

NBER WORKING PAPER SERIES

PROPPING UP THE WAGE FLOOR:
COLLECTIVE LABOR SUPPLY WITHOUT UNIONS

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

We thank Siwan Anderson, Doug Bernheim, Lucas Coffman, Stefano DellaVigna, Henry Farber, Andrew Foster, Patrick Francois, Pat Kline, and Michael Kremer for helpful comments and conversations. Shreoshee Mukherjee, Arnesh Chowdhury, Piyush Tank, Medha Aurora, and Sayan Kundu provided terrific research 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. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2019 by Emily Breza, Supreet Kaur, and Nandita Krishnaswamy. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Propping Up the Wage Floor: Collective Labor Supply without Unions
Emily Breza, Supreet Kaur, and Nandita Krishnaswamy
NBER Working Paper No. 25880
May 2019, Revised August 2020
JEL No. D71,E24,J22,J31,J43,J50,O15,O17

ABSTRACT

Social norms have the potential to alter the functioning of economic markets. We test whether norms shape the aggregate labor supply curve by leading decentralized individuals to maintain wage floors in their local labor markets. We partner with existing employers who create new jobs for workers in informal spot labor markets. Unemployed workers would like to find work, and prefer to do so even at wages below the prevailing wage rather than remain unemployed. However, they rarely do so when this choice is observable to other workers. In contrast, social observability does not affect labor supply at the prevailing wage. Consistent with the idea that norms could have aggregate implications, measures of social cohesion correlate with downward wage rigidity and business cycle volatility across India.

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

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

Supreet Kaur
Department of Economics
University of California, Berkeley
Evans Hall
Berkeley, CA 94720
and NBER
supreet@berkeley.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.socialsciscenceregistry.org/trials/1290/history/>

I Introduction

We study the potential for social norms to shape aggregate behavior in economic markets. Broad literatures in economics, psychology, political science, and sociology document that norms embody a shared understanding of acceptable actions (e.g. Fehr and Schurtenberger [2018]; Cialdini and Goldstein [2004]; Crawford and Ostrom [1995]; Bendor and Swistak [2001]). This implies that norms can serve as an equilibrium coordination device—enabling collective behaviors among a large decentralized group of individuals, even in the absence of any explicit coordination or formal institutions.

We examine this idea within the context of the labor market. Specifically, we document the presence of implicit collusion among workers to maintain wage floors in their local labor markets—altering aggregate labor supply without any formal labor organization. Consistent with theories of how norms are sustained in equilibrium, we provide positive evidence for the role of social sanctions, which are used to deter self-interested deviations by workers.¹

This mechanism has potential applicability in any setting with meaningful social interaction. This makes the labor market a particularly relevant domain: most labor arrangements involve repeated inter-personal interaction among workers. Consistent with this, collective behaviors—mass walkouts with foregone wages, punishment of rate busters, and productivity compression—have been documented across history and contexts, even without formal unions. More broadly, the presence of cartel-like behavior among decentralized individuals has been qualitatively observed in a range of markets, from NASDAQ traders and real estate agents in the US (Christie and Schultz [1994]; Hsieh and Moretti [2003]), to local traders and contractors in poor countries (Bergquist [2018]; Houde et al. [2018]). We posit that social norms can sustain such behavior—enabling agents to maintain market power in equilibrium. Moreover, such forces may be particularly relevant in poor countries, where the communal nature of village economies and high levels of social capital may make it more likely that norms arise and can be sustained effectively through the threat of punishment.

We test whether norms shape aggregate labor supply in informal markets for casual daily labor in India. Such markets are ubiquitous in poor countries, serving as an employment channel for hundreds of million workers in India alone (National Sample Survey [2010]). Casual labor markets are decentralized and informal: daily contracts are arranged bilaterally between individual workers and employers, and there are no unions or formal modes of labor organization (e.g. Rosenzweig [1988]). Despite this, they exhibit features consistent with potential distortions, such as downwardly rigid wages and high unemployment rates in the

¹A nearly universal observation in social psychology is that individuals engage in punishment or social disapprobation when they see norm violations (e.g. Cialdini and Goldstein [2004]; Henrich et al. [2006]). This is consistent with a theoretical literature in economics that models norms as a coordination device, sustained through social punishment (Kandori [1992]; Ellison [1994]; Jackson, Rodriguez-Barraquer and Tan [2012]).

agricultural lean season (e.g. Dreze and Mukherjee [1989]; Kaur [2019]; Breza, Kaur and Shamdasani [2019]).

We hypothesize that, in our setting, there is a norm against accepting jobs below the prevailing wage, and this norm is consequential for workers' labor supply during times of unemployment. We provide evidence from two experimental exercises conducted during the agricultural lean season. First, we use a field experiment to document that although underlying labor supply is positive below the prevailing wage, it is suppressed due to extrinsic concerns: workers will not accept jobs below the prevailing wage because they do not want others to learn they have done so. Second, we provide evidence for sanctions: workers are willing to pay to punish those who have accepted wage cuts.

For the field experiment, we partner with existing agricultural employers, who make job offers to workers in their respective local labor markets. Employers follow the typical process for agricultural hiring: offering jobs by approaching workers at their homes. In this setting, there is a commonly known prevailing daily wage for each type of agricultural task in the village—providing a clear benchmark wage for job offers.²

To test for effects on labor supply, we induce two types of cross-cutting variation during employers' hiring process. (i) First, the job is offered at a random wage level: at the prevailing wage, or 10% below the prevailing wage. (ii) Second, we vary the extent to which the wage level is publicly observable: whether the job is offered inside the worker's home or outside on the street where neighbors (who are typically other workers) can overhear the offer.³

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. Under our hypothesis, observability will only matter under norm violations (i.e. under wage cuts); it should not matter when jobs are offered at the prevailing wage. Finding this prediction in our results helps assuage such concerns.

All offered jobs correspond to actual employment opportunities on the employer's land—so that our data reflect real employment decisions by workers. Treatment randomization is at the village level, so that all workers within a given village receive the same wageXobservability condition. In addition, the workers in our experiment (i.e. those who are offered jobs) are sampled randomly from the village population of workers. Aside from inducing variation in hiring, we are not involved in any other aspect of the employment relationship: employers

²We provide direct evidence for the presence of a prevailing wage that is common knowledge.

³The 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. Condition (iii) enables us to measure full residual labor supply below the prevailing wage. However, in order to understand the effects of other workers observing the wage, we can run analyses comparing conditions (i) and (ii).

supervise workers as usual, provide them food, etc. The experiment is conducted across 183 labor markets (i.e. villages) with 183 partnering employers (one in each village), with jobs extended to 502 workers.

At the prevailing wage, the average take-up rate of jobs is 26%, with no detectable differences across the observability conditions. Given a 43% baseline involuntary unemployment rate, this reflects a robust level of take-up.⁴

In contrast, when a worker is offered a job below the prevailing wage, take-up depends crucially on whether his decision is publicly observable to other workers. When a wage cut is offered in private, take-up remains a robust 18%.⁵ However, the willingness to accept a wage cut falls by 13.6 percentage points (78%) when offers are observable to other laborers (significant at 1% level). When restricting the sample to workers who are actually in the agricultural labor force—those who consider agricultural labor as their primary or secondary occupation—these results become even starker. Only 1.8% of agricultural workers accept wage cuts in public (relative to an estimated 26% in private).

Overall, this pattern indicates there is substantial underlying labor supply below the prevailing wage, but workers will not accept such jobs due to social concerns. This distortion on individual labor supply is economically meaningful. Workers forego 26-49% of weekly agricultural wage earnings to avoid being seen as breaking the social norm.⁶

We provide three supplemental exercises to further characterize our results. First, in places with more involuntary unemployment, one would expect more scope for residual labor supply below the prevailing wage. Consistent with this, a larger share of workers accept jobs with wage cuts in private when unemployment is higher. However, workers rarely accept wage cuts in public regardless of the unemployment level. In fact, workers are less likely to accept public wage cuts when unemployment is higher than when it is lower (difference not statistically significant). This indicates that norms are sufficiently strong to deter self-interested deviations even when workers are desperate for jobs. In addition, while only speculative, these findings are potentially consistent with the idea that places with stronger norms may themselves generate more involuntary unemployment.

Second, we document that social observability is more consequential when information about job offers is more likely to be transmitted to others: the decline in take-up of wage

⁴Even under the prevailing wage, we would not expect take-up to be 100%. Because we sampled randomly from the labor force in each village, we did not prescreen on any characteristics. Workers may decline the job because they already found other work, have a home production activity, or because their reservation wage is higher than the prevailing wage.

⁵This is a residual labor supply estimate (the fraction of workers who will take up a wage cut conditional on not having found a job elsewhere), and is therefore an upper bound on the fraction of workers for whom the wage cut is above their reservation wage.

⁶These estimates are based on exit surveys conducted 1 week after the day of work, in which we ask workers to recount the employment on each day over the past week.

cuts in public is especially concentrated in such villages. This is consistent with previous work showing that network structure is central to spreading information and reputation (e.g. Alatas et al. [2016]; Breza and Chandrasekhar [2019]). Third, we provide evidence against an employer bargaining interpretation of our results. For this, we exploit both variation in our treatment design and also the worker’s employment history; for example, the pattern of our findings holds regardless of whether the worker is likely to work for the employer again in the future (based on their recurring pre-existing relationship).

