Customer Discrimination and Quality Signals: A Field Experiment with Healthcare Shoppers

Alex Chan*

Stanford University

(Click here for the most recent version)

September 26, 2021

Abstract

This paper provides evidence that customer discrimination in the market for doctors can be largely accounted for by statistical discrimination. I evaluate customer preferences in the field with an online platform where cash-paying consumers can shop and book a provider for medical procedures based on an experimental paradigm called validated incentivized conjoint analysis (VIC). Actual paying customers evaluate doctor options they know to be hypothetical to be matched with a customized menu of real doctors, preserving incentives. Racial discrimination reduces patient willingness-to-pay for black and Asian providers by 12.7% and 8.7% of the average colonoscopy price respectively; customers are willing to travel 100–250 miles to see a white doctor instead of a black doctor, and somewhere between 50–100 to 100–250 miles to see a white doctor instead of an Asian doctor. Providing signals of provider quality reduces this willingness-to-pay racial gap by about 90%, which suggests that statistical discrimination is an important cause of the gap. Actual booking behavior allows cross-validation of incentive compatibility of stated preference elicitation via VIC. (*JEL* Codes: J71, I11, L15, L86, M31)

Keywords: Discrimination; Field Experiment; Healthcare Consumers; Conjoint Analysis

*Email: chanalex@stanford.edu. I cannot thank my advisors, Loren Baker and Al Roth, enough for their guidance and encouragement. I also thank Nick Bloom, Dave Chan, Matthew Gentzkow, David Klinowski, Negar Matoorian, Michelle Mello, Thayer Morrill, Muriel Niederle, Heather Royer,

1 Introduction

Discrimination by customers on the basis of race can arise when selecting sellers of the goods and services they will use. Customer discrimination of this type may be particularly notable when consumers are seeking services from high-skilled professionals. Customer discrimination can pose serious problems, including threatening equality in labor markets and leading to inefficiencies in human capital investments and labor market participation. It is important to identify the existence and the underlying sources of discrimination so that we can discuss market design and optimal policies to address customer discrimination (List (2004); Gneezy et al. (2012); Fisman and Luca (2016)).

Existing reports have raised concerns about customer discrimination in medicine, where patients seeking care often select from multiple physicians or other healthcare providers. Anecdotal evidence from interviews with doctors has documented reports of discrimination by patients toward black and Asian doctors.[1] In addition, a large body of circumstantial evidence is consistent with the presence of discrimination: minority physicians are underrepresented in the U.S. medical workforce (Bergen Jr (2000), Lett et al. (2019), Merchant and Omary (2010)), and minority physicians earn less than white physicians.[2] Such evidence is consistent with a reduction in the marginal revenue product of minority professionals associated with customer discrimination.[3] The result is

Colin Sullivan, Kurt Sweat, and Kenneth Train for helpful comments, and thank all the doctors and patients who participated in detailed interviews. Special thanks go to the partnering platform's CEO and Director of Business Development. This paper also benefitted from comments from participants in a number of seminars and conferences. The study was registered on the AEA RCT registry, ID AEARCTR-0006815. This project was supported by the Gerhard Casper Stanford Graduate Fellowship and grant number T32HS026128 from the Agency for Healthcare Research and Quality. The content is solely the responsibility of the author and does not necessarily represent the official views of the Agency for Healthcare Research and Quality.

[1] See Filut et al. (2020) for a recent systematic review of the literature on discrimination toward physicians of color. In Filut et al. (2020), the authors noted that there is "almost no attempt to collect data on patient interactions... The lack of curiosity regarding experiences of physicians of color with discrimination from patients may reflect underlying assumptions that these physicians care solely or predominantly patients of color which have their root in U.S. history."

[2] In the U.S., black physicians earn 13.5% less than white physicians, and Asian physicians earn 7.8% less (Grisham (2017)). Also, P Ly et al. (2016) found that in the U.S. black male physicians have an adjusted median annual income that is 26% lower than that of their white counterparts while black female physicians have an adjusted median annual income that is 7% lower than there white counterparts.

[3] Other factors and doctor attributes can also affect the marginal revenue product of service professionals via customer preferences. Some papers document correlations between consumer preferences and various factors like communication skills (Fung et al. (2005); Hirpa et al. (2020)), distance (Schmitt et al. (2003); Yoon et al. (2019); Mooney et al. (2000)), and other quality measures (Harris (2003); Santos et al. (2017); Kolstad and Chernew (2009); Razzouk et al. (2004); Luft et al. (1990)). Li and Hubner (2019) conduct an experiment using paid online subjects who are not actual shoppers to engage in hypothetical and non-incentivized scenarios in order to evaluate the impact of online ratings on their

reduced access to a diverse set of healthcare providers. Research has suggested the positive impacts of a diverse workforce (Alsan et al. (2019)). Despite the importance of understanding customer discrimination in healthcare, there have been few analyses that provide direct evidence about discrimination or its structure.[4]

In this paper, I conduct a field experiment in which I study whether paying customers discriminate against black and Asian physicians when they choose healthcare providers.[5] In addition, the experiment is designed to allow me to obtain information about whether any discrimination observed is associated with statistical discrimination (Arrow et al. (1973)) or taste-based discrimination (Becker et al. (1971)). Statistical discrimination arises when customers use easily observable characteristics such as race to infer, accurately or inaccurately, the expected quality of doctors. Statistical discrimination can be eliminated when enough information about the quality of a doctor is shared with customers so that they can rely on information other than race to update their beliefs about the doctor. Taste-based discrimination theory suggests that customers choose as if there were a disutility from associating with a particular racial group. Even with perfect information, taste-based discrimination cannot be eliminated. A critical test for whether discrimination is statistical is to see whether the introduction of an economically relevant signal closes the racial gap.[6]

To test for discrimination, I observe paying customers as they shop for a doctor to perform a colonoscopy via an established online platform. This platform has over two hundred thousand, mostly self-paying, subscribed customers (prospective patients) who

choice of doctors and found that the subjects prefer doctors with higher ratings. In a related question of physicians referring patients to surgeons, Sarsons (2017) found that the gender of the surgeon influences the way signals about the surgeon's quality are interpreted.

[4]One reason for the dearth of studies is that it is hard to observe the full choice set faced by the customer along with all the individual doctor attributes pertinent to customer choice. Another reason is that there might not be enough variation in doctor race within the choice sets even when the data on provider options are available. Moreover, the race of a doctor can be correlated with many other factors that can explain customer choice.

[5]In the experiment, I do not differentiate between the various Asian ethnic groups and refer to East Asians (Chinese, Japanese, Korean, Okinawan, Taiwanese, Tibetan), Southeast Asians (Bruneian, Burmese, Cambodian, Filipino, Hmong, Indonesian, Laotian, Malaysian, Mien, Singaporean, Timorese, Thai, Vietnamese), and South Asians (Bangladeshi, Bhutanese, Indian, Maldivians, Nepali, Pakistani, Sri Lankan) as "Asians."

[6]Previous research has used this approach to disentangle taste-based and statistical discrimination. For example, "dressing for success" in List and Gneezy (2014), signaling low search costs by declaring "I'm getting a few price quotes" (Gneezy et al. (2012)), or having reviews posted on a guest's profile page on Airbnb (Cui et al. (2020)). In these experiments, the "quality signal" is introduced through an outfit signalling high socioeconomic status, a verbal statement signalling low search costs, and reviews signalling guest "quality" respectively.

can submit a request for a medical procedure and receive anonymized bids from providers across the United States. In this field setting, I show that prospective patients discriminate against black and Asian doctors, and that discrimination fades when information about doctor quality is provided to the shopper.

My experimental design involves a choice-based conjoint analysis where customers make choices from menus of hypothetical options described in terms of independently varied levels of attributes along various dimensions designed to mimic market experience (Ben-Akiva et al. (2019)). Actual customers shopping for a colonoscopy were recruited to the study. Subjects reviewed hypothetical doctor profiles for their procedure with exogenous variation in price, travel distance, gender, and race. Based on their responses to the conjoint survey, subjects were offered 10 actual physician options for booking. Subjects were asked to evaluate menus of hypothetical options with the promise to match the subject to a menu of actual options based on their evaluations of the menus of hypothetical options.[7] I also observed the actual booking choices made by patients among the actual doctors offering colonoscopies, to validate the estimated discrete choice model.

This is a novel way to incentivize a choice-based conjoint survey and I will refer to this paradigm as validated incentivized conjoint analysis (VIC). The main innovation of VIC is its method for incentivized preference elicitation that includes cross-validation using actual booking data from the same shopping instance.

I randomize the study population into two groups, one that receives quality information and one that does not. I calculate estimates of willingness-to-pay for attributes based on the results of the conjoint estimates. All attributes including quality are randomized and balanced across different races for the hypothetical doctor profiles in the VIC.

This experimental setting allows me to conduct two primary tests. In the first, using the patients randomized to receive no quality information, I can assess whether there is evidence for discrimination as measured by a willingness-to-pay gap between minority and white doctors. Second, I can compare the racial gap in willingness-to-pay among patients who did not receive quality information against the racial gap in willingness-to-pay for those who were randomized to receive quality information, to assess whether the provision

[7]Hypothetical bias can arise when the experimental design fails to prompt individuals to reveal their true preferences. If incentives are not present or not strong enough to prompt subjects to reveal their true type, subjects could answer conjoint surveys strategically, provide random answers, or deviate from their true preferences for other reasons. It is worth noting that previous methodological work in economics (List et al. (2006)) has found no evidence of hypothetical bias when estimating marginal attribute values as I do in the present paper (unlike "levels" or total willingness-to-pay).

of quality information changes the result. This is a test for statistical discrimination.

I report two main findings. First, there is customer discrimination against black and Asian doctors. Second, the majority of the observed discrimination is statistical. Customers are willing to pay a significant premium to have their colonoscopy done by a white doctor rather than a minority doctor. In particular, willingness-to-pay for black and Asian doctors is lower than the willingness-to-pay for white doctors by 12.7% and 8.7% of the average colonoscopy price of $2122, respectively. These racial gaps in willingness-to-pay suggest that customers are willing to travel 100–250 miles just to see a white doctor for the procedure instead of a black doctor, and somewhere between 50–100 to 100–250 miles to see a white doctor for the procedure instead of an Asian doctor. However, when quality information is provided for each doctor option, the willingness-to-pay gap for non-white providers drops significantly (to 1.2% and 1.1% of the average price for black and Asian doctors, respectively). I validate the estimated discrete choice model using actual shopping behavior of the study subjects and find that actual doctor booking coincides with one of three options predicted to be most likely chosen in 82% of the cases.

I then address several additional issues. First, could the results be due to information burden or distraction? I find that differences between treatment and control of the magnitudes described above seem to be present only among customers who later reported that the doctor quality signal gave them adequate information on doctor quality. This and other evidence rules out the hypothesis that observed differences are merely due to distraction given that the treatment group menus present more doctor attributes.[8]

Second, the data indicate that without quality signals, black and Asian doctors suffer a significant willingness-to-pay penalty about two to three times larger than the willingness-to-pay penalty for 1-star (lowest quality level) relative to 5-star doctors (highest quality level). Drawing on standard models of discrimination (Aigner and Cain (1977); Bohren et al. (2019); Fang and Moro (2011)), I show how risk-averse customers holding inaccurate statistical models of doctor quality in relation to race can account for this empirical observation. I also discuss how the exact breakdown between taste-based and statistical discrimination might depend heavily on modeling assumptions. It is clear, however, that inaccurate statistical discrimination plays a prominent role in the observed empirical

[8]The lack of a difference in willingness-to-pay for various longer travel distances between the treatment and control groups offers the most cogent evidence for the absence of a distraction effect. I discuss this in Section 3.3.2 in more detail.

pattern.

To my knowledge, this study is the first to get inside the black box of customer discrimination toward healthcare providers and to shed light on the underlying mechanisms. Customer discrimination is potentially a factor for the persistence of major issues in the medical labor markets such as discriminatory employment practices, the leaky pipeline for minority medical students, and the wage gap between white and minority providers (Freeman et al. (2016); Barr et al. (2008); Courey (2020); McGregory Jr (2013)). The finding that signals of doctor quality for customers drastically lower discrimination rules out the hypothesis that the observed customer discrimination is mainly due to taste-based discrimination. More importantly, the finding also suggests provider quality information may be a potential solution to reduce labor market discrimination for minority doctors on average.

The present analysis also provides new information about the consumer behavior and preferences of the 30 million uninsured Americans (Gunja and Collins (2019)). This group consumes healthcare in widely disparate circumstances, and little has been done to understand the preferences that guide their choices. As far as I know, this is the first study that estimates demand for healthcare among uninsured Americans with respect to prices, travel distance, and provider attributes.

1.1 Related Literature

This paper builds on previous work on labor market discrimination, quality signals in market design, and the design of discrete choice experiments.

Customer discrimination in labor market settings outside of healthcare has been documented in a growing literature.[9] These papers look at observed labor market outcomes such as wages to infer customer discrimination. The wage discount estimates associated with non-white workers in these studies are similar to my estimates.[10] Unlike these pre-

[9]For example, see Kahn and Sherer (1988), Nardinelli and Simon (1990), Neumark et al. (1996), Holzer and Ihlanfeldt (1998), Leonard et al. (2010), Combes et al. (2016), Bar and Zussman (2017), Stone and Warren (1999), and Burdekin and Idson (1991).

[10]Customer preferences as a factor driving the racial wage gap is a small but growing literature. In an interesting study design, Nardinelli and Simon (1990) found that baseball cards of non-white players sell for about 10% less than the cards of white players of comparable ability. Some studies looked at the equilibrium impact of customer discrimination on prices and labor market outcomes: Kahn and Sherer (1988) find that customer demand contributed to black baseball players getting paid 20% less, Holzer and Ihlanfeldt (1998) find that black service employees are paid roughly 17%–26% less at establishments with only black customers than at those with only white customers (this could be evidence of customer

vious studies, this paper sheds light on customer discrimination in the healthcare labor market, an important sector that represents over 13% of the U.S. labor force or 19 million jobs accounting for $1.0 trillion in annual payroll in 2018 (Census-Bureau (2018)). Also, my randomized field experiment setting allows me to add to the literature that rely mostly on pre-existing observational data.

A body of previous experimental work on discrimination finds evidence that significant proportions of the discrimination observed in various product markets and non-market interactions[11] are statistical. The range of estimated market price differentials associated with minority sellers when there is statistical discrimination roughly brackets the estimates of willingness-to-pay differentials in the present paper.[12] By estimating willingness-to-pay directly rather than inferring it from market prices, I am able to do so unburdened with additional assumptions about market structure. This approach allows me to get directly at the demand function and preferences of customers.

A literature on quality signals and market design suggests that provision or restriction of information can affect the level of discrimination in various labor markets. For instance, there is empirical evidence that racial minorities were harmed when certain quality signals for hiring purposes were disallowed. "Ban the Box" policies that bar employers from requesting criminal records of prospective employees, for example, seem to have widened the black-white gap in callbacks from prospective employers (Agan and Starr (2018)).[13] My results suggest that part of the differential treatment is driven by individual customers and not just corporate policies or human resource agents who might have a lower stake

discrimination if white customers are more likely to discriminate or discriminate more against black service professional), and Bar and Zussman (2017) found that the price quotes obtained by Arab service contractors are about 17% lower that those obtained by their Jewish counterparts in Israel.

[11]For example, Doleac and Stein (2013), Ayres et al. (2015), Edelman et al. (2017), Bartik and Nelson (2019), Cui et al. (2020), and Laouénan and Rathelot (2020) also leveraged online market settings, varied perceived race via skin color on an image or name of seller, and estimated race differentials in the price obtained by the seller and the probability of the seller being chosen by the buyer. For non-market interactions, see Bertrand and Duflo (2017), Gneezy et al. (2012), or List and Gneezy (2014) for well documented field experimental results.

[12]Doleac and Stein (2013) found that iPods held by a hand with darker skin generated offer prices that were 11%–12% lower on Craig's List, while Ayres et al. (2015) found that baseball cards held by a hand with darker skin sold for 20% less on eBay. Edelman et al. (2017) and Laouénan and Rathelot (2020) both looked at online marketplaces for sellers of services (like this paper) and found discrimination against minority sellers to the tune of 3% lower charges, but the latter went a step further and showed that the entire price discount vanishes when quality signals are exogenously introduced.

[13]Other examples include employer credit check bans (Ballance et al. (2020); Bartik and Nelson (2019)) and pre-employment drug testing (Wozniak (2015)). What is less studied, though, is when the hiring decisions are made by individual customers, as is the case for many medical providers, lawyers, and other professional service providers.

in the decision.

VIC makes a methodological contribution to the literature on stated preference elicitation, which aims to overcome hypothetical bias not only by offering a menu of real choices that reflects hypothetical choices, but also by examining the real choices that result. In particular, it builds on the innovations in Sullivan (2021), Kessler et al. (2019), Elías et al. (2019), List et al. (2006), List and Lucking-Reiley (2002), and others. The most common way researchers incentivize a conjoint choice task is by promising to deliver the product choice from one of the menus at the end of the experiment and deducting the corresponding price of the product from the payment the subject received at the beginning of the experiment (e.g. List et al. (2006)). Kessler et al. (2019) adopted a related approach and incentivized their choice experiment by promising to match employers to a menu of actual resumes based on their evaluations of the hypothetical resumes. While some of these pioneering works have evaluated hypothetical bias of pure hypothetical choice experiments against incentivized experiments, they have not attempted to directly assess the incentive compatibility of the incentivized choice experiments themselves against actual consumer choice within the same choice experiment. Going beyond previous approaches, VIC incorporates analysis of the actual booking choice of buyers paying entirely out-of-pocket in the menu of actual doctors from a real shopping environment so that one can validate whether behavior in the incentivized conjoint setting maps well into real shopping behavior.

The paper is organized as follows. Section 2 describes the unique context that enables the VIC design and details the experiment and empirical strategy. Sections 3 and 4 describe the data and present the key results. Section 5 outlines a simple model that distinguishes between the various sources of discrimination and helps interpret the results. Section 6 discusses the implications of this study for the literature on healthcare consumerism, discrimination, and stated preference elicitation. Section 7 concludes.

2 Experimental Design

This section presents implementation of VIC combining the incentives and ecological validity of a real shopping environment that is familiar to its users with the control of a laboratory. Section 2.1 describes the context of a popular online platform for healthcare

shoppers and how active customers were recruited as subjects for this experiment. In section 2.2, I outline incentives for reporting preferences in the conjoint survey and cross validate the incentive compatibility of survey design directly within the same shopping instance using booking data. Section 2.3 describes the doctor attributes and clarifies the nature of the quality signal that defines the difference between the treatment and control conditions. The hypothetical doctor profiles and the choice sets are described in Section 2.4. Section 2.5 outlines the estimation approach.

2.1 Online Shoppers for Medical Procedures and the Market Setting

I deploy VIC in partnership with an online platform in the U.S. that links self-pay patients seeking non-emergency care with doctors and facilities, much the way Priceline or Hotwire connects travelers and hotels.[14] This platform charges its some 265,000 paid subscribing customers a nonrefundable $25 fee to put a procedure out for bid, this fee is in addition to the payment for the actual procedure paid directly to the providers. Doctors pay a fee of $50[15] to bid on one request for bids that come as a package deal that include the physician's fee, the facility fee, and the anesthesiologist's fee if applicable. The procedures most shopped for on this platform are colonoscopy and imaging (including MRI, CT scans etc.). Roughly six thousand providers bid regularly. Actual doctor identity and contact information are concealed from the customer as the customer evaluates bids that include price and zip code of the provider.

After clicking "Accept this Bid" for a specific doctor, the customer promptly receives the contact information for the doctor who submitted the bid, while the doctor receives the contact of the customer. The doctors have contractually committed to the upfront price equal to their accepted bid, and the customer will pay the doctor directly after booking with the doctor. The customer and doctor can arrange for the procedure to take place at a date acceptable to both, typically within six months. This platform facilitated more than 7,500 transactions resulting in over $70 million in charges in 2018. As far as I know, this market has not been explored by the economics and health policy literature

[14] Due to terms in the Data Use Agreement, the company's identity will remain private. The business model and all figures on prices, customer base, participating doctors, and transaction volumes are provided by the company and cross-checked by the author through news coverage published in Forbes, The Wall Street Journal, Washington Post, Fox News, NBC, CNN, ABC and CBS.

[15] The doctors can also opt to pay an annual $250 fee for multiple bids.

that mainly relies on administrative data covering insured but not self-pay patients.

The VIC experimental design builds on the typical shopping experience of the customers of this platform where they evaluate anonymized doctor options, to naturally introduce choice sets where features characterizing each doctor option are restricted. The VIC is offered as a pilot tool that is similar in many ways to the usual shopping experience but different in that there are hypothetical menus of doctors to evaluate before being offered an actual menu of doctors. Furthermore, the customers gain access to a larger set of the doctors than usually listed on the platform through the pilot as the doctors who can potentially be matched include those usually listed on the platform as well as those listed with two of the platform's competitors.

The platform sent invitation emails to its subscription base in two waves. In both waves, prospective patients were invited to use a pilot tool that gathers prices from hundreds of platform doctors as well as doctors from the platform's two main competitors to "Shop for Your Next Provider for Colonoscopy, MRI or Knee Replacement" and that "the use of this pilot tool will be a one-time-only offer for each customer....[and] we'll waive the $25 fee for the appointment you book through the pilot tool" (see Appendix Figure 11 for the censored version of this email). The first wave was early January 2021 and the second at the end of February 2021. A recruitment tweet with similar content was also sent out weekly between 1 December 2020 and 31 March 2021.[16]

2.2 Validated, Incentivized Conjoint Design

After being informed of the instructions (see Appendix Section A.3) and collecting basic information about the outside options and whether they had previously had a colonoscopy, each customer was presented with a series of menus of doctors including a "None of the above" option. This is called choice-based conjoint elicitation (Ben-Akiva et al. (2019), McFadden (2017)). Each doctor offered in each menu would be described in terms of price and levels of attributes. Customers were asked to choose their most preferred option in each menu. As illustrated in Figure 1, customers were offered menus of doctors

[16]Customers can only participate once in the pilot (experiment), and each customer participates independently without the ability to observe or influenced by what other customers are doing on the pilot. I also used location and background information from the subscription database (where possible) to validate that no co-located customers have participated together. In the analysis, I assume that the Stable Unit Treatment Value Assumption (SUTVA) is met. In other words, I assume no-interference between customers, and that I can model a customer's treatment status and not be concerned with the treatment status of every customer in the sample.

who provide colonoscopies, with each doctor anomymized to resemble the platform's typical shopping experience but described in terms of price, travel distance from the customer's zip code, and a profile of the doctor including a blurred picture and gender.[17] The information provided for each doctor option is presented in a layout that would be familiar to the usual customer.

Subjects are asked to evaluate menus of hypothetical options with the understanding that the options are hypothetical and that more accurate evaluations will maximize the value of their participation. In my experiment, each customer evaluates fourteen menus with five doctor options and a "None of the above" option. Their participation incentive is that they will be matched to a specific menu of ten real booking options out of a pool of doctors aggregated in the partnering platform site as well as its two main competitors (more than doubling the usual number of doctors available to the customer). Eight of the options are doctors with the highest predicted choice probability based on particular individual preferences parameters estimated from the choice-based conjoint analysis data with mixed logit using Markov Chain Monte Carlo Hierarchical Bayes (MCMC HB) estimation (Ben-Akiva et al. (2019)). The remaining two options are randomly drawn from the remaining 331 possible doctors who offer colonoscopy services (Figure 2). By analyzing the actual booking choice in this menu of ten doctors, I can validate whether behavior in the incentivized conjoint setting maps well into real behavior.

A shortcoming of conjoint methods is that they are less reliable when the items are unfamiliar or incompletely described (McFadden (2017)). A benefit of my experimental setting is customer familiarity with the shopping experience during the experiment. Doctor information provided as customer evaluate options in the "pilot" is similar to the doctor information typically provided when customers shop on this platform. This makes the present conjoint analysis more reliable.

2.3 Doctor Attributes

Five attributes characterize each doctor option. The doctor options vary in race, gender, price, travel distance from the customer zip code, and quality. Race and gender are indicated by a blurry profile picture,[18] discussed in detail in Subsection 2.3.1 below.

[17] This is the baseline configuration. In the treatment group, quality "stars" are also provided. See below.

[18] Gender is also explicitly stated in the doctor profiles.

Each attribute of the doctors is selected from a realistic range based on real doctors listed on the platform and its competitor sites (see Appendix Table 3 for details on the attribute levels).

In the hypothetical doctor options, prices can take on one of ten evenly-spaced levels between \$1704.00 and \$2541.60, or the 10^{th} and 90^{th} percentiles of the actual prices for a colonoscopy from the full set of doctors on the platform and the two competitor sites. The customers are told (truthfully) that "[t]his price includes a screening or diagnostic colonoscopy with or without specimens/polyps removal by biopsy or brushing. Fees for facility and physician are include in the price. TRAVEL costs are NOT included in this price" (see Figure 1). Distance from the customer zip code can take on any of five levels: "0-10 miles", "10-50 miles", "50-100 miles", "100-250 miles," or "More than 250 miles." Finally, gender can be male or female while the profile picture has a blurred photo suggestive of either white, black, or Asian races.

2.3.1 Indicating Doctor Race

Race is the central characteristic of interest in the present paper. I manipulate perceptions of the doctor race by using profile pictures that clearly indicate the race of the hypothetical doctor. The resolution of each profile picture is reduced so that the doctor's race shows through but their attractiveness and other features are obscured (similar to Fong and Luttmer (2009)).[19]

To further ensure that detected outcome differences are due to race, I indirectly verify whether profile pictures in the VIC are similar in terms of attractiveness and other features. Besides blurring the photos, I assess the underlying unblurred photos' attractiveness. I use a popular facial symmetry detection algorithm deployed in previous publications evaluating attractiveness and economic outcomes (Dietl et al. (2020)).[20] I found no significant difference in attractiveness estimated by these measures across races (white,

[19]Recall that doctors are anonymized in their bids in the default shopping experience on this platform, and so the customers are reminded that "[w]e do not share detailed provider information until you have booked your appointment, therefore: We've blurred the details of each provider... Instead of the actual photo of each provider, our staff matched every one of the hundreds of actual providers to one of 24 blurred provider pictures that looks most like them" before they completed the conjoint survey and before they made the final actual booking (see Appendix Section A.4).

[20]I also rated all the profile pictures for "attractiveness" using a deep neural network based on a pre-trained ResNet-50 architecture (He et al. (2016)) that is fine-tuned with a dataset of hand-labeled photos. The results are similar to the validation using facial symmetry. Facial symmetry is used for validating similarity in attractiveness across races as it is a better proxy for racial-neutral attractiveness.

black, and Asian).[21]

2.3.2 Treatment versus Control: With versus Without Quality Signal

The last doctor attribute of quality, "stars," plays a key role for the test of statistical discrimination. Customers are randomized into the treatment or the control group. The groups differ in attributes they see in doctor profiles. While the control group sees doctor price, distance, gender, and profile picture as in Figure 1, the treatment group also sees stars that indicate the quality of the doctor in each profile (see Figure 3 for a side-by-side comparison). These stars range from 1 to 5 with 5 stars being the highest quality.

Customers in the treatment group see an additional informational screen explaining the stars. This screen explains that doctors with higher ratings tend to have better outcomes (fewer complications and lower patient mortality) and shows a table of ratings and other indicators of quality such as years of experience, graduation from a top 20 medical school, board certification, and customer reviews from three popular sites (Healthgrades, WebMD and Vitals). In the pool of real doctors 38% with 5 stars graduated from a top 20 medical school, and average 30 years of experience and a 4.7/5.0 patient review score from popular sites. In contrast, none of the doctors with 1 star graduated from a top 20 medical school, and they average less than 20 years of experience and a 2.2/5.0 patient review score. After these customers made their choices for all menus, I surveyed the subjects' perception of quality score by asking them "[t]o what extent do you agree or disagree with the following statement: "A provider with a higher Comprehensive Quality Score is a better provider than a provider with a lower Comprehensive Quality Score"." The subjects responded to this question using a Likert scale (1 to 5).

2.4 Choice Sets Creation and Variation

The present implementation of VIC asked each customer to evaluate 14 different menus of doctor options varying multiple doctor attributes to allow for good statistical identification of the valuation of separate attributes.[22] Conjoint designs that allow considerable linearly independent variation in the levels of different attributes and a considerable span

[21] Regression results from this validation technique is shown in Appendix Table 4

[22] Previous choice-based conjoint surveys used by rigorous studies in economics have presented subjects with 15 or fewer menus. For instance, the example in Ben-Akiva et al. (2019) gathered data with 11 menus per subject.

of attribute levels are well documented in various textbooks on conjoint analysis (Rao et al. (2014), Raghavarao et al. (2010)). I adopted an off-the-shelf fractional factorial choice design to optimize balance, overlap, and other characteristics. See Data Appendix for the order of questions for each respondent, the order of alternatives in each choice set for each respondent, and the placement of respondents into blocks in this paper. While price is always the last attribute, the order of the other attributes presented for doctor profiles are randomized across customers.

I check whether balance is achieved across the exogenously varied doctor attribute levels between the treatment and control groups. The p-values for the test of equality of the levels for each one of the hypothetical doctor attributes (race, sex, price, distance) between the treatment and control groups are 0.9 or higher, consistent with a properly executed conjoint design (Rao et al. (2014)).[23]

2.5 Discrete Choice Model

The goal of the empirical analysis is to estimate preferences using conjoint data. A generalized multinomial logit model with mixed (random) parameters and heteroskedasticity is estimated to allow the coefficients to vary across customers and exploit the information from the same individual making a particular sequence of choices out of the fourteen menus.

In the present discrete choice model, the customer maximizes a linear indirect utility function. Customer i's utility from choosing doctor d is

$$U_{id} = \delta_{id} + \eta_i(E_i - p_d) + e_{id} \tag{1}$$

where δ_{id} is marginal utility i gets from doctor d, η_i is marginal utility i gets from money (the price coefficient), E_i is the upper-limit of expenditure for the procedure in the allotted budget, and p_d is the price of doctor d. The term e_{id} is an unobserved random component of utility and is assumed to be distributed according to a type-1 extreme value distribution.

[23] See Appendix Table 5 for the actual table of p-values.

Customer i's utility from choosing the outside option, labelled as option 0, is

$$U_{i0} = \eta_i E_i + e_{i0}. \tag{2}$$

All customers solve a straightforward utility maximization problem: customer i chooses doctor d if and only if,

$$U_{id} \geq U_{id'}, \forall d' \in D \cup \{0\} \tag{3}$$

where D is the set of doctors.

Assuming distributional support for the error term on the real line then allows this model to rationalize any pattern of customer choices in the conjoint and obtain analytical expressions for individual choice probabilities of customers. Integrating over regions of the error space that coincide with a choice having the highest utility, I obtain choice probability of customer i choosing doctor d with the familiar multinomial logit form:

$$P_i(d) = \frac{\exp[\delta_{id} - \eta_i p_d]}{1 + \sum_j \exp[\delta_{ij} - \eta_i p_j]} \tag{4}$$

I further assume that marginal utility δ_{id} is a linear function of doctor attributes:

$$\delta_{id} = \mathbf{a}_d^{'} \lambda_i \tag{5}$$

where $\mathbf{a}_d$ denotes the vector of attributes of doctor d including race, gender, distance, and quality (for treatment group) and λ_i is the vector of customer specific attribute coefficients λ_{ia}.

The main empirical results and hypothesis testing are based on maximum simulated likelihood estimators (Ben-Akiva et al. (2019)).[24] More estimation details are presented in the Appendix (Section A.6).

[24]To allow for segmentation of customers by their preferences and allow more parameters (attributes and levels) to be estimated with smaller amounts of data collected from each customer, I also proceed with hierarchical Bayes estimation as is commonly practiced in choice-based conjoint data analysis in the marketing literature (Ben-Akiva et al. (2019), Allenby and Rossi (2006)). Recall that hierarchical Bayes estimation was also used to estimate the individual level parameters to generate match scores to identify the 10 real doctor options for customers on an on-going basis as the experiment proceeded.

2.5.1 Estimating Willingness-to-pay and Evaluating the Impact of Treatment

Customer utility in willingness-to-pay or money metric form is estimated to obtain more interpretable and comparable estimates for various doctor attributes (McFadden (2017)).

I estimate willingness-to-pay (WTP):

$$\mathbb{E}_i[WTP_{ia}] = \frac{\mathbb{E}_i[\lambda_{ia}]}{\eta} \tag{6}$$

where WTP_{ia} denotes the willingness-to-pay for attribute a, $\mathbb{E}_i[\lambda_{ia}]$ is the estimated coefficient from the marginal utility function estimation in (5), and η is the price coefficient. The main results are estimates of this metric.

To evaluate the differences of the effect on customer choice of the same attribute (e.g. black doctor) between the treatment and control group, slightly different models for the treatment and control groups are estimated with the former having additional covariates. The effects are translated into willingness-to-pay from the logit coefficients as shown above in Equation 6. The estimates for the impact of an attribute between the treatment and control group are evaluated with a test statistic:

$$t = \frac{WTP_a^{Control} - WTP_a^{Treatment}}{\sqrt{(SE[WTP_a^{Control}])^2 + (SE[WTP_a^{Treatment}])^2}} \tag{7}$$

where $SE[WTP_a^T]$ is the standard error of WTP_a estimated from the model for treatment condition T.

3 Data and Results

3.1 Sample Population and Check for Balance

The 229 customers shopping for a colonoscopy used the pilot tool between 1 December 2020 and 31 March 2021, with 224 completing the conjoint survey. The randomization approach is a simple Bernoulli Trial where the unit-level probability of treatment is 0.5.[25]

[25]Think of this as the case where as customers arrive at the experiment, a coin is flipped and a "heads" outcome sends the customer to the control group and a "tails" outcome sends the customer to the treatment group.

Of those 224 in the sample,[26] 104 were randomly assigned to the treatment condition and responded to menus with options that include quality "stars" while the other 120 customers were assigned the control condition.

Table 1 presents summary statistics, with the full sample in columns 1 and 2. On average, customers are clustered in the age range between 45 and 64 and much more likely to be uninsured than the overall U.S. population. The U.S. uninsured population is about 10% as opposed to 90% in this sample. The six states most represented in the sample are Texas, Illinois, Washington, North Carolina, and Florida-Indiana in a tie. The full sample population stated that they will be paying cash out of pocket for the colonoscopy, and most have had a colonoscopy in the past 10 years. On average, these customers also looked at more than four colonoscopy options other than the platform, and have an outside option with a price that is within the range of actual colonoscopy prices offered by doctors included in this experiment.

Health economics and health policy research have focused almost entirely on the insured population given the availability of administrative data, leaving about 30 million Americans understudied. By capturing a largely uninsured population, this paper provides a contribution that complements the rest of the literature. Furthermore, customer decision making could potentially be more fully expressed for uninsured cash-paying patients without the constraints and influences of co-pay, co-insurance, and a provider network determined by an insurance company.

The last column of table 1 investigates the balance between treatment and control groups. The control and treatment groups are very well balanced on all measured dimensions, as one would expect from the random assignment to treatments.

