

Does Identity Affect Labor Supply?

Suanna Oh

Impressum:

CESifo Working Papers

ISSN 2364-1428 (electronic version)

Publisher and distributor: Munich Society for the Promotion of Economic Research - CESifo GmbH

The international platform of Ludwigs-Maximilians University's Center for Economic Studies and the ifo Institute

Poschingerstr. 5, 81679 Munich, Germany

Telephone +49 (0)89 2180-2740, Telefax +49 (0)89 2180-17845, email office@cesifo.de

Editor: Clemens Fuest

<https://www.cesifo.org/en/wp>

An electronic version of the paper may be downloaded

- from the SSRN website: www.SSRN.com
- from the RePEc website: www.RePEc.org
- from the CESifo website: <https://www.cesifo.org/en/wp>

Does Identity Affect Labor Supply?

Abstract

Does identity influence economic behavior in the labor market? I investigate this question in rural India, focusing on the effect of caste identity on job-specific labor supply. In a field experiment, laborers choose whether to take up various job offers, which differ in associations with specific castes. Workers are less willing to accept offers that are linked to castes other than their own, especially when those castes rank lower in the social hierarchy. Workers forego large payments to avoid job offers that conflict with their caste identity, regardless of whether these decisions are made in private.

JEL-Codes: D910, J240, O100.

Keywords: identity, labor supply, caste, occupational choice.

Suanna Oh
Paris School of Economics
48 boulevard Jourdan
France – 75014 Paris
suanna.oh@psemail.eu

December 14, 2021

I gratefully acknowledge guidance and support from Eric Verhoogen, Supreet Kaur, Jonas Hjort, Jack Willis, and Suresh Naidu. I thank Michael Best, Leonardo Bursztyn, Ernesto Dal Bó, Stefano DellaVigna, Alejandro Favela Nava, Evan Friedman, Florian Grosset, Simon Haenni, Seema Jayachandran, Yoon Jo, Matt Lowe, Karen Macours, Edward Miguel, Cristian Pop-Eleches, Martin Rotemberg, Yogita Shamdasani, Rodrigo Soares, Mahesh Srinivasan, Christopher Walters, David Yanagizawa-Drott, Yue Yu, and various seminar participants for helpful comments. Sneha Subramanian, Manvi Govil, Arnesh Chowdhury, and Santosh Kumar Sahoo provided excellent research assistance. I am grateful to the Center for Development Economics and Policy, the Program for Economic Research at Columbia University, and the National Science Foundation (SES-1919139) for generous funding, and to J-PAL South Asia and IFMR for operational support. This research was reviewed by the Columbia University IRB and the IFMR Human Subjects Committee. The experiment was registered on the AEA RCT registry, ref. AEARCTR-0003493. All errors are my own. I declare that I have no relevant or material financial interests that relate to the research described in this paper.

Identity refers to individuals’ concepts of “who they are.” These concepts are typically oriented around social categories (e.g. woman, parent, worker), which people use to process and organize information about the world. The social categories may be associated with behavioral prescriptions, dictating how the members of the social category ought to think and act. Violating such prescriptions could lead to emotional or psychological costs (Akerlof and Kranton 2000). This way, concerns about protecting identity may compel someone to follow her internal rules of behavior, even when no one can observe her or when it is economically costly to do so. A nascent literature in economics, as well as long-standing ones in other social sciences, investigates how identity concerns influence individual behavior as well as market outcomes.¹

Theories of identity postulate that in the labor market, people may avoid otherwise desirable job opportunities which conflict with their identity.² If people failed to pursue certain careers despite their potential aptitude due to concerns about identity, it would result in inefficient allocation of talent in the economy (Hsieh et al. 2019; Cassan, Keniston, and Kleineberg (2021)). Furthermore, identity-based occupational preferences could interact with other channels of misallocation, such as discrimination, with amplified consequences for occupational distribution and aggregate productivity.

However, there exists limited empirical evidence on the extent to which—and how—identity concerns affect how people choose jobs. A social group that dislikes some occupations tends to differ from other groups along many dimensions, such as training and outside options. Moreover, engaging in certain jobs could influence how others perceive the group, i.e. social image, in addition to any internal feelings. For these reasons, it is challenging to establish the effect of identity using observational or survey data alone. Isolating the impact using an experiment is also complicated, because it is difficult for researchers to randomly assign deeply-ingrained identities or radically alter existing perceptions about occupations.

I address these challenges by drawing on unique features of the Indian caste system and provide the first experimental test of how identity affects job-specific labor supply. Casual daily-wage laborers in rural Odisha, India, evaluate various real job offers,

¹For overviews of identity theories from social psychology, see Stryker ([1982] 2002), Burke and Stets (2009), and Stets and Serpe (2013). For reviews in economics, see Hoff and Stiglitz (2016) and Shayo (2020).

²For influential early theories on the links between self-concept, occupational images and preferences, see Super et al. (1963) and Gottfredson (1981). For more recent reviews, see Skorikov and Vondracek (2011) and Patton and McMahon (2014).

which differ in caste associations. I show that workers are less willing to take up offers that are linked to castes other than their own, especially when those caste rank lower in the social hierarchy. These effects are invariant to whether or not workers' decisions are publicized, suggesting that identity—rather than social image—is the main driver of the effects.

The experiment design is aimed at isolating the impact of identity violations on job take-up. An ideal experiment would randomly determine whether a job involves an identity-conflicting task or not while keeping constant all other features, such as the associated effort costs and amenities. I attempt to approximate such job offers in the context of a casual labor market in India. In this setting, workers tend to engage in bilateral wage contracts with employers that are short in duration, typically one day, to work on various unskilled tasks (Kaur 2019).

My approach is to find pairs of tasks so that the tasks within each pair are as similar as possible but differ in their caste associations. The historic Indian caste system provides links between castes and jobs, which often extend to simple manual tasks associated with the jobs. This allows me to identify pairs of tasks so that the tasks within each pair are as similar as possible but differ in their caste associations. This way, I can observe workers' take-up decisions for *identity tasks*, those predicted to evoke identity-violations, as well as for *paired control tasks* that are absent of caste associations, with the difference attributable to the effects of identity.

The main concern with this approach is that an identity task and its paired control task, while similar, could involve slightly different costs of effort. To address this, I construct the offers such that they involve working on a default manufacturing task as well as one extra task, which could be an identity or control task. This allows me to observe how workers react to offers that require spending just a little bit of time on identity or control tasks, while holding constant the total working time. Thus, the difference in take-up would reflect the difference in fixed costs of spending any time at all on the extra tasks, separate from effort costs. Social theories of identity suggest that engaging at all in identity-inconsistent jobs could constitute a violation of one's internal rule of behavior, evoking psychological costs. Hence, a large gap in take-up would indicate a significant difference in fixed costs between identity and control tasks, suggesting strong identity concerns.

One remaining concern is that the gap in take-up associated with the fixed costs may still capture some costs that are unrelated to caste identity concerns. For ex-

ample, workers might be inherently averse to trying out new tasks for the first time and they might lack experience with certain tasks. To address this, I focus on a difference-in-differences (DiD) comparison: the difference *across caste groups* in the differences in take-up rates between identity and control tasks. Under the assumption that the difference in inherent fixed costs of working on identity tasks compared to paired control tasks is similar across castes, this would yield a credible estimate for a lower bound on the impact of identity concerns on labor supply.

Furthermore, I use the fact that castes can be ranked according to social status to test that working on a job associated with a lower-status group leads to stronger identity concerns, as suggested by theories of identity. Precisely, I examine whether the estimated gap in take-up between identity and control tasks is indeed larger when identity tasks are associated with groups with relatively lower status.

To implement this design, I collect two surveys to measure locally prevalent perceptions about the caste system. The first survey documents how tasks are linked to castes. The second survey records how the castes that make up the experimental sample rank in the social hierarchy. Based on the survey results, I select three identity tasks (e.g. washing clothes, which is associated with the Dhoba caste) and the corresponding paired control tasks (e.g. washing farming tools, which has no association).

The main experiment involves 630 male casual laborers and elicits their willingness to take up one-day job offers.³ All offers involve working on a common default task of producing paper bags, which is not associated with any caste. The offers also entail working on one additional task in a private space. The offers are constructed to vary only in two dimensions—the type of extra task and the share of total time required to work on it. The job offers are the same in all other aspects, including the fixed daily wage, employer, worksite location, total working time of five hours, and other characteristics, so that preference over these attributes cannot influence job take-up decisions.⁴

To truthfully elicit workers' responses, I use a choice exercise based on the Becker-

³The experiment involves only male workers as it focuses on the effect of caste (rather than gender) identity effects. In addition, there are practical difficulties with hiring female workers who typically cannot work outside villages without accompanying male family members.

⁴The offers are designed such that working on extra tasks does not require any prior training or experience. Workers are also explicitly told that they are one-time offers and workers' decisions or outcomes will not influence their future job prospects.

DeGroot-Marschak (BDM) method. Each worker is presented with a set of all potential job offers and is asked to indicate whether he would accept or decline each one. After the worker indicate all decisions, one offer is randomly selected and his choice for this offer is implemented. In addition to having the procedure carefully explained in order to ensure good understanding, the worker is explicitly encouraged to consider each offer separately, treating it as a single, take-it-or-leave-it offer, and give a simple honest answer.

The results indicate that workers' willingness to take up job offers decreases significantly when they are predicted to involve conflicts of identity. When workers are offered tasks that do not correspond to their castes, take-up rate is significantly lower for the identity tasks compared to control tasks. In contrast, when workers' castes are directly associated with identity tasks, the take-up rates are similar for identity and control tasks. When the castes associated with identity tasks rank higher than the workers' own, the take-up rate for identity tasks is still 23 percentage points (pp) lower compared to control tasks. This gap increases by an additional 28 pp when castes associated with identity tasks rank lower. This latter effect is especially strong for those who are caste-sensitive, i.e. those who express strong support for observing caste norms in a follow-up survey.

Part of these effects, however, might be driven by concerns about social image, rather than identity. To distinguish the additional effect of social image, I randomize whether or not worker decisions are publicized; and I find similar effects across the privacy treatments. This suggests that many workers are intrinsically motivated to behave in ways that are deemed appropriate for their castes. Because workers are already strongly motivated by identity, concerns for social image—even if present—may have little additional effect on take-up decisions, an explanation that is supported by the follow-up survey answers.

The experiment is designed to rule out many factors that typically confound the relationship between identity concerns and labor supply. In addition, any alternative explanation would need to address why 1) take-up rates appear to drop as soon as workers spend any time at all on extra tasks; 2) such falls are larger when tasks are associated with castes different from the workers' own, even compared to other tasks that involve similar skills; and 3) such decreases are larger when the associated castes have relatively lower social status. Workers' intrinsic desire to behave consistently with their caste identity can explain all of these findings.

I run a supplementary experiment to directly quantify the wage workers are willing to forego in order to avoid engaging in tasks associated with other castes. A new set of 106 workers are hired for a one-day job of producing paper bags, the default task. They are then unexpectedly given a chance to switch to a different task for part of the remaining working time. As in the main experiment, each worker is asked to evaluate many switching offers, which involve similar variations in the type of extra task and the time required to work on it. A key difference is that the switching offers might provide a bonus payment (varying from Rs. 30 to Rs. 3000) on top of the default daily wage of Rs. 300. The largest bonus amount is nearly a whole month’s earnings in the agricultural lean season during which the experiment takes place.

I find that 43% of workers are willing to forego at least ten times their daily wage in order to avoid spending as little as ten minutes on tasks associated with other castes. This is 28 pp greater than the analogous rate for control tasks and invariant to whether workers’ decisions are publicized. Notably, among those who agree to working on identity tasks, the large majority is willing to accept a bonus that is one tenth of daily wage or less. These polarized reactions even among workers of the same caste suggest that working on caste-inconsistent tasks is not inherently difficult or unpleasant, but for some, constitutes an unthinkable violation of identity.

These findings highlight just how important identity can be in the context of the labor market. Jobs that offer decent wages and working conditions are not valuable if they hinder workers from upholding their identity. Not just in India, but in many countries, some social groups remain under-represented in certain occupations, often despite policy efforts and changing economic conditions.⁵ It is possible that identity is an important driver of this phenomenon, which points to the need for more empirical evidence on identity-based frictions in the labor market.

This study is part of a growing empirical literature which examines the impact of identity on economic decision-making (Shayo 2020). One main theme in the literature is the role of identity in explaining in-group conformity, i.e. the tendency for individuals to follow the prototypical behaviors of their social groups.⁶ In this vein, lab

⁵Cassan, Keniston, and Kleineberg (2021) show that in India, 17% of male workers are engaged in their traditional occupations, while the rate would be 9% if one were to randomly allocate workers to existing occupations. In many countries, women are under-represented in male-stereotypical occupations (Cejka and Eagly 1999).

⁶Another theme in the literature concerns in-group bias—why some social groups engage in adversarial relationships or identity-based discrimination (Charness and Chen 2020; Hjort 2014; Lowe 2021).

experimental studies show that priming subjects about their specific identities make them behave more like their groups in terms of patience (Benjamin, Choi, and Strickland 2010), dishonesty (Cohn, Fehr, and Maréchal 2014), and public contributions (Benjamin, Choi, and Fisher 2016). Hoff and Pandey (2014) find that priming caste identities lowers cognitive performance among low-caste students in India. Other empirical studies show that identity considerations affect consumption behavior (Forman 2008; Atkin, Colson-Sihra, and Shayo 2021), selfish behavior (Falk 2021), and political expression (Bursztyn et al. 2020b). My study adds to this literature by establishing identity effects on labor supply.

By highlighting the impact of identity concerns on job-specific labor supply, my paper also contributes to the extensive literature on occupational choice. Psychologists have long emphasized the importance of the interactions between self-concept and occupational images for shaping job preferences (Super et al. 1963); Gottfredson (1981). The canonical economic models of occupational choice typically omits the discussion of these factors (Roy 1951; Topel and Ward 1992; Keane and Wolpin 1997; Neal 1999). In gender economics, however, there are active investigations of norms and image concerns as the potential drivers of gender division of labor as well as gender segregation of occupations. Gender norms prescribing whether it is appropriate for women to work outside of household appear to be key determinants of female labor force participation (Alesina, Giuliano, and Nunn 2013; Bursztyn, González, and Yanagizawa-Drott 2020a; Jayachandran 2021). Similarly, social prescriptions on which jobs are suitable for men or women may determine individuals' occupational choice (West and Zimmerman 1987; Cejka and Eagly 1999). One method of testing for this effect is to experimentally vary occupational perceptions (Delfino 2021; Del Carpio and Guadalupe 2021). I use a different approach, which directly tests theories of identity by exploiting the historical associations between real tasks and castes and the hierarchical structure of the caste system.

Finally, this work extends the literature on the misallocation of talent in the economy by establishing conflicts of identity as a potential source. My findings suggest that some people may fail to pursue certain careers despite their potential aptitude or existing skills due to identity concerns. Existing studies examine the impacts on aggregate productivity of different barriers that hinder talented individuals from pursuing their comparative advantage. These barriers include discrimination (Hsieh et al. 2019), exposure to innovation (Bell et al. 2019), access to capital (Goraya

2019), and constraints on working hours (Erosa et al. 2017). Notably, Hsieh et al. (2019) develop a general equilibrium Roy model of occupational choice to study the effect of the convergence in the occupational distribution of US social groups between 1960 and 2010. Their model can account for the effects of discrimination, barriers to forming human capital, as well as group-specific occupational preference. Cassan, Keniston, and Kleineberg (2021) build a similar framework to quantify the effect of misallocation resulting from workers choosing occupations based on caste identity. My paper not only provides empirical motivations for such investigations, but also underscores specific ways in which identity-based occupational preference operates, which may help to refine these frameworks.

The rest of the paper is organized as follows. Section 1 presents some key ideas from theories of identity and a simple theoretical framework. Section 2 discusses the surveys which collect information on castes and tasks used in the experiment. Section 3 describes the experimental design and the empirical specification. Section 4 discusses the results on identity conflict and job take-up. Section 5 presents the supplementary experiment design and findings, and Section 6 concludes.

1 Conceptualizing identity

1.1 Theories of identity and social image

Social psychologists posit that sense of identity is a powerful motivator of human behavior. Stets and Serpe (2013) relate identities to roles individuals occupy in the social structure, groups they identify with and belong to, and unique ways in which they see themselves. These identities often overlap and are always embedded in social and cultural contexts.⁷ People are intrinsically motivated to uphold their identity. Their behaviors are governed by “a strong need to maintain conformity between actions or even feelings” and their identities (Bénabou and Tirole (2006)). Failing to achieve congruence can evoke negative emotions, such as “anxiety and discomfort in oneself” (Akerlof and Kranton 2000).⁸ Hence, identity can curtail behavior even in private situations.

⁷This characterization draws from both identity theory and social identity theory, which to some extent have developed separately in the literature. Stryker and Burke (2000), Hogg, Terry, and White (1995), and Owens, Robinson, and Smith-Lovin (2010) discuss the similarities and differences between the strands.

⁸Stryker (2004) reviews the literature on the links between affect and identity theory.

The concepts of identity and social image are closely interlinked: people care about their own conception of who they are (alternatively referred to as self-image, self-identity, or intrinsic motivations) as well as other people’s perception of them (reputation, social-identity, or social pressure). Maintaining a positive social image could be valuable due to affect, e.g., social esteem or shame, or for instrumental reasons, e.g. social punishments (Bénabou and Tirole (2006). Bursztyn and Jensen (2017) describe the conditions under which one could empirically isolate the effect of social image concerns.

Two ideas from theories of identity are particularly important to my experimental design, which aims to test for identity effects in the labor market.⁹ First, the concept of violation is frequently featured in models of identity. In Akerlof and Kranton (2000), breaching a behavioral prescription results in a loss of utility. In Bénabou and Tirole (2011), an individual infers her own values or types from her past actions, and hence are motivated to take actions consistent with her identity as signals to her future self. Notably, if she engages in calculating and weighing the costs and benefits of violating her identity, the memory of this contemplation can later serve as a negative signal about her type. This could cause her to avoid merely thinking about transgressions, making them “priceless,” i.e. one vows to never do them regardless of any pecuniary benefits. These models suggest that those facing identity concerns may seek to avoid working in identity-inconsistent jobs even for a short time, as it still constitutes a breach. Exploiting time variations on different work tasks could help test whether certain jobs provoke concerns about identity.

Second, those belonging to a social category may be averse to adopting the characteristics and practices of others, particularly if the other categories have lower social status (Tajfel and Turner 1979). Shayo (2020) models this as a utility function increasing in the status of one’s group and decreasing in the distance between the prototypical attributes of the group and own attributes.¹⁰ This implies that workers

⁹There are extensive discussions on the complex ways in which image concerns affect behavior, such as those relating to multiple identities or identity change (Burke and Stets 2009). Carvalho and Pradelski 2021 builds a model of occupational choice that incorporates these considerations. Here I only focus on the two most relevant ideas for the experimental design.