Our core results establish that while social observability has no discernible effect on labor supply under the prevailing wage, it drastically reduces labor supply under wage cuts. We posit these observability effects stem from an attempt to avoid social disapprobation. However, other reputational mechanisms could produce this pattern—for example, shame from appearing financially desperate enough to accept a wage cut.

In the second part of the paper, we provide positive evidence for social sanctions. We first use a survey to document that 80% of workers believe that accepting a job below the prevailing wage will result in sanctions from others. These range from a decrease in referrals to exclusion from social activities like drinking after work.

To provide revealed preference evidence on sanctions, we design a costly punishment game with a random subset of workers (“players”) who were *not* offered jobs. Each player is paired with an anonymous worker. Using a 2x2 design, the player is told whether his paired worker: (i) lives in the player’s own village or a village that is far away, and (ii) accepted a job at the prevailing wage or at 10% below the prevailing wage. The player can give up some of his endowment to reduce the endowment of his paired worker, providing a mechanism to punish the worker at a cost to himself.⁷

As expected, there is no punishment of workers who accept jobs at the prevailing wage. In contrast, those who accepted a wage cut are punished 37% of the time. Conditional on punishment, the magnitude corresponds to 37% of average daily labor market earnings (requiring players to give up 7.4% of typical daily earnings in order to undertake the sanction). 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]). Moreover, these patterns are similar if the paired worker is from a distant village—even though that worker’s action has no scope to affect the player’s own labor market. This is particularly consistent with individuals viewing norm violations in moral or general terms.

If social norms enable workers to resist downward competitive pressures on wages, this

⁷Note that we never use the term “prevailing wage” to describe the worker’s actions; the player is simply told the worker accepted a job at wage X (e.g. Rs. 200). In addition, to obfuscate the reason for the exercise, players play a total of 4 rounds of the game - 2 with the wage scenarios and 2 non-wage scenarios that depict a positive, neutral, or negative action taken by one’s partner in another non-labor domain.

could help us understand equilibrium wage adjustment. In the third part of the paper, we undertake a supplementary suggestive exercise using observational data across India. Specifically, our results indicate that the lack of labor supply below the prevailing wage stems from social pressure, and is maintained through the threat of social sanctions. A potential implication of this mechanism is that wage floors may be more likely to arise in places with higher social cohesion because there is potentially more scope for norm enforcement—for example, information flows better through the network (making violations more public) and it is easier to levy sanctions (Jackson et al. [2012]; Breza and Chandrasekhar [2019]). We exploit the fact that in India, caste homogeneity is a strong proxy for social cohesion (Munshi and Rosenzweig [2006]; Mazzocco and Saini [2012]; Munshi and Rosenzweig [2016]).⁸

Using the wage rigidity test from Kaur (2019), we examine whether districts with higher cohesion—proxied by caste homogeneity among agricultural workers—exhibit more rigidities. We find that areas with higher social cohesion exhibit substantively greater downward wage rigidity in response to labor demand shocks. Such areas also have corresponding 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 social conformity can have aggregate implications for markets. Of course, this correlation with wage rigidity is only suggestive—it does not necessarily denote a causal relationship.

Finally, despite such unemployment effects, wage floors could benefit workers by increasing the surplus captured from employers. A simple back of the envelope exercise suggests workers likely benefit from implicit collusion: worker surplus is 64% higher relative to the competitive equilibrium. However, we cannot make any direct claims about efficiency. For example, worker collusion could increase efficiency if employers have monopsony power.

Our findings advance the literature on norms and social influence. First, we provide evidence for the idea that norms can shape the equilibrium in economic markets—in our case, by altering the aggregate labor supply curve—by enabling coordination on the same action among decentralized agents. Second, we offer field evidence for social sanctions as a force that helps sustain norms in equilibrium. This augments a rich body of laboratory observations that when individuals are seen as violating a norm, others will sanction them—even in the absence of direct benefit of instrumental value to themselves (e.g. Gächter and Fehr [1999]; Charness and Rabin [2002]; Henrich et al. [2006]). This helps us understand why norms can be long-lived in societies, and accords with economic theories of how decentralized group behavior can be sustained (Kandori [1992]; Ellison [1994]; MacLeod [2007]). Third, our work relates to the social observability literature. A growing number of stud-

⁸While in our field experiment, workers within a village are fully homogenous by caste, this is not true for India as a whole.

ies document that individuals alter their behavior when it is observable to others in a range of settings—including charitable giving, voting, education, and effort in the workplace (Bandiera, Barankay and Rasul [2005]; Mas and Moretti [2009]; DellaVigna, List and Malmendier [2012]; Bursztyn and Jensen [2015]; DellaVigna et al. [2016]; Bursztyn and Jensen [2017]; Bursztyn, Fujiwara and Pallais [2017]).⁹ We tie the relevance of social observability explicitly to the presence of norms, and show that an underlying mechanism that drives social conformity (at least in our setting) is fear of sanctions.

Our findings also advance the literature on wage adjustment and labor market distortions in poor countries. Early work in development economics focused heavily on the idea that downwardly rigid wages contribute to high involuntary unemployment levels (e.g. Lewis [1954]). Recent empirical evidence points to the continued relevance of these ideas today (Kaur [2019]; Breza et al. [2019]). However, the presence of rigidities in this setting has been a long-standing puzzle in development. Micro-foundations proposed in the previous literature, such as nutrition efficiency wages, have not withstood empirical scrutiny (Rosenzweig [1988]). To date, there is scant empirical evidence supporting any micro-foundation for why wage floors should arise in this setting. Our study offers the first piece of positive evidence in support of any micro-foundation for this phenomenon.¹⁰

Finally, our study has bearing on the labor literature on unions (see Farber and Saks [1980]; Dickens et al. [2007]). While the prevalence of formal unions has declined over time, our findings indicate that informal versions of these forces have potential relevance for the labor market. As we argue in the Conclusion section, the presence of informal unions in our setting suggests that some of the considerations historically attached to formal unions may apply more broadly in the labor market—e.g., wage rigidity or productivity compression within establishments.

The paper proceeds as follows. In Sections II and III, we describe the setting and experimental design. We detail the data in IV. We present the results of the main field experiment in Section V, and evidence on sanctions in Section VI. Section VII examines the relationship between correlates of social cohesion and wage rigidity levels across India. Section VIII uses a back-of-the-envelope calculation to estimate worker surplus from maintaining wage floors. Section IX concludes with a discussion of the potential generality of our mechanism for other settings.

⁹In related work, Bursztyn, González and Yanagizawa-Drott (2018) document that altering perceptions of norms can affect labor force participation. See Bursztyn and Jensen (2017) for a review of the literature on social observability.

¹⁰Our findings are consistent with Osmani (1990), who offers a model based on informal worker collusion.

II Context

The core field experiment takes place in 183 villages in four districts of rural Odisha, India. 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 tend to be bilaterally arranged by individual employers and are also 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.

During the agricultural lean season, which spans approximately half the year and is when we conduct our field activities, employment rates are low. Workers in our study have a daily employment rate below 50%, and report not being able to find work when they wanted it an average of 12 days the past month. These unemployment rates create the potential for substantial willingness to supply labor below the market wage during the lean season.

In our study areas, the village constitutes a prominent boundary for the labor market. Agricultural employers hire daily-wage workers primarily from within or close to their village. For example, workers report that 70% of worker-days in agriculture are for work 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 employers draw a majority of hired labor. These neighborhoods are densely-populated, and residents tend to share both a common occupation and caste identity: 84% of males engage regularly in daily-wage agricultural labor, and the median number of *jatis* (subcastes) per “labor colony” is one in our sample. Such an environment, where workers are embedded in the same risk-sharing and information networks, may help to enable worker collusion and sanctions.

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 public (in the street or in a central square of the village) or at the worker’s house. In addition to the prevailing wage, the timings of the work and breaks, and the inclusion of any meals are commonly understood (and expected) by the workers.

A clear prevailing wage for agricultural labor exists in our study areas. There is little variation in the agricultural wage paid to workers of the same gender within a village at a given point in time (Figure I(a)). In addition, the vast majority of workers agrees about the

value of the prevailing wage for a given task (Figure I(b)). In our experiment, the prevailing wage serves as a clear benchmark for wage job offers.

We hypothesize that, in our setting, there is a norm against accepting jobs below the prevailing wage. Figure II provides motivational evidence for this idea. About 80% of workers state that it is “unacceptable” or “very unacceptable” for an unemployed worker to offer to work below the prevailing wage. In addition, about 80% of workers also state that workers would become angry at a worker who accepts a job at a rate below the prevailing wage—indicating the potential role of sanctions in deterring deviations.

