NBER WORKING PAPER SERIES

SOCIAL NORMS AS A DETERMINANT OF AGGREGATE LABOR SUPPLY

Emily Breza
Supreet Kaur
Nandita Krishnaswamy

Working Paper 25880
http://www.nber.org/papers/w25880

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2019, Revised February 2024

We thank Lauren Falcao Bergquist, Leonardo Bursztyn, Lucas Coffman, Stefano DellaVigna, Ernst Fehr, Andrew Foster, Pat Kline, Michael Kremer, and Jeremy Magruder for helpful comments and conversations. Arnesh Chowdhury, Shreoshee Mukherjee, Piyush Tank, Medha Aurora, and Sayan Kundu provided terrific research assistance. We also thank Kristina Hallez and Abigail Powers for editorial assistance. We gratefully acknowledge financial support from the National Science Foundation (Grant #1658937), the Agricultural Technology Adoption Initiative, and the Institute for Research on Labor and Employment (UC Berkeley). The project was registered in the AEA RCT Registry, ID 0001290. A previous version of this paper was circulated under the title "Scabs: The Social Suppression of Labor Supply". The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

Social Norms as a Determinant of Aggregate Labor Supply
Emily Breza, Supreet Kaur, and Nandita Krishnaswamy
NBER Working Paper No. 25880
May 2019, Revised February 2024
JEL No. D71,E24,J22,J31,J43,J50,O15,O17

ABSTRACT

In developing countries, the individuals that participate in the same localized market often share social ties—creating scope for collective behaviors that can generate market power. We test whether large groups of decentralized workers implicitly cooperate to prevent downward pressure on wages, using a field experiment with existing employers in 183 local labor markets in rural India. Only 1.8% of agricultural workers are willing to accept jobs below the prevailing wage despite high unemployment, but this number jumps to 26% when this choice is not observable to other workers—indicating substantial distortion in the aggregate labor supply curve. In contrast, social observability does not affect labor supply at the prevailing wage. In addition, workers are willing to pay to sanction those who accept wage cuts. Consistent with aggregate implications, measures of social cohesion correlate with downward wage rigidity and its unemployment effects across India. In line with our experimental evidence, sellers in other decentralized spot market settings in India and Kenya state they would be unwilling to adjust prices downwards, and would face strong social and economic repercussions if they do so. In developing countries, market power may be more widespread than previously believed.

Emily Breza
Harvard University
Littauer Center, M28
1805 Cambridge Street
Cambridge, MA 02138
and NBER
ebreza@fas.harvard.edu

Supreet Kaur
Department of Economics
University of California, Berkeley
Evans Hall
Berkeley, CA 94720
and NBER
supreet@berkeley.edu

Nandita Krishnaswamy
School of International and Public Affairs
Columbia University
420 W 118th St #1410
New York, NY 10027
nk2530@columbia.edu

A data appendix is available at
http://www.nber.org/data-appendix/w25880
A randomized controlled trials registry entry is available at
https://www.socialscienceregistry.org/trials/1290/history/

1. Introduction

In developing countries, many rural markets are organized as informal spot markets, where large numbers of individuals contract in decentralized transactions. This includes, for example, markets for labor, land, irrigation, and bullock rental. Explicit organization, such as the presence of unions, is rare. In principle, this embodies an ideal set of conditions for competitive markets.

However, in this setting, the individuals who operate in the same local market are also often bound by strong social and economic ties. For example, in village labor markets, workers are dependent on each other for socialization, informal insurance, and job referrals. We argue this creates the scope for market power to arise: groups can informally enforce norms around behavior, deterring self-interested deviations through the threat of sanctions (e.g. Osmani, 1990; Kandori, 1992; Ellison, 1994; Olson, 2009; Fehr and Schurtenberger, 2018).[1,2]

This paper tests whether a large set of decentralized individuals—spanning the workers in an entire labor market—implicitly coordinate to prevent downward pressure on the wage, by restricting labor supply at low wages. Specifically, we hypothesize that during times of unemployment, at least some workers would prefer working below the prevailing wage rather than remaining jobless, but do not do so due to the threat of sanctions from community members. As a result, the aggregate labor supply curve is distorted—shaped by social pressure rather than solely determined by underlying preferences for leisure—making it appear perfectly elastic at the prevailing wage.

We use a field experiment to test this hypothesis within the context of rural labor markets in Odisha, India. We obtain estimates of "private" versus "public" labor supply: whether the willingness to work for local employers at a given wage rate varies with community knowledge. In addition, we directly test whether unemployed workers face sanctions if they accept wage cuts. We run our test during the agricultural lean season, when unemployment is relatively high. Note that finding effects of social pressure on aggregate labor supply does not necessarily imply wage distortions; for example, this may even be efficiency enhancing in the presence of monopsony power on employers' side.

To implement the field experiment, we partner with existing agricultural employers, who make job offers to workers in their respective local labor markets. All jobs correspond to actual employment on the employer's land, so that our data reflect real employment decisions

[1]We use the term "norm" in a broad sense, to denote a set of decision rules that are known to all agents. We do not take a stance on whether the norm is also internalized into utility as "moral" behavior.

[2]The idea that social capital can enable groups to sustain cooperation has also been highlighted, for example, in work on public goods provision and informal insurance (e.g. Ostrom, 1990; Bloch et al., 2008; Jackson et al., 2012; Putnam et al., 2000).

by workers. We conduct the experiment across 183 labor markets (i.e. villages) with 183 partnering employers (one in each village), with jobs extended to 502 workers. Employers follow the typical process for agricultural hiring: offering jobs by approaching workers at their homes. In our setting, there is a commonly known prevailing daily wage for each type of agricultural task in the village—providing a clear benchmark for wage deviations.

To test for the effects of social pressure on labor supply, we induce two types of cross-cutting variation during employers' hiring process. First, the job is offered at a random wage level: at the prevailing wage, or 10% below the prevailing wage. Second, we vary the extent to which the wage level is publicly observable. In the "private" condition, the job is offered inside the worker's home and consequently not directly observable to others in the community. In the "public" condition, the employer offers the job outside on the street where neighbors, who are typically other workers, can overhear the offer, and could then presumably also tell other workers in the community.[3] Note that this design uncovers the *residual* labor supply facing employers: the fraction of workers who will take up a job at a given wage, conditional on not having found work elsewhere. As a result, take-up in our experiment provides a lower bound on underlying labor supply at any given wage.

A natural concern about our observability treatments is whether they may change worker beliefs about other aspects of the job, confounding their interpretation. We use the prevailing wage job offers as a placebo test against this concern. We predict that observability will only matter under norm violations (i.e. under wage cuts); it should not matter for jobs offered at the prevailing wage. This prediction distinguishes our hypothesis from this confound.

Treatment randomization is at the village level, so that all workers within a given village receive the same wage × observability condition. In addition, the workers in our experiment (i.e. those who are offered jobs) are sampled randomly from the village population of potential workers. Aside from inducing variation in hiring and payment, we are not involved in any other aspect of the employment relationship: employers supervise workers as usual, provide them food, etc. We verify that employment took place by visiting employers' farms at the start and end of each workday.

At the prevailing wage, mean take-up of private job offers is 26% across all potential workers in the village. Among those who report agricultural labor as their primary or secondary occupation (potential "agricultural workers"), 33% accept prevailing wage jobs. Importantly, take-up under private vs. public vs. employer wage offers is similar, with no detectable

[3]The full design randomizes between three levels of social observability: (i) the wage offer is publicly observable (i.e. to other workers), (ii) it is observable to the employer only but not other workers, and (iii) a fully private treatment where even the hiring employer does not know the exact wage level.

differences by observability (p=0.816). This suggests that, when offers accord with the wage norm, observability in and of itself does not affect workers' employment decisions.

In contrast, when a worker is offered a job below the prevailing wage, take-up depends crucially on whether his decision is observable to other workers. When a wage cut is offered in private, take-up remains a robust 18% among all potential workers. However, the willingness to accept a wage cut falls by 13.6 percentage points (74%) when offers are observable to other laborers (p=0.019). These results become even starker when restricting the sample to potential workers who are actually in the agricultural labor force: only 1.8% of agricultural workers accept wage cuts in public, relative to an estimated 26% in private (p=0.009).

The two observability conditions therefore paint dramatically different pictures of the aggregate labor supply curve. In private, when workers make decisions without the fear of social observability, labor supply appears upward sloping, with substantial willingness to work below the prevailing wage—even under our lower bound measure. In contrast, when other workers would learn about the contract, there is scant take-up of wage cuts—so that labor supply appears almost perfectly elastic at the prevailing wage.

The resultant distortion on individual labor supply is economically meaningful. Because unemployment is high, workers who decline wage cuts do not easily find other jobs. Relative to the private wage cut arm, those assigned to the public wage cut arm have a 59% lower employment rate (p=0.010) and 49% lower total wage earnings (p=0.090) on average over the week following the job offer. In other words, workers forego substantial economic benefits in order to avoid being seen by others as accepting low wages.

Consistent with our hypothesized mechanism, our effects are stronger in villages where individuals believe that information about accepting wage cuts would be more widely transmitted to other workers (i.e. where information diffusion appears stronger). In addition, we argue our findings are not driven by an employer bargaining motive—where workers refuse public wage cuts because they do not want to reveal their reservation wage to future employers. For example, workers robustly take up wage cuts despite this being observed by their hiring employer, who is a regular employer in the labor market; take-up only plummets under observability by other workers. Moreover, the pattern of our findings holds regardless of whether the worker is likely to work for the hiring employer again in the future, or whether the employer is relatively larger and more influential in the village.

Our field experiment demonstrates that social observability alters labor supply, but does not directly document sanctions. This is not possible precisely because sanctions are off the equilibrium path: when others can observe their choice, agricultural workers almost never accept wage cuts. In addition, among those that do accept wage cuts across treatment arms, for ethical reasons, we (unexpectedly) give them the remaining 10% of the wage at the end of the experiment, so that they can credibly report having received the prevailing wage.

Consequently, we collect two additional sets of data to examine whether accepting wage cuts leads to sanctions. First, we provide survey evidence on the nature of social disapprobation. 90% of workers state that others would be angry with someone who accepted a wage cut, while only 2% say others would be happy for him because he managed to find a job. The resultant sanctions enumerated by respondents range from a decrease in referrals for future work, to exclusion from social activities like drinking in the evenings. Other social motivations, such as shame, are rarely mentioned. Moreover, when asked about the consequences of accepting a wage *above* the prevailing wage, these patterns reverse: the percentage saying that others would be angry and happy is 2% and 95%, respectively. This is consistent with a norm against accepting wage cuts, rather than wage deviations in general.[4]

Second, to construct a revealed preference test for whether workers desire to sanction those who accept wage cuts, we build on our field experiment to design a costly punishment game. In the game, the decision-makers are randomly-chosen workers ("players") who were *not* offered jobs. Each player is paired with an anonymous worker who was offered a job with a local employer, as per our field experiment protocols. Using a 2x2 design, the player is told whether his paired worker: (i) accepted a job at the prevailing wage or at 10% below the prevailing wage (with wage levels communicated in Rupees), and (ii) lives in the player's own village or a village that is far away. The player can give up some of his endowment to reduce the endowment of his paired worker, thereby "punishing" the worker. We obfuscate the reason for the exercise by including non-job related scenarios in the game.

As expected, there is no punishment of workers who accept jobs at the prevailing wage. In contrast, 37% of players choose to pay money to punish those who accepted a wage cut. Sanctioning occurs even if the paired worker is from a distant village—where the worker's action cannot impact the player's own labor market—suggesting that individuals view violations in moral terms, consistent with internalized norms (Charness and Rabin, 2002). While this does not rule out other ways in which social observability could affect labor supply in our experiment, as a whole, we view the above as providing positive evidence that accepting wage cuts leads to retaliation by other workers.

[4]This is consistent with workers' stated beliefs that the wages earned by any one laborer can contribute to changes in the equilibrium wage for all workers. Note that these patterns are also consistent with Breza et al. (2018b): in our sample, workers express a preference for wage equality among co-workers performing the same task together for a given employer, but relative pay concerns are less salient across the decentralized daily spot contracts with different employers that are typical in this setting.

An implication of our findings is that places with stronger worker collective action may be more successful at preventing downward pressure on wages.[5] In the final part of the paper, we explore this idea further through a supplementary suggestive exercise using observational data across India. We use caste homogeneity among workers to proxy for social connectedness and the ability to levy sanctions—which are likely to make collective action easier to sustain.[6] Using the wage rigidity test from Kaur (2019), we find that areas where workers are more socially cohesive exhibit substantively greater downward wage rigidity in response to labor demand shocks (generated by rainfall). Such areas also have correspondingly greater reductions in employment in instances where downward wage adjustment is needed, leading to larger boom and bust cycles in employment. These patterns are consistent with the idea that implicit cooperation could have aggregate implications for markets. Of course, this correlation with wage rigidity is only suggestive—it does not necessarily denote a causal relationship.

While the above is consistent with the prediction that a norm against accepting wage cuts could exacerbate the employment effects of transitory fluctuations, wage floors could still benefit workers overall. Under monopsony power, they could increase both worker surplus and average employment levels. Even in the absence of monopsony power, we show in a simple back-of-the-envelope exercise that wage floors can allow workers to claim substantially more surplus from employers, with minor deadweight loss (see e.g., Lee and Saez, 2012).

We conclude by presenting survey evidence from other market sectors in Kenya and India: urban labor stand workers, taxi drivers, food vendors, and butcher shops. In each sector, most sellers state they would not accept a price below the prevailing price even when in need of money. Moreover, doing so would result in consequential sanctions from others in their local market, including potential loss of livelihood. While only speculative, these patterns open the door to the possibility that social forces may be relevant for understanding price compression, price rigidities, and market power in developing countries more broadly.

Our study contributes to three sets of literatures. First, a small body of work tests for market power in developing countries—primarily by measuring price effects in response to price or firm entry shocks (Bergquist and Dinerstein, 2020; Casaburi and Reed, 2021; Houde et al., 2020; Busso and Galiani, 2019; Atkin et al., 2018; Atkin and Donaldson, 2015; Iacovone and McKenzie, 2019). For example, this work finds imperfect competition among sellers who work daily within the same local market (Bergquist and Dinerstein, 2020; Busso and

[5]This, in turn, could affect employment levels. We see speculative signs of this possibility in our field experiment data: the difference in take-up of wage cuts in public relative to private (a potential proxy for the strength of social pressure) is positively correlated with baseline involuntary unemployment levels.
[6]We are unable to examine caste heterogeneity in our field experiment, because workers within a village are largely homogeneous by caste (jati).

Galiani, 2019), but not among traders who work separately across large sets of villages (Casaburi and Reed, 2021). While this literature diagnoses market power, it generally does not take a stance on mechanisms. We complement this work by providing the first empirical evidence for a micro-foundation that could give rise to collective action in settings with large numbers of atomistic, decentralized agents: social repercussions among individuals who interact repeatedly.[7] This, along with other explanations such as spatial frictions, offers guidance on when market power may be more likely to arise, and helps rationalize prior findings in the literature. Our study also implies that market power may potentially be more widespread than has been previously realized, given that the features of our setting—individuals with social or economic ties who operate in the same localized market—are common in a range of developing country settings.

Second, our findings advance the literature on labor markets in poor countries. We document a way in which social forces alter the aggregate labor supply curve, offering evidence on a novel channel that could affect labor market functioning. Why wage rigidity and unemployment should arise in this setting remains a long-standing puzzle in the development literature (Lewis, 1954; Leibenstein, 1957; Kaur, 2019; Breza et al., 2021). Theoretical work has proposed myriad micro-foundations, from implicit insurance to nutrition efficiency wages (e.g. Azariadis, 1975; Shapiro and Stiglitz, 1984; Dasgupta and Ray, 1986), but there remains limited evidence on their potential relevance for the labor market equilibrium.

The idea that workers may implicitly cooperate to maintain wage floors has been discussed in qualitative accounts (Rudra, 1982; Dreze et al., 1986; Walker and Ryan, 1990), and was modeled by Osmani (1990), but has not been empirically tested. More generally, there remains scant empirical evidence for any particular mechanism that could potentially interfere with equilibrium wage adjustment in this setting. Our study offers the first direct evidence for one such potential mechanism, which of course may not be mutually exclusive with other channels. Moreover, by documenting that "informal unions" can arise in unorganized settings, our findings have potential relevance for the broader labor literature on wage adjustment (e.g. Mathewson, 1969; Solow, 1990).[8]

Third, a growing body of work documents the relevance of social observability and social norms (e.g. Bandiera et al., 2005; Mas and Moretti, 2009; DellaVigna et al., 2012; Bursztyn and Jensen, 2015; DellaVigna et al., 2016; Bursztyn et al., 2017; Karing, 2021). While these studies show the relevance of social influence, there has been limited empirical work

[7]Note that social repercussions encompass various channels—including social punishment, reputation, and shame—all of which could deter self-interested deviations by agents.

[8]Mathewson (1969) offers qualitative historical accounts of such behavior among unorganized US workers. Solow (1990) argues that because repeat interactions among co-workers are an inherent feature of the workplace, social pressure likely alters labor supply within firms in rich countries today.

documenting the mechanism through which norms deter violations, though the role of social sanctions has been hypothesized as a likely reason (see Bursztyn and Jensen (2017) for an excellent review). Consistent with this view, a rich body of laboratory studies documents that when individuals are seen as violating a norm, others will sanction them—even in the absence of direct benefit to themselves (e.g. Gachter and Fehr, 1999; Charness and Rabin, 2002; Henrich et al., 2006; Fehr and Schurtenberger, 2018). We build on this literature by offering positive evidence that the threat of sanctions plays a role in affecting high-stakes field behavior—labor supply—with meaningful employment consequences. Moreover, we expand this literature by documenting that social influence can generate cooperative behavior among decentralized individuals at the level of an entire labor market. This indicates that the effect of norms on individual behavior could aggregate up to impacts on market equilibria.

2. Context and Motivation

2.1. **Context: Village Labor Markets.** Markets for casual daily labor are ubiquitous across poor countries. They constitute an employment channel for hundreds of millions of workers in India alone, and account for the 98% of the country's hired agricultural labor (NSS, 2010). Casual labor markets are characterized by a high degree of decentralization and informality (e.g Rosenzweig, 1988). Contracts are typically bilaterally arranged by individual employers and workers, and are short in duration (typically 1-3 days). The intermittent nature of any individual employer's labor demand, along with the relatively large number of workers and employers in a village, leads to frequent resorting of hiring relationships. Minimum wage laws are ignored, and formal unions or other formal labor market institutions are virtually non-existent.

We conduct our study during lean months during the agricultural year. During this time, workers in our sample report a daily employment rate below 50%, and being involuntarily unemployed for 12 days in the past month. This creates the potential for substantial willingness to supply labor below the market wage.

In our study areas, the village constitutes a prominent boundary for the labor market. For example, workers report that 70% of days in agricultural labor occur within the village, and 97% within 5 kilometers of the village. All of our field activities are conducted in the village "labor colony", from which agricultural employers draw the majority of their hired labor. These are densely populated neighborhoods where 84% of males engage regularly in daily-wage agricultural labor.

The process through which agricultural workers are typically hired looks similar across our study villages. The employer typically travels to the labor colony one or two evenings before the intended work will take place. He offers jobs to workers who may or may not have a prior job history with the employer. Moreover, the offers may be made in the street or in a central square of the village, or at the worker's house.

There is a clear prevailing wage for agricultural labor in our study areas, with limited variation in the wage paid to workers of the same gender within a village at a given point in time (Figure 1(a)). In addition, the vast majority of workers agrees about the value of the prevailing wage for a given task (Figure 1(b)). Similarly, there is a convention around work hours and in-kind payments that is commonly understood by workers and employers. In our experiment, the prevailing wage serves as a clear benchmark for wage job offers.

However, there is no explicit coordination among workers on wage norms. For example, in our sample, 89% of workers state that there is no meeting where workers in the village collectively discuss the wage, and 97% state that there is no meeting between workers and landowners to bargain over the wage for the season (Appendix Figure B.1). Consequently, to the extent that we see coordinated behavior among workers in our experiment, it would be consistent with the potential for norms to sustain cooperation among decentralized individuals.

2.2. **Motivation: Labor Supply Norms.** We hypothesize that, in our setting, there is a norm against accepting jobs below the prevailing wage. Figure 2 provides motivational evidence for this idea. Panel A shows that 84% of workers state that it is "unacceptable" or "very unacceptable" for an unemployed worker to offer to work below the prevailing wage.[9] Panel B shows that 90% of respondents believe that other workers would become angry at an individual who accepts a job at a rate below the prevailing wage—indicating the potential role of sanctions in deterring deviations. Only 8% of respondents believe that others would not care, and 2% believe others would be happy that the worker was able to find a job.[10]

In contrast, respondents express no disapprobation if a worker accepts a wage increase. 99% of respondents find it somewhat or highly acceptable for a worker to accept a job that pays a premium over the prevailing wage (Figure 2 Panel A). In addition, 95% of workers think that others would be happy for an individual who gets paid more than the prevailing wage (Panel B). These patterns are consistent with a norm against accepting wage cuts, rather than wage deviations in general. Consequently, to maximize statistical power, we focus our empirical tests on understanding labor supply behavior and social sanctions when work is offered at versus below the prevailing wage.

[9] These scenarios follow the approach developed by Kahneman et al. (1986) and adopted for this setting by Kaur (2019).

[10] The responses in Figure 2 are drawn from workers in non-study villages in study districts. Appendix Figure B.2 presents comparable evidence from an untreated "hold-out" sample in the study villages. The responses are quite similar across samples.

Note that these patterns are not inconsistent with fairness norms for pay equality. For example, Breza et al. (2018b) document that wage inequality is generally not tolerated among workers doing the same job in the same production unit, but do not compare themselves to those in other production units at the same worksite. Consistent with this, in our data, when respondents are asked, "Is it acceptable to take a job where you work directly alongside people paid a lower wage?" 92% indicate that the behavior is somewhat or highly unacceptable. However, such relative pay concerns are less salient when considering the wages of others in the decentralized daily spot contract setting of the labor market as a whole, where different workers are constantly re-matched with different employers.

Rather, workers seem to primarily view others' wage contracts in instrumental terms—with the belief that the wages accepted by some workers have the potential to exert pressure on the future wages of all workers in the same local labor market (Appendix Figure B.3). For example, in the scenario where a laborer works for a wage increase, 83% of respondents say they would use that knowledge to ask for higher wages themselves from future employers. In contrast, in the scenario where a worker accepts a wage below the prevailing wage, 60% of respondents say employers will use this to attempt to pay other workers lower wages in the future. Assessing the potential accuracy of such beliefs—and the wage determination process in general—is beyond the scope of our paper. However, from workers' perspectives, these beliefs help rationalize the presence of strong norms against accepting wage cuts in this setting.

3. Research Design

3.1. **Hypotheses.** We use the term "norm" to denote a specific desirable behavior, which is common knowledge among agents.[11] In our setting, we hypothesize that not accepting wages below the prevailing wage is a social norm held among workers in the local labor market, and violations of this norm result in sanctions. Consequently, while there exists underlying latent labor supply below the prevailing wage, it is suppressed.

Denote the prevailing village wage as w. We predict:

H1.) The true private opportunity cost of working (in the absence of the norm) is less than w for a subset of workers. This implies that some unemployed workers would be willing to accept work at wages below w.

[11]We adapt this definition from Kandori (1992), who defines a social norm as the "specification of desirable behaviour together with sanction rules in a community". Note that we do not take a stance on whether the norm is simply a decision rule, or internalized into utility as "moral" behavior. We provide a suggestive test for this distinction in Section 6 below.

H2.) Social pressure prevents workers from supplying labor below w. This implies that when other workers can observe an individual's job take-up decision, workers will be less likely to accept work below w.

H3.) Violations of the norm result in sanctions. That is, a mass of workers will sanction others who have accepted work below w.

Note that H2 distinguishes intrinsic altruism (i.e., workers themselves think it is not acceptable to take wage cuts) from external pressure (i.e., workers do not accept wage cuts due to concerns about others seeing them as doing so) (e.g. Bénabou and Tirole, 2006).

We next turn to our core labor supply experiment, which tests Hypotheses 1 and 2. This enables us to test our primary hypothesis: social considerations shape labor supply behavior in a decentralized market setting. In Section 6, we unpack the specific nature of these social considerations by more directly investigating Hypothesis 3.

3.2. **Experimental Design: Overview.** Our primary goal is to test whether social pressure meaningfully distorts aggregate labor supply by causing workers to withhold labor below the prevailing wage. To ascertain this, we seek to construct a design that varies the scope for social pressure, enabling us to measure the resultant effect on labor supply. Our design has five primary elements. We provide an overview of these elements in this section, with the specific experimental protocols detailed in Section 3.3.

