

NBER WORKING PAPER SERIES

THE MORALE EFFECTS OF PAY INEQUALITY

Emily Breza
Supreet Kaur
Yogita Shamdasani

Working Paper 22491
<http://www.nber.org/papers/w22491>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2016

We thank James Andreoni, Dan Benjamin, Stefano DellaVigna, Pascaline Dupas, Edward Glaeser, Robert Gibbons, Uri Gneezy, Seema Jayachandran, Lawrence Katz, Peter Kuhn, David Laibson, Ulrike Malmendier, Bentley MacLeod, Sendhil Mullainathan, Mark Rosenzweig, Bernard Salanie, and Eric Verhoogen for their helpful comments. Arnesh Chowdhury, Mohar Dey, Piyush Tank, and Deepak Saraswat provided outstanding research assistance. We gratefully acknowledge operational support from JPAL South Asia and financial support from the National Science Foundation, the IZA Growth and Labor Markets in Low Income Countries (GLM-LIC) program, and the Private Enterprise Development for Low Income Countries (PEDL) initiative. The project was registered in the AEA RCT Registry, ID 0000569. 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 They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by Emily Breza, Supreet Kaur, and Yogita Shamdasani. 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.

The Morale Effects of Pay Inequality
Emily Breza, Supreet Kaur, and Yogita Shamdasani
NBER Working Paper No. 22491
August 2016
JEL No. D03,E24,J3,O15

ABSTRACT

The idea that worker utility is affected by co-worker wages has potentially broad labor market implications. In a month-long experiment with Indian manufacturing workers, we randomize whether co-workers within production units receive the same flat daily wage or different wages (according to baseline productivity rank). For a given absolute wage, pay inequality reduces output and attendance by 0.24 standard deviations and 12%, respectively. These effects strengthen in later weeks. Pay disparity also lowers co-workers' ability to cooperate in their self-interest. However, when workers can clearly observe productivity differences, pay inequality has no discernible effect on output, attendance, or group cohesion.

Emily Breza
Graduate School of Business
Columbia University
3022 Broadway, Uris Hall 821
New York, NY 10027
and NBER
ebreza@gsb.columbia.edu

Yogita Shamdasani
Department of Economics
Columbia University
1022 International Affairs Building
420 West 118th Street
New York, NY 10027
yjs2112@columbia.edu

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

A randomized controlled trials registry entry is available at
<https://www.socialscienceregistry.org/trials/569/history/>

1. INTRODUCTION

In traditional agency models, workers care about only their own wage levels when making labor supply decisions. However, a long tradition in economic thought—as well as in psychology, sociology, and organizational behavior—has advanced the notion that individuals also care about their pay *relative* to that of their co-workers.¹ This paper tests the empirical validity of this view. In addition, it examines workers’ underlying notion of equity: specifically, whether perceived justifications for pay differences mediate fairness concerns.²

Relative pay comparisons between workers have the potential to influence various labor market features. For example, they could help explain why wage compression—when wages vary less than the marginal product of labor—appears prevalent in both poor and rich countries (Frank 1984, Dreze and Mukherjee 1989, Lazear 1989, Fang and Moscarini 2005, Charness and Kuhn 2007). They have also been proposed as a micro-foundation for wage rigidity (Akerlof and Yellen 1990). In addition, relative pay concerns could affect how workers sort into firms—for example, leading some firms to specialize by worker ability, or mediating how productivity dispersion maps to earnings inequality (Frank 1984, Song et al. 2015, Card et al. 2013). Relatedly, they could influence the decision to contract labor within or across firm boundaries (e.g., Nickerson and Zenger 2008).³

The existing behavioral economics literature has focused on the idea that agents’ own internal preferences over relative pay can affect their utility. This idea has two related implications for worker behavior. First, relative pay can act as a compensating differential—affecting the willingness to accept employment at a given absolute wage. Second, such utility effects could impact effort, for example, in the presence of reciprocity.⁴ Work in social psychology and sociology suggests that this exclusive focus on internal preferences may provide an incomplete view. These literatures emphasize that pay disparity can foster “resentment, lack of cooperation, sabotage, and lack of team potency” (Shaw 2015). These forces have the potential to create social conflict, amplifying effects on labor supply and output across all workers. For example, even if higher-paid workers individually derive positive utility from being paid more than others, their outcomes could be worsened due to discontent or hostility from their co-workers—who, mechanically, must be paid relatively less (Schmitt and Marwell 1972, Deutsch 1986, Lazear 1989, Levine 1991, Duffy et al. 2012).

¹In economics, see, e.g., Marshall (1890), Veblen and Almy (1899), Keynes (1936), Duesenberry (1949), Hicks (1963), Easterlin (1974), Hamermesh (1975). In psychology, seminal work by Adams (1963) develops a theory of pay inequity, building on Festinger (1954)’s cognitive dissonance theory. See Cook and Hegtvedt (1983) and Shaw (2014) for reviews of work in social psychology, sociology, and organizational behavior.

²A long standing debate in this literature is on conceptualizing equity—whether individuals compare pay in levels, or ratios of pay to productivity (e.g., Adams 1963, Thibaut and Kelley 1959, Baron and Kreps 2013).

³See Fehr et al. (2009) for a general discussion of how fairness considerations may affect the labor market.

⁴Behavioral agency theory predicts that workers may reduce effort in response to fairness violations, especially under incomplete contracting (Kőszegi 2014). Fehr et al. (2009) argue that because most jobs contain some incomplete contracting, maintaining morale is essential for effort provision. This view is espoused by Bewley (1999), who documents managers’ perceptions that relative pay is important for worker motivation.

Figure 1 documents perceptions among Indian workers that pay inequality creates such morale effects. Respondents predict that, for a given absolute wage level, a worker who is paid less than his peers would be less likely to accept employment (Question 1) and would reduce effort (Question 2). In addition, 94% of respondents state that pay disparity will lead to conflict among co-workers (Question 3). While only suggestive, these perceptions match those of respondents in other contexts such as the U.S. (e.g., Bewley 1999).

In this paper, we use a field experiment to test whether workers care about relative pay. We construct a design that enables comparisons of workers who earn the same absolute wage, but differ in their co-workers’ wages. In addition, we incorporate variation in the perceived justification for pay differences. We examine the effects of relative pay on labor supply: the extensive margin decision to accept work (attendance) and effort provision (output). We also explicitly test for impacts on group cohesion among co-workers.⁵

In the experiment, 378 workers in Odisha, India are employed full-time for one month in seasonal manufacturing jobs—a prominent source of local employment. They work in small factories, which are organized into distinct production units, with 3 workers per unit. All unit members produce the same exact product (e.g. rope), while every unit within a factory produces a different product (e.g. rope, brooms, incense sticks). Each unit sits together in a separate physical space in the factory, both during work and lunch breaks. Thus, the other workers in one’s unit constitute a natural and salient reference group for pay comparisons. Note that production is an individual activity, with no joint production of any kind.

All workers are paid a flat daily wage for attendance, in accordance with the typical pay structure in the area. The decision to come to work is therefore incentivized: being absent results in a loss of earnings. However, there is incomplete contracting on effort: workers have some latitude to select effort levels, which are reflected in output. To obtain measures of worker output, we bear substantial overhead costs—in the form of extra staff and record keeping—to quantify each worker’s daily individual production.

We induce exogenous variation in co-worker pay by randomly assigning units to one of four different pay structures. In the *Pay disparity* condition, each unit member is paid a different wage— w_{High} , w_{Med} , or w_{Low} —in accordance with his respective productivity rank within the unit (determined by baseline productivity levels). These pay differences are fairly modest: the difference between each wage level is less than 5%. In the three *Compressed* pay conditions, all three unit members are paid the exact same wage, which we randomly assign to be w_{High} , w_{Med} , or w_{Low} .⁶ This allows us to compare, for example, workers with

⁵From the firm’s perspective, there are of course potential benefits of differential pay, such as incentive effects. Optimal pay policies would balance such benefits against potential morale costs. Our goal in this paper is not to evaluate this tradeoff. Rather, our design focuses on isolating the more basic question of whether workers care about relative pay.

⁶At the beginning of the baseline (i.e. “training”) period, workers are told that they will receive a wage increase on a pre-specified date, and that the size of this increase may depend on their baseline productivity. Once they are randomized into their wage treatment on this date, no additional future wage changes are possible. This shuts down dynamic incentive effects; see below for a discussion of this.

the same average baseline productivity level who both earn an absolute wage of w_{Low} , but differ in whether they are paid less than their peers (under the *Pay disparity* treatment) or the same as their peers (under the *Compressed Low* wage treatment). Note that managers maintain pay secrecy; any learning about peer wages is through self-disclosure.

To test whether perceived justifications mediate morale effects, we incorporate two additional sources of variation into our design. First, while wage levels are fixed, underlying baseline productivity is continuous. This induces variation in the extent to which pay differences among co-workers overstate productivity differences.⁷ Second, co-worker output is far more observable for some tasks relative to others. By randomizing units to production tasks, we generate variation in the observability of co-worker productivity.⁸

Even though managers maintain pay secrecy, workers learn about co-worker wages. At the end of the month, 86.6% of workers can accurately report the wages of both co-workers in their product units. In contrast, they know the wages of those in other product units only 7.2% of the time. This is consistent with the presumption that workers primarily compare their pay with that of co-workers in their own unit.

For a given absolute pay level, output declines by 0.33 standard deviations (22%) on average when a worker is paid less than both his co-workers.⁹ This is accompanied by a 12 percentage point decrease in attendance. In addition, holding fixed the level of absolute pay, we find little evidence that performance improves if a worker is paid more than his peers. In fact, in units where there is *Pay disparity*, the highest and median wage workers also have substantively lower attendance than their counterparts on *Compressed* units (10 and 13 percentage points, respectively). We also see no evidence of positive impacts on output for these workers. Overall, we estimate that workers give up 9.3% of their earnings to avoid a workplace where they are paid differently than their peers. The negative effects of *Pay disparity* on labor supply persist over the duration of the employment period and appear to strengthen in later weeks.

Perceived justifications play an important role in mitigating these morale effects. Each of our two sources of variation in perceived justifications yields the same pattern of effects. First, when co-workers' baseline productivity levels are farther apart—so that differences in productivity swamp differences in wages—workers in *Pay disparity* units have the same attendance and output as their counterparts on *Compressed* units. Second, in production

⁷Workers are randomly assigned to units upon recruitment, making this a random source of variation.

⁸We quantified the observability of each of the ten tasks ex ante in a pilot. A different sample of workers—all of whom were in product units with *Compressed* wages—were asked after three weeks of work to rank their output relative to that of their unit-mates. We use the mean accuracy of these responses for a given task as the observability value for that task. This observability value is uncorrelated with other differences across tasks, such as output dispersion or growth. In the experiment, we stratify wage treatments by production task, ensuring variation in task observability within each treatment cell.

⁹This estimate compares low-rank workers in *Pay disparity* (wage w_{Low}) with low-rank workers in *Compressed* units where everyone is paid w_{Low} . In general, all relative pay effects are identified off pairwise comparisons of workers in *Pay disparity* with workers in *Compressed* units who have the same absolute wage level and same baseline productivity rank.

tasks where workers can easily see that their higher paid peers are more productive than themselves, we also find no negative effect of pay disparity. These findings indicate that fairness violations from *Pay disparity* are only triggered when the rationale for pay differences is not extremely clear to workers.

If relative pay effects operate through emotions such as resentment or envy of co-workers, this could generate hostility and reduce social cohesion among unit members. This could help explain, for example, why workers in *Pay disparity* give up substantial earnings to avoid coming to work—even those who are paid relatively more than their peers. We designed two cooperative endline games, conducted on the last day of work, to investigate effects on group cohesion. The games have no benefit for the firm, ruling out motives such as reciprocity or beliefs about the firm. In both games, workers are paid group piece rates for performance. This enables us to test whether *Pay disparity* induces a group dynamic that subsequently affects co-workers' ability to cooperate in their own self-interest.

In the first game, workers are organized into their product units, and build towers out of raw materials. Each unit is paid a piece rate based on the height of its tower. *Pay disparity* units build towers that are 17% shorter on average than those with *Compressed* wages. However, in units where pay differences are justified—based on baseline productivity differences or task observability—*Pay disparity* units perform as well as *Compressed* units.

In the second set of activities, workers play cooperative puzzle games in pairs. Workers are randomly paired with someone from either their own product unit or from another unit. *Compressed* unit workers perform better on these games when paired with someone from their own product unit than when paired with a “stranger” (someone from another unit). In stark contrast, *Pay disparity* workers perform 28% worse when they are paired with someone from their own unit than with a stranger. In addition, when in mixed pairs, there is no evidence that *Pay disparity* workers perform worse than *Compressed* pay workers—their decrease in performance arises only when paired with someone from their own unit. Finally, when pay disparity is clearly justified, we cannot reject that *Pay disparity* workers perform similarly to *Compressed* workers—regardless of with whom they are paired.

Data from an endline survey also suggests a decrease in social cohesion. Specifically, after over a month of working together, *Pay disparity* workers report fewer social network connections—defined as a willingness to borrow or lend, seek or give advice, or visit one-another's homes—than *Compressed* workers. Survey responses also provide suggestive evidence about fairness perceptions. Relatively lower paid workers in *Pay disparity* units are more likely to state that their wages were set unfairly in relation to their unit-mates than low rank workers in *Compressed* units. In contrast, *Pay disparity* workers with relatively higher wages do not believe their wages are unfair; however, they are substantially less likely to report being happy at endline. While only suggestive, this is consistent with a decrease in attendance for higher-paid workers being driven by resentment or hostility in the work environment.

This study builds on the literature on relative pay comparisons in the workplace. Two recent field experiments have examined relative pay concerns. First, [Card et al. \(2012\)](#) document that University of California employees report higher job dissatisfaction on surveys when they find out they are paid less than their co-workers. Second, [Cohn et al. \(2012\)](#) show that random relative pay cuts matter more than absolute pay cuts for effort; these effects persist strongly over a six-hour period. Our results are consistent with those of both studies.¹⁰ In addition, our work relates to the broader literature on fairness preferences and effort provision, such as gift-exchange ([Akerlof 1982](#)).¹¹

Our study advances this literature. We find substantial impacts of relative pay comparisons on output and labor supply, with workers giving up earnings to avoid a workplace with pay disparity. We also document deleterious impacts of pay inequality on group cohesion and cooperation. This dimension has been emphasized in social psychology and sociology—albeit with limited identified evidence—but largely ignored in the empirical behavioral economics literature to date. In addition, in our experiment, relative wages are not arbitrary, but rather reflect productivity differences; this is important if justifications can undo fairness violations ([Falk et al. 2008](#), [Bracha et al. 2015](#)) and matches why pay disparity may arise in the labor market. We provide the first piece of field evidence that perceived justifications play an essential role in mediating morale effects.¹² This has bearing on understanding, for example, why wage compression may arise in some settings or occupations and not in others. Finally, workers make decisions for a job from which they derive their primary source of income over the one-month study period. Given low baseline levels of income and employment, the decision to give up earnings is therefore meaningful. The time horizon also indicates that impacts do not disappear once the novelty of wage changes wears off ([Gneezy and List 2006](#), [Levitt and List 2007](#)).

¹⁰In addition, [Cullen and Pakzad-Hurson \(2015\)](#) use a field experiment with TaskRabbit workers to show that relatively lower pay lowers productivity. Two recent studies also examine relative pay using natural experiments: [Rege and Solli \(2013\)](#) find an increase in job separations after pay disclosures in Norway, and [Dube et al. \(2015\)](#) document an increase in worker quits following differential pay increases in the US. In earlier work, laboratory studies have used gift exchange games to explore effects of different pay levels ([Charness and Kuhn 2007](#), [Gatcher and Thoni 2010](#), [Bartling and von Siemens 2011](#), [Bracha et al. 2015](#)) and ranks ([Brown et al. 2008](#), [Clark et al. 2010](#), [Kuziemko et al. 2014](#)), with mixed results. Finally, a number of other papers are also consistent with a relationship between relative pay and worker satisfaction or behavior ([Levine 1993](#), [Pfeffer and Langton 1993](#), [Clark and Oswald 1996](#), [Hamermesh 2001](#)). A set of related studies examines the effects of revealing performance ranks to piece rate workers, indicating that social comparisons may matter for effort even in the absence of incomplete contracting ([Blanes i Vidal and Nossol 2011](#), [Bandiera et al. 2013](#), [Barankay 2016](#)).

¹¹In contrast to our focus on relative pay, gift exchange posits reference dependence in a worker's own absolute wage (past or expected). A large body of work finds gift exchange in the lab (e.g., [Fehr et al. 1993](#)). Field experiments have found little evidence for sustained positive reciprocity and mixed evidence for negative reciprocity over one-day periods ([Gneezy and List 2006](#), [Kube et al. 2013](#), [Esteves-Sorenson and Macera 2015](#), [DellaVigna et al. 2016](#)). [List \(2009\)](#) and [Charness and Kuhn \(2011\)](#) present literature reviews.

¹²Consistent with this, [Pfeffer and Langton \(1993\)](#) document that wage dispersion within academic departments is negatively correlated with faculty satisfaction and co-authorships. They argue this correlation is weaker in fields with more developed scientific paradigms and productivity- or experience-based pay.