The above evidence is consistent with a clear decision rule that can be easily followed by workers when making labor supply decisions: do not accept jobs below the prevailing wage. This accords with the observation laid out by Elinor Ostrom that maintaining decentralized group behavior in informal environments requires clear decisions rules (e.g. Ostrom, Walker and Gardner [1992]). While the decision of a clear prevailing wage is specific to certain settings—e.g., casual labor markets in poor countries, hourly or fixed wage workers within establishments in both rich and poor countries—the notion of clear decision rules around which norms can be constructed is more general (see Section IX).

III Research Design

III.A Hypotheses

Denote the prevailing village wage as w . We hypothesize that while there exists underlying latent labor supply below the prevailing wage, it is suppressed due to social norms that are reinforced through sanctions. Specifically, 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 .
- 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, individuals 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]). An important role for social observability would point to the relevance of external forces.

We next turn to our core labor supply experiment, which tests Hypotheses 1 and 2. In Section VI, we present survey evidence and our social punishment lab-in-the-field experiment to investigate Hypothesis 3.

III.B Experimental Design: Labor Supply

We partnered with 183 existing agricultural employers (one per study village), who made job offers to a total of 502 male laborers. In this section, we outline the experimental design; the next section provides more detail on protocols.

Our design builds on the usual procedures for making job offers. The hiring employer approaches the worker at his home in the labor colony to offer him one day of employment, as is typical in this setting. The employers live in the same village, are known to the workers, and hire for familiar standard tasks (weeding, field preparation, etc.). The standard amenities of the job—start and end times, breaks, and meals provided by the employer—are implicitly understood according to convention in the village for a given task, and so are not explicitly discussed; more generally, there is no bargaining, but rather employers make a take-it-or-leave-it offer. The hiring takes place two days before the work day.¹¹

We partner with employers to induce variation in the hiring process. After this, we are in no way involved in the employment relationship: on the day of work, the employer supervises the workers and gives them any meals etc. Employers in our study are compensated with a lump sum payment for working with us, and are blind to treatment status. Workers that receive job offers are selected randomly from the labor force of residents in the labor colony.

Figure III details the cross-cutting 2x3 treatment design. The first dimension of variation, the wage offer, is randomized to be either the prevailing wage (w), or a 10% cut to the prevailing wage ($w - 10\%$). The second dimension of experimental variation is the degree of social observability and takes one of three conditions – Public, Employer Only, and Private (described below). Treatment is randomized at the village level, so that all workers within a given village receive the same wageXobservability condition.

Aside from the above, we hold all aspects of implementation fixed across treatments. In all cases, the employer, accompanied by two JPAL-South Asia enumerators, approaches the pre-selected worker at home. The hiring protocol follows 3 basic steps. (i) The employer describes the job task, location and timings for the work on the specified date. He 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 the speaking with the enumerator. This creates a natural handoff from the employer to the enumerator. (ii) The

¹¹This is the modal gap between hiring and workday for agricultural jobs in this setting.

enumerator conveys the wage level to the worker. (iii) The worker then tells the employer if he wants to accept the job or not.¹² As is common, the hiring process is short, with the entire exchange lasting about five minutes. These protocols are fixed across all treatments. Below, we describe how we embed differences in social observability within this set-up.

In the *Public* condition, the job offer is simply made in the street *in front* of the worker’s home. Therefore, the employer and any neighbors or passers-by can overhear the job offer. Given how the villages in our study area are organized, those who overhear the job offer are almost always other residents of the labor colony. In the *Employer Only* condition, the job offer is made *inside* the worker’s home. In this case, the employer overhears the wage offer, but other members of the community do not.

To estimate the effect of making the $w - 10\%$ wage offer observable to other workers, we can therefore compare take-up in the *Employer Only* treatment to take-up in the *Public* treatment (cell D vs. E in Figure III). However, ex ante, this test might be too strong: the employer is also a member of the village and could potentially spread wage information to other workers. This could lead to depressed take-up relative to the worker’s true reservation wage, biasing the tests of both H1 and H2 toward zero.

Thus, an ideal conceptual test of H1 and H2 is a treatment where only the worker knows the wage offer. This motivates our *Private* condition. This treatment is very similar to the *Employer Only* condition, with the job offer made inside the worker’s home. However, after the employer describes the job and hands things over to the enumerator (i.e. after step (i)), he wanders out of earshot so that he does not hear the wage level conveyed. This handoff appears natural since the employer has said the enumerator wants to do a “survey”. During the enumerator’s exchange with the worker, which only lasts 1-3 minutes, the enumerator assures him that only the worker will know his own wage, and that it will be kept confidential from the employer.

To test H1—a mass of individuals is willing to work below the prevailing wage—we will test whether take-up in the *Private* treatment at wage $w - 10\%$ (cell F) is statistically distinguishable from zero. Note that the take-up in the *Employer Only* treatment (cell E), for reasons described above, likely provides a lower bound on the private willingness to work (H1). In order to test H2—whether social observability decrease the take up of work at $w - 10\%$ —we can test whether take up in *Private-Public* (cell F-D) is positive and statistically distinguishable from zero. Again, *Employer Only-Public* (cell E-D) likely provides a stronger test and a lower bound for the difference in take-up between private and public. As we show below, in practice, our predicted results hold under either of these comparisons.

One natural concern is that the social observability treatments may inadvertently convey other information, changing worker beliefs across treatment cells—confounding their inter-

¹²We include the script used to make job offers in Appendix XIV.

pretation. 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). Finding this prediction in our results helps assuage such concerns. In addition, we can also use a differences-in-differences regression specification to absorb any level shifters across the observability treatments, comparing (F-D)-(C-A) in our test of H2, for example. Finally, in our *Public* and *Employer Only* treatments, the mechanics of the wage offers look almost identical to common practice in the village, and we can use the results from these two observability conditions alone to obtain a lower bound on our hypothesized effects.

III.C Treatment Implementation

Our experiment took place during non-peak (lean) production periods between 2015 and 2017. 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. 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 was between 30 and 100 households (the average labor colony in our sample has 46 households (Table I)). The lower bound ensured that by making 2-3 job offers per village, there was not direct communication with a large fraction of the village. The upper bound on village size ensured a level of comparability across study villages. Finally, there needed to be a clear notion of the prevailing wage at the time we came to the village.¹³

In each study village, we constructed a listing of 20 employers who had regular demand for agricultural labor. We then recruited an employer from the middle of the land size distribution, who had work to be done on his land in the subsequent two weeks. We did not partner with employers who held political positions. Employers were incentivized to participate with a lump sum payment and a wage subsidy. We set the wage subsidy so that the employer’s contribution to the wage was independent of treatment, permitting employers to be blind to the wage in the *Private* conditions. In the *Public* and *Employer Only* conditions, employers only learned the offered wage during the “survey” portion of the job offer (step (ii)).

At the time of the employer listing, enumerators also constructed a listing of households in the labor colony. Using this list, we pre-randomized which households would be approached for the job offer and the order in which they were visited. During hiring, the job was offered

¹³The prevailing wage criterion only disqualified a few villages where the wage was in flux at the start of the new agricultural year (just before the time of monsoon arrival).

to the male household head, else to one prime-age male household member who was home at the time. Door-locked households (where no men were home when we visited) were skipped. In order to minimize the total amount of information injected into any village, the employer made a maximum of 2-3 job offers per village (with the number of offers pre-determined based on task).

To verify that the *Public* treatment was indeed public, one of the JPAL enumerators recorded the number of people within earshot of the participant at the time the offer was made. An average of five individuals observed each job offer in the public treatments.

On the day of work, enumerators conducted a spot check in the morning to verify that the worker was present on the employer’s field, and to collect the employer’s contribution to the workers’ wages. At the end of the scheduled work day, enumerators returned to the field to give each worker his daily wage in an envelope (to keep the amount private from the employer). These procedures were identical across treatment arms.

IV Data

Our primary outcome is whether the approached worker worked for the employer. This information comes directly from the enumerator’s observation of workers in the employer’s fields.

The enumerators also conducted exit surveys after the scheduled work day with all main study participants (the approached workers). The survey contained 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. Surveys were also administered at the same time to a holdout, untreated sample of workers in each study village. Given our randomization procedure, these workers are comparable to the main study participants. Finally, the enumerators conducted exit surveys with the partner employer to record information about any in-kind or cash transfers made to the workers, and the perceived effort of each worker.

IV.A Summary Statistics and Covariate Balance

Balance We present summary statistics in Table I. 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.

Of workers who were 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, with individuals reporting any work on only 29% of days in the prior month. 32% have worked for the partnering employer in the past.

We also report information about village-level labor market conditions using survey responses from the untreated holdout sample. The level of involuntary unemployment is high (43% of days in the past month) and is comparable with the levels reported in by the main study participants. NREGA work was rare in our study villages at the time of our experiment (0.5% of the past ten days).

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. This accords with natural rates of sampling variation in a balance table. Regardless, we show in Appendix Table A.III that results are very similar when including controls for these covariates. Moreover, we show in Sections V.A and V.C 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 I 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.I). 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.

V Evidence: Labor Supply

V.A Take-up of the Job

Our primary tests of Hypotheses 1 and 2 are presented graphically in Figure IV. Panel A of Figure IV 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 II 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 the agricultural workers. Columns (2) and (3) include task fixed effects and year X month fixed effects. All standard errors are clustered at the village-level. The dependent variable is job take-up, measured as an indicator for a worker reporting to the job on the day of work.