3.2 Racial Discrimination and Other Factors Influencing Customer Choice

Table 2 presents the main results. Using the approach outlined in Section 2.5.1 to estimate the control group, money metric willingness-to-pay is lower by 12.7% of the average colonoscopy price for black than white doctors, and lower by 8.7% of the average

[26]This paper uses data from colonoscopy shoppers. The broader experiment included 407 customers who were shopping to get a colonoscopy, an MRI or a knee replacement. 177 shopped for an MRI and 1 shopped for a knee replacement. The MRI data is used for a different paper and also not suitable for the purposes of this paper since the MRI providers are facilities not individuals with clearly identified race and demographics. I was unable to include the shopper for knee replacement since there is only one.

colonoscopy price for Asian than white doctors. In other words, in the absence of any information about doctor quality, customers were willing to pay $270 more to have their colonoscopy performed by a white doctor than a black one, and $186 more for a white doctor than for an Asian one. These results are presented in the first column. Both of these racial differentials are statistically significant (p-values < 0.001). Such statistical results have post-study probability (PSP) of > 0.95 if the fraction of true associations in the field of interest for this paper is ≥ 0.1.[27] Given that perceived race is independently varied across doctor options and menus, the empirical evidence suggests that race causally affects the likelihood of being chosen by customers. In other words, in an environment with no direct quality signals, there is customer discrimination in this market for doctors.

Without quality signals, a female doctor suffers a statistically significant $16 willingness-to-pay penalty for gender. This is significant at conventional levels.

Working down column 1 of Table 2, the willingness-to-pay estimates suggest that travel distance also affects customer choice. Overall, willingness-to-pay drops or becomes more negative as the doctor is further away from the customer. The willingness-to-pay for a doctor 10-50 miles away relative to a doctor who is 0-10 miles away is lower by $37, while having to potentially travel 50-100 miles lowers willingness-to-pay $118. Likewise, if the doctor is 100-250 miles or more than 250 miles away from the customer, the willingness-to-pay figures drop by $263 and $589 relative to a doctor who is 0-10 miles away. All distances have willingness-to-pay differences relative to 0-10 miles of distance with statistical significance of p-value less than 0.001. Some of these coefficients might be biased upwards due to the coincidence of the experiment with the COVID-19 pandemic. This might have especially enlarged the coefficient on the "More than 250 miles" as air

[27] List (2021) proposed that studies may be assessed on the basis of their positive predictive value, as measured by the post-study probability (PSP) that a research finding that has achieved formal statistical significance is in fact true. The basic ingredients of this metric are the Type I (α) and II (β) error rates, and the fraction of true associations among all associations tested in a given field (π). The PSP is given by:

$$PSP = \frac{(1-\beta)\pi}{(1-\beta)\pi + \alpha(1-\pi)}. \tag{8}$$

I report all PSP results in this paper with the level of power fixed at 80% or $\beta = 0.2$.

I am not aware of good data that provides the prior on whether the association between race and willingness-to-pay in absence of quality signals is true. I offer an indirect reference point for this prior by obtaining expert forecasts of this particular empirical result (discrimination against black and Asian doctors) using the Social Science Prediction Platform (DellaVigna et al. (2019)). Based on the responses of 14 economists, 100% predicted that there will be a discrimination effect strictly larger than zero for black race for doctors as I find here; 50% predicted that there will be a discrimination effect strictly larger than zero for Asian race for doctors as I find here.

travel was less preferred than during normal times.[28] External validity of the results on travel distance require further evidence.

To put the size of the discrimination into perspective, the willingness-to-pay penalty for black doctors is approximately equal to the size of the willingness-to-pay penalty for having to travel 100-250 miles without direct quality signal, and the willingness-to-pay penalty for Asian doctors somewhere between 50-100 to 100-250 miles without quality signal. This observation is confirmed by an F-test of equality between the coefficients for black race and 100-250 miles distance.

3.3 What Difference does a Quality Signal Make?

Figure 4 presents the willingness-to-pay for different levels of stars in the treatment group. The pattern is what one would expect to see if customers value quality represented by stars. The willingness-to-pay increases by $72 when a doctor moves from one to five stars, while the jump from one to three stars is $16. It is worth noting that, on average for the population exposed to the quality signal, the hypothesis that the willingness-to-pay for a five and a four star doctor are the same is not rejected. At other levels, more stars generally increase willingness-to-pay. The relationship between stars and willingness-to-pay seems to be nonlinear in Figure 4. Finally, the average Likert scale rating from 5 ("Strongly Agree") to 1 ("Strongly Disagree") is 4 ("Agree") for the self-reported exit survey question "To what extent do you agree or disagree with the following statement: 'A provider with a higher Comprehensive Quality Score is a better provider than a provider with a lower Comprehensive Quality Score'," suggesting the average customer in the treatment group agrees that stars indicate better providers.

Column 2 of Table 2 presents the willingness-to-pay estimates for each of the attributes for the control group. The coefficients for various travel distances are statistically similar in magnitudes with those recovered from the control group who did not observe the quality signal.[29] In fact, evaluating the test statistics for each of these comparisons as described in equation (7), the hypotheses that each of the pairs of coefficients for each attribute above are the same within each pair except 10–50 miles is not rejected at all

[28] Note that it can also be argued that cheaper flights and better hotel availability during the COVID-19 pandemic could have lowered travel expenses.

[29] The difference in the coefficients for 10–50 miles is not significant at the 5% level but significant at the 10% level.

conventional levels.

While willingness-to-pay for some attributes is similar with or without the introduction of a quality signal, this does not hold for the coefficients capturing the racial discrimination against black and Asian doctors (as well as discrimination against female doctors[30]). Instead of being $>$ \$180 and significantly negative as shown on column 1 of Table 2, the willingness to pay for black and Asian doctors are now \$26 and \$24 lower than white doctors, 10% and 13% of the estimate before. The test statistic from (7) applied to the willingness-to-pay for black and Asian races yields p-values < 0.001. The post-study probability for the drop in racial gap being true is > 0.95 if the fraction of true associations among all tested associations in studies looking at whether signal provision reduces statistical discrimination is ≥ 0.1,[31] suggesting that racial discrimination is significantly attenuated when a quality signal is introduced relative to when there are no direct quality signals.

This analysis suggests that customers who lack individualized information about quality might be relying on race-based assumptions about doctor quality. Racial discrimination decreased substantially for both black and Asian doctors when the quality signal was introduced. A surprising result is that the reduction in willingness-to-pay for both black and Asian relative to white doctors is larger than dropping from a 5 star doctor to a 1 star doctor. The test statistic from (7) rejects the hypotheses that the penalty in willingness-to-pay for being a black or Asian doctor is the same as the penalty in willingness-to-pay for a doctor dropping from 5 stars to 1 star. This result compels evaluating forms of statistical discrimination beyond the simplest in subsection 5.2.[32]

3.3.1 Heterogeneity across Customer Demographics

The present findings may change based on characteristics of the customer. If discrimination is only driven by homophily or in-group bias, the quality signals should not have

[30]With the quality signal, the willingness-to-pay for female doctors are statistically indistinguishable from the that for male doctors.

[31]I am not aware of good data that provides the prior on whether the drop in willingness-to-pay with the introduction of quality signals is true. I offer an indirect reference point for this prior by obtaining expert forecasts of this particular empirical result using the Social Science Prediction Platform (DellaVigna et al. (2019)). Based on the responses of 14 economists, 78.6% predicted a non-zero drop in willingness-to-pay for black doctors when the quality signal is introduced as I find here; 42.9% predicted a non-zero drop in willingness-to-pay for Asian doctors when the quality signal is introduced as I find here.

[32]This is an approach related to the discussion in Bohren et al. (2019).

lowered the willingness-to-pay penalty for minority doctors. Moreover, it is interesting to investigate whether statistical discrimination differs by racial concordance between customer and doctor. Figure 6 presents evidence of such heterogeneity.[33] While black customers also discriminate against black doctors to the tune of a $78 willingness-to-pay penalty when there is no quality signal, such a figure is much smaller than the same figure for white customers against black doctors, which is more than four times larger. Looking at black patient willingness-to-pay for a black doctor when quality signals are introduced, there is a slightly positive willingness-to-pay edge for black over white doctors. This can be evidence for a preference for racial concordance that might include benefits like improved trust (Alsan et al. (2019), Alsan and Wanamaker (2018), Alsan and Eichmeyer (2021), Idan et al. (2020)). The small sub-sample of black customers limits precision of these estimates making the evidence of preference for racial concordance only suggestive.

As with racial concordance, gender concordance can be a favorable characteristic in the eyes of a customer. Figure 7 shows the willingness-to-pay for a female doctor relative to a male doctor, broken down by whether the customers themselves are female or male. Female customers are willing to pay $85 (no signal) and $39 (with signal) more for a female doctor while male customers are willing to pay $82 (no signal) and $37 (with signal) more for a male doctor. These figures are approximately 4% and 2% of the average price and are significant at conventional levels.

3.3.2 Could Customers just be Distracted by the Presence of One More Doctor Attribute?

Recall that the treatment group is exposed to 5 doctor attributes while the control group is exposed to only 4 (no quality stars). The lower sensitivity to race could be due to competition for customer attention by the higher number of attributes. One way to evaluate the concerns with distraction is to look at whether coefficients for variables unlikely used as proxies for quality differ between treatment and control groups. Assuming customers do not use the furthest travel distance buckets as a proxy for doctor quality in the control group, comparisons of the coefficients on travel distances between treatment and control can be used to test whether distraction was significant. In the exit survey, customers

[33] I look at heterogeneity by patient race by evaluating the white and black sub-populations only. The segments of Asian patients, Native American patients and patients of other races do not have enough observations in my sample to yield precise enough estimates to draw reliable insights.

subjectively estimated average doctor quality of doctors based on their distances and I found that distance is not correlated with customers' subjective assessment of provider quality. The pairwise differences between the average subjective quality for each pair of distance levels have p-values > 0.9. Figure 9 shows that the coefficients for "50-100 miles", "100-250 miles" and "More than 250 miles" are all similar for the treatment when 5 doctor attributes are shown and control when 4 doctor attributes are shown, lending support to the hypothesis that the drop in the racial willingness-to-pay estimates when a quality signal was introduced was not driven by distraction by more attributes.

Another way to evaluate this issue is to look at customers who are exposed to 5 attributes but did not intake information from the quality signal itself to see whether they are just like the control condition. Figure 8 shows how the main results from Figure 5 varies by customers who selected "Disagree" or "Strongly disagree" to the statement "A provider with a higher Comprehensive Quality Score is a better provider than a provider with a lower Comprehensive Quality Score" in the exit survey. Those who chose "Disagree" and "Strongly disagree" are non-believers of the signals and the rest "believers". For "believers" the results are similar to the entire population (See Figure 5). However, for "non-believers" the willingness-to-pay penalty did not go down to 1%. Instead the penalty for black doctors went down to 7% of average price while the penalty for Asian doctors went down to 3%.

It is hard to conclude whether "non-believers" were truly not affected by the information embedded in the stars.[34] It is difficult to accurately disentangle how much of the drop in willingness-to-pay penalty for black and Asian doctors is due to the presence of a distraction from an additional attribute in doctor profiles. Nonetheless, the two main empirical results that there is discrimination towards black and Asian doctors and that most of it is statistical discrimination, are still true comparing within the treatment group between "believers" and "non-believers". The penalties for black and Asian doctors are statistically significant at conventional levels, as are the differences of these penalties between the "believers" and "non-believers." The willingness-to-pay penalty for black doctors for "believers" is 8.9% of that for "non-believers" and the willingness-to-pay penalty for Asian doctors for "believers" is 26.8% of that for "non-believers."

[34]In Appendix Table 8, one can see that "non-believers" place a statistically significantly higher willingness-to-pay for doctors with 5 stars relative to doctors with 1 or 2 stars despite the fact that the estimates are very imprecisely estimated.

4 Validation of VIC Estimates

This section compares predicted choice with actual choice to validate the VIC study design.

Recall that a menu of actual doctors was presented to each customer after they responded to the 14 hypothetical choice tasks as part of the VIC. Every day during the experimental period, a new mixed logit model (including all features described in Section 2.3 except perceived race variable) was estimated[35] to obtain a new set of individual-level parameters. Every doctor is scored based on choice probability for every customer who completed the conjoint survey that day with the model of the day, and the 10 options of real doctors are just the 8 with the highest choice probability plus 2 randomly drawn from the remaining options without replacement. After the experiment was closed, a new mixed logit model was estimated including all variables mentioned in Section 2.3 with hierarchical Bayes and individual-specific choice probabilities calculated for each doctor in the menus offered to every customer ("an ex post model"). A rank is assigned for each doctor within each customer menu based on their ex post predicted choice probability specific to that customer.

Figure 10 shows the percentage of actual booking choices by the 188 customers who booked with one of the 10 doctors matched to them out of 224 who completed the conjoint by that option rank within their menu. It is not surprising that most bookings are with doctors who had the highest predicted (fitted) choice probability rank while none of the bookings are with the randomly chosen options that tend to have much lower predicted choice probability than the other 8. The top-3 accuracy score and top-5 accuracy score for the fitted choice probability are 82% and 93% respectively. In general, actual choice of an option is more likely the higher ranked the option is within its menu.

The discrete choice model estimated with the VIC seems able to predict actual choice behavior consistent with the VIC being incentive-compatible.

[35] Estimations were done with hierarchical Bayes using the Allenby-Train procedure (Allenby (1997), Train (2009), Ben-Akiva et al. (2019)) with 8,000 burn-in iterations and 8,000 iterations after burn-in.

5 Interpretation of Results

This section sets up a theoretical framework to link the empirical results to models of discrimination. To illustrate the effects of introducing a quality signal, the decision problem of customers who may be risk averse is evaluated.[36] One particular model may not explain the results. Rather, assumptions about customer beliefs and preferences leading to behavior consistent with a few stylized facts documented in the previous sections are developed.

Two key stylized facts will be addressed in this section. First, the data indicate that without quality signals, black and Asian doctors suffer a significant willingness-to-pay penalty about two to three times larger than the willingness-to-pay penalty for 1-star relative to 5-star doctors (see Table 2). Second, the willingness-to-pay penalty for black and Asian doctors drops by about $80\% - 90\%$ when a quality signal is introduced (see Figure 5).

5.1 A Model of Discrimination

I first present a different model from the preceding discrete choice model in order to facilitate a clear interpretation. Specifically, I assume the customer observes the group identity (i.e. race) of doctors $j \in \{B, W\}$. Doctor quality q is assumed to equal the value of the doctor's marginal product to the customer, and it is drawn from a normal distribution[37] $N(\mu_j, \sigma_j^2)$. I consider noisy unbiased signals of doctor quality, $\theta = q + \epsilon$, where ϵ is a zero-mean error that is normally distributed according to $N(0, \sigma_{\epsilon_j}^2)$.

Following Aigner and Cain (1977) and Fang and Moro (2011), I assume that customers are risk averse with (money-metric) utility:

$$V_c(q, j) = \alpha_j - \beta \exp(-\gamma q) + \upsilon_c, \tag{9}$$

[36] Theoretical models for statistical discrimination that incorporate risk-averse employers have been proposed (e.g., Aigner and Cain (1977)). However, the labor economics literature has favored risk-neutral employers as a more realistic model when firms are doing the hiring (Fang and Moro (2011)). The customers making the decisions in the present paper are individual prospective patients, not firms. Therefore, I do not assume away risk aversion for their utility functions.

[37] Typical (accurate) statistical discrimination models assume that the customer knows this true underlying distribution of quality by identity group. As I will discuss in Section 5.1.1, the assumption that customers hold inaccurate statistical models will be necessary to reconcile the models with the key stylized facts listed at the beginning of Section 5

where υ_c is a utility shifter idiosyncratic to customer c (iid with mean 0) and α_j captures the preference for doctors from group j independent of quality. The difference $(\alpha_W - \alpha_B)$ is referred to as taste-based discrimination against doctors of group B relative to group W. Expected utility from booking a doctor with signal θ is given by

$$\mathbb{E}[V_c(q,j)|\theta,j] = \alpha_j - \beta \exp[-\gamma\, \mathbb{E}(q|\theta,j) + \frac{\gamma}{2} Var(q|\theta,j)]. \tag{10}$$

A customer's willingness-to-pay for a doctor is equal to the expected utility conditional on the quality signal and the group identity. Quality and signal are jointly normally distributed. Using the properties of the conditional normal distribution,[38] the willingness-to-pay conditional on the signal, θ, and group identity of the doctor, j, is

$$WTP(\theta,j) = \mathbb{E}[V_c(q,j)|\theta,j] = \alpha_j - \beta \exp[-\gamma(\frac{\sigma_j^2}{\sigma_j^2+\sigma_{\epsilon_j}^2}\theta + \frac{\sigma_{\epsilon_j}^2}{\sigma_j^2+\sigma_{\epsilon_j}^2}\mu_j) + \frac{\gamma}{2}\frac{\sigma_j^2\sigma_{\epsilon_j}^2}{\sigma_j^2+\sigma_{\epsilon_j}^2})]. \tag{13}$$

5.1.1 No signal versus quality signals

Consider the customer decision problem under the control and treatment conditions. Recall that signals of doctor quality are unbiased with a noise term normally distributed according to $N(0,\sigma_{\epsilon_j}^2)$. Under the control condition with no quality signals sent, one can think of the customer as having a signal where $\sigma_{\epsilon_j}^2 \to \infty$. It follows that willingness-to-pay can be written as

$$\lim_{\sigma_{\epsilon_j}^2 \to \infty} WTP(\theta,j) = \alpha_j - \beta \exp[-\gamma\mu_j + \frac{\gamma}{2}\sigma_j^2]. \tag{14}$$

Recall that μ_j and σ_j^2 are the mean and variance of the underlying distribution of doctor quality for group j. If a customer thinks that the two groups have the same underlying quality distribution, any willingness-to-pay differences of that customer that are observed

[38] If the quality and signal are jointly normally distributed, the conditional distribution of q given θ is normal with mean equal to

$$\mathbb{E}(q|\theta,j) = \frac{\sigma_j^2}{\sigma_j^2+\sigma_{\epsilon_j}^2}\theta + \frac{\sigma_{\epsilon_j}^2}{\sigma_j^2+\sigma_{\epsilon_j}^2}\mu_j \tag{11}$$

and variance equal to

$$Var(q|\theta,j) = \frac{\sigma_j^2\sigma_{\epsilon_j}^2}{\sigma_j^2+\sigma_{\epsilon_j}^2}. \tag{12}$$

(DeGroot (2005)).

under the control condition will be attributable to taste-based discrimination ($\alpha_W - \alpha_B$) within the context of this paper (as the second term in (14) will be differenced out).

On the other hand, customers can observe a quality signal under the treatment condition. This is modeled as a signal with a finite $\sigma^2_{\epsilon_j}$. For the sake of simplicity and a clear contrast with the control group, I assume that under the treatment condition $\sigma^2_{\epsilon_j} = 0$ from the customer's perspective. I will discuss some implications of $\sigma^2_{\epsilon_j}$ being finite but not equal to zero below. For now, under the treatment condition where $\sigma^2_{\epsilon_j} = 0$, willingness-to-pay can be written as:

$$\lim_{\sigma^2_{\epsilon_j} \to 0} WTP(\theta, j) = \alpha_j - \beta \exp[-\gamma\theta], \tag{15}$$

where θ is the quality signal (number of stars). Conditional on quality signal θ, when $\sigma^2_{\epsilon_j} = 0$, any willingness-to-pay differences of the customer that are observed under the treatment condition will be attributable to taste-based discrimination ($\alpha_W - \alpha_B$).

Finally, using this set-up and the stylized fact that black and Asian doctors suffer a significant willingness-to-pay penalty about two to three times larger than the willingness-to-pay penalty for 1-star relative to 5-star doctors without quality signals, the assumption of risk aversion is justified. If I assume risk neutrality, where utility is a linear function of expected quality, the customer will have to believe that the unconditional mean quality of white doctors is higher than the unconditional mean quality of black and Asian doctors by a magnitude two to three times larger than the difference in quality between a 5-star and a 1-star doctor for the model to be consistent with data. Such customer priors about the unconditional mean quality of doctors of different racial groups depart substantially from the actual differences in mean of doctor quality by race (measured by stars) in our pool of doctors. The assumption of risk averse customers seems more credible.

5.2 Inaccurate Statistical Discrimination as Driver of Willingness-to-pay Penalty

Next, the sources of customer discrimination are unpacked with the theoretical framework above and the other stylized fact mentioned in the beginning of this Section.

Access to the underlying quality distributions by identity group for the doctors in this study allows me to learn the true distribution of quality for each identity group. The

mean qualities in terms of stars for white, black, and Asian doctors are 2.93, 2.91, and 2.92 with variances 1.55, 1.51, and 1.52. With the normality assumption, the underlying "true" distribution of the quality of doctors is similar across all racial groups in the pool of doctor options faced by the customers in this experiment. This means that the data suggest that $\mu_i = \mu_j$ and $\sigma_i = \sigma_j$ $\forall i, j \in \{white, black, Asian\}$.

Taste-based discrimination cannot be the whole story: If there is no statistical discrimination, willingness-to-pay differences between white and minority doctors in the control group (differences in willingness-to-pay in (14)) and in the treatment group (differences in (15)) would both equal $(\alpha_W - \alpha_B)$. However, the willingness-to-pay penalty for minority doctors in the control condition where no quality signal is provided is about eight to ten times larger than the willingness-to-pay penalty for minority doctors in the treatment condition where the quality signal is introduced (see Figure 5). Standard information manipulation experiments as in List (2004) have interpreted this stylized fact as evidence that most observed discrimination is statistical in nature.

What can one say about the nature of this statistical discrimination here? Many of the previous papers trying to parse the types of discrimination assume that decision makers hold accurate beliefs about the true distribution of quality. For instance, they might compare the hiring decisions to the true underlying distribution of the quality for these decisions (for example, Arnold et al. (2018)). Taking a similar approach, I should find that the willingness-to-pay differences between white and minority doctors in the control group and in the treatment group are both equal to $(\alpha_W - \alpha_B)$. The estimated willingness-to-pay penalties are not consistent with this equality. Because actual quality information for the doctors included in this study is available, the possibility that consumers held accurate beliefs about quality with race as a proxy could be ruled out.

Next, I examine whether inaccurate beliefs can explain the observed disparities in willingness-to-pay, a direction proposed by Bohren et al. (2019) and others.

Inaccurate statistical discrimination stems from incorrect beliefs held by the customers. While accurate statistical discrimination is sometimes referred to as "efficient discrimination" inaccurate statistical discrimination can give rise to inefficiencies that policymakers might want to address through information provision. In the theoretical framework, allowing for inaccurate beliefs is modeled as allowing customers to hold subjective beliefs $\hat{\mu}_j$ and $\hat{\sigma}_j$ about the mean and variance of doctor quality for group j

where $\hat{\mu_j}$ and $\hat{\sigma_j}$ are not required to equal μ_j and σ_j. In other words, customers are allowed to have a mis-specified model of the quality distribution. Consequently, the willingness-to-pay in the control condition where no quality signal is provided is now given by $\alpha_j - \beta \exp[-\gamma\hat{\mu_j} + \frac{\gamma}{2}\hat{\sigma_j}^2]$ instead of the expression in (14).

In the simple case outlined in the theoretical framework above, we still obtain willingness-to-pay penalty for doctors from group B relative to doctors from group W under the treatment condition (with quality signal) as $(\alpha_W - \alpha_B)$ from (15). With inaccurate beliefs, the willingness-to-pay penalty for doctors from group B relative to doctors from group W under the control condition is given by

$$\lim_{\sigma^2_{\epsilon_W} \to \infty} WTP(\theta, W) - \lim_{\sigma^2_{\epsilon_B} \to \infty} WTP(\theta, B) = \underbrace{(\alpha_W - \alpha_B)}_{\text{taste-based}} + \underbrace{\beta(\exp[-\gamma\hat{\mu_B} + \frac{\gamma}{2}\hat{\sigma_B}^2] - \exp[-\gamma\hat{\mu_W} + \frac{\gamma}{2}\hat{\sigma_W}^2])}_{\text{statistical}} \tag{16}$$

where the first term represents taste-based discrimination and the second term gives the statistical discrimination. While accurate beliefs of the groups having similar actual quality distribution require the statistical discrimination term to equal zero, we can get a statistical discrimination term that is positive and large relative to the taste-based discrimination term under inaccurate beliefs. If I assume inaccurate beliefs and if this statistical discrimination term is large enough (say 5 to 10 times larger than the first), it is possible to observe willingness-to-pay penalty for black and Asian doctors dropping by about $80 - 90\%$ when a quality signal is introduced. Using my simple theoretical framework, I can reconcile a model of customer discrimination largely driven by inaccurate beliefs rather than accurate statistical or preference-based sources with the data from this experiment.

5.3 Imperfect Signal Precision and Residual Willingness-to-pay Penalty for Minorities

Next, consider the implications of relaxing the assumption that the quality signal has perfect precision ($\epsilon_j^2 = 0$) under the treatment condition regarding ability to interpret the empirical estimates in this paper. To do this, consider the subjective beliefs about

signal precision denoted as $\hat{\sigma_{\epsilon_j}}^2$ for identity group j, that can differ by identity group of the doctor. The willingness-to-pay penalty for doctors from group B relative to doctors from group W under the treatment condition, where a quality signal is provided, is given by

$$
\begin{aligned}
WTP(\theta, W) - WTP(\theta, B) = & \underbrace{(\alpha_W - \alpha_B)}_{\text{taste-based}} \\
& \underbrace{\begin{aligned}
& + \beta\{\exp[-\gamma(\frac{\hat{\sigma_B}^2}{\hat{\sigma_B}^2 + \hat{\sigma_{\epsilon_B}}^2}\theta + \frac{\hat{\sigma_{\epsilon_B}}^2}{\hat{\sigma_B}^2 + \hat{\sigma_{\epsilon_B}}^2}\hat{\mu_B}) + \frac{\gamma}{2}\frac{\hat{\sigma_B}^2\hat{\sigma_{\epsilon_B}^2}}{\hat{\sigma_B}^2 + \hat{\sigma_{\epsilon_B}}^2})] \\
& - \exp[-\gamma(\frac{\hat{\sigma_W}^2}{\hat{\sigma_W}^2 + \hat{\sigma_{\epsilon_W}}^2}\theta + \frac{\hat{\sigma_{\epsilon_W}}^2}{\hat{\sigma_W}^2 + \hat{\sigma_{\epsilon_W}}^2}\hat{\mu_W}) + \frac{\gamma}{2}\frac{\hat{\sigma_W}^2\hat{\sigma_{\epsilon_W}}^2}{\hat{\sigma_W}^2 + \hat{\sigma_{\epsilon_W}}^2})]\}.
\end{aligned}}_{\text{statistical}}
\end{aligned}
\tag{17}
$$

In the previous subsection, the observed willingness-to-pay penalty for minority doctors in the treatment condition ("residual") is taste-based discrimination. While the observed discriminatory behavior was decomposed into taste-based discrimination and inaccurate statistical discrimination with perfect signal precision, such decomposition might not be viable without additional assumptions about unobservable customer beliefs.

If one assumes customers subjectively believe that minority doctors have lower mean quality ($\hat{\mu_B} < \hat{\mu_W}$), higher variability in quality ($\hat{\sigma_B} > \hat{\sigma_W}$) and weakly lower signal precision given the same quality scoring system ($\hat{\sigma_{\epsilon_B}}^2 \geq \hat{\sigma_{\epsilon_W}}^2$), the statistical discrimination term in (17) will be positive. In this case, it is possible that the willingness-to-pay penalty for minority doctors in both the treatment and control conditions are entirely due to statistical discrimination (the statistical discrimination terms of (17) and (16)) while taste-based discrimination is equal to zero ($\alpha_W - \alpha_B = 0$). On the other hand, if customers subjectively believe minority doctors have higher mean quality ($\hat{\mu_B} > \hat{\mu_W}$), lower variability in quality ($\hat{\sigma_B} < \hat{\sigma_W}$) and weakly higher signal precision given the same quality scoring system ($\hat{\sigma_{\epsilon_B}}^2 \leq \hat{\sigma_{\epsilon_W}}^2$), the statistical discrimination term in (17) will be negative with the statistical discrimination favoring the minority group when the imprecise signal is introduced. A negative and large enough statistical discrimination term can explain the stylized fact where the observed overall willingness-to-pay penalty for minority doctors is much smaller in the treatment condition compared to the control condition, even if the willingness-to-pay penalty in the control group is entirely accounted for by taste-

based discrimination. Although the former case is more reasonable given the previous literature on customer discrimination, it is clear that any other case with different mixtures of taste-based and statistical discrimination could be consistent with observed data in this experiment.[39] The data tell us: taste-based discrimination alone cannot be the whole story with statistical discrimination showing up prominently in either the control or treatment or both conditions taking on either sign. It is therefore worth noting that interpretation of the data will ultimately be tied to assumptions about signal precision and subjective beliefs.

While acknowledging the limitations described in this subsection, it is under the reasonable assumptions in Subsections 5.1 and 5.2 that the data is interpreted consistent with the claim that most of the observed discrimination represents statistical discrimination.

6 Discussion

6.1 Implications for Healthcare Labor Markets

Since Arrow (1963), economists have been aware consumers face difficulties determining the quality of services provided by doctors because of asymmetric information. In this paper, customers of healthcare are found to use race as a proxy for quality in light of asymmetric information. Most doctors in the U.S. are employed by hospitals and groups, not directly negotiating prices with customers. However, lower willingness-to-pay associated with minority doctors could make them less attractive to employers. Furthermore, we need to ensure minority doctor entry is not deterred in specialties where trust is important for better population health outcomes (Alsan et al. (2019)). Such specialties are also likely those in which customer discrimination can play the largest role as in primary care or oncology. Specialty choice is responsive to wages (Nicholson and Propper (2011)). Should the willingness-to-pay penalties translate into lower wages for minority doctors, diversity of the workforce in these specialties will be affected.

A key result in this paper is that the provision of a quality signal increases the relative willingness-to-pay for black and Asian doctors on average. While the magnitude of the willingness-to-pay penalties is similar to the raw wage gaps for these minority groups relative to the average white doctors, it is hard to directly map the impact of informa-

[39] Depending on the relative sizes of the statistical discrimination terms in Equations 16 and 17.

tion provision to actual effects on the racial wage gap or labor force participation by minorities in medicine due to the industrial organization of medicine and unknown general equilibrium effects. Nevertheless, the prominence of statistical discrimination found in this paper strongly suggests that the provision of provider quality information can effectively reduce much of the discrimination in healthcare labor markets by customers.[40]

While quality signals decrease discrimination to black and Asian doctors as a group by $80\%-90\%$, doctors within each racial group with poor quality ratings might find it harder to get customers at the same price as their peers with better quality ratings. And while the average quality scores of black and Asian doctors do not differ from white doctors for the labor pool in the marketplace of the present experiment, black and Asian doctor ratings could be biased downwards due to racial preferences from majority patients in the broader healthcare market (Cooper et al. (2003)). In this case, quality ratings might help perpetuate racial biases held by customers. What one can confidently conclude from this paper is that provision of unbiased quality signals can reduce racial disparities. Caution must be applied, however, when drawing welfare conclusions when adopting information provision as a policy.

The present results have implications for the general problems of customer discrimination and for discrimination in the healthcare sector. As more online marketplaces emerge for healthcare, clear market design implications are provided in this paper for platforms and publicly run marketplaces that aspire to provide a discrimination-free environment. Online platforms create new markets by eliminating search friction and facilitating transactions. Even online, service professionals face discrimination. While the study setting is a specialized market compared to the majority of insured healthcare customers, the unique nature of the setting enabled study design to causally investigate the important

[40] A few limitations should be noted. First, the customer sample here is not representative of the US population by age and insurance status, nor may it be representative of the uninsured as most people in that category are probably not savvy enough to use a tool like this. Second, preferences regarding colonoscopy doctors might be more sensitive than preferences concerning a less invasive kind of care. These limitations compel caution against over-generalization of the present results. Finally, the result that the willingness-to-pay penalty for black and Asian doctors are about two to three times larger than the willingness-to-pay gap between 1-star and 5-star doctors prevents explaining the results with traditional statistical discrimination models with risk neutral customers. As discussed in Section 5.3, I cannot definitively establish the respective roles of statistical and taste-based discrimination under the more flexible models. Skeptical readers might even claim that as VIC informs respondents that the menu options (including the race of the doctors) are hypothetical, one should not describe the observed results as evidence for some form of traditional statistical discrimination at play. I will not be able to explore potential forms of behavioral statistical discrimination to satisfy such skeptical readers. However, even the skeptical readers will agree that I found evidence that provision of quality information can reduce discrimination in healthcare labor markets.

topic of customer discrimination. Previous work had suggested that by "reducing the salience of race, platforms could reduce discrimination" (Fisman and Luca (2016)). In light of the evidence presented in this paper, market designs that favor better information flow regarding quality of providers can also reduce discrimination. And since both white and minority patients might have a preference for doctor racial concordance,[41] the more welfare-improving design might be closer to the treatment in my experiment than suppressing provider race information in online platforms.

6.2 A Methodological Point on the Validated Incentivized Conjoint Approach

This paper introduced the novel methodology VIC to overcome hypothetical bias while using choice-based conjoint designs. Conjoint analyses simulate market conditions set up via an experimental design while choice-based conjoint studies offer superior recovery of preferences compared to a direct stated preference elicitation (Louviere and Hensher (1983)). The value of conjoint analysis is most prominent in many markets where there is very little variation in key product attributes (e.g., prices in the short run) with the set of existing products representing only a very sparse set of points in the product characteristic space (Allenby et al. (2019)). The biggest advantage of conjoint analysis is that it enables the researcher to avoid "endogeneity" problems as all variation in both product attributes and price is exogenous and usable to estimate preference.

Choice-based conjoint analysis is one of the most popular tools in marketing research with an estimated 14,000 studies conducted yearly by various industry and academic research bodies (Allenby et al. (2019)). Despite this, the uptake in economics is limited. Economics experiments generally aspire to have incentive compatibility properties so that the subjects reveal their true preferences due to stakes aligned with their decisions. Hypothetical bias limits the uptake of conjoint analysis by economists. There are more conjoint studies conducted every year by researchers and businesses than randomized clinical trials conducted by medical researchers. Further development of the VIC method will provide a useful empirical tool for economists to complement methods using observational data.

[41] See the suggestive evidence in Section 3.3.1. Also, evidence from Alsan and Wanamaker (2018) and Alsan et al. (2019) would suggest that racial concordance between patient and doctor can enhance trust between them, especially if the patient is black.

Practically, many tech companies can avail themselves of VIC to improve the performance of their offerings when accurate predictions of customer preferences is critical to their success. While usage data can tell platforms a lot about customer preferences, periodically offering VIC surveys to customers to adjust presented contents can help companies such as social media platforms obtain information orthogonal to their usual data sources much like the company that partnered with me on this experiment.

