

SCABS: THE SOCIAL SUPPRESSION OF LABOR SUPPLY

EMILY BREZA[§], SUPREET KAUR^{*}, AND NANDITA KRISHNASWAMY[‡]

Preliminary and Incomplete

ABSTRACT. Social norms can serve as a powerful force for conformity, producing collective behaviors among decentralized individuals. We test for this force in the labor market: whether norms prevent workers from supplying labor at wage cuts, generating cartel behavior without explicit collusion. We partner with 183 existing employers, who offer jobs to 502 workers in informal spot labor markets in India. Unemployed workers are privately willing to accept jobs below the prevailing wage, but rarely do so when this choice is observable to other workers. In contrast, social observability does not affect labor supply at the prevailing wage. Workers give up 49% of average weekly earnings to avoid being seen as breaking the social norm. In addition, workers pay to punish anonymous laborers who have accepted wage cuts—indicating that cartel behavior is reinforced through the threat of social sanctions. Finally, consistent with the idea that social conformity could have aggregate implications, measures of social cohesion correlate with downward wage rigidity and business cycle volatility across India.

Date: This Version: November 8, 2018.

We thank Siwan Anderson, Doug Bernheim, Lucas Coffman, Stefano DellaVigna, Henry Farber, Andrew Foster, Patrick Francois, Pat Kline, and Michael Kremer for helpful comments and conversations. Shreoshee Mukherjee, Arnesh Chowdhury, Piyush Tank, Medha Aurora, and Sayan Kundu provided exceptional research assistance. We gratefully acknowledge financial support from the National Science Foundation, the Agricultural Technology Adoption Initiative, and the Institute for Research on Labor and Employment (UC Berkeley).

[§]Department of Economics, Harvard University; NBER; JPAL.

^{*}Department of Economics, University of California, Berkeley; NBER; JPAL.

[‡]USC Dornsife INET, Department of Economics, University of Southern California.

1. INTRODUCTION

Social psychology argues for the role of social influence in affecting individual behavior (e.g. [Asch, 1956](#); [Freedman and Fraser, 1966](#); [Milgram, 1974](#)). Under this view, group expectations or norms can be a powerful force for conformity among individuals. Consequently, social norms have the potential to enable collective behaviors among a large, decentralized group of individuals—without any formal coordination or institutions.

We study the potential for this force to shape aggregate behavior in economic markets. Specifically, we document the presence of implicit collusion among workers to maintain wage floors in the labor market—generating collective labor supply behavior without any formal labor organization or coordination. We provide positive evidence for the role of social sanctions, which are used to deter self-interested deviations by workers, in sustaining this cartel-like behavior. This accords with a nearly universal observation in social psychology that when an individual is seen as violating a norm, others are prone to engage in punishment or social disapprobation, serving to curtail norm violations (e.g. [Henrich et al., 2006](#); [Cialdini and Goldstein, 2004](#)).¹

This mechanism has potential applicability in any setting with meaningful social interaction. This makes the labor market a particularly relevant economic domain: the workplace inherently entails close and repeated interaction among workers, making it not only a productive environment, but also inherently a social one. This feature provides the ingredients for norms to develop and be sustained among workers. This is consistent with the fact that collective behaviors—such as the coordinated restriction of output, suppression of rate busters, mass walkouts with foregone wages—have been observed across history and contexts, even in the absence of formal unions or labor organization. More broadly, the presence of cartel-like behavior in the absence of explicit collusion has been observed in a range of markets with decentralized agents, from NASDAQ traders ([Christie and Schultz, 1994](#)) and real estate agents in the US ([Hsieh and Moretti, 2003](#)), to taxi drivers and vegetable vendors in poor countries.

In this paper, we seek to provide documentation of this phenomenon by testing for the role of social influence on labor supply. Specifically, we test whether norms

¹This view of norms is consistent with a theoretical literature in economics that models norms as an equilibrium coordination device. In such models, social punishment is a way to sustain collective behavior in equilibrium—either through a decision rule in a repeated game ([Kandori, 1992](#); [Ellison, 1994](#); [Jackson et al., 2012](#)) or through the internalization of norms into utility (e.g. [MacLeod, 2007](#)). These two modes of enforcement are not mutually exclusive, and suggest that collective behaviors could be locally maintained through repeated interaction, or broadly maintained even in one-shot encounters.

against accepting wage cuts constrain laborers' willingness to accept work below the prevailing wage during times of unemployment. The setting for our study is informal markets for casual daily labor in India. Such markets are ubiquitous in poor countries, serving as the primary channel for hired employment for hundreds of million workers in India alone (National Sample Survey 2010).

Figure 1 provides initial suggestive support for the presence of norms against accepting wage cuts in this setting, using the approach developed in [Kahneman et al. \(1986\)](#). In a survey conducted with Indian agricultural laborers, over 80% of respondents said it was "unacceptable" or "very unacceptable" for an unemployed worker to offer to work at a rate below the prevailing wage (Panel A). In addition, about 80% of respondents stated that other workers would become angry with an individual who accepted work below the prevailing wage (Panel B). While only speculative, these responses suggest that laborers view working for a lower wage, even when unemployed, as a violation which could result in social disapprobation or sanctions.

We hypothesize that during times of unemployment, (at least some) workers find it privately optimal to take up jobs at wages lower than the prevailing wage, but are less likely to do so because this would be perceived as a norm violation, resulting in sanctions from co-villagers. To test this hypothesis, we proceed in two steps. First, we test for social influence on labor supply. The main part of our study is a field experiment in which existing employers make job offers to unemployed workers under varying levels of social observability. Second, we document evidence of social punishment for norm violations. Using a supplementary exercise, we test for workers' willingness to destroy surplus in order to sanction those who have accepted wage cuts. We describe each of these components in turn below.

In this setting, agricultural employers typically hire laborers in their village in one-day spot contracts to work on their land for a given cultivation activity (e.g. weeding). To implement the field experiment, we partner with such employers. We induce two types of variation during their hiring process. (i) First, the job is offered at a random wage level: at the prevailing wage, or 10% below the prevailing wage.² (ii) Second, we vary the extent to which the wage level is publicly observable: whether the job offer is made inside the worker's home or outside on the street where neighbors (who are typically other workers) can overhear the offer.

²In this setting, there is a prevailing daily wage for each type of agricultural task. We provide direct evidence for this below.

All offered jobs correspond to actual employment opportunities on the employer’s land—so that our data reflects real employment decisions by workers. Treatment randomization is at the village level, so that all workers within a given village receive the same wageXobservability condition. In addition, the workers in our experiment (i.e. those who are offered jobs) are sampled randomly from the village population of laborers.³ Aside from hiring, we are not involved in any other aspect of the employment relationship: employers supervise workers as usual, provide them food, etc. The experiment is conducted across 183 villages (i.e. labor markets) with 183 distinct partnering employers (one in each village), with jobs extended to 502 workers.

We predict that, when assured privacy, at least some unemployed workers will choose to accept jobs below the prevailing wage. However, if other villagers can observe their decision, this will dampen their willingness to accept a job at this lower wage. In contrast, we predict that observability will not decrease take-up of jobs at the prevailing wage—since taking up these jobs does not constitute a norm violation.

Our results are consistent with these predictions. At the prevailing wage, the average take-up rate of jobs is 26%, and we cannot reject that this take-up rate is the same regardless of whether the job offer is publicly observable by others.⁴ In contrast, when a worker is offered a job below the prevailing wage, take-up depends crucially on whether his decision is publicly observable. When the lower wage is offered in private, take-up remains a robust 18%. However, this falls by 13.6 percentage points when low-wage offers are observable (significant at 1% level). When restricting the sample to workers who are in the agricultural labor market—defined as those who consider agricultural labor as their primary or secondary occupation—these results become even starker. Only 1.8% of agricultural workers accept wage cuts in public, an estimated 24.6 percentage point decline relative to take-up of wage cuts in private.

This distortion on individual labor supply is economically meaningful. The experiment was conducted during the lean season, when workers typically only find a few days of employment in a week. Consequently, passing up one day of work at a 10% wage cut is equivalent to foregoing 49% of average weekly earnings in our sample. This is a large magnitude, especially given that workers report skipping meals and

³As is typical, hiring is done by employers who approach laborers at their homes to make job offers. We randomly select among workers who are home at the time of hiring. We provide several robustness checks to compare these workers with the overall population.

⁴Note that, even under the prevailing wage, we would not expect take-up to be 100%. Recall that we sampled randomly from the labor force in each village when making job offers. Workers may decline the job because they have another work activity already lined up, or because their reservation wage is higher than the prevailing wage.

struggling for cash during this time. Our experimental results suggest that, among agricultural workers in our sample, 24.6% (i.e. those who would have accepted the job in private but do not do so in public) choose to forego these earnings in order to avoid being seen as violating the social norm.

The experimental results document that social observability reduces the willingness to accept wage cuts, but has no relevance for jobs at the prevailing wage. While we hypothesize that this stems from an attempt to avoid social disapprobation, other reputational mechanisms could also produce this pattern—for example, shame from being seen as financially desperate enough to accept a wage cut. In the second part of the paper, we provide positive evidence for social sanctions from accepting wage cuts.

We begin by providing suggestive evidence from a survey on worker perceptions. The majority of respondents (who did not participate in the field experiment) say that if a worker accepted a job below the prevailing wage, this would result in some sort of sanctions by other workers in the village. For example, 48% of respondents said that others would impede the worker’s future labor market opportunities through a decrease in referrals, and 20% said others would reduce social interaction like drinking together with the worker.

To provide more concrete revealed preference evidence on sanctions, we design a supplementary exercise. In another set of villages—drawn from the same population as our study villages—we partner with employers to make private job offers to a random subset of laborers (“workers”) at varying wage rates within each village. We then play a costly punishment game with another random subset of laborers (“players”) in each village who were *not* offered jobs. Each player is paired with an anonymous worker and told that the worker is either in the player’s own village, or in a village that is geographically far away. To implement the game, the player is told that his paired worker accepted a job at either (a) the prevailing wage or (b) 10% below the prevailing wage. The player can then give up some of his endowment to reduce the endowment of his paired worker.

As expected, we find that there is no punishment of workers who accept jobs at the prevailing wage. In contrast, when paired with a worker who accepted a wage cut, players punish the worker 37% of the time. When players do punish, the amount of money they deduct from those who violated the norm corresponds to 37.2% of average daily labor market earnings in our sample. In order to impose this punishment on their partner, the amount that players forego from their own endowment corresponds

to 7.4% of typical daily earnings. Finally, we find that the desire to punish norm violations is not limited to workers in one’s own village. Players also punish workers from distant villages who accepted a wage cut—even though that worker’s action has no scope to affect the player’s own labor market. These results are consistent with the literature on social preferences, which indicates that individuals will be willing to destroy their own surplus to punish those who have engaged in norm violations (Charness and Rabin, 2002).⁵

Finally, we use a descriptive exercise to explore whether the mechanism documented in our field experiment has potential relevance for wage rigidity. Using observational data from across India, we examine whether areas with lower social cohesion—where it may be more difficult to levy sanctions and maintain conformity in behavior—exhibit less rigidity. We proxy for social cohesion by exploiting variation in caste heterogeneity among agricultural laborers across Indian districts. Using the wage rigidity test from Kaur (2018), we find areas with higher social cohesion exhibit substantively more downward wage rigidity, with correspondingly higher levels of employment fluctuations (i.e. boom and bust cycles in employment).⁶ These patterns are consistent with the idea that social conformity can have aggregate implications for markets. Of course, this correlation with wage rigidity is only suggestive—it does not necessarily denote a causal relationship.

Our findings relate to the economics literature on norms and social influence. Economists have also explored these forces in the domain of social preferences, such as charitable giving, voting, and environmental conservation (DellaVigna et al., 2012, 2016). This work indicates that such “altruistic” behaviors are partly motivated by social pressure. We build on these insights to examine social pressure in a core economic domain: the labor market. There is emergent evidence on the impact of such forces on the labor market (Bursztyn and Jensen, 2017; Bursztyn et al., 2018). In this paper, we document that social pressure suppresses labor supply for jobs below the prevailing wage, generating conformity across workers within a village. These forces are sufficiently strong to lead workers to give up substantial amounts of money

⁵Our findings are also consistent with contagious punishment models (Ellison, 1994), in which norms are an equilibrium strategy that is enforced through decentralized sanctions. We should note, however, that the willingness to punish those in other labor markets—where the deviating party’s actions have no scope to affect one’s own payoffs—is particularly consistent with villagers viewing norm violations in moral or general terms.

⁶In contrast, areas with higher social cohesion do not respond differently to agricultural shocks per se. This suggests that the wage rigidity patterns are not simply due to differences in the agricultural production function across different areas.

during periods of unemployment. These findings suggest that social influence may have general relevance for economic behaviors and markets.

In addition, our findings relate to the literature on wage adjustment and labor market distortions in poor countries. Early work in development economics focused heavily on the observation that wages in poor countries appear downwardly rigid, potentially contributing to high levels of involuntary unemployment (Lewis, 1954; Eckaus, 1955; Leibenstein, 1957). Recent empirical evidence documents that downward nominal wage rigidity continues to be relevant in village labor markets today, with consequences for unemployment levels (Kaur, 2018). The presence of rigidities in this setting has been a long-standing puzzle in the development literature. A substantial body of theoretical work has proposed various micro-foundations for rigid wages (Shapiro and Stiglitz, 1984; Dasgupta and Ray, 1986). However, many of these proposed micro-foundations, such as nutrition efficiency wages, have not withstood empirical scrutiny (Rosenzweig, 1988). To date, there is scant empirical evidence supporting any micro-foundation for why wage floors should arise in this setting. Osmani (1990) offers a model based on informal worker collusion that theoretically reconciles the different stylized facts about wage adjustment in this setting. Our study provides the first empirical test of this mechanism.