Panel A of Figure IV shows that, on average across the three observability conditions, 26% of workers accept the job at w . This magnitude is reasonable when compared with our exit survey, which finds a baseline involuntary unemployment rate of 43%. There are no statistically significant differences among the three observability conditions at w . This supports the claim that there is no detectable effect of the social observability protocols on job take-up when the prevailing wage norm is not being violated (for example, through changing workers' beliefs), although the estimates are somewhat imprecise. We find slightly higher levels of job take-up at w in the subsample of agricultural workers (30% job take-up). Again, we see no detectable differences by observability. Having verified our predicted insensitivity to social observability, 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. We consider the take-up rate for the *Private* $w - 10\%$ job. In the full sample, 17.5% of workers are willing to accept a wage cut in *Private*—indicating robust labor supply below the prevailing wage. It is also suggestive of a positive private residual labor supply elasticity, as we would anticipate, though the difference in take-up is not quite significant at conventional levels (p-value = 0.118). The take-up rate for the *Private* $w - 10\%$ job among agricultural workers is even higher, at 21.1%.

Recall that a stronger test of H1 is whether there is any $w - 10\%$ take-up in the *Employer Only* condition. We again find that take-up is robustly positive, even when the employer is present. While take-up is 5.1 pp lower than in the *Private* setting, the difference is statistically insignificant (p-value = 0.416). The employer's knowledge (alone) of the wage rate does not appear to shift workers' labor supply in quantitatively meaningful ways. The same pattern holds for agricultural workers.

In contrast, take-up plummets when wage cuts are offered in public. This is especially stark for members actually in the agricultural labor force: only 1.8% of them accept wage cuts in public. Our evidence in support of H2 comes from estimating the treatment effect of social observability on take-up rates for jobs offered at $w - 10\%$. We can first compare take-up rates between *Public* and *Private* wage cut treatments. We find that participants are 13.6 pp less likely to take up the job at a wage cut when it is offered in public relative to when it is private to the worker. This is 78% lower than the 17.5% take-up rate in the

Private condition (at $w - 10\%$). This difference is even more stark for agricultural workers, who reduce their labor supply by 24.6 pp when wage cuts are offered in *Public* rather than *Private*. Overall, increased social observability (i.e. the *Public* condition), which had no impact on take-up at the prevailing wage, significantly reduces labor supply only for jobs offered at below the prevailing wage.

We can also run a similar test of H2 using a differences-in-differences specification with the w treatments. This allows us to net out any potential level-shifters between the *Private* and *Public* observability conditions. Unsurprisingly, the results are quite similar. In the full sample, we estimate that social observability at $w - 10\%$ causes a decrease in take-up of 17.3 pp (p-value = 0.0481). In the agricultural worker subsample, this difference grows to 28.4 pp (p-value = 0.0086).

We hypothesize that workers behave as an informal union and sanction one another for deviating from the wage norm. Therefore, the spread of information to other workers should be what matters. Importantly, the *Public* treatments reveal the wage offer both to other workers and to the hiring employer. Therefore, we can also test H2 by comparing the job take-up rates at $w - 10\%$ between the *Public* treatment and *Employer Only* treatment. However, as we noted above, this test may be too strong. The impact of social observability in front of other laborers might be attenuated if the employer, who is part of the same broader village network, communicates information about norm violations to other workers.

Relative to the *Employer Only* treatment, the increased social observability of the *Public* treatment causes a statistically significant decline in take-up of $w - 10\%$ jobs of 8.4 pp (p-value=0.0865) among all labor colony residents. This decline grows to 17.0 pp (p-value=0.0107) for agricultural workers. These findings indicate that the presence of the employer itself doesn't drive our social observability results. Rather, consistent with H2, observability matters specifically in front 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. Below, we present further discussion and evidence against an employer-bargaining interpretation.

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

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.IV 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.

V.B 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 were—we conducted a daily recall of all employment activities in the days since the job offer was made as a part of the exit survey.

We present our findings in Table III. Each observation is a worker-day. In Cols. (1) and (2), we focus on the probability of working for a wage and earnings on the day of work as reported in the daily recall module. Participants offered the job in *Public* at a wage of $w - 10\%$ are 16.1 pp less likely to work for a wage on the day of work than their counterparts in the *Private* $w - 10\%$ treatment. This corresponds closely to the gap in take-up rate for the job from our “administrative” records. Accordingly, they earn Rs. 32.4 less for the day, on average. For a sense of magnitude, the Rs. 32.4 loss is 71% of mean earnings on the day of work in the *Private* $w - 10\%$ group. As a proportion of the mean earnings of the untreated holdout sample (not offered the job), the magnitude of earnings losses is even larger.

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).¹⁴ We then test the total impact of the job offer on weekly earnings, including our experimental day of work (Cols. (5)-(6)).¹⁵ 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, with a decline in estimated average daily earnings of Rs. 11.8. 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 earnings for the *Public* $w - 10\%$ group relative to the *Private* $w - 10\%$ group. In a more conservative

¹⁴Note, 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.

¹⁵We weight the data to account for missing days in the worker grid due to the timing of the exit survey.

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,¹⁶ those who are offered jobs at below the prevailing wage in public lose 26% of average weekly earnings, relative to those offered the same jobs in private. This indicates that workers are willing to give up a large fraction of weekly earnings in order to avoid being seen as breaking the social norm. It is important to note that while informative, earnings are not a direct measure of welfare (surplus). We discuss worker welfare in Section VIII.

V.C Involuntary Unemployment

Our empirical test only has power when there is involuntary unemployment. Suppose the market cleared, so that all agents willing to supply labor at $w - 10\%$ were able to find jobs at w . Then, even if our hypothesis were true, we would not expect robust take-up in any of the $w - 10\%$ treatments. In other words, the private wage cut treatment elicits not the underlying actual labor supply level, but rather residual labor supply: the percentage of workers who are willing to accept jobs at a wage cut given the market employment rate at w .

Figure V examines treatment effects 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.V, Cols. (1)-(2) present the regressions associated with Panels A and B of the figure, respectively. To economize on power, we pool all of the social observability treatments for wage w , and we pool the *Employer* and *Private* treatments for wage $w - 10\%$.

Consistent with the above intuition, takeup of private wage cuts ($w - 10\%$, Pooled Private) is higher when the village has a higher unemployment rate. Specifically, private takeup of wage cuts is 33.3 pp at $w - 10\%$ when involuntary unemployment levels are high (above median). In contrast, it is significantly (16.6 pp, p-value=0.0187) lower in villages with below-median involuntary unemployment (Figure V Panel A and Appendix Table A.V Col (1)). This difference is starker among agricultural workers, with a 23.1 pp lower takeup rate in low involuntary unemployment areas (p-value=0.009), relative to a mean of 39.3 pp in high involuntary unemployment areas (Figure V Panel B and Appendix Table A.V Col (2)).

How our treatment effects—the difference in takeup of private vs. public wage cuts—vary with unemployment levels is conceptually more ambiguous, with two potential opposing forces. (i) On one hand, because take-up is bounded below by zero, our ability to detect

¹⁶This is due to the noise in the point estimate in Col (3).

differences in public vs. private takeup will be higher when unemployment is higher, due to the intuition above. (ii) On the other hand, when workers have been unemployed for longer, they may be more desperate for work, and therefore more willing to incur social sanctions by accepting wage cuts in public.

We find stronger support for H2 in villages with above-median levels of involuntary unemployment. Namely, in the full sample, take-up of the job falls by 19.6 pp (p-value=0.00354) in high unemployment villages in the *Public* treatment relative to the *Private* treatment. This treatment effect of social observability is 16.0 pp smaller in magnitude in low unemployment villages (p-value=0.0808); the resulting total drop in take-up between the *Public* and *Private* wage cut treatments is only 3.6 pp and is statistically indistinguishable from zero. 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. In such villages, the difference in take-up of wage cuts in public is 37.2 pp lower than in private (p-value<0.001). Again, this large treatment effect is almost completely offset in low unemployment villages (coefficient=0.355, p-value=0.00252).

We also test separately for heterogeneous treatments effects using individual- rather than village-level involuntary unemployment in Cols. (3)-(4).¹⁷ 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.

Our findings indicate the effects of social observability are higher in areas with higher involuntary unemployment. Moreover, counter to force (ii) above, takeup levels of wage cuts are not higher when unemployment is higher; in fact, agricultural workers are *less* likely to take up public wage cuts in areas with higher unemployment. One possible interpretation of these patterns is that unemployment itself is influenced by the strength of the social norm. Specifically, if the areas with stronger social norms against accepting wage cuts are exactly those which have more unemployment (because the union leads to a more distorted labor market), then this could generate the various patterns we observe in Figure V.

V.D Information Flow

Our view of the mechanism underlying the labor supply results is that workers reduce their willingness to take a wage cut when they worry that their decisions are observable to other villagers. Therefore, in more diffusive villages - where more individuals are likely to learn of worker take-up decisions - we hypothesize that the main treatment effects will be larger.