¹⁰Consistent with the model, Atkin, Colson-Sihra, and Shayo (2021) show that when the status of a religious group increases, more households adopt food consumption patterns that are characteristic of the group. For a review on status and in-group bias, see Bettencourt et al. (2001). This way, the literatures on identity and social image are also tied to those on status and social norms. Bernheim (1994), Akerlof (1980), and Jones (1984) describe models in which desire for status, reputation, or conformity leads to the development of social norms.

may avoid jobs associated with other social groups, especially those jobs linked to lower-status groups. Hence, one approach of testing for identity effects is to use tasks associated with specific social groups which vary in social status.

1.2 Worker’s job take-up decision

I present a conceptual framework for a worker’s decision problem of job take-up to motivate the experimental design.¹¹ A worker, indexed by i and belonging to social category c_i , considers one-day job offers that involve working on two tasks—a default task and an extra task.¹² The worker expects to spend t_0 working on the default task and t_k working on extra task k . Thus, he will work for a total amount of time $T = t_0 + t_k$. All offers involve the same wage and the same total amount of time, so I use M_i and L_i to indicate worker i ’s utility for money and leisure, respectively.

Suppose I observe that the worker refuses an offer when the extra task is associated with a social category other than c_i , but accepts an offer whose extra task is associated with c_i . This is not enough to establish that the worker has identity concerns, because the identity-conflicting task may be more difficult or boring to work on. I take this to mean that the time-dependent costs of working on the extra tasks may be different. To isolate the effects of identity, I focus on measuring the fixed, or time-independent, costs of working on tasks. Theories of identity suggest that spending even a tiny bit of time on an identity-inconsistent task could still constitute an identity violation and lead to feelings of anxiety or shame. That is, working on an identity-conflicting task may involve a large fixed cost, which I measure by observing worker reactions as the time the worker has to spend on the extra task approaches zero.¹³

Formally, I assume that the worker’s utility from a job offer involving extra task k is described by:

$$U_{ik}(c_i, t_k) = \underbrace{M_i}_{\text{Money}} + \underbrace{L_i}_{\text{Leisure}} - \sum_{g \in \{0, k\}} \underbrace{[V_{ig}(c_i, t_g) + \mathbb{1}[t_g > 0] \cdot F_{ig}(c_i)]}_{\text{Variable Cost}}. \quad (1)$$

¹¹Carvalho and Pradelski 2021 presents a more general model of identity and occupational choice, which shows how identity-specific norms evolve as a function of representation.

¹²This decision here does not involve consideration of multiple identities, as the offers are expected to differ only in one identity dimension, i.e. caste. Specifically, the background surveys are used to verify that tasks do not involve conflicts of gender identity. I also omit the cost of working on different numbers of tasks, since all job offers in the experiment involve working on two tasks.

¹³Identity conflicts can of course increase the time-dependent costs of working on such a task as well. Therefore this approach does not measure the total effect of identity concerns on labor supply, but provides a convincing way of obtaining a lower bound.

The utility costs of working on any extra task k (and similarly for default task 0) are of two types. $V_{ik}(c_i, t_k)$ refers to the variable cost that varies with the time spent on the task, and $F_{ik}(c_i)$ indicates the fixed cost which gets incurred if the worker spends any amount of time on the task, i.e. if $t_k > 0$.¹⁴ Assume that the variable cost equals zero when the worker does not spend any time on the task (when $t_k = 0$) and is continuous in the time spent on the task (formally stated as Assumption 2 in Appendix Section B.1).¹⁵ Then by observing worker take-up decisions as t_k approaches 0, I estimate the impact on labor supply that is driven by the fixed cost of working on the extra task k .

However, the fixed cost may still not be entirely driven by identity concerns, i.e. some tasks may involve high fixed costs, independent of identity concerns. For example, the worker may be averse to trying out tasks for the first time. Accordingly, the fixed cost of working on task k is written as:

$$F_{ik}(c_i) = f_k + \beta_{ik} \cdot I_k(c_i), \quad (2)$$

where f_k is the “inherent” fixed cost that does not depend on c_i and $\beta_{ik} \cdot I_k(c_i)$ is the cost associated with identity violations.¹⁶ The indicator $I_k(c_i)$ takes the value of 1 when task k is associated with a social group other than c_i , and so β_{ik} gives the cost of the identity violation.¹⁷

The question is how to establish that β_{ik} is positive. An ideal experiment would randomly assign identity violations, $I_k(c_i)$. For instance, one could manipulate whether a task involves identity conflict, or create identity concerns within some people and not others. However, one would have to do this without changing beliefs about any other kind of costs associated with the job, which may be difficult to achieve. Finding a setting where required assumptions are naturally satisfied could be also challenging.

The approach taken in this paper is to use control tasks and multiple social groups in order to make a difference-in-differences comparison and estimate the share of

¹⁴Gilboa, Minardi, and Wang (2021) axiomatize a utility representation which also introduces discontinuities near zero for value-attached consumption.

¹⁵One functional form that satisfies this assumption is a linear function of t_k . This functional form turns out to be a good approximation in the case of the experimental tasks. Alternate functional forms are tested in robustness checks.

¹⁶Similarly, the variable cost could involve an inherent component, e.g. some tasks are naturally more difficult to perform. Assuming a linear function of time, $V_{ik}(c_i, t_k) = [v_k + \alpha_{ik} \cdot I_k(c_i)] \cdot t_k$.

¹⁷If the work has multiple identities, f_k would include all other kinds of fixed costs that do not depend on the particular social category, c_i .

workers who have identity concerns. I find tasks that involve similar skills but sharply differ in their identity associations. Importantly, the tasks are such that even if an identity-inconsistent task and its paired control task involve slightly different inherent costs f_k , it is plausible that the difference in these costs is similar across all workers, regardless of their identities.

To illustrate, suppose there are two large groups of workers belonging to social categories A and B who evaluate job offers as discussed above. The extra task involved in the offers is either b or u . Task b is associated with group B whereas task u has no association. Hence, working on b could cause identity-related concerns in A workers. Suppose that the following assumption holds:

Assumption 1. *Let $\theta_i = U_{ik}(c_i, 0)$ represent the net utility from taking up the offer that only involves the default task. The distributions of f_b , f_u , and θ_i are such that*

$$P[f_b > \theta_i | i \in A] - P[f_u > \theta_i | i \in A] \leq P[f_b > \theta_i | i \in B] - P[f_u > \theta_i | i \in B].$$

By construction, whenever $f_k > \theta_i$, the worker rejects any offer involving task k . Hence, this assumption concerns a difference-in-differences. The left-hand gives the difference in the probability of turning down task b relative to task u due to inherent fixed costs for group A . The right-hand side gives the similar expression for group B . Hence, this gives a sense in which the inherent difference in costs between task b and task u is weakly lower for group A than for group B . Under this assumption, the difference-in-differences comparison of take-up rates provides a lower-bound on the share of A workers who have identity concerns, i.e. for whom $\beta_{ik} > 0$. Appendix Section B.1 presents a more formal version of this discussion.

In addition, having multiple social groups ranked by status leads to more specific predictions. Theories suggest that workers are more averse to working on jobs associated with lower-status groups. In this case, the fixed cost may be described more precisely as:

$$F_{ik}(c_i) = f_k + \beta_{ik}^d \cdot I_k^d(c_i) + \beta_{ik}^l \cdot I_k^l(c_i) \cdot I_k^d(c_i), \quad (3)$$

where $I_k^d(c_i)$ is an indicator that equals one when task k is associated with a social group different from c_i , and $I_k^l(c_i)$ is an indicator for when this group has a lower status than c_i . Under similar assumptions, one could estimate lower bounds for the

shares of workers with positive β_{ik}^d and β_{ik}^l . If $\beta_{ik}^l > 0$, then this suggests that identity concerns are even larger when the extra task is associated with a lower status-group.

Finally, working on identity-inconsistent jobs can also lead to concerns about social image. Both being seen performing identity-violating jobs, and being perceived as the type willing to do so, can have negative social consequences. Hence some workers could face additional utility costs from taking up identity-violating jobs when their decisions are observable to others. By randomly varying whether worker decisions are publicized, one could measure to what extent workers have this type of social image concern.

2 Background surveys on caste

2.1 Caste system in India

The Indian labor market with the historical caste system provides an ideal setting for studying the effect of identity on labor supply. The caste system is composed of around 4,000 endogamous communities called *jatis* or castes. Caste membership, determined at birth, constitutes an important part of people’s identity and influences social, economic, and political spheres of Indian life even today (Deshpande 2011; Jodhka 2017; Mosse 2018).

Two features of the caste system are important for the experimental design. First, there exist historical associations between castes and occupations, some of which strongly persist in the current labor market (Iversen and Raghavendra 2006; and Guérin, D’espallier, and Venkatasubramanian 2015). The links between castes and occupations (e.g. Dhoba caste in Odisha and washer) often carry over to those between castes and simple manual tasks (e.g. washing clothes), which allows me to construct one-day jobs associated with specific castes. Second, caste is essentially a system of social hierarchy (Shah et al. 2017). In the current political system, Other Backward Class (OBC), Scheduled Castes (SC), and Scheduled Tribes (ST) are clusters of castes that are officially recognized by the Indian government for their historical disadvantage. Beyond these categories, individual castes form even finer layers of social order. This hierarchical structure enables me to examine worker responses to jobs linked to lower caste status.

In order to document the locally prevalent views on the caste-task links as well as the caste hierarchy, I conducted two surveys separately from the experiments.

It is crucial to understand the local context, since there are substantial geographic variations in how caste groups are distributed and perceived (Marriot 1958; Munshi 2019). Appendix Section B.2 provides a more detailed description of the caste system.

2.2 Survey procedure

Surveys and experiments were conducted in the state of Odisha during 2018-2019. During the initial scouting phase, the survey team visited more than 580 villages to get a sense of how castes groups are distributed across villages. Then a subset of villages were selected for surveys, while the remaining villages were reserved for experiment activities. The surveys utilized stratified convenience sampling; the goal was to survey around 10 male participants per caste from 15 different caste groups, while overweighing OBC and SC castes.¹⁸ Appendix Section B.3 describes the resulting sample compositions.

Task Survey (N=151) collected information about caste-task links, participants' experiences with tasks, and knowledge about castes. From a list of manual tasks, participants indicated for each task whether a particular caste performs it and the extent to which they have performed it.¹⁹ In addition, participants saw a list of caste groups present in Odisha and reported whether they knew of each caste.²⁰

Based on the results from Task Survey, seven caste groups comprising six SC castes and one OBC caste were selected for the experiment. Three of the SC castes were chosen due to their strong connections to manual tasks. The others were also drawn from SC and OBC, so that all groups would be broadly similar in terms of wealth and status, i.e. the chosen groups did not include castes of high socioeconomic status such as the Brahmins. Notably, only the castes that were known to over 70% of the survey participants were included. This was so that subjects in surveys as well as experiments would be similarly aware of the caste connections and rankings.²¹

¹⁸The survey and experiment samples only involved male workers to avoid both the confounding effects of gender identity and the practical difficulties with employing female workers in this setting. OBC and SC castes were overweighed in the surveys since the experiment sample involve workers from these categories.

¹⁹The list of tasks was prepared based on qualitative interviews prior to the surveys. The participants also indicated whether a task is gender-specific so that the experiments would only involve tasks that are not associated with the female gender.

²⁰The list of castes residing in Odisha was taken from the Additional Rural Incomes Survey & Rural Economic and Demographic Survey (ARIS/REDS) 2006 codebook.

²¹All SC castes that met this knowledge threshold were included. Since there were many OBC castes that met this condition, only one OBC caste was chosen for the experiment.

Rank Survey (N=209) documented how the seven castes are placed in the caste hierarchy. The survey sample enlisted only those who knew of all seven castes. Participants were provided with cards that had caste names written on them — with surveyors also reading out the names — and asked to arrange the cards according to their perceptions of caste hierarchy. They could place multiple names on the same level to indicate equal status of castes, but in practice, this happened rarely. To check whether the perceptions of hierarchy vary with specific contexts, I randomly gave the participants one of three types of instructions, which asked for the caste rankings to be based on: 1) general perceptions, 2) the practice of taking cooked food, or 3) the practice of taking water. The latter two practices were chosen because higher castes’ not accepting food or water from the hands of lower castes is among the most common behavioral rules attached to the caste hierarchy (Marriott 1958; Mahar 1960). The different types of instruction resulted in similar rankings, as discussed below.²²

2.3 Caste perception and ranking

Summarizing the results from Task and Rank Surveys, Table 1 shows caste ranking and caste-task associations. The table is organized such that connected castes and tasks are placed within the same rows, close to each other.

In Column 1, seven castes appear sorted according to their rank. Because participants from the Rank Survey knew of all the castes and ranked them without missing values, rank scores in Column 2 are just simple averages of the assigned ranks. Since there are variations in individual opinions about how these castes are ranked, Appendix Table A1 Column 1 compares the rank scores using an OLS regression, controlling for people’s tendency to inflate their own castes’ ranks. Testing the equality of coefficients for any two adjacent castes shows that the null is rejected at the 1% level. Furthermore, this ranking is consistent across all three versions of instructions, as shown in Columns 2-4. Hence, I take the ranking shown in Table 1 to be the representative view on the caste hierarchy and use it for analysis.²³

The remaining part of Table 1 describes how the castes are associated with simple

²²Afterwards, participants additionally ranked nine other castes, that were either SC castes not included in the experiment or other castes that participated Task Survey. The participants added the nine cards into the rank formation, skipping over any caste names they do not know.

²³An alternate approach would be to use individual-specific beliefs, which may correlate more strongly with their job take-up decisions. However, individual beliefs may be more biased. Appendix Table A1 shows that individuals tend to inflate the ranking of their own castes. In addition, eliciting the beliefs prior to the experiment can bias the take-up decisions, and vice versa.

manual tasks. The tasks are divided in different categories based on their connections to castes. Column 3 lists three tasks that have strong caste associations, which I refer to as “identity tasks” hereon. The share of survey participants who made these links are reported in Column 4. Specifically, 72% of the participants stated that washing clothes is specifically performed by the Dhoba caste. 97% and 84% associated mending leather shoes and sweeping latrines with Mochi and Hadi, respectively.²⁴

Column 5 shows three “paired control tasks,” which require similar skills as identity tasks but do not have caste associations. While these tasks also involve washing, mending, or sweeping, no participant linked these tasks to Dhoba, Mochi, or Hadi, respectively. Column 6 reports the share of participants that connects the tasks to any SC caste, including those outside of the experiment. Still, only mending grass (floor) mats is associated with some SC castes according to 15% of the sample, with the answers varying across subjects. Appendix Table A2 provides more information on all of the tasks above tasks as well as some additional control tasks, specifically regarding their caste and gender associations. Only the tasks that do not have strong associations with the female gender are used in the experiment so as to not introduce any gender identity effects.

I assign relative status based on these average perceptions about the tasks and castes. A task is considered a same-ranked task if it appears in the same row as the worker’s caste, and otherwise a different task. If a task appears in a row above (or below) the worker’s caste, the task is called a higher (or lower) task. For example, for a Mochi worker, mending leather shoes and mending grass mats are same-ranked tasks, washing clothes and washing farming tools are higher tasks, and sweeping latrines and sweeping animal sheds are lower tasks. Based on these assignments, I can examine how workers’ decision to take up job offers vary with tasks’ relative status, and compare across identity tasks and paired control tasks.

3 Experiment on job offer take-up

The goal of the main experiment is to truthfully elicit workers’ willingness to take up job offers that differ in caste associations. Consistent with the conceptual framework, the experiment is designed to capture changes in take-up rates that are due to spend-

²⁴Risley discusses some of these castes in his anthropological works from the early 20th century and reports similar ordering of the castes as well as caste-task associations (1908; 1892). The typical occupations for Dhoba, Mochi, and Hadi are reported to be washer, leather worker/cobbler, and scavenger, respectively.

ing small amounts of time on various tasks. This helps determine whether workers have strong reservations about engaging at all in caste-inconsistent tasks.

3.1 Setting and recruiting

The experiment sample is composed of 630 male workers drawn from 141 villages in Odisha, who primarily derive income from casual daily-wage labor. Wage laborers typically engage in agricultural work during peak planting and harvesting seasons and perform short-term contractual work in unskilled manufacturing or construction during the remaining lean periods. The experiment took place during the lean periods, namely October-November of 2018 and June-August of 2019.

The experiment involved offering one-day jobs at a manufacturing site, where the primary task is producing paper bags. Paper bags are commonly used in this setting by market or roadside vendors to store nuts or snacks. The general set-up and operation of the work sites were similar to those developed by Breza, Kaur, and Shamdasani (2018). All the produced bags were sold to local wholesale traders.

The jobs and recruiting process for the experiment were meant to resemble real labor market conditions as much as possible. Employers in this setting tend to recruit workers by visiting workers' villages, providing job descriptions, and making offers at market prevailing daily wage rates. Workers who agree to the offered terms may start work that day or on a prearranged, upcoming date. The recruiters and surveyors followed similar steps, as described below.

To construct a sample stratified by caste, the recruiters first visited a set of villages containing target caste groups and advertised an upcoming work opportunity. Potential participants were informed about location, the main work task, duration of employment, and compensation. Interested workers answered basic questions regarding eligibility and provided their contact information. Workers were deemed eligible if they satisfied the following criteria: (i) male household heads aged between 18 and 55, (ii) worked regularly as wage laborers, and (iv) spent less than 2 of the last 6 months away from the village. As most workers' castes could be identified from their last names, this process helped create for each caste, a list of eligible workers who are interested in a one-day job of making paper bags. Hence the sample excludes those who have better outside options than this offer as well as those whose main job is associated with any caste-specific task.

Then the surveyors visited the villages again over the next few days for the ex-

periment. Since there were a larger number of identified workers than needed for the experiment, surveyors approached them in a random order, typically surveying about 6 and no more than 12 workers per village.²⁵ Prior to the surveyors' visits, villages were randomized into two privacy conditions, as explained in the next subsection. The breakdowns of the sample by caste and privacy conditions are reported in Appendix B.3.

3.2 Choice exercise procedure

It is critical that the experiment is able to elicit workers' true willingness to take up job offers. In a simple survey, for instance, workers may exaggerate their dislike for certain offers since they are not foregoing real wages, or agree to many offers to signal their eagerness. Hence the experiment involves a procedure based on the Becker-DeGroot-Marschak (BDM) mechanism (Becker, DeGroot, and Marschak 1964). Each worker sees a list of potential job offers, and for each offer, indicates whether he would take up or decline the offer. After he indicates all choices, one offer is randomly selected and his stated choice for it is implemented. Because there is some chance that any stated choice is actually implemented, this mechanism is incentive compatible.²⁶ In addition, the experiment involves explicitly asking workers to consider each offer as a single take-it-or-leave-it offer and giving a "simple honest answer" about what they prefer, in order to further encourage truth-telling. In the main analysis, I treat workers' answers as reflecting their true willingness to take up job offers, and later discuss some potential issues with the approach.