Finally, our study has bearing on the labor literature on formal and informal unions. While a large literature has sought to understand formal unions in developed countries (see Farber and Saks, 1980; Dickens et al., 2007), there has been less work on the informal versions of these forces. Casual labor markets in poor countries display many of the same characteristics that are often rationalized by unions in developed markets: wage rigidity and wage compression (Kaur, 2018; Breza et al., 2016; Dreze et al., 1986). Consequently, documenting the presence of informal unions in our setting suggests that some of the considerations historically attached to formal unions may apply more broadly in the labor market.

While our design enables us to understand whether workers' labor supply is affected by collusive pressure, our evidence does not allow us to make predictions about what wage levels would exist in equilibrium in the absence of such pressure. Such predictions would require understanding the demand side of the market, which is outside the scope of our study. In addition, our findings will of course be specific to the five Indian districts (and 183 villages) in which our study was conducted. However, the features of our setting, such as wage rigidity and low employment rates, are mirrored

across India and in other developing countries (Kaur, 2018; Beegle et al., 2015). Providing the first piece of evidence for a potential micro-foundation for wage rigidity in such a setting would advance the literature and suggest an exploration of this micro-foundation in other locations as well.

The paper proceeds as follows. In Section 2, we describe the setting, research hypotheses, and experimental design. We present the results of the main experiment in Section 3, and the costly punishment game in Section 4. Section 5 discusses potential threats to validity. Section 7 concludes.

2. SETTING AND EXPERIMENTAL DESIGN

2.1. Overview of Setting. The study takes place rural Odisha, one of India’s most underdeveloped states. Markets for casual daily wage labor are extremely active, and provide the primary source of wage labor earnings for residents in the area. A large proportion of workers in construction, unskilled manufacturing, and other factories are hired through these labor markets. In the data collected by Breza et al. (2016), there appears to be a wage floor. This floor coincides with the prevailing wage in agriculture. Concretely, denote as W the prevailing wage in agriculture in a village. When workers from that village work in the non-agricultural sector – largely in jobs that take place outside the village – the wage they earn is at or above W in 98% of cases. This is despite the fact that unemployment levels appear high. Employment rates (in terms of total worker-days across all sectors) are below 50%. A striking 80% of workers report being involuntarily unemployed at least one day in the past two weeks. Understanding the source of the wage floor inside the village therefore has potential bearing on understanding determinants of the wage in the labor market as whole.

2.2. Research Hypotheses. As before, we denote the prevailing village wage as W . If worker collusion contributes to downward wage rigidity at W , then we hypothesize that during times of high unemployment, (at least some) workers would find it privately optimal to take up jobs at wages lower than the prevailing wage, but do not do so because this would result in sanctions from co-villagers. Specifically, we predict:

- H1.) The true private opportunity cost of working for a subset of individuals is less than W – i.e., workers will be privately willing to accept work at wages below W .
- H2.) When other workers can observe an individual’s job take-up decision, workers will be less likely to accept work below W .

H3.) Workers will sanction others who have accepted work below W .⁷

We construct a design to test these hypotheses, and rule out confounding factors.

2.3. Experimental Design. Our experiment takes place in five rural districts in Odisha, India. In each study village, we partner with a local agricultural employer (i.e. landowner in that village). We induce experimental variation in the wage rate offered by the employers, and in the observability of these offers to other laborers in the community. Partner employers typically hire daily-wage laborers for tasks like weeding and field preparation. Our experiment involves measuring the job take-up of each worker approached by the employer under the different treatment conditions below. Importantly, workers in the experiment make decisions about real jobs, working for an actual local employer who is typically familiar to the workers.

The research design requires that we have full experimental control over offered wages and the observability of those offers. In exchange for this control, we subsidize the cost of the labor for the employers. The size of the employer’s contribution is the same regardless of the size of the wage offer. This allows us to make employment offers in some treatment arms, described below, without the employer knowing the wage.⁸ Note that because we care about the labor supply side only, internal validity is not affected by the fact that the employer is being compensated for his cooperation (see Section VII).

Our full experimental design is presented in Figure 2. The core experiment follows a 2x2 design (treatment cells A-D). We incorporate 2 supplemental treatment cells (D and E) to allow for additional tests. Note that randomization is at the village (i.e. labor market) level, so that only one treatment cell will be implemented in any given village.

The first dimension of exogenous variation in the core 2x2 design sets the wage offer at either W or $W-10\%$. The second dimension changes the observability of the wage offer. In the core design, there are two observability conditions:

⁷We are agnostic as to how the community sustains the costly punishment of deviations by workers. Punishment may be a community norm or may be enforced through mechanisms similar to Ellison (1994) and Kandori (1992).

⁸Following the completion of work, wage payments are made to workers by members of the research staff. The employer pays his contribution directly to the research field team and never makes any payments directly to the workers. For ethical reasons, the employer in each village is aware of the possibility that we offer a wage below the prevailing wage.

- i.) Fully Private: Employment offers are made in private, inside the worker's home. The employer does not enter the worker's home, and the research team never informs the employer of the wage.
- ii.) Fully Public: Employment offers are made in public, on the street in front of the worker's home. The employer and any other passers-by can hear the terms of the wage offer.⁹

Treatment cells A and B give rise to basic tests of H1. First, the take-up rate in cell B measures whether there is any willingness to work below the prevailing wage when the offer is made in private. Under H1, take-up in cell B should be strictly positive. Second, a comparison of job take-up rates in treatment cells A versus B measures the fall in willingness to work attributable to a 10% lower wage, when job offers are made in private. Note that if workers do not believe that the information will be kept completely private, then this will result in a downwardly-biased estimate of an individual's true willingness to work below the prevailing wage, making it harder to validate H1. We return to this in our discussion of threats to validity in Section 5.

The difference in take-up rates between treatment cells B and D offers a basic test of hypothesis H2. This comparison identifies how much an individual's willingness to work below the prevailing wage falls when the take-up decision is made observable to the community.

By examining the impact of observability on take-up at the prevailing wage—i.e. the differences between cells A and C—we can validate whether observability itself has any impact on job take-up even when no community norm is being violated. We predict observability will have no impact in this case (i.e. A and C will have the same take-up rate). This helps us rule out a story where changes between B and D are not due to community pressure around wage cuts, but some other level shifter. Similarly, we can perform our test of H2 in a differences-in-differences framework by examining $(D-B) - (C-A)$ — thereby partialling out any level shifters in from observability itself.

Note that in moving from fully private to fully public, there is a change both in whether community members at-large learn about the wage rate and also whether the employer himself learns about the wage rate. Consequently, one interpretation of any differential take-up between cells B and D could stem from a desire to avoid

⁹Villages in the study districts are typically quite compact, with small dwellings that share adjacent walls and no real doors. When an employer visits the home of a worker, it is not uncommon for curious neighbors to overhear the wage offer. It is important to note that in the study villages, laborers and large employers live in distinct neighborhoods. Thus in most cases, all passers-by who overhear the wage offers will be individuals whose primary occupation is also wage labor and not landowners (i.e. employers).

having the employer learn the worker’s reservation wage, which may affect future bargaining dynamics with employers. This story would not necessarily rely on co-worker sanctions. To help distinguish between these two interpretations, we introduce a third source of variation in the observability of the wage offer.

- iii.) Partially Private: Employment offers are made in private, inside the worker’s home. However, the employer does enter the worker’s home and overhears the wage offer.

The difference (B-F) captures the aversion to taking a wage below the prevailing rate in front of an employer, while the difference (F-D) captures the aversion to accepting a wage below the prevailing rate in front of other laborers. Of course, if an employer is aware of the wage, then information transmission through the village may lead workers to learn it as well. Consequently, to the extent that we observe a take-up difference between B and F, we cannot disentangle whether it results from employer knowledge only, or the indirect channel of employer knowledge spreading to workers. We acknowledge that, to the extent that the purpose of collusive behavior is to enforce a collective bargaining outcome with employers, the difference between D and F is not necessarily well defined. Regardless, we view this as a useful additional source of variation.

2.4. Context and Protocols.

Context. This experiment takes place in 183 villages in four districts of Odisha, India. Agricultural production in these districts focuses mainly on paddy, which is both seasonal and labor-intensive. Over 70% of survey respondents are primarily engaged in agriculture, with 53% listing daily-wage agricultural labor as their main occupation. 91% of all respondents engage in daily-wage agricultural labor. There is strong baseline evidence of wage rigidity and wage compression in this area. In the survey conducted by [Kaur \(2018\)](#) in this area of Odisha, 100% of laborers and employers reported that they could not recall a year when the prevailing nominal wage in the village was lower than the wage in the previous year. This is consistent with the distribution of wage changes across India as a whole (Figure ??).¹⁰ In addition, the baseline survey evidence collected by [Breza et al. \(2016\)](#) indicates that there tends to be very little variation in wages within a village (Figure 3). Over 80% of agricultural

¹⁰The figure shows some areas where nominal wage cuts may occur. This could be driven by measurement error and compositional changes (see [Kaur \(2018\)](#)). More generally, the occurrence of wage cuts is not inconsistent with the presence of wage rigidity.

workers in a village receive the modal village wage. This is consistent with the presence of a clear wage norm that can be easily followed by laborers when making labor supply decisions.

In our setting, the village constitutes a prominent boundary for the labor market. Agricultural employers hire daily-wage laborers solely from within or close to their village. For example, in our pilot surveys, laborers report that 70% of worker-days in agriculture involve work within the village, and 97% of agricultural work-days take place within 5 kilometers of the village. In addition, workers within a village (i.e. those whose primary source of earnings is wage labor) tend to live the same tightly packed area of the village (referred to as the “labor colony”). This is common in India as laborers within a village generally stem from low-caste groups, and live in designated areas. Such an environment may be expected to make worker collusion and sanctions easier. Indeed, prior work suggests that the presence of strong within-village ties, risk-sharing, and job search networks allow sanctions to have a significant clout against those who violate a village norm.¹¹

We take advantage of a few distinct features of production in our study area. First, as is typical in subsistence farming, paddy production has lean periods in which employment declines (particularly between February and June, and again between September and November), which allows us to document that informal unions hold even during the eight months of the year when the opportunity cost of turning down a job is potentially high.¹² Second, the labor-intensive nature of paddy production in Odisha results in the ubiquitous use of casual daily-wage labor. Third, because of the uniformity of the crop produced in the region, we can work with local employers to offer nearly-identical agricultural jobs in different villages. We can work with everyone who participates in the agricultural daily-wage labor market, without selecting for people who have special skills or knowledge in a particular type of production or crop, which is helpful for the external validity of our results.

These features of village agricultural labor markets in Odisha (and, indeed, in many places elsewhere in India and the rest of the developing world)—siloe labor

¹¹See, for example, Townsend (1994), Chandrasekhar, Kinnan, and Larreguy (2014), Karlan, Mobius, and Szeidl (2014), and Dhillon, Iverson, and Torsvik (2013).

¹²We pause the experiment during the four months of peak demand in the agricultural labor market. The labor market is most likely to clear at this time, and workers’ alternative to taking up a job with us is taking a job with another employer at the prevailing wage with very high certainty. At these times of the year, a worker’s reservation wage of employment is likely to be the prevailing wage, leading us to expect zero takeup at below the prevailing wage in both public and private. While conducting the experiment during these times would be fascinating, for cost and sample size reasons, we stick to the remaining eight months of the year for our study.

markets, a clear and consistent prevailing wage, and a relatively homogeneous skill and knowledge base among those who participate in the agricultural labor market—make this an effective context in which to study the micro-foundations of the wage rigidity and wage compression we observe.

Protocols. Village Selection: We sample 183 villages in rural areas (i.e. at least 20 km from a town) across four districts in Odisha, India. We limited the sample to villages that have forty or more households in the labor colony.¹³ The sample means for key village and individual-level variables are presented in Table 1, along with balance across treatments.

Employer Selection: Once a village has been selected, we conduct a preliminary visit in which we ask an informant to list 20 employers in the village and tell us how much land they own and cultivate. After using this information to understand the distribution of land size in the village, we then recruit a mid-sized employer willing to hire up to three workers to work on his land within the next week.

Determining the Prevailing Wage: At this time, we also elicit the prevailing wage from the informant, and ensure that it has remained constant within the past month. In figure

Employers are told that the wage rate offered for a job on their land may be above or below the prevailing wage. They are not given any information about the level of observability of the job offers. Their contribution to the wage that will be paid to each worker is always Rs. 100 (approximately US\$.1.6) per worker hired. This ensures that their incentives to hire are not changed differentially across treatments, and enables us to keep wage offers blind to the employer in the fully private treatments. Our field staff accompanies the employer during hiring, as discussed below. After that point, we are not involved in the employment arrangement and the employer supervises the worker on his land as usual etc.

Treatment Assignment: Once an employer has been selected in the village, we assign the village to one of six treatment cells, in accordance with the sampling weights assigned to each treatment.¹⁴ Our unit of randomization is the village. Thus

¹³We use a floor on the size of the village to ensure heterogeneity across villages in the level of information spread, particularly in private treatments. In smaller villages, information may consistently be transmitted to all households.

¹⁴Specifically, we generated a random treatment assignment order in accordance with our desired sampling weights. Villages in the sample were then sequentially assigned to the next treatment assignment on this list as we moved through the study areas. In the second half of the sample, we also stratified by whether the labor colony population size in the village is above or below median

within a village, all workers are recruited under the same wage and observability condition.

Participant Selection and Hiring: We select one informant from the labor colony of the village to create a comprehensive census of daily-wage agricultural workers. We then partner with the employer to offer jobs to two-three workers (depending on the task, which was specified before treatment assignment and based on the area). We approached a random subset of workers in the labor colony with job offers, moving on to the next randomly selection household in case the worker was not home.

In accordance with local practice, job offers are made two days in advance, and at dusk, when the majority of workers are home. In our survey of employers in the area, employers in 60% of villages typically hire two days in advance of when they would like to complete the work, and less than two days in advance in 85% of villages. The norm in every village is to hire less than four days before the day of work.