Naturally, the methodology has limitations. Because VIC is a type of conjoint design, similar limitations apply. First, VIC can only estimate demand and cannot be used to compute market equilibrium outcomes such as market prices (wages for doctors in this paper) or equilibrium product positioning in characteristic space. Without supply assumptions and cost information, one is limited about statements on equilibrium outcomes and welfare. Second, conjoint surveys can present an unnatural environment with respect to the typical shopping experience. Furthermore, there is the burden on the researcher to adequately specify all product attributes and describe them to the respondents in a clear and meaningful way. In the present unique setting, this second limitation is less concerning as the conjoint design and the way options are presented to customers are very similar to the typical shopping experience on the platform.

As VIC informs respondents that the menu options are hypothetical and might be used in research, one might worry about experimenter demand effects. This concern is likely not significant as customers are shopping in the privacy of their own home to identify the best match (doctor) to whom they will be paying cash out of their own pocket. As discussed above, the individual choice models estimated with the responses to the hypothetical menus predicted the actual booking choice nicely.

7 Conclusion

Racial gaps in hiring, promotion, and pay persist in many industries even for high skilled professionals. This paper identifies a mechanism that contributes to these gaps. Using VIC on actual healthcare shoppers trying to book a doctor without a strong quality signal on doctor quality, customers are willing to pay a significant premium to have their colonoscopy done by a white doctor rather than a black or Asian doctor. Willingness-to-pay is lowered by 12.7% of the average price for black doctors and lower by 8.7% for Asian

doctors. But when quality information in form of "stars" is provided for each provider option, the willingness-to-pay gap for non-white providers dropped to 1.2% of the average price for black doctors and and 1.1% for Asian doctors.

Results are consistent with inaccurate statistical discrimination where customers of healthcare use mis-specified statistical models of provider quality. While the data suggest the majority of discrimination is statistical but not taste-based, it is worth noting that we must be careful to interpret this as a lack of animus against minorities. Customers may develop inaccurate beliefs due to such animus. I can say that a large part of the discrimination is statistical and that credible information on the relevant quality distributions can be provided to compel customers to adjust their inaccurate beliefs, but I cannot settle whether animus exists.

Decomposing the nature of customer discrimination towards doctors and other service professionals has pivotal implications for policy. If the discrimination is due to customers forming inaccurate statistical models about minority provider service quality, then interventions that can revise and correct those inaccurate beliefs would have impact.

Lastly, VIC methodology can broadly study individual preferences in a way economists find reliable, incentive compatible and cross-validated. Just as randomized controlled trials have come to be widely adopted by economists, choice-based conjoint surveys can become as much a part of the empirical economist arsenal as they are for marketing scholars. The novel design proposed with VIC can be a step in that direction.

References

Agan, Amanda and Sonja Starr, "Ban the box, criminal records, and racial discrimination: A field experiment," *The Quarterly Journal of Economics*, 2018, *133* (1), 191–235.

Aigner, Dennis J and Glen G Cain, "Statistical theories of discrimination in labor markets," *Ilr Review*, 1977, *30* (2), 175–187.

Allenby, G, "An introduction to hierarchical Bayesian modeling," in "Tutorial notes, Advanced Research Techniques Forum, American Marketing Association" 1997.

Allenby, Greg M and Peter E Rossi, "Hierarchical bayes models," *The handbook of marketing research: Uses, misuses, and future advances*, 2006, pp. 418–440.

_ **, Nino Hardt, and Peter E Rossi**, "Economic foundations of conjoint analysis," in "Handbook of the Economics of Marketing," Vol. 1, Elsevier, 2019, pp. 151–192.

Alsan, Marcella and Marianne Wanamaker, "Tuskegee and the health of black men," *The quarterly journal of economics*, 2018, *133* (1), 407–455.

_ **and Sarah Eichmeyer**, "Persuasion in Medicine: Messaging to Increase Vaccine Demand," Technical Report, National Bureau of Economic Research 2021.

_ **, Owen Garrick, and Grant Graziani**, "Does diversity matter for health? Experimental evidence from Oakland," *American Economic Review*, 2019, *109* (12), 4071–4111.

Arnold, David, Will Dobbie, and Crystal S Yang, "Racial bias in bail decisions," *The Quarterly Journal of Economics*, 2018, *133* (4), 1885–1932.

Arrow, Kenneth et al., "The theory of discrimination, discrimination in labor markets," *Achenfelter, A. Ress (eds.), Princeton–New Jersey*, 1973.

Arrow, Kenneth J., "Uncertainty and the Welfare Economics of Medical Care," *American Economic Review*, 1963, *53* (5), 941–973.

Ayres, Ian, Mahzarin Banaji, and Christine Jolls, "Race effects on eBay," *The RAND Journal of Economics*, 2015, *46* (4), 891–917.

Ballance, Joshua, Robert Clifford, and Daniel Shoag, ""No more credit score": Employer credit check bans and signal substitution," *Labour Economics*, 2020, *63*, 101769.

Bar, Revital and Asaf Zussman, "Customer discrimination: evidence from Israel," *Journal of Labor Economics*, 2017, *35* (4), 1031–1059.

Barr, Donald A, Maria Elena Gonzalez, and Stanley F Wanat, "The leaky pipeline: Factors associated with early decline in interest in premedical studies among underrepresented minority undergraduate students," *Academic Medicine*, 2008, *83* (5), 503–511.

Bartik, Alexander W and Scott T Nelson, "For Online Publication: Appendix for Deleting a Signal: Evidence from Pre-Employment Credit Checks," 2019.

Becker, Gary S et al., "The Economics of Discrimination," *University of Chicago Press Economics Books*, 1971.

Ben-Akiva, Moshe Emanuel, Daniel McFadden, Kenneth Train et al., *Foundations of stated preference elicitation: Consumer behavior and choice-based conjoint analysis*, Now, 2019.

Bertrand, Marianne and Esther Duflo, "Field experiments on discrimination," *Handbook of economic field experiments*, 2017, *1*, 309–393.

Bohren, J Aislinn, Kareem Haggag, Alex Imas, and Devin G Pope, "Inaccurate statistical discrimination," Technical Report, National Bureau of Economic Research 2019.

Burdekin, Richard CK and Todd L Idson, "Customer preferences, attendance and the racial structure of professional basketball teams," *Applied Economics*, 1991, *23* (1), 179–186.

Census-Bureau, U.S., "2018 Census," U.S. Department of Commerce May 2018.

Combes, Pierre-Philippe, Bruno Decreuse, Morgane Laouenan, and Alain Trannoy, "Customer discrimination and employment outcomes: theory and evidence from the french labor market," *Journal of Labor Economics*, 2016, *34* (1), 107–160.

Cooper, Lisa A, Debra L Roter, Rachel L Johnson, Daniel E Ford, Donald M Steinwachs, and Neil R Powe, "Patient-centered communication, ratings of care, and concordance of patient and physician race," *Annals of internal medicine*, 2003, *139* (11), 907–915.

Courey, Gabriel, "Gender and Racial Wage Differentials in Nonprofit Hospitals," *LABOUR*, 2020, *34* (4), 373–398.

Cui, Ruomeng, Jun Li, and Dennis J Zhang, "Reducing discrimination with reviews in the sharing economy: Evidence from field experiments on Airbnb," *Management Science*, 2020, *66* (3), 1071–1094.

DeGroot, Morris H, *Optimal statistical decisions*, Vol. 82, John Wiley & Sons, 2005.

DellaVigna, Stefano, Devin Pope, and Eva Vivalt, "Predict science to improve science," *Science*, 2019, *366* (6464), 428–429.

Dietl, Helmut, Anil Özdemir, and Andrew Rendall, "The role of facial attractiveness in tennis TV-viewership," *Sport Management Review*, 2020, *23* (3), 521–535.

Doleac, Jennifer L and Luke CD Stein, "The visible hand: Race and online market outcomes," *The Economic Journal*, 2013, *123* (572), F469–F492.

Edelman, Benjamin, Michael Luca, and Dan Svirsky, "Racial discrimination in the sharing economy: Evidence from a field experiment," *American Economic Journal: Applied Economics*, 2017, *9* (2), 1–22.

Elías, Julio J, Nicola Lacetera, and Mario Macis, "Paying for kidneys? a randomized survey and choice experiment," *American Economic Review*, 2019, *109* (8), 2855–88.

Fang, Hanming and Andrea Moro, "Theories of statistical discrimination and affirmative action: A survey," *Handbook of social economics*, 2011, *1*, 133–200.

Filut, Amarette, Madelyn Alvarez, and Molly Carnes, "Discrimination toward physicians of color: a systematic review," *Journal of the National Medical Association*, 2020, *112* (2), 117–140.

Fisman, Raymond and Michael Luca, "Fixing discrimination in online marketplaces," *Harvard business review*, 2016, *94* (12), 88–95.

Fong, Christina M and Erzo FP Luttmer, "What determines giving to Hurricane Katrina victims? Experimental evidence on racial group loyalty," *American Economic Journal: Applied Economics*, 2009, *1* (2), 64–87.

Freeman, Brandi Kaye, Alden Landry, Robert Trevino, David Grande, and Judy A Shea, "Understanding the leaky pipeline: perceived barriers to pursuing a career in medicine or dentistry among underrepresented-in-medicine undergraduate students," *Academic Medicine*, 2016, *91* (7), 987–993.

Fung, Constance H, Marc N Elliott, Ron D Hays, Katherine L Kahn, David E Kanouse, Elizabeth A McGlynn, Mark D Spranca, and Paul G Shekelle, "Patients' preferences for technical versus interpersonal quality when selecting a primary care physician," *Health services research*, 2005, *40* (4), 957–977.

Gneezy, Uri, John List, and Michael K Price, "Toward an understanding of why people discriminate: Evidence from a series of natural field experiments," Technical Report, National Bureau of Economic Research 2012.

Grisham, Sarah, "Medscape physician compensation report 2017," *Medscape website. https://www. medscape. com/slideshow/compensation-2017-overview-6008547. Published April*, 2017, *5*.

Gunja, Munira Z and Sara R Collins, "Who are the remaining uninsured, and why do they lack coverage," *Commonwealth Fund. https://www. commonwealthfund. org/publications/issue-briefs/2019/aug/who-are-remaining-uninsured-and-why-do-they-lack-coverage*, 2019.

Harris, Katherine M, "How do patients choose physicians? Evidence from a national survey of enrollees in employment-related health plans," *Health services research*, 2003, *38* (2), 711–732.

He, Kaiming, Xiangyu Zhang, Shaoqing Ren, and Jian Sun, "Deep residual learning for image recognition," in "Proceedings of the IEEE conference on computer vision and pattern recognition" 2016, pp. 770–778.

Hirpa, Meron, Tinsay Woreta, Hilena Addis, and Sosena Kebede, "What matters to patients? A timely question for value-based care," *Plos one*, 2020, *15* (7), e0227845.

Holzer, Harry J and Keith R Ihlanfeldt, "Customer discrimination and employment outcomes for minority workers," *The Quarterly Journal of Economics*, 1998, *113* (3), 835–867.

Idan, Edna, Anlu Xing, Javarcia Ivory, and Marcella Alsan, "Sociodemographic correlates of medical mistrust among African American men living in the East Bay," *Journal of health care for the poor and underserved*, 2020, *31* (1), 115–127.

Jr, Richard Carey McGregory, "An analysis of black-white wage differences in nursing: wage gap or wage premium?," *The Review of Black Political Economy*, 2013, *40* (1), 31–37.

Jr, Stanley S Bergen, "Underrepresented minorities in medicine," *JAMA*, 2000, *284* (9), 1138–1139.

Kahn, Lawrence M and Peter D Sherer, "Racial differences in professional basketball players' compensation," *Journal of Labor Economics*, 1988, *6* (1), 40–61.

Kessler, Judd B, Corinne Low, and Colin D Sullivan, "Incentivized resume rating: Eliciting employer preferences without deception," *American Economic Review*, 2019, *109* (11), 3713–44.

Kolstad, Jonathan T and Michael E Chernew, "Quality and consumer decision making in the market for health insurance and health care services," *Medical care research and review*, 2009, *66* (1_suppl), 28S–52S.

Laouénan, Morgane and Roland Rathelot, "Can information reduce ethnic discrimination? evidence from airbnb," *American Economic Journal: Applied Economics*, 2020.

Leonard, Jonathan S, David I Levine, and Laura Giuliano, "Customer discrimination," *The Review of Economics and Statistics*, 2010, *92* (3), 670–678.

Lett, Lanair Amaad, H Moses Murdock, Whitney U Orji, Jaya Aysola, and Ronnie Sebro, "Trends in racial/ethnic representation among US medical students," *JAMA network open*, 2019, *2* (9), e1910490–e1910490.

Li, Siyue and Austin Hubner, "The impact of web-based ratings on patient choice of a primary care physician versus a specialist: randomized controlled experiment," *Journal of medical Internet research*, 2019, *21* (6), e11188.

List, John, "A Course in Experimental Economics," 2021.

List, John A, "The nature and extent of discrimination in the marketplace: Evidence from the field," *The Quarterly Journal of Economics*, 2004, *119* (1), 49–89.

_ **and David Lucking-Reiley**, "The effects of seed money and refunds on charitable giving: Experimental evidence from a university capital campaign," *Journal of political Economy*, 2002, *110* (1), 215–233.

_ **, Paramita Sinha, and Michael H Taylor**, "Using choice experiments to value non-market goods and services: evidence from field experiments," *Advances in economic analysis & policy*, 2006, *5* (2).

List, John and Uri Gneezy, *The why axis: Hidden motives and the undiscovered economics of everyday life*, Random House, 2014.

Louviere, Jordan J and David A Hensher, "Using discrete choice models with experimental design data to forecast consumer demand for a unique cultural event," *Journal of Consumer research*, 1983, *10* (3), 348–361.

Luft, Harold S, Deborah W Garnick, David H Mark, Deborah J Peltzman, Ciaran S Phibbs, Erik Lichtenberg, and Stephen J McPhee, "Does quality influence choice of hospital?," *Jama*, 1990, *263* (21), 2899–2906.

Ly, Dan P, Seth A Seabury, and Anupam B Jena, "Differences in incomes of physicians in the United States by race and sex," *BMJ: British Medical Journal*, 2016, *353*.

McFadden, Daniel, "Stated preference methods and their applicability to environmental use and non-use valuations," in "Contingent Valuation of Environmental Goods," Edward Elgar Publishing, 2017.

Merchant, Juanita L and M Bishr Omary, "Underrepresentation of underrepresented minorities in academic medicine: the need to enhance the pipeline and the pipe," 2010.

Mooney, Cathleen, Jack Zwanziger, Ciaran S Phibbs, and Susan Schmitt, "Is travel distance a barrier to veterans' use of VA hospitals for medical surgical care?," *Social science & medicine*, 2000, *50* (12), 1743–1755.

Nardinelli, Clark and Curtis Simon, "Customer racial discrimination in the market for memorabilia: The case of baseball," *The Quarterly Journal of Economics*, 1990, *105* (3), 575–595.

Neumark, David, Roy J Bank, and Kyle D Van Nort, "Sex discrimination in restaurant hiring: An audit study," *The Quarterly journal of economics*, 1996, *111* (3), 915–941.

Nicholson, Sean and Carol Propper, "Medical workforce," in "Handbook of health economics," Vol. 2, Elsevier, 2011, pp. 873–925.

Raghavarao, Damaraju, James B Wiley, and Pallavi Chitturi, *Choice-based conjoint analysis: models and designs*, CRC Press, 2010.

Rao, Vithala R et al., *Applied conjoint analysis*, Springer, 2014.

Razzouk, Nabil, Victoria Seitz, and Janet Marie Webb, "What's important in choosing a primary care physician: an analysis of consumer response," *International Journal of Health Care Quality Assurance*, 2004.

Santos, Rita, Hugh Gravelle, and Carol Propper, "Does quality affect patients' choice of doctor? Evidence from England," *The Economic Journal*, 2017, *127* (600), 445–494.

Sarsons, Heather, "Interpreting signals in the labor market: evidence from medical referrals," *Job Market Paper*, 2017.

Schmitt, Susan K, Ciaran S Phibbs, and John D Piette, "The influence of distance on utilization of outpatient mental health aftercare following inpatient substance abuse treatment," *Addictive Behaviors*, 2003, *28* (6), 1183–1192.

Stone, Eric W and Ronald S Warren, "Customer discrimination in professional basketball: evidence from the trading-card market," *Applied Economics*, 1999, *31* (6), 679–685.

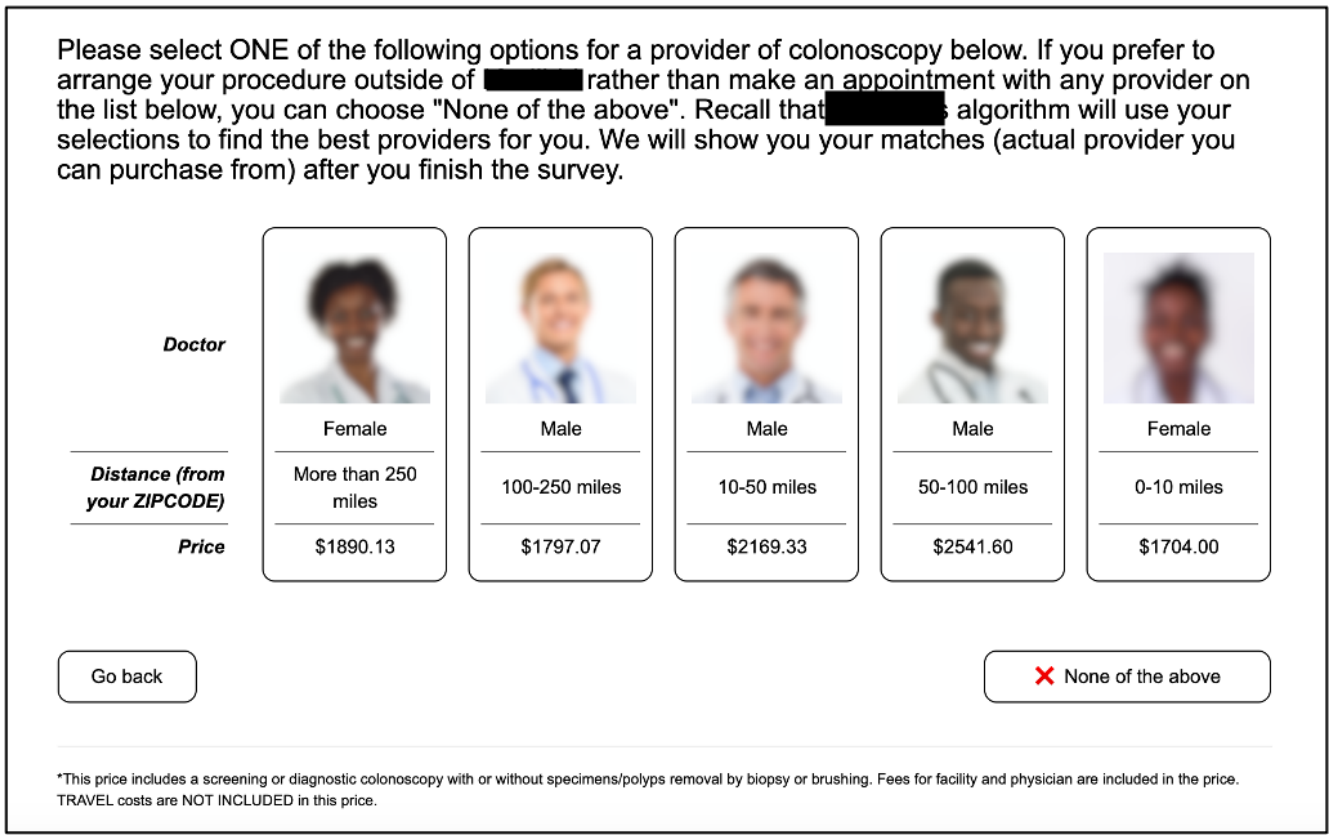
Sullivan, Colin D, "Eliciting preferences over life and death: experimental evidence from organ transplantation," 2021.

Train, Kenneth E, *Discrete choice methods with simulation*, Cambridge university press, 2009.

Wozniak, Abigail, "Discrimination and the effects of drug testing on black employment," *Review of Economics and Statistics*, 2015, *97* (3), 548–566.

Yoon, Jean, Megan E Vanneman, Sharon K Dally, Amal N Trivedi, and Ciaran S Phibbs, "Veterans' reliance on VA care by type of service and distance to VA for nonelderly VA-Medicaid dual enrollees," *Medical care*, 2019, *57* (3), 225.

Figure 1: Single Example of Conjoint Menu of Doctors



Notes: This figure shows an example of a menu of hypothetical doctor options, with 5 doctor options and a "None of the above" option, that form the basis of the choice experiments used in this paper. Each customer/respondent will encounter 14 menus like this one but with exogenously varied doctor features like price, profile picture, travel distance and sex of the doctor. The customer will have to click one of the options for each menu to proceed in the experiment. The name of the online platform that partnered with me on this project is blacked out in this figure to maintain privacy for the platform per the Data Use Agreement.

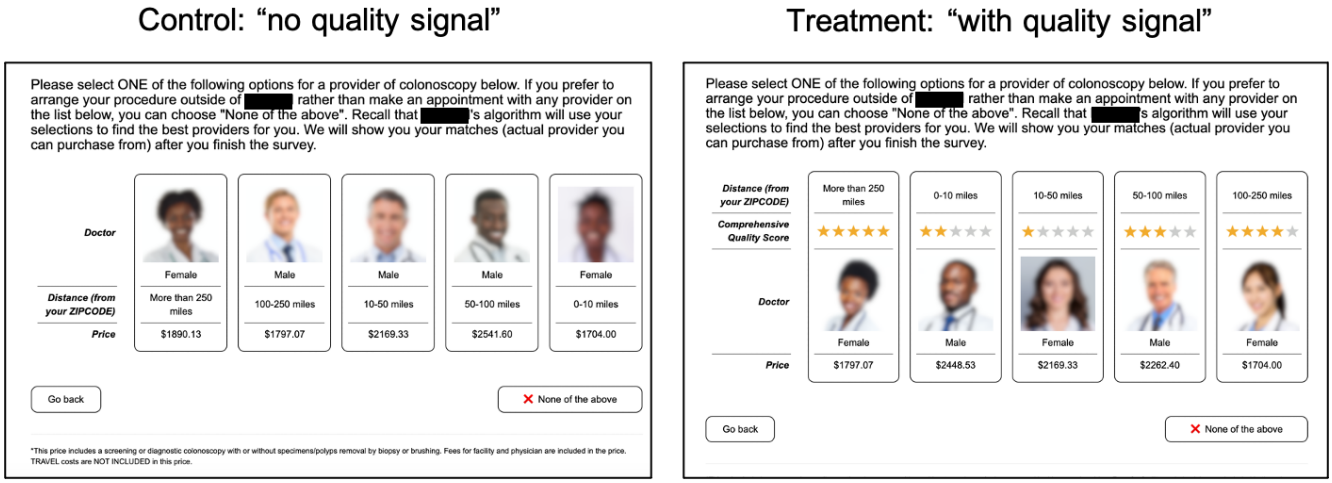
Figure 2: Matching Conjoint Responses to Actual Doctor Options



10 Actual Doctor options
- ✓ 8 with highest match score
- ✓ 2 randomly selected
- ✓ Options ordering scrambled

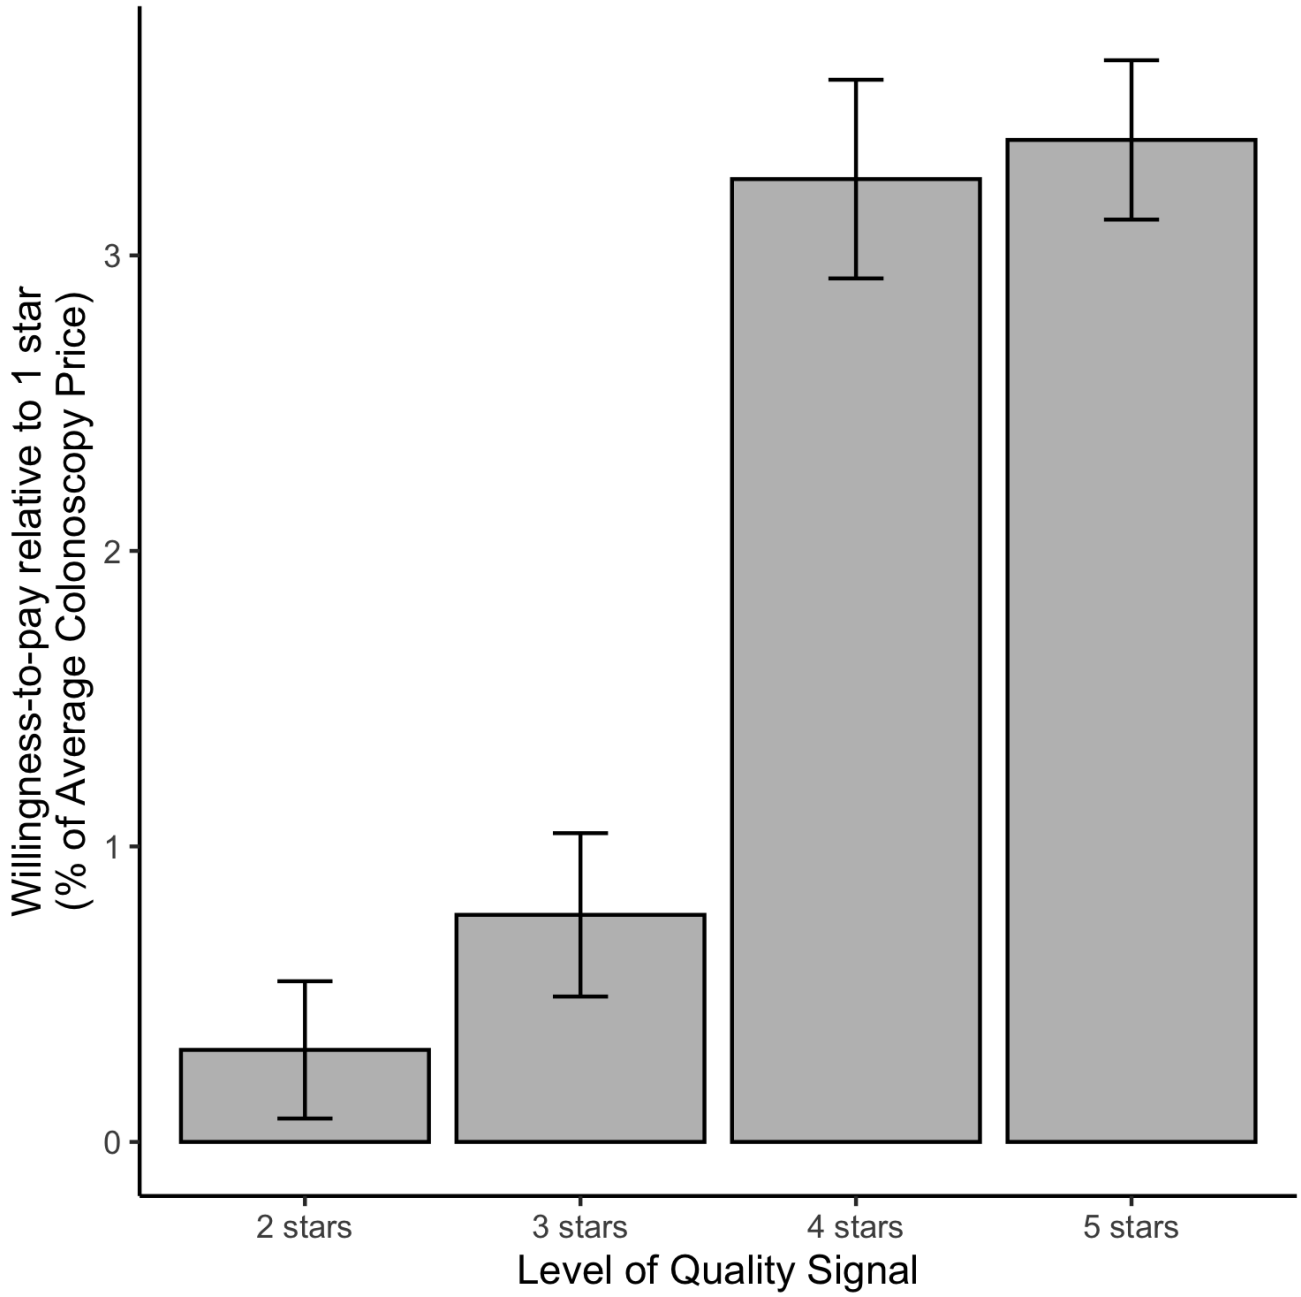
Notes: This figure illustrate the mechanism with which the conjoint survey with hypothetical doctor profiles is incentivized. Based on the customer's choice across the 14 menus of hypothetical doctor options, they will be matched to 10 actual doctors that they can purchase from. For each customer, 8 of the actual doctors will be the ones with the highest predicted choice probability based on the estimated preference parameters of that customer using their responses to the 14 menus while 2 actual doctors will be randomly drawn from the remaining doctors. On the menus in this figure, the name of the online platform that partnered with me on this project is blacked out in this figure to maintain privacy for the platform per the Data Use Agreement.

Figure 3: Single Menu of Doctor Options for Under Control and Treatment Conditions



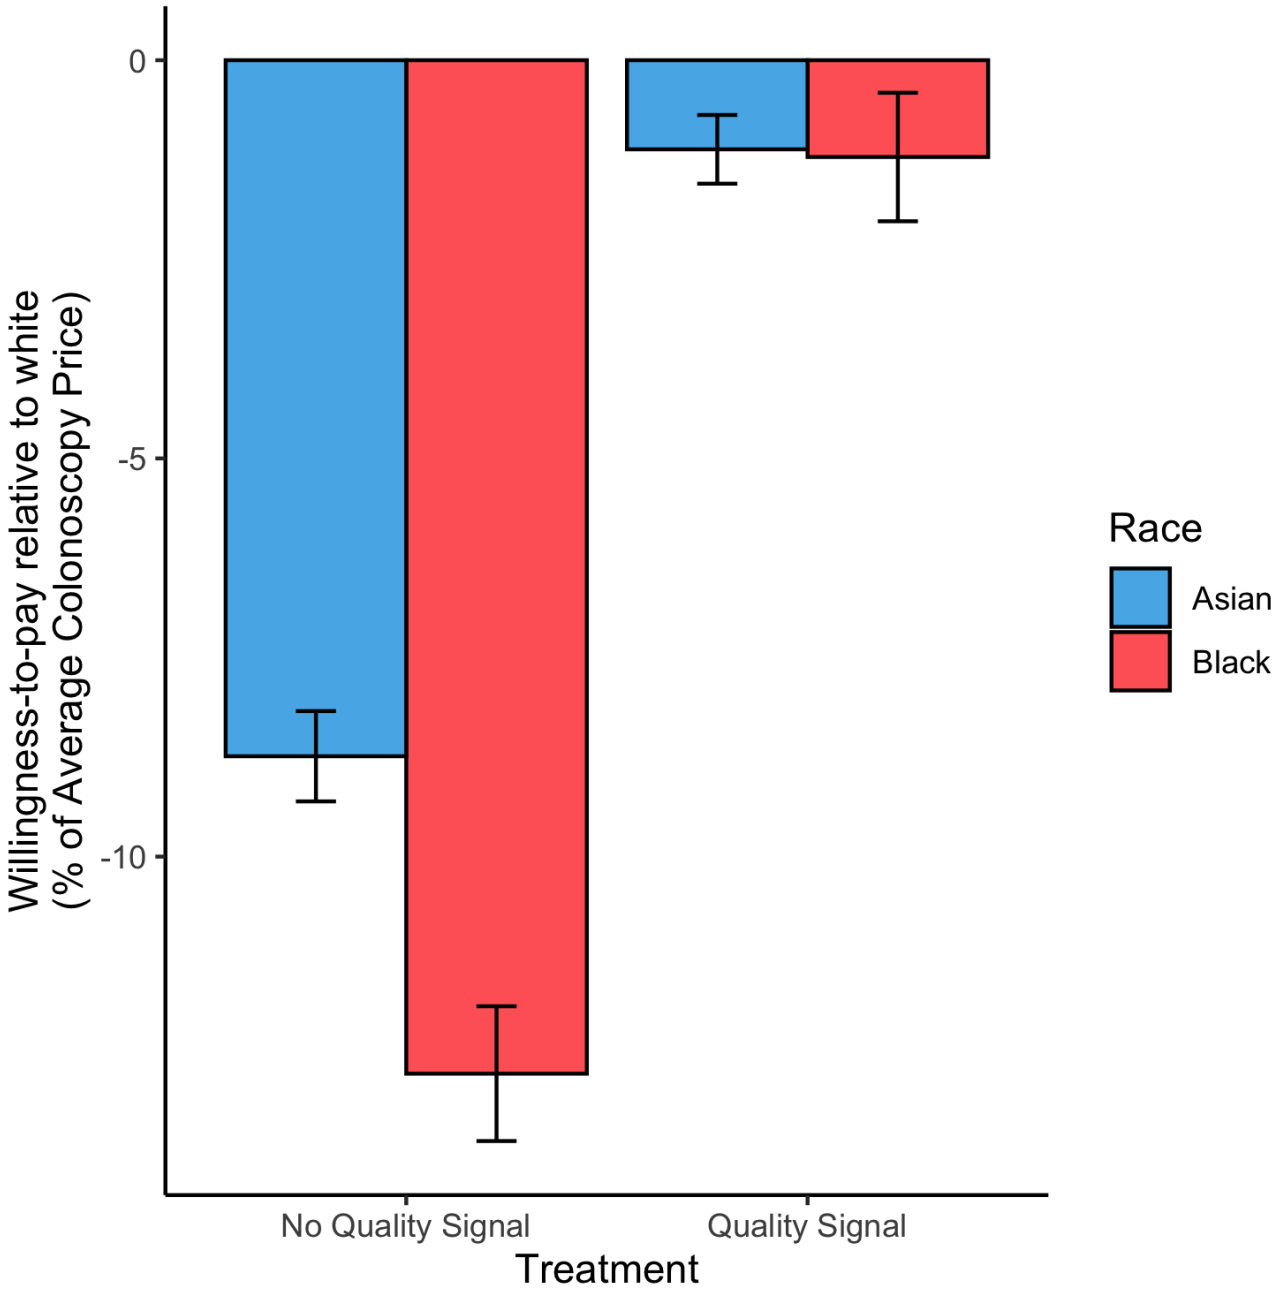
Notes: This figure shows the difference between the menus of hypothetical doctors that will be presented by each customer in the control group (left panel) and the treatment group (right panel). The treatment group choice menus feature an additional doctor attribute: the quality "stars" as shown on the right panel. Except for the price, the ordering from top to bottom of each doctor attribute (profile picture and sex, travel distance, and stars) is randomized across individuals. The name of the online platform that partnered with me on this project is blacked out in this figure to maintain privacy for the platform per the Data Use Agreement.