This prediction echoes findings in the social networks literature. Breza and Chan-

¹⁷Individual-level unemployment is constructed using the percentage of days workers report being involuntarily unemployed in the last 30 days in the main worker sample. High involuntary unemployment is an indicator for above-median involuntary unemployment in this measure.

drasekhar (2019) argue that individuals behave better in the context of achieving self-set savings goals when information about their savings is likely to spread more widely in the village network. Alatas et al. (2016) show that communities with more diffusive networks have better information about community member characteristics. Finally, in their theory of social enforcement, Jackson et al. (2012) show that networks with locally dense structures are better suited to maintaining cooperation.

We asked two questions about village diffusiveness in our untreated holdout sample: a) the extent to which workers learn about the wages at which others accept agricultural work and b) how many others would find out if a worker accepted an agricultural job at below the prevailing wage.¹⁸ 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 is smaller for these low-diffusiveness villages.

In Col. (1) of Table IV, we consider the diffusiveness of general wage information and find that publicizing a low wage offer in highly diffusive villages leads to an approximately 20.0 pp decline in take-up rates (p-value=0.00344). However, in low diffusiveness villages, this large treatment effect is almost completely offset (coeff=17.0 pp, p-value=0.0701), leading to no measurable differences in take-up. The same pattern holds when, in Col. (2), we instead use the second measure of diffusiveness that more directly addresses norm violations. Finally, in Col. (3), we show that the magnitudes of both the decline in take-up rates (coeff= -30.8 pp, p-value<0.001) and the offsetting effect in low diffusiveness villages (coeff=21.4 pp, p-value=0.0614) are larger for the subset of agricultural workers.

These findings indicate larger effects of social observability when more individuals are likely to learn about norm violations. More generally, the level of diffusiveness could be correlated with other features, such as 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. Consequently, we interpret these results as stemming from the direct effects of diffusiveness and the features correlated with it—an interpretation that is consistent with our proposed mechanism.

V.E Threats to Validity and Interpretation

Next, we discuss several potential threats to the interpretation of our main results.

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 that fact

¹⁸Measuring network structure in each study village was beyond the scope of our field activities.

from the employers in the village in order to extract more private surplus. One implication is that workers only worry about others in the village observing their acceptance of wage cuts insofar as that information will be transmitted to employers.

The fact that the presence of the employer does not appear to lead to a meaningful reduction in 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.¹⁹ In Figure VI and Appendix Table A.VI, we examine whether the treatment effects hold when the worker is more likely to work for the hiring employer in the future (proxied by whether they have a pre-existing relationship). If the employer bargaining story were true, then we should expect large drops in job take-up rates moving from *Private* to *Employer* when the worker is more likely to work for the employer in the future. We find no evidence of this pattern. In fact, the point estimates are positive rather than negative, though noisy. Thus, employer bargaining does not appear 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 individual reservation wages.

Side Payments and Adjustment on Other Margins. Our hypotheses are constructed under the assumption 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 making cash side-payments, by shortening the work day, by increasing the value of in-kind perks, or by allowing workers to shirk. The results are presented in Appendix Table A.VII.

We first consider adherence to wage treatment status: whether workers in the $w - 10\%$ treatments indeed received lower wages than the prevailing wage. Overall, across all treatments, only 5% 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 (wage reported by the worker in the exit survey – the treatment wage level) (Col. 1), length of workday (Col. 2), or meals provided (Col. 3). The coefficients in Cols. (1)-(3) are small and statistically indistinguishable from zero.

Turning to worker selection and quality, we find that if anything, the $w - 10\%$ treatment workers are less likely to have worked for the employer in that past, though the coefficient is far from significant (Col. 4).²⁰ We also find very small, though imprecise, coefficients

¹⁹However, this test is admittedly imperfect. If workers worried about information spreading to other employers, it is conceivable that other workers might still transmit that information to them.

²⁰Even if the set of workers willing to accept an offer at $w - 10\%$ were somewhat negatively selected, in these markets all workers have been employed at the prevailing wage and are thus capable of satisfying the basic job requirements.

on the employer’s assessment of worker effort during the day (Col. 5).²¹ These patterns are similar when we restrict the sample of $w - 10\%$ treatments to only the *Private* arm, where the employer made his assessment not knowing the wage rate of the workers (Col. 6). Furthermore, even if there were compensating differentials that partially closed the gap between w and $w - 10\%$, then they would most likely cause an increase in take-up across all $w - 10\%$ treatments.

Relatedly, given that the employer sought out each worker in the experimental sample, it is possible that the offers of employment decreased worker search costs. Again, this should be true across all wage offers and social observability treatments. Moreover, search costs alone cannot explain why taking the job would be viewed as an action worthy of punishment.

Threats to Implementation. We consider two concerns that might impact our interpretation. First, 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.

Second, 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 I(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.VIII).

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, workers state that while some desperate workers may occasionally take the risk of accepting a wage cut, they would not do so regularly because they cannot actually trust the employer to keep the deal confidential from other workers. The employer’s incentive is to advertise acceptance of wage

²¹If lower wages select for lower quality workers (e.g., if the outside option of on-farm work is more valuable for higher ability workers), then the small positive coefficient on the work day rating is inconsistent with negative effort effects (e.g., stemming from negative gift exchange or efficiency wages), or indicates that monitoring effort by employers is difficult.

cuts to other workers in an attempt to break the wage floor or at least elicit supply at low wage offers from others. Consequently, workers fear that other workers would come to know what they have done, making acceptances risky. Second, and importantly, employers state that if a landowner gains a reputation for offering below prevailing wage jobs, he may benefit with some wage savings in the lean season. However, he runs a big risk that village workers will retaliate against him in the peak planting or harvest season—when work is abundant and workers have their pick of employers—by not coming to work for him. This would cause devastating losses to his annual income, making this too risky to do more than occasionally. This accords 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 against individual members.

VI Evidence: Sanctions

We posit that the observability effects stem from an attempt to avoid social disapprobation. However, other reputational mechanisms could produce patterns consistent with H1 and H2. For example, individuals could feel shame from being seen as financially desperate enough to accept a wage cut. Rather than attempt to fully rule out all other reasons why social observability may matter, we take on the more modest aim of providing positive evidence for H3—that individuals sanction norm violators. Note that if the only consequence of accepting a wage cut is that others will view the worker as financially vulnerable, this should not lead others to punish the worker.

VI.A Survey Responses

We begin by describing worker beliefs about the consequences of accepting wage cuts. In a survey with our untreated holdout sample (i.e. workers who were not offered jobs), we asked “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. We compile these responses into categories in Figure VII.

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.

VI.B Costly Punishment Exercise

To obtain direct, revealed-preference evidence on sanctions, we use a costly punishment game in a supplementary lab-in-the-field exercise. In a different set of 13 villages—drawn from the same village population as our main experiment—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. 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. The player and his anonymous partner are both given an endowment of Rs. 100. The player can “punish” his partner, reducing his endowment, by giving up some of his own endowment. Specifically, for every Rs. 5 that is removed from the partner’s endowment, the player must give up Rs. 1 of his own endowment. To make visualization easy, we implement this by placing 2 trays in front of the player, with Rs. 100 on each tray. The player then removes money from his tray and his partner’s tray, in 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 player is partnered with a worker in his own village or with a worker in a village that is geographically far away but has similar characteristics as the worker’s village (e.g. crop type, work composition, wage levels). 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 $2 \times 2 = 4$ 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.²² 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 the lab to decide that working at some wage levels deserves punishment and not others.

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

Figure VIII 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 punish the worker about 40% of the time. In addition, the desire to punish norm violations is not limited to actions in one’s own village. Players also punish workers from distant villages in similar frequencies—even though that worker’s action has no scope to affect the player’s own labor market.

Table V presents these results in regression form. Col. (1) shows that, on average, the punishment probability increases by 42 percentage points when the “partner” accepts a wage lower than the prevailing level (statistically significant at the 1% level). Col. (2) shows that this effect size is of very similar magnitude and is statistically indistinguishable when the “partner” lives in a different labor market 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

²²The 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.III. 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.

their partner, the amount that players forego from their own endowment, conditional on punishment, corresponds to 8.6% of typical daily earnings.

These results are consistent with the literature on social preferences, which indicates that individuals will be willing to destroy their own surplus to punish those who have engaged in norm violations (Charness and Rabin [2002]). Our findings are also potentially consistent with contagious punishment models (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—indicates that workers view norm violations in moral or general terms, suggesting 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. This helps shed light on why norms are effective in enabling workers to maintain wage floors despite large scope for self-interested deviations.

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). Moreover, many of the sanctions described in the survey evidence above may not even be costly for co-villagers to impose. This suggests that 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.

The above evidence indicates that decentralized sanctions help enable adherence to the norm. Additional survey evidence also suggests that labor supply behavior is decentralized and lacks formal organization (Appendix Figure B.I). For example, 89% of workers state that there is no village-level meeting for all or most workers in the village to discuss the wage, and 97% state that there is no meeting between workers and landowners to bargain over the wage for the season. This supports our view that there is a lack of explicit or organized collective action in this setting, highlighting the role of social norms (and the sanctions that preserve the norm) as an equilibrium coordination device.