While we find that pay disparity alters labor supply and group cohesion in our setting, one cannot draw conclusions about optimal pay structure. One potential benefit to firms of differential pay is dynamic incentives: workers know that if they work hard now, it could lead to higher pay in the future. Our design intentionally shuts down this channel because after the baseline period, there is no further chance of wage changes (or even of future employment). This is important for our goal: cleanly isolating whether workers care about relative pay. It is also a realistic feature of seasonal and other contract jobs—a common form of employment among the workers in our study. However, in choosing the optimal pay structure, a firm would weigh any potential costs of differential pay (e.g. morale effects) against the potential benefits (e.g. dynamic incentive or selection effects). Our findings indicate that workers’ relative pay concerns could affect this calculus.

Section 2 below presents a brief framework. We lay out the experiment design in Section 3. Section 4 outlines the empirical strategy and Section 5 presents the results. In Section 6, we discuss potential alternate explanations and threats to validity. Section 7 concludes.

2. FRAMEWORK

We adapt the framework of DellaVigna et al. (2016), in which workers’ social preferences affect effort provision. We modify their approach to allow peer wages to affect morale.

We assume that a worker i receives a wage offer w_i from the firm and makes two decisions: a) whether to work, $s_i \in \{0, 1\}$; and b) if $s_i = 1$, how much effort $e_i \geq 0$ to exert. Effort is not contractible by the firm. If the worker chooses not to work (i.e., $s_i = 0$), then he receives a stochastic outside option, $R_i = R + \varepsilon_i$. His payoff from working is:

$$V(w_i, w_R, e_i; \theta_i) = w_i - c(e_i; \theta_i) + M(w_i, w_R) e_i,$$

where $c(e_i; \theta_i)$ is a convex effort cost, θ_i is a worker-specific productivity parameter, and $M(\cdot)$ is a morale effect term that depends on the worker’s own wage and a reference value, w_R . The worker chooses to supply his labor if the value from working is greater than his outside option, that is: $s_i = 1 (\arg \max_e V(w_i, w_R, e_i; \theta_i) \geq R_i)$. Note that this formulation leads to potential effects on both effort e_i and attendance s_i . A decrease in $M(\cdot)$ leads to lower benefits from exerting effort and will decrease e_i . Given the restriction $e_i \geq 0$, if $M(\cdot)$ falls far enough, the participation constraint will no longer be satisfied, leading to $s_i = 0$.¹³

We conceptualize relative pay concerns as reference-dependence in utility, where the worker’s reference value is a function of co-worker pay (and productivity):

$$w_R = r(\mathbf{w}_{-i}, \theta_i, \theta_{-i}).$$

We incorporate this into the worker’s morale effect term:

$$M(w_i, w_R) = \alpha f(w_i - w_R | w_i < w_R) + \beta f(w_i - w_R | w_i > w_R) + g(w_i).$$

¹³An alternate way to generate effects on s_i would be to separately allow relative pay differences to shift $V(\cdot)$ in a way that is independent of e_i .

This formulation captures the idea that morale is at least partially determined by relative wage differences. The function $f(\cdot)$ is monotonically increasing in the gap between the worker’s wage and his reference value, and $f(0) = 0$. We allow for asymmetric effects of differences between own and reference wages. Specifically, if the worker is paid less than the reference wage, $w_i < w_R$, the morale component of his utility increases by $\alpha f'(w_i - w_R)$ per unit of effort, relative to the case where $w_i = w_R$. Consequently, the direction of the effect on both e_i and s_i is pinned down by the sign of α . Similarly, if $w_i > w_R$, the direction of the effect on e_i and s_i (relative to when $w_i = w_R$) is pinned down by the sign of β . Finally, $g(w_i)$ captures peer-independent contributors toward morale, such as gift exchange.¹⁴

Note that this is a reduced form specification of worker utility. The parameters α and β capture both innate preferences and social dynamics in the work place. Prior work conceptualizing relative pay comparisons predicts that $\alpha < 0$, that is, individuals dislike being paid less than their peers (i.e., [Adams 1963](#), [Akerlof and Yellen 1990](#)). The prediction on the sign and relative magnitude of β , however, varies. Preferences for status or advantageous inequality generate $\beta > 0$. However, if workers have inequity aversion (e.g., [Fehr and Schmidt 1999](#)), or if pay disparity generates strife or a breakdown in social relations in the workplace, then this could generate $\beta < 0$ (e.g., [Shaw 2014](#)).

In this simple framework, changes in $1(w_i < w_r)$ and $1(w_i > w_r)$, holding fixed w_i , will affect both the probability of accepting work and effort if α and β are non-zero. In the experiment, we generate variation in \mathbf{w}_{-i} (i.e. co-worker wages). This enables us to construct test cases where the value of $1(w_i < w_r)$ or $1(w_i > w_r)$ is clear. For example, we posit that $w_i < w_r$ if a worker is paid strictly less than his co-workers, but not when he is paid the same as all his co-workers. We infer the sign of α and β by examining the impact on labor supply of changes in $1(w_i < w_r)$ and $1(w_i > w_r)$, respectively.

In addition, we use experimental manipulations to test whether workers care only about wage differences in levels (i.e. $w_R = r(\mathbf{w}_{-i})$), or whether relative productivity also enters into workers’ assessments of pay equity ($w_R = r(\mathbf{w}_{-i}, \theta_i, \theta_{-i})$). To accomplish this, we vary the composition of co-workers, in terms of their productivity. Conceptually, this enables us to change θ_{-i} , holding fixed θ_i , generating orthogonal variation with respect to w_i and \mathbf{w}_{-i} . The impact of this on behavior provides insight into workers’ underlying notions of equity.

3. EXPERIMENTAL DESIGN AND PROTOCOLS

3.1. Experimental Design. We explore relative pay concerns using a field experiment with manufacturing workers. In our setting, workers are paid a flat daily wage for attendance. Consequently, as in the framework above, extensive margin changes in labor supply (i.e. attendance) affect earnings, whereas there is incomplete contracting on effort.

¹⁴Alternately, $g(w_i)$ could capture other drivers of effort including monitoring norms for effort provision.

3.1.1. *Reference group.* We must first be able to isolate, for each worker, a clear reference group of peers for pay comparisons. In the factories, we organize workers into production units of 3 workers each. All workers within a unit produce the *same* exact product, while every unit in a factory produces a *different* product. Production is strictly an individual activity. Because each worker’s two unit members are the only other people at the factory making the same product, they constitute the most salient reference group for wage comparisons. We empirically validate this intuition in Section 5.1.¹⁵

3.1.2. *Wage treatments.* We construct wage treatments to generate variation in co-worker pay, holding own absolute pay fixed. Using baseline productivity data (see below), we rank each worker as the lowest, medium, or highest productivity worker within his respective unit. Each unit is then randomized into one of four wage structures, as shown in Table 1:

- *Pay disparity:* Each worker is paid according to his baseline productivity rank within the unit. In each unit, the lowest rank worker receives w_L , the middle rank receives w_M , and the highest rank worker receives w_H . Thus, all three co-workers in the unit always have a different wage from each other.
- *Compressed Low:* All unit members are paid the same daily wage of w_L .
- *Compressed Medium:* All unit members are paid the same daily wage of w_M .
- *Compressed High:* All unit members are paid the same daily wage of w_H .

The difference between each of the three wage levels is fairly modest: less than 5%.

This design enables us to compare workers who have the same expected productivity and are paid the same absolute wage, but differ in their wage relative to their co-workers. For example, low rank workers in *Compressed Low* are paid w_L —the same wage as the other workers in their unit. The low rank workers in *Pay disparity* also receive an absolute wage of w_L , but earn strictly less than their co-workers. If being on a unit where one is paid less than one’s peers reduces morale, low rank workers in *Pay disparity* units will have worse performance than their counterparts in *Compressed Low* units. The design generates three such pair-wise comparisons, illustrated in each respective row of Table 1.

3.1.3. *Perceived justifications.* We incorporate two sources of variation in the perceived justification for pay differences. First, while wage differences are fixed at discrete intervals, underlying baseline productivity differences among co-workers are continuous. This generates variation in the ratio of {wage difference}/{productivity difference} within and across wage treatments. Because workers are randomly organized into their product units, these baseline productivity differences are exogenously determined.

Second, we exploit the extent to which workers can observe co-worker productivity. Ten different products are manufactured in the factories. These products differ substantively in how easy it is to observe the output of one’s unit-mates—for example, whether piles of

¹⁵This view is consistent with Card et al. (2012), who find that University of California workers are most likely to compare pay with those in the same department.

output build up quickly, clearly highlighting production differences across co-workers. To ex ante quantify the observability of each production task, we used a pilot that was conducted with a different sample of workers before the experiment began. These pilot participants were paid equal wages and asked to rank their productivity relative to their co-workers at the end of three weeks. We use the accuracy of responses to this exercise as a proxy for output observability for each product. We stratify wage treatments by production task, enabling us to test for the effects of observability within and across wage treatments.

We examine whether variation in perceived justifications mediates the effects of *Pay disparity*. This provides insight into workers' notions of what constitutes fairness in wages (e.g. Adams 1963). Figure 2 summarizes the sources of variation in the experiment design.

3.2. Context, Protocols, and Social Cohesion Tests.

3.2.1. *Context.* The experiment takes place across three factory sites in India, located in semi-rural areas surrounding the city of Bhubaneswar, Odisha. Workers are employed in low skill manufacturing of local products such as rope, brooms, incense sticks, candle wicks, disposable plates, and floor mats. Using a partnership with local contractors—who provide expertise in production, set quality and training standards, and sell all output in the local wholesale market—we manage the operational implementation of the experiment.

In this area, laborers work in agriculture during peak planting and harvesting, which comprise about four months of the year. In the remaining lean agricultural months, they typically seek short-term contract employment in non-agricultural jobs, such as manufacturing and construction. In our experiment, workers are employed in such seasonal contract jobs full-time for one month. During the period of employment, the job constitutes the primary source of household earnings. These earnings are particularly valuable because of high unemployment in the agricultural lean season. For example, before recruitment, the average worker reports involuntarily unemployment 15 of the past 30 days (Table 2).

In accordance with the typical pay structure in the area, workers are paid a flat daily wage for each day they come to work.¹⁶ There are no strict production minimums, but workers can be fired for excessive absences (more than three days in a row) or disruptive behavior. All workers have experience with flat daily wage pay. In addition, they are familiar with the concept of performance pay—for example, because piece rates are often used during harvesting—and about 45% have worked under piece rates themselves.¹⁷ All workers in the sample are adult males; see Table 2 for demographic descriptives.

¹⁶In general, firms in the area pay both flat daily wages and piece rates for low skill manufacturing. This often corresponds to differences in quality grades, because quality for these tasks is difficult to quantify and costly to monitor. Flat wages are more common across firms as a whole.

¹⁷In addition, factories will sometimes pay workers differential flat wages based on experience. In construction, workers of different skill grades will usually receive differential flat wages.

3.2.2. *Timeline & Protocols.* The experiment was run in 14 one-month rounds across three factory sites. Figure 3 details the within-round timeline. For each round, workers are recruited from a different local labor market (set of villages) within a 15 km radius of the factory. The position is advertised as a one-time contract job. Each round requires hiring 30 workers; if more than 30 workers apply for the job, we randomly select among applicants.¹⁸

On the first day of work, workers are randomly assigned into one of the ten product units, with three workers per unit.¹⁹ Each worker’s unit (and therefore production task) remains the same throughout their employment. The workers within a unit sit together in a designated area—they work together and also eat lunch together in their production area. Each unit is physically separated from other units—often in separate rooms—due to space requirements for each production activity and for storing inputs nearby.

Production is strictly an individual activity—all workers produce their product from beginning to end on their own, with no joint production within units. We hire numerous additional staff, who sit in the factories and count each worker’s individual production of goods of salable quality (e.g. number of individual incense sticks or paper bags) in a separate back storage room. This enables us to obtain daily measures of worker output.

Workers initially undergo a “training period” on how to produce their assigned product. Typically after day four, output has reached a level of quality that can be sold in the market. However, we elongate the training period to two weeks to obtain stable baseline productivity measures for each worker once experience trends begin to flatten out. During training, all workers are paid a daily wage that corresponds to the prevailing daily wage in the area (e.g. Rs. 250/day or ~\$4/day). On the first day of work, they are told that their post-training wage may depend on their productivity during the training period. In addition, on Day 10, each worker is given private feedback on his own productivity rank within his unit (as part of a routine check-in with the manager).²⁰

On Day 14, each product unit is randomized into one of the four wage treatments in Table 1.²¹ For the wage assignments, each worker’s productivity rank within his unit is defined based on performance in the final three days of the training period (Days 11-13).²² All workers get a pay raise relative to the training wage—with w_L , w_M , and w_H

¹⁸Four of the rounds involved fewer than 30 workers; the remaining rounds had 30 workers exactly.

¹⁹On the first day of work, we also conduct a very short baseline survey to capture demographic characteristics including age, literacy, and basic employment history.

²⁰This helps underscore to workers that we are paying attention to productivity. It also helps ensure our subsequent wage treatment effects are not confounded by information revelation about relative ranks.

²¹An implication of our design is that within a factory, different units have differing pay structures and average pay levels. This is not odd since every unit within a factory produces a unique product and is associated with a distinct contractor.

²²The three days before the wage change captures production levels after experience trends have flattened out. Results are robust to dropping the small number of units where ranks are different within this window relative to the full training period.

corresponding to a roughly 4%, 8%, and 12% pay increase, respectively.²³ Consequently, even w_L constitutes a wage premium above workers’ outside option. On Day 14, each worker is told his own post-training wage privately by his manager. Managers maintain pay secrecy, so that workers are not informed of their co-workers’ wages. Consequently, to the extent workers learn about relative wages, it is through self-disclosure among co-workers.²⁴

After wage treatments are implemented, workers continue to work in their assigned units for an additional 20 days—allowing us to observe subsequent output and attendance. On the final day (Day 35), we conduct tests for social cohesion and administer an endline survey.

To ensure understanding of the timeline, each worker is given an individualized calendar on Day 1. The calendar highlights the date that training will end (and wages will increase) as well as the last day of the work contract. On Day 1, the training wage amount (Rs. 250) is filled in on the calendar for each of the training days, with the remainder of the days left blank. On Day 14, the worker’s post-training wage is written on the calendar for each of the remaining days until the end of employment. Each worker sees his individual calendar each day upon arrival to work. This makes it clear to workers that: there will be a wage increase on Day 14 (which may depend on baseline productivity); the post-training wage is fixed, with no further wage changes; and the employment period ends on a pre-set date.

3.2.3. Endline Tests for Social Cohesion. As discussed above, the impact of pay disparity can operate through both individual preferences over relative pay as well as group-level dynamics such as social cohesion and conflict. To the extent that pay disparity generates feelings such as envy or resentment, these two channels are inherently related.

To explicitly test for effects on social cohesion, we design a set of cooperative games on the last day of employment, as part of a “fun farewell”. In these games, workers are paid piece rates for group-level performance. There is no benefit to the firm from the games—ruling out motives such as retaliation against the firm in driving behavior.

We implement two types of activities, which we describe in detail in Section 5.5. In the first activity, workers play in their assigned product units, enabling us to examine general effects on team cooperation. In the second activity, we induce variation in whether workers are paired with someone from their own product unit versus a “stranger” (someone from another unit). This enables us to distinguish general changes in worker disgruntlement from dynamics that are specific to working with one’s co-unit members. These activities allow us to test whether *Pay disparity* affects workers’ ability to cooperate in their own self-interest.

Finally, we use an endline survey to obtain additional suggestive evidence on mechanisms. We ask workers about social ties at endline—for example, whether, after working together

²³For example, if the training wage (corresponding to the prevailing wage in the area) was Rs. 250/day, then w_L , w_M , and w_H would be Rs. 260, Rs. 270, and Rs. 280, respectively.

²⁴It would be odd for an employer to publicly list wages. We verify in Section 5.1 that workers shared wage information. To the extent that sharing was imperfect, this will dampen our observed treatment effects.

for over a month, they would visit the home of or borrow money from their co-workers. We also collect self-reports on the perceived fairness of wages and worker happiness.

4. EMPIRICAL STRATEGY

4.1. Outcomes. Attendance is a binary variable capturing whether worker i is present on day t . We also record the continuous, total output of salable products by each worker i on each day t .²⁵ We code raw production as zero when a worker is absent. We standardize output within each task—using the mean and standard deviation for that task in the final three days of the training period.²⁶ This enables us to pool across tasks and measure output using consistent units: standard deviations from the mean.

Table 2 presents baseline values of attendance and standardized production in the full training period. Attendance is approximately 95%. Both baseline attendance and productivity are balanced across the *Compressed* and *Pay disparity* units. Appendix Figure A5 documents, within *Compressed* units, both the stability of relative productivity ranks over time as well as the stability of output levels even in the post period.

