papers and Proceedings*, 2010, *100* (2), 470–474.

Butler, Alexander W., Erik J. Mayer, and James Peter Weston, "Racial Discrimination in the Auto Loan Market," *Working Paper*, 2021.

Cain, Glen G., "The Economic Analysis of Labor Market Discrimination: A Survey," in Orley Ashenfelter and Richard Laynard, eds., *Handbook of Labor Economics*, Vol. 1 1986.

CFPB, "Using publicly available information to proxy for unidentified race and ethnicity: A methodology and assessment," Technical Report, Consumer Financial Protection Bureau, Washington, D.C. Summer 2014.

______ , "Fair Lending Report of the Consumer Financial Protection Bureau," Technical Report, Consumer Financial Protection Bureau, Washington, D.C. April 2015.

Charles, Kerwin Kofi and Erik Hurst, "The Transition to Home Ownership and the Black-White Wealth Gap," *Review of Economics and Statistics*, 2002, *84* (2), 281–297.

______ **and Jonathan Guryan**, "Prejudice and Wages: An Empirical Assessment of Becker's The Economics of Discrimination," *Journal of Political Economy*, October 2008, *116* (5), 773–809.

______ **, Erik Hurst, and Melvin Stephens**, "Rates for Vehicle Loans: Race and Loan Source," *American Economic Review: Papers and Proceedings*, 2008, *98* (2), 315–320.

Coate, Stephen and Glenn C. Loury, "Antidiscrimination Enforcement and the Problem of Patronization," *American Economic Review*, May 1993, *83* (2), 92–98.

Cohen, Mark A., "Imperfect Competition in Auto Lending: Subjective Markups, Racial Disparity, and Class Action Litigation," *Review of Law and Economics*, 2012, *8* (1), 21–58.

Davis, Delvin, "The State of Lending in America and its Impact on U.S. Households: Auto Loans," Technical Report, Center for Responsible Lending, Durham, N.C. December 2012.

______ **and Joshua M. Frank**, "Under The Hood: Auto Loan Interest Rate Hikes Inflate Consumer Costs and Loan Losses," Technical Report, Center for Responsible Lending, Durham, N.C. April 2011.

Dobbie, Will, Andres Liberman, and Daniel Paravisini, "Measuring Bias in Consumer Lending," *The Review of Economic Studies*, 2020, *In Press.*

Elliott, Marc N., Allen Fremont, Peter A. Morrison, Philip Pantoja, and Nicole Lurie, "A New Method for Estimating Race/Ethnicity and Associated Disparities Where Administrative Records Lack Self-Reported Race/Ethnicity," *Health Services Research*, 2008, *43* (5, part 1), 17221736.

Goldberg, Pinelopi Koujianou, "Dealer Price Discrimination in New Car Purchases: Evidence from the Consumer Expenditure Survey," *Journal of Political Economy*, 1996, *104* (3), 622–654.

Grunewald, Andreas, Jonathan Lanning, David Low, and Tobias Salz, "Auto Dealer Loan Intermediation: Consumer Behavior and Competitive Effects," *NBER Working Paper 28136*, 2020.

Lang, Kevin and Jee-Yeon K. Lehmann, "Racial Discrimination in the Labor Market: Theory and Empirics," *Journal of Economic Literature*, December 2012, *50* (4), 959–1006.

______ **and Michael Manove**, "Education and Labor Market Discrimination," *American Economic Review*, 2011, *101* (4), 1467–1496.

______ **,** ______ **, and William T. Dickens**, "Racial Discrimination in Labor Markets with Posted Wage Offers," *American Economic Review*, September 2005, *95* (4), 1327–1340.

Lanning, Jonathan A., "Productivity, Prejudice, and Lost Profits During Baseball's Integration," *Journal of Economic History*, 2010, *70* (4), 965–989.

______ , "A Search Model With Endogenous Job Destruction and Discrimination: Why Equal Wage Policies May Not Eliminate Wage Disparity," *Labour Economics*, 2014, *26* (1), 55–71.

Levitt, Steven D., "Testing Theories of Discrimination: Evidence from Weakest Link," *Journal of Law and Economics*, October 2004, *47*, 431–452.

List, John A., "The Nature and Extent of Discrimination in the Marketplace: Evidence From the Field," *Quarterly Journal of Economics*, 2004, *119* (1), 49–89.

Morse, Adair and Karen Pence, "Technological Innovation and Discrimination in Household Finance," *NBER Working Paper 26739*, 2020.