First, we measure labor supply impacts by examining effects on job take-up decisions. For these decisions to be informative about actual labor supply behavior, the jobs must correspond to the actual employment opportunities that workers regularly face. This is especially because norms may be context-specific. Consequently, rather than offering jobs ourselves, we partner with existing employers, who make job offers to workers in their respective local labor markets. We conduct the experiment across 183 local labor markets (i.e. villages). In each village, we partner with an agricultural employer who makes job offers to workers in his respective village. Each of our 183 partnering employers regularly hires laborers in his village, and all offers are for actual, regular jobs on the employer's land.

Second, we randomize the wage at which jobs are offered. In some villages, jobs are offered at the prevailing wage, providing a benchmark for the level of labor supply at w. In other villages, jobs are offered at 10% below the prevailing wage ($w - 10\%$)—providing an estimate of the willingness to work under a wage cut, relative to the benchmark of supply at w.

Third, to vary the scope for social pressure to affect labor supply decisions, we randomize whether other community members are aware of the wage at which the job is offered. Specifically, we cross-randomize workers into one of three observability conditions (see Figure 3 for an overview of this cross-cutting design). In the *Public* condition, other workers and the hiring employer directly observe the offered wage, while in *Employer only*, only the hiring employer knows the wage. However, ex ante, the difference between these conditions may underestimate the effect of worker pressure: the employer is a member of the village and could potentially spread wage information to other workers.

In contrast, an ideal conceptual test of H1 and H2 requires a condition where the worker's labor supply decision is based solely on his true internal reservation wage, without social influence considerations. This motivates our *Private* condition, where only the worker himself knows the wage. We implement these observability conditions by having the employer extend the job offer for work on his land, but then having a local research staff member convey the actual wage rate across all experimental arms (see details below). This design enables us to measure how labor supply changes with the scope for social pressure—comparing *Private* with *Public*, as well as comparing *Employer only* with *Public* to obtain a lower bound for social pressure effects.

However, a natural concern about the observability treatments is that they may inadvertently change worker beliefs or preferences about the employment contract, through a mechanism other than social observability—confounding their interpretation. We use the prevailing wage job offers as a placebo test against this concern. Under our hypothesis, observability will only matter under norm violations (i.e. under wage cuts, $w - 10\%$); it should not matter when jobs are offered at the prevailing wage (i.e. w). We test this prediction in the data to assuage this concern. In addition, we directly assess various specific potential confounds in Sections 5.5 and 7.1 below.

Fourth, treatment is randomized at the village level, so that all workers within a given village receive the same wage $\times$ observability condition (see Figure 3). While village-level randomization increases the operational complexity of the experiment, for example by requiring us to recruit large numbers of partnering employers, it offers important benefits. Notably, it helps ensure the privacy of the *Private* condition—both by enabling us to make a small number of offers within any given local labor market, and to refrain from advertising *Public* offers in places where *Private* offers are being made.

Finally, we ensure labor supply decisions are made by a representative sample of agricultural workers in the local labor market. We consequently work with employers to randomly sample from the potential labor force when extending offers.[12] This maximizes the relevance of our

[12]Recall that since there is constant re-matching between workers and employers, the employer has some familiarity with most workers in the local labor market—so that being approached by the employer is not unusual from a random worker's perspective.

treatment effects for understanding distortions in the aggregate labor supply curve facing employers in the village.

3.3. **Implementation and Protocols.**

Process of Making Job Offers. We build on the usual procedure for making job offers in this setting: the employer approaches a worker at his home, and extends an offer for work. Because the employers live in the same village and regularly hire labor, they are usually known to the workers. The job offers are for familiar standard agricultural tasks, such as weeding or field preparation, which workers are used to performing. As is normal in this setting, there is no bargaining; employers make a take-it-or-leave-it offer. The hiring takes place two days before the work day—the modal gap between hiring and workday for agricultural jobs in this setting.

The hiring protocol follows the below four basic steps. Only step (ii) differs based on treatment status, as described below. All other steps are exactly the same across treatment arms. When making offers, the employer is accompanied by two enumerators (i.e. research staff members) from the local region.[13]

(i) The employer approaches the pre-selected worker at his home in the early evening, when laborers are typically home. He describes the job task and location for the work on the specified date. Because the standard job amenities—timing, breaks, and meals—are implicitly understood according to convention in the village for a given task, these are not explicitly discussed. The employer then introduces the enumerators as individuals "from a research institute" in the state capital who are studying agriculture and who would like to do a brief survey with the worker. The employer then tells the worker to let him know his work decision after speaking with the enumerator. This creates a natural hand-off from the employer to the enumerator.

(ii) Depending on the observability condition, this step occurs either out on the street in front of the worker's home, or inside the worker's home. In the *Public* condition, all individuals remain out in the street, where neighbors or passers-by can hear and observe them. Given how the villages in our study area are organized, these onlookers are almost always other residents of the labor colony, who typically engage in agricultural labor themselves. The enumerators recorded an average of 5 onlookers during *Public* job offers on average. Note that the *Public* condition most closely resembles the typical level of visibility for job offers in our context.

[13] Because it can be a long walk from the employer's house to the labor colony (e.g. 1 kilometer), it is not unusual for an employer to enter the labor colony with someone else, like a friend or relative, who can keep him company on the walk and while hiring.

In the *Employer only* condition, both the employer and an enumerator step inside the home or a private area adjacent to the home, while the other enumerator stays outside to ensure onlookers are not within hearing range. In the *Private* condition, only an enumerator enters the house or private area with the worker, while the employer and other enumerator remain outside. Since the worker is to take a "survey", this justifies needing to step to a private area in the latter two conditions, so that there is no reason for onlookers to suspect wage information is being conveyed.

Across treatment arms, the exchange with the enumerator only lasts 2-3 minutes. The enumerator introduces himself again and tells the worker the wage rate for the job (e.g. Rs. 200), according to the worker's assigned wage treatment condition. In addition, in the *Private* condition, he reiterates that only the worker will know his own wage, and that it will be kept confidential from the employer. Across treatment arms, he explains that because enumerators need to return to speak further with the employer at the end of the workday, the employer will deputize the research team to deliver the payment to the worker.[14] After this, to maintain the premise of a survey, the enumerator implements a short "survey" that lasts about 1 minute, with simple questions all workers can easily answer (e.g. "How many laborers are needed to undertake land-levelling (a common task) on one acre of land?").

(iii) All parties come together and the worker tells the employer if he accepts the job. As is common, this entire hiring process is short, with the entire exchange lasting about five minutes from start to end. The script used to make job offers is given in Appendix C.

(iv) On the day of work (2 days later), the worker reports to the employer's field for work if he accepted the job. An enumerator visits the field on the morning of the work day to spot check that the worker showed up and is working. However, we are in no way involved in the work itself. The employers supervises the workers as usual, give them tea or meals, etc, without our presence. At the end of the day, a research staff member returns and hands the worker an envelope containing his payment, so that no one else can observe the amount inside. He also schedules a time to revisit the worker at his house, where we do a follow-up survey. At this revisit, for the wage cut arms, we also give workers who accepted the job an unexpected compensation top-up, so that, for ethical reasons, even these workers can truthfully report having earned the prevailing wage.

Our design necessarily introduces some deviations from the natural hiring process, particularly in the *Private* treatment. In Section 7.1 below, we discuss potential resultant external

[14]This is consistent with the usual timing of payment for casual daily labor, which occurs in cash at the end of the workday. It is not unusual for an employer to deputize a son or someone else to run and make payment to the workers.

validity concerns and argue our design enables us to draw valid inferences on whether social observability distorts labor supply.

Village Sample. The 183 study villages in our sample frame were chosen to satisfy four main criteria, none of which are particularly restrictive for rural areas in our study setting.[15] First, village residents must be engaged in agricultural production. Second, among laborers, employment within the village was a large component of workers' labor earnings. These first two criteria essentially just exclude villages that were in close commuting distance to a town or large factory. Third, the labor colony population is between 30 and 100 households (the average labor colony in our sample has 46 households (Table 1)). The lower bound ensures that by making 2-3 job offers per village, there is not direct communication with a large fraction of the village. The upper bound on village size ensures a level of comparability across study villages. Finally, there needs to be a clear notion of the prevailing wage at the time we came to the village. Because we conduct the experiment during non-peak (lean) production periods, this criterion was almost always satisfied. This only disqualified a few villages at the very end of the lean season, where the wage was in flux at the start of the new agricultural year (just before the time of monsoon arrival).

Employer Sample and Recruitment. In each study village, we construct a listing of 20 employers who have regular demand for agricultural labor. We then recruit an employer from the middle of the land size distribution, who has work to be done on his land in the subsequent two weeks. We exclude employers who hold political positions, such as on the local panchayat (village government).

Employers are incentivized to participate with a lump sum payment and a wage subsidy that offsets the cost of hiring the workers. This both incentivizes participation, and also compensates employers in the event that they face complaints from workers for having tried to offer a below-prevailing wage job (see related discussion in Section 7.1).

When they are recruited, employers are informed that the wage rate may vary, and may possibly be set below the prevailing wage. However, they are not told their specific treatment status. More generally, the information provided to employers at the time of recruitment is exactly the same across treatment arms. We set the wage subsidy so that the employer's contribution to the wage is independent of treatment, permitting employers to be blind to the wage in the *Private* condition. In the *Public* and *Employer only* conditions, employers only learn the offered wage during the "survey" portion of the job offer (step (ii) above).

[15]Since we could only run field operations during the lean months in each given region where we worked, our sample size is limited by the number of villages we could cover by the end of the main lean season in 2017 (i.e. June 2017).

Worker Sample. At the time of the employer listing, enumerators also construct a listing of households in the labor colony. Using this list, we pre-randomize which households will be approached for the job offer and the order in which they are visited. During hiring, the job is offered to the male household head; if the household head is not present, then it is offered to another prime-age male household member. Door-locked households, where no men are home when the employer arrives to make job offers, are skipped. In order to minimize the total amount of information injected into any village, the employer extends a maximum of 2-3 job offers per village, with the number of offers pre-determined based on task. In total, job offers are made to 502 male laborers across all villages.

Because we do not pre-screen workers on any dimension, not all our sampled workers are in the agricultural labor force. While almost all workers (94%) in the labor colony report working in casual daily labor as their primary or secondary occupation, not all undertake agricultural labor. In our sample, 84% report agricultural labor as their primary or secondary occupation (Table 1). In addition, note that all workers in our sample are hired regularly by employers to work at the prevailing wage—and so are qualified for the jobs that they are offered. However, we take this broad approach to sampling to ensure we are drawing workers from the full potential labor force.

3.4. **Predictions.** Note that our design does not measure total labor supply at a given wage. Rather, it captures the *residual* labor supply facing employers—the fraction of workers who will take up a job at a given wage, conditional on not having found work elsewhere. Workers who have found (or expect to find) a prevailing wage job will refuse a lower wage one, even if the latter is substantially above their reservation wage. Consequently, take-up in our experiment provides a lower bound on underlying labor supply at any given wage.

Note that, as a result, our test only has power to detect positive labor supply below the prevailing wage in the presence of involuntary unemployment. If all workers could find jobs, then no would accept work at a $w - 10\%$, even if labor supply is positive at this level. This motivates our decision to conduct the experiment during the lean season, when unemployment is high.

We refer to Figure 3 to detail how our design maps to our key predictions. Our primary hypothesis is that social observability will reduce take-up of wage cuts. To test this, we predict that take-up of *Public* wage cuts (cell A) will be substantially lower than take-up of *Private* wage cuts (cell C). We therefore measure whether the difference in take-up between cells C vs. A is significantly different from zero.

As discussed in Section 3.2 above, take-up of wage cuts in the *Employer only* condition (cell B) likely provides a lower bound on the private willingness to work, since the employer could convey wage information to other workers. Consequently, we can also measure the difference in take-up of wage cuts in cells B vs. A to estimate a lower bound on the effects of social observability on labor supply.

In addition, we predict that social observability only matters under norm violations (i.e. under wage cuts), but should have no impact on take-up at the prevailing wage. Consequently, we predict that acceptance of jobs at the prevailing wage will be statistically indistinguishable across the three observability conditions. Moreover, we can also use a differences-in-differences regression specification to absorb any level-shifters across the observability treatments, comparing (C-A)-(F-D) in our test of H2, for example.

4. Data

4.1. **Data Sources.** Our primary outcome is whether the approached worker worked for the employer. This information comes directly from our spot checks verifying that workers actually worked in the employer's fields.

We supplement this with three survey-based data sources. First, the enumerators also conduct exit surveys after the scheduled work day with all workers who are offered jobs. The survey contains basic demographics, the worker's primary and secondary occupations, and a full time-use listing of all employment activities and compensation (both cash and in-kind) for the previous 10 days. Second, at this time, similar surveys are also administered to an untreated "holdout" sample of 886 workers—who were not approached for jobs—in each study village. These workers are also randomly selected from the labor force of potential workers, and are therefore comparable to the main study participants. Finally, the enumerators conduct exit surveys with the partnering employer to record information about any in-kind or cash transfers made to the workers, and the employer's perceived effort and work quality of each worker.

4.2. **Summary Statistics and Covariate Balance.**

Summary Statistics. We present summary statistics in Table 1. To test for balance across the treatment arms, we regress each covariate on treatment indicators and use a Wald test of joint significance of the regression coefficients.

Among workers who are offered the job at $w - 10\%$ in *Private*, 94% of participate in the casual daily labor market as their primary or secondary occupation, with 84% reporting casual labor in agriculture to be their primary or secondary occupation. 59% do not own land, highlighting the importance of the casual labor market for their livelihoods. Lean season employment rates are low: individuals report having any work on only 30% of days, and report being involuntarily unemployed on 43% of days in the prior month. 32% have directly worked for the partnering employer in the past; however, as discussed above, almost all workers are familiar with the employer since he is a member of their village.

We also report information about village-level labor market conditions using survey responses from the untreated holdout sample. These respondents are similar to the workers in our study. For example, the level of involuntary unemployment is comparable with the levels reported by the main study participants, at 43%. Employment through the government's workfare program, NREGS, was rare in our study villages at the time of our experiment (0.5% of the past ten days).

Balance. We achieve balance across treatments for almost all variables, with two exceptions: individual involuntary unemployment and whether the respondent engages in non-agricultural daily labor (Table 1, Col. 5). This accords with natural rates of sampling variation in a balance table. Regardless, we show in Appendix Table A.3 that our experimental results are very similar when including controls for these covariates. Moreover, we show in Sections 5.1 and 5.4 that our hypotheses hold even for the subset of workers who engage primarily in agricultural labor, and those with low involuntary unemployment—providing further reassurance that this sampling variation is not important for our results.

Table 1 also confirms that there were approximately five onlookers in each of the *Public* treatments. There are no detectable differences in the number of onlookers between the w and $w - 10\%$ *Public* treatment arms.

Attrition. Our main outcome—whether each worker worked for the employer—is directly observed by our enumerators for all workers. Consequently, we do not need to worry about differential attrition in our core take-up regressions. However, we do use responses to our exit surveys as covariates and in heterogeneity analyses in several places. Our survey completion rate for the *Private* $w - 10\%$ group was 88%, and response rates are balanced across treatment arms (Appendix Table A.1). We completed an average of 5.4 untreated holdout sample surveys in *Private* $w - 10\%$ villages. These survey rates are also balanced across village treatment status.

5. Evidence: Labor Supply

5.1. **Effects on Job Take-up.** Our primary tests of Hypotheses 1 and 2 are presented graphically in Figure 5. Panel A of Figure 5 includes the full experimental sample. Because we did not pre-screen workers in any way, this includes some respondents who do not participate regularly in the village agricultural labor market. Panel B restricts the sample to the 84% of workers who are actually in the agricultural labor force: those who report agriculture wage work as their primary or secondary occupation (henceforth, "agricultural workers"). Table 2 presents the results in regression form along with all of the relevant statistical tests.

Cols. (1)-(2) report OLS regression results for the main experimental sample, where *Private* $w - 10\%$ is the omitted category. Col. (3) restricts the sample to agricultural workers. Cols. (2) and (3) include year $\times$ month (i.e. strata) fixed effects and task fixed effects. All standard errors are clustered at the village-level. The dependent variable is job take-up, measured as an indicator for the laborer working on the day of work. Figure 5 reports the raw mean for the *Public* wage cut arm, and reports estimated means for the remaining treatment arms using the regression results from Cols. (2)-(3) of Table 2.

Panel A of the figure indicates that, in the private condition, 26% of workers accept the job at w. There are no statistically significant differences among the three observability conditions at w (Test 4, p=0.816). As expected, we find higher levels of job take-up at w among the sample of potential agricultural workers, who take-up jobs at w 33% of the time in *Private*. Given the baseline unemployment rate of 46% for agricultural laborers and the fact that this is a residual labor supply measure, this magnitude is reasonable. Again, we see no detectable differences by observability—agricultural workers take-up prevailing wage jobs 36% of the time in the *Employer only* arm and 37% of the time in the *Public* arm (Test 4, p=0.904). This mitigates the concern that the different observability conditions alter workers' desire to take up jobs through a channel other than social observability—for example, by changing beliefs about the job. Having verified our predicted insensitivity to social observability at w, we pool the three w treatments for power in subsequent analyses.

Next, we turn to H1, our hypothesis that workers are privately willing to supply labor at below the prevailing wage. Among all potential workers, Figure 5 shows that 18.4% are willing to accept a wage cut in *Private*, indicating robust underlying labor supply below the prevailing wage. Moreover, while these point estimates suggest a positive residual labor supply elasticity—with increased job take-up at higher wages—we cannot reject that *Private* take-up of jobs at w and $w - 10\%$ is the same (p=0.232). Similarly, among agricultural workers, the take-up of *Private* wage cuts is 26.4%.

However, consistent with H2, the take-up of wage cuts plummets when this choice is observable to other workers. Among all potential workers, only 4.8% accept $w - 10\%$ offers in *Public*. This is a 13.6 percentage point (pp), or 74%, drop in take-up relative to the *Private* wage cut arm (p=0.019). This drop is even starker for those in the agricultural labor force: only 1.8% of agricultural workers are willing to accept public wage cuts. This corresponds to a 24.6 pp, or 93%, reduction in labor supply relative to *Private* wage cuts for these workers (p <0.001).

We can also run a similar test of H2 using a differences-in-differences specification with the w treatments. This corresponds to estimating (C-A)-(F-D) using the notation of Figure

3. This allows us to net out any potential level-shifters between the *Private* and *Public* observability conditions. Unsurprisingly, the results are quite similar. For example, among agricultural workers, the effect of social observability on the take-up of wage cuts is 28.4 pp (Test 2, p=0.009).

As is apparent in Figure 5, take-up in the *Employer only* condition is statistically indistinguishable from that in the *Private* condition. Rather, acceptance of job offers only plummets under the *Public* condition. As discussed in Section 3.2, we can estimate a lower-bound of the effect of worker observability by comparing the *Employer only* condition with the *Public* condition. Among agricultural workers, this difference is 17.0 pp—corresponding to a stark 90% drop in take-up when from moving from *Employer only* to*Public* (Test 3, p=0.0107). These findings indicate that the presence of the employer itself doesn't drive our social observability results. Rather, consistent with H2, the decline in take-up is driven by the presence of other *workers*. Given these results, in many of the subsequent empirical exercises, we pool the $w - 10\%$ treatments across *Private* and *Employer only* for power. In Section 5.5 below, we present further discussion and evidence against an employer-bargaining interpretation of our results.

We provide several robustness checks. First, in Appendix Table A.2, we use randomization inference to check our main empirical conclusions. The results are very similar to those presented in Table 2.

Second, recall that we could only make job offers to households where a male worker was home at the time of our visit. One might therefore worry that while our surveyors were in the village, information about the treatments quickly spread, impacting whether or not we could find workers in the households on our randomized list. Appendix Table A.4 presents additional robustness to address this concern. In Cols. (1)-(2), we restrict our analysis to the first household and first two households where job offers were made. While results are noisier due to the reduced sample size, the results are qualitatively similar. In Col. (3) we code any "doorlock" household as having zero take-up and run an intent-to-treat regression. While this mechanically dampens the observed take-up levels across all treatments, and thus predictably decreases statistical power, the results again remain qualitatively similar.

Overall, the two observability conditions paint dramatically different pictures of the aggregate labor supply curve. In private, when workers make decisions without the fear of social observability, labor supply appears upward sloping, with substantial willingness to work below the prevailing wage—even under our lower bound measure. In contrast, when other workers would learn about the contract, there is scant take-up of wage cuts—so that labor supply appears almost perfectly elastic at the prevailing wage.

5.2. **Effects on Earnings.** We next quantify the implications of our treatments on worker earnings. While our labor supply effects are large in magnitude, there are two reasons why they may overstate how much workers are giving up by refusing job offers. First, those who turn down jobs in the experiment may have found alternate employment. Second, there may be inter-temporal substitution in labor supply across days; e.g., a worker who declines a public wage cut could be more likely to work on a later day that week. To quantify earnings effects—and therefore how consequential take-up decisions are—we conducted a daily recall of all employment activities in the week since the job offer was made as a part of the exit survey.

We present our findings in Table 3. Each observation is a worker-day. In Cols. (1) and (2), we restrict observations to the specific day for which work was offered in the experiment. The dependent variables in these columns are whether the individual worked for a wage in agriculture (for our partnering employer or any other employer) and total agricultural wage earnings, respectively. The omitted category in the regressions is being assigned to a non-*Public* (i.e. *Private* or *Employer only*) wage cut condition.[16] Participants offered a *Public* wage cut are 16.1 pp less likely to have any agricultural wage employment on that day than those offered a non-*Public* wage cut (Col. 1, p=0.002). This corresponds closely to the gap in take-up rate for the job from our "administrative" records. Accordingly, they earn Rs. 32.4 less that day on average (Col. 2, p=0.004). This magnitude corresponds to 71% of mean earnings that day in the *Private* $w - 10\%$ group. As a proportion of the mean earnings of the untreated holdout sample (not offered the job)—which represents the counterfactual for what workers would earn if they were not part of our experiment—the magnitude of earnings losses is even larger, representing 196% of a typical day's earnings.

Next, we investigate whether inter-temporal substitution mitigates earnings losses from the one-day offer. In Cols. (3)-(4), we find no evidence that those who forego working under the *Public* wage cut are more likely to work on other days in the week (specifically, the 1 day before and 5 days after the day of work).[17] We then test the total impact of the job offer on total weekly employment and earnings, including our experimental day of work (Cols. (5)-(6)).[18] Relative to the workers who are offered a wage cut in *Private*, those offered a wage cut in *Public* have a 6 pp lower employment rate over the course of the week (Col. 5, p=0.010). This translates into a reduction in estimated average daily earnings of Rs. 11.8

[16]In several places throughout our remaining results, we economize on power by pooling all the social observability treatments for wage w and pooling the *Employer only* and *Private* treatments for wage $w-10\%$.
[17]Note, this result is not surprising if workers are not on their labor supply curves—as indicated by the high reported levels of involuntary unemployment at baseline (e.g. Breza et al., 2021).
[18]Because our employment recall grid covers slightly different recall periods across workers, we weight the data in the full-week regressions to account for missing days.

(Col. 6, p=0.090). Given the results in Cols. (3)-(4), these declines in weekly earnings are driven by declining the job on the experimental day of work.

The coefficients in Cols. (5)-(6) translate to a 49% loss in average weekly wage earnings for the *Public* wage cut arm relative to the *Private* wage cut arm. In a more conservative calculation, in which we assume there is no decrease in the probability of employment on any given day other than our single day of work, those who are offered public wage cuts lose 26% of average weekly earnings, relative to those offered the same jobs in private. Relative to workers in the untreated holdout sample, these magnitudes correspond to 71% and 38% of weekly earnings, respectively. These findings indicate that workers are willing to give up a large fraction of weekly earnings in order to avoid being seen as breaking the social norm.

5.3. **Heterogeneity: Information Flow.** Under our hypothesized mechanism, workers rarely accept *Public* wage cuts because they do not want to incur sanctions from other workers. Prior work highlights that such reputational implications depend on the extent to which information diffuses through the network (Alatas et al., 2016). This implies that our treatment effects should be larger when the *Public* condition is more likely to result in other community members learning about norm violations. We construct a suggestive test for this prediction by examining heterogeneity by the level of information diffusiveness in villages. For example, in more diffusive villages, information may be more likely to spread from onlookers to other community members.

We asked two questions about village diffusiveness in our untreated holdout sample of workers: i) the extent to which workers learn about the wages at which others accept agricultural work and ii) how many others would find out if a worker accepted an agricultural job at below the prevailing wage.[19] We aggregate responses to each of these questions at the village level and create two indicators for whether a village has below-median information flow. We predict that the magnitude of the treatment effect will be smaller in villages when information is less likely to diffuse to other community members.

