Essays in Behavioral Development Economics

Suanna Seung-yun Oh

Submitted in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy
under the Executive Committee
of the Graduate School of Arts and Sciences

COLUMBIA UNIVERSITY

2020

Abstract

Essays in Behavioral Development Economics

Suanna Seung-yun Oh

This dissertation analyzes how cultural and behavioral frictions affect decision-making in labor markets of developing economies. It studies factors that have received relatively little attention in economics—namely concerns about preserving identity, cognitive strain from financial stress, and gender norms—and examines their impacts on labor supply and productivity. Field experiments in the state of Odisha, India are used to provide direct empirical evidence on these relationships.

Chapter 1 investigates how identity—one's concept of self—influences economic behavior in the labor market, focusing on the effect of caste identity on labor supply. In the experiment, casual laborers belonging to different castes choose whether to take up various real job offers. All offers involve working on a default manufacturing task and an additional task. The additional task changes across offers, is performed in private, and differs in its association with specific castes. Workers' average take-up rate of offers is 23 percentage points lower if offers involve working on tasks that are associated with castes other than their own. This gap increases to 47 pp if the castes associated with the relevant offers rank lower than workers' own in the caste hierarchy. Responses to job offers are invariant to whether or not workers' choices are publicized, suggesting that the role of identity itself—rather than social image—is paramount. Using a supplementary experiment, I show that 43% of workers refuse to spend ten minutes working on tasks associated with other castes, even when offered ten times their daily wage. Results indicate that identity may be an

important constraint on labor supply, contributing to misallocation of talent in the economy.

Chapter 2—joint work with Supreet Kaur, Sendhil Mullainathan, and Frank Schilbach—tests for a direct causal impact of financial strain on worker productivity. The experiment randomly varies timing of income receipt among laborers who earn piece rates for manufacturing tasks: some workers receive their wages on earlier dates, altering when cash constraints are eased while holding overall wealth constant. Workers increase productivity by 5.3% on average in the days after cash receipt. The impacts are concentrated among poorer workers in the sample, who increase output by over 10%. This effect of cash on hand on productivity is not explained by mechanisms such as gift exchange, trust in the employer, or nutrition. The chapter also presents positive evidence that productivity increases are mediated through lower attentional errors in production, indicating a role for improved cognition after cash receipt. Finally, directing workers' attention to their finances via a salience intervention produced mixed results—consistent with concerns about priming highlighted in the literature. Results indicate a direct relationship between financial constraints and worker productivity and suggest that psychological channels mediated through attention play a role in this relationship.

Chapter 3 examines whether gender norms lead women to hold back their potential in the labor market. While the existing literature has shown that women tend to earn less than their husbands, there is limited direct evidence on whether women actively avoid earning more than their spouses and the determinants of such behavior. The experiment engages married couples working as casual laborers in a short-term manufacturing job that pays piece-rate on output. The experiment provides women an extra hour to work without this difference being salient, making it likely that they could earn more than their husbands. After husbands finish piece-rate production, women are randomized into one of three conditions in which 1) the wife is informed of her husband's production and expects both spouses to learn how much each spouse has produced, 2) the wife is informed of her husband's production and expects that only she will learn how much each spouse has produced, or 3) both spouses are only informed of their total joint production. Results show that women in the last two conditions achieve on average one hour's worth of production more than that of their husbands, suggesting that women do not face intrinsic concerns about earning more than their husbands. However, this productivity gap substantially decreases when husbands are expected to learn about individual production. This finding suggests that norms in marriage may be an important factor contributing to gender inequality in the labor market.

Table of Contents

List of Figures

List of Tables

Acknowledgements

I would like to extend heartfelt gratitude to those around me who made completing this dissertation possible.

My advisors provided invaluable support through every step of its progress. While all of them generously shared their time and insights, I would like to express especial gratitude to: Eric Verhoogen for enabling me to see the big picture and not lose the sight of it, Jonas Hjort for providing a lucid and cogent voice to deal with the complexities, Supreet Kaur for sending me to India and teaching me everything I know about running field experiments, Jack Willis for finding and filling in the gaps I made, Suresh Naidu for offering endless creative ideas for tackling my problems, and Kiki Pop-Eleches for recognizing the values of my ideas before I did. I would feel incredibly proud if I could one day become as great a mentor to someone as they were to me.

I would also like to thank the professors who provided my early training and made me want to become an economist: Seema Jayachandran at Northwestern University; Loren Brandt, Martin Burda, and Dwayne Benjamin at the University of Toronto; Marjorie McElroy and Thomas Nechyba at Duke University. The dissertation benefitted from interactions with Doug Almond, Michael Best, Ernesto Dal Bó, Frederico Finan, Emi Nakamura, Ted Miguel, Sendhil Mullainathan, Frank Schilbach, Rodrigo Soares, Jon Steinsson, and Miguel Urquiola.

I feel immensely lucky that the following people helped me run my field experiments in Odisha: Medha Aurora, Arnesh Chowdhury, Manvi Govil, Santosh Kumar Sahoo, Sneha Subramanian, and other staff members at the Abdul Latif Jameel Poverty Action Lab (J-PAL) South Asia. They were not only competent and driven, but were endlessly patient with my fumblings in the field.

Funding for the dissertation was generously provided by the Center for Development Economics and Policy (CDEP), Columbia Economics Program for Economic Research (PER), Columbia University Graduate School of Arts and Sciences, the National Science Foundation, the Pershing Square Venture Fund for Research on the Foundations of Human Behavior, the Trudy and Paul Woodruff Fellowship, and the Weiss Family Program for Research in Development Economics.

I was never lonely on this road thanks to the following people. Yoon-joo Jo, Nandita Krishnaswamy, and Yogita Shamdasani gave me tips for avoiding rough patches and finding places to rest. My classmates proved to be excellent company: Ashish Aggarwal, David Alfaro Serrano, Claudia Allende S.C., Juan Herreño, Han Huynh, Lorenzo Lagos, Cameron LaPoint, (RC) Xi Zhi Lim, Janet Martin, Xavier Moncasi, Lan Nguyen, Andy Pham, Divya Singh, Wonmun Shin, Mengxue Wang, Yue Yu, Jon Zytnick, and others. Amy Devine, Alejandro Favela Nava, and Florian Grosset were especially helpful during the final stretch. Bomi Choi, Lesley Choi, and Clare Park offered counsel whenever I was in need.

My family in Canada, South Korea, and the US provided infinite love and support. My partner and colleague, Evan Friedman, worked tirelessly as my most scrupulous editor and unwavering advocate.

Dedication

Dedicated to Mi-kyung Kim and Gi-sub Oh

for making me dream bigger and reach further.

Chapter 1: Does Identity Affect Labor Supply?

1.1 Introduction

People care about "who they are," both in their own view (identity) and in the perception of others (social image). These distinct but related concepts arise from social categories (e.g. men, women, high caste) and the behavioral prescriptions attached to them (Akerlof and Kranton 2000). A nascent literature in economics, as well as long-standing ones in other social sciences, investigates how concerns about identity and social image—image concerns—influence individual behaviors and market outcomes.[1] Findings in these literatures suggest that people may avoid otherwise desirable opportunities that evoke worries about upholding their identity or social image.[2]

However, the extent to which—and how—identity and social image affect economic behaviors in the labor market is not well understood. While image concerns could plausibly cause some groups to prefer certain occupations, groups tend to differ from others along many dimensions, including training and outside options. For this reason, it is difficult to establish the effects of image concerns from observational or survey data alone. It is also difficult, however, to isolate the effect in an experiment because researchers cannot randomly assign ingrained identities or radically change existing perceptions about specific occupations.

I address these challenges by exploiting unique features of the Indian caste system and provide the first experimental test of how identity and social image affect job-specific labor supply. Offering real jobs to casual laborers in rural Odisha, India, I show that image concerns have a large negative impact on workers' willingness to take up certain jobs. Workers are 23 percentage points (pp) more

[1] For overviews of identity theories from psychology and sociology, see Burke and Stets (2009), Stryker and Burke (2000), Hogg, Terry, and White (1995), and Owens, Robinson, and Smith-Lovin (2010), for example. For reviews in economics, see Hoff and Stiglitz (2016) and Bursztyn and Jensen (2017).

[2] For example, Gottfredson (1981), West and Zimmerman (1987) and Cejka and Eagly (1999) discuss how gendered perceptions of jobs affect occupational preference.

likely to turn down job offers that involve spending as little as ten minutes on tasks (out of five hours of total working time) if the tasks are associated with castes other than their own. The take-up rate gap increases to 47 pp when the castes associated with the relevant job offers are perceived as having lower social status than the individuals' own. These effects are invariant to whether or not workers' decisions are publicized, suggesting that identity—rather than social image—is the main driver of these effects.[3]

Social theories of identity suggest that performing tasks associated with other social groups may constitute a violation of identity. Concerns about such violations may be greater if the groups associated with the tasks have lower status (Tajfel and Turner 1979). In India, caste constitutes a central part of people's identity, and the social hierarchy of castes is commonly recognized. In addition, castes have historical links to specific occupations, which often extend to simple tasks associated with those occupations. Hence, these features allow me to construct job offers that involve working on such tasks and develop predictions of whether the offers involve conflicts of caste identity.

To obtain concrete information on caste-task associations and the caste hierarchy, I conduct two surveys separately from the experiment. The first survey allows me to identify a set of manual tasks to be used in the experiment, and document their associations with specific castes. The second survey is used to establish the ranking of castes selected for the experiment.

The experiment elicits 630 workers' willingness to take up job offers that involve spending some time on different manual tasks. All potential job offers involve working on a common default task of producing paper bags, which is not associated with any caste. The offers also entail working privately on an additional task. 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.[4] I can therefore side-step the concern that worker preferences

[3] The tasks that may involve caste associations are always performed in private.

[4] No task requires formal training or prior experience. Workers are also explicitly told that the offers are one-time offers and they will not influence their future job prospects.

for these attributes may vary across castes, which has been difficult to address in existing research.

To truthfully elicit worker preference for job offers, workers are asked to participate in a choice exercise based on the Becker-DeGroot-Marschak (BDM) procedure. Each worker is presented with a set of potential job offers and is asked to indicate whether he would take up or decline each one.[5] In addition, he is explicitly encouraged to consider them separately, and decide over each one as if it were a single, take-it-or-leave-it offer. After the worker indicates all decisions, one offer is randomly selected and his choice for this offer is implemented.

I assess how workers' willingness to take up job offers varies depending on the caste association of the extra task included. I hone in on the effect of spending only a brief time on the task, exploiting the across-offer variation in the time allotted to it. This allotment can be as much as ninety minutes, or as little as ten minutes. The fall in take-up from working on the extra task can be decomposed into changes at the intensive and extensive margin, i.e. due to spending longer time vs. spending any time at all on the task. Identity is expected to have a large effect on the latter, since spending any amount of time would still imply breaking one's internal rule of behavior (Akerlof and Kranton 2000). Hence, the discrete drops in take-up due to spending any time on caste-specific tasks instead of others would point to the effects of identity.

The resulting experimental data show that workers' willingness to take up job offers decreases significantly when they are predicted to involve conflicts of identity. I compare the take-up rates of offers involving "identity tasks" (tasks associated with specific castes, such as washing clothes) to those involving control tasks (similar to above but without any caste associations, such as washing farming tools). Among workers whose castes are closely associated with identity tasks, the take-up rates are similar across both task categories. Among the other workers, the take-up rate of offers involving identity tasks is much lower. The estimated take-up gap is 23 pp when the castes associated with identity tasks rank higher than the workers' own. The gap increases by an additional 24 pp when those castes rank lower. This second effect is larger for those who are caste-sensitive, i.e. those who express strong support for observing caste norms in a follow-up survey.

[5]Only male workers participate in the experiment due to practical difficulties. Many female workers are averse to traveling to work sites without male family members.

Notably, the large and statistically significant changes in take-up are present when workers are required to spend only ten minutes on extra tasks and vary little with any additional time. This indicates that the effects are due to the costs of engaging at all in identity tasks—consistent with the predicted effects of identity violations. It is unlikely that these patterns are driven by differential effort costs across tasks, which are expected to cause a continuous change in the take-up rate with the time spent.

This design provides a novel strategy for estimating the impact of identity on labor supply separate from the effect of social image. To distinguish the additional effect of social image, I randomize whether or not worker decisions are publicized; and I find similar effects across these 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.

It is difficult to find an alternate explanation for the constellation of findings. Any explanation would need to address why 1) take-up rates appear to drop as soon as workers spend any time on extra tasks, but vary little with additional time; 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 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.[6] Then they are 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.

[6]The focus of the supplementary experiment was not to verify the result with relative caste status variations so a smaller sample with two caste groups was used.

300. The largest bonus is ten times their daily wage, and is close to a whole month's earnings in the agricultural lean seasons during which the experiment takes place. As in the main experiment, one of the switching offers is randomly selected, and the worker's choice for it is implemented.

I find that 43% of workers are willing to forego as much as Rs. 3000 in order to avoid spending ten minutes on tasks associated with other castes in private. This is 29 pp greater compared to the take-up rate of offers involving control tasks, which do not have any caste association. Again, this difference is invariant to whether or not worker decisions are publicized. These findings suggest that identity can motivate workers to completely avoid certain jobs even at large economic costs.

This paper builds on and contributes to three literatures. First, my findings add to the literature on occupational choice by establishing the role of identity and social image.[7] These channels have been largely overlooked in economics, despite a large literature in sociology and psychology discussing their potential importance (e.g. Gottfredson 1981; West and Zimmerman 1987; Cejka and Eagly 1999). A number of theoretical studies (Akerlof and Kranton 2000; Bénabou and Tirole 2006, 2011) that espouse the need to account for these factors in economic decision making motivate the experiment.

Second, the study highlights identity as a channel that could contribute to the misallocation of talent in the economy. My findings suggest that some people may fail to pursue certain careers despite their potential aptitude due to concerns about identity. In addition to its direct impact on labor supply, identity-based occupational preferences could also interact with other well-studied channels of misallocation, such as discrimination.[8] The existing models on allocation of talent that do not take these mechanisms into account may over-attribute changes in aggregate productivity to certain channels.[9] Cassan, Keniston, and Kleineberg (2019) develops a structural general equilibrium Roy model of occupational choice that accounts for caste identity effects in India. This study can

[7] E.g. Topel and Ward (1992); Acemoglu and Autor (2011); Goldin (2014); Adda, Dustmann, and Stevens (2017).

[8] For example, existing studies show that some groups, such as high caste groups or men in occupations where they are over-represented, discriminate against other social groups that try to enter into their occupations (Schultz 1998; Padavic 1991; and Goldin (1990). One motivation behind this behavior may be the desire to reduce competition for the jobs that their own group members prefer due to identity reasons.

[9] Studies that focus on other sources of misallocation of talent include Hsieh et al. (2019), Erosa et al. (2017), Bell et al. (2019), and Goraya (2019).

complement such an approach.

More broadly, this study is part of a rapidly growing strand of work in economics which focuses on how culture and social contexts affect individual decision making (Hoff and Stiglitz 2016). Related studies show that norms and expectations surrounding social categories affect decision making in the lab as well as in field settings.[10] This paper is closely related to the studies that use field experiments to examine the role of social image and norms in the labor market (Breza, Kaur, and Krishnaswamy 2019; Bursztyn, Gonzalez, and Yanagizawa-Drott 2018; Karing 2018). To my knowledge, it is the first study to provide empirical evidence on the effect of identity on labor supply.

The rest of the paper is organized as follows. Section 2 presents some key ideas from theories of identity and builds a simple theoretical framework, which informs the experimental design. Section 3 describes the surveys which collect information on castes and tasks used in the experiment. Section 4 explains the empirical strategy for identifying identity effects. Section 5 describes the sample and procedures of the job take-up experiment, and Section 6 discusses the results. Section 7 presents the supplementary experiment design and findings. Section 8 concludes.

1.2 Conceptualizing identity

1.2.1 Theories of identity and social image

Psychologists and sociologists posit that identity and social image are powerful motivators of human behavior. While they have been discussed considerably less in economics, a number of theoretical studies suggest how to incorporate insights from other disciplines into economic models (Akerlof and Kranton 2000; Bénabou and Tirole 2006, 2011). The two concepts are often discussed

[10]For example, these studies examine outcomes such as cognitive performance (Hoff and Pandey 2006, 2014), dishonesty (Cohn, Fehr, and Maréchal 2014), contributions (Bursztyn et al. 2017; Benjamin, Choi, and Fisher 2016), investment in education (Fryer and Torelli 2010; Austen-Smith and Fryer 2005), and women's labor marker outcomes (Alesina, Giuliano, and Nunn 2013; Bertrand, Kamenica, and Pan 2015; Bursztyn, Fujiwara, and Pallais 2017). Focusing on the role of identity, a large number of studies show that interventions changing identity salience affect decision making in the lab, e.g. Benjamin, Choi, and Strickland (2010). A few studies use other experimental techniques to study identity effects, e.g. Bursztyn et al. (2017) and Falk (2017). However, little is known about identity effects on any field outcome.

together: people care about their own conception of "who they are" (variously referred to as identity, self-image, self-identity, and intrinsic motivations) as well as other people's perception of them (referred to as social image, reputation, and social identity). As previously mentioned, I refer to the former as identity, and the latter as social image.[11]

On the one hand, people are intrinsically motivated to uphold their identity. Bénabou and Tirole (2006) describe this as "a strong need to maintain conformity between actions or even feelings and ... identities they seek to uphold." In a related model (Bénabou and Tirole 2011), individuals infer their own values (types) from their past actions, and hence are motivated to take actions consistent with their identities for their future selves. Akerlof and Kranton (2000) emphasize that people abide by the prescriptions because failing to do so "evokes anxiety and discomfort in oneself." Both studies suggest the ways in which identity can curtail behavior, even in private settings.[12]

On the other hand, people seek to maintain a positive social image in front of others such as family, friends, and neighbors. The value of social-image can be material (e.g. social sanctions, loss of reputation decreasing future payoffs) and/or purely affective (e.g. social esteem or shame as a hedonic good).[13] A number of recent empirical studies examine the role of social image in economic decision making.[14] For example, some studies show that people behave differently depending on whether their actions are observable (e.g. Jakiela and Ozier 2016; Breza et al. 2019). Another study shows that correcting people's misconceptions about others' values changes their behaviors (Bursztyn et al. 2018).

I now discuss two ideas from social theories that are particularly important for establishing identity effects in field settings. Although existing literature describes a number of potential ways of studying social image, they are largely silent on identity. The theoretical ideas below suggest

[11]There are extensive discussions in social psychology on the complex ways in which identity and social image affect behaviors, such as those relating to multiple identities, desire for conformity vs. individuality, etc. Here I only focus on the most fundamental and relevant ideas.

[12]For example, Bursztyn et al. (2017) show that some Pakistani men are willing to forgo one-fifth of a day's wage in an anonymous and private setting to uphold their Anti-American political identity.

[13]Similarly, the reasons why people care about other people's observance of behavioral prescriptions can be material or affective. Dominant groups may strive to maintain the status quo system of prescribed behaviors that benefits them (Tajfel and Turner 1979). Akerlof and Kranton (2000) describe how seeing other people's violation of prescriptions can cause negative emotions in oneself.

[14]Bursztyn and Jensen (2017) provide a review of the literature on social-image.

how identity can affect decision making in the labor market, motivating some potential strategies for capturing its effects using an experiment.

First, those belonging to a social group may be averse to adopting the characteristics and practices of other groups, particularly if the other groups have lower social status.[15] Tajfel and Turner (1979) theorize that the utility people derive from identifying with a social category increases with its status.[16] This implies that in the labor market, workers may be averse to engaging in tasks associated with other groups as this conflicts with their identity; this aversion may be especially strong if they perceive their groups to be of higher status than those associated with the tasks.

Second, the concept of violation is central to understanding how identity affects behavior. People care about whether their internal rules of behavior have been breached at all—not just whether there have been persistent deviations or complete abandonment of identity. The literature on personality development emphasizes the negative emotions one experiences when internal rules of behavior are broken. Motivated by this, Akerlof and Kranton (2000) build a model of identity in which any violation of such rules results in a loss in utility. In the model by Bénabou and Tirole (2011), an individual inferring her type from past actions may remember whether she contemplated violating her rules of behavior, and this memory can serve as a negative signal about her type. This could compel her to avoid even the mere thought of breaking such rules, making them "priceless", i.e. they become what one "would never do" regardless of any pecuniary benefits. These models imply that those facing identity concerns regarding specific jobs would avoid engaging in them even for a short time. Furthermore, they may refuse to put a price on undertaking them, i.e. avoid them regardless of the offered wage.

The two ideas above suggest potential strategies for establishing identity effects. One could focus on capturing the effects of identity violations, i.e. examine worker preference for jobs that require spending very little time on tasks that involve conflicts of identity. Such tasks should have

[15]Consistent with this idea, Atkin et al. (2019) show that in India, when the status of a religious group increases, more households adopt food consumption patterns that are characteristic of the group.

[16]For a related review, see Bettencourt et al. (2001). This way, the literature on identity and social image is also tied to that 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.

associations with different social groups whose status can be clearly measured. These strategies motivate the framework below.

1.2.2 Worker's job take-up decision

I present a simple conceptual framework for a worker's decision problem of job take-up. The worker considers whether to take up a one-day job, which involves working on two tasks. As in a standard economic model, the worker's utility depends on wage, total working time, and task-specific costs of effort. The novel feature is that it also factors in the costs of engaging at all in each task, which can vary depending on the worker's social category. Suppose the worker expects the working conditions as well as the take-up decisions to be private information.

Worker preferences are described by:

$$U(c_i, w, \mathbf{t}) = \underbrace{M_i(w)}_{\text{Money}} + \underbrace{L_i(1-T)}_{\text{Leisure}} - \sum_{g=j,k} [\underbrace{V_{ig}(c_i, t_g)}_{\text{Variable Cost}} + \mathbb{1}[t_g > 0] \cdot \underbrace{F_{ig}(c_i)}_{\text{Fixed Cost}}]. \tag{1.1}$$

An individual worker, indexed by i and belonging to social category c_i, considers a job offer that involves working on two tasks—a default task j and an extra task k.[17] The worker expects to spend fraction t_k (and t_j) of his day working on task k (and j), and thus spend a total fraction $T = t_k + t_j$ working. $M_i(w)$ indicates worker i's utility from the daily wage w, and $L_i(1-T)$ indicates the utility from leisure, which is a function of the total non-working time. The utilities from wage and leisure are offset by the sum of the utility costs from working on each task involved in the job.

The utility costs of working on task k are of two types: $V_{ig}(c_i, t_g)$, which describes the variable effort cost that depends on the time spent on the task, and $F_{ig}(c_i)$ which indicates the fixed utility cost of engaging at all in the task. The fixed utility cost only gets incurred if the worker spends any time on the task, i.e. if $t_k > 0$, and does not depend on the amount of time spent on it. By assumption (formally stated as Assumption 1.10.1 in Appendix Section 1.10.1), the variable effort cost is zero when the worker does not spend any time on the task (when $t_k = 0$) and is continuous in

[17] Since all job offers in the experiment involve working on two tasks, I do not consider the utility costs of working on fewer or more tasks.

the time spent on the task. One functional form that satisfies this assumption is a linear function of t_k.[18] With these assumptions, the utility costs of working on task k (and similarly for task j) are written as

$$\begin{aligned} V_{ik}(c_i, t_k) &= [v_{ik} + \alpha^d_{ik} \cdot I^d_k(c_i) + \alpha^l_{ik} \cdot I^l_k(c_i) \cdot I^d_k(c_i)] \cdot t_k \\ F_{ik}(c_i) &= f_{ik} + \beta^d_{ik} \cdot I^d_k(c_i) + \beta^l_{ik} \cdot I^l_k(c_i) \cdot I^d_k(c_i). \end{aligned} \tag{1.2}$$

Both $V_{ik}(c_i, t_k)$ and $F_{ik}(c_i)$ contain components that do not depend on the worker's social category, namely v_{ik} and f_{ik}. The motivation is as follows. The variable cost component that is independent of identity concerns, v_{ik}, may be large because, for example, this task is too difficult, tiresome, or boring to spend time on. The independent fixed cost f_{ik} may be large because, for example, the worker could be averse to trying out this task for the first time or expect it to involve initial unpleasantness.[19]

The remaining cost components depend on worker's social category, c_i. The indicator $I^d_k(c_i)$ takes the value of 1 when task k is associated with a category that is different from c_i. The indicator $I^l_k(c_i)$ takes the value of 1 when the category associated with task k is not only different from c_i, but also has lower status than c_i.[20] The utility effects of the relationships between task associations and the worker's category are represented by the parameters α^d_{ik}, α^l_{ik}, β^d_{ik}, and β^l_{ik}.

The theoretical discussion in Section 1.2.1 suggests that concerns about identity would have a large effect on the fixed utility costs of working on task k, because people care about whether their internal rules are violated. Hence, I focus the discussion here on β^d_{ik} and β^l_{ik}, although the effects related to variable utility costs are also examined with the experimental data. If working on a task associated with a different social category constituted an identity violation, this would increase the fixed utility cost of engaging in it, as represented by a positive value of β^d_{ik}. If this utility effect is larger when the category associated with the task has a lower status, then this would be captured by

[18] This functional form turns out to be appropriate for the experimental tasks involved in the study. Other functional forms are tested during robustness checks.

[19] For instance, the worker could be worried that initially touching job-related objects may greatly increase the risk of contacting germs.

[20] The implicit assumption is that social categories never have the same social status.

a positive value of β^{l}_{ik}.

Although it is not possible to directly measure these utility parameters, under some specific assumptions, one can estimate the lower bound on the share of workers for whom the parameters are positive. Here I provide a quick overview on this approach, giving a more formal discussion in Appendix Section 1.10.1.

Suppose there are two groups of workers belonging to social categories A and B, all of whom are willing to take up the job offer of only working on the default task. They are asked to evaluate two job offers like the one described above, against the workers' individual outside options. The first offer involves spending some time on extra task b, which is associated with category B which has a higher social status than A, while the second offer involves spending some time on extra task u, which is unassociated with any category. This means the utilities from wage, leisure, outside option, and costs of working on the default task are the same across the two offers. In addition, suppose the offers involve spending very short amounts of time on the extra tasks. Given Assumption 1.10.1, the variable effort costs of working on task b and u are close to zero in this case. Then any decreases in the take-up rates of the two offers must be due to the fixed utility costs of working on task b and u.

The shares of workers who decline the offers are given by:

$$\begin{aligned}
\delta_{A,b} &= \sum_{i \in A} \mathbb{1}[f_{ib} + \beta^{d}_{ib} > \theta_i]/N_A \\
\delta_{A,u} &= \sum_{i \in A} \mathbb{1}[f_{iu} > \theta_i]/N_A \\
\delta_{B,b} &= \sum_{i \in B} \mathbb{1}[f_{ib} > \theta_i]/N_B \\
\delta_{B,u} &= \sum_{i \in B} \mathbb{1}[f_{iu} > \theta_i]/N_B
\end{aligned} \tag{1.3}$$

where θ_i refers to the utility from taking up the job of only working on the default task and N refers to the size of the worker group. There are a number of different assumptions that would allow one to infer the lower bound on the share of workers for whom β^{d}_{ib} is positive by comparing observed $\hat{\delta}$'s. One assumption that is sufficient for this purpose and is likely to hold in the experimental setting is the following.