The experiment proceeded as follows. Each worker had a private conversation with a surveyor. First, a worker went through a practice exercise which was designed to help him understand the BDM mechanism. Specifically, he was offered a chance to buy different combinations of packaged foods (e.g. 40 grams of mustard seeds and 60 grams of sugar). While the offers differed in types or quantities of products, all of them involved the same price. The worker could choose to accept—purchase the combination—or decline each offer. After he indicated all decisions, one offer was randomly selected and implemented according to his choice.

Second, the surveyor described the set of all potential job offers. The offers were

²⁵An exception was made for the Kela workers who are concentrated in a smaller number of villages. So one village contains 23 Kela workers.

²⁶Specifically, it is incentive compatible under risk neutrality. For more details on the mechanism and its use in experiments, see Fudenberg, Levine, and Maniadis (2012).

exactly the same in most aspects, including fixed daily wage of Rs. 300, total working time of five hours, work site location, employer, and so on. All offers required spending the majority of working time on the default task of producing paper bags, and the remainder on an extra task. Job offers varied only in the *type* of extra task and in the *time* required to work on it. The worker’s choice set always included eight extra tasks: three identity tasks, three paired control tasks, and two pure control tasks (ref. Table 1). There were also four different time requirements for the extra tasks: 10 minutes, 30 minutes, 1 hour, and 1.5 hours. To test for any order effect, the order in which tasks were discussed was randomized in four different ways. Time requirement was randomly presented either in an ascending or descending order.²⁷

The surveyor explained each offer in detail and showed photos depicting the tasks, such as those in Appendix Figure A1. The worker was told that none of the extra tasks required any prior experience or training. In addition, for the two tasks involving specific skills, i.e. mending leather shoes and grass mats, the worker was told that he would only be assisting experienced trainers. Importantly, the extra task was to be performed in a private space where other workers could not observe his performance. Throughout the first two stages, the surveyor asked multiple questions to verify the worker’s understanding and provided more explanations in case the worker failed to answer correctly.

Third, the worker was asked to go over the list of job offers and indicate whether he would take up or decline each offer. The surveyor re-iterated that only one offer would be randomly selected and given out. The randomization was justified with the explanation that the employer was looking for people to complete all the different tasks and wanted to be fair while giving out a limited number of job offers. Since the employer was interested in finding out which tasks people are willing to perform, the surveyor encouraged the worker to answer honestly about what they prefer.²⁸

Fourth, one offer per worker was randomly selected, and his choice for this offer was implemented. All offers had the same chance of being selected and one offer was chosen based on the worker’s rolling dice and drawing scratch cards. If the worker

²⁷The time lengths were chosen to create as much variation as possible while making the jobs sound realistic and be practical given the constraints at the work sites. The set of control tasks was also randomized across workers: one task was stitching and the other was either making ropes or deshelling peanuts (ref. Appendix Table A1).

²⁸It was not costless for the worker to accept an offer and change his mind later. He was aware that if an accepted offer was selected, the work site managers would visit his village over the following days to coerce him to complete the job.

had chosen to take up this offer, he could complete the job within the next three days and receive Rs. 300. If the worker had refused the offer, the worker would not receive any other job offer.²⁹

Finally, the worker was asked to complete a follow-up survey. If the worker visited the work site to complete the job, he was asked to do the survey at the work site. Otherwise, the worker was offered to complete the survey and receive a gift worth Rs. 50. This compensation was offered so that those who did not end up with job offers would still be willing to complete the survey. The gift was not mentioned until this step, so that the worker would not factor this into his decisions during the choice exercise. The follow-up survey completion rate was high at 87%.

While the above procedures were common to all workers, the worker received different instructions during the second step of the choice exercise depending on whether the worker's village was randomized into public or private conditions. Each village was scheduled to host a focus group meeting in the days following the exercise. Local agricultural practices were to be discussed in these meetings, and many village members, including those who did not participate in the experiment, were invited to attend. If the worker was in the public condition, he was told that all his choices during the job offer exercise would be openly discussed during these meetings, irrespective of his attendance. If the worker was in the private condition, he was assured that his choices would remain private information, except for his willingness to wash farming tools, a control task.³⁰ Hence, the two conditions were designed to differ only in the observability of the worker's decisions, not in that of his job performance nor his beliefs about other focus group activities.

This design was registered on the AEA RCT registry, ref. AEARCTR-0003493.

3.3 Regression specification

To examine how workers' responses to job offers vary with predicted presence of identity violations, I estimate a linear model by OLS using observations at the worker-

²⁹57% of workers received offers that they were willing to take up, and 67% of them completed the jobs. The completion rate is not high, as absenteeism is prevalent in this region among casual contract jobs (Krishnaswamy 2019). Later I perform robustness checks using completed jobs.

³⁰The justification was that discussing local agricultural practices would involve talking about people's willingness to wash farming tools.

task-time level:

$$\begin{aligned}
Y_{ikt} = & \sigma^d \text{different}_{ik} + \lambda^d \text{different}_{ik} \cdot \text{identity}_k \\
& + \sigma^l \text{lower}_{ik} + \lambda^l \text{lower}_{ik} \cdot \text{identity}_k \\
& + T'_{tk} \tau + P'_k \rho + X'_i \nu + \epsilon_{ikt}.
\end{aligned} \tag{4}$$

In this basic specification, the dependent variable Y_{ikt} is an indicator for whether worker i is willing to accept the job offer that requires spending time t_k on task k . The independent variables, identity_k , different_{ik} , and lower_{ik} , are indicators for whether task k is an identity task, a different task, and/or a lower task, as defined in Section 2.3. P_k is a vector of task-specific indicators, and X_i is a vector of worker-level or caste-level indicators. This specification controls for task-specific linear time trends, with T_{tk} referring to the amount of time assigned to task k in minutes. Standard errors are clustered at the worker-task level, since identity concerns are predicted to vary at this level. I show robustness to alternate specifications, e.g. controlling for alternate time trends or worker characteristics, with the results virtually unchanged.

The key coefficients of interest are λ^d and λ^l , relating to the effect of caste-inconsistency on (the desire for) take-up. These coefficients measure the gaps in take-up between the offers involving identity tasks and those involving paired control tasks, separately across relative status. Specifically, λ^d estimates how much larger the take-up gap is for higher tasks, compared to that for same-ranked tasks, and λ^l estimates the additional increase in the gap for lower tasks. Due to the task-specific time controls, these estimates takes out any effects on take-up from spending longer time on the tasks.

These coefficients can measure the shares of workers who have identity concerns under some specific assumption. Assumption 1 implies that the inherent fixed cost of working on an identity task compared to its paired control task is not higher when tasks are higher or lower. The regression specification above relies on a related, weaker assumption, which only holds conditional on task- and caste-specific effects. That is, some tasks could involve higher fixed costs for everyone, and some caste groups could face higher fixed costs when working on any task, but the gaps in inherent fixed costs would not differ across caste groups. If true, the coefficients λ^d and λ^l would provide lower bounds on the effect of caste-identity violations on take-up.

The assumption behind this specification may be reasonable for a number of rea-

sons. First, the tasks are specifically chosen to involve minimal skills and most tasks resemble common household chores such as washing and sweeping. It is difficult to imagine why people would have strong aversions to, for example, washing clothes compared to washing tools, if not for some internal psychological costs surrounding caste-based practices. Second, even if some identity tasks were inherently more difficult or unpleasant to perform, these differences would matter for the estimates only to the extent to which they affected the fixed utility costs. For instance, if one caste group had extensive experience with washing clothes and enjoyed working on it more, this effect would be captured by the time controls rather than the key coefficients. Third, because the assumption specifically concerns the differences-in-differences in fixed costs across task category and relative status, it is likely to hold even when caste groups differ in various aspects such as wealth and education. For the assumption to be violated, some characteristics would have to correlate with caste ranking in such a way that their impacts on the costs of working on identity tasks do not monotonically increase with ranking, but differ based on how the tasks' connected castes rank against own castes. While it is not immediately clear which worker characteristics could satisfy this property, I discuss some potential threats to the assumption after presenting the experimental results.

4 Results: conflicts of identity lower job take-up

The experimental results show that workers are often averse to taking up job offers associated with castes other than their own, and especially so when those castes rank lower than theirs. This is the case, even when workers make such decisions in private. The set of findings, along with workers' stated opinions, indicate that some workers face strong concerns about violating caste-based behavioral rules.

4.1 Visualizing offer take-up rates

I first use plots to examine the basic patterns in the raw data. The data is at the worker-task-time level, and the outcome is whether worker accepts a job offer involving a specific amount of time on a specific extra task. Figure 1 plots the average take-up rate against the time required on extra tasks, separately by task category and relative status. Plot headings indicate task category (paired control vs. identity tasks) and markers indicate relative status as defined in Section 2.3: same-ranked, higher, or lower.

Circular markers have similar positions in both panels, indicating that take-up rates are similar between (the offers involving) paired control tasks and identity tasks. For higher tasks, drawn with triangular markers, take-up rates of identity tasks are significantly lower compared to those of paired control tasks. Finally, as indicated by rectangular markers, the gaps in take-up increase further for lower tasks. The connected lines appear approximately linear and parallel to one another. Since the total working time and wage are fixed across the offers, the negative slopes of these lines suggest that workers in general prefer to spend time on the default task compared to the extra tasks. However, the changes in take-up associated with time variations are small compared to the distances across the lines. This implies that take-up decisions vary mainly due to the costs of engaging at all in different extra tasks, rather than the costs of working longer on any task. As predicted by the theories of identity, workers are more likely to turn down offers that involve working on identity-inconsistent tasks, especially if those tasks are associated with lower-ranked castes.

Whereas in Figure 1, the averages are calculated by pooling across tasks and castes, Appendix Figure A2 plots the results separately by task and caste-level. For ease of visualization, here I group the seven castes into four levels based on how many tasks are considered lower: level 4 contains Hadi, which is the lowest ranked caste, level 3 contains Mochi and Pana, level 2 contains Dhoba and Kela, and level 1 contains Kaibarta and Sundhi. This way, the same markers refer to the same caste levels across panels. For identity and paired control tasks (center and right columns), I plot the take-up rates for the same-ranked tasks separately as green lines with hollow circular markers. I also indicate relative task status by using dashed lines when tasks are considered lower.

The figure shows that higher castes are generally associated with lower take-up, as blue and light-blue lines tend to lie above the other lines. However, in the case of control asks (left column), the lines are close to each other and even overlap. In the case of identity tasks (right column), two deviations are notable. First, take-up rates are higher when tasks have historical associations with own castes. For example, even though Dhoba and Kela are similarly ranked (level 2), Dhoba workers have higher take-up for washing clothes (top right, hollow markers) compared to Kela (triangle markers). Similarly, Mochi workers are much more willing to mend leather shoes (middle right, hollow markers) compared to Pana (square markers). In the case of paired control tasks, however, the take-up rates are similar (top center) or the

pattern is reversed (middle center), indicating that workers are only more willing to do the identity tasks that have direct connections to their castes. Second, whether tasks are higher or lower matter more in the case of identity tasks relative to paired control tasks; the panels in the last column show greater distance between the solid and dashed lines, compared to those in the center column. Overall, these patterns confirm the findings from Figure 1.

4.2 Regression estimates of the impact on take-up

I use regression analysis to examine how take-up varies with predicted presence of identity violations. Table 2 reports the results from running ordinary-least-squares (OLS) regressions based on the empirical specification in Equation 4. Columns 1-2 correspond to the basic specification controlling for task and caste (or worker) fixed effects and linear time trends. The key coefficients are those on *Different* \times *Identity* and *Lower* \times *Identity*, which measure how much lower the take-up rate is for identity tasks compared to paired control tasks. When tasks are same-ranked, take-up rates are similar for identity and paired control tasks, as shown in the table footer. When tasks are higher, workers are 23 percentage points (pp) more likely to refuse offers for identity tasks compared to paired control tasks. This gap increases by an additional 28 pp in the case of lower tasks. These coefficients are statistically significant at 1% level and consistent with the patterns shown in Figure 1.

I check whether these results can be explained by worker differences in age, education, or wealth. A worker is expected to evaluate each job offer against his outside option of having no offers, and the regression in Column 2 controls for worker fixed effects. Hence, any individual characteristic that affects workers' general willingness to take up one-day jobs should not change the estimates. Nonetheless, it is possible that some characteristic would make workers more averse to performing certain tasks, e.g. wealthier workers may dislike wash clothes more. The summary statistics reported in Appendix Table A3 show that workers in higher caste levels tend to be older, more educated, and wealthier than the rest. To address this issue, Table 2 Column 4 controls for the interactions of task-specific dummies with survey measures of age, education, and wealth.³¹ Column 5 instead controls for the binaries for whether

³¹Wealth PCA score is generated by performing principle component analysis with the variables reported in rows 5-12 of Appendix Table A3. Specifically, the analysis uses inverse hyperbolic sine of last month's income and ten indicators for asset ownership: sewing machine, bicycle, motorcycle, fridge, radio, tv, mobile phone, stove, and watches.

age, education, or wealth is greater than the median, interacted with task dummies. The results are robust to adding these controls.

The ease of finding jobs outside of the experiment might be particularly important for take-up decisions and could be proxied by the number of paid work days in the past week. Additionally controlling for task-level interactions with this, however, does not alter the key coefficients shown in Columns 4-5 (results not shown). In addition, Appendix Table A4 shows that the results are robust to adding more controls, including quadratic time trends, and surveyor, question order, and choice set fixed effects. The results are also similar when I exclude those who score low on comprehension questions, or those who exhibit choice inconsistency, i.e. refusing some offer and accepting another offer involving longer time on the same task.³² Hence it seems unlikely that the results are driven by caste-differences in understanding of the procedures.

4.2.1 Using alternate caste rankings

The initial online registration of the design mis-specified the ranking of two castes, Kaibarta and Dhoba. This was because the registration—and the launch of the main experiment—happened before the Rank Survey was fully completed.³³ The registered ranking, which was based on field interviews and partial survey data, deviates from the final survey ranking by classifying some tasks as being different rather than lower for those two castes.³⁴

Although the registered ranking is likely to be inaccurate, given the heterogeneity in survey answers, it may be useful to examine how the results change with alternate rankings. Appendix Table A5 Columns 1-3 show results from running the main regressions in Table 2 using the registered ranking. The coefficient on *Different* \times *Identity* is larger and the one on *Lower* \times *Identity* is smaller, with the sum unchanged.³⁵ I also inspect how these results change when I partially correct the ranking. Appendix

³²Section 3.2 describes how question order and choice set are randomized across workers. 25% of workers who score 5 or fewer out of 7 comprehension questions correctly, i.e. perform worse than the median worker, are classified as having low comprehension. 17% of workers have at least one choice inconsistency across all offers.

³³The launch of the main experiment was rushed due to time constraints associated with agricultural seasons.

³⁴The tasks mis-classified this way include washing clothes for Kaibarta, and mending shoes and sweeping latrines for Kela, as well as the corresponding paired control tasks.

³⁵The sum of the two coefficients is not statistically significantly different from the sum shown in Table 2.

Figure A3 Panel A shows the full variations in respondents’ reported rankings. Despite heterogeneity in answers (especially for Kela), it is apparent that Kaibarta ranks higher than Dhoba, and Kela ranks higher than Hadi. Appendix Table A5 Columns 4-6 show that the results based on these partial corrections are similar to those using the final ranking. Appendix Table A6 also show results dropping one caste at a time, and the estimates are similar when Kaibarta or Kela caste is excluded from the sample.

4.2.2 Offer randomization and job completion

Since the main data is the worker-task-time level, I look at how the results change when I use fewer observations per worker. In Appendix Table A7 Column 1, the dependent variable an indicator constructed at the worker-task level, which equals one if the worker accepts any of the offers involving the task. The results are nearly identical to those using the full data. In Column 2, I run the same regression, using just one observation per worker by only using the randomly selected offer. While the coefficient and standard error on *Different* \times *Identity* become larger, the overall results are again quite similar. These findings indicate that randomization for selecting job offers was implemented successfully and that the results are robustly statistically significant.

In addition, I examine how completion of jobs compares to stated willingness to take up job offers. While 57% of workers received job offers which they indicated willingness to take up, only 67% of them actually completed these jobs. Because workers were allowed to complete jobs within three days following the offer, it was possible that some did not complete jobs due to unforeseen events or change of mind. As worker absenteeism is prevalent in this region (Krishnaswamy 2019), the low completion rate is not usual, but it is important to see how completion varies with task category and relative status. Appendix Table A7 Column 3 shows that the coefficients on *Different* \times *Identity* for job completion is much larger compared to those in Columns 1-2.³⁶ This indicates that workers are more likely to renege on the agreed offers if the offers are associated with castes other than their own. The job completion result being in line with the predicted effects of identity concerns lends credibility to the elicitation method in the experiment: it is inconsistent with workers pretending to dislike certain offers under the belief that it would never result in actual loss of a

³⁶Column 4 shows that the rate of completing the follow-up survey does not vary across task category and relative status.

job offer. It is still possible that some workers understated their aversion to caste-conflicting jobs. However, an alternate explanation would be that workers stated their preferences truthfully, but further reflecting on the offers or discussing them with others convinced them to refuse the jobs in the end.

4.2.3 Heterogeneity analysis

If the identity channel explained the findings, one could expect that the effects would be larger among those who are more traditional, i.e. have stronger preference for following caste-based norms or behavioral rules. To test this idea, I categorize some workers as being caste sensitive based on their opinions. During the follow-up survey, workers listened to seven vignettes about characters violating various caste norms—related to job, marriage, and food sharing—and stated whether they approve of those behaviors.³⁷ As the median worker expresses disapproval on four out of seven questions, 40% of the workers who have traditional views on strictly more than four scenarios are designated as caste sensitive. As expected, caste sensitivity is positively correlated with being older and less educated.³⁸

Traditional workers are even more likely to turn down the job offers associated with lower identity tasks. In Appendix Table A8 Column 1, the coefficient on $Traditional \times Different \times Identity$ is small and not statistically significant, indicating that caste sensitive workers are similarly averse to taking up offers associated with higher castes as the remaining workers. The coefficient on $Traditional \times Lower \times Identity$, however, is negative and statistically significant at the 1% level. This means that caste sensitive workers are especially unwilling to take up offers associated with lower castes. The results are qualitatively similar when I use a caste-sensitivity measure based on only the vignettes unrelated to jobs or only the remaining ones (results not shown). Columns 2-3 show that older or less educated workers are also more averse to taking up jobs involving lower identity tasks. These results are in line with the idea that those who hold more traditional views about caste norms would be more likely to have concerns about taking up caste-inconsistent jobs. However, it is notable that even among relatively less traditional workers, predicted presence of identity concerns is strongly associated with lower offer take-up.

³⁷The questions are listed in Appendix Section B.4; four of them are also used in the Task Survey.