In Col. (1) of Table 4, we examine heterogeneity by the diffusiveness of general wage information (measure i). In highly diffusive villages, take-up under *Public* wage cuts is 20 pp lower than take-up of non-*Public* wage cuts (p=0.003). However, in low diffusiveness villages, this large treatment effect is almost completely offset—the interaction term is 17.0 pp (p=0.070)—so that we cannot reject that there is no difference in overall take-up of wage cuts across the *Public* and non-*Public* conditions. The same pattern holds in Col. (2), where

[19]Measuring network structure in each study village was beyond the scope of our field activities.

we instead use the second measure of diffusiveness that more directly pertains to norm violations. In Col. (3), we examine heterogeneity among potential agricultural workers only, and find an even stronger (albeit similar) pattern of results. Finally, take-up rates under private wage cuts are similar across high and low diffusiveness villages: the estimates on the level effect of "Low info spread village" are close to zero in most columns. However, these estimates are imprecise, and we cannot rule out the potential for meaningful differences in take-up across such villages.

The findings are consistent with larger effects of social observability when more individuals are likely to learn about norm violations. This echoes prior results in the social networks literature: more extensive information flow has a disciplining effect on agents (Breza and Chandrasekhar, 2019), improving incentives for social cooperation (Bloch et al., 2008). However, there are two important caveats to these results. First, in our context, the level of diffusiveness could also be correlated with the network's ability to sanction—for example, if the kinds of villages where information spreads faster are also those where individuals are more reliant on each other for risk sharing. This latter case is not a problem for our hypothesized mechanism, since it would simply reflect a related way in which social observability creates stronger incentives for cooperation. Second, variation in diffusiveness may be correlated with other characteristics that could affect labor supply behavior. Consequently, we interpret this as correlational (not causal) evidence, offering suggestive support for our proposed mechanism.

5.4. **Heterogeneity: Involuntary Unemployment.** Figure 6 presents the results on take-up differentially by underlying involuntary unemployment rates. The village's involuntary unemployment rate is defined as the average proportion of the past ten days that workers in the untreated holdout sample report they would have preferred a job at the prevailing wage—a measure of labor market slack. We examine effects using a median cut of this variable. Appendix Table A.5, Cols. (1)-(2) present the regressions associated with Panels A and B of the figure, respectively.

Because take-up corresponds to residual labor supply, we would expect acceptance of wage cuts to be higher when there is more involuntary unemployment; otherwise, individuals would expect to obtain prevailing wage jobs, and would turn down wage cuts. Consistent with this, take-up of non-*Public* wage cuts is higher when the village has a higher unemployment rate: mean take-up in Figure 6 is 26.1 pp when involuntary unemployment is above-median, and is 16.6 pp lower when they are below-median (difference p=0.0187). This difference is starker among agricultural workers, with a 23.1 pp lower take-up rate in low involuntary unemployment areas (p-value=0.009), relative to a mean of 37.2 pp in high involuntary unemployment areas.

How our treatment effects—the difference in take-up of *Public* vs. non-*Public* wage cuts—vary with unemployment levels is conceptually more ambiguous, with potentially opposing forces. On one hand, because higher unemployment may increase the desperation for work, laborers may be more willing to incur sanctions by accepting *Public* wage cuts—dampening treatment effects. On the other hand, if the places with higher unemployment are precisely those where stronger collective action leads to more rigid wage floors, then the willingness to take-up *Public* wage cuts may be lower in such places—leading to larger treatment effects.

Overall, we find that treatment effects are higher when unemployment is higher. Specifically, in the full sample, in high unemployment villages, take-up wage cuts is 19.6 pp lower when they are made in *Public* (p-value=0.0035). In contrast, in villages where unemployment is low, this treatment effect of social observability is 16.0 pp smaller (p-value=0.0808). The patterns are even more pronounced in the subsample of agricultural workers. In high unemployment villages, take-up of wage cuts in *Public* is a stark zero.[20]

Our findings indicate the effects of social observability are higher in areas with higher involuntary unemployment, and that workers in areas with higher unemployment are *less* likely to take up *Public* wage cuts. Speculatively, these patterns could be consistent with stronger wage norms leading to more distorted labor markets. We return to this idea in Section 8.

5.5. **Alternative Explanation: Employer Bargaining.** An alternate mechanism that might generate predictions similar to H1 and H2 is employer bargaining. Under this theory, workers may have heterogeneous reservation wages, such that those with reservation wage below w may prefer to hide their private value from employers in order to avoid being paid lower wages in the future.

Note that this is distinct from a *collective* bargaining interpretation of our results—where the wage norm exists for workers to collectively obtain market power with employers. While workers state that explicit collective bargaining does not occur (see Section 2.1), an implicit version of such a group motive is consistent with our hypothesized mechanism. Consequently, in this section, we focus on the confound of workers wanting to shield their own individual reservation wage due to a private bargaining motive.

[20]We also test separately for heterogeneous treatments effects using individual- rather than village-level involuntary unemployment in Appendix Table A.5, Cols. (3)-(4). Individual-level unemployment is constructed using the percentage of days a worker reports being involuntarily unemployed in the last 30 days. High involuntary unemployment is an indicator for an above-median value of this measure. We show robustness of the village unemployment heterogeneity results controlling for individual-level unemployment in Cols. (5)-(6). We find evidence of similar patterns in both cases.

The fact that the presence of the employer does not significantly reduce labor supply at $w-10\%$—especially among agricultural workers, who actually participate in this labor market—is already seemingly at odds with the employer bargaining mechanism. Recall from Figure 5 that, among agricultural workers, take-up of wage cuts in *Private*, *Employer only*, and *Public* is 26.4%, 18.8%, and 1.8%, respectively. Consequently, take-up of wage cuts only plummets when other *workers* can observe the wage offer.

Consequently, for this confound to be consistent with our results, one would need to believe that a worker's drastic decline in take-up is driven by a concern that other workers would inform other employers of his reservation wage; moreover, the scope for this information transmission to other employers would need to be larger from other workers than from the hiring employer himself. In this context, this seems unlikely for two reasons. First, the hiring employer is more likely to actively communicate with other employers than workers are—both by virtue of where they live and due to higher communication within individuals of the same caste. Second, as discussed in Section 2.2, workers believe that spreading information about acceptance of wage cuts could reduce wages for all workers in the labor market (Appendix Figure B.3, Panel A)—so that other workers would not have a dynamic incentive to spread such information to employers. Finally, it is common knowledge among workers and employers that workers are desperate for work during lean months; it is reasonable to assume that employers already know that many workers' true reservation wage is likely below the prevailing wage.

Moreover, this channel also does not accord with workers' own stated views. We asked a separate sample of 110 untreated workers, in a separate set of villages, to state the most important reason why a worker might refuse to accept a wage cut (Appendix Figure B.4). Not a single one of the respondents stated a fear that the hiring employer or other employers would try to pay less in the future as the primary reason.

Finally, to further investigate the employer bargaining mechanism, we use a set of supplementary analyses. Specifically, we identify two sets of situations where revealing one's reservation wage may be particularly costly to future employment prospects.

First, if workers don't want future employers to know their reservation wage, then they should be more likely to decline wage cuts when the hiring employer is more likely to employ them in the future. We proxy for this by whether the worker has worked for the employer in the past, since this pre-existing relationship positively predicts future employment. However, the pattern of our results is extremely similar regardless of whether the worker is more likely to work for the employer again (Figure 7, Panel A).[21] In both cases, the level of take-up

[21] We report the corresponding regression estimates in Appendix Table A.6. Because the employer bargaining mechanism operates through the value of future employment, we focus our analysis on agricultural workers.

of wage cuts in *Employer only* is similar in magnitude. If the employer bargaining story were true, then we should expect large drops in job take-up rates moving from *Private* to *Employer only* when the ongoing employment relationship is more valuable. However, we see no evidence of this in the data; in fact, the point estimate is small and positive rather than negative, though insignificant (difference=0.0016, p=0.991). Moreover, even among these employers for whom the worker is likely to work again, we see a 15.6 pp (89%) decline in take-up when moving from the *Employer only* wage cut to *Public* wage cut condition (p=0.0725). Consequently, even for such employers, take-up falls sharply in the presence of other *workers*.

Second, we examine heterogeneity by whether the worker views the employer as influential in the village. More influential employers can be expected to be more prominent and socially well-connected in the village—so that they would be in a better position to spread information to other employers. However, as before, we see no evidence that workers are less likely to accept wage cuts in *Employer only* when the employer is more influential; the point estimate actually goes in the opposite direction but is insignificant (Figure 7, Panel B). Overall, the patterns are very similar to those in Panel A.[22]

Consequently, we argue that employer bargaining is unlikely to be a primary driver of our results. Furthermore, it is unclear why such a mechanism would trigger social sanctions from other workers if agents only worry about revealing their own reservation wages.

6. Evidence: Social Sanctions

The evidence from our core field experiment supports the idea that social considerations alter the aggregate labor supply curve in villages. This is a novel finding that advances the literatures on market power and labor markets in developing countries—irrespective of the specific social considerations that motivate individuals, from sanctions to shame. However, in this section, we investigate these social repercussions.

In our main field experiment, it is not possible to directly measure social repercussions such as punishment, precisely because such outcomes should be off the equilibrium path. Under our hypothesis, the threat of sanctions alone will deter at least some laborers from working for a wage cut, and those who choose to accept the wage cut may be differentially susceptible to community sanctions. Consistent with this, in *Public*, when other workers can observe their behavior, agricultural workers almost never accept wage cuts, with only a 1.8% acceptance

[22]Specifically, counter to an employer barganing story, relative to when they are less influential, more influential employers see less of a decline in the take-up of wage cuts when moving from *Private* to *Employer only* (difference=0.0566, p=0.696), Appendix Table A.6, Col. 2). Moreover, among the subset of more influential employers, there remains a suggestive large drop in take-up when moving from *Employer only* to *Public* (16.4 pp or 89%, p=0.105).

rate. Moreover, recall that across treatment arms, workers who did accept the wage cut were ultimately paid the prevailing wage, mitigating the potential for sanctions ex post.

Consequently, we proceed in two steps. First, we present survey evidence from workers about why they would not accept wage cuts and the consequences of doing so. Second, to examine hypothesis H3, we build on our field experiment to construct a revealed-preference test for sanctions using a costly punishment game. Overall, our main goal is to test the prediction that norm violations result in sanctions. This provides positive evidence for the relevance of social disapprobation, but is not intended to exclude the potential relevance of other social considerations that could also affect labor supply behavior.

6.1. **Survey Evidence.** Workers who privately accept work below the prevailing wage may refuse public wage cuts due to several types of distinct social considerations. They may face sanctions from others, such as social or economic repercussions. Alternately, they may to avoid feelings of shame from others perceiving them as financially desperate.

We begin by asking workers their views on why we observe so few public norm violations in the data. In surveys with 110 workers who were not in our field experiment sample, we asked respondents to provide us with the most important reason an individual would reject a job below the prevailing wage (Appendix Figure B.4).[23] 75% of workers select a reason consistent with not wanting to violate the norm: 30% worry that others would get angry and 45% select not wanting to go against the collective group of workers.[24] The remaining 25% of workers state that the wage is insufficient to live (i.e. below their reservation wage). Shame, feelings of insufficiency, or changes in the wage offered by future employers are never chosen by any worker. However, because these are the primary reason, this does not rule out sanctions or these various factors as additional motives for other workers.

To ascertain the consequences workers might face for accepting work below the prevailing wage, we asked the untreated holdout sample of workers: "Suppose a worker 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. As we document in Figure 2 above, 90% of respondents say that others would be angry with him—indicating a high level of basic disapprobation.

[23]We first presented individuals with a vignette about a hypothetical worker and asked whether that worker would accept a job paying below the prevailing wage. For the vast majority of individuals who suggested that the laborer would reject the offer, we further asked them to elaborate on the most important reason.

[24]This latter motive may reflect the social repercussions of going against the group, or may also reflect an internalized notion of morality. However, note that the latter interpretation should apply equally to the *Private* and *Public* wage cut conditions—depressing take-up across all wage cut arms.

In Figure 4, Panel A, we augment this by detailing the other repercussions enumerated by workers. 59% of respondents stated 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 group of workers for an outside job. 56% of respondents said that a worker 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 agreed with the notion that accepting wage cuts results in monetary punishments—for example, refusal to help a worker with a financial emergency.

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

Thus, workers believe that breaking the wage norm is likely to result in costly social sanctions. These sanctions have both social costs (e.g. receiving anger from those who are central to one's social life) and also important economic costs (e.g. losing future job opportunities). Further, the community is well-positioned to take action to prevent non-compliant behavior.

6.2. **Costly Punishment Experiment.** Next, we build on our field experiment to directly test whether workers will sanction those who accept wage cuts. Even though in practice many of the sanctions that are described in the surveys above may not actually be costly to those who levy them, we design a game in which sanctioning is costly, in order to provide a stronger test for social disapprobation.

We conduct costly punishment experiments in a different, but comparable set of 13 villages from our field experiment sample. We again partner with employers to make job offers to a random subset of workers. 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 approach 6-8 workers with job offers in each village (decided *ex ante*). This larger number of offers guarantees that in each village, at least some workers have accepted a wage cut. Note that, as before, these are all offers for real jobs on employers' land. This sets up the backdrop for the costly punishment exercise.

We then recruit another (random) subset of 8-12 workers in each village who were *not* offered jobs. These other workers, 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. We tell the player that he and his anonymous partner each have an endowment of Rs. 100. The player can "punish" his partner at a cost to himself. Specifically, for every Rs. 5 that the player removes from his 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 the 1:5 proportion, until he is satisfied with the final allocation.

We randomize two features of the partner's characteristics. First, we vary whether the partner is a worker in the player's own village or in a village that is geographically far away but has similar characteristics as the worker's village (e.g. crop type, work composition, and wage level). Note that in this latter case, the worker's job acceptance decision has no direct consequences for the player, since the worker'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. Note that we never use the term "prevailing wage" but rather simply state the wage level at which the worker accepted a job (e.g. Rs. 200). The sample is weighted so that there is an equal number of observations in each of the 2x2=4 cells.

Furthermore, to obfuscate the reason for the exercise, we add in two "placebo" rounds of the game, in which the partner undertakes a positive, negative, or neutral action in a non-labor domain. The player's payoff is determined by a random die roll, selecting one of the four rounds.[25] These "placebo" exercises and the fact that the enumerators simply state the wage levels at which the partner accepted a job help assuage concerns of experimenter demand effects. If individuals punish their partnered workers, then it can only be because they draw on beliefs formed outside our experiment to decide that working at some wage levels deserves punishment and not others.

[25]The player is told that in each round, he is paired with a different potential partner. Whether all 4 partners are in his village or a different village is fixed across rounds. The additional rounds are played by the player before the wage acceptance scenarios. The positive, negative, or neutral actions are: giving a gift of a bag of grain, an employer withholding wage payments to workers for two months, and traveling to the city for work, respectively. Participant responses to these scenarios are presented in Appendix Figure B.5. The costly punishment game is played in the evening after job offers are made, but before the day of employment occurs. After all players in a village have finished the game, we announce that those workers 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.

If accepting a wage cut violates the social norm, then individuals may be willing to demonstrate their social disapprobation by destroying their own surplus in order 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 employment 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 give up money to 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 partners from distant villages in similar frequencies—even though that partner's action has no scope to affect the player's own labor market.

Table 5 presents these results in regression form. The omitted category is a dummy for whether the partner accepted a job at the prevailing wage. Col. (1) shows that, on average, the punishment probability increases by 42 pp when the partner accepts a wage lower than the prevailing level (p <0.001). 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 from the player's own. 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. 14.6 lower (on a base of Rs. 100).

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; Fehr and Schurtenberger, 2018). Our findings are also potentially consistent with contagious punishment models (Kandori, 1992; Ellison, 1994), in which norms are an equilibrium strategy that is enforced through decentralized sanctions. However, the willingness to punish those in other labor markets—where there is no scope for equilibrium effects on one's own payoffs—suggests that workers view norm violations in moral or general terms, consistent with some internalization of the norm into preferences. As MacLeod (2007) argues, even a mild amount of such internalization greatly expands the range of parameter values over which norm behavior can be sustained in equilibrium. However, note that a global punishment strategy is also potentially consistent with traditional contagious punishment models.

In addition, note that while a deviator from the norm may only receive a small punishment from any specific individual, the total punishment at the village level may be much more substantial (when aggregated across individuals). For example, if any worker who violates the norm faces sanctions from 30% of his community members, the total impact is likely to be meaningful. Moreover, because many of the sanctions described in the survey evidence may not be costly for co-villagers to impose, the willingness to pay to punish that we measure in the costly punishment exercise is likely a lower bound on the scope for social punishment.

Finally, note that usage of costly punishment indicates that not all of the social observability effects are driven by shame. For example, if accepting a wage cut signals financial desperation, and this is the only reason we see declines in *Public* take-up in our field experiment, then it is unclear why we should also see financial punishments of those who reveal themselves to be economically vulnerable in our costly punishment game.

7. Threats to Validity and Interpretation

7.1. **Threats to Validity.** We integrate the discussion of several key potential confounds into earlier sections of the paper. We discuss employer bargaining in Section 5.5 above. In Section 2.2, we offer evidence that the norm is specifically against accepting wage *cuts*, rather than any wage deviations (i.e. wage increases). In addition, in Section 6, we discuss the possibility that the social observability effects may be driven by shame rather than the threat of sanctions. In this section, we discuss other potential threats to the interpretation of our findings.

Unravelling through Side Deals outside Experiment. Given that we find robust take-up of wage cuts under our *Employer only* condition, this raises a natural question of why employers don't undermine the village equilibrium by cutting private side deals with unemployed workers. While this is not a threat to internal validity, this relates to the external validity of our findings. In interviews, workers and employers state that such side deals do occur from time to time, but cite two reasons why they remain uncommon.

First, employers state that if they develop a reputation for repeatedly offering low wage jobs, this would result in sanctions from workers.In Appendix Figure B.6 we present evidence from employer surveys in Panel A and worker surveys in Panel B. 74% of employers and 60% of workers indicate that other workers will become angry with the employer. Moreover, 86% of employers and 95% of workers state that workers may withhold their labor from offending employers during periods of peak labor demand (e.g., during rice transplanting and harvest); this would be extremely costly for employers. Finally, 73% of employers indicate that they would have to hire workers from outside the village, which requires paying a higher wage rate. Note that such potential punishments accord with the view in the labor literature that one of the functions of unions is to use their collective power to enforce work conditions and prevent transgressions by the employer against individual members.

Second, and more speculatively, workers state that while some workers may occasionally take the risk of accepting a wage cut, they would not do so regularly because a reputation for being a norm violator may reach other workers. This stems from the fear that employers may advertise acceptance of wage cuts to other workers in an attempt to elicit supply at low wages from others.

The discussion above raises one additional question: why would employers be willing to participate in our experiment if they also faced potential sanctions? First, employers were compensated for participating in our study with a hiring subsidy and lump sum payment. Second, the employers were accompanied by an enumerator who was understood to be from outside of the village. This potentially made it easier for employers to signal to workers that this was a one-off occurrence, enabling them to avoid strong reputational effects.

Side Payments, Job Amenities, and Worker Ability. Our design is constructed under the premise that working with the partnering employer at a wage of $w - 10\%$ is indeed less valuable than working at w. Using our exit surveys, we can check whether employers tried to compensate workers for the low offer wage by altering other amenities of the job: cash side-payments, shorter hours, in-kind payments, or by allowing workers to reduce effort. The results are presented in Appendix Table A.7.

Overall, there was high reported adherence to the intended wage: (the wage reported by the worker in the exit survey) − (the treatment wage level). Across treatments, under 2% of workers report a wage different than what was offered at the time of hiring. Moreover, comparing the w and $w-10\%$ treatments, there is no systematic difference in wage deviations (Col. 1), length of workday (Col. 2), or meals provided (Col. 3). The coefficients in Cols. (1)-(3) are typically small and statistically indistinguishable from zero.

In addition, we see little evidence that worker effort—as assessed by the hiring employer—was lower among those who worked under wage cuts (Col. 5). This remains true if we restrict the sample of $w-10\%$ treatments to only the *Private* condition, where the employer made his assessment not knowing the wage rate of the workers (Col. 6). The estimated coefficients in Cols. (5)-(6) are actually positive—counter to a compensating differentials or morale effects story—though extremely small and magnitude and insignificant. Moreover, to explain our take-up results, employers would need to be substantially more likely to increase job perks under *Private* wage cuts than *Public* ones, but no more likely to do so across observability conditions in the prevailing wage treatments, and workers would need to know to expect this pattern ex ante; it is unclear why this should be the case.

Finally, we also do not see evidence that the workers who select into jobs under wage cuts are systematically different from those who select into jobs under the prevailing wage. They are not more likely to have worked for the hiring employer in the past (Col. 4). In addition, the results in Cols. (5)-(6) indicate they are also not likely to be lower quality workers.[26]

Outside intervention in hiring process. Our experimental protocols necessarily require us to intervene in the hiring process by accompanying the employer and handing the worker the payment. As we discuss In Section 3.3, it is not unusual in this setting for an employer to bring someone with him while hiring, or to deputize someone he knows to make a payment to workers. However, our experimental protocols still introduce novelty, particularly in the *Private* condition, where we keep the wage private from the employer—raising potential external validity concerns.

We argue there are a few reasons why such concerns do not undermine the interpretation of our experiment. First, our core predictions hold if we simply compare the *Employer only* and *Public* treatments, both of which are reasonably similar to typical hiring processes. Recall that the estimated decline in take-up between these conditions is 8.4 pp (63%) for all potential workers (p=0.0865), and 17.0 pp (90%) for the workers that are actually in the agricultural labor force (p=0.0107). Second, the main goal of our experiment is to isolate whether social pressure reduces labor supply at wage cuts. Even if our hiring process introduced a level shifter (increasing or decreasing labor supply overall), this would not undermine the validity of the comparative static between *Private* and *Public* wage cuts. The fact that we see no differences in take-up across observability conditions among the prevailing wage offers validates the assumption that any labor supply effect from our involvement should be a level shifter; if it interacted differentially with the observability conditions, then these differences should show up in the w wage offers as well. Third, recall that the entire hiring interaction—from when the employer first approaches the worker to when the worker conveys his decision—lasts about 5 minutes only. We are in no way involved in the substantive employment relationship itself, which proceeds as usual on the employer's land for an entire day. Moreover, given the high involuntary unemployment in lean months, the untreated

[26] In addition, given that the employer sought out each worker in the experimental sample, it is possible that the offers of employment decreased worker search costs. However, this should be true across all wage offers and social observability treatments.

hold-out sample only finds wage work on 30% of days (i.e. 9 days) per month. Consequently, a worker who refuses a day of work at $w-10\%$ is giving up 10% of his typical *monthly* wage earnings. These high stakes help with external validity by ensuring that the decision for workers is extremely consequential.

Lack of privacy of Private wage offers. What if workers worried that the private wage offers would not remain private? Recall that we capped the number of offers at three per village. In addition, the fact that we observe robust take-up in the *Private* wage cut treatment (opposed to near-zero in *Public* $w-10\%$) validates our premise that at least some workers believed that confidentiality would be maintained in the *Private* condition. To the extent that workers did not believe the assurance of confidentiality, our take-up estimates are a lower bound.

Information conveyed through Public wage offers. The public treatments may have provided workers with information about the prevailing wage—e.g., through potential comments from onlookers. This information, in turn, could have depressed take-up of public jobs at $w-10\%$. However, we have already shown that the prevailing wage is generally known in the study villages (Figure 1(b)). Moreover, we document that among workers who were approached for job offers, reports of the prevailing wage are not systematically different across treatment cells (Appendix Table A.8).

7.2. **Is the Wage Floor In the Workers' Best Interest?** In this section, we discuss whether, as a group, workers might benefit from the existence of the wage norm. If employers exert monopsony power, then coordination among workers leads to an increase in worker surplus and an increase in efficiency in many cases (see Online Appendix C and Figure D.3 for further discussion).

To analyze the alternate scenario, where employers do not behave monopsonistically, we conduct a simple back-of-the-envelope exercise. Specifically, we estimate the counterfactual levels of wages and employment in the absence of the wage floor, and compute the change in worker surplus and deadweight loss from moving to the wage floor equilibrium. Any such exercise requires a number of strong assumptions. We consider a static, 1-sector environment, where following Lee and Saez (2012), the workers with the highest reservation wages are rationed first under the wage floor. Online Appendix C provides further details. In this highly stylized world, we find that the counterfactual equilibrium wage in the absence of distortions is 7% lower than the observed wage, and employment is 7% higher.

Under the above (strong) assumptions, workers benefit from the wage floor: workers' surplus is 64% higher relative to the equilibrium with no wage floor. We should also note that 96% of the gains to the workers come at the expense of employer surplus and only 4% from deadweight loss. Our calculations, albeit speculative, indicate that the ability to set a wage floor helps workers extract more total surplus for themselves.[27] Moreover, because our setting is a daily labor market—with different workers obtaining jobs on different days—these gains are likely, at least to some extent, to be shared across workers.

8. Correlation with Wage Rigidity

While wage floors may benefit workers on average by helping them claim more surplus from employers, norms against accepting wage cuts could nonetheless exacerbate the employment effects of transitory fluctuations. 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 use data collected by Kaur (2019) in a different set of Indian villages that span 6 districts in two states.