4.2. Empirical Specification. To test our key predictions, we compare outcomes between individuals in the *Pay disparity* and *Compressed* units, holding fixed a worker’s production rank and wage. Recall from Table 1 that the most direct comparisons are between the Low rank *Pay disparity* worker with the Low rank *Compressed Low* worker, the Medium rank *Pay disparity* worker with the Medium rank *Compressed Medium* worker, and the High rank *Pay disparity* worker with the High rank *Compressed High* worker. We refer to this set of six rank-treatment cells, which are circled in Table 1, as the “relevant group”. To use all of the variation in our experimental data, we use a differences-in-differences strategy that incorporates the pre-period production information.

The most basic differences-in-differences approach restricts the sample to the “relevant” group of six rank-treatment cells:

$$\begin{aligned}
 y_{it} = & \alpha_1 [Post_t \times PayDisp_i \times Low_i] + \alpha_2 [Post_t \times PayDisp_i \times Med_i] & (4.1) \\
 & + \alpha_3 [Post_t \times PayDisp_i \times High_i] + \alpha_4 [Post_t \times Low_i] + \alpha_5 [Post_t \times Med_i] \\
 & + \alpha_6 [Post_t \times High_i] + \lambda_i + \tau_t + \eta_1 x_{kt} + \eta_2 x_{kt}^2 + \varepsilon_{it}.
 \end{aligned}$$

In all specifications, i indexes the worker, k indexes the task, and t indexes the day-round. $Post_t$ is a binary indicator that equals 1 on the days after the wage treatment takes effect within the round. $PayDisp_i$ is an indicator for being a member of a *Pay Disparity* unit. The variables Low_i , Med_i and $High_i$ are binary indicators for productivity rank, as defined in Section 3.2.2. Any time-invariant unit or worker characteristics are absorbed by

²⁵During training, workers learn how to create output so that it meets the quality standards required for it to be salable. In Appendix Table A2 we explore treatment effects on production quality.

²⁶Recall this corresponds to the days used to compute productivity ranks. Standardized output over the full training period is 0.2 standard deviations lower on average than during these final three days (Table 2).

the worker fixed effects, λ_i , while any time trends are captured by the day-by-round fixed effects, τ_t . Finally, x_{kt} and x_{kt}^2 allow for task-specific quadratic experience trends.

The key treatment effects of interest are captured by α_1 , α_2 , and α_3 . The coefficient α_1 measures the average change in outcomes in the post period (relative to the training period) for Low rank workers in *Pay disparity* units relative to those in *Compressed L* units. Similarly, α_2 captures the average treatment effect for Medium rank workers in *Pay disparity* vs. *Compressed M* units. α_3 captures the average treatment effect for High rank workers in *Pay disparity* vs. *Compressed H* units.

Note that by restricting the sample to the “relevant” workers, the above specification ignores observations from the “irrelevant” workers (the complement of the “relevant” group). These additional observations can be used to help estimate the controls: round-by-day fixed effects and experience trends. To improve precision, we augment Equation (4.1):

$$\begin{aligned} y_{it} = & \alpha_1 [Post_t \times PayDisp_i \times Low_i] + \alpha_2 [Post_t \times PayDisp_i \times Med_i] & (4.2) \\ & + \alpha_3 [Post_t \times PayDisp_i \times High_i] + \alpha_4 [Post_t \times Low_i] + \alpha_5 [Post_t \times Med_i] \\ & + \alpha_6 [Post_t \times High_i] + Irrel'_{it}\theta + Neigh'_{it}\gamma + \lambda_i + \tau_t + \eta_1 x_{kt} + \eta_2 x_{kt}^2 + \varepsilon_{it}. \end{aligned}$$

We define $Irrel_i$ as an indicator for whether a worker is in the “irrelevant” group (not among the six “relevant” rank-treatment cells). In Specification (4.2), $Irrel'_{it}$ is a vector of interactions of $Irrel_i$ with $Post_t$ and each rank:]. These interactions fully absorb the change in outcomes for “irrelevant” workers in the post period; the main coefficients of interest— α_1 , α_2 , and α_3 —are therefore still estimated off of only the “relevant” group. Consequently, this improves statistical power without affecting identification.

In addition, we define $Neigh_i$ as an indicator for whether a unit is located in the same physical room or production space as another unit, facilitating conversation across units.²⁷ In Equation (4.2), the vector $Neigh'_{it}$ contains interactions of $Neigh_i$ with $Post_t$, the treatment and rank of worker i , and the treatment of $Neigh_i$.²⁸ Thus, the coefficients α_1 , α_2 , and α_3 are identified off units with no neighbors. In neighborless units, it is extremely clear that one’s two product unit co-workers comprise the relevant reference group.²⁹

To maximize statistical power, we use Equation (4.2) as our main specification. Our results are similar if we exclude “irrelevant” workers or omit the $Neigh'_{it}$ controls. Appendix Table A1 documents robustness to these and other changes to our estimating equation.

We also examine the pooled effect of *Pay disparity* across Low, Medium, and High ranks:

$$y_{it} = \alpha [Post_t \times PayDisp_i] + Irrel'_{it}\theta + Neigh'_{it}\gamma + \lambda_i + \tau_t + \eta_1 x_{kt} + \eta_2 x_{kt}^2 + \varepsilon_{it}. \quad (4.3)$$

²⁷Usually, units were in separate physical rooms so this wasn’t possible. Seating charts available on request.

²⁸Due to random assignment, $Neigh_i$ is exogenous to both treatment status and worker characteristics.

²⁹Our experiment isn’t powered to investigate whether workers respond to the wage levels of neighboring units; such a question would require sufficient variation in the number of instances of neighboring units and in their wages. We leave this for future work.

The main coefficient of interest, α , is still identified off of comparisons within the “relevant” cells. It therefore captures the average treatment effect of *Pay disparity* compared to *Compressed* pay (holding fixed absolute pay) across all worker ranks.³⁰

Finally, we test whether perceived justifications mediate the effects of pay inequality. As described in Section 3.1, we built two sources of heterogeneity in perceived justifications into the design: productivity differentials and output observability. To examine their role, we fully interact the variables encoding treatment status, rank, and $Post_t$ in Equations (4.2) and (4.3) with each of the justification measures.

5. RESULTS

5.1. Knowledge of Co-worker Wages. Given that managers maintain pay secrecy throughout the experiment, a necessary condition for the wage treatments to affect productivity is that workers actually learn co-workers’ wages. We verify this using the endline survey. Panel A of Table 3 reports that among *Compressed* unit workers, 95.8% correctly report the wage of at least one co-worker, and 90.9% correctly report the wages of both co-workers. Workers from *Pay disparity* units are less accurate—with 87.1% and 74.2% of respondents reporting the correct wage of at least one or both co-workers, respectively.³¹ These findings indicate that overall, there is substantial learning within production units, but *Compressed* unit workers have statistically significantly more accurate beliefs. This suggests that it may be awkward for co-workers to discuss pay under pay dispersion.

For the treatments to have power, it is also important that workers consider their product unit as their reference group.³² In the endline survey, each worker was also asked to list the wages of each member of the incense stick, candle wick, floor mat, and paper bag production units—each of which could be assigned to either *Pay disparity* or *Compressed* pay, depending on the round.³³ In Panel B of Table 3, we examine how well individuals learn about the wages of other units. Strikingly, the majority of workers have no opinion about the wages of any of the workers on these other units. Approximately 87% of respondents are not willing to wager a guess about the wages on *Compressed* units, and only 8.6% of respondents can accurately list all of the wages. It appears even harder to learn about the wages on *Pay disparity* units, with only 1.9% of respondents able to accurately list the wages of all workers on such units. These findings support our assumption that workers primarily compared their wages with their two other unit-mates.

5.2. Effects of Pay Disparity. Figure 4 provides an overview of the patterns in the underlying production data. It plots average standardized production on each day for each

³⁰All variables are defined exactly as in Equation 4.2. Again, $Post_t$ is absorbed by the time fixed effects, τ_t .

³¹The accuracy of beliefs is similar across Low, Medium, and High rank workers. It is also worth mentioning that only four workers on *Pay disparity* units believe that both their co-workers earned the same wage.

³²If workers instead compared themselves to those in other units, this should decrease the potency of our treatments and make it harder to find treatment effects.

³³For this exercise, we exclude observations where individuals were asked about their own units.

of the 3 sets of “relevant” pairwise comparisons. Among Low rank workers in the baseline period, production in the *Compressed Low* and *Pay disparity* units shows a common trend as workers gain experience. The treatment (“post”) period begins on day 0, when each worker is told his post-training wage. Within about 5 days (i.e. by the first payday following the wage change), differences in output start to emerge, with workers on *Pay disparity* units (who are paid less than their peers) reducing output relative to the *Compressed L* units. This delay in the onset of treatment effects is consistent with non-immediate diffusion of pay information among unit members. In addition, there is no evidence of positive effects of being paid relatively more than one’s peers: High rank workers in *Pay disparity* units appear to perform worse than those in *Compressed H* units. This pattern also generally holds for Medium rank workers in *Pay disparity* versus *Compressed M* units.

Panel B of Table 4 presents regression estimates of treatment effects separately for each rank. Low wage workers in *Pay disparity* units reduce output by 0.385 standard deviations relative to those in *Compressed Low* units (significant at the 1% level). This effect is robust to the inclusion of individual fixed effects (Col. 2, our preferred specification corresponding to Equation (4.2)), indicating that $\alpha_1 < 0$. The 0.332 standard deviation decrease in Col. 2 is equivalent to a 22% reduction in output relative to the *Compressed Low* (control) treatment mean. In addition, we find little evidence that performance improves when workers are paid more than their peers. In fact, there are large, but statistically insignificant decreases in production for High rank workers in *Pay disparity* units relative to their counterparts on *Compressed High* units. We find similar results for the Medium rank workers on the *Pay disparity* units relative to those in *Compressed Medium* units. In contrast, we see little evidence for changes in quality (Appendix Table A2).³⁴

In addition, *Pay disparity* reduces attendance for all ranks. Treatment effects for Low, Medium, and High wage earners are -12, -12.9, and -10.4 percentage points, respectively (all significant at 5% level, Col. 4), off a base of 94% attendance among *Compressed* units.

Because we code a worker’s raw production as zero when he is absent, treatment effects on standardized production combine the extensive margin attendance effects with any intensive margin effort effects. We attempt to decompose these effects in two different ways. First, we regress output conditional on positive attendance in Col. 5. Conditional on coming to work, Low rank workers in *Pay disparity* units produce 0.2 standard deviations less than those in *Compressed Low* units (p-value 0.077). Because this is a selected sample, this result must be interpreted with caution. However, if the most aggrieved workers are those who decide to skip work, this coefficient is likely to underestimate the true intensive margin effect. Second, we use a back-of-the-envelope calculation. Mean output conditional on attendance for Low rank workers in *Compressed Low* units is 1.64. The attendance effect of

³⁴It is difficult to quantify quality in our setting—one of the reasons for the prevalence of flat daily pay. For a subset of rounds, tasks and days, we had management rate the quality of each worker’s output for that day on a scale of 1-5. We do not see evidence for a change in these subjective quality ratings. It is possible, however, that this subjective measure is too crude and noisy to enable us to examine such effects.

Pay disparity for the Low rank is -12 percentage points. If the full effect on production were coming through attendance, we would predict an output decrease of $-0.12 \times 1.64 = -0.197$ standard deviations. This corresponds to about half (59%) of the total effect on output. While both these decomposition approaches have their limitations, they suggest that in our setting, being paid less than one’s peers likely negatively affects effort as well as attendance.

In contrast, the effects for the Medium and High wage earners appear to operate primarily through attendance. For these workers, there is little evidence of conditional intensive margin effort effects in Col. 5. Furthermore, the back-of-the-envelope decomposition suggests that the attendance effects can fully account for the (statistically insignificant) coefficients on standardized output. This suggests that Medium and High wage earners are less likely to show up to work, but there is no evidence that they reduce effort while at the workplace. While we did not have strong ex ante predictions for these workers, these effects may indicate a reduced desire to work alongside disgruntled co-workers—consistent with resentment or hostility at work. We further explore this idea below.

Given that the treatment effects of *Pay disparity* are negative in sign for all three ranks, the pooled specification of Equation 4.3 may provide improved statistical power. These results are presented in Panel A. Overall, individuals in *Pay disparity* units decrease production by 0.311 standard deviations and reduce attendance by 11.5 percentage points relative to individuals of the same rank and absolute wage level in *Compressed* units. We also report pooled regression results for all of the analyses that follow.

Finally, in Appendix Table A1, we show that the results in Table 4 are robust to a number of alternative specifications, including dropping “irrelevant” workers, omitting controls for the presence of other units in close proximity, and estimating treatment effects using only post-treatment observations. We also provide suggestive evidence in Appendix Figure A2 that attendance effects are driven both by increases in the number of days missed by the chronically absent and in the likelihood of workers missing only a few days of work.

Earnings Consequences of Attendance Effects. Because employees are not compensated when absent, the attendance effects indicate that workers are willing to give up full-time earnings to avoid a workplace where they are paid differently than their peers. However, the foregone factory earnings may overstate this value if workers can find outside casual employment when absent. To bound this value, we use survey data from individuals who applied for the factory jobs, but were not hired in the randomization protocol. These respondents report finding paid work on only 23.7% of days. Their average total cash plus in-kind wage, conditional on working, is Rs. 271. Thus, in the experiment, workers are paid an average daily wage premium of Rs. 39.

If absentee workers find outside jobs at similar rates as the non-selected workers, then the attendance treatment effect of *Pay disparity* causes workers to forego an average of

Rs. 403 in earnings.³⁵ This corresponds to 9.3% of potential earnings in the factory in the post period. Furthermore, under the extremely conservative assumption that absentees find work at the prevailing wage every day they are absent, they are still giving up the average wage premium of Rs. 39 each day. These foregone earnings amount to 7.1% of wages earned over the same period by individuals who were not selected for the factory jobs.³⁶

Effects Across Time. Table 5 examines the evolution of treatment effects over time. We separately estimate Equation 4.2 over three parts of the post period—the days before the first payday, between the first and second payday, and the days after the second payday—after wage treatments took effect.³⁷ Focusing on the pooled regressions in Panel A, the effects of *Pay disparity* are only evident after the first pay day, post wage change. This is consistent with non-immediate diffusion of wage information. It also matches supervisors’ informal observations that workers were more likely to discuss wages on paydays, which occurred each Friday. Furthermore, we find little evidence that the effects wane over time.

5.3. Perceived Justifications. We use the two (pre-registered) sources of heterogeneity built into our design to investigate if perceived justifications mediate morale effects.

5.3.1. Baseline Productivity Differences. We first examine treatment effects when workers’ higher-ranked peers are substantially more productive than themselves. For each worker, we compute average baseline output conditional on attendance; this corresponds to what is observed by workers before wages change.³⁸ We then compare this baseline productivity average with that of the worker’s next-higher ranked co-worker (i.e. for each low rank worker, the difference with his medium-ranked co-worker; for each medium rank worker, the difference with his high-ranked co-worker). Note that this variable is not defined for high rank workers. Recall that because workers are randomly assigned to units, baseline productivity differentials are exogenously determined.³⁹

We define the productivity difference between a worker and his higher-ranked peer to be “large” when it is above 0.37 standard deviations, the mean difference among the “relevant” *Compressed* workers. In Columns 1-2 of Table 6, we interact this binary variable

³⁵The wage premium is calculated using the average factory cash wage plus an in-kind wage of Rs. 40 (the monetary value of lunch), which is similar to the in-kind wage paid in the market. We assume a wage of 0 on days that workers cannot find work; on days they can find paid work, they only forego the wage premium. We use the pooled treatment effect on attendance of 11.7 percentage points and calculate that *Pay disparity* causes workers to miss 1.64 extra days of work (based on the modal number of workdays in the post period).

³⁶Non-selected workers earn an average of Rs. 899 over this period.

³⁷The full pre-treatment period is included in all regressions.

³⁸Specifically, we use the final six days of training—which are used to convey ranks to workers and compute ranks for wage treatments. Results are robust to using only the final three days for this calculation.

³⁹Baseline output differences among co-workers could also reflect team dynamics or culture. Note this does not affect identification, because baseline characteristics are orthogonal to the wage treatments. However, this potentially broadens the interpretation of the heterogeneous treatment effects presented here. The fact that baseline productivity ranks have stable predictive power in other domains, such as the group cohesion games in Section 5.5, suggests that these ranks do capture some inherent differences among co-workers.

with treatment status. Panel A displays pooled treatment effects for Low and Medium ranks combined, while Panel B decomposes the effects by rank. There are large negative impacts of *Pay disparity* when productivity differences are small: Low and Medium rank workers decrease output by 0.358 standard deviations and attendance by 16.7 percentage points (both significant at 1% level). However, large baseline productivity differences almost fully mitigate these negative effects. When a worker is substantially less productive than his higher-paid peer, the total treatment effect of *Pay disparity* is indistinguishable from zero. Note that the $Post \times Perceived\ justification$ term in Panel A reveals no detectable changes in the behavior of *Compressed* units when productivity differences are large. These heterogeneous effects appear (though under-powered) for each rank separately in Panel B.