Figure 4: Willingness-to-pay "Premium" by Number of Stars Relative to Doctors with 1 Star as Quality Signal



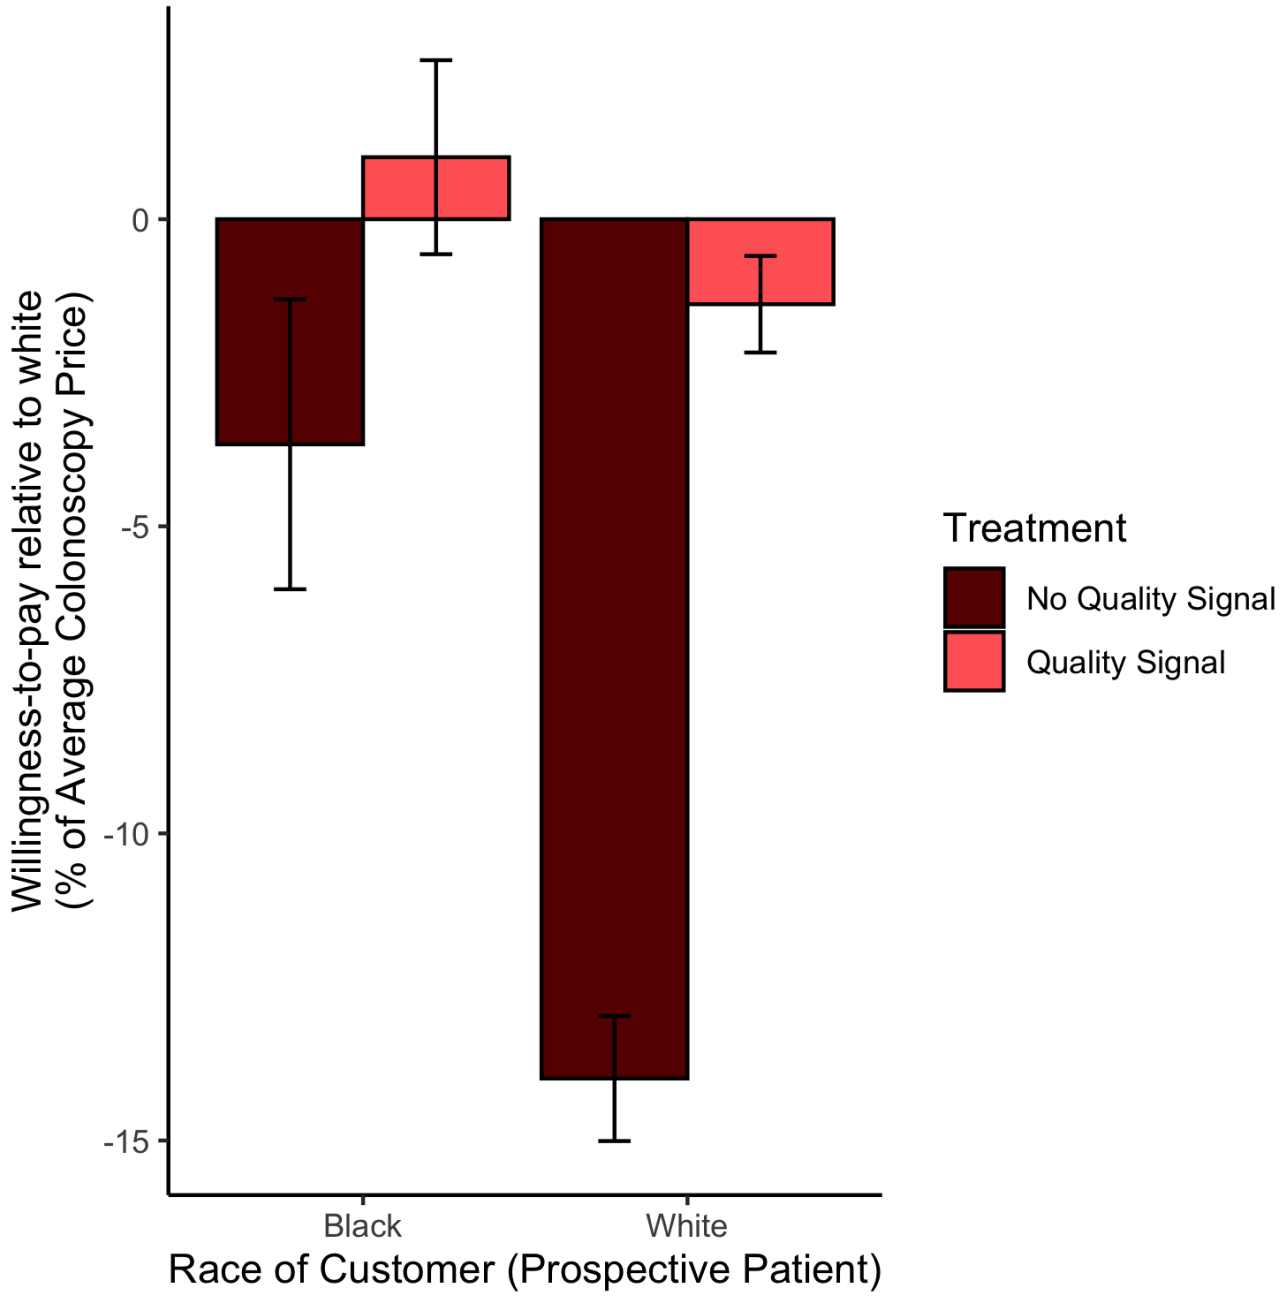
Notes: This figure presents the estimated parameter results for different star-levels (relative to 1 star) from the primary choice model from Section 2.5 expressed in terms of willingness-to-pay calculated as outlined in Section 2.5.1, using the treatment group data. The height of the bars represents willingness-to-pay coefficients as a percentage of the average price of a colonoscopy in my sample ($2122.80). The error bars show the 95% confidence intervals.

Figure 5: Willingness-to-pay "Penalty" by Race Relative to White Doctors



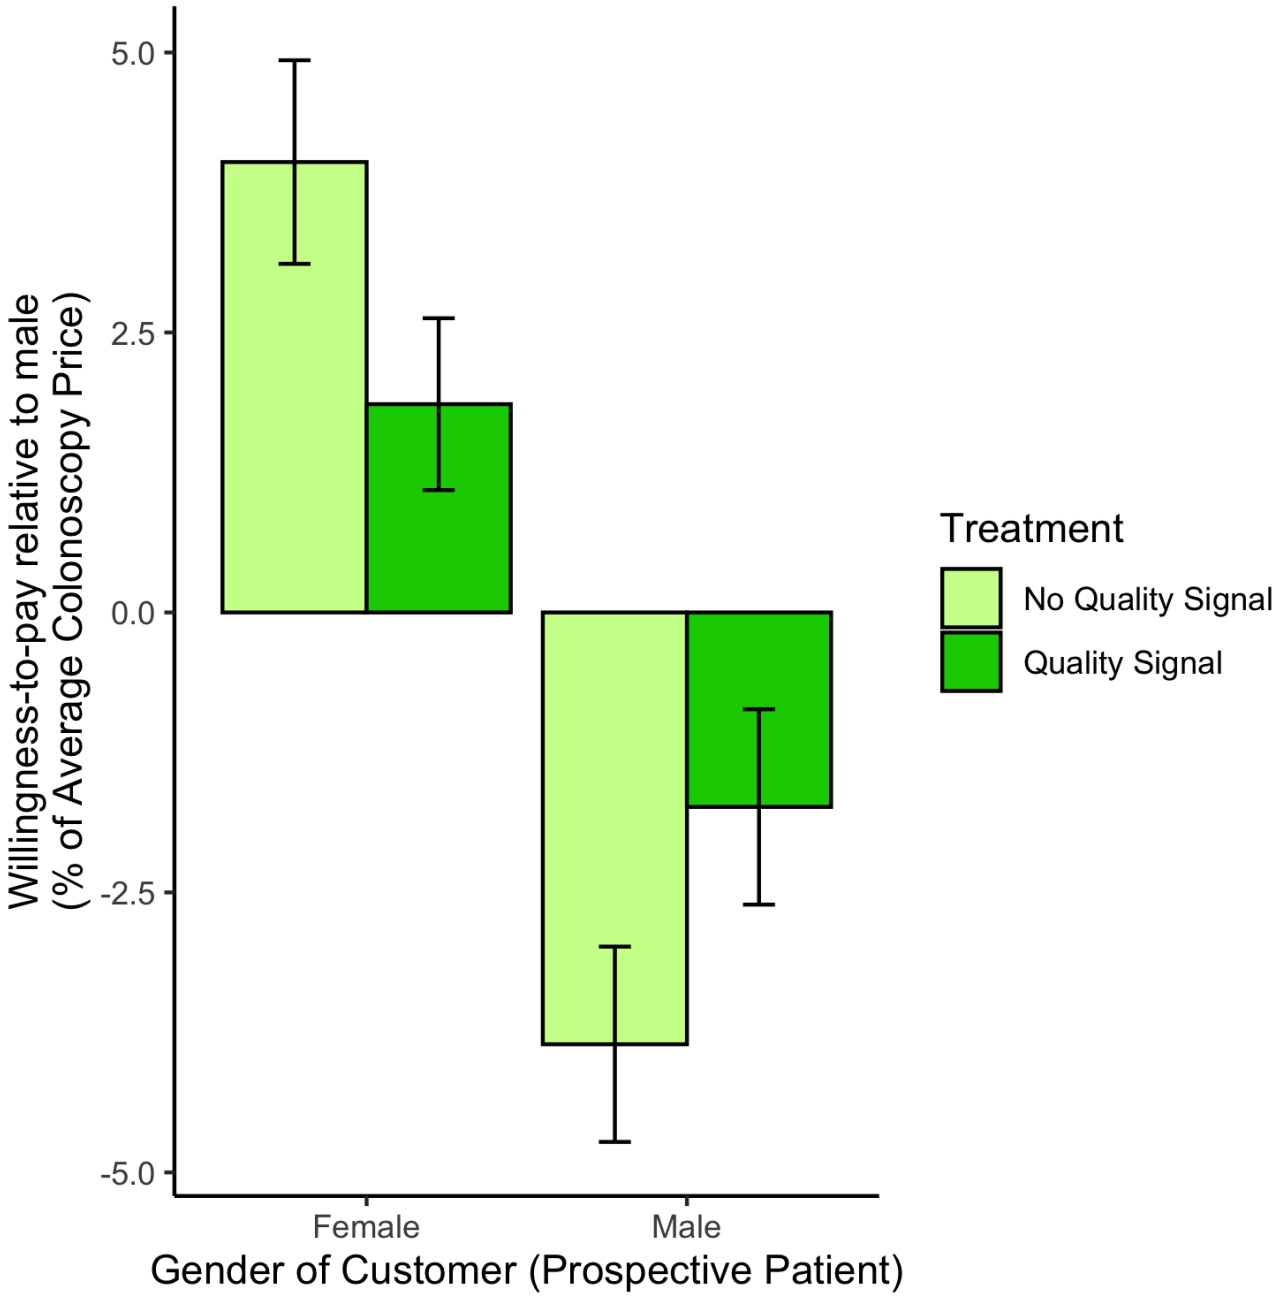
Notes: This figure presents the estimated parameter results for black and Asian race (relative to white race) from the primary choice model from Section 2.5 expressed in terms of willingness-to-pay calculated as outlined in Section 2.5.1, using the control group and treatment groups. The height of the bars represents willingness-to-pay coefficients as a percentage of the average price of a colonoscopy in my sample ($2122.80). The error bars show the 95% confidence intervals.

Figure 6: Comparison of Black and White Customers' Responses



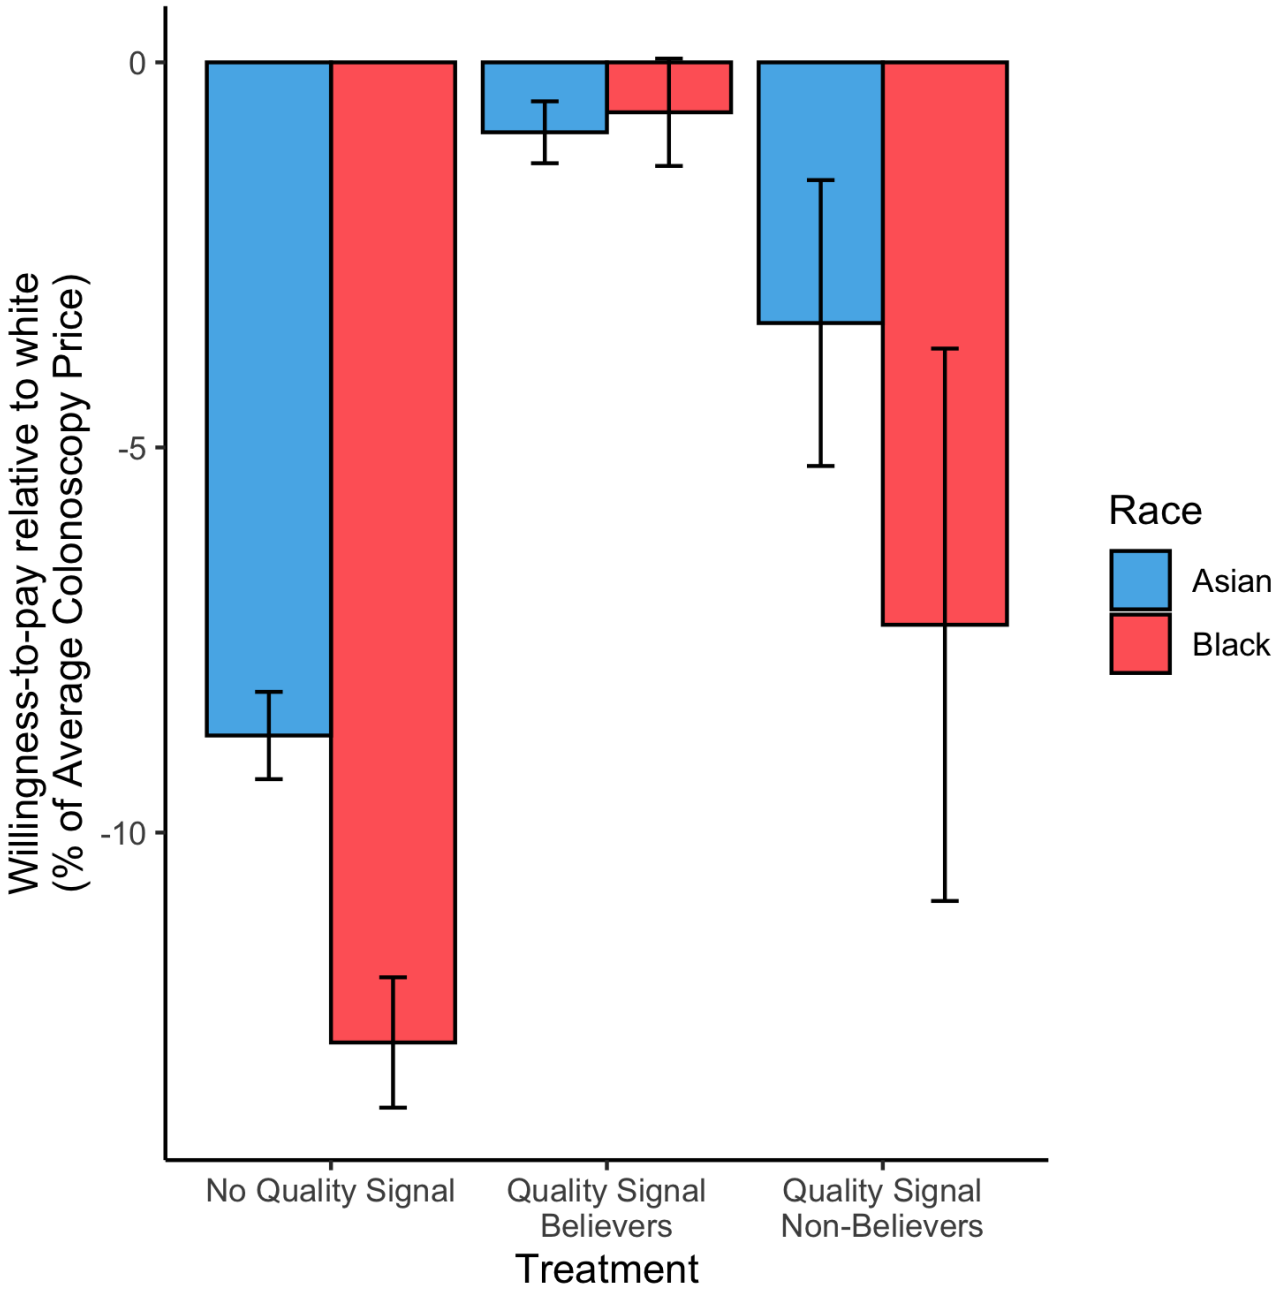
Notes: This figure presents the estimated parameter results for black race (relative to white race) from the primary choice model from Section 2.5 expressed in terms of willingness-to-pay calculated as outlined in Section 2.5.1, using the control group and treatment groups segmented by the customer/respondent's race (black or white). The height of the bars represents willingness-to-pay coefficients as a percentage of the average price of a colonoscopy in my sample ($2122.80). The error bars show the 95% confidence intervals.

Figure 7: Comparison of Female Doctor Willingness-to-pay "Penalty" Relative to Male by Customer Gender



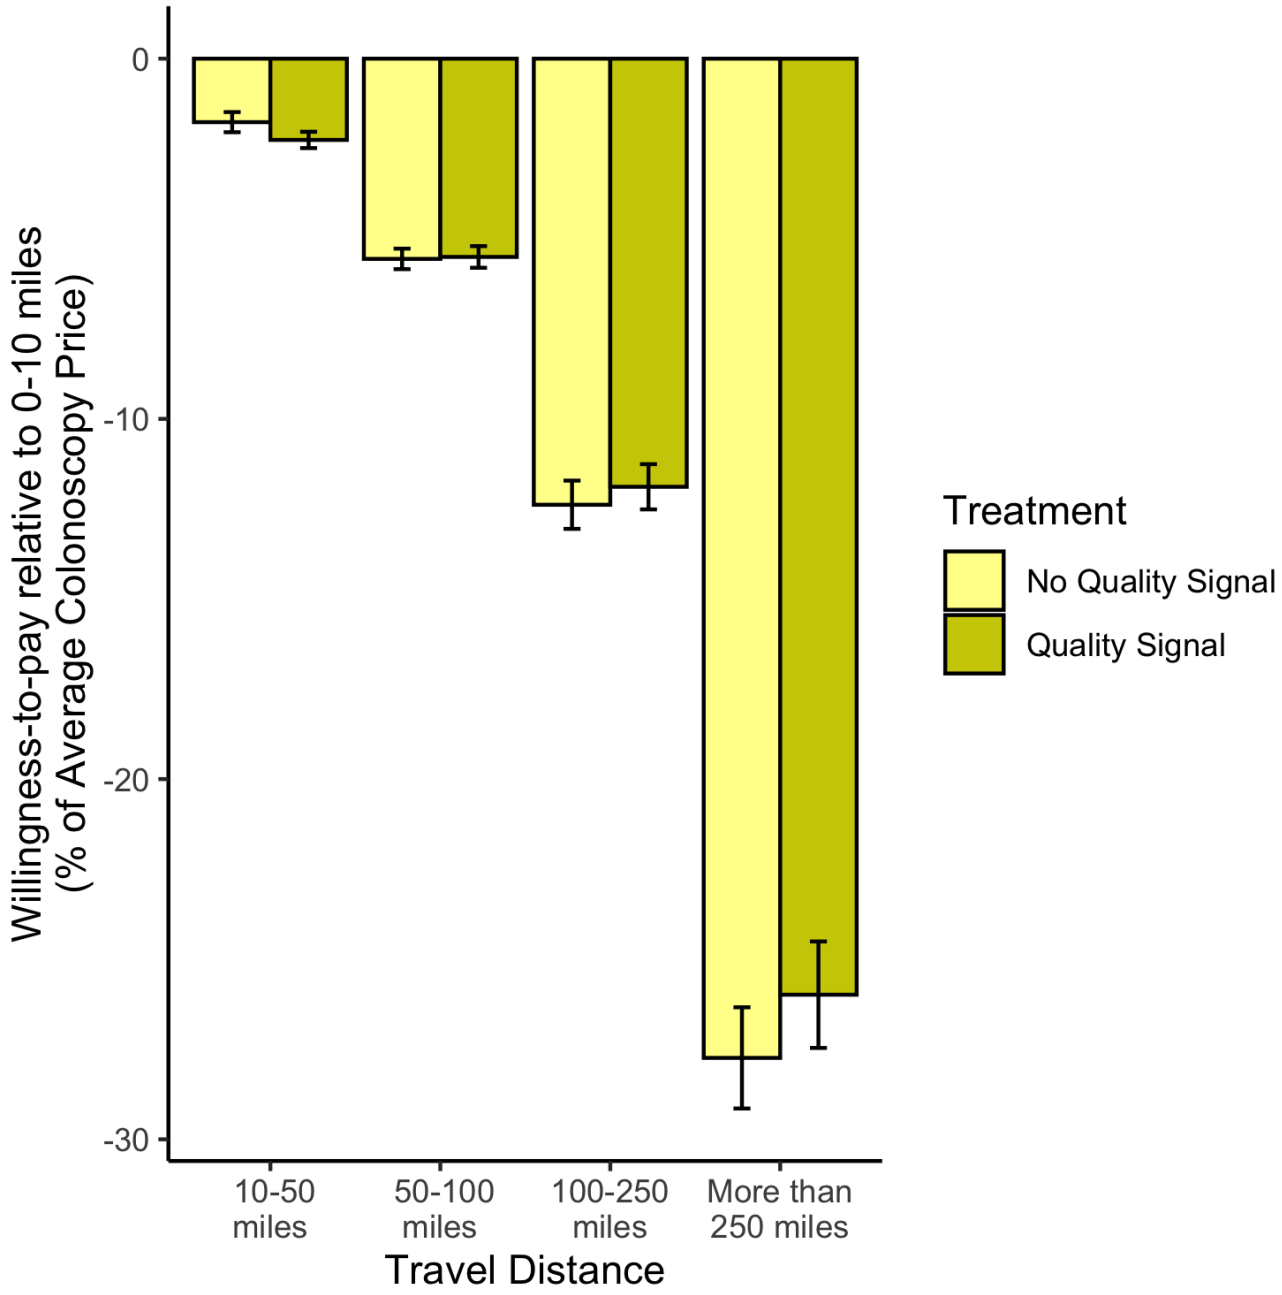
Notes: This figure presents the estimated parameter results for female gender (relative to male) from the primary choice model from Section 2.5 expressed in terms of willingness-to-pay calculated as outlined in Section 2.5.1, using the control group and treatment groups segmented by the customer/respondent's gender. The height of the bars represents willingness-to-pay coefficients as a percentage of the average price of a colonoscopy in my sample ($2122.80). The error bars show the 95% confidence intervals.

Figure 8: Comparison of "Believers" of the Quality Signals versus "Non-Believers"



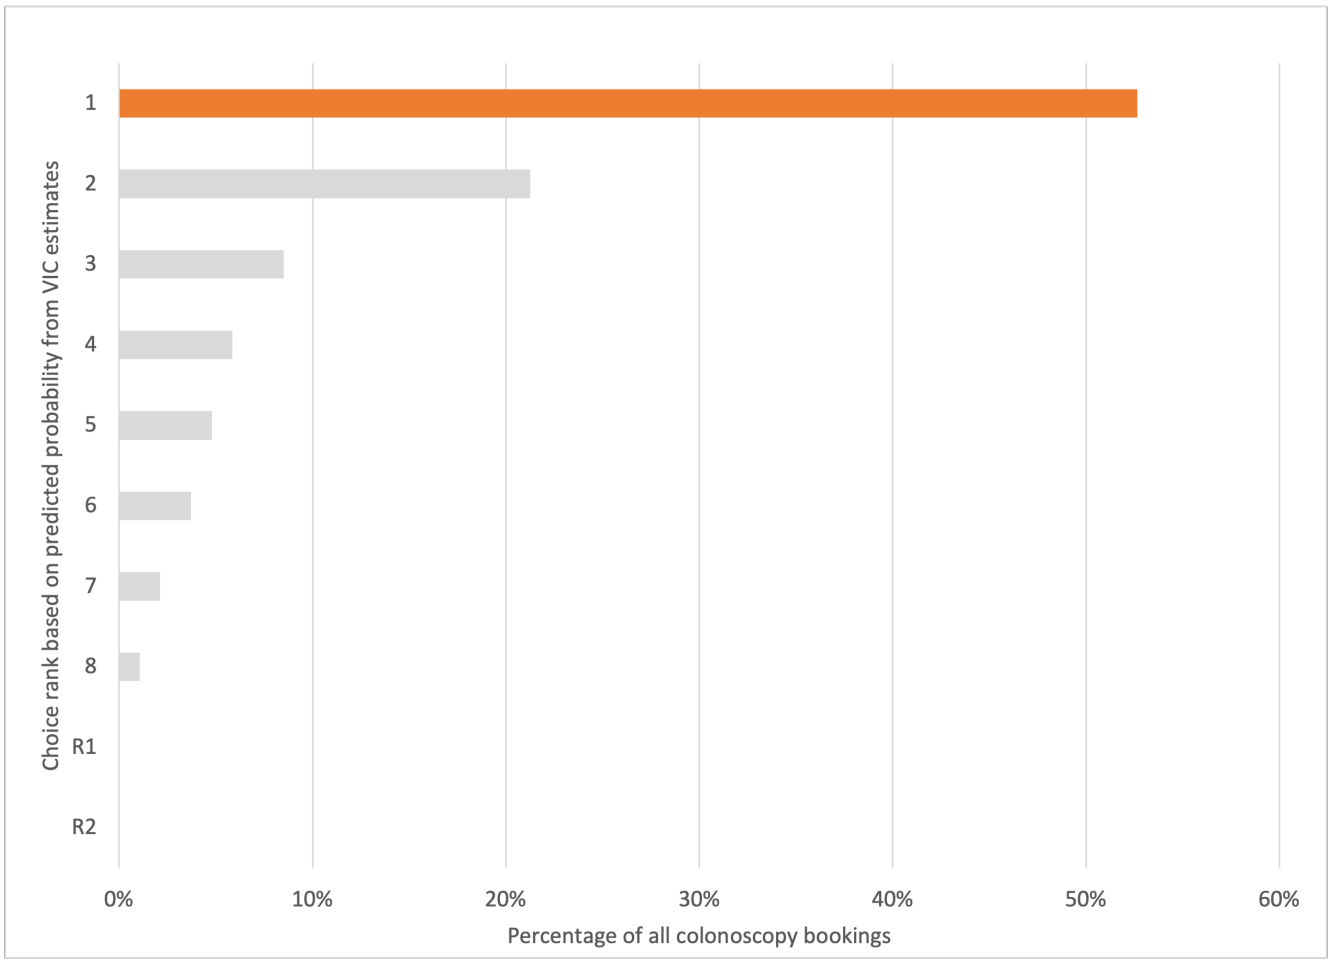
Notes: This figure presents the estimated parameter results for black race (relative to white race) from the primary choice model from Section 2.5 expressed in terms of willingness-to-pay calculated as outlined in Section 2.5.1, using the control group, the "Believers" sample in the treatment group, and the "Non-Believers" sample in the treatment group. "Non-Believers" are defined as customers who responded that they "Strongly Disagree" or "Disagree" that "A provider with a higher Comprehensive Quality Score is a better provider than a provider with a lower Comprehensive Quality Score". The height of the bars represents willingness-to-pay coefficients as a percentage of the average price of a colonoscopy in my sample ($2122.80). The error bars show the 95% confidence intervals.

Figure 9: Comparison of Willingness-to-pay for Different Travel Distances With and Without Quality Signals



Notes: This figure presents the estimated parameter results for different travel distances (relative to "0-10 miles") from the primary choice model from Section 2.5 expressed in terms of willingness-to-pay calculated as outlined in Section 2.5.1, using the control and treatment group data. The height of the bars represents willingness-to-pay coefficients as a percentage of the average price of a colonoscopy in my sample ($2122.80). The error bars show the 95% confidence intervals.

Figure 10: Actual Booking Choice by Predicted Choice Probability Using Individual Coefficients from Conjoint



Notes: This figure presents the percentage of actual doctor booking by the rank (within the actual menu of 10 booking options) of doctor based on predicted choice probability individual coefficients estimated from choices made in the 14 hypothetical menus in the validated incentivized conjoint set-up. Individual coefficients are estimated with a mixed logit model using hierarchical Bayes estimation. The rank number 1 through 8 are the highest ranking doctors for each customer, R1 and R2 are two random doctors chosen from the remaining doctors in the pool and offered alongside the 8 to form the actual menu of 10 doctors.

Table 1: Summary Statistics and Balance Checks

	Total		Control		Treatment		
	mean	s.d.	mean	s.d.	mean	s.d.	Difference p-value
Female	0.46	0.50	0.43	0.50	0.49	0.50	0.395
Black	0.10	0.30	0.12	0.32	0.09	0.28	0.461
Asian	0.06	0.23	0.06	0.24	0.06	0.23	0.984
Hispanic	0.10	0.30	0.08	0.28	0.12	0.32	0.424
White	0.83	0.38	0.80	0.40	0.86	0.35	0.274
Age under 45	0.07	0.25	0.05	0.22	0.09	0.28	0.277
Age between 45 and 54	0.45	0.50	0.47	0.50	0.42	0.50	0.515
Age between 55 and 64	0.46	0.50	0.45	0.50	0.48	0.50	0.647
High School Graduate	0.80	0.40	0.80	0.40	0.81	0.40	0.886
College Graduate	0.29	0.46	0.29	0.46	0.30	0.46	0.917
Married	0.50	0.50	0.50	0.50	0.51	0.50	0.886
Currently employed	0.54	0.50	0.53	0.50	0.57	0.50	0.528
Currenly Self-employed	0.20	0.40	0.17	0.37	0.24	0.43	0.171
Uninsured	0.90	0.30	0.89	0.31	0.91	0.28	0.587
# of outside options evaluated	4.42	3.17	4.19	2.85	4.68	3.48	0.249
Price of outside option	2535	660	2563	560	2502	758	0.489
Colonoscopy in past 10 years	0.79	0.41	0.80	0.40	0.77	0.42	0.578
Booked one of matched options	0.84	0.37	0.82	0.39	0.87	0.34	0.324
Patients	224		120		104		
N	18,816		10,080		8,736		

Notes: This table reports the background characteristics of the 224 customers in the main sample, pooled and by treatment group. "Female" indicates the share of female sex; "Black," "Asian," "Hispanic," and "White" indicate the shares of customers belonging to each of these categories. Age data was recorded in intervals, "Age under 45," "Age between 45 and 54," and "Age between 55 and 64" indicate shares of customers in these age buckets. "High School Graduate" indicate the share of customers who did not graduate from high school, "College Graduate" indicate indicate the share of customers who reported that they have a bachelor's or associate degree. "Married" indicate the share of customers who are currently married (not single, divored, separated, or widowed). "Currently employed" indicate the share who are employed part-time or full-time (not including self-employment while "Currently Self-employed" indicate the share who are self-employed. "Uninsured" indicates the share of customers who selected "NOT COVERED by any health insurance plan" in the survey. "# of outside options evaluated" indicates the number of colonoscopy providers that the customer has explored outside of the pilot's tool, while "Price of outside option" is the self-reported price that the customer expects to pay out-of-pocket (in USD). "Colonoscopy in past 10 years" indicates the share of customers who reported that they have had a colonoscopy in the past 10 years. "Colonoscopy in past 10 years" indicates the share of customers who reported that they have had a colonoscopy in the past 10 years. "Booked one of matched options" indicates the share who booked one of the 10 real doctor options matched to them through the experiment. Table shows averages ("mean") and standard deviations ("s.d."). The Difference p-value column reports the p-value for the test of equality between the treatment and control groups. Stars indicate whether this difference is significant. I do not find any statistically significant differences for all these differences.

Table 2: Choice Model Parameter Estimates in Willingness-to-Pay

Attribute	Control	Treatment	t-Statistic
a	$WTP_a^{Control}$	$WTP_a^{Treatment}$	
Black	-270.133	-25.803	-19.295
	(9.164)	(8.739)	
Asian	-185.529	-23.757	-20.970
	(6.139)	(4.672)	
Female	-15.598	0.529	-1.669
	(7.227)	(6.413)	
10-50 miles	-37.409	-47.888	2.689
	(3.028)	(2.454)	
50-100 miles	-117.994	-116.837	-0.289
	(2.879)	(3.255)	
100-250 miles	-262.882	-252.300	-1.064
	(6.587)	(6.798)	
More than 250 miles	-588.866	-551.620	-1.687
	(13.133)	(16.004)	
5 stars		71.981	
		(2.920)	
4 stars		69.167	
		(3.641)	
3 stars		16.312	
		(2.994)	
2 stars		6.612	
		(2.516)	
N	10,080	8,736	

Notes: This table presents the estimated parameter results for the primary choice model from Section 2.5 expressed in terms of willingness-to-pay calculated as outlined in Section 2.5.1. All willingness-to-pay coefficients are in dollar units with standard errors for parameters given in parentheses. Marginal utility is assumed to be a linear function of doctor attributes and the parameters represents the willingness-to-pay relative to the reference attribute level. For race, the reference level is white race for the doctor. For gender, the reference level is male. For distance, the reference level is "0-10 miles" distance between the doctor's zip code and the customer's zip code. For quality stars, 1 star is the reference level. The t-statistic is calculated as described in Section 2.5.1.

A Appendix

A.1 Experimental Details

In this Section, I describe the design of my experiment in detail, including recruitment materials (A.2), and the VIC survey construction (A.3).

A.2 Recruitment Materials

The partnering platform sent emails to its subscription base to offer an opportunity to use a pilot tool to shop for one of three medical procedures: colonoscopy, MRI or knee replacement. The pilot tool allows each customer a one-time opportunity to get provider options with prices without having to pay the usual fee, rather than a replacement for the usual channel. The recruitment email for customers, shown in Figure 11, was sent to the subscribed customer base to recruit customers who want to "Shop for Your Next Provider for Colonoscopy, MRI or Knee Replacement" and that "the use of this pilot tool will be a one-time-only offer for each customer....[and] we'll waive the $25 fee for the appointment you book through the pilot tool." This email was sent out twice in the beginning of 2021 spaced 7 weeks apart. A Tweet was also posted weekly from the week of November 30 2020 to the week of March 29 2021 (Figure 12).

[redacted] Pilot: Shop for Your Next Provider for Colonoscopy, MRI or Knee Replacement

[redacted]@[redacted]

Thu 1/14/2021 4:08 PM

To: [redacted]

Dear Customer,

[redacted] is piloting [redacted]–Choice to help you find the best doctor for your needs. We will survey your preferences in this tool and identify providers that you can purchase your procedure from. With this pilot tool, we search the web for hundreds of providers across [redacted], [redacted], and more, giving you more options than any individual website to choose from. Are you shopping for a colonoscopy, MRI or knee replacement? If so, you can use this pilot tool. Currently, the use of this pilot tool will be a one-time-only offer for each customer.

After you complete a few questions that will take about 15 minutes, [redacted] will then find the best providers across all major websites for you based on your responses, using [redacted]'s newly developed matching algorithm. We will email you with 10 real providers within 48 hours after you finish. The more carefully you complete these questions, the better we will be able to match you with real provider options.

If you participate in this pilot, we'll waive the $25 fee for the appointment you book through the pilot tool. When you are ready to book a provider for the procedure you are looking for, please take the tool by clicking the link here:

https:[redacted].com/[redacted]

The pilot tool will be available until February 28, 2021.

Happy Shopping!