Appendix Figure B.7 shows that workers recognize that if a laborer agrees to work below the prevailing wage, he likely increases his own individual chances of employment (Question 1). 84% of workers also believe that others would be 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). Consequently, workers appear to believe that a sufficient number of deviations from the social norm could undermine the wage floor in the village.

Our experimental results indicate that social pressure, and specifically the threat of social sanctions, helps prevent 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, jobs), and information will flow less well through the network (making it harder to learn about and enforce deviations) (e.g. Jackson et al., 2012; Breza et al., 2018a). Consequently, in areas with less social cohesion, wage floors may be more difficult to sustain.

We use observational data from across India to test one potential implication of this idea: whether stronger social cohesion leads to more downward wage rigidity.[28] In India, caste is a

[27] When we relax the assumption of efficient rationing, we still find that worker surplus increases up to 72% when labor demand is relatively inelastic ($\eta < 2$).

[28] In order to maximize their surplus, workers could distort wages above market clearing levels, but the wages themselves could still adjust to shocks. However, as we describe above, the norm that operationalizes union behavior is that individuals should not accept jobs below the prevailing wage (e.g. Figure 2). This rule effectively creates downward wage rigidity. This highlights a way in which norms are a crude technology, which must be operationalized through simple decision rules that are observable. In contrast, formalized groups can write state contingent rules for their members, enabling more optimal outcomes. This point has been made, for example, in Elinor Ostrom's work on key ingredients through which informally organized groups provide public goods (e.g. Ostrom et al., 1992).

strong proxy for in-group and social cohesion.[29] In our experiment, the caste composition of workers is extremely cohesive: all agricultural workers in a given village belong to the same caste (Scheduled Caste or Scheduled Tribe), and the median number of subcastes in a village is 1. This high level of social cohesion likely helps 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.[30]

We use the wage rigidity test developed by Kaur (2019), which measures how wages and employment respond to transitory labor demand shocks (generated exogenously by rainfall). The core result of Kaur (2019) is that lagged positive shocks generate ratcheting in the labor market. Specifically, wages adjust upward in response to positive rainfall shocks, but in the following year, when the positive shock has dissipated, they do not adjust back down. 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.[31]

Panel A, Col. (1) of Table 6 replicates the core result for wages from Kaur (2019):

$$y_{idt} = \alpha_0 + \alpha_1 Pos_{dt} + \alpha_2 Pos_{d,t-1} NonPos_{dt} + \delta_d + \rho_t + \varepsilon_{idt},$$

where y_{idt} is worker i's log average nominal daily wage in district d in year t, Pos_{dt} is a binary indicator for having a positive shock (rainfall above the 80th percentile of the district's usual rainfall distribution), and $Pos_{d,t-1}NonPos_{dt}$ is a binary indicator for having a positive shock in the previous year and no positive shock this year. The omitted category in the regression is no positive shock this year or last year.

[29]Munshi and Rosenzweig (2006), Mazzocco and Saini (2012), and Rosenzweig and Stark (1989) link caste networks to investments in human capital, job choice, marriage markets, and risk sharing. Munshi and Rosenzweig (2016) argue that caste groups provide high levels of social insurance largely because of their ability to self-monitor and enforce collective punishments.

[30]The NSS measures four caste categories: Scheduled Caste, Scheduled Tribe, Other Backward Caste, and General Caste. Prior work has shown that subcaste in particular is a dominant indicator (Mazzocco and Saini, 2012). Since the NSS only captures these caste categories, our analysis relies on the assumption that members of the same caste show stronger in-group cohesion than those of different castes.

[31]The 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, here, we focus on lagged positive shocks as it is the core test to look at both wage and employment effects. In addition, pooling to look at current and lagged positive shocks increases power in the analysis relative to the 6-cell specification in the original paper.

Relative to no shock, a positive shock this year increase wages by 6.3% (α_1). 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 (α_2).

We examine whether these wage rigidity effects are mediated by 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 effect, 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 (i.e. low social cohesion). In areas with high social cohesion, there is strong wage rigidity: lagged positive shocks lead to a 10% increase (p-value=0.015) in current wages. However, in areas with low cohesion, the interaction term of -8.26% (p-value=0.100) offsets the level effect, and we cannot reject that lagged shocks have no predictive power for future wages. In contrast, we do not see a strong interaction of social cohesion with 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.[32] 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 wage rigidity. Col. (1) replicates the basic employment test; in the above specification, y_{idt} is the number of days in the past week that worker i did any work in agriculture. Employment rises in response to current positive shocks (α_1). 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 (α_2).

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 or 13%. However, in areas with low cohesion, we cannot

[32] More generally, under our hypothesis, we would expect the interaction with ”Positive shock this year” to be either zero, or possibly positive (if employers are more willing to adjust wages upward fully when there is no implicit union, since they know wages will be able to fall again in the future without having to worry about the ratcheting effect).

reject that there is no employment effect of lagged shocks: the interaction term of 0.189 (p-value of 0.032) 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.

Appendix Figure B.8 shows the underlying distribution of each caste Herfindahl index across Indian districts. Appendix Table A.9 shows robustness of these results to alternate specifications, such as the linear Herfindahl and linear in ranks in the Herfindahl.

These 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 business cycle volatility from rainfall. While only suggestive, the results in Table 6 are consistent with the view that social cohesion, and its resultant ability to sustain stronger social norms, could have aggregate labor market implications.

9. Conclusion: Relevance of Social Norms for Market Power in Other Contexts

We conclude with a discussion of the potential generality of our findings for understanding market power in other contexts. Our proposed mechanism hinges on two features: a clearly defined norm (with observable violations) and a mechanism to impose social sanctions (in order to prevent individual deviations). These features can arise naturally in many settings with repeated interpersonal interaction.

We present survey evidence from five additional types of markets in India and Kenya where these features are present. Specifically, we surveyed workers at urban labor stands in Bangalore, India and Busia, Kenya. In addition, we surveyed taxi (boda-boda) drivers, food (roti) vendors, and butchers in Busia.[33] Note that, unlike our experimental context, most of these markets are located in a city or town. For each example, we survey approximately 30 sellers, typically sampled from 3-6 distinct local market locations for each market type. In each of these markets, participants come to the same location to sell their labor or wares and interact with the same group of workers or vendors regularly (in most cases, daily). Moreover, there tends to be a uniform price for each unit of work or goods within a given local market (e.g., one day of work or 1 kg of beef).

In Figure 9, we document that market participants have strong norms against under-cutting one another. While making a sale at the prevailing price is universally accepted (Panel A),

[33]Due to a clerical error, we do not have data for all survey questions among casual laborers in India.

almost all respondents across markets view the practice of making sales below the prevailing price as either Unacceptable or Highly Unacceptable on a four-point scale (Panel B). Consistent with this view, very few respondents would be likely to accept an offer below the prevailing price, even on a slow day (Panel C).

In Figure 10 we show that workers expect to receive costly sanctions from selling below the market price. In Panel A, the majority of market participants view negative consequences from such behavior to be Highly Likely, with very few respondents deeming the likelihood of negative consequences to be Somewhat or Highly Unlikely. In Panel B, we present the types of consequences that sellers who violate the norm may face. Between 73% and 98% of believe that others in the market will become angry with them. Sellers also predict that there would be costs to the their economic livelihood: 43%-83% believe that the other sellers in the local market would try to prevent the individual from being able to sell there in the future (e.g. via social pressure or "ganging up" on them); 23%-50% believe that others would try to actively steer customers away from the individual; and between 63%-97% believe that the other local sellers would spread information about the individual's norm violation to other local markets to make it difficult for the individual to conduct business elsewhere as well. We also find evidence that in some types of markets (especially labor stands), violators may be subject to social sanctions, such as socializing in the evenings, as well as the withdrawal of informal credit.

These findings are all broadly consistent with our field experiment setting of rural workers in Odisha, India and provide suggestive evidence that the mechanisms we document have broader applicability in developing country contexts. For example, the communal nature of village economies may help explain price homogeneity and price rigidity in a range of factor markets—from bullock and tractor rentals to sharecropping (e.g. Shaban, 1987). Moreover, our survey evidence above indicates that such outcomes are not necessarily limited to rural settings, but also may be prevalent in urban markets across a range of goods. Across all these settings, high levels of social capital and repeat interaction in networks creates scope for norms to arise and be sustained through social sanctions—generating the potential for market power in decentralized spot markets. While only speculative, one implication of this idea is that, through the process of development, the anonymity that arises in markets involving firms rather than individuals has the potential to generate more competitive outcomes.

However, given that a substantive component of economic exchange—for example, in the labor market—must invariably involve repeat interpersonal interaction, our findings suggest that the relevance of social norms for market outcomes is not limited to developing countries. For example, in *The Labor Market as a Social Institution*, Robert Solow (1990) argues that in rich countries, social norms within establishments shape collective labor supply. At most firms, workers work alongside their peers for hours every day; they also have the ability to levy sanctions on peers that are meaningful for success at work (e.g., help with tasks, undermining a peer with the manager) and also meaningful for experienced utility (e.g., eating lunch together, chatting at the coffee machine, warmth in routine exchanges). He posits that the resultant implicit social pressures on workers—such as pressure to not undercut co-worker wages by accepting wage cuts, or to not exceed co-workers' output—are important for understanding outcomes such as wage rigidity, equilibrium unemployment, and productivity compression. Consistent with this, collective behaviors among workers—e.g., mass walkouts with foregone wages, the coordinated restriction of output, punishment of rate busters—have been qualitatively documented across history and contexts in settings with no formal unions (Mathewson, 1969; Clark, 1984, 1987).

More broadly, the presence of cartel-like behavior among decentralized individuals has been qualitatively observed in a range of US markets. For example, NASDAQ traders historically only quoted buy and sell prices in denominations of even-eighths (i.e. 2/8, 4/8, etc); avoiding odd-eights increased their income by widening bid-ask spreads (Christie and Schultz, 1994). However, in its formal criminal investigation of this behavior, the US Department of Justice was unable to make a case for explicit coordination or collusion. Rather, interviews indicated the presence of a norm: traders implicitly understood this convention and felt that deviations would be punished through sanctions by fellow traders—e.g., exclusion from social events like after-work drinks, and awkwardness or hostility when facing a trader on the other side of a transaction.[34] Similarly, in many US cities, all real estate agents charged a 6% commission on housing sales, regardless of circumstance (Hsieh and Moretti, 2003). Here as well, agents have to interact with each other on each deal, and are also dependent on agent networks for referrals—making social disapprobation potentially both socially and financially costly.

The above discussion is, of course, only speculative. However, the strength of our findings—along with the potential generality of our proposed mechanism—offers impetus to understand the extent to which social norms affect collective behaviors and market equilibria in a range of other settings.

References

Abdulai, A. and C. L. Delgado (1999): "Determinants of nonfarm earnings of farm-based husbands and wives in Northern Ghana," *American Journal of Agricultural Economics*, 81, 117–130.

[34]We thank Doug Bernheim for pointing us to this example.

Alatas, V., A. Banerjee, A. G. Chandrasekhar, R. Hanna, and B. A. Olken (2016): "Network structure and the aggregation of information: Theory and evidence from Indonesia," *American Economic Review*, 106, 1663–1704.

Atkin, D. and D. Donaldson (2015): "Who's getting globalized? The size and implications of intra-national trade costs," .

Atkin, D., B. Faber, and M. Gonzalez-Navarro (2018): "Retail globalization and household welfare: Evidence from mexico," *Journal of Political Economy*, 126, 1–73.

Azariadis, C. (1975): "Implicit Contracts and Underemployment Equilibria," *Journal of Political Economy*, 83, 1183–1202.

Bandiera, O., I. Barankay, and I. Rasul (2005): "Social preferences and the response to incentives: Evidence from personnel data," *The Quarterly Journal of Economics*, 120, 917–962.

Bardhan, P. K. (1979): "Labor supply functions in a poor agrarian economy," *The American Economic Review*, 69, 73–83.

Bénabou, R. and J. Tirole (2006): "Incentives and prosocial behavior," *American Economic Review*, 96, 1652–1678.

Bergquist, L. F. and M. Dinerstein (2020): "Competition and entry in agricultural markets: Experimental evidence from Kenya," *American Economic Review*, 110, 3705–47.

Bloch, F., G. Genicot, and D. Ray (2008): "Informal insurance in social networks," *Journal of Economic Theory*, 143, 36–58.

Breza, E. and A. G. Chandrasekhar (2019): "Social Networks, Reputation, and Commitment: Evidence From a Savings Monitors Experiment," *Econometrica*, 87, 175–216.

Breza, E., A. G. Chandrasekhar, and A. Tahbaz-Salehi (2018a): "Seeing the forest for the trees? An investigation of network knowledge," Tech. rep., National Bureau of Economic Research.

Breza, E., S. Kaur, and Y. Shamdasani (2018b): "The morale effects of pay inequality," *The Quarterly Journal of Economics*, 133, 611–663.

——— (2021): "Labor Rationing," *American Economic Review*, 111, 3184–3224.

Bursztyn, L., T. Fujiwara, and A. Pallais (2017): "'Acting wife': Marriage market incentives and labor market investments," *American Economic Review*, 107, 3288–3319.

Bursztyn, L. and R. Jensen (2015): "How does peer pressure affect educational investments?" *The Quarterly Journal of Economics*, 130, 1329–1367.

——— (2017): "Social image and economic behavior in the field: Identifying, understanding, and shaping social pressure," *Annual Review of Economics*, 9, 131–153.

Busso, M. and S. Galiani (2019): "The causal effect of competition on prices and quality: Evidence from a field experiment," *American Economic Journal: Applied Economics*, 11, 33–56.

Casaburi, L. and T. Reed (2021): "Using individual-level randomized treatment to learn about market structure," *American Economic Journal: Applied Economics*.

Charness, G. and M. Rabin (2002): "Understanding social preferences with simple tests," *The Quarterly Journal of Economics*, 117, 817–869.

Christie, W. G. and P. H. Schultz (1994): "Why do Nasdaq market makers avoid odd-eighth quotes?" *The Journal of Finance*, 49, 1813–1840.

Clark, G. (1984): "Authority and efficiency: the labor market and the managerial revolution of the late nineteenth century," *The Journal of Economic History*, 44, 1069–1083.

——— (1987): "Why isn't the whole world developed? Lessons from the cotton mills," *The Journal of Economic History*, 47, 141–173.

Dasgupta, P. and D. Ray (1986): "Inequality as a determinant of malnutrition and unemployment: Theory," *The Economic Journal*, 1011–1034.

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.

DellaVigna, S., J. A. List, U. Malmendier, and G. Rao (2016): "Voting to tell others," *The Review of Economic Studies*, 84, 143–181.

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*.

Ellison, G. (1994): "Cooperation in the prisoner's dilemma with anonymous random matching," *The Review of Economic Studies*, 61, 567–588.

Fehr, E. and I. Schurtenberger (2018): "Normative foundations of human cooperation," *Nature Human Behaviour*, 2, 458.

Gachter, S. and E. Fehr (1999): "Collective action as a social exchange," *Journal of Economic Behavior & Organization*, 39, 341 – 369.

Goldberg, J. (2016): "Kwacha gonna do? Experimental Evidence about labor supply in rural Malawi," *American Economic Journal: Applied Economics*, 8, 129–49.

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.

Houde, J.-F., T. Johnson, M. Lipscomb, and L. Schechter (2020): "Imperfect competition and sanitation: Evidence from randomized auctions in Senegal," *Unpublished Manuscript*.

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.

Iacovone, L. and D. J. McKenzie (2019): "Shortening Supply Chains: Experimental Evidence from Fruit and Vegetable Vendors in Bogotá," *World Bank Policy Research*

Working Paper.

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.

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.

Kandori, M. (1992): "Social norms and community enforcement," *The Review of Economic Studies*, 59, 63–80.

Karing, A. (2021): "Social signaling and childhood immunization: A field experiment in Sierra Leone," *Working Paper*.

Kaur, S. (2019): "Nominal Wage Rigidity in Village Labor Markets," *American Economic Review*, 109, 3585–3616.

Lee, D. and E. Saez (2012): "Optimal minimum wage policy in competitive labor markets," *Journal of Public Economics*, 96, 739–749.

Leibenstein, H. (1957): *Economic backwardness and economic growth : Studies in the theory of economic development.*, New York: Wiley.

Lewis, W. A. (1954): "Economic Development With Unlimited Supplies of Labour," *The Manchester School*, 22, 139–191.

MacLeod, W. B. (2007): "Can Contract Theory Explain Social Preferences?" *The American Economic Review*, 97, 187–192.

Mas, A. and E. Moretti (2009): "Peers at work," *American Economic Review*, 99, 112–45.

Mathewson, S. B. (1969): *Restriction of output among unorganized workers*, Southern Illinois University Press.

Mazzocco, M. and S. Saini (2012): "Testing efficient risk sharing with heterogeneous risk preferences," *American Economic Review*, 102, 428–68.

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.

——— (2016): "Networks and misallocation: Insurance, migration, and the rural-urban wage gap," *American Economic Review*, 106, 46–98.

NSS, N. S. S. (2010): "Employment and Unemployment Situation in India 2007– 08: NSS Report 531." Tech. rep., Ministry of Statistics and Programme Implementation Government of India, New Delhi.

Olson, M. (2009): *The logic of collective action*, vol. 124, Harvard University Press.

Osmani, S. R. (1990): "Wage determination in rural labour markets," *Journal of Development Economics*, 34, 3–23.

Ostrom, E. (1990): *Governing the commons: The evolution of institutions for collective action*, Cambridge university press.

Ostrom, E., J. Walker, and R. Gardner (1992): "Covenants with and without a sword: Self-governance is possible," *American Political Science Review*, 86, 404–417.

Putnam, R. D. et al. (2000): *Bowling alone: The collapse and revival of American community*, Simon and schuster.

Rosenzweig, M. R. (1978): "Rural wages, labor supply, and land reform: A theoretical and empirical analysis," *American Economic Review*, 68, 847–861.

——— (1988): "Labor markets in low-income countries," in *Handbook of Development Economics*, Elsevier, vol. 1, 713–762.

Rosenzweig, M. R. and O. Stark (1989): "Consumption smoothing, migration, and marriage: Evidence from rural India," *Journal of Political Economy*, 97, 905–926.

Rudra, A. (1982): "Extraeconomic constraints on agricultural labour results of an intensive survey in some villages near Santiniketan, West Bengal," Asian Employment Programme Working Paper, ARTEP, ILO, Bangkok.

Shaban, R. A. (1987): "Testing between competing models of sharecropping," *Journal of Political Economy*, 95, 893–920.

Shapiro, C. and J. E. Stiglitz (1984): "Equilibrium unemployment as a worker discipline device," *The American Economic Review*, 433–444.

Solow, R. M. (1990): *The labor market as a social institution*, Royer lectures, Cambridge, MA: B. Blackwell.

Walker, T. S. and J. G. Ryan (1990): *Village and household economies in India's semi-arid tropics*, Baltimore: Johns Hopkins University Press.

10. Tables

Table 1. Covariate Balance

COVARIATES	(1)	(2)	(3)	(4)	(5)	(6)
Treatments	Wage Cut Private	Wage Cut Employer	Wage Cut Public	Prevailing All	Joint Sig F-Stat	Obs
INDIVIDUAL-LEVEL						
Age	44.337	2.890	1.893	0.881	0.180	442
	(12.639)	(2.032)	(1.911)	(1.806)		
		[0.157]	[0.323]	[0.626]		
1.(Caste: Scheduled Tribe)	0.333	-0.0179	-0.106	-0.0259	0.842	444
	(0.474)	(0.0827)	(0.0671)	(0.0755)		
		[0.829]	[0.116]	[0.731]		
1.(Caste: Scheduled Caste)	0.667	0.0169	0.100	0.0255	0.382	444
	(0.474)	(0.0825)	(0.0670)	(0.0756)		
		[0.838]	[0.136]	[0.737]		
1.(Casual Laborer)	0.944	0.0417	0.0604	0.0486	0.177	446
	(0.230)	(0.0346)	(0.0269)	(0.0267)		
		[0.229]	[0.0260]	[0.0700]		
1.(Casual Laborer-Agriculture)	0.844	-0.0691	-0.0398	-0.00587	0.840	446
	(0.364)	(0.0662)	(0.0620)	(0.0530)		
		[0.298]	[0.522]	[0.912]		
1.(Doesn't Own Land)	0.678	-0.104	-0.0720	-0.0156	0.876	446
	(0.470)	(0.0897)	(0.0745)	(0.0705)		
		[0.249]	[0.335]	[0.825]		
Individual Employment Rate- Past 30 days	0.295	-0.0194	0.0331	0.0521	0.054	427
	(0.206)	(0.0326)	(0.0343)	(0.0315)		
		[0.553]	[0.336]	[0.100]		
Involuntary Unemployment Rate- Past 30 Days	0.428	-0.0130	-0.0880	-0.0642	0.071	427
	(0.229)	(0.0425)	(0.0350)	(0.0356)		
		[0.759]	[0.0130]	[0.0730]		
1.(Employer Experience)	0.318	0.188	0.0346	0.114	0.248	426
	(0.468)	(0.0987)	(0.0999)	(0.0871)		
		[0.0590]	[0.730]	[0.194]		
1.(Completed Exit Survey)	0.874	0.0178	0.0284	-0.00678	0.389	502
	(0.334)	(0.0495)	(0.0458)	(0.0499)		
		[0.719]	[0.537]	[0.892]		
Number of Onlookers	5.162	.	5.162	-0.447	0.481	189
	(2.646)	.	(2.646)	(0.631)		
		.	.	[0.481]		
VILLAGE-LEVEL						
Number of Households	45.639	-3.283	-3.880	1.366	0.172	172
	(13.135)	(3.603)	(3.266)	(3.057)		
		[0.364]	[0.237]	[0.656]		
Involuntary Unemployment Rate - Past 30 Days	0.398	0.0289	0.0243	0.00691	0.449	170
	(0.104)	(0.0280)	(0.0256)	(0.0240)		
		[0.304]	[0.344]	[0.773]		
Info Flow: Wage	0.752	-0.0202	0.0457	0.0612	0.489	182
	(0.306)	(0.0743)	(0.0679)	(0.0631)		
		[0.786]	[0.502]	[0.333]		
Info Flow: Norm Violation	0.622	-0.0674	0.0442	0.0413	0.486	182
	(0.345)	(0.0837)	(0.0765)	(0.0710)		
		[0.422]	[0.564]	[0.562]		
1.(NREGA ever in village)	0.914	-0.110	-0.0492	0.0335	0.213	180
	(0.284)	(0.0710)	(0.0650)	(0.0612)		
		[0.125]	[0.450]	[0.585]		
Prop Days NREGA - Recall	0	0.0106	-0.00248	0.00288	0.410	180
	(0)	(0.0108)	(0.00989)	(0.00931)		
		[0.327]	[0.802]	[0.758]		
Number of Villages	37	34	40	72		37
Number of Observations	103	88	108	203		103

Notes: Col. (1) provides covariate means and standard deviations for the reference group, Wage cut: Private. Cols. (2)-(4) report regression coefficients relative to the reference group. All specifications include Year × Month and task fixed effects. First panel shows individual characteristics. Involuntary unemployment is prop. days in past 30 that worker wanted work but could not find it. Observations weighted by the number of individuals in the regression sample by village. Number of onlookers is recorded only for public treatments. Mean num. onlookers for $W - 10\%$, public is shown in Col. (3). Col. (4) presents the regession coefficient comparing the public W relative to $W - 10\%$. Second panel shows village covariates measured in untreated holdout sample in study villages. Measures of information flow are based on i) average perceived likelihood that knowledge of the wages of others would spread and ii) average perceived likelihood that knowledge of norm violations would spread. Standard errors (clustered at the village level) are reported in parentheses. P-values are reported in brackets. P-values from Wald tests of joint significance of all treatment arms (relative to the reference group) are reported in Col. (5). Variation in sample sizes (Col. (6)) comes from non-response in the exit survey and differences in survey modules across rounds.