In Appendix Table A3, we show that the results are robust to alternate specifications, including adding interactions of the worker’s own baseline productivity level with an indicator for post treatment.⁴⁰ Appendix Table A4 explores a range of productivity thresholds, defined as the percent difference in baseline output between a worker and his next higher-rank co-worker.⁴¹ Productivity differences need to be quite large (on the order of 20%) for us to detect a significant offsetting response to *Pay disparity*. While this is substantially larger than the underlying difference in wages, it may simply reflect observability: it is difficult for workers to recognize that their peers are more productive unless differences are big.⁴²

5.3.2. *Output Observability.* We next check for mediating effects of whether co-worker output is observable. We quantify observability using data from 3-week pilot rounds with a different sample of workers (conducted before the start of this experiment). These workers were all paid *Compressed* wages, and were never told their relative productivity ranks. On the last day of the pilots, we asked workers to rank their co-workers by productivity.⁴³ We use the mean accuracy of these responses for a given task as a proxy for output observability. There is substantial variation across tasks in accuracy rates—from essentially 0 to 0.88 (Figure 5). We define tasks with accuracy rates above the sample median of 0.5 as “observable.”

Columns 3-4 of Table 6 present heterogeneous treatment effects by task observability. The results mirror those in Columns 1-2. Panel A shows pooled effects across all three ranks.

⁴⁰One potential concern is that when a Low rank worker has a large baseline difference with his Medium rank peer, this may mean his absolute productivity level is extremely low—leaving him little room to fall in response to *Pay disparity*, or indicating something else about his “type”. In Cols. 3-4 of Appendix Table A3, we show robustness to adding controls so that the bottom decile of Low rank workers are not used to estimate effects. Cols. 5-6 add full interactions with own baseline productivity; here, the heterogeneous effects are estimated off differences with one’s higher-rank peer, holding fixed own productivity. While this demands a lot of statistical power, the results are robust to the inclusion of even these controls.

⁴¹Note that this scaling of productivity differentials is highly correlated with that used in the main analysis.

⁴²Alternately, this finding matches the predictions of Fang and Moscarini (2005), in which self-serving bias leads workers to assume they are not less productive than peers unless this difference is extremely apparent.

⁴³Two production tasks were added after piloting was completed; we followed a similar procedure with a separate sample of workers to quantify observability for these two additional tasks.

When co-worker output is difficult to observe, *Pay disparity* sharply lowers production and attendance. In contrast, there are no detectable total effects when co-worker output is easy to observe; the triple interactions are large and statistically significant. We also find no evidence that task observability differentially affects the performance of *Compressed* units post wage change. The same patterns emerge in Panel B, when we estimate heterogeneous effects separately by rank. We show in Appendix Table A5 that the results are robust to sequentially dropping each of the ten tasks.

As with any heterogeneous treatment effects analysis, one might be concerned that observability is correlated with other task characteristics. Appendix Figure A3 plots the distribution of output for each production task separately and shows that there is no discernible correlation between, for example, a task’s observability value and output dispersion. More generally, for other correlated task characteristics to pose a problem, they would need generate our pattern of results through a different mechanism. One possible trait might be the amount of social cohesion naturally developed during the training period. In Appendix Table A6, however, we find that observability is not correlated with turnover or attendance in the training period. Further, the fact that the pattern of results for task observability matches that in the productivity differences regressions bolsters our interpretation.

Note that observability and productivity differences derive from independent sources of variation – productivity differences from the randomization of individuals to units and observability differences from the randomization of units to tasks. Unsurprisingly, Appendix Table A7 verifies that there is substantial non-overlap among these two measures. Finally, we find no evidence to suggest that observable tasks simply have more natural productivity dispersion, making it easier to learn about productivity differences. In Col. (4) of Appendix Table A6, we show that observability is uncorrelated with the steepness of the learning curve.

5.3.3. *Aggregate Perceived Justifications.* Finally, for parsimony and power, we define an aggregate perceived justifications indicator. This binary indicator equals 1 if any justification is present—i.e., if an individual has a large productivity difference with his next-higher ranked peer, or he is on an observable production task.⁴⁴ Not surprisingly, heterogeneous treatment effects using this indicator are similar to those above (Table 6, Columns 5-6).⁴⁵

5.4. **Effects of Pay Disparity at the Production Unit-Level.** Our key empirical tests are based on comparing workers who earn the same absolute wage, but differ in co-worker wages. However, our design also enables us to examine the effects of *Pay disparity* at the production unit level. Note that this exercise uses a different source of variation than the analyses presented above; all workers in each production unit are used for identification. Table 7 presents the results.⁴⁶ Panel A reports comparisons between *Pay disparity* units and

⁴⁴For high rank workers, this indicator always takes the value of task observability.

⁴⁵Appendix Table A7 verifies that the productivity difference results and observability results are identified off of different sources of variation. There is substantial non-overlap among these two measures in the sample.

⁴⁶This table uses the same specification as Equation 4.3, but eliminates all controls for “irrelevant” workers.

all *Compressed* units combined. Panel B compares *Pay disparity* units with only *Compressed Low* units. This latter test is strong given that all workers on *Pay disparity* units are paid weakly more than those on *Compressed Low* units. Columns 1-2 present average treatment effects, while Columns 3-8 present heterogeneous treatment effects.⁴⁷

On average, *Pay disparity* units have lower rates of attendance relative to all *Compressed* units pooled and to *Compressed Low* units. The average effects on production are also negative, but not statistically different from zero. When output differences are not large or co-worker output is not observable, compared to *Compressed Low* units, *Pay disparity* units have 0.306 standard deviations lower output and 8 percentage points lower attendance on average (Panel B, Cols. 7-8). However, when pay differences are clearly justified, no negative treatment effects are detectable. We again show that *Compressed* units behave no differently in the post period when there are large differences in productivity or when output is easily observable. These results suggest that in our setting, *Pay disparity* is indeed detrimental to unit-level outcomes—even compared to *Compressed Low* units which have weakly lower absolute wages for all workers—when perceived justifications are not present.⁴⁸

5.5. Effects on Group Cohesion and Dynamics. The above results confirm our initial hypothesis that workers dislike disadvantageous inequality in pay: workers who are paid less than their peers reduce output and attendance. Further, we find striking evidence that even workers in a position of advantageous inequality forego wages to avoid the workplace. Recall that workers spend the entire day with their unit co-workers—sitting in close proximity and also eating lunch together at their stations. If *Pay disparity* fosters resentment between workers, then it might create a negative or even hostile work environment that workers would pay to avoid. Emotions such as acrimony or envy may also erode the ability of peers to cooperate even when it is in their own self-interest. Such forces are of course impossible to examine with just the output and attendance data alone. Consequently, we designed tests to enable us to directly examine social cohesion and cooperation.

5.5.1. *Teamwork Games.* We developed two sets of collaborative games that require teamwork. Workers played these games at endline on the last day of work, as part of a “fun

⁴⁷Note that observability is already defined at the production unit level. We define large productivity differentials at the unit level as an indicator for whether both the Low and Medium rank workers have large productivity differentials with the worker of next higher rank.

⁴⁸While not our primary goal, we can use wage variation across *Compressed* units to explore whether absolute pay affects performance. We examine this using differences-in-differences regressions in Appendix Table A8. If units reciprocate higher wages with higher productivity, then we should expect output to be increasing in pay. Further, any positive peer effects from harder working co-workers should amplify such effects. However, consistent with previous literature, we find no evidence of positive gift exchange on productivity at any time horizon—in the days immediately following wage changes or in the longer-run. Relative to *Compressed Low* units, *Compressed Medium* and *Compressed High* units do increase attendance in the short run—consistent with a simple substitution effect on labor supply. However, even these effects disappear in the full sample period. Note that ours is not an ideal test for gift exchange because workers were all expecting some wage increase, and thus we are not able to identify the worker’s ex ante reference wage.

farewell” day. Workers were paid their full wage plus piece rates based on performance in these games.⁴⁹ Importantly, there was clearly no benefit to the firm from worker effort on the games—ruling out reciprocity as a motive for performance.

In the first game, workers were asked to build a tower with the other members of their assigned product units. Each unit was given the same set of raw materials (e.g., cardboard, pens, rubber bands, playing cards) and was asked to build as tall a tower as possible within a 25 minute time limit. The payment schedule was a linear piece rate for the tower’s height (in cm), paid equally to each unit-member.

Pay disparity units build towers that are 9.376 cm (17.5%) shorter than those of *Compressed units* on average (Table 8, Col. 1). However, as before, these negative effects are concentrated in cases where *Pay disparity* is not clearly justified (Cols. 2-4). When it is, we detect no difference between the performance of *Pay disparity* and *Compressed* units.

Do these results indicate general disgruntlement among *Pay disparity* workers, or are they driven by a breakdown in within-unit cooperation? To understand this, we constructed a second set of games in which workers solved two types of cooperative puzzles in pairs of two. In the “Spot the Difference” game, each person in the pair received a printed sheet with nearly-identical pictures on each sheet. The workers had to compare and circle any difference in the pictures on the two sheets (Appendix Figure A4, Panel A). In the “Symbol Matching” game, each pair member was given a sheet with a grid of symbols. Workers had to match symbols—circling all instances where the same symbol appeared in the same grid position in both of their respective sheets (Appendix Figure A4, Panel B).⁵⁰ In both games, both members of the pair received the same payment, a piece rate for every correct answer that was circled on *both* workers’ respective sheets. For these partnered games, we constructed pairs by reshuffling workers across product units. Each worker played four iterations each of Spot the Difference and Symbol Matching. We randomized pair construction so that in 50% of cases, a worker played with someone from his own product unit, and 50% with someone from another product unit for each of the games.

Table 9 presents effects on the cooperative pairwise puzzles. If both members of a pair are from the same unit, they score 1.105 points (27%) lower if that unit had *Pay disparity* versus *Compressed* pay (Col. 1). The F-test p-value of 0.0536 in Column 1 suggests that *Pay disparity* workers actually perform better playing with a stranger (i.e. someone from another product unit) than playing with a unit co-worker. In addition, note that *Pay disparity* workers do not perform worse than *Compressed* workers in general—the point estimate for at least one *Pay disparity* worker in pair is actually positive (though insignificant). Rather,

⁴⁹At the end of the day, we randomly selected one of the games for each worker, and the worker received his piece rate earnings for that game only. Workers were told we would randomly select one game for payment in advance. Note that these endline games were only played in the final 8 rounds of the experiment (80 product units only).

⁵⁰We thank Heather Schofield for providing us with the Symbol Matching game grids (see Schofield 2014.)

these workers only perform worse when they are paired with a unit co-worker.⁵¹ Finally, Columns 3-4 provide evidence that treatment effects are larger in the absence of perceived justifications, and perceived justifications mitigate the negative treatment effect. However, the mitigating effects of justifications are not statistically distinguishable from zero.

It is worth noting that the worker productivity rankings in the main experiment have predictive power in these endline games. On average, a pair with a Low or Medium rank worker scores 15.8% and 13.8% less, respectively, than a pair with a High rank worker (Col. 2). This suggests that the baseline rankings capture, in part, some stable differences in ability or effort across workers.

Overall, these findings indicate a decrease in *Pay disparity* workers' ability to cooperate in their own self-interest. This does not appear to result from general disgruntlement: low performance only arises when they must work with the other people in their own product units. While in our study, workers engage in individual production, the results suggest that the effects of pay inequality may be exacerbated in settings with team production.

5.5.2. *Network Formation.* Tension within *Pay disparity* groups might impede the formation of relationships outside of work, and could represent an additional manifestation of a lack of group cohesion. In the endline survey, we asked each employee whether he would engage in the following with each co-worker in his unit: seek or dispense advice; socialize; and borrow or lend.⁵² We code a worker as having a network link with a co-worker if he states that he and his co-worker would engage in any of these activities.

In Table 10, the outcome variable in Columns 1-2 measures the number of co-workers with whom the worker reports a link, with possible values of 0, 1, or 2. The outcome in Columns 3-4 indicates whether the worker has a link with both of his unit co-workers. Consistent with the results of the team cohesion games, workers in *Pay disparity* units appear substantially less likely to interact with one another outside of work. On average, *Pay disparity* workers report .23 fewer friends from the production unit (a 30% decrease, Col. 1) and are 12.7 percentage points less likely to be friends with both co-workers (a 48% decrease, Col. 3). These effects are larger and statistically significant in magnitude when no

⁵¹Because the endline games were conducted on the last day of work—when all workers received their final pay for the contract job, attendance was high. In the tower game, units just played with whichever workers were present. In the cooperative puzzles, if a worker was absent, then the person who had been paired with that worker for a pair-game sat out during that round and their score for that pair-game is coded as 0. Appendix Table A9 verifies that the endline game results are not driven by differential absence across units. First, as expected on the last day, treatment status did not affect worker attendance (Cols. 1-2). Cols. 3-4 replicate results for the tower game restricting the sample to units where all three workers. Cols. 5-6 replicate results for only those pair-games where both members of a pair were present. The results are similar to those in Tables 8 and 9.

⁵²This method is based on four of the network elicitation questions used in Banerjee et al. (2013). We also asked about familial relationships. Given that family relationships cannot change over time, we control for whether a worker is related to any co-workers in his unit in the regressions.

justification is present; as before, there are no detectable treatment effects when disparity is clearly justified (Cols. 2 and 4).

5.5.3. *Fairness and Happiness.* In the endline survey, we asked workers to assess how fair their wages were in relation to those of their peers. Panel A of Table 11 presents the results. Overall, *Pay disparity* units are no more likely to believe that their wages were set unfairly in relation to their unit co-workers than the *Compressed* units. However, this pooled result hides substantial heterogeneity. Compared to their *Compressed* unit counterparts, *Pay disparity* Low rank workers are significantly more likely to state that their wages were set unfairly in relation to their co-workers (Col. 4). Moreover, *Pay disparity* High wage workers are significantly less likely to state that their wages were set unfairly (Col. 2). Taken seriously, the survey responses of the High rank workers suggest that inequity aversion does not play a very meaningful role in explaining the main attendance results, as no fairness violation appears to be triggered by the higher levels of pay.

In addition, we also asked respondents to report their happiness on a ten-step ladder.⁵³ Panel B of 11 presents the results on this proxy of subjective well-being. Overall, *Pay disparity* units report being less happy than their *Compressed* unit counterparts (Col. 1). Despite believing that their wages are fair, High workers on *Pay disparity* units report the largest reductions in happiness (Col. 2).

5.6. **Discussion.** We find that *Pay disparity* causes all workers in our setting to decrease attendance and forego earnings, regardless of whether they are paid more or less than their co-workers. Moreover, the results paint a consistent picture of a drop in social cohesion within the production unit. *Pay disparity* undermines cooperation between workers on the same unit, causes a decrease in social ties outside of work, and leads to reports of more unhappiness at endline. These findings are consistent with the view in the social psychology and sociology literature that inequality in the workplace can lead to a break-down in social cohesion and an increase in social conflict.

This evidence also helps to shed light on why *Pay disparity* workers decrease attendance even when they are paid more than their co-workers. Our endline results suggest that while high rank workers consider their wages fair, the condition of *Pay disparity*, nonetheless, makes them less happy. This could reflect the fact that working in close proximity to aggrieved workers (i.e., those earning less than their unit-mates) could be unpleasant due to hostility, resentment, or social awkwardness. Our finding that co-workers in *Pay disparity* units are less likely to discuss their wages (Table 3) suggests that *Pay disparity* may also give rise to awkwardness that all workers in a unit might seek to avoid.

⁵³The survey question was taken directly from the World Values Survey (WVS). See Frey and Stutzer (2002), Luttmer (2005), and Benjamin et al. (2012) for previous work exploring the links between relative income and social status with happiness.

6. THREATS TO VALIDITY

6.1. Internal validity concerns. We now consider whether an explanation other than relative pay concerns could produce our findings.

Peer effects in attendance. Recall that in the endline games, *Pay disparity* workers perform worse with their own co-workers than with someone from another unit, with whom they have not worked at all. This implies that the decreased performance of *Pay disparity* units in the endline games does not arise from the fact that they have had less time to work together (due to lower attendance) than the *Compressed* units. However, could the main attendance results still be partly driven by traditional peer effects in attendance, leading us to overstate the treatment effect of relative pay? We explore this in Appendix Table A11, taking advantage of the randomization of workers to units. Because some villages have much higher attendance rates than others, the villages of one’s co-workers strongly predicts co-worker attendance (Col. 1).⁵⁴ If there were a causal attendance peer effect, then the villages of one’s co-workers should also predict one’s *own* attendance. However, this is not the case. Co-worker attendance—identified using the random composition of co-workers’ villages—does not predict one’s own attendance (Cols. 2-3).⁵⁵

Career concerns. We stress to workers that this is a one-time contract job—a common form of employment in this setting. Despite this, when a worker in *Pay disparity* observes he is paid less than his co-workers, he may believe the firm is less likely to re-hire him and therefore may decrease effort. However, such beliefs should be more likely when it is clear that the worker is much less productive than his peers. In contrast, we find decreased performance when workers perceive that they are close in productivity to their higher-paid peers. In addition, a basic career concerns story would not predict that workers would give up earnings while they have the job, or that higher-paid workers would also reduce attendance. It also cannot explain the social cohesion results.