[redacted]

Figure 11: Recruitment email sent to partner platform's customer base

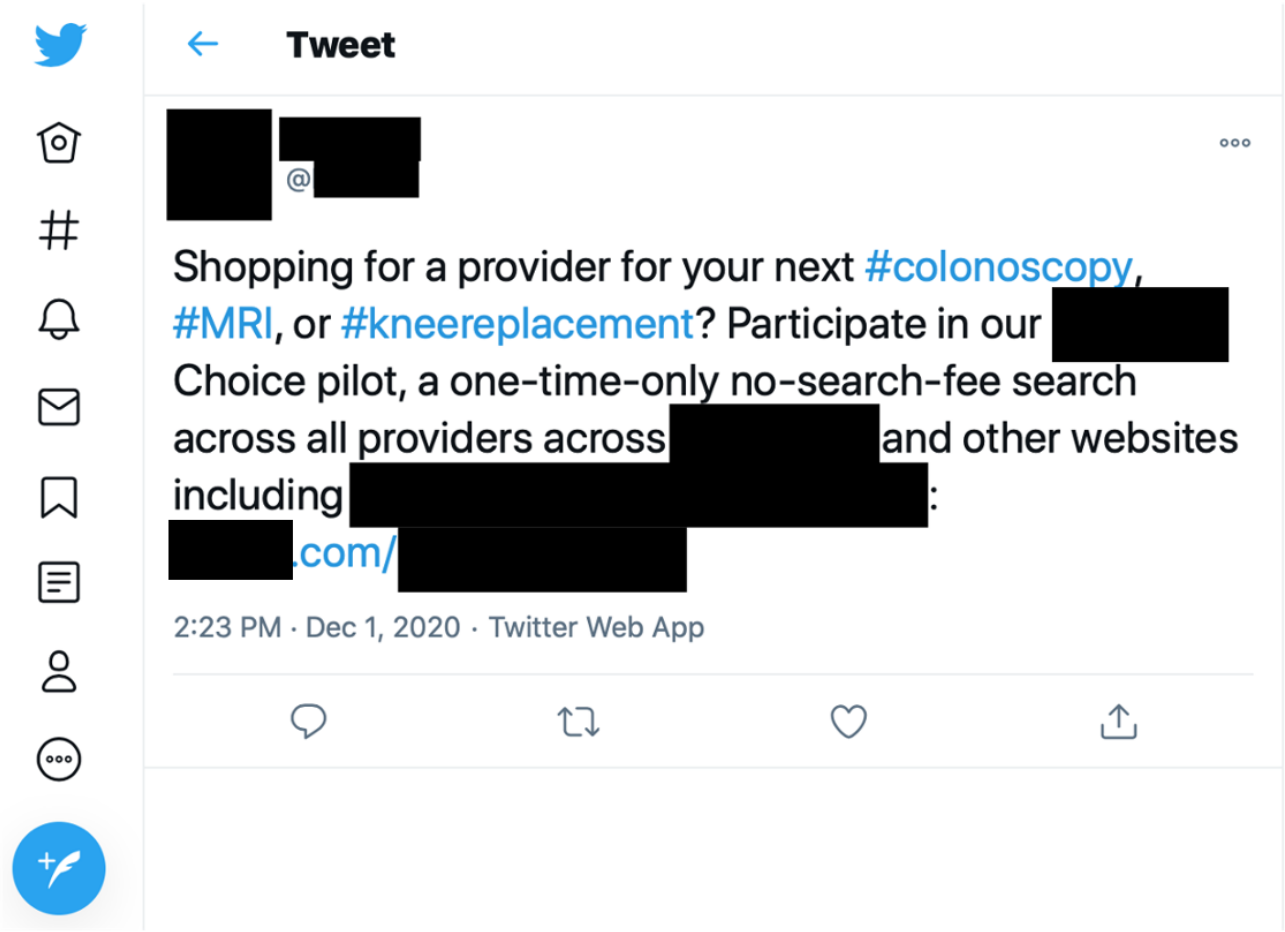


Figure 12: Recruitment "tweet" sent from partner platform's Twitter account

A.3 Experiment Design

Here, I outline the experiment design including the entry and exit surveys as well as the conjoint instruments. The entry survey provides information on the "pilot" shopping tool, and gather information from the customers on things like the procedure that is being shopped for, insurance status, and outside options. After the entry survey, the customers who will be redirected to the conjoint survey and randomized (via Bernoulli draws as subjects arrive) into either the treatment or control group. After the completion of the conjoint survey, the customer will complete an exit survey where demographic information is gathered. Questions about the customers' own demographic information like race and gender are deliberately delayed until after the customer have evaluated all the doctor profiles in the conjoint survey.

This appendix Section should serve to provide additional details about the actual

experience of subjects of this experiment as they shop via the "pilot" tool. I will share the screens broken down into the following four subsections.

A.3.1 Entry Survey

I constructed the survey tool using Qualtrics software for customers who clicked the URL to participate in the "pilot" to shop for their provider to access from a web browser. Upon opening the survey link, respondents must initial on the instructions page (see Figure 13) to continue.

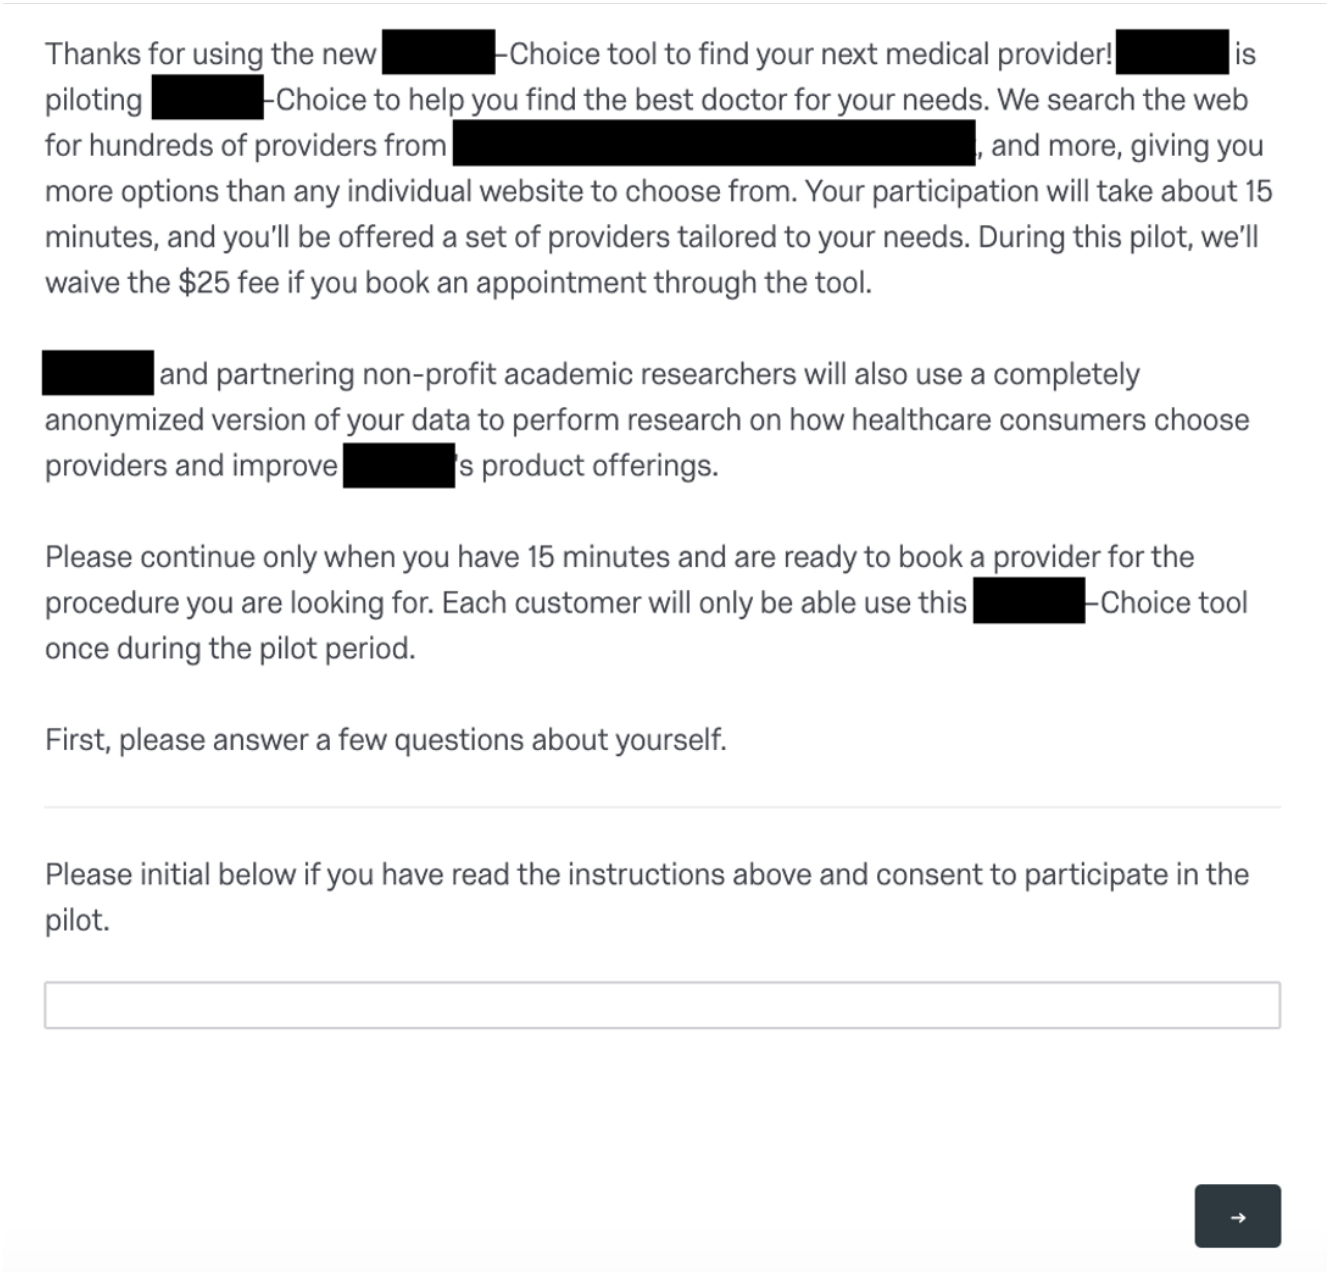


Figure 13: Instruction page

Then, the subjects are asked to answer a series of questions listed below before they are redirected to the conjoint survey itself. These question gather information about the customer on (in the order of appearance): source of information about the pilot (Figure 19), location (Figure 15), age (Figure 16), education (Figure 17), marital status (Figure 18), employment status (Figure 19), insurance status (Figure 20), method of payment (Figure 21), and what procedure is being shopped for (Figure 22).

If the customer selected the option to indicate that they are shopping for a procedure other than colonoscopy, MRI, or knee replacement, they will be re-directed to a screen to indicate that they are not eligible for the "pilot" and then redirected to the regular shopping landing page on the online platform's website.

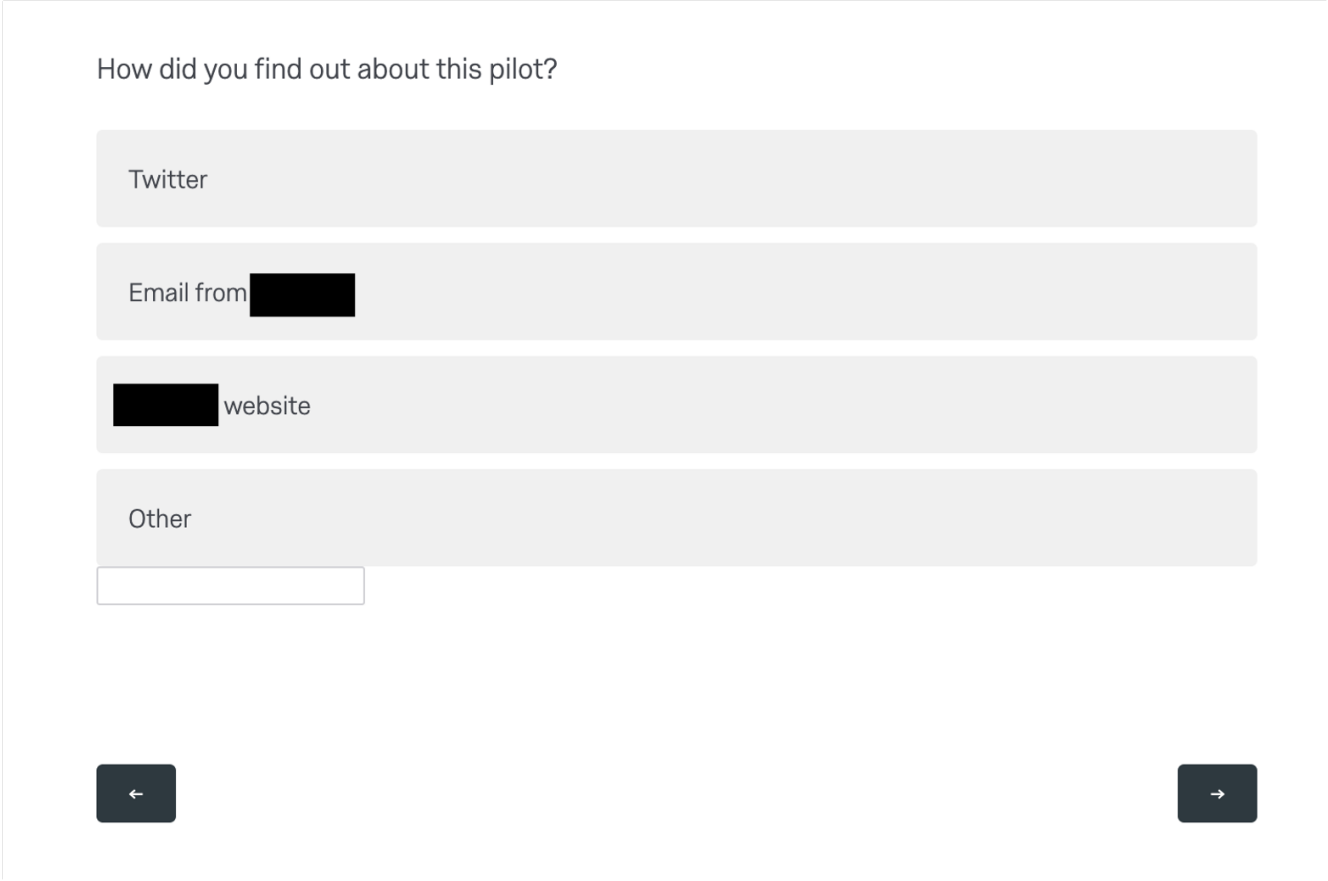


Figure 14: Entry survey question about how the subject heard about the "pilot"

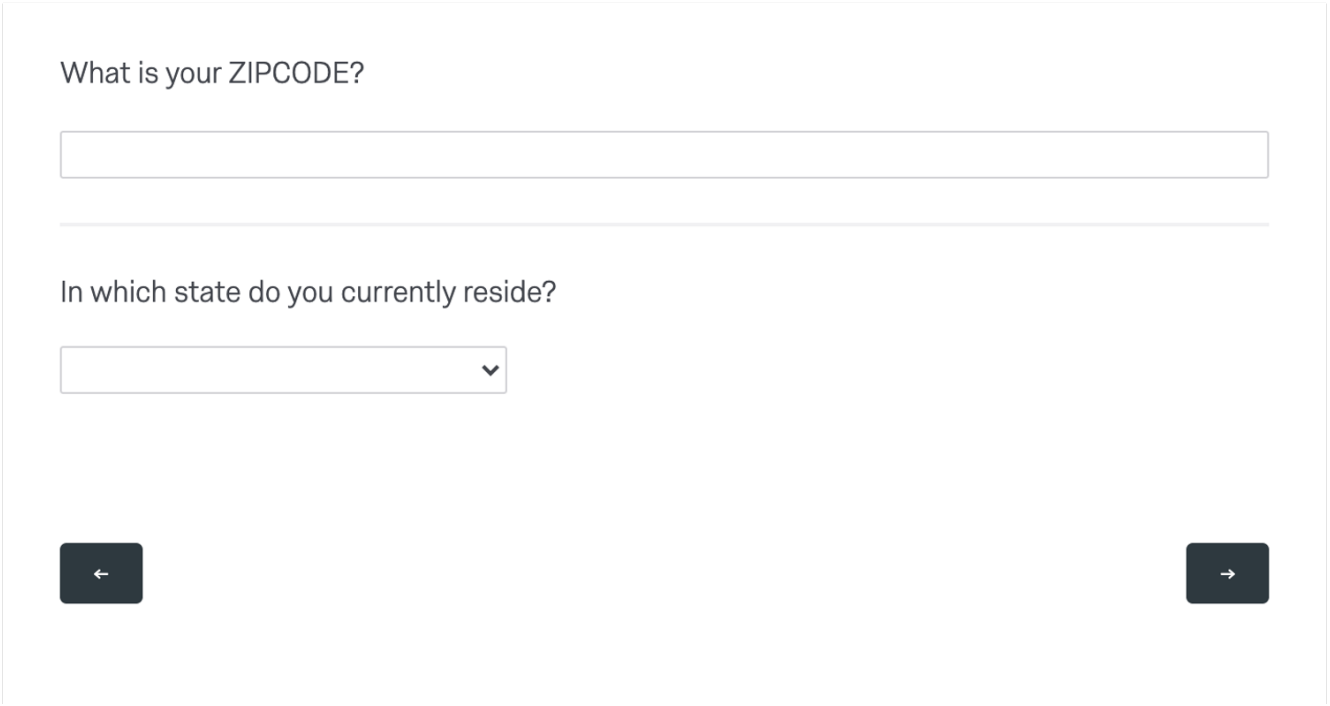


Figure 15: Entry survey question about how the subject's address

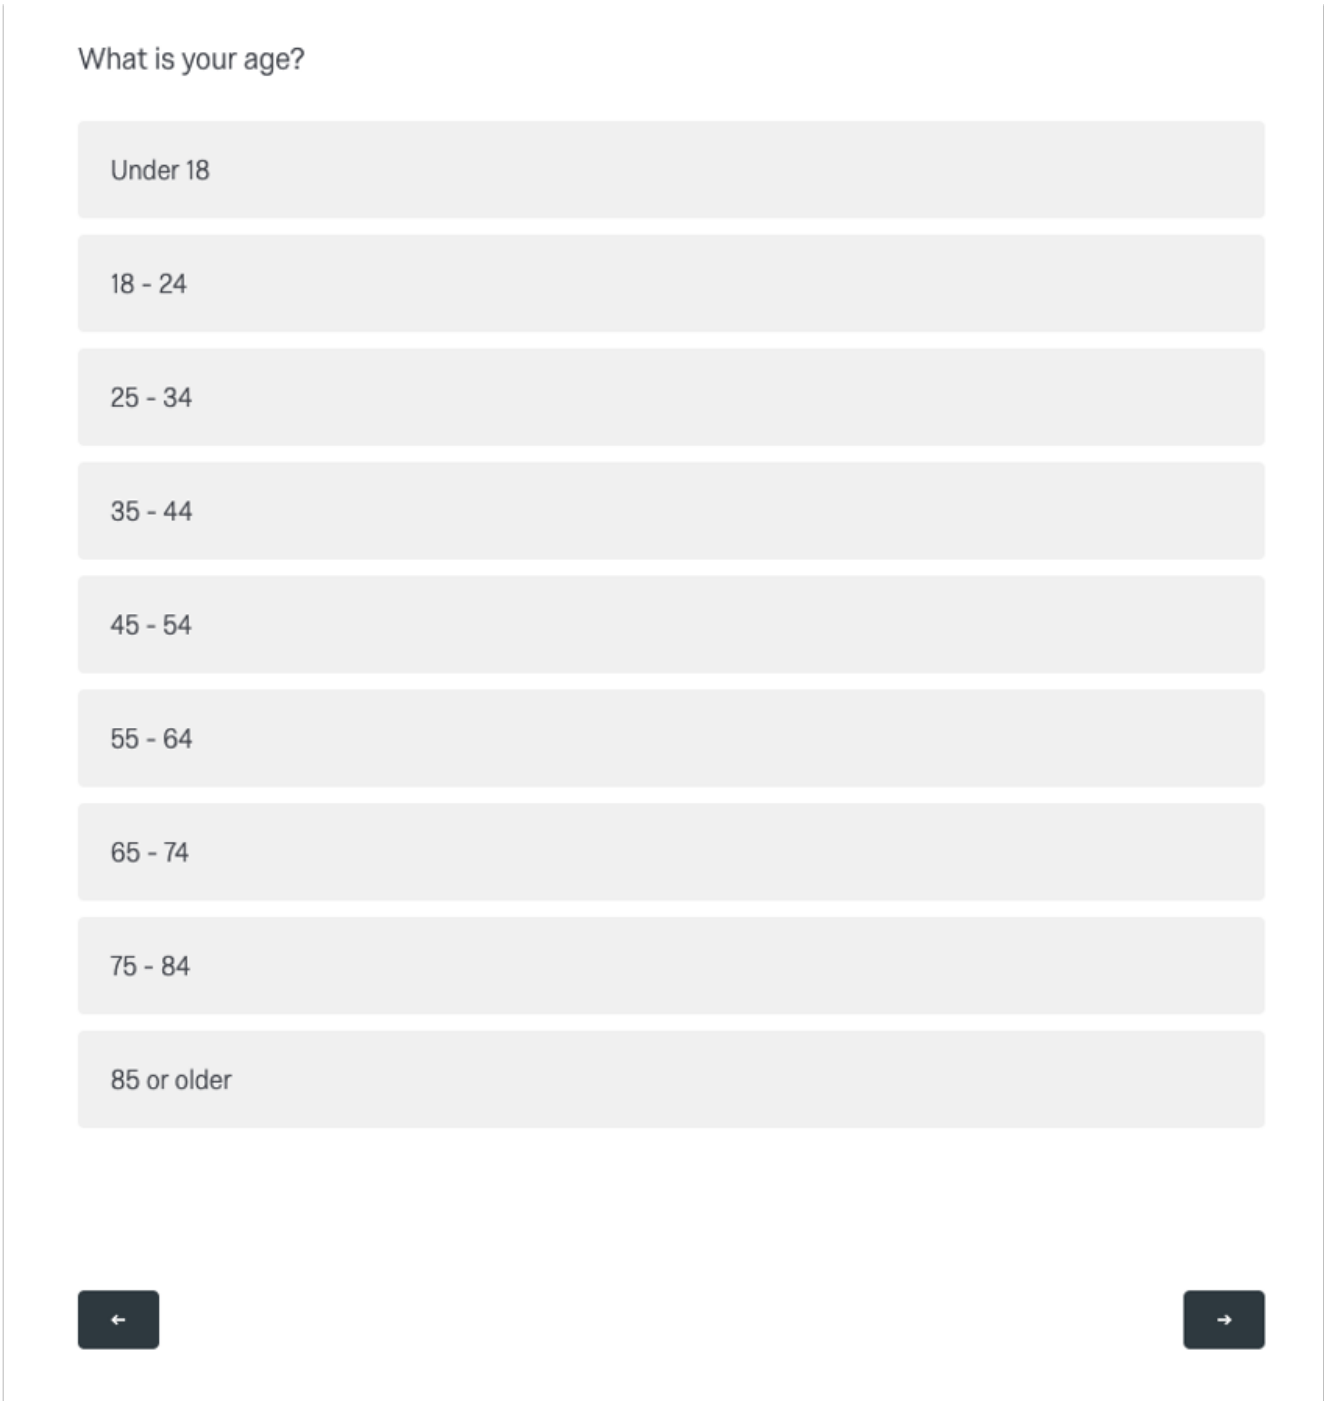


Figure 16: Entry survey question about how the subject's age

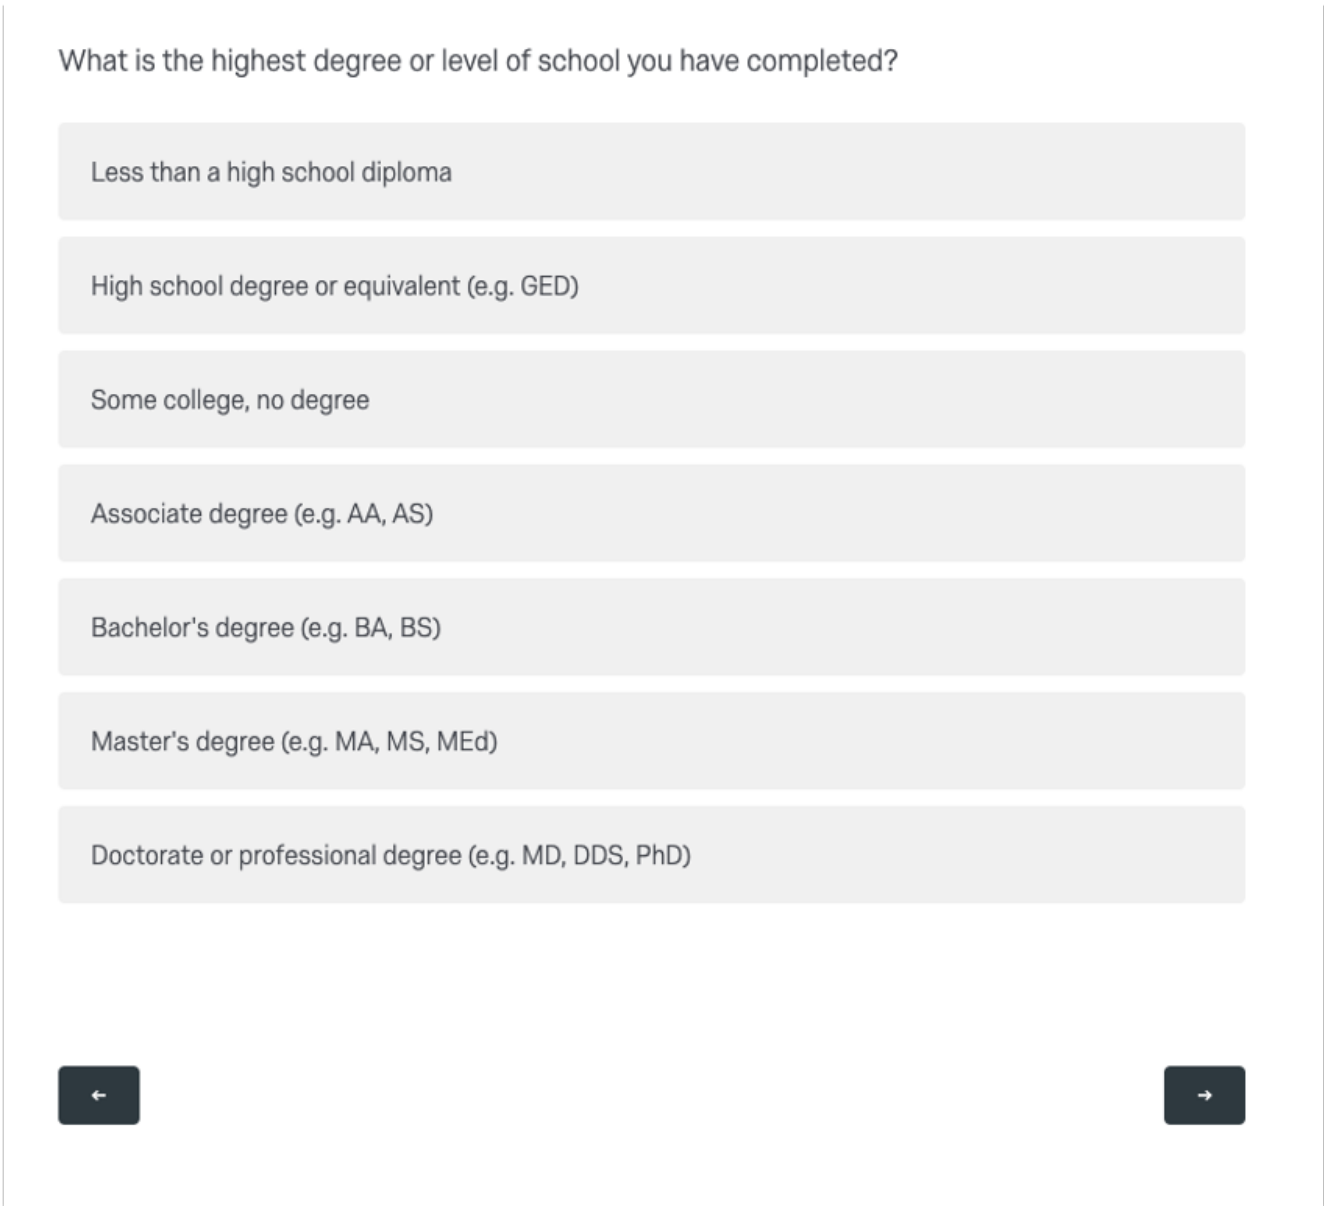


Figure 17: Entry survey question about how the subject's highest level of education