In all hiring, the employer informed the worker he wanted to hire laborers to work on his land. The employer then introduces one of our field staff, saying "this person is here with me, and will give you some more details". Across all treatments, the field staff person accompanying the employer then relays the wage level to the worker. This enables us to keep which information is being conveyed constant across all the observability treatments.

Hiring - Public observability: Job offers are made outside the participant's home, which generally lead to others in the labor colony observing the job offer and take-up decision. Field staff do not interact with or provide the onlookers with any information directly. As we document below, on average there were 4 onlookers present when public hirings happened; these would typically have been other residents of the labor colony, i.e. laborers.

Hiring - Private observability: Job offers are made in the participant's home. After his initial conversation with the worker, the employer wanders away with a staff member out of earshot, while a second staff member continues the conversation with the worker and informs him of the exact wage level.

Hiring - Partially private observability: Job offers are made in the participant's home, but the employer remains present for the entire conversation.

Confirmation and Day of Work: On the day before work is scheduled, the employer and field staff confirm the work with those who accepted the job. This is the common practice by employers in our study area. On the day of work, the employer meets the

for the block (a geographical subunit of a district). We did not perform this stratification for the first half of the sample due to an oversight.

workers in his fields. The work itself proceeds as it would normally: the employer supervises the work, provides in-kind benefits like tea and lunch, without our staff present. Members of the research team do verify that the workers who agreed to the job actually work a full day. They also deliver the physical wage payments to workers at the end of the day across all treatments. This enables us to hold total wages confidential from the employer.

Surveys. After the workday is complete, we conduct a variety of surveys to provide further support for our hypotheses, including surveys with the employer, approached workers, and randomly chosen workers from the labor colony who were not approached for job offers.

3. RESULTS

3.1. Take-up of the Job. Panel A of Figure 5 presents the raw job take-up rates in the public and private treatments across wage offers made at the prevailing wage and at a 10% discount to the prevailing wage. In all the analysis, the outcome variable for job take-up equals 1 if the worker accepted the job: i.e. showed up to and completed the work. The figure shows that when the wage is set below the prevailing rate, take-up falls substantially when the job offers are made in public instead of private. However, when the wage is set at the prevailing rate, if anything, public offers lead to a weak increase in take-up rates. Table 2 presents the results in regression form across treatment cells. Cols. (1)-(2) report OLS regression results for the main experimental sample, where Fully private: Prevailing wage - 10% is the omitted category.

For job offers at a 10% wage cut, take-up falls by 13.6 percentage points when an offer is public versus fully private (Col. (2)). The results are similar with and without controls, with p-values of this difference ranging from 0.019-0.032 in Cols. (1)-(2) (p-values reported at bottom of table).

In contrast, for jobs made at the prevailing wage, take-up rates are positive in sign and statistically indistinguishable under public and private. This is consistent with no role for observability when the social norm is not being violated (i.e. under the prevailing wage), but an important role for social observation when workers are contemplating whether they will take up jobs below the prevailing wage.

In addition, our design enables us to gain some suggestive evidence on whether the difference between Fully Private versus Public is driven by the presence of the employer rather than presence of other workers. Our results suggest that in the presence of the employer (under Partially Private (Employer)), take-up of low wage

offers declines—the difference with Fully Private is negative but insignificant—but remains substantively higher than take-up under Public offers. In Col. (2), we can reject that (Partially Private W-10% = Public W-10%) at the 10% level (p-values reported at the bottom of the table). This test has limited statistical power due to the smaller sample size in the Partially Private treatment cells. As discussed above, it is also not perfectly interpretable as the incremental impact of co-worker observability, since the employer could also indirectly spread information to others in the village. However, these results support our presumption that pressure from other workers plays an important role in depressing labor supply below the prevailing wage.

Finally, note that the absolute magnitude of difference between Fully Private at the prevailing wage and Fully private at a 10% cut is 6-8 percentage points, but statistically not different than zero in any column. This is consistent with an underlying labor supply curve that is likely upward sloping, but not highly elastic around the prevailing wage. However, in the presence of social pressure (i.e. under the Public treatments), observed labor supply drops substantially below the prevailing wage—behavior that can reinforce a wage floor at the prevailing wage.

As explored in detail below, we hypothesize that the collective behavior of the informal union is stronger for “insiders” in the village union. We define an “insider” to be a worker who engages regularly in the agricultural wage market and show our main results for only this group of individuals. Panel B of Figure 5 shows the raw take-up rates for the subset of workers who self-identify in the exit survey as working in wage agricultural labor as a primary or secondary occupation. Interestingly, take-up rates in private look very similar to those in the full sample at both wage rates. However, we observe that when the below-prevailing wage offer is made public, take-up falls almost to zero. Col. (3) of Table 2 shows the results of the OLS regression specification, restricting the sample to this same group of “insiders.” Indeed, the key treatment effect of public wage offers holding fixed the below-prevailing wage rate increases in magnitude to 24.6% (significant at the 1% level). Moreover, for this group of insiders, the treatment effect is much larger when the low wage offer is made in front of other workers in comparison to when the offer is made only in front of the employer (25.6% versus 7.58%). This difference is also significant at the 1% level. These results provide preliminary evidence that workers who are members of the informal union are especially concerned with violating the village wage norms in front of other workers.

A potential concern with our design is that when public offers begin in the village, information may spread about the wage rate at which offers are being made. This, in turn, could affect which workers are available to be approached by the employer for a job offer (for the subsequent second and third offers made in the village). Note that information spread affecting the take-up decision is not a problem by itself. However, it is important that the treatments do not induce a type of differential selection into receiving a job offer. To ensure that this is not a problem, in Cols.(1)-(2) of Appendix Table A.2, we restrict our analysis to the first household and first two households, respectively, approached in our randomized list. While results are noisier due to the reduced sample size, the results are qualitatively similar.

Recall that during hiring, if a household that was approached for a job was not home, the employer moved to another household—approaching no more than six households total (regardless of how many people were home). The process of approaching three households in each village was sufficiently quick that we do not think any aspect of the wage level would have led to differential door locks. However, as a robustness check against this concern, in Col. (3) of Appendix Table A.2, we code any household that was not home as zero take-up. While this mechanically dampens the observed take-up levels across all treatments, and consequently predictably decreases statistical power, the results remain qualitatively similar. For job offers below the prevailing rate, the take-up difference between public and fully private offers is statistically significant at the 10% level. In addition, for jobs at the prevailing wage, take-up levels are similar across the observability levels.

3.2. Earnings. We then explore the impact of differing take-up behavior under the various treatment arms on participants' wage work and earnings. We hypothesize that suppression of labor supply due to the village norm (which can be enforced when job offers are made in public) will result in lower probability of working for a wage and lower earnings for participants on the day of work. However, if all participants who decided to turn down the job at W-10% in public (who were willing to work at the prevailing wage) were able to find alternative employment, then we should observe no difference for these outcomes for participants offered the job at W-10% in public and in private. Further, to rule out that inter-temporal substitution of work can sufficiently mitigate earnings losses from not taking up the job at below the prevailing wage, we also test labor supply and earnings over two longer periods: a) the day before work + five days after the day of work, and b) the day before work, the day of work, and five days after the day of work.

We test these hypotheses in Table 3. In Cols. (1) and (2), we first focus on probability of working for a wage and earnings on the day of work. We find that participants offered the job in Public at a wage of W-10% are 16pp less likely to work for a wage on the day of work, and correspondingly, earn Rs. 32 less for the day, on average, than their counterparts in the Private W-10% treatment. For a sense of magnitude, the Rs. 32 loss is 71% of the mean earnings on the day of work in the Private W-10% group. As a proportion of the mean earnings of the control group (not offered the job), the magnitude of earnings losses is even larger.

Second, in Cols. (3) and (4), we confirm that inter-temporal substitution of work cannot mitigate the earnings losses on the day of work. We find no significant change in earnings or in the probability of working in a 7-day window around (but excluding) the day of work, indicating that there is no spillover effect on other days in the same week. In Cols. (5)-(6), we include the day of work in this window to test the average impact of the treatment for the week. We find that, in that 7-day window (including the day of work), participants in the Public W-10% treatment are on average 7 pp less likely to work on any given day than those in the Private W-10% treatment, and earn Rs. 11.82 less per day. This translates to 49% loss in average weekly earnings relative to the group offered the job at W-10%, but in private.

3.3. Heterogeneous Treatment Effects. We next turn to two tests for heterogeneous treatment effects. First, we ask whether the treatment effects are any different statistically between insiders in the informal union and outsiders. We hypothesize that insiders, who rely on casual agricultural labor markets as an important income source, have more to lose from violating the village wage norm in front of other workers. We look for support for this hypothesis in Table ???. Here we define outsider in two different ways using each worker's exit survey responses: first, as above, an outsider is any individual who does not participate in the agricultural labor market as a primary or secondary occupation; second, we define an outsider as an individual who participates in the non-agricultural labor market as a primary or secondary occupation. Because we undersampled the partially private treatment cells, we pool Fully Private and Partially Private together for each wage. Cols. (1)-(2) present the main treatment effects, but using the pooled specification. Col. (3) presents heterogeneous treatment effect estimates using the first outsider definition, and Col. (4) does the same for the second definition. We find that insiders respond to a public low-wage offer with an 18-22 percentage point decrease in job take-up rates (both significant

at the 1% level). Moreover, this treatment effect is detectably smaller for outsiders, supporting our insider vs. outsider hypothesis.

Our view of the mechanism underlying our main results is that workers reduce their take-up of below-prevailing wage jobs when they worry that their decisions are observable to others in the village. Therefore, in villages where more individuals are likely to learn of worker take-up decisions, we hypothesize that the main treatment effects will be larger. We first verify the observability of our public treatments by showing that there are, on average, 4 onlookers within earshot of a public hiring treatment¹⁵. There is no statistically significant difference in the number of onlookers in villages in which the offered wage is W and villages in which it is $W - 10\%$ (Table 1).

In addition, following the completion of hiring in all villages, we returned to all but one of them to capture village-level characteristics that were not recorded in our initial exit survey. In this survey exercise we asked approximately five randomly-chosen workers per village a series of questions, some of which pertained to information flow in the village. We asked each individual the extent to which laborers learned about the wages at which others accept agricultural work; second, we asked how many others would find out if a worker accepted an agricultural job at below the prevailing wage. We aggregate responses at the village level and create an indicator for whether a village has below-median information flow. We predict that the magnitude of the treatment effect will be relatively smaller for these low-diffusiveness villages.

In Table 4 we explore heterogeneous treatment effects based on this measure of low information flow. Again, we augment the pooled average treatment effects specification. Publicizing a low wage offer in highly diffusive villages leads to an approximately 20 percentage point decline in take-up rates (significant at the 1% level) (Col. 1). However, in low diffusiveness villages, this large treatment effect is completely offset, leading to no measurable differences in take-up rates between public and private low-wage offers. These findings are consistent with our proposed mechanism. The same pattern holds when we instead use an alternate measure of diffusiveness that more directly addresses norm violation: the proportion of people in the village who would find out if a worker were to accept an agricultural job at below the prevailing wage. Finally, in Col. (3) of Table 4, we show that the magnitudes of both the decline in take-up rates and the offsetting effect in low diffusiveness villages are larger

¹⁵Only four public hiring interactions in our main sample were conducted with no onlookers. This information was recorded for each hiring interaction by one of our hiring staff.

for insiders in the union (participants with agriculture as a primary or secondary occupation).

4. EVIDENCE FOR SANCTIONS - COSTLY PUNISHMENT

The field experiment above establishes that social observability decreases labor supply at wage cuts. In this section, we provide positive evidence that accepting wage cuts results in sanctions. This helps distinguish the mechanism behind the observability effects from other types of reputational concerns, such as shame.

4.1. Survey Responses. We first use survey evidence to show that wage-setting happens in a decentralized manner, which highlights the role of social norms (and the sanctions that preserve the norm) as an equilibrium coordination device. In follow-up surveys with control group workers (Figure A.2), 89% of workers agree that there is no village-level meeting for all or most laborers in the village to discuss the wage, and 97% state that there is no meeting between laborers and landowners to bargain over the wage for the season.

We then tabulate worker perceptions of sanctions. In a survey with agricultural workers who did not participate in the experiment (i.e. were not offered jobs), we elicited beliefs about the consequences of accepting wage cuts. Respondents were asked “Suppose a laborer accepts work at a rate lower than the prevailing wage. What will be the reaction of other workers?” Respondents could agree with as many options as they wanted, or could give their own. We compile these responses into categories in Figure 7.

59% of respondents state that others would impede that worker’s future labor market opportunities. For example, a common source of non-agricultural employment is contractors, who come into the village and deputize a worker to round up a larger number of workers for an outside job. 56% of respondents said that a laborer who accepted the wage cut would not be included in such an opportunity. In addition, 17% of respondents said that a worker who accepts wage cuts would be excluded from social activities, such as drinking together. In contrast, only 1% of respondents rarely agreed with the notion that accepting wage cuts results in financial punishments—for example, refusal to help a laborer with a financial emergency in the future.

Workers also expressed a belief that social pressure is generally successful in preventing such actions to begin with (Panel B). 66% of workers stated that others would try to convince the worker not to accept a job at a wage cut. In addition, we asked all workers “If others try to convince such a worker not to take the job, will he still

do it?” 87% of workers said “No”, indicating their view that a worker would not go against group pressure.

Of course, such survey evidence is only suggestive. To obtain more direct revealed preference evidence on sanctions, we use a costly punishment game in a supplementary lab-in-the-field exercise.

4.2. Costly Punishment Exercise. In another set of 13 villages—drawn from the same population as our study villages—we again partner with employers to make job offers to a random subset of workers at varying wage rates within each village. These offers are always made in private. Each worker is first offered a job at 10% below the prevailing wage, and if he says no, is asked if he would be willing to work for the employer at the prevailing wage. We typically approached 10-15 workers with job offers in each village (with the number of workers per village decided *ex ante*). This larger number of offers guaranteed that in each village, at least some workers have accepted a wage cut. This sets up the backdrop for the costly punishment exercise.