VII Correlation with Wage Rigidity

In this section, we explore whether the mechanism documented in our field experiment has potential relevance for wage rigidity. To motivate this link, we first use data collected by Kaur (2019) in a different set of Indian villages that span 6 districts in two states.

Appendix Figure B.II 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). 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.²³ In India, caste is a strong proxy for in-group and social cohesion.²⁴ 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.²⁵

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. However, in the following year, when the positive shock has dissipated, wages do not adjust back down—

²³In 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 II). 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]).

²⁴Munshi 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.

²⁵The 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.

they remain ratcheted up. Because of this distortion on the wage, agricultural employment falls: it is lower than it would have been in absence of the lagged positive shock.²⁶

Panel A, Col. (1) of Table VI 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.

Relative to having 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.²⁷ In Col. (3), we use an alternate

²⁶The 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.

²⁷More 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

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 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.IV shows the underlying distribution of each caste Herfindahl index across Indian districts. Appendix Table A.IX 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. A causal analysis is beyond the scope of our paper. While only suggestive, the results in Table VI are consistent with the view that social cohesion, and its resultant ability to sustain stronger social norms, could have aggregate implications by leading large numbers of individuals to coordinate on the same strategy.

VIII Surplus Estimation

One key question is whether workers benefit from adhering to the wage norm. On one hand, the norm might help workers to exert market power, extracting surplus from employers. On the other hand, it is possible that the norm originated under different labor market conditions, and could actually make workers worse off today. After all, wage floors may increase wages for the average worker at the cost of increased involuntary unemployment.

We conduct a simple back-of-the-envelope exercise to estimate the counterfactual market-

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).

clearing level of wages and employment in the absence of the wage floor, and compute the change in worker surplus 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. We also assume that employers do not behave monopsonistically.

Figure IX(a) illustrates the distortionless competitive equilibrium (L^*, W^*) , and Figure IX(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 Kaur (2019).

We present our calculations in Table VII. 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 (e.g. Rosenzweig [1978]; Bardhan [1979]; Abdulai and Delgado [1999]; Goldberg [2016]). 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 A.X). Within these ranges, there are gains to employer surplus from the wage floor under efficient rationing of between 8% and 98%.²⁸ As in Lee and Saez (2012), workers gain more when labor supply is relatively elastic

²⁸With the exception of the case in which labor demand is extremely elastic ($\eta = 4$) and labor supply is

and labor demand is relatively inelastic.

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 V(b), and provide estimates in Appendix Table A.XI. Even under this conservative assumption, worker surplus increases up to 72% when labor demand is relatively inelastic ($\eta < 2$).

Finally we note that, contrary to our assumption, if employers do exert monopsony power, then this would lead to even greater gains in surplus for workers under the wage floor. This makes our main estimates a lower-bound on workers’ potential surplus gains. Moreover, in that case, acting like an informal union might even be efficiency-improving.

IX Conclusion: Potential Generality of Mechanism

We find evidence that workers would privately like to supply labor below the prevailing wage, but do not do so when their take-up decisions are publicly observable. This indicates that social pressure dampens labor supply below the prevailing wage, supporting the presence of wage floors in village labor markets. Our findings provide documentation of a way in which norms against accepting wage cuts distort the aggregate labor supply curve.

The forces we document have potential generality beyond our specific setting. 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.

This makes the labor market an especially relevant economic domain. In the US labor market, norms have the potential to arise at the establishment level, where social interaction is high: workers see their peers for eight hours per day every day. In addition, peers have the ability to levy sanctions 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). Consistent with this view, in *The Labor Market as a Social Institution*, Robert Solow (1990) argues that the social norms that naturally arise in the workplace are inherent to what distinguishes labor from other commodity markets. 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. Indeed, collec-

extremely inelastic ($\nu = 0.5$).

tive behaviors among workers—e.g., mass walkouts with foregone wages, the coordinated restriction of output, punishment of rate busters—have been documented across history and contexts in settings with no formal unions (Clark [1984]; Clark [1987]).

More broadly, the presence of cartel-like behavior among decentralized individuals has been qualitatively observed in a range of 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-eighths 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, qualitative interviews revealed that this behavior reflected a norm: traders implicitly understood that this pricing behavior should be followed. Moreover, traders felt that deviations would be punished through social 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.²⁹ 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 impact of social norms on markets may be particularly relevant in poor countries. Given the small scale of most enterprises, the "firm" is often a single individual—for example, a vendor in a market, whose stall is surrounded by that of other vendors who he sees each day. This may help explain price collusion or rigidity in product markets—from taxi drivers to vegetable vendors to agricultural traders (e.g. Bergquist [2018]). Similarly, the communal nature of village economies may help explain price homogeneity in a range of factor markets—from bullock and tractor rentals to sharecropping (e.g. Shaban [1987]). More broadly, the high levels of social capital in networks in poor countries may make it more likely that norms arise and can be sustained effectively through punishment.

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 social norms are a potentially relevant force for equilibria in a variety of contexts.

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

²⁹We thank Doug Bernheim for pointing us to this example.

wage rigidity and wage compression.

EBREZA@FAS.HARVARD.EDU; DEPARTMENT OF ECONOMICS, HARVARD UNIVERSITY AND NBER
CORRESPONDING AUTHOR, SUPREET@BERKELEY.EDU; DEPARTMENT OF ECONOMICS, UNIVERSITY OF CALIFORNIA AT BERKELEY AND NBER

KRISHNASWAMY.NANDITA@COLUMBIA.EDU; SCHOOL OF INTERNATIONAL AND PUBLIC AFFAIRS,
COLUMBIA UNIVERSITY

REFERENCES

- Abdulai, Awudu and Christopher L Delgado, “Determinants of nonfarm earnings of farm-based husbands and wives in Northern Ghana,” *American Journal of Agricultural Economics*, 81 (1999), 117–130.
- Alatas, Vivi, Abhijit Banerjee, Arun G Chandrasekhar, Rema Hanna, and Benjamin A Olken, “Network structure and the aggregation of information: Theory and evidence from Indonesia,” *American Economic Review*, 106 (2016), 1663–1704.
- Bandiera, Oriana, Iwan Barankay, and Imran Rasul, “Social preferences and the response to incentives: Evidence from personnel data,” *The Quarterly Journal of Economics*, 120 (2005), 917–962.
- Bardhan, Pranab K, “Labor supply functions in a poor agrarian economy,” *The American Economic Review*, 69 (1979), 73–83.
- Bénabou, Roland and Jean Tirole, “Incentives and prosocial behavior,” *American Economic Review*, 96 (2006), 1652–1678.
- Bendor, Jonathan and Piotr Swistak, “The evolution of norms,” *American Journal of Sociology*, 106 (2001), 1493–1545.
- Bergquist, Lauren Falcao, “Pass-through, competition, and entry in agricultural markets: Experimental evidence from Kenya,” *Working Paper*, (2018).
- Breza, Emily and Arun G Chandrasekhar, “Social Networks, Reputation, and Commitment: Evidence From a Savings Monitors Experiment,” *Econometrica*, 87 (2019), 175–216.
- _____, Supreet Kaur, and Yogita Shamdasani, “The morale effects of pay inequality,” *The Quarterly Journal of Economics*, 133 (2017), 611–663.
- _____, _____, and _____, “Labor Rationing: Revealed Preference Estimates from Hiring Shocks,” *Working Paper*, (2019).
- Bursztyn, Leonardo, Alessandra L. González, and David Yanagizawa-Drott, “Misperceived social norms: Female labor force participation in Saudi Arabia,” (2018).

- and Robert Jensen, “How does peer pressure affect educational investments?,” *The Quarterly Journal of Economics*, 130 (2015), 1329–1367.
- and — , “Social image and economic behavior in the field: Identifying, understanding, and shaping social pressure,” *Annual Review of Economics*, 9 (2017), 131–153.
- , Thomas Fujiwara, and Amanda Pallais, “‘Acting wife’: Marriage market incentives and labor market investments,” *American Economic Review*, 107 (2017), 3288–3319.
- Charness, Gary and Matthew Rabin, “Understanding social preferences with simple tests,” *The Quarterly Journal of Economics*, 117 (2002), 817–869.
- Christie, William G and Paul H Schultz, “Why do Nasdaq market makers avoid odd-eighth quotes?,” *The Journal of Finance*, 49 (1994), 1813–1840.
- Cialdini, Robert B and Noah J Goldstein, “Social influence: Compliance and conformity,” *Annu. Rev. Psychol.*, 55 (2004), 591–621.
- Clark, Gregory, “Authority and efficiency: the labor market and the managerial revolution of the late nineteenth century,” *The Journal of Economic History*, 44 (1984), 1069–1083.
- , “Why isn’t the whole world developed? Lessons from the cotton mills,” *The Journal of Economic History*, 47 (1987), 141–173.
- Crawford, Sue ES and Elinor Ostrom, “A grammar of institutions,” *American Political Science Review*, 89 (1995), 582–600.
- DellaVigna, Stefano, John A List, and Ulrike Malmendier, “Testing for altruism and social pressure in charitable giving,” *The Quarterly Journal of Economics*, 127 (2012), 1–56.
- , — , — , and Gautam Rao, “Voting to tell others,” *The Review of Economic Studies*, 84 (2016), 143–181.
- Dickens, William T, Lorenz Goette, Erica L Groshen, Steinar Holden, Julian Messina, Mark E Schweitzer, Jarkko Turunen, and Melanie E Ward, “How wages change: Micro evidence from the International Wage Flexibility Project,” *Journal of Economic Perspectives*, 21 (2007), 195–214.
- Dreze, Jean P. and Anindita Mukherjee, “Labour contracts in rural india: Theories and evidence,” in “The balance between industry and agriculture in economic development” (Springer 1989): pp. 233–265.
- Ellison, Glenn, “Cooperation in the prisoner’s dilemma with anonymous random matching,” *The Review of Economic Studies*, 61 (1994), 567–588.