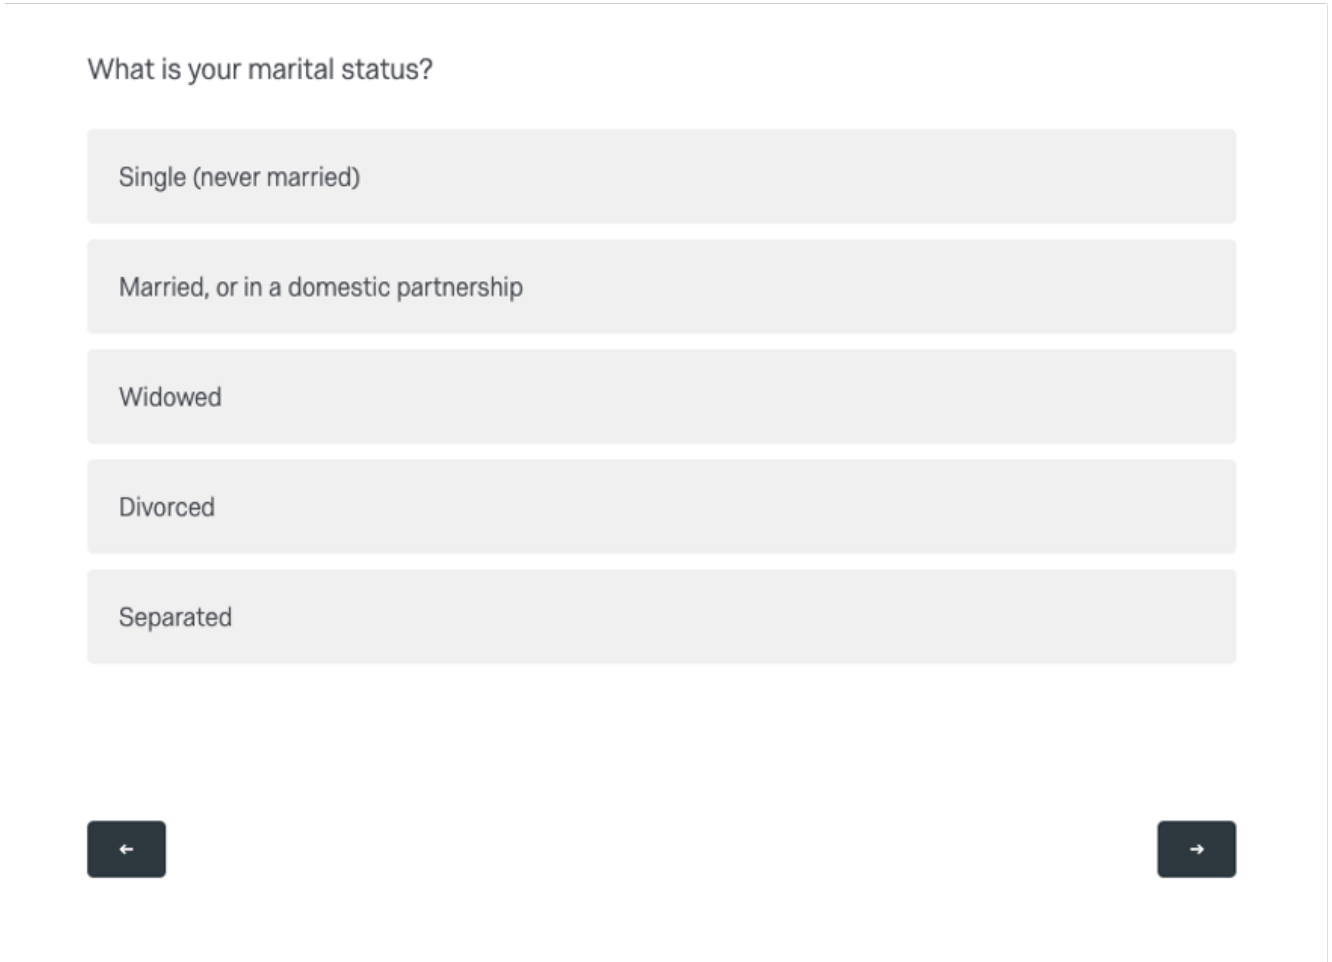


Figure 18: Entry survey question about how the subject's marital status

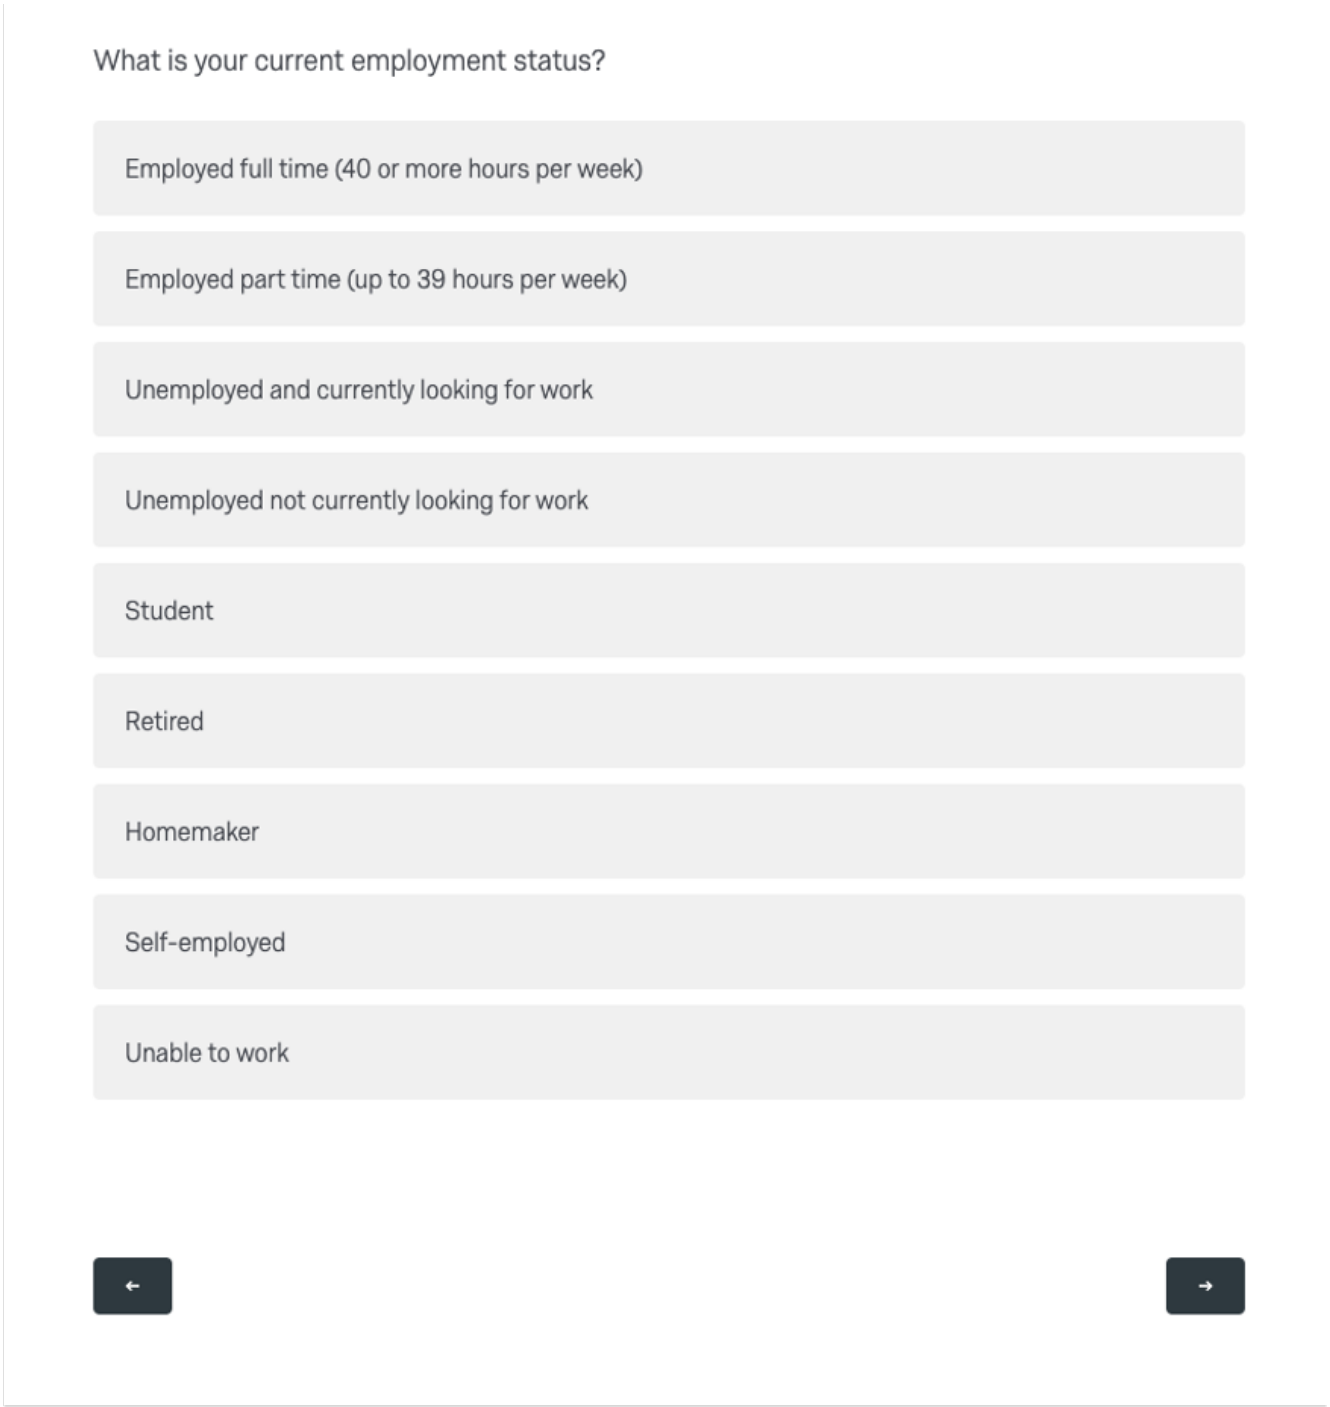


Figure 19: Entry survey question about how the subject's employment status

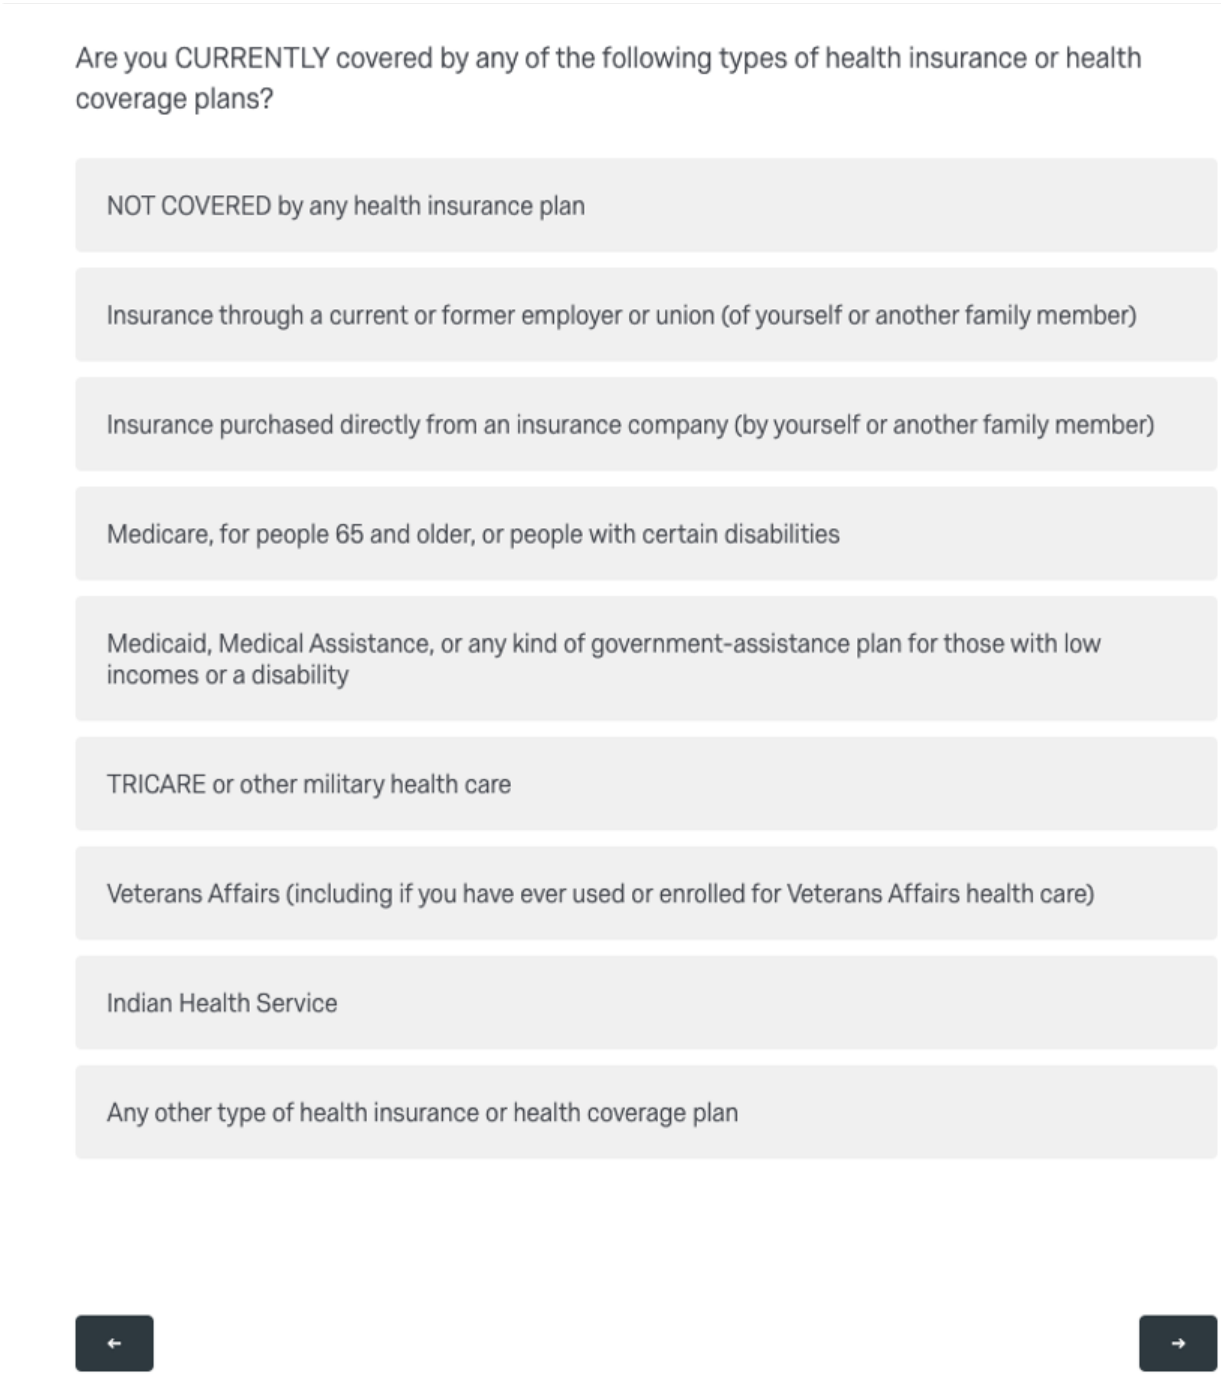


Figure 20: Entry survey question about how the subject's insurance status

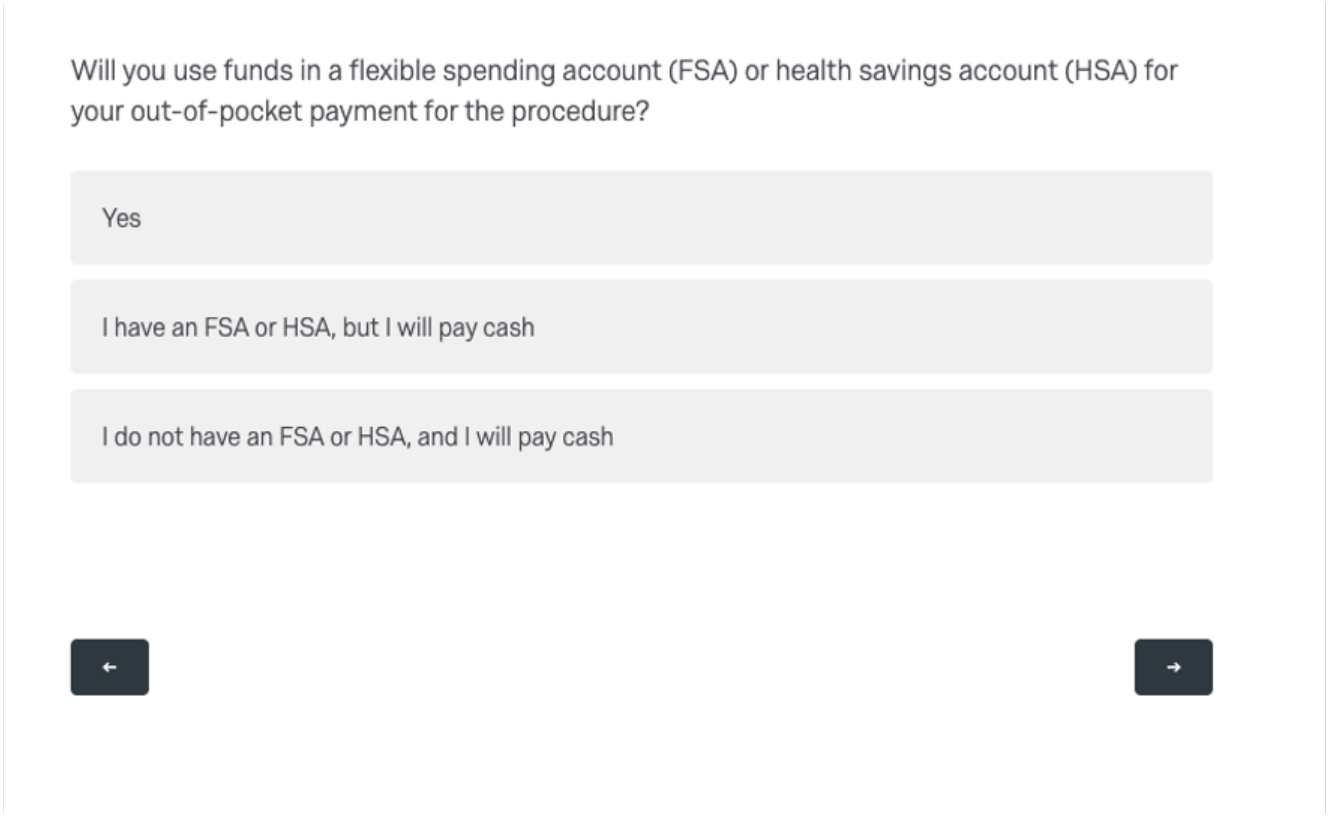


Figure 21: Entry survey question about the subject's method of payment

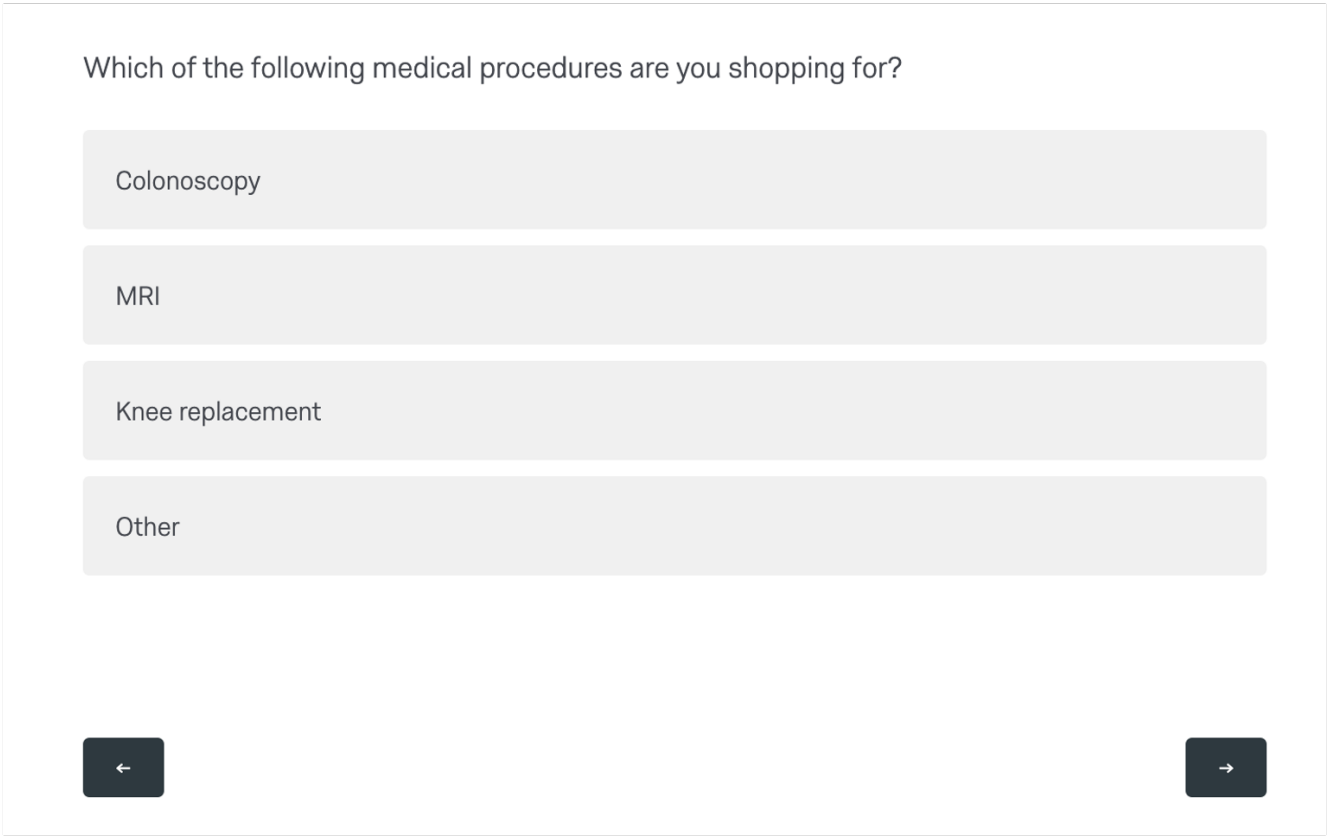


Figure 22: Entry survey question about which procedure the subject is shopping for

The subjects in the experiment in this present paper are the ones who had selected "colonoscopy" as the medical procedure that they are shopping for. They will then answer a few questions before getting redirected to the conjoint survey itself.

These last few questions inquire the subjects about their outside options: the time/effort used to identify a provider for the colonoscopy prior to using the "pilot" tool (Figure 23), number of options explored outside of the "pilot" (Figure 24), and the expected price and distance of the customer's outside options (Figure 25). After these questions, the subjects are shown a screen that describes the nature of the conjoint survey and explained how the responses to the conjoint survey maps to actual provider options that will be shared with them shortly after the completion of the survey (Figure 26).

There are customers who participated in the "pilot" to shop for either an MRI or a knee replacement. As mentioned in the main text of the paper, I do not use that data in the present study. The MRI customer data is used for a separate paper, while there is only one customer who used the pilot for knee replacement.

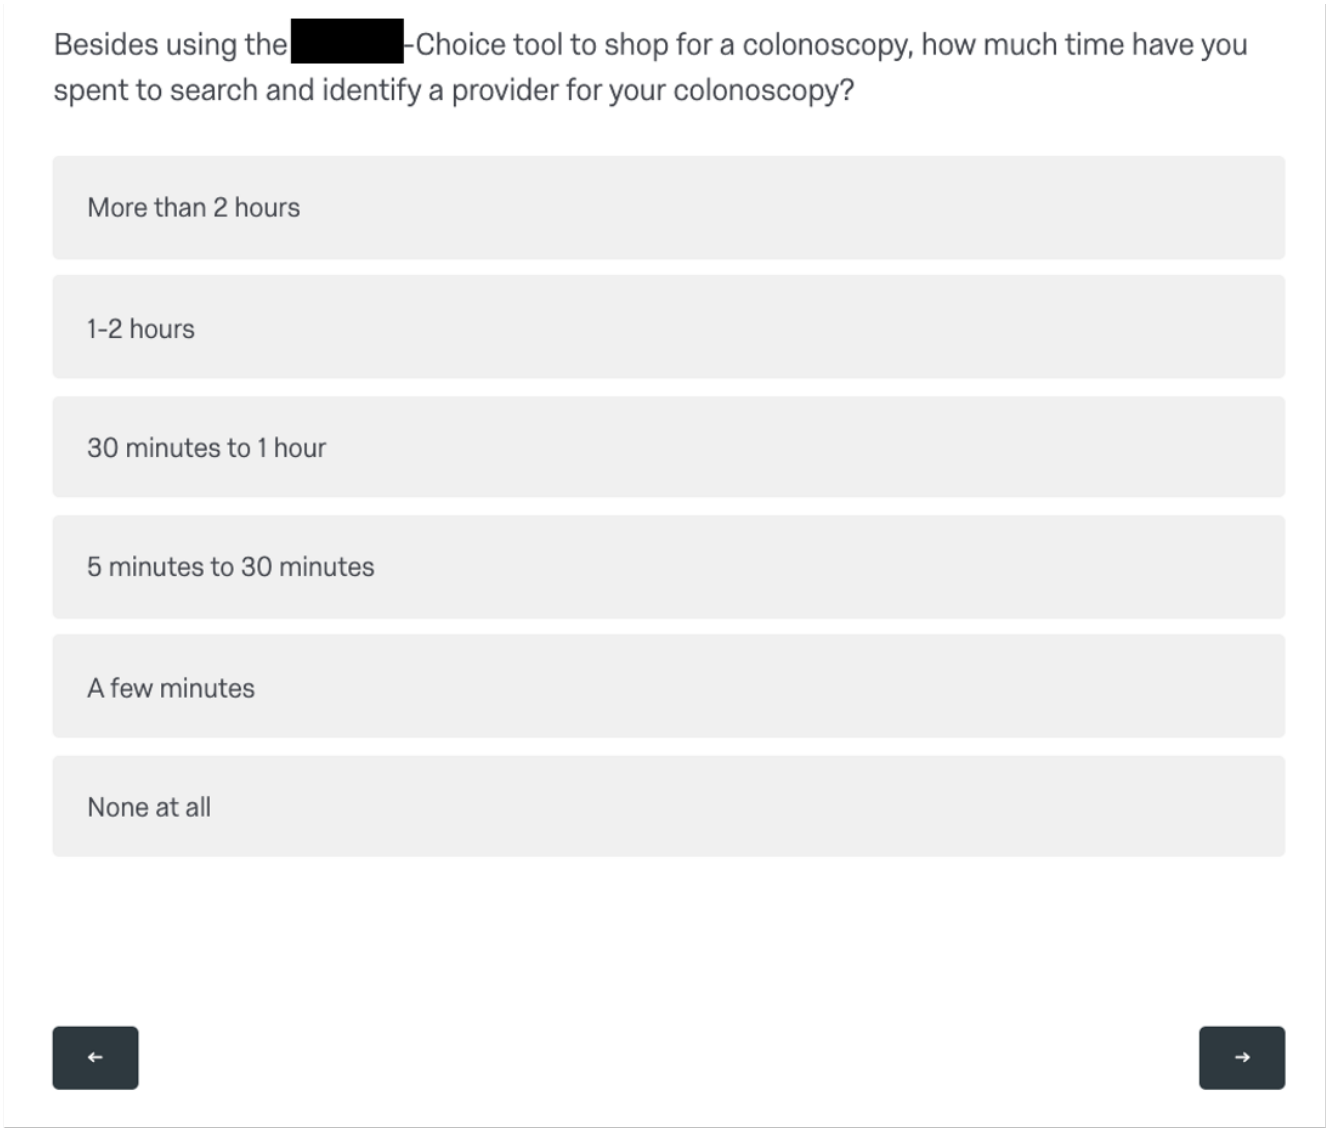


Figure 23: Entry survey question about how the subject's search effort before using the "pilot" tool

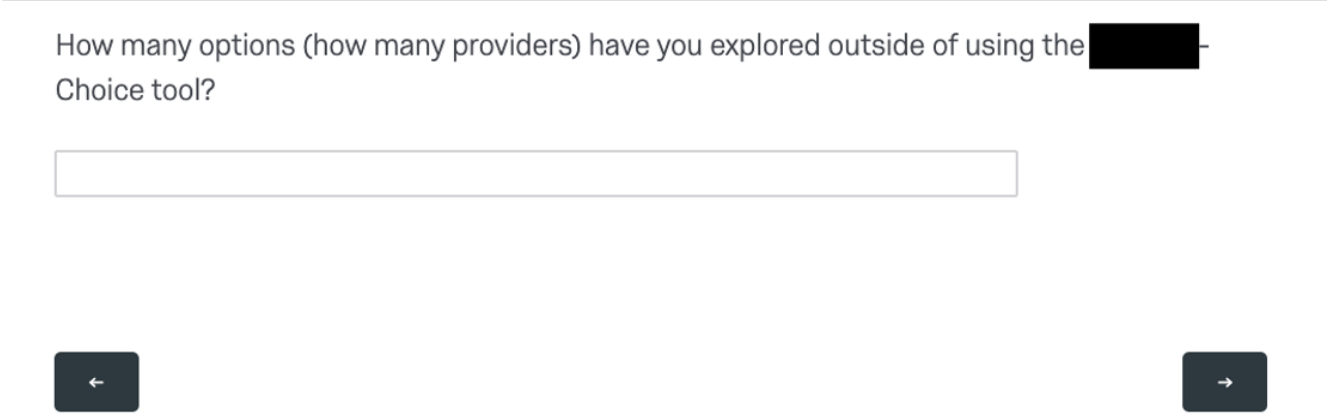


Figure 24: Entry survey question about the number of doctor options explored outside of the "pilot"

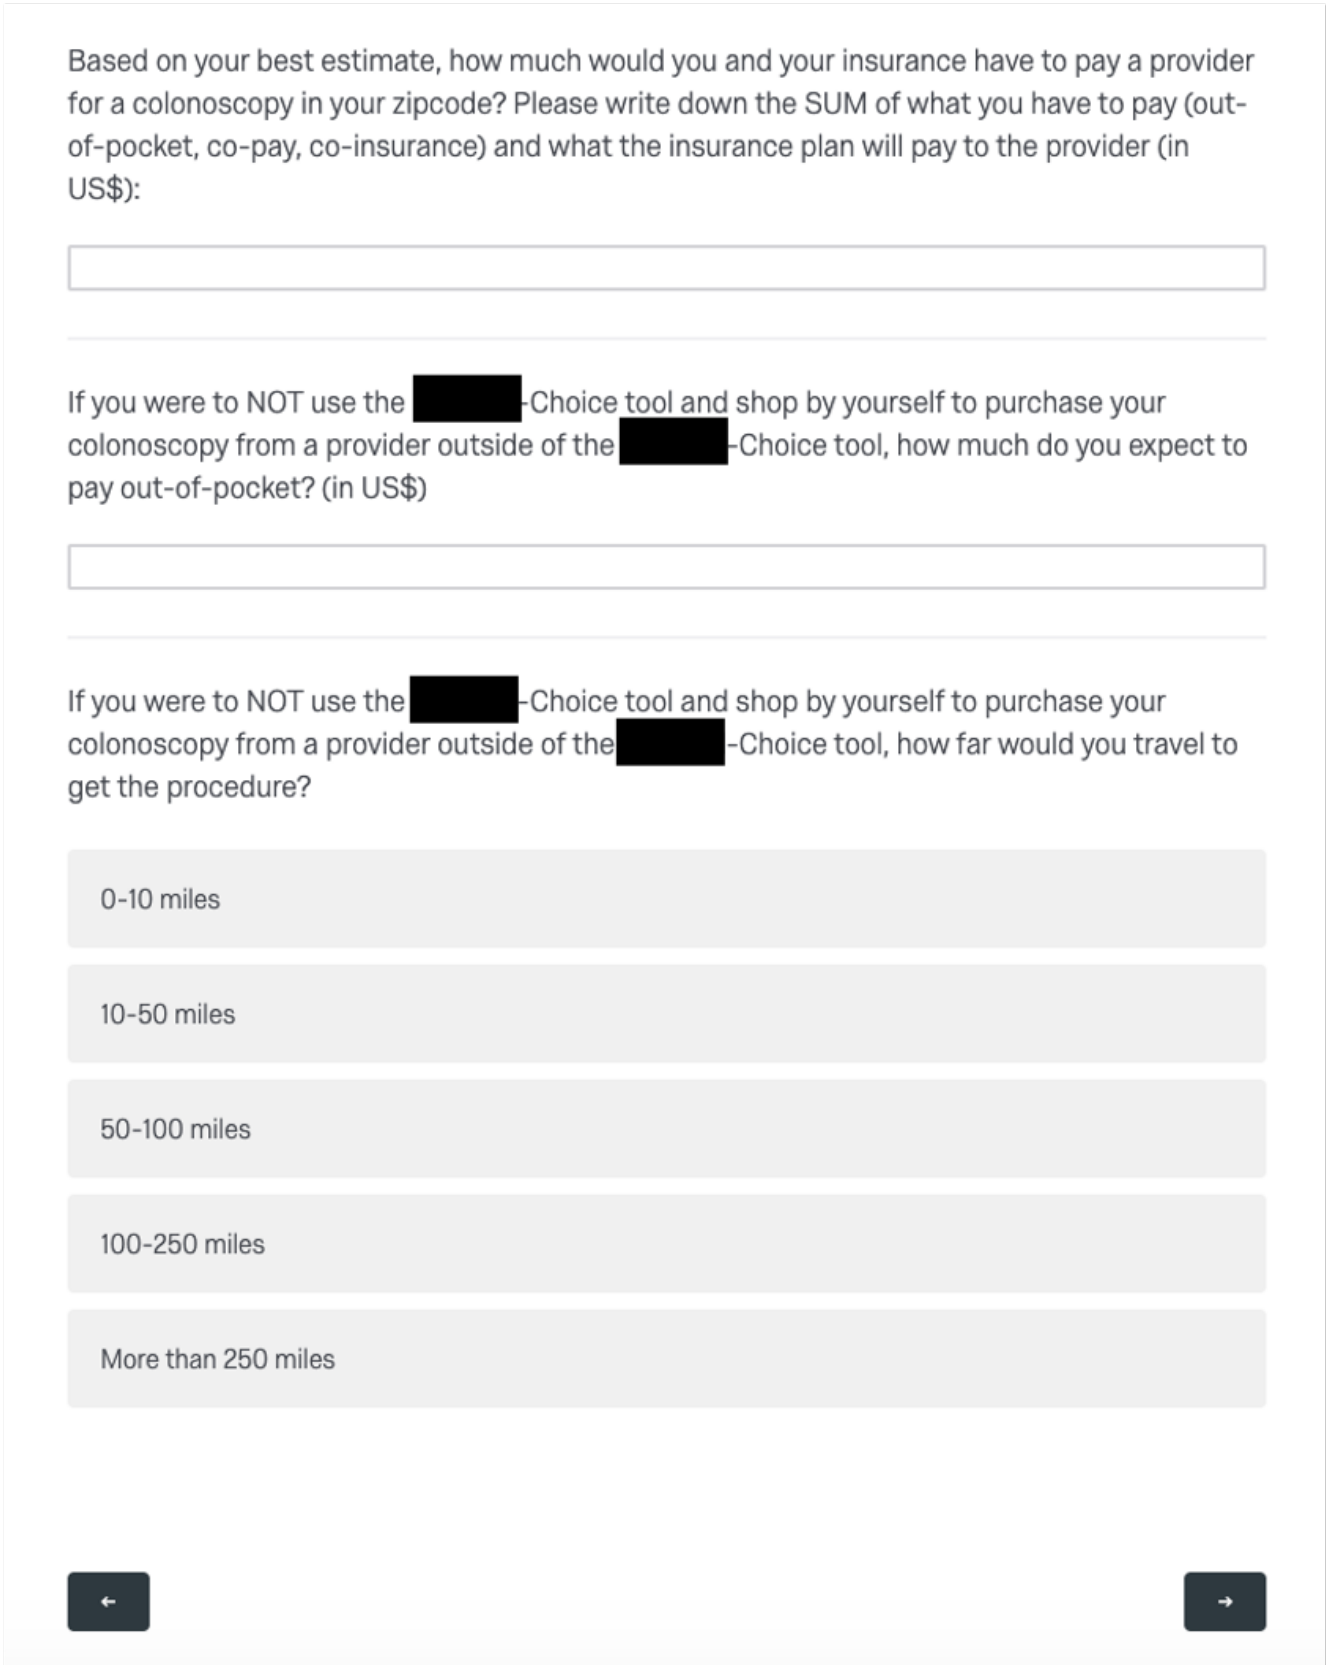


Figure 25: Entry survey question about the doctor options explored outside of the "pilot" (e.g. price and distance of the outside option)

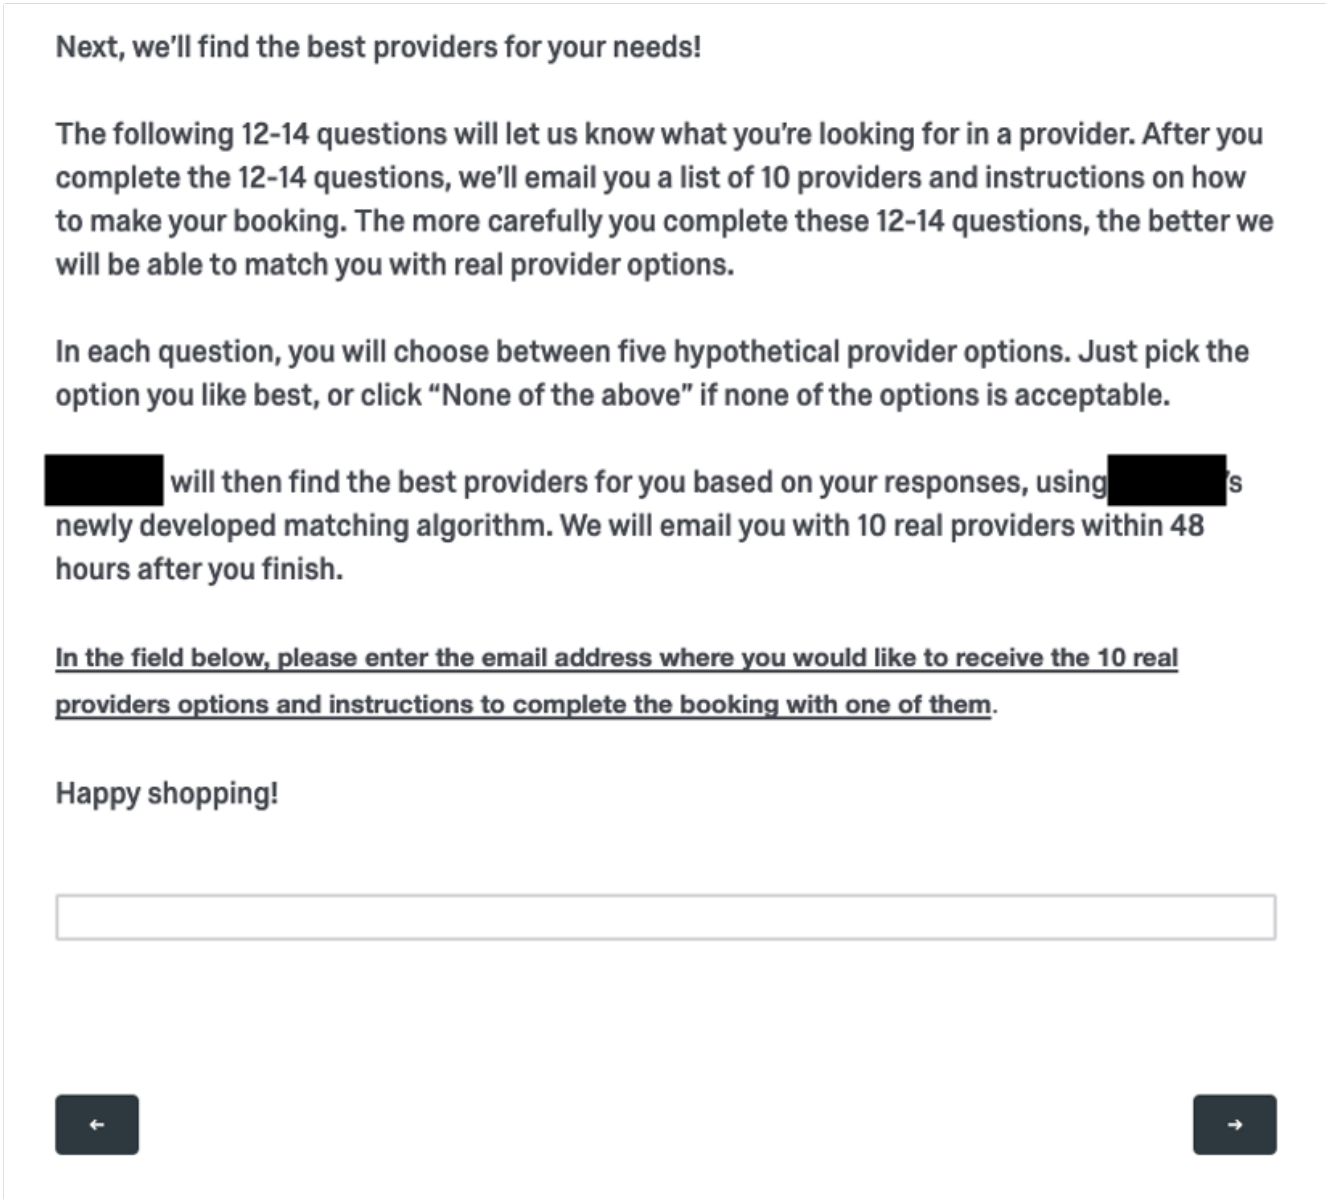


Figure 26: Entry survey screen to record email and reiterate instructions of the "pilot"

A.3.2 Conjoint Survey

The conjoint survey where customers evaluated 14 hypothetical doctors generates the data for the main analysis in the paper.