Table 2. Main Results - Effects on Labor Supply

	(1) Worked	(2) Worked	(3) Worked
[1] Wage cut: Public	-0.122	-0.136	-0.246
	(0.0564)	(0.0573)	(0.0644)
	[0.0316]	[0.0188]	[0.000181]
[2] Wage cut: Employer	-0.0657	-0.0516	-0.0758
	(0.0611)	(0.0633)	(0.0788)
	[0.284]	[0.416]	[0.337]
[3] Prevailing wage: Private	0.0609	0.0791	0.0663
	(0.0703)	(0.0659)	(0.0819)
	[0.387]	[0.232]	[0.419]
[4] Prevailing wage: Public	0.119	0.116	0.104
	(0.0808)	(0.0713)	(0.0856)
	[0.144]	[0.105]	[0.226]
[5] Prevailing wage: Employer	0.0364	0.0690	0.0935
	(0.0775)	(0.0886)	(0.0992)
	[0.639]	[0.437]	[0.347]
Observations	502	502	363
Task and Year x Month FE		✓	✓
Sample	All Workers	All Workers	Agri. Workers
Dep Var Mean (Wage cut: Private)	0.175	0.175	0.211
Test 1: [3]=[4]	0.460	0.589	0.658
Test 2: omitted-[1] = [3]-[4]	0.0629	0.0481	0.00858
Test 3: [1]=[2]	0.143	0.0865	0.0107
Test 4: [3]=[4]=[5]	0.609	0.816	0.904
Test 5: [4]=[5]	0.331	0.608	0.918
Test 6: {[3],[4],[5]}=0	0.196	0.118	0.228

Notes: This table presents the effect of each treatment on our main outcome of interest: take-up of the job offer. The dependent variable is whether the laborer worked for the employer on the day of work. In all specifications, the omitted category is the Wage cut: Private treatment. Cols. (2) and (3) include year $\times$ month (i.e., strata) fixed effects and task fixed effects. Cols. (1) and (2) include the full sample. Col. (3) restricts the sample to workers who report 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. P-values are reported in brackets. *Test 1* compares (w, Private) with (w, Public). *Test 2* is the differences-in-differences test. *Test 3* compares ($w - 10\%$, Employer) with ($w - 10\%$, Private). *Test 4* is a test of equality of the three w treatments. *Test 5* compares (w, Public) with (w, Employer), and *Test 6* is a joint test for whether the w treatments are distinguishable from ($w - 10\%$, Private).

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	-32.42	-0.0376	-6.794	-0.0646	-11.82
	(0.0510)	(11.13)	(0.0278)	(7.019)	(0.0249)	(6.942)
	[0.00190]	[0.00405]	[0.177]	[0.334]	[0.0102]	[0.0903]
Prevailing wage (pooled)	0.0937	27.97	0.0170	3.747	0.0230	6.690
	(0.0515)	(13.07)	(0.0247)	(6.167)	(0.0252)	(6.399)
	[0.0706]	[0.0338]	[0.491]	[0.544]	[0.363]	[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	✓	✓	✓	✓	✓	✓
Dep Var Mean (Omitted)	0.222	45.49	0.0781	17.96	0.110	24.09
Dep Var Mean (Holdout)	0.0743	16.53	0.0743	16.53	0.0743	16.53

Notes: This table presents treatment effects on worker earnings, measured using an employment recall grid in the exit survey. Each observation is a worker-day. In Cols. (1), (3), and (5), the dependent variable is an indicator for whether the respondent worked in agricultural wage labor that day. In Cols. (2), (4), and (6), the dependent variable is the total wage (cash + in kind) earned on that day in agricultural work. In all specifications, the omitted category is being assigned to a non-*Public* (i.e. *Private* or *Employer only*) wage cut arm. All specifications include year × month (i.e., strata) fixed effects and task fixed effects. Cols. (1)-(2) restrict the sample to the day of work offered as part of the experiment. Cols. (5)-(6) include all the days from the day before the experimental day of work up to 5 days after the experimental day of work; this includes the experimental day of work. Cols. (3)-(4) use the same sample as Cols. (5)-(6) but exclude the experimental day of work. In Cols. (1)-(4), observations are weighted by the number of experimental subjects in each village. In Cols. (5)-(6), 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. In all regressions, standard errors are clustered at the village level and are reported in parentheses. P-values are reported in brackets.

Table 4. Heterogeneous Treatment Effects: Village Information Spread

	(1) Worked	(2) Worked	(3) Worked
Wage cut: Public	-0.200	-0.186	-0.308
	(0.0675)	(0.0646)	(0.0745)
	[0.00344]	[0.00442]	[5.50e-05]
Wage cut: Public x Low info spread village	0.170	0.150	0.214
	(0.0932)	(0.0921)	(0.114)
	[0.0701]	[0.105]	[0.0614]
Low info spread village	-0.0732	-0.0380	-0.0263
	(0.0667)	(0.0621)	(0.0796)
	[0.274]	[0.541]	[0.742]
Observations	499	499	361
Task and Year x Month FE	✓	✓	✓
Low info definition	Wage info	Norm violation	Norm violation
Sample	All Workers	All Workers	Agri. Workers
Dep Var Mean (Omitted)	0.204	0.200	0.214

Notes: This table presents heterogeneous treatment effects by village-level diffusiveness, using two measures captured from surveys with the untreated holdout survey respondents. In Col. (1), "Low info spread village" is an indicator for below-median knowledge of the wages of others. In Cols. (2)-(3), it is 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 worked for the employer. In all columns, the omitted category is being assigned to the non-*Public* (i.e. *Private* or *Employer only*) wage cut arm in villages where "Low info spread village" equals zero. All specifications include year $\times$ month (i.e. strata) and task fixed effects. Col. (3) restricts the sample to those who indicated in the exit survey 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. P-values are reported in brackets.

TABLE 5. Sanctions - Costly Punishment Game

	(1) Any Punishment	(2) Any Punishment	(3) Any Punishment	(4) Any Punishment	(5) Partner's Payoff
Worker Accepted Job Below Prevailing Wage	0.420	0.393	0.393	0.436	-14.57
	(0.0447)	(0.0632)	(0.0647)	(0.103)	(4.425)
	[0]	[5.99e-09]	[1.23e-08]	[4.66e-05]	[0.00132]
Worker Accepted a Job Below Prevailing Wage x Distant Labor Market		0.0494	0.0494	-0.00310	5.569
		(0.0894)	(0.0916)	(0.137)	(4.551)
		[0.582]	[0.590]	[0.982]	[0.224]
Worker Lives in Distant Labor Market		0.0143	0.0133	0.00737	-0.701
		(0.0143)	(0.0185)	(0.0294)	(1.259)
		[0.320]	[0.472]	[0.803]	[0.579]
Observations	262	262	262	131	131
Village FE			✓	✓	✓
First Round Only				✓	✓
Dep Var Mean: Worker Accepted 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 workers) in 13 villages (separate from the 183 used 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 (omitted category). Each observation is a player $\times$ round. 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 of Rs. 100 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. Cols. (3)-(5) include village fixed effects. Standard errors are clustered by player and reported in parentheses. P-values are given in brackets.

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)
Panel A - Dependent variable: Log Agricultural Wage			
Positive shock last year	0.0532	0.102	0.0971
	(0.022)	(0.042)	(0.033)
	[0.017]	[0.015]	[0.003]
Positive shock last year x Low worker cohesion		-0.0826	-0.0899
		(0.050)	(0.038)
		[0.100]	[0.019]
Positive shock this year	0.0633	0.0800	0.0751
	(0.018)	(0.038)	(0.039)
	[0.000]	[0.036]	[0.055]
Positive shock this year x Low worker cohesion		-0.0242	-0.0181
		(0.042)	(0.043)
		[0.561]	[0.675]
Observations (worker-days)	59243	59243	59243
Panel B - Dependent variable: Agricultural Employment			
Positive shock last year	-0.135	-0.234	-0.172
	(0.055)	(0.078)	(0.080)
	[0.014]	[0.003]	[0.031]
Positive shock last year x Low worker cohesion		0.189	0.0716
		(0.088)	(0.107)
		[0.032]	[0.505]
Positive shock this year	0.157	0.133	0.131
	(0.062)	(0.083)	(0.091)
	[0.011]	[0.109]	[0.153]
Positive shock this year x Low worker cohesion		0.0394	0.0469
		(0.114)	(0.123)
		[0.729]	[0.702]
Dep Var Mean	1.74	1.74	1.74
Observations (workers)	632324	623861	631909

Notes: This table presents the effect of current and lagged productivity (rainfall) shocks on wages and employment and examines the heterogeneity of the effect by two measures of worker cohesion. Under wage rigidity, lagged positive shocks will generate ratcheting, increasing wages and decreasing employment the next year. Source: National Sample Survey data (1986-2007). Positive shock = 1{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. Low worker cohesion is defined as a) a below-median Herfindahl index of caste (i.e. higher caste heterogeneity) among workers who report any daily-wage labor (Col. 2), and b) a below-median Herfindahl index of caste (i.e. higher caste heterogeneity) among all workers who report agricultural labor as their primary or secondary occupation (Col. 3). 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. All regressions include year and district fixed effects. Standard errors in parentheses clustered by region-year; p-values in brackets.

11. Figures

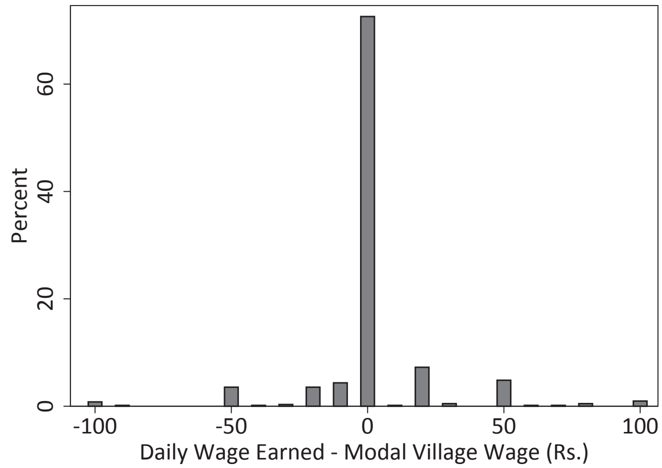

(a) Distribution of Wages Inside the Village. *Source:* Breza, Kaur, and Shamdasani (2019).

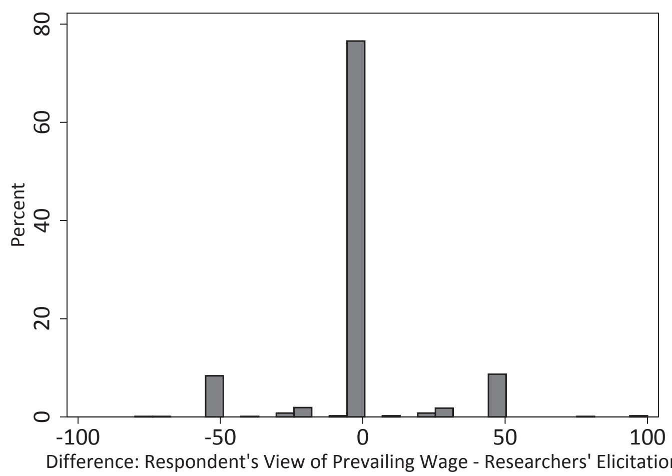

(b) Distribution of Wages Inside the Village: Study Sample

Figure 1. **Dispersion of Agricultural Wages.**
Notes: Panel A displays the distribution of (wage reported by worker) − (mode of wage reported by all workers in village). Respondents were asked to list their work activity and wage level on each of the past 10 days. Observations are at the worker-day level. Source: Data collected by Breza et al. (2018b). N = 377 worker-days in 26 villages in Odisha. Panel B displays the distribution of (prevailing daily agricultural wage for men in village reported by untreated holdout sample of workers (N=626 workers)) − (mode of male agricultural wage in village). The latter was elicited from employers and village resident who helped with the worker census.

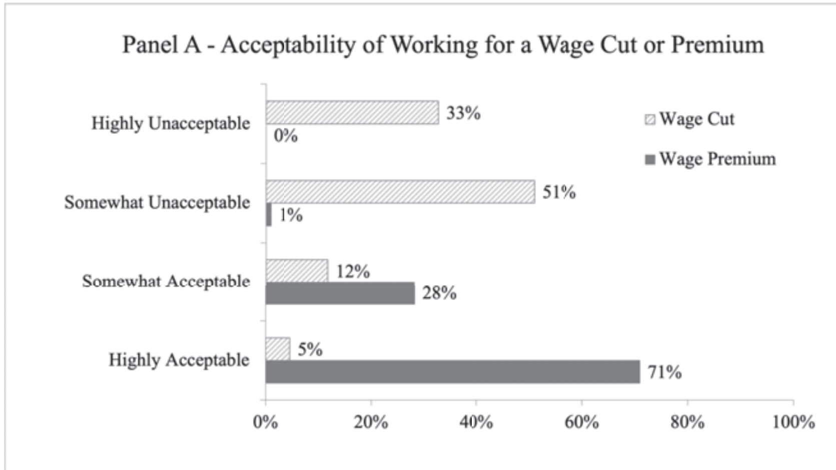

(a) Acceptability of Working for a Wage Cut or Premium.

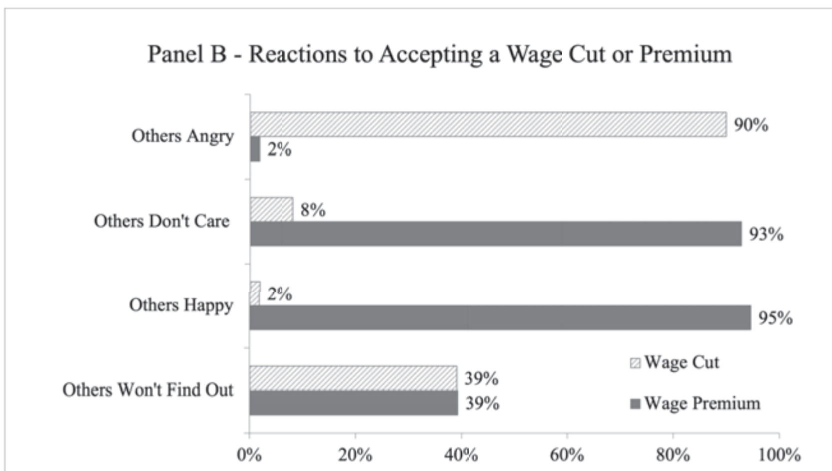

(b) Reactions to Accepting a Wage Cut or Premium.

Figure 2. **Motivational Evidence - Norms around Wage Cuts and Increases.**
Notes: These figures graph the survey responses from 110 respondents in 10 non-study villages in the study districts to questions about the acceptability and consequences of taking a job at a wage cut or wage premium. We restrict the sample to participants from villages where jobs were offered at the prevailing wage in our experiment.

Social Observability	Wage Level: w - 10%	Wage Level: w	
Public	A (0.2)	D (0.16)	Job offer made on street in front of worker's home
Employer only	B (0.2)	E (0.08)	Job offer made on premises of worker's home
Private	C (0.2)	F (0.16)	Job offer made on premises of worker's home: employer out of earshot for wage

FIGURE 3. **Experimental Design.** *Notes*: Numbers in parentheses indicate randomization weights on each treatment cell.

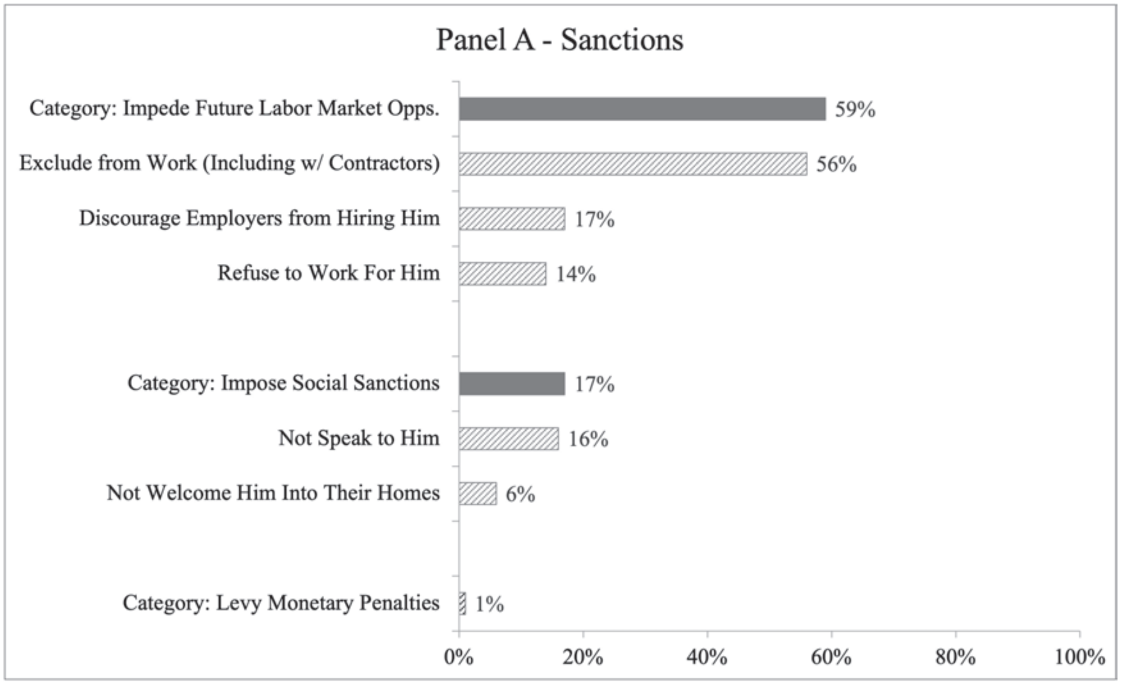

(a) Sanctions

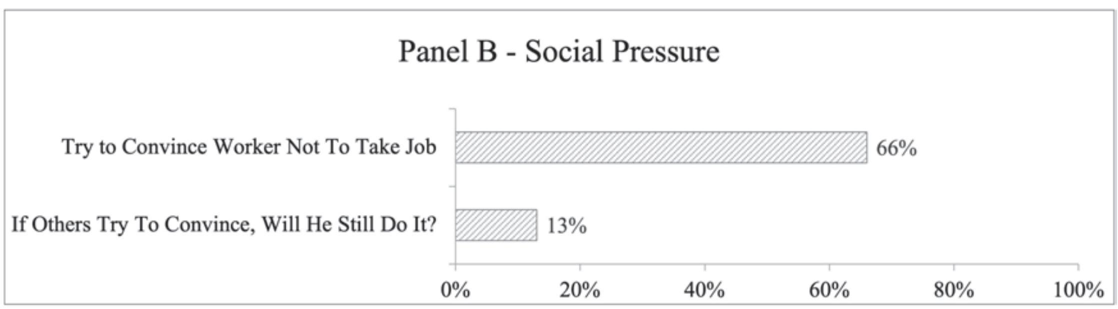

(b) Social Pressure

FIGURE 4. **Survey Evidence - Sanctions for Accepting Wage Cuts.** *Notes*: This figure graphs exit survey responses from 1,448 untreated holdout sample workers to the question: “Suppose a worker 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), the second question, “If other workers try to convince him not to take the job, will he still do it?” was asked of all respondents (regardless of their answer to the question in panel (a)).

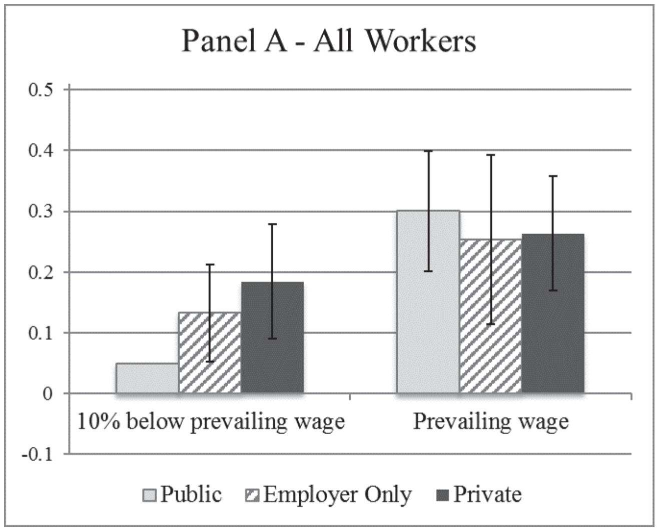

(a) All Workers

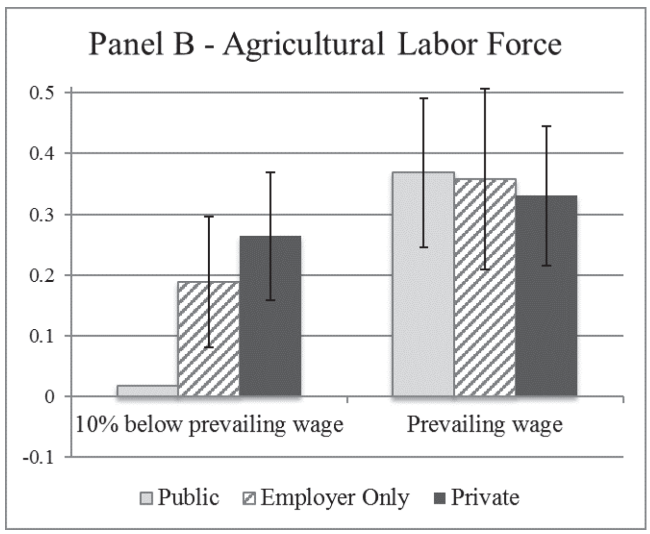

(b) Agricultural Workers Only

Figure 5. **Main Results - Labor Supply.**
Notes: The figure graphs take-up rates for the job offer under each of the 6 treatment arms. The take-up rate for the job is defined as the proportion of workers who work for the employer in the given arm. Panel A uses the entire sample (N=502 participants); Panel B restricts the sample to those who report their primary or secondary occupation to be agricultural labor (N=363 participants). Each panel plots the raw mean take-up rate for the *Public* wage cut arm, and plots estimated means for the remaining treatment arms using the regression results from Cols. (2)-(3) of Table 2. All regressions include year $\times$ month (i.e. strata) and task fixed effects. 90% CIs are constructed using standard errors (clustered at the village-level) from a test of the difference between the take-up rate in that treatment arm and the take-up rate for the *Public* wage cut job offers.

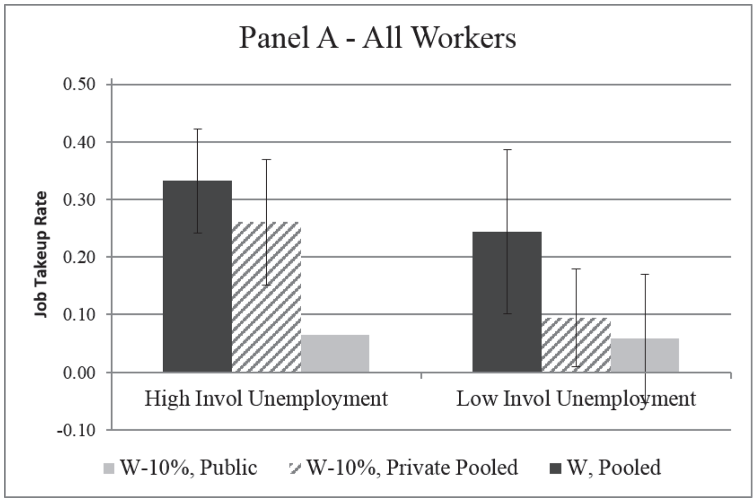

(a) All Workers

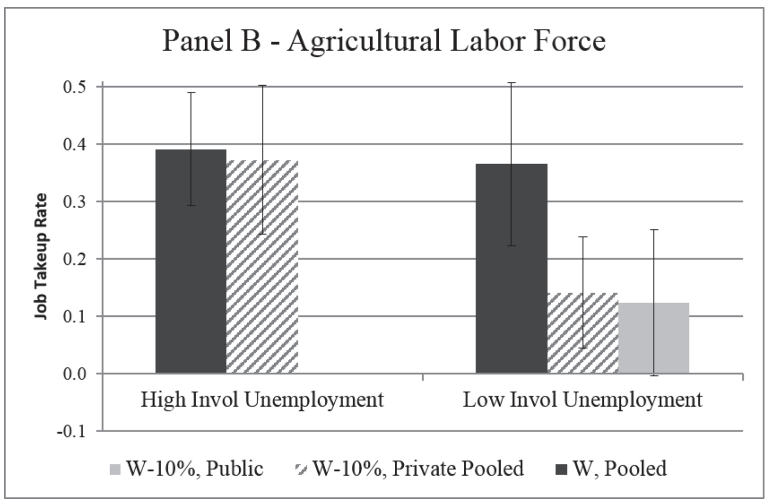

(b) Agricultural Workers Only

Figure 6. **Heterogeneity: Level of Involuntary Unemployment.**
Notes: These figures present heterogeneous treatment effects by below- and above-median village involuntary unemployment. Village involuntary unemployment is measured by the proportion of days in the past 10 days that untreated holdout sample respondents would have preferred a prevailing-wage job to what they did that day. The take-up rate is defined as the proportion of workers who work for the employer in the given treatment arm. Panel A uses the entire sample (N=493 participants); Panel B restricts the sample to workers who report their primary or secondary occupation to be agricultural labor (N=363 participants). 90% CIs are constructed using standard errors (clustered at the village-level) from a test of the difference between the take-up rate for that treatment arm and the take-up rate for the *Public* wage cut job offers in villages with above-median involuntary unemployment. These results are also presented in the form of regressions in Cols. (1) and (2) of Table A.5. All regressions include task and year $\times$ month fixed effects.