Learning about one’s outside option. It is also unlikely that *Pay disparity* causes workers to infer that they could find higher wages at another firm. Even the lowest post-treatment wage, w_L , is above workers’ outside options. For example, Appendix Figure A1 documents that only 1.7% of workers state the prevailing wage in their local labor market to be above our training wage of Rs. 250. In addition, given that employment rates are low in the lean season (Table 2), it would be surprising for workers to be absent to search for alternate work, rather than waiting for the current contract to end. Finally, this hypothesis cannot explain why the higher-paid workers also attend less or why *Pay disparity* units do not cooperate well in the social cohesion games.

⁵⁴This is due to factors including distance to the factory, local social events, and local labor demand shocks.

⁵⁵In addition, while one’s own education predicts output levels, co-workers’ education does not predict own output in the post period. This indicates limited peer effects in output within the context of our experiment (e.g., Mas and Moretti 2009, Bandiera et al. 2010, Guiteras and Jack 2014).

Discouragement effects / self-signaling. Our results are not consistent with discouragement effects from the *Pay disparity* wages revealing a worker’s type (i.e., productivity). Similar arguments to those above make this explanation difficult to reconcile with our results. Namely, we find no effects of *Pay disparity* in the cases with perceived justifications, when it’s easiest for a worker to learn about his low productivity. In addition, we disclose information about relative productivity ranks to workers on Day 10—before treatment. We find little evidence that this disclosure has a differential impact on those who learn they are relatively less productive than their coworkers (Appendix Table A10).

Belief formation about employer. Finally, it is unlikely that *Pay disparity* affected performance because workers found differential pay unusual. In our setting, all workers are aware of performance pay regimes such as piece rates for harvesting, and about half of workers have worked under piece rates themselves. Indeed, there are no differential effects of *Pay disparity* for workers who have worked on piece rates versus those who have not. In addition, among local employers that pay flat daily wages, we observe differential pay based on experience or worker skill-levels.⁵⁶ Finally, such an explanation could not explain the impact of *Pay disparity* on workers’ ability to cooperate in the endline games, which had nothing to do with the employer.

6.2. External validity concerns. *Pay structure.* Because output in our production tasks is in principle measurable, firms could consider paying performance incentives like piece rates. However, doing so depends on the cost of monitoring. In the experiment, we bear considerable expense to hire extra staff to measure each worker’s output daily (e.g. count every single incense stick). Indeed, different local firms produce the same retail good under piece rates or under flat wages—with implications for the quality grade sold by the firm.⁵⁷

Dynamic incentives and pay policy. One potential benefit to firms of differential pay is dynamic incentives: workers know that if they work hard now, it could lead to higher pay in the future. Our study design explicitly shuts down this channel because after the training period, there is no possibility of a wage change.⁵⁸ The objective of our study is not to identify the optimal pay policy for firms, but rather, to isolate whether workers care

⁵⁶It is not surprising that pay differentials are often based on clear rules or formulas based on experience, skill level, or piece rates—our findings indicate the importance of such clear justifications.

⁵⁷For example, the contractors with whom we worked sold the output produced as part of the experiment at a relatively high quality grade (and therefore price premium). As another example, some employers pay piece rates while others in the same village pay fixed daily wages to harvest a given crop—again, with implications for quality. There is a large literature on such multitasking problems.

⁵⁸Note that fixed wage setting based on prior (or expected) productivity is not uncommon in many settings. As discussed above, workers in our study are routinely employed in one-time seasonal contract jobs, where the pay is fixed for the duration of the contract. In other contexts, firms often set the pay of short-term consultants based on expected productivity. Even for salaried workers, pay is usually based on ex-ante expectations, with stickiness throughout a worker’s tenure at the firm (e.g. Fehr et al. 2009)—this is not adjusted with new information on ex-post performance, but rather re-negotiated at infrequent intervals. More generally, explicit incentives like piece rates based on ex-post output are not very common in our setting (e.g. Dreze and Mukherjee 1989, Kaur 2015) or in the US (MacLeod and Parent 1999).

about relative pay. Optimal pay policy would depend on weighing the potential costs of differential pay (e.g. morale reductions) against the potential benefits (e.g. selection or increased effort). Our findings indicate that relative pay concerns could affect this calculus.

Magnitude of effects and context. Of course, as with any empirical study, the effect sizes we document are specific to our setting. While we cannot predict the magnitude of effects in other populations, a growing body of work in economics lends credence to the view that concerns about relative pay matter in a range of field settings including the US and Europe (Card et al. 2012, Cohn et al. 2012, Rege and Solli 2013, Dube et al. 2015).

7. CONCLUSION

The pattern of our findings broadly supports the view of pay disparity articulated in the psychology, sociology, and organizational economics literatures. These literatures posit that the workplace is a social organism, where the relational aspects of pay cannot be divorced from its economic value. They predict that being paid less than one’s peers is a disamenity; this accords with our finding that relatively lower paid workers substantially reduce output and attendance. These literatures also predict that discontent among some workers can break down social cohesion and cooperation, fostering social conflict and altering the dynamics of the group more broadly. This is consistent with our findings that *Pay disparity* workers are unable to cooperate at endline—doing worse when they work with co-workers from their own unit than with strangers.⁵⁹ The fact that even relatively higher-paid workers decrease attendance—coupled with their happiness responses on the endline survey—also provides some suggestive evidence in support of this view. If lower-paid workers were discontent or resentful, then working and eating lunch alongside them may have been socially awkward or unpleasant for their relatively higher paid peers, dampening their desire to go to work. Such externalities are in accordance with Frank (1984)’s observation that “Status is, like Coase’s social costs, a reciprocal phenomenon...[O]ne person’s gain in status can occur only at the expense of a loss in status for others.” This underscores the reason why our empirical results do not allow us to isolate workers’ individual internal preferences for relative pay. The effects we document are a reduced form combination of internal preferences and group-level dynamics.

An important theme throughout our results is that negative morale effects disappear when the reason for differential pay is extremely transparent. This could help advance our understanding of when pay compression is more likely to arise in the labor market. Specifically, in settings where it is difficult to quantify individual productivity, it will be challenging for an employer to justify pay differences. For example, in our experiment, managers informed each worker of his relative rank during the training period. Despite this, workers did not find pay differences acceptable unless, for example, output differences

⁵⁹This finding also has bearing on the more general idea that inequality may generate discord in society.

greatly overstated pay differences (Appendix Table A4).⁶⁰ This may help explain why, in the labor market, pay compression appears to be especially common among flat (e.g. hourly or daily) wage workers, but not among performance pay workers. Performance pay, such as piece rates or commissions, is possible only when it is feasible to quantify important dimensions of output. In such settings, the pay structure embeds a clear and justifiable mapping between output and pay; consequently, pay differentials across workers are less likely to violate fairness norms. In contrast, flat wages tend to arise in those settings where important dimensions of output are hard to quantify. Thus, even though managers may have a sense of which workers are more productive than others on average, setting pay according to expected output can be difficult to justify from the worker’s perspective. This suggests pay compression may be optimal in flat wage occupations, even when signals of differential productivity are available to employers—such as tollbooth attendants, supermarket cashiers, or agricultural day laborers (Frank 1984, Dreze and Mukherjee 1989). Furthermore, to the extent that wage compression plays a role in generating wage rigidity (e.g., Akerlof and Yellen 1990), our findings may have some relevance for understanding why wages appear to be more rigid among flat hourly workers relative to salaried workers with bonuses (e.g., Kahn 1997). More generally, while such implications are only suggestive, our perceived justifications results have potential relevance for understanding when relative pay concerns may affect wage structure and labor market outcomes.

In addition, our findings suggests that firms may have several potential tools at their disposal to manage morale in the presence of pay dispersion. For example, increased transparency in output could more easily allow firms to motivate workers through disparate wages—generating aggregate output benefits not just through a reduction in moral hazard, but also through improved morale. Firms could also potentially alter the organizational structure of the workplace itself—through job titles, physical co-location of similar workers, or the construction of teams or “units” (as we did in the experiment)—to affect who a worker views as being in her reference group.

While speculative, the above discussion suggests a variety of ways through which relative pay concerns could influence labor market phenomena—such as the returns to human capital investment, wage compression, wage rigidity, firm boundaries, and unemployment. Theoretical and empirical explorations of such relationships could enrich our view of agency theory and the functioning of the labor market.

REFERENCES

Adams, J. S.: 1963, Towards an understanding of inequity., *The Journal of Abnormal and Social Psychology* **67**(5), 422.

⁶⁰This matches the predictions of Fang and Moscarini (2005), who argue that self-serving bias will lead workers to feel aggrieved unless it’s clear that their peers are more deserving than themselves.

- Akerlof, G.: 1982, Labor contracts as partial gift exchange, *The Quarterly Journal of Economics* **97**(4), 543–569.
- Akerlof, G. A. and Yellen, J. L.: 1990, The fair wage-effort hypothesis and unemployment, *The Quarterly Journal of Economics* **105**(2), 255–283.
- Bandiera, O., Barankay, I. and Rasul, I.: 2010, Social incentives in the workplace, *The Review of Economic Studies* **77**(2), 417–458.
- Bandiera, O., Barankay, I. and Rasul, I.: 2013, Team incentives: evidence from a firm level experiment, *Journal of the European Economic Association* **11**(5), 1079–1114.
- Banerjee, A., Chandrasekhar, A. G., Duflo, E. and Jackson, M. O.: 2013, The diffusion of microfinance, *Science* **341**(6144), 1236498.
- Barankay, I.: 2016, Rank incentives: Evidence from a randomized workplace experiment, *Working Paper*.
- Baron, J. N. and Kreps, D. M.: 2013, Employment as an economic and a social relationship, in R. Gibbons and J. Roberts (eds), *Handbook of Organizational Economics*, Princeton University Press, pp. 315–341.
- Bartling, B. and von Siemens, F. A.: 2011, Wage inequality and team production: An experimental analysis, *Journal of Economic Psychology* **32**(1), 1–16.
- Benjamin, D. J., Heffetz, O., Kimball, M. S. and Rees-Jones, A.: 2012, What do you think would make you happier? what do you think you would choose?, *American Economic Review* **102**(5), 2083–2110.
- Bewley, T. F.: 1999, *Why wages don't fall during a recession*, Cambridge: Harvard University Press.
- Blanes i Vidal, J. and Nossol, M.: 2011, Tournaments without prizes: Evidence from personnel records, *Management Science* **57**(10), 1721–1736.
- Bracha, A., Gneezy, U. and Loewenstein, G.: 2015, Relative pay and labor supply, *Journal of Labor Economics* **33**(2), 297–315.
- Brown, G. D., Gardner, J., Oswald, A. J. and Qian, J.: 2008, Does wage rank affect employees? well-being?, *Industrial Relations: A Journal of Economy and Society* **47**(3), 355–389.
- Card, D., Heining, J. and Kline, P.: 2013, Workplace heterogeneity and the rise of west german wage inequality, *The Quarterly journal of economics* **128**(3), 967–1015.
- Card, D., Mas, A., Moretti, E. and Saez, E.: 2012, Inequality at work: The effect of peer salaries on job satisfaction, *American Economic Review* **102**(6), 2981–3003.
- Charness, G. and Kuhn, P.: 2007, Does pay inequality affect worker effort? experimental evidence, *Journal of Labor Economics* **25**(4), 693–723.
- Charness, G. and Kuhn, P.: 2011, Lab labor: What can labor economists learn from the lab?, *Handbook of Labor Economics* **4**, 229–330.
- Clark, A. E. and Oswald, A. J.: 1996, Satisfaction and comparison income, *Journal of Public Economics* **61**(3), 359–381.

- Clark, A., Masclet, D. and Villeval, M.-C.: 2010, Effort and comparison income: Experimental and survey evidence, *Industrial and Labor Relations Review* **63**(3), 407–426.
- Cohn, A., Fehr, E., Herrmann, B. and Schneider, F.: 2012, Social comparison in the workplace: Evidence from a field experiment, *IZA Discussion Paper No. 5550*.
- Cook, K. S. and Hegtvedt, K. A.: 1983, Distributive justice, equity, and equality, *Annual review of sociology* **9**(1), 217–241.
- Cullen, Z. and Pakzad-Hurson, B.: 2015, Equal pay for equal work? evidence from the renegotiation of short-term contracts online, *Working Paper, Stanford University*.
- DellaVigna, S., List, J., Malmendier, U. and Rao, G.: 2016, Estimating social preferences and gift exchange at work, *NBER Working Paper No. 22043*.
- Deutsch, M.: 1986, Cooperation, conflict, and justice, *Justice in social relations*, Springer, pp. 3–18.
- Dreze, J. and Mukherjee, A.: 1989, Rural labor markets in india, *The Balance Between Industry and Agriculture in Economic Development: Manpower and Transfers*, Vol. 3, London: Macmillan Press, pp. 223–265.
- Dube, A., Giuliano, L. M. and Leonard, J. S.: 2015, Fairness and frictions: The impact of unequal raises on quit behavior, *Working Paper*.
- Duesenberry, J. S.: 1949, *Income, saving, and the theory of consumer behavior*, Vol. 87, Cambridge: Harvard University Press.
- Duffy, M. K., Scott, K. L., Shaw, J. D., Tepper, B. J. and Aquino, K.: 2012, A social context model of envy and social undermining, *Academy of Management Journal* **55**(3), 643–666.
- Easterlin, R. A.: 1974, Does economic growth improve the human lot? some empirical evidence, in P. A. David and M. W. Reder (eds), *Nations and Households in Economic Growth: Essays in Honor of Moses Abramovitz*, Elsevier.
- Esteves-Sorenson, C. and Macera, R.: 2015, Gift exchange in the workplace: Addressing the conflicting evidence with a careful test, *Working Paper*.
- Falk, A., Fehr, E. and Fischbacher, U.: 2008, Testing theories of fairness? intentions matter, *Games and Economic Behavior* **62**(1), 287–303.
- Fang, H. and Moscarini, G.: 2005, Morale hazard, *Journal of Monetary Economics* **52**(4), 749–777.
- Fehr, E., Goette, L. and Zehnder, C.: 2009, A behavioral account of the labor market: The role of fairness concerns, *Annual Review of Economics* **1**(1), 355–384.
- Fehr, E., Kirchsteiger, G. and Riedl, A.: 1993, Does fairness prevent market clearing? an experimental investigation, *The Quarterly Journal of Economics* **108**(2), 437–459.
- Fehr, E. and Schmidt, K. M.: 1999, A theory of fairness, competition, and cooperation, *Quarterly journal of Economics* **114**(3), 817–868.
- Festinger, L.: 1954, A theory of social comparison processes, *Human relations* **7**(2), 117–140.
- Frank, R.: 1984, Are workers paid their marginal products?, *American Economic Review* **74**(4), 549–571.


- Frey, B. S. and Stutzer, A.: 2002, What can economists learn from happiness research?, *Journal of Economic Literature* **40**(2), 402–435.
- Gatcher, S. and Thoni, C.: 2010, Social comparison and performance: Experimental evidence on the fair wage–effort hypothesis, *Journal of Economic Behavior and Organization* **76**(3), 531–543.
- Gneezy, U. and List, J. A.: 2006, Putting behavioral economics to work: Testing for gift exchange in labor markets using field experiments, *Econometrica* **74**(5), 1365–1384.
- Guiteras, R. P. and Jack, B. K.: 2014, Incentives, selection and productivity in labor markets: Evidence from rural malawi, *NBER Working Paper No. 19825* .
- Hamermesh, D.: 2001, The changing distribution of job satisfaction, *Journal of Human Resources* **36**(1), 1–30.
- Hamermesh, D. S.: 1975, Interdependence in the labour market, *Economica* **42**(168), 420–429.
- Hicks, J.: 1963, *The Theory of wages*, London: Macmillan.
- Kahn, S.: 1997, Evidence of nominal wage stickiness from microdata, *The American Economic Review* **87**(5), 993–1008.
- Kaur, S.: 2015, Nominal wage rigidity in village labor markets, *NBER Working Paper No. 20770* .
- Keynes, J. M.: 1936, *The general theory of interest, employment and money*, London: Macmillan.
- Kószegei, B.: 2014, Behavioral contract theory, *Journal of Economic Literature* **52**(4), 1075–1118.
- Kube, S., Maréchal, M. A. and Puppe, C.: 2013, Do wage cuts damage work morale? evidence from a natural field experiment, *Journal of the European Economic Association* **11**(4), 853–870.
- Kuziemko, I., Buell, R., Reich, T. and Norton, M.: 2014, Last-place aversion: Evidence and redistributive implications", *Quarterly Journal of Economics* **129**(1), 105–149.
- Lazear, E. P.: 1989, Pay equality and industrial politics, *Journal of Political Economy* **97**(3), 561–580.
- Levine, D. I.: 1991, Cohesiveness, productivity, and wage dispersion, *Journal of Economic Behavior & Organization* **15**(2), 237–255.
- Levine, D. I.: 1993, What do wages buy?, *Administrative Science Quarterly* **38**(3), 462–483.
- Levitt, S. D. and List, J. A.: 2007, What do laboratory experiments measuring social preferences reveal about the real world?, *The Journal of Economic Perspectives* **21**(2), 153–174.
- List, J. A.: 2009, Social preferences: Some thoughts from the field, *Annu. Rev. Econ.* **1**(1), 563–579.
- Luttmer, E. F.: 2005, Neighbors as negatives: Relative earnings and well-being, *The Quarterly Journal of Economics* **120**(3), 963–1002.