Specifically, we then recruit another (random) subset of 8-12 laborers in each village who were *not* offered jobs. These other laborers, who we will refer to as “players”, are the ones who actually participate in the costly punishment game. Each player is paired with an anonymous worker (the “partner”) who received a job offer. The player and his anonymous partner are both given an endowment of Rs. 100. The player can “punish” his partner, reducing his endowment, by giving up some of his own endowment. Specifically, for every Rs. 5 that is removed from the partner’s endowment, the player must give up Rs. 1 of his own endowment. To make visualization easy, we implement this by placing 2 trays in front of the player, with Rs. 100 on each tray. The player then removes money from his tray and his partner’s tray, in accordance with the above proportion, until he is satisfied with the final allocations.

To conduct our test, we randomize two features of the partner’s characteristics. First, we randomly vary whether the player is partnered with a worker in the player’s own village, or is partnered with a worker in a village that is geographically far away. Note that in this latter case, the worker’s job acceptance decision has no direct consequences for the player, since the partner’s actions take place in a different labor market. Second, the player is told that his paired worker accepted a job at either (a) the prevailing wage or (b) 10% below the prevailing wage. The sample is weighted so that there is an equal number of observations in each of the $2 \times 2 = 4$ cells.

Furthermore, in order to obfuscate the reason for the exercise, we add in two "placebo" rounds of the game, which are played by the player before the above conditions.¹⁶ The player's payoff is determined by a random roll of the dice, in which one of his four rounds is implemented.¹⁷

If accepting a wage cut violates the social norm, then the literature on social preferences indicates that individuals may be willing to destroy their own surplus to punish those who have engaged in norm violations. In contrast, we do not expect to see punishment among workers who accept work at the prevailing wage—providing a helpful benchmark.

Figure 8 shows the estimated level of punishment under each scenario. As expected, there is virtually no punishment of workers who accept jobs at the prevailing wage. In contrast, when paired with a worker who accepted a wage cut from their own labor market, players punish the worker about 40% of the time. In addition, the desire to punish norm violations is not limited to actions in one's own village. Players also punish workers from distant villages in similar frequencies—even though that worker's action has no scope to affect the player's own labor market.

Table 5 presents these results in regression form. Col. (1) shows that, on average, the punishment probability increases by 42 percentage points when the "partner" accepts a wage lower than the prevailing level (statistically significant at the 1% level). Col. (2) shows that this effect size is of very similar magnitude and is statistically indistinguishable when the "partner" lives in a different labor market versus the player's own labor market. Cols. (3)-(4) show that these results are robust to village fixed effects and to considering only the first experimental round pertaining to the "partner's" labor supply decisions. Finally, Col. (5) shows that "partners" who accept a job below the prevailing wage from the same labor market receive payoffs that are about Rs. 15 smaller (on a base of Rs. 100).

¹⁶In each round, the player's paired partner is a different individual. In each of these earlier rounds, the paired worker undertakes a positive, negative, or neutral action: giving a gift of a bag of grainbaking someone a cake, stealing someone's bike, and traveling to the city for work.

¹⁷Note that the costly punishment game is played in the evening after job offers are made, but before the day of employment occurs. After the game is played, we announce that those laborers who do get jobs will receive the full prevailing wage (regardless of their initial response at the time of the wage offer). This enables us to fully preserve the anonymity of workers' take up decisions and prevent any sanctions outside the game.

When players do punish, the amount of money they deduct corresponds to 42.8% of average daily labor market earnings in our sample. In order to impose this punishment on their partner, the amount that players forego from their own endowment, conditional on punishment, corresponds to 8.6% of typical daily earnings.

These results are consistent with the literature on social preferences, which indicates that individuals will be willing to destroy their own surplus to punish those who have engaged in norm violations (Charness and Rabin, 2002). Our findings are also consistent with contagious punishment models (Ellison, 1994), in which norms are an equilibrium strategy that is enforced through decentralized sanctions. We should note, however, that the willingness to punish those in other labor markets—where the deviating party’s actions have no scope for equilibrium effects on one’s own payoffs—is particularly consistent with villagers viewing norm violations in moral or general terms.

5. THREATS TO VALIDITY

5.1. Internal Validity. We discuss some potential confounds that could contaminate the interpretation of the results of the main take-up field experiment.

Information Spread: One might worry that the private wage offers do not remain private. This is a valid concern, a priori, for several reasons. First, if the number of employment offers is high relative to village size, then even in the private offer condition, the wage offers will essentially become public. This would likely bias take-up at $W-10\%$ in private toward zero. As mentioned above, we limit the number of job offers to a small number in each village. The fact that we observe robust take-up in the private wage cut treatment (as opposed to close to zero take-up in the public wage cut treatment) validates our premise that at least a portion of workers believed that confidentiality would be maintained in the private treatment. To the extent that workers did not believe their take-up decision would remain confidential, this suggests our take-up estimates are a lower bound.

Number of onlookers: Our design rests on the idea that the presence of onlookers during public job offers will affect take-up behavior, because it directly enables observability by other laborers in the labor colony. We construct a direct measure of observability of the job offer by counting, for each public hiring interaction, the number of people within earshot of the participant at the time the offer was made.

We find that each job offer made under the public treatment arm is observed by, on average, 4 individuals¹⁸ (Table 1).

We conducted the hiring process in precisely the same way across all public treatment arms. In addition, to assuage any concerns that the difference in take-up rate under the two public treatments is driven by differences in observability of the job offers at the prevailing wage W and at $W - 10\%$, we validate that the number of onlookers was similar under both the Public treatments under the two different wage rates (Table 1).

Information about the prevailing wage: One potential concern with our design is that the public treatments provide workers with information about the prevailing wage—e.g., through potential comments from onlookers. This information, in turn, could depress take-up of public jobs below the prevailing wage. This is not consistent with this setting: the prevailing wage is general knowledge, as validated in our exit survey (Figure 4). As further evidence in support of this idea, in Appendix Table A.4, we document that among workers who were approached for job offers, reports of the prevailing wage are not systematically different across treatment cells. Importantly, there is no evidence that knowledge of the prevailing wage is different among Public and Fully Private treatments.

Signaling Poverty: It is possible that in some villages, only the poorest households might absolutely need to take a job below the prevailing wage. Thus, other households that might prefer to take such a job in private might worry that doing so in public might send a signal about their wealth to the community. If individuals experience disutility from being classified as very poor, such a mechanism could explain a fall in take-up at $W-10\%$ when the wage offer is public. The difference between cells A versus C, and G versus I, respectively, provides a possible suggestive test against such an explanation. If projecting status and wealth is desirable, workers should be marginally more likely to reject job offers in public versus private in these conditions as well, providing a helpful, albeit imperfect, placebo test. In addition, in exit surveys, the majority of workers state that accepting job offers below the prevailing wage would result in anger and sanctions from others—consistent with our hypothesized mechanism. Finally, our costly punishment game results provide positive support for sanctions. If accepting a wage cut is only costly because it is a sign of financial

¹⁸While we do not know if each of these onlookers is from within the village, our treatment effects suggest that at least some onlookers had a significant impact on participants' take-up decisions.

destitution, then it is unclear why workers would punish such individuals by taking money away from them.

Side payments: In our exit surveys we checked whether employers tried to compensate workers for the low offer wage by making side transfers, and do not find evidence for this. Furthermore, if such behavior were to exist, then it would most likely cause an increase in take-up across all W-10% treatment cells. Thus side payments can't rationalize our hypotheses.

5.2. External Validity. The magnitudes of our estimates are, of course, specific to labor markets in the study districts in Odisha, India, during the agricultural lean season. Our primary goal is to provide evidence that in our setting, villagers belong to informal unions and use social sanctions to enforce adherence to a village wage. While the shape and form of informal unions may vary across settings, we have reason to believe that the phenomenon of interest is not limited to rural Odisha. Several papers provide descriptive evidence consistent with worker co-ordination in setting and enforcing wages across rural labor markets in South Asia (Kaur, 2018; Osmani, 1990; Dreze et al., 1986). There is also reason to believe that similar informal institutions may exist in other settings. For example, Prothero (1912) points to a similar phenomenon in the early stages of industrialization in England. Further, the implications of informal unions – such as wage rigidity and low employment rates – are phenomena that are observed in many developing country contexts even outside of South Asia (Beegle et al., 2015).

6. CORRELATION WITH WAGE RIGIDITY

In this section, we explore whether the mechanism documented in our field experiment has potential relevance for wage rigidity. To motivate this link, we first document that workers believe that acceptance of a wage cut by one individual can affect the prevailing wage—generating a potential externality from an individual's labor supply onto the equilibrium wage for all workers. Appendix Figure A.3 documents worker beliefs, using data collected by Kaur (2018) in a different set of Indian villages that span 6 districts in two states. Workers recognize that if an individual worker agrees to work below the prevailing wage, he likely increases his own individual chances of employment (Question 1). 84% of workers in this broader sample also believe that other workers would get angry with such behavior, suggesting the relevance of our mechanism more broadly within India. In addition, 74% of workers believe that such behavior could lead other employers to try to pay lower wages for

future work (Question 3). This suggests that, according to worker beliefs, a sufficient number of deviations from the social norm could undermine the wage floor in the village.

The core hypothesis of our paper is that social pressure, and specifically the threat of social sanctions, is what prevents such deviations from occurring. In areas with less social cohesion, it may be harder to levy meaningful social sanctions: workers will be less socially integrated, potentially less reliant on each other (e.g., for leisure, marriage networks, job networks), and information will flow less well through the network (making it harder to learn about deviations and enforce them across the network). Consequently, a potential implication of our hypothesis is that, in areas with less social cohesion, there will be less downward wage rigidity.

We explore this potential implication using observational data from across India. In India, caste is a strong proxy for in-group and social cohesion (Munshi and Rosenzweig, 2006, 2016; Mazzocco and Saini, 2012). In our experimental sample, the composition of laborers is extremely cohesive by caste: all agricultural workers in a given village belong to the same caste (Scheduled Caste or Scheduled Tribe), and the median number of subcaste groups within a village is 1. This indicates a high level of social cohesion in our sample—helping explain the strength of our experimental results, with virtually no agricultural workers willing to accept wage cuts in public.

We exploit the fact that across India, the level of caste cohesion varies substantively. We use the National Sample Survey (NSS) household data (all employment rounds from 1983-2009, covering all of the 600+ districts in India). We measure caste heterogeneity by constructing a Herfindahl index of the caste composition of agricultural workers.¹⁹

To test whether social cohesion correlates with wage rigidity, we use the wage rigidity test developed by Kaur (2018). This paper tests how wages and employment respond to transitory labor demand shocks (generated exogenously by rainfall). The core result in the paper is that lagged positive shocks generate ratcheting in the labor market. Specifically, wages adjust upward in response to positive rainfall shocks. However, in the following year, when the positive shock has dissipated and rainfall is back at its normal levels, wages do not adjust back down—they remain ratcheted

¹⁹The NSS measures four caste categories: Scheduled Caste, Scheduled Tribe, Other Backward Caste, and General Caste.

up. Because of this distortion on the wage, agricultural employment falls: it is lower than it would have been in absence of the lagged positive shock.²⁰

In Col. (1) of Table 6, we replicate a simple version of this core test Kaur (2018). The omitted category in the regression is no positive shock this year or last year. Panel A, Col. (1) of Table 6 replicates the core result for wages from Kaur (2018). Relative to having no shock, wages rise robustly by 6.3% in response to having a positive shock this year. In addition, consistent with rigidities, lagged positive shocks also positively predict current wages: wages are 5.3% higher if there was a positive shock last year than if there had been no lagged shock.

We examine whether these wage rigidity effects are mediated by the level of social cohesion, proxied by caste heterogeneity. For this test, we add interactions of caste heterogeneity to the shock covariates. Such analysis is, of course, only suggestive. As with any heterogeneous treatment effects, our social cohesion proxies may be correlated with other factors, and may themselves be endogenously determined. We consequently view this as a descriptive exercise, not a causal one.

In Col. (2), we proxy for social cohesion by constructing a Herfindahl index of caste heterogeneity among those who are observed as doing any agricultural wage labor in the district. We interact each shock covariate with a dummy for a below median value of the index—indicating a diversity of castes among agricultural wage earners. In areas with high social cohesion, there is strong wage rigidity: lagged positive shocks lead to a 10% increase in current wages. However, in areas with low cohesion, the interaction term of -0.0826 offsets the level effect (significant at the 10% level), and we cannot reject that lagged shocks have no predictive power for future wages. In contrast, we do not see a strong interaction effect by social cohesion for current positive shocks; this serves as a placebo test, and suggests that places with high vs. low caste cohesion do not simply have different agricultural production functions. In Col. (3), we use an alternate definition for the cohesion proxy measure: the Herfindahl index of caste heterogeneity among all individuals who state that their primary or secondary occupation is agricultural wage labor. The results are similar to those in Col. (2).

In Panel B, we examine whether this correlation tracks the employment effects of rigidity. Col. (1) replicates the basic employment test. Employment rises in response

²⁰The paper also examines downward wage adjustment in response to negative shocks. However, there is no clean test for employment effects for negative shocks; the paper focuses on lagged positive shocks to look for employment effects. Consequently, we focus on the effect of lagged positive shocks here, which is the core test across outcomes in the paper

to current positive shocks. However, the following year, when wages are ratcheted above market clearing levels, employment is lower than it would have been in the absence of the lagged positive shock—consistent with boom and bust cycles. We add interactions with the proxy for social cohesion among agricultural wage earners in Col. (2). In areas with high social cohesion, lagged positive shocks lead to a decrease in weekly employment of 0.234 days. However, in areas with low cohesion, we cannot reject that there is no employment effect of lagged shocks: the interaction term of 0.189 (p-value of 0.03) almost fully offsets the level effect. This is consistent with the fact that there is no lasting ratcheting effect on the wage from lagged shocks in Panel A. As before, there is no significant interaction effect of social cohesion with current positive shocks.