³⁸The correlation coefficient is 0.13 with being older, i.e. age being greater than the median, and is 0.09 with being less educated, i.e. years of education being less or equal to the median.

4.3 The role of social image concerns

I investigate the potential role of social image in explaining the results. Some workers may not have any intrinsic identity concerns about taking up caste-inconsistent job offers, but still want to avoid being perceived by others as being willing to do so.³⁹ If so, the estimated effects on take-up would be larger among the workers who expect their take-up decisions to be publicized. As described in Section 3.2, workers were randomized into private or public conditions at the village level. Appendix Table A9 Columns 1-2 show that worker characteristics are balanced across the two groups.

To see how worker responses to job offers differ due to the observability of decisions, I re-run the main regressions, this time with the key covariates interacted with the indicator for the public condition. Table 3 shows that the results are similar across the two privacy conditions. The coefficients on *Different* \times *Identity* and *Lower* \times *Identity* are similar to those in Table 2 and statistically significant at the 1% level. The coefficients on their interactions with *Public*, however, are close to zero and not statistically significant. Standard errors are clustered at the village level since the treatments are randomized at this level.⁴⁰

It is worth noting that this finding does not imply that workers do not face worries about other people’s judgments or reactions. It is likely that the workers with social image concerns also have strong identity concerns about taking up caste-conflicting jobs. Then they would turn down those job offers regardless of whether their decisions are publicized, leading to similar results across the treatments.

There are some caveats to this interpretation. Under both conditions, workers disclosed their take-up decisions and opinions to surveyors, and hence could have social image concerns towards surveyors. For this to drive the results, however, workers would have to be highly conscientious about their images towards surveyors—whom they may never interact with again—and yet face no additional pressure about upholding images towards friends and neighbors. It is also possible that workers did not believe that their decisions would be kept private; surveyors might fail to keep their promises or it could be costly to lie to others who would ask about their decisions.

³⁹In this setting, networks are important to livelihoods (Munshi 2019), and therefore social image concerns could be particularly salient. Recent field experimental studies in similar settings show that social image concerns affect various outcomes, such as savings (Jakiela and Ozier 2016), child vaccination (Karing 2021), and wage floor (Breza, Kaur, and Shamdasani 2019).

⁴⁰Appendix Table A10 reports the regression results with standard errors clustered at the worker time task level, consistent with Table 2. The results are similar regardless of the clustering level.

Still, there are a few reasons why the identity channel seems to be the most likely explanation for the findings.

First, the privacy variation here is similar to that used in another study in the same setting, which finds social image effects with regards to wage-related norms. Breza, Kaur, and Krishnaswamy (2019) show that workers' willingness to take up jobs at wages below the market prevailing rate increases when workers are told their decisions will be kept confidential. Given that their study involves a similar sample composed of daily-wage laborers in Odisha, it seems that the method should be effective at altering worker expectations about privacy to some degree. Yet, the results in this experiment hardly change across the treatments.

Second, workers' stated reasons for refusing job offers are more in line with concerns about identity. During the follow-up survey, workers who turned down all offers involving a particular task are asked why they turned down those offers. Figure 2 shows the shares of answers that relate to only identity (e.g. feeling ashamed, task for lower castes), only social image (e.g. unacceptable to family or neighbors), both identity and social image, or neither (e.g. task is difficult to do, never done the task before).⁴¹ Among those who turn down the offers involving identity tasks, half of the workers mention reasons only concerning identity and a quarter bring up factors relating to both identity and social image. Only a small share of workers (7%) just talk about other people's judgments and reactions as motivations.

Third, in this setting, people's personal opinions regarding caste norms appear similar to their beliefs about other people's opinions. The privacy treatments would capture social image effects if some workers wished to privately take up certain offers but believed that others would disapprove.⁴² To test whether this condition is plausible, the Task Survey asked four of the seven vignette questions which describe characters violating various caste norms.⁴³ A randomly selected half of the participants were asked whether they approve of the characters' actions in their personal

⁴¹Surveyors did not read out any options but marked all applicable options based on workers' free-form answers. If workers mentioned feeling ashamed or embarrassed in front of others, this answer would be categorized as being unacceptable to others. Some people refuse control tasks citing reasons relating to identity or social image, claiming the tasks are too menial for them.

⁴²Bašić and Verrina (2021) provide a detailed discussion on how one's private perceptions about what is appropriate to do can differ from what the society finds appropriate, and how both types of perceptions affect behaviors.

⁴³These questions are listed as Q1-Q4 in Appendix Section B.4. Two questions are related to the practice of taking up lower-caste jobs.

opinions. The rest were asked whether they thought their friends and neighbors would approve. Figure A4 shows that on each of the questions, the shares of participants expressing disapproval is remarkably similar regardless of how the questions were asked. This consistency between first- and second-order beliefs can explain why the privacy treatments have no effect on workers' decisions.

4.4 Alternate explanations

I discuss here whether an explanation other than identity concerns could produce the above findings.

4.4.1 Worker experience

One potential factor that could drive the patterns in take-up decision is worker's prior experience with tasks. The idea is that some caste groups may be more experienced with specific identity tasks, e.g. due to the tasks being more commonly performed among certain caste-networks. In addition, people may be generally averse to working on any tasks for the first time. Then, the gaps in take-up would be attributable to the differences in familiarity, rather than the concerns about violating caste-based behavioral rules. Since the experimental data does not contain information on workers' personal experience, I use the two background surveys to look at the variations in experience.⁴⁴ This provides two advantages: 1) the background surveys use a different sample living in the same broad region, so unlike the follow-up survey, people's answers cannot be influenced by job take-up decisions and outcomes; and 2) the background survey sample comprises ten castes—four of which overlap with the experiment sample—and cover a wider range of castes in terms of ranking and socioeconomic status. Hence, the variation in experience is potentially larger than in the experiment sample and can be estimated more precisely.

The surveys contain information on whether people have performed each task in own household, for friends or neighbors without wage (outside of household), or for wage. Appendix Table A2 shows that most people have experiences with doing the washing and sweeping tasks at home, but few people have experience with performing any task for wage. To see how these experience levels vary across task category and relative status, I estimate a linear model similar to Equation 4 using observations at the worker-task level, omitting linear time controls.

⁴⁴Appendix Section B.3 provides more details on this sample is constructed.

Appendix Table A11 Columns 7-8 show that people are more likely to have ever performed the tasks associated with their own castes, i.e. same-ranked identity tasks as compared to the paired control tasks. When tasks are not same-ranked, the experience gap becomes completely offset, as indicated by the large negative coefficients on *Different* \times *Identity*. If this experience is crucial to the offer take-up decisions, then in the main regression of take-up, the coefficient on *Different* \times *Identity* could be an overestimate of the identity effect. Even so, it is notable that the coefficients on *Lower* \times *Identity* are small and statistically insignificant for all types of experience. When tasks are not same-ranked, there is no evidence that experience gaps could differ depending on relative status. Therefore, in the main regression of take-up, the coefficient on *Lower* \times *Identity* is still unlikely to be biased upwards and can provide a lower bound on the effect of caste-identity concerns on labor supply.

Overall this analysis suggests some caution with interpreting the key coefficients. Nevertheless, I generally discuss both key coefficients as capturing identity effects. One justification is that the experience gaps would bias the estimates only under some strict conditions, e.g. workers are averse to working on tasks for the first time even when it requires no special skill or training, their performance does not affect wage or future job opportunities, and no one can observe their performance. In addition, in the supplementary experiment, I collect data on experience directly with the sample and show that controlling for workers' personal experience does not affect the estimated gaps in willingness to engage in identity vs. control tasks.

4.4.2 Other explanations

Expectations about the employer. One may be concerned that workers form different expectations about employers depending on the type of job offers. This is unlikely since the offers are explicitly described as one-time offers coming from the same employer providing work at the same location. One may also wonder whether the results are driven by workers fearing caste-based discrimination. However, as discrimination is typically practiced against lower caste workers, fear of discrimination has difficulty explaining why given an offer associated with a specific caste, groups that rank higher would be more averse to taking up the offer. In addition, the results seem unrelated to beliefs about surveyors' castes, as they cannot be identified just based on appearance.⁴⁵

⁴⁵The follow-up survey shows that 71% of workers have difficulty forming any guess about surveyors' castes and another 25% are not confident about their guesses.

Surveyor demand effect. Another concern is that workers may believe that the employer (or the surveyor) prefers to hire specific castes depending on the job offer. To reduce this concern, the surveyors were careful not to bring up any mention of caste during the choice exercise. This was possible because the recruiting team obtained information on workers' last names—and thereby identified their castes—prior to soliciting them for the experiment. In addition, the employer was described as searching for workers to complete various extra tasks and wanting honest answers about the workers' preferences. Hence in order for this concern to be valid, workers would have to form beliefs inconsistent with the job description and yet be willing to give up potential job offers due to those beliefs.

Status. Status may be important for labor supply decisions in general. Even when offers only involve control tasks—without any caste associations—some workers turn them down stating that it would be shameful to take up menial jobs. However, this study's main estimates come from comparing jobs with and without caste associations across relative status, controlling for worker and task fixed effects, e.g. taking out the effect from a task being commonly perceived as degrading. Hence, the main findings specifically relate to the effect of caste status; workers are less willing to take up offers associated with groups with caste status different from their own, and especially so with lower caste status.

Untouchability and religion. The historic and currently illegal practice of untouchability socially segregates groups such as those belonging to Scheduled Castes and Scheduled Tribes. The practice is rooted in Hinduism and governs various domains, including sharing food/water, seating arrangements at weddings, entering places of worship, etc. and are still adopted by some in social spheres (Shah et al. 2017). However, the general tendency to segregate SC/ST cannot explain the results as six of the seven experimental castes belong to SC. In addition, the offered jobs do not involve any inter-personal interactions among workers. Hence, untouchability could explain the results only in the sense that the workers' religious beliefs may prescribe how they ought to treat caste-associated jobs, which is analogous to the channel of caste identity.

5 Supplementary experiment: pricing identity violations

The results thus far indicate that identity is an important factor constraining workers' labor supply decisions. To the workers in the sample, a one-job represents a valuable income-earning opportunity during agricultural lean seasons; they report finding only about two days of paid work in the week prior to the experiment. Yet, 85% of workers turn down at least one job offer involving caste-inconsistent tasks.

One may wonder whether these workers would be eager to take up caste-inconsistent jobs as soon as the jobs offered higher wages. Theories of identity suggest that violating identity-based behavioral prescriptions could lead to severe psychological costs, and some will be willing to incur large monetary losses to protect their identities. Some actions may become taboo-like in the sense that identity-concerned individuals would refuse to take those actions regardless of price (Bénabou and Tirole 2011). This means some workers may act as if they would never take up caste-conflicting jobs, regardless of offered wages. The supplementary experiment aims to test for this behavior and quantify the costs of identity violations by varying the wages associated with job offers.

5.1 Experimental procedures

The supplementary experiment involves a new set of 106 workers belonging to Kaibarta and Pana castes—which are not associated with any experimental task. Workers get started on a one-day job of producing paper bags, the default control task. Then they individually talk to surveyors, who inform them about a chance to switch to working on a different task for part of the remaining working time. As in the main experiment, the switching offers involve variations in the extra task's type and time requirement, with privacy conditions randomized across workers.⁴⁶ Notably, these switching offers involve a bonus wage payment on top of their daily wage of Rs. 300. The amount of bonus is drawn from the following list: 0, 30, 60, 90, 120, 180, 240, 300, 1500 or 3000. Hence the maximum bonus amount is ten times the daily wage, close to one month's wage income during agricultural lean seasons. Workers go over the entire

⁴⁶The options are slightly reduced to save time. The offers involve the same identity tasks and paired control tasks, in addition to one pure control task of moving bricks—chosen to be the most physically exacting task. The time options are 10 minutes, 30 minutes, or one hour.

set of potential switching offers, each linked to the extra wage list, indicating their willingness to take up a given switching offer for a given extra wage amount. After workers indicate all of their choices, a combination of offer and wage is randomly selected, and the worker’s choice for this combination is implemented.

This design departs from the previous one in a number of ways. While the main experiment focused on making the decisions seem as similar as possible to the ones workers make everyday in the labor market, the amounts of bonus offered may make the choices here seem removed from reality. However, making workers get started at the worksite first—which involve meeting supervisors, seeing the worksites, learning to make paper bags, etc.—should help workers consider these decisions seriously. The outside options from refusing offers are different across the two experiments, since the workers keep their default job in the latter one. However, an alternative design that involves offering higher wages in the first experiment, as well as the current design, both provide bounds on the utility costs of spending some time on an extra task instead of the default task in wage amounts.⁴⁷ The supplementary sample only involves two castes, because the goal is no longer to establish that identity effects exist using variations in relative status. This design focuses on testing whether workers are willing to forego very large amounts of wages to avoid caste-inconsistent tasks, which would suggest concerns about violating internal rules of behavior.

5.2 Results: responses to extra wage

I first plot the average take-up rates of offers when the offers are linked to two specific bonus amounts. In Panel A of Figure 3, the dashed lines plot take-up rates against the time required on the extra tasks when bonus is 10% of daily wage. In both panels, the lines are downward-sloping, suggesting that the variable costs of working on the extra tasks is greater than those for the default task and therefore take-up falls with longer time requirements. The take-up rates are always lower for identity tasks, consistent with the idea that working on identity tasks involves larger fixed costs due to identity concerns. The solid lines plot responses when bonus is ten times daily wage. These lines appear flatter, suggesting that the variable costs matter less for take-up decisions when the offers involve the maximum amount of bonus. Despite this large bonus, 43% of workers still refuse to spend ten minutes on identity tasks,

⁴⁷Moreover, these bounds would be the same if the utility in money is linear over the relevant wage range.

as compared to the 15% refusal rate for paired control tasks.

I then examine how workers tendency to refuse all switching offers differ across task type using OLS regressions. In Table 4, the dependent variable is whether a worker refuses all offers involving a particular task regardless of time requirement and bonus amount. *Paired tasks* is an indicator for both identity tasks and paired control tasks, so the omitted category is the pure control task. The coefficient on *Identity tasks* measures the difference in refusal of identity tasks compared to paired control tasks. Consistent with Figure 3, Columns 1-2 show that this estimated gap is 28 pp when controlling for caste- or worker-specific fixed effects. The results are similar when the regression additionally controls for the interactions of task-specific dummies with survey measures of age, education, and wealth (Column 3). The results are again invariant to whether workers' decisions are publicized (Columns 4-6).⁴⁸ In Appendix Section B.5, I discuss how these results are robust to adding more control variables related to workers' experience and comprehension. The appendix also provides some discussion on the validity of the BDM elicitation method, i.e. how to interpret the results given that workers evaluated multiple offers instead of just one.

I also inspect the bonus wage amounts at which workers are willing to accept working on extra tasks. Figure 3 Panel B plots the minimum amount at which workers accepts any offer involving extra tasks. Those who refuse all offers regardless of wage offered are put into the bar labeled " $\geq 3K$." The distributions in both panels look clearly bimodal, with most workers represented in 0–30 or in $>3k$. The share of workers who are only willing to switch at some extra wage larger than 10% of daily wage is only about a fifth of the sample. These results suggest that there are roughly two types of workers: identity-concerned workers who are decidedly averse to working on caste-inconsistent tasks vs. those who are unconcerned about this behavioral rule. For the latter, identity tasks do not seem particularly difficult or unpleasant to work on, as they are willing to accept just a small bonus or even none at all. Crucially, the former workers could be resistant to ever taking up caste-inconsistent jobs even if those jobs became widely available and offered high wages.

It may be surprising that some workers reject even working on paired control tasks. Looking at the joint distribution of decisions offers some potential explanations.

⁴⁸One caveat is that worker characteristics are not well balanced across privacy conditions in the supplementary experiment, as shown in Appendix Table A9. This may be due to the small sample size with just 12 villages under randomization.

The tendency to refuse a task, i.e. refuse all offers regardless of time and bonus, is concentrated among a subset of workers. In Appendix Table A12, I divide workers into different groups based on the number of refusals of identity tasks and report how many in a given group reject each task—excluding seven workers who refuse all tasks. Those refusing control tasks mostly belong to the group that refuse one of more identity tasks (Column 2), in particular those who refuse all three identity tasks (Column 4). In addition, the refusal rate of control tasks is mainly driven by the task of sweeping animal sheds—a task typically perceived as a job for poor people. Therefore, one potential explanation is that workers with strong caste identity concerns also tend to face status concerns (unrelated to caste), which lead them to turn down sweeping animal sheds. Another is that the unusual nature of the offers in the supplementary experiment—which provides a high monetary incentive for violating one’s identity—may make some workers refuse to engage in thinking about even the control offers.⁴⁹ Given these considerations, only the gap between take-up of identity and control tasks may be interpreted as the lower bound on the effect of caste identity violations.

Finally, I investigate whether survey answers help predict which workers have strong identity concerns. In Appendix Section B.5, I discuss how most variables related to workers’ demographics and wealth, typically found in other surveys, appear to have limited predictive power. This could be partly due to the small sample size. Even so, it is plausible that factors that shape workers’ personal values and principles, e.g. religious or migratory experiences which the surveys here do not cover, are pertinent to determining whether workers are averse to caste-inconsistent jobs. Future studies could investigate why individuals adopt or abandon particular identity-based prescriptions and explain how extreme divergence in behavior could arise even within the same social group.

6 Conclusion

This study finds that concerns about caste identity importantly constrain labor supply decisions of casual workers in Odisha, India. Despite having interest in an one-day manufacturing job, many workers are averse to taking up a similar job when it requires spending just ten minutes on caste-inconsistent tasks. This tendency is present even when the castes linked to the tasks rank relatively higher than the workers’ own castes, but is stronger when they rank lower. Nearly half of the workers are willing to

⁴⁹This loosely relates to the notion of taboo thoughts (Bénabou and Tirole (2006)).

forego ten times their daily wage—nearly a month’s wage income in the agricultural lean season—in order to avoid working on identity-violating tasks, claiming that they would never engage in such jobs regardless of wage offered.

These findings show that many workers deeply care about whether jobs are consistent with their identity. While workers may be individually making the best choices for themselves at the moment, the implications for society could be complicated. For instance, if some workers had great talent or skills for certain occupations but avoided them due to identity conflicts, it would lead to inefficient allocation of talent in the economy. In addition, the historical or social conditions that have created the links between jobs and identities may themselves be considered discriminatory or unjust. The current study is not aimed at directly estimating such costs, as the experiment involves limited sets of castes and jobs.⁵⁰

Nevertheless, the study sheds light on some fundamental ways in which identity concerns shape people’s decision-making. First, internal feelings or psychological concerns about identity are powerful motivators, even absent of other job disamenities. The workers in the experiment refuse job offers even when there is little to worry about skills, working environment, or employer discrimination. Also crucially, workers’ decisions are not influenced by whether or not their decisions are kept hidden from others, suggesting that few are motivated only by social image. In such a case, policy interventions that inform workers about others’ opinions on caste-conflicting jobs would have little effect on changing labor supply decisions. Second, identity concerns can make some workers entirely avoid certain occupations, even those associated with groups of higher social status. This could serve as a channel of poverty trap, if workers under harsh economic conditions still avoid well-paying jobs that conflict with their sense of identity. Finally, even among those sharing a single identity—well recognized both by self and others, as in the case of caste—there could be extremely divided reactions regarding specific behavioral prescriptions. It may be important to account for such heterogeneous preferences when building models of labor supply or considering policy interventions. To this end, future research that provides evidence on the processes through which jobs become associated with identities and on what causes some people to strongly respond to such connections could be informative.