Assumption 1.2.1. *(Distributions of fixed cost differences)*

$$(f_{ib} - f_{iu})_{\{i\in A\}} \overset{\text{dist}}{=} (f_{ib} - f_{iu})_{\{i\in B\}}$$

This condition implies that without concerns about identity, workers would be similarly averse to engaging at all in task b compared to engaging at all in task u. If the above assumption also holds regarding another social category that has a higher social status than B, then the lower bound on the share of workers for whom β^l_{ib} is positive can also be estimated.

An additional note is that one can easily incorporate the utility costs associated with social image into this framework. Assuming that concerns for social image can additionally increase the utility costs of working on a task, $F_k(n_{ik}, c_i)$ contains additional components as follows:

$$\begin{aligned} F_{ik}(c_i) = \quad & f_{ik} + \beta^d_{ik} \cdot I^d_k(c_i) + \beta^l_{ik} \cdot I^l_k(c_i) \cdot I^d_k(c_i) \\ & + x_k \gamma^d_{ik} \cdot I^d_k(c_i) + x_k \gamma^l_{ik} \cdot I^l_k(c_i) \cdot I^d_k(c_i) \end{aligned} \tag{1.4}$$

where x_k is a parameter indicating whether the worker's take-up decision is observable or not. Hence, to study the effect of social image on labor supply, one could randomly vary the observability of workers' decisions.

In Appendix Section 1.10.1, I provide more details on the approach discussed here. The next section describes the tasks and groups on which I will apply this framework. Section 1.4 discusses the related empirical strategy for identification.

1.3 Background surveys on castes

The Indian labor market, with the historical caste system, provides an ideal setting for testing the theoretical framework. Ideally, the labor market should contain a number of jobs or tasks which are associated with different social groups, and these groups should form a distinctive social hierarchy. In India, caste membership defines an important part of people's identity, with different castes conceived as embodying particular characteristics and values (Hoff and Pandey 2014). Notably, caste provides a system of social hierarchy. Caste status determines how one ought to interact with others, even with respect to everyday practices such as sharing food or water (Marriott 1958; Mahar 1960). Furthermore, there are historical associations between castes and occupations, many of which are still widely recognized (Desai and Debey 2012). These associations often carry over to simple daily tasks that relate to those occupations such as washing clothes. These features of the Indian caste system can be used to build different worker-task combinations for testing identity effects.

I conducted two surveys separately from the experiment; they were designed to document the locally prevalent views on the associations between castes and tasks, as well as on the caste hierarchy. There is a substantial geographic variation in the availability of certain castes, and therefore in the knowledge and perceptions regarding them (Munshi 2019; Marriot 1958), which makes the local context important. The surveys and the experiment were conducted in the Nayagarh, Dhenkanal, and Khordha districts in the state of Odisha. More detailed background information on the Indian caste system is provided in Appendix Section 1.10.2.

1.3.1 Survey procedures

The Task Survey (N=151, 15 caste groups) was designed to collect information about the associations between castes and simple manual tasks. The list of tasks was generated based on qualitative interviews conducted prior to the survey. For each task, the participants indicated whether a particular caste performs it, and the extent to which they have personally performed it.[21] In addition, I collected their knowledge of local castes. A list of caste groups residing in Odisha was taken from the Additional Rural Incomes Survey & Rural Economic and Demographic Survey (ARIS/REDS) 2006 codebook. For each caste on the list, participants were asked whether they knew of the caste and discussed the caste's historic occupations.

Based on the survey findings, seven caste groups were selected for the experiment. Three caste

[21] They also indicated whether a task is gender-specific. This was useful for selecting experimental tasks that do not involve conflicts of gender identity. The survey and experiment sample only involved male workers due to difficulties with employing female workers in this setting.

groups had strong connections to the manual tasks and were well known to the participants. All three belonged to Scheduled Castes (SC)—officially designated groups of historical disadvantage, formerly known as the untouchables. To make the remaining groups in the experimental sample comparable in terms of wealth and outside options, they were drawn from Scheduled Castes (SC) and Other Backward Castes (OBC). All selected groups were known to over 70% of the survey participants; six of them were known to over 90%.[22]

The Ranking Survey (N=209, 15 caste groups) was designed to document how the seven castes in the experimental sample place in the caste hierarchy. Those who knew of all seven castes were recruited for this survey. The participants were provided with cards, and on each of them there was a caste name written.[23] They were asked to arrange the cards according to the caste hierarchy, placing any names horizontally if they occupy the same position in the hierarchy. After ranking the seven castes, the participants additionally ranked nine other castes.[24] The participants inserted the nine name cards into the hierarchy, skipping over any castes they do not know.

When defining the caste hierarchy, participants randomly received one of three different instructions. The first version directly asked about the caste hierarchy. The other two asked the rankings to be based on the practice of accepting cooked food or the practice of accepting water—as higher castes are not to accept such items from lower castes. These two practices are among the most common behavioral rules attached to the caste hierarchy (Marriott 1958; Mahar 1960).

1.3.2 Caste ranking and associations with tasks

Table 1.1 summarizes key information from the Task and Ranking Surveys. The table is organized such that castes and tasks that have connections are placed within the same rows.

Column (1) shows the list of the castes selected for the experiment sample, sorted according to their average rank in the caste hierarchy. Because the Ranking Survey participants ranked all seven

[22]This sample includes all six SC castes that meet the knowledge threshold of 70%. There are many castes in OBC that meet this condition, so just one group that was perceived as having similar wealth and status as those in SC was selected. More detailed information about SC, OBC, and other categories are provided in Appendix Section 1.10.2.

[23]The surveyors also guided participants with the name cards in case they were illiterate.

[24]These castes participated in the Task Survey or were SC castes not meeting the knowledge threshold.

castes without missing values, the rank scores are generated by simply averaging across reported caste rankings. This ranking is similar across different districts and versions of instructions used.[25] Further details on the consistency of the ranking are provided in Appendix Section 1.10.3.

The three tasks with strong caste associations are listed in Column (3), and are referred to as identity tasks hereon. Column (4) shows the share of survey participants who report such associations. For instance, 72% of the participants report that washing clothes is specifically performed by the Dhoba caste.[26]

Three tasks that are especially similar to identity tasks but do not have caste associations are listed in Column (5), and are referred to as paired control tasks. No one associates these tasks with the relevant castes in Column (1); therefore, Column (5) shows the share of participants that report associations between these tasks with any SC caste. 15% of participants associate mending grass (floor) mats with a number of different SC castes, and no one associates the other two tasks with any SC caste. Four additional tasks without caste associations are designated as a default task or pure control tasks, and are described in Appendix Table 1.7.[27]

Participants report varying degrees of familiarity with these tasks. Columns (6)-(9) of Appendix Table 1.7 show that while most people have experiences with washing clothes (98%) and washing farming tools (89%), few have ever mended leather shoes (19%) or grass mats (10%). More people have performed sweeping animal sheds (81%) compared to sweeping latrines (51%), which is plausible since many households do not have latrines in this setting. Notably, most respondents have performed washing and sweeping tasks in their household, but have rarely done so for other people or as paid work.

[25] Some of these castes are mentioned in the anthropological work by Risley (1908) from the early 20th century, which report similar relative status of those groups.

[26] These task associations closely follow the traditional connections between castes and occupations. For example, Risley (1892) report that the typical occupations for Dhoba, Mochi, and Hadi are washer, leather work/cobber, and scavenger, respectively.

[27] Gender associations of all tasks are reported in Appendix Table 1.7 as well.

1.4 Empirically identifying identity effects

The caste groups and tasks presented in Table 1.1 can be used to test for identity effects. In an experiment based on the conceptual framework (in Section 1.2.2), the latter can be used as extra tasks involved in job offers. To compare worker responses to job offers involving these tasks and infer the share of workers facing costs of caste identity violations, the differences in independent fixed utility costs for these tasks must be distributed similarly across the different caste groups (Assumption 1.2.1).

I examine whether these assumptions appear consistent with the data patterns from the Task and Ranking Surveys. Even though the experimental tasks are simple manual activities and job offers would not require any training, workers' fixed utility costs of engaging in them could still depend on their prior experiences. I examine how participants' prior experience with tasks vary across castes using the following empirical specification:

$$\begin{aligned} Y_{ik} =& \sigma^d \, different_{ik} + \sigma^l \, different_{ik} \cdot identity_k \\ &+ \nu^d \, lower_{ik} + \nu^l \, lower_{ik} \cdot identity_k \\ &+ \eta \, identity_k + \phi \, purecontrol_k + \rho_k P_k + \chi_i X_i + \epsilon_i . \end{aligned} \tag{1.5}$$

The dependent variable Y_{ik} is a measure of worker i's task-relevant experience for task k. The covariates $identity_k$ and $purecontrol_k$ are the indicators for task k's category, as specified in Table 1.1.

The relative status variables $different_{ik}$ and $lower_{ik}$ are defined according to the ranking in Table 1.1. For any worker i that belongs to caste c_i, task k is considered to be a same-ranked task if it appears in the same row as c_i. For example, for a Dhoba worker, both washing clothes and washing farming tools are considered same-ranked tasks. Same-ranked tasks are in the omitted category of the specification. If task k appears in a different row from the worker's caste, the task is considered a different task and $different_{ik}$ takes the value of 1. If task k appears in a row beneath the worker's caste, the task is considered a lower task and $lower_{ik}$ takes the value of 1; thus, it would be 0 if task k is in a row above the worker's caste (a higher task). The other covariates, P_k and X_i, are used to control for task or worker-related fixed effects. P_k is a vector of task-specific indicator variables, and X_i is a vector of dummies indicating each caste or each worker. P_k and X_i are not included in the basic regression; when they are included, $identity_k$ and $purecontrol_k$ are omitted.

Table 1.2 shows the results of OLS regressions with three different measures of experience with selected tasks. Standard errors are clustered at the worker-task level, since identity effects are predicted to vary at this level.[28] Column (1) shows the basic regression result, which suggests that there are task and caste-specific differences in experience levels even across paired control tasks, i.e. everyone has more experience with certain tasks and some caste groups have more experience with all tasks. Controlling for task and caste/worker fixed effects in Columns (2) and (3), the results show that experience with paired control tasks do not vary across caste groups in a statistically significant way.

However, workers are 22 pp less likely to be experienced with identity tasks associated with other castes, as suggested by the coefficient on Different $\times$ Identity in Columns (2)-(3). The results in the remaining columns indicate that this is driven by the differences in wage-paying experience. Workers whose castes are directly associated with tasks are 27 pp more likely to have wage-paying experience. Caste groups do not differ in their non-wage-paying experience (i.e. ever performed within their household or for friends/neighbors). This is plausible given that people tend to perform washing or sweeping at home, but only those castes directly associated with identity tasks may be willing to such tasks for a wage. Notably, among those who are not directly associated with the tasks, there appear to be no statistically significant differences in experience levels, regardless of their relative ranking.[29]

Given these results, the preferred empirical specification for examining workers' take-up deci-

[28]Standard errors are clustered at the worker level during robustness checks, which leads to similar results.

[29]The sample used for the experience level analysis includes 10 castes, 4 of which are selected for the experiment. Although this sample includes some castes that are ranked higher than all of the experimental castes, the experience levels still do not differ across whether the tasks are considered higher vs. lower.

sions is as follows:

$$\begin{aligned} Y_{ikt} =& \sigma^d \; different_{ik} + \sigma^l \; different_{ik} \cdot identity_k \\ & + \nu^d \; lower_{ik} + \nu^l \; lower_{ik} \cdot identity_k \\ & + \tau_k \; time_{tk} + \rho_k P_k + \chi_i X_i + \epsilon_i. \end{aligned} \tag{1.6}$$

The dependent variable Y_{itk} is now worker i's take-up decision for the offer that requires spending time t_k on task k. The covariates $different_{ik}$, $lower_{ik}$, P_k, and X_i are all defined in the same way. Given the importance of task- and caste-specific factors in determining worker's experience level, the preferred specification controls for task and caste(/worker) fixed effects. The specification also accounts for the effect of the time required on the extra task, $time_{tk}$. The preferred specification controls for task-specific linear time trends; other functional forms are tested during robustness checks.

Given Assumption 1.2.1, the coefficients on Different × Identity and Lower × Identity would indicate the changes in offer take-up due to concerns about identity. However, one may be concerned that this assumption does not hold when comparing those whose castes are directly associated with the tasks against the rest. For example, having wage-paying experience with a task might reduce the aversion to engaging in it (in terms of fixed utility costs). If so, the coefficient on Different × Identity would be biased upwards; it would overestimate the identity effect related to engaging in tasks associated with other castes.

If the assumption held only for groups that are not directly associated with the extra tasks, the coefficient on Lower × Identity will still measure the change in take-up due to identity concerns. It would indicate the *additional* impact of engaging in tasks associated with castes that rank lower—rather than higher—than one's own. Thus, it can serve as the *lower bound* on the overall effect of identity on job take-up.

1.5 Experiment on job offer take-up

1.5.1 Sample construction and recruiting

The experimental sample is composed of 630 male household heads aged 18-55, who primarily derive income from casual daily-wage labor.[30] The sample is stratified by caste and randomized privacy condition. The breakdown of the sample is described in Appendix 1.10.4.

Casual laborers in this setting tend to work in agriculture during peak planting and harvesting seasons, and get short-term contract jobs in unskilled manufacturing or construction in the remaining lean periods. Potential employers often recruit directly by visiting workers' villages. They provide a job description and offer wages at the market prevailing daily wage rate. Workers who agree to the offered terms start working that day or on a prearranged, upcoming date. The experiment takes place during agricultural lean periods, when many workers are unemployed, often involuntarily (Breza, Kaur, and Shamdasani 2018).

The recruiting process for the experiment exactly mimics the natural labor market procedures. A scouting team first identifies villages with casual laborers belonging to the target caste groups. The identified villages are randomized into two privacy conditions: private vs. public. Surveyors visit the identified villages in the following days to discuss and make job offers. Only those who express interest in a one-day job of making paper bags are asked to participate in the choice exercise, which effectively results in a slightly different job offer.[31] This means all workers in the experiment prefer having a job making paper bags to their default outside option, although this answer is not incentivized.

The motivation for this recruiting process and criteria is threefold: 1) workers' decisions during the choice exercise mimic their day-to-day labor supply decisions; 2) workers in the sample have similar outside options, i.e. their opportunity costs are smaller than the utilities of taking up the hypothetical job of making paper bags; and 3) it provides a simple reference point to compare

[30] I restrict my sample to male workers, for the reasons discussed earlier.

[31] These paper bags are commonly used in the markets to store nuts and snacks, and the produced paper bags are sold to wholesale vendors. The setup and operation of the worksites are similar to those described in (Breza, Kaur, and Shamdasani 2018).

different take-up rates. For any sub-sample, the average take-up rate of an offer implies a decrease from the perfect take-up rate (i.e. rate of 1) of the hypothetical job.

1.5.2 Choice exercise and experimental procedures

To elicit workers' true willingness to take up job offers, a procedure based on the Becker-DeGroot-Marschak (BDM) mechanism is used. (Becker, DeGroot, and Marschak 1964). In the choice exercise, each worker is presented with an entire set of potential job offers—all involving the same daily wage of Rs. 300—and is asked to indicate whether he would take up each offer. After he chooses over all offers, one offer is randomly selected and his stated choice for it is implemented. This mechanism is incentive compatible under risk neutrality.[32] To further encourage truth-telling, workers are explicitly asked to consider each one as a single take-it-or-leave-it offer and give a "simple honest answer" about what they prefer.[33] In the analysis, I consider each worker's choices as reflecting his true willingness to take up offers.

The experiment proceeds as follows. Each choice exercise is conducted privately between a surveyor and a worker. Workers first go through a practice exercise which is designed to help them understand the BDM procedure. During the practice exercise, workers are offered the opportunity to buy different combinations of packaged food (e.g. mustard seeds and sugar), all involving at the same fixed price. Workers can choose to accept—purchase the combination—or decline each offer. One offer is randomly selected and implemented, and surveyors verify workers' understanding of the process.

Then, workers go through the choice exercise with job offers. Surveyors describe the set of all potential job offers and highlight how only one of their choices will be enforced. All job offers are the same in most aspects, including wage, total working time of five hours, worksite location, etc.

[32] For more details on the mechanism and its use in experiments, see Fudenberg, Levine, and Maniadis (2012).

[33] The exercise is justified to workers by explaining that the employer is looking for people to complete different tasks in addition to making paper bags and thus wants to collect information on what kind of jobs people are willing to do. The employer is also described as wanting to be fair when assigning different job offers. Finally, workers are told it would be costly if they accept an offer and change their mind later. If a worker accepts an offer that is randomly selected, surveyors will visit his village multiple times over the following days to compel him to complete the agreed upon job.

They require spending their majority of working time on the default task of producing paper bags, and the remainder on an extra task in private.

Job offers vary only in the *type* of extra task and in the *time* required to work on it. All workers' choice sets include eight extra tasks: three identity tasks, three paired control tasks, and two pure control tasks (ref. Table 1.1).[34] All tasks are described as not requiring any prior experience.[35] There are four different time requirements for the extra tasks: 10 minutes, 30 minutes, 1 hour, and 1.5 hours.[36] Surveyors describe each offer in detail and show photos depicting the tasks, similar to those in Appendix Figure 1.4.

Although the extra task is always performed in private, workers' take-up decisions may be publicized depending on their randomized privacy condition. Each village is scheduled to host a focus group meeting in the days following the exercise. Local agricultural practices are discussed in these meetings, and many village members (including those who did not participate in the experiment) are invited to attend. Those in the public condition are told that all their choices during the job offer exercise will be openly discussed during these meetings, irrespective of their attendance. Those in the private condition are assured that their choices will remain private information, except for their willingness to wash farming tools, a control task.[37] Hence, the two conditions should differ only in the observability of workers' decisions, not the observability of their job performance nor their expectations about the focus group activities.

Third, workers go over the list of job offers and choose whether to take up each offer. To test for any order effect, the order in which tasks appear on this list is randomized in four different ways across workers. Time requirement is also randomly chosen to be presented in an ascending or descending order. All offers have the same chance of being randomly selected.[38]

Finally, one offer per worker is randomly selected, and his choice for this offer is implemented.

[34] One pure control task is always stitching. The other is randomly chosen to be making ropes or deshelling peanuts.

[35] The tasks that could involve special skills, such as mending leather shoes or grass mats are described as assisting an experienced worker.

[36] These 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 worksites.

[37] The justification is that discussing local agricultural practices would involve talking about people's willingness to wash farming tools.

[38] Workers roll dice and draw scratch cards to select offers.

If the worker chose to take up the offer, he can complete the job within the next three days and receive Rs. 300.[39] He is also asked to complete a follow-up survey at the worksite. If the worker declined his selected offer, he does not get any other offer. However, he is still asked to complete the follow-up survey to receive a gift worth Rs. 50.[40] This compensation is mentioned at the very end, so that they do not factor this into their outside options during the choice exercise.

1.6 Results: conflicts of identity lower job take-up

The experimental results reveal that many workers are averse to taking up job offers associated with other castes, and especially so when those castes rank lower than their own. My set of findings are well in line with the explanation that those workers face strong concerns about violating caste identity.

1.6.1 Visualizing offer take-up rates

I first use plots to examine patterns in the raw data. Figure 1.1 plots the average take-up rates of job offers against the time required on extra tasks, separately for paired control tasks and identity tasks. The three connected lines show the differences in take-up rates according to the relative status of the extra task: same-ranked, higher, or lower, as defined in Section 1.4. When tasks are same-ranked, take-up rates are similar between (offers involving) paired control tasks and identity tasks. When tasks are higher, take-up rates for paired control tasks increase slightly while take-up rates for identity tasks fall by about 20 pp. When tasks are lower, take-up rates fall for both paired control tasks and identity tasks, but the decrease for identity tasks is almost twice as large at about 40 pp. All connected lines appear linear and parallel with small negative slopes, i.e. spending longer time on an extra task appears to have little effect on take-up rates. Given that the take-up rate is one when time requirement on extra tasks is zero minutes (according to workers' stated preference

[39] 57% of workers receive offers that they are willing to take up, and 67% of those complete their jobs. Absenteeism is prevalent in this region (Krishnaswamy 2019), especially for casual contract jobs. More details on job completion are provided in the next section.

[40] Those who do not complete their offered jobs in the next three days are also asked to do the survey for the same compensation. The follow-up survey completion rate is high (87%) and discussed in the next section.

during recruiting), there appear to be discrete drops in take-up rates associated with spending any time at all on an extra task.

Appendix Figure 1.5 shows a similar plot for each extra task type, with four connected lines representing four disjoint sets of castes. This provides more detailed information compared to Figure 1.1, which plots pooled averages over different tasks and castes. Even within the same groups of workers, take-up rates vary widely depending on the type of extra task. For example, among the workers belonging to castes other than Hadi (associated with sweeping latrines), 93% are willing to deshell peanuts for ten minutes, but only 25% are willing to sweep latrines for ten minutes. Higher caste groups tend to have lower take-up rates for some control tasks as well as identity tasks.[41] Again, the discrete drops in take-up rates are larger for identity tasks compared to paired control tasks when tasks are considered lower.

1.6.2 Regression analysis

A regression analysis with the experimental data confirms the findings from the plots. Table 1.3 reports the results from running ordinary-least-squares (OLS) regressions based on the empirical specification in Equation 1.6. Column (1) shows the basic specification result, and Columns (2)-(3) the preferred specification results that control for task and caste(/worker) fixed effects. All specifications control for task-specific linear time trends. Hence the main coefficients can be interpreted as the changes in take-up rates that are due to engaging at all in extra tasks, i.e. due to spending close to zero minutes on extra tasks.

The coefficients on Different $\times$ Identity and Lower $\times$ Identity have large magnitudes, similar across all specifications, and are always statistically significant at 1% level. The other coefficients are smaller in magnitude and less consistent. When offers involve tasks that are not same-ranked, take-up rates are lower for identity tasks compared to paired control tasks. This gap is 23 pp when tasks are higher, and increases by an additional 24 pp when tasks are lower. Hence the total estimated effect of identity violation on take-up is 47 pp, with 24 pp as the unbiased estimate of the

[41] The control tasks, such as mending grass mats and sweeping animal sheds, are typically perceived as menial according to field interviews.

lower bound.

Using a different caste ranking. Replicating the analysis with a different ranking pre-registered on the RCT registry results in similar findings with a smaller point estimate for the lower bound. The registered ranking is based on field interviews conducted prior to the Ranking Survey. Under this ranking, one identity task for the Kaibarta caste and two identity tasks for the Kela caste are considered higher instead of lower—and similarly for paired control tasks.[42] Appendix Table 1.8 reports the results based on the registered ranking. The total estimated effect of identity violation on take-up is 45 pp with 13 pp as the unbiased estimate of the lower bound. The change is due to some identity tasks associated with a large negative drop in take-up being categorized as higher tasks under this ranking. Since both rankings give qualitatively similar results and the Task Survey provides a more objective measure of caste hierarchy based on a larger survey sample, the remaining analysis uses the ranking based on the Task Survey data.

Identity vs. social image. I examine whether the main findings are better explained by concerns for identity or social image. If workers have similar intrinsic willingness to work on different extra tasks but face concerns about other people's judgments or reactions, the estimated effects on take-up should be concentrated among the workers who expect their take-up decisions to be publicized. In Table 1.3, columns (4)-(6) show how the main results differ depending on the randomized privacy condition. The main covariates are interacted with *Public*, an indicator variable equal to one if the worker is are in the public condition.

The estimated effects are invariant to the randomized privacy condition. The coefficients on Different $\times$ Identity and Lower $\times$ Identity have similar magnitudes as before and are still statistically significant at 1% level, while their interactions with the Public indicator are small and not statistically significant. One may think that the experimental variation is not effective in creating different expectations about observability, i.e. workers in the private condition may also expect their decisions to become known to others. In such a case, the results would indicate a strong effect of social image on labor supply. While this is plausible, a better explanation for the results may be that workers

[42] A more detailed discussion on caste ranking consistency is provided in Appendix Section 1.10.3.

facing concerns about social image regarding certain jobs already consider those jobs to violate their identity. This can explain the lack of additional effects from publicizing worker decisions and is likely given the following reasons.

First, the experimental variation used here is similar to that used in another study in the same setting, which finds social image effects. Breza et al. (2019) show that workers' willingness to take up jobs at wages below the market prevailing rate increases when workers are told their wages will be kept confidential. Given that their study also involved a similar sample of casual laborers, it seems unlikely that this study did not alter worker expectations about confidentiality at all.

Second, workers' stated reasons for refusing job offers are more in line with concerns about identity. During the follow-up survey, workers are asked why they turned down all offers involving a particular task. Figure 1.2 plots which share of answers bring up only the reasons related to identity (e.g. feel ashamed of oneself, lower caste work), only the reasons related to social image (e.g. unacceptable to family or neighbors), both, or neither (e.g. task is difficult, never done the task before). Among those who turn down offers involving identity tasks, the share mentioning only the reasons related to social image are small (7%), compared to the shares mentioning only the identity-related reasons (50%) or both types of reasons (25%).[43]

Third, in this setting, people's personal opinions on caste norms are very similar to what they believe about other people's opinions on caste norms. The Task Survey asked four vignette questions describing characters violating various caste norms.[44]. Randomly selected half of the participants were asked whether they approve of the characters' actions in their personal opinion. The rest were asked whether they think their friends and neighbors would approve. Figure 1.6 shows that the shares of participants who support following caste norms are similar regardless of how the questions are asked. The effect of social image on take-up would be present only if workers personally find it desirable to take up certain offers but believe others will not approve. The general consistency in people's personal opinions and beliefs about others suggests that workers may personally find it

[43]For paired control tasks, many people also bring up feeling ashamed in oneself, suggesting some of these tasks may be too menial, i.e. involve a violation of status identity.

[44]These questions are listed as Q1-Q4 in Appendix Section 1.10.5. Two questions are related to the practice of taking up lower-caste jobs.

undesirable to take up certain offers and believe others will agree with his decisions.

These considerations point to identity as the major driver of the findings. One caveat in this study is that workers always disclose their take-up decisions and opinions to surveyors. If workers were mainly concerned about surveyors' judgments of them, it could result in similar findings. However, this would mean that workers face as severe concerns about judgments by some surveyors as those by their friends and neighbors.

Heterogeneity by caste sensitivity and age. If the identity channel explained the findings, the effects should be larger among those workers who strongly believe in following caste norms. During the follow-up survey, workers were asked seven vignette questions that describe characters violating caste norms and stated their personal opinions on whether they approve of those behaviors.[45] With their answers, I generated a caste-sensitivity score using a principal component analysis (PCA) and categorized those who score above the median as caste-sensitive.

Caste sensitivity is associated with a higher likelihood of turning down job offers associated with lower castes. In Columns (1)-(2) of Table 1.4, the coefficient on Different $\times$ Identity is -0.25, similar to the original estimate, and the coefficient on Caste sensitive $\times$ Different $\times$ Identity is small and not statistically significant. Therefore worker responses to offers associated with other castes are similar regardless of their caste sensitivity. The coefficient on Lower $\times$ Identity is -0.17, which is smaller than the original estimate, and the coefficient on Caste sensitive $\times$ Lower $\times$ Identity is -0.14, and both coefficients are statistically significant at 1% level. These coefficients indicate that caste sensitive workers are especially unwilling to take up offers associated with lower castes. These results are robust to using alternate definitions of caste sensitivity, as shown in Appendix Table 1.9.