- MacLeod, W. B. and Parent, D.: 1999, Job characteristics, wages and the employment contract, *Federal Reserve Bank of Saint Louis Review* **81**(3), 13–27.
- Marshall, A.: 1890, *Principles of Economics*, London, Macmillan.
- Mas, A. and Moretti, E.: 2009, Peers at work, *The American Economic Review* **99**(1), 112–145.
- Nickerson, J. A. and Zenger, T. R.: 2008, Envy, comparison costs, and the economic theory of the firm, *Strategic Management Journal* **29**(13), 1429–1449.
- Pfeffer, J. and Langton, N.: 1993, The effect of wage dispersion on satisfaction, productivity, and working collaboratively: Evidence from college and university faculty, *Administrative Science Quarterly* **38**(3), 382–407.
- Rege, M. and Solli, I.: 2013, Lagging behind the joneses: The impact of relative earnings on job separation, *Working Paper* .
- Schmitt, D. R. and Marwell, G.: 1972, Withdrawal and reward reallocation as responses to inequity, *Journal of Experimental Social Psychology* **8**(3), 207–221.
- Schofield, H.: 2014, The economic costs of low caloric intake: Evidence from india, *Working Paper* .
- Shaw, J. D.: 2014, Pay dispersion, *Annu. Rev. Organ. Psychol. Organ. Behav.* **1**(1), 521–544.
- Shaw, J. D.: 2015, Pay dispersion, sorting, and organizational performance, *Academy of Management Discoveries* **1**(2), 165–179.
- Song, J., Price, D. J., Guvenen, F., Bloom, N. and von Wachter, T.: 2015, Firming up inequality, *NBER Working Paper No. 21199* .
- Thibaut, J. W. and Kelley, H. H.: 1959, *The social psychology of groups.*, New York, John Wiley.
- Veblen, T. and Almy, C.: 1899, *The Theory of the Leisure Class: An Economic Study in the Evolution of Institutions*, Macmillan & Company, Limited.


FIGURES

FIGURE 1. Perceived Impact of Pay Disparity on Labor Supply and Cohesion


1) A laborer accepts a job from an external contractor at Rs. 240/day. The work begins the next day. At night, he learns that 2 other workers in the village have been offered the same job at:
 A) Rs. 240/day.
 B) Rs. 270/day.
 Do you think the worker would show up to work the next day?



2) There is a local soda factory that hires workers to load bottles onto trucks. Sunil accepts a contract to work at the factory for Rs. 240/day. The next morning, he arrives at work. He learns that other workers in his neighborhood have been hired for the same job at:
 A) Rs. 240/day.
 B) Rs. 260/day.
 How carefully will Sunil do the work relative to his usual level of effort?

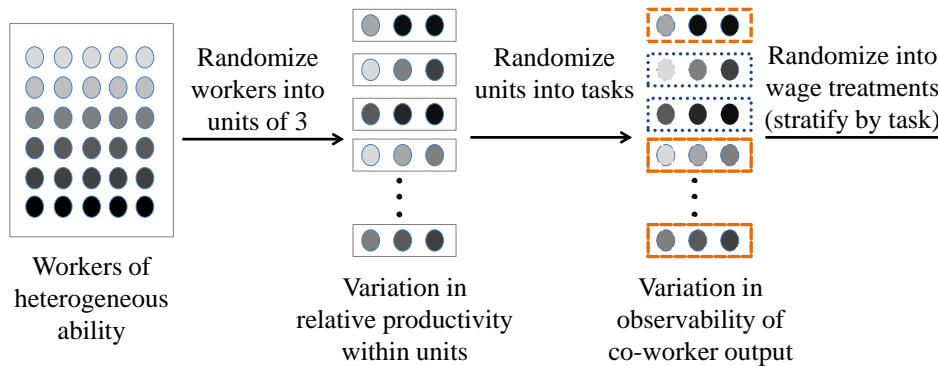


3) 3 people from a village get hired to work on a construction site together. The prevailing wage is Rs. 250. The contractor pays them:
 A) Rs. 250/day.
 B) Different wages based on their quality: Rs. 250/day, Rs. 270/day, and Rs. 290/day.
 How well will they work together?



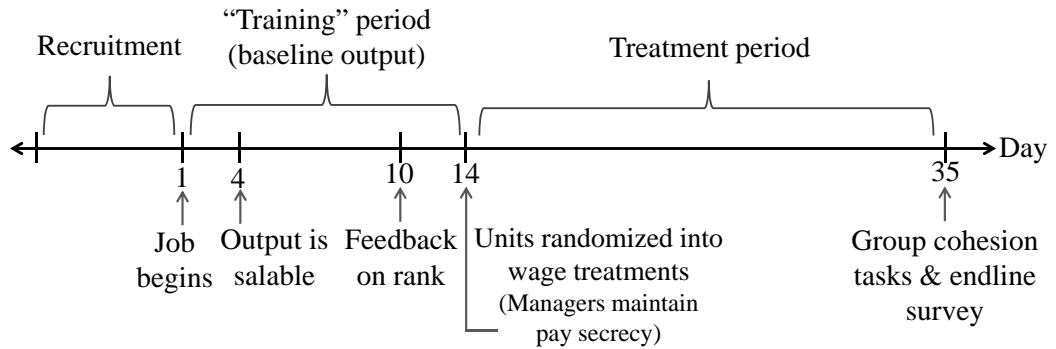
Notes: Surveys conducted with 200 laborers drawn from the same population as the experiment participants. Each respondent was asked 1 question out of each pair above. The x-axis in each chart shows the percentage of respondents choosing a given answer.

FIGURE 2. Randomization Design



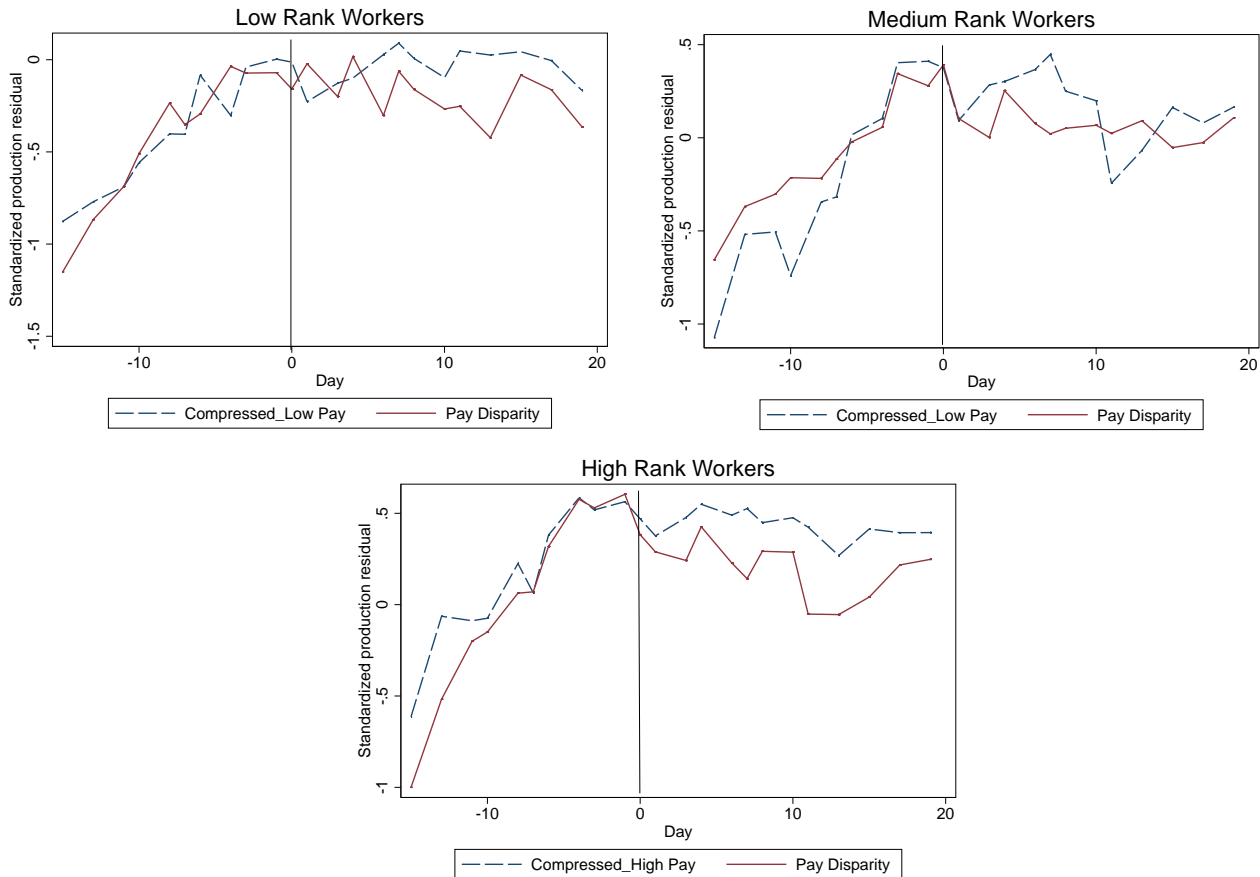
Notes: Three steps of randomization design. All unit, task and treatment assignments are made on day one of the experiment by the researchers. Factory management not informed of treatment status until the day of the wage change.

FIGURE 3. Time Line




Notes: This figure shows the timeline of activities in each of the 14 production rounds.

FIGURE 4. Effects of Pay Disparity on Worker Output




Notes: Y-axis shows residuals of standardized output after removing individual fixed effects in the pre-period and dummies for festivals. Each plot compares workers of same rank who earn the same wage, but are in *Compressed* vs. *Pay disparity* units. Day=0 is when wage treatments took effect.

FIGURE 5. Task Observability: Actual vs. Survey Correlations



Notes: Figure plots the correlation between actual productivity rankings and perceived rankings by the workers (reported after 3 weeks of work). Note that this data come from four pilot rounds, using a separate sample of workers, where all workers were paid the same wage and were not informed about their production rankings. In our analysis, we split the tasks at the median level of observability (0.5 correlation).

FIGURE 6. Daily wage distribution in rural Orissa



Notes: Wage data from 10-day recall surveys of 1207 workers conducted in 63 villages. Villages were randomly chosen from the same sample frame as those from which workers were recruited for the main experiment. The figure plots the distribution of reported daily wages minus the modal wage in the village.

TABLES

TABLE 1. Wage Treatments and Relevant Comparisons

Worker Rank	Pay Disparity	Compressed Low	Compressed Medium	Compressed High
Low productivity	W_{Low}	W_{Low}	W_{Medium}	W_{High}
Medium productivity	W_{Medium}	W_{Low}	W_{Medium}	W_{High}
High productivity	W_{High}	W_{Low}	W_{Medium}	W_{High}

Notes: Table presents randomization design. The key comparison groups in each row are highlighted. These six cells combined form the “relevant” group of workers.

TABLE 2. Summary Statistics

	Compressed (1)	Pay Disparity (2)	P-value of Difference (3)
Married	0.738 (0.440)	0.800 (0.402)	0.341
Number of Children	1.531 (1.490)	1.698 (1.502)	0.418
Hindu	0.996 (0.062)	1.000 (0.000)	0.319
Owens land	0.522 (0.500)	0.461 (0.501)	0.381
Sharecrops land	0.565 (0.497)	0.490 (0.502)	0.270
Has missed at least one meal in last 30 days	0.112 (0.315)	0.105 (0.309)	0.881
Number of days can't find work in last 30 days	14.612 (6.838)	15.260 (6.972)	0.474
Number of days worked in past 10 days	2.707 (2.736)	2.441 (2.585)	0.461
Number of days worked inside village in past 10 days	1.938 (2.436)	1.843 (2.512)	0.757
Total wage earnings over past 10 days	543.768 (708.571)	438.559 (587.394)	0.164
Typical wage in village during work period	216.464 (33.190)	211.505 (35.138)	0.383
Has experience working with piece rates	0.413 (0.493)	0.473 (0.502)	0.391
Baseline production (full training period)	-0.218 (0.688)	-0.199 (0.635)	0.870
Bseline attendance (full training period)	0.950 (0.082)	0.960 (0.074)	0.319

Notes: Responses taken from baseline surveys conducted on first day of work. Means and standard deviations are shown in cols (1) and (2) for Compressed and Pay Disparity units, respectively. Col (3) displays the p-values of the comparisons of means across Compressed and Pay Disparity production units obtained from a simple univariate regression, where standard errors are clustered by production unit. N=378.

TABLE 3. Knowledge of Co-worker Wages

	Compressed (1)	Pay Disparity (2)	P-value of Difference (3)
<i>Panel A — Own Unit</i>			
Indicator for knows correct wage of both co-workers	0.909 (0.288)	0.742 (0.440)	0.003***
Indicator for knows correct wage of at least one co-worker	0.958 (0.200)	0.871 (0.337)	0.031**
Indicator for has an opinion of wage of both co-workers	0.921 (0.271)	0.871 (0.337)	0.203
Indicator for has an opinion of wages of at least one co-worker	0.966 (0.181)	0.923 (0.265)	0.206
<i>Panel B — Other Units</i>			
Indicator for knows correct wage of all unit members	0.086 (0.280)	0.019 (0.138)	0.000***
Indicator for has an opinion of wages of all unit members	0.117 (0.321)	0.143 (0.351)	0.237
Indicator for has an opinion of wages of some unit members	0.134 (0.341)	0.174 (0.380)	0.070*

Notes: Responses taken from endline surveys in which workers were asked to list the wages of the members of their own production units and also the members of a set of four fixed product units. Means and standard deviations are shown in cols (1) and (2) for Compressed and Pay Disparity units, respectively. Col (3) displays the p-values of the comparisons of means across Compressed and Pay Disparity units obtained from a simple univariate regression, where standard errors are clustered by production unit. Panel A describes beliefs about a worker's own production unit. Panel B summarizes beliefs about the wages of individuals on other units: Compressed units in col (1) and Pay Disparity units in col (2). Observations in Panel B only include individuals who are evaluating the wages of units other than their own. N=358.

TABLE 4. Effects of Pay Disparity

	Output (std dev.) (1)	Output (std dev.) (2)	Attendance (3)	Attendance (4)	Output Attendance (5)
<i>Panel A — Pooled Treatment Effects</i>					
Post x Pay disparity	-0.311*** (0.110)	-0.242** (0.097)	-0.115*** (0.026)	-0.117*** (0.025)	-0.0893 (0.094)
<i>Panel B — Treatment Effects Separately by Rank</i>					
Post x Pay disparity x Low wage	-0.385*** (0.134)	-0.332** (0.128)	-0.113** (0.055)	-0.120** (0.053)	-0.204* (0.114)
Post x Pay disparity x Med wage	-0.262 (0.201)	-0.226 (0.187)	-0.126** (0.056)	-0.129** (0.060)	-0.0608 (0.138)
Post x Pay disparity x High wage	-0.288 (0.199)	-0.172 (0.181)	-0.106** (0.053)	-0.104** (0.052)	-0.00901 (0.152)
Individual fixed effects?	No	Yes	No	Yes	No
Post-treatment Compressed mean	-0.099	-0.099	0.939	0.939	0.015
N	8375	8375	8375	8375	7678

Notes : Difference in differences regressions. Panel A pools the treatment effects across the low, medium, and high rank workers, while Panel B shows the treatment effects separately by rank. Post is an indicator that equals 1 if the day is after workers have been randomized into wage treatments, and 0 during the baseline training period. Col (5) limits to observations where the worker was present. Regressions include day*round fixed effects, task-specific quadratic experience trends, and controls for neighboring production units. All coefficients are identified off comparisons of workers who earn the same absolute wage and have the same productivity rank within their production unit (see regression specification in text). Standard errors clustered by production unit.

TABLE 5. Effects Over Time

	Output (std dev.)			Attendance		
	Before first payday in post period (1)	Between first and second paydays in post period (2)	After second payday in post period (3)	Before first payday in post period (4)	Between first and second paydays in post period (5)	After second payday in post period (6)
<i>Panel A — Pooled Treatment Effects</i>						
Post x Pay disparity	-0.0530 (0.123)	-0.309*** (0.114)	-0.294** (0.133)	-0.0260 (0.037)	-0.133*** (0.029)	-0.149*** (0.042)
<i>Panel B — Treatment Effects Separately by Rank</i>						
Post x Pay disparity x Low wage	-0.0607 (0.187)	-0.325** (0.151)	-0.461*** (0.174)	0.0493 (0.087)	-0.103* (0.062)	-0.224*** (0.079)
Post x Pay disparity x Med wage	0.0541 (0.169)	-0.343* (0.195)	-0.231 (0.283)	-0.0400 (0.049)	-0.153** (0.063)	-0.126 (0.089)
Post x Pay disparity x High wage	-0.144 (0.235)	-0.259 (0.248)	-0.199 (0.203)	-0.0753 (0.069)	-0.144** (0.066)	-0.101 (0.069)
Individual fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes
Post-treatment Compressed Mean	0.081	-0.028	-0.264	0.973	0.945	0.916
N	4565	4670	5654	4565	4670	5654

Notes: Difference in differences regressions. Post is an indicator that equals 1 if the day is after workers have been randomized into wage treatments, and 0 during the baseline training period. All specifications include all of the pre-wage change observations. Col (1) and (4) only include post observations before the first payday. Col (2) and (5) only include post observations after the first but before the second payday. Col (3) and (6) only include post observations after the second payday. Regressions include individual fixed effects, day*round fixed effects, task-specific quadratic experience trends, and controls for neighboring production units. All coefficients are identified off comparisons of workers who earn the same absolute wage and have the same productivity rank within their production unit (see regression specification in text). Standard errors clustered by production unit.