There are two slightly different version of the conjoint survey, and customers will only see one version. Which version of the conjoint survey the customer would see depends on whether they are randomized in the treatment or the control group.

The conjoint survey constitute screens where customer are informed about the doctor attributes they will see on the doctor profiles, followed by the 14 menus, and then the exit survey (described in the next subsection).

I present the conjoint survey screen-by-screen as a subject would experience it. If the screens for the treatment and control groups differ, I will show the control group version first followed by the treatment group version before moving to the next screen.

First, both treatment arms will see an introduction screen (Figure 27), followed by a screen introducing the attributes of the doctor options in the menus the customer would see. The one for the control group (Figure 28) is different from the one for the treatment group (Figure 29), as the latter has one more attribute than the former.

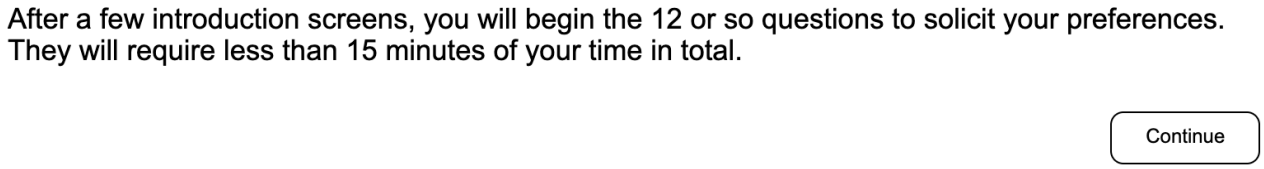


Figure 27: Launch screen for conjoint survey

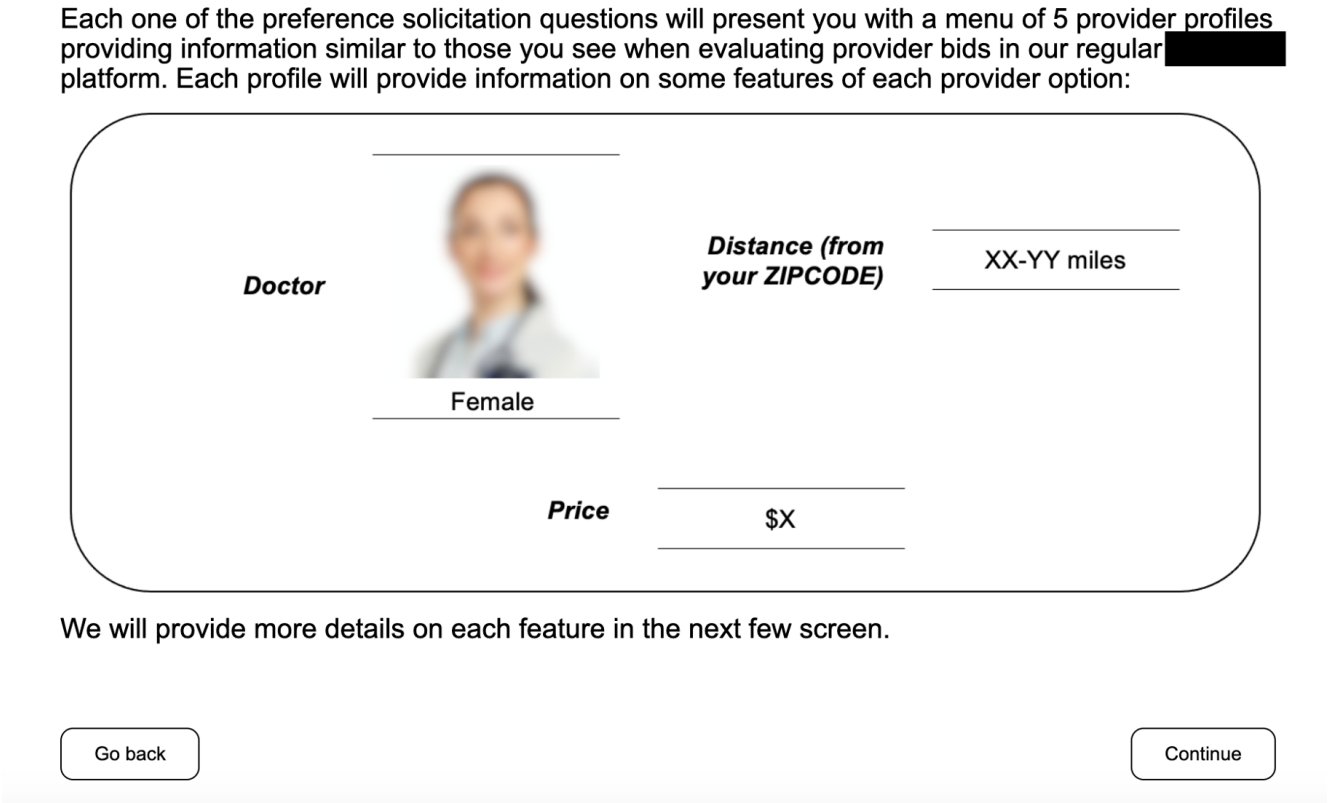


Figure 28: Screen introducing attributes for control group

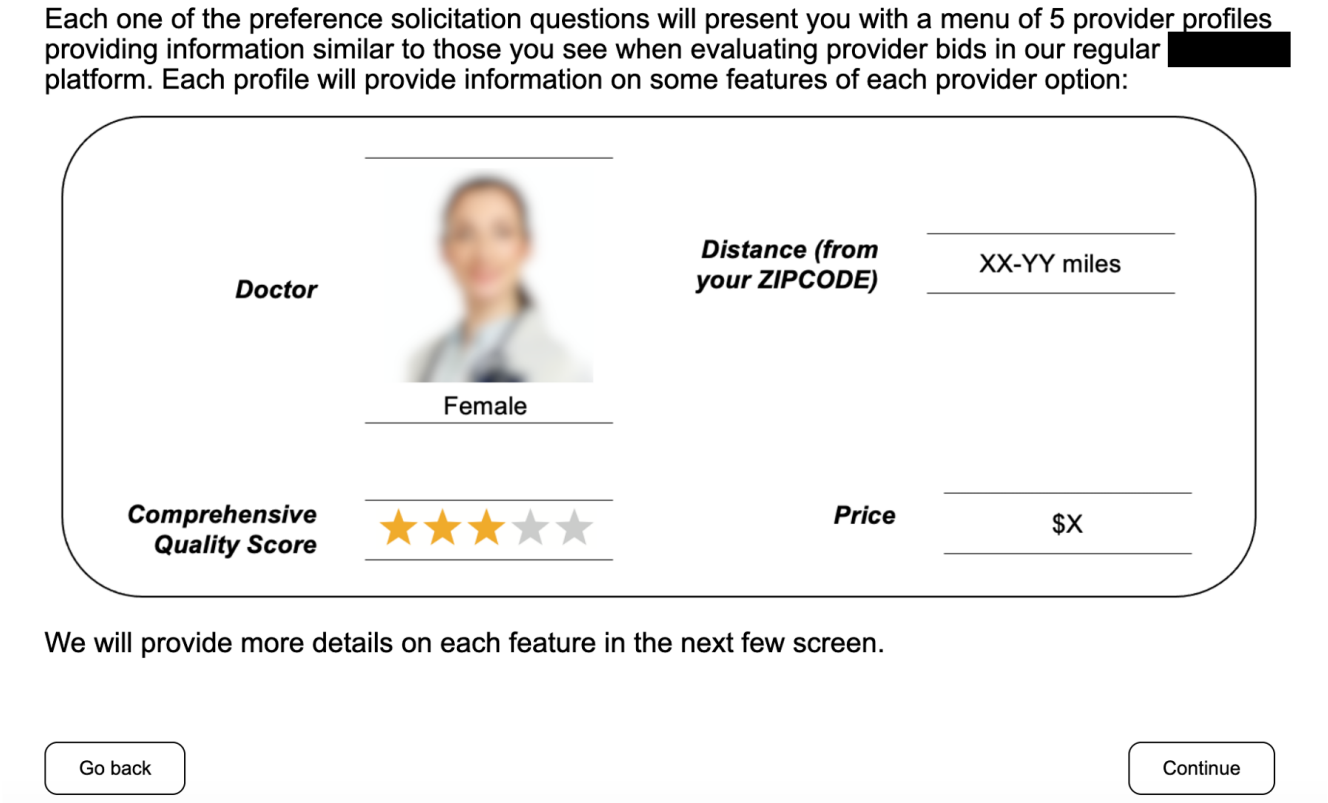


Figure 29: Screen introducing attributes for treatment group

Information about each attribute will be presented with its own screen. The order with which these screens are presented is randomized across subjects (e.g. one subject might see the information about price first, distance second and the doctor profile picture and gender last but another subject might see distance first and other attributes in a different order). Before the subject can move on from one screen to the next, they have to pass an attention test by picking the correct color (yellow, green, and blue). The ordering of these colors (ordering of the buttons) are randomized across subjects and screens.

Both treatment and control group subjects will see a screen clarifying the doctor attributes of "price," "doctor" (which include a doctor profile picture and gender), and "distance from your ZIPCODE." The treatment group will also see an additional screen clarifying the doctor attribute of "Comprehensive Quality Score," the quality signal.

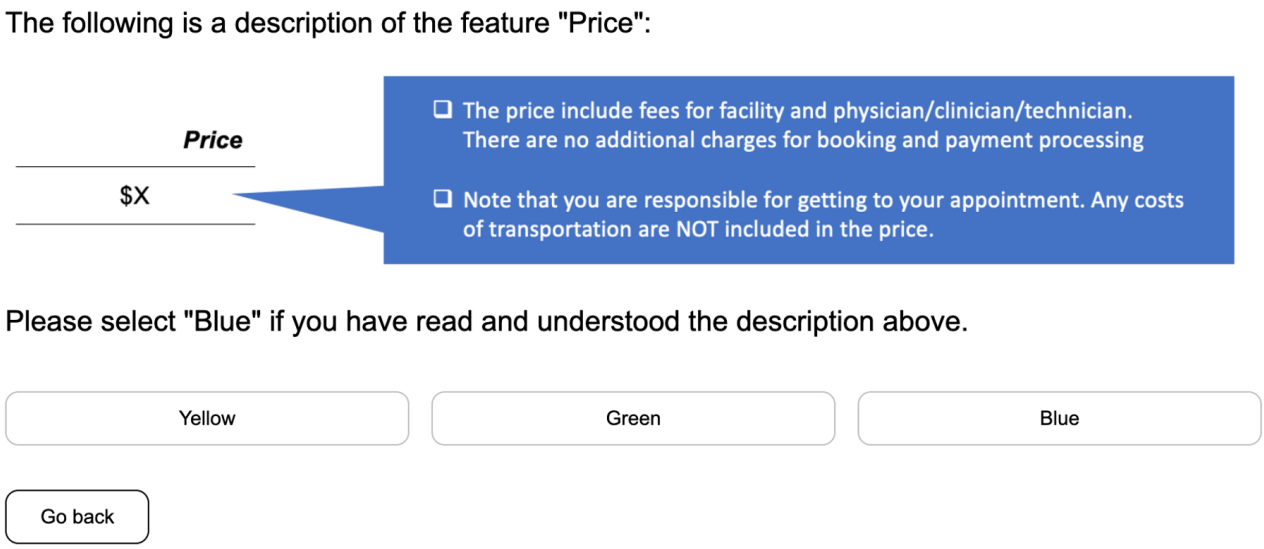


Figure 30: Information screen for doctor attribute of "Price"

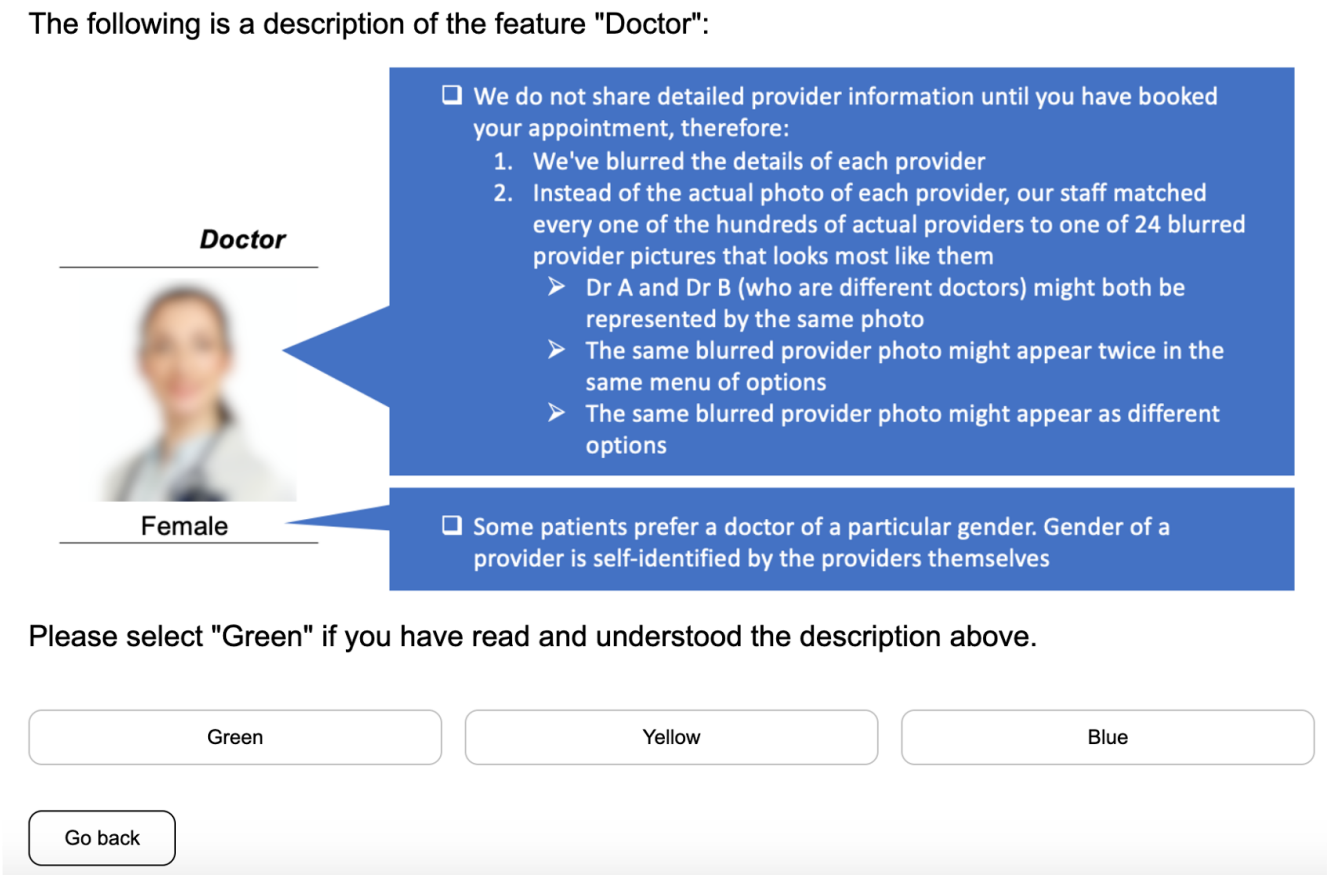


Figure 31: Information screen for doctor attribute of "Doctor": including the profile picture and gender

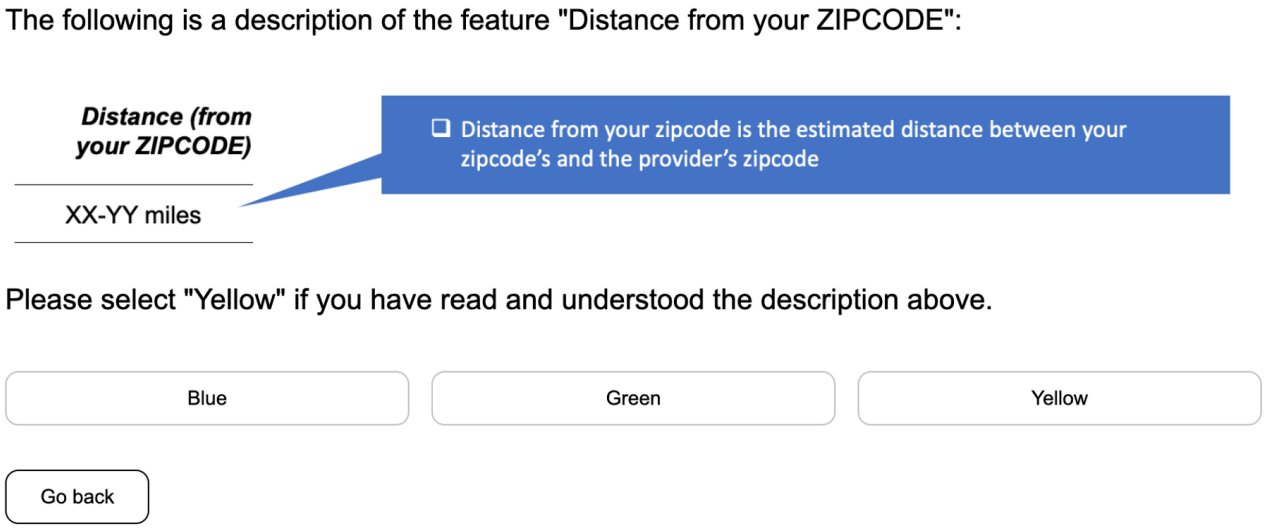


Figure 32: Information screen for doctor attribute of Travel Distance

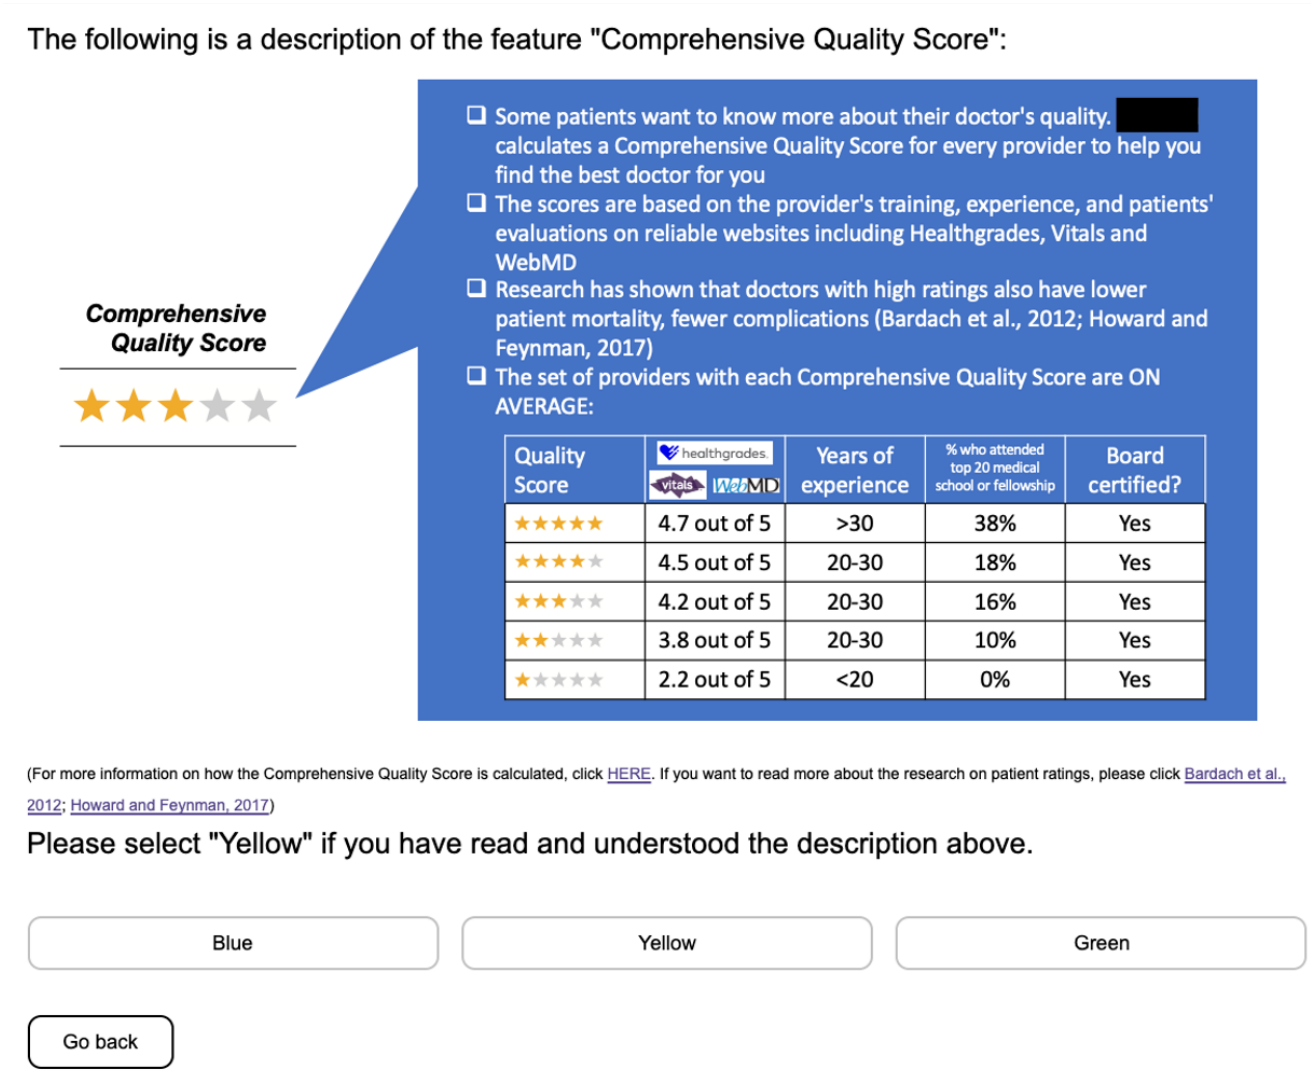


Figure 33: Information screen for doctor attribute of "Comprehensive Quality Score"

After the screens to help customers comprehend the doctor attributes, a screen alerts the customers that the "preference solicitation questions" will start, this is the last screen before the customers evaluate the 14 menus of hypothetical doctor profiles (see Figure 34).

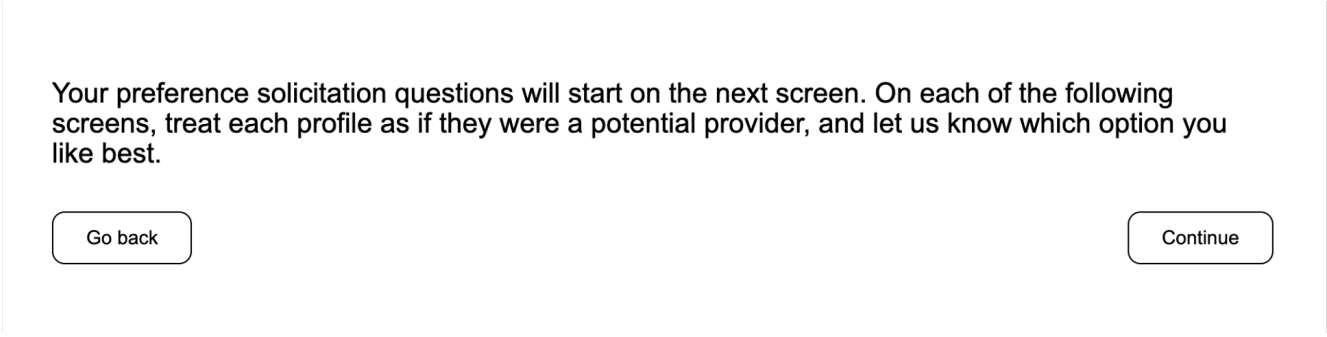


Figure 34: Screen alerting subjects that conjoint survey is about to start

In the key part of the conjoint survey (and VIC), the customers evaluate 14 menus. Each menu will have 5 doctor options and a button for "None of the above." The ordering of the attributes from top to bottom of a doctor profile is randomized across subjects but price is always presented as the bottom attribute. The ordering of attributes from top to bottom is the same within each subject's set of 14 menus. An example of the control group's menus is presented in Figure 35. An example of the treatment group's menus is presented in Figure 36.

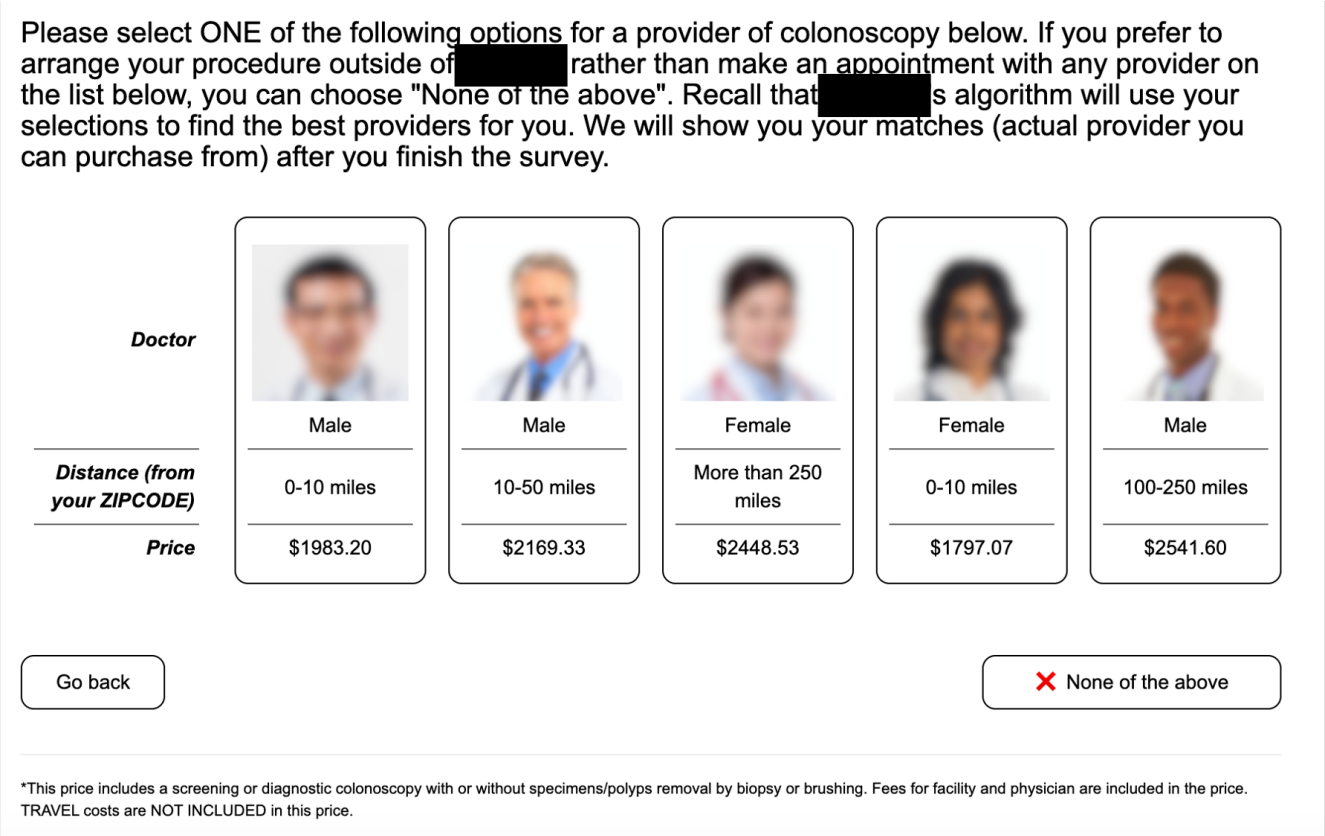


Figure 35: Information screen for doctor attribute of "Comprehensive Quality Score"

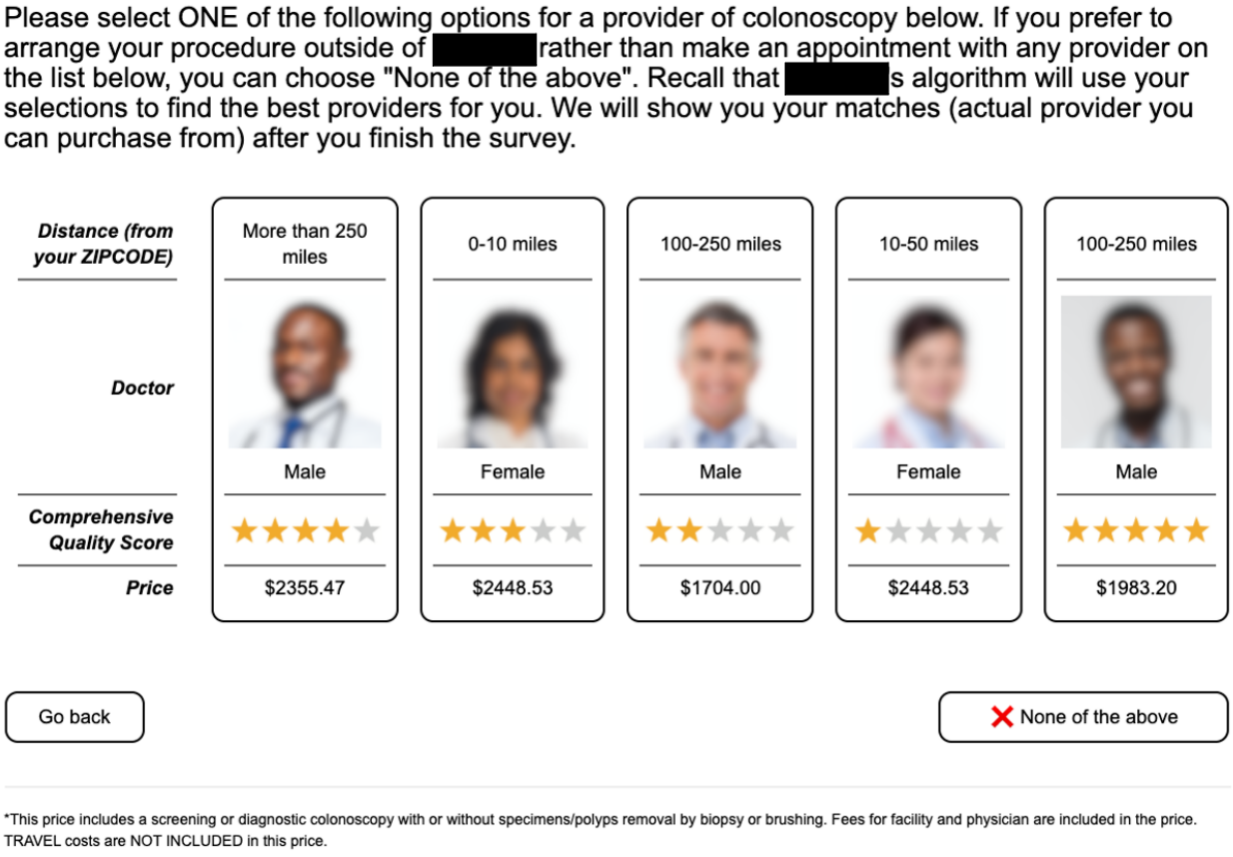


Figure 36: Information screen for doctor attribute of "Comprehensive Quality Score"

A.3.3 Exit Survey

After finishing the evaluation of the 14 menus, each customer is asked to describe their approach in a free text box (Figure 37). The provision of comments is optional.

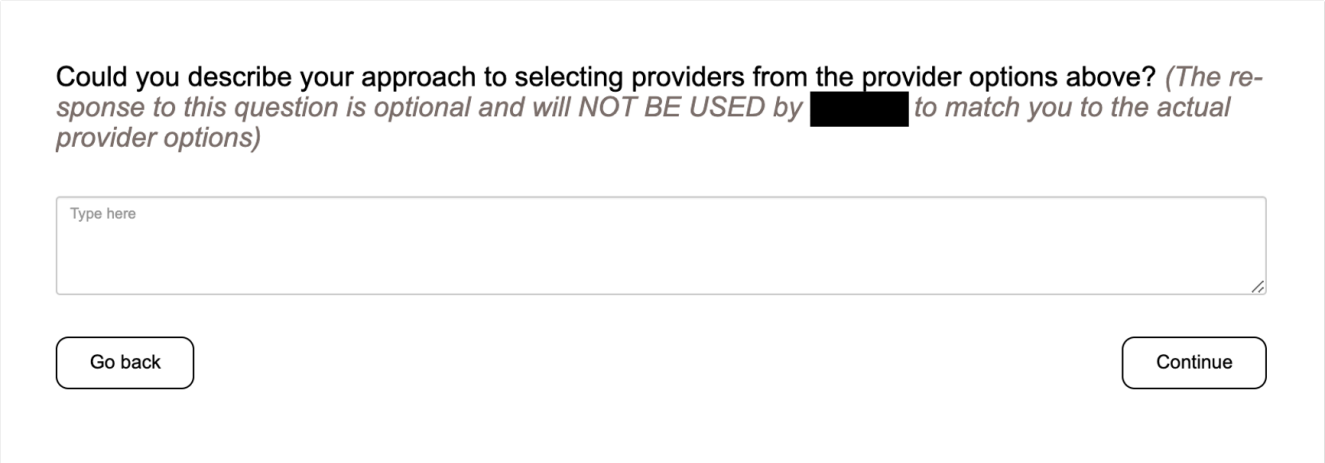


Figure 37: Screen asking subjects for free-text description of their approach

After the comment box, subjects from both treatment and control groups are asked to subjectively evaluate the quality of providers from various distances from their zipcode based on their own experience (Figure 38).