A heterogeneity analysis using worker age gives similar results, since caste sensitivity is positively correlated with age (ρ = 0.15). In columns (3)-(4) of Table 1.4, the coefficient on Older $\times$ Lower $\times$ Identity is -0.10 and statistically significant at 5% level, while the coefficients on other interacted variables are not statistically significant. Caste sensitivity is negatively correlated with

[45] The questions are listed in Appendix Section 1.10.5; four of them are from the Task Survey. For example, the questions describe characters getting jobs associated with other castes, serving food to higher caste people, and marrying outside of own caste.

education and wealth.[46] Appendix Table 1.10 shows that less-educated workers are especially unwilling to take up offers associated with lower castes, and the results do not vary much by wealth. Overall, the heterogeneous effects by caste sensitivity show the clearest and most robust patterns.

1.6.3 Alternate explanations

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

Worker education and wealth. Although the heterogeneity analysis above indicates that the results are not driven by more educated or wealthier workers, I consider whether the task-specific effects of education or wealth can explain the results. This would be the case, for example, if higher caste workers are more educated and being more educated is negatively correlated with having task-relevant skills for identity tasks. The summary statistics reported in Appendix Table 1.11 show that workers in the two highest-ranked castes are more educated and wealthier than the rest, although the economic implications of these differences may be small. Workers not in those two castes tend to be statistically indistinguishable.

The main findings are robust to controlling for the effects of education and wealth measures. I modify the regression specifications to control for the interactions of education or wealth measures with task-specific dummies. If being more educated or wealthier has a negative effect on take-up for specific tasks, this specification would control for those effects. The results in Table 1.5 Columns (1)-(5) indicate that that the education and wealth measures play a limited role in explaining the main findings.

Other job opportunities. I do not find evidence that caste groups differ in their access to other job opportunities. Having greater access to other jobs could decrease workers' willingness to engage in specific tasks. However, the summary statistics in Appendix Table 1.11 show that workers report on average getting 2.5 days of paid work in the previous week, and the number is marginally lower only for those in the Mochi and Pana castes (ranked 5th and 6th). In addition, individual workers'

[46] The correlation coefficients with years of education and being wealthier than the median are -0.08 and -0.06, respectively.

opportunity costs are held fixed when workers evaluate different job offers, and the results are robust to controlling for worker fixed effects.

Status. The findings may be explained by differences in worker status; workers may be less willing to take up offers associated with groups that have different social status from their own, and especially so when those groups have relatively lower status. Figure 1.2 shows that some workers report concerns about feeling ashamed or reactions from others as reasons for refusing to work on paired control tasks. This finding is in line with the idea that worker status and how menial a task is can matter for job take-up, independent of caste associations. According to this explanation, the results still document the effects of identity on job take-up, but the relevant factor would be status identity instead of caste identity.

Untouchabitity. The historic and currently illegal practice of untouchability socially segregates groups now known as Scheduled Castes and delegates them the activities that deal with emissions of the human body. This practice is unlikely to be driving the results since six of the caste groups belong to Scheduled Castes. However behavioral rules (norms) that are specific to each caste can be driving the results, and this would be consistent with attributing the results to identity effects.

Expectations about the employer. One may be concerned that workers form different expectations about employers who would hire them to perform different extra tasks, e.g. some may be more discriminatory. This is unlikely given the offers are explicitly advertised as one-time offers coming from the same employer providing work at the same location. In addition, discrimination has been more prevalent against low caste workers (Mosse 2018). Fear of discrimination has difficulty explaining why higher caste workers are especially averse to taking up jobs associated with lower castes.

Surveyor demand effect. Although the recruiting team asked around for worker names and castes during the process of scouting villages, the surveyors who conducted the choice exercises in the days following were careful not to bring up any discussions of caste. The Task and Ranking surveys were also conducted in separate areas from where the experiment was conducted. In addition, this explanation requires that although the employer was described as looking for people to complete extra tasks, the workers in the experiment must still believe that the employer (or the surveyor) wants to only hire specific castes for certain tasks. It seems unlikely that workers would forego job opportunities based on such speculations.

1.6.4 Robustness checks

I examine the randomization outcomes for job offers and discuss job and survey completion rates. I also show that the main findings are robust to using different regression specifications and control variables in Appendix Section 1.10.6.

The randomization process for selecting job offers was implemented successfully. In Appendix Table 1.12, columns (1) and (2) replicate the main take-up results only using the job offers that were randomly selected at the end of the choice exercise. The coefficients on Different × Identity and Lower × Identity are similar to those in Table 1.3, although less statistically significant due to the smaller sample size.

Job completion results also go in the direction of the main findings. Overall, 57% of workers receive offers that they are willing to take up, and 67% of them complete their jobs. The completion rate is not very high, as absenteeism is prevalent in this region (Krishnaswamy 2019). Columns (3) and (4) show how job completion varies with extra task category and relative status. The coefficients on Different × Identity is larger compared to those in columns (1) and (2), showing that job completion is even lower compared to take-up for offers associated with other castes. This additional effect on completion might be also due to identity or social image, e.g. workers thinking more deeply or talking to others about the implications of taking up specific offers. The survey completion rate is high (87%) and does not depend on extra task category and relative status, as shown in columns (5) and (6).

In Appendix Section 1.10.6, I show that the main findings are robust to the following: excluding any one caste, controlling for different time trends, and clustering standard errors in another way. The results are also robust to controlling for other variations in workers' decision making environment—surveyors, orders in which offers are discussed, and choice sets–as well as excluding inconsistent decisions that may involve worker mistakes.

1.7 Supplementary experiment: pricing identity violations

My findings show that identity is an important factor constraining workers' labor supply decisions. Many workers in the sample are willing to forego a valuable income-earning opportunity in order to avoid spending ten minutes on some tasks associated with other castes. The same workers report finding only about 2.4 days of paid work in the week prior to the experiment and generally have little wealth. These results suggest that the utility costs of violating identity could be very high for some workers.

The supplementary experiment aims to directly quantify the utility costs of identity violations in monetary terms; it examines how much workers need to be offered in wages in order for them to engage in tasks associated with other castes. Using job offers similar to those in the first experiment, I document whether workers are willing to take up the offers when offered more wages.

Referring back to the framework, the fixed utility cost of engaging in task k takes the following form:

$$F_k(n_{ik}, c_i) = \quad f_k(n_{ik}) + \beta_k^d \cdot I^d(k, c_i) + \beta_k^l \cdot I^l(k, c_i) \cdot I^d(k, c_i). \tag{1.7}$$

Comparing across worker decisions, one could in principle put bounds on the fixed utility cost of engaging at all in task k relative to task j, $F_k(n_{ik}, c_i) - F_j(n_{ij}, c_i)$, in terms of utilities from money.

As discussed earlier, theories of identity suggest that some workers may be unwilling to engage at all in tasks associated with other castes, regardless of the offered wage. In other words, the costs of engaging in such tasks, β_k^d, may be so high that workers would effectively never perform them in the labor market. This experiment is designed to test whether workers refuse to work on certain tasks despite being offered high wages, which would be strongly indicative of a high value of β_k^d.

1.7.1 Supplementary experimental procedures

A new set of 106 workers belonging to the Kaibarta and Pana castes are recruited. These two castes do not have any associations with the tasks used in this experiment, since its aim is to document worker reactions to tasks that are associated with other caste groups.[47]

Workers get started on a one-day job of producing paper bags, the default task. Then they individually sit with surveyors who inform them about a chance to switch to working on some other extra task for part of the total working time. Similarly to the main experiment, these switching offers involve variations in the extra task's type and time requirement. There are seven different types of extra tasks: three identity tasks, three paired control tasks, and one pure control task. Identity tasks and paired control tasks are the same as those used in the first experiment, defined in Table 1.1. The time required on these tasks may be 10 minutes, 30 minutes or one hour.[48] The extra task is always to be performed in private, but workers hear different scripts about whether or not their choices are going to be publicized to their neighbors depending on their randomized privacy conditions.

A notable difference from the main experiment is that these switching offers can include a bonus payment on top of their daily wage of Rs. 300. The amount of extra wage for any switching offer is to be randomly drawn from the following list: 0, 30, 60, 90, 120, 180, 240, 300, 1500 or 3000. The amount of Rs. 3000 is close to a month's worth of wage earnings during agricultural lean seasons.

After doing a practice choice exercise, workers participate in the offer choice exercise. They go over the entire 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 one switching offer and one extra wage amount is randomly selected, and the worker's choice for this combination is implemented.

An alternate design could involve promising job offers as in the first experiment, but offering much higher wages. Both sets would provide bounds on the differences in utility costs of engaging

[47] The two caste groups are still compared in the results section below. The sample breakdown is described in Appendix Section 1.10.4.

[48] To reduce the length of this choice exercise, only one pure control task and three time variations are used. The pure control task is moving bricks, a task frequently performed in construction by casual laborers. It was chosen because it would not involve any identity concerns but could require high effort cost, which would allow for useful comparisons.

in extra tasks (and these bounds would be the same if the utility in money was increasing linearly over the relevant wage range). However, because the supplementary experiment requires offering workers much larger sums of money (e.g. ten times the prevailing market wage), it is critical to convince workers that the offers are real. This is more likely to hold under the procedures I used because they involved workers coming to worksites, meeting their supervisors, and being promised a base wage prior to being offered large sums.[49]

1.7.2 Results: responses to extra wage

The results indicates a striking divide in worker responses regarding switching offers. Panel A of Figure 1.3 plots the average take-up rates of offers against the time required on the extra tasks, separately for an extra wage of Rs. 30 and Rs. 3000. The graph for paired control tasks shows that over 60% of workers are willing to switch to working on these tasks for ten minutes when Rs. 30 is offered. When the time requirement increases to one hour, the rate falls, consistent with the idea that working on these tasks involve time-dependent variable effort costs. When Rs. 3000 is offered, over 80% of workers are willing to switch to these tasks for ten minutes. The slope becomes more flat, indicating that for those who refuse Rs. 3000, the variable costs of working matters little.

According to Appendix Figures 1.7, the task-specific trends in take-up are such that as time requirement goes to zero minutes, take-up rates would get close to the perfect rate of one for most control tasks (except for sweeping animal sheds). This is consistent with the assumption that variable effort costs of working on tasks are close to none when workers spend very little time on them. This also suggests the three control tasks—moving bricks, washing farming tools, and mending grass mats—do not involve any identity concerns. 20% of workers refuse to sweep animal sheds for ten minutes even when offered Rs. 3000. As this task is typically considered a menial task (according to field interviews), this task may involve other types of identity concerns (e.g. status identity).

Going back to Panel A of Figure 1.3, the graph for identity tasks shows that around 40% of

[49]These procedures also had logistical advantages, which allowed me to conduct more choice exercises per day.

workers are willing to work on identity tasks when offered Rs. 30. At the same time, 43% of workers refuse to switch to identity tasks even when offered Rs. 3000.

Panel B of the same figure highlights this stark division in worker responses. The histograms in this panel plot the minimum extra wage amounts at which workers agree to switching offers. Those who refuse all offers regardless of wages offered are plotted in the bin, "≥ 3000." I focus on the offers that involve spending ten minutes on extra tasks, as these offers are expected to involve the smallest variable effort costs. Both histograms are double-peaked, first at Rs. 0 - 30 and then at Rs. 3000, and the shares of workers who demand something in between are relatively small. Notably, the bar at Rs. 3000 is more than twice as tall for identity tasks compared to paired control tasks. These patterns are also clearly shown in the task-specific graphs in Appendix Figures 1.8. These patterns suggest that some workers incur small fixed utility costs of engaging in the extra tasks, whereas the majority of the remaining workers incur extremely large utility costs.

As the workers who turn down switching offers regardless of extra wage amounts are likely facing concerns about identity (or social image), the results in Table 1.6 examine the shares of such workers using regressions. Columns (5) and (6) show that the share of workers who demand more than Rs. 3000 for switching (i.e. refuse all offers) is 29 pp larger for identity tasks compared to paired control tasks. Furthermore, this estimate does not vary with whether workers' decisions are publicized, as indicated by the results in Columns (7) and (8).[50]

These findings show that some workers refuse to work on tasks associated with other castes even when they are offered ten times daily wage to do so. This is consistent with the idea presented by Bénabou and Tirole (2011) that concerns about identity can lead to taboo-like behaviors. Workers' reported reasons for refusing offers even at Rs. 3000 are also consistent with this idea. All workers who refused offers involving moving bricks (a pure control ask) said the task is too difficult for them (particularly due to health problems). On the other hand, the reasons workers cited for turning down offers associated with other castes relate to feeling shame in themselves, caste-related concerns, and simple lack of will (e.g. I would never do this task).

[50] In addition, Appendix Table 1.13 shows that the share of workers who refuse all offers is 6 pp larger when tasks are associated with lower castes, although not statistically significant.

1.8 Conclusion

This project's findings indicate that caste identity can constrain labor supply decisions in rural Odisha, India. In the job take-up experiment, the average take-up rate of the offers associated with other caste groups is 23 pp lower than the offers associated with the individual's own caste. This gap increases by an additional 24 pp if the groups associated with the relevant offers rank lower than the individual's own in the caste hierarchy. This latter increase is especially larger among those who strongly believe in following caste norms. These provide the first experimentally documented, quantitative estimates of identity effects on labor supply. Responses to job offers are invariant to whether worker decisions are publicized, which strongly suggests that identity—rather than social image—is the main motivating factor.

These findings highlight a channel through which occupational opportunity can become unevenly distributed across groups. The results from the supplementary experiment suggest that some workers will completely avoid certain jobs despite those jobs offering much higher wages. Such responses are in line with the idea that the value of identity could be "priceless," a key implication of the theoretical model on identity by Bénabou and Tirole (2011). If some groups of workers have the same inherent talent or existing skill sets as the rest, but avoid certain occupations due to concerns for identity, it would cause misallocation of talent in the economy. This identity channel has typically been omitted in the existing economic models on misallocation of talent such as that of Hsieh et al. (2019). This paper points towards the importance of accounting for the role of identity in studying inefficiency in the economy.

1.9 Figures and tables

Figure 1.1: Raw take-up rates

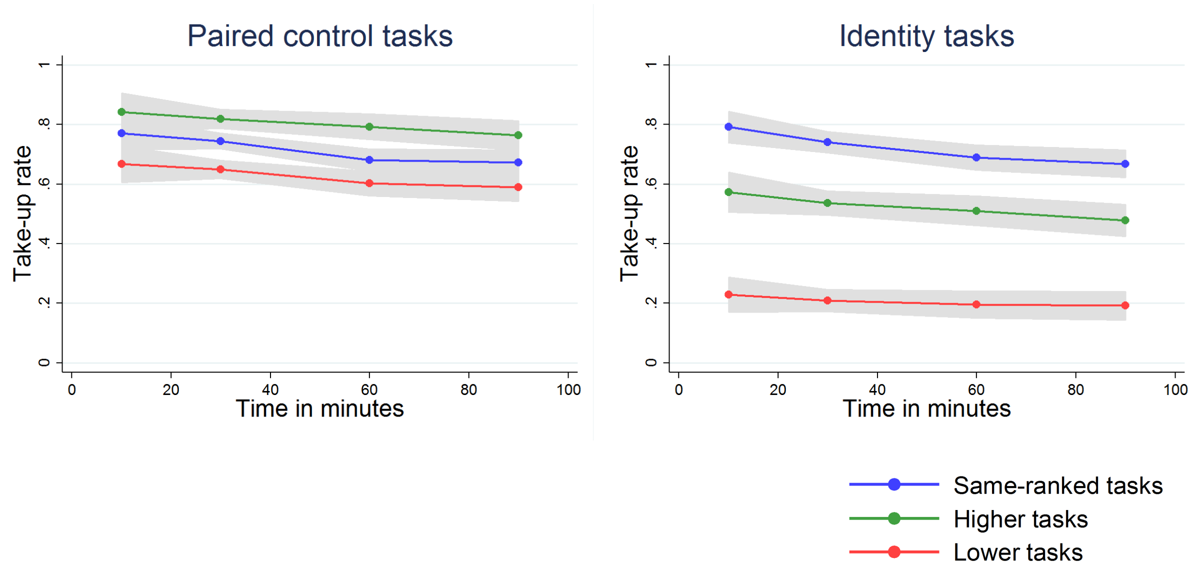

Notes. The average take-up rates of the job offers are plotted against the amount of time required on the extra tasks. The plotted average take-up rates are calculated separately by task category (paired control tasks on the left vs. identity tasks on the right), and also by relative task status as indicated by the three separate lines in each graph. The relative task status (same ranked, higher, or lower tasks) is determined based on the rank scores in Table 1.1).

Figure 1.2: Reasons for turning down offers

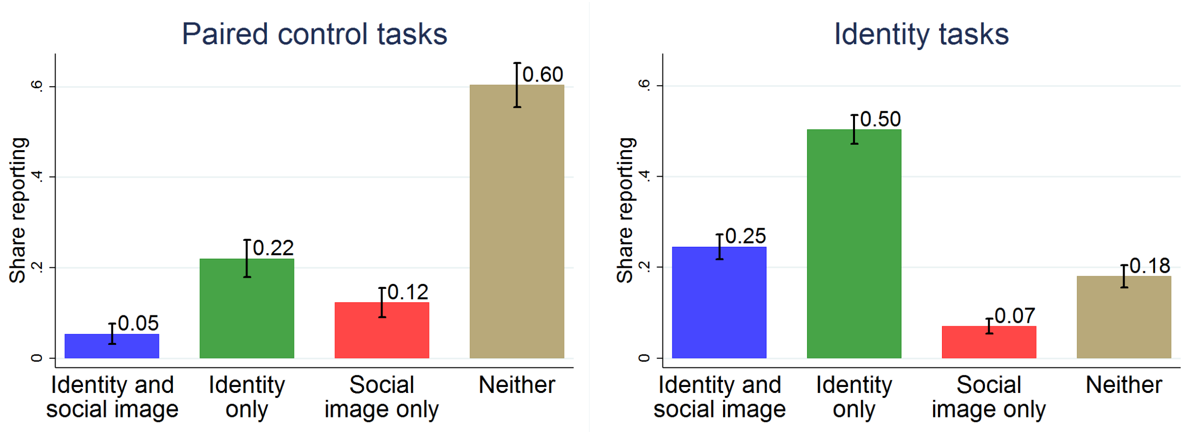

Notes. This figure plots the shares of workers reporting particular reasons for refusing job offers. Observations are at the $worker \times task$ level, i.e. worker answers regarding a set of offers involving a particular extra task type, and hence include 391 answers for the paired control tasks and 953 answers for the identity tasks. Reasons relating to identity include: this is lower caste work, I would feel ashamed of myself, and this is against cultural practice. Reasons relating to social image mention reactions from other people: family/neighbors will find it unacceptable, they will be embarrassed or upset, and so on.

Figure 1.3: Worker responses to switching offers

Panel A: take-up rates of switching offers

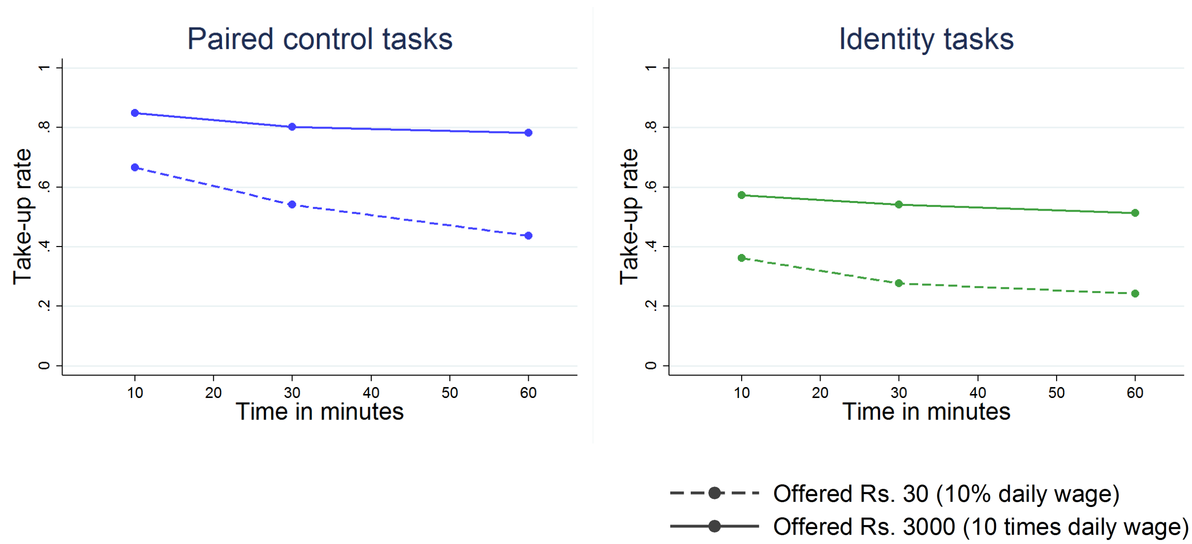

Panel B: Distributions of extra wage demanded for spending 10 minutes on extra tasks

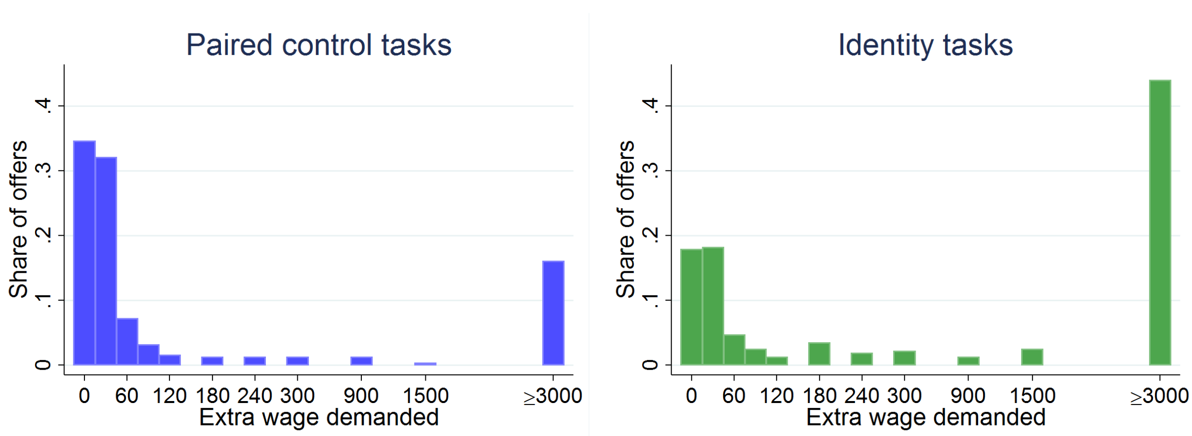

Notes. In *Panel A*, the average take-up rates of switching offers in the supplementary experiment are plotted against the amount of time required on extra tasks. The plotted average take-up rates are calculated separately by task category, and also by size of extra wage offer as indicated by the two separate lines in each graph. In *Panel B*, the minimum extra wage offer at which workers agree to spend 10 minutes on extra tasks are plotted, separately by task category. The bars at 3000 include the cases in which workers switch for Rs. 3000 (0.9% and 1.2% of the offers respectively), as well as the cases in which workers refuse all offers.

Table 1.1: 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 simple tasks. The caste names in Column (1) are sorted according to the mean rank scores of the castes, which are reported in Column (2). In Column (3), the "identity" tasks that have specific caste associations are listed in the same rows as the caste names, e.g. Dhoba is associated with washing clothes. Column (4) reports the share of the survey participants who report such connections. Column (5) lists the "paired control" tasks that involve similar skills as those Column (3) in the same rows. No participant reported connections between the paired control tasks with the specific castes in the same rows, so Column (6) shows instead the share of participants who report association between the paired control tasks with any Scheduled Caste (SC). Using this table, the relative status of the tasks are determined. Given any caste group, the tasks that are in the same rows are called "same-ranked tasks", whereas those that appear in the other rows are called "different tasks". In particular, those that appear in the higher (lower) rows than the caste group are considered "higher (lower) tasks". All pure control tasks are categorized as higher tasks.

Table 1.2: Experiences with tasks

Dependent var. =	Ever performed			Performed without wage			Performed for wage		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Different task	-0.255***	-0.070	-0.073	-0.279***	-0.084	-0.087	-0.017	-0.020	-0.020
	(0.096)	(0.099)	(0.097)	(0.097)	(0.101)	(0.103)	(0.074)	(0.079)	(0.079)
Different × Identity	-0.044	-0.225*	-0.222*	0.118	-0.052	-0.049	-0.283**	-0.277**	-0.276**
	(0.120)	(0.116)	(0.120)	(0.142)	(0.141)	(0.140)	(0.127)	(0.130)	(0.127)
Lower task	-0.026	0.008	0.009	-0.012	0.006	0.009	-0.068**	-0.012	-0.012
	(0.064)	(0.058)	(0.058)	(0.065)	(0.060)	(0.061)	(0.032)	(0.034)	(0.034)
Lower × Identity	-0.107	0.050	0.050	-0.114	0.050	0.048	0.068**	0.036	0.036
	(0.090)	(0.063)	(0.063)	(0.090)	(0.065)	(0.065)	(0.032)	(0.035)	(0.036)
Identity task	0.100			-0.050			0.200		
	(0.094)			(0.120)			(0.123)		
Pure control tasks	-0.150**			-0.144**			-0.061*		
	(0.059)			(0.060)			(0.031)		
Mean: same-ranked									
control tasks	0.850	0.850	0.850	0.850	0.850	0.850	0.100	0.100	0.100
identity tasks	0.950	0.950	0.950	0.800	0.800	0.800	0.300	0.300	0.300
Task FE	No	Yes	Yes	No	Yes	Yes	No	Yes	Yes
Caste FE	No	Yes	No	No	Yes	No	No	Yes	No
Worker FE	No	No	Yes	No	No	Yes	No	No	Yes
R-squared	0.042	0.399	0.506	0.035	0.387	0.492	0.080	0.115	0.189
Observations	1,004	1,004	1,004	1,004	1,004	1,004	1,004	1,004	1,004

Notes. During the Task Survey, participants described to what extent they have performed the tasks that are categorized in Table 1.1. Each regression reports a difference-in-differences style estimate of how workers' prior experience varies with task category (identity vs. paired control) and relative task status (e.g. different, lower). The omitted category includes the same-ranked tasks, and the dependent variable means for the same-ranked tasks are reported in the table footer. Some specifications additionally control for task and caste(/worker) fixed effects, as indicated in the table footer. Standard errors are clustered at the $worker \times task$ level and shown in parentheses.
* $p < .1$, ** $p < .05$, *** $p < .01$

Table 1.3: Predicted identity violations and job take-up

Dependent var. = Willingness to take up job offer						
	(1)	(2)	(3)	(4)	(5)	(6)
Different task	0.059*	-0.053	-0.053**	0.054	-0.058	-0.053
	(0.031)	(0.033)	(0.025)	(0.044)	(0.045)	(0.034)
Different × Identity	-0.251***	-0.233***	-0.233***	-0.242***	-0.223***	-0.223***
	(0.046)	(0.046)	(0.037)	(0.064)	(0.065)	(0.051)
Lower task	-0.124***	0.065**	0.065***	-0.094***	0.096***	0.086***
	(0.022)	(0.028)	(0.022)	(0.029)	(0.034)	(0.028)
Lower × Identity	-0.205***	-0.238***	-0.238***	-0.221***	-0.253***	-0.253***
	(0.033)	(0.035)	(0.026)	(0.045)	(0.046)	(0.035)
Identity task	0.000			-0.012		
	(0.038)			(0.053)		
Public × Different				0.010	0.010	0.000
				(0.062)	(0.060)	(0.048)
Public × Different × Identity				-0.018	-0.019	-0.019
				(0.091)	(0.091)	(0.072)
Public × Lower				-0.059	-0.060	-0.040
				(0.041)	(0.041)	(0.035)
Public × Lower × Identity				0.032	0.030	0.030
				(0.062)	(0.061)	(0.046)
Public × Identity				0.023	0.026	0.026
				(0.075)	(0.075)	(0.061)
Mean: same-ranked						
control tasks	0.717	0.717	0.717	0.717	0.717	0.717
identity tasks	0.722	0.722	0.722	0.722	0.722	0.722
Time controls	Yes	Yes	Yes	Yes	Yes	Yes
Task FE	No	Yes	Yes	No	Yes	Yes
Caste FE	No	Yes	No	No	Yes	No
Worker FE	No	No	Yes	No	No	Yes
R-squared	0.200	0.223	0.498	0.202	0.225	0.498
Observations	20,160	20,160	20,160	20,160	20,160	20,160

Notes. Each regression reports a difference-in-differences style estimate of how worker willingness to take up job offers varies with task category and relative task status, similarly to those in Table 1.2. *Public* indicates that worker is in the public condition. The coefficients on *Pure control*, *Public*, and their interaction variable are not displayed. All regressions control for task-specific linear time trends. Standard errors are clustered at the $worker \times task$ level and shown in parentheses. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 1.4: Heterogeneity by caste sensitivity and age

Dependent var. = Willingness to take up job offer	(1)	(2)	(3)	(4)
Different task	-0.034	-0.024	-0.061	-0.064**
	(0.043)	(0.033)	(0.043)	(0.033)
Different × Identity	-0.251***	-0.251***	-0.213***	-0.213***
	(0.055)	(0.043)	(0.053)	(0.041)
Lower task	0.070*	0.056*	0.022	0.038
	(0.038)	(0.030)	(0.038)	(0.030)
Lower × Identity	-0.173***	-0.173***	-0.167***	-0.167***
	(0.051)	(0.037)	(0.049)	(0.036)
Caste sensitive × Different	-0.017	-0.039		
	(0.050)	(0.044)		
Caste sensitive × Different × Identity	0.038	0.038		
	(0.053)	(0.039)		
Caste sensitve × Lower	-0.025	0.004		
	(0.043)	(0.038)		
Caste sensitive × Lower × Identity	-0.138**	-0.138***		
	(0.065)	(0.049)		
Caste sensitive	0.078**			
	(0.039)			
Older × Different			0.058	0.073
			(0.051)	(0.045)
Older × Different × Identity			-0.065	-0.065
			(0.054)	(0.041)
Older × Lower			0.046	0.008
			(0.042)	(0.038)
Older × Lower × Identity			-0.103	-0.103**
			(0.066)	(0.050)
Older			0.051	
			(0.040)	
Time Controls	Yes	Yes	Yes	Yes
Task FE	Yes	Yes	Yes	Yes
Caste FE	Yes	No	Yes	No
Worker FE	No	Yes	No	Yes
R-squared	0.230	0.502	0.234	0.504
Observations	17,632	17,632	17,632	17,632

Notes. This table reports results on heterogeneity by caste sensitivity and age. The follow-up survey contained seven vignette questions on caste norms, listed in Appendix Section 1.10.5. *Caste sensitive* indicates that worker expressed stronger support for abiding by caste norms, i.e. his sensitivity PCA score is greater than the median. *Older* indicates that worker age is greater than the median. The specifications are similar to those in Columns (5)-(6) in Table 1.3. Standard errors are clustered at the $worker \times task$ level and shown in parentheses. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 1.5: Effects not explained by education or wealth

Dependent var. = Willingness to take up job offer	(1)	(2)	(3)	(4)	(5)
Different task	-0.041	-0.041	-0.038	-0.040	-0.038
	(0.026)	(0.027)	(0.027)	(0.027)	(0.027)
Different × Identity	-0.233***	-0.233***	-0.240***	-0.234***	-0.240***
	(0.039)	(0.039)	(0.039)	(0.039)	(0.039)
Lower task	0.062***	0.064***	0.058**	0.060**	0.060**
	(0.024)	(0.024)	(0.024)	(0.024)	(0.024)
Lower × Identity	-0.245***	-0.247***	-0.238***	-0.244***	-0.238***
	(0.028)	(0.028)	(0.029)	(0.029)	(0.029)
Task FE interactions	High edu.	Years of edu.	High wealth	Wealth PCA score	High edu. and high wealth
Time Controls	Yes	Yes	Yes	Yes	Yes
Task FE	Yes	Yes	Yes	Yes	Yes
Caste FE	No	No	No	No	No
Worker FE	Yes	Yes	Yes	Yes	Yes
R-squared	0.502	0.503	0.501	0.501	0.503
Observations	17,600	17,600	17,632	17,632	17,600

Notes. This table examines whether the job take-up results are robust to controlling for worker differences in education and wealth. Each specification is similar to that in Column (3) of Table 1.3, but additionally controls for the interaction variables between task-specific dummies and variables describing worker characteristics. *High age* and *High wealth* indicate worker age and wealth PCA score are greater than their respective medians. Standard errors are clustered at the $worker \times task$ level and shown in parentheses. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 1.6: Predicted identity violations and extra wage demand

Dependent var. =	Demand more than Rs. 30 (10% daily wage)				Demand more than Rs. 3000 (10 times daily wage)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Identity	0.322***	0.322***	0.302***	0.302***	0.291***	0.291***	0.272***	0.272***
	(0.037)	(0.028)	(0.049)	(0.034)	(0.032)	(0.024)	(0.040)	(0.029)
Public × Identity			0.041	0.041			0.039	0.039
			(0.063)	(0.038)			(0.061)	(0.041)
Public			0.067				0.145***	
			(0.041)				(0.032)	
Mean: control tasks	0.468	0.468	0.468	0.468	0.170	0.170	0.170	0.170
Time controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Caste FE	Yes	No	Yes	No	Yes	No	Yes	No
Worker FE	No	Yes	No	Yes	No	Yes	No	Yes
R-squared	0.106	0.592	0.114	0.592	0.159	0.600	0.191	0.601
Observations	2,226	2,226	2,226	2,226	2,226	2,226	2,226	2,226

Notes. This table shows how the minimum extra wage worker demands for spending time on the extra tasks differ across identity and control tasks. The dependent variable in Columns (1) and (2) indicates that the minimum wage worker demands is greater than Rs. 30, whereas that in Columns (3) and (4) indicates that the demanded wage is greater than the maximum offer of Rs. 3000, i.e. workers turns down all offers. Columns (3), (4), (7) and (8) examine whether the results differ by the privacy condition. Standard errors are clustered at the $worker \times task$ level and shown in parentheses. * $p < .1$, ** $p < .05$, *** $p < .01$

1.10 Appendix A. Supplementary notes

1.10.1 Notes on the conceptual framework

Section 1.2.2 outlines how to get the lower bounds on the shares of workers for whom β_{ik}^{d} and β_{ik}^{l} are positive by observing workers' offer take-up decisions.

Suppose worker i evaluates two job offers like the one described in the framework, each against the worker's outside option. The worker will take up the job offer involving additional task k if the utility from the offer exceeds that of his outside option, O_i. The worker's take-up decision is determined as follows:

$$takeup_{ik}(t_k) = \begin{cases} 1, & \text{if } M_i(w) + L_i(1-T) \\ & \quad -[V_{ij}(c_i, t_j) + F_{ij}(c_i) + V_{ik}(c_i, t_k) + F_{ik}(c_i)] > O_i \\ 0, & \text{otherwise} \end{cases} \tag{1.8}$$

The variable utility cost of working on task k (or any other task) is assumed to be zero when the worker does not spend any time on that task. It is also assumed to be continuous in time as follows.

Assumption 1.10.1. *The variable cost function $V_{ik}(c_i, t_k) : R \times [0,1] \rightarrow R$ for any task k is continuous in t from the right at 0, from the left at 1, and from both sides for all $t \in (0,1)$. Furthermore, $V_{ik}(c_i, 0) = \lim_{t_k \to 0^+} V_{ik}(c_i, t_k) = 0$.*

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 $V_{ij}(c_i, T-\epsilon) \approx V_{ij}(c_i, T)$ for all $\epsilon < \bar{\epsilon}$. That is, when a worker spends almost no time on task k, the time-varying utility cost of working on task k would be close to 0. In addition, the time-varying utility cost of working on the default task in this case would similar to that of spending the entire working time on it.

Substituting for t_k with ϵ in Equation 1.8 and rearranging:

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

Suppose the worker also evaluates a job offer which involving spending the entire working time on the default task j. The take-up decision of this offer is given by:

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

Comparing $takeup_{ik}(0)$ and $takeup_{ik}(\epsilon)$, one can see that the difference in the two take-up decisions would be almost entirely due to the fixed utility cost of working on task k, $F_{ik}(c_i)$. Let θ_i be the utility from taking up the job of only working on the default task:

$$\theta_i \equiv M_i(w) + L_i(1-T) - O_i - [V_{ij}(c_i, T) + F_{ij}(c_i)].$$

Worker i takes up the job of only working on the default task but declines the job of spending ϵ amount of time on task k when the following holds:

$$F_{ik}(c_i) > \theta_i \geq 0.$$

Hence the shares of workers in groups A and B who decline the offers involving tasks b and u

are given by:

$$
\begin{aligned}
\delta_{A,b} &= \sum_{i\in A} \mathbb{1}[f_{ib} + \beta^d_{ib} > \theta_i]/N_A \\
\delta_{A,u} &= \sum_{i\in A} \mathbb{1}[f_{iu} > \theta_i]/N_A \\
\delta_{B,b} &= \sum_{i\in B} \mathbb{1}[f_{ib} > \theta_i]/N_B \\
\delta_{B,u} &= \sum_{i\in B} \mathbb{1}[f_{iu} > \theta_i]/N_B.
\end{aligned}
\tag{1.9}
$$

There are a number of different assumptions that would allow one to infer the lower bound on the share of workers for whom β^d_{ib} is positive by comparing observed $\hat{\delta}$'s. For example, suppose one of the following held:

$$
\begin{aligned}
(f_{ib})_{\{i\in A\}} &\stackrel{\text{dist}}{=} (f_{iu})_{\{i\in A\}} \\
(f_{ib} - \theta_i)_{\{i\in A\}} &\stackrel{\text{dist}}{=} (f_{ib} - \theta_i)_{\{i\in B\}}.
\end{aligned}
\tag{1.10}
$$

Then one could estimate the lower bound by either by comparing take-up rates across the two types of offers within group A or by comparing take-up rates of only the offers involving task b across the two groups. However, these assumptions may be difficult to satisfy with tasks and worker groups in the real world. Assumption 1.2.1 is less-restrictive and is likely to hold in the experimental setting.

1.10.2 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. Caste members maintain close intra-group connectedness through the tradition of endogamy—strictly marrying within castes—and caste networks continue to influence many spheres of Indian life even to this day (Munshi 2019).

The hierarchy embedded in the caste system is easily recognizable in political, economic, and social spheres. 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's). As in the traditional hierarchy, FC's are considered to be above Other Backward Castes (OBC's), which are in turn above Scheduled Castes (SC's, formerly the untouchables) and Scheduled Tribes (ST's, marginalized indigenous groups).

Within each of these official categories, castes form an even finer layers of social hierarchy. 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 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 (Marriott 1958; Mahar 1960). These practices serve as frequent reminders of individuals' caste identities as well as their castes' relative social positions.

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 beteween 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 (Munshi and Rosenzweig 2006, 2016; Madheswaran and Attewell 2007; Thorat and Attewell 2007).[51] Other channels through

[51] 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 discrimination.

which caste influences labor market behaviors may include stereotype threat (Hoff and Pandey 2006, 2014), willingness to punish norm violations (Hoff, Kshetramade, and Fehr 2011), and in-group favoritism (Rao 2019; Lowe 2019). This paper suggests people's desire to uphold their caste identity may be another critical channel through which caste affects people's labor supply decisions.[52]

1.10.3 Consistency of caste ranking

The ranking of castes reported in Table 1.1 is consistent across the different versions of instructions. Appendix Table 1.14 shows the results from running OLS regressions of the reported rank on the indicator for each caste name, controlling for whether the ranked caste is the same as the participant's own. Participants tend to rank their own caste higher, but the ranking is the same across all three versions of the instruction. For the pooled regression in Column (1), t-tests reject that the coefficients are the same for any two adjacent castes in the ranking. In Column (2)-(4), where sub-samples are used, only the tests for the Kela castes (comparisons with Dhoba and Mochi) are sometimes not rejected (0.1<p<0.18).

Across the two districts where the survey was conducted, there are greater variations regarding the position of some castes. Column (5) of Appendix Table 1.14 shows that one cannot reject Sundhi and Kaibarta are of the same rank in Nayagarh. In addition, Column (6) shows that the rank scores of Kela and Mochi are statistically indistinguishable in Dhenkanal (p=0.92). Due to these variations, the robustness checks include doing the main analysis, dropping one caste at a time, to ensure that the results are not sensitive to including the Kela caste.

The experimental design registered on the RCT registry included information on caste rankings that also differed slightly from that based on the Ranking Survey. Due to some time constraints associated with agricultural seasons, the Ranking Survey and some initial rounds of the experiment were conducted at the same time, after the Task Survey was completed. The registered ranking was therefore based on field interviews and some pilot ranking data, which resulted in misspecifying

For a comprehensive review, see Munshi (2019).

[52] A number of news articles report that people avoid working as barbers (Gowda 2011-08-20) or sanitary workers (Mohanty and Dwivedi 2018-07-10) due to the strong caste associations of these jobs, despite the large existing demand for such workers.

the positions for two castes, Kaibarta and Kela, whose rank scores are also noisier according to the Ranking Survey. The main analysis therefore report results using both versions of rankings. They yield qualitatively similar results with different magnitudes for estimated identity effects.

1.10.4 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 and surveyors were more successful with completing surveys in certain villages, so there are small deviations from targets.

The sample for the supplementary experiment is described below.

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

1.10.5 Vignette questions related to caste sensitivity

The following questions were used during the follow-up survey to determine caste sensitivity. Participants answered on a 5-point-scale indicating their 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?

1.10.6 Additional robustness checks

My main findings are robust to dropping any one caste. Appendix Table 1.15 shows results are broadly similar regardless of which caste is dropped, and the coefficient on Lower $\times$ Identity is in particular consistent across all specifications.

The results are robust to controlling for alternate time trends. Appendix Table 1.16 Columns (1)-(4) control for linear or quadratic time trends specific to each task given each caste group. The results change very little. Columns (5) and (6) show how the time trends change depending on task category and relative status, by interacting linear time trends with the main covariates. Take-up generally falls slightly with longer time requirement, and varies little by task category and relative status. Take-up actually falls even less with time for lower identity tasks. These results confirm that the large differences in take-up decisions are due to the costs of engaging at all in the different tasks, which are expected to vary due to concerns about identity.

Using an alternate clustering method does not change the main findings. Appendix Table 1.17 shows that the main results do not change when the standard errors are clustered at the worker level.

Finally, the main findings are robust to controlling for other variations in workers' decision making environment. The experiment involves 12 surveyors, 4 orders in which tasks are discussed, 2 orders in which time requirements are discussed, and 2 potential choice sets (only one of 2 pure control tasks are randomly presented). I test whether these variations having different effects on take-up for identity tasks can explain the results. Appendix Table 1.18 Columns (1)-(3) shows that the results are robust to addressing the effects of these variations. In addition, I consider that workers may make mistakes during the choice exercise. A worker's choices regarding a set of offers involving a particular extra task are considered inconsistent if he accepts some offer, while rejecting another offer with a shorter time requirement on the task. Column (4) shows that the results are robust to only using decisions that do not involve such inconsistency.

1.11 Appendix B. Additional figures and tables

Figure 1.4: 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 1.5: Take-up rates by task type

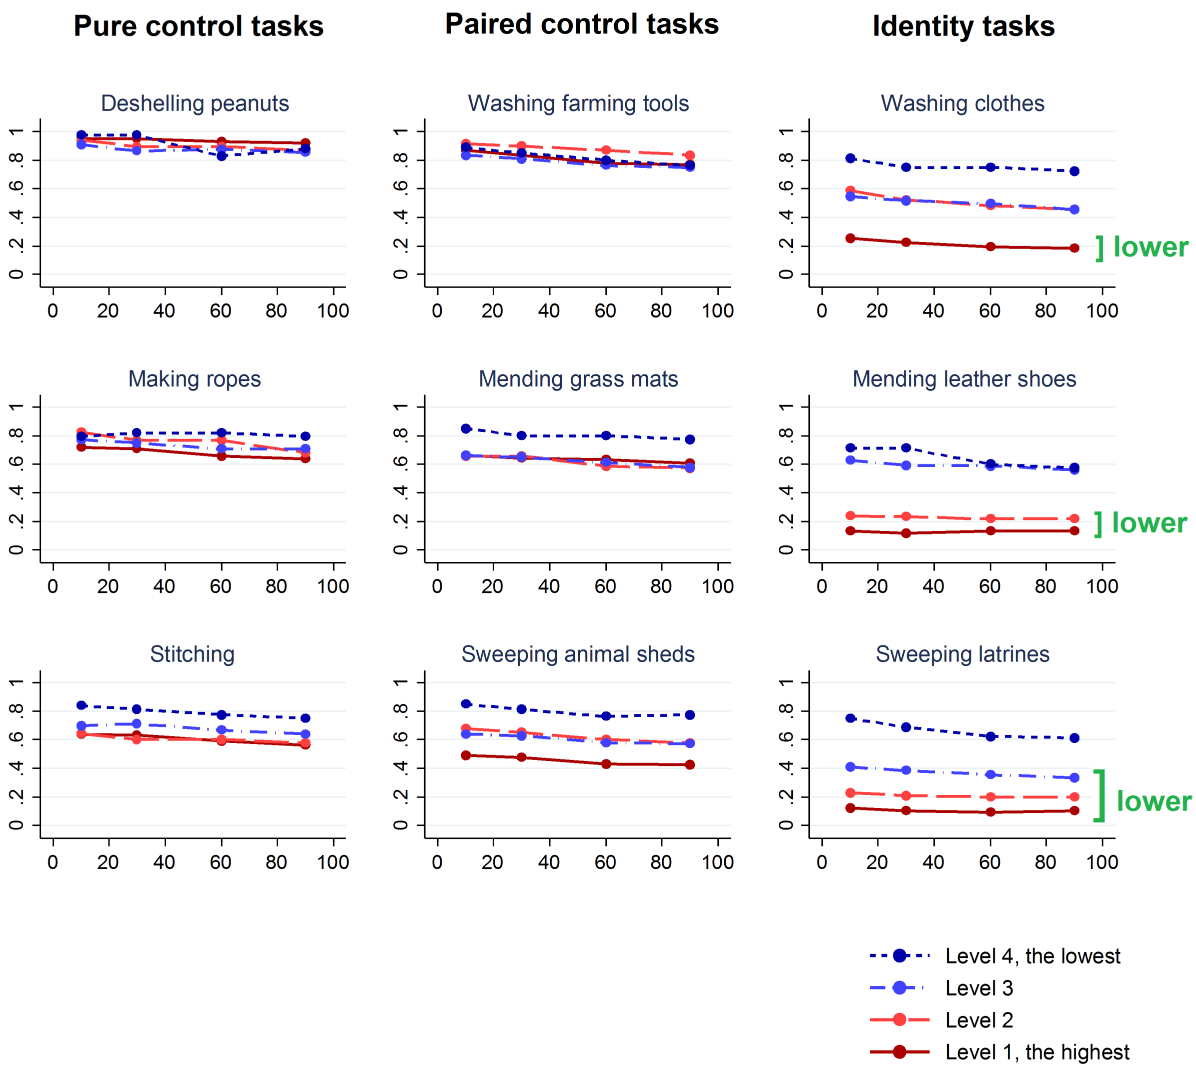

Notes. The average take-up rates of job offers are plotted against the amount of time required on the extra tasks. The plotted average take-up rates are calculated separately for each extra task type and shown in separate graphs. The lines of the same color and pattern in all the graphs refer to the same caste groups. Level 1 refer to Kaibarta and Sundhi, for whom all identity tasks are considered lower tasks. Level 2 refer to Dhoba and Kela, for whom two of the identity tasks are lower. Likewise, Level 3 refer to Mochi and Kela, and Level 4 Hadi. In Column 3, the group for which the identity tasks are considered lower tasks are marked with green brackets.

Figure 1.6: Caste-sensitive opinions of oneself vs. others

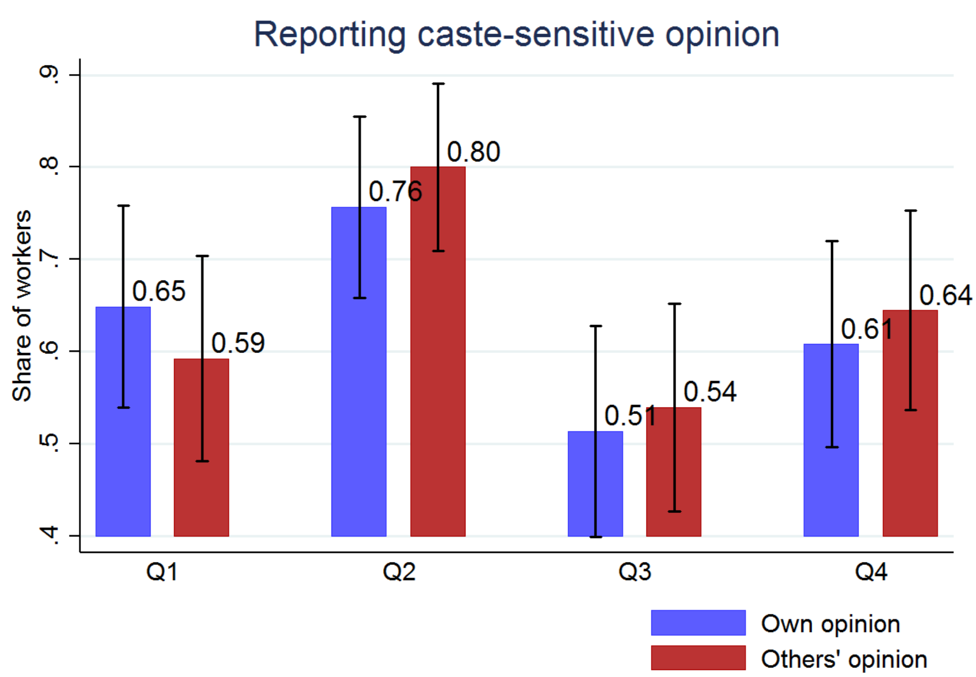

Notes. This figure plots the share of Task Survey participants who express caste-sensitive opinions, either of their own or of their friends and neighbors. There were four vignette questions describing characters violating various caste norms, listed as Q1-Q4 in Appendix Section 1.10.5. Randomly selected 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 graph plots the share of participants who express opinions in favor of abiding by caste norms with 95% confidence intervals.

Figure 1.7: Extra wage offer take-up rates by task type

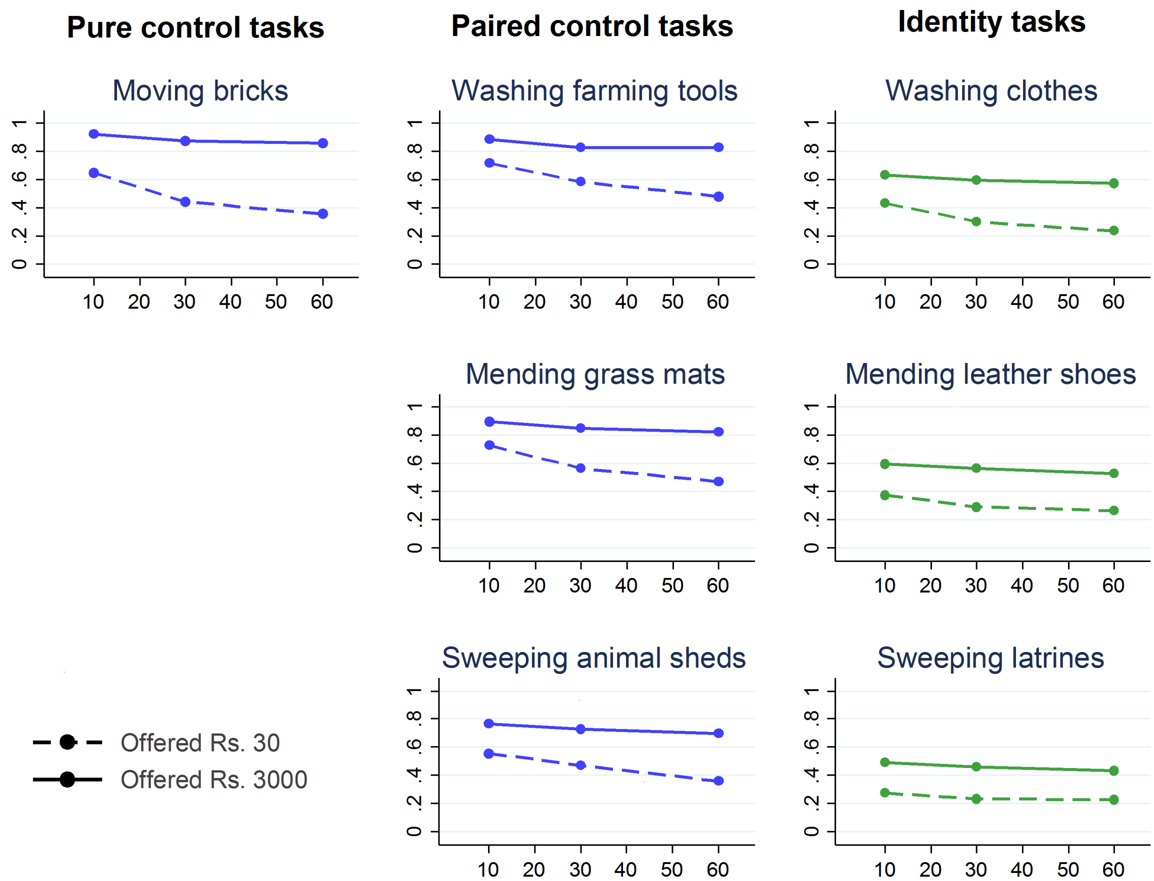

Notes. As in Panel A of Figure 1.3, the average take-up rates of switching offers are plotted against the amount of time required on the extra tasks, separately for each extra task type.

Figure 1.8: Distributions of extra wage demand by task type

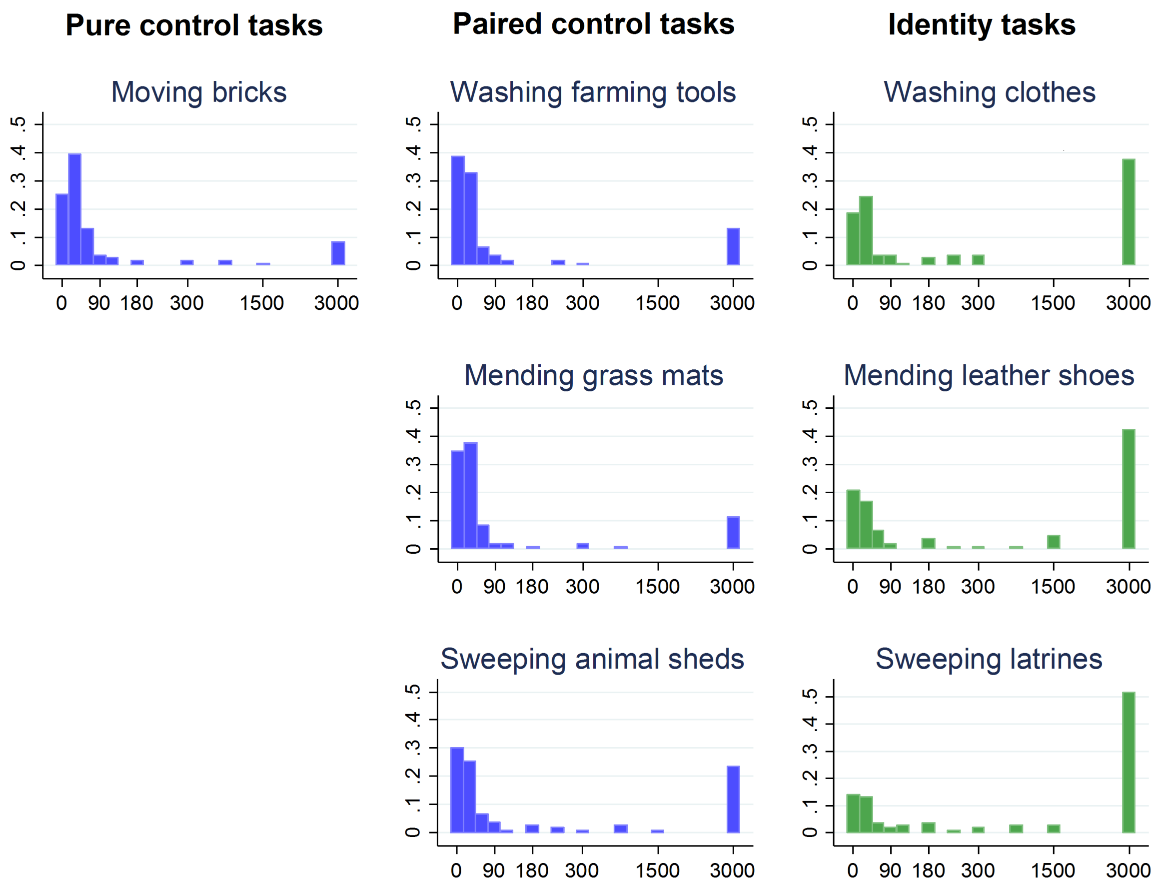

Notes. As in *Panel B* of Figure 1.3, the minimum extra wage at which workers take up offers that involve spending 10 minutes on extra tasks are plotted as histograms, separately for each extra task type.

Table 1.7: Task associations and experiences

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

Notes. This table summarizes the results from the Task Survey, which pertain to the caste and gender associations of the tasks listed in the row headings and the participants' prior experiences with those tasks. Columns (1)-(5) report the shares of participants who associate the tasks with the groups named in the column headings. Columns (6)-(9) report the shares who have prior experiences with the tasks as described in the column headings. Participants may be counted in the shares in Columns (7)-(9) multiple times. The bottom panel shows the results for the additional tasks that were chosen not to be part of the experiment due to having strong associations with women or other caste groups.

Table 1.8: Using registered (different) caste ranking

Dependent var. = Willingness to take up job offer	(1)	(2)	(3)
Different task	0.047	-0.013	-0.013
	(0.030)	(0.032)	(0.024)
Different × Identity	-0.339***	-0.330***	-0.330***
	(0.044)	(0.044)	(0.035)
Lower task	-0.141***	0.018	0.018
	(0.022)	(0.029)	(0.022)
Lower × Identity	-0.097***	-0.120***	-0.120***
	(0.031)	(0.033)	(0.025)
Identity task	-0.001		
	(0.038)		
Mean: same-ranked control tasks	0.717	0.717	0.717
Mean: same-ranked identity tasks	0.722	0.722	0.722
Time controls	Yes	Yes	Yes
Task FE	No	Yes	Yes
Caste FE	No	Yes	No
Worker FE	No	No	Yes
R-squared	0.191	0.218	0.493
Observations	20,160	20,160	20,160

Notes. This table is similar to Table 1.3, but the covariates regarding caste relative status are defined according to the ranking pre-registered on the RCT registry. Compared to the ranking based on the Task Survey, the originally registered ranking—which was based on field interviews—differently categorizes one identity task for the Kaibarta caste and two identity tasks for the Kela caste in terms of *Lower task*. Standard errors are clustered at the $worker \times task$ level and shown in parentheses. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 1.9: Alternate definitions of caste sensitivity

Dependent var. = Willingness to take up job offer	(1)	(2)	(3)	(4)
Different task	-0.027	-0.019	0.001	-0.011
	(0.035)	(0.031)	(0.053)	(0.044)
Different × Identity	-0.240***	-0.263***	-0.283***	-0.277***
	(0.044)	(0.041)	(0.057)	(0.050)
Lower task	0.065**	0.036	0.058	0.064
	(0.032)	(0.029)	(0.048)	(0.040)
Lower × Identity	-0.184***	-0.168***	-0.081	-0.131**
	(0.040)	(0.035)	(0.061)	(0.051)
Caste sensitive × Different	-0.029	-0.060	-0.012	-0.027
	(0.044)	(0.045)	(0.012)	(0.030)
Caste sensitive × Different × Identity	0.013	0.080**	0.013	0.037
	(0.039)	(0.040)	(0.011)	(0.028)
Caste sensitve × Lower	-0.014	0.057	0.000	-0.003
	(0.038)	(0.038)	(0.011)	(0.026)
Caste sensitive × Lower × Identity	-0.097**	-0.183***	-0.041***	-0.090***
	(0.049)	(0.050)	(0.014)	(0.034)
Alternate definition	4 or more sensitive views	5 or more sensitive views	Number of sensitive views	Sensitivity pca score
Time Controls	Yes	Yes	Yes	Yes
Task FE	Yes	Yes	Yes	Yes
Caste FE	No	No	No	No
Worker FE	Yes	Yes	Yes	Yes
R-squared	0.502	0.502	0.502	0.502
Observations	17,632	17,632	17,632	17,632

Notes. This table shows heterogeneity results using by caste sensitivity, using alternate definitions for *Caste sensitive*. The follow-up survey contained seven vignette questions describing characters violating various caste norms (listed in Appendix Section 1.10.5). The table footer describes how *Caste sensitive* is defined using these questions. Standard errors are clustered at the $worker \times task$ level and shown in parentheses. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 1.10: Heterogeneity by education and wealth

Dependent var. = Willingness to take up job offer	(1)	(2)	(3)	(4)
Different task	-0.055*	-0.051	-0.026	-0.043
	(0.032)	(0.040)	(0.032)	(0.027)
Different × Identity	-0.216***	-0.213***	-0.214***	-0.233***
	(0.041)	(0.046)	(0.042)	(0.039)
Lower task	0.079***	0.078**	0.060**	0.063***
	(0.028)	(0.034)	(0.029)	(0.024)
Lower × Identity	-0.301***	-0.348***	-0.250***	-0.245***
	(0.035)	(0.044)	(0.037)	(0.029)
High SES × Different	0.031	0.001	-0.032	-0.017
	(0.045)	(0.006)	(0.045)	(0.015)
High SES × Different × Identity	-0.035	-0.004	-0.048	-0.006
	(0.040)	(0.006)	(0.040)	(0.013)
High SES × Lower	-0.033	-0.002	0.003	-0.000
	(0.039)	(0.005)	(0.039)	(0.012)
High SES × Lower × Identity	0.118**	0.019***	0.025	0.009
	(0.049)	(0.007)	(0.050)	(0.015)
High SES definition	High edu.	Years of edu.	High wealth	Wealth pca score
Time Controls	Yes	Yes	Yes	Yes
Task FE	Yes	Yes	Yes	Yes
Caste FE	No	No	No	No
Worker FE	Yes	Yes	Yes	Yes
R-squared	0.501	0.502	0.501	0.500
Observations	17,600	17,600	17,632	17,632

Notes. This table shows heterogeneity results by education and wealth. Each column is similar to Column (4) of Table 1.4, but uses a different definition of *High SES*, as defined in the table footer. *High education* and *High wealth* indicate that the worker's years of education and wealth PCA score are greater than their respective medians. Standard errors are clustered at the $worker \times task$ level and shown in parentheses. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 1.11: Summary of worker characteristics

	Mean for Level 4	Diff. for Level 3	Diff. for Level 2	Diff. for Level 1
Age	37.440	-0.641	3.163	5.013
	(1.078)	(1.268)	(1.316)*	(1.258)***
Years of education	4.707	0.268	-0.508	1.481
	(0.402)	(0.475)	(0.500)	(0.476)**
Family size	5.053	0.337	0.049	-0.171
	(0.195)	(0.242)	(0.263)	(0.234)
Share of working members	0.373	-0.102	0.002	-0.033
	(0.021)	(0.025)***	(0.026)	(0.025)
Mud house	0.387	-0.123	-0.034	-0.169
	(0.056)	(0.066)	(0.068)	(0.065)**
Owns land	0.373	-0.002	0.031	0.335
	(0.056)	(0.068)	(0.069)	(0.067)***
Land size in acres (if owns land)	0.977	-0.228	-0.294	0.025
	(0.257)	(0.288)	(0.282)	(0.273)
Last month income in Rs.	5350	1794	-29	856
	(287)	(495)***	(402)	(446)
Paid work days last week	2.813	-0.719	0.046	-0.559
	(0.259)	(0.304)*	(0.301)	(0.307)
Number of assets owned	3.307	0.096	-0.287	0.861
	(0.184)	(0.220)	(0.223)	(0.212)***
Wealth PCA score	-0.478	0.409	-0.130	1.359
	(0.167)	(0.207)*	(0.215)	(0.202)***
Number of caste sensitive views	3.760	-0.181	-0.010	0.656
	(0.207)	(0.249)	(0.251)	(0.247)**
Caste sensitivity PCA score	1.214	-0.088	-0.054	0.210
	(0.085)	(0.101)	(0.102)	(0.101)*

Notes. This table summarizes the worker-level variables related to demographics, wealth, and caste sensitivity using the follow-up survey data. Each row reports the coefficients from regressing the row heading variable on the caste level indicators, which are defined in the notes for Appendix Figure 1.5. Robust standard errors are shown in parentheses. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 1.12: Completion rates of actually selected offers

Dependent var. =	Accepted job		Completed job		Completed survey	
	(1)	(2)	(3)	(4)	(5)	(6)
Different task	-0.015	-0.076	0.145	0.076	0.012	0.023
	(0.121)	(0.127)	(0.130)	(0.138)	(0.080)	(0.088)
Different × Identity	-0.267*	-0.284*	-0.493***	-0.491***	-0.028	-0.026
	(0.152)	(0.152)	(0.166)	(0.168)	(0.098)	(0.100)
Lower task	-0.064	0.086	-0.005	0.132	-0.042	-0.011
	(0.075)	(0.093)	(0.087)	(0.104)	(0.049)	(0.063)
Lower × Identity	-0.251**	-0.270**	-0.234**	-0.247**	-0.071	-0.070
	(0.107)	(0.105)	(0.114)	(0.115)	(0.071)	(0.070)
Mean: same-ranked cont. tasks	0.737	0.737	0.316	0.316	0.895	0.895
Mean: same-ranked iden. tasks	0.857	0.857	0.750	0.750	0.964	0.964
Time controls	No	No	No	No	No	No
Task FE	Yes	Yes	Yes	Yes	Yes	Yes
Caste FE	No	Yes	No	Yes	No	Yes
Worker FE	No	No	No	No	No	No
R-squared	0.174	0.213	0.110	0.126	0.033	0.072
Observations	629	629	629	629	629	629

Notes. This table shows the job take-up and completion results only using the offers that were randomly selected at the end of the choice exercise, i.e. using one randomly selected offer per worker. The dependent variables indicate whether worker chose to take up the offer, completed the job, or completed survey. Time controls are not included since time requirements were not determined for some workers; after extra task type was randomly selected, it was obvious that some workers refused all offers involving that task so they did not continue with randomization. One observation is missing for a worker who had to leave the exercise without getting an offer. Standard errors are clustered at the $worker \times task$ level and shown in parentheses. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 1.13: Comparing two caste groups in the supplementary experiment

Dependent var. = Demand more than Rs. 3000 (10 times daily wage)				
	(1)	(2)	(3)	(4)
Lower task	-0.020	-0.020	-0.159**	-0.091
	(0.065)	(0.047)	(0.060)	(0.050)
Lower × Identity	0.062	0.062	0.056	0.056
	(0.079)	(0.050)	(0.081)	(0.056)
Public × Lower			0.285***	0.146**
			(0.054)	(0.052)
Public × Lower × Identity			0.012	0.012
			(0.083)	(0.055)
Kaibarta caste (higher ranked)	0.188***		0.185***	
	(0.045)		(0.044)	
Mean: higher cont. tasks	0.090	0.090	0.090	0.090
Mean: higher iden. tasks	0.310	0.310	0.310	0.310
Time controls	Yes	Yes	Yes	Yes
Task FE	Yes	Yes	Yes	Yes
Caste FE	Yes	No	Yes	No
Worker FE	No	Yes	No	Yes
R-squared	0.164	0.606	0.221	0.610
Observations	2,226	2,226	2,226	2,226

Notes. This table shows difference-in-differences style estimates of how task association and relative task status are linked to worker refusal to take up all switching offers (even those providing Rs. 3000 in extra wage). Standard errors are clustered at the $worker \times task$ level and shown in parentheses. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 1.14: Consistency of caste rank scores

Dependent var. = Reported rank		Question type			District	
	(1)	(2)	(3)	(4)	(5)	(6)
Sundhi	0.576***	0.518***	0.718***	0.491**	0.129	0.995***
	(0.104)	(0.167)	(0.182)	(0.196)	(0.137)	(0.147)
Dhoba	2.234***	2.157***	2.309***	2.241***	1.970***	2.477***
	(0.101)	(0.199)	(0.174)	(0.150)	(0.114)	(0.159)
Kela	2.624***	2.573***	2.678***	2.620***	2.240***	2.983***
	(0.110)	(0.176)	(0.183)	(0.218)	(0.152)	(0.153)
Mochi	3.080***	2.983***	3.199***	3.056***	3.218***	2.967***
	(0.107)	(0.191)	(0.181)	(0.187)	(0.169)	(0.132)
Pana	3.707***	3.714***	3.760***	3.647***	3.810***	3.612***
	(0.093)	(0.160)	(0.166)	(0.164)	(0.133)	(0.132)
Hadi	5.123***	5.047***	5.322***	5.000***	4.950***	5.282***
	(0.087)	(0.157)	(0.122)	(0.174)	(0.123)	(0.122)
Own caste	-0.757***	-0.730***	-0.946***	-0.602***	-0.822***	-0.689***
	(0.112)	(0.187)	(0.211)	(0.181)	(0.176)	(0.151)
Constant	1.556***	1.602***	1.503***	1.562***	1.709***	1.410***
	(0.052)	(0.088)	(0.088)	(0.097)	(0.079)	(0.065)
Sample	All	General	Food-based	Water-based	Nayagarh	Dhenkanal
R-squared	0.674	0.663	0.697	0.664	0.735	0.639
Observations	1,463	490	497	476	700	763

Notes. This table shows the caste ranking results from the Caste Survey. Each observation is a rank score that a participant assigned to a caste. The rank score is regressed on the indicator for each caste group as well as the indicator for whether the ranked caste is the same as the participant's caste. Columns (2)-(4) show the results by the survey question type and Columns (5)-(6) show the results by district.

Table 1.15: Robustness: dropping any one caste

Dependent var. = Willingness to take up job offer							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Different task	-0.046*	-0.041	-0.011	-0.067**	-0.111***	-0.061**	-0.016
	(0.026)	(0.026)	(0.033)	(0.027)	(0.027)	(0.027)	(0.031)
Different × Identity	-0.232***	-0.249***	-0.346***	-0.198***	-0.042	-0.183***	-0.412***
	(0.037)	(0.037)	(0.050)	(0.038)	(0.039)	(0.039)	(0.044)
Lower task	0.055**	0.047*	0.069***	0.076***	0.055**	0.067***	0.086***
	(0.026)	(0.024)	(0.023)	(0.024)	(0.022)	(0.026)	(0.024)
Lower × Identity	-0.234***	-0.194***	-0.220***	-0.266***	-0.239***	-0.306***	-0.224***
	(0.030)	(0.029)	(0.027)	(0.027)	(0.027)	(0.031)	(0.030)
Mean: same-ranked							
control tasks	0.717	0.717	0.629	0.717	0.826	0.717	0.674
identity tasks	0.722	0.722	0.798	0.722	0.637	0.722	0.750
Dropped caste	Kaibarta	Sundhi	Dhoba	Kela	Mochi	Pana	Hadi
Time controls	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Task FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Caste FE	No	No	No	No	No	No	No
Worker FE	Yes	Yes	Yes	Yes	Yes	Yes	Yes
R-squared	0.488	0.490	0.497	0.499	0.508	0.508	0.502
Observations	16,576	17,536	17,120	17,568	18,240	16,320	17,600

Notes. The table shows that the main findings are robust to dropping any one caste. The table footer indicates which caste groups is excluded in each regression. Standard errors are clustered at the $worker \times task$ level and shown in parentheses. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 1.16: Robustness: using alternate time trends

Dependent var. = Willingness to take up job offer	(1)	(2)	(3)	(4)	(5)	(6)
Different task	-0.068*	-0.068**	-0.083**	-0.083**	-0.072**	-0.072**
	(0.035)	(0.030)	(0.037)	(0.034)	(0.035)	(0.029)
Different × Identity	-0.231***	-0.231***	-0.230***	-0.230***	-0.233***	-0.233***
	(0.049)	(0.042)	(0.051)	(0.048)	(0.048)	(0.041)
Lower task	0.083***	0.083***	0.096***	0.096***	0.071**	0.071***
	(0.030)	(0.026)	(0.032)	(0.029)	(0.029)	(0.023)
Lower × Identity	-0.278***	-0.278***	-0.285***	-0.285***	-0.276***	-0.276***
	(0.037)	(0.030)	(0.040)	(0.034)	(0.036)	(0.028)
Time (in hours)					-0.079***	-0.079***
					(0.016)	(0.016)
Time × Identity					-0.013	-0.013
					(0.024)	(0.024)
Time × Different					0.024	0.024
					(0.017)	(0.017)
Time × Different × Identity					0.000	0.000
					(0.027)	(0.027)
Time × Lower					-0.007	-0.007
					(0.009)	(0.010)
Time × Lower × Identity					0.048***	0.048***
					(0.015)	(0.015)
Task-caste specific time control type	Linear	Linear	Quadratic	Quadratic	None	None
Task FE	Yes	Yes	Yes	Yes	Yes	Yes
Caste FE	Yes	No	Yes	No	Yes	No
Worker FE	No	Yes	No	Yes	No	Yes
R-squared	0.242	0.517	0.247	0.522	0.223	0.498
Observations	20,160	20,160	20,160	20,160	20,160	20,160

Notes. The table shows the main findings are robust to using alternate time controls. The specifications in Columns (1)-(4) are similar to those in Columns (2)-(3) in Table 1.3, but control for task-caste specific time trends as described in the table footer. Columns (5)-(6) show how the linear time trends vary with task category and relative status. Standard errors are clustered at the $worker \times task$ level and shown in parentheses. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 1.17: Robustness: clustering errors at the worker level

Dependent var. = Willingness to take up job offer						
	(1)	(2)	(3)	(4)	(5)	(6)
Different task	0.059**	-0.053**	-0.053**	0.054	-0.058	-0.053
	(0.028)	(0.026)	(0.026)	(0.039)	(0.037)	(0.035)
Different × Identity	-0.251***	-0.233***	-0.233***	-0.242***	-0.223***	-0.223***
	(0.043)	(0.043)	(0.044)	(0.058)	(0.058)	(0.059)
Lower task	-0.124***	0.065***	0.065***	-0.094***	0.096***	0.086***
	(0.026)	(0.022)	(0.022)	(0.032)	(0.030)	(0.027)
Lower × Identity	-0.205***	-0.238***	-0.238***	-0.221***	-0.253***	-0.253***
	(0.029)	(0.031)	(0.032)	(0.038)	(0.040)	(0.041)
Identity task	0.000			-0.012		
	(0.035)			(0.048)		
Public × Different				0.010	0.010	0.000
				(0.054)	(0.054)	(0.050)
Public × Different × Identity				-0.018	-0.019	-0.019
				(0.085)	(0.084)	(0.086)
Public × Lower				-0.059	-0.060	-0.040
				(0.044)	(0.044)	(0.036)
Public × Lower × Identity				0.032	0.030	0.030
				(0.052)	(0.052)	(0.053)
Public × Identity				0.023	0.026	0.026
				(0.070)	(0.070)	(0.071)
Mean: same-ranked						
control tasks	0.717	0.717	0.717	0.717	0.717	0.717
identity tasks	0.722	0.722	0.722	0.722	0.722	0.722
Time controls	Yes	Yes	Yes	Yes	Yes	Yes
Task FE	No	Yes	Yes	No	Yes	Yes
Caste FE	No	Yes	No	No	Yes	No
Worker FE	No	No	Yes	No	No	Yes
R-squared	0.200	0.223	0.498	0.202	0.225	0.498
Observations	20,160	20,160	20,160	20,160	20,160	20,160

Notes. This table replicates Table 1.3 using an alternate clustering method. Standard errors are clustered at the *worker* level and shown in parentheses. * $p < .1$, ** $p < .05$, *** $p < .01$

Table 1.18: Robustness: other specification changes

Dependent var. = Willingness to take up job offer				
	(1)	(2)	(3)	(4)
Different task	-0.057**	-0.052**	-0.053**	-0.052**
	(0.025)	(0.025)	(0.025)	(0.026)
Different × Identity	-0.221***	-0.235***	-0.233***	-0.241***
	(0.037)	(0.037)	(0.037)	(0.038)
Lower task	0.065***	0.065***	0.065***	0.065***
	(0.022)	(0.022)	(0.022)	(0.023)
Lower × Identity	-0.235***	-0.237***	-0.238***	-0.236***
	(0.026)	(0.026)	(0.026)	(0.027)
Specification change	Surveyor controls	Task and time order controls	Choice set controls	Drop inconsistent choices
Time controls	Yes	Yes	Yes	Yes
Task FE	Yes	Yes	Yes	Yes
Caste FE	No	No	No	No
Worker FE	Yes	Yes	Yes	Yes
R-squared	0.503	0.500	0.499	0.524
Observations	20,160	20,160	20,160	19,364

Notes. The regressions in this table include additional control variables or have different sample restrictions, as specified in the table footer. The experiment involves 12 surveyors, 4 orders in which tasks are discussed, 2 orders in which time requirements are discussed, and 2 potential choice sets (only one of 2 pure control tasks are randomly presented). Dummy variables representing these variations are interacted with the indicator for identity tasks. Columns (1)-(3) control for these additional control variables. The sample used in Column (4) excludes the cases where worker decisions involve choice inconsistency regarding offers involving specific tasks. Standard errors are clustered at the $worker \times task$ level and shown in parentheses. * $p < .1$, ** $p < .05$, *** $p < .01$

Chapter 2: Does Financial Strain Lower Productivity?

2.1 Introduction

The idea that poverty itself could lower one's income has implications for understanding the persistence of poverty. The focus of the past literature has been on various investments that raise output, including complementary inputs (such as machines or fertilizer), education (schooling or training), and health (nutrition, bed nets, or vaccines). More recent research at the intersection of psychology and economics brings forward another, potentially complementary possibility: Poverty may impact economic behavior by inducing psychological effects such as stress, worries, or sadness.[1] As a consequence, available cash on hand *itself* can make individuals more productive, holding constant any impacts of potential investments facilitated by the availability of cash on hand. However, there remains scant direct evidence this mechanism affects field behaviors and earnings specifically. The main focus of this paper is to experimentally investigate the plausibility of such a relationship.

We run an experiment to test for a causal effect of cash on hand on productivity among small-scale manufacturing workers in rural India. The context is well-suited for studying the impact of financial strain on worker productivity. Workers in this area face severe financial constraints, which are particularly binding during the lean season, the time of our experiment. At baseline, the vast majority of workers report outstanding loans and high levels of worries about their finances. Workers' earnings from the worksite — primarily pieces rates for their output — comprise their primary source of income during the two-week experiment. They are thus highly motivated to exert effort since their productivity is directly linked to their earnings and thus to their consumption and expenditures.

[1] An extensive literature has investigated these channels, including [70, 71, 72, 73, 74, 75].

The study features two experimental variations. First, our main experimental manipulation comprised of temporarily easing workers' cash constraints while holding piece rates and overall wealth constant: Some workers received their wages at the end of their two-week contract while others were paid out a share of their pay five days earlier. The sizable early payments were equivalent to almost one month's worth of typical labor earnings in the lean season. We find that this cash-on-hand intervention meaningfully eased financial constraints. Within three days of receipt of the early cash payments, workers were 40 percentage points more likely to pay off loans to moneylenders and other debtors. Second, following the previous literature, we cross-randomized a psychological priming intervention making financial strain salient to a random subset of workers. These workers were asked in the morning of one workday to recount their outstanding loans and think about how they would come up with a large sum in an emergency. When receiving this intervention, a subset of workers had already been paid a share of their earnings.

Easing workers' cash constraints made them more productive. From the next day after being paid early, worker productivity increased by 5.3% relative to the control group. Productivity increased throughout the work day and persisted for the remaining days of the treatment period. The productivity impacts were concentrated among poorer workers, as measured by several wealth indicators such as land ownership or housing quality. Early payment increased worker productivity for poorer workers by about 10%. These impacts are remarkable given the relatively minor experimental variation in payment schedules and the low wage elasticity of productivity in real-effort settings [76]. Other interventions such as commitment devices at work [77], inducing environmental noise [78], or increasing sleep [79] each caused significantly smaller productivity effects.

The impacts of easing workers' financial constraints only materialized once cash payments were actually made. The timing of payments was scheduled and announced several days in advance, allowing us to explicitly test for announcement effects. We find no evidence of impacts of announcing early payments to workers. That is, learning about their future payment schedule did not alter workers' productivity. Instead, productivity increased only once workers actually received their payment. This result suggests that receiving cash itself was crucial for causing the observed effects, rather than changes in expected future payment streams, consistent with the findings by [72].

We test for and find positive evidence of attention as one underlying channel of the observed treatment effects. Throughout the experiment, we collected markers of workers' attention and focus. These unincentivized measures capture errors causing workers to exert additional effort for the same piece rate. For example, we measure workers' mistakes during production, which they then had to undo, thus slowing them down. Early payment decreased attentional errors by 0.11 standard deviations, an impact that was concentrated among poorer workers, whose attentional errors fell by 0.21 standard deviations. These effects do not rule out other psychological pathways. However, they indicate that the productivity effects are (at least in part) mediated through improved attention during the production process. Indeed, various factors—including worries, stress, or happiness—could affect attentiveness at work.

Our experimental allows us to rule out two sets of alternative explanations for the impact of early payment on productivity, as discussed in detail in Section 2.3.4. First, a set of confounds concerns altered beliefs or perceptions toward the employer, reciprocity, and trust. Such impacts are unlikely to explain the observed effects given the lack of announcement effects and previous work in this literature. Second, we address the possibility of productivity-enhancing investments by workers due to the early payment. Through the design of the experiment, we can rule out investments in physical capital and other longer-run investments such as training. In addition, inconsistent with a nutrition channel, the treatment effects last throughout the work day and we find no impacts of the treatment on caloric intake based on detailed measures of breakfast consumption.

Finally, we find opposing impacts of the psychological priming intervention. Previous work examines the effects of financial priming on laboratory measures of cognition [72]. In contrast, in our setting, workers' actions can affect the problem that is brought top of mind. The hypothesized negative cognition effect on productivity could be offset by a positive motivational effect from making concerns salient, as has been documented in studies of reminders [e.g. 80]. We hypothesized that while the overall level effect of priming is ambiguous, the negative effects would be relatively more important when workers were more financially strained (i.e. before cash receipt). Consistent with this hypothesis, priming leads workers to increase productivity by 7% when they are cash-rich; in addition, there is an offsetting -7% relative productivity effect of priming when workers are cash-poor relative to when they are cash-rich. Unlike the effects of directly manipulating cash receipt, these effects are not driven by poorer individuals; rather, they stem from those who reported worrying more about outstanding loans about at baseline. These patterns highlight challenges with priming interventions discussed in the literature. Put differently, rather than using attention as a treatment (through primes) to uncover mechanisms, our results suggest it might be more effective to use attention as an outcome.

We view the primary contribution of our paper as establishing a direct relationship between cash on hand and worker productivity. This evidence is distinct from traditional theories of how income could affect productivity and earnings [e.g. 81]. Our findings are consistent with findings in other studies that income transfers to the poor can boost labor supply, productivity, and earnings [e.g. 82, 83]. While such studies have focused on traditional physical and human capital channels, our results suggest that psychological pathways could also have the potential to contribute to such effects. This evidence is consistent with work showing that transfer programs affect psychological outcomes such as mental health, [e.g. 71, 84]. Our results, while only an initial proof of concept, provide impetus to explore these possibilities more formally in future work.

Our study also contributes to the nascent but growing literature on the psychology of poverty. Existing work has examined the relationship between poverty and outcomes such as stress [85, 70], cognitive function, and decision making [86, 72, 87, 73, 88], and mental health [71]. Existing work has focused on examining this relationship using laboratory measures and tests or survey self-reports. Our findings lend credence to the view that psychological mechanisms have the potential to affect economic field behaviors. This evidence suggests that cash transfers, social safety programs, and other policies to reduce financial strain among the poor might have important productivity impacts.

The remaining parts of this paper are organized as follows. Section 2.2 provides background information on the study and its experimental design. Section 2.3 discusses the impact of the Early-Pay Treatment on workers' expenditure patterns and their work performance, as well as potential confounds. Section 2.4 discusses positive evidence of one specific channel through which the hypothesized psychological effects might operate: workers' cognition. Section 2.5 concludes.

2.2 Experimental Design

Our experimental design features three key ingredients. First, the experiment takes place in a setting with the potential to impact productivity via psychological channels. In the experiment, we hire piece-rate workers to complete a cognitively challenging production task for which we collect precise measures of hourly productivity for two weeks. Second, we create experimental variation in workers' (perceived) financial situation by varying workers' payment schedules while holding overall payments (approximately) constant across treatment groups. Third, our design allows for the testing of potential confounds in addition to providing positive evidence of the hypothesized psychological channel, attention.

2.2.1 Measuring Worker Productivity

The experiment took place in a low-skill manufacturing environment in rural Odisha, India, using the infrastructure developed by [89]. As part of the field experiment, 408 male workers were employed full-time over the course of 14 rounds. The majority of the rounds (9 rounds) had a standard schedule lasting for 2 weeks with 5 working hours per day.[2] Such seasonal contract jobs are common during the agricultural lean season. The jobs at the study worksites were the primary source of earnings for workers and a regular job from their perspective. The output produced by workers was sold in a local wholesale market.

Work task. Workers produced disposable plates by stitching together sal tree leaves, as depicted in Figure 2.1. This task is relatively cognitively demanding as it requires considerable attention to

[2]Some of the rounds were shorter in length and had different daily work schedules. These changes are described in detail in the appendix.

stitch together the leaves in a way that satisfies the required quality standards. In accordance with quality standards set by partnering contractors, leaf plates were required to (i) meet a minimum size requirement, (ii) have no gaping holes, (iii) have all leafstalks (petioles) covered by other leaves, and (iv) have the inner center parts placed underneath the outer rings of the plates. Workers were paid a flat base wage of Rs. 200 for attendance plus a piece rate of Rs. 3 per completed leaf plate that satisfied the quality standards.[3]

Output measures. We collected hourly measures of output. At the end of each work hour, staff collected completed leaf plates from each worker. Workers were allowed to continue and complete unfinished leaf plates in the subsequent hour. The main measure of output in our study is the number of completed leaf plates, which can be further divided into rejected and accepted output. Workers quickly learned to meet the required standards such that over 96% of leaf plates were accepted by the fourth day of work. The empirical analysis below assumes accepted leaf plates. However, given the high acceptance rates, using the completed number of leaf plates yields nearly identical results.

Attentional errors. In addition to the incentivized measures of production, we collected three unincentivized markers of worker efficiency ("attentional errors") for a subset of work hours. These measures comprise (i) the number of leaves per plate, (ii) the number of stitches per plate, and (iii) the number of "double holes", i.e. instances in which it took a worker several attempts to connect a given set of leaves. For each of these markers, high values indicate inefficiencies that may occur as a result of lapses in worker attention. That is, workers need to increase the required time and effort to complete the leaf plate to receive a given piece-rate compensation for the plate. For instance, conditional on passing the quality threshold (which is the case for nearly all plates), a higher number of leaves per plate means that a worker needs to stitch together a higher number of leaves to be paid for a given plate.

[3]The piece rate in round 1 was Rs. 2 and the base rates were Rs. 180 and Rs. 175 in rounds 2 and 3, respectively.

2.2.2 Inducing Variation in Financial Constraints

Figure 2.2 provides an overview of the timeline of the experiment. The main goal of the experiment was to create and examine an exogenously-induced reduction in financial constraints. To accomplish this goal, we varied the timing of earnings payout across workers while holding piece rates and flat base payments across workers constant, thus altering workers' short-run financial constraints while holding their overall wealth approximately constant.[4]

Common features across workers. With the standard schedule, workers worked for 12 days at their worksite from 9 am to 2 pm.[5] All payments occurred at the end of work days. Workers were informed of their output for each day throughout the experiment, limiting any uncertainty about the outstanding payment amount. On day 1, all workers were paid a flat wage of Rs. 250 for their training and work on this day to foster trust in the worksite among workers. While larger or additional early payments would have been desirable to foster further trust, they would have eased financial constraints among all workers, thus limiting the potential for the experimental variation to create meaningful differences in financial constraints.

In addition, all workers were told on day 1 that they would be paid all of their wages at the latest by day 12. On this day, all workers were paid any outstanding payments they had not been paid for until then. Workers were also paid a completing bonus (Rs. 300) if they attended all of days 6 through 11, so as to avoid selective attendance issues. This completion bonus effectively shuts down any potential labor supply responses to the treatment. Accordingly, the experiment is designed to isolate the impact of cash on hand on worker productivity while holding labor supply constant.

Payment schedule variation. The key experimental variation of the experiment was paying some workers earlier than others. On day 1, we informed all workers everyone would be paid for their

[4]There are two reasons for (expected) wealth differences across treatment groups. First, the Early-Pay Group might save some interest by paying back loans or credits following the early payment. Second, productivity differences due to the early payment translate into differences in worker pay by the end of the experiment.

[5]The deviations from the standard schedule are described in Appendix 2.7.1.

work by day 12, but the payment schedules would differ across workers, with some workers receiving part of the payments earlier than others. On the morning of day 5, workers were given full information about their individual payment schedules. Workers in the control group were told that they would be paid on day 12 (as promised on day 1). Moreover, they were informed that some workers at their worksite would be paid on an earlier day.

In contrast, workers in the Early-Pay Group were told that they would be paid on day 8 (Early Group I) or 9 (Early Group II) of the experiment for their wages earned by the previous day—i.e. wages from days 2 to 7 for Early Group I and days 2 to 8 for Early Group II—and the remaining amount on day 12. While payments were made in private at the end of each day, all workers were aware of payments when they occurred at their worksite.

Anticipated payday variation vs. randomized cash drop. As an alternative way to induce variation in financial constraints we could have randomized unconditional cash transfers across workers. We implemented the current design of anticipated payday variation because it features several advantages. First, it is more realistic than randomized cash drops in this setting since unconditional cash transfers are less common in our study area than in sub-Saharan Africa or other parts of the world. Second, anticipated payday variation is commonplace in developing and developed countries, thus boosting the external validity of our study [90, 91]. Third, the payday variation holds workers' wealth (approximately) constant, thus limiting the potential for effort or labor supply responses due to wealth or income effects.[6]

However, the current design also entails some drawbacks. First, the empirical test comparing workers across treatment groups around the early payday is only powerful if the psychological effects ensue only when workers actually received their payments, as opposed to when they only anticipate receipt of payment. Second, since the employer (our staff) delivered the news of payment variation across workers, the Early-Pay Treatment features an ancillary component of potential changes in workers' relationship with their employer, including gift giving, fairness, and trust in the

[6] Empirical studies have found limited evidence of cash transfers negatively or positively impacting labor supply [92, 93]. However, we would have been underpowered to confirm this evidence in our setting.

employer paying workers as promised. We discuss and address each of these issues in Section 2.3.4.

2.2.3 Recruitment, Sample Description, and Balance Checks

Recruitment. We recruited our study subjects from rural villages in Odisha, where a large number of villagers are engaged in daily wage labor. The study focused exclusively on male workers since it is more culturally appropriate for them to take jobs outside of their village for an extended period of time. A few days prior to the start of a new round of experiment, the recruiters visited the target villages and advertised the upcoming work opportunity through door-to-door visits and fliers. Potential study participants were informed about the location and purpose of the study, the tasks that they would be asked to do, the duration of the study, and their potential compensation. They were also provided with the contact information of the recruiting staff for any questions.

The day before the experimental round, recruiters revisited the villages so that the interested villagers could sign up to participate in the study. During the sign-up process, recruiters used a number of screening questions to determine eligibility. Male workers meeting the following criteria were eligible to participate in the study: (i) aged between 18 and 55, (ii) fluent in Odiya (the local language), (iii) regularly working as daily wage laborers, (iv) having been present in their home villages for more than half of the time during the preceding 6 months, and (v) no prior experience producing leaf plates. In addition, the recruiters verified that potential participants were willing to come to work for the entire duration of the round.

Since usually the number of interested villagers exceeded the worksite capacity per experimental round, we randomly selected approximately 30 participants from the sign-up list to be invited to participate in the study. In addition, 5 back-up participants were selected so that in case some subjects dropped out of the study during the first three days of a round (i.e. before the randomization), they could be replaced with back-up participants.

Sample description and balance checks. The main experiment sample comprises of 408 male workers from 14 experimental rounds with about 30 workers each.[7] 90.7 percent of workers started on the first day of the rounds, and 98.7 percent stayed until the last day of the rounds. Overall, daily attendance was high at 98.7 percent (excluding the days the workers had not yet joined or dropped out of the study before the randomization). 10.1 percent of workers had at least one day of unexpected absence, with 2 percent having 2 or 3 days of absences during their participation in the study. The sample includes 4,094 worker-days and 3,949 non-absent worker-days, with 5 to 7 hourly productivity measures per day.

During workers' first day of work, we conducted a short baseline survey to collect basic demographic and wealth information from workers. Basic sample characteristics and balance checks comparing the Control Group to the Early-Pay Groups are reported in Table 2.1. A typical worker in our sample was about 40 years old, had 4 to 5 years of education and was primarily employed in daily wage labor.

Workers' responses suggest relatively low wealth levels and severe financial constraints, as expected to be case especially during the lean season. Over 70 percent of workers live in houses that contain mud structures, indicating low wealth, and over half have outstanding credits at stores for food and basic household consumption. When asked how concerned they were about their (future) finances, 86 percent of workers indicated any worries and 70 percent reported being very worried about their finances. 68 percent of workers reported outstanding loans, including 18 percent of workers indicating loans from moneylenders charging high interest rates, suggesting their lack of access to other sources of credit.

The baseline characteristics do not statistically differ between the Early-Pay Group and the Control Group, which indicates a successful randomization procedure.

[7] This excludes 21 people who dropped out in the first four days before the payment schedules were announced.

2.3 The Impact of Early Pay

2.3.1 Expenditure Patterns

The early payments provided meaningful amounts of liquidity to workers. These payments comprised almost one month's typical wages during the lean season (over Rs. 1,000 for most workers), given that typical workers worked 8.7 days of wage paying work in the month preceding the experiment. The majority (83%) of workers had outstanding loans at baseline, with a median amount of Rs. 6,000 of debt in our study population. Accordingly, the early payment relieved some of the pressure from indebtedness, but did not eliminate debt for most workers.

Consistent with most workers in the sample facing meaningful credit constraints, the early treatment induced variation in workers' expenditure patterns (Table 2.2). Within the first couple of days of being paid, workers were 40 percentage points more likely to pay off any loan or credit, corresponding to an additional Rs. 278 of repaid loans and credits, an increase of over 200 percent relative to the control group mean of Rs. 122.

Workers also reported changes in other expenditure patterns, most meaningful among those an increase in food expenditures of Rs. 68 relative to a mean of Rs. 270. While these impacts indicate a clear need to consider potential impacts through nutrition channels (which we consider in Section 2.3.4), it is worth noting here that the expenses reported in this survey were often family expenditures (as opposed to individual expenditures) that did not necessarily translate into workers' increased short-run nutritional intake.

2.3.2 Labor Supply Responses

In addition to potential productivity impacts, cash on hand may impact the type of jobs workers engage in, or the number of days or hours worked. Our experiment was designed to capture potential impacts of increased cash on hand on worker productivity while keeping the extensive and intensive margins of labor supply constant.

The completion bonus induced a high overall attendance (98.7 percent), thus limiting the extent of any potential labor-supply response to the experimental variation in payments. Indeed, we find only minimal and statistically insignificant impacts of the early payment on the number of days worked. By the design of the experiment, there was no intensive-margin response to the treatment given that workers came and left their worksite jointly.

2.3.3 Productivity Impacts

Before considering the impacts of the Early-Pay Treatment on productivity, it is worth noting that increasing worker productivity at real-world workplaces is challenging. Similar to other real-effort work settings, effort and productivity were relatively inelastic to wage variation [76].

The Early-Pay Treatment increased worker productivity by 5.3 percent following the early payment (Table 2.3 columns 1 and 2). This treatment effect is economically meaningful, both relative to the wage elasticity in this setting and compared to the productivity impacts of other interventions [94]. For instance, the observed treatment effects of relieving workers' cash constraints were larger than the impacts of offering commitment devices at work [77], exposing workers to considerable levels of noise or heat [78, 95], or increasing their night-sleep and offering them the opportunity to nap [79].

The treatment effects concentrate almost exclusively among poorer workers, as illustrated by heterogeneous treatment effects with respect to baseline wealth (remaining columns of Table 2.3). Wealth measures include whether the individual reported being a landowner, living in a non-mud house, not having outstanding food credits, and whether he reported the ability to come up with Rs. 1,000 easily in case of an emergency, a standard measure of financial health [96]. For each of these indicator variables, we observed large impacts of the Early-Pay Treatment on worker productivity for the poorer part of the sample and significantly lower (and often close to zero) impacts for the richer part of the distribution.

There are two potentially complementary interpretations for the stronger impacts among poorer workers. First, the poor might be experiencing more financial strain (e.g. loans, worries about finances) to start with, thus providing more opportunities for the greater impact of any given intervention that might reduce such strain. Alternatively, most workers in the sample were financially strained, but the intervention was more meaningful for poorer workers since it was larger compared to their wealth.

The impacts on worker productivity proved persistent over several days and throughout the day (Table 2.4). Worker productivity in the Early-Pay Groups increased on each of the three days following the early payment. Similarly, the patterns of heterogeneous treatment effects with respect to wealth persisted for each of these days. Finally, productivity impacts occurred throughout the day, including the last two hours of the workday (see Table 2.7, discussed below).

2.3.4 Potential Confounds

Ancillary Components of Bundled Treatment

Since the employers administered the Early-Pay Treatment, two ancillary components could have contributed to the observed treatment effects on productivity. First, the Early-Pay Treatment may have changed workers' feelings toward their employer, which could have affected their work performance via gift-exchange or fairness concerns. Second, the Early-Pay Treatment could have specifically impacted workers' trust in their employer.

Gift exchange and fairness. Being paid early might have caused the early payment group to feel more positive about their employer. Conversely, the Control Group might have felt unfairly treated by their employer. Several pieces of evidence contradict the hypothesis that such effects caused significant impacts on worker productivity. First, the literature on gift exchange at the workplace has largely found no and/or short-lived effects, especially in field settings [97, 76, 98, 99, 100].

Second, if gift exchange or fairness concerns were important considerations in this setting, we would expect there to be measurable impacts immediately following the announcement. However, we do not find any evidence of positive announcement effects on productivity on day 5 and/or day 6 of the study. The estimates in Table 2.5 show no evidence of significant announcement effects, neither when considering the entire announcement period (columns 1 and 2) nor the announcement day only (i.e. on day 5; columns 3 and 4) nor the two work days following the announcement (columns 5 and 6).

One might explain the absence of such effects by the fact that there were no actual payment differences across workers (unlike in [89]). Moreover, while we did not collect data of workers' demand for the different payment regimes, evidence from other settings suggests that at least some workers prefer more infrequent payments as a method of commitment savings [101], such that the direction of any potential effects was a priori unclear.

Third, we do not find evidence of negative impacts on control-group workers after cash payments were made at the worksite (Table 2.6). Comparing workers in the Control Group with workers in the Early Payment Group II (who were paid at the end of day 9) revealed no productivity differences on the day following payments to workers in the Early Payment Group I (who were paid at the end of day 8). Finally, any explanation involving gift exchange and fairness would also need to address why the effects would concentrate among poorer workers, which might be possible ex post but not what one might expect ex ante.

Trust in the employer. The Early-Pay Treatment could have increased workers' trust in their employers' assurances of future payments. Such increased trust would have increased workers' perceived expected piece rate—the probability of being paid and the piece rate—and thus potentially increasing both effort and output. Several reasons lead us to believe that the observed treatment effects are not explained by such considerations.

First, we designed the payment on day 1 so as to build workers' trust in their employer. Second, as described above, some workers in the Early-Pay Group were paid on day 8; others were paid on day 9 (at the same worksite). If trust in the employer were a major concern among workers, then we would expect workers who were to be paid on day 9 to display an increase in trust towards their employer following the payment of day 8 payees. However, we found no impacts on day 9 for workers who were going to be paid later that day compared to the Control Group (Table 2.6).

Investment Channels

In general, cash on hand can have a variety of impacts on worker productivity, ranging from physical capital (e.g. machines, fertilizer) to human capital (e.g. training, schooling) and health investments (e.g. bed nets, nutrition). By design, the results of our experiment cannot be explained via effects of investments in physical or human capital since there was no scope for workers to bring any of their own physical capital to the worksite. Moreover, any human-capital investments would have taken much longer to come to fruition than the horizon of the experiment.

Nutrition. A long literature in development economics considers the impact of nutrition on worker productivity [81]. We find some evidence of workers increasing their food expenditures following the early payment, as discussed above (Table 2.2). However, meaningful impacts of the Early-Pay Treatment on worker productivity via nutritional channels are unlikely. We consider two categories of potential pathways.

A first potential channel could be biological changes for malnourished workers due to increased food intake. However, according to the biological and medical literatures on the impacts of increased food intake, such changes do not occur overnight. Consistent with this view, [102] finds evidence of increased earnings among workers only starting a week after increasing their caloric intake.

A second potential channel could occur via potential impacts of increased breakfast intake due to blood-sugar spikes. We find clear evidence against such effects. We collected direct measures of breakfast consumption following the early payday. We find no evidence of increased breakfast on any of the dimensions of our survey, including whether workers had breakfast, how much, and what they ate (columns 4 through 8 of Table 2.7). A possible explanation for this lack of impacts on breakfast consumption patterns appears to be the fact that almost all workers (98 percent) in the Control Group reported eating breakfast (thus leaving not much room at the extensive margin), and almost everyone (94 percent) reported eating a particular rice dish that is common in the area (often involving vegetables).

Moreover, we would expect any impacts of blood sugar spikes due to increased breakfast consumption to wear off by the end of the work day. However, we find persistent impacts of the Early-Pay Treatment throughout the day, including the last couple of hours or the workday, i.e. 5 to 7 hours after eating breakfast (columns 1 through 3 of Table 2.7).

2.4 Psychological Channels

Having documented evidence of direct impacts of cash on hand on worker productivity that is not explained by the ancillary components of the treatment or investment impacts, we now investigate the underlying psychological impacts of cash constraints more closely. We first provide positive evidence of impacts on worker attention by considering the above-described measures of attentional errors as a potential channel of reduced worker productivity. Second, we consider the impact of a salience intervention that seeks to bring worries about financial strain to the top of workers' minds and its interaction with the Early-Pay Treatment on worker productivity.

2.4.1 Attentional Errors

The manufacturing task in our setting is a relatively cognitively demanding production task, in particular when compared to other manual labor such as carrying sand bags or loading trucks. To consider whether cognitive impacts contributed to the productivity impacts of the Early-Pay Treatment, we collected detailed measures of three markers of inattention, which may have influenced workers' efficiency of production ("attentional errors"), as described in detail in Section 2.2.1. Each of these measures indicates work patterns that occur easily as consequences of attention lapses. Such patterns are inefficient as they increase the time and effort per leaf plate (for the same piece rate). The measures are not incentivized. In fact, workers are not even aware that we collected these measures, such that we would not expect any changes in these measures as a consequence of alternative explanations for impacts (e.g. gift exchange).

The early treatment reduced workers' attentional errors by 0.07 to 0.1 standard deviations, as measured by a normalized index of attentional errors (columns 1 and 2 of Table 2.8). Mirroring the impacts of the Early-Pay Treatment on productivity, the impacts are almost exclusively concentrated among the poorer half of the sample (columns 3 to 5).

The observed impacts of the Early-Pay Treatment on attentional errors suggest that impacts on attention are one contributing mechanism of the observed treatment effects on worker productivity. Such attentional impacts could be explained by cash on hand reducing worries and thus distractions during work hours, as hypothesized by [86]. However, this evidence could also potentially be consistent with other psychological channels such as stress, mental health, sleep, happiness, or motivation, that operate in the same way, i.e that are concentrated among the poor and mediated through attentional errors.

2.4.2 Salience Treatment

In addition to considering the impact of the Early-Pay Treatment, we also investigate the role of attention by investigating potential impacts of directly focusing workers' attention on their financial situation. To do so, we implemented a salience intervention that was cross-randomized to the Early-Pay Treatment. Some workers received the salience treatment on day 6 of the study, others on day 10 of the study, and others not at all. In this intervention, surveyors told workers during the first hour of their work about another (fictional) worker's financial strain and then asked them about their own finances. Workers then returned back to their work. The exercise is similar in spirit to the mall study in [72].

Directing workers' attention can have two potentially opposing effects. First, since attention is limited, drawing workers' attention to their finances might divert valuable attentional resources from the work task and thus reduce worker productivity. Second, however, focusing workers' attention on their finances might raise workers' perceived marginal value of a dollar. Such impacts, resembling reminder effects in [80], might increase worker effort and thus increase worker productivity. Importantly, previous work on scarcity—such as [72]—had only limited scope for a positive channel.

To test for differential effect before vs. after payment, the salience treatment was randomized to be conducted before or after cash payment. Some workers received the salience intervention on day

6 (i.e. before any early payments occurred), others received it on day 10 (i.e. after some of those workers had received early payments), and others received no salience intervention. This design allowed us to test whether potentially stronger impacts of induced financial worries on cognition before the early payment caused differential effects of the intervention for workers who had been paid compared to workers who had not been paid.

In our setting, the overall impact of the salience intervention on worker productivity is positive (column 1 of Table 2.9). On days after receiving the salience intervention in the morning, workers were about 3.2 percent more productive compared to the remaining study sample. This result suggests that the motivational effect of focusing workers' attention on their finances is stronger than adverse effects of diverting attentional resources away from work.

The positive effects are entirely concentrated among workers who have already been paid. When workers were cash-poor (i.e. before their first major payment), we found no evidence of any (positive or negative) impact of the salience treatment on worker productivity, suggesting that the two opposing effects described above cancel each other out. In contrast, after workers were paid, the salience intervention increased worker productivity by up to 7 percent (columns 2 and 3 of Table 2.9).

The heterogeneous treatment effects (columns 4 and 5 of Table 2.9) highlight the difficulties in targeting salience interventions precisely. Our intention was to target poverty with the salience intervention, but it appears that we may have instead made loans more salient. This interpretation may explain the lack of heterogeneous treatment effects with respect to wealth and the clear evidence of heterogeneous impacts with respect to existing loans. More generally, it is difficult to raise the salience of only one particular issue of interest.[8]

Finally, while we find some evidence of increased attentional errors following the salience treatment, the evidence is only suggestive. Overall, we find a complicated set of results of the impact of the salience treatment on worker productivity, which highlights the caution warranted

[8]Moreover, salience interventions are likely to be non-monotonic (e.g. in baseline wealth or worries). For instance, on the one hand, one can only bring worries top of mind if they exist, e.g. impacts of making financial strain might be larger for people with more severe financial strain. On the other hand, the underlying concerns may already be top of mind for people who are very strained, leading to smaller treatment effects for people with more severe financial strain.

with salience or priming interventions, as also pointed out by Kahneman (2012).

2.5 Conclusion

This paper tests for a direct relationship between financial constraints and productivity. We provide evidence that even relatively minor improvements in workers' financial situations can have relatively large impacts on their productivity. When workers have less cash on hand, they produce fewer plates, make more errors per plate, and earn less in total. This evidence suggests that financial constraints by themselves may be detrimental for earnings, beyond potential impacts through investments in complementary inputs, human capital, or health. We also provide some evidence that attention is one mediating mechanism. We find that relaxing workers' financial constraints also reduced attentional errors.

Given the impacts, it seems worth revisiting other contexts in search of similar direct effects. For instance, [103] document large seasonal variation in earnings among farmers in Zambia. [82] and [83] find large and persistent impacts of bundled treatments to support the ultra-poor. Such impacts are often attributed to neoclassical explanations such as credit constraints. Our evidence suggests that direct effects of changes in financial strain may have contributed to the observed impacts in these settings.

Finally, our findings may have some implications for policy. The direct impact of financial strain on worker productivity is a parameter of interest for various policies, including unemployment insurance, basic income, or conditional and unconditional cash transfers. Importantly, the observed direct effects of reducing financial strain may occur *in addition* to any investment effects economists usually consider.

2.6 Figures and tables

Figure 2.1: Leaf plate

Notes: This figure shows a sal tree leaf plate akin to the ones produced as part of the experiment. In accordance with quality standards set by parterning contractors, leaf plates were required to (i) meet a minimum size requirement, (ii) have no gaping holes, (iii) have all leafstalks (petioles) covered by other leaves, and (iv) have the inner center parts placed underneath the outer rings of the plates.

Figure 2.2: Experimental Design

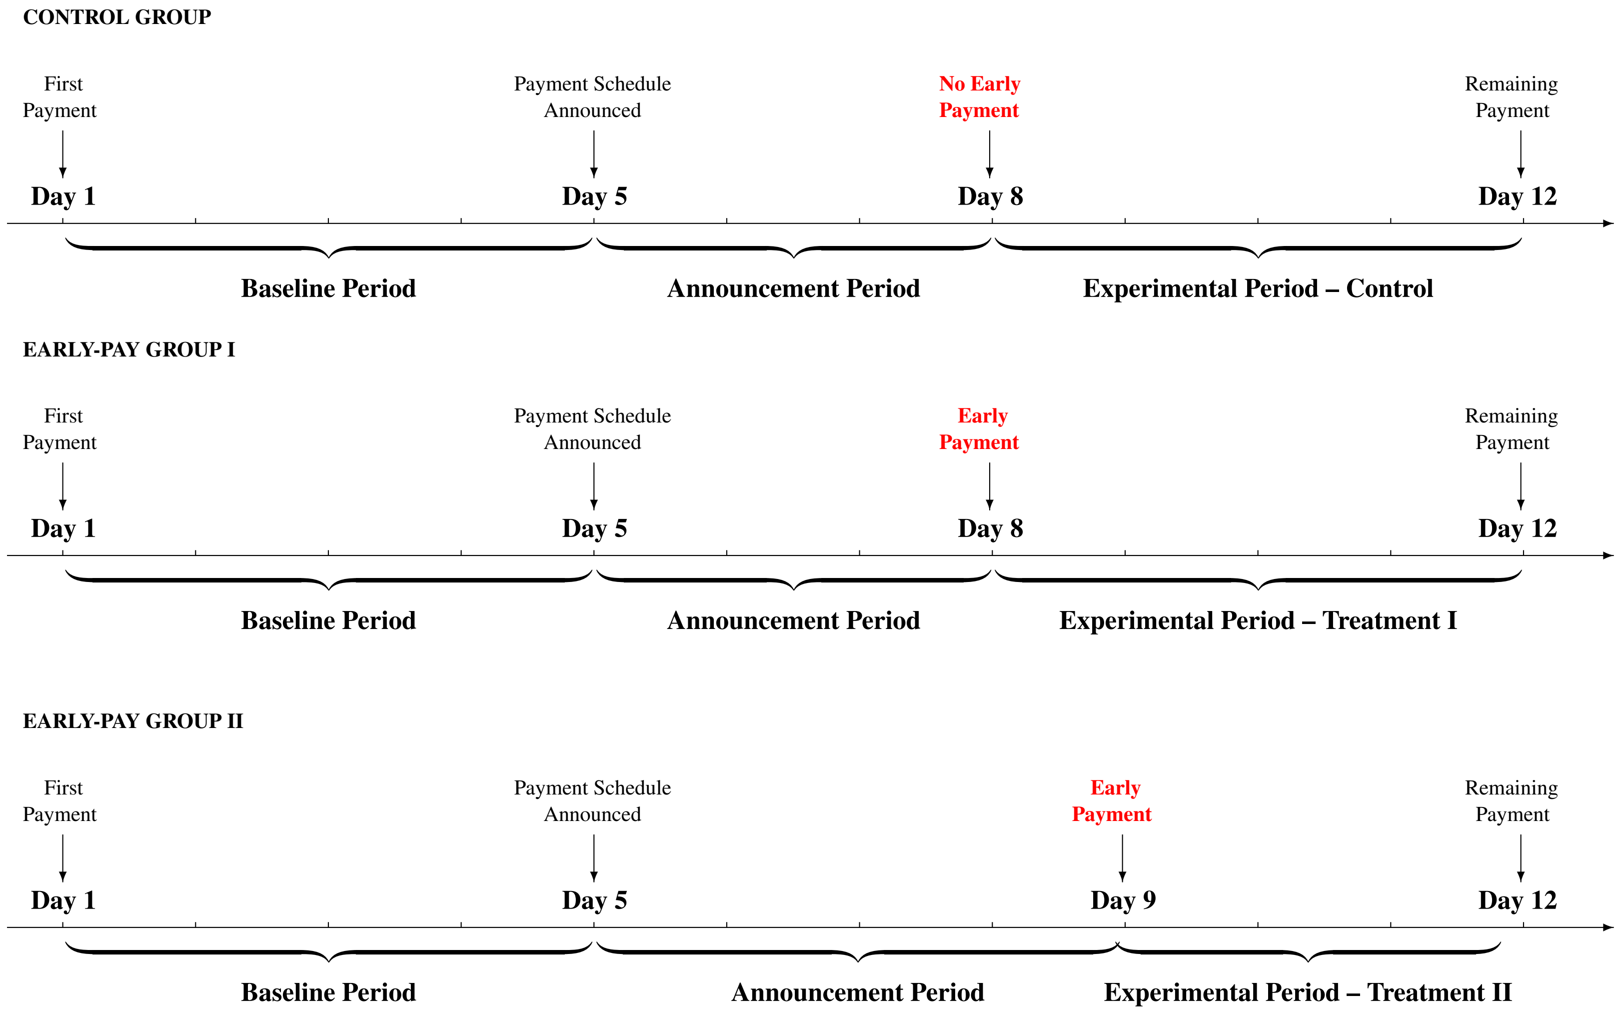

Notes: This figure shows the experimental design of the study. In the Control Group (upper part of the figure), workers were paid on days 1 and 12. In the Early-Pay Group I (center part of the figure), workers were paid on days 1, 12, and additionally on day 8. In the Early-Pay Group II (center part of the figure), workers were paid on days 1, 12, and additionally on day 9.

Table 2.1: Balance of Worker Characteristics

	Late payment group mean	Diff. for early payment group	N
Age	39.188	0.180	404
	(0.658)	(0.893)	
Years of education (top-coded at 13)	4.694	-0.322	406
	(0.256)	(0.343)	
Can read newspaper in Odiya	0.630	0.010	403
	(0.036)	(0.048)	
Primarily does daily wage labor	0.751	-0.049	403
	(0.032)	(0.045)	
Income quartile	2.393	0.027	404
	(0.080)	(0.109)	
Ate meat during previous week	0.579	-0.079	405
	(0.037)	(0.050)	
Very worried about future finances	0.691	0.030	352
	(0.036)	(0.049)	
Has loans	0.683	0.048	406
	(0.034)	(0.046)	
Has loans from moneylender	0.175	-0.005	407
	(0.028)	(0.038)	
Very worried about any loan	0.410	0.007	406
	(0.036)	(0.049)	
Total loan amount	11,223	1,371	378
	(1,157)	(1,660)	
Owns land	0.564	0.000	401
	(0.037)	(0.050)	
Non-mud house	0.238	-0.021	403
	(0.032)	(0.042)	
No food credits	0.459	0.005	407
	(0.037)	(0.050)	
Can get 1K in emergency	0.355	-0.066	404
	(0.035)	(0.047)	
Wealth index (avg of the 4 vars above)	0.405	-0.021	401
	(0.018)	(0.026)	
1st principal factor above median	0.530	-0.058	401
(using the same 4 vars)	(0.037)	(0.050)	

Notes: This table shows the tests of differences in baseline worker characteristics. Each characteristic is regressed on “Cash”, which refers to whether an individual is in one of the early-payment groups. Standard errors are clustered by worker.

Table 2.2: Impact of Early Payment on Expenditure Patterns

	Dependent Variable - Expenditure Category							
	Paid off any loan or credit (1)	Loans/ Credits (2)	Food (3)	Medical (4)	Agricultural (5)	Tobacco/ Alcohol (6)	Other (7)	Total (8)
Cash x Post	0.401***	277.9***	67.88***	16.47	-17.11	0.260	-125.3	220.1
	(0.0439)	(58.06)	(23.77)	(12.27)	(13.90)	(4.634)	(132.7)	(150.8)
Dependent var mean	0.18	121.98	269.94	31.55	28.33	34.01	285.63	771.43
N: workers	401	401	401	401	401	401	401	401

Notes: This table shows the impact of the early-payment treatment on expenditure patterns. “Cash” refers to whether an individual was part of one of the two early-payment groups. “Post” indicates whether a worker had received his early payment. Individuals were surveyed about their expenditure patterns during the preceding three days. All regressions control for individual fixed effects. Standard errors are clustered by worker.

Table 2.3: The Impact of Early Payment on Worker Productivity

Dependent variable: Log hourly output

			Wealth proxy					
			Owns land	Non-mud house	No food credits	Can get 1K in emergency	1st principal factor	Wealth index (avg)
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Cash x Post	0.0533**	0.0535**	0.0907**	0.0693**	0.0858**	0.0693**	0.129***	0.110***
	(0.020)	(0.020)	(0.029)	(0.022)	(0.027)	(0.024)	(0.034)	(0.027)
Cash x Post x Wealth			-0.0650	-0.0820*	-0.0719*	-0.0688	-0.125**	-0.200**
			(0.040)	(0.049)	(0.041)	(0.044)	(0.040)	(0.073)
Salience controls?	N	Y	N	N	N	N	N	N
N: worker-hours	22523	22523	22470	22470	22470	22470	22470	22470

Notes: This table shows the impact of the Early-Payment Treatment on worker productivity. “Cash” refers to whether an individual is in one of the Early-Payment Groups. “Post” indicates whether a worker had received his early payment. Regressions control for individual, day in study, hour of day, and round times work hour fixed effects, as well as for whether individuals had received the salience intervention previously. Standard errors are clustered by work.

Table 2.4: Persistence of Early Pay Impacts

Dep. variable: Log hourly production	**Wealth proxy**	
	Wealth Index	First PC
Cash X Post 1 Day	0.116***	0.137***
	(0.022)	(0.026)
Cash X Post 1 Day X High Wealth	-0.108***	-0.085***
	(0.031)	(0.029)
Cash X Post 2 Days	0.070***	0.085***
	(0.021)	(0.026)
Cash X Post 2 Days X High Wealth	-0.092***	-0.068**
	(0.031)	(0.027)
Cash X Post 3 Days	0.086***	0.120***
	(0.025)	(0.032)
Cash X Post 3 Days X High Wealth	-0.159***	-0.133***
	(0.054)	(0.041)
Observations	22013	22471
R-squared	0.346	0.344

Notes: This table shows the impact of the early-payment treatment on worker productivity. This table shows the same regressions as in columns (7) and (8) in the previous table, splitting up the "Post" coefficient into three days following the early payment. "Cash" refers to whether an individual is in one of the early-payment groups. "Post" indicates whether a worker had received his early payment. All regressions control for individual, day in study, hour of day, and round times work hour fixed effects, as well as for whether individuals have received the salience intervention previously. Standard errors are clustered by worker.

Table 2.5: The Impact of Announcing Early Payment

Dependent variable: Log hourly production

	(1)	(2)	(3)	(4)	(5)	(6)
Cash x Post-announcement	-0.000687	0.00176				
	(0.021)	(0.021)				
Cash x Post-announcement (Day 1)			-0.0302	-0.0339		
			(0.021)	(0.022)		
Cash x Post-announcement (Days 2+)			0.00892	0.0122		
			(0.023)	(0.023)		
Cash x Post-announcement (Days 1-2)					-0.0000374	0.00705
					(0.021)	(0.021)
Cash x Post-announcement (Days 3+)					-0.00106	-0.00144
					(0.024)	(0.024)
Cash x Post-treatment		0.0541		0.0558		0.0542
		(0.029)		(0.029)		(0.029)
Observations	Pre-treatment	All observations	Pre-treatment	All observations	Pre-treatment	All observations
R-squared	0.341	0.341	0.346	0.344	0.341	0.342
N	14983	22523	14983	22523	14983	22523

Notes: This table shows the impact of announcing the early payment on worker productivity. The post-announcement period is defined as "1" for days when the payment schedule had been made (i.e. starting on day 5) until the day of the early payment (i.e. day 8 or 9), and 0 otherwise. Standard errors are clustered by worker.

Table 2.6: The Impact of Early Payment on Workers Who Did Not Get Paid

Dependent variable: Log hourly production

	(1)	(2)	(3)
Later Cash x 1 day before payday	-0.00518	0.0145	
	(0.025)	(0.026)	
Cash x 1 day before payday			0.0158
			(0.024)
Cash x 2 days before payday			-0.0176
			(0.024)
Cash x 3 days before payday			0.0238
			(0.023)
Cash x Post		0.0530	0.0518
		(0.020)	(0.026)
Sample	Pre-treatment period	All observations	All observations
R-squared	0.349	0.342	0.344
N	14983	22523	22523

Notes: This table shows the Impact of the Early Payment on workers who were *not* paid. Columns 1 and 2 shows regressions that consider the difference in performance on day 9 between workers in the Early-Pay Group II (i.e. who were paid at the end of day 9) to the Control Group (who were paid on day 12). Column 3 estimates potential payday effects, i.e. it considers worker performance during the days before their payment. In all rounds (except for round 2), work on the payday itself did not count toward the payment, e.g. workers paid on day 8 were paid for their work until day 7, thus mitigating potential payday effects.

Table 2.7: The Impact of the Early Payment via Nutrition Channels

	Dependent variable: Output			Dependent variable: Breakfast measures				
	Log hourly production			Ate any breakfast	Ate rice	Amount of rice	Ate vegetables	Ate lentils
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Cash x Post	0.0513**	0.0481**	0.0548***	-0.00674	-0.00155	-3.823	-0.0207	0.0189
	(0.021)	(0.021)	(0.020)	(0.0130)	(0.0244)	(7.125)	(0.0414)	(0.0159)
Cash x Post x Last 2 hours of day	0.00341	-0.00237						
	(0.015)	(0.016)						
Cash x Post x Last 1 hour of day			-0.0125					
			(0.019)					
p-value: cash effect + interaction	0.0116	0.0396	0.0839					
Dep var mean - Control group	1.36	1.36	1.36	0.98	0.94	180.63	0.76	0.03
N	22523	22523	22523	320	320	320	320	320

Notes: This table shows the impact of the early-payment treatment on worker productivity during the last two hours of the day (columns 1 through 3) and on different measures of breakfast consumption (columns 4 through 8).

Table 2.8: The Impact of Early Payment on Attentional Errors

	Dependent variable				
	Attention (normalized index)	Attention index above median	Attention (normalized index)	Attention (normalized index)	Attention index above median
	(1)	(2)	(3)	(4)	(5)
Cash x Post	-0.110**	-0.0741**	-0.211**	-0.183**	-0.114**
	(0.040)	(0.025)	(0.073)	(0.058)	(0.035)
Cash x Post x Wealth			0.278*	0.163**	0.0879*
			(0.146)	(0.078)	(0.049)
Dependent var mean	-0.049	0.491	-0.049	-0.049	0.491
N: worker-hours	15265	15265	15227	15227	15227

Notes: This table shows the impact of the early-payment treatment on attentional errors. "Cash" refers to whether an individual is in one of the early-payment groups. "Post" indicates whether a worker had received his early payment. All regressions control for individual, day in study, hour of day, and round times work hour fixed effects, as well as for whether individuals have received the salience intervention previously. Standard errors are clustered by worker.

Table 2.9: The Impact of the Salience Intervention on Worker Productivity

Dependent variable: Log hourly production					
	(1)	(2)	(3)	(4)	(5)
Salience	0.0315**	0.0731***	0.0663***	0.0593**	0.0282
	(0.013)	(0.018)	(0.022)	(0.026)	(0.026)
Salience x Pre-cash		-0.0703***	-0.0517**	-0.0591*	-0.00808
		(0.023)	(0.026)	(0.034)	(0.032)
Salience x High wealth index				0.0181	
				(0.037)	
Salience x Pre-cash x High wealth index				0.00434	
				(0.047)	
Salience x High loan amount					0.0932**
					(0.036)
Salience x Pre-cash x High loan amount					-0.117**
					(0.047)
Cash treatment controls?	No	No	Yes	Yes	Yes
R-squared	0.340	0.341	0.343	0.343	0.344
N	22523	22523	22523	22523	21137

Notes: This table shows the impact of the salience intervention on worker productivity. “Cash" refers to whether an individual is in one of the early-payment groups. “Post" indicates whether a worker had received his early payment. All regressions control for individual, day in study, hour of day, and round times work hour fixed effects, as well as for whether individuals have received the salience intervention previously. Standard errors are clustered by worker.

Figure 2.3: Relationship between Productivity and Attentional Errors

	Dependent Variable						
	Number leaves	Number of undone mistakes	Number of stitches	Attention (normalized index)	Attention index above median	Attention (normalized index)	Attention index above median
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Baseline productivity	-0.179***	-0.239***	-0.259***	-0.176***	-0.0708***		
	(0.0245)	(0.0401)	(0.0458)	(0.0224)	(0.00902)		
Baseline prod. quartile = 2						-0.136*	-0.0155
						(0.0803)	(0.0422)
Baseline prod. quartile = 3						-0.376***	-0.103**
						(0.0823)	(0.0437)
Baseline prod. quartile = 4						-0.704***	-0.258***
						(0.0811)	(0.0417)
N: worker-hours	15265	11620	11620	15265	15265	15265	15265

Notes: This table shows the cross-sectional relationship between worker productivity and attentional errors. Standard errors are clustered by worker.

2.7 Appendix C. Supplementary notes

2.7.1 Deviations in Work and Payment Schedules

The standard schedule refers to the 12-day, 5-hour work schedules with a base rate of Rs. 200 and a piece rate of Rs. 3 per plate, implemented for rounds 4 to 11 of the study. While most rounds had consecutive work days, some rounds had one-day breaks in the first half of the rounds due to local events and religious festivals. Specifically, there were one-day breaks after day 5 in round 2, day 2 of round 3, and day 3 of round 12.

Rounds 1-3, which were conducted in March-June of 2017, had a number of deviations from the standard schedule and wage rates which were later finalized and then implemented during March-June of 2018.

First, in the earlier rounds, each workday consisted of 7 hours of work and a lunch break, rather than 5 continuous hours of work without lunch. Both types of workday schedules are common in the local region. Since some workers expressed their preferences for shorter work days due to hot weather during the lean season, the daily schedules were updated in 2018. Workers with the 5-hour schedules still received a snack at the end of each day.

Second, rounds 2 and 3 had deviations in weekly schedules. Round 2 was shorted by one day, effectively removing day 12 of the standard schedule and giving final payments on day 11. This was due to a local festival that coincided with day 12 of this round. Round 3 had early-payment days pushed back to days 9 and 10. This change was similar to inserting one additional regular work day after day 5, and removing day 12 from the standard schedule. Workers in this round initially predicted a large number of absences for day 6 due to a local event, so the payment days were pushed back, but the event did not take place so the worksite had its regular operation.

Third, the wage rates and payment lags had minor differences. In round 1, workers received a flat wage of Rs. 230 for Day 1, and a lower piece rate of Rs. 2 per accepted leaf plate. Rounds 2 and 3 had lower base wage rates of Rs. 180 and Rs. 175 respectively. In all three rounds, workers in the Early-Pay Group received wages earned two days prior to the early payment day (e.g. wages from days 2-6 for those with an early payment on day 8). Finally, all the workers received the bonus payment of 350 Rs. if they attended all five mandatory days in the later days of rounds.

The later rounds (rounds 12-14) were shortened in order to avoid running the experiment during the transplanting season. Round 12 was shorted by one day, effectively removing day 5 from the standard schedule and announcing the individual payment schedules in the morning on day 6.

Round 13-14 were shorted to 6 days. There was no flat wage payment on day 1 in these rounds. The Early-Pay Groups received payments on day 3 (group I) and 4 (group II). In order to make the size of the early payments comparable to the other rounds, the workers in the Early-Pay Groups were paid the flat wage for day 1, the wages they earned until that day (e.g. wages for day 2-3 for group I), in addition to a bonus of Rs. 200. The late-pay group received all wages, including the first day wage and the bonus, on the last day.

Chapter 3: Gender Norms in Marriage and Female Labor Productivity

3.1 Introduction

Do gender norms cause women to hold back their potential in the labor market? Using US data, Bertrand, Kamenica, and Pan (2015) show that the distribution of the share of income earned by the wife has a sharp drop-off at the point where the wife's income exceeds the husband's. They argue that this pattern is best explained by gender norms that cause an aversion to the situation in which the wife earns more than her husband. Such norms can govern the various ways in which concerns surrounding marriage and marital relationships affect labor-market outcomes. For instance, recent studies show that unmarried female MBA students reduce their career investments in front of other classmates (Bursztyn, Fujiwara, Pallais 2017) and being promoted to top jobs results in a large increase in the probability of divorce for women but not for men (Folke and Rickne 2020). A key related question is whether gender norms can also cause women to hold back their productive potential in the labor market.

This project uses a field experiment in rural India to test whether married women reduce own productivity on manufacturing tasks in order to avoid out-earning their husbands. The experiment engages married couples working as casual laborers in a two-day job of producing paper bags. In addition to base wage, workers are paid piece-rate wage on their paper bag production on the second day. On this day, men are required to make paper bags in the morning and engage in other activities in the afternoon, and women have a reversed schedule that includes one additional hour of paper bag production in the afternoon. Before starting on paper bag making, women are randomized into one of three conditions in which: 1) wife is not informed of her husband's production and told that the couple will learn only their joint production; 2) wife is informed of her husband's production and told that only she will learn the couple's final individual production; or 3) wife is informed of her husband's production and told that both spouses will learn the couple's final individual production.

The results indicate that women's production decreases only when women expect their husbands to find out about individual production. Women in the first two conditions achieve on average one hour's worth of production more than that of their husbands, suggesting that women do not have intrinsic concerns about out-producing their husbands. However, this productivity gap disappears in the last condition, making women's average productivity similar to that of men's. This could be explained by women slowing down on purpose because their husbands would suffer a loss of identity from violating gender norms. The relevant norms here may concern relative income between spouses, or rather, relative performance on gendered tasks. I discuss how to potentially disentangle such mechanisms by using similar sets of experiments with specific variations in design.

This paper relates to the literature that discuss the effect of gender norms on marriage and women's labor market outcomes. Some studies link women's relatively higher income compared to their husbands' income to a greater probability of marital disruption (Heckert, Nowak, and Snyder 1998; Jalovaara 2003; Bertrand, Kamenica, and Pan 2015. If these findings were driven by gender norms, the norms could also cause women to hold back their potential in the labor market. However, there has not been experimental evidence establishing this behavior.

Testing whether gender norms in marriage constrain women's labor market performance is important for expanding the discussion on gender inequality and poverty. Many low-income countries are characterized by gender inequality in the labor market as well as strong gender norms (Jayachandran 2015). Policies that are aimed at increasing women's income may be less effective if women are concerned about earning more than their husbands. Such norms can also provide a partial explanation for why some policies that increase women's income lead to increased reports of marital dissatisfaction or domestic violence. In addition, women married to low-earning husbands are potentially subject to more severely binding constraints on their earning potential, making it difficult for the households to increase their joint incomes. Finally, even in higher-income countries, the gender norms governing relative income and performance could exacerbate gender wage gap and segregation across occupations (Gottfredson 1981; West and Zimmerman 1987). Therefore, this study's finding contributes to the academic and policy discussions on gender norms and labor market inequality.

3.2 Experiment design

3.2.1 Sample construction and recruiting

The study takes place in rural Odisha, one of India's more under-developed states. Casual laborers in this setting tend to work in agriculture during peak planting and harvesting seasons, and get short-term contract jobs in unskilled manufacturing or construction in the remaining lean periods. Female casual laborers tend to work in the same jobs as male workers, but it is common to find gender wage gaps as well as gender-division of tasks within a job. For example, in agriculture, women are actively involved in transplanting and weeding, but rarely in plowing or well digging (Dewan 2015). Hence while some tasks are perceived as being gender-neutral, other tasks are predominantly performed by men or women.

The experimental sample is composed of 57 married couples, both aged between 18 and 55, who primarily derive income from casual daily-wage labor and have no prior experience in the experimental tasks. The sample is recruited using the same process that actual employers tend to use in this setting. The recruiters directly visit workers' villages, provide job descriptions, and offer transportation to the worksites. Eligible couples who agree to the offered terms participate in the experiment.

3.2.2 Procedures

The experiment tests whether gender norms surrounding marriage make some women reduce their productivity at work. Married couples are recruited for a two-day job of manufacturing a simple household good, paper bag. On the first day, workers receive training and are paid fixed daily wage. On the second day, in addition to base wage, workers are paid piece-rate wage on production. Based on market quality standards, supervisors determine the number of "accepted" paper bags produced by each household that and pay piece-rate wage on this amount.

Importantly, schedules for the second day differ across workers, with only part of the working time being dedicated to paper bag production. Men work on paper bags in the morning while women do so in the afternoon. When not engaged in paper bag production, workers may be trained on another task or take surveys. Because men finish paper bag production first, their productivity level can serve as a benchmark for their wives. Before women begin making paper bags in the afternoon, they are randomized into different treatment arms as follows:

- Control: women do not learn about their husbands' production and are reminded that they will only learn about household joint production at the end of the day.
- Single: women learn exactly how many accepted paper bags their husbands produced in the morning. They are told that only women will learn about individual productivity.
- Both: as above, women learn exactly how many accepted paper bags their husbands produced in the morning. However, they are told that both spouses will learn about individual productivity.

To ensure that workers cannot easily determine their own productivity and infer information about their spouses, a number of measures are undertaken. First, there is a chance that some products are not accepted due to the aforementioned quality control process. Second, supervisors frequently take away the finished products from workers' individual sitting spots to store them elsewhere. Third, workers' sitting spots are partitioned with dividers so that they cannot compare their production speed against others.

For the women not in the Control group, however, different measures are taken so that they can easily compare their production against their husbands'. Their husbands' accepted products are kept in a pile in their sitting area with a memo indicating the total amount. Women's finished products are also piled next to the husbands' piles. In addition, after each hour of production, the products are checked for quality control and women receive updated information about how many accepted paper bags they have produced so far. These procedures are designed to give them control over the final production outcome.

Notably, women are given one hour advantage in terms of the time used for producing paper bags. Therefore, on average, women should be able to produce more paper bags than their husbands and earn more money.[1]. However, this difference in allocated time are not made salient to the workers, i.e. the number of hours is not explicitly mentioned to the workers and there are no clocks on the walls at the worksites. Hence it should not be straightforward for the workers to determine that women have a clear advantage in production.

3.3 Results and discussion

Comparing productivity gaps between women and men across treatment groups reveals that women in the Control and Single groups out-produce their husbands by about one hour's worth, but this productivity gap is significantly reduced under the Both condition. Table 3.1 reports the OLS regression results from regressing the productivity gap (women's accepted product - men's accepted product) on indicators for treatment conditions. The preferred specification in Column 2 that controls for previous day's productivity gap shows that women produce 8-12 more paper bags under the Control and Single conditions. However, under Both, the productivity gap decreases by 9 units, and this change is statistically significant at 1% level.

Hence I find that women's productivity decreases compared to their husbands' only when women expect to their husbands to gain knowledge of relative productivity. This is consistent with the main hypothesis that gender norms in marriage cause some women to purposefully lower their productivity at work. Furthermore, this result suggests that women do not have intrinsic identity concerns about out-earning their husbands, but wish to avoid doing so in front of their husbands. This could be because they correctly believe that their husbands would suffer a loss of identity-related utility when women earn higher income than them.[2]

However, this result can also be explained by other related mechanisms, or by some other factors unrelated to gender norms. To disentangle the different mechanisms, running similar experiments

[1] A set of 'dry-run' sessions simply engaged men and women in paper bags making for two days and observed their production. Men and women were similarly productive during these sessions, as shown in Figure 3.1

[2] Oh 2020 provides a detailed discussion on the distinction between identity and social-image concerns.

with some specific alternations in design could be helpful. For example, I could reverse the role of men and women to observe how men's productivity changes across the same treatment conditions. Other changes could include, whether afternoon productions are one hour longer or shorter, whether the production task is perceived to be female-oriented, whether workers are compared to random partners instead of spouses, and what the wage schemes are.

I discuss below some potential explanations and how additional sets results could distinguish them from others.

Gender norms about performance. If the result was entirely driven by norms about relative income, giving men higher base wage would reduce the effect on women's productivity. However, if the norms broadly governed performance in different domains, women may still avoid out-performing their husbands even when husbands are expected to make higher income. Furthermore, women may act as if it is acceptable for them to out-earn their husbands when the task is perceived to be female-oriented.[3]

Identity vs. social-image. It would be important to distinguish whether the result is driven by men's identity concerns–intrinsic desire to act according to the norms–or by social-image concerns–following the rules in front their wives. Men facing identity concerns would try to work extra hard if they were given one hour less to work in the afternoon under all non-Control conditions. Men facing only social-image concerns would increase productivity only in front or their wives or neighbors.

Misperceptions. It may be important to also distinguish whether workers have correct beliefs about their spouses' preferences. It is possible that men do not have any concerns about breaking gender norms, but women believe they do. This can be tested by eliciting women's beliefs and comparing them to their husbands' production outcomes.

Matching effort. It could be that some women prefer to put in just as much effort as their husbands at work but believe that their husbands would discount their effort unless there is clear information about individual productivity. However, this would be less plausible if women do not increase productivity when they have one-hour disadvantage under the Both and Public conditions.

[3] I ran a separate survey to collect information on gender perceptions of tasks. Figure 3.2 shows that very few people associate women with making paper bags whereas most associate women with making leaf mats and wicks.

Also, it would have difficulty explaining how the results could be different if workers are engaged in a female-oriented task.

Differences in preference for competition. Some studies suggest men perform better under competitive environments. However, this environment does not resemble a competition in which winners are distinguished and get greater rewards. Furthermore, if the results were to be different on female-oriented tasks, this explanation would still mean that women would only prefer to not complete and "win" against their husbands in specific domains.

Control over resources. Couples may not make allocation decisions jointly and hence bargaining considerations could affect their production outcomes at work. The current results show that women choose to make less money if the individual earnings are explicitly attributed to them, which is slightly difficult to reconcile with a bargaining story. Collecting more survey information could be useful for learning about how household bargaining interacts with the specific gender norms at play.

3.4 Conclusion

This project tests whether some women engage in self-handicapping behavior at work due to concerns about violating gender norms. The initial findings suggest that women do not have intrinsic identity concerns about out-performing their husbands at work but decrease own productivity when they expect their husbands to find out about relative productivity. This finding relates to the important discussion on how gender norms surrounding marital relationships could exacerbate gender inequality in the labor market. In addition, it provides motivation for futures studies to disentangle the potential mechanisms behind this finding.

3.5 Figures and tables

Figure 3.1: Gender difference in productivity

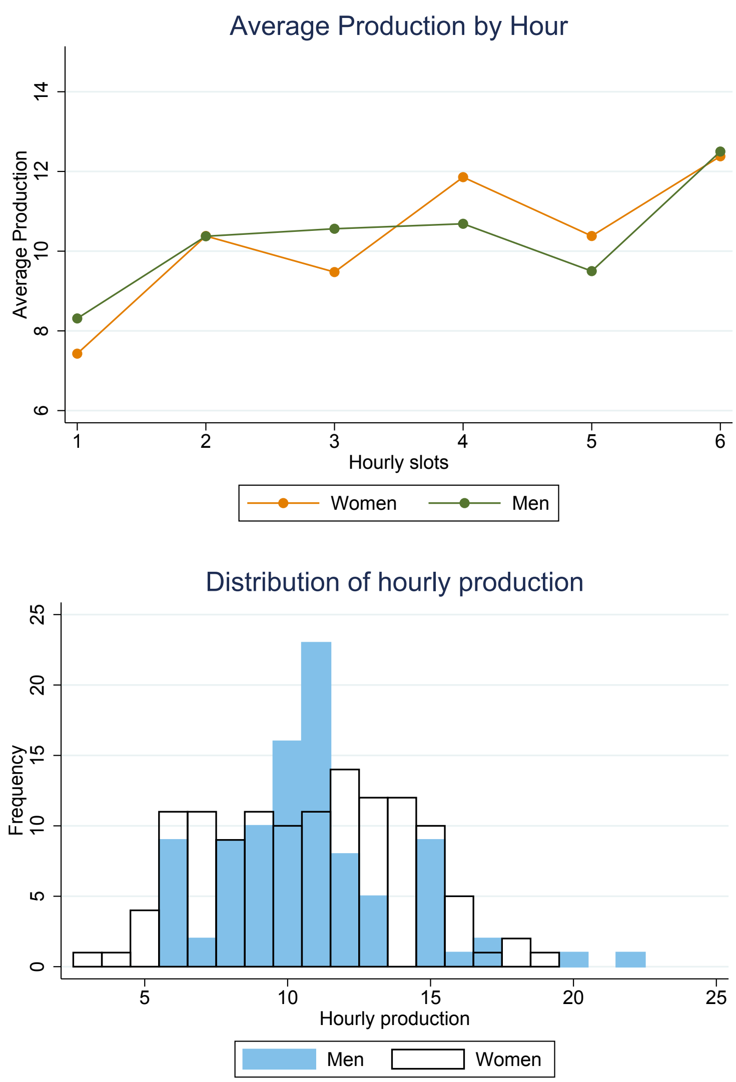

Notes. A set of 'dry-run' sessions engaged 37 men and women in paper bag making for two days. The figure plots productivity data from the second day. The top panel shows that average accepted production in each hourly plot similarly increases over time for men and women. The bottom panel shows that the distributions of hourly production are similar across gender.

Figure 3.2: Gendered perceptions of tasks

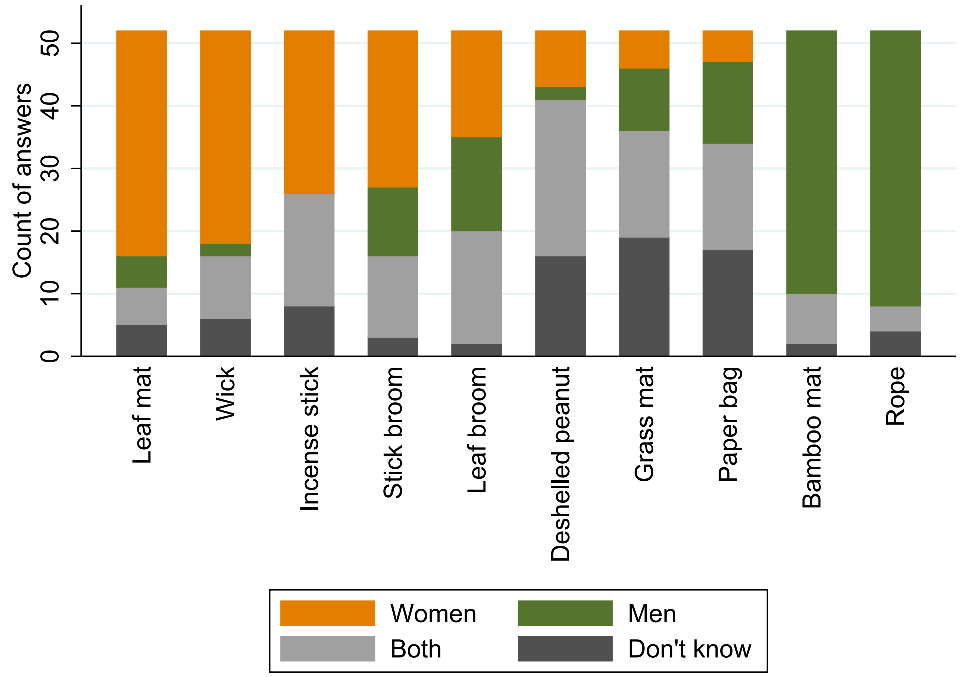

Notes. A pilot survey asked 54 subjects to look at a list of tasks and asked if each task is predominantly performed by men, women, or both. This figure plots the share of subjects reporting different answers by task.

Table 3.1: Productivity gap by treatment

	Productivity gap (1)	Productivity gap (2)
Control	-6.396	-4.247
	(7.681)	(3.969)
Both	-11.422**	-9.115***
	(4.927)	(3.208)
Productivity gap on Day 1		1.031***
		(0.125)
Constant	7.857**	12.765***
	(3.440)	(2.444)
R-squared	0.071	0.726
Observations	57	57

Notes. OLS estimates. Productivity gap is defined as a woman's accepted production net of her husband's accepted production. The omitted category is an indicator for being in the Single condition. Robust standard errors are reported.
* $p < .1$, ** $p < .05$, *** $p < .01$

References

[1] G. A. Akerlof and R. E. Kranton, "Economics and identity", Quarterly Journal of Economics, vol. 115, no. 3, pp. 715–753, 2000.

[2] P. J. Burke and J. E. Stets, Identity Theory. New York: Oxford Univ. Press, 2009.

[3] S. Stryker and P. J. Burke, "The past, present, and future of an identity theory", Social Psychology Quarterly, vol. 63, no. 4, pp. 284–297, 2000.

[4] M. A. Hogg, D. J. Terry, and K. M. White, "A tale of two theories: A critical comparison of identity theory with social identity theory", Social Psychology Quarterly, vol. 58, pp. 255–69, 1995.

[5] T. J. Owens, D. T. Robinson, and L. Smith-Lovin, "Three faces of identity", Annual Review of Sociology, vol. 36, pp. 477–499, 2010.

[6] K. Hoff and J. E. Stiglitz, "Striving for balance in economics: Towards a theory of the social determination of behavior", Journal of Economic Behavior & Organization, vol. 126, pp. 25–57, 2016.

[7] L. Bursztyn and R. Jensen, "Social image and economic behavior in the field: Identifying, understanding, and shaping social pressure", American Economic Review, vol. 9, pp. 131–153, 2017.

[8] L. S. Gottfredson, "Circumscription and compromise: A developmental theory of occupational aspirations", Journal of Counseling Psychology, vol. 28, no. 6, pp. 545–579, 1981.

[9] C. West and D. H. Zimmerman, "Doing gender", Gender & Society, vol. 1, no. 2, pp. 125–151, 1987.

[10] M. A. Cejka and A. H. Eagly, "Gender-stereotypic images of occupations correspond to the sex segregation of employment", Personality and Social Psychology Bulletin, vol. 25, no. 4, pp. 413–423, 1999.

[11] H. Tajfel and J. C. Turner, "An integrative theory of intergroup conflict", in The Social Psychology of Intergroup Relations, W. G. Austin and S. Worchel, Eds., Brooks-Cole, 1979, ch. 3, pp. 33–47.

[12] R. H. Topel and M. P. Ward, "Job mobility and the careers of young men", Quarterly Journal of Economics, vol. 107, no. 2, pp. 439–479, 1992.

[13] D. Acemoglu and D. Autor, “Skills, tasks and technologies: Implications for employment and earnings”, in Handbook of labor economics, vol. 4, Elsevier, 2011, pp. 1043–1171.

[14] C. Goldin, “A grand gender convergence: Its last chapter”, American Economic Review, vol. 104, no. 4, pp. 1091–1119, 2014.

[15] J. Adda, C. Dustmann, and K. Stevens, “The career costs of children”, Journal of Political Economy, vol. 125, no. 2, pp. 293–337, 2017.

[16] R. Bénabou and J. Tirole, “Incentives and prosocial behavior”, American Economic Review, vol. 96, no. 5, pp. 1652–1678, 2006.

[17] ——, “Identity, morals, and taboos: Beliefs as assets”, Quarterly Journal of Economics, vol. 126, no. 2, pp. 805–855, 2011.

[18] V. Schultz, “Reconceptualizing sexual harassment”, Yale Law Journal, vol. 107, pp. 1683–1805, 1998.

[19] I. Padavic, “The re-creation of gender in a male workplace”, Symbolic Interaction, vol. 14, no. 3, pp. 279–294, 1991.

[20] C. Goldin, Understanding the Gender Gap: An Economic History of American Women. New York: Oxford University Press, 1990.

[21] C.-T. Hsieh, E. Hurst, C. I. Jones, and P. J. Klenow, “The allocation of talent and us economic growth”, Econometrica, vol. 87, pp. 1439–1474, 5 2019.

[22] A. Erosa, L. Fuster, G. Kambourov, and R. Rogerson, “Hours, occupations, and gender differences in labor market outcomes”, 2017.

[23] A. Bell, R. Chetty, X. Jaravel, N. Petkova, and J. V. Reenen, “Who becomes an inventor in america? the importance of exposure to innovation”, Quarterly Journal of Economics, vol. 134, no. 2, 647–713, 2019.

[24] S. S. Goraya, “How does caste affect entrepreneurship? birth vs worth”, 2019.

[25] G. Cassan, D. Keniston, and T. Kleineberg, “A division of laborers: Identity and efficiency in india”, 2019.

[26] K. Hoff and P. Pandey, “Discrimination, social identity, and durable inequalities”, American Economic Review Papers & Proceedings, vol. 96, no. 2, pp. 206–211, 2006.

[27] ——, “Making up people–the effect of identity on performance in a modernizing society”, Journal of Development Economics, vol. 106, pp. 118–131, 2014.

[28] A. Cohn, E. Fehr, and M. A. Maréchal, "Business culture and dishonesty in the banking industry", Nature, vol. 516, pp. 86–89, 2014.

[29] L. Bursztyn, M. Callen, B. Ferman, S. Gulzar, A. Hasanain, and N. Yuchtman, "Political identity: Experimental evidence on anti-americanism in pakistan", Mimeo, 2017.

[30] D. J. Benjamin, J. J. Choi, and G. Fisher, "Religious identity and economic behavior", American Economic Review, vol. 98, no. 4, pp. 617–637, 2016.

[31] R. G. F. Jr. and P. Torelli, "An empirical analysis of 'acting white'", Journal of Public Economics, vol. 94, pp. 380–396, 2010.

[32] D. Austen-Smith and R. G. F. Jr., "An economic analysis of 'acting white'", Quarterly Journal of Economics, vol. 120, no. 2, pp. 551–583, 2005.

[33] A. Alesina, P. Giuliano, and N. Nunn, "On the origins of gender roles: Women and the plough", Quarterly Journal of Economics, vol. 128, no. 2, pp. 469–530, 2013.

[34] M. Bertrand, E. Kamenica, and J. Pan, "Gender identity and relative income within households", Quarterly Journal of Economics, vol. 130, no. 2, pp. 571–614, 2015.

[35] L. Bursztyn, T. Fujiwara, and A. Pallais, "'acting wife': Marriage market incentives and labor market investments", American Economic Review, vol. 107, no. 11, pp. 3288–3319, 2017.

[36] D. J. Benjamin, J. J. Choi, and A. J. Strickland, "Social identity and preferences", American Economic Review, vol. 100, no. 4, pp. 1913–1928, 2010.

[37] A. Falk, "Facing yourself: A note of self-image", IZA DP No. 10606, 2017.

[38] E. Breza, S. Kaur, and N. Krishnaswamy, "Scabs: The social suppression of labor supply", NBER Working Paper No. 25880, 2019.

[39] L. Bursztyn, A. González, and D. Yanagizawa-Drott, "Misperceived social norms: Female labor force participation in saudi arabia", NBER Working Paper No. 24736, 2018.

[40] A. Karing, "Social signaling and childhood immunization: A field experiment in sierra leone", Innovations for Poverty Action (IPA) Working Paper, 2018.

[41] P. Jakiela and O. Ozier, "Does africa need a rotten kin theorem? experimental evidence from village economies", Review of Economic Studies, vol. 83, no. 1, pp. 231–268, 2016.

[42] D. Atkin, E. Colson-Sihra, and M. Shayo, "How do we choose our identity? a revealed preference approach using food consumption", Mimeo, 2019.

[43] B. A. Bettencourt, K. Charlton, N. Dorr, and D. 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, vol. 127, no. 4, pp. 520–542, 2001.

[44] B. D. Bernheim, "A theory of conformity", Journal of Political Economy, vol. 102, no. 5, pp. 841–877, 1994.

[45] G. A. Akerlof, "A theory of social custom, of which unemployment may be one consequence", Quarterly Journal of Economics, vol. 94, no. 4, pp. 749–775, 1980.

[46] S. R. G. Jones, The Economics of Conformism. Oxford: Blackwell Pub, 1984.

[47] M. Marriott, "Caste ranking and community structure in five regions of india and pakistan", Bulletin of the Deccan College Research Institute, vol. 19, no. 1/2, pp. 31–105, 1958.

[48] P. M. Mahar, "A ritual pollution scale for ranking hindu castes", Sociometry, vol. 23, no. 3, pp. 292–306, 1960.

[49] S. Desai and A. Dubey, "Caste in 21st century india: Competing narratives", Economic and Political Weekly, vol. 46, no. 11, pp. 40–49, 2012.

[50] K. Munshi, "Caste and the indian economy", Journal of Economic Literature, vol. 57, no. 4, pp. 781–834, 2019.

[51] H. H. Risley, The People of India. Calcutta: Thacker, Spink and Co., 1908.

[52] ——, The Tribes and Castes of Bengal. Calcutta: Bengal Secretariat Press, 1892.

[53] E. Breza, S. Kaur, and Y. Shamdasani, "The morale effects of pay inequality", Quarterly Journal of Economics, vol. 133, no. 2, pp. 611–663, 2018.

[54] G. M. Becker, M. H. DeGroot, and J. Marschak, "Measuring utility by a single-response sequential method", Behavioral Science, vol. 9, no. 3, pp. 226–232, 1964.

[55] D. Fudenberg, D. K. Levine, and Z. Maniadis, "On the robustness of anchoring effects in wtp and wta experiments", American Economic Journal: Microeconomics, vol. 4, no. 2, pp. 131–145, 2012.

[56] N. Krishnaswamy, "Missing and fired: Worker absence, labor regulation, and firm outcomes", Working paper, 2019.

[57] D. Mosse, "Caste and development: Contemporary perspectives on a structure of discrimination and advantage", World Development, vol. 110, pp. 422–436, 2018.

[58] D. Gupta, Interrogating Caste: Understanding Hierarchy and Difference in Indian Society. New Delhi: Penguin Books, 2000.

[59] N. B. Dirks, Castes of Mind: Colonialism and the Making of Modern India. Princeton: Princeton University Press, 2001.

[60] S. Bayly, Caste, Society and Politics in India from the Eighteenth Century to the Modern Age. Cambridge: Cambridge University Press, 2001.

[61] K. Munshi and M. Rosenzweig, "Traditional institutions meet the modern world: Caste, gender and schooling choice in a globalizing economy", American Economic Review, vol. 96, no. 4, pp. 1225–1252, 2006.

[62] ——, "Networks and misallocation: Insurance, migration, and the rural-urban wage gap", American Economic Review, vol. 106, no. 1, pp. 46–98, 2016.

[63] S. Madheswaran and P. Attewell, "Caste discrimination in the indian urban labour market: Evidence from the national sample survey", Economic and Political Weekly, vol. 42, no. 41, pp. 4146–4153, 2007.

[64] S. Thorat and P. Attewell, "The legacy of social exclusion: A correspondence study of job discrimination in india", Economic and Political Weekly, vol. 42, no. 41, pp. 4141–4145, 2007.

[65] K. Hoff, M. Kshetramade, and E. Fehr, "Caste and punishment: The legacy of caste culture in norm enforcement", The Economic Journal, vol. 121, no. 556, F449–F475, 2011.

[66] G. Rao, "Familiarity does not breed contempt: Diversity, discrimination and generosity in delhi schools", American Economic Review, vol. 109, no. 3, pp. 774–809, 2019.

[67] M. Lowe, "Types of contact: A field experiment on collaborative and adversarial caste integration", Working paper, 2019.

[68] C. Gowda, "Barbers and hairstylists", The Hindu, 2011-08-20.

[69] R. Mohanty and A. Dwivedi, "What would urban sanitation look like without caste?", The Wire, 2018-07-10.

[70] J. Haushofer and E. Fehr, "On the Psychology of Poverty", Science, vol. 344, no. 6186, pp. 862–867, 2014.

[71] J. Haushofer and J. Shapiro, "The Short-Term Impact of Unconditional Cash Transfers to the Poor: Experimental Evidence from Kenya", Quarterly Journal of Economics, vol. 131, no. 4, pp. 1973–2042, 2016.

[72] A. Mani, S. Mullainathan, E. Shafir, and J. Zhao, “Poverty Impedes Cognitive Function”, Science, vol. 341, no. 6149, pp. 976–980, 2013.

[73] L. S. Carvalho, S. Meier, and S. W. Wang, “Poverty and Economic Decision-Making: Evidence from Changes in Financial Resources at Payday”, American Economic Review, vol. 106, no. 2, pp. 260–284, 2016.

[74] A. Oswald, E. Proto, and D. Sgroi, “Happiness and productivity”, Journal of Labor Economics, vol. 33, no. 4, pp. 789–822, 2015.

[75] D. Fehr, G. Fink, and K. Jack, “Poverty, Seasonal Scarcity and Exchange Asymmetries: Evidence from Small-Scale Farmers in Rural Zambia”, mimeo, 2019.

[76] S. DellaVigna, J. List, U. Malmendier, and G. Rao, “Estimating Social Preferences and Gift Exchange with a Piece-Rate Design”, mimeo, 2019.

[77] S. Kaur, M. Kremer, and S. Mullainathan, “Self-Control at Work”, Journal of Political Economy, vol. 123, no. 6, pp. 1227–1277, 2015.

[78] J. T. Dean, “Noise, Cognitive Function, and Worker Productivity”, mimeo, 2018.

[79] P. Bessone, G. Rao, F. Schilbach, H. Schofield, and M. Toma, “Sleepless in Chennai: The Consequences of Increasing Sleep among the Urban Poor”, mimeo, 2019.

[80] D. Karlan, M. McConnell, S. Mullainathan, and J. Zinman, “Getting to the Top of Mind: How Reminders Increase Saving”, Management Science, vol. 62, no. 12, pp. 3393–3411, 2016.

[81] P. Dasgupta and D. Ray, “Inequality as a Determinant of Malnutrition and Unemployment”, Economic Journal, vol. 96, pp. 1011–1034, 1986.

[82] A. Banerjee, E. Duflo, N. Goldberg, D. Karlan, R. Osei, W. Pariente, J. Shapiro, B. Thuysbaert, and C. Udry, “A Multifaceted Program Causes Lasting Progress for the Very Poor: Evidence from Six Countries”, Science, vol. 348, no. 6236, 2015.

[83] O. Bandiera, R. Burgess, N. Das, S. Gulesci, I. Rasul, and M. Sulaiman, “Labor Markets and Poverty in Village Economies”, Quarterly Journal of Economics, vol. 132, no. 2, pp. 811–870, 2017.

[84] A. Finkelstein, S. Taubman, B. Wright, M. Bernstein, J. Gruber, J. P. Newhouse, H. Allen, K. Baicker, and O. H. S. Group, “The Oregon Health Insurance Experiment: Evidence from the First Year”, Quarterly Journal of Economics, vol. 127, no. 3, pp. 1057–1106, 2012.

[85] M. Chemin, J. de Laat, and J. Haushofer, “Poverty and Stress: Rainfall Shocks Increase Levels of the Stress Hormone Cortisol”, mimeo, 2013.

[86] S. Mullainathan and E. Shafir, Scarcity: Why Having Too Little Means So Much. New York: Macmillan, 2013.

[87] A. K. Shah, E. Shafir, and S. Mullainathan, “Scarcity Frames Value”, Psychological Science, vol. 26, no. 4, pp. 402–412, 2015.

[88] Q. Ong, W. Theseira, and I. Y.H. Ng, “Reducing Debt Improves Psychological Functioning and Changes Decision-making in the Poor”, Proceedings of the National Academy of Sciences, vol. 116, no. 15, pp. 7244–7249, 2019.

[89] E. Breza, S. Kaur, and Y. Shamdasani, “The Morale Effects of Pay Inequality”, Quarterly Journal of Economics, vol. 133, no. 2, pp. 611–663, 2018.

[90] J. M. Shapiro, “Is There a Daily Discount Rate? Evidence From the Food Stamp Nutrition Cycle”, Journal of Public Economics, vol. 89, no. 2-3, pp. 303–325, 2005.

[91] J. Kaminski, L. Christiaensen, and C. L. Gilbert, “The End of Seasonality? New Insights from Sub-Saharan Africa”, World Bank Policy Research Working Paper No. 6907, 2014.

[92] A. V. Banerjee, R. Hanna, G. E. Kreindler, and B. A. Olken, “Debunking the Stereotype of the Lazy Welfare Recipient: Evidence from Cash Transfer Programs”, World Bank Research Observer, vol. 32, pp. 155–184, 2017.

[93] S. Baird, D. McKenzie, and B. Özler, “The Effects of Cash Transfers on Adult Labor Market Outcomes”, IZA Journal of Development and Migration, vol. 8, no. 1, pp. 1–20, 2018.

[94] E. B. Dean, F. Schilbach, and H. Schofield, “Poverty and Cognitive Function”, in The Economics of Poverty Traps, C. B. Barrett, M. R. Carter, and J.-P. Chavas, Eds., Chicago: University of Chicago Press, 2018.

[95] A. Adhvaryu, N. Kala, and A. Nyshadham, “The Light and the Heat: Productivity Co-Benefits of Energy-Saving Technology”, NBER Working Paper No. 24314, 2018.

[96] A. Lusardi, D. J. Schneider, and P. Tufano, “Financially Fragile Household: Evidence and Implications”, Brookings Papers on Economic Activity, vol. Spring, pp. 83–134, 2011.

[97] U. Gneezy and J. A. List, “Putting Behavioral Economics to Work: Testing for Gift Exchange in Labor Markets using Field Experiments”, Econometrica, vol. 74, no. 5, pp. 1365–1384, 2006.

[98] C. Esteves-Sorenson, “Gift Exchange in the Workplace: Addressing the Conflicting Evidence with a Careful Test”, Management Science, vol. 64, no. 9, pp. 4365–4388, 2017.

[99] J. de Ree, K. Muralidharan, M. Pradhan, and H. Rogers, “Double for Nothing? The Effects of Unconditional Teacher Salary Increases in Indonesia”, Quarterly Journal of Economics, vol. 133, no. 2, pp. 923–1039, 2018.

[100] D. Gilricht, M. Luca, and D. Malhotra, “When 3 + 1 > 4: Gift Structure and Reciprocity in the Field”, Management Science, vol. 62, no. 9, pp. 2639–2650, 2016.

[101] L. Casaburi and J. Willis, “Time vs. State in Insurance: Experimental Evidence From Contract Farming in Kenya”, mimeo, 2018.

[102] H. Schofield, “The Economic Costs of Low Caloric Intake: Evidence From India”, mimeo, 2014.

[103] G. Fink, K. Jack, and F. Maxiye, “Seasonal Liquidity, Rural Labor Markets and Agricultural Production”, NBER Working Paper No. 24564, 2018.

[104] O. Folke and J. Rickne, “All the single ladies: Job promotions and the durability of marriage”, American Economic Journal: Applied Economics, vol. 12, no. 1, pp. 260–287, 2020.

[105] D. A. Heckert, T. C. Nowak, and K. A. Snyder, “The impact of husbands’ and wives’ relative earnings on marital disruption”, Journal of Marriage and the Family, vol. 60, no. 3, pp. 690–703, 1998.

[106] M. Jalovaara, “The joint effects of marriage partners’ socioeconomic positions on the risk of divorce”, Demography, vol. 40, no. 1, pp. 67–81, 2003.

[107] S. Jayachandran, “The roots of gender inequality in developing countries”, Annual Review of Economics, vol. 7, no. 1, pp. 63–88, 2015.

[108] S. Dewan, “Closing the gender wage gap in indian agriculture”, in Global Wage Debates: Politics or Economics?, G. Randolph and K. Panknin, Eds., JustJobs Network, 2015, ch. 8, pp. 153–165.

[109] S. Oh, “Does identity affect labor supply?”, Working paper, 2020.