⁵⁰Cassan, Keniston, and Kleineberg (2021) is one study that uses a structural model to quantify the effects of distortions from caste-based occupational choice on aggregate output. My study may help to refine such models by providing empirical evidence on the structure of identity costs.

References

- Akerlof, George A.**, “A Theory Of Social Custom, of which Unemployment May Be One Consequence,” *Quarterly Journal of Economics*, 1980, *94* (4), 749–775.
- **and Rachel E. Kranton**, “Economics and Identity,” *Quarterly Journal of Economics*, 2000, *115* (3), 715–753.
- Alesina, Alberto, Paola Giuliano, and Nathan Nunn**, “On the Origins of Gender Roles: Women and the Plough,” *Quarterly Journal of Economics*, 2013, *128* (2), 469–530.
- Atkin, David, Eve Colson-Sihra, and Moses Shayo**, “How Do We Choose Our Identity? A Revealed Preference Approach Using Food Consumption,” *Journal of Political Economy*, 2021, *129* (4), 1193–1251.
- Bašić, Zvonimir and Eugenio Verrina**, “Personal Norms—and Not Only Social Norms—Shape Economic Behavior,” *Working paper*, 2021.
- Bayly, Susan**, *Caste, Society and Politics in India from the Eighteenth Century to the Modern Age*, Cambridge: Cambridge University Press, 2001.
- Becker, Gordon M., Morris H. DeGroot, and Jacob Marschak**, “Measuring Utility by a Single-response Sequential Method,” *Behavioral Science*, 1964, *9* (3), 226–232.
- Bell, Alex, Raj Chetty, Xavier Jaravel, Neviana Petkova, and John Van Reenen**, “Who Becomes an Inventor in America? The Importance of Exposure to Innovation,” *Quarterly Journal of Economics*, 2019, *134* (2), 647–713.
- Bénabou, Roland and Jean Tirole**, “Incentives and Prosocial Behavior,” *American Economic Review*, 2006, *96* (5), 1652–1678.
- **and –**, “Identity, Morals, and Taboos: Beliefs as Assets,” *Quarterly Journal of Economics*, 2011, *126* (2), 805–855.
- Benjamin, Daniel J., James J. Choi, and A. Joshua Strickland**, “Social Identity and Preferences,” *American Economic Review*, 2010, *100* (4), 1913–1928.
- , – , **and Geoffrey Fisher**, “Religious Identity and Economic Behavior,” *American Economic Review*, 2016, *98* (4), 617–637.
- Bernheim, B. Douglas**, “A Theory of Conformity,” *Journal of Political Economy*, 1994, *102* (5), 841–877.

- Bettencourt, B. Ann, Kelly Charlton, Nancy Dorr, and Deborah L. Hume**, “Status Differences and In-group Bias: A Meta-analytic Examination of the Effects of Status Stability, Status Legitimacy, and Group Permeability.,” *Psychological Bulletin*, 2001, *127* (4), 520–542.
- Breza, Emily, Supreet Kaur, and Nandita Krishnaswamy**, “Scabs: The Social Suppression of Labor Supply,” *NBER Working Paper No. 25880*, 2019.
- , – , and **Yogita Shamdasani**, “The Morale Effects of Pay Inequality,” *Quarterly Journal of Economics*, 2018, *133* (2), 611–663.
- Burke, Peter J. and Jan E. Stets**, *Identity Theory*, New York: Oxford University Press, 2009.
- Burszty, Leonardo, Alessandra González, and David Yanagizawa-Drott**, “Misperceived Social Norms: Women Working Outside the Home in Saudi Arabia,” *American Economic Review*, 2020, *110* (10), 2997–3029.
- and **Robert Jensen**, “Social Image and Economic Behavior in the Field: Identifying, Understanding, and Shaping Social Pressure,” *Annual Review of Economics*, 2017, *9*, 131–153.
- , **Michael Callen, Bruno Ferman, Saad Gulzar, Ali Hasanain, and Noam Yuchtman**, “Political identity: Experimental evidence on anti-Americanism in Pakistan,” *Journal of the European Economic Association*, 2020, *18* (5), 2532–2560.
- Carpio, Lucía Del and Maria Guadalupe**, “More Women in Tech? Evidence from a Field Experiment Addressing Social Identity,” *Management Science*, 2021, *0* (0).
- Carvalho, Jean-Paul and Bary Pradelski**, “Identity and Underrepresentation: Interactions between Race and Gender,” *Working paper*, 2021.
- Cassan, Guilhem, Daniel Keniston, and Tatjana Kleineberg**, “A Division of Laborers: Identity and Efficiency in India,” *Working paper*, 2021.
- Cejka, Mary Ann and Alice H. Eagly**, “Gender-stereotypic Images of Occupations Correspond to the Sex Segregation of Employment,” *Personality and Social Psychology Bulletin*, 1999, *25* (4), 413–423.
- Charness, Gary and Yan Chen**, “Social Identity, Group Behavior, and Teams,” *Annual Review of Economics*, 2020, *12*, 691–713.
- Cohn, Alain, Ernst Fehr, and Michel André Maréchal**, “Business Culture and Dishonesty in the Banking Industry,” *Nature*, 2014, *516*, 86–89.

- Delfino, Alexia**, “Breaking Gender Barriers: Experimental Evidence on Men in Pink-Collar Jobs,” 2021. Working paper.
- Desai, Sonalde and Amaresh Dubey**, “Caste in 21st Century India: Competing Narratives,” *Economic and Political Weekly*, 2012, 46 (11), 40–49.
- Deshpande, Ashwini**, *The grammar of caste: Economic discrimination in contemporary India*, Oxford University Press, 2011.
- Dirks, Nicholas B.**, *Castes of Mind: Colonialism and the Making of Modern India*, Princeton: Princeton University Press, 2001.
- Erosa, Andrés, Luisa Fuster, Gueorgui Kambourov, and Richard Rogerson**, “Hours, occupations, and gender differences in labor market outcomes,” *NBER Working Paper No. w23636*, 2017.
- Falk, Armin**, “Facing Yourself—A Note on Self-image,” *Journal of Economic Behavior & Organization*, 2021, 186, 724–734.
- Forman, Chris, Anindya Ghose, and Batia Wiesenfeld**, “Examining the Relationship between Reviews and Sales: The Role of Reviewer Identity Disclosure in Electronic Markets,” *Quarterly Journal of Economics*, 2008, 119 (3), 291–313.
- Fudenberg, Drew, David K. Levine, and Zacharias Maniadis**, “On the Robustness of Anchoring Effects in WTP and WTA Experiments,” *American Economic Journal: Microeconomics*, 2012, 4 (2), 131–145.
- Gilboa, Itzhak, Stefania Minardi, and Fan Wang**, “Consumption of Values,” *Working paper*, 2021.
- Goraya, Sampreet Singh**, “How does Caste Affect Entrepreneurship? Birth vs Worth,” *Barcelona GSE Working Paper No. 1104*, 2019.
- Gottfredson, Linda S.**, “Circumscription and Compromise: A Developmental Theory of Occupational Aspirations,” *Journal of Counseling Psychology*, 1981, 28 (6), 545–579.
- Guérin, Isabelle, Bert D’espallier, and G. Venkatasubramanian**, “The Social Regulation of Markets: Why Microcredit Fails to Promote Jobs in Rural South India,” *Development and Change*, 2015, 46 (6), 1277–1301.
- Gupta, Dipankar**, *Interrogating Caste: Understanding Hierarchy and Difference in Indian Society*, New Delhi: Penguin Books, 2000.
- Hjort, Jonas**, “Ethnic Divisions and Production in Firms,” *Quarterly Journal of Economics*, 2014, 129 (4), 1899–1946.


- Hoff, Karla and Joseph E. Stiglitz**, “Striving for Balance in Economics: Towards a Theory of the Social Determination of Behavior,” *Journal of Economic Behavior & Organization*, 2016, 126, 25–57.
- **and Priyanka Pandey**, “Making Up People—The Effect of Identity on Performance in a Modernizing Society,” *Journal of Development Economics*, 2014, 106, 118–131.
- **, Mayuresh Kshetramade, and Ernst Fehr**, “Caste and Punishment: The Legacy of Caste Culture in Norm Enforcement,” *The Economic Journal*, 2011, 121 (556), F449–F475.
- Hogg, Michael A., Deborah J. Terry, and Katherine M. White**, “A Tale of Two Theories: A Critical Comparison of Identity Theory with Social Identity Theory,” *Social Psychology Quarterly*, 1995, 58, 255–69.
- Hsieh, Chang-Tai, Erik Hurst, Charles I. Jones, and Peter J. Klenow**, “The Allocation of Talent and US Economic Growth,” *Econometrica*, 2019, 87, 1439–1474.
- Iversen, Vegard and P. S. Raghavendra**, “What the Signboard Hides: Food, Caste and Employability in Small South Indian Eating Places,” *Contributions to Indian Sociology*, 2006, 40 (3), 311–341.
- Jakiela, Pamela and Owen Ozier**, “Does Africa Need a Rotten Kin Theorem? Experimental Evidence from Village Economies,” *Review of Economic Studies*, 2016, 83 (1), 231–268.
- Jayachandran, Seema**, “Social Norms as a Barrier to Women’s Employment in Developing Countries,” *IMF Economic Review*, 2021, 69 (3), 576–595.
- Jodhka, Surinder S**, *Caste in Contemporary India*, Routledge India, 2017.
- Jones, Stephen R. G.**, *The Economics of Conformism*, Oxford: Blackwell Pub, 1984.
- Karing, Anne**, “Social Signaling and Childhood Immunization: A Field Experiment in Sierra Leone,” *Working paper*, 2021.
- Kaur, Supreet**, “Nominal Wage Rigidity in Village Labor Markets,” *American Economic Review*, 2019, 109 (10), 3585–3616.
- Keane, Michael P. and Kenneth I. Wolpin**, “The Career Decisions of Young Men,” *Journal of Political Economy*, 1997, 105 (3), 473–522.
- Krishnaswamy, Nandita**, “Missing and Fired: Worker Absence, Labor Regulation, and Firm Outcomes,” *Working paper*, 2019.

- Lowe, Matt**, “Types of Contact: A Field Experiment on Collaborative and Adversarial Caste Integration,” *American Economic Review*, 2021, 111 (6), 1807–1844.
- Madheswaran, S. and Paul Attewell**, “Caste Discrimination in the Indian Urban Labour Market: Evidence from the National Sample Survey,” *Economic and Political Weekly*, 2007, 42 (41), 4146–4153.
- Mahar, Pauline M**, “A Ritual Pollution Scale for Ranking Hindu Castes,” *Sociometry*, 1960, 23 (3), 292–306.
- Marriott, McKim**, “Caste Ranking and Community Structure in Five Regions of India and Pakistan,” *Bulletin of the Deccan College Research Institute*, 1958, 19 (1/2), 31–105.
- Mosse, David**, “Caste and Development: Contemporary Perspectives on a Structure of Discrimination and Advantage,” *World Development*, 2018, 110, 422–436.
- Munshi, Kaivan**, “Caste and the Indian Economy,” *Journal of Economic Literature*, 2019, 57 (4), 781–834.
- **and Mark Rosenzweig**, “Traditional Institutions Meet the Modern World: Caste, Gender and Schooling Choice in a Globalizing Economy,” *American Economic Review*, 2006, 96 (4), 1225–1252.
- **and –**, “Networks and Misallocation: Insurance, Migration, and the Rural-Urban Wage Gap,” *American Economic Review*, 2016, 106 (1), 46–98.
- Neal, Derek**, “The Career Decisions of Young Men,” *Journal of Labor Economics*, 1999, 17 (2), 237–261.
- Owens, Timothy J., Dawn T. Robinson, and Lynn Smith-Lovin**, “Three Faces of Identity,” *Annual Review of Sociology*, 2010, 36, 477–499.
- Patton, Wendy and Mary McMahon**, *Career Development and Systems Theory: Connecting Theory and Practice*, Rotterdam: SensePublishers, 2014.
- Rao, Gautam**, “Familiarity Does Not Breed Contempt: Diversity, Discrimination and Generosity in Delhi Schools,” *American Economic Review*, 2019, 109 (3), 774–809.
- Risley, Herbert H.**, *The Tribes and Castes of Bengal*, Calcutta: Bengal Secretariat Press, 1892.
- , *The People of India*, Calcutta: Thacker, Spink and Co., 1908.
- Roy, Andrew Donald**, “Some Thoughts on the Distribution of Earnings,” *Oxford Economic Papers* 3, 1951, (2), 135–146.

- Shah, Alpa, Jens Lerche, Richard Axelby, Dalel Benbabaali, Brendan Donegan, Raj Jayaseelan, and Thakur Vikramditya**, *Ground Down by Growth: Tribe, Caste, Class and Inequality in 21st Century India*, Pluto Press, 2017.
- Shah, Ghanshyam, Harsh Mander, Sukhadeo Thorat, Satish Deshpande, and Amita Baviskar**, *Untouchability in rural India*, Sage Publications, 2006.
- Shayo, Moses**, “Social Identity and Economic Policy,” *Annual Review of Economics*, 2020, *12*, 355–389.
- Skorikov, Vladimir B. and Fred W. Vondracek**, “Occupational Identity,” in Seth J. Schwartz, Koen Luyckx, and Vivian L. Vignoles, eds., *Handbook of Identity Theory and Research*, New York: Springer, 2011, pp. 693–714.
- Stets, Jan E. and Richard T. Serpe**, “Identity Theory,” in John DeLamater and Amanda Ward, eds., *Handbook of Social Psychology*, 2 ed., Springer, 2013, chapter 2, pp. 31–60.
- Stryker, Sheldon**, *Symbolic Interactionism: A Social Structural Version*, Caldwell: Blackburn Press, 2002. Originally published in 1982.
- , “Integrating Emotion into Identity Theory,” in Jonathan H. Turner, ed., *Theory and Research on Human Emotions (Advances in Group Processes, Vol. 21)*, Bingley: Emerald Group Publishing Limited, 2004, pp. 1–23.
- **and Peter J. Burke**, “The Past, Present, and Future of an Identity Theory,” *Social Psychology Quarterly*, 2000, *63* (4), 284–297.
- Super, Donald E, Reuben Starishevsky, Norman Matlin, and Jan Pieere Jordaan**, *Career Development: Self-concept Theory*, New York: College Entrance Examination Board, 1963.
- Tajfel, Henri and John C. Turner**, “An Integrative Theory of Intergroup Conflict,” in William G. Austin and Stephen Worchel, eds., *The Social Psychology of Intergroup Relations*, Brooks-Cole, 1979, chapter 3, pp. 33–47.
- Thorat, Sukhadeo and Paul Attewell**, “The Legacy of Social Exclusion: A Correspondence Study of Job Discrimination in India,” *Economic and Political Weekly*, 2007, *42* (41), 4141–4145.
- Topel, Robert H. and Michael P. Ward**, “Job Mobility and the Careers of Young Men,” *Quarterly Journal of Economics*, 1992, *107* (2), 439–479.
- West, Candace and Don H. Zimmerman**, “Doing Gender,” *Gender & Society*, 1987, *1* (2), 125–151.

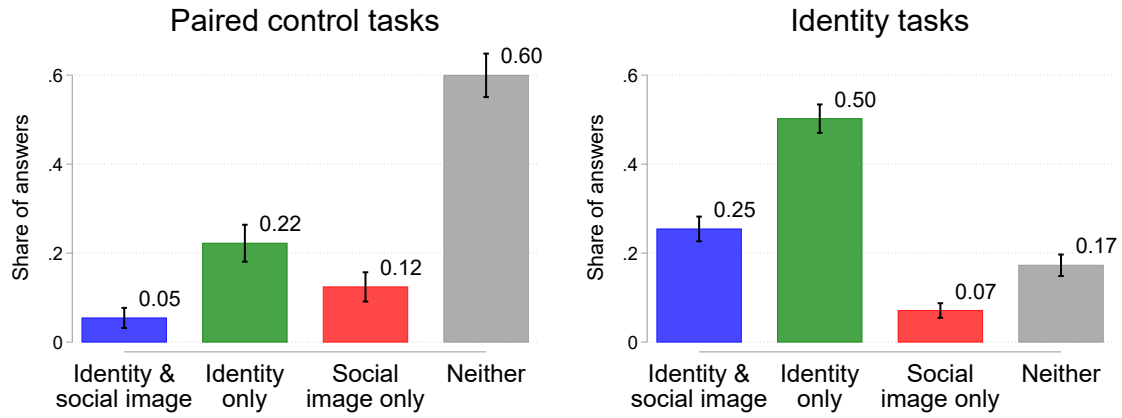
7 Figures and tables

Figure 1: Willingness to take up job offers and caste associations



Notes. This figure plots workers' willingness to take up job offers against the amount of time required on extra tasks. Average take-up rates are calculated separately by task category (paired control tasks on the left vs. identity tasks on the right), and by relative task status, as indicated by the three connected lines in each panel (same-ranked, higher, and lower). The relative task status is determined based on the rank scores in Table 1).


Figure 2: Reasons for turning down job offers




Notes. This figure plots workers' willingness to take up job offers against the amount of time required on extra tasks. The average take-up rates are calculated separately by task category (paired control tasks in the left panel vs. identity tasks in the right), and also by relative task status, as indicated by three lines in each panel (same-ranked, higher, and lower). The relative task status is determined based on the rank scores in Table 1).

Figure 3: Willingness to switch to working on extra tasks

Panel A: Take-up rates of switching offers



Panel B: Minimum additional wage required for working on extra tasks



Notes. Panel A plots workers' willingness to switch to working on extra tasks against the amount of time required on extra tasks. Average take-up rates are calculated separately by task category (paired control vs. identity tasks), and by the amount of additional wage offered for switching (Rs. 30 vs. 3000). Panel B plots the minimum additional wage at which workers agree to switch to working on extra tasks, separately by task category. Minimum additional wage of >3K means that workers refuse to switch regardless of wage, including when the maximum bonus of Rs. 3000 is offered.