TABLE 6. Mediating Effects of Perceived Justifications

Dependent variable	Definition of Perceived Justification Indicator					
	Baseline output difference between a worker and his higher ranked peer is large		Co-worker output is highly observable		Baseline differences are large or co-worker output is highly observable	
	Output (1)	Attendance (2)	Output (3)	Attendance (4)	Output (5)	Attendance (6)
<i>Panel A — Pooled Treatment Effects</i>						
Post x Pay disparity	-0.358*** (0.133)	-0.167*** (0.039)	-0.384*** (0.131)	-0.153*** (0.031)	-0.445*** (0.147)	-0.181*** (0.037)
Post x Pay disparity x Perceived justification	0.292* (0.173)	0.159** (0.061)	0.395** (0.161)	0.0996** (0.046)	0.429*** (0.154)	0.134*** (0.045)
Post x Perceived justification	0.0483 (0.107)	-0.0500 (0.032)	-0.0518 (0.103)	0.0332 (0.028)	-0.0342 (0.094)	0.0160 (0.025)
<i>Panel B — Treatment Effects Separately by Rank</i>						
Post x Pay disparity x Low wage	-0.448*** (0.147)	-0.168*** (0.061)	-0.513*** (0.160)	-0.158** (0.071)	-0.718*** (0.191)	-0.239*** (0.091)
Post x Pay disparity x Low wage x Perceived justification	0.467** (0.231)	0.181** (0.087)	0.512** (0.220)	0.121 (0.077)	0.712*** (0.230)	0.225** (0.096)
Post x Pay disparity x Med wage	-0.270 (0.224)	-0.170** (0.075)	-0.248 (0.227)	-0.157** (0.068)	-0.283 (0.280)	-0.187** (0.088)
Post x Pay disparity x Med wage x Perceived justification	0.127 (0.267)	0.150 (0.094)	0.0890 (0.293)	0.0876 (0.118)	0.0829 (0.311)	0.104 (0.115)
Post x Pay disparity x High wage			-0.386 (0.242)	-0.139* (0.071)	-0.372 (0.242)	-0.137* (0.072)
Post x Pay disparity x High wage x Perceived justification			0.582** (0.269)	0.0884 (0.078)	0.564** (0.269)	0.0845 (0.082)
R-squared	0.436	0.171	0.436	0.172	0.437	0.173
Number of observations (worker-days)	8375	8375	8375	8375	8375	8375

Notes: Panels A and B show comparisons of each worker in Pay disparity production units with the relevant worker (who has the same rank and absolute earnings level) in the Compressed production units. In Cols. (1)-(2), the Perceived justification indicator equals 1 if the baseline productivity difference between a worker and his higher-ranked co-worker (for low and medium rank workers) is above the Compressed group mean. In Cols. (3)-(4), this indicator equals 1 if the observability correlation for the worker's production task (computed using a separate baseline sample) is above the mean. In Cols. (5)-(6), this indicator equals 1 if the baseline productivity difference between a worker and his higher-ranked co-worker (for low and medium rank workers) is above the Compressed group mean or if the observability correlation for the worker's production task is above the mean. Regressions include individual fixed effects, day*round fixed effects, task-specific quadratic experience trends, and controls for neighboring production units. Standard errors clustered by production unit.

TABLE 7. Effects of Pay Disparity: Unit Level Variation

Dependent variable	<i>Definition of Perceived Justification Indicator</i>							
	Output (1)	Attendance (2)	Co-worker output is highly observable		Large baseline output difference between co-workers		Co-worker output is highly observable or baseline differences are large	
			Output (3)	Attendance (4)	Output (5)	Attendance (6)	Output (7)	Attendance (8)
<i>Panel A — Comparison with all Compressed Units</i>								
Post x Pay disparity	-0.134 (0.092)	-0.0735*** (0.026)	-0.273** (0.122)	-0.116*** (0.032)	-0.125 (0.120)	-0.0797*** (0.027)	-0.273** (0.125)	-0.111*** (0.033)
Post x Pay disparity x Perceived justification			0.383** (0.159)	0.119** (0.047)	0.241 (0.177)	0.0961* (0.051)	0.379** (0.160)	0.111** (0.047)
Post x Perceived justification			-0.0375 (0.100)	0.0140 (0.028)	0.0345 (0.100)	0.0413 (0.030)	-0.0241 (0.085)	0.0289 (0.024)
<i>Panel B — Comparison with Compressed Low Units Only</i>								
Post x Pay disparity	-0.148 (0.108)	-0.0597* (0.035)	-0.317** (0.139)	-0.0845** (0.042)	-0.136 (0.122)	-0.0605 (0.039)	-0.306** (0.144)	-0.0800* (0.045)
Post x Pay disparity x Perceived justification			0.465*** (0.163)	0.0892* (0.051)	0.127 (0.219)	0.106 (0.070)	0.432** (0.169)	0.0848 (0.053)
Post x Perceived justification			-0.122 (0.125)	0.0465 (0.035)	0.147 (0.157)	0.0371 (0.052)	-0.0784 (0.124)	0.0577 (0.035)
R-squared	0.432	0.165	0.433	0.167	0.432	0.166	0.433	0.167
Number of observations (worker-days)	8375	8375	8375	8375	8375	8375	8375	8375

Notes: Differences in differences regressions examining unit-level performance across wage treatments. Specifications are similar to the pooled individual regressions, but use variation from all workers, not only the "relevant" ones. Panel A examines the average performance of all workers on Pay disparity units relative to all workers in Compressed units. In Panel B, Pay disparity units are compared to the Compressed Low units only; these regressions include a dummy for being on a Compressed Medium or Compressed High unit, and an interaction of the dummy with the Perceived justification indicator (so the omitted category is Compressed Low). In Cols. (3)-(8), the Perceived justification indicator takes the same value for all workers in a unit. In Cols. (2)-(3), the Perceived justification indicator equals 1 if the worker was assigned to a production task where co-worker output is highly observable. In Cols. (5)-(6), this indicator equals 1 if the baseline productivity difference between each worker and his higher-ranked co-worker (for both the low and medium rank workers) is above the Compressed group mean. In Cols. (5)-(6), this indicator equals 1 if either the observability or large baseline productivity difference equals 1. All regressions include individual fixed effects, day*round fixed effects, task-specific quadratic experience trends, and controls for neighboring production units. Standard errors clustered by production unit.

TABLE 8. Effects on Group Cohesion: Endline Games 1

<i>Dependent variable: Tower height</i>				
	(1)	(2)	(3)	(4)
Pay disparity	-9.376*** (3.487)	-18.89** (8.068)	-20.49*** (5.532)	-25.72*** (6.923)
Pay disparity x Observable task		17.81* (9.472)		
Pay disparity x Large productivity difference			17.11** (6.815)	
Pay disparity x Perceived justification (aggregate)				19.66*** (7.445)
Dependent variable mean	53.97	53.97	53.97	53.97
R-squared	0.291	0.397	0.397	0.410

Notes: Observable task equals 1 if the observability correlation for the unit's production task (computed using a separate baseline sample) is above the mean. Large productivity difference equals 1 if the difference between a worker and his higher paid peer is above average (>0.375 standard deviations) for at least one member of the unit. Perceived justification (aggregate) is a boolean that equals 1 if either of the perceived justification measure dummies -- Observable task or Large productivity difference -- equals 1. All regressions include round fixed effects. N=80 production units. Games were only run in later rounds of the experiment. Robust standard errors.

TABLE 9. Effects on Group Cohesion: Endline Games 2

<i>Dependent variable: Number correct</i>				
	(1)	(2)	(3)	(4)
Both workers from same unit x Pay disparity	-1.105** (0.498)	-1.082** (0.499)	-1.480** (0.617)	-1.467** (0.616)
Both workers from same unit	0.350 (0.295)	0.498 (0.300)	0.502* (0.300)	0.502 (0.319)
At least one worker on Pay disparity unit	0.234 (0.413)	0.222 (0.406)	0.203 (0.504)	0.185 (0.504)
Both workers from same unit x Pay disparity x Perceived justification			0.598 (0.558)	0.613 (0.561)
At least one worker on Pay disparity unit x Perceived justification			0.0305 (0.489)	0.0221 (0.484)
At least one low rank worker in pair		-0.645** (0.261)	-0.659** (0.260)	
At least one medium rank worker in pair		-0.564* (0.286)	-0.587** (0.286)	
F-test p-value: Sum of first two coefficients	0.0536	0.135	0.0593	0.0639
Fixed effects for all rank combinations?	No	No	No	Yes
Dependent variable mean	4.085	4.085	4.085	4.085
R-squared	0.123	0.129	0.130	0.132

Notes: Dependent variable is the number of correct matches made by the pair. Perceived justification equals 1 if a worker on a pay disparity unit was assigned to an observable production task, or was substantially less productive than his higher paid peer, and equals 0 otherwise. This variable always equals 0 for workers in Compressed units. Regressions include fixed effects for the order in which a game was played during the day and the game station (location in worksite), as well as education controls (a dummy for whether both workers in the pair completed primary school, and a dummy for whether both can write a sentence in the local language). N=1904 pair-game observations. Games were only run in later rounds of the experiment. Standard errors are clustered by production unit.

TABLE 10. Social Cohesion Outside Work: Network Formation

Dependent variable	Number of within-unit network links (1)	Number of within-unit network links (2)	Indicator for network link with both co-workers in unit (3)	Indicator for network link with both co-workers in unit (4)
Pay disparity	-0.236* (0.124)	-0.349** (0.169)	-0.127* (0.070)	-0.173** (0.085)
Pay disparity x Perceived justification		0.208 (0.245)		0.0845 (0.137)
Compressed mean	0.779	0.779	0.265	0.265

Notes: Network links are measured from endline surveys. We define a link if the worker reports that he would go to the co-worker or the co-worker would come to him for: advice, borrowing money, or visiting one another's houses. In cols. (2) and (4), the perceived justification indicator equals 1 if the observability correlation for the worker's production task is above the mean or if any worker in the unit has a large productivity difference with his higher-ranked co-worker. Regressions include an indicator for having unit co-workers who are relatives. Standard errors clustered by production unit. N=358.


TABLE 11. Endline Survey: Fairness and Happiness

Sample restriction	All workers (1)	High rank only (2)	Medium rank only (3)	Low rank only (4)
<i>Panel A — Believes wage set fairly relative to co-workers</i>				
Pay disparity	-0.0143 (0.091)	0.363*** (0.120)	-0.163 (0.170)	-0.587*** (0.186)
Pay disparity x Perceived justification	0.0154 (0.118)	-0.108 (0.190)	0.111 (0.228)	0.470* (0.252)
R-squared	0.0840	0.244	0.203	0.239
<i>Panel B — Above-median happiness (World Values Survey)</i>				
Pay disparity	-0.292*** (0.111)	-0.448*** (0.163)	-0.241 (0.172)	-0.146 (0.168)
Pay disparity x Perceived justification	0.298** (0.133)	0.349 (0.225)	0.307 (0.220)	0.241 (0.205)
R-squared	0.190	0.315	0.292	0.278
N	358	121	119	118

Notes: OLS regressions using endline survey responses. Panel A dependent variable is whether the worker viewed his wages as “fair” or “very fair” (relative to a 5 point scale). Panel B dependent variable is an indicator for above-median happiness (World Values Survey). Perceived justification indicator equals 1 if the observability correlation for the worker's production task is above the mean or if the baseline productivity difference between a worker and his higher-ranked co-worker (for low and medium rank workers) is above the Compressed group mean. All specifications include task and round fixed effects. Col (1) includes all observations. Cols (2) - (4) restrict the samples to the high, medium, and low rank workers. Standard errors clustered by production unit in Col (1). Robust standard errors are used in the single-rank specifications.


APPENDIX A. SUPPLEMENTAL FIGURES AND TABLES

FIGURE A1. Distribution of Village Prevailing Wages: Endline Survey Responses




Notes: Histogram plots distribution of the prevailing wage, as reported by workers.. Responses from endline survey responses of 356 individuals from 39 villages. Note only 1.7% of workers indicate a prevailing wage above Rs. 250, which is the training wage in all rounds. The modal village wage is always \leq Rs.250.

FIGURE A2. Total Attendance on Pay Disparity and Compressed Units




Notes: Kernel density plots of worker attendance rates during the post-wage change period among “relevant” workers, separately for Pay disparity and Compressed production units. Attendance rate is measured as the fraction of days worker was present in the post-wage change period.


FIGURE A3. Density of Raw Production by Task, Scaled by Mean




(A) Correlation = 0




(B) Correlation = 0.15




(C) Correlation = 0.39




(D) Correlation = 0.41




(E) Correlation = 0.42




(F) Correlation = 0.50




(G) Correlation = 0.52



(H) Correlation = 0.62



(I) Correlation = 0.75




(J) Correlation = 0.88

Notes: Density plots of raw production, scaled by mean. Subfigures ordered by task observability.


FIGURE A4. Cooperative Puzzle Games - Examples

Panel A: Spot the Difference - Example

Sheet 1, Player 1



Sheet 2, Player 2



Panel B: Symbol Matching - Example

Sheet 1, Player 1


β	ϕ	÷	Φ	÷	↓	∞	ϕ	ξ	T	μ	λ	†	↓	β
ς	≡	≡	ϕ	†	↓	∞	T	ε	μ	¶	∞	÷	β	ς
Z	ξ	÷	¶	λ	ξ	∞	β	μ	Φ	β	∞	≡	ζ	ξ
¶	¶	↓	↓	β	∞	ϕ	ξ	Λ	Λ	∞	ς	ς	¶	¶
†	ς	ε	μ	μ	ε	∞	T	β	T	↓	≡	T	∞	Φ
β	ϕ	↓	ζ	Φ	ξ	ϕ	¶	β	Φ	ς	∞	≡	λ	ε
ϕ	ξ	¶	ς	∞	β	θ	∞	†	↓	↓	∞	λ	ϕ	ξ
ε	μ	÷	↓	÷	↓	∞	T	ς	∞	¶	ξ	μ	÷	
T	ϕ	ζ	Φ	ξ	Φ	¶	β	ς	ε	θ	T	ϕ	Z	
ς	↓	≡	ς	†	∞	Λ	T	β	÷	λ	†	¶	¶	ς

Sheet 2, Player 2

μ	ξ	÷	†	÷	ε	β	ξ	∞	ξ	ζ	β	μ	ξ	ξ
Φ	ξ	≡	θ	ε	↓	∞	†	Ω	ε	θ	Φ	ξ	λ	θ
↓	ξ	ς	ξ	ς	ξ	∞	θ	†	T	∞	Ω	ς	↓	
Ω	ζ	∞	T	T	β	ϕ	θ	Z	ς	θ	T	Λ	ζ	ς
Λ	↓	ϕ	≡	↓	Ω	∞	ϕ	θ	ϕ	T	†	ϕ	ϕ	Λ
ς	μ	ξ	T	ζ	μ	ϕ	ζ	Λ	Φ	↓	θ	Φ	÷	μ
β	↓	ξ	≡	θ	≡	≡	θ	Λ	÷	ς	∞	ζ	ε	β
Ω	ξ	ξ	β	∞	Λ	∞	Λ	ς	β	ξ	ϕ	θ	Λ	ξ
ε	∞	ϕ	T	†	↓	Φ	ϕ	Ω	∞	ε	∞	ϕ	T	†
†	†	∞	ς	÷	ς	≡	ϕ	Λ	β	λ	β	Φ	¶	†

Notes: Examples of the cooperative pair games. Each worker in a pair would receive one of the sheets. Workers had to compare their respective sheets, and circle items that were different (Spot the Difference) or matched (Symbol Matching) on both their sheets.

FIGURE A5. Stability of Relative Productivity Ranks Across Time



Notes: The y-axis shows the residual of standardized output after removing individual fixed effects in the pre-period (before day 0) and dummies for festival days. The figures plot, for each day of the experiment, the average of the residuals for each group of workers restricting to members of the Compressed production units. Day=0 is the day wage treatments took effect (i.e. when workers were told their post-training wage).