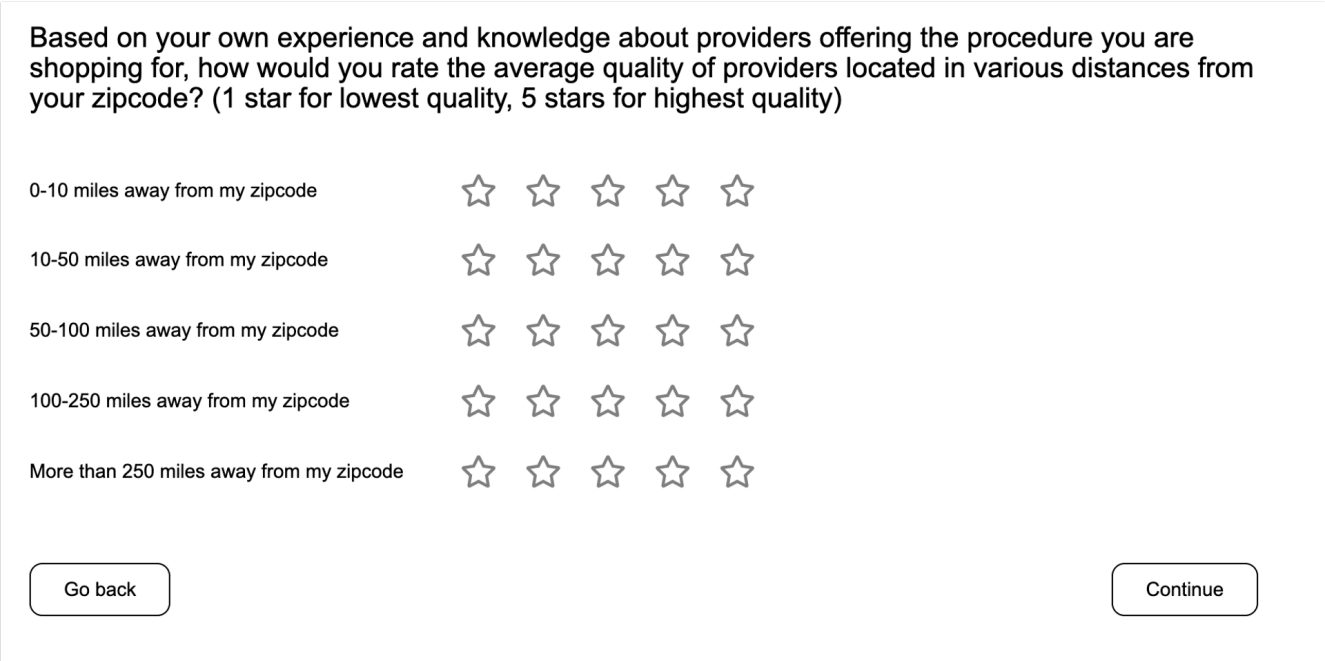


Figure 38: Screen asking customer's expected correlation between distance and quality

For subjects in the treatment group, they are asked 3 Likert scale questions for their opinion on the Comprehensive Quality Score (Figures 39, 40, and 41).

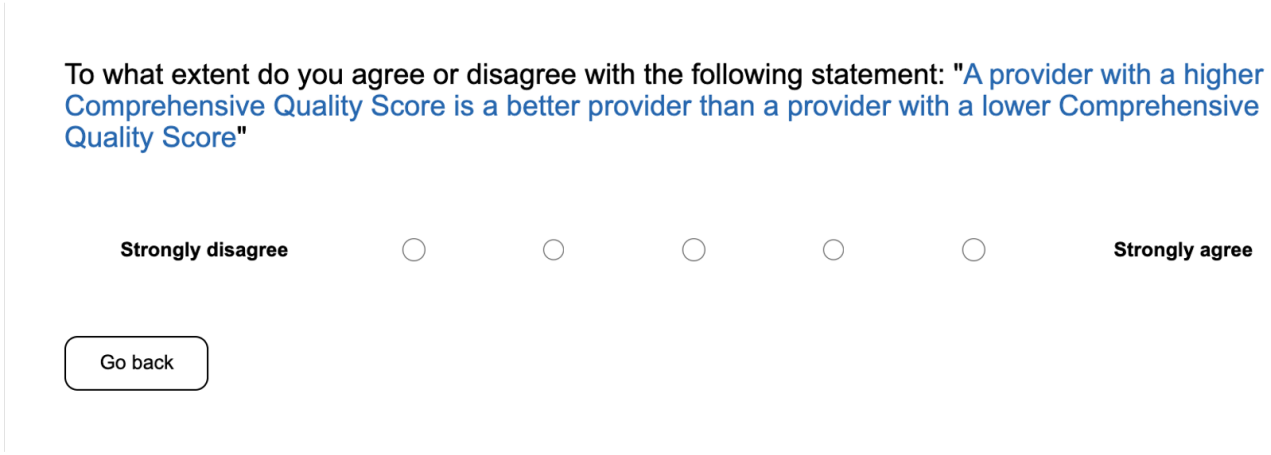


Figure 39: First Likert scale question on quality score for treatment group subjects

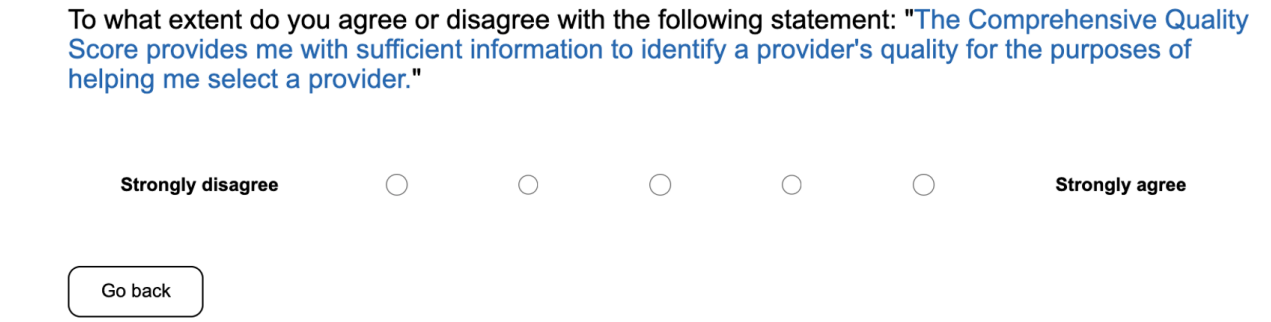


Figure 40: Second Likert scale question on quality score for treatment group subjects

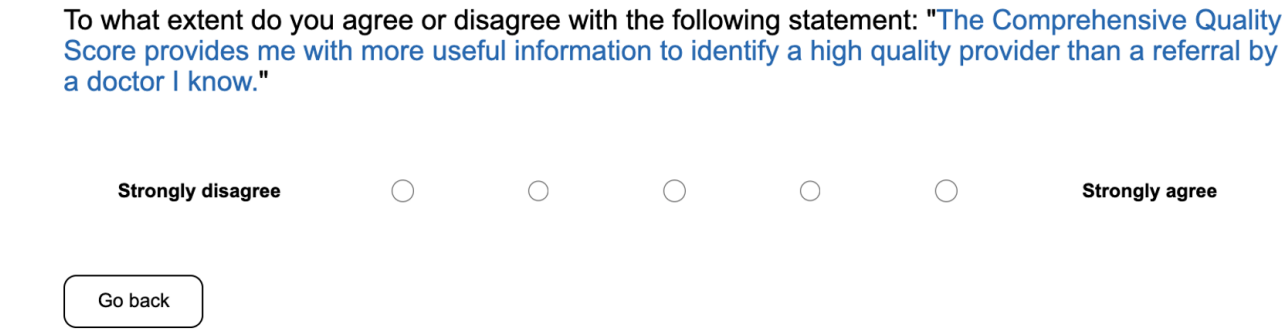


Figure 41: Third Likert scale question on quality score for treatment group subjects

Before the subjects finish the conjoint survey, I gathered information on their gender and race as well. This happen across 3 consecutive screens (Figures 42, 43 and 44). This

information is gathered for both treatment and control group.



Figure 42: Exit survey question about gender

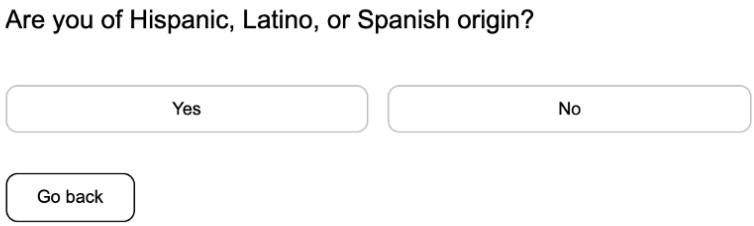


Figure 43: Exit survey question about ethnicity

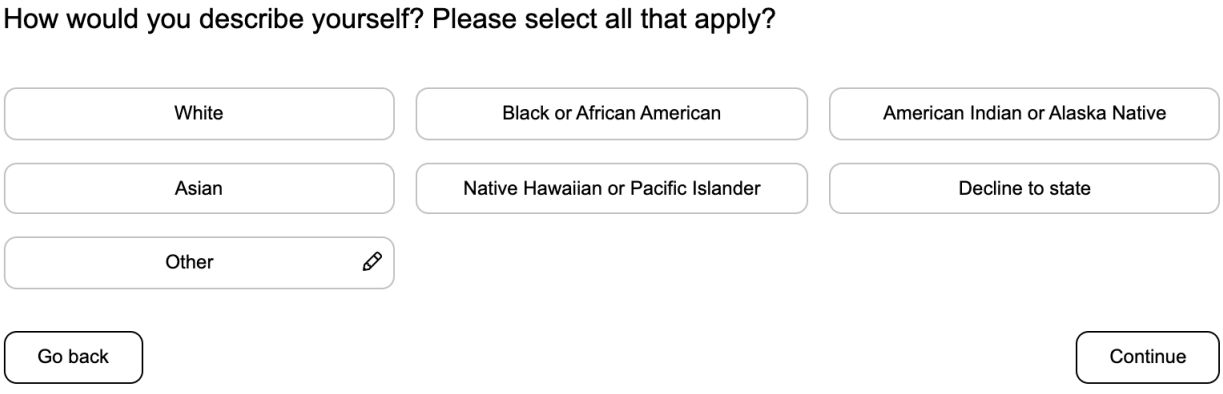


Figure 44: Exit survey question about race

A.3.4 Actual Booking Options

Within 48 hours after each subject finished the entry survey, conjoint survey, and exit survey, they will received a customized email (see Figure 45) listing the 10 actual doctors matched to them. The subjects were told that the offers were only valid for the next 10 business days and have to click on a link to get redirected to an online booking agreement form to complete the booking.

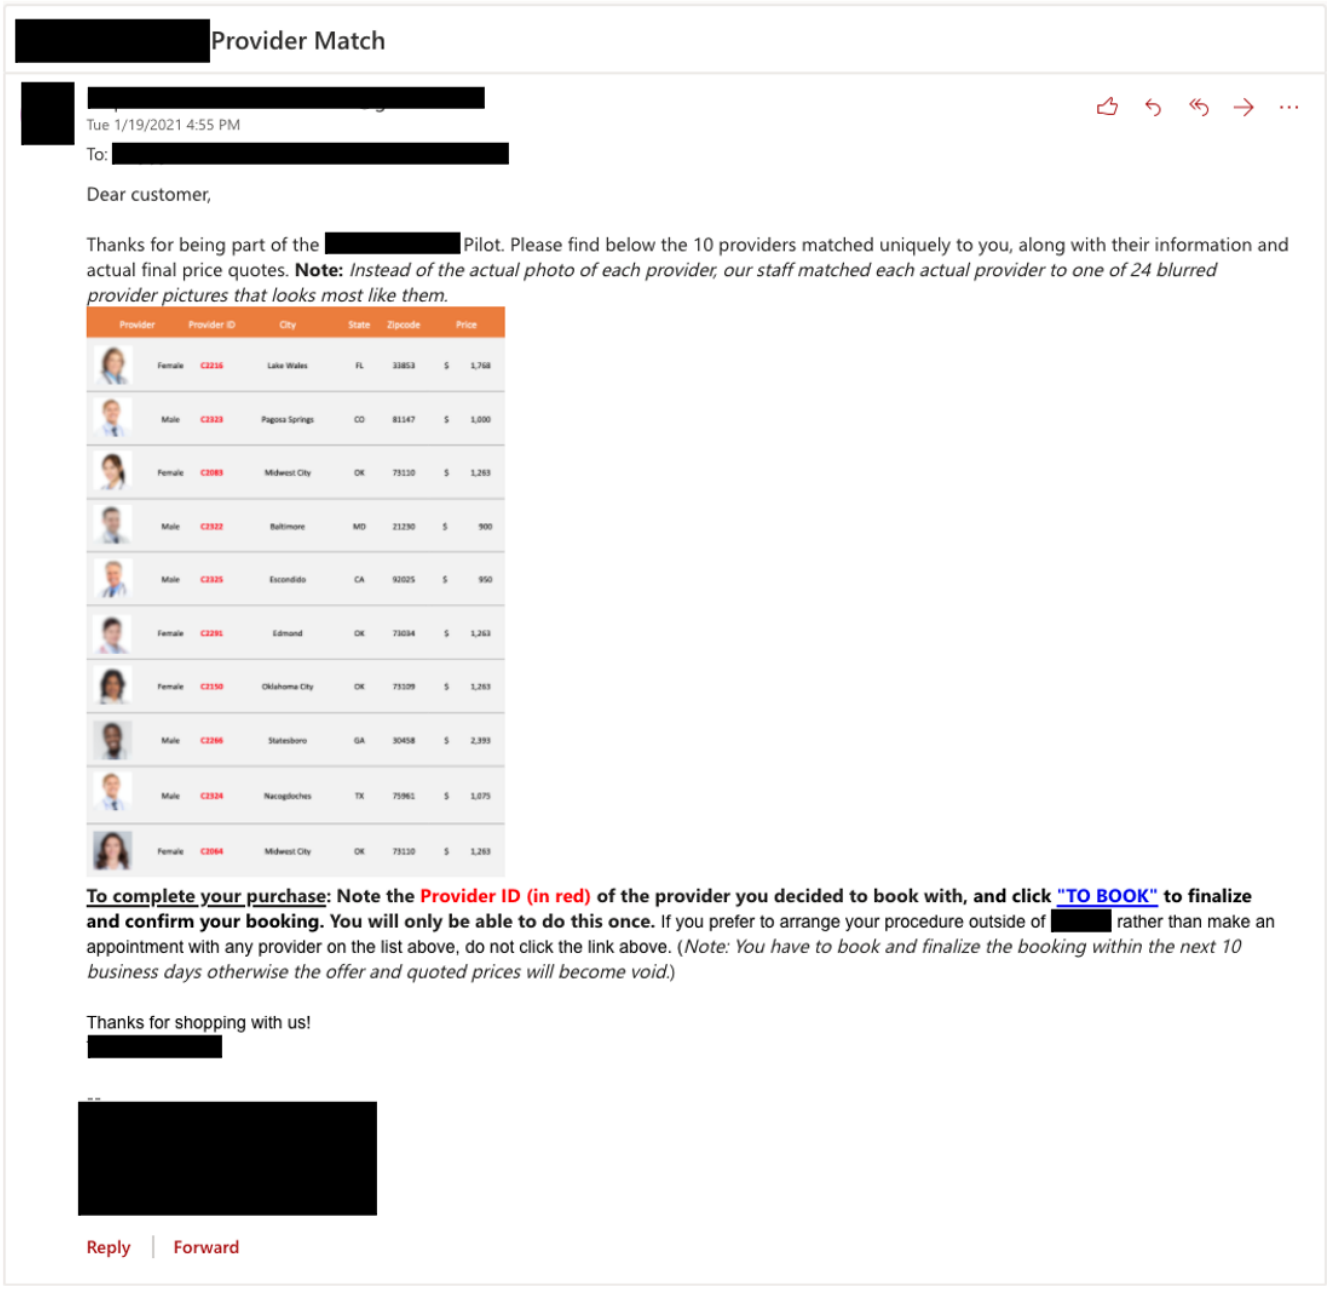


Figure 45: Email to customer with actual doctor options for booking

A.4 Hypothetical Doctor Profile Pictures

A hypothetical doctor profile picture appears as an element on each profile in the menus (Figure 46). I manipulate perceptions of the doctor race by using profile pictures that clearly indicate the race of the hypothetical doctor. The resolution of each profile picture is reduced so that the doctor's race shows through but their attractiveness and other features are obscured. The profile pictures are chosen such that they have a similar degree of facial symmetry across races (see Table 4).

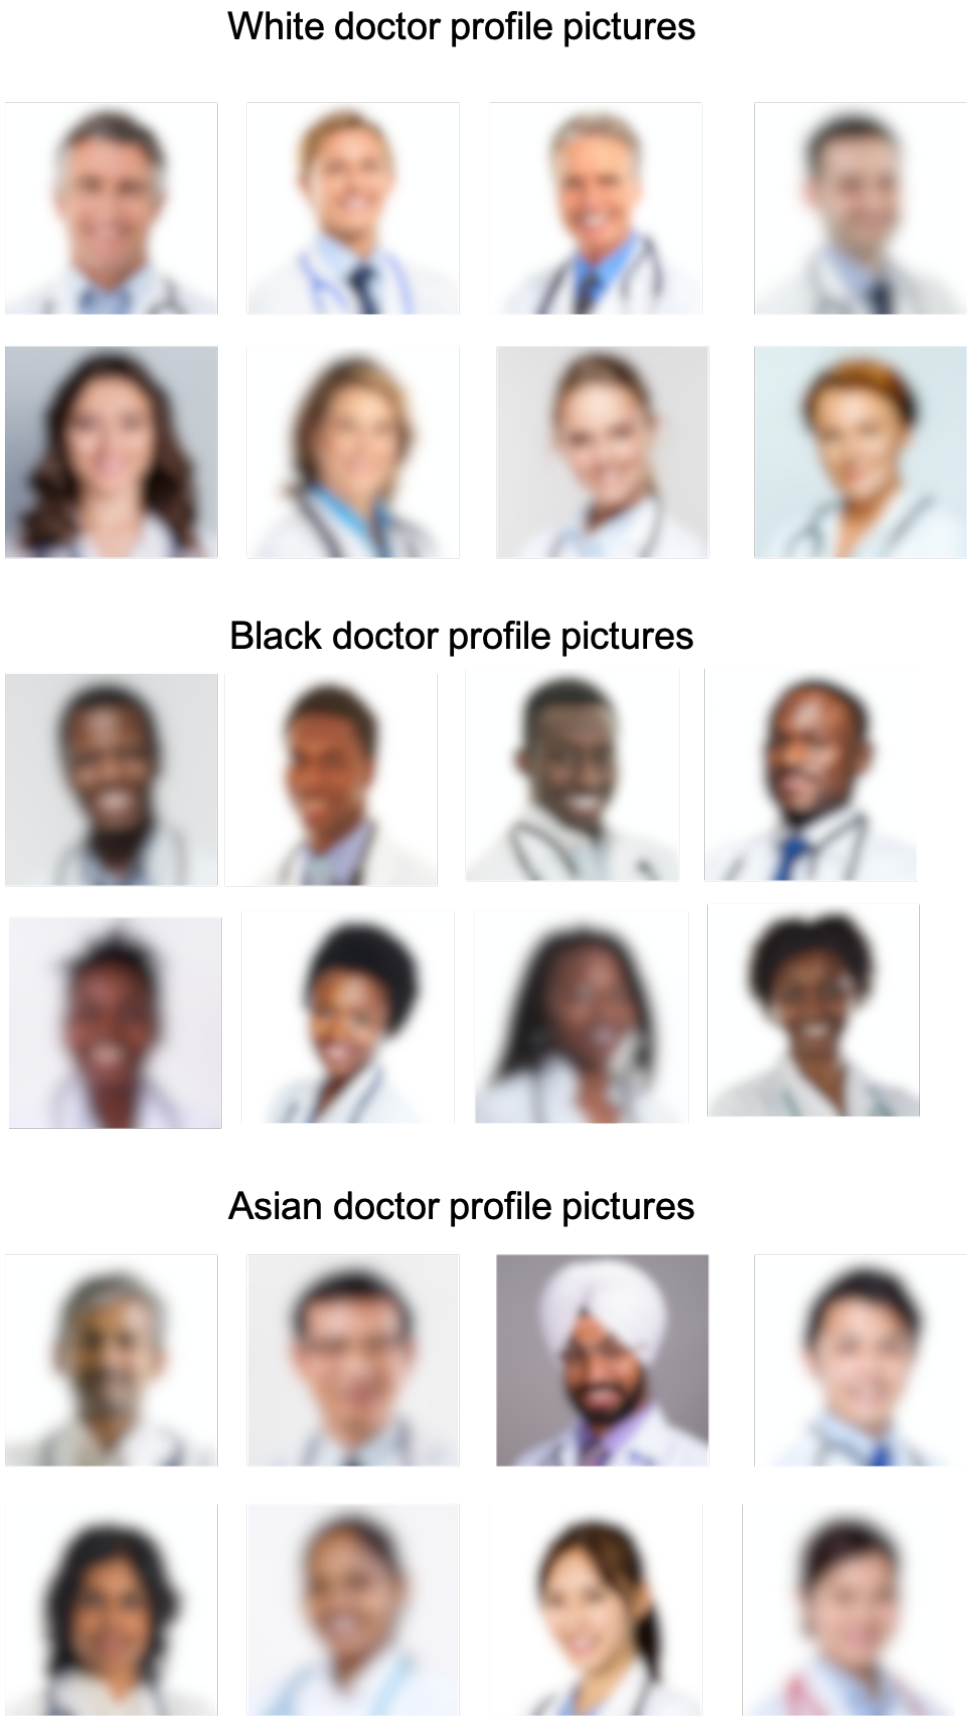


Figure 46: Blurred doctor profile pictures indicative of race

A.5 Customer Descriptions on Choice Approach

In this Section, I present some word clouds for the free text description submitted by the subjects from the control and treatment groups. "Price" is prominent in both control and treatment groups. The word cloud for the control group is shown in Figure 47 and the word cloud for the treatment group is shown in Figure 48.



Figure 47: Word cloud based on customer description on their choice approach in the control group



Figure 48: Word cloud based on customer description on their choice approach in the treatment group

A.6 Maximum Simulated Likelihood Estimators

In Section 2.5, I obtain choice probability of customer i choosing doctor d with the familiar multinomial logit form:

$$P_i(d) = \frac{\exp[\delta_{id} - \eta_i p_d]}{1 + \sum_j \exp[\delta_{ij} - \eta_i p_j]} \tag{18}$$

and I further assume that marginal utility δ_{id} is a linear function of doctor attributes:

$$\delta_{id} = \mathbf{a}_d^{'} \lambda_i \tag{19}$$

where $\mathbf{a}_d$ denotes the vector of attributes of doctor d including race, gender, distance, and quality (for treatment group). λ_i is the vector of attribute coefficients λ_{ia}. To simplify notation, I will call the vector of attribute coefficient plus the price coefficient λ and refer to the observable attributes and price of customer i's doctor option d in menu m as x_{idm}. I will also assume that this discrete choice model is parameterized with deep parameters θ.

I estimate a mixed logit model to allow the coefficients in the model to vary across customers. Considering that the customers make several choices (14 menus, denoted as $m \in \{1, 2, ..., 14\}$), the probability of a particular sequence of choices by customer i is given by:

$$S_i = \int \prod_{m=1}^{14} \prod_{i=0}^{D} [\frac{\exp[x_{idm}^{'} \lambda]}{1 + \sum_j \exp[x_{jdm}^{'} \lambda]}]^{y_{idm}} f(\lambda|\theta) d\lambda, \tag{20}$$

where $f(\lambda|\theta)$, which depends on deep parameters θ, is the density function of λ. And $y_{idm} = 1$ if customer i chose alternative d in menu m and $y_{idm} = 0$ otherwise.

The main empirical results and hypothesis testing are based on maximum simulated likelihood estimators (Ben-Akiva et al. (2019)). In the present paper, I assumed that marginal utility is a linear function of doctor attributes and the λ parameters are estimated by maximizing the simulated log-likelihood function using the data from customers 1 through N:

$$SLL = \sum_{i=1}^{N} \ln\{\frac{1}{R} \sum_{r=1}^{R} \prod_{m=1}^{14} \prod_{i=0}^{D} [\frac{\exp[x_{idm}^{'} \lambda_i^{[r]}]}{1 + \sum_j \exp[x_{jdm}^{'} \lambda_i^{[r]}]}]^{y_{idm}}\}, \tag{21}$$

where $\lambda_i^{[r]}$ is the r^{th} draw for customer i from the distribution of λ.

B Supplementary Tables

Table 3: Attribute Levels Used for Doctor Profiles in the Conjoint Surveys

Doctor Attribute	Levels	*Remarks*
Price	\$1704.00	10^{th}*-percentile of actual prices in real doctor sample*
	\$1797.07	
	\$1890.13	
	\$1983.20	
	\$2076.27	
	\$2169.33	
	\$2262.40	
	\$2355.47	
	\$2448.53	
	\$2541.60	90^{th}*-percentile of actual prices in real doctor sample*
Doctor Demographic	Black Male	*4 different profile pictures*
	Black Female	*4 different profile pictures*
	White Male	*4 different profile pictures*
	White Female	*4 different profile pictures*
	Asian Male	*4 different profile pictures*
	Asian Female	*4 different profile pictures*
Travel Distance	0-10 miles	*Distance b/w customer and provider zipcodes*
	10-50 miles	*Distance b/w customer and provider zipcodes*
	50-100 miles	*Distance b/w customer and provider zipcodes*
	100-250 miles	*Distance b/w customer and provider zipcodes*
	250 miles or more	*Distance b/w customer and provider zipcodes*
Comprehensive Quality Score	1 star	*Lowest quality level*
	2 stars	
	3 stars	
	4 stars	
	5 stars	*Highest quality level*

Notes: This table lists all the attributes and levels of the attributes that are used to design the conjoint survey doctor profiles. There are 6000 different attribute level combinations that are possible for the doctor profiles. Subjects cannot be exposed to every one of these profiles. In the experiment, I deploy allow considerable linearly independent variation in the levels of different attributes and a considerable span of attribute levels with an off-the-shelf fractional factorial choice design to optimize balance, overlap, and other characteristics.

Table 4: Comparing Facial Symmetry Across Race for Doctor Profile Pictures

	(1) Facial Symmetry	(2) Facial Symmetry
Black	0.375	0.375
	(7.291)	(7.303)
Asian	1.750	1.750
	(7.291)	(7.303)
Female		5.750
		(5.963)
N	24	24
R^2	0.003	0.047
White Average	59.250	59.250
Black Average	59.625	59.625
Asian Average	61.000	61.000

Standard errors in parentheses
$^{+}$ $p < 0.1$, * $p < 0.05$, ** $p < 0.01$

Notes: This table shows the balance on facial symmetry (as a proxy of race-neutral attractiveness) across races in the doctor profiles used in the conjoint surveys. The facial symmetry score is calculated based on the unblurred pictures. I estimate correlation between race of doctor profile picture and facial symmetry with the model $Symmetry_i = \beta_0 + \beta_1 1\{Black\} + \beta_2 1\{Asian\} + \epsilon_i$, in column (1). Standard errors for coefficients are given in parentheses. I control for gender of the doctor profile pictures in the estimation model for column (2).

Table 5: Balance Check for Conjoint Survey Profiles

	Total		Control		Treatment		
	mean	s.d.	mean	s.d.	mean	s.d.	Difference p-value
Black	0.34	0.47	0.34	0.47	0.34	0.47	0.981
White	0.33	0.47	0.33	0.47	0.33	0.47	0.964
Asian	0.33	0.47	0.33	0.47	0.33	0.47	0.945
Female	0.50	0.50	0.50	0.50	0.50	0.50	0.976
Black Male	0.17	0.37	0.17	0.37	0.17	0.37	0.988
Black Female	0.17	0.37	0.17	0.37	0.17	0.37	0.963
White Male	0.17	0.37	0.17	0.37	0.17	0.38	0.740
White Female	0.16	0.37	0.17	0.37	0.16	0.37	0.695
Asian Male	0.16	0.37	0.16	0.37	0.16	0.37	0.696
Asian Female	0.17	0.37	0.17	0.37	0.17	0.38	0.635
Price	2122.79	266.95	2122.42	267.47	2123.22	266.37	0.852
Distance 0-10 Miles	0.20	0.40	0.20	0.40	0.20	0.40	0.954
Distance 10-50 Miles	0.20	0.40	0.20	0.40	0.20	0.40	0.961
Distance 50-100 Miles	0.20	0.40	0.20	0.40	0.20	0.40	0.902
Distance 100-250 Miles	0.20	0.40	0.20	0.40	0.20	0.40	0.901
Distance Over 250 Miles	0.20	0.40	0.20	0.40	0.20	0.40	0.915
Patients	224		120		104		
N	18,816		10,080		8,736		

Notes: This table reports the attribute levels of the 18,816 hypothetical profiles across 224 customers in the conjoint survey, pooled and by treatment group. "Black," "Asian," and "White" indicate the shares of doctor profiles belonging to each of these categories; "Female" indicates the share of female sex for the doctor profiles; likewise, "Black Male," "Black Female," "White Male," "White Female," "Asian Male," and "Asian Female" indicate the shares of doctor profiles belonging to each of these categories. "Price" is the average price of the doctor profiles. The "Distance" rows indicate the shares of doctor profiles belonging to each of these distance categories. The Difference p-value column reports the p-value for the test of equality between the treatment and control groups. Stars indicate whether this difference is significant. As one would expect from a well executed conjoint design, I do not find any statistically significant differences for all these differences.

Table 6: Choice Model Parameter Estimates in Willingness-to-Pay Broken Down by Race Segments (Black and White)

Attribute	Black Patients		White Patients	
	Control	Treatment	Control	Treatment
a	$WTP_a^{C,black}$	$WTP_a^{T,black}$	$WTP_a^{C,white}$	$WTP_a^{T,white}$
Black	-77.808	21.425	-296.996	-29.386
	(25.554)	(16.819)	(11.028)	(8.752)
Asian	-101.558	-13.154	-198.336	-24.864
	(17.263)	(12.581)	(7.123)	(4.791)
Female	-22.694	-27.515	-16.428	3.783
	(22.839)	(15.671)	(8.610)	6.543
10-50 miles	-58.704	-34.627	-35.265	-49.459
	(10.540)	(6.558)	(3.837)	2.390
50-100 miles	-127.319	-111.704	-117.167	-117.029
	(13.293)	(16.653)	(3.462)	3.584
100-250 miles	-252.136	-269.950	-257.211	-253.020
	(24.383)	(32.056)	(7.089)	7.537
More than 250 miles	-595.366	-567.257	-577.198	-551.991
	(53.173)	(70.206)	(15.416)	16.617
5 stars		74.228		72.035
		(8.927)		2.917
4 stars		61.797		69.867
		(9.957)		3.815
3 stars		9.300		16.389
		(9.504)		3.218
2 stars		0.793		7.058
		(6.076)		3.320
N	1,176	756	8,064	7,476

Notes: This table reports the results for the 23 black and 185 white customers in sample, by treatment group. This table presents the estimated parameter results for the primary choice model from Section 2.5 expressed in terms of willingness-to-pay calculated as outlined in Section 2.5.1. All willingness-to-pay coefficients are in dollar units with standard errors for parameters given in parentheses. Marginal utility is assumed to be a linear function of doctor attributes and the parameters represents the willingness-to-pay relative to the reference attribute level. For race, the reference level is white race for the doctor. For gender, the reference level is male. For distance, the reference level is "0-10 miles" distance between the doctor's zip code and the customer's zip code. For quality stars, 1 star is the reference level.

Table 7: Choice Model Parameter Estimates in Willingness-to-Pay Broken Down by Gender Segments (Female and Male)

Attribute	Female Patients		Male Patients	
	Control	Treatment	Control	Treatment
a	$WTP_a^{C,female}$	$WTP_a^{T,female}$	$WTP_a^{C,male}$	$WTP_a^{T,male}$
Black	-270.147	-32.374	-270.125	-19.494
	(13.210)	(12.160)	(14.410)	(10.589)
Asian	-193.795	-21.263	-180.106	-26.151
	(8.918)	(7.028)	(7.115)	(4.989)
Female	85.370	39.480	-81.852	-36.866
	(9.845)	(8.903)	(9.444)	(7.899)
10-50 miles	-55.739	-54.229	-25.381	-41.799
	(4.653)	(3.336)	(3.886)	(3.313)
50-100 miles	-122.466	-113.761	-115.060	-119.790
	(6.004)	(4.142)	(3.651)	(4.931)
100-250 miles	-292.051	-255.389	-243.743	-249.335
	(13.927)	(9.884)	(7.775)	(9.528)
More than 250 miles	-670.749	-568.394	-535.135	-535.517
	(29.230)	(23.010)	(14.738)	(19.739)
5 stars		69.831		74.045
		(4.610)		(3.378)
4 stars		70.319		68.060
		(5.809)		(4.145)
3 stars		16.728		15.913
		(4.074)		(3.977)
2 stars		14.708		-1.161
		(3.819)		(3.745)
N	4,368	4,284	5,712	4,452

Notes: This table reports the results for the 103 female and 121 male customers in full sample, by treatment group. This table presents the estimated parameter results for the primary choice model from Section 2.5 expressed in terms of willingness-to-pay calculated as outlined in Section 2.5.1. All willingness-to-pay coefficients are in dollar units with standard errors for parameters given in parentheses. Marginal utility is assumed to be a linear function of doctor attributes and the parameters represents the willingness-to-pay relative to the reference attribute level. For race, the reference level is white race for the doctor. For gender, the reference level is male. For distance, the reference level is "0-10 miles" distance between the doctor's zip code and the customer's zip code. For quality stars, 1 star is the reference level.

Table 8: Choice Model Parameter Estimates in Willingness-to-Pay Broken Down by Believers versus Non-Believers

Attribute	Control	Treatment (sub-samples)	
		Believers	Non-Believers
a	WTP_a^C	$WTP_a^{T,Believ}$	$WTP_a^{T,NBeliev}$
Black	-270.133	-13.763	-155.009
	(9.164)	(8.188)	(37.968)
Asian	-185.529	-19.275	-71.850
	(6.139)	(4.401)	(19.695)
Female	-15.598	1.283	-7.563
	(7.227)	(6.262)	(29.227)
10-50 miles	-37.409	-46.538	-62.371
	(3.028)	(2.566)	(15.098)
50-100 miles	-117.994	-113.204	-155.825
	(3.082)	(3.014)	(20.520)
100-250 miles	-262.883	-250.106	-275.847
	(7.254)	(6.590)	(41.856)
More than 250 miles	-588.866	-542.791	-646.376
	(15.216)	(15.224)	(92.269)
5 stars		73.655	54.013
		(2.847)	(13.976)
4 stars		71.470	44.447
		(3.913)	(9.282)
3 stars		16.389	15.488
		(3.198)	(14.223)
2 stars		8.277	-11.259
		(2.997)	(11.522)
N	10,080	7,812	924

Notes: This table reports the results for the 224 customers in the main sample, by treatment group and by whether a treatment group subject "believes" the quality signal. Customers are categorized based their their response to the Likert scale question: "A provider with a higher Comprehensive Quality Score is a better provider than aprovider with a lower Comprehensive Quality Score" in the exit survey. Those who chose "Disagree" and "Strongly disagree" are non-believers of the signals and the rest "believers." This table presents the estimated parameter results for the primary choice model from Section 2.5 expressed in terms of willingness-to-pay calculated as outlined in Section 2.5.1. All willingness-to-pay coefficients are in dollar units with standard errors for parameters given in parentheses. Marginal utility is assumed to be a linear function of doctor attributes and the parameters represents the willingness-to-pay relative to the reference attribute level. For race, the reference level is white race for the doctor. For gender, the reference level is male. For distance, the reference level is "0-10 miles" distance between the doctor's zip code and the customer's zip code. For quality stars, 1 star is the reference level.