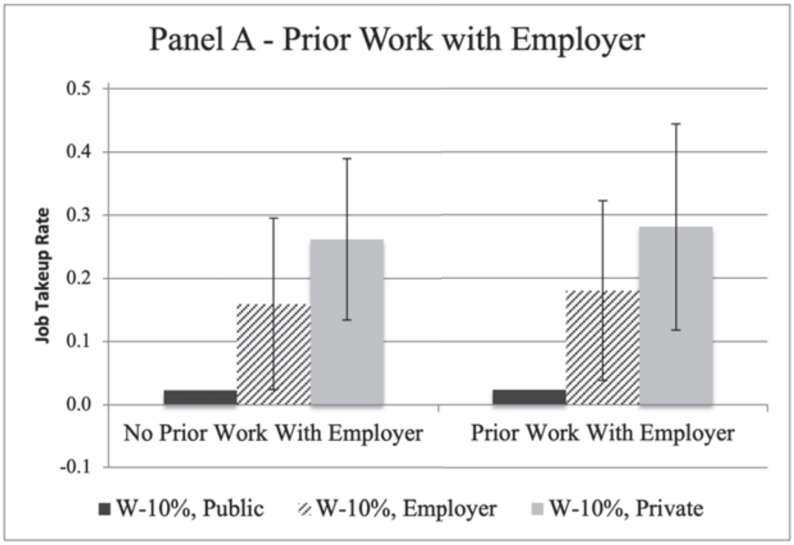

(a) Prior Work with Employer

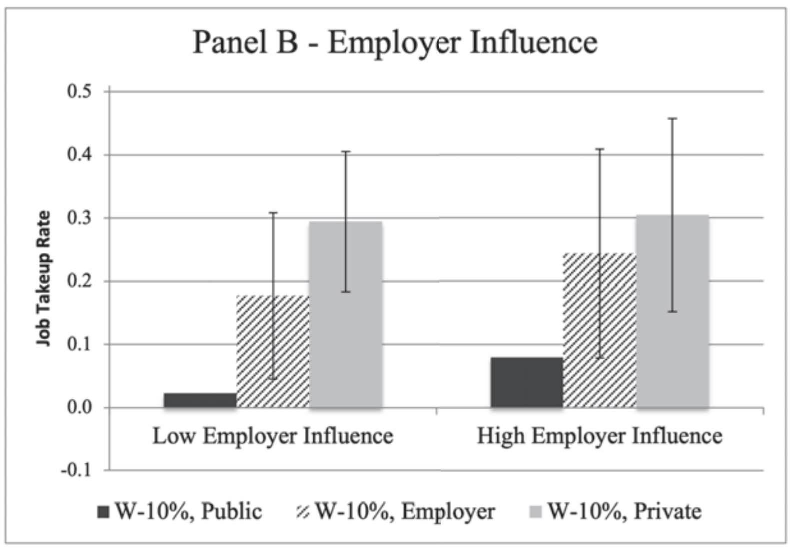

(b) Employer Influence

FIGURE 7. **Heterogeneity: Employer's Relevance for Worker's Future Job Prospects.**
Notes: This figure presents heterogeneous treatment effects by whether the worker is relatively more likely to work with the hiring employer again in the future (as proxied by having worked with the employer in the past, Panel A) and by whether the worker views the employer as influential in the village (Panel B). The take-up rate is defined as the proportion of workers who work for the employer in the given treatment arm. Both panels restrict the sample to workers in the agricultural labor force (who report their primary or secondary occupation to be agriculture) and who responded to the exit survey (N=350). 90% CIs in each panel are constructed using standard errors (clustered at the village-level) from a test of the difference between the take-up rate for that treatment arm and the take-up rate for *Public* wage cut. These results are also presented in the form of regressions in Cols. (1) and (2) of Table A.6. All regressions include year $\times$ month (i.e. strata) and task fixed effects.

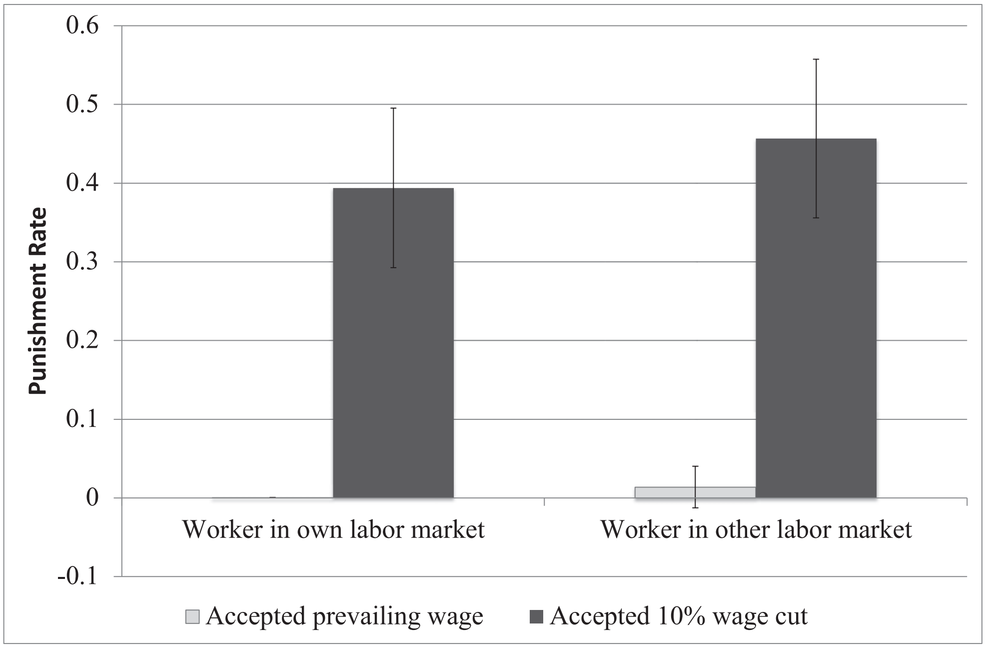

Figure 8. **Sanctions: Costly Punishment.**
Notes: This figure graphs the results of the costly punishment game from N=131 participants in 13 villages (villages are different from those in the main experimental sample). Each participant (player) was anonymously paired with a worker in his own village or in a distant village, and given various scenarios about his paired worker. The figure plots the proportion of times players reduced the endowment of the anonymous worker they were paired with under the two 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. 90% CIs are constructed using standard errors (clustered by player) from a test of the difference between that coefficient and the omitted category (worker in own labor market who accepted a job at the prevailing wage). The plotted rates of punishment correspond to Col. (3) of Table 5, and control for village fixed effects.

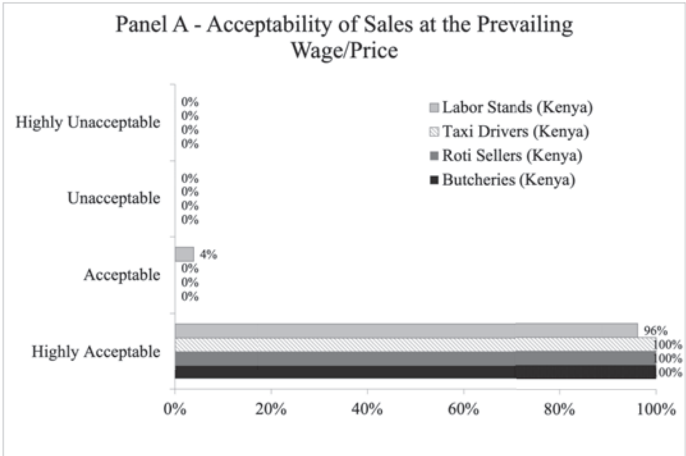

(a) Acceptability of sales at the prevailing wage/price

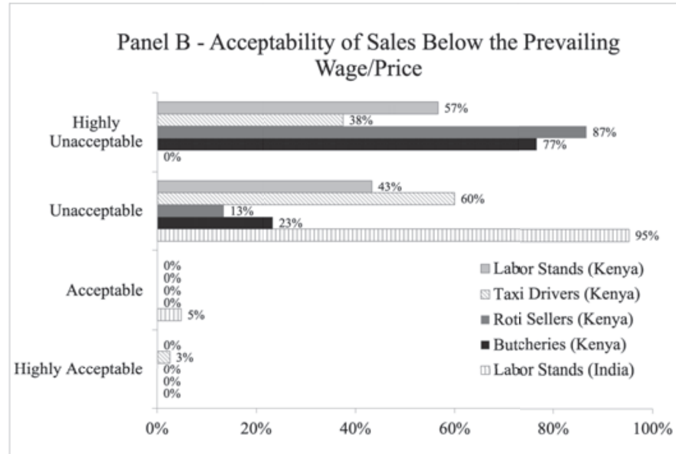

(b) Acceptability of sales below the prevailing wage/price

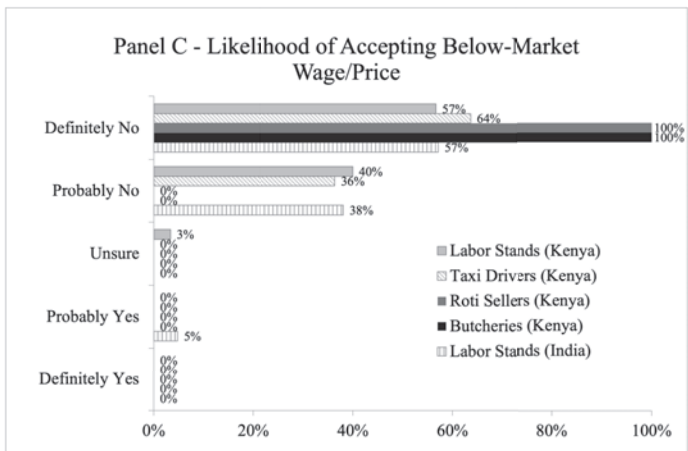

(c) Likelihood of accepting a below-market wage/price

Figure 9. **Norms Against Price Cuts in Other Market Settings.**
Notes: Survey responses from urban workers at labor stands in Bangalore, India (N=30) and Busia, Kenya (N=30), and taxi (boda-boda) drivers (N=40), roti sellers (N=30), and butcher stalls (N=30), all in Busia, Kenya. Responses to the following questions are plotted. Panel A: "Suppose one worker/driver/vendor takes a job/sells their products for «prevailing price». Do you think this behaviour is acceptable?" Panel B: "Suppose one worker/driver/vendor takes a job/sells their products for «prevailing price - 10%». Do you think this behaviour is acceptable?" Panel C: "Suppose it has been a slow day, and many people are trying to find work/customers; you would really like to earn some money. Suppose an employer/customer approaches you. He says that if you work for «prevailing price - 10%», he will hire you instead of other workers/buy your product over others. Would you agree to this lower price?"

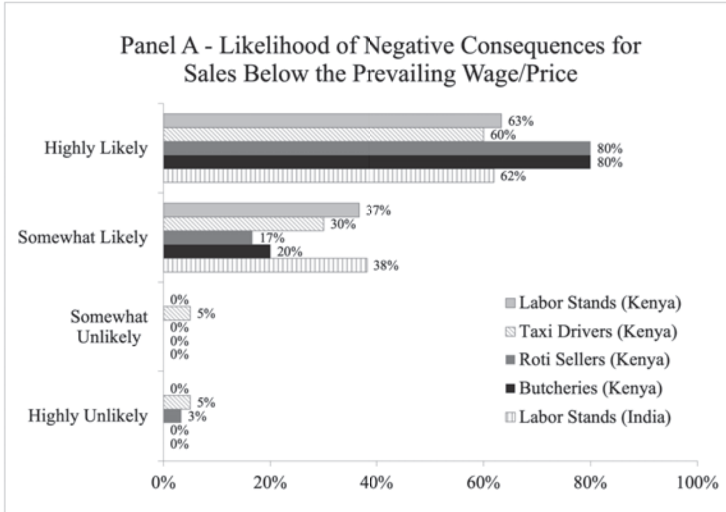

(a) Likelihood of negative consequences for sales below prevailing wage/price

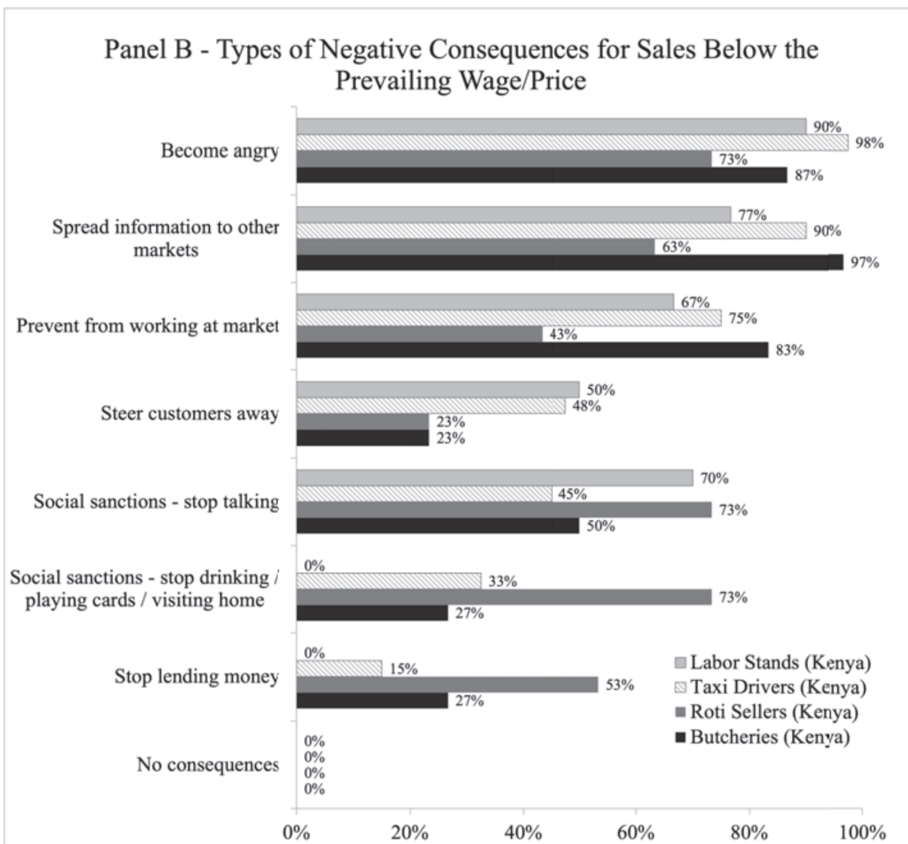

(b) Types of negative consequences for sales below prevailing wage/price

FIGURE 10. **Social Sanctions for Price Cuts in Other Settings**
Notes: Survey responses presented from workers at urban labor stands in Bangalore, India (N=21) and urban labor stand workers (N=30), taxi (boda-boda) drivers (N=40), roti sellers (N=30), and butcher stalls (N=30), in Busia, Kenya. Panel A: "Suppose a worker/driver/vendor from the area accepts a job/makes a sale for «prevailing price - 10%». How likely is it that he would face negative consequences from other workers/drivers/vendors in the area?" Panel B: "Since this worker/driver/vendor took a job/made a sale at «prevailing price - 10%», what are the reactions he will face from the other workers/drivers/vendors?" Respondents were read a list of options and could select any that apply, and also volunteer their own.

Coordination without Organization: Collective Labor Supply in Decentralized Spot Markets

Online Appendices

Emily Breza
Supreet Kaur
Nandita Krishnaswamy

Appendix A. Appendix Tables

Table A.1. Survey Attrition and Untreated Holdout Sample Composition

	(1) Has Endline Survey	(2) Num Control in Village
Wage cut: Public	0.0342	0.407
	(0.0514)	(0.376)
	[0.506]	[0.280]
Wage cut: Employer	0.0124	0.104
	(0.0525)	(0.355)
	[0.813]	[0.770]
Prevailing wage: Private	0.0383	-0.154
	(0.0525)	(0.397)
	[0.467]	[0.698]
Prevailing wage: Public	-0.0857	0.214
	(0.0662)	(0.433)
	[0.197]	[0.621]
Prevailing wage: Employer	0.0696	0.834
	(0.0554)	(0.486)
	[0.211]	[0.0876]
Observations	502	502
Task and Year x Month FE	✓	✓
Sample	All Workers	All Workers
Dep Var 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 by a member of the main experimental sample, by treatment. Col. (2) reports the number of untreated holdout sample surveys conducted in each village. In all columns, the omitted category is the Wage cut: Private treatment. All specifications include year × month and task fixed effects. 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. P-values are reported in brackets.

Table A.2. Main Results With Randomization Inference

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

Notes: This table presents our main specifications from Table 2 with p-values based on randomization inference. 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.3. Main Results: Robustness to Controls For Imbalance

	(1) Worked	(2) Worked
Wage cut: Public	-0.177	-0.243
	(0.0634)	(0.0644)
	[0.00588]	[0.000230]
Wage cut: Employer	-0.0934	-0.101
	(0.0691)	(0.0760)
	[0.178]	[0.187]
Prevailing wage: Private	0.0224	0.0320
	(0.0708)	(0.0760)
	[0.752]	[0.674]
Prevailing wage: Public	0.0881	0.0908
	(0.0779)	(0.0859)
	[0.260]	[0.292]
Prevailing wage: Employer	0.0407	0.0796
	(0.0924)	(0.0966)
	[0.660]	[0.411]
Observations	427	350
Task and Year x Month FE	✓	✓
Sample	All Workers	Agri. Workers
Dep Var Mean (Wage cut: Private)	0.205	0.216
Test Prevailing wage: Private = Prevailing wage: Public	0.372	0.482
Test Wage cut: Private - Public = Prev. wage: Private - Public	0.0141	0.00476
Test Wage cut: Employer = Wage cut: Public	0.140	0.0328
Test Prev. wage: Private = Employer = Public	0.663	0.763
Test Prevailing wage: Employer = Prevailing wage: Public	0.619	0.913

Notes: This table presents results from our primary specification, with controls for variables imbalanced across treatments, as a robustness check. In all specifications, the dependent variable is an indicator for whether the worker worked for the hiring employer. All specifications include controls for participating in the non-agricultural casual labor market and number of days employed for a wage in the past 30 (our measure of individual employment). In all columns, the omitted category is the Wage cut: Private treatment. All specifications include year $\times$ month fixed effects and task fixed effects. Col. (1) includes the full sample. Col. (2) restricts the sample to workers who indicated in the exit survey 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. P-values are reported in brackets.

Table A.4. Main Results: Sample Robustness

	(1) Worked	(2) Worked	(3) Accepted Offer
Wage cut: Public	-0.126	-0.122	-0.0817
	(0.0820)	(0.0645)	(0.0474)
	[0.127]	[0.0611]	[0.0869]
Wage cut: Employer	0.0260	-0.0374	-0.0377
	(0.0911)	(0.0702)	(0.0493)
	[0.775]	[0.595]	[0.445]
Prevailing wage: Private	0.0664	0.0788	0.0598
	(0.100)	(0.0754)	(0.0598)
	[0.508]	[0.297]	[0.318]
Prevailing wage: Public	0.136	0.0966	0.0793
	(0.102)	(0.0776)	(0.0514)
	[0.183]	[0.215]	[0.125]
Prevailing wage: Employer	0.126	0.137	0.0629
	(0.131)	(0.105)	(0.0746)
	[0.335]	[0.197]	[0.400]
Observations	188	359	545
Sample Restriction	First HH	First Two HHs	Intended Sample
Task and Year x Month FE	✓	✓	✓
Dep Var Mean (Wage cut: Private)	0.158	0.173	0.213
Test Prevailing wage: Private = Prevailing wage: Public	0.506	0.824	0.725
Test Wage cut: Private - Public = Prev. wage: Private - Public	0.139	0.171	0.170
Test Wage cut: Employer = Wage cut: Public	0.0628	0.161	0.241

Notes: This table presents results from our primary specification, restricted to various samples as a robustness check. In all specifications, the dependent variable is an indicator for whether the worker worked for the hiring employer. In all columns, the omitted category is the Wage cut: Private treatment. All specifications include year × month fixed effects and task fixed effects. In Col. (1), the sample is restricted to the first household approached in each village, and in Col. (2), the sample is restricted to the first two households approached in each village. In Col. (3), the sample is restricted to the intended main experimental 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). 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. P-values are reported in brackets.

TABLE A.5. Heterogeneous Treatment Effects: Individual Unemployment History

	(1) Worked	(2) Worked	(3) Worked	(4) Worked	(5) Worked	(6) Worked
Wage cut: Public	-0.196	-0.372	-0.213	-0.287	-0.298	-0.436
	(0.0664)	(0.0792)	(0.0742)	(0.0693)	(0.0880)	(0.0839)
	[0.00354]	[5.24e-06]	[0.00463]	[5.36e-05]	[0.000877]	[5.84e-07]
Prevailing wage (pooled)	0.0712	0.0192	0.133	0.124	0.0735	0.0313
	(0.0645)	(0.0773)	(0.0693)	(0.0798)	(0.0761)	(0.0889)
	[0.271]	[0.804]	[0.0575]	[0.123]	[0.336]	[0.725]
Low Village Unemployment	-0.166	-0.231			-0.192	-0.238
	(0.0699)	(0.0875)			(0.0733)	(0.0850)
	[0.0187]	[0.00895]			[0.00950]	[0.00576]
Wage cut: Public x Low Village Unemployment	0.160	0.355			0.200	0.358
	(0.0909)	(0.116)			(0.0989)	(0.114)
	[0.0808]	[0.00252]			[0.0449]	[0.00201]
Prevailing wage (pooled) x Low Village Unemployment	0.0780	0.205			0.134	0.198
	(0.0948)	(0.124)			(0.104)	(0.125)
	[0.412]	[0.101]			[0.201]	[0.114]
Low Individual Unemployment			-0.0901	-0.0988	-0.0822	-0.0892
			(0.0500)	(0.0605)	(0.0484)	(0.0557)
			[0.0734]	[0.105]	[0.0912]	[0.111]
Wage cut: Public x Low Individual Unemployment			0.132	0.154	0.119	0.140
			(0.0711)	(0.0685)	(0.0706)	(0.0687)
			[0.0648]	[0.0262]	[0.0927]	[0.0426]
Prevailing wage (pooled) x Low Individual Unemployment			-0.00102	0.0310	-0.0144	0.0193
			(0.0752)	(0.0877)	(0.0753)	(0.0846)
			[0.989]	[0.724]	[0.849]	[0.819]
Observations	493	363	427	350	427	350
Sample	All Workers	Agri. Workers	All Workers	Agri. Workers	All Workers	Agri. Workers
Task and Year x Month FE	✓	✓	✓	✓	✓	✓
Dep Var Mean (Wage cut: Private, High unempl.)	0.333	0.393	0.262	0.282	0.262	0.282
Take-up Wage cut: Public, High unempl.	0.0611	0	0.0882	0.0370		
Take-up Wage cut: Public, Low unempl.	0.0308	0.0444	0.0500	0.0217		
Test 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 defined as an indicator for below-median unemployment at the village level. Unemployment is measured as the number of days in the 10-day recall that the respondent reports preferring work at the prevailing wage to whatever esle they did that day. Only the untreated holdout sample surveys are used to compute this village-level measure. 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. In all specifications, the dependent variable is an indicator for whether the worker accepted the job and worked for the employer. In all specifications, the omitted category is the Wage cut: Private treatment for the workers with high (individual or village or both) unemployment. All specifications include year $\times$ month fixed effects and task fixed effects. Cols. (1), (3), and (5) include the full sample. Cols. (2), (4), and (6) restrict the sample to workers whose primary or secondary occupation is agricultural labor. 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. P-values are reported in brackets.

Table A.6. Heterogeneous Treatment Effects: Employer's Relevance for Worker's Future Job Prospects

VARIABLES	(1) Worked	(2) Worked
Wage cut: Public	-0.239	-0.271
	(0.0778)	(0.0677)
	[0.00249]	[9.35e-05]
Wage cut: Employer	-0.103	-0.117
	(0.0953)	(0.0907)
	[0.284]	[0.197]
Prevailing wage (pooled)	0.0376	0.0739
	(0.0880)	(0.0749)
	[0.670]	[0.326]
Wage cut: Public x Employer Experience	-0.0184	0.0464
	(0.118)	(0.0899)
	[0.876]	[0.606]
Wage cut: Employer x Employer Experience	0.00161	0.0566
	(0.143)	(0.144)
	[0.991]	[0.696]
Prevailing wage (pooled) x Employer Experience	0.107	0.0193
	(0.135)	(0.109)
	[0.430]	[0.860]
Employer Experience	0.0193	0.00985
	(0.108)	(0.0831)
	[0.858]	[0.906]
Observations	350	350
Employer Experience	Prior Work	Influential
Task and Year x Month FE	✓	✓
Sample	Agri. Workers	Agri. Workers
Test Wage cut: Public + Public x Experience = 0	0.0105	0.0168
Test Wage cut: Employer + Employer x Experience = 0	0.395	0.627
Test Wage cut: Pub. + Pub. x Exp. = Wage cut: Empl. + Empl. x Exp.	0.0725	0.105
Dep Var Mean (Omitted)	0.188	0.196

Notes: This table presents heterogeneous treatment effects by measures of how likely the hiring employer is to affect the worker's future job prospects. In Col. (1), "Employer Experience", is an indicator for the worker having ever worked for the hiring employer in the past (which positively predicts likelihood of future work). In Col. (2), "Employer Experience", is an indicator for whether the worker views the employer as particularly influential in the village. In all specifications, the dependent variable is an indicator for whether the worker worked for the employer. In all columns, the omitted category is being assigned to a non-*Public* (i.e. *Private* or *Employer only*) wage cut arm for the low employer experience group only. All specifications include year × month fixed effects and task fixed effects, and restrict the sample to workers who report agricultural labor as their 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. P-values are reported in brackets.