Table 1: Caste ranking and associations with tasks

Caste	Rank score	Identity tasks (Caste-associated tasks)	Share associating task w. caste	Paired control tasks	Share associating task w. any SC
(1)	(2)	(3)	(4)	(5)	(6)
Kaibarta	1.48	-	-	-	-
Sundhi	2.07	-	-	-	-
Dhoba	3.71	Washing clothes	0.72	Washing farming tools	0
Kela	4.14	-	-	-	-
Mochi	4.59	Mending leather shoes	0.97	Mending grass mats	0.15
Pana	5.19	-	-	-	-
Hadi	6.60	Sweeping latrines	0.84	Sweeping animal sheds	0

Notes. This table summarizes the survey results on caste ranking and the associations between castes and tasks. The caste names in Column 1 are sorted according to the average ranks assigned to castes, reported in Column 2. In the remaining columns, caste and tasks that have connections are placed within the same rows, close to each other. Column 3 lists identity tasks which have strong caste associations and Column 4 reports the share of the survey participants who report these connections. Column 5 lists paired control tasks that involve similar skills as identity tasks. Column 6 shows the share of participants who report association between the paired control tasks with any Scheduled Caste (SC). A task's relative status is determined based on this table. For any worker, a task is considered a same-ranked task if it appears in the same row as the worker's caste, and otherwise a different task. If a task appears in a row above (or below) the worker's caste, the task is called a higher (or lower) task.

Table 2: Identity inconsistency and job offer take-up

	Willing to take up job offer				
	(1)	(2)	(3)	(4)	(5)
Identity \times Different	-0.231*** (0.049)	-0.231*** (0.042)	-0.234*** (0.044)	-0.242*** (0.044)	-0.237*** (0.044)
Identity \times Lower	-0.278*** (0.037)	-0.278*** (0.030)	-0.285*** (0.032)	-0.254*** (0.034)	-0.260*** (0.034)
Different tasks	-0.068* (0.035)	-0.068** (0.030)	-0.052* (0.031)	-0.050 (0.032)	-0.053* (0.032)
Lower tasks	0.083*** (0.030)	0.083*** (0.026)	0.075*** (0.028)	0.062** (0.029)	0.065** (0.029)
Mean: same-ranked tasks					
Identity tasks	0.722	0.722	0.731	0.731	0.731
Control tasks	0.717	0.717	0.728	0.728	0.728
Fixed effects included	Task, Caste	Task, Worker	Task, Worker	Task, Worker	Task, Worker
Answered follow-up survey			Yes	Yes	Yes
Demographic controls				Linear	Binary
Observations	20,160	20,160	17,632	17,632	17,632

Notes. This table shows how willingness to take up job offers varies with predicted presence of identity violations. Each column reports the result of an OLS regression of take-up on variables related to task category and relative task status, following Equation 4. All regressions control for task-caste-specific linear time trends, task fixed effects, as well as caste or worker fixed effects. In Column 3-5, the sample is restricted to those who answered the follow-up survey. The specification in Column 4 additionally controls for interactions between task-dummies and three demographic variables: age, years of education, and wealth PCA score. Column 5 instead controls for the interactions with binaries indicating higher (above median) age, education, and wealth. Standard errors are clustered at the worker times task level and shown in parentheses. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 3: Role of social image concerns

	Willing to take up job offer				
	(1)	(2)	(3)	(4)	(5)
Identity \times Different	-0.221*** (0.071)	-0.221*** (0.072)	-0.203*** (0.073)	-0.203*** (0.071)	-0.201*** (0.072)
Identity \times Lower	-0.292*** (0.050)	-0.292*** (0.051)	-0.307*** (0.054)	-0.277*** (0.054)	-0.281*** (0.053)
Different	-0.073** (0.034)	-0.067** (0.034)	-0.066** (0.033)	-0.069** (0.033)	-0.069** (0.032)
Lower	0.113*** (0.035)	0.103*** (0.034)	0.095** (0.037)	0.084** (0.037)	0.086** (0.037)
Public \times Identity \times Different	-0.020 (0.082)	-0.020 (0.084)	-0.060 (0.088)	-0.075 (0.090)	-0.069 (0.088)
Public \times Identity \times Lower	0.026 (0.058)	0.026 (0.059)	0.041 (0.062)	0.044 (0.059)	0.041 (0.060)
Public \times Different	0.010 (0.059)	-0.001 (0.048)	0.027 (0.048)	0.037 (0.049)	0.031 (0.048)
Public \times Lower	-0.058 (0.046)	-0.038 (0.037)	-0.038 (0.038)	-0.040 (0.036)	-0.038 (0.036)
Fixed effects included	Task, Caste	Task, Worker	Task, Worker	Task, Worker	Task, Worker
Answered follow-up survey			Yes	Yes	Yes
Demographic controls				Linear	Binary
Observations	20,160	20,160	17,632	17,632	17,632

Notes. This table shows how willingness to take up job offers varies with predicted presence of identity violations, depending on whether worker decisions are publicized. The regressions are similar to those in Table 2, but the key covariates are interacted with an indicator for the public condition. All regressions control for task-caste-specific linear time trends, task fixed effects, as well as caste or worker fixed effects. The coefficients on *Public \times Identity*, *Public \times Paired control*, and *Public* are not displayed. Standard errors are clustered at the village level and shown in parentheses. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 4: Caste inconsistency and refusal of all offers involving a task

	Refuse all offers regardless of bonus					
	(1)	(2)	(3)	(4)	(5)	(6)
Identity tasks	0.277*** (0.033)	0.277*** (0.024)	0.313*** (0.048)	0.272*** (0.049)	0.272*** (0.053)	0.318*** (0.090)
Paired tasks	0.075** (0.032)	0.075** (0.032)	0.082 (0.062)	0.031* (0.014)	0.031* (0.015)	0.036 (0.036)
Public × Identity				0.010 (0.056)	0.010 (0.061)	-0.008 (0.066)
Public × Paired				0.091 (0.056)	0.091 (0.060)	0.088 (0.071)
Fixed effects included	Caste	Worker	Worker	Caste	Worker	Worker
Answered follow-up survey			Yes			Yes
Demographic controls			Binary			Binary
Observations	742	742	735	742	742	735

Notes. This table shows how much more workers are likely to refuse the offers involving identity tasks compared to control tasks, using the supplementary experiment data. The dependent variable is a worker-task-level indicator for whether worker refuses all offers regardless of time amount and bonus wage. Paired tasks refer to both identity and paired control tasks, so that the omitted category is the pure control task of moving bricks. The dependent variable mean for the pure control task is 7 percent. Standard errors are clustered at the worker times task level in Columns 1-3 or at the village level in Columns 4-6.

For Online Publication

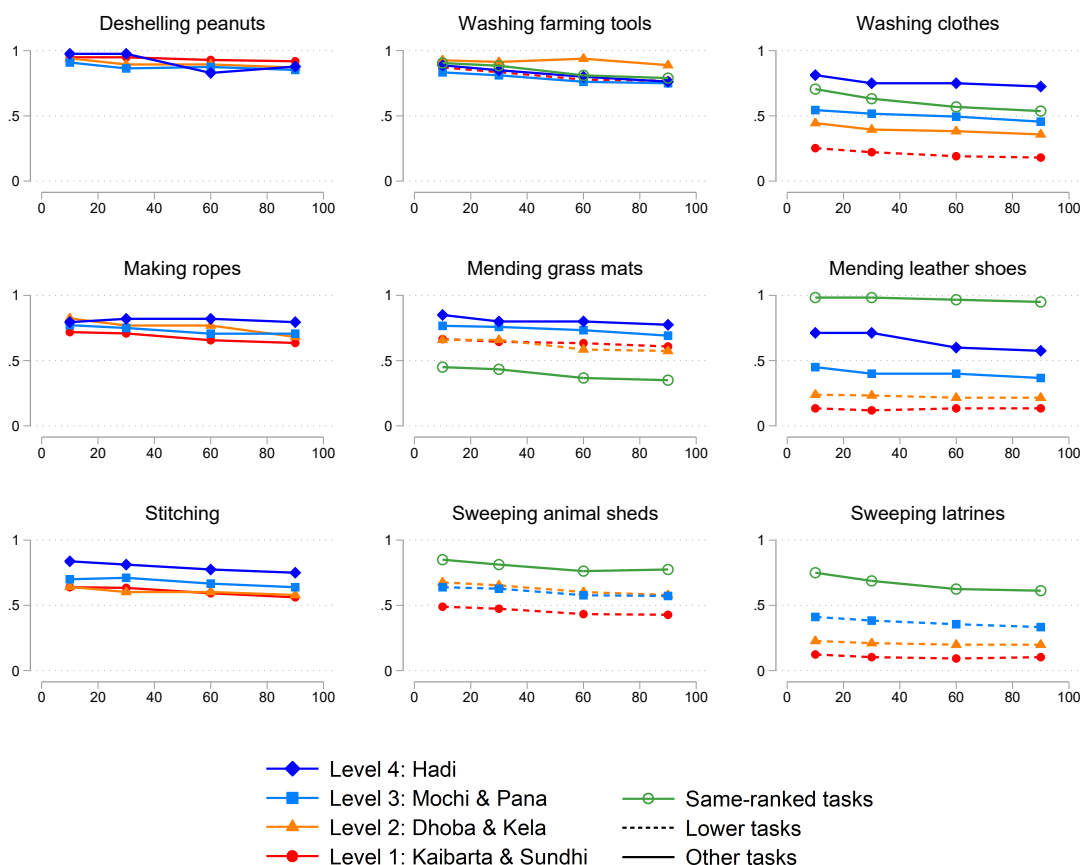
A Appendix figures and tables

Figure A1: Descriptive pictures of tasks



Notes. During the job take-up exercise, workers were provided descriptive pictures of the extra tasks, such as these in this figure. The examples here depict washing clothes, sweeping animal sheds, mending grass mats, and mending leather shoes.

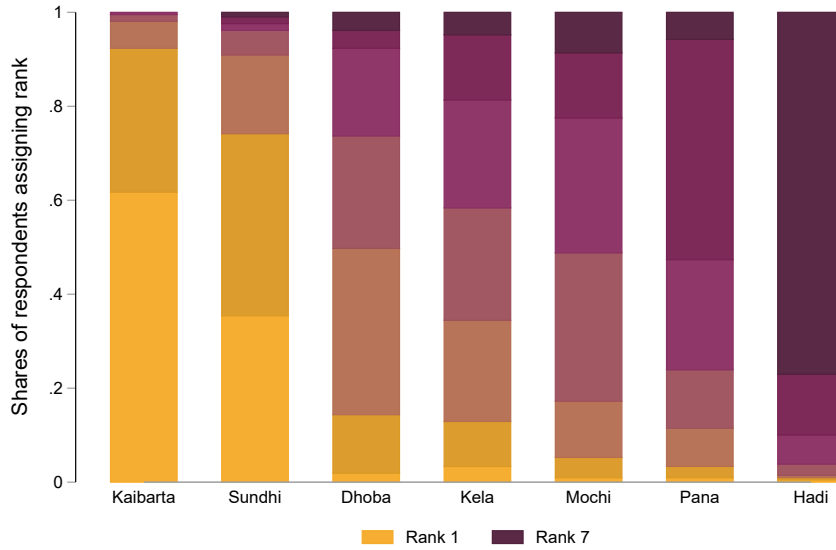
Figure A2: Willingness to take-up by task




Notes. This figure plots workers' willingness to take up job offers against the amount of time required on extra tasks, with each panel plotting the results for a specific task. The left column is for pure control tasks, the center column for paired control tasks, and the right column for identity tasks. Castes are grouped into four different levels based on caste ranking: level 4 contains Hadi, which is the lowest ranked caste, level 3 contains Mochi and Pana, level 2 contains Dhoba and Kela, and level 1 contains Kaibarta and Sundhi. In the left column, the four connected lines concern different caste levels. The plots in the center and right columns have two modifications. First, the results for the same-ranked tasks are graphed separately as the green lines with hollow circular markers. Second, when tasks are considered lower, the graphs are drawn with dashed lines instead of solid lines. For example, in the panels for washing farming tools and washing clothes, the green lines are for the Dhoba caste and the orange lines with triangular markers are only for the Kela caste; and since Kaibarta and Sundhi rank higher than Dhoda, their take-up rates are plotted with dashed lines.

Figure A3: Ranks assigned to castes

Panel A: Variation in reported ranks




Panel B: Comparison against castes with task associations



Notes. This figure shows provides additional information from the Rank Survey. Panel A plots the distribution of ranks assigned to the seven castes involved in the experiment. Lighter colors indicate higher ranks. Panel B shows how the ten castes used in analyzing experience levels in Section 4.4.1 rank against the three castes with task associations. The bars indicate the shares of respondents that rank a given caste lower than Dhoba, Mochi, and Hadi, respectively.

Figure A4: Caste-sensitive opinions of oneself vs. others



Notes. This figure plots the share of Task Survey participants who express caste-sensitive opinions, either of their own or of their friends and neighbors. There were four vignette questions describing characters violating various caste norms, listed as Q1-Q4 in Appendix Section B.4. Randomly selected half of the participants were asked in their personal view whether they approve of the characters' actions. The rest were asked whether their friends and neighbors would approve of such actions. The figure shows the share of participants who express disapproval for not following caste norms, along with 95% confidence intervals.

Table A1: Consistency of caste rank scores

	Rank assigned to caste			
	(1)	(2)	(3)	(4)
Sundhi	0.573*** (0.104)	0.518*** (0.167)	0.705*** (0.183)	0.491** (0.196)
Dhoba	2.234*** (0.100)	2.157*** (0.199)	2.296*** (0.173)	2.250*** (0.149)
Kela	2.620*** (0.110)	2.573*** (0.176)	2.666*** (0.184)	2.619*** (0.218)
Mochi	3.076*** (0.107)	2.983*** (0.191)	3.186*** (0.182)	3.055*** (0.187)
Pana	3.703*** (0.093)	3.714*** (0.160)	3.746*** (0.165)	3.647*** (0.164)
Hadi	5.120*** (0.087)	5.047*** (0.157)	5.309*** (0.123)	5.000*** (0.174)
Own caste	-0.766*** (0.111)	-0.730*** (0.187)	-0.937*** (0.207)	-0.634*** (0.181)
Instruction type	All types	General	Food-related	Water-related
Mean rank for Kaibarta	1.48	1.53	1.42	1.50
P-val: equality of ranks				
Sundhi = Dhoba	0.00	0.00	0.00	0.00
Dhoba = Kela	0.01	0.14	0.17	0.16
Kela = Mochi	0.00	0.12	0.02	0.10
Mochi = Pana	0.00	0.00	0.02	0.01
Pana = Hadi	0.00	0.00	0.00	0.00
Observations	1,463	490	497	476

Notes. This table shows how respondents assigned ranks to seven experimental castes during the Rank Survey. All columns show results from the OLS regressions of assigned ranks on caste-level dummies, as well as an indicator for whether the ranked caste coincides with the respondent's caste. Column 1 show the results using all observations and Columns 2-4 show the results by instruction type. Standard errors are clustered at the respondent level.

Table A2: Task associations and experiences

	Caste association		Gender association			Previously performed			
	Any caste (1)	Any SC (2)	Men (3)	Women (4)	Both (5)	In own HH (6)	Outside HH (7)	For wage (8)	Ever (9)
Washing clothes	0.74	0.73	0.01	0.19	0.79	0.97	0.00	0.02	0.98
Washing farming tools	0.04	0.00	0.70	0.01	0.27	0.84	0.01	0.11	0.89
Mending leather shoes	0.99	0.99	0.86	0.00	0.13	0.18	0.00	0.00	0.18
Mending grass mats	0.28	0.15	0.32	0.05	0.39	0.10	0.01	0.01	0.10
Sweeping latrines	0.85	0.85	0.51	0.08	0.38	0.51	0.01	0.02	0.51
Sweeping animal sheds	0.04	0.00	0.10	0.17	0.73	0.80	0.01	0.01	0.81
Making paper bags	0.09	0.01	0.05	0.15	0.65	0.10	0.00	0.00	0.10
Deshelling peanuts	0.03	0.01	0.05	0.15	0.66	0.71	0.01	0.05	0.74
Making ropes	0.07	0.03	0.67	0.01	0.27	0.31	0.01	0.01	0.33
Stitching	0.05	0.01	0.06	0.08	0.85	0.58	0.00	0.01	0.58
Making leaf mats	0.83	0.75	0.04	0.45	0.45	0.02	0.00	0.00	0.02
Making leaf brooms	0.73	0.67	0.15	0.12	0.69	0.15	0.00	0.02	0.15
Making bamboo mats	0.71	0.67	0.47	0.04	0.47	0.42	0.01	0.07	0.45
Making stick brooms	0.43	0.40	0.13	0.12	0.69	0.40	0.01	0.01	0.41
Making incense sticks	0.03	0.01	0.03	0.41	0.51	0.03	0.01	0.06	0.09
Making candle wicks	0.13	0.00	0.01	0.52	0.37	0.49	0.03	0.01	0.51

Notes. This table summarizes the results from the Task Survey, pertaining to the caste and gender associations of tasks and respondents' prior experiences with tasks. Columns 1-2 report the shares of participants who associate the tasks with any caste or with any scheduled caste. Columns 3-5 show the share of respondents who associate the tasks with men, women or both genders. Columns 7-9 show the shares of respondents who have previously performed the task in own household, performed for friends or neighbors without wage, performed for wage, or any of the above. Participants can report multiple experience levels as applicable. The bottom panel shows the results for additional tasks which are not part of the experiment due to their strong associations with women or other caste groups.

Table A3: Summary of worker characteristics

	Mean for Level 4	Diff. for Level 3	Diff. for Level 2	Diff. for Level 1
Age	37.440 [9.365]	-0.641 (1.268)	3.163** (1.316)	5.013*** (1.258)
Years of education	4.707 [3.490]	0.268 (0.475)	-0.508 (0.500)	1.442*** (0.477)
Able to read	0.653 [0.479]	0.083 (0.065)	-0.096 (0.068)	0.191*** (0.062)
Family size	5.053 [1.692]	0.337 (0.242)	0.049 (0.263)	-0.171 (0.234)
Share of working members	0.373 [0.184]	-0.102*** (0.025)	0.002 (0.026)	-0.033 (0.025)
Mud house	0.387 [0.490]	-0.123* (0.066)	-0.034 (0.068)	-0.169*** (0.065)
Semi-mud house	0.320 [0.470]	-0.075 (0.064)	-0.153** (0.062)	-0.177*** (0.061)
Owns land	0.373 [0.487]	-0.002 (0.068)	0.031 (0.069)	0.335*** (0.067)
Land size in acres	0.365 [0.956]	-0.087 (0.124)	-0.089 (0.123)	0.345*** (0.133)
Last month income in Rs.	5,350 [2,474]	1,794*** (494.67)	-29.359 (402.31)	856.25* (446.46)
Paid work days last week	2.813 [2.246]	-0.719** (0.304)	0.046 (0.301)	-0.559* (0.307)
Number of assets owned	3.307 [1.602]	0.096 (0.220)	-0.287 (0.223)	0.861*** (0.212)
Wealth PCA score	-0.327 [1.438]	0.209 (0.199)	-0.211 (0.211)	1.139*** (0.196)
Number of caste sensitive views	3.760 [1.800]	-0.181 (0.249)	-0.010 (0.251)	0.656*** (0.247)

Notes. This table summarizes the work-level data on workers' age, education, wealth, and caste sensitivity, gathered from the follow-up survey. Column 1 shows the variable means for the lowest ranked caste (Hadi). Columns 2-4 show the coefficients and p-values from regressing each variable on the indicator variables for level 3 (Mochi and Pana), level 2 (Dhoba and Kela), and level 1 (Kaibarta and Sundhi). Standard deviations are reported in brackets and robust standard errors in parentheses.