TABLE A1. Effects of Pay Disparity: Robustness to Alternate Specifications

	Output (std dev.)	Attendance	Output (std dev.)	Attendance	Output (std dev.)	Attendance	Output (std dev.)	Attendance
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
<i>Panel A — Pooled Treatment Effects</i>								
Post x Pay disparity	-0.222*** (0.070)	-0.0961*** (0.019)	-0.272*** (0.086)	-0.124*** (0.022)	-0.242** (0.097)	-0.117*** (0.025)	-0.272** (0.111)	-0.115*** (0.028)
<i>Panel B — Treatment Effects Separately by Rank</i>								
Post x Pay disparity x Low wage	-0.279** (0.110)	-0.100** (0.043)	-0.400*** (0.120)	-0.115** (0.049)	-0.332** (0.128)	-0.120** (0.053)	-0.292** (0.134)	-0.0987* (0.054)
Post x Pay disparity x Med wage	-0.180 (0.127)	-0.0995** (0.045)	-0.278 (0.194)	-0.148** (0.065)	-0.226 (0.187)	-0.129** (0.060)	-0.291 (0.198)	-0.146** (0.062)
Post x Pay disparity x High wage	-0.205 (0.140)	-0.0882** (0.037)	-0.145 (0.174)	-0.111** (0.052)	-0.172 (0.181)	-0.104** (0.052)	-0.238 (0.202)	-0.103** (0.052)
Sample	Relevant	Relevant	Relevant	Relevant	Full	Full	Full	Full
Include Pre-treatment observations?	Yes	Yes	Yes	Yes	Yes	Yes	No	No
Include "Irrelevant" workers?	No	No	No	No	Yes	Yes	Yes	Yes
Neighbor Controls?	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Individual fixed effects?	Yes	Yes	Yes	Yes	Yes	Yes	No	No
Post-treatment Compressed Mean	-0.099	0.939	-0.099	0.939	-0.099	0.939	-0.099	0.939
N	4307	4307	4307	4307	8375	8375	5283	5283

Notes : Difference in differences regressions are presented in Cols (1)-(6). Cols (7)-(8) only include post-training, post-treatment observations. Panel A pools the treatment effects across the low, medium, and high rank workers, while Panel B shows the treatment effects separately by rank. Post is an indicator that equals 1 if the day is after workers have been randomized into wage treatments, and 0 during the baseline training period. Regressions include day*round fixed effects, and task-specific quadratic experience trends. Cols (3)-(8) include controls for neighboring production units. All coefficients are identified off comparisons of workers who earn the same absolute wage and have the same productivity rank within their production unit (see regression specification in text), i.e., from the so-call "relevant" workers. Standard errors clustered by production unit.

TABLE A2. Effects on Output Quality

	Quality rating (1)	High quality rating (2)	Proportion rejected (3)
<i>Panel A — Pooled</i>			
Post x Pay disparity	-0.0128 (0.026)	-0.0907 (0.092)	-0.0123 (0.017)
<i>Panel B — Separately by Rank</i>			
Post x Pay disparity x Low wage	0.00705 (0.050)	-0.0471 (0.173)	0.0187 (0.021)
Post x Pay disparity x Med wage	-0.0361 (0.046)	-0.183 (0.164)	-0.0396 (0.039)
Post x Pay disparity x High wage	-0.0182 (0.026)	-0.109 (0.110)	-0.00563 (0.018)
Post-treatment Compressed mean	0.823	0.842	0.012
N	3868	3868	1669

Notes : Difference in differences regressions. The sample is restricted to days when a worker was present. Cols. (1)-(2) show effects on supervisors' subjective daily assessment of the quality of each worker's output (collected in only a subset of days and rounds). Quality rating is a proportion: the rating score divided by the maximum possible rating (either 3 or 5). High quality rating is a boolean for whether the worker's quality rating on that day was above 0.8. The dependent variable in Col. (3) is the proportion of output that was rejected due to substandard quality; this is only measured for the two production tasks where quality standards were quantifiable (candle wicks and incense sticks). Panel A pools effects across the low, medium, and high rank workers, while Panel B shows effects separately by rank. Post is an indicator that equals 1 if the day is after workers have been randomized into wage treatments, and 0 during the baseline training period. Regressions include individual fixed effects, day*round fixed effects, task-specific quadratic experience trends, and controls for neighboring units. Note that caution must be used when interpreting these coefficients, since the sample conditions on attendance. Standard errors clustered by production unit.

TABLE A3. Perceived Justifications: Robustness of Relative Productivity Results

Dependent variable	Controls for own baseline productivity x post x rank		Controls for baseline productivity lowest 10% x post x treatment		Controls for own baseline productivity x post x treatment x rank	
	Output	Attendance	Output	Attendance	Output	Attendance
	(1)	(2)	(3)	(4)	(5)	(6)
Post x Pay disparity	-0.353** (0.137)	-0.158*** (0.042)	-0.413*** (0.134)	-0.179*** (0.042)	-0.456*** (0.144)	-0.178*** (0.042)
Post x Pay disparity x Perceived justification	0.299* (0.177)	0.152** (0.063)	0.292* (0.171)	0.162*** (0.061)	0.221 (0.178)	0.129* (0.068)
R-squared	0.436	0.173	0.438	0.172	0.439	0.175
Number of observations (worker-days)	8375	8375	8375	8375	8375	8375

Notes: Table shows comparisons of each worker in Pay disparity teams with the relevant worker (who has the same rank and absolute earnings level) in the Compressed teams. The Perceived justification indicator equals 1 if the baseline productivity difference between a worker and his higher-ranke co-worker (for low and medium rank workers) is above the Compressed group mean. Cols. (1)-(2) include time- and rank- varying controls for baseline productivity. Cols. (3)-(4) include time- and treatment- varying controls for whether the worker's baseline productivity was in the lowest 10% of the distribution (for Low workers). Cols. (5)-(6) add time-, treatment-, and rank- varying controls for the worker's baseline productivity. Regressions include individual fixed effects, day*round fixed effects, task-specific quadratic experience trends, and controls for neighboring production units. Standard errors clustered by production unit.

TABLE A4. Perceived Justifications: Robustness to Alternative Cutoffs for Relative Productivity Thresholds

Dependent variable	Definition of Threshold for Relative Productivity Fraction							
	>10% difference		>15% difference		>20% difference		>25% difference	
	Output (1)	Attendance (2)	Output (3)	Attendance (4)	Output (5)	Attendance (6)	Output (5)	Attendance (6)
Post x Pay disparity	-0.270 (0.170)	-0.122** (0.052)	-0.321** (0.145)	-0.156*** (0.040)	-0.355*** (0.135)	-0.177*** (0.037)	-0.351** (0.138)	-0.173*** (0.039)
Post x Pay disparity x Perceived justification	0.00679 (0.187)	-0.00936 (0.076)	0.151 (0.164)	0.0790 (0.064)	0.332** (0.166)	0.178*** (0.060)	0.313* (0.175)	0.171** (0.067)
R-squared	0.434	0.170	0.435	0.170	0.436	0.172	0.436	0.172
Number of observations (worker-days)	8375	8375	8375	8375	8375	8375	8375	8375

Notes: Regressions show comparisons of each worker in Pay disparity teams with the relevant worker (who has the same rank and absolute earnings level) in the Compressed production units, heterogeneously by Perceived justifications. Results are pooled across the Low and Medium rank workers. Perceived justifications indicator equals 1 if the difference between a worker and his higher-ranked peer (for low and medium rank workers) is greater than X%, where X is stated at the top of each column. Regressions include individual fixed effects, day*round fixed effects, task-specific quadratic experience trends, and controls for neighboring production units. Standard errors clustered by production unit.

TABLE A5. Perceived Justifications: Task Observability Robustness

Sample restriction: drop tasks	Task 1	Task 2	Task 3	Task 4	Task 5	Task 6	Task 7	Task 8	Task 9	Task 10
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
<i>Panel A — Output (std. dev)</i>										
Pay disparity	-0.359*** (0.125)	-0.381*** (0.130)	-0.423*** (0.137)	-0.385*** (0.129)	-0.391*** (0.135)	-0.115 (0.121)	-0.399*** (0.129)	-0.533*** (0.145)	-0.516*** (0.149)	-0.378** (0.151)
Pay disparity x Perceived justification	0.393** (0.180)	0.386** (0.164)	0.453*** (0.171)	0.373** (0.167)	0.430** (0.183)	0.177 (0.156)	0.313* (0.160)	0.563*** (0.166)	0.513*** (0.177)	0.351* (0.184)
<i>Panel B — Attendance</i>										
Pay disparity	-0.143*** (0.032)	-0.152*** (0.034)	-0.151*** (0.035)	-0.152*** (0.032)	-0.153*** (0.034)	-0.144*** (0.036)	-0.156*** (0.033)	-0.205*** (0.029)	-0.139*** (0.041)	-0.125*** (0.029)
Pay disparity x Perceived justification	0.141*** (0.049)	0.101** (0.049)	0.0901* (0.049)	0.103** (0.047)	0.105* (0.056)	0.0871* (0.047)	0.0672 (0.055)	0.155*** (0.046)	0.0885* (0.053)	0.0719 (0.047)
N	7459	7638	7453	7638	7511	7500	7482	7573	7565	7556

Notes: Table shows specification for main observability regressions, but sequentially drops each production task from the regression. Comparisons are between each worker in Pay disparity production units with the relevant worker (who has the same rank and absolute earnings level) in the Compressed production units. Treatment effects are pooled across the Low, Medium and high Rank workers. The Perceived justification indicator equals 1 if the observability correlation for the worker's production task (computed using a separate baseline sample) is above the mean. Regressions include individual fixed effects, day*round fixed effects, task-specific quadratic experience trends, and controls for neighboring production units. Standard errors clustered by production unit.

TABLE A6. Absence of Correlation Between Task Observability Measure and Training Outcomes

	Attendance (1)	Attendance (2)	Did not complete training (3)	Output growth (std dev.) (4)
Observable task	-0.0107 (0.011)	-0.0103 (0.008)	-0.00476 (0.011)	-0.0125 (0.151)
Sample	All hired workers	Training completors	All hired workers	Training completors
Frequency of data	Daily	Daily	Worker-level	Worker-level
Pre-treatment Mean	0.914	0.954	0.019	0.233
N	3557	3092	415	378

Notes: OLS regressions are presented in Cols (1)-(4) to measure pre-treatment correlates with observability. The regressor in all specifications is an indicator for whether the production task has above-median observability. Cols (2) and (4) restrict the analysis to only workers who completed the training period and were thus available to be randomized into treatments. Cols (1)-(2) use daily attendance as the dependent variable. Col (3) captures whether a worker terminated his employment before the end of the training period. Col (4) measures output growth in the training period by taking the difference between standardized output on days 8 and 3. This captures the steepness of the learning curve for each worker. Regressions include round fixed effects. Standard errors clustered by production unit.

TABLE A7. Sample Overlap Between Perceived Justification Measures

	Below mean observability (Total 51.59%)	Above mean observability (Total 48.41%)
Below mean production difference (Total 69.44%)	34.52%	34.92%
Above mean production difference (Total 30.56%)	17.06%	13.49%

Notes: Tabulations of overlap between the two sources of variation for perceived justifications: baseline productivity differences between a worker and his higher ranked peer, and the observability of co-worker output. The binary splits for each measure shown here are the same as those used in the tables in the analysis. The table shows the percentage of observations in each cell.

TABLE A8. Effects of Higher Absolute Pay among Compressed Units

	Output (std dev.)			Attendance		
	First two days post wage change (1)	First week post wage change (2)	Full sample (3)	First two days post wage change (4)	First week post wage change (5)	Full sample (6)
Post x Compressed Medium	0.00218 (0.144)	0.0875 (0.129)	0.0184 (0.138)	0.0411 (0.044)	0.0570* (0.033)	0.0225 (0.039)
Post x Compressed High	-0.00447 (0.109)	0.0274 (0.099)	-0.00973 (0.099)	0.0691* (0.037)	0.0768** (0.031)	0.0321 (0.034)
Post-treatment Compressed Low Mean	-0.212	-0.237	-0.292	0.891	0.894	0.899
N	2807	3620	6107	2807	3620	6107

Notes : Difference in differences regressions restricting the sample to only Compressed teams. Post is an indicator that equals 1 if the day is after workers have been randomized into wage treatments, and 0 during the baseline training period. Regressions use unit-level variation, so all Compressed wage unit workers are included. Omitted category is Compressed low wage units. All specifications include the full training period. Cols (1) and (4) include only the first two post wage change days. Cols (2) and (5) include only the first work week post wage change. Col (3) and (6) include the full sample. Regressions include individual fixed effects, day*round fixed effects, task-specific quadratic experience trends, and controls for neighboring production units. Standard errors clustered by production unit.

TABLE A9. Robustness: Effects on Group Cohesion - Conditional on Attendance

Dependent variable			Game: Tower building in teams		Game: Cooperative puzzles in pairs	
	Attendance	Attendance	Tower height	Tower height	Number correct	Number correct
	(1)	(2)	(3)	(4)	(5)	(6)
Pay disparity	-0.0106 (0.031)	-0.00889 (0.034)	-10.18** (4.354)	-12.98** (6.241)		
Compressed_Medium pay		0.0198 (0.044)		-7.400 (6.429)		
Compressed_High pay		-0.0150 (0.053)		-0.936 (8.362)		
Both workers from same unit x Pay disparity					-0.939** (0.431)	
Both workers from same unit					0.404* (0.227)	
At least one worker in pair was on Pay disparity unit					0.203 (0.325)	
At least one low rank worker in pair					-0.604** (0.236)	-0.714*** (0.252)
At least one medium rank worker in pair					-0.387* (0.219)	-0.398* (0.216)
Dependent variable mean	0.921	0.921	54.17	54.17	4.695	4.695
Observations	240	240	65	65	1632	1632
R-squared	0.170	0.172	0.318	0.338	0.203	0.105

Notes : This table replicates results for cooperative games, conditional on attendance. Cols 1-2 include the full sample of workers who participated in endline games (run only in the later rounds of the experiment). Cols. 3-4 limit analysis to units where all 3 unit members were present on the day of endline games. Cols. 1-4 include round fixed effects. Cols. 5-6 limit analysis to pair-games where both workers in an assigned pair were present the day of endline games. Col. (5) includes fixed effects for the order in which a game was played during the day and the game station (location in worksite) as well as education controls. Standard errors are clustered by production unit.

TABLE A10. Response to Disclosure of Productivity Ranks in Training Period

	Output (std dev.) (1)	Output (std dev.) (2)	Output (std dev.) (3)	Attendance (4)
Post rank disclosure x Low rank (pre disclosure)	0.00346 (0.089)	0.00801 (0.093)	-0.000635 (0.094)	0.00146 (0.024)
Post rank disclosure x Medium rank (pre disclosure)	0.00828 (0.058)	-0.00984 (0.059)	-0.00891 (0.060)	0.00303 (0.020)
Experience x task controls?	No	Quadratic trends	Fixed effects	Fixed effects
Mean of dependent variable: 3 days prior to ranking treatment	-0.322	-0.322	-0.322	0.969
F-test p-value: Post ranking x Low = Post ranking x Medium	0.943	0.814	0.911	0.935
N	3060	3060	3060	3060

Notes: Difference in differences regressions presented to test whether workers of different ranks respond differently to the information about their relative productivity that was revealed four days (on average) before the end of the training period. Here, rank is calculated based on production during the three days prior to the release of the ranking information. Regressions only contain data from the training period, and Post rank disclosure is an indicator for the four days (on average) following disclosure before the end of the training period. All regressions include individual fixed effects as well as worksite x day fixed effects. The coefficients are all relative to the output and attendance of the person with the highest rank. Regressions include individual fixed effects and day*round fixed effects. The specifications that include experience x task fixed effects are identified using variation in the lengths of the training periods across rounds. Standard errors clustered by production unit.

TABLE A11. Lack of Evidence of Attendance Peer Effects

Dependent variable	First Stage	Reduced Form	IV (2SLS)
	Co-worker attendance (1)	Own Attendance (2)	Own Attendance (3)
Predicted attendance of co-workers	0.600** (0.278)	0.0035 (0.169)	
Actual attendance of co-workers (instrumented)			0.0063 (0.279)
F-stat: First stage excluded instrument	4.65		
Number of observations (worker-days)	6020	6020	6020

Notes: Table shows peer effects regressions instrumenting actual peer attendance with predicted peer attendance. Predicted attendance is calculated using the average attendance of all workers from the same village excluding the worker and others from his production unit (i.e., a jack-knife or leave-one-out strategy). Observations limited to workers on Compressed pay units. Column 1 presents the first stage regression of actual peer attendance on predicted peer attendance. Column 2 presents the reduced form regression of own attendance on predicted peer attendance. Column 3 presents the IV estimation of the effects of peer attendance on own attendance. All specifications include fixed effects for round, calendar month, and day of round, and indicators for being assigned to a unit with one or more co-workers from the same village. Standard errors clustered at the production unit level.

TABLE A12. Supplemental Survey: Fairness of Different Wage Regimes

Outcome: Percentage of respondents saying the scenario is "Unfair" or "Very Unfair"

	Pay type	
	A) flat daily wage (1)	B) piece rate (2)
1 A brick factory in an area pays its workers a: A) flat daily wage B) piece-rate for laying bricks. Balu and Mohit are both laborers who have worked in the factory for 6 months. Balu earns Rs. 20 more per day.	98%	0%
2 A brick factory in an area pays its workers a: A) flat daily wage B) piece-rate for laying bricks. Balu and Mohit are both laborers who have worked in the factory for 6 months. Balu, <i>who makes more bricks than Mohit</i> , earns Rs. 20 more per day.	16%	4%

Notes: Survey conducted with workers who were not a part of the experiment, but drawn from a similar population in the area where the experiment was conducted. N=200 workers. Each worker was only asked one of the 4 questions shown in the table. Respondents were given one of the 4 scenarios, and asked to rate them as "Completely fair", "Acceptable", "Unfair", or "Very unfair".