TABLE A.7. Side Payments, Job Amenities, and Worker Effort

	(1) Received vs. Offered Cash Wage (%)	(2) Length of work (mins)	(3) Number of meals included	(4) Hired before	(5) Work day rating	(6) Work day rating
Wage Cut	0.0243	-2.863	-0.280	-0.207	0.0206	0.0469
	(0.0259)	(16.91)	(0.291)	(0.195)	(0.297)	(0.402)
	[0.354]	[0.866]	[0.340]	[0.294]	[0.945]	[0.908]
Observations	70	74	74	77	74	61
Task and Year x Month FE	✓	✓	✓	✓	✓	✓
Dep Var Mean (Omitted)	0	313.3	0.690	0.652	1.178	1.178

Notes: This table presents statistics on job amenities, and worker effort and quality on the day of work. The dependent variable in Col. (1) is the percent difference between the total cash wage received (including any side transfers) and the offered wage. Total cash wage received, length of the work day (Col. (2)), and meals received from the employer (Col. (3)) are measured in the exit survey. Workers' prior experience with employer is measured by an indicator for having been hired before by the employer (Col. (4)), and is also collected in the exit survey. Worker quality (Cols. (5) and (6)) is reported for the day of work on a rating scale of 1-4 by the employer. In all specifications, the omitted category is Prevailing wage (Pooled). All specifications include year × month and task fixed effects. The sample is restricted to all workers who came to the job on the day of work. Cols. (1)-(5) consider the full sample of workers, while in Col (6) we only consider the *Private* $w - 10\%$ treatment compared against all of the w treatments. Standard errors are clustered at the village level and reported in parentheses. P-values are presented in brackets.

Table A.8. Exit Survey Reports of Village Prevailing Wage

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

Notes: This table presents statistics on the accuracy of the worker census participant's *ex-ante* report of the prevailing wage, relative to respondents' *ex-post* reports in the 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 worker census participant 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 worker census participant'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. All specifications include year $\times$ month fixed effects and task fixed effects. The sample is restricted to all experimental sample workers who responded to our exit survey. Standard errors are clustered at the village level and are reported in parentheses. P-values are reported in brackets.

Table A.9. Robustness of Wage Rigidity Results to Definition of Low Worker Cohesion

	Proxy for Low Worker Cohesion									
	Wage Labor: Caste Herfindahl					*Agri Labor Force: Caste Herfindahl*				
	Below median	Below median	Linear	Linear in ranks	Bottom tercile	Below median	Below median	Linear	Linear in ranks	Bottom tercile
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel A - Dependent variable: Log Agricultural Wage										
Positive shock last year	0.102	0.110	0.0580	0.0601	0.0480	0.0971	0.106	0.0662	0.0565	0.0583
	(0.042)	(0.042)	(0.023)	(0.023)	(0.037)	(0.033)	(0.033)	(0.027)	(0.022)	(0.032)
Positive shock last year	-0.0826	-0.0821	-0.208	-0.112	-0.0319	-0.0899	-0.0898	-0.183	-0.110	-0.0443
x Low worker cohesion	(0.050)	(0.050)	(0.141)	(0.071)	(0.047)	(0.038)	(0.038)	(0.145)	(0.066)	(0.042)
Positive shock this year	0.0800	0.0910	0.0569	0.0639	0.0540	0.0751	0.0870	0.0531	0.0653	0.0503
	(0.038)	(0.038)	(0.021)	(0.024)	(0.063)	(0.039)	(0.039)	(0.028)	(0.022)	(0.054)
Positive shock this year	-0.0242	-0.0273	0.138	-0.00253	0.0195	-0.0181	-0.0221	0.107	-0.0155	-0.000865
x Low worker cohesion	(0.042)	(0.041)	(0.169)	(0.075)	(0.068)	(0.043)	(0.043)	(0.181)	(0.072)	(0.056)
Observations (worker-days)	59243	59243	59243	59243	59243	59243	59243	59243	59243	59243
Panel B - Dependent variable: Agricultural Employment										
Positive shock last year	-0.234	-0.269	-0.133	-0.142	-0.183	-0.172	-0.206	-0.160	-0.143	-0.223
	(0.078)	(0.079)	(0.055)	(0.056)	(0.114)	(0.080)	(0.083)	(0.063)	(0.055)	(0.087)
Positive shock last year	0.189	0.193	0.429	0.387	0.259	0.0716	0.0766	0.358	0.311	0.180
x Low worker cohesion	(0.088)	(0.089)	(0.331)	(0.183)	(0.135)	(0.107)	(0.109)	(0.448)	(0.201)	(0.122)
Positive shock this year	0.133	0.0860	0.158	0.147	0.235	0.131	0.0853	0.127	0.145	0.124
	(0.083)	(0.089)	(0.062)	(0.062)	(0.129)	(0.091)	(0.098)	(0.075)	(0.061)	(0.122)
Positive shock this year	0.0394	0.0420	-0.142	0.135	0.0675	0.0469	0.0503	0.394	0.246	0.137
x Low worker cohesion	(0.114)	(0.115)	(0.496)	(0.231)	(0.162)	(0.123)	(0.123)	(0.600)	(0.243)	(0.162)
Observations (workers)	623861	623861	623861	623861	623861	631909	631909	631909	631909	631909

Notes: This table presents the sensitivity of the wage rigidity tests in Table VII to varying the definition of low social cohesion. Sample, variable definitions, and controls are as in Table VII. Cols. (1) and (6) replicate the results from the main table in the paper. Cols. (2) and (7) add controls for a longer history of lagged positive shocks (2 and 3 years ago) to the main specifications. The definitions of worker cohesion in the remaining columns are as follows. Cols. (3) and (8): linear control for the district Caste Herfindahl index (reversed so that a higher value means less cohesion). Cols. (4) and (9): linear control for the rank of the district's Caste Herfindahl index (reversed so that a higher value means less cohesion). Cols. (5) and (10): An indicator for the worker's district being in the bottom tercile by caste Herfindahl index (compared to omitted category of top tercile). Standard errors are clustered by region-year and reported in parentheses.

Appendix B. Appendix Figures

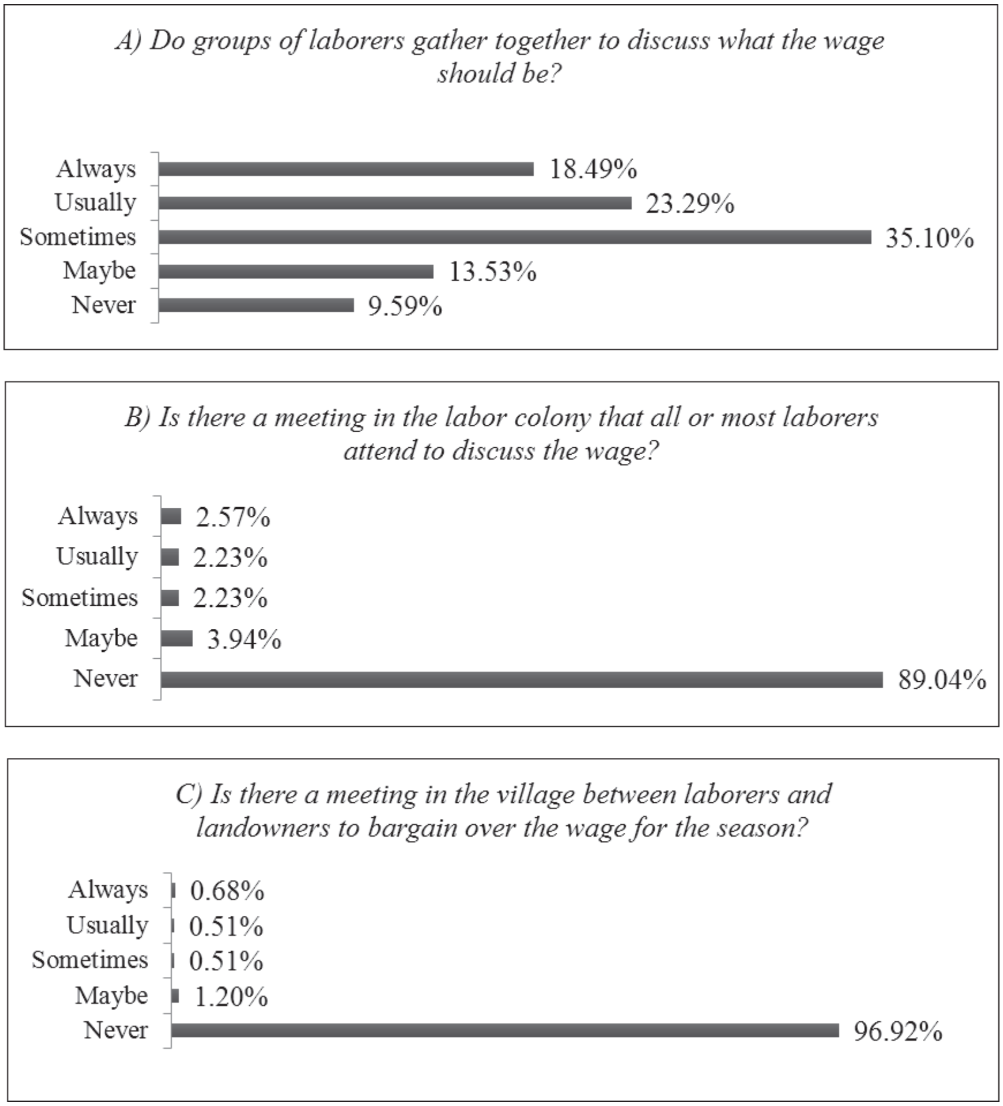

Figure B.1. Collective Action: Wage Setting in the Village. *Notes*: This figure graphs responses to two survey questions about collective action in wage-setting within the village. Data are from the sample of untreated holdout sample households from experiment villages surveyed following the completion of the experiment. Data are from N = 584 male casual workers in 183 villages.

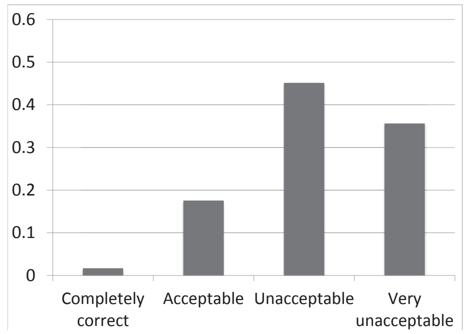

(a) Acceptability of Working for a Wage Cut. *Suppose it is the lean season. The prevailing wage is Rs. 200. To increase his chance of finding work, a worker tells farmers that he would be willing to work any day that week at Rs. 180. Is the worker's behavior acceptable?*

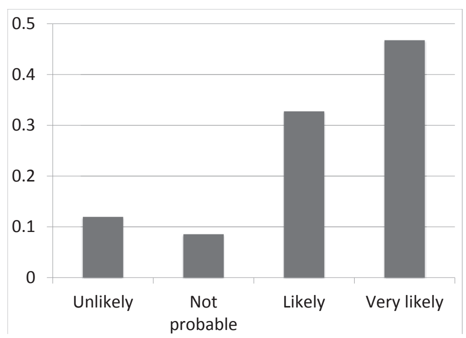

(b) Reactions to Accepting a Wage Cut. *If a worker accepts work at a rate lower than the prevailing wage, how likely is it that the other workers in the village become angry?*

FIGURE B.2. Motivational Evidence - Norms Governing Wage Cuts. *Notes*: These figures graph the exit survey responses from 370 untreated holdout sample 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 workers were offered jobs at the prevailing wage in our experiment.

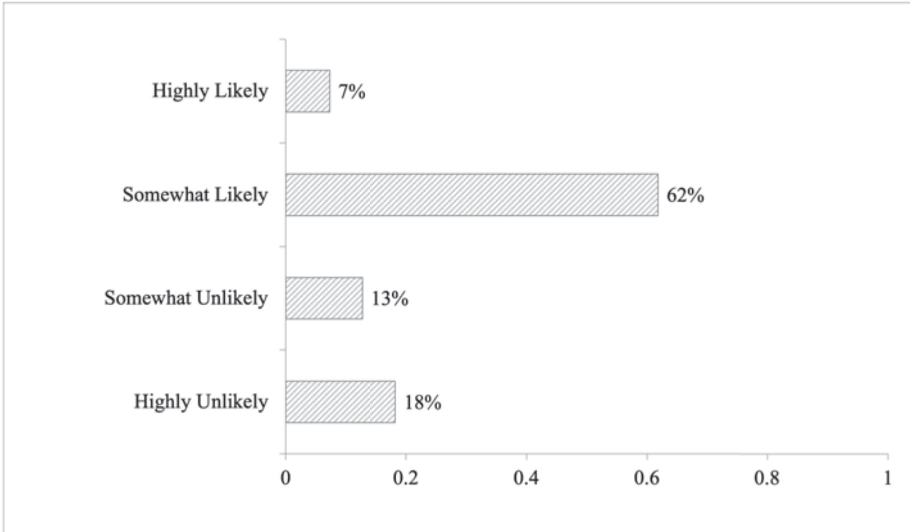

(a) Perceived Wage Consequences from Workers Accepting Wage Cuts. *Suppose a worker accepts a job at* $w - 10\%$*. According to you, how likely is it that the farmers will now try to offer work to other labourers in the villages at* $w - 10\%$*?*

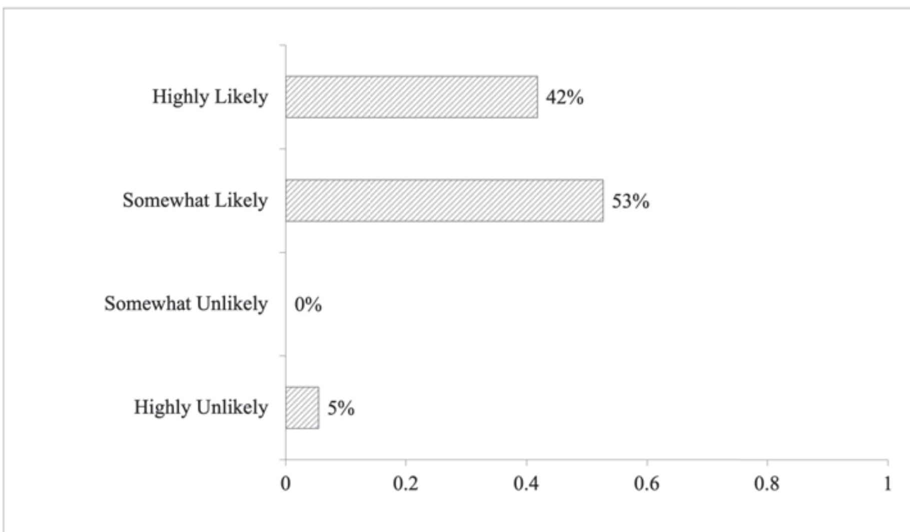

(b) Perceived Wage Consequences from Workers Accepting Wage Premium. *Suppose a worker accepts a job at* $w + 10\%$*. According to you, how likely is it that other laborers in the village would ask for a higher wage than w from other farmers?*

FIGURE B.3. Perceived Consequences for Wage Adjustment: Wage Cuts and Premia. *Notes*: Figure displays survey results from 110 workers in 10 non-study villages in study districts.

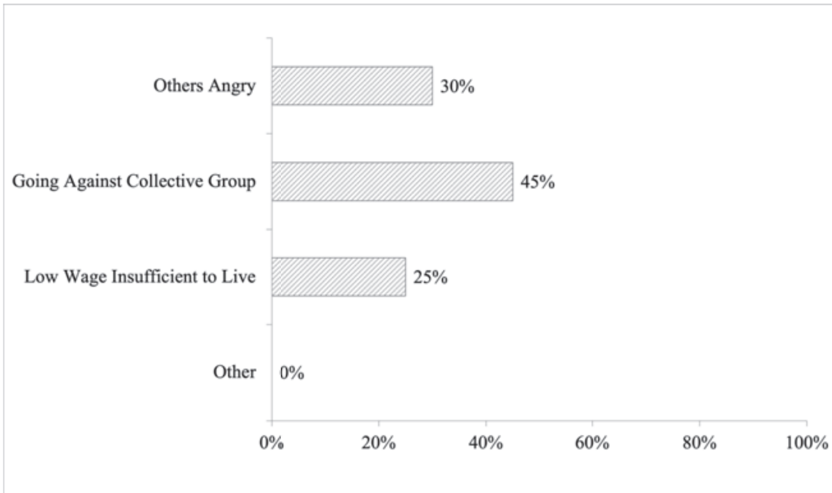

Figure B.4. Most Important Reason for Rejecting a Wage Cut. *Notes*: This figure graphs responses to a survey question eliciting the most important reason why a worker might reject a job that pays Rs. 20 below the prevailing wage. Data from from N = 110 male casual workers in 10 non-study villages in study districts.

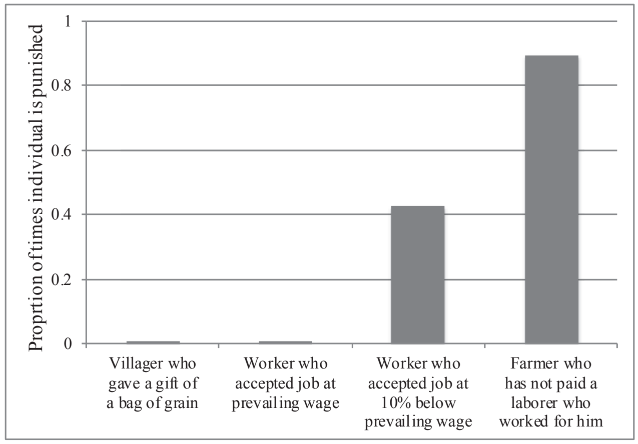

(a) Punishment across Scenarios

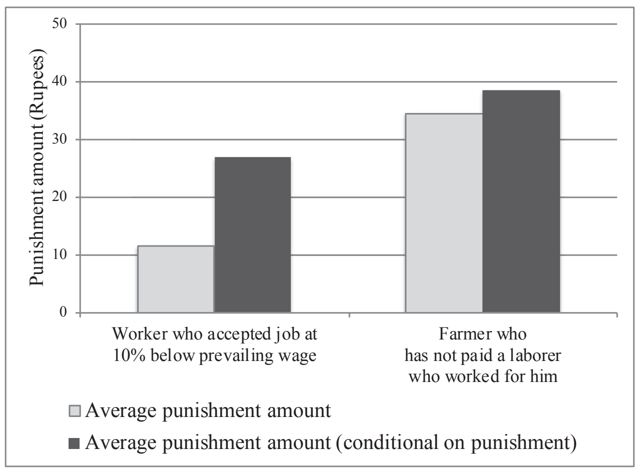

(b) Punishment Amounts

Figure B.5. Laboratory Games: All Scenarios. *Notes*: This figure shows responses in the costly punishment game 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.

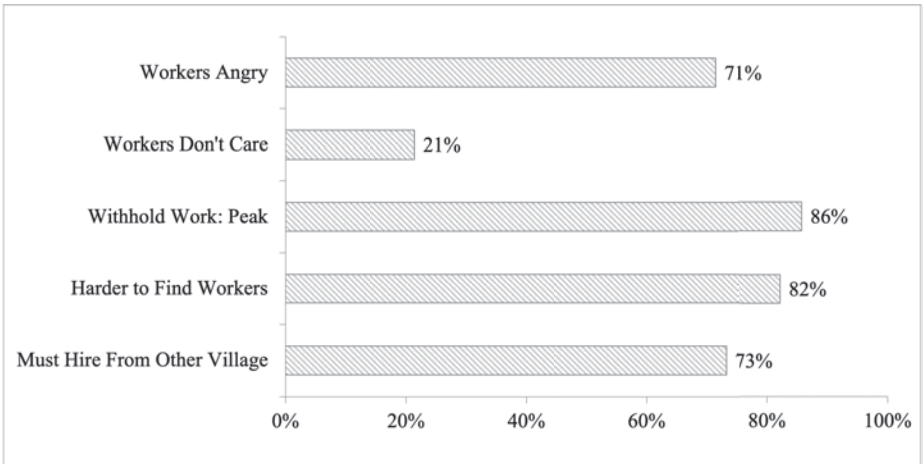

(a) Employer Surveys.

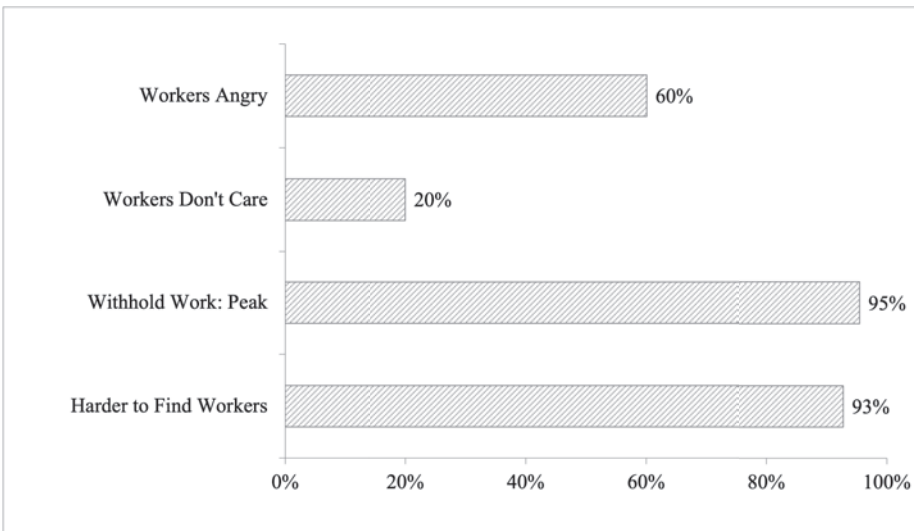

(b) Worker Surveys.

FIGURE B.6. Consequences to Employers from Offering Jobs Below Prevailing Wage. *Notes*: Figure displays survey results pertaining to perceived consequences to employers who offer work at Rs. 20 below the prevailing wage. Panel A presents responses from the perspective of 56 employers. Panel B presents responses form the perspective of 110 workers. Employer and worker surveys were conducted in 10 non-study villages in study districts.

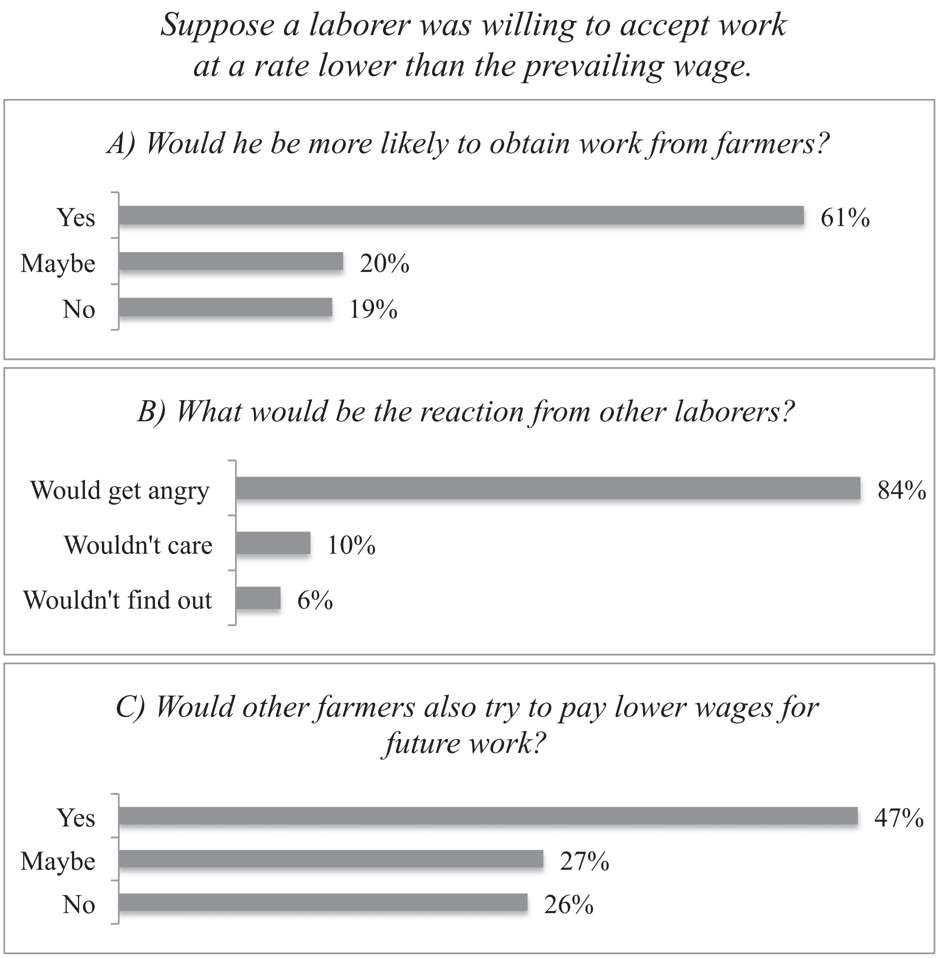

Figure B.7. Worker Beliefs: Impacts of Accepting Wage Cuts. *Notes*: This figure graphs responses to three survey questions about employer and worker responses to a worker accepting a job at below the prevailing wage. Data from Kaur (2019), from N = 196 male casual workers in 34 villages across 6 districts in the states of Odisha and Madhya Pradesh.