- Farber, Henry S. and Daniel H. Saks, “Why Workers Want Unions: The Role of Relative Wages and Job Characteristics,” *Journal of Political Economy*, 88 (1980), 349–369.
- Fehr, Ernst and Ivo Schurtenberger, “Normative foundations of human cooperation,” *Nature Human Behaviour*, 2 (2018), 458.
- Gächter, Simon and Ernst Fehr, “Collective action as a social exchange,” *Journal of Economic Behavior & Organization*, 39 (1999), 341 – 369.
- Goldberg, Jessica, “Kwacha gonna do? Experimental Evidence about labor supply in rural Malawi,” *American Economic Journal: Applied Economics*, 8 (2016), 129–49.
- Henrich, Joseph, Richard McElreath, Abigail Barr, Jean Ensminger, Clark Barrett, Alexander Bolyanatz, Juan Camilo Cardenas, Michael Gurven, Edwins Gwako, Natalie Henrich et al., “Costly punishment across human societies,” *Science*, 312 (2006), 1767–1770.
- Houde, Jean-Francois, Terence Johnson, Molly Lipscomb, and Laura Schechter, “Measuring and Addressing Market Power issues in the Desludging Market in Senegal,” *Working Paper*, (2018).
- Hsieh, Chang-Tai and Enrico Moretti, “Can free entry be inefficient? Fixed commissions and social waste in the real estate industry,” *Journal of Political Economy*, 111 (2003), 1076–1122.
- Jackson, Matthew O, Tomas Rodriguez-Barraquer, and Xu Tan, “Social capital and social quilts: Network patterns of favor exchange,” *American Economic Review*, 102 (2012), 1857–97.
- Kandori, Michihiro, “Social norms and community enforcement,” *The Review of Economic Studies*, 59 (1992), 63–80.
- Kaur, Supreet, “Nominal Wage Rigidity in Village Labor Markets,” *American Economic Review (Forthcoming)*, (2019).
- Lee, David and Emmanuel Saez, “Optimal minimum wage policy in competitive labor markets,” *Journal of Public Economics*, 96 (2012), 739–749.
- Lewis, W Arthur, “Economic Development With Unlimited Supplies of Labour,” *The Manchester School*, 22 (1954), 139–191.
- MacLeod, W. Bentley, “Can Contract Theory Explain Social Preferences?,” *The American Economic Review*, 97 (2007), 187–192.
- Mas, Alexandre and Enrico Moretti, “Peers at work,” *American Economic Review*, 99 (2009), 112–45.
- Mazzocco, Maurizio and Shiv Saini, “Testing efficient risk sharing with heterogeneous risk preferences,” *American Economic Review*, 102 (2012), 428–68.

- Munshi, Kaivan and Mark Rosenzweig, “Traditional institutions meet the modern world: Caste, gender, and schooling choice in a globalizing economy,” *American Economic Review*, 96 (2006), 1225–1252.
- and —, “Networks and misallocation: Insurance, migration, and the rural-urban wage gap,” *American Economic Review*, 106 (2016), 46–98.
- Osmani, S. R., “Wage determination in rural labour markets,” *Journal of Development Economics*, 34 (1990), 3–23.
- Ostrom, Elinor, James Walker, and Roy Gardner, “Covenants with and without a sword: Self-governance is possible,” *American Political Science Review*, 86 (1992), 404–417.
- Rosenzweig, Mark R., “Rural wages, labor supply, and land reform: A theoretical and empirical analysis,” *American Economic Review*, 68 (1978), 847–861.
- Rosenzweig, Mark R., “Labor markets in low-income countries,” in “Handbook of Development Economics,” Vol. 1 (Elsevier 1988): pp. 713–762.
- Rosenzweig, Mark R and Oded Stark, “Consumption smoothing, migration, and marriage: Evidence from rural India,” *Journal of Political Economy*, 97 (1989), 905–926.
- Shaban, Radwan Ali, “Testing between competing models of sharecropping,” *Journal of Political Economy*, 95 (1987), 893–920.

X Tables

TABLE I: Covariate Balance

COVARIATES	(1)	(2)	(3)	(4)	(5)	(6)
<i>Treatments</i>	Wage Cut Private	Wage Cut Employer	Wage Cut Public	Prevailing All	Joint Sig F-Stat	Obs
INDIVIDUAL-LEVEL						
Age	44.337 (12.639)	2.890 (2.032)	1.893 (1.911)	0.881 (1.806)	0.180	442
1.(Caste: Scheduled Tribe)	0.333 (0.474)	-0.0179 (0.0827)	-0.106 (0.0671)	-0.0259 (0.0755)	0.842	444
1.(Caste: Scheduled Caste)	0.667 (0.474)	0.0169 (0.0825)	0.100 (0.0670)	0.0255 (0.0756)	0.382	444
1.(Casual Laborer)	0.944 (0.230)	0.0417 (0.0346)	0.0604 (0.0269)	0.0486 (0.0267)	0.177	446
1.(Casual Laborer-Agriculture)	0.844 (0.364)	-0.0691 (0.0662)	-0.0398 (0.0620)	-0.00587 (0.0530)	0.840	446
1.(Doesn't Own Land)	0.678 (0.470)	-0.104 (0.0897)	-0.0720 (0.0745)	-0.0156 (0.0705)	0.876	446
Individual Employment Rate- Past 30 days	0.295 (0.206)	-0.0194 (0.0326)	0.0331 (0.0343)	0.0521 (0.0315)	0.054	427
Involuntary Unemployment Rate- Past 30 Days	0.428 (0.229)	-0.0130 (0.0425)	-0.0880 (0.0350)	-0.0642 (0.0356)	0.071	427
1.(Employer Experience)	0.318 (0.468)	0.188 (0.0987)	0.0346 (0.0999)	0.114 (0.0871)	0.248	426
1.(Completed Exit Survey)	0.874 (0.334)	0.0178 (0.0495)	0.0284 (0.0458)	-0.00678 (0.0499)	0.389	502
Number of Onlookers	5.162 (2.646)	.	5.162 (2.646)	-0.447 (0.631)	0.481	189
VILLAGE-LEVEL						
Number of Households	45.639 (13.135)	-3.283 (3.603)	-3.880 (3.266)	1.366 (3.057)	0.172	172
Involuntary Unemployment Rate - Past 30 Days	0.398 (0.104)	0.0289 (0.0280)	0.0243 (0.0256)	0.00691 (0.0240)	0.449	170
Info Flow: Wage	0.752 (0.306)	-0.0202 (0.0743)	0.0457 (0.0679)	0.0612 (0.0631)	0.489	182
Info Flow: Norm Violation	0.622 (0.345)	-0.0674 (0.0837)	0.0442 (0.0765)	0.0413 (0.0710)	0.486	182
1.(NREGA ever in village)	0.914 (0.284)	-0.110 (0.0710)	-0.0492 (0.0650)	0.0335 (0.0612)	0.213	180
Prop Days NREGA - Recall	0 (0)	0.0106 (0.0108)	-0.00248 (0.00989)	0.00288 (0.00931)	0.410	180
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 for each treatment arm relative to the reference group. All specifications include Year X Month and task fixed effects. First panel focuses on individual-level characteristics. Involuntary unemployment is measured as prop days in past 30 that the worker wanted work but could not find it. Observations are weighted by the number of individuals in the regression sample from each village. Number of onlookers variable is recorded only for public treatments. Mean number of onlookers for hiring conducted at $W - 10\%$ in public is shown in Col. (3). Col. (4) presents the regression coefficient comparing the public W relative to $W - 10\%$ treatments. Second panel focuses on village-level covariates measured in the untreated holdout sample. Village unemployment is based on prop of days in past 30 that individuals wanted work but could not find it. Measures of information spread are based on the average perceived likelihood that knowledge of the wages of others would spread and average perceived likelihood that knowledge of norm violations in the labor market would spread. Standard errors (clustered at the village level in Panel 1) 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 slight differences in the surveys across experimental rounds.