These descriptive findings indicate that areas with low social cohesion exhibit larger levels of downward wage adjustment in response to labor market conditions. Consequently, areas with high levels of social cohesion exhibit not only more wage rigidity, but also higher levels of business cycle volatility. A causal analysis of such forces is beyond the scope of our paper. While only suggestive, the results in Table 6 are consistent with the view that social cohesion, and its resultant ability to lead to stronger social norms, could have aggregate implications by leading large numbers of workers to coordinate on the same strategy.

7. CONCLUSION

We find evidence that workers would privately like to supply labor below the prevailing wage, but do not do so when their take-up decision is publicly observable. This supports the hypothesis that collective pressure dampens labor supply below the prevailing wage, supporting the presence of wage floors in village labor markets. Our findings provide documentation of a way in which norms against accepting wage cuts distort labor supply behavior, with large impacts on the foregone earnings of unemployed workers.

Finding evidence that co-worker pressure dampens labor supply below the prevailing wage—even during times of high unemployment—provides impetus for exploring this mechanism in other settings. If this mechanism is indeed more generally applicable, then this can inform our understanding of the role of norms in shaping labor market outcomes, such as wage rigidity and wage compression.


REFERENCES

- ASCH, S. E. (1956): "Studies of Independence and Conformity: 1. A Minority of One Against a Unanimous Majority," *Psychological Monographs*, 70. 1
- BEEGLE, K., E. GALASSO, AND J. GOLDBERG (2015): "Direct and Indirect Effects of Malawi's Public Works Program on Food Security," *Working Paper*. 1, 5.2
- BREZA, E., S. KAUR, AND Y. SHAMDASANI (2016): "The Morale Effects of Pay Inequality," *NBER Working Paper 22491*. 1, 2.1, 2.4
- BURSZTYN, L., A. L. GONZÁLEZ, AND D. YANAGIZAWA-DROTT (2018): "Misperceived social norms: Female labor force participation in Saudi Arabia," Tech. rep., National Bureau of Economic Research. 1
- BURSZTYN, L. AND R. JENSEN (2017): "Social image and economic behavior in the field: Identifying, understanding, and shaping social pressure," *Annual Review of Economics*, 9, 131–153. 1
- CHARNESS, G. AND M. RABIN (2002): "Understanding social preferences with simple tests," *The Quarterly Journal of Economics*, 117, 817–869. 1, 4.2
- CHRISTIE, W. G. AND P. H. SCHULTZ (1994): "Why do Nasdaq market makers avoid odd-eighth quotes?" *The Journal of Finance*, 49, 1813–1840. 1
- CIALDINI, R. B. AND N. J. GOLDSTEIN (2004): "Social influence: Compliance and conformity," *Annu. Rev. Psychol.*, 55, 591–621. 1
- DASGUPTA, P. AND D. RAY (1986): "Inequality as a determinant of malnutrition and unemployment: Theory," *The Economic Journal*, 1011–1034. 1
- DELLAVIGNA, S., J. A. LIST, AND U. MALMENDIER (2012): "Testing for altruism and social pressure in charitable giving," *The quarterly journal of economics*, 127, 1–56. 1
- DELLAVIGNA, S., J. A. LIST, U. MALMENDIER, AND G. RAO (2016): "Voting to tell others," *The Review of Economic Studies*, 84, 143–181. 1
- DICKENS, W. T., L. GOETTE, E. L. GROSHEN, S. HOLDEN, J. MESSINA, M. E. SCHWEITZER, J. TURUNEN, AND M. E. WARD (2007): "How Wages Change: Micro Evidence from the International Wage Flexibility Project," *The Journal of Economic Perspectives*, 21, 195–214. 1
- DREZE, J., L. LERUTH, AND A. MUKHERJEE (1986): "Rural Labour Markets in India: Theories and Evidence," in *8th World Congress of the International Economic Association, New Delhi, December*. 1, 5.2
- ECKAUS, R. S. (1955): "The factor proportions problem in underdeveloped areas," *The American Economic Review*, 539–565. 1


- ELLISON, G. (1994): "Cooperation in the prisoner's dilemma with anonymous random matching," *The Review of Economic Studies*, 61, 567–588. 1, 5, 7, 4.2
- FARBER, H. S. AND D. H. SAKS (1980): "Why Workers Want Unions: The Role of Relative Wages and Job Characteristics," *Journal of Political Economy*, 88, 349–369. 1
- FREEDMAN, J. L. AND S. C. FRASER (1966): "Compliance without pressure: the foot-in-the-door technique." *Journal of personality and social psychology*, 4, 195. 1
- HENRICH, J., R. MCELREATH, A. BARR, J. ENSMINGER, C. BARRETT, A. BOLYANATZ, J. C. CARDENAS, M. GURVEN, E. GWAKO, N. HENRICH, ET AL. (2006): "Costly punishment across human societies," *Science*, 312, 1767–1770. 1
- HSIEH, C.-T. AND E. MORETTI (2003): "Can free entry be inefficient? Fixed commissions and social waste in the real estate industry," *Journal of Political Economy*, 111, 1076–1122. 1
- JACKSON, M. O., T. RODRIGUEZ-BARRAQUER, AND X. TAN (2012): "Social capital and social quilts: Network patterns of favor exchange," *American Economic Review*, 102, 1857–97. 1
- KAHNEMAN, D., J. L. KNETSCH, AND R. THALER (1986): "Fairness as a constraint on profit seeking: Entitlements in the market," *The American Economic review*, 728–741. 1
- KANDORI, M. (1992): "Social norms and community enforcement," *The Review of Economic Studies*, 59, 63–80. 1, 7
- KAUR, S. (2018): "Nominal Wage Rigidity in Village Labor Markets," *American Economic Re.* 1, 2.4, 10, 5.2, 6, A.3
- LEIBENSTEIN, H. (1957): *Economic backwardness and economic growth : Studies in the theory of economic development.*, New York: Wiley. 1
- LEWIS, W. A. (1954): "Economic Development With Unlimited Supplies of Labour," *The Manchester School*, 22, 139–191. 1
- MACLEOD, W. B. (2007): "Can Contract Theory Explain Social Preferences?" *The American Economic Review*, 97, 187–192. 1
- MAZZOCCO, M. AND S. SAINI (2012): "Testing efficient risk sharing with heterogeneous risk preferences," *American Economic Review*, 102, 428–68. 6
- MILGRAM, S. (1974): "Obedience to authority; an experimental view." . 1
- MUNSHI, K. AND M. ROSENZWEIG (2006): "Traditional institutions meet the modern world: Caste, gender, and schooling choice in a globalizing economy," *American*

- Economic Review*, 96, 1225–1252. 6
- (2016): “Networks and misallocation: Insurance, migration, and the rural-urban wage gap,” *American Economic Review*, 106, 46–98. 6
- OSMANI, S. R. (1990): “Wage determination in rural labour markets,” *Journal of Development Economics*, 34, 3–23. 1, 5.2
- PROTHERO, R. E. (1912): *English Farming, Past and Present*, Longmans, Green. 5.2
- ROSENZWEIG, M. R. (1988): *Labor markets in low-income countries*, Elsevier, vol. 1, 713–762. 1
- SHAPIRO, C. AND J. E. STIGLITZ (1984): “Equilibrium unemployment as a worker discipline device,” *The American Economic Review*, 433–444. 1

FIGURES



(a) Acceptability of Taking a Wage Cut. *Suppose it is the lean season. The prevailing wage is Rs. 200. To increase his chance of finding work, a laborer tells farmers that he would be willing to work any day that week at Rs. 180. Is the laborer's behavior acceptable?*



(b) Sanctions for Accepting Wage Cuts. *If a laborer accepts work at a rate lower than the prevailing wage, how likely is it that the other laborers in the village become angry?*

FIGURE 1. Survey Evidence

Note: These figures graph the exit survey responses from (N= 370) participants to questions about the acceptability of wage cuts and about other workers' responses to a worker taking a wage cut. We restrict the sample to participants from villages in which the participating employers offered jobs at prevailing wages.

		Wage Level		
		w	$w-10\%$	
Social Observability	Public	A	D	Job offer made on street in front of worker's home
	Employer only	B	E	Job offer made inside worker's home
	Private	C	F	Job offer made inside worker's home: employer out of earshot for wage

FIGURE 2. Experimental Design




FIGURE 3. Distribution of Wages Inside the Village. *Source:* Breza, Kaur, and Shamdasani 2017.

Note: This figure graphs the dispersion from the mode of the prevailing wage reported by untreated holdout sample households in a village. Wage reports are taken from exit surveys conducted in similar villages in Odisha in Breza, Kaur, and Shamdasani 2017.

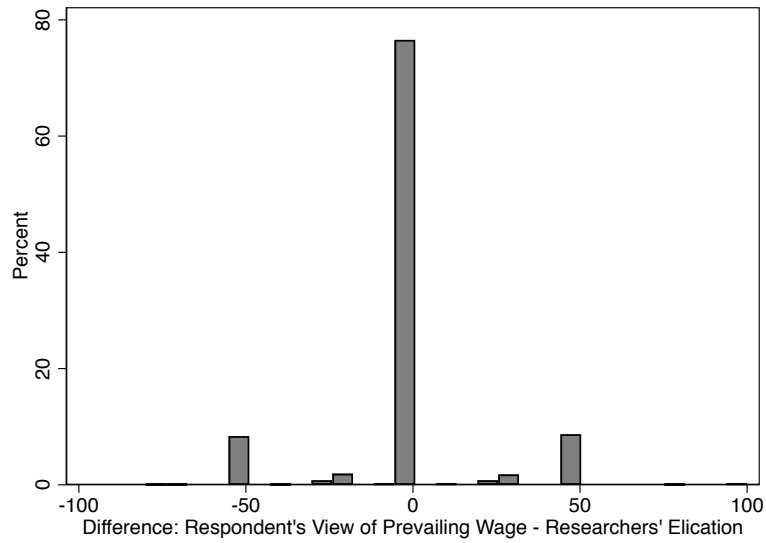




FIGURE 4. Untreated Holdout Group Reports of Prevailing Wage, Normalized by Informant Report.

Note: This figure graphs the distribution of the difference between the prevailing wage reported by untreated holdout households in a village and the prevailing wage reported by the informant in that village.



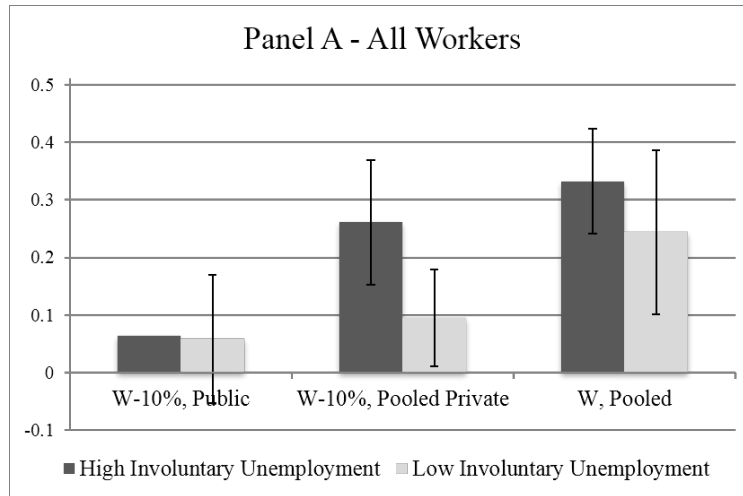
(a) All Workers




(b) Agricultural Workers Only

FIGURE 5. Job Take-Up by Treatment

Note: These figures graph the take-up rates for the job offer under different treatment arms. Job offers in each village are made either in public, with only the employer present, or in private, and are offered either at the prevailing wage, or at 10% below the prevailing wage. Each bar represents the take-up rate for the job as defined by attendance on the day of work. Panel A uses the entire sample ($N=502$ participants) while Panel B restricts the sample to casual daily wage laborers ($N=363$ participants), who report their primary or secondary occupation to be agriculture. All robust 90% CIs are constructed using standard errors from a test of the difference between the take-up rate for that treatment arm and the take-up rate for the public job offers at 10% below the prevailing wage. These results are also presented in the form of regressions in Col. (1) of Table 2. In Table 2, we include task and yearXmonth fixed effects in Col. (2) with the full sample and in Col. (3) with the sample restricted to casual daily wage laborers only.



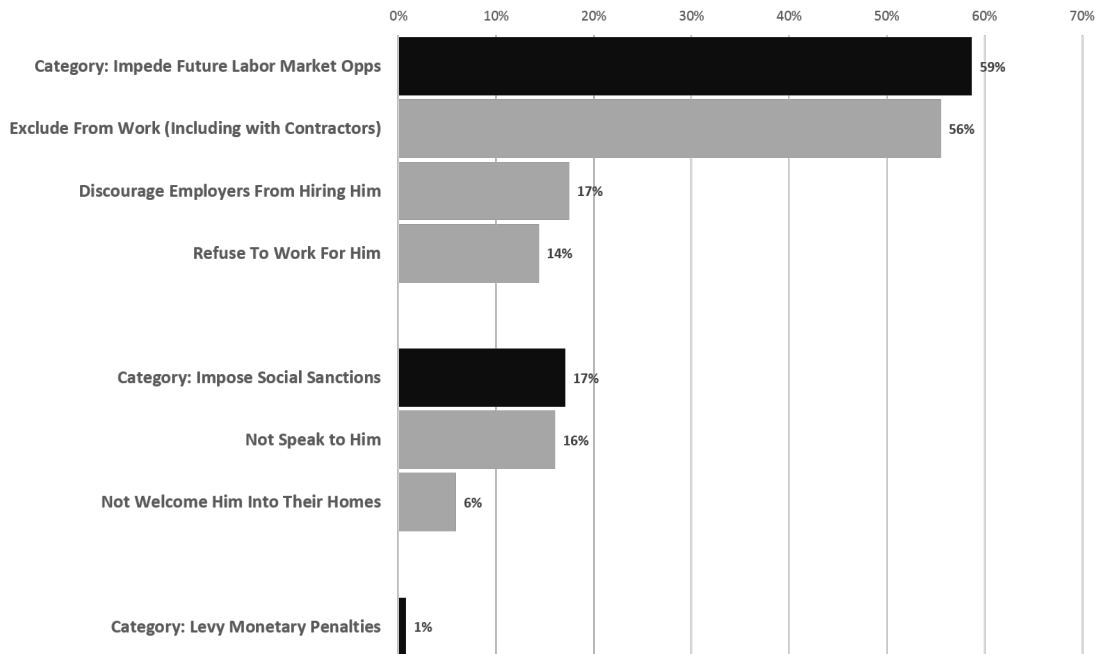
(a) All Workers




(b) Agricultural Workers Only

FIGURE 6. Job Take-Up by Level of Involuntary Unemployment

Note: These figures present heterogeneous treatment effects by below- and above-median involuntary village unemployment. Involuntary unemployment is measured by the mean number of days untreated holdout sample respondents who wanted work in the past 30 days but could not find it, as reported in the untreated holdout sample survey. Job offers in each village are made either in public, with only the employer present, or in private, and are offered either at the prevailing wage, or at 10% below the prevailing wage. Each bar represents the take-up rate for the job as defined by attendance on the day of work. Panel A uses the entire sample ($N=502$ participants) while Panel B restricts the sample to casual daily wage laborers ($N=363$ participants), who report their primary or secondary occupation to be agriculture. All robust 90% CIs are constructed using standard errors from a test of the difference between the take-up rate for that treatment arm and the take-up rate for the public job offers at 10% below the prevailing wage in villages with above-median involuntary unemployment.