Morton, Fiona Scott, Florian Zettelmeyer, and Jorge Silva-Risso, "Consumer Information and Discrimination: Does the Internet Affect the Pricing of New Cars to Women and Minorities?," *Quantitative Marketing and Economics*, 2003, *1* (1), 6592.

Sasaki, Masaru, "An Equilibrium Search Model with Coworker Discrimination," *Journal of Labor Economics*, April 1999, *17* (2), 377–407.

Simoiu, Camelia, Sam Corbett-Davies, and Sharad Goeli, "The problem of inframarginality in outcome tests for discrimination," *The Annals of Applied Statistics*, 2017, *11* (3), 1193–1216.

Spriggs, William, "Is now a teachable moment for economists? An open letter to economists from Bill Spriggs," Technical Report, Howard University June 2020.

Szymanski, Stefan, "A Market Test for Discrimination in the English Professional Soccer Leagues," *Journal of Political Economy*, June 2000, *108* (3), 590–603.

Voicu, Ioan, "Imperfect Competition in Auto Lending: Subjective Markups, Racial Disparity, and Class Action Litigation," February 2016. Available at SSRN: http://dx.doi.org/10.2139/ssrn.2763826.

This preprint was prepared with the AAS LaTeX macros v5.2.

A. Economic Theories of Discrimination

A.1. Becker-Style Theory of Taste-Based Discrimination in Markup

Adding a distaste parameter to make the model more consistent with a Becker-type taste-based model is straightforward. In a simple Becker model, distaste would enter into the dealer's utility in a linear fashion, yielding:

$$U_k = r_k + f_k(j) - d_k(j) \tag{A1}$$

where r_k is the non-financing rents earned by dealer k, $f_k(j)$ is the financing rents earned by dealer k on a deal with a customer of type j (either ϕ_k or m_i, depending on whether the dealer imposes markup or accepts the flat fee from the lender), and $d(j)_k$ is the level of distaste dealer k has for customers of type j, which takes the value 0 for favored types, and d for disfavored types.[49] From this it is trivial to derive that, for a common value of r_k, in equilibrium:

$$f_k^*(w) = f_k^*(b) - d_k^*. \tag{A2}$$

That is, the observed financing rents extracted by dealers will be higher on deals with disfavored types (in this case Black) than on similar deals with favored types (in this case White) by an amount equal to the distaste of the marginal discriminator (i.e. the dealer exactly indifferent to financing a deal with a disfavored customer versus a favored customer). Note that the formulation above assumes a common value of r, and a single type of customer within each race. In practice, there may be many other factors affecting the dealer's utility and/or capacity to assign markup. As such, it may be easier and more accurate to conceptualize this finding as indicating that, ceteris paribus, Black customers will face higher markups in equilibrium than similarly situated White customers. Correspondingly, as the array in 2 implies the realized values of m_i will exceed those of ϕ_i, Equation A1 implies that disfavored types will both be more likely to experience non-zero markup, and will experience higher average markups.

---

[49] Obviously more complex functional forms of utility could also be applied. However, as the dealer receives monotonically increasing returns from markup, and pays no explicit cost for said markup, a linear utility assumption seems appropriate.

While the marginal level of prejudice directly impacts the markup in the formulation above, it is interesting to note that the average level of prejudice has no direct impact on the relative levels of markup. This is a familiar result in economics. In effect, the levels of the supply and demand curves to the left of equilibrium are relevant for the surplus derived from the market-clearing price, but do not affect the price itself. Similarly, within the population of dealers who are unlikely to interact with disfavored types, the distribution of prejudice should be irrelevant to the market outcomes for disfavored types. However, for a given distribution of prejudice, a change in the share of the population constituted by disfavored types will affect the market outcome. This is, effectively, a shift in the demand for financing by disfavored type. Thus the Becker-style taste-based models provide two sharp predictions about the relationship between prejudice and markup; the average markup disparity will be positively affected by the marginal level of prejudice in a market, and should have a negligible relationship with the average level of prejudice in that market.