Table A4: Job take-up results with alternate specifications

	Willing to take up job offer					
	Progressively add more controls				Restrict sample	
	(1)	(2)	(3)	(4)	(5)	(6)
Identity \times Different	-0.230*** (0.048)	-0.223*** (0.048)	-0.226*** (0.048)	-0.226*** (0.048)	-0.211*** (0.049)	-0.247*** (0.045)
Identity \times Lower	-0.285*** (0.034)	-0.281*** (0.034)	-0.280*** (0.034)	-0.280*** (0.034)	-0.267*** (0.035)	-0.300*** (0.032)
Different	-0.083** (0.034)	-0.085** (0.034)	-0.085** (0.034)	-0.085** (0.034)	-0.067* (0.035)	-0.067** (0.033)
Lower	0.096*** (0.029)	0.095*** (0.029)	0.095*** (0.029)	0.095*** (0.029)	0.082*** (0.030)	0.095*** (0.027)
Controls added	Quadratic time trends	Surveyor FE	Question order FE	Choice set FE		
Excluded from sample					Low com- prehension	Choice in- consistency
Observations	20,160	20,160	20,160	20,160	15,104	16,736

Notes. The regressions in this table are similar to those in Table 2, but include additional control variables or have different sample restrictions, as specified in the table footer. Column 1 controls for task-caste-specific quadratic time trends. The choice exercise involved 12 surveyors, 4 different ways in which tasks are presented, time requirements randomly presented in ascending or descending order, and one of two pure control tasks randomly being presented. Columns 2-4 control for the indicators related to these variations interacted with the dummy for identity tasks. During the practice and job take-up choice exercises, surveyors asked seven comprehension questions, and if worker did not answer correctly, explained the relevant procedure again up to three more times. On their first attempts, the median worker answered 6 questions correctly. Column 5 excludes 25% of workers who answered 5 or fewer questions correctly. Choice inconsistency refers when a worker refuses an offer involving a particular task and also accepts another offer involving a longer amount of time on the same task. Column 6 excludes 17% of workers who exhibit at least one case of choice inconsistency across all offers.

Table A5: Job take-up results using alternate rankings

	Willing to take up job offer					
	Registered ranking			Partially corrected ranking		
	(1)	(2)	(3)	(4)	(5)	(6)
Identity \times Different	-0.337*** (0.045)	-0.337*** (0.040)	-0.337*** (0.042)	-0.265*** (0.048)	-0.265*** (0.042)	-0.272*** (0.044)
Identity \times Lower	-0.159*** (0.036)	-0.159*** (0.029)	-0.146*** (0.033)	-0.242*** (0.036)	-0.242*** (0.030)	-0.221*** (0.034)
Different	-0.021 (0.034)	-0.021 (0.029)	-0.017 (0.030)	-0.053 (0.036)	-0.053* (0.030)	-0.037 (0.032)
Lower	0.029 (0.030)	0.029 (0.026)	0.031 (0.029)	0.068** (0.031)	0.068** (0.027)	0.049 (0.031)
Fixed effects included	Task, Caste	Task, Worker	Task, Worker	Task, Caste	Task, Worker	Task, Worker
Answered follow-up survey			Yes			Yes
Demographic controls			Linear			Linear
Observations	20,160	20,160	17,632	20,160	20,160	17,632

Notes. This table shows how willingness to take up job offers vary with predicted presence of identity violations, using alternate caste rankings. Columns 1-3 use pre-registered ranking, which mis-specifies the ranking for Kaibarta and Kela, as explained in Section 4.2.1. Columns 4-6 use partially corrected ranking that places Kaibarta above Dhoba, and Kela above Hadi. The regressions are the same as Columns 1, 2, and 4 in Table 2. All regressions control for task-caste-specific linear time trends, task fixed effects, as well as caste or worker fixed effects. Standard errors are clustered at the worker times task level.

Table A6: Job take-up results excluding one caste at a time

	Willing to take up job offer						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Identity × Different	-0.226*** (0.043)	-0.250*** (0.042)	-0.327*** (0.057)	-0.200*** (0.043)	-0.056 (0.045)	-0.180*** (0.045)	-0.406*** (0.050)
Identity × Lower	-0.290*** (0.035)	-0.229*** (0.034)	-0.260*** (0.032)	-0.305*** (0.032)	-0.283*** (0.031)	-0.351*** (0.036)	-0.254*** (0.036)
Different	-0.065** (0.031)	-0.053* (0.030)	-0.025 (0.040)	-0.076** (0.031)	-0.124*** (0.032)	-0.081** (0.032)	-0.039 (0.038)
Lower	0.083*** (0.030)	0.059** (0.029)	0.086*** (0.027)	0.086*** (0.029)	0.075*** (0.027)	0.092*** (0.031)	0.105*** (0.029)
Excluded caste group	Kaibarta	Sundhi	Dhoba	Kela	Mochi	Pana	Hadi
Observations	16,576	17,536	17,120	17,568	18,240	16,320	17,600

Notes. The table shows that the main results are robust to dropping any one caste. The regressions are analogous to that in Table 2 Column 2. The table footer indicates which caste groups is excluded in each regression. Standard errors are clustered at the worker times task level.

Table A7: Completion rates of actually selected offers

	Willing to take up job offer		Completion	
	Any offer involving task (1)	Randomly selected offer (2)	One-day job (3)	Follow-up survey (4)
Identity \times Different	-0.236*** (0.047)	-0.284* (0.152)	-0.491*** (0.168)	-0.026 (0.100)
Identity \times Lower	-0.282*** (0.036)	-0.270** (0.105)	-0.247** (0.115)	-0.070 (0.070)
Different	-0.068** (0.033)	-0.076 (0.127)	0.076 (0.138)	0.023 (0.088)
Lower	0.093*** (0.028)	0.086 (0.093)	0.132 (0.104)	-0.011 (0.063)
Mean: same-ranked tasks				
Identity tasks	0.772	0.857	0.750	0.964
Control tasks	0.770	0.737	0.316	0.895
Observations	5,040	629	629	629

Notes. This table shows how the results change when different outcome measures are used. The dependent variables are indicators for the following: whether worker accepts any of the offers involving the task (Column 1); whether worker accepts the randomly selected offer (Column 2) whether worker completed the one-day job from the randomly selected offer (Column 3), and whether worker completed the follow-up survey (Column 4). All regressions control for task and caste fixed effects. Column 1 outcome is constructed at the worker-task level, and the remaining outcomes are at the worker level. Hence, standard errors are also clustered at these respective levels.

Table A8: Heterogeneity in job offer take-up

	Willing to take up job offer		
	Caste sensitive (1)	Older (2)	Less educated (3)
Identity \times Different	-0.251*** (0.051)	-0.244*** (0.054)	-0.239*** (0.059)
Identity \times Lower	-0.213*** (0.038)	-0.208*** (0.040)	-0.211*** (0.043)
Different	-0.034 (0.036)	-0.056 (0.038)	-0.044 (0.042)
Lower	0.050 (0.032)	0.046 (0.033)	0.059 (0.037)
Traditional \times Different \times Identity	0.039 (0.071)	0.001 (0.073)	0.010 (0.069)
Traditional \times Lower \times Identity	-0.173*** (0.049)	-0.104** (0.049)	-0.142*** (0.048)
Traditional \times Different	-0.049 (0.048)	0.030 (0.048)	-0.016 (0.048)
Traditional \times Lower	0.060 (0.038)	0.021 (0.038)	0.039 (0.038)
Observations	17632	17632	17632

Notes. This table shows how willingness to take up job offers varies with predicted presence of identity violations, depending on whether workers are expected to hold more traditional opinions. The regressions are similar to those in Table 3, with the key covariates interacted with different proxies for traditional views (instead of the indicator for the public condition). *Caste sensitive* indicates that worker expressed stronger support for observing caste norms in the follow-up survey, i.e. the number of caste sensitive views is greater than the median value of four. *Older* means worker's age is greater than the median and *Less educated* means worker's years of education is not greater than the median. All regressions control for task-caste-specific linear time trends, task fixed effects, as well as worker fixed effects. Standard errors are clustered at the worker times task level.

Table A9: Balance of worker characteristics

	Main experiment data		Supplementary data	
	Mean for Private	Diff. for Public	Mean for Private	Diff. for Public
Age	40.267 [8.887]	-1.243 (1.247)	38.660 [8.976]	3.340 (2.106)
Years of education	4.996 [3.480]	0.125 (0.396)	5.849 [3.682]	-1.195** (0.509)
Can read Odiya	0.718 [0.451]	-0.022 (0.050)	0.717 [0.455]	-0.044 (0.056)
Family size	5.092 [1.744]	0.043 (0.185)	5.302 [1.814]	-0.494 (0.330)
Share of working members	0.343 [0.175]	-0.017 (0.020)	0.387 [0.190]	-0.035 (0.060)
Mud house	0.286 [0.453]	0.011 (0.054)	0.415 [0.497]	-0.184* (0.096)
Semi-mud house	0.187 [0.391]	0.031 (0.046)	0.151 [0.361]	-0.016 (0.039)
Owens land	0.527 [0.500]	-0.091 (0.067)	0.642 [0.484]	-0.045 (0.148)
Land size in acres	0.375 [0.723]	0.076 (0.117)	0.445 [0.587]	0.036 (0.218)
Last month income in Rs.	5,934 [4,089]	328.08 (482.62)	10000 [7,093]	-3,800** (1,378)
Paid work days last week	2.363 [2.010]	0.177 (0.280)	3.170 [2.268]	-0.054 (0.437)
Number of assets owned	3.599 [1.515]	-0.181 (0.204)	3.528 [1.324]	-0.182 (0.310)
Wealth PCA score	0.084 [1.451]	-0.149 (0.200)	0.029 [1.162]	-0.102 (0.238)
Number of caste sensitive views	3.756 [1.729]	0.268 (0.202)	3.019 [1.886]	0.308 (0.406)

Notes. This table checks the balance of worker characteristics across randomized privacy conditions. Column 1 shows the variable means for the private treatment group. Columns 2 shows the coefficients and p-values of a regression at the worker-level of each variable on the indicator for the public treatment. Column 3-4 are similar but use the supplementary experiment data. Standard deviations are reported in brackets. Standard errors are clustered at the village level and shown in parentheses.

Table A10: Social image results with alternate clustering of standard errors

	Willing to take up job offer				
	(1)	(2)	(3)	(4)	(5)
Identity \times Different	-0.221*** (0.064)	-0.221*** (0.053)	-0.203*** (0.055)	-0.203*** (0.055)	-0.201*** (0.055)
Identity \times Lower	-0.292*** (0.047)	-0.292*** (0.038)	-0.307*** (0.041)	-0.277*** (0.043)	-0.281*** (0.042)
Different	-0.073 (0.045)	-0.067* (0.037)	-0.066* (0.038)	-0.069* (0.038)	-0.069* (0.038)
Lower	0.113*** (0.035)	0.103*** (0.031)	0.095*** (0.034)	0.084** (0.035)	0.086** (0.035)
Public \times Identity \times Different	-0.020 (0.086)	-0.020 (0.065)	-0.060 (0.068)	-0.075 (0.068)	-0.069 (0.068)
Public \times Identity \times Lower	0.026 (0.060)	0.026 (0.044)	0.041 (0.047)	0.044 (0.047)	0.041 (0.047)
Public \times Different	0.010 (0.058)	-0.001 (0.045)	0.027 (0.046)	0.037 (0.046)	0.031 (0.047)
Public \times Lower	-0.058 (0.040)	-0.038 (0.035)	-0.038 (0.037)	-0.040 (0.037)	-0.038 (0.037)
Fixed effects included	Task, Caste	Task, Worker	Task, Worker	Task, Worker	Task, Worker
Answered follow-up survey			Yes	Yes	Yes
Demographic controls				Linear	Binary
Observations	20,160	20,160	17,632	17,632	17,632

Notes. This table shows how willingness to take up job offers vary with predicted presence of identity violations, depending on whether worker decisions are publicized. The regressions are similar to those in Table 3, but here standard errors are clustered at the worker times task level.

Table A11: Experiences with tasks

	In own household		Outside household		For wage		Ever performed	
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Different \times Identity	-0.042 (0.141)	-0.041 (0.139)	-0.053 (0.047)	-0.053 (0.046)	-0.277** (0.130)	-0.277** (0.127)	-0.228** (0.116)	-0.227* (0.120)
Lower \times Identity	0.036 (0.064)	0.038 (0.064)	-0.011 (0.019)	-0.011 (0.019)	0.036 (0.035)	0.036 (0.036)	0.055 (0.062)	0.057 (0.062)
Different tasks	-0.101 (0.101)	-0.105 (0.100)	0.006 (0.015)	0.007 (0.016)	-0.020 (0.079)	-0.020 (0.079)	-0.064 (0.099)	-0.068 (0.097)
Lower tasks	0.024 (0.059)	0.027 (0.059)	0.006 (0.020)	0.006 (0.019)	-0.012 (0.034)	-0.012 (0.034)	0.006 (0.058)	0.010 (0.058)
Mean for same-ranked tasks								
Identity tasks	0.800	0.800	0.050	0.050	0.300	0.300	0.950	0.950
Control tasks	0.850	0.850	0.000	0.000	0.100	0.100	0.850	0.850
Fixed effects included	Task, Caste	Task, Worker	Task, Caste	Task, Worker	Task, Caste	Task, Worker	Task, Caste	Task, Worker
Observations	999	999	999	999	999	999	999	999

Notes. This table shows how survey participant's experience with tasks vary with the task's association with castes. During the Task Survey, participants described the extent to which they have performed the tasks listed in Table 1. The outcomes are indicators for whether participant has previously performed the task in own household in Columns 1-2, performed for friends and neighbors without wage in Columns 3-4, performed for wage in Columns 5-6, and performed any of the above in Columns 7-8. The table presents OLS regression estimates of how experience outcomes vary with task category (identity, paired control) and relative task status (different, lower). The omitted category is same-ranked tasks, and the dependent variable means for same-ranked tasks are reported in the table footer. All regressions additionally control for task and caste/worker fixed effects. Standard errors are clustered at the worker times task level.

Table A12: Number of refusals for each task within worker-subgroups

	Refuse any identity task		Refuse all identity tasks	
	Refuse 0 (1)	Refuse 1+ (2)	Refuse 2- (3)	Refuse 3 (4)
A. Control tasks				
Moving bricks	1	1	1	1
Washing farming tools	0	5	0	5
Mending grass mats	1	3	1	3
Sweeping animal sheds	0	18	2	16
B. Identity tasks				
Washing clothes	0	32	4	28
Mending leather shoes	0	36	8	28
Sweeping latrines	0	47	19	28
Total	46	53	71	28

Notes. This table shows how the decision to refuse a task, i.e. turn down all offers involving the task regardless of time amount and bonus, correlates within workers. Workers are divided into two subgroups based on whether they refuse any identity task (Column 1 vs. 2) or whether they refuse all identity tasks (Column 3 vs. 4). The table reports how many workers within the subgroups refuse each of the tasks. Seven workers who turn down all tasks are omitted from these reports.

Table A13: Role of experience and comprehension

	Refuse all offers regardless of bonus				
	(1)	(2)	(3)	(4)	(5)
Identity tasks	0.259*** (0.034)	0.269*** (0.024)	0.296*** (0.047)	0.297*** (0.058)	0.303*** (0.051)
Paired tasks	-0.054 (0.058)	0.011 (0.057)	-0.002 (0.079)	-0.025 (0.092)	-0.004 (0.083)
Performed in own HH	-0.138*** (0.031)	-0.071*** (0.025)	-0.114*** (0.029)	-0.095*** (0.033)	-0.114*** (0.032)
Performed outside HH	-0.082 (0.051)	-0.004 (0.060)	-0.008 (0.065)	0.020 (0.072)	0.018 (0.069)
Performed for wage	-0.139*** (0.053)	-0.077 (0.055)	-0.088 (0.055)	-0.064 (0.062)	-0.038 (0.057)
Fixed effects included	Caste	Worker	Worker	Worker	Worker
Answered follow-up survey	Yes	Yes	Yes	Yes	Yes
Demographic controls			Binary	Binary	Binary
Excluded from sample				Low comprehension	Choice inconsistency
Observations	735	735	735	511	644

Notes. This table shows that the results on offer refusal is robust to controlling for workers' experience and comprehension. The supplemental follow-up survey contains information on whether workers have previously performed tasks in own household, performed for friends and neighbors without wage, and/or performed for wage. Columns 1-3 are similar to Appendix Table 4 Columns 1-3, but the regressions additionally control for the indicators of these different experience levels. Column 4 excludes those who scored fewer comprehension questions than the median worker and Column 5 excludes those who exhibit any choice inconsistency. Standard errors are clustered at the worker times task level.

Table A14: Predicting which workers have identity concerns

	Refuse any identity task			Refuse all identity tasks		
	(1)	(2)	(3)	(4)	(5)	(6)
Age	0.010 (0.007)	0.007 (0.007)	0.009 (0.007)	0.005 (0.006)	0.003 (0.006)	0.004 (0.006)
Years of education	0.003 (0.019)	0.006 (0.019)	0.010 (0.017)	-0.000 (0.016)	0.002 (0.016)	0.004 (0.015)
Share of working members	-0.174 (0.257)	-0.112 (0.253)	-0.062 (0.244)	-0.151 (0.261)	-0.098 (0.261)	-0.070 (0.251)
Mud house	0.076 (0.114)	0.082 (0.114)	0.087 (0.110)	0.025 (0.106)	0.030 (0.106)	0.033 (0.105)
Semi-mud house	0.083 (0.164)	0.072 (0.156)	0.118 (0.156)	0.002 (0.156)	-0.008 (0.155)	0.019 (0.156)
Owns land	-0.014 (0.118)	0.004 (0.119)	0.070 (0.120)	-0.085 (0.121)	-0.071 (0.121)	-0.033 (0.126)
Land size in acres	-0.074 (0.089)	-0.063 (0.086)	-0.110 (0.082)	-0.110 (0.074)	-0.101 (0.074)	-0.127 (0.079)
Last month income	-0.038** (0.017)	-0.036** (0.017)	-0.041* (0.024)	-0.015 (0.036)	-0.013 (0.038)	-0.016 (0.046)
Paid work days last week	-0.029 (0.021)	-0.028 (0.021)	-0.021 (0.021)	-0.023 (0.019)	-0.023 (0.020)	-0.019 (0.020)
Number of assets owned	-0.053 (0.039)	-0.030 (0.044)	-0.026 (0.042)	0.010 (0.040)	0.030 (0.042)	0.032 (0.042)
Kaibarta caste	0.158 (0.105)	0.177* (0.106)	0.170 (0.103)	0.183 (0.111)	0.199* (0.111)	0.195* (0.110)
Comprehension score		-0.064 (0.041)	-0.050 (0.039)		-0.055 (0.038)	-0.047 (0.037)
Number of caste sensitive views			0.067** (0.026)			0.038 (0.027)
R-squared	0.144	0.167	0.220	0.138	0.156	0.175
Observations	105	105	105	105	105	105

Notes. This table shows how the decision to refuse offers involving identity tasks is correlated with various survey measures. The dependent variable is a worker-level indicator for refusing any identity task, i.e. turn down all offers involving an identity task (Columns 1-3), or an indicator for refusing all identity tasks (Columns 4-6). These outcomes are regressed on the variables related to worker characteristics. Comprehension score refers to the number of correct answers on the first attempt of the comprehension questions. Standard errors are clustered at the worker level.