(a) Sanctions



(b) Social Pressure

FIGURE 7. Survey Evidence - Sanctions for Accepting Wage Cuts

Note: This figure graphs exit survey responses from 1,448 untreated holdout sample households (who did not participate in our experiment) to the question: “Suppose a laborer accepts work at a rate lower than the prevailing wage. What will be the reaction of other workers?” Respondents were able to select as many responses as were applicable, and had the option of providing their own response. Responses were then aggregated into the categories shown in (a). In (b), we graph the proportion of respondents who answered ‘Yes’ to the question “If others try to convince him not to take the job, will he still do it?” to elicit workers’ anticipation of the effectiveness of the verbal sanctions.




FIGURE 8. Sanctions: Costly Punishment Game

Note: This figure graphs the results of the costly punishment game. Each participant (player) was anonymously paired with a worker in his own village or in a distant village, and given various scenarios about his paired worker. The figure plots the proportion of times players punished their paired worker under the 2 employment scenarios: (i) the worker accepted a job at the prevailing wage, or (ii) the worker accepted a job at a wage 10% below the prevailing wage. $N=131$ participants in 31 villages (villages are different from those in the main experimental sample). The plotted coefficients correspond to Col. (3) of Table 5.




FIGURE 9. Equilibrium wages and employment under different market structures

TABLES

SCABS: THE SOCIAL SUPPRESSION OF LABOR SUPPLY
TABLE 1. Covariate Balance

VARIABLES	(1)	(2)	(3)	(4)	(5)	(6)
<i>Treatments</i>	Wage Cut Private	Wage Cut Employer	Wage Cut Public	Prevailing All	Joint Sig F-Stat	Obs
INDIVIDUAL-LEVEL						
Age	44.337 (12.639)	2.890 (2.032)	1.893 (1.911)	0.881 (1.806)	0.180	442
Caste: Scheduled Tribe	0.333 (0.474)	-0.0179 (0.829)	-0.106 (0.0671)	-0.0259 (0.0755)	0.842	444
Casual Laborer	0.944 (0.23)	0.0417 (0.0346)	0.0604 (0.0269)	0.0486 (0.0267)	0.177	446
Casual Laborer - Agriculture	0.844 (0.364)	-0.0691 (0.0662)	-0.0398 (0.0620)	-0.00587 (0.0530)	0.840	446
Casual Laborer - Non-Agriculture	0.398 (0.492)	0.225 (0.0734)	0.161 (0.0614)	0.210 (0.0660)	0.001	502
Not in Agricultural Labor Market	0.136 (0.344)	0.0672 (0.0603)	0.0435 (0.0577)	0.00699 (0.0481)	0.375	502
Doesn't Own Land	0.592 (0.494)	-0.0835 (0.0851)	-0.0502 (0.0752)	-0.0260 (0.0694)	0.976	502
Individual Employment	8.841 (6.172)	-0.582 (0.978)	0.994 (1.030)	0.0841 (0.0699)	0.054	427
Low Individual Unemployment	0.523 (0.502)	-0.0858 (0.0851)	0.156 (0.0733)	0.114 (0.0871)	0.011	427
Employer Experience	0.318 (0.468)	0.188 (0.0987)	0.0346 (0.0999)	0.0645 (0.0845)	0.248	426
Employer Experience - Past Year	0.239 (0.429)	0.162 (0.0947)	0.0978 (0.0966)	-0.0716 (0.0769)	0.318	426
Considers Employer Not Influential	0.593 (0.494)	-0.0790 (0.102)	-0.0875 (0.0867)	1.563 (0.944)	0.713	383
Number of Onlookers	.	.	4.099 (0.472)	3.651 (0.472)	.	189
Took Exit Survey	0.874 (.334)	0.0178 (0.0495)	0.0284 (0.0458)	-0.00678 (0.0499)	0.389	502
VILLAGE-LEVEL						
Number of Households in Labor Colony	45.639 (13.135)	-3.283 (3.603)	-3.880 (3.266)	1.366 (3.057)	0.195	172
Low Village Unemployment	0.577 (0.501)	-0.0966 (0.144)	-0.124 (0.132)	-0.123 (0.126)	0.766	180
Low Information Spread Village	.541 (0.505)	0.0676 (0.134)	-0.0773 (0.123)	-0.0871 (0.114)	0.242	182
Number of Villages	37	34	40	72		
Number of Observations	103	88	108	203		

Notes: Table reports results from a regression of covariates of interest on treatment arms, restricting to the main experimental sample. The omitted category is households offered the wage cut (W-10%) in private. Means and standard deviations of each dependent variable for this omitted group are provided in Column (1). Columns (2)-(4) report coefficients for each treatment arm relative to the omitted treatment group. Standard errors relative to the omitted group are clustered at the village level and are reported below each coefficient in parentheses. P-values of tests of significance of coefficients in Columns (2)-(4) relative to this omitted category are reported below each coefficient in brackets. The Wald test of joint significance of all treatment arms (relative to the omitted category), p-value reported in Column (5). Variation in sample sizes (reported in Column (6)) comes from non-response in the exit survey and slightly different questions being asked in the exit survey across experimental rounds. The last two rows of the table report the number of villages and observations in each treatment arm. Regressions include yearxmonth and task fixed effects. Observations are weighted by the number of individuals in each village. In the first panel, coefficients are from regressions of individual-level characteristics collected in the exit survey. These variables include: respondent age, indicators for scheduled caste status, participating in the casual labor market, for participating in the casual agricultural daily labor market, for not participating in the casual agricultural daily labor market, and for not owning land, below-median proportion of days in the prior 30 that the participant wanted paid work but could not find paid work, indicators for having previous experience with the employer ever, and in the past year, and rating of the employer as 1, 2 or 3 on a 4-point influence scale, number of days in the prior 30 in which the worker earned a positive wage, number of onlookers during hiring, and an indicator for completing the exit survey. In the second panel, coefficients are from regressions at the village-level. These variables include: the number of households in the village, an indicator for below-median information flow as measured in the mop-up survey, and below-median village average of number of days the untreated holdout sample individuals would have preferred to work at the prevailing wage instead of their actual timeuse. The number of onlookers variable is recorded only for public treatments. The coefficient for the number of onlookers for hiring conducted at *W* in public is shown in Column (4). The test statistic in brackets for the number of onlookers in Column (4) is from a test of the equality of the number of observers in the two public treatments.

TABLE 2. Main Results

VARIABLES	(1) Worked	(2) Worked	(3) Worked
Wage cut: Public	-0.122 (0.0564)	-0.136 (0.0573)	-0.246 (0.0644)
Wage cut: Employer	-0.0657 (0.0611)	-0.0516 (0.0633)	-0.0758 (0.0788)
Prevailing wage: Private	0.0609 (0.0703)	0.0791 (0.0659)	0.0663 (0.0819)
Prevailing wage: Public	0.119 (0.0808)	0.116 (0.0713)	0.104 (0.0856)
Prevailing wage: Employer	0.0364 (0.0775)	0.0690 (0.0886)	0.0935 (0.0992)
Observations	502	502	363
Task and Year x Month FE		✓	✓
Sample	Main	Main	Ag. laborers
Depvar Mean (Wage cut: Private)	0.175	0.175	0.211
<i>Test</i> Wage cut: Private = Wage cut: Public	0.0316	0.0188	0.000181
<i>Test</i> Prevailing wage: Private = Prevailing wage: Public	0.460	0.589	0.658
<i>Test</i> Wage cut: Private - Public = Prev. wage: Private - Public	0.0629	0.0481	0.00858
<i>Test</i> Wage cut: Employer = Wage cut: Public	0.143	0.0865	0.0107
<i>Test</i> Wage cut: Private = Prevailing wage: Private	0.387	0.232	0.419
<i>Test</i> Prev. wage: Private = Employer = Public	0.609	0.816	0.904
<i>Test</i> Prevailing wage: Employer = Prevailing wage: Public	0.331	0.608	0.918
<i>Test</i> Prevailing wage: All = Wage cut: Private	0.196	0.118	0.228

Notes: This table presents the effect of each treatment on our main outcome of interest, the take-up rate for the job offer. In all specifications, the dependent variable is an indicator for whether the laborer accepted the job and worked for the employer. In all columns, the omitted category is the Wage cut: Private treatment. Cols. (1) and (2) include the full sample. Col. (3) restricts the sample to workers who answered the exit survey and who indicated that they engage in agricultural labor as a primary or secondary occupation. Observations are weighted by the number of experimental subjects in each village. Standard errors are clustered at the village level and are reported in parentheses.

TABLE 3. Earnings Results

VARIABLES	(1) Wage work	(2) Wage earnings	(3) Wage work	(4) Wage earnings	(5) Wage work	(6) Wage earnings
Wage cut: Public	-0.161 (0.0510) [0.00190]	-32.42 (11.13) [0.00405]	-0.0376 (0.0278) [0.177]	-6.794 (7.019) [0.334]	-0.0646 (0.0249) [0.0102]	-11.82 (6.942) [0.0903]
Prevailing wage (pooled)	0.0937 (0.0515) [0.0706]	27.97 (13.07) [0.0338]	0.0170 (0.0247) [0.491]	3.747 (6.167) [0.544]	0.0230 (0.0252) [0.363]	6.690 (6.399) [0.297]
Observations	428	428	1,303	1,303	1,731	1,731
Period	Work day	Work day	Ex work day	Ex work day	Full Week Weighted	Full Week Weighted
Sample	Endline recall	Endline recall	Endline recall	Endline recall	Endline recall	Endline recall
Task and Year x Month FE	✓	✓	✓	✓	✓	✓
Depvar Mean (Omitted)	0.222	45.49	0.0781	17.96	0.110	24.09

Notes: This table presents the effects of our job offers on individual earnings, derived from the employment recall grid performed in the exit survey. Each observation represents a day of recall. In Cols. (1), (3), and (5), the dependent variable is an indicator for whether the respondent worked that day for a wage. In Cols. (2), (4), and (6), the dependent variable is the total wage (cash + in kind) earned on that day in agricultural work. Cols (1)-(2) only consider responses for the day on which work was completed for the experiment in the village. Cols (3)-(4) consider the day before the work day and up to five days following the day of work, excluding the day of work. Cols.(5)-(6) include the day before work occurred in the village, the day work occurred in the village, and up to five days after the day of work. Non-workday observations are weighted to account for missing grid days in the worker exit survey (due only to the timing of the survey). Variation across respondents comes from the timing of when the exit surveys were conducted across households and villages. Observations are weighted by the number of experimental subjects in each village. Standard errors are clustered at the village level and are reported in parentheses.

TABLE 4. Heterogeneous Treatment Effects: Village Information Spread

VARIABLES	(1) Worked	(2) Worked	(3) Worked
Wage cut: Public	-0.200 (0.0675)	-0.186 (0.0646)	-0.308 (0.0745)
Prevailing wage (pooled)	0.0794 (0.0717)	0.0564 (0.0567)	0.0467 (0.0717)
Wage cut: Public x Low info spread village	0.170 (0.0932)	0.150 (0.0921)	0.214 (0.114)
Prevailing wage (pooled) x Low info spread village	0.0521 (0.0913)	0.115 (0.0844)	0.146 (0.106)
Low info spread village	-0.0732 (0.0667)	-0.0380 (0.0621)	-0.0263 (0.0796)
Observations	499	499	361
Task and Year x Month FE	✓	✓	✓
Low info definition	Wage info	Norm violation	Norm violation
Sample	Main	Main	Ag. laborers
Depvar Mean (Omitted)	0.204	0.200	0.214

Notes: This table presents heterogeneous treatment effects by village-level diffusiveness, as measured in the mop-up survey. In Col. (1), the heterogeneous variable of interest is an indicator for below-median knowledge of the wages of others. In Cols. (2)-(3), we use an indicator for below-median spread of information about other workers accepting a job below the prevailing wage. In all specifications, the dependent variable is an indicator for whether the worker signed up for the job and showed up for work. Col. (3) restricts the sample to only agricultural laborers. In this table, we pool Wage cut: Private and Wage cut: Employer. We also pool all of the Prevailing wage treatments together. In all columns, the omitted category is the Wage cut: Private pooled treatment for high info spread villages only. Observations are weighted by the number of experimental subjects in each village. Standard errors are clustered at the village level and are reported in parentheses.