An interesting idiosyncrasy in this market is that, depending on how d is conceptualized, this setup could also provide dueling incentives to a dealer when negotiating markup for a disfavored customer. If d is simply the negative enhancement to a dealer from selling a car to a disfavored consumer, then the incentives are clear and unambiguous: increase all flexible aspects of the deal yielding positive utility (e.g. price of vehicle, markup, add-ons, etc.) to offset the negative utility from the dealer's distaste. However, if the dealer bears the "cost" of his/her distaste simply from negotiating with the customer, then the dealer will still want to increase all flexible aspects of the deal yielding positive utility, but will be less likely to risk assigning a markup level that would result in the customer walking away. That is, because the distaste is paid throughout the negotiation process, the dealer may not assign as high a markup, since the customer walking away from the deal would result in negative, as opposed to zero utility in the case where the distaste parameter is realized only at the conclusion of the deal. This is unlike the typical labor context envisioned by most economic models of discrimination. However, this should only attenuate the realized disparity facing disfavored customers (as dealers would be willing to initiate more deals at lower markups than in a situation where the distaste is only realized at a deals consummation), biasing the predictions of the model downward. The prediction that markup for disfavored types will increase as the marginal level of distaste (prejudice) increases, but should not be related to the average level of distaste, will form the basis of the empirical test for taste-based discrimination to follow.

A.2. Statistical Model of Discrimination in Markup

A standard statistical discrimination framework in this market could take the form of a dealer making different inferences about a customer's a level of financial sophistication from both a buyer's race/ethnicity and a signal of his/her financial literacy.[50] In a "classic" formulation:

$$E[\sigma_i \mid s_i, j] = \alpha_j * s_i + (1 - \alpha_j) * \mu(\sigma_j), \tag{A3}$$

where $E[\sigma_i]$ is the dealer's expectation about buyer i's financial sophistication, s_i is the signal sent by buyer i, j is the buyer's type (e.g. race, ethnicity), $\alpha_j \in [0, 1]$ is the weight the dealer places on signals from type j buyers, and $\bar{\sigma_j}$ is the average level of financial sophistication over buyers of type j.[51] For empirical purposes it is often useful to make the simplifying assumption that $\mu(s_j) = \mu(\sigma_j)$, which allows the above to be written in terms likely to be observable to the econometrician. Specifically:

$$E[\sigma_i \mid s_i, j] = \alpha_j * s_i + (1 - \alpha_j) * \mu(s_j). \tag{A4}$$

I will impose this assumption that the distribution of signals corresponds to the distribution of "true" financial sophistication within type, and that these distributions are known to the market.

[50] It should be noted that there are many, "less standard" models of statistical discrimination that could also be applied to this market including ones with investment in financial literacy, self-fulfilling stereotypes, etc. For simplicity's sake, only the "classic" formulation is considered here.

[51] In most models of statistical discrimination, it is assumed that the average characteristics of the different types are known, and accurately assessed by the discriminating agents. This ensures a rationalizable model and outcome. However, Coate and Loury (1993) show that inaccurate stereotypes can actually induce behavior among disfavored types that rationalize the once-inaccurate stereotypes. It is also possible that, in a system with limited information and slow learning, inaccurate stereotypes can persist even if behavior doesn't change in a way that would rationalize the stereotypes.

In this framework, the s term in Equation 3 would be replaced by an expectation, such that

$$E[\bar{m}_i] = f(p_i, E[\sigma_i|j], r_k, \epsilon_{i,k}), \tag{A5}$$

or, ignoring the structure imposed by Equation A3,

$$E[\bar{m}_i] = f(s_i, j, n_i, r_k, \epsilon_{i,k}). \tag{A6}$$

To the extent that dealers expect the markup a buyer will tolerate to be inversely related to that buyer's financial sophistication (i.e. $\frac{dE[\bar{m}]}{dE[\sigma]} < 0$), this type of statistical discrimination could lead to disparities for minority buyers relative to non-Hispanic White buyers, conditional on the same signal of financial sophistication. This would result if

$$\frac{dE[\sigma]}{ds_w} > \frac{dE[\sigma]}{ds_{nw}}, \tag{A7}$$

which implies

$$\frac{dE[\bar{m}]}{ds_w} > \frac{dE[\bar{m}]}{ds_{nw}} \tag{A8}$$

(where w and nw indicate White and non-White). The disparities described in Equation A7 could result if $\alpha_n \neq \alpha_{nw}$ and/or if $\mu(\sigma_n) \neq \mu(\sigma_{nw})$. In the case where $\alpha_{nw} < \alpha_w$ we would expect markups for minority buyers more tightly distributed around their mean compared to White buyers with similar signals. In the case where $\mu(\sigma_{nw}) < \mu(\sigma_w)$ we would expect an average markup for minority borrowers that is higher than for White borrowers who send the same signal. The former result will be evaluated with suggestive, graphical evidence, while the latter result will form the basis of the empirical test for statistical discrimination to follow.