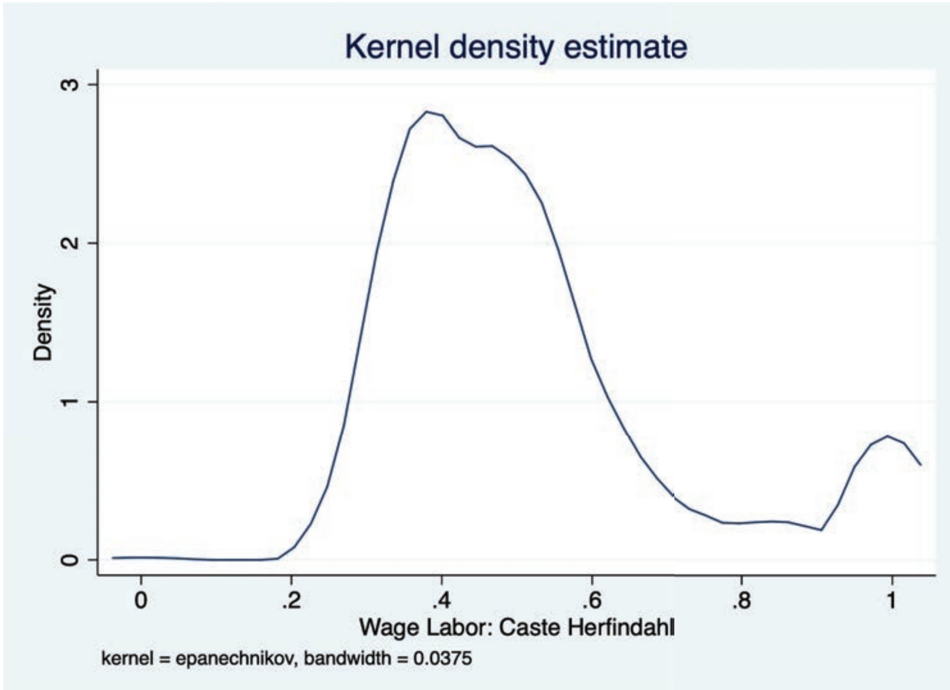

(a) Positive Wage Labor Force: Caste Herfindahl

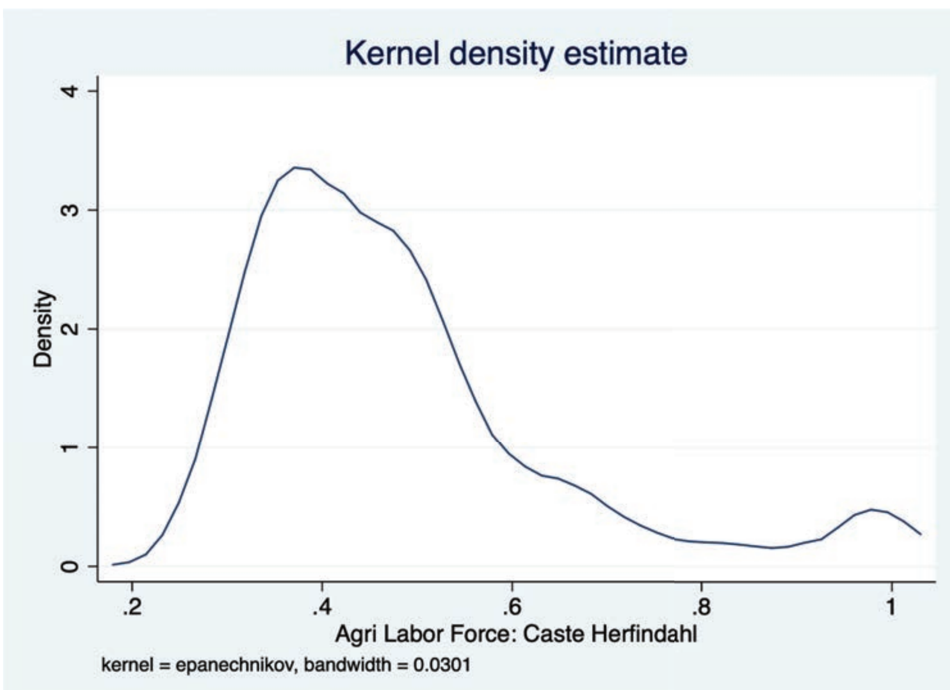

(b) Agri Wage Labor Force: Caste Herfindahl

Figure B.8. Kernel Density of Caste Herfindahl Measures. *Notes*: This figure shows the distribution of the Herfindahl index of caste across districts for a) all workers who report any daily agricultural labor for a wage, and b) for all workers who report agricultural labor as their primary or secondary occupation. Data are from the National Sample Survey data (1986-2007).

Appendix C. Hiring Script

The following is representative of the hiring script that was used across all experimental rounds:

Employer: I have some work on my field for which I'd like to hire you for one day. The work is *[insert specific task]* on my field for one day on *[insert the date for the work]*.

These people here with me are from a research institute in Bhubaneswar. They would like to do a very quick survey with you to understand agricultural practices better. Can you please take a minute to chat with them — they will give you additional details. Anything you would like to discuss with them will remain confidential and will not be shared with me. After chatting with one of them you can let me know if you are available to work for me. *[Employer indicates for the enumerator to take over the dialogue.]*

Protocol based on observability condition:

- *Public condition: All parties remain in street outside worker's home*
- *Employer only condition: Employer, worker, and one enumerator step into premises of worker's house or other private space out of earshot.*
- *Private condition: Worker and one enumerator step into premises of worker's house or other private space out of earshot.*

Enumerator: Before the survey, I would like to clarify some of the details of the job offer the employer just gave you. You will get *[insert wage here]* for this work. The employer needs three people for this task on *[employer name]*'s field on *[scheduled day]*. This job is being offered to you and two others in this hamlet for this task. Because we will need to return on *[insert the date for the work]* to do some more surveys, we can also deliver your payment to you on behalf of the employer at the end of the day. This in no way affects the work itself — you would work on the employer's land doing *[insert specific task]* as usual, and the employer will supervise you as usual. We will not be there except to save the employer some time by handing you your payment at the end of the day. So, you should consider this as you would a normal job for the employer, and make your decision accordingly.

(Private condition only:) We will not inform anyone else how much you have been offered for this job. Only you will know your wage (the employer will not know it). Remember that we will hand you your payment at the end of the workday, so the employer will not see it.

Do you understand? Do you have any questions? Whether or not you want to take this job is completely up to you.

Employer rejoins the conversation

Employer: Would you like to take the job?

If laborer accepts: Employer: Be ready to work at [insert time]AM in the morning i.e. half an hour before the work starts. The work will happen at *[insert location of the employer's farm/field]*.

Employer steps away to give worker "privacy" to answer "survey" questions.

Enumerator: We have a few questions about agriculture to ask you related to our study. Would you be able to answer them?
1. For one acre land, what is the optimal number of laborers required for land levelling?
2. Post crop cutting, what is the total number of paddy straw bundles one laborer can get in one day?
3. How many paddy shoot bundles can one laborer pick in one day (contractual piece rate labor)?
4. When you go for daily wage labor, does the employer provide you with the farming tools or do you carry them along?
Thank you for your time.

[For enumerator: Afterwards, record on a scale of 1-3 whether the person seemed engaged and knowledgeable in answering questions, with 1=not engaged/knowledgeable, and 3=very engaged/knowledgeable.]

Appendix D. Worker Surplus Calculations

Figure 1(a) illustrates the distortionless competitive equilibrium (L^*, W^*), and Figure 1(b) illustrates the distorted wage floor equilibrium (L^F, W^F). Workers are better off under a wage floor if the increase in average wages for those who are employed is large enough to offset the portion of the worker surplus that becomes deadweight loss.

We proceed by estimating the demand and supply curves and assume that both are locally linear. To estimate the labor supply curve, we use the data directly from our field experiment, namely the levels of take-up and wages from the private $w-10\%$ and the pooled w treatments. To estimate labor demand, we observe that the equilibrium level of employment under a wage floor (L^F, W^F) is determined by the demand curve and estimate those quantities using the employment levels and earned wages reported in our untreated holdout sample. We also use the labor demand elasticity estimated in **?**.

We present our calculations in Table D.1. We find that the counterfactual equilibrium wage in the absence of distortions is 7% lower than the observed wage, and employment is 7% higher. Under the above assumptions, workers benefit from the wage floor: workers' surplus is 64% higher relative to the competitive equilibrium. We should also note that 96% of the gains to the workers come at the expense of employer surplus and only 4% from deadweight loss. Our calculations, albeit crude, indicate that the ability to set a wage floor helps workers extract more total surplus for themselves. Moreover, because our setting is a daily labor market—with different workers obtaining jobs on different days—these gains are likely, at least to some extent, to be shared across workers.

We then examine our key assumptions. The labor supply elasticity estimate of 3.89 from our experiment is a *residual* labor supply elasticity from workers who were unemployed at the time of the job offer. Thus, this elasticity is an upper bound on the true labor supply elasticity, which is consistent with its large magnitude relative to other estimates in similar contexts Rosenzweig (1978); Bardhan (1979); Abdulai and Delgado (1999); Goldberg (2016)e.g.. Similarly, the labor demand elasticity measure is based on data from the peak agricultural season, and may be inelastic relative to the true labor demand elasticity in our lean season experiment. We analyze the sensitivity of our main estimates to a range of possible elasticities between 0.5 and 4 (Appendix Table D.2). Within these ranges, there are gains to employer surplus from the wage floor under efficient rationing of between 8% and 98%.[1] As in Lee and Saez (2012), workers gain more when labor supply is relatively elastic and labor demand is relatively inelastic.

[1] With the exception of the case in which labor demand is extremely elastic ($\eta = 4$) and labor supply is extremely inelastic ($\nu = 0.5$).

Next, we relax our assumption that rationing is efficient. We present a version of our results in which workers with the highest surplus (lowest reservation wages) from the job are rationed first under the wage floor, which we term "inefficient rationing". Logically, this implies a greater loss in worker surplus from the rationing of limited jobs than in the base case. We illustrate this scenario in Appendix Figure 2(b), and provide estimates in Appendix Table D.3. Even under this conservative assumption, worker surplus increases up to 72% when labor demand is relatively inelastic ($\eta < 2$).

Finally, we consider how our surplus calculation would change under the case of monopsony wage-setting by employers. We again focus on efficient rationing. Figure 3(a) details the canonical monopsonist's wage-setting problem. Unlike in the competitive case, when the monopsonist raises the wage to attract a marginal worker, he also internalizes the increase in average wages. So, the monopsonist chooses L^M such that the marginal cost of labor equals the marginal revenue product of labor (aka, the demand curve). The wage W^M is set at the minimum wage required to attract L^M workers. In this case, the wage and employment levels are both set below the market clearing wage and employment levels. Worker surplus is lower under monopsony, and there is also deadweight loss.

Now suppose that the workers are able to impose a wage floor $W^F > W^M$. Denote $\bar{W}$ as the wage where the marginal cost curve intersects the demand curve. If $W^F \leq \bar{W}$ (as in Figures 3(b) and 3(c), the monopsonist has the incentive to hire all the way until $L = S(W^F)$. Hiring any less would mean that he could earn positive profits on the marginal worker, as hiring one more worker doesn't affect the average wage. In both cases depicted in the figures, $L^F > L^M$, $W^F > W^M$, and deadweight loss is reduced. Thus, workers are strictly better off, and the resulting equilibrium is more efficient. If the wage floor is set higher, where $W^F > \bar{W}$ as depicted in Figure 3(d), the case looks more similar to the competitive equilibrium, where $L^F < L^M$, $W^F > W^M$, and deadweight loss increases.

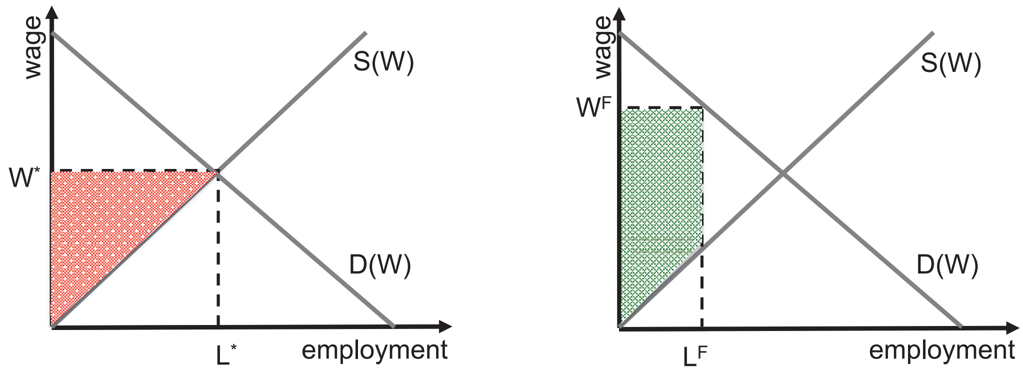

(a) Competitive Equilibrium (b) Wage Floor, Efficient Rationing

Figure D.1. Equilibrium wages and employment under the competitive equilibrium and under efficient rationing. Note: The shaded regions of each panel represent worker surplus. Panel A illustrates the competitive equilibrium with wage W^*. Panel B illustrates the case of a wage floor with efficient rationing.

Table D.1. Surplus Estimation

Labor Supply - W	0.29
Labor Supply - W-10%	0.17
Labor Demand - W (Breza et al. [2019])	0.20
Labor Supply Elasticity	3.89
Labor Demand Elasticity (Kaur [2019])	1
L*	0.21
W*	186.90
Percent Difference between Prevailing Wage W and W*	7.01%
Percent Difference between Labor Supply at W and (L*,W*)	-6.55%
Worker Surplus, Competitive Equilibrium	4.08
Worker Surplus, Wage Floor	6.69
Increase in WS compared to Comp Eq.	2.61
Percent Increase in WS compared to Comp Eq.	63.81%
Employer Surplus, Competitive Equillibrium	22.71
Deadweight Loss	0.10

Notes: This table presents estimates of worker surplus under wage floors relative to the (estimated) competitive equilibrium. The level of labor supply at various wage rates is derived from take-up rates for one day of agricultural work in our main experiment. The labor demand elasticity measure is derived in **?** using National Sample Survey data. The prevailing wage, W, is taken to be Rs. 200, the modal prevailing wage among villages in our experiment. The level of labor demand at the prevailing wage is based on survey evidence collected as part of a labor market rationing experiment by **?**, set in 64 other villages in the same districts of Odisha as this experiment.

Table D.2. Surplus Estimation Sensitivity Analysis

Labor Supply in the Competitive Equilibrium, L*

LS Elasticity (Rows)/LD Elasticity(Columns)	**0.5**	**1**	**2**	**4**
0.5	0.24	0.26	0.28	0.30
1	0.23	0.24	0.26	0.28
2	0.22	0.23	0.24	0.26
4	0.21	0.22	0.23	0.24

Wage in the Competitive Equilibrium, W*

LS Elasticity (Rows)/LD Elasticity(Columns)	**0.5**	**1**	**2**	**4**
0.5	110.26	135.34	158.52	175.84
1	144.33	155.13	167.67	179.26
2	168.36	172.17	177.56	183.83
4	183.02	184.18	186.08	188.78

Total Surplus in Competitive Equilibrium

LS Elasticity (Rows)/LD Elasticity(Columns)	**0.5**	**1**	**2**	**4**
0.5	97.95	79.40	70.74	66.74
1	68.35	48.97	39.70	35.37
2	53.95	34.17	24.49	19.85
4	46.91	26.98	17.09	12.24

Deadweight Loss in Surplus Under Wage Floor Equilibrium, W

LS Elasticity (Rows)/LD Elasticity(Columns)	**0.5**	**1**	**2**	**4**
0.5	3.29	4.74	6.08	7.08
1	1.02	1.64	2.37	3.04
2	0.29	0.51	0.82	1.18
4	0.08	0.14	0.26	0.41

Percent Increase in Worker Surplus Under Wage Floor Equilibrium, W

LS Elasticity (Rows)/LD Elasticity(Columns)	**0.5**	**1**	**2**	**4**
0.5	43.89%	23.17%	7.76%	-1.94%
1	66.22%	43.89%	23.17%	7.76%
2	85.24%	66.22%	43.89%	23.17%
4	98.49%	85.24%	66.22%	43.89%

Notes: This table presents various outcomes of interest in our surplus calculations for a range of labor demand and supply elasticities. We consider labor supply and the wage under the competitive equilibrium, deadweight loss under the wage floor equilibrium, and worker and employer surplus under both equilibria. The level of labor demand at the prevailing wage is based on survey evidence collected as part of a labor market rationing experiment by **?**, set in 64 other villages in the same districts of Odisha as this experiment.

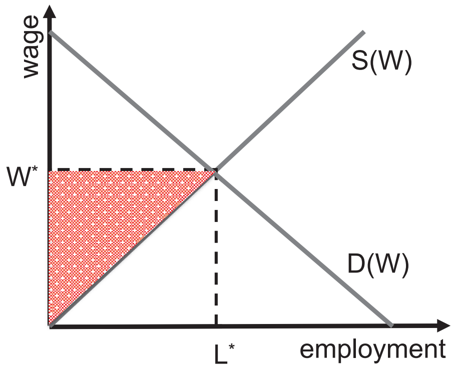

(a) Competitive Equilibrium

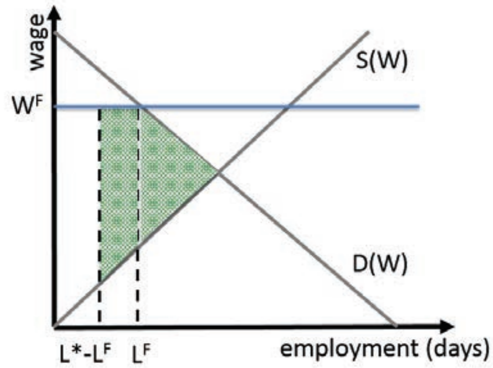

(b) Wage Floor, Inefficient Rationing

Figure D.2. Equilibrium wages and employment under the competitive equilibrium and under inefficient rationing. Note: The shaded regions of each panel represent worker surplus. Panel A illustrates the competitive equilibrium with wage W^*. Panel B illustrates the case of a wage floor with inefficient rationing.

TABLE D.3. Surplus Estimation Sensitivity Under Inefficient Rationing

Labor Supply in the Competitive Equilibrium, L* (Identical to Efficient Rationing)

LS Elasticity (Rows)/LD Elasticity(Columns)	**0.5**	**1**	**2**	**4**
0.5	0.24	0.26	0.28	0.30
1	0.23	0.24	0.26	0.28
2	0.22	0.23	0.24	0.26
4	0.21	0.22	0.23	0.24

Wage in the Competitive Equilibrium, W* (Identical to Efficient Rationing)

LS Elasticity (Rows)/LD Elasticity(Columns)	**0.5**	**1**	**2**	**4**
0.5	110.26	135.34	158.52	175.84
1	144.33	155.13	167.67	179.26
2	168.36	172.17	177.56	183.83
4	183.02	184.18	186.08	188.78

Total Surplus in Competitive Equilibrium (Identical to Efficient Rationing)

LS Elasticity (Rows)/LD Elasticity(Columns)	**0.5**	**1**	**2**	**4**
0.5	97.95	79.40	70.74	66.74
1	68.35	48.97	39.70	35.37
2	53.95	34.17	24.49	19.85
4	46.91	26.98	17.09	12.24

Deadweight Loss in Surplus Under Wage Floor Equilibrium, W, and Inefficient Rationing

LS Elasticity (Rows)/LD Elasticity(Columns)	**0.5**	**1**	**2**	**4**
0.5	10.17	14.66	18.81	21.91
1	4.55	7.33	10.56	13.55
2	2.08	3.66	5.91	8.51
4	0.98	1.83	3.22	5.20

Percent Increase in Worker Surplus Under Wage Floor Equilibrium, W

LS Elasticity (Rows)/LD Elasticity(Columns)	**0.5**	**1**	**2**	**4**
0.5	25.77%	0.82%	-17.33%	-28.54%
1	44.77%	13.96%	-13.75%	-33.69%
2	60.93%	27.84%	-9.67%	-42.90%
4	72.18%	39.49%	-6.01%	-56.92%

Notes: This table presents various outcomes of interest in our surplus calculations for a range of labor demand and supply elasticities under inefficient rationing. The outcomes of interest include labor supply and wage under the competitive equilibrium, deadweight loss under the wage floor equilibrium, and worker and employer surplus under both equilibria. Here, we assume "inefficient rationing" of jobs in the following sense: workers with the highest potential surplus from the jobs are rationed out of the market in the wage floor equilibrium.

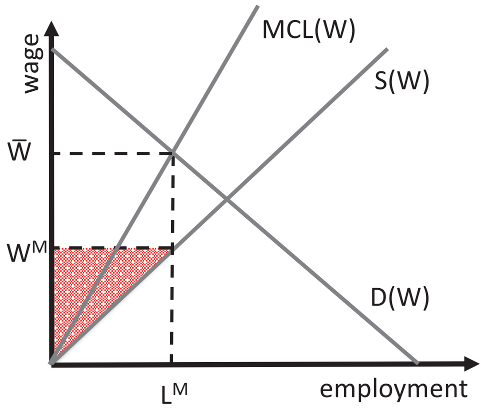

(a) Equilibrium Under Monopsony Wage-Setting

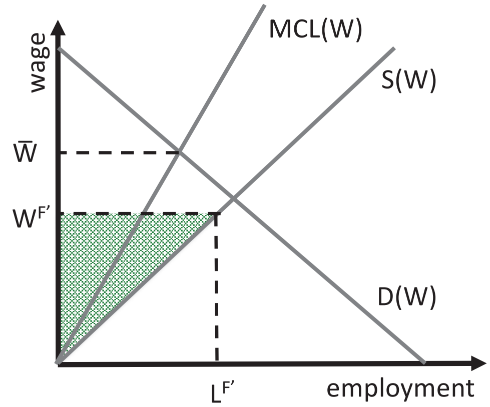

(b) Wage Floor, Case 1: $W^{F'} < W^{*}$, Efficiency Gain

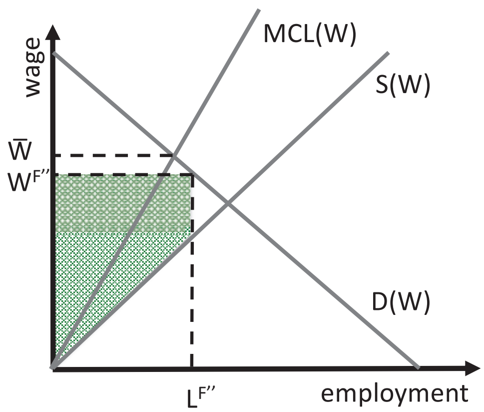

(c) Wage Floor, Case 2: $W^{F''} > W^{*}$, Efficiency Gain

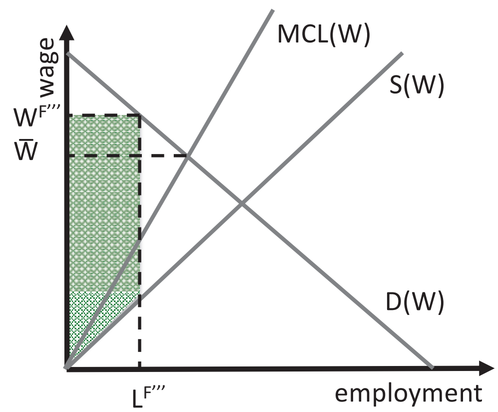

(d) Wage Floor, Case 3: $W^{F'''} > W^{*}$, Efficiency Loss

Figure D.3. Equilibrium wages and employment under the competitive equilibrium and under inefficient rationing. Note: The shaded regions of each panel represent worker surplus. Panel A illustrates the monopsony equilibrium with wage W^{M}. Panel B illustrates the case where the wage floor is set below the competitive equilibrium wage, but higher than the monopsony wage $W^{M} < W^{F'} < W^{*}$, generating an increase in worker surplus and efficiency relative to Panel A. Panel C illustrates a case where the wage floor overshoots the competitive equilibrium wage $W^{M} < W^{*} < W^{F''}$, but still generates a gain in worker surplus and efficiency. In Panel D, the wage floor overshoots the competitive equilibrium by so much that deadweight loss actually decreases relative to Panel A.