Determining the Extent of Taste-Based and Accurate Statistical Discrimination: Evidence from a Field Experiment in India*

Asad Islam#
Debayan Pakrashi⁑
Liang Choon Wang‡
Yves Zenou†

May 2024

Abstract

We conduct a controlled field experiment to examine the extent of taste-based and accurate statistical discrimination in the demand for healthcare in India. By eliciting patients' preference rankings for doctors of different castes and experiences, we find that for 57% of patients, it is difficult to distinguish between taste-based and statistical discrimination. Most preference rankings suggest patients hold inaccurate beliefs about doctors. Thus, policies aimed at correcting inaccurate beliefs could be particularly effective in addressing discrimination.

JEL codes: I14, J15, O15

Keywords: Field experiment, taste-based discrimination, statistical discrimination, affirmative action, healthcare.

---

* We thank the editor as well as three anonymous referees for very helpful comments that help improve the paper.
# Monash University, CDES and Department of Economics. Email: asadul.islam@monash.edu
⁑ Indian Institute of Technology, Kanpur. Email: pakrashi@iitk.ac.in.
‡ Monash University, Department of Economics. Email: liang.c.wang@monash.edu.
† Monash University, Department of Economics, and CEPR. Email: yves.zenou@monash.edu.

1. Introduction

The economics literature identifies two primary sources of discrimination: taste-based and statistical. Taste-based discrimination arises when individuals have aversions to certain population groups. In contrast, statistical discrimination occurs in environments with imperfect information, where individuals form beliefs based on limited and potentially inaccurate signals. Empirically distinguishing whether an individual's behavior aligns more closely with one of these sources is challenging, as differential treatments across groups can be consistent with both theories. Since statistical discrimination is driven by beliefs, understanding whether these beliefs are accurate or inaccurate is crucial for developing effective anti-discrimination policies. The primary contribution of this paper is the application of a direct revealed preference approach to differentiate between statistical and taste-based discrimination. Additionally, it assesses the extent to which statistical discrimination, driven by accurate beliefs, influences demand in the healthcare sector.

We conduct a controlled field experiment across 40 localities in Uttar Pradesh, India, to examine whether taste-based discrimination or statistical discrimination is the main driver of patients' discriminatory behaviors against physicians of different castes. Indeed, in India, caste is a salient dividing line and can factor into patients' healthcare decisions (see Section 3.1 below). Uttar Pradesh is a suitable setting to study this question because it is one of the Indian states with the largest concentration of population from the "backward-caste" background, who also receive caste-based reservations for medical college seats. Through a simple theoretical framework that describes the experimental setting, we first show that taste-based discriminators, whether they have homophily (in-group) or heterophily (out-group) caste preference, can exhibit preference rankings that completely overlap with a subset of statistical discriminators' preference rankings. Because of this overlap, we can determine the extent of taste-based discrimination.

We use a correspondence method to elicit 3,128 patients' rankings of four doctors of two different caste groups and two different years of experience. A patient can have one out of 24 feasible preference rankings of doctors. Our theoretical framework shows that if we do not impose any parametric assumptions about taste while maintaining that preference relations satisfy five fundamental axioms of preferences, at most only eight out of these 24 rankings are possible under taste-based discrimination whereas all 24 of them are possible under statistical discrimination. Focusing on the most preferred doctors as the measure of choices, we find that mean choices by caste of patients are consistent with both theories of discrimination. Thus, it can be difficult to empirically distinguish whether *individuals' choices* are primarily driven by taste-based or belief-based discrimination with the typical data collected in most field experiments. Using individual-level preference rankings, we find that 43% of patients have rankings that are unambiguously consistent with statistical discrimination. The results highlight that in most cases (57%), preference rankings of taste-based discriminators and statistical discriminators are indistinguishable. We further demonstrate that this conclusion is robust to the potential presence of caste-blind patients.

Since statistical discrimination is a belief-based explanation, the natural question that arises is whether statistical discriminators hold accurate or inaccurate beliefs. As there can only be at most one preference ranking that is consistent with accurate beliefs when the true quality of doctors differs by caste and experience, we can provide the lower bound of patients who hold inaccurate or biased beliefs. As the preference ranking most frequently observed in the sample represents 21% of all responses, at least 79% of patients hold inaccurate or biased beliefs. Among the 57% of preference rankings that are potentially driven by taste-based discrimination, at least 63% of them overlap with those consistent with inaccurate beliefs. We also examine whether inaccurate beliefs are likely due to the presence of the stereotype that general-category doctors provide better quality healthcare, which may arise due to the caste-based reservations for backward-caste individuals to pursue a medical degree. We find that the general-category caste stereotype is unlikely to be the main source of these inaccurate beliefs. Given that most patients have preference rankings that are consistent with inaccurate beliefs, our findings imply that policies designed to correct inaccurate beliefs may have strong potentials to address caste discrimination.

Policies aimed at dispelling misinformation and improving information exchanges between caste groups are not only useful for addressing inaccurate beliefs, but are also potentially useful for reducing social distance between groups and encouraging intergroup interaction. To the extent the information interventions that encourage intergroup exchanges can address taste-based discrimination, policies that address inaccurate beliefs may also be used to address taste-based discrimination. In particular, such policies may work more effectively on individuals who do not hold animosity towards the caste group they discriminate against. Thus, we further estimate the share of patients whose caste preferences are less robust and are potential receptive to information interventions by using Ferschtman and Gneezy's (2001) and List's (2004) notion of group animosity revealed in incentivized dictator games. We find that the vast majority of patients, even if they taste-based discriminate doctors, do not have a robust caste preference and are potentially receptive to information interventions aimed at addressing caste discrimination.

There is a large literature focusing on detecting discrimination and examining the sources of discrimination using laboratory experiments, field experiments, natural experiments, and non-experimental approaches. In particular, correspondence methods have been the primary approach used to investigate discrimination in a variety of settings, including employment (Bertrand & Mullainathan, 2004; Banerjee et al., 2009; Guiletti et al., 2019), housing (Ewens et al., 2014), product markets (Gneezy et al., 2012; Doleac and Stein, 2013; Zussman, 2013; Siddique et al., 2023), financial markets (Bayer et al., 2018), education (Hanna & Linden, 2012), and along different dimensions, including race, ethnicity, caste, gender, age, disability, sexual orientation, obesity, and religion.[2] However, as noted by Charles and Guryan (2011) and Neumark (2018), it remains a challenge to empirically differentiate behaviors consistent with taste-based discrimination and statistical discrimination.[3]

This paper helps advance the research on discrimination. Past studies that reject taste-based discrimination as the main driver of differential treatments across groups typically make specific assumptions about preferences, such as functional form and group bias assumptions. They then focus on how the aggregate patterns of *choices* vary with experimentally manipulated information in order to make inferences about whether the *average* individual's discriminatory behavior is more consistent with taste-based discrimination or statistical discrimination (e.g., List 2004, Ewens et al. 2014, Bohren et al. 2019). We show that if we only maintain that preferences must satisfy five fundamental axioms of preferences as all past studies that reject taste discrimination do, while not imposing additional assumptions, then it is difficult to rule out the role of taste-based discrimination in driving the aggregate patterns of choices. More generally, our contribution is to show customer discrimination in linking a correspondence study to a real transaction by separating

[2] See Altonji and Blank (1999), Guryan and Charles (2013), Lang and Lehmann (2012), Neumark (2018), and Lang and Kahn-Lang Spitzer (2020) for general overviews, Anderson et al. (2006) and Lane (2016), for tests of discrimination in the laboratory, Baert (2018), for correspondence experiments, and Riach and Rich (2002), Rich (2014), and Bertrand and Duflo (2017), for field experiments.

[3] See Altonji and Pierret (2001) and Arcidiacono et al. (2010) for testing in a non-structural way statistical discrimination in the context of ability in the labor market. There is also an entire literature on identifying prejudiced based discrimination by looking at performance. Most notably, Anwar and Fang (2006) on police stops and Arnold et al. (2018) on bail decisions. However, these papers do not separately identify statistical discrimination, prejudice-based discrimination or both.

belief-based and preference-based motives for discrimination while relaxing several of the strong assumptions usually required.

Our focus on preference rankings circumvents a major shortcoming of past studies on discrimination in which an *individual* discriminator's preference rankings over the *choice set* were unobserved. The technique of eliciting individuals' preference rankings was previously implemented in the laboratory setting to test theories about correlated beliefs (Cason et al., 2020) and self-regarding and other-regarding preferences (Levati et al., 2014). We apply this technique to a field setting in the *final stage of real transactions* to bound taste-based discrimination and statistical discrimination with accurate beliefs. We demonstrate that it is possible to bound the extent of taste-based discrimination by excluding individuals' preference rankings that do not satisfy five fundamental axioms of preferences. Importantly, this technique also allows us to demonstrate that few statistical discriminators hold accurate beliefs.

Our axiomatic approach departs from the recent literature that focuses on testing whether discriminatory behaviors are on *average* more consistent with biased beliefs, unbiased beliefs, or taste-based discrimination (e.g., Bohren et al., 2019; Bohren et al., 2024; Chan, 2023).[4] By not imposing assumptions on functional form of utility and group preference (e.g., all subjects have distaste against the same group), we show it is much harder to separate preference-based explanation from belief-based explanation. Importantly, we are able to show that in most cases statistical discrimination with inaccurate beliefs and taste-based discrimination are indistinguishable. Despite this challenge, we also provide suggestive evidence that the majority of discriminators do not hold animosity towards the discriminated groups and are potentially receptive to information interventions aimed at addressing both inaccurate beliefs and taste-based discrimination.

2. The benchmark model

We outline the two major economic theories of discrimination in the context of our field experiment and present the situations in which the predictions of statistical discrimination completely overlap with the predictions of taste-based discrimination and the situations in which they do not. In order to match our field experiment, we assume only two types of castes for both the doctors and the patients: general category and backward caste, and two different levels (years)

[4] See Onuchic (2024) for a recent theory overview.

of experience for the doctors: low and high. We denote the caste of a doctor by $c = c_G, c_B$, where c_G corresponds to the general category and c_B corresponds to backward caste, and the caste of a patient by $c^p = c_G^p, c_B^p$, where the superscript p refers to the patient. We denote the experience of a doctor by $e = e_H, e_L$, with $e_H > e_L$.

2.1. Taste-based discrimination

According to Becker's (1957) theory of taste-based discrimination, prejudiced employers (or workers or consumers) dislike employing (or working with, or purchasing from) people from a certain group (e.g., race, gender, caste, etc.).

Here, we model taste-based as a consumer choice problem. The patient (consumer) chooses a doctor consumption bundle, x, which contains both the social closeness or proximity (Φ) to the doctor's social group relative to the patient's social identity and the perceived quality of healthcare (q) provided by the doctor, among the patient's consumption set X of all possible doctors. Social proximity between the patient and a doctor is fully characterized by strength to which the patient relates or identifies with the doctor's caste group. The idea that social proximity measures taste for caste is consistent with Becker's (1957) notion of taste where the cost of interacting with the socially distant group is higher as well as Akerlof's (1997) idea that social distance determines attitudes toward discrimination. We denote $\Phi(c^p, c)$, the social proximity of a patient from caste c^p to a doctor from caste c. Furthermore, the perceived quality of healthcare provided by a doctor is fully characterized by the years of experience e the doctor has been practicing medicine. On average, we expect the perceived quality to be increasing in years of experience, but there can be individual patients who perceive a doctor with fewer years of experience as one with better quality.[5] We thus allow for the possibility that $q'(e) > 0$ for some patients, but $q'(e) < 0$ for other patients. A consumption bundle x $\in X$ is thus represented by a vector x $\in \mathbb{R}_+^2$.

[5] What matters to patients is likely not the experience of a doctor per se, but the quality of healthcare, such as the accuracy of diagnosis, the efficacy of the prescribed medicine, etc., provided by the doctor. We assume that the quality of healthcare can be fully described by the doctor's experience as it is the dimension that we manipulate in the experiment. In reality, patients may use experience as a noisy signal of quality $\mathbb{E}(q|e)$ to predict the quality of healthcare to be provided by the doctor. We can accommodate this possibility in the model by having the patient to first use the experience of the doctor to predict the quality of healthcare that the doctor will provide (with some errors) before choosing doctors with different levels of predicted quality and caste backgrounds. In order to put aside the possibility of statistical discrimination, we must however assume patients do not utilize the information related to the caste of a doctor when predicting the quality of healthcare that the doctor will provide.

We assume that patients' preferences for doctors with various levels of social proximity and perceived healthcare quality satisfy five fundamental axioms of consumer choice: completeness, transitivity, continuity, strict monotonicity, and strict convexity. It follows that a patient's preference relation between two doctors with different levels of social proximity and perceived healthcare quality can be represented by a real-valued utility function: $\mathrm{U}: \mathbb{R}_{+}^{2} \rightarrow \mathbb{R}, \forall \mathrm{x}^{0}, \mathrm{x}^{1} \in \mathbb{R}_{+}^{2}$, such that $\mathrm{U}(\mathrm{x}^{0}) \geq U(\mathrm{x}^{1}) \Leftrightarrow \mathrm{x}^{0} \gtrsim \mathrm{x}^{1}$. We denote this utility function for a patient c^{p} by $U(\Phi(c^{p}, c), q)$ and assume that it is strictly increasing in both arguments.

2.1.1. Homophily versus heterophily taste-based discrimination

We consider two possible manifestations of caste preference. The first is patients with a preference bias for own caste or homophily (McPherson et al., 2001; Currarini et al., 2009), so that the patients identify more with people from the same caste group and there is a cost of interacting with a doctor from a different caste. This is a standard assumption that the racial discrimination literature typically makes (at least for the average individual). Thus, for a backward-caste patient with homophily preference, the social proximity to a backward-caste doctor is higher than to that of a general-category doctor, $\Phi(c_{B}^{p}, c_{B}) > \Phi(c_{B}^{p}, c_{G})$. Likewise, for a general-category patient with homophily preference, the social proximity to a backward-caste doctor is lower than to that of a general-category doctor, $\Phi(c_{G}^{p}, c_{B}) < \Phi(c_{G}^{p}, c_{G})$.

The second type of manifestation is patients with out-group preference bias or heterophily. For this type of patients, they identify more with people from a different caste group and there is a cost of interacting with a doctor from the same caste. We do not expect the majority of individuals to be this type, but it is possible that there are some of them and our theoretical framework allows for such a possibility. For a backward-caste patient with heterophily preference, we have $\Phi(c_{B}^{p}, c_{B}) < \Phi(c_{B}^{p}, c_{G})$. Likewise, for a general-category patient with heterophily preference, we have $\Phi(c_{G}^{p}, c_{B}) > \Phi(c_{G}^{p}, c_{G})$.

2.1.2. Rankings from the perspective of a backward-caste patient

On the basis of the axioms and assumptions specified above for taste-based discriminators, there are four possible preference rankings of doctors for a backward-caste patient c_{B}^{p} who has homophily preference, given their choice set of doctors with two levels of experience e_{L} and e_{H}

and two caste backgrounds c_B and c_G. Similarly, there are also four possible preference rankings of doctors for a backward-caste patient who has heterophily preference.

For a homophily backward-caste patient who views more years of experience as an indication of better healthcare quality, the best possibility is a doctor $c_B e_H$, i.e., a doctor from a backward-caste group c_B and with a high experience level e_H, because the social proximity and quality of this doctor are both the highest. For this patient, the worst possibility is a doctor $c_G e_L$, as the social proximity and quality of the doctor are both the lowest. Figures 1A and 1B illustrate the two possible rankings for a homophily backward-caste patient given an arbitrary utility function $U\left(\Phi\left(c_B^p, c\right), q\right)$ that satisfies the axioms and assumptions of consumer choice specified above. Note that, in the first ranking (Figure 1A), the patient exhibits stronger preference for social proximity than for quality, while in the second ranking (Figure 1B), the patient exhibits stronger preference for quality than for social proximity. We show these results using an additively separable utility function that generates all rankings consistent with taste-based discrimination in Online Appendix A.

[Figure 1 here]

For a homophily backward-caste patient who views fewer years of experience as an indication of better healthcare quality, the best possibility is a doctor $c_B e_L$ and the worst possibility is a doctor $c_G e_H$. Figures 1C and 1D illustrate the two possible rankings for a homophily backward-caste patient given an arbitrary utility function $U\left(\Phi\left(c_B^p, c\right), q\right)$ that satisfies the axioms and assumptions of consumer choice specified above.

For a heterophily backward-caste patient who views more years of experience as an indication of better healthcare quality, the best possibility is a doctor $c_G e_H$, as both the social proximity and quality of this doctor are the highest. The worst possibility for this patient is a doctor $c_B e_L$ as both the social proximity and quality of this doctor are the lowest. In contrast, a heterophily backward-caste patient who views fewer years of experience as an indication of better healthcare quality, the best possibility is a doctor $c_G e_L$. The worst possibility for this patient is a doctor $c_B e_H$. Figure 2 illustrates all four possible rankings for a heterophily backward-caste patient given an arbitrary utility function $U\left(\Phi\left(c_B^p, c\right), q\right)$ that satisfies the axioms and assumptions of consumer choice specified above.

[Figure 2 here]

In sum, there are 8 out of 24 possible rankings for a backward-caste patient who taste-based discriminates doctors by caste (Appendix A lists the rankings).

2.1.3. Rankings from the perspective of a general-category patient

There are 8 possible preference rankings of doctors for a general-category patient whose preferences satisfy the axioms and assumptions specified above. The preference rankings of a homophily general-category patient who views more (fewer) years of experience as an indication of health quality are identical to the preference rankings of a heterophily backward-caste patient who views more (fewer) years of experience as an indication of healthcare quality. This is because both patients view a general-category doctor being socially closer than a backward-caste doctor and more (fewer) years of experience as an indication of better healthcare quality. Similarly, the preference rankings of a heterophily general-category patient who views more (fewer) years of experience as an indication of healthcare quality are identical to the preference rankings of a homophily backward-caste patient who views more (fewer) years of experience as an indication of healthcare quality.

In sum, there are 8 out of 24 possible rankings for a general-category patient who taste-based discriminates doctors by caste. These 8 possible rankings are identical to those for a backward-caste patient who taste-based discriminates doctors by caste. Therefore, there are 16 out of 24 possible rankings that are inconsistent with taste-based discrimination.

2.2. Statistical discrimination

Phelps (1972) and Arrow (1973) pioneered statistical discrimination theory. The theory posits that, in the absence of direct information about quality, a decision maker would use group averages (beliefs) to make inferences. For instance, labor market discrimination may exist because employers do not know with certainty workers' productivity and, therefore, may base their employment decisions on the workers' visible features, such as group identity or race, as long as these features are correlated with the unobserved productivity. This type of discrimination can result in self-fulfilling behavior from the disadvantaged groups. For example, Verdier and Zenou

(2004) show that, if all agents, including blacks themselves, believe with no reason that blacks are more criminal than whites, blacks can become more criminal than whites because, based on wrong beliefs, employers pay them less, which forces them to reside far away from job centers, which leads blacks to rationally commit more crime than whites.[6]

As in the case of taste-based discrimination, we assume that a statistically discriminating patient's preference relations satisfy five fundamental axioms consumer choice: completeness, transitivity, continuity, strict monotonicity, and strict convexity. However, for this statistically discriminatory patient, the caste of a doctor does not influence the social proximity of the doctor and that all doctors irrespective of caste have the same social proximity. Thus, it is only the quality of healthcare that influences their preferences for doctors. Because the actual quality of healthcare, q, to be delivered by a doctor is unobserved to the patient before the transaction takes place, she uses the information about the doctor's caste, c, and experience, e, to predict this quality, $\mathbb{E}(q|e,c)$. It is important to note that what matters to this patient is not the experience of a doctor per se but the quality of healthcare or treatment provided by the doctor. Their consumption bundle, x, contains only the expected quality of healthcare, $\mathbb{E}(q|e,c)$. The patient's preference relation between two doctors is *now* represented by a real-valued utility function: $\mathrm{U}: \mathbb{R}_+ \rightarrow \mathbb{R}, \forall \mathrm{x}^0, \mathrm{x}^1 \in \mathbb{R}_+$, such that $\mathrm{U}(\mathrm{x}^0) \geq U(\mathrm{x}^1) \Leftrightarrow \mathrm{x}^0 \gtrsim \mathrm{x}^1$.

We follow Phelps (1972) and Aigner and Cain (1977) to model statistical discrimination for our case. The experience e of a doctor from caste group c now provides a signal of the doctor's quality q with an error (noise) ε so that:

$$e = q + \varepsilon \tag{1}$$

where $\varepsilon \sim N\left(0, \sigma^2_{\varepsilon,c}\right)$ and $q \sim N\left(\beta_c, \sigma^2_{q,c}\right)$. It is assumed that $cov(q,\varepsilon) = 0$. Thus, $\mathbb{E}(e_c) = \beta_c$ and $Var(e_c) = \sigma^2_{q,c} + \sigma^2_{\varepsilon,c}$. Each patient infers the expected value of the doctor quality q from the noisy signal e (experience) using the available information, including the caste of the doctor c. In order to choose (rank) a doctor, each patient forms $\mathbb{E}(q|e,c)$. Since q and e are jointly normally distributed, for each caste of doctor $c = c_B, c_G$, we have:[7]

$$\hat{q}_c \equiv \mathbb{E}(q|e,c) = (1-\gamma_c)\beta_c + \gamma_c e_c \tag{2}$$

where $0 < \gamma_c < 1$ is given by:

[6] For a substantive survey on the theory of statistical discrimination, see Fang and Moro (2011).

[7] Observe that under statistical discrimination, there is no bias on average, i.e., $\mathbb{E}(\hat{q}_c|e,c) = \beta_c$.

$$\gamma_c = \frac{\sigma^2_{q,c}}{\sigma^2_{q,c}+\sigma^2_{\varepsilon,c}} = \frac{Cov(q_c,e_c)}{Var(e_c)} \tag{3}$$

where $Cov(q_c, e_c) > 0$ as experience is a positive signal of quality according to equation (1). In other words, for a given caste of doctor, c, a doctor with higher experience is perceived to be providing a higher quality healthcare. On the other hand, if experience is a negative signal of quality, then $Cov(q_c, e_c) < 0$.

Equation (2) says that $\hat{q}_c \equiv \mathbb{E}(q|e, c)$, the conditional distribution of q given e and c, follows a normal distribution with mean equal to a weighted average of the signal e_c and the unconditional group mean β_c. If the signal e_c is very noisy, i.e., the variance of ε, $\sigma^2_{\varepsilon,c}$, is large, the expected conditional value of doctor's quality is close to β_c, the population average of caste group c, regardless of the signal's value. In other words, when experience is not informative of quality, the patient uses the average quality of healthcare provided by the doctor's caste group to make inferences about a particular doctor's quality. On the other hand, if the signal is very precise, i.e., $\sigma^2_{\varepsilon,c}$ close to zero, then the signal e_c provides an accurate estimate of the doctor's quality. γ_c is often interpreted as the "reliability" of the signal since the higher is γ_c, the less noisy and thus more precise is the signal e_c.

2.2.1. Belief differences across caste groups

The choice of a doctor from a patient of caste c^p will depend on $\mathbb{E}(e_c) = \beta_c$, the signal e_c, and γ_c, the "reliability" of the signal. Different cases may arise based on what the patients' beliefs are.

We are agnostic about how beliefs about doctors from different castes may arise for an *individual* patient in the first place. One can imagine that everyone starts with a common prior about the quality of doctor from various caste groups, but the idiosyncratic experiences with different types of doctors over a patient's lifetime lead to the patient having different posterior beliefs about different types of doctors. In other words, their beliefs are shaped by the draws of doctors (from different caste groups and with different years of experience) they have encountered. These draws are unlikely to be independently and identically distributed, and the posterior beliefs that they lead to are unlikely to map to the posterior beliefs of the average patient.[8]

[8] Ewens et al. (2014) show that for racial discrimination in the US rental apartment market, when the draws are correlated within neighborhoods in which landlords are leasing apartments, the differences in variances across groups may arise since the variance includes all the pairwise correlations of observations. Bohren et al. (2019) examine whether belief differences across groups get updated as new information arrives in the context of gender discrimination

Consider an assumption typically made in the labor market discrimination literature, where the signal of labor productivity is assumed to be noisier for minority workers to explain why minority individuals with strong test performance may be discriminated on average (e.g., Aigner & Cain, 1977). In the case of an individual patient that discriminates doctors based on caste, the equivalence would be to assume that the patient have the beliefs that $\beta_{c_G} = \beta_{c_B}$, $\sigma^2_{q,c_G} = \sigma^2_{q,c_B}$, $Cov(q_c, e_c) > 0$, and $\sigma^2_{\varepsilon,c_B} > \sigma^2_{\varepsilon,c_G}$. These beliefs imply that $0 < \gamma_{c_B} < \gamma_{c_G} < 1$, so that the signal about a doctor's experience is less informative the quality of healthcare for the backward-caste c_B doctor than the general-category c_G doctor. In this standard case, three possible rankings of doctors may arise:

$$c_G e_H \succ c_G e_L \succ c_B e_H \succ c_B e_L$$
$$c_G e_H \succ c_B e_H \succ c_G e_L \succ c_B e_L$$
$$c_G e_H \succ c_B e_H \succ c_B e_L \succ c_G e_L$$

In the above three possible preference rankings of the four doctors, the first two of them are identical to the preference rankings that a taste-based discriminator (either a homophily general-category or a heterophily backward-caste patient) who views greater experience as an indicator of better healthcare quality may have. Only the last of the three rankings is unique for a statistical discriminator. Panels A to B in Figure 3 illustrate these three possible rankings. We show the relationship between $\hat{q}_c$ and e_c for simplicity given that $\mathrm{U}: \mathbb{R}_+ \to \mathbb{R}, \forall \mathrm{x}^0, \mathrm{x}^1 \in \mathbb{R}_+$ such that $\mathrm{U}(\mathrm{x}^0) \geq U(\mathrm{x}^1) \Leftrightarrow \mathrm{x}^0 \succsim \mathrm{x}^1$ for a caste-based statistical discriminator.

[Figure 3 here]

It is highly unlikely for all patients to share the beliefs that $\beta_{c_G} = \beta_{c_B}$, $\sigma^2_{q,c_G} = \sigma^2_{q,c_B}$, $Cov(q_c, e_c) > 0$, and $\sigma^2_{\varepsilon,c_B} > \sigma^2_{\varepsilon,c_G}$, even though it might be the case that the *average* patient has these beliefs. More generally, the idiosyncratic experiences of individual patients may give rise to all kinds of beliefs, including $\beta_{c_G} > \beta_{c_B}$, $\beta_{c_G} < \beta_{c_B}$, $\sigma^2_{q,c_G} > \sigma^2_{q,c_B}$, $\sigma^2_{q,c_G} < \sigma^2_{q,c_B}$, $Cov(q_c, e_c) < 0$, $\sigma^2_{\varepsilon,c_B} < \sigma^2_{\varepsilon,c_G}$, and $\sigma^2_{\varepsilon,c_B} > \sigma^2_{\varepsilon,c_G}$. Indeed, it is even possible for a patient to simultaneously have

in online evaluations of user-generated mathematics discussions. They find that without prior evaluations, women face significant discrimination, but the direction of discrimination reverses following a sequence of positive evaluations. Their findings imply that the beliefs are biased.

the belief that $Cov(q_{c_G}, e_{c_G}) > 0$ and the belief that $Cov(q_{c_B}, e_{c_B}) < 0$, for example. Without data on the history of encounters that each patient has had, it is impossible to estimate these beliefs at the *individual* patient level. Nonetheless, we can *infer* preference rankings that are consistent with the beliefs of statistical discriminators and identify those preference rankings that are distinguishable from the preference rankings of taste-based discriminators.

In total, all 24 feasible rankings are consistent with statistical discrimination. Thus, 8 of them completely overlap with the preference rankings of taste-based discriminators.

3. Context and experimental design

In this section, we provide some background on the caste system in India and then explain the way we implemented the field experiment.

3.1. Caste system in India

The caste groups in India are broadly categorized into the General Category (GC), the Scheduled Castes (SCs) or Dalits, Scheduled Tribes (STs) and Other Backward Classes (OBCs). For centuries, caste dictated almost every aspect of Hindu religious and social life. In recent decades, the influence of caste has somewhat declined, especially in cities where different castes live side-by-side and interact economically and socially. Despite the changes, caste identities remain strong, and surnames provide identification of castes.

After the independence of India, discriminating a person based on caste was legally forbidden. In 1950, the Indian government launched Affirmative Action (AA),[9] which is known as reservation policy in India, to promote equal opportunities for SCs and STs in areas of government jobs, government-funded education, and politics (Deshpande, 2012). Following the Mandal Commission's recommendations, quotas for government jobs were extended to OBCs in the early 1990s. In 1992, the Supreme Court of India put a cap on reservation and ruled that reservations should not exceed 50%. In 2006, educational quotas for OBCs were established through the 93rd educational amendment. Government-funded colleges and universities allot seats according to caste-based quotas, which assign 7.5% to STs, 15% to SCs, and 27% to OBCs

[9] For general overviews on the pros and cons of AA policies, see Holzer and Neumark (2000, 2006) and Arcidiacono and Lovenheim (2016). The literature examining the unintended consequences of AA policies under statistical discrimination dated back to the seminal work of Coate and Loury (1993).

(Deshpande, 2012). However, there is evidence indicating that individuals from SCs, STs, and OBCs continue to face discrimination, stigmatization, exclusion and rejection (Madheswaran and Attewell, 2007; Thorat and Attewell, 2007; Banerjee et al., 2009; Siddique, 2011; Islam et al., 2021).

3.2. The field experiment

We conducted a field experiment to test for the presence and source of caste discrimination in the demand for healthcare in the Kanpur Nagar district of Uttar Pradesh (UP), India. UP has the most population and also the largest concentration of backward caste people among all Indian states. Caste-based issues and policies have historically dominated the state's politics.

The field experiment took place in 40 localities across the Kanpur Nagar district between August and October 2017. Online Appendix B lists these localities. We selected these locations because their demographic and social characteristics are representative of the overall demographic and social characteristics of the UP state. A total of 3,128 adults participated in the field experiment. Table 1 shows that the average demographic and social economic characteristics of the participants are broadly similar to the demographic and social economic characteristics of individuals in UP state.

[Table 1 here]

We implemented the field experiment in four stages. In the first stage, participants registered their interests and expressed preferences for different types of doctors presented to them. In the second stage, participants answered a short survey questionnaire. In the third stage, participants were assigned to doctors and appointments. In the fourth stage, participants received the health services.

In the first stage, we randomly approached households in each locality to advertise for an upcoming, free-of-charge health check service offered by a mobile clinic. Due to safety and ethical concerns, individuals with potential urgent and life-threatening diseases or injuries were advised to seek immediate medical attention at the local hospital, instead of waiting for the upcoming health check. At the point of registration, we requested participants to express their preferences over four doctors with different profiles listed on a sign-up sheet (see Appendix C for an example).

In India, mobile medical units are a common practice in places where medical facilities are inadequate, or in areas populated by low-income households. Thus, the main advantage of our field experiment is that it occurred in a "natural" environment since people in these areas have used such services. It is also common for patients to register their interest for an upcoming service and express their preference. We therefore believe that the participants did not know that they were taking part in a caste discrimination study and they acted the way they normally would.

The sign-up sheet on which participants expressed their preferences showed a two-by-two matrix containing information for four doctors with different profiles: (i) a doctor with a general-category surname and a high number of years of experience ($c_G e_H$); (ii) a doctor with a backward-caste surname and a high number of years of experience ($c_B e_H$); (iii) a doctor with a general-category surname and a low number of years of experience ($c_G e_L$); and (iv) a doctor with a backward-caste surname and a low number of years of experience ($c_B e_L$).[10] The backward-caste surnames that appear on the sign-up sheet belong to either SC or OBC group. We did not include ST surnames given that less than 1% of participants are from a ST background.[11] Similarly, the high number of years of experience is either 12 years or eight years but never both. The low number of years of experience is always four years. We randomized the order in which each type of doctor appeared in the matrix. We did not disclose the first name of the doctor but only the initial. Participants were randomly assigned to either a female-doctor group or a male-doctor group and the sign-up sheet indicates the gender of the doctors. This design is to ensure that there are only two dimensions (caste group and experience level) that the four doctors differ.

Participants were instructed to rank the four doctors from their most desired (rank 1) to their least desired (rank 4), without the possibility of an equal rank. They were also explained that they had a higher chance of getting the more preferred doctor than the less preferred doctor. All participants were assigned a doctor that matched one of the four profiles.[12] Thus, the elicited

[10] The sign-up sheet also indicated that, in case the participant was not assigned to any of the listed doctors, an alternative doctor would be provided. This statement is included to ensure that we are not deceptive when the doctors assigned are different from those the listed.

[11] The general-category surnames used are: Bajpai, Dixit, Mishra, and Pandey. The backward-caste surnames used are: Kanaujiya, Katiyar, Kureel, Pal, Sonkar, Valmiki, Vishwakarma, and Yadav. Online Appendix C reports the likelihoods of these surnames being correctly associated with the corresponding caste groups based on a sample of 320 individuals residing in Kanpur Nagar district. For all surnames, the probabilities of them being correctly associated with either the general category or backward caste groups are at least 90%.

[12] 43% of patients ended up getting their most preferred doctor profile, 38% patients ended up getting their second most preferred doctor profile, 12% patients ended up getting their third most preferred doctor profile, 7% patients ended up getting their least preferred doctor profile.

rankings are incentivized. There are several reasons why we did not allow participants to rank doctors equally. First, the fact that a participant chooses one doctor over another is completely consistent with the theoretical notion of indifference between two doctors, because when a patient is indifferent between two doctors the patient chooses one at random. Second, by forcing participants to give strict ranking, we prevent the situations in which participants with weak preferences give equal ranking out of social desirability concern. Third, as long as we empirically detect ranking differences by doctor type, any measurement errors due to indifferences are differenced out on average.[13]

In stage two, the participants filled out a short demographic and social economic survey. The survey collects information about their age, gender, caste identity, caste identity, educational attainment, religious affiliation, etc. The short survey also includes questions about their attitudes toward individuals of different castes. By surveying them after the elicitation exercise, we minimized any potential priming effect. The correspondence study effectively concluded by the end of stage two.

In stage three, we informed the participants about the doctor they were assigned to and the location and time of their upcoming health-check appointment. In stage four, the mobile clinic arrived in the locality to deliver service. The service was delivered within one week of registration.

4. Results

4.1. Evidence of caste discrimination

We examine whether on average there is any evidence of caste discrimination first by pooling the responses of all patients and ignore the preference rankings of doctors at the individual patient level. We focus on the share of first rank that each caste-experience doctor-type receives in the full sample, before splitting the sample by the caste of patients. By focusing on first preferences, we illustrate why it can be difficult for choice data to help identify the primary source of discrimination.

[Figure 4 here]

[13] Some patients may be indifferent between some or all of the four doctors. In this case, they may randomize the rankings. The randomization may have implications on the estimated extent of taste-based discrimination. We examine the sensitivity of our results to such possibilities in Section 4.3 and Section 5.

Panel A of Figure 4 shows that backward-caste, high-experience ($c_B e_H$) doctors are ranked first the most (46%); general-category, high-experience ($c_G e_H$) doctors are ranked first the second most (43%); general-category, low-experience ($c_G e_L$) doctors are ranked first the third most (7%); and backward-caste, low-experience ($c_B e_L$) doctors are ranked first the least (4%) among the four doctors. If each patient was instructed to select only one doctor and they always selected their most preferred, the inferred average ranking of doctors is $c_B e_H \succ c_G e_H \succ c_G e_L \succ c_B e_L$, which is consistent with statistical discrimination and inconsistent with taste-based discrimination. The beliefs that are consistent with this preference ranking under statistical discrimination are either $\sigma^2_{\varepsilon,c_G} > \sigma^2_{\varepsilon,c_B}$, $\beta_{c_G} < \beta_{c_B}$, or both. These belief differences are not standard in the discrimination literature, which usually shows that the minority groups are discriminated; however, our results can be rationalized in the UP context because *backward-caste individuals are the majority group*. Specifically, if signals of doctors are positively correlated in the sample that each patient observed based on their past experience and the majority of patients (i.e., backward-caste patients) tend to encounter backward-caste doctors, then $\sigma^2_{\varepsilon,c_G} > \sigma^2_{\varepsilon,c_B}$ (see e.g., Ewens et al. 2014).

If there were no caste discrimination, for the same level of experience, there should not be any difference between general-category and backward-caste doctors. The difference in the share of first-ranked doctor is statistically significant between general-category and backward-caste doctors at each level of experience based on a simple t-test of the two proportions ($p < 0.05$). The significant differences also mean that by not offering patients the possibility to rank doctors equally is not a major concern for detecting discrimination, as we would not have detected the statistical differences had patients randomized doctors due to indifference. Without information about the discriminators, most field experiments would conclude that statistical discrimination is the main driver of differential treatments (choices) by caste of doctors based on the evidence in panel A of Figure 4.

As we also collect information about the group identities of discriminators, the natural next step is to split the sample by the caste of patients to examine the first-ranked doctors. Panel B of Figure 4 shows that for *backward-caste patients*, the inferred average ranking of doctors is $c_B e_H \succ c_G e_H \succ c_B e_L \succ c_G e_L$, which is consistent with both theories. Panel B of Figure 4 shows that for *general-category patients*, the inferred average ranking of doctors is $c_G e_H \succ c_B e_H \succ c_G e_L \succ c_B e_L$, which is also consistent with both theories. If patients taste-based discriminate on average, then the inferred rankings based on the patterns shown in panel B suggest that patients in both groups have homophily caste preferences, on average. If patients statistically discriminate on average, then the inferred rankings based on the patterns shown in panel A suggest that backward-caste patients hold the beliefs $\sigma^2_{\varepsilon,c_G} > \sigma^2_{\varepsilon,c_B}$, $\beta_{c_G} < \beta_{c_B}$, or both on average, while general-category patients hold the beliefs $\sigma^2_{\varepsilon,c_G} < \sigma^2_{\varepsilon,c_B}$, $\beta_{c_G} > \beta_{c_B}$, or both on average.

4.2. Bounding taste-based discrimination using individual preference rankings

As the aggregate-level choice data by caste of patients are consistent with both taste-based discrimination and statistical discrimination, we now turn to the preference rankings of doctors at the individual patient level to provide an upper bound of taste-based discrimination by examining the share of rankings that are consistent with both theories of discrimination. If this upper bound is less than 50% of rankings, then we can conclude that statistical discrimination is likely the main source of caste discrimination.

Table 2 reports the share of patients who report each of the 24 possible rankings of doctors and quantifies these differences by the type of discrimination and by caste of patients. Looking at the shares among backward-caste patients, 33.4% of them have preference rankings of doctors consistent with homophily taste-based discrimination and statistical discrimination, 23.2% of them have preference rankings of doctors consistent with both heterophily taste-based discrimination and statistical discrimination, and a little over 43.4% of them have preference rankings of doctors consistent with only statistical discrimination. Looking at the shares among general-category patients, 44% of them have preference rankings of doctors consistent with homophily taste-based discrimination and statistical discrimination, 15.4% of them have preference rankings of doctors consistent with both heterophily taste-based discrimination and statistical discrimination, and a little over 40.6% of them have preference rankings of doctors consistent with only statistical discrimination. Overall, both general-category and backward-caste patients are more likely to have preference rankings consistent with both homophily taste-based discrimination and statistical discrimination (36.3%) as opposed to preference rankings consistent with both heterophily taste-based discrimination and statistical discrimination (21.1%). More importantly, 42.7% of patients have preference rankings that are consistent with statistical discrimination only. Thus, the upper bound of taste-based discrimination can be as high as 57.4%.

Given the large extent of overlap between the two potential sources of discrimination, without imposing additional assumptions, such as homophily preference, we cannot conclude that statistical discrimination is the main driver of discrimination.

[Table 2 here]

4.3. The possibility of caste-blind and experience-blind patients

Given that our elicitation method prohibits patients from expressing indifferences, when patients are indifferent between two doctors they are forced to indicate one is being more preferred over the other. For example, when caste-blind patient who prefers more years of experience to fewer years of experience, the elicited rankings $c_B e_H \succ c_G e_H \succ c_B e_L \succ c_G e_L$ and $c_B e_H \succ c_G e_H \succ c_G e_L \succ c_B e_L$ are equivalent for them. However, the first case is classified as uniquely statistical and the second case is classified as consistent with both taste-based and statistical discrimination. If the patient provides the second ranking, then we incorrectly estimate the upper bound of taste-based discrimination. Thus, the prohibition of indifference in our elicitation method can potentially affect the inferred extent of taste-based discrimination.

To assess the sensitivity of the inferred extent of taste-based discrimination to the potential presence of caste-blind patients in our experiment, we may first consider that such caste-blind patients have preferences for experience. Due to the prohibition of ties in the elicitation exercise, they might randomize doctors for a given level of experience. Thus, among the preference rankings consistent with taste-based discrimination, rankings r_2, r_4, r_6, and r_8 in Table 2 can potentially be provided by caste-blind patients who randomize the order of doctors for a given level of experience. Similarly, among the preference rankings consistent with only statistical discrimination, rankings r_9, r_{10}, r_{11}, and r_{12} in Table 2 can potentially be provided by caste-blind patients who randomize the order of doctors who have the same experience. If we treat these preference rankings as potentially consistent with no caste discrimination and exclude them, then among the remaining preference rankings we get 54% of preference rankings that are consistent with taste-based discrimination. This share is similar to the 57% when we treat all preference rankings as having caste discrimination. Thus, our results are not sensitive to the potential of caste-blind patients who have preferences for experience.

It is also possible that there exist some discriminatory patients who have no preferences for experience, even though the strong relationship between first-ranked and experience shown in Figure 4 suggests such possibility is likely to be very small. If these experience-blind patients are taste-based discriminators, they might randomize the doctors for a given caste. Such randomization may make their preference rankings consistent with some of those classified as uniquely statistical discrimination. Among the preference rankings consistent with statistical discrimination only, rankings r_{15}, r_{16}, r_{21}, and r_{22} in Table 2 can potentially be provided by experience-blind patients who randomize the order of doctors for a given caste. If we treat these preference rankings as the results of randomization by experience-blind taste-based discriminators, then preference rankings that are consistent with taste-based discrimination can be as high as 67.7%.

4.4. Inaccurate beliefs

Our statistical discrimination model does not distinguish between accurate and inaccurate beliefs. Bohren et al. (2024) argue that it is useful to categorize statistical discrimination into those driven by accurate beliefs and those driven by inaccurate or biased beliefs. Such categorization can be potentially helpful for designing policy interventions. For example, if inaccurate beliefs are the drivers of caste discrimination, information provision regarding the correct distributions may help reduce caste discrimination.

To identify whether beliefs are accurate or not, past experiments typically examine performance on tasks, such as knowledge tests and mathematical problems, which tend to have objective answers or solutions. Inaccurate or biased beliefs are identified when the elicited beliefs about performance differ from the actual performance. It is generally difficult to objectively measure the quality of care provided by doctors (van den Heuvel et al. 2013). As we do not observe the true quality and signal (experience) distributions of doctors in India, it is difficult to identify the exact share of statistical discriminators with accurate beliefs, even if we have measures of their beliefs. Nonetheless, if the true quality of the four doctor types differs, then there can only be one preference ranking out of the 24 preference rankings that is consistent with the ranking of true quality. Furthermore, if a patient's preference ranking is consistent with the true quality ranking, it is not necessarily the case that the patient has accurate beliefs because the errors in their beliefs may not be large enough to change the preference order. Thus, we provide the upper (lower) bound of statistical discriminators with accurate (inaccurate) beliefs on the basis of preference rankings elicited in our field experiment.

We can infer the upper bound of statistical discriminators who hold accurate beliefs in our experiment from Table 2. As at most one of the 24 preference rankings can be consistent with accurate statistical discrimination, the maximum share of statistical discriminators with accurate beliefs ranges between 0.26% (r_{19}) and 20.97% (r_4). Thus, at least 79% of patients in the experiment have preference rankings consistent with those of statistical discriminators who hold inaccurate or biased beliefs.

Using the information in Table 2, we can also further classify preference rankings that are consistent with taste-based discrimination into two broad types: those consistent with accurate beliefs and those consistent with inaccurate beliefs. Among preference rankings that are consistent with taste-based discrimination (i.e., 57.4% of all), the maximum share of those that may also be consistent with accurate beliefs ranges between 1.2% (r_6) and 36.6% (r_4). As more than 63% of taste-based discrimination is also consistent with statistical discrimination with inaccurate beliefs, preference rankings of taste-based discriminators tend to resemble the preference rankings of statistical discriminators with inaccurate beliefs.

In sum, our experiment reveals that the vast majority of statistical discriminators (at least 79%) hold inaccurate or biased beliefs, and taste-based discrimination is indistinguishable from statistical discrimination with inaccurate beliefs most (at least 63%) of the time.

4.5. Stereotypes

According to Bordalo et al. (2016, 2019), stereotypes are inaccurate beliefs that manifest as an exaggeration of the truth. In the context of caste discrimination, such exaggeration can take the form of exaggerating the quality of general-category doctors due to the presence of caste-based reservations for backward-caste individuals to pursue a medical degree, for example. Caste-based reservations have been implemented in government-funded higher educational institutions for decades. Government-funded colleges and universities allot seats according to caste-based quotas, which assign 7.5% to STs, 15% to SCs, and 27% to OBCs (Deshpande, 2012). Candidates from backward-caste groups have to take entrance examinations for higher educational institutions, but they compete only among themselves to fill the reserved seats for their caste groups. Depending on the extent of competition within each caste group, admission requirements adjust to fill the reserved seats. These adjustments can lead to a gradation of admission scores, with general-category students facing the highest qualifying scores, and backward-caste students facing lower qualifying scores (Bertrand et al., 2010; Bagde et al., 2016; Deshpande, 2012; Frisancho & Krishna, 2016). Given that general-category students face the highest qualifying scores, the stereotype that general-category doctors provide better quality of healthcare may arise. Given that the vast majority of statistical discriminators hold inaccurate beliefs, we now investigate the extent to which these inaccurate beliefs can be driven by general-category caste stereotype.

To see how the presence of general-category stereotype can influence the preference rankings of doctors, consider the case when the true quality of doctors is such that a general-category doctor has better average quality than a backward-caste doctor for a given level of experience. The preference ranking r_2 in Table 2, $c_G e_H \succ c_B e_H \succ c_G e_L \succ c_B e_L$, is consistent with this true quality ranking. The existence of a general-category stereotype can lead to the second true-quality ranked and third true-quality ranked doctors being swapped in the preference order such that the preference ranking becomes $c_G e_H \succ c_G e_L \succ c_B e_H \succ c_B e_L$ (i.e., r_1 in Table 2). In this case, r_1 can be the results of inaccurate beliefs that fit Bordalo et al.'s (2016, 2019) definition of stereotype. The reason is that it has a "kernel of truth" to it given that the general-category doctor with more years of experience is ranked better than the backward-caste doctor with the same years of experience while the general-category doctor with a low level of experience is ranked better than the backward-caste doctor with the same years of experience. Thus, both r_1 and r_2 can be due to the presence of general-category stereotype because if the extent of stereotype is not sufficiently large to swap the order of the doctors, the preference ranking is identical to the true quality ranking. More generally, for a given preference ranking that is assumed to be consistent with the true-quality ranking, preference rankings that are consistent with the assumed true quality ranking as well as those that are not can be driven by general-category stereotype.

[Table 3]

Table 3 lists the preference rankings that can be consistent with general-category stereotype in column 2 for each assumed true quality ranking. For example, consider the case when r_3 is consistent with the assumed true quality ranking. The truth here is that backward-caste doctors have better quality than general-category doctors for both levels of experience. If general-category stereotype exists, a patient may express the preference ranking r_4. In this case, the more experienced general-category doctor is preferred to the less experienced backward-caste doctor, but the more experienced backward-caste doctor remains as the most preferred and the less experienced general-category doctor remains as the least preferred. Thus, although r_4 is consistent with inaccurate beliefs, it has a kernel of truth to it. Similarly, r_9 and r_{10} are also consistent with general-category stereotype when r_3 is consistent with the assumed true quality ranking. In contrast, r_1 is also consistent with inaccurate beliefs and an inflated quality of general-category doctors, but it lacks a kernel of truth to it because all backward caste doctors are ranked at the bottom when r_3 is consistent with the true quality ranking. Similarly, r_{11} is also consistent with inaccurate beliefs and an inflated quality of general-category doctors, but we do not classify it as consistent with a general-category stereotype because the quality-experience relationship within a doctor's caste group is inconsistent with that revealed in r_3.

Drawing upon the data from Table 2, column 4 of Table 3 reports the maximum share of patients whose preference rankings can be considered as consistent with having general-category stereotypical beliefs. Note that all preference rankings can be consistent with inaccurate beliefs because a preference ranking consistent with the assumed true quality ranking can also be driven by inaccurate beliefs when the cardinal differences between doctors are not sufficiently large to affect the ordinal ranking. The maximum share of patients with general-category stereotype ranges between 1% and 82%. In only two cases of the assumed truth, r_3 and r_4, out of the 24 possible cases, there are more than 50% of patients whose preference rankings are potentially driven by general-category stereotype. In the case of r_3, the assumed truth is that backward-caste doctors provide better average quality of healthcare than general-category doctors irrespective of their years of experience. The systematic reviews by Ajmi and Aase (2021) and Choudhry et al. (2005) show that the experience of doctors is associated with healthcare quality across a wide range of countries. Thus, it seems highly unlikely that r_3 reflects the truth. In the case of r_4, backward-caste doctors provide better average quality of healthcare than general-category doctors do at each level of experience. Although we do not have the data to support such quality ranking, it is entirely possible. When the true quality ranking is such, the preferences of the majority of patients are consistent with general-category stereotype.[14] Overall, the results here suggest that general-

[14] It is the case even if we assume that r_4 is fully driven by accurate beliefs. In this case, the share of patients with general-category stereotypical beliefs falls from 82% to 61% given that 21% of patients report r_4.

category stereotype is unlikely to be the primary source of inaccurate beliefs in the majority of cases.

4.6. The extent of robust caste preferences

Our results reveal that we cannot tell whether inaccurate beliefs or caste preferences are the primary source of caste discrimination against doctors in India. Nonetheless, policies aimed at dispelling misinformation and improving information exchanges between caste groups are not only useful for addressing inaccurate beliefs, but may also be useful for reducing social distance between groups and encouraging intergroup interaction. If the individuals targeted by the policies are hostile towards the discriminated groups, then it can be challenging to implement such policies. Thus, estimating the share of individuals with potential animosity towards the discriminated caste group can help identify the extent to which information interventions may be deployed.

We classify participants as potentially bearing animosity towards the discriminated caste group by using an incentivized measure of group animosity obtained from a lab-in-the-field experiment on a subset of patients (see Online Appendix D for details). Specifically, we adopt Ferschtman and Gneezy's (2001) and List's (2004) notion of group animosity revealed in dictator games. When a patient gives strictly (weakly) less to the discriminated group than the favored group (as inferred from the caste preference in the field experiment) in the dictator games, we argue that the patient potentially holds strong (weak) animosity towards the discriminated group. We may view individuals who taste-based discriminate doctors of a particular caste group and also give strictly (weakly) less to people from that group as having a robust caste preference and the remaining individuals as having non-robust caste preferences. Using the relative amounts allocated to different groups in dictator games to measure relative caste preferences are also consistent with the Fehr and Schmidt's (1999) notion of other-regarding preferences that are shaped by social proximity. For patients with non-robust caste preferences, if they report preference rankings of doctors consistent with taste-based discrimination, we treat them as effectively non-tasted-based discriminators. If the share of patients who have preference rankings consistent with taste-based discrimination falls below 50% after we reclassify patients with non-robust caste preferences as potential non-taste-based discriminators, then there is the potential for policies aimed at dispelling misinformation, improving information exchanges, and encouraging intergroup interaction to be effective for the majority of individuals.

[Table 4 here]

Table 4 shows that among the subset of patients who participated in the dictator games, roughly 60.2% of them have preference rankings of doctors that are indistinguishable between taste-based and statistical discrimination This share is similar to the 57.4% for the full sample of patients. Using the notion of strong (weak) group animosity, we find that roughly 41% (60.3%) of patients who have preference rankings consistent with taste-based discrimination are potentially holding animosity against the discriminated group. After reclassifying the remaining patients who have non-robust caste preferences and preference rankings consistent with taste-based discrimination as patients who are effectively non-taste-based discriminators, the extent of taste-based discrimination decreases to between 24.7% and 36.3%. Thus, although we cannot clearly distinguish statistical discrimination and taste-based discrimination for the majority of patients, policies that involve information interventions are potentially useful for addressing discrimination for a large segment of the population.

5. Robustness

We perform several robustness checks to examine whether the results are sensitive to a number of assumptions made. The results are reported and discussed in Online Appendix E. Here, we briefly summarize the key points.

First, we have assumed that patients are incentivized to report their preferences accurately. However, because patients were told that there was a possibility that they would be assigned a doctor different from those listed on the sign-up sheet, the incentive for some patients to report their preferences accurately might be weakened. As a result, they might randomize the orders of doctors. In the most extreme case, all patients randomize the orders. As one third of all possible rankings (8/24) are consistent with both theories, while two thirds of all possible rankings (16/24) are uniquely statistical, randomization is more likely to turn patients' preference rankings from those that are consistent with both theories into those that are consistent with statistical discrimination only. To examine the implication of this extreme form of randomization, we apply differential weights so that preference rankings that are consistent with taste-based discrimination get twice the weight that preference rankings that are uniquely statistical do (see Online Appendix

E.1). As lower weights are placed on preference rankings that are consistent with statistical discrimination only, the share of patients with preference rankings consistent with both theories becomes even higher (72.9%). As the share of preference rankings that are consistent with accurate statistical discrimination is no more than 26.7%, the majority of statistical discriminators hold inaccurate beliefs. Similarly, as no more than 36.5% of taste-based discrimination overlaps with accurate statistical discrimination, taste-based discrimination is generally indistinguishable from inaccurate statistical discrimination. Only when r_3 is assumed to be consistent with the true quality ranking, more than 50% of patients whose preference rankings are potentially driven by general-category stereotype. As r_3 is highly unlikely to reflect the truth, the conclusion that general-category stereotype is unlikely to be the primary source of inaccurate beliefs remains unchanged. Importantly, we still find that only a minority of patients have robust caste preferences.

Second, we examine whether our results are sensitive to allowing taste-based discriminatory patients to have preferences for the gender of a doctor, which is an additional unobserved good or characteristic of a doctor that is correlated with experience but not theoretically modelled under taste-based discrimination (see Online Appendix E.2). We find that the estimated upper bound of taste-based discrimination is similar by gender of doctor. Thus, our results are robust to Heckman's (1998) critique about using experiments to detect taste-based discrimination.

6. Conclusion

This paper highlights the empirical challenges in separating the predictions of taste-based discrimination from the predictions of statistical discrimination using a simple theoretical framework that models discrimination of patients against doctors based on their castes and years of experience. The framework informs our approach to determine the extent of taste-based discrimination using the revealed preferences of patients. Specifically, by eliciting individual patients' preference rankings of doctors in a controlled field experiment conducted in Uttar Pradesh, India, we first show that when we focus on the most-preferred doctor to examine discriminatory choices made by patients, the aggregate level data can be consistent with both (homophily) taste-based discrimination and statistical discrimination. The results highlight one of the major challenges in separating taste-based and belief-based explanations that most field experiments face when only choice data are available. We then demonstrate that even if we examine individual patients' preference rankings of doctors, in 57% of patients, we cannot distinguish whether they taste-based discriminate or statistically discriminate. The results highlight that if we only allow preferences to satisfy some fundamental axioms of preferences without imposing additional assumptions on preferences, it is difficult to rule out the role of taste in driving most discriminatory behaviors.

Because statistical discrimination is a belief-based explanation, we examine whether the variation in preference rankings is primarily consistent with accurate beliefs or inaccurate beliefs. Given that at most one preference ranking can be consistent with the truth and the most commonly observed preference ranking in the sample represents only 21% of all rankings reported by the patients, most statistical discriminators are likely to hold inaccurate beliefs. Similarly, we find that most taste-based discrimination is indistinguishable from inaccurate statistical discrimination. As caste-based reservations in higher educational institutions have been implemented in India to help improve representation of backward-caste individuals in medical professions for decades, we further examine whether the stereotype that general-category doctors are of higher quality may drive the inaccurate beliefs. Our results suggest that such a stereotype is unlikely to be the primary source of inaccurate beliefs in the majority of cases.

As most statistical discrimination in our experiment is consistent with inaccurate beliefs, policies aimed at dispelling misinformation and encouraging intergroup information exchanges are potentially useful for addressing caste discrimination. Such policies may also be useful for addressing taste-based discrimination and more implementable among discriminators who do not bear animosity towards the discriminated group. To examine the extent to which these policies may be useful for addressing inaccurate beliefs as well as taste-based discrimination, we reclassify patients who are unlikely to be hostile towards the discrimination groups as potential non-taste-based discriminators. Using incentivized measures of animosity based on giving to different caste groups in dictator games to define animosity, we estimate that information interventions are potentially useful for addressing discrimination among the majority of patients.

More generally, although our controlled field experiment does not vary the availability of quality signal, our methodology that exploits revealed preference rankings to identify the extent of taste-based discrimination and accurate statistical discrimination can be applied in many other settings to understand the primary source of discrimination along different dimensions, such as race, ethnicity, gender, religion, age, etc. Given the recent major events around the world, such as the Black Lives Matter movement and Covid-19 related racism, identifying the primary source of discrimination and the extent of accurate beliefs can be a first step towards implementing appropriate policy responses to address the underlying problems.

References

1. Aigner, D. J., & Cain, G. G. (1977). Statistical theories of discrimination in labor markets. *ILR Review*, *30*(2), 175-187.
2. Ajmi, S. C., & Aase, K. (2021). Physicians' clinical experience and its association with healthcare quality: a systematised review. *BMJ Open Quality, 10(4)*, e001545.
3. Akerlof, G. A. (1997). Social distance and social decisions. *Econometrica*, *65*(5), 1005-1027.
4. Altonji, J. G., & Blank, R. M. (1999). Race and gender in the labor market. In: O. Ashenfeher and D. Card (Eds.), *Handbook of Labor Economics*, Vol. *3*, Amsterdam: Elsevier Science, pp. 3143-3259.
5. Altonji, J.G., & Pierret, C.R. (2001). Employer learning and statistical discrimination. *The Quarterly Journal of Economics*, *116*(1), 313-50.
6. Anderson, L., Fryer, R., & Holt, C. (2006). Discrimination: Experimental evidence from psychology and economics. In: W. Rogers (Ed.), *Handbook on the Economics of Discrimination*, Cheltenham: Edward Elgar Publishing, 97-115.
7. Anwar, S., & Fang, H. (2006). An alternative test of racial prejudice in motor vehicle searches: Theory and evidence. *American Economic Review, 96*(1), 127-51.
8. Arcidiacono, P., Bayer, P., & Hizmo, A. (2010). Beyond signaling and human capital: Education and the revelation of ability. *American Economic Journal: Applied Economics*, *2*, 76-104
9. Arcidiacono, P., & Lovenheim, M. (2016). Affirmative Action and the quality-fit trade-off. *Journal of Economic Literature*, *54*(1), 3-51.
10. Arnold, D., Dobbie, W., & Yang, C. (2018). Racial bias in bail decisions. *Quarterly Journal of Economics*, *133*(4), 1885-1932.
11. Arrow, K. (1973). The theory of discrimination. In O. Ashenfelter & A. Rees (Eds.), *Discrimination in Labor Markets.* Princeton, NJ: Princeton University Press, pp. 3-33.
12. Baert, S. (2018). Hiring discrimination: An overview of (almost) all correspondence experiments since 2005. *Audit Studies: Behind the Scenes with Theory, Method, and Nuance*, 63-77.
13. Bagde, S., Epple, D., & Taylor, L. (2016). Does affirmative action work? Caste, gender, college quality, and academic success in India. *American Economic Review*, *106*(6), 1495-1521.
14. Banerjee, A., Bertrand, M., Datta, S., & Mullainathan, S. (2009). Labor market discrimination in Delhi: Evidence from a field experiment. *Journal of Comparative Economics*, *37*(1), 14-27.
15. Bayer, P., Ferreira, F., & Ross, S. L. (2018). What drives racial and ethnic differences in high-cost mortgages? The role of high-risk lenders. *Review of Financial Studies*, *31*(1), 175-205.
16. Bertrand, M., Hanna, R., & Mullainathan, S. (2010). Affirmative action in education: Evidence from engineering college admissions in India. *Journal of Public Economics*, *94*(1-2), 16-29.
17. Bertrand, M., & Duflo, E. (2017). Field experiments on discrimination. *Handbook of Economic Field Experiments*, *Vol. 1*, 309-393.

18. Bertrand, M., & Mullainathan, S. (2004). Are Emily and Greg more employable than Lakisha and Jamal? A field experiment on labor market discrimination. *American Economic Review*, *94*(4), 991-1013.
19. Bhadra, M. (2011). Indian women in medicine: an enquiry since 1880. *Indian Anthropological Association*, *41*(1), 17-43.
20. Bohren, J. A., Imas, A., & Rosenberg, M. (2019). The dynamics of discrimination: Theory and evidence. *American Economic Review*, *109*(10), 3395-3436.
21. Bohren, J. A., Haggag, K., Imas, A., & Pope, D. G. (2024). Inaccurate statistical discrimination: An identification problem. *Review of Economics and Statistics*, forthcoming.
22. Bordalo, P., Coffman, K., Gennaioli, N., & Shleifer, A. (2016). Stereotypes. *Quarterly Journal of Economics, 131(4)*, 1753-1794.
23. Bordalo, P., Coffman, K., Gennaioli, N., & Shleifer, A. (2019). Beliefs about gender. *American Economic Review, 109(3)*, 739-73.
24. Cason, T. N., Sharma, T., & Vadovič, R. (2020). Correlated beliefs: Predicting outcomes in 2× 2 games. *Games and Economic Behavior*, *122*, 256-276.
25. Chan, A. (2023). Discrimination against doctors: A field experiment. Unpublished manuscript.
26. Charles, K. K., & Guryan, J. (2011). Studying discrimination: Fundamental challenges and recent progress. *Annual Review of Economics*, *3*(1), 479-511.
27. Choudhry, N. K., Fletcher, R. H., & Soumerai, S. B. (2005). Systematic review: the relationship between clinical experience and quality of health care. *Annals of Internal Medicin*e, *142(4)*, 260-273.
28. Coate, S., & Loury, G. C. (1993). Will affirmative-action policies eliminate negative stereotypes?. *American Economic Review*, 1220-1240.
29. Cooper, R.W., DeJong, D.V., Forsythe, R. and Ross, T.W. (1990). Selection criteria in coordination games: Some experimental results. *American Economic Review*, 80(1), 218-233.
30. Currarini, S., Jackson, M. O., & Pin, P. (2009). An economic model of friendship: Homophily, minorities, and segregation. *Econometrica*, *77*(4), 1003-1045.
31. Deshpande, A. (2012). Social Justice through affirmative action in India: An assessment. In *Capitalism on Trial: Explorations in the Tradition of Thomas Weisskopf*, Northampton, MA: Edward Elgar Publishing.
32. Doleac, J. L., & Stein, L. C. (2013). The visible hand: Race and online market outcomes. *Economic Journal*, *123*(572), F469-F492.
33. Ewens, M., Tomlin, B., & Wang, L. C. (2014). Statistical discrimination or prejudice? A large sample field experiment. *Review of Economics and Statistics*, *96*(1), 119-134.
34. Fang, H., & Moro, A. (2011). Theories of statistical discrimination and affirmative action: A survey. *Handbook of Social Economics*, *Vol. 1*, 133-200.
35. Fehr, E., & Schdmit, K. M. (1999). A theory of fairness, competition, and cooperation. *Quarterly Journal of Economics*, *114*(3), 817-868.
36. Ferschtman, C., & Gneezy, U. (2001). Discrimination in a segmented society: An experimental approach. *Quarterly Journal of Economics*, *116*(1), 351-377.
37. Frisancho, V., & Krishna, K. (2016). Affirmative action in higher education in India: Targeting, catch up and mismatch. *Higher Education*, 71(5), 611-649.
38. Gneezy, U., List, J., & Price, M. K. (2012). Toward an understanding of why people discriminate: Evidence from a series of natural field experiments. National Bureau of Economic Research (NBER) Working Paper No. 17855.

39. Giulietti, C., Tonin, M., & Vlassopoulos, M. (2019). Racial discrimination in local public services: A field experiment in the United States. *Journal of the European Economic Association*, *17*(1), 165-204.
40. Guryan, J., & Charles, K. K. (2013). Taste-based or statistical discrimination: the economics of discrimination returns to its roots. *Economic Journal*, *123*(572), F417-F432.
41. Hanna, R. N., & Linden, L. L. (2012). Discrimination in grading. *American Economic Journal: Economic Policy*, *4*(4), 146-68.
42. Heckman, J. J. (1998). Detecting Discrimination. *Journal of Economic Perspectives, 12(2),* 101–116.
43. Holzer, H., & Neumark, D. (2000). Assessing affirmative action. *Journal of Economic literature*, *38*(3), 483-568.
44. Holzer, H. J., & Neumark, D. (2006). Affirmative action: What do we know? *Journal of Policy Analysis and Management*, *25*(2), 463-490.
45. Islam, A., Pakrashi, D., Vlassopoulos, M., and Wang, L. C. (2021). Stigma and misconceptions in the time of the COVID-19 pandemic: A field experiment in India. *Social Science and Medicine*, 278, 113966.
46. Lane, T. (2016). Discrimination in the laboratory: A meta-analysis of economics experiments. *European Economic Review*, *90*, 375-402.
47. Lang, K., & Lehmann, J.Y.K. (2012). Racial discrimination in the labor market: Theory and empirics. *Journal of Economic Literature*, *50*(4), 959-1006.
48. Lang, K., & Kahn-Lang Spitzer, A. (2020). Race discrimination: An economic perspective. *Journal of Economic Perspectives*, *34*(2), 68-89.
49. Levati, M. V., Nicholas, A., & Rai, B. (2014). Testing the single-peakedness of other-regarding preferences. *European Economic Review*, *67*, 197-209.
50. List, J. A. (2004). The nature and extent of discrimination in the marketplace: Evidence from the field. *Quarterly Journal of Economics*, *119*(1), 49-89.
51. Madheswaran, S., & Attewell, P. (2007). Caste discrimination in the Indian urban labour market: Evidence from the National Sample Survey. *Economic and Political Weekly*, 42(41), 4146-4153.
52. McPherson, M., Smith-Lovin, L., & Cook, J. M. (2001). Birds of a feather: Homophily in social networks. *Annual Review of Sociology*, *27*(1), 415-444.
53. Neumark, D. (2018). Experimental research on labor market discrimination. *Journal of Economic Literature*, *56*(3), 799-866.
54. Onuchic, P. (2024). Recent contributions to theories of discrimination. Unpublished manuscript.
55. Phelps, E. S. (1972). The statistical theory of racism and sexism. *American Economic Review*, *62*(4), 659-661.
56. Riach, P. A. & Rich, J. (2002). Field experiments of discrimination in the market place. *Economic Journal*, *112*(483), F480-F518.
57. Rich, J. (2014). *What Do Field Experiments of Discrimination in Markets Tell Us? A Meta Analysis of Studies Conducted since 2000*. IZA Discussion Paper No. 8584.
58. Siddique, Z. (2011). Evidence on caste-based discrimination. *Labour Economics*, 18, S146-S159.
59. Siddique, A., Vlassopoulos, M., & Y. Zenou (2023), “Market competition and discrimination,” *European Economic Review* 152, 104361.

60. Thorat, S., & Attewell, P. (2007). Legacy of social exclusion: A correspondence study of job discrimination in India. *Economic and Political Weekly*, 42(41), 4141-4145.
61. van den Heuvel, J., Niemeijer, G. C., & Does, R. J. (2013). Measuring healthcare quality: the challenges. *International Journal of Health Care Quality Assurance; 26(3)*, 269-78.
62. Verdier, T., & Zenou, Y. (2004). Racial beliefs, location, and the causes of crime. *International Economic Review*, *45*(3), 731-760.
63. Zussman, A. (2013). Ethnic discrimination: Lessons from the Israeli online market for used cars. *Economic Journal*, *123*(572), F433-F468.

FIGURES

Figure 1: Backward-caste homophily patients' preferences for and rakings of doctors with different levels of social proximity and quality of healthcare

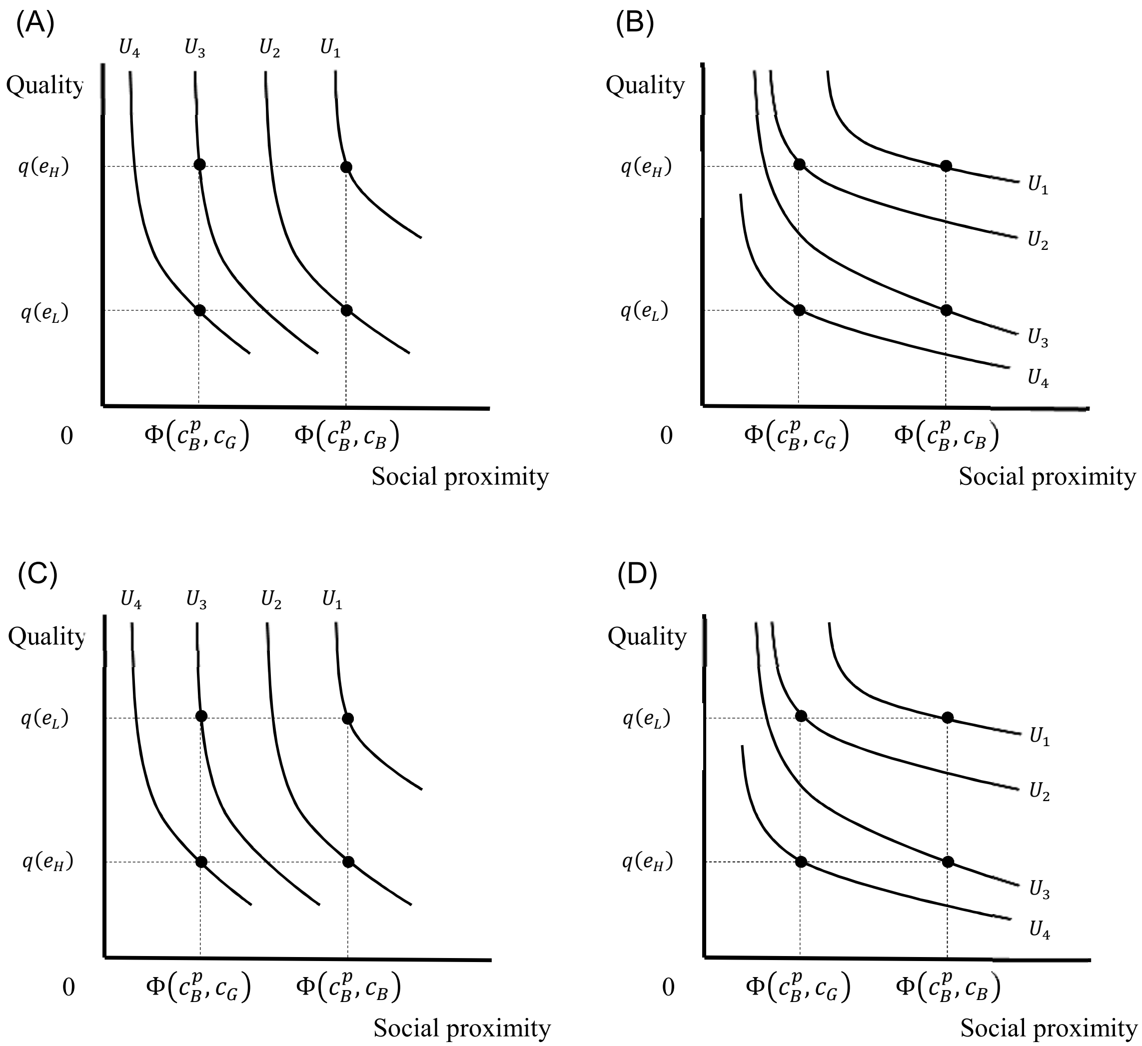

Notes: In panels A and B, the backward-caste homophily patient prefers more years of experience to fewer years of experience. In panels C and D, the backward-caste homophily patient prefers fewer years of experience to more years of experience. All four panels are identical for general-category heterophily patients.

Figure 2: Backward-caste heterophily patients' preferences for and rakings of doctors with different levels of social proximity and quality of healthcare

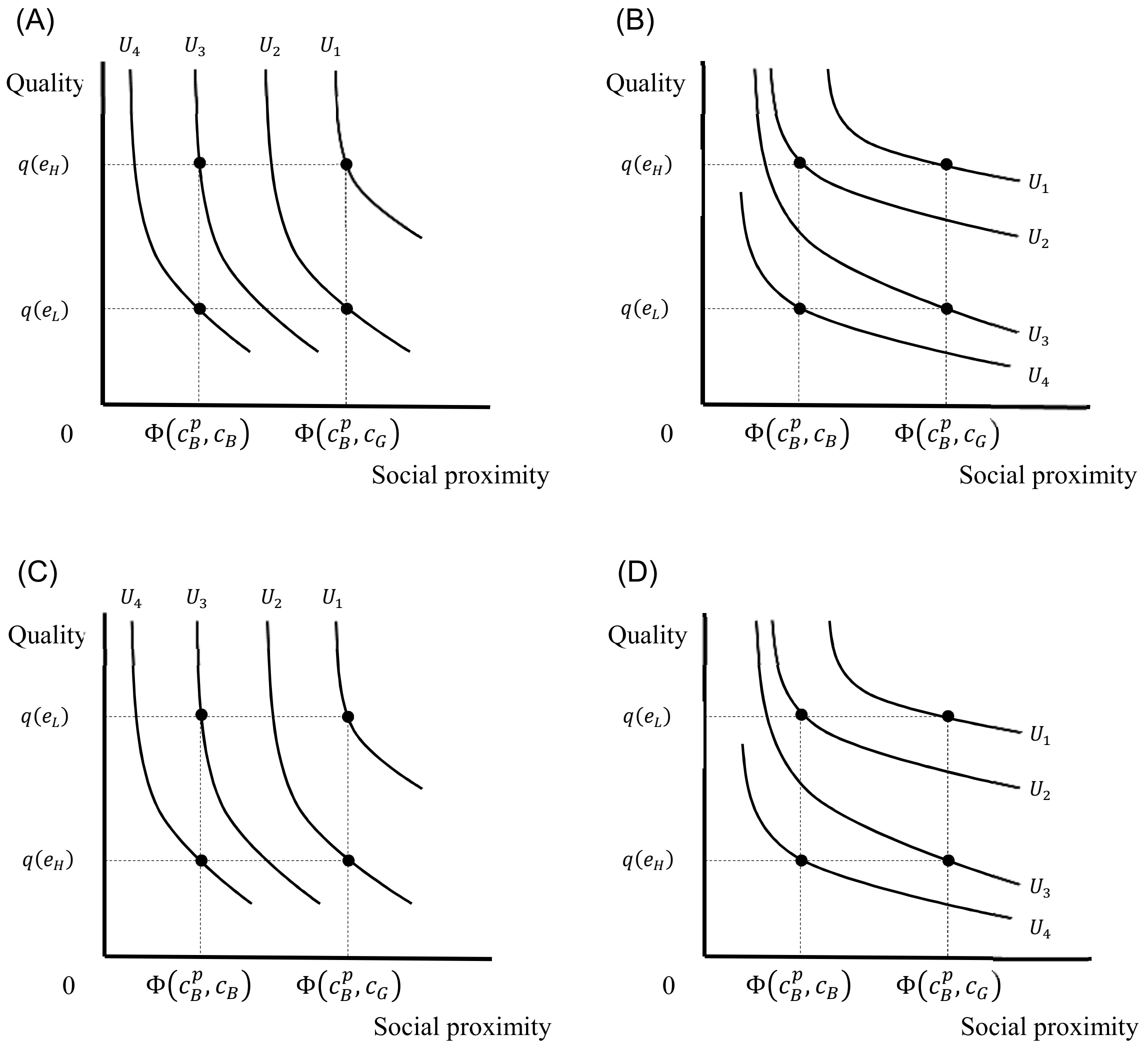

Notes: In panels A and B, the backward-caste heterophily patient prefers more years of experience to fewer years of experience. In panels C and D, the backward-caste heterophily patient prefers fewer years of experience to more years of experience. All four panels are identical for general-category homophily patients.

Figure 3: The relationship between expected quality of healthcare and years of experience when patients statistically discriminate

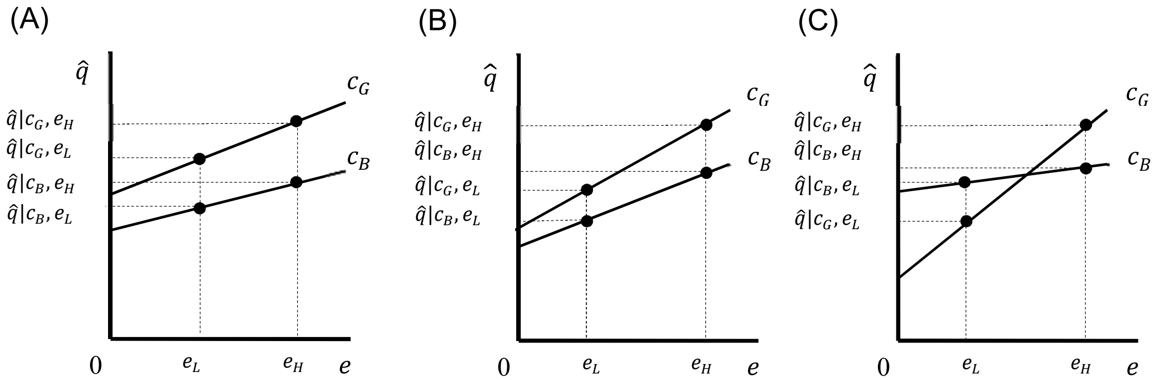

Notes: In all cases, the beliefs are assumed to be $\beta_{c_G} = \beta_{c_B}$, $\sigma^2_{q,c_G} = \sigma^2_{q,c_B}$, $Cov(q_c, e_c) > 0$, and $\sigma^2_{\varepsilon,c_B} > \sigma^2_{\varepsilon,c_G}$.

Figure 4: Share of first rank by doctor type and caste of patient

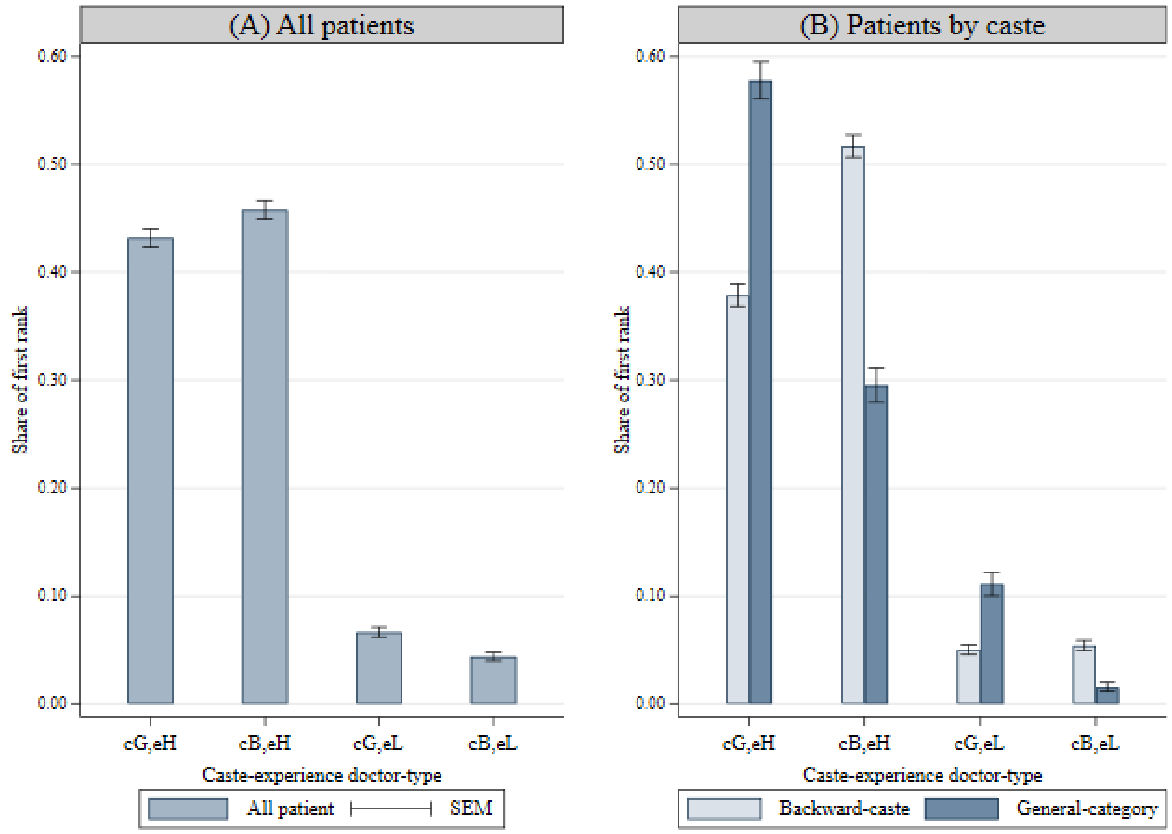

Notes: First rank is a dummy variable that takes the value of one if the doctor is chosen as the most preferred. cG,eH = general-category high-experience doctor; cB,eH = backward-caste high-experience doctor; cG,cL = general-category low-experience doctor; and cB,eL = backward-caste low-experience doctor. The standard error bar denotes the standard error of the mean.

TABLES

Table 1: Descriptive statistics

	Uttar Pradesh	Experimental Sample	
	Mean	Mean	Std. Dev.
Male	0.51	0.51	0.50
Age	38.0	37.8	14.3
General category	0.27	0.27	0.44
Hindu	0.80	0.80	0.40
College educated	0.08	0.11	0.32
Below poverty line	0.29	0.34	0.47
Urban resident	0.34	0.34	0.48

Notes: The field experiment sample include 3,128 participants. All statistics for Uttar Pradesh were sourced from NSS 68th Round, 2011-2012, except the below poverty line figure which came from World Bank (2016).

Table 2: Distribution of preference rankings consistent with different theories

	Backward-caste	General-category	All
A. Taste-based or statistical discrimination			
$r_1: c_G e_H \succ c_G e_L \succ c_B e_H \succ c_B e_L$	2.53	22.97	7.99
$r_2: c_G e_H \succ c_B e_H \succ c_G e_L \succ c_B e_L$	18.72	16.63	18.16
$r_3: c_B e_H \succ c_B e_L \succ c_G e_H \succ c_G e_L$	7.16	0.84	5.47
$r_4: c_B e_H \succ c_G e_H \succ c_B e_L \succ c_G e_L$	23.47	14.11	20.97
$r_5: c_G e_L \succ c_G e_H \succ c_B e_L \succ c_B e_H$	1.31	3.59	1.92
$r_6: c_G e_L \succ c_B e_L \succ c_G e_H \succ c_B e_H$	0.61	0.84	0.67
$r_7: c_B e_L \succ c_B e_H \succ c_G e_L \succ c_G e_H$	1.61	0.00	1.18
$r_8: c_B e_L \succ c_G e_L \succ c_B e_H \succ c_G e_H$	1.18	0.48	0.99
B. Statistical discrimination			
$r_9: c_G e_H \succ c_B e_H \succ c_B e_L \succ c_G e_L$	12.22	8.37	11.19
$r_{10}: c_B e_H \succ c_G e_H \succ c_G e_L \succ c_B e_L$	17.06	13.40	16.08
$r_{11}: c_G e_L \succ c_B e_L \succ c_B e_H \succ c_G e_H$	0.61	0.24	0.51
$r_{12}: c_B e_L \succ c_G e_L \succ c_G e_H \succ c_B e_H$	0.79	0.72	0.77
$r_{13}: c_G e_H \succ c_B e_L \succ c_B e_H \succ c_G e_L$	0.35	0.24	0.32
$r_{14}: c_G e_H \succ c_B e_L \succ c_G e_L \succ c_B e_H$	1.18	0.72	1.05
$r_{15}: c_G e_H \succ c_G e_L \succ c_B e_L \succ c_B e_H$	2.88	8.85	4.48
$r_{16}: c_B e_L \succ c_B e_H \succ c_G e_H \succ c_G e_L$	1.09	0.00	0.80
$r_{17}: c_B e_L \succ c_G e_H \succ c_G e_L \succ c_B e_H$	0.39	0.12	0.32
$r_{18}: c_B e_L \succ c_G e_H \succ c_B e_H \succ c_G e_L$	0.35	0.24	0.32
$r_{19}: c_G e_L \succ c_B e_H \succ c_B e_L \succ c_G e_H$	0.22	0.36	0.26
$r_{20}: c_G e_L \succ c_B e_H \succ c_G e_H \succ c_B e_L$	0.57	0.72	0.61
$r_{21}: c_G e_L \succ c_G e_H \succ c_B e_H \succ c_B e_L$	1.70	5.38	2.69
$r_{22}: c_B e_H \succ c_B e_L \succ c_G e_L \succ c_G e_H$	3.18	0.24	2.40
$r_{23}: c_B e_H \succ c_G e_L \succ c_G e_H \succ c_B e_L$	0.35	0.60	0.42
$r_{24}: c_B e_H \succ c_G e_L \succ c_B e_L \succ c_G e_H$	0.48	0.36	0.45
C. Summary			
Homophily taste-based or statistical	33.42	44.02	36.25
	(0.99)	(1.72)	(0.86)
Heterophily taste-based or statistical	23.17	15.43	21.10
	(0.88)	(1.25)	(0.73)
Uniquely statistical discrimination	43.41	40.55	42.65
	(1.04)	(1.70)	(0.88)

Notes: Taste-based discrimination is indistinguishable from statistical discrimination. High experience may signal better or worse quality of healthcare service when classifying with which theory a particular ranking is consistent. Standard errors are reported in parentheses.

Table 3: Preference rankings that are consistent with general-category stereotype

True-quality consistent preference ranking	Preference rankings consistent with general-category stereotype	Maximum share with accurate beliefs	Maximum share with stereotypical beliefs
r_1	r_1	7.99	7.99
r_2	r_2, r_1	18.16	27.92
r_3	$r_3, r_2, r_4, r_9, r_{10}$	5.47	75.71
r_4	r_4, r_9, r_{10}	20.97	82.01
r_5	r_5	1.92	1.92
r_6	r_6, r_5	0.67	2.60
r_7	$r_7, r_6, r_8, r_{11}, r_{12}$	1.18	4.16
r_8	r_8, r_{11}, r_{12}	0.99	3.28
r_9	r_9, r_1, r_2	11.19	40.63
r_{10}	r_{10}, r_1, r_2	16.08	47.24
r_{11}	r_{11}, r_5, r_6	0.51	3.11
r_{12}	r_{12}, r_6	0.77	1.45
r_{13}	r_{13}, r_{14}, r_{15}	0.32	5.87
r_{14}	r_{14}, r_{15}	1.05	5.58
r_{15}	r_{15}	4.48	4.48
r_{16}	$r_{16}, r_{13}, r_{14}, r_{17}, r_{18}$	0.80	2.83
r_{17}	r_{17}, r_{14}	0.32	1.37
r_{18}	$r_{18}, r_{13}, r_{14}, r_{17}$	0.32	2.02
r_{19}	r_{19}, r_{20}, r_{21}	0.26	3.57
r_{20}	r_{20}, r_{21}	0.61	3.32
r_{21}	r_{21}	2.69	2.69
r_{22}	$r_{22}, r_{19}, r_{20}, r_{23}, r_{24}$	2.40	4.18
r_{23}	r_{23}, r_{20}	0.42	1.03
r_{24}	$r_{24}, r_{10}, r_{19}, r_{20}, r_{23}$	0.45	17.90

Notes: Column 1 lists the 24 possible preference rankings when each of them is consistent with the assumed true quality ranking of doctors. Column 2 lists the preference rankings that are consistent with a general-category stereotype for a given preference ranking assumed to be consistent with the true quality ranking. Column 2 includes the preference rankings that are consistent with the assumed true quality rankings because general-category stereotype may not be sufficiently large to swap the rank order of doctors. Column 3 reports the maximum share of patients with preference rankings consistent with accurate beliefs based on the assumed true quality ranking. Column 4 reports the maximum share of patients with preference rankings that are consistent with general-category stereotype. See Table 2 for details of rankings.

Table 4: Distribution of preference rankings consistent with different theories by caste of patient and classification of non-robust caste preferences

	Backward-caste	General-category	All
A. Before reclassification of caste preferences			
Homophily taste-based or statistical discrimination	34.92	47.58	38.17
	(2.52)	(4.50)	(2.22)
Heterophily taste-based or statistical discrimination	24.58	14.52	21.99
	(2.28)	(3.18)	(1.89)
Non-taste-based or statistical discrimination	40.50	37.90	39.83
	(2.60)	(4.37)	(2.23)
B. After reclassification based on strong animosity			
Homophily taste-based or statistical discrimination	22.35	13.71	20.12
	(2.20)	(3.10)	(1.83)
Heterophily taste-based or statistical discrimination	1.96	12.10	4.56
	(0.73)	(2.94)	(0.95)
Non-taste-based discrimination	75.70	74.19	75.31
	(2.27)	(3.95)	(1.97)
C. After reclassification based on weak animosity			
Homophily taste-based or statistical discrimination	29.89	22.58	28.01
	(2.42)	(3.77)	(2.05)
Heterophily taste-based or statistical discrimination	6.98	12.10	8.30
	(1.35)	(2.94)	(1.26)
Non-taste-based discrimination	63.13	65.32	63.69
	(2.55)	(4.29)	(2.19)

Notes: The sample size is 482 patients who participated in both the field experiment and dictator games. Panel A show shows that the data based on the sub-sample of 482 patients are similar to the full sample. In panel B, a patient is defined to have animosity against the discriminated caste group inferred from the preference ranking of doctors in the field experiment when they give less to the discriminated caste group in the dictator games. In panel C, a patient is defined to have animosity against the discriminated caste group inferred from the preference ranking of doctors in the field experiment when they give equal or less to the discriminated caste group in the dictator games. Standard errors are reported in the parentheses.

Online Appendix

A. Rankings of doctors under taste-based discrimination

A.1. Homophily taste-based discrimination

We present a simple additively separable utility function for patients of caste $c = c_B, c_G$ to illustrate the eight possible rankings under taste-based discrimination when quality of healthcare is permitted to be positively or negatively correlated with experience. We note that these rankings are true for any utility function that satisfies the five axioms of preferences specified in Section 2.1.

The utility function for a backward-caste patient c_L^p choosing a doctor of caste $c = c_B, c_G$ is given by:

$$U\big(\Phi\big(c_B^p, c\big), q\big) = q(e) - 1_{\theta_G} \tag{A.1}$$

where $1_{\theta_G} = \theta_G > 0$ if the doctor is from a general-category background and zero otherwise. The indicator 1_{θ_G} indicates whether there is any social distance between the doctor's caste group and the patient's caste group.

If $q'(e) > 0$ (i.e., doctors with more years of experience provide better quality health service), the possible rankings for a backward-caste patient compatible with (A.1) are:

$$c_B e_H \succ c_B e_L \succ c_G e_H \succ c_G e_L$$
$$c_B e_H \succ c_G e_H \succ c_B e_L \succ c_G e_L$$

In the first ranking, we need to give the condition for which $c_B e_L \succ c_G e_H$ (the other inequalities are always true by definition since $e_H > e_L$). The condition is:

$$\theta_G > q(e_H) - q(e_L) \tag{A.2}$$

In the second ranking, we need to give condition for which $c_G e_H \succ c_B e_L$ (the other inequalities are always true by definition since $e_H > e_L$). The condition is:

$$\theta_G < q(e_H) - q(e_L) \tag{A.3}$$

If $q'(e) < 0$ (i.e., doctors with fewer years of experience provide better quality health service), the possible rankings for a backward-caste patient compatible with (A.1) are:

$$c_B e_L \succ c_B e_H \succ c_G e_L \succ c_G e_H$$
$$c_B e_L \succ c_G e_L \succ c_B e_H \succ c_G e_H$$

In the first ranking, we need to give the condition for which $c_B e_H \succ c_G e_L$ (the other inequalities are always true by definition since $e_L < e_H$). The condition is:

$$\theta_G > q(e_L) - q(e_H) \tag{A.4}$$

In the second ranking, we need to give condition for which $c_G e_H \succ c_B e_L$ (the other inequalities are always true by definition since $e_L < e_H$). The condition is:

$$\theta_G < q(e_L) - q(e_H) \tag{A.5}$$

The utility function for a general-category patient c_G^p choosing a doctor of caste $c = c_B, c_G$ is given by:

$$U\left(\Phi\left(c_G^p, c\right), q\right) = q(e) - 1_{\theta_B} \tag{A.6}$$

where $1_{\theta_B} = \theta_B > 0$ if the doctor is from a general-category background and zero otherwise.

If $q'(e) > 0$ (i.e., doctors with more years of experience provide better quality health service), the possible rankings for a general-category patient compatible with (A.6) are:

$$c_G e_H \succ c_G e_L \succ c_B e_H \succ c_B e_L$$
$$c_G e_H \succ c_B e_H \succ c_G e_L \succ c_B e_L$$

In the first ranking, we need to give the condition for which $c_G e_L \succ c_B e_H$ (the other inequalities are always true by definition since $e_H > e_L$). The condition is:

$$\theta_B > q(e_H) - q(e_L) \tag{A.7}$$

In the second ranking, we need to give condition for which $c_B e_H \succ c_G e_L$ (the other inequalities are always true by definition since $e_H > e_L$). The condition is:

$$\theta_B < q(e_H) - q(e_L) \tag{A.8}$$

If $q'(e) < 0$ (i.e., doctors with fewer years of experience provide better quality health service), the possible rankings for a general-category patient compatible with (A.6) are:

$$c_G e_L \succ c_G e_H \succ c_B e_L \succ c_B e_H$$
$$c_G e_L \succ c_B e_L \succ c_G e_H \succ c_B e_H$$

In the first ranking, we need to give the condition for which $c_G e_H \succ c_B e_L$ (the other inequalities are always true by definition since $e_L < e_H$). The condition is:

$$\theta_B > q(e_L) - q(e_H) \tag{A.9}$$

In the second ranking, we need to give condition for which $c_B e_L \succ c_G e_H$ (the other inequalities are always true by definition since $e_L < e_H$). The condition is:

$$\theta_B < q(e_L) - q(e_H) \tag{A.10}$$

To summarize, with taste-based discrimination, there are eight possible rankings in total. The four possible rankings for a backward-caste homophily patient that are compatible with (A.1) are:

$$\begin{gathered} c_B e_H \succ c_B e_L \succ c_G e_H \succ c_G e_L \\ c_B e_H \succ c_G e_H \succ c_B e_L \succ c_G e_L \\ c_B e_L \succ c_B e_H \succ c_G e_L \succ c_G e_H \\ c_B e_L \succ c_G e_L \succ c_B e_H \succ c_G e_H \end{gathered}$$

The four possible rankings for a general-category homophily patient that are compatible with (A.6) are:

$$\begin{gathered} c_G e_H \succ c_G e_L \succ c_B e_H \succ c_B e_L \\ c_G e_H \succ c_B e_H \succ c_G e_L \succ c_B e_L \\ c_G e_L \succ c_G e_H \succ c_B e_L \succ c_B e_H \\ c_G e_L \succ c_B e_L \succ c_G e_H \succ c_B e_H \end{gathered}$$

A.2. Heterophily taste-based discrimination

The utility function for a backward-caste patient c_B^p choosing a doctor of caste $c = c_B, c_G$ is given by:

$$U\left(\Phi\left(c_B^p, c\right), q\right) = q(e) - 1_{\theta_B} \tag{A.11}$$

where $1_{\theta_B} = \theta_B > 0$ if the doctor is from a backward-caste background and zero otherwise.

If $q'(e) > 0$, the possible rankings for a backward-caste patient compatible with (A.11) are:

$$\begin{gathered} c_G e_H \succ c_G e_L \succ c_B e_H \succ c_B e_L \\ c_G e_H \succ c_B e_H \succ c_G e_L \succ c_B e_L \end{gathered}$$

If $q'(e) < 0$, the possible rankings for a backward-caste patient compatible with (A.11) are:

$$c_G e_L \succ c_G e_H \succ c_B e_L \succ c_B e_H$$

$$c_G e_L \succ c_B e_L \succ c_G e_H \succ c_B e_H$$

The utility function for a general-category patient c_G^p choosing a doctor of caste $c = c_B, c_G$ is given by:

$$U\left(\Phi\left(c_G^p, c\right), q\right) = q(e) - 1_{\theta_G} \tag{A.12}$$

where $1_{\theta_G} = \theta_G > 0$ if the doctor is from a general-category background and zero otherwise.

If $q'(e) > 0$, the possible rankings for a general-category patient compatible with (A.12) are:

$$c_B e_H \succ c_B e_L \succ c_G e_H \succ c_G e_L$$
$$c_B e_H \succ c_G e_H \succ c_B e_L \succ c_G e_L$$

If $q'(e) < 0$, the possible rankings for a general-category patient compatible with (A.12) are:

$$c_B e_L \succ c_B e_H \succ c_G e_L \succ c_G e_H$$
$$c_B e_L \succ c_G e_L \succ c_B e_H \succ c_G e_H$$

B. Locations of field experiment

The areas covered in our study include Ratanpur, Lodhar, Kursauli, Maksudabad, Tikra, Singhpur, Hora, Paigupur, Pachor, Mandhana, Kukradev, Tikkanpurwa, Bairy, Mharajpur, Loharkheda, Pargahi, Guraha, Sandeela, Shadipur, Naurangabad, Baikunthpur, Sakshupurwa, Iswaringanj, Hradaypur, Parapratappur, Chandula, Pokharpurwa, Naramau, Karsaitpur, Madarpur, Indra Nagar, Kalyanpur Khud, Devi Shai Nagar, Sahab Nagar, Jai Prakash Nagar, Loharanbhatta, Fazalganj, Barasirohi, Mirjapur, and Maswanpur.

C. Sign-up sheet

In section C.1 below, we provide an example of a sign-up sheet that participants saw when they were registering for the upcoming health check sessions. In section C.2, we provide evidence that the majority of individuals residing in the localities of the field experiment can correctly associate the surnames with their caste groups.

C.1. An example

IMPORTANT: We only provide non-urgent medical examination. If you think you have an urgent medical issue that needs attention, please visit the nearby hospital at:

Your full name:

Your phone number:

Your gender:

Your date of birth:

The patient's full name:

The patient's gender:

The patient's date of birth:

The following four profiles of male/female doctors represent the doctors that may treat you for the upcoming medical appointment. Please nominate your preference rankings of doctors by writing 1, 2, 3, or 4 in the box next to each profile. Write 1 for the most preferred, 2 for the second most preferred, 3 for the third most preferred, and 4 for the least preferred. You have a higher chance of getting the more preferred doctor.

Doctor profile	Preference rank	Doctor profile	Preference rank
C. Pandey (MBBS 2005)		F. Sonkar (MBBS 2005)	
S. Kureel (MBBS 2013)		S. Bajpai (MBBS 2013)	

NOTE: In the case that none of the above profiles is available, we will ensure a doctor is provided.

We will inform you the date/time of the appointment within the next two days.

C.2. Recognizability of caste-sounding surnames

Although the caste-sounding surnames that we use on the sign-up sheets are common surnames in Uttar Pradesh, it is possible that the general population may not correctly associate them with corresponding caste groups. We conducted a phone survey on 320 individuals in Kanpur Nagar district in Uttar Pradesh to examine the likelihood of each surname being associated with the correct corresponding caste group. Table C1 reports the results. Overall, in all cases, more than 90% of the respondents correctly associated the surname with either the general-category or a backward caste group correctly.

Table C1: Distribution of caste groups being associated with each caste-sounding surname

Surnames	GC	OBC	SC	ST	Don't know	Correctly identified	Correct at the broad caste group
A. General category (GC)							
Bajpai	0.94	0.06	0.00	0.00	0.00	0.94	0.94
						(0.01)	(0.01)
Dixit	0.89	0.11	0.00	0.00	0.00	0.89	0.89
						(0.02)	(0.02)
Mishra	0.93	0.07	0.00	0.00	0.00	0.93	0.93
						(0.01)	(0.01)
Pandey	0.98	0.03	0.00	0.00	0.00	0.98	0.98
						(0.01)	(0.01)
B. Other backward classes (OBC)							
Katiyar	0.09	0.80	0.11	0.00	0.00	0.80	0.91
						(0.02)	(0.02)
Pal	0.00	0.91	0.09	0.00	0.00	0.91	1.00
						(0.02)	(0.00)
Vishwakarma	0.09	0.80	0.11	0.00	0.00	0.80	0.91
						(0.02)	(0.02)
Yadav	0.07	0.93	0.00	0.00	0.00	0.93	0.93
						(0.01)	(0.01)
C. Scheduled caste (SC)							
Kanaujiya	0.00	0.09	0.91	0.00	0.00	0.91	1.00
						(0.02)	(0.00)
Kureel	0.00	0.02	0.88	0.10	0.00	0.88	1.00
						(0.02)	(0.00)
Sonkar	0.00	0.01	0.91	0.08	0.00	0.91	1.00
						(0.02)	(0.00)
Valmiki	0.00	0.02	0.90	0.08	0.00	0.90	1.00
						(0.02)	(0.00)

Notes: Sample size is 320. Correctly identified means that the respondent associates the surname with the correct caste group. Correct at the broad caste group means that the respondent classifies the surname consistent with the general category or a backward caste group, but the backward caste group may not necessarily match with the exact group (OBC, SC or ST). Standard errors reported in parentheses.

D. Lab-in-the-field experiment

In late October 2017, we invited a random subset of the initial field-experimental participants in 30 randomly selected localities to participate in an incentivized lab-in-the-field experiment. The lab-in-the-field experiment includes four dictator money-giving games. The four games correspond to four different groups of partners: general-category, backward-caste, above the poverty line (APL), and below the poverty line (BPL).

In total, 482 subjects participated in the lab-in-the-field experiment. In each dictator game, each participant, who had an endowment of 100 Rupees, decided how much to keep from this endowment (a number between 0 and 100 (inclusive) Rupees), given that what was not kept went to a randomly drawn "partner" from one of the four social groups of participants. If the participant is from a backward-caste group, we reminded the participant that the backward-caste partner was from the same caste group. The anonymous partner was randomly drawn from our field experiment and the allocation was later given to them.

Steps were taken to minimize experimenter demand effect and social desirability bias. At the beginning of each game, an envelope with the group identity of an anonymous partner written on the envelope was drawn from a set of four envelopes. The participant was then given the group identity of the anonymous partner and the envelope with 100 Rupees (10 x Rs10 notes). The experimenter then instructed the participants to go to quiet corner to allocate whatever amount they wished for themselves and put the remaining amount in the envelope they wanted to give to the anonymous partner. They were also informed that, once they finished the task, they would drop the envelope in a bag full of similar-looking envelopes that the experimenter placed in a different corner. In the inside of each envelope, each participant's unique ID is written, so the amount can be linked to their responses in the field experiment. By letting them allocate the money in a quiet corner and drop each of the envelopes in a bag full of similar-looking envelopes away from the scrutiny of the experimenter, we minimize experimenter demand effect and social desirability bias.

Table D1: Descriptive statistics

	Field experiment (n = 3,128)		Lab-in-the-field experiment (n = 482)	
	Mean	Std. Dev.	Mean	Std. Dev.
Male	0.51	0.50	0.53	0.50
Age	37.8	14.3	38.5	15.0
General category	0.27	0.44	0.26	0.44
Hindu	0.80	0.40	0.76	0.43
College educated	0.11	0.32	0.09	0.28
Below poverty line	0.34	0.47	0.34	0.47
Urban resident	0.34	0.48	0.33	0.47

Table D1 shows that the characteristics of these 482 participants in the lab-in-the-field experiment are similar to the 3,128 participants in the initial field experiment. The main findings in Table 2 are also replicated when we restrict the sample to these 482 participants (Table D2).

Thus, these 482 participants behave, on average, similar to the 3,128 participants in the field experiment.