TABLE 5. Costly Punishment Game

VARIABLES	(1) Any Punishment	(2) Any Punishment	(3) Any Punishment	(4) Any Punishment	(5) Partner's Payoff
Partner Accepts a Job Below Prevailing Wage	0.420 (0.0447)	0.393 (0.0632)	0.393 (0.0647)	0.436 (0.103)	-14.57 (4.425)
Partner Accepts a Job Below Prevailing Wage x Different Village		0.0494 (0.0894)	0.0494 (0.0916)	-0.00310 (0.137)	5.569 (4.551)
Partner lives in Different Village		0.0143 (0.0143)	0.0133 (0.0185)	0.00737 (0.0294)	-0.701 (1.259)
Observations	262	262	262	131	131
Village FE			✓	✓	✓
First Round Only				✓	✓
Depvar Mean: Partner Accepts Job at Prevailing Wage	0.00763	0.00763	0.00763	0	100

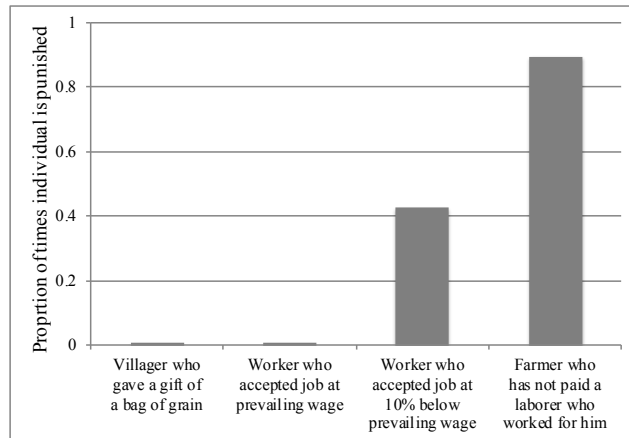
Notes: This table presents results from our costly punishment lab game exercise from N=131 participants (i.e. agricultural laborers) in 31 villages (villages are different from those in the main experimental sample). Each participant ("player") was anonymously paired with either another worker in his village or in a distant village, and given various scenarios about his paired worker. A player could take away money from his paired worker's endowment by giving up money from his own endowment. The table reports results under the 2 employment scenarios: (i) the worker accepted a job at the prevailing wage, or (ii) the worker accepted a job at a wage 10% below the prevailing wage. OLS regressions. The dependent variable in Cols. (1)-(4) is a dummy for whether the player punished the other worker at all; in Col. (5) it is the payoff of the anonymous partner (his initial endowment minus the amount deducted by the participant). Each player plays these two scenarios in random order; Cols. (4)-(5) report results only from the first of these two rounds. Standard errors clustered by player.

TABLE 6. Wage Rigidity: Correlation with Social Cohesion


	Proxy for Low Worker Cohesion		
		Wage Labor: Caste Herfindahl (Below Median)	Agri Labor Force: Caste Herfindahl (Below Median)
	(1)	(2)	(3)
<i>Panel A - Dependent variable: Log Agricultural Wage</i>			
Positive shock last year	0.0532 (0.022)	0.102 (0.042)	0.0971 (0.033)
Positive shock last year x Low worker cohesion		-0.0826 (0.050)	-0.0899 (0.038)
Positive shock this year	0.0633 (0.018)	0.0800 (0.038)	0.0751 (0.039)
Positive shock this year x Low worker cohesion		-0.0242 (0.042)	-0.0181 (0.043)
Observations (worker-days)	59243	59243	59243
<i>Panel B - Dependent variable: Agricultural Employment</i>			
Positive shock last year	-0.135 (0.055)	-0.234 (0.078)	-0.172 (0.080)
Positive shock last year x Low worker cohesion		0.189 (0.088)	0.0716 (0.107)
Positive shock this year	0.157 (0.062)	0.133 (0.083)	0.131 (0.091)
Positive shock this year x Low worker cohesion		0.0394 (0.114)	0.0469 (0.123)
Observations (workers)	632324	623861	631909

Notes: This table presents the effect of current and lagged productivity shocks on wages and employment (testing the ability of wages to fall after a positive shock has dissipated, a test for wage rigidity) and examines the heterogeneity of the effect by two measures of worker cohesion. The analysis uses National Sample Survey data (1986-2007). Low worker cohesion is defined as a) a below-median Herfindahl index of caste for all workers who engage in daily-wage labor within the village, indicating higher caste heterogeneity (Col. (1)), and b) a below-median Herfindahl index of caste for all workers who report agriculture as their primary or secondary occupation, indicating higher caste heterogeneity within the agricultural work force (Col. (2)). The dependent variable in Panel A is the log of the daily agricultural wage, and in Panel B is the number of days of agricultural employment (in wage labor or on one's own farm) in the past week. Positive shock is defined as rainfall above the 80th percentile of the district's usual rain distribution. Positive shock this year is a dummy for a positive shock in the current year. Positive shock last year is a dummy that equals one if the district had a positive shock last year and did not have a positive shock in the current year. The interaction terms are the binary variables defined at the top of each column. Regressions include year and district fixed effects. Standard errors clustered by region-year.

APPENDIX FIGURES



(a) Punishment across Scenarios



(b) Punishment Amounts

FIGURE A.1. Laboratory Games: All Scenarios

Notes: This figure shows punishment responses by N=131 lab game participants (players) to various scenarios about the behavior of the anonymous partner. Panel A shows the proportion of times players punished their anonymous partners under 4 different scenarios about partner behavior: (i) A villager who gave a gift of a bag of grain when it was needed; (ii) A worker who accepted a job at the prevailing wage (pooled across partners in own and other villages); (iii) A worker who accepted a job at 10% below the prevailing wage (pooled again across own and other villages); (iv) A farmer who hired a worker two months ago but still has not paid him. Panel B shows the amount (in rupees, out of a maximum possible of Rs. 100) deducted from the partner's payoff under scenarios (iii) and (iv), unconditional and conditional on punishment.

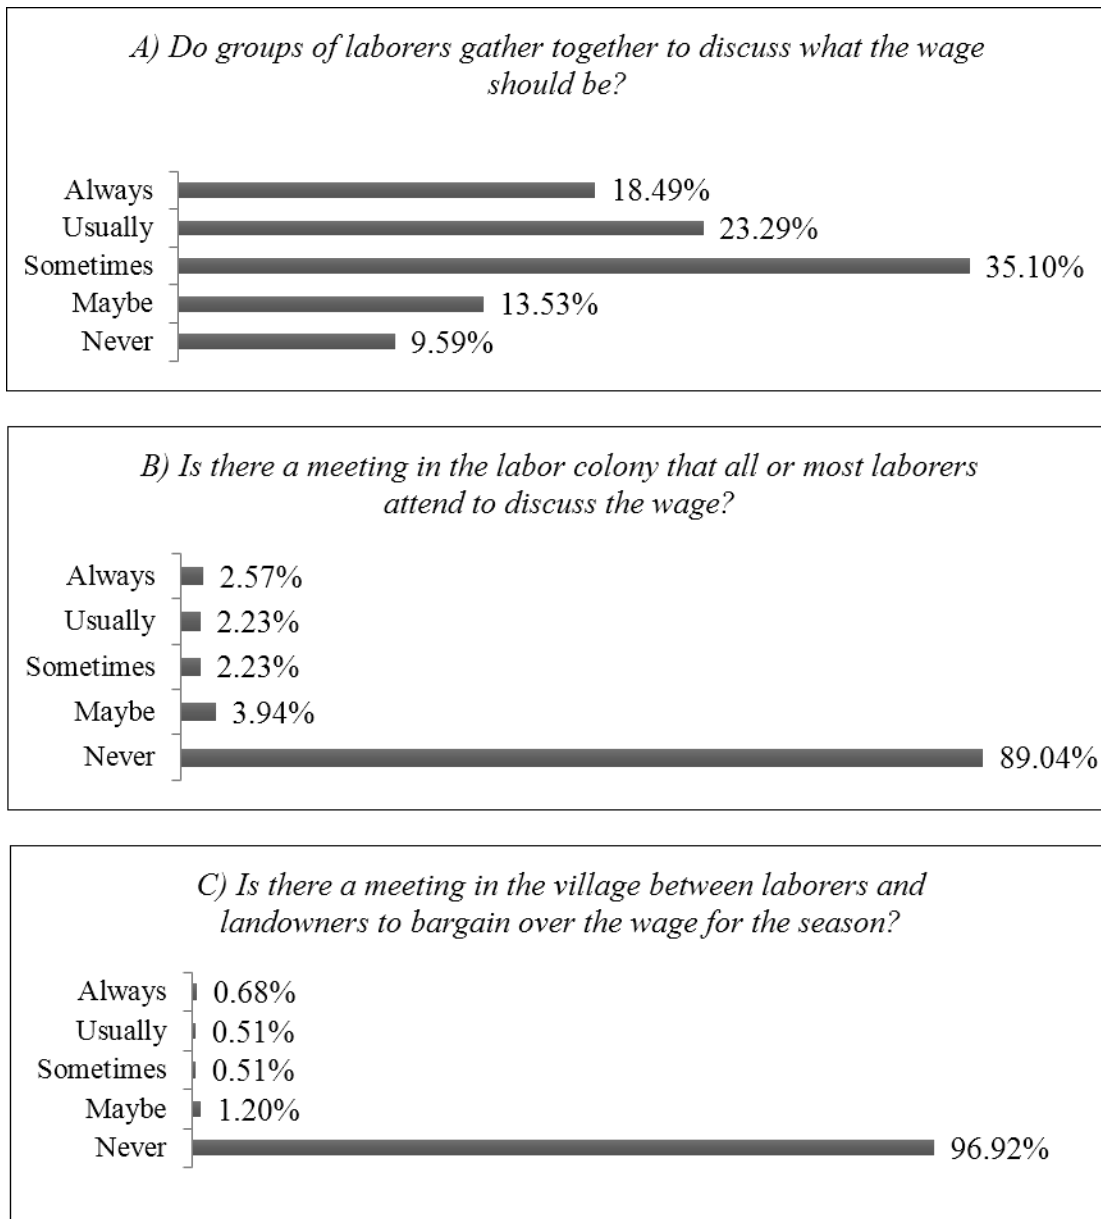


FIGURE A.2. Collective Action: Wage Setting in the Village

Notes: This figure shows responses to two survey questions about collective action in wage-setting within the village. Data are from the sample of untreated holdout sample households from experiment villages surveyed following the completion of the experiment. $N = 584$ male casual laborers in 183 villages.

Suppose a laborer was willing to accept work at a rate lower than the prevailing wage.

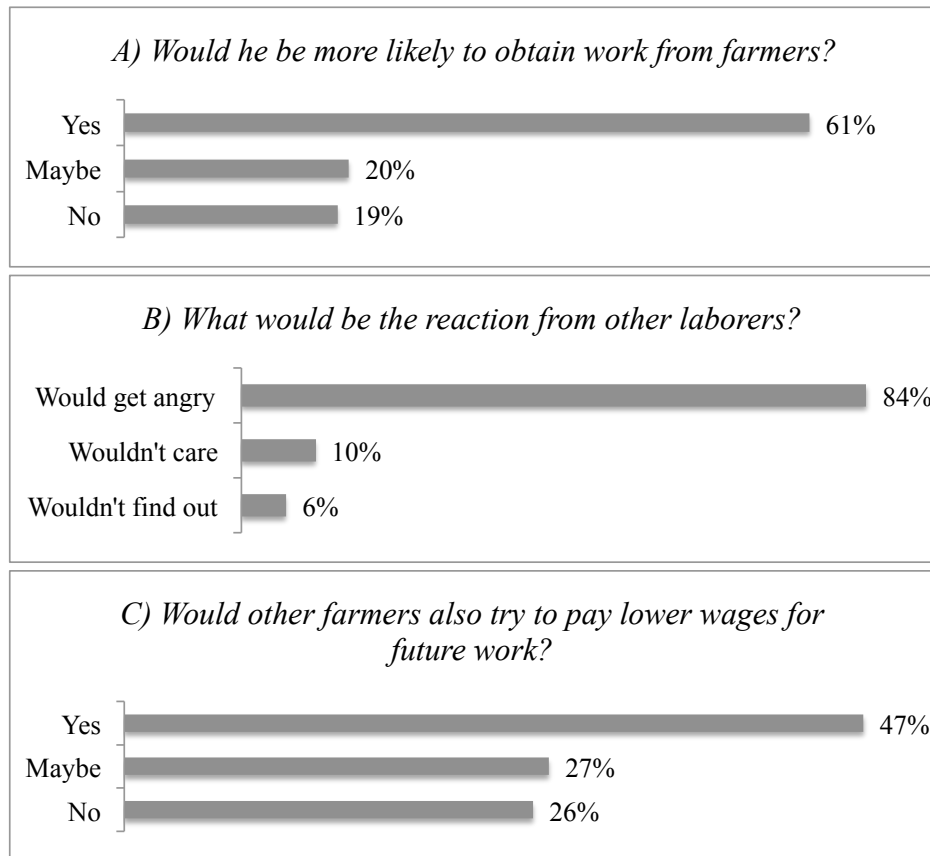


FIGURE A.3. Worker Beliefs: Impacts of Accepting Wage Cuts

Notes: This figure shows responses to three survey questions about employer and worker responses to a worker accepting a job at below the prevailing wage. Data from Kaur (2018). N = 196 male casual laborers in 34 villages across 6 districts in the states of Odisha and Madhya Pradesh.