It is important to note this is a somewhat simple form of statistical discrimination. Unlike in the models tested in labor markets, there is no private learning about buyers over time, no opportunity for more information to be signaled to the market over time by the buyers themselves, no self-fulfilling stereotypes, etc.[52] While this may seem like a limitation relative to the more complex

[52]See Cain (1986), Altonji and Blank (1999), and Lang and Lehmann (2012) for examples of the more complex models.

models applied to the labor market, I argue this is a more appropriate model for this market. As financing a car is a discrete transaction that typically takes place over a matter of hours, there is minimal opportunity for dealers to learn about buyers' true levels of financial sophistication beyond the initial signals. As auto purchases are often in response to unanticipated event (e.g. previous car breaking down) and as the financing portion of the transaction is often a lesser consideration relative to make, model, price, etc., it is unlikely many buyers work to improve their signals of sophistication (e.g. credit scores) in anticipation of the transaction. In short: if dealers setting markup based on signals of financial sophistication, they are likely to base this markup on the observed signals at the time of the transaction only, and those signals are unlikely to have been differentially manipulated by buyers. This indicates both that a static model of statistical discrimination without learning is likely appropriate.

A.3. A Model of Search Over Markup with Discrimination

The model used here is similar to the one developed by Black (1995), which is one of the first and most cited search models with discrimination. In that model a small number of prejudice firms were shown to have a persistent impact on the wage outcomes of disfavored workers. The basic result follows from the fact that, if any portion of the market is unavailable to disfavored types, these types have lower expected values of (non-targeted) searching, resulting in lower reservation wages and equilibrium wage offers by even non-prejudice, profit-maximizing firms.[53] Unlike in Becker-type models, this result does not rely on the interaction of any disfavored type with any prejudiced firm; the mere existence of these firms is enough to depress the equilibrium wages. This model

[53]Lang et al. (2005) employ a different approach, with targeted search that arises via wage posting. Such a model may be useful in a context with subvented loans, which are loans subsidized by lenders in an attempt to entice customers. These loans are often offered and advertised by captive lenders associated with a specific manufacturer in an attempt to sell particular vehicles. As subvented loans are excluded from this analysis, the practice of markup is generally unknown to the market, and non-subvented interest rates are rarely advertised, I focus here on non-targeted search.

was nested in a later formulation by Lanning (2014), which added non-discriminating firms that do not exploit the presence of prejudice to lower their wage offers, and endogenous job destruction that places Becker-type competitive pressure on prejudice firms. Yet, even in the presence of these mitigating forces, discrimination persists in equilibrium.

Both models described above start with a share of the market composed of prejudiced firms that will not disfavored workers at any reasonable wage. In the context of this paper, that could correspond to a share of prejudiced finance and insurance employees who refuse to secure loans for disfavored workers (or who would only be willing to secure loans with untenable markups). This seems unlikely, given that these employees are not owners of the business, and only enter the transaction after a viable sale has been negotiated. As such, any unwillingness to offer loans would undermine the entire sale and lose the dealership profits. Only in a case where a discriminatory owner allowed this to occur would F&I employees be able to indulge their discriminatory tastes on the extensive margin; even in the case of such prejudiced ownership, we would expect the discrimination to occur prior to the financing stage. Unlikely though it may be, it is possible that minority customers believe that some prejudice dealers exist who are unwilling to secure financing for these customers. This would reduce the expected value of continued search and, if that fact is known and exploited by dealers, would result in the same outcomes predicted by the model.

The model developed here derives reservation markups that differ across customer types based on the different expected values of continued searching for jobs. In the context of this paper, that would correspond to higher reservation markups for disfavored types based on the lower expected value of search (in the form of higher expected markups) for those types. For the sake of consistency with the literature, it is easier to conceptualize the object of the search to be a discount relative to the maximum markup, i.e. $\gamma = M - m$, where M is the maximum allowable markup and γ is the discount from that markup. Making the target of search a negative function of markup allows it to be a positive-utility good, which is consistent with the formulation of search models with discrimination (which have almost exclusively been developed in the labor context as a search over wages). The value of search for a favored (in this case, White) type in such a model could be written as:

$$S_i^w = \theta_p \int_0^\Gamma \gamma f_p(\gamma) d\gamma + \theta_n \int_0^\Gamma \gamma f_n(\gamma) d\gamma + \theta_s \int_0^\Gamma \gamma_w f_s(\gamma_w) d\gamma_w, \tag{A9}$$

where S_i^w is the value of a search for worker i of race w; $\theta_{p,n,s}$ are the shares of dealerships that are prejudice, non-discriminatory (i.e. that offer same markup to all borrowers regardless of race), and strategic-discriminators (i.e. that vary their markups based on race, in proportion to the value of future searches for members of that race); and $\int_0^\Gamma \gamma f_k(\gamma) d\gamma$ is the expected value of markup discount γ from a firm of type k (with a common, maximum allowable markdown of Γ, which can be greater than M if lenders allow dealers to provide net markdowns for the customer[54]). Noting that a searcher will never accept a markup discount below her reservation level, the above becomes:

$$S_i^w = \theta_p \int_{\bar{\gamma}_i}^\Gamma \gamma f_p(\gamma) d\gamma + \theta_n \int_{\bar{\gamma}_i}^\Gamma \gamma f_n(\gamma) d\gamma + \theta_s \int_{\bar{\gamma}_i}^\Gamma \gamma_w f_s(\gamma_w) d\gamma_w + [F_p(\bar{\gamma}_i) + F_n(\bar{\gamma}_i) + F_s(\bar{\gamma}_i)]\beta S_i^w, \tag{A10}$$

where $\bar{\gamma}_i$ is searcher i's reservation markup discount, and $F_{p,n,s}(\bar{\gamma}_i)$ are the chances that a searchers receives a markup discount below her reservation value, and searches again next period, with discounted expected value βS. Assuming stationary (i.e. the expected value of S_i^w is constant with respect to time) yields:

$$S_i^w = \frac{\theta_p \int_{\bar{\gamma}_i}^\Gamma \gamma f_p(\gamma) d\gamma + \theta_n \int_{\bar{\gamma}_i}^\Gamma \gamma f_n(\gamma) d\gamma + \theta_s \int_{\bar{\gamma}_i}^\Gamma \gamma_w f_s(\gamma_w) d\gamma_w}{1 - [F_p(\bar{\gamma}_i) + F_n(\bar{\gamma}_i) + F_s(\bar{\gamma}_i)]\beta}. \tag{A11}$$

This is the equilibrium value of search for white searchers.

$$S_i^b = \frac{\theta_n \int_{\bar{\gamma}_i}^\Gamma \gamma f_n(\gamma) d\gamma + \theta_s \int_{\bar{\gamma}_i}^\Gamma \gamma_b f_s(\gamma_b) d\gamma_b}{1 - [\theta_p + F_n(\bar{\gamma}_i) + F_s(\bar{\gamma}_i)]\beta} \tag{A12}$$

Equations A11 and A12 show that the expected discount received by Black searches is lower than that of White searches. This results from both a larger numerator in Equation A11 (where

[54]Note that in many cases it is possible for dealers to provide markdowns for customers. Most often these markdowns are granted in "special" cases, e.g. when the customer has a competing offer, when the dealer is "bundling" multiple loans of varying quality for a specific lender, etc. In practice these net markdowns are uncommon in the data, comprising around one percent of observations.

there is a positive expected value from type p dealers), and a larger denominator in Equation A12 (where the entire weight of θ_p yields no positive value, and in expectation results in a discounted future search). In equilibrium the reservation markup discounts are equal to the expected value of a search, meaning disfavored searchers (in this case Black) set lower reservation markup discounts, and are subsequently willing to accept higher markups.

Black (1995) considers the equilibrium wage gap (analagous to the markup discount gap here) in the above system of equations when there are only two firm types (prejudice and strategic discriminators, i.e. $\theta_n = 0$), and finds the gap decreases as the share of Black searchers increases (i.e. $\frac{d(m^w - m^b)}{dBlack} < 0$), and/or as the share of prejudiced firms decreases (i.e. $\frac{d(m^w - m^b)}{dtheta_p} > 0$). Lanning (2014) extends the model by adding non-discriminating firms and endogenous firm destruction, and shows the equilibrium gap similarly depends on the share of Black searchers and prejudiced firms, but also on the share of unprejudiced firms in the market, with the gap decreasing as the share unprejudiced increases (i.e. $\frac{d(m^w - m^b)}{d\theta_n} < 0$). These are the predictions that I will empirically test with the data.