TABLE II: Main Results - Effects on Labor Supply

	(1)	(2)	(3)
	Worked	Worked	Worked
Wage cut: Public	-0.122 (0.0564) [0.0316]	-0.136 (0.0573) [0.0188]	-0.246 (0.0644) [0.000181]
Wage cut: Employer	-0.0657 (0.0611) [0.284]	-0.0516 (0.0633) [0.416]	-0.0758 (0.0788) [0.337]
Prevailing wage: Private	0.0609 (0.0703) [0.387]	0.0791 (0.0659) [0.232]	0.0663 (0.0819) [0.419]
Prevailing wage: Public	0.119 (0.0808) [0.144]	0.116 (0.0713) [0.105]	0.104 (0.0856) [0.226]
Prevailing wage: Employer	0.0364 (0.0775) [0.639]	0.0690 (0.0886) [0.437]	0.0935 (0.0992) [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
<i>Test</i> Prevailing wage: Private = Prevailing wage: Public	0.460	0.589	0.658
<i>Test</i> Wage cut: Private - Public = Prev. wage: Private - Public	0.0629	0.0481	0.00858
<i>Test</i> Wage cut: Employer = Wage cut: Public	0.143	0.0865	0.0107
<i>Test</i> Prev. wage: Private = Employer = Public	0.609	0.816	0.904
<i>Test</i> Prevailing wage: Employer = Prevailing wage: Public	0.331	0.608	0.918
<i>Test</i> Prevailing wage: All = Wage cut: Private	0.196	0.118	0.228

Notes: This table presents the effect of each treatment on our main outcome of interest, the take-up rate for the job offer. 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 X Month and task fixed effects. Cols. (1) and (2) include the full sample. Col. (3) restricts the sample to workers who indicated in the exit survey that they engage in agricultural labor as a primary or secondary occupation. 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 III: Earnings Results

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

Notes: This table presents the effects of our job offers on individual earnings, derived from the employment recall grid performed in the exit survey. Each observation 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 the Wage cut: Private treatment. All specifications include yearXmonth and task fixed effects. Cols. (1)-(2) restrict the sample to the day of work offered in the village 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 IV: Heterogeneous Treatment Effects: Village Information Spread

	(1) Worked	(2) Worked	(3) Worked
Wage cut: Public	-0.200 (0.0675) [0.00344]	-0.186 (0.0646) [0.00442]	-0.308 (0.0745) [5.50e-05]
Prevailing wage (pooled)	0.0794 (0.0717) [0.269]	0.0564 (0.0567) [0.321]	0.0467 (0.0717) [0.515]
Wage cut: Public x Low info spread village	0.170 (0.0932) [0.0701]	0.150 (0.0921) [0.105]	0.214 (0.114) [0.0614]
Prevailing wage (pooled) x Low info spread village	0.0521 (0.0913) [0.569]	0.115 (0.0844) [0.175]	0.146 (0.106) [0.171]
Low info spread village	-0.0732 (0.0667) [0.274]	-0.0380 (0.0621) [0.541]	-0.0263 (0.0796) [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. In Col. (1), the heterogeneous variable of interest is an indicator for below-median knowledge of the wages of others reported by the untreated holdout survey respondents. In Cols. (2)-(3), we use an indicator for below-median spread of information about other workers accepting a job below the prevailing wage, again reported by the untreated holdout survey respondents. In all specifications, the dependent variable is an indicator for whether the worker accepted the job and worked for the employer. In all columns, the omitted category is the Wage cut: Private pooled treatment for high info spread villages only. We also pool all of the Prevailing wage treatments together. In this table, we pool Wage cut: Private and Wage cut: Employer. All specifications include Year X Month 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 V: Sanctions - Costly Punishment Game

	(1)	(2)	(3)	(4)	(5)
	Any Punishment	Any Punishment	Any Punishment	Any Punishment	Partner's Payoff
Worker Accepted Job Below Prevailing Wage	0.420 (0.0447) [0]	0.393 (0.0632) [5.99e-09]	0.393 (0.0647) [1.23e-08]	0.436 (0.103) [4.66e-05]	-14.57 (4.425) [0.00132]
Worker Accepted a Job Below Prevailing Wage x Distant Labor Market		0.0494 (0.0894) [0.582]	0.0494 (0.0916) [0.590]	-0.00310 (0.137) [0.982]	5.569 (4.551) [0.224]
Worker Lives in Distant Labor Market		0.0143 (0.0143) [0.320]	0.0133 (0.0185) [0.472]	0.00737 (0.0294) [0.803]	-0.701 (1.259) [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. Each observation is a player X 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 VI: Wage Rigidity: Correlation with Social Cohesion

	Proxy for Low Worker Cohesion		
		Wage Labor: Caste Herfindahl (Below Median)	Agri Labor Force: Caste Herfindahl (Below Median)
	(1)	(2)	(3)
<i>Panel A - Dependent variable: Log Agricultural Wage</i>			
Positive shock last year	0.0532 (0.022) [0.017]	0.102 (0.042) [0.015]	0.0971 (0.033) [0.003]
Positive shock last year x Low worker cohesion		-0.0826 (0.050) [0.100]	-0.0899 (0.038) [0.019]
Positive shock this year	0.0633 (0.018) [0.000]	0.0800 (0.038) [0.036]	0.0751 (0.039) [0.055]
Positive shock this year x Low worker cohesion		-0.0242 (0.042) [0.561]	-0.0181 (0.043) [0.675]
Observations (worker-days)	59243	59243	59243
<i>Panel B - Dependent variable: Agricultural Employment</i>			
Positive shock last year	-0.135 (0.055) [0.014]	-0.234 (0.078) [0.003]	-0.172 (0.080) [0.031]
Positive shock last year x Low worker cohesion		0.189 (0.088) [0.032]	0.0716 (0.107) [0.505]
Positive shock this year	0.157 (0.062) [0.011]	0.133 (0.083) [0.109]	0.131 (0.091) [0.153]
Positive shock this year x Low worker cohesion		0.0394 (0.114) [0.729]	0.0469 (0.123) [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 shocks on wages and employment (testing whether wages to fall after a positive shock has dissipated, a test for wage rigidity) and examines the heterogeneity of the effect by two measures of worker cohesion. 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.

TABLE VII: 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 Equilibrium	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 Kaur (2019) 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 Breza et al. (2019), set in 64 other villages in the same districts of Odisha as this experiment.

XI Figures




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



(b) Difference Between Untreated Holdout Group Reports and Worker Census Participant Report of Prevailing Wage.

FIGURE I: Dispersion of Agricultural Wages

Note: 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. (2017). $N = 377$ worker-days in 26 villages in Odisha similar to (but distinct from) those in the experimental sample. Panel B displays the distribution of: (prevailing daily agricultural wage for men in village reported by untreated holdout sample of workers in experimental villages ($N=626$ workers)) – (our ex ante elicitation of the village male prevailing wage, upon which the experimental wage offer levels were assigned). The latter was elicited through asking employers and the “worker census participant” in the labor colony during the time of ex ante listings in each experimental village.



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

(b) Sanctions for Accepting Wage Cuts. *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 II: Motivational Evidence - Norms Against Accepting Wage Cuts

Note: 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.


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

FIGURE III: Experimental Design

Note: Numbers in parentheses indicate intended randomization weights on each treatment cell.




(a) All Workers




(b) Agricultural Workers Only

FIGURE IV: Main Results - Labor Supply

Note: These figures graph take-up rates for the job offer under each of the 6 treatment arms. The take-up rate for the job is defined as attendance on the day of work. Panel A uses the entire sample (N=502 participants); Panel B restricts the sample to casual daily wage workers (N=363 participants), who report their primary or secondary occupation to be agricultural labor. All robust 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 job offers at 10% below the prevailing wage. These results are also presented in the form of regressions in Cols. (2)-(3) of Table II. All regressions include task and yearXmonth fixed effects.



(a) All Workers




(b) Agricultural Workers Only

FIGURE V: Job Take-Up by Level of Involuntary Unemployment

Note: 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 their actual timeuse, as reported in the untreated holdout sample survey. The take-up rate for the job is defined as attendance on the day of work. Panel A uses the entire sample ($N=493$ participants); Panel B restricts the sample to casual daily wage workers ($N=363$ participants), who report their primary or secondary occupation to be agricultural labor. All robust 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 job offers at 10% below the prevailing wage 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.V. All regressions include task and yearXmonth fixed effects.




(a) All Workers




(b) Agricultural Workers Only

FIGURE VI: Job Take-Up by Level of Prior Experience With Employer

Note: These figures present heterogeneous treatment effects by whether the worker has worked with the participating employer in the past. The take-up rate for the job is defined as attendance on the day of work. Panel A uses the entire sample of workers with exit survey responses (N=426 participants); Panel B restricts the sample to casual daily wage workers (N=350 participants), who report their primary or secondary occupation to be agriculture and who responded to the exit survey. All robust 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 job offers at 10% below the prevailing wage for that level of experience. These results are also presented in the form of regressions in Cols. (1) and (2) of Table A.VI. All regressions include task and yearXmonth fixed effects.



(a) Sanctions



(b) Social Pressure

FIGURE VII: Survey Evidence - Sanctions for Accepting Wage Cuts

Note: 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 others 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)).



FIGURE VIII: Sanctions: Costly Punishment Game

Note: 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. All robust 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 V, and control for village fixed effects.




FIGURE IX: Equilibrium wages and employment under different market structures