APPENDIX TABLES

TABLE A.1. Main Results With Randomization Inference

VARIABLES	(1) Worked	(2) Worked	(3) Worked
Wage cut: Public	-0.122 [0.035]	-0.136 [0.032]	-0.246 [0]
Wage cut: Employer	-0.0657 [0.346]	-0.0516 [0.448]	-0.0758 [0.349]
Prevailing wage: Private	0.0609 [0.414]	0.0791 [0.334]	0.0663 [0.413]
Prevailing wage: Public	0.119 [0.157]	0.116 [0.191]	0.104 [0.272]
Prevailing wage: Employer	0.0364 [0.687]	0.0690 [0.538]	0.0935 [0.309]
Observations	502	502	363
Task and Year x Month FE		✓	✓
Sample	Main	Main	Ag. laborers
Depvar Mean (Wage cut: Private)	0.175	0.175	0.211
<i>Test</i> Prevailing wage: Private = Prevailing wage: Public	0.505	0.637	0.621
<i>Test</i> Wage cut: Employer = Wage cut: Public	0.124	0.107	0.012
<i>Test</i> Prevailing wage: Employer = Prevailing wage: Public	0.439	0.646	0.930
<i>Test</i> Prevailing wage: All = Wage cut: Private	0.195	0.137	0.244

Notes: This table presents the effect of each treatment on our main outcome of interest, the take-up rate for the job offer. In all specifications, the dependent variable is an indicator for whether the laborer accepted the job and worked for the employer. In all columns, the omitted category is the Wage cut: Private treatment. Cols. (1) and (2) include the full sample. Col. (3) restricts the sample to workers who answered the exit survey and who indicated that they engage in agricultural labor as a primary or secondary occupation. Observations are weighted by the number of experimental subjects in each village. Randomization inference p-values are reported in square brackets below each coefficient, and at the bottom of the table for relevant tests. Inference for the coefficients was carried out with 1000 permutations of treatments (at the village level), permuting over the treatment of interest and the omitted treatment, the Wage Cut: Private category.

TABLE A.2. Main Results: Sample Robustness

VARIABLES	(1) Worked	(2) Worked	(3) Accepted Offer
Wage cut: Public	-0.126 (0.0820)	-0.122 (0.0645)	-0.0817 (0.0474)
Wage cut: Employer	0.0260 (0.0911)	-0.0374 (0.0702)	-0.0377 (0.0493)
Prevailing wage: Private	0.0664 (0.100)	0.0788 (0.0754)	0.0598 (0.0598)
Prevailing wage: Public	0.136 (0.102)	0.0966 (0.0776)	0.0793 (0.0514)
Prevailing wage: Employer	0.126 (0.131)	0.137 (0.105)	0.0629 (0.0746)
Observations	188	359	545
Sample Restriction	First HH	First Two HHs	Include Door Locks
Task and Year x Month FE	✓	✓	✓
Depvar Mean (Wage cut: Private)	0.158	0.173	0.213
<i>Test</i> Wage cut: Private = Wage cut: Public	0.127	0.0611	0.0869
<i>Test</i> Prevailing wage: Private = Prevailing wage: Public	0.506	0.824	0.725
<i>Test</i> Wage cut: Private - Public = Prev. wage: Private - Public	0.139	0.171	0.170
<i>Test</i> Wage cut: Employer = Wage cut: Public	0.0628	0.161	0.241
<i>Test</i> Wage cut: Private = Prevailing wage: Private	0.508	0.297	0.318

Notes: This table presents results from our primary specification, restricted to various samples as a robustness check. In Col. (1), sample restricted to the first household approached in each village, and in column 2, sample restricted to the first two households approached in each village. In Col. (3), sample restricted to the intended main sample households in the village, including households where no respondent was home. In these cases, we code the outcome variable “Accepted Job” as 0 (job refusal). In all specifications, the dependent variable is an indicator for whether the worker signed up for the job and showed up for work. In all columns, the omitted category is the Wage cut: Private treatment. Standard errors are clustered at the village level and are reported in parentheses. Observations are weighted by the number of experimental subjects in each village.

TABLE A.3. Survey Attrition and Untreated Holdout Sample Composition

VARIABLES	(1)	(2)
	Has Exit Survey	Num Untreated Holdout Surveys in Village
Wage cut: Public	0.0342 (0.0514)	0.407 (0.37)
Wage cut: Employer	0.0124 (0.0525)	0.104 (0.355)
Prevailing wage: Private	0.0383 (0.0525)	-0.154 (0.397)
Prevailing wage: Public	-0.0857 (0.0662)	0.214 (0.433)
Prevailing wage: Employer	0.0696 (0.0554)	0.834 (0.486)
Observations	502	502
Task and Year x Month FE	✓	✓
Sample	Main	Main
Depvar Mean (Wage cut: Private)	0.879	5.364

Notes: This table reports survey attrition and untreated holdout sample composition by treatment arm. Col. (1) reports the likelihood of successfully completing an exit survey with an member of the main experimental sample, by treatment. The outcome variable in Col. (2) is the number of untreated holdout sample surveys conducted in the experimental household's village. In all columns, the omitted category is the Wage cut: Private treatment. Standard errors are clustered at the village level and are reported in parentheses. Observations are weighted by the number of experimental subjects in each village.

TABLE A.4. Exit Survey Reports of Village Prevailing Wage

VARIABLES	(1) 1(Agree)	(2) Difference	(3) Abs. Difference
Wage cut: Public	0.0442 (0.0713) [0.536]	-1.126 (3.246) [0.729]	-1.291 (3.051) [0.673]
Wage cut: Employer	0.0333 (0.0841) [0.692]	-1.900 (4.011) [0.636]	-1.266 (3.557) [0.722]
Prevailing wage: Private	0.123 (0.0771) [0.112]	-1.598 (4.109) [0.698]	-2.557 (3.718) [0.493]
Prevailing wage: Public	0.0579 (0.0856) [0.499]	2.640 (4.505) [0.559]	-0.109 (4.084) [0.979]
Prevailing wage: Employer	0.122 (0.0918) [0.185]	-0.675 (6.082) [0.912]	-3.194 (4.805) [0.507]
Observations	431	431	431
Sample	Main	Main	Main
Task and Year x Month FE	✓	✓	✓
Depvar Mean	0.800	5.650	8.875
<i>Test</i> Wage cut: Private = Wage cut: Public	0.536	0.729	0.673
<i>Test</i> Prevailing wage: Private = Prevailing wage: Public	0.399	0.369	0.561

Notes: This table presents statistics on the accuracy of the informant's report of the prevailing wage, relative to respondents' reports in the exit survey. Sample restricted to all experimental subjects who responded to our exit survey. In Col. (1), the dependent variable is an indicator for whether the respondent reports the same prevailing wage in the exit survey as the village informant reported prior to the intervention. In Col. (2), the dependent variable is the difference between the respondent's view of the prevailing wage and the informant's report. In Col. (3), the dependent variable is the absolute value of this difference. In all columns, the omitted category is the Wage cut: Private treatment. Standard errors are clustered at the village level and are reported in parentheses. P-values are reported in brackets.

TABLE A.5. Heterogeneous Treatment Effects: Experience with the Hiring Employer

VARIABLES	(1) Worked	(2) Worked
Wage cut: Public	-0.202 (0.0708)	-0.239 (0.0778)
Wage cut: Employer	-0.110 (0.0790)	-0.103 (0.0953)
Prevailing wage (pooled)	0.0232 (0.0794)	0.0376 (0.0880)
Wage cut: Public x Employer Experience	0.0636 (0.111)	-0.0184 (0.118)
Wage cut: Employer x Employer Experience	0.0494 (0.119)	0.00161 (0.143)
Prevailing wage (pooled) x Employer Experience	0.140 (0.123)	0.107 (0.135)
High Employer Experience	-0.0206 (0.0966)	0.0193 (0.108)
Observations	426	350
Task and Year x Month FE	✓	✓
Sample	Main	Ag. laborers
<i>Test</i> Wage cut: Public + Public x Experience = 0	0.163	0.0105
<i>Test</i> Wage cut: Employer + Employer x Experience = 0	0.557	0.395
<i>Test</i> Wage cut: Pub. + Pub. x Exp. = Wage cut: Empl. + Empl. x Exp.	0.316	0.0725
Depvar Mean (Omitted)	0.183	0.188

Notes: This table presents heterogeneous treatment effects by previous work experience with the participating employer, as measured in the worker exit survey. The specifications use an indicator for the worker having ever worked for the hiring employer in the past. Col.(3) restricts the sample to workers who indicated in the exit survey that they engage in agricultural labor as a primary or secondary occupation. In all columns, the omitted category is the Wage cut: Private pooled treatment for the low employer experience group only. Observations are weighted by the number of experimental subjects in each village. Standard errors are clustered at the village level and are reported in parentheses.

TABLE A.6. Heterogeneous Treatment Effects: Individual Unemployment History

VARIABLES	(1) Worked	(2) Worked	(3) Worked	(4) Worked	(5) Worked	(6) Worked
Wage cut: Public	-0.196 (0.0664) [0.00354]	-0.372 (0.0792) [5.24e-06]	-0.213 (0.0742) [0.00463]	-0.287 (0.0693) [5.36e-05]	-0.298 (0.0880) [0.000877]	-0.436 (0.0839) [5.84e-07]
Prevailing wage (pooled)	0.0712 (0.0645) [0.271]	0.0192 (0.0773) [0.804]	0.133 (0.0693) [0.0575]	0.124 (0.0798) [0.123]	0.0735 (0.0761) [0.336]	0.0313 (0.0889) [0.725]
Low Village Unemployment	-0.166 (0.0699) [0.0187]	-0.231 (0.0875) [0.00895]			-0.192 (0.0733) [0.00950]	-0.238 (0.0850) [0.00576]
Wage cut: Public x Low Village Unemployment	0.160 (0.0909) [0.0808]	0.355 (0.116) [0.00252]			0.200 (0.0989) [0.0449]	0.358 (0.114) [0.00201]
Prevailing wage (pooled) x Low Village Unemployment	0.0780 (0.0948) [0.412]	0.205 (0.124) [0.101]			0.134 (0.104) [0.201]	0.198 (0.125) [0.114]
Low Individual Unemployment			-0.0901 (0.0500) [0.0734]	-0.0988 (0.0605) [0.105]	-0.0822 (0.0484) [0.0912]	-0.0892 (0.0557) [0.111]
Wage cut: Public x Low Individual Unemployment			0.132 (0.0711) [0.0648]	0.154 (0.0685) [0.0262]	0.119 (0.0706) [0.0927]	0.140 (0.0687) [0.0426]
Prevailing wage (pooled) x Low Individual Unemployment			-0.00102 (0.0752) [0.989]	0.0310 (0.0877) [0.724]	-0.0144 (0.0753) [0.849]	0.0193 (0.0846) [0.819]
Observations	493	363	427	350	427	350
Sample	Main	Ag. Lab.	Main	Ag. laborers	Main	Ag. laborers
Task and Year x Month FE	✓	✓	✓	✓	✓	✓
Depvar Mean (Wage cut: Private, High unempl.)	0.333	0.393	0.262	0.282	0.262	0.282
<i>Take-up</i> Wage cut: Public, High unempl.	0.0611	0	0.0882	0.0370		
<i>Take-up</i> Wage cut: Public, Low unempl.	0.0308	0.0444	0.0500	0.0217		
<i>Test</i> Wage cut: Public, High - Low unempl.	0.928	0.114	0.474	0.165		

Notes: This table presents heterogeneous treatment effects by individual and village unemployment. Low village unemployment (Cols. (1), (2), (5), and (6)) is measured as below-median village average of days in the past 10 that the untreated holdout sample group reports preferring working for a prevailing wage to their timeuse on that day, aggregated to the village-level from untreated holdout sample surveys. Low individual unemployment (Cols. (3)-(6)) is defined as below-median days in the past 30 that the worker reports wanting work but is unable to find any, as measured in the worker exit survey. Cols. (2), (4), and (6) restrict the sample to workers who indicated in the exit survey that they engage in agricultural labor as a primary or secondary occupation. In all columns, the omitted category is the Wage cut: Private treatment at $W - 10\%$ for the workers with high (individual or village) unemployment. Observations are weighted by the number of experimental subjects in each village. Standard errors are clustered at the village level and are reported in parentheses.

TABLE A.7. Worker quality, amenities, and selection on the day of work: All treatments

VARIABLES	(1) Length of work (mins)	(2) Work day rating	(3) Hired before	(4) Number of meals included	(5) Total cash wage paid
Wage cut (pooled)	-24.11 (20.96)	-0.109 (0.219)	-0.00701 (0.122)	0.0283 (0.245)	-29.76 (8.920)
Constant	313.3 (16.52)	1.178 (0.141)	0.652 (0.0721)	0.690 (0.172)	219.8 (7.494)
Depvar Mean (Omitted)	313.3	1.178	0.652	0.690	219.8

Notes: This table presents statistics on worker selection into treatments according to quality, experience with the employer, and cash and kind components of the wage. The sample is restricted to all individuals who came to the job on the day of work. The omitted category is Prevailing wage (Pooled). Worker quality is reported for the day of work on a rating scale of 1-4 by the employer. Workers' prior experience with employer is measured by an indicator for having been hired before by the employer. Cash and kind payments are measured by the total cash wage and the number of meals provided by the employer, respectively, that are reported by the participant in the exit survey. Standard errors are clustered at the village level and reported in parentheses. N=74.

TABLE A.8. Worker quality, amenities, and selection on the day of work: Private treatments

VARIABLES	(1) Length of work (mins)	(2) Work day rating	(3) Hired before	(4) Number of meals included	(5) Total cash wage paid
Wage cut: Private	-3.542 (21.20)	-0.180 (0.325)	0.0441 (0.168)	0.0972 (0.359)	-17.01 (10.44)
Constant	284.4 (17.31)	1.118 (0.211)	0.706 (0.107)	0.625 (0.221)	208.1 (7.565)
Depvar Mean (Omitted)	284.4	1.118	0.706	0.625	208.1

Notes: This table presents statistics on worker selection into treatments according to quality, experience with the employer, and cash and kind components of the wage. The sample is restricted to all individuals who came to the job on the day of work in one of the private treatments. The omitted category is Prevailing wage: Private. Worker quality is reported for the day of work on a rating scale of 1-4 by the employer. Workers' prior experience with employer is measured by an indicator for having been hired before by the employer. Cash and kind payments are measured by the total cash wage and the number of meals provided by the employer, respectively, that are reported by the participant in the exit survey. Standard errors are clustered at the village level and reported in parentheses. N=34.