Table D2: Distribution of preference rankings consistent with different theories by caste of patient

	Backward-caste	General-category	All
A. Taste-based or statistical discrimination			
r_1: $c_G e_H \succ c_G e_L \succ c_B e_H \succ c_B e_L$	1.68	23.39	7.26
r_2: $c_G e_H \succ c_B e_H \succ c_G e_L \succ c_B e_L$	21.23	20.97	21.16
r_3: $c_B e_H \succ c_B e_L \succ c_G e_H \succ c_G e_L$	7.54	0.00	5.60
r_4: $c_B e_H \succ c_G e_H \succ c_B e_L \succ c_G e_L$	24.58	14.52	21.99
r_5: $c_G e_L \succ c_G e_H \succ c_B e_L \succ c_B e_H$	0.56	2.42	1.04
r_6: $c_G e_L \succ c_B e_L \succ c_G e_H \succ c_B e_H$	1.12	0.81	1.04
r_7: $c_B e_L \succ c_B e_H \succ c_G e_L \succ c_G e_H$	2.23	0.00	1.66
r_8: $c_B e_L \succ c_G e_L \succ c_B e_H \succ c_G e_H$	0.56	0.00	0.41
B. Statistical discrimination			
r_9: $c_G e_H \succ c_B e_H \succ c_B e_L \succ c_G e_L$	10.89	9.68	10.58
r_{10}: $c_B e_H \succ c_G e_H \succ c_G e_L \succ c_B e_L$	15.08	13.71	14.73
r_{11}: $c_G e_L \succ c_B e_L \succ c_B e_H \succ c_G e_H$	1.12	0.00	0.83
r_{12}: $c_B e_L \succ c_G e_L \succ c_G e_H \succ c_B e_H$	1.12	0.81	1.04
r_{13}: $c_G e_H \succ c_B e_L \succ c_B e_H \succ c_G e_L$	0.56	0.81	0.62
r_{14}: $c_G e_H \succ c_B e_L \succ c_G e_L \succ c_B e_H$	1.12	0.00	0.83
r_{15}: $c_G e_H \succ c_G e_L \succ c_B e_L \succ c_B e_H$	2.23	5.65	3.11
r_{16}: $c_B e_L \succ c_B e_H \succ c_G e_H \succ c_G e_L$	0.56	0.00	0.41
r_{17}: $c_B e_L \succ c_G e_H \succ c_G e_L \succ c_B e_H$	0.84	0.00	0.62
r_{18}: $c_B e_L \succ c_G e_H \succ c_B e_H \succ c_G e_L$	0.84	0.00	0.62
r_{19}: $c_G e_L \succ c_B e_H \succ c_B e_L \succ c_G e_H$	0.56	0.00	0.41
r_{20}: $c_G e_L \succ c_B e_H \succ c_G e_H \succ c_B e_L$	0.28	0.81	0.41
r_{21}: $c_G e_L \succ c_G e_H \succ c_B e_H \succ c_B e_L$	2.23	5.65	3.11
r_{22}: $c_B e_H \succ c_B e_L \succ c_G e_L \succ c_G e_H$	2.23	0.81	1.87
r_{23}: $c_B e_H \succ c_G e_L \succ c_G e_H \succ c_B e_L$	0.56	0.00	0.41
r_{24}: $c_B e_H \succ c_G e_L \succ c_B e_L \succ c_G e_H$	0.28	0.00	0.21
C. Summary			
Homophily taste-based or statistical	34.92	47.58	38.17
	(2.52)	(4.50)	(2.22)
Heterophily taste-based or statistical	24.58	14.52	21.99
	(2.28)	(3.18)	(1.89)
Uniquely statistical discrimination	40.50	37.90	39.83
	(2.60)	(4.37)	(2.23)

Notes: The sample size is 482. Taste-based discrimination is indistinguishable from statistical discrimination. High experience may signal better or worse quality of healthcare service when classifying with which theory a particular ranking is consistent. Standard errors are reported in parentheses.

Table D3 reports the mean amount given to different groups by caste of patient. In general, both general-category and backward-caste patients tend to give more to backward-caste individuals than general-category individuals. However, general-category patients tend to give slightly more to general-category individuals than backward-caste patients do, while backward-caste patients tend to give more to backward-caste individuals than general-category patients do.

Table D3: Mean amount given to different groups by caste of patient

	Backward-caste	General-category	All
Amount given to backward-caste individuals	47.63	44.19	46.74
	[22.53]	[23.24]	[22.74]
Amount given to general-category individuals	30.61	32.50	31.10
	[19.75]	[17.33]	[19.16]
Sample size	358	124	482

Notes: Standard deviations reported in brackets.

E. Robustness

E.1. Potential randomization due to weakened incentivization

Table E1.1: Distribution of preference rankings consistent with different theories

	Backward-caste	General-category	All
A. Taste-based or statistical discrimination			
$r_1: c_G e_H \succ c_G e_L \succ c_B e_H \succ c_B e_L$	3.23	28.81	10.16
$r_2: c_G e_H \succ c_B e_H \succ c_G e_L \succ c_B e_L$	23.91	20.86	23.08
$r_3: c_B e_H \succ c_B e_L \succ c_G e_H \succ c_G e_L$	9.14	1.05	6.95
$r_4: c_B e_H \succ c_G e_H \succ c_B e_L \succ c_G e_L$	29.98	17.70	26.66
$r_5: c_G e_L \succ c_G e_H \succ c_B e_L \succ c_B e_H$	1.67	4.50	2.44
$r_6: c_G e_L \succ c_B e_L \succ c_G e_H \succ c_B e_H$	0.78	1.05	0.85
$r_7: c_B e_L \succ c_B e_H \succ c_G e_L \succ c_G e_H$	2.06	0.00	1.50
$r_8: c_B e_L \succ c_G e_L \succ c_B e_H \succ c_G e_H$	1.50	0.60	1.26
B. Statistical discrimination			
$r_9: c_G e_H \succ c_B e_H \succ c_B e_L \succ c_G e_L$	7.80	5.25	7.11
$r_{10}: c_B e_H \succ c_G e_H \succ c_G e_L \succ c_B e_L$	10.89	8.40	10.22
$r_{11}: c_G e_L \succ c_B e_L \succ c_B e_H \succ c_G e_H$	0.39	0.15	0.33
$r_{12}: c_B e_L \succ c_G e_L \succ c_G e_H \succ c_B e_H$	0.50	0.45	0.49
$r_{13}: c_G e_H \succ c_B e_L \succ c_B e_H \succ c_G e_L$	0.22	0.15	0.20
$r_{14}: c_G e_H \succ c_B e_L \succ c_G e_L \succ c_B e_H$	0.75	0.45	0.67
$r_{15}: c_G e_H \succ c_G e_L \succ c_B e_L \succ c_B e_H$	1.84	5.55	2.84
$r_{16}: c_B e_L \succ c_B e_H \succ c_G e_H \succ c_G e_L$	0.70	0.00	0.51
$r_{17}: c_B e_L \succ c_G e_H \succ c_G e_L \succ c_B e_H$	0.25	0.08	0.20
$r_{18}: c_B e_L \succ c_G e_H \succ c_B e_H \succ c_G e_L$	0.22	0.15	0.20
$r_{19}: c_G e_L \succ c_B e_H \succ c_B e_L \succ c_G e_H$	0.14	0.23	0.16
$r_{20}: c_G e_L \succ c_B e_H \succ c_G e_H \succ c_B e_L$	0.36	0.45	0.39
$r_{21}: c_G e_L \succ c_G e_H \succ c_B e_H \succ c_B e_L$	1.09	3.38	1.71
$r_{22}: c_B e_H \succ c_B e_L \succ c_G e_L \succ c_G e_H$	2.03	0.15	1.52
$r_{23}: c_B e_H \succ c_G e_L \succ c_G e_H \succ c_B e_L$	0.22	0.38	0.26
$r_{24}: c_B e_H \succ c_G e_L \succ c_B e_L \succ c_G e_H$	0.31	0.23	0.28
C. Summary			
Homophily taste-based or statistical	42.69	55.21	46.08
	(0.83)	(1.36)	(0.71)
Heterophily taste-based or statistical	29.59	19.35	26.82
	(0.76)	(1.08)	(0.63)
Uniquely statistical discrimination	27.72	25.43	27.10
	(0.75)	(1.19)	(0.63)

Notes: Preference rankings consistent with both theories are weighted twice as much as those uniquely statistical when weights are applied. Taste-based discrimination is indistinguishable from statistical discrimination. High experience may signal better or worse quality of healthcare service when classifying with which theory a particular ranking is consistent. Standard errors are reported in parentheses.

Table E1.2: Preference rankings that are consistent with general-category stereotype

True-quality consistent preference ranking	Preference rankings consistent with general-category stereotype	Maximum share with accurate beliefs	Maximum share with stereotypical beliefs
r_1	r_1	10.16	10.16
r_2	r_2, r_1	23.08	33.24
r_3	$r_3, r_2, r_4, r_9, r_{10}$	6.95	74.02
r_4	r_4, r_9, r_{10}	26.66	43.99
r_5	r_5	2.44	2.44
r_6	r_6, r_5	0.85	3.29
r_7	$r_7, r_6, r_8, r_{11}, r_{12}$	1.50	4.43
r_8	r_8, r_{11}, r_{12}	1.26	2.08
r_9	r_9, r_1, r_2	7.11	40.35
r_{10}	r_{10}, r_1, r_2	10.22	43.46
r_{11}	r_{11}, r_5, r_6	0.33	3.62
r_{12}	r_{12}, r_6	0.49	1.34
r_{13}	r_{13}, r_{14}, r_{15}	0.20	3.71
r_{14}	r_{14}, r_{15}	0.67	3.51
r_{15}	r_{15}	2.84	2.84
r_{16}	$r_{16}, r_{13}, r_{14}, r_{17}, r_{18}$	0.51	1.78
r_{17}	r_{17}, r_{14}	0.20	0.87
r_{18}	$r_{18}, r_{13}, r_{14}, r_{17}$	0.20	1.27
r_{19}	r_{19}, r_{20}, r_{21}	0.16	2.26
r_{20}	r_{20}, r_{21}	0.39	2.10
r_{21}	r_{21}	1.71	1.71
r_{22}	$r_{22}, r_{19}, r_{20}, r_{23}, r_{24}$	1.52	2.61
r_{23}	r_{23}, r_{20}	0.26	0.65
r_{24}	$r_{24}, r_{10}, r_{19}, r_{20}, r_{23}$	0.28	11.31

Notes: Preference rankings consistent with both theories are weighted twice as much as those uniquely statistical when weights are applied. Column 1 lists the 24 possible preference rankings when each of them is consistent with the assumed true quality ranking of doctors. Column 2 lists the preference rankings that are consistent with a general-category stereotype for a given preference ranking assumed to be consistent with the true quality ranking. Column 2 includes the preference rankings that are consistent with the assumed true quality rankings because general-category stereotype may not be sufficiently large to swap the rank order of doctors. Column 3 reports the maximum share of patients with preference rankings consistent with accurate beliefs based on the assumed true quality ranking. Column 4 reports the maximum share of patients with preference rankings that are consistent with general-category stereotype. See Table 2 and Table D1.1 for details of rankings.

Table E1.3: Distribution of preference rankings consistent with different theories by caste of patient and reclassification of caste preferences

	Backward-caste	General-category	All
A. Before reclassification of caste preferences			
Homophily taste-based or statistical	43.78	58.71	47.67
	(2.63)	(4.44)	(2.28)
Heterophily taste-based or statistical	30.82	17.91	27.46
	(2.44)	(3.46)	(2.04)
Uniquely statistical discrimination	25.39	23.38	24.87
	(2.30)	(3.82)	(1.97)
B. After reclassification of strong animosity			
Homophily taste-based or statistical	28.02	16.92	25.13
	(2.38)	(3.38)	(1.98)
Heterophily taste-based or statistical	2.45	14.93	5.7
	(0.82)	(3.21)	(1.06)
Uniquely statistical discrimination	69.53	68.16	69.17
	(2.44)	(4.20)	(2.11)
C. After reclassification of weak animosity			
Homophily taste-based or statistical	37.48	27.86	34.97
	(2.56)	(4.04)	(2.17)
Heterophily taste-based or statistical	8.76	14.93	10.36
	(1.50)	(3.21)	(1.39)
Uniquely statistical discrimination	53.77	57.21	54.66
	(2.64)	(4.46)	(2.27)

Notes: The sample size is 482 patients who participated in both the field experiment and dictator games. Preference rankings consistent with both theories are weighted twice as much as those uniquely statistical when weights are applied. Panel A show shows that the data based on the sub-sample of 482 patients are similar to the full sample. In panel B, a patient is defined to have animosity against the discriminated caste group inferred from the preference ranking of doctors in the field experiment when they give less to the discriminated caste group in the dictator games. In panel C, a patient is defined to have animosity against the discriminated caste group inferred from the preference ranking of doctors in the field experiment when they give equal or less to the discriminated caste group in the dictator games.

E.2. Preferences for attributes other than caste and experience

In the taste-based discrimination model, we assume that a patient has only preferences for the social proximity of a doctor in terms of caste and the quality of healthcare, but nothing else. In reality, patients may have preferences for other attributes of a doctor. A good example is the gender of a doctor. Now imagine that, in our experiment, patients are shown only two of the three attributes in their consumption bundle, x. Given that they also have preferences for the gender of a doctor, it is plausible to think that they may use the doctor's experience to infer the doctor's gender because female doctors are increasingly more represented in the medical profession in India

(Bhadra, 2011). This possibility is similar to the Heckman's (1998) critique about using experiments to detect taste-based discrimination. Note that, to put aside the possibility of caste-based statistical discrimination, we still assume that these patients do not use the doctor's caste to help make inferences about gender. In this case, differences in the predicted gender of a doctor across patients will influence how they rank the four doctors. The possible rankings under taste-based discrimination for the case when patients try to infer the unobserved attribute from an observed attribute in $x \in \mathbb{R}_{+}^{3}$ may thus include more than the eight possible rankings we have highlighted for the case when $x \in \mathbb{R}_{+}^{2}$.

Table E2.1: Robustness of results to genders of doctors

	Female Doctors	Male Doctors
Homophily taste-based or statistical discrimination	35.68	36.82
	(1.21)	(1.22)
Heterophily taste-based or statistical discrimination	21.20	21.00
	(1.03)	(1.03)
Uniquely statistical discrimination	43.11	42.18
	(1.25)	(1.25)

Notes: The sample is based on responses of all patients. Taste-based discrimination is indistinguishable from statistical discrimination. High experience may signal better or worse quality of healthcare service when classifying with which theory a particular ranking is consistent. Standard errors are reported in parentheses.

If the share of (uniquely) statistical discriminators does not vary considerably when additional correlated attributes are taken into consideration, then our method of bounding taste-based discrimination is unlikely to yield biased estimates due to taste-based discriminators using observable attributes of a doctor to infer other unobservable attributes of the doctor for which they have preferences.

We thus examine if the findings differ between the case when all the four doctors are females and the case when all the four doctors are males. The results are reported in Table E2.1. The share of discriminators who have preference rankings consistent with only statistical discrimination is similar regardless of whether they are presented with female or male doctors. Therefore, our results are robust to the possibility that taste-based discriminators use a doctor's experience to make inferences about other unobserved attributes of the doctor.