B Appendix exposition

B.1 Notes on the conceptual framework

Section 1.2 outlines how to estimate the share of workers with identity concerns by observing workers' offer take-up decisions. Here the approach is described with greater detail.

Consider worker i evaluating different job offers as in Section 1.2. He prefers to take up the job offer involving default task 0 and extra task k if and only if the utility from taking up the offer exceeds that from his outside option. The outside option, represented by O_i , could involve working in another job or taking leisure. The worker's take-up decision is given by:

$$takeup_{ik}(c_i, t_k) = \begin{cases} 1, & \text{if } M_i + L_i - (V_{i0}(c_i, t_0) + F_{i0}(c_i)) \\ & - (V_{ik}(c_i, t_k) + F_{ik}(c_i)) > O_i \\ 0, & \text{otherwise.} \end{cases} \quad (5)$$

Suppose the variable utility cost of working on task k (or any other task) satisfies the following.

Assumption 2. *The variable cost function $V_{ik}(c_i, t_k) : R \times [0, 1] \rightarrow R$ is continuous in t from the right at 0, from the left at 1, and from both sides for all $t \in (0, 1)$. In addition, $V_{ik}(c_i, 0) = \lim_{t_k \rightarrow 0^+} V_{ik}(c_i, t_k) = 0$.*

Then, being slightly informal, one can find $\bar{\epsilon} > 0$ such that $V_{ik}(c_i, \epsilon) \approx V_{ik}(c_i, 0) = 0$ and $V_{i0}(c_i, T - \epsilon) \approx V_{i0}(c_i, T)$ for all $\epsilon < \bar{\epsilon}$. That is, when a worker spends very little time on task k , the time-varying utility cost of working on task k would be close to nothing. In addition, the time-varying utility cost of working on the default task would be similar to that of spending the entire working time on the default task.

Now, one can compare the offer of spending a small amount of time on extra task k to the offer of only working on the default task. Let θ_i be the net utility from taking up the latter offer.

$$\theta_i \equiv U_{ik}(c_i, 0) = M_i + L_i - O_i - [V_{i0}(c_i, T) + F_{i0}(c_i)]. \quad (6)$$

To see the take-up decision for the former offer, substitute for t_k with ϵ in Equation 5 and rearrange:

$$takeup_{ik}(c_i, \epsilon) \approx \begin{cases} 1, & \text{if } M_i + L_i - O_i - [V_{i0}(c_i, T) + F_{i0}(c_i) + F_{ik}(c_i)] \geq 0 \\ & \text{i.e. if } \theta_i - F_{ik}(c_i) \geq 0 \\ 0, & \text{otherwise.} \end{cases}$$

Similarly, to see the take-up decision for the latter offer, substitute t_k with 0 in Equation 5:

$$takeup_{ik}(c_i, 0) = \begin{cases} 1, & \text{if } M_i + L_i - O_i - [V_{i0}(c_i, T) + F_{i0}(c_i)] \geq 0 \\ & \text{i.e. if } \theta_i \geq 0 \\ 0, & \text{otherwise.} \end{cases}$$

Hence, the difference in the two take-up decisions would be attributable to $F_{ik}(c_i)$, the fixed utility cost of working on task k . Worker i declines the former and accepts the latter if:

$$F_{ik}(c_i) = f_k + \beta_{ik} \cdot I_k(c_i) > \theta_i \geq 0. \quad (7)$$

where the first equality follows from Equation 2.

Then, a number of different approaches are possible for testing whether β_{ik} is positive. One way would be to compare two tasks k and u , when they have the

same inherent fixed costs and yet only k is inconsistent with the worker's identity.⁵¹ Specifically, the tasks would be such that $I_k(c_i) = 1$, $I_u(c_i) = 0$, and $f_k = f_u$. If the worker only declines the offer with k , Equation 7 shows that $\beta_{ik} > \theta_i - f_k \geq 0$, i.e. this worker has identity concerns about working on task k .

Alternatively, one could compare two similar workers i and h such that task k is inconsistent only with worker i 's identity. The workers would be such that $I_k(c_i) = 1$, $I_k(c_h) = 0$, and $\theta_i - f_k = \theta_h - f_k$. If only worker i declines the offer, this would again indicate that $\beta_{ik} > \theta_i - f_k \geq 0$. However, in real life, it is difficult to find two tasks or two workers that satisfy these assumptions.

Instead, one could compare across groups of workers and tasks, as described in Section 1.2. Suppose there are two large groups of workers belonging to social categories A and B , who are willing to work on a job that only involves the default task. They evaluate two job offers that involve spending a small amount of time on extra tasks b and u . Task b is associated with group B whereas task u has no association. Thus only the former poses an identity violation for group A . The shares of workers in groups A and B who decline the offers involving tasks b and u are given by:

$$\begin{aligned}
\delta_{A,b} &= \sum_{i \in A} \mathbb{1}[f_b + \beta_{ib} > \theta_i] / N_A \\
\delta_{A,u} &= \sum_{i \in A} \mathbb{1}[f_u > \theta_i] / N_A \\
\delta_{B,b} &= \sum_{i \in B} \mathbb{1}[f_b > \theta_i] / N_B \\
\delta_{B,u} &= \sum_{i \in B} \mathbb{1}[f_u > \theta_i] / N_B.
\end{aligned} \tag{8}$$

Suppose Assumption 1, reproduced below, is true.

⁵¹Some studies take this approach by experimentally associating the same job with different identities (Delfino 2021; Del Carpio and Guadalupe 2021).

Assumption 1. *The distributions of f_b , f_u , and θ_i are such that*

$$P[f_b > \theta_i | i \in A] - P[f_u > \theta_i | i \in A] \leq P[f_b > \theta_i | i \in B] - P[f_u > \theta_i | i \in B].$$

This is satisfied, for example, if shifting the distributions of f_b , f_u , and θ_i for group A by the same amount gives the respective distributions for group B .

Let $\tilde{\delta}_{A,b}$ represent the shares of workers in groups A who decline the offers involving tasks b if no one in A faced any identity concerns:

$$\tilde{\delta}_{A,b} = \sum_{i \in A} \mathbb{1}[f_b > \theta_i] / N_A \quad (9)$$

In such a hypothetical world without identity concerns, the difference in differences of the shares, $(\tilde{\delta}_{A,b} - \delta_{A,u}) - (\delta_{B,b} - \delta_{B,u})$, is weakly negative in expectation.

Therefore, if the actual observed difference, $\Delta\delta := (\delta_{A,b} - \delta_{A,u}) - (\delta_{B,b} - \delta_{B,u})$ was strictly positive, it would indicate that the share of workers in A with positive β_{ib} is greater than $\Delta\delta$. Since for some workers β_{ib} could be positive but too small to add to $\Delta\delta$, this provides a lower bound on the share of workers who face identity concerns. I discuss whether the necessary assumptions seem realistic in the experimental setting in Section 3.3.

By randomizing whether workers' decisions are publicized to their neighbors, one could also estimate the share of workers who have social image concerns. The fixed utility cost function incorporating these concerns is written as:

$$F_{ik}(c_i) = f_k + \beta_{ik} \cdot I_k(c_i) + x_k \gamma_{ik} \cdot I_k(c_i) \quad (10)$$

where x_k is an indicator for whether worker i 's decision is observable.⁵² The social

⁵²This implicitly assumes that f_{ik} does not change with observability and that there is no cost

image costs described here specifically relate to the one-time costs of being perceived as the type to willingly engage in identity-inconsistent jobs. There could be additional costs, for example, if workers expect that their job performances will be also observable by neighbors.

B.2 The caste system in India

The historic caste system, dating as far back as 1500-500 BCE, comprises four hierarchical classes or *varnas*, the Brahmins, Kshatriyas, Vaishyas, and Shudras. The social group at the bottom of this hierarchy was excluded from the *varnas* altogether, and were called the untouchables. Each *varna* and the untouchables are further divided into many discrete communities called *jatis* or castes. There exist approximately 4,000 castes, whose members tend to live in small clusters scattered over potentially large regions (Munshi 2019).

The hierarchy embedded in the caste system is easily recognizable in political, economic, and social spheres of modern India (Deshpande 2011; Jodhka 2017). The modern Indian government endorses an affirmative action program, formally acknowledging the historical disadvantage some groups have faced compared to the other “forward” castes (FC). As in the traditional hierarchy, FC is considered to be above Other Backward Class (OBC), which is in turn above Scheduled Castes (SC, formerly the untouchables) and Scheduled Tribes (ST, marginalized indigenous groups).

Within each of these official categories, castes form an even finer layers of social hierarchy (Marriott 1958; Mahar 1960). The Hindu religious notions of purity and pollution determines which castes rank higher and thus are able to access or perform the more exclusive and prized ritual services. The system further imposes various behavioral prescriptions regarding how different castes ought to interact. Individuals associated with being observed by employers and surveyors.

belonging to higher castes are prohibited from making contact with—e.g. receiving water from, sharing cooked food with, or entering the houses of—those from lower castes. These practices serve as frequent reminders of individuals’ caste identities as well as their castes’ relative social positions (Shah 2006).

Another notable feature of the caste system is the historic links between castes and occupations. Some scholars (Gupta 2000) trace their origins to occupational guilds from the feudal period (7th to 12th century), whereas others argue that the British colonial government (19th to 20th century) either created or rigidly reinforced the connections between castes and jobs (Dirks 2001, Bayly 2001). These links effectively sustained a system of labor division in which individuals performed their caste-designated jobs for many generations.

Although a large number of people have abandoned their traditional jobs for new opportunities that arrived with modern developments, caste continues to play an important role in the Indian labor market (Mosse 2018; Desai and Dubey 2012). A number of studies examine the effects of caste-based networks or discrimination on labor market outcomes.⁵³ The behavioral channels through which caste could influence labor market outcomes include stereotype threat (Hoff and Pandey 2014), willingness to punish norm violations (Hoff, Kshetramade, and Fehr 2011), and in-group favoritism (Rao 2019; Lowe 2021). This paper suggests people’s desire to uphold caste identity may be another mechanism through which caste affects occupational preference and labor supply.

⁵³For example, Munshi and Rosenzweig study the influence of caste networks on schooling and job choice (2006) and migration decisions (2016). Madheswaran and Attewell (2007) and Thorat and Attewell (2007) study caste-based hiring discrimination. For a review in economics, see Munshi (2019).

B.3 Sample breakdown

The sample for the main experiment is stratified by caste and randomized privacy condition, as shown below.

	Public	Private	Total
Kaibarta	55	57	112
Sundhi	41	41	82
Dhoba	51	44	95
Kela	46	35	81
Mochi	30	30	60
Pana	59	61	120
Hadi	40	40	80
Total	322	308	630

The pre-registered targets were 120 for castes that are not associated with any experimental tasks (i.e. Kaibarta, Sundhi, Kela, and Pana), and 80 for the rest (i.e. Dhoba, Mochi, and Hadi). Due to the logistical difficulty of locating certain caste groups and time constraints, the targets were revised down for Sundhi (80), Kela (80), and Mochi (60). Privacy condition was randomized at the village level. Within each day, surveyors could not coordinate on the number of completed surveys exactly, so there are small deviations from targets for some groups.

The sample breakdown for the supplementary experiment is as follows:

	Public	Private	Total
Kaibarta	25	25	50
Pana	27	29	56
Total	52	54	106

The background surveys were conducted in a subset of villages that did not overlap with the experimental villages but are located in the same districts. The targeted castes were selected based on the Additional Rural Incomes Survey & Rural Eco-

conomic and Demographic Survey (ARIS/REDS) 2006 data and field scouting data, over-weighing OBC and SC castes. If certain caste groups were not present in the set of villages reserved for a particular survey, they were omitted from the sample. Overall each survey sample covered 15 different castes with some overlaps with the experimental castes.

	Task Survey	Rank Survey	Total
Bauri	7	7	14
Bhoi	11	7	18
Brahman	11	18	29
Chamar	0	6	6
Dhoba	10	21	31
Duma	10	0	10
Gauda	10	0	10
Gokha	10	6	16
Gudia	10	0	10
Hadi	10	21	31
Kaibarta	11	21	32
Kandara	11	7	18
Kela	0	10	10
Khandayat	10	0	10
Mali	0	12	12
Mochi	0	13	13
Pana	10	21	31
Sundhi	0	18	18
Teli	9	0	9
Tanla	10	21	31
Total	150	209	359

The analysis of experience levels in Section 4.4.1 uses ten castes that are part of the Task Survey sample which were also assigned rank scores during the Rank Survey. Figure A3 Panel B shows how these castes compare against the task-associated castes. The average rank scores determine how the castes are ranked, as in the main analysis. Specifically, because more than half the participants assign Bauri, Kandara, and Pana rank scores that are lower than Mochi's, these castes are considered lower-ranked than

Mochi. The three castes as well as Kela and Gokha are considered lower-ranked than Dhoba. Based on these relative status measures, I examine how the experience gaps with identity and control tasks vary across caste groups.

B.4 Vignette questions related to caste sensitivity

The following questions were used during the Task Survey as well as the follow-up surveys to measure caste sensitivity. Participants answered on a 5-point-scale indicating the strengths of approval or disapproval.

1. Sameer Jena went to Khorda recently to find work. There he met Sarveshwara Barik, who has been a barber in the area for 10 years. Sarveshwara has been looking for someone to take over the work and offered Sameer the job. Do you think it is acceptable for Sameer to become a barber even though he is from a higher caste?
2. Tukuna Naika is from the Hadi caste. He is currently looking for work in villages around him. Recently a contractor offered him work in his catering business, where Tukuna will be required to serve food to guests at functions. Do you feel it is acceptable for Tukuna to perform this task?
3. Shantilatha Sahoo is currently in the last year of college. She goes to college with a friend Nilakanth Sethi. They have been friends ever since childhood and Shantilatha likes Nilakanth very much. She wants to marry him but her village finds this relationship unacceptable as Shantilatha is from a higher caste and Nilakanth is from a lower caste. Do you think it is acceptable for a higher caste woman to marry a lower caste man?
4. Gagan Dalai has not been finding enough work in his village recently. He is very worried for his family. A contractor had recently come to the village and offered

him 7 days' work in another village. The contractor offered him Rs.350/day for cleaning sewage tanks. Gagan refused the job as it is lower caste work. Do you think Gagan did the right thing?

5. Kartik Behera and Tuna Naika are both agricultural laborers. They work together for the same landlord and in the evenings they come back to the village together. Once, when they were returning to the village, Tuna offered some home-made sweets to Kartik. A senior village member saw this and reprimanded Kartik for eating the sweets because Tuna Naika is of a lower caste. Do you think it's wrong for a higher caste person to accept home-cooked food from a lower caste person?
6. Bindusagar Behera and Rabi Naika have been friends since childhood. Whenever Rabi went to meet Bindusagar, he was not allowed to enter Bindusagar's house. They would talk outside Bindusagar's house. Now Bindusagar is getting married and he has invited Rabi to be a part of the marriage festivities. During the wedding, Rabi sits separately to eat (according to his caste). Do you think these village norms are acceptable as Rabi is from a lower caste?
7. Nerua Naika has recently finished secondary school and is looking for a job. He lives near Ramesh Maharana who is a carpenter. Ramesh offers to train Nerua in carpentry so that he can work with him. Do you think Nerua should try to work as a carpenter although he is from a lower caste?

B.5 Supplementary experiment results

The supplementary experiment results indicate that workers are more likely to refuse all offers involving identity tasks as opposed to control tasks. Appendix Table A13 shows how workers' experience with tasks or their comprehension of the experimental

procedures affects these results. The follow-up survey for the supplementary experiment asked questions regarding to what extent workers have performed each task. The coefficients on the different experience variables show that workers who have performed a specific task in own household—and to some extent for wage—are less likely to refuse the offers involving that task. However, the estimated gap in refusal between identity and paired control tasks is similar even when these controls are added to regressions (Columns 1-3). Hence, workers’ having relatively limited experience with identity tasks does not seem to explain why workers are more averse to performing identity tasks. In addition, the results are robust to excluding those who score low on comprehension questions (Column 4) or those who exhibit choice inconsistency (Column 5).

One might wonder if workers could incur identity-related utility costs from simply expressing willingness to work on identity tasks, rather than actually working on them. If so, the cost of accepting an offer might be incurred “now” while the benefit of acceptance would be weighed by the probability of having the offer randomly selected. In such a case, however, it is also plausible that such costs are not additive, i.e. once a worker agrees to some offer involving an identity task, it could be costless to accept other offers involving the same task or even other identity-violating tasks. In addition, workers who refuse any identity task on average turn down 3.2 tasks entirely. This means that in order to avoid admitting willingness to work on extra tasks, these workers give up a bonus of Rs. 252 in expectation, which is 84% of their daily wage.

More broadly, this concern relates to whether the BDM-based elicitation gives accurate bounds on how much wage workers are willing to give up, i.e. whether workers would give different answers when they are given just one offer involving ten times their daily wage. In practice, workers rarely expressed regret once a high bonus

offer involving an identity task was randomly selected. In addition, when the survey specifically asked if offered even more money, whether they would agree to doing a task which they refused, 99% said they would refuse such offers regardless of wage.

While there appears to be a clear division in worker types based on their reactions towards caste-inconsistent jobs, it does not seem straightforward to categorize them based on their characteristics. In Appendix Table A14, I use two proxies for whether a worker has identity concerns, namely refusing any identity task and refusing all identity tasks. The table reports the results from regressing them on a number of variables describing worker characteristics. The variables commonly found in other surveys, such as those related to age, education, and wealth, generally do not have statistically significant coefficients.⁵⁴ The coefficient on belonging to the Kaibarta caste is most robustly statistically significant at the 10% level. While having caste sensitive views is positively correlated with refusal (Column 3), the key remaining question is what kind of factors affect such opinions as well as related behaviors.

⁵⁴While last month's income is negatively correlated with refusal in Columns 1-3, the coefficient is not statistically significant in